

Looking forward to lepton-flavor-violating ALPs

Lorenzo Calibbi,^a Diego Redigolo,^{b,c} Robert Ziegler^d and Jure Zupan^e

^a*School of Physics, Nankai University,
Tianjin 300071, China*

^b*INFN Sezione di Firenze and Department of Physics and Astronomy, University of Florence,
Via G. Sansone 1, I-50019 Sesto Fiorentino, Italy*

^c*CERN, Theory Division,
CH-1211 Geneva 23, Switzerland*

^d*Institute for Theoretical Particle Physics, Karlsruhe Institute of Technology,
Engesserstrasse 7, D-76128 Karlsruhe, Germany*

^e*Department of Physics, University of Cincinnati,
Cincinnati, Ohio 45221, U.S.A.*

E-mail: calibbi@nankai.edu.cn, d.redigolo@gmail.com,
robert.ziegler@kit.edu, zupanje@ucmail.uc.edu

ABSTRACT: We assess the status of past and future experiments on lepton flavor violating (LFV) muon and tau decays into a light, invisible, axion-like particle (ALP), a . We propose a new experimental setup for MEG II, the *MEGII-fwd*, with a forward calorimeter placed downstream from the muon stopping target. Searching for $\mu \rightarrow ea$ decays MEGII-fwd is maximally sensitive to LFV ALPs, if these have nonzero couplings to right-handed leptons. The experimental set-up suppresses the (left-handed) Standard Model background in the forward direction by controlling the polarization purity of the muon beam. The reach of MEGII-fwd is compared with the present constraints, the reach of Mu3e and the Belle-II reach from $\tau \rightarrow \ell a$ decays. We show that a dedicated experimental campaign for LFV muon decays into ALPs at MEG II and Mu3e will be able to probe the ALP parameter space in an unexplored region well beyond the existing astrophysical constraints. We study the implications of these searches for representative LFV ALP models, where the presence of a light ALP is motivated by neutrino masses, the strong CP problem and/or the SM flavor puzzle. To this extent we discuss the majoron in low-scale seesaw setups and introduce the LFV QCD axion, the LFV axiflavor and the leptonic familon, paying particular attention to the cases where the LFV ALPs constitute cold dark matter.

KEYWORDS: Beyond Standard Model, Effective Field Theories, Cosmology of Theories beyond the SM

ARXIV EPRINT: [2006.04795](https://arxiv.org/abs/2006.04795)

Contents

1	Introduction	1
2	Notation and summary	4
3	ALPs in $\mu^+ \rightarrow e^+ + \text{invis. decays}$	10
3.1	Past searches at TRIUMF: Jodidio et al. and TWIST	13
3.2	New searches at PSI: MEGII-fwd and Mu3e-online	16
3.3	The potential at $\mu^- \rightarrow e^-$ conversion experiments	20
4	ALPs in $\mu \rightarrow e + \gamma + \text{invis. decays}$	21
4.1	Past searches at Crystal Box	21
4.2	Possibilities for future search at MEG II	24
5	ALPs in τ decays	24
5.1	Past searches at ARGUS	24
5.2	Future searches at Belle-II	25
6	Bounds from astrophysics and cosmology	26
6.1	Bounds from stellar cooling	26
6.2	ALP Dark Matter	30
7	LFV ALP models	34
7.1	The LFV QCD axion	35
7.2	The LFV axiflavor	40
7.3	The leptonic familon	44
7.4	The majoron	47
8	Conclusions	52

1 Introduction

Probes of the Standard Model (SM) based on rare processes with charged leptons are set to improve substantially in the next decade. The muon beam experiments MEG II [1], Mu3e [2, 3], COMET [4] and Mu2e [5] will collect unprecedented datasets using $\mathcal{O}(10^{15} - 10^{17})$ muons each. Similarly, Belle-II is expected to collect roughly 5×10^{10} $\tau^+\tau^-$ pairs [6], exceeding by more than an order of magnitude the datasets at Belle and BaBar. The standard New Physics (NP) targets for these experiments are rare lepton flavor violating

(LFV) transitions¹ induced by dimension-6 NP operators with SM particles on the external legs. The NP operators are suppressed by the heavy NP scale Λ so that the corresponding LFV branching ratios scale as $\text{BR} \propto 1/\Lambda^4$. Assuming $\mathcal{O}(1)$ Wilson coefficients for the dimension-6 operators the reach on the scale Λ is expected to exceed 10^8 GeV during the ongoing experimental campaign [7].

This can be contrasted with LFV decays into a light axion-like particle (ALP), a , in which case the LFV experiments probe much higher NP scales. The $\mu \rightarrow ea$, $\tau \rightarrow \mu a$ or $\tau \rightarrow ea$ decays are induced by dimension-5 operators so that the LFV branching ratios scale as $\text{BR} \propto 1/f_a^2$, where f_a is the ALP decay constant. The projected bounds on LFV decays then translate to a reach on f_a that, as we will show below, could exceed 10^{10} GeV, assuming $\mathcal{O}(1)$ flavor violating couplings. These scales are among the highest we could probe with ground-based experiments and are well above the present astrophysical constraints induced by the coupling of the light ALP to electrons. This conclusion is the main result of our paper and is presented in figure 1.

More broadly, in this paper we summarize the status of LFV ALP searches. The experimental difficulty is that the $\ell \rightarrow \ell' a$ decays look very similar to the SM decays, resulting in a single visible object plus missing energy (since the ALP is long-lived on detector scales in a large region of the allowed parameter space). As a consequence, the $\ell \rightarrow \ell' a$ decays are not covered by the standard LFV searches and require dedicated experimental strategies/setups. We show that the experimental strategies to improve the coverage of the LFV ALP parameter space depends crucially i) on the ALP mass, and ii) on the chiral structure of the ALP couplings to the SM. For this reason we explore the full range of ALP masses from an effectively massless ALP with $m_a \ll m_e$ to a massive one up to $m_a \lesssim m_\tau$, as well as all the possible chiral structures of ALP couplings to the SM leptons.

At present, the best bounds on f_a from the $\mu \rightarrow ea$ decays give $f_a \gtrsim 10^9$ GeV [8]. The reach can be substantially improved with the next generation of experiments, mainly due to the increased integrated luminosities. For instance, a combination of the experiment by Jodidio et al. from 1986 [9] and the TWIST experiment from 2015 [10] gives the best present bounds on $\text{BR}(\mu \rightarrow ea)$ based on merely $10^7 - 10^8$ stopped muons. This luminosity is at least seven orders of magnitude less than those expected at MEG II and Mu3e.

Taking full advantage of the available datasets will require adjusted experimental approaches. In section 3.2 we put forward a new experimental strategy to improve the bound on $\text{BR}(\mu^+ \rightarrow e^+ a)$ using MEG II. The main idea is to mimic the 1986 experiment by Jodidio et al., utilizing that μ^+ is polarized antiparallel to the beam direction, up to depolarization effects. We study the feasibility of a forward detector configuration in MEG II which we call *MEGII-fwd* (“fwd” for forward), where a calorimeter is placed in the forward direction relative to the muon beam. In this configuration the $\mu^+ \rightarrow e^+ a$ decay can be detected by searching for a positron of maximal energy emitted in the direction opposite to the polarization of μ^+ . For a highly polarized muon beam the SM background from $\mu^+ \rightarrow e^+ \nu \bar{\nu}$ is strongly suppressed in this part of the phase space, while the $\mu^+ \rightarrow e^+ a$

¹These are $\mu \rightarrow e\gamma$, $\mu \rightarrow eee$ and $\mu \rightarrow e$ conversion for the muon and $\tau \rightarrow \ell\gamma$, $\tau \rightarrow \ell\ell\ell$, $\tau \rightarrow \ell\rho$, $\tau \rightarrow \ell\pi$ for the tau, where $\ell = e, \mu$. A more complete list of the LFV transitions and a theory summary can be found in ref. [7].

decay is allowed for an LFV ALP with nonzero right-handed couplings to the SM leptons. We estimate the reach of this setup for two weeks of dedicated run at MEG II and for different configurations of the magnetic field, which will be crucial to control the polarization of the muon beam and the positron yield in the forward direction. We also compare the sensitivity of our proposal to the one that could be obtained at Mu3e by performing an online analysis of the positron spectrum obtained from triggerless data acquisition [11]. Our conclusion is that the two experiments are complementary and, given their timelines, MEGII-fwd has a chance to explore new ALP parameter space several years before Mu3e.

Accessing this new portion of ALP parameter space could unveil connections between ALP Dark Matter (DM) and lepton flavor violation. Indeed, for ALP masses below 1–10 eV, the ALP is a viable DM candidate, if it is produced non-thermally through the misalignment mechanism. The reach on LFV coupling from MEGII-fwd and Mu3e is then complemented by on-going and future experiments sensitive to the couplings of ALP DM to photons and electrons.

While the bulk of our analysis is based on a model-independent parametrization of ALP couplings at low energies, we do explore the implications of the projected experimental sensitivities for a number of different NP scenarios, where the presence of a light ALP can be motivated by the strong CP problem, Dark Matter (DM), the SM flavor puzzle or neutrino masses.

First, we show how LFV decays arise naturally in QCD axion models of the DFSZ-type [12, 13]. In these models the axion solves the strong CP problem which then connects the axion mass, m_a , and decay constant, f_a , giving $m_a \propto 1/f_a$. Furthermore, such a QCD axion can make up for the total amount of DM in the Universe [14–17]. The future reach on $\mu^+ \rightarrow e^+ a$ decays could explore new axion parameter space in the mass range $30 \text{ meV} \gtrsim m_a \gtrsim 0.4 \text{ meV}$, where the upper bound is given by astrophysical constraints. Interestingly, this high mass range requires either a non-standard cosmology [18] or special axion initial conditions [19] in order for the axion to be the DM. The same mass range presents severe experimental challenges for axion DM detection through flavor diagonal couplings. While these could be possibly overcome by ongoing axion-mediated force experiments [20] or new experimental ideas [21, 22], the QCD axion could also well be first observed through its flavor violating couplings. For related studies where flavor violating couplings are (only or also) in the quark sector, see [23–33].

Second, we discuss the reach of LFV decays in the parameter space of the familon. This is the pseudo-Goldstone boson associated with the spontaneous breaking of the lepton flavor symmetry which could explain the charged lepton hierarchies via the Froggatt-Nielsen mechanism [34]. In this setup we show how the strength of LFV decays is correlated with the texture of neutrino masses. A similar construction could also simultaneously address the strong CP problem and the flavor puzzle in the quark sector as proposed in refs. [24, 33] (and was therefore dubbed as “axiflavor”). In the axiflavor setup we show how the flavor violation in the quark sector could be naturally suppressed in a U(2) flavor model similar to the one presented in ref. [35]. This would leave LFV decays as the main experimental signature to hunt for. Similarly, our updated LFV sensitivities will probe the parameter space of axion or relaxion models which try to address the flavor puzzle [36–40].

As a final example of a well-motivated LFV ALP setup we discuss a class of majoron models where the breaking of lepton family number is decoupled from the spontaneous breaking of the lepton number [41]. In this context the future reach of Mu3e and Belle II will explore new parameter space beyond the present astrophysical bounds.

The paper is organized as follows. We start by setting up the notation in section 2, followed by the discussion of $\mu \rightarrow ea$ searches in section 3, with section 3.1 devoted to the review of past searches, while the future proposals, both our proposal for MEGII-fwd as well as the prospects at Mu3e, are presented in section 3.2. Comments on $\mu \rightarrow e$ conversion are given in section 3.3. Section 4 contains a short summary of searches for $\mu \rightarrow e\gamma a$ decays, while in section 5 we summarize the bounds and prospects for tau decays. In section 6.1 we discuss the astrophysical constraints coming from star cooling and SN1987a, comparing these with the reach of LFV decays. In section 7 we discuss several models where LFV violation arises naturally: in section 7.1 the LFV DFSZ axion, in section 7.2 the LFV axiflavor, in section 7.3 the leptonic familon, and in section 7.4 the majoron. Finally, our conclusions are presented in section 8.

2 Notation and summary

The ALP is a (pseudo-) Nambu-Goldstone boson (PNGB) and couples derivatively to the SM fermions. The interaction Lagrangian is thus given by

$$\mathcal{L}_{\text{eff}} = \sum_i \frac{\partial_\mu a}{2f_a} \bar{\ell}_i C_{\ell_i \ell_i}^A \gamma^\mu \gamma_5 \ell_i + \sum_{i \neq j} \frac{\partial_\mu a}{2f_a} \bar{\ell}_i \gamma^\mu (C_{\ell_i \ell_j}^V + C_{\ell_i \ell_j}^A \gamma_5) \ell_j, \quad (2.1)$$

where $C_{\ell_i \ell_i}^A$ is a real diagonal matrix² and $C_{\ell_i \ell_j}^{V,A}$ are hermitian matrices in flavor space, while the summation is over $i, j = 1, 2, 3$, and we ignore possible axion couplings to quarks. For ALP mass we take $m_a < m_\tau$, where m_a could be well below the electron mass. The decay constant f_a is related to the spontaneous breaking scale of the symmetry the ALP is associated with. We do not assume any relations between the couplings in eq. (2.1), and discuss the experimental bounds and prospects separately. For these 6+3 couplings we also introduce the short-hand notation

$$F_{\ell_i \ell_j}^{V,A} = \frac{2f_a}{C_{\ell_i \ell_j}^{V,A}}, \quad F_{\ell_i \ell_j} = \frac{2f_a}{\sqrt{|C_{\ell_i \ell_j}^V|^2 + |C_{\ell_i \ell_j}^A|^2}}. \quad (2.2)$$

When kinematically allowed, the couplings in eq. (2.1) give rise to LFV decays with the (invisible) ALP in the final state.³ The corresponding total decay width is given by

$$\Gamma(\ell_i \rightarrow \ell_j a) = \frac{1}{16\pi} \frac{m_{\ell_i}^3}{F_{\ell_i \ell_j}^2} \left(1 - \frac{m_a^2}{m_{\ell_i}^2}\right)^2, \quad (2.3)$$

²The diagonal vector couplings were set to zero, $C_{\ell_i \ell_i}^V = 0$, via fermion field redefinitions that are anomalous only under $SU(2)_L$, and thus affect only the ALP couplings to electroweak gauge bosons. These couplings can be constrained by their loop-induced contributions to $K \rightarrow \pi a$ [49], which result for an invisible axion in the very weak constraint $f_a/C_{ee}^V \geq 0.7 \text{ TeV}$.

³We note in passing, that while we do not study the phenomenology of the LFV neutrino decays, $\nu_i \rightarrow \nu_j a$, the typical decay time for this process is shorter than the age of the Universe for the ALP decay constants under consideration. This has interesting phenomenological consequences on neutrino cosmology that will be testable in future large scale structure surveys [50, 51].

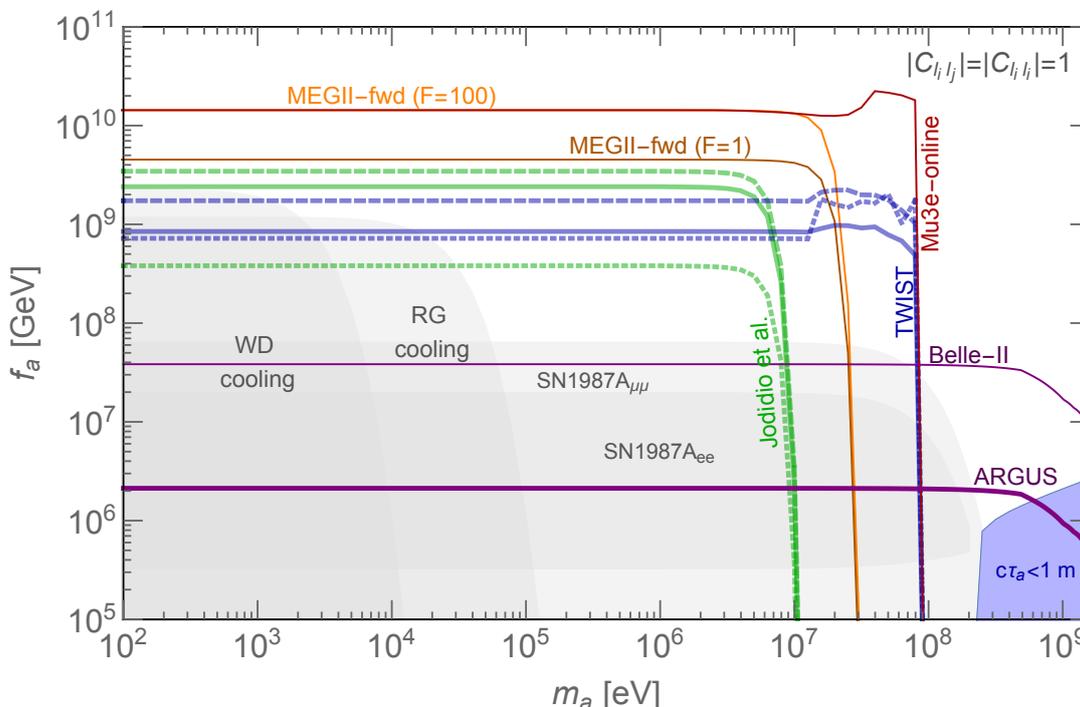


Figure 1. Summary of the present bounds and future projections for an ALP with generic couplings to leptons, i.e., we set $C_{\ell\ell'}^V = C_{\ell\ell'}^A = 1$ for all the couplings in eq. (2.1). For the isotropic case we set $C_{\mu e}^V = 0$ and $C_{\mu e}^A = 1$ (the opposite choice leads to the same results). In the $V \pm A$ case we set $C_{\mu e}^V = \pm C_{\mu e}^A = 1$. The *gray shaded* regions are excluded by the astrophysical bounds from star cooling due to C_{ee}^A and by SN1987A due to C_{ee}^A and $C_{\mu\mu}^A$, see section 6.1. We present these bounds for the isotropic case. The *blue shaded* region corresponds to a prompt/displaced ALP. The *green solid* line is the exclusion due to the bound on $\mu^+ \rightarrow e^+a$ by Jodidio et al., assuming an isotropic ALP [9]. The *green dotted (dashed)* line is our recast of this bound for the $V - A$ ($V + A$) case. The sensitivity in the $V - A$ case is worse since then the signal is suppressed in the forward direction as much as the background. The *blue solid (dotted, dashed)* lines are the bounds from the TWIST experiment on isotropic ($V - A$, $V + A$) ALP [10]. The *dark orange thin solid* line is the MEGII-fwd projection for an isotropic ALP with no magnetic focusing while for the *orange thin solid* line we assumed that focusing increases the luminosity in the forward direction by a factor of 100, cf. section 3.2 for details. The *dark red thin solid* line is the Mu3e projection from [42], for the isotropic ALP. The sensitivity for the other chiral structures is expected to be similar since there is no background suppression in this setup. The *purple solid* line is the bound from the $\tau \rightarrow \mu a$ search by the ARGUS collaboration [43], and does not depend on the chirality of the ALP couplings. The *purple thin* line is the projected reach at Belle-II, see section 5 for details. The bound on $\mu^+ \rightarrow e^+a\gamma$ from Crystal Box is subdominant, see section 4, and is not displayed.

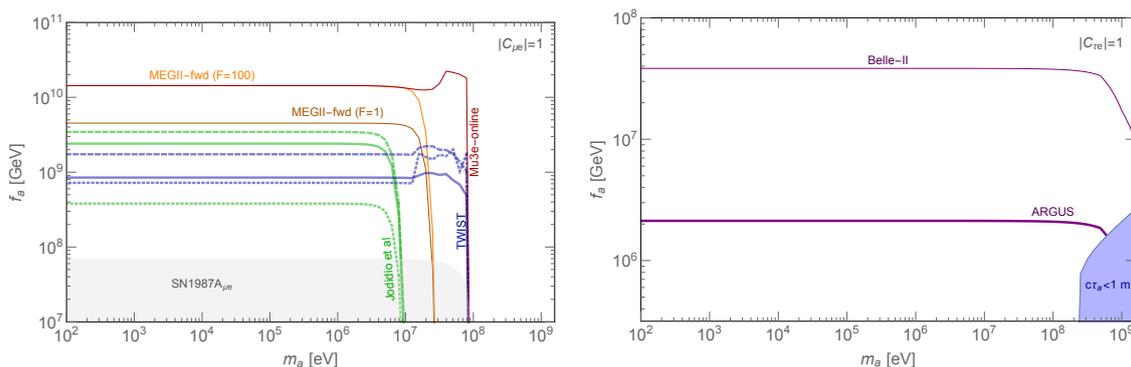


Figure 2. Summary of the bounds and future projections for LFV ALPs. The color coding is the same as in figure 1. *Left:* We set $C_{\mu e} = 1$, while all the other couplings in eq. (2.1) are set to zero. For comparison we also show the bound on $\mu \rightarrow ea$ from SN1987A derived in section 6.1 even though it is subdominant relative to the existing bounds from ground based experiments. *Right:* The only nonzero coupling is $C_{\tau\mu} = 1$. The plot for $C_{\tau e} = 1$ is similar, and is not displayed for brevity.

where for simplicity we neglected the mass of the final-state lepton. The differential decay rate reads (in the same $m_{\ell_j} = 0$ limit)

$$\frac{d\Gamma(\ell_i \rightarrow \ell_j a)}{d\cos\theta} = \frac{m_{\ell_i}^3}{32\pi F_{\ell_i\ell_j}^2} \left(1 - \frac{m_a^2}{m_{\ell_i}^2}\right)^2 \left[1 + 2P_{\ell_i} \cos\theta \frac{\text{Re}(C_{\ell_i\ell_j}^V C_{\ell_i\ell_j}^{A*})}{|C_{\ell_i\ell_j}^V|^2 + |C_{\ell_i\ell_j}^A|^2}\right], \quad (2.4)$$

where θ is the angle between the polarization vector, $\hat{\eta}$, of the decaying lepton ℓ_i and the momentum of the final state lepton ℓ_j , while P_{ℓ_i} is the polarization of the decaying leptons. The convention used for P_{ℓ_i} is such that for the phenomenologically most important case of $\mu^+ \rightarrow e^+ a$ decays we have $P_{\ell_i} = \hat{\eta} \cdot \hat{z}$, where \hat{z} is the beam axis. The μ^+ are predominantly polarized antiparallel to the beam direction, thus $P_{\ell_i} < 0$ and θ is the angle between $-\hat{z}$ and the momentum of the positron, cf. figure 4 (left).

The total width of the ALP can be computed as a function of its mass by summing the different partial decay widths

$$\Gamma_{\text{tot}}(m_a) = \Gamma(a \rightarrow \gamma\gamma) + \sum_{i,j=1,2} \Gamma(a \rightarrow \ell_i\ell_j) + \sum_{i,j=1,2,3} \Gamma(a \rightarrow \nu_i\nu_j). \quad (2.5)$$

Since we restricted the ALP mass to $m_a < m_\tau$, only the decays to photons, neutrinos, electrons, and possibly muons are kinematically open. The corresponding partial decay widths are

$$\Gamma(a \rightarrow \ell_i\ell_j) = \frac{m_a}{8\pi} \left[\left(\frac{m_{\ell_i} - m_{\ell_j}}{F_{ij}^V}\right)^2 z_+ + \left(\frac{m_{\ell_i} + m_{\ell_j}}{F_{ij}^A}\right)^2 z_- \right] \sqrt{z_+ z_-}, \quad (2.6)$$

$$\Gamma(a \rightarrow \gamma\gamma) = \frac{\alpha_{\text{em}}^2 E_{\text{eff}}^2 m_a^3}{64\pi^3 f_a^2}, \quad (2.7)$$

Present best limits				
Process	BR Limit	Decay constant	Bound (GeV)	Experiment
Star cooling	—	F_{ee}^A	4.6×10^9	WDs [44]
	—	$F_{\mu\mu}^A$	1.3×10^8	SN1987A $_{\mu\mu}$ [45, 46]
	4×10^{-3}	$F_{\mu e}$	1.4×10^8	SN1987A $_{\mu e}$ (section 6.1)
$\mu \rightarrow e a$	$2.6 \times 10^{-6*}$	$F_{\mu e} (V \text{ or } A)$	4.8×10^9	Jodidio et al. [9]
$\mu \rightarrow e a$	$2.5 \times 10^{-6*}$	$F_{\mu e} (V + A)$	4.9×10^9	Jodidio et al. [9]
$\mu \rightarrow e a$	$5.8 \times 10^{-5*}$	$F_{\mu e} (V - A)$	1.0×10^9	TWIST [10]
$\mu \rightarrow e a \gamma$	$1.1 \times 10^{-9*}$	$F_{\mu e}$	$5.1 \times 10^{8\#}$	Crystal Box [47]
$\tau \rightarrow e a$	$2.7 \times 10^{-3**}$	$F_{\tau e}$	4.3×10^6	ARGUS [43]
$\tau \rightarrow \mu a$	$4.5 \times 10^{-3**}$	$F_{\tau\mu}$	3.3×10^6	ARGUS [43]
Expected future sensitivities				
Process	BR Sens.	Decay constant	Sens. (GeV)	Experiment
$\mu \rightarrow e a$	$7.2 \times 10^{-7*}$	$F_{\mu e} (V \text{ or } A)$	9.2×10^9	MEGII-fwd*
$\mu \rightarrow e a$	$7.2 \times 10^{-8*}$	$F_{\mu e} (V \text{ or } A)$	2.9×10^{10}	MEGII-fwd**
$\mu \rightarrow e a$	$7.3 \times 10^{-8*}$	$F_{\mu e} (V \text{ or } A)$	2.9×10^{10}	Mu3e [42]
$\tau \rightarrow e a$	$8.3 \times 10^{-6**}$	$F_{\tau e}$	7.7×10^7	Belle II
$\tau \rightarrow \mu a$	$2.0 \times 10^{-5**}$	$F_{\tau\mu}$	4.9×10^7	Belle II

Table 1. The present model independent 95% C.L. best bounds on leptonic ALP couplings $F_{\ell\ell'}^{V,A}$, eq. (2.2), are given in the upper part of the table, with future projections listed in the lower part. The bounds assume m_a below the mass resolution of the experiments considered here (see figure 1 for modifications when m_a is sizable). These follow from 90% C.L. (*) and 95% C.L. (**) bounds on branching ratios in the 2nd column, rescaled using $Z_{95}/Z_{90} = 1.3$ when necessary. The MEGII-fwd projections are obtained for two different sets of assumptions: MEGII-fwd* assumes $\delta x_e = 10^{-2}$ and $\langle P_\mu \rangle - 1 = 10^{-2}$ with no focusing, while MEGII-fwd** in contrast sets the focusing to $F = 100$, roughly what was achieved in the 1986 experiment by Jodidio et al. [9], cf. section 3.2 for details. The Belle II projection for $\tau \rightarrow \mu a$ is rescaled from the Belle MC simulation in ref. [48], while the one for $\tau \rightarrow e a$ is rescaled directly from the ARGUS result [43]. (#) The Crystal Box bound on $F_{\mu e}$ can vary between $(5.1 - 8.3) \times 10^8$ GeV depending on the assumed positron energy loss, cf. eq. (4.9).

where $z_\pm = 1 - (m_{\ell_i} \pm m_{\ell_j})^2/m_a^2$, so that for $m_a \gg m_{\ell_i, \ell_j}$ we have $z_\pm \rightarrow 1$. In the limit $m_i = m_j$ the result reduces to

$$\Gamma(a \rightarrow \ell_i \ell_i) = \frac{m_a}{2\pi} \left(\frac{m_{\ell_i}}{F_{ij}^A} \right)^2 \sqrt{1 - \frac{4m_{\ell_i}^2}{m_a^2}}. \quad (2.8)$$

The ALP decays to neutrinos are often suppressed, so that in the bulk of the paper we set $\Gamma(a \rightarrow \nu_i \nu_j) = 0$ (the majoron is an important exception, see section 7.4). The coupling to photons, E_{eff} , depends on the UV physics as well as on the IR derivative couplings of ALP

to the SM leptons running in the loop,

$$E_{\text{eff}} = E_{\text{UV}} + \sum_f C_f^A B(\tau_f), \quad B(\tau) = \tau \arctan^2 \frac{1}{\sqrt{\tau-1}} - 1. \quad (2.9)$$

Here, $\tau_f = 4m_f^2/m_a^2 - i\epsilon$, and the summation is over the SM leptons, $f = e, \mu, \tau$. Note that the loop function in (2.9) tends to $B(\infty) = 0$ for heavy fermions, and thus the contributions due to the derivative ALP couplings decouple in the heavy fermion limit. The anomaly contribution is encoded in the Wilson coefficient, E_{UV} , and depends on the structure of the UV model. We use the following normalization for the effective ALP-photon Lagrangian⁴

$$\mathcal{L}_{\text{eff}} = E_{\text{UV}} \frac{\alpha_{\text{em}}}{4\pi} \frac{a}{f_a} F \tilde{F}, \quad (2.10)$$

such that for the QCD axion $E_{\text{UV}} = E/2N$, where E and N are the electromagnetic and color anomaly coefficients of the Peccei-Quinn symmetry. For example, in the DFSZ-II model for the QCD axion [12, 13], in which the charged leptons couple to the same Higgs as the up-quarks, one has $E_{\text{UV}} = 1/3$. In section 7 below, we give four explicit examples of LFV ALP models. For these we have $E_{\text{UV}} = \{2/3, 10/9, (10 \div 24), 0\}$ for the LFV QCD axion (section 7.1), the LFV axiflavor (section 7.2), the anarchic LFV familon (section 7.3) and the majoron, (section 7.4), respectively. If the ALP also couples to quarks and gluons, there are additional contributions to the effective photon coupling in eq. (2.9), both from heavy quarks as well as from pions running in the loop (see ref. [52] for complete expressions). From now on we fix $E_{\text{UV}} = 1$, unless specified otherwise, since its precise value does not affect most of the physics discussed in this paper.

As shown in figure 1, we focus in this paper on the region of parameter space where the ALP is long-lived on detector scales. As we will motivate extensively in section 7, we believe that this region is the most appealing from a theoretical perspective. For the discussion of phenomenologically interesting decay channels in the displaced and prompt regions, we refer the reader to refs. [53–55] and to the recent MEG limit on LFV light particles decaying to two photons, $\mu \rightarrow eX$, $X \rightarrow \gamma\gamma$ [56]. ALP masses in the range $2m_e \leq m_a \leq m_\mu - m_e$ are constrained by the search for the $\mu \rightarrow 3e$ decays at the SINDRUM experiment [57]. Indeed, a tiny fraction of ALPs from prompt $\mu \rightarrow ea$ decays leads to $a \rightarrow e^+e^-$ decays within the detector, even for large ALP decay lengths. Assuming 1 cm for the effective decay length measurable in the instrumented volume, the SINDRUM upper bound of $\text{BR}(\mu \rightarrow 3e) \leq 10^{-12}$ is expected to constrain f_a to be above 10^7 GeV for $C_{\mu e} = C_{ee} = 1$, which is much weaker than the other bounds discussed here. The bound on f_a could be substantially improved by the Mu3e experiment [58] given the larger luminosity. However, both the precise SINDRUM exclusion and the Mu3e projection would require to take into account the effect of further experimental cuts on the vertex quality and it is beyond the scope of this paper. A complementary probe of LFV couplings in the region of heavier ALP masses is the muonium-antimuonium oscillation, which would be induced by s - and t -channel exchanges of ALPs with couplings to μe . While the resulting bounds

⁴Our conventions are $\epsilon_{0123} = 1$ and $\tilde{F}_{\mu\nu} = 1/2 \epsilon_{\mu\nu\rho\sigma} F^{\rho\sigma}$.

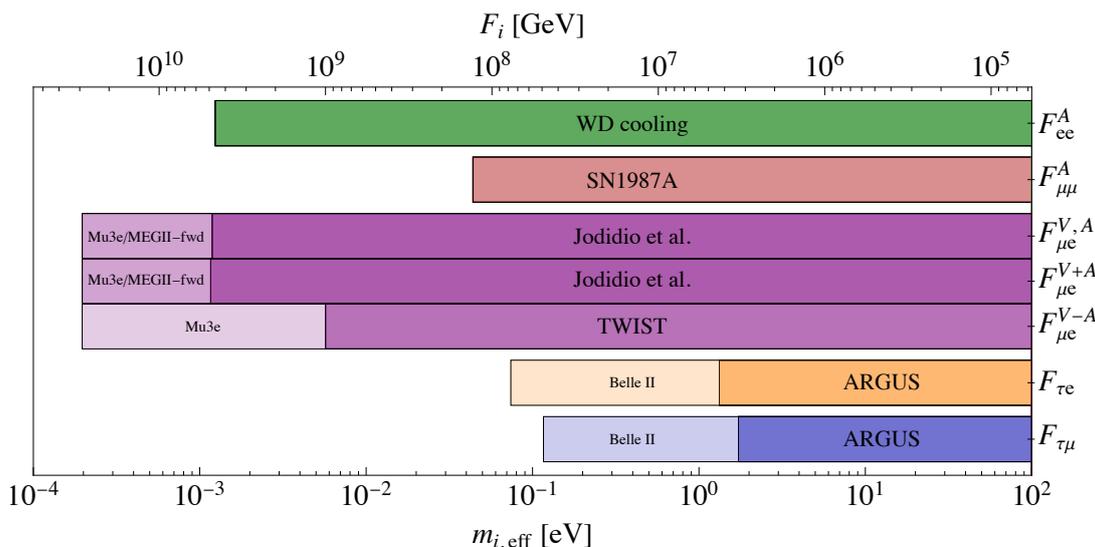


Figure 3. Current bounds and future projections for different couplings F_i of an effectively massless ALP, also reported in table 1. On the lower axis we indicate the corresponding values for the effective axion mass defined by $m_{i,\text{eff}} = 4.7 \text{ eV} \times 10^6 \text{ GeV}/F_i$.

on f_a are several order of magnitude below the ranges shown in figures 1 and 2,⁵ they are stringent enough to effectively rule out the LFV ALP explanations of possible deviations in $(g - 2)_e$ and $(g - 2)_\mu$ [54, 55, 60].

In the numerical analyses throughout the paper all the axion couplings are assumed to be real to simplify the discussion. The interpretations of the present LFV experimental results and future projections in terms of bounds on $F_{\ell_i \ell_j}$ are summarized in figure 1, assuming all the lepton couplings in eq. (2.1) to be $\mathcal{O}(1)$. Figure 2 shows instead the same constraints for the case when only a single LFV coupling is taken to be nonzero. In figures 1 and 2 we also show the typical reach of astrophysical bounds on the ALP decay constant coming from star cooling and SN1987A observations (see section 6.1 for details). In table 1 and figure 3 we summarize the current best bounds and future projections for an effectively massless m_a , i.e. lighter than the typical mass resolution of the experiments considered here. This is the ALP mass range that applies to most of the concrete models discussed in section 7. In the subsequent sections we discuss in detail the observables and the experiments from which these constraints were derived.

⁵For heavier ALPs, $m_a \gtrsim m_\mu$, we can integrate out the ALP to generate the muonium-antimuonium oscillation EFT operators. Translating the results of ref. [59] to our notation gives

$$\frac{1}{1.9 \text{ TeV}} > \left| \frac{1}{F_{\mu e}^A} \pm \frac{1}{F_{\mu e}^V} \right| \left(\frac{m_\mu}{m_a} \right), \quad \frac{1}{3.8 \text{ TeV}} > \left| \frac{1}{(F_{\mu e}^A)^2} - \frac{1}{(F_{\mu e}^V)^2} \right|^{1/2} \left(\frac{m_\mu}{m_a} \right).$$

The constraints for light ALP, $m_a \lesssim m_\mu$, are obtained by taking $m_\mu/m_a \rightarrow 1$ in the two expressions above (see also similar results for heavy meson mixings in the limit of light ALP in ref. [23]). In the future these bounds could be improved for m_a few GeV at Belle II by searching for $e^+e^- \rightarrow ee\mu\mu$ events [60].

3 ALPs in $\mu^+ \rightarrow e^+ + \text{invis. decays}$

We first summarize the status and prospects to search for the two body $\mu^+ \rightarrow e^+ a$ rare decays, where a is invisible, i.e., it decays outside the detector. The challenge of this measurement is to distinguish $\mu^+ \rightarrow e^+ a$ from the background distribution of the SM $\mu^+ \rightarrow e^+ \nu \bar{\nu}$ decay.

The $\mu^+ \rightarrow e^+ a$ decay produces a monochromatic positron line in the muon rest frame at the positron momentum

$$|\vec{p}_e^{\text{line}}| = \sqrt{\left(\frac{m_\mu^2 - m_a^2 + m_e^2}{2m_\mu}\right)^2 - m_e^2}, \quad (3.1)$$

or in terms of the positron energy, $E_e^{\text{line}} \simeq m_\mu/2$ for $m_a \ll m_\mu$. The angular distribution of the positrons depends on the initial muon polarization and the chiral structure of the ALP interactions. We discuss three representative cases that lead to distinct angular distributions:

- the *isotropic ALP* has either $C_{\mu e}^V = 0$ or $C_{\mu e}^A = 0$. The angular distribution of the final state positrons is isotropic in the muon rest frame and independent of the muon polarization, cf. eq. (2.4).
- the *left/right-handed ALP* couples only to the left/right-handed SM fermions, i.e., $C_{\mu e}^V = -C_{\mu e}^A$ for the left-handed and $C_{\mu e}^V = +C_{\mu e}^A$ for the right-handed ALP. The positron angular distribution is $\propto (1 \mp P_\mu \cos \theta)$ for the left(right)-handed ALP.

The SM $\mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ three-body decay proceeds through an off-shell W^+ and produces the so-called Michel spectrum

$$\frac{d^2\Gamma(\mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu)}{dx_e d\cos\theta} \simeq \Gamma_\mu [(3 - 2x_e) - P_\mu(2x_e - 1)\cos\theta]x_e^2, \quad (3.2)$$

with $\Gamma_\mu \simeq m_\mu^5 G_F^2 / (192\pi^3) = 3 \times 10^{-10}$ eV the total muon decay width, and θ the angle between muon polarization vector and the positron momentum in the muon rest frame, see figure 4 (left). The positron energy fraction $x_e = 2E_e/m_\mu$ takes the values $0 \leq x_e \leq 1$ (when neglecting the positron mass). In writing eq. (3.2) the NP scale was taken to be well above the weak scale, as is the case for NP models we are interested in, so that the three-body muon decay is the SM one. For further details on the SM muon properties we refer the reader to the two excellent reviews, refs. [61, 62].

For an unpolarized muon beam, $\langle P_\mu \rangle = 0$, the SM background in eq. (3.2) peaks at $E_e^{\text{line}} = m_\mu/2$ which corresponds to $x_e = 1$. That is, the peak of the SM background for unpolarized muons coincides with the $\mu^+ \rightarrow e^+ a$ positron line for a massless ALP. Luckily, this is not the situation encountered at the intense muon facilities. The muon flux in low energy muon beamlines such as those at TRIUMF or PSI is dominated by the muons produced from pions decaying at rest, at the surface of the production target. These muons are 100% polarized in the direction opposite to the muon momentum, i.e., $\langle P_\mu \rangle = -1$ in the notation of eq. (3.2). As a consequence, muon facilities produce very intense μ^+ fluxes of

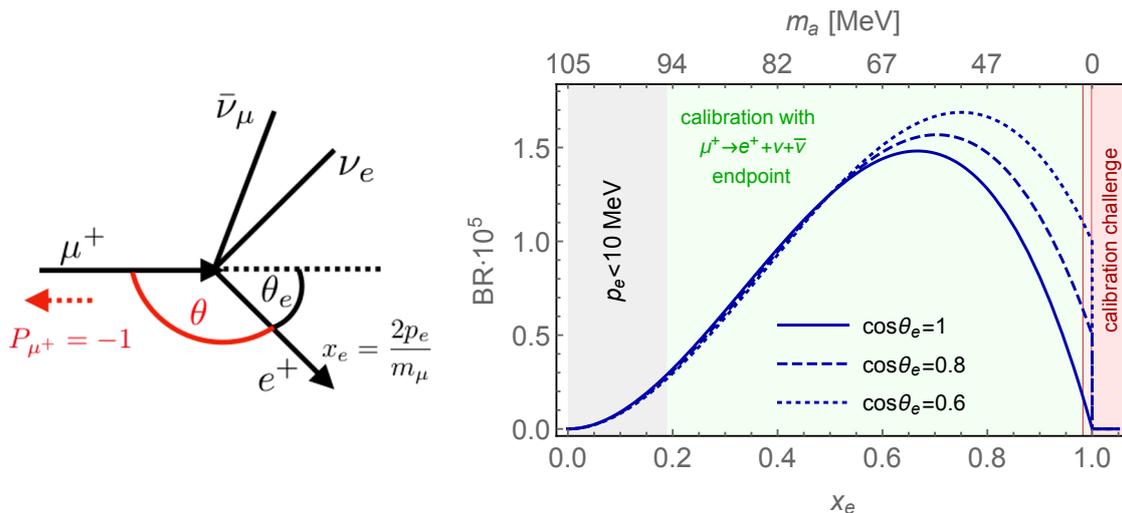


Figure 4. *Left:* Cartoon summarizing the relevant kinematical variables in the SM decay of a polarized μ^+ . *Right:* Summary of the experimental strategies to hunt for $\mu^+ \rightarrow e^+ a$. The three blue lines show the distribution of the positron spectrum in the SM decay $\mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ for fixed angle $\cos \theta_e = 1, 0.8, 0.6$ up to a typical angular resolution of $\delta \theta_e = \pm 5 \times 10^{-3}$. The muon beam is assumed to be 100% polarized $\langle P_\mu \rangle = -1$ and as a consequence the positron distribution drops after having a maximum at x_{\max} (see eq. (3.3)). The value of the branching ratio at the end point depends on the angular position and resolution (see eq. (3.4)). In the gray region the positron momentum is below 10 MeV and typically not recorded, in the green region $\mu^+ \rightarrow e^+ a$ can be searched as a spike on the Michel background. In the red region looking for the signal spike of $\mu^+ \rightarrow e^+ a$ requires to overcome challenges in the calibration of the instrument. The two different red lines correspond to two different momentum resolutions as in eq. (3.6): 0.13% in Jodidio et al. [9] and 1.8% in TWIST [10] (see text for details).

almost 100% polarized muons. The final polarization at the stopping point varies between 80% and 100% depending on the size of the depolarization effects at the production point, during the propagation and at the stopping target.

Polarization of the muon can significantly reduce the SM background. This is illustrated in figure 4 (right), which shows the Michel spectra as functions of x_e for fixed $\langle P_\mu \rangle = -1$ and three representative values of $\cos \bar{\theta}_e = 1, 0.8, 0.6$, where $\theta_e \equiv \pi - \theta$ is the angle between the positron momentum and the muon beamline, see figure 4 (left).⁶ For $0 < \cos \bar{\theta}_e \leq 1$ the non-zero polarization moves the position of the maximum of the Michel spectrum to

$$x_{\max} = \frac{3 - \langle P_\mu \rangle \cos \bar{\theta}_e}{3(1 - \langle P_\mu \rangle \cos \bar{\theta}_e)} \stackrel{\langle P_\mu \rangle = -1}{=} \frac{3 + \cos \bar{\theta}_e}{3(1 + \cos \bar{\theta}_e)}, \quad (3.3)$$

i.e., away from the massless ALP positron line, while for $\cos \bar{\theta}_e \leq 0$ it remains at $x_{\max} = 1$.

⁶Technically, the Michel spectra were integrated over small angular bins, $\cos \bar{\theta}_e \pm \delta \cos \theta_e$, where $\delta \theta_e = 5 \times 10^{-3}$, which is the typical angular resolution of these experiments.

The SM decay rate at the position of the massless ALP positron line, $x_e = 1$, is

$$\frac{d}{dx_e} \text{BR}(\mu \rightarrow e \nu_e \bar{\nu}_\mu) \Big|_{x_e=1, \theta_e=\bar{\theta}_e} = 2\delta \cos \theta_e (1 + \langle P_\mu \rangle \cos \bar{\theta}_e)^{\langle P_\mu \rangle - 1} 2\delta \cos \theta_e (1 - \cos \bar{\theta}_e). \quad (3.4)$$

The SM background for fully polarized muons is exactly zero for $\cos \bar{\theta}_e = 1$, i.e., for positrons emitted in the forward direction relative to the muon beam, up to terms quadratic in the angular resolution. However, in order to have appreciable signal rates, one cannot work in the exact forward limit, $\cos \bar{\theta}_e = 1$, but need to accept events in some range around $\bar{\theta}_e = 0$, which also makes the SM background nonzero. Furthermore, any suppression in the average muon polarization increases the SM background linearly with $\delta P_\mu = \langle P_\mu \rangle + 1$. For the ALP positron line the background suppression is linear in the momentum uncertainty, δx_e , at least within the naive assumption of Gaussian smearing, making this uncertainty dominant compared to the one on the angle, $\delta \theta_e$.

With judicious choice of cuts one can optimize the signal to background ratios. The only two discriminants between the SM background and the ALP signal are the momentum of the final state positron, p_e , and the angle θ_e . The reach of different $\mu^+ \rightarrow e^+ a$ searches can be understood schematically in the cut-and-count scheme, giving

$$\text{BR}(\mu^+ \rightarrow e^+ a) \lesssim Z \sqrt{\frac{a_{\text{bkd}}}{a_{\text{sig}}^2 N_{\mu^+}} + \epsilon_{\text{sys}}^2}, \quad (3.5)$$

where $Z = 1.28(1.64)$ for 90% (95%) C.L. intervals, taking the limit of Gaussian statistics and using limits for one-sided intervals. Here, N_{μ^+} is the number of μ^+ at a given experiment, a_{bkd} (a_{sig}) is the background acceptance (signal efficiency), while ϵ_{sys} encodes the systematic uncertainties. The a_{bkd} depends linearly on the momentum resolution, assuming the background is roughly constant in a given momentum bin. Whenever $\epsilon_{\text{sys}} \gtrsim 1/\sqrt{N_{\mu^+}}$ the reach in the branching ratio saturates independently on the muon luminosity.

Figure 4 (right) summarizes the experimentally accessible regions that can be used for $\mu^+ \rightarrow e^+ a$ searches. Due to the momentum threshold for soft positrons the ALPs with masses above $m_a \gtrsim 95 \text{ MeV}$ are unaccessible. For massive ALPs lighter than 95 MeV a standard bump hunt over the Michel spectrum can be performed, after the instrument is calibrated with the Michel spectrum endpoint. For an almost massless ALP, i.e., for masses below the momentum resolution of the experiment

$$m_a \lesssim \sqrt{\frac{\delta p_e}{p_e}} \cdot m_\mu, \quad (3.6)$$

the positron line gets close to the endpoint of the Michel spectrum. As a consequence, a shift in the endpoint of the SM spectrum due to the momentum resolution or other uncertainties on the experimental setup would result in a spike which would be impossible to distinguish from the signal. Reducing the large systematics on the endpoint of the Michel spectrum then requires alternative ways to calibrate the instrument in this mass region. In figure 4 we show the two mass regions where this is required for the momentum resolutions of the two TRIUMF experiments we describe here.

In section 3.1 we first review the two experimental searches which were performed at TRIUMF: i) the 1986 experiment by Jodidio et al. [9] and ii) the 2015 TWIST experiment [10]. The two searches had very different philosophies. The setup by Jodidio et al. tried to suppress the SM background as much as possible, so that $a_{\text{bkd}} \ll a_{\text{sig}}$ in eq. (3.5). This was achieved by using a highly polarized muon beam and measuring in the forward region with a tight angular cut and excellent momentum resolution. Crucially, magnetic focusing of the positrons improves the signal acceptance above the naive geometric one, and the sensitivity in eq. (3.5) is maximized. The TWIST experiment relies instead on larger integrated luminosity and uses wide angular acceptance. As we show below, combining the results of the two experiments gives complete coverage of different ALP masses and chiral structures of couplings for leptons.

In section 3.2 we discuss the future prospects for $\mu^+ \rightarrow e^+ a$ at PSI: at MEG II and Mu3e [58]. MEG II is set to soon start its physics run, with the primary goal to improve the sensitivity to $\mu^+ \rightarrow e^+ \gamma$ decays. The primary goal of Mu3e is to improve on $\mu^+ \rightarrow e^+ e^- e^+$, with data taking scheduled to start in the relatively near future. Both experiments use the $\pi E5$ beamline which provides roughly $10^8 \mu/\text{sec}$ with muon momenta of 28 MeV, and can also be used for $\mu^+ \rightarrow e^+ a$ searches. The two experimental proposals to hunt for $\mu^+ \rightarrow e^+ a$, MEGII-fwd that is part of this paper, and Mu3e-online, are very different in spirit and could be complementary given that Mu3e will presumably start the physics run after MEG II will have collected the planned muon luminosity.

3.1 Past searches at TRIUMF: Jodidio et al. and TWIST

The 1986 experiment by Jodidio et al. [9] used two datasets, the “spin held” sample of $1.8 \times 10^7 \mu^+$, and “spin precessed” sample of $1.4 \times 10^7 \mu^+$. The two datasets both passed the trigger requirements, but differed in the magnetic field configurations that were used in the experiment. The spin held sample had much higher purity of polarized muons and was used to perform the $\mu^+ \rightarrow e^+ a$ search. The spin precessed sample was used to calibrate the end point of the Michel spectrum in order to reduce the systematic uncertainty. The extremely high purity of polarized muons in the spin held sample was achieved through the use of high purity metal foils as targets, which highly suppressed the muonium production, while the strong magnetic field of 1.1 T parallel to the muon beam line suppressed the muon spin precession. The measured averaged muon polarization at the stopping point was

$$\langle P_\mu \rangle = -0.99863 \pm 0.00088, \tag{3.7}$$

where we assumed the SM values for the muon decaying parameters and combined in quadrature the statistical and systematic uncertainties.

The $\mu^+ \rightarrow e^+ a$ analysis searched for a positron line in the “spin held” data. The cuts

$$\cos \theta_e > 0.975, \quad x_e > 0.97, \tag{3.8}$$

selected the region of phase space in which the SM three-body decay is heavily suppressed for polarized muons, see eq. (3.4). The positrons emitted in the beam direction, i.e., at $\cos \theta_e \simeq 1$, were measured downstream in the spectrometer after they bent by more than 90° by a magnetic field.

First, we check how well the simple cut-and-count scheme reproduces the bound derived by the experimental collaboration for the massless isotropic ALP,

$$C_{\mu e}^V = 0 \text{ or } C_{\mu e}^A = 0 : \quad \text{BR}(\mu \rightarrow e a) < 2.6 \times 10^{-6} \quad (90\% \text{ CL}) \quad [9]. \quad (3.9)$$

This will prove useful when deducing the projected sensitivities in the next section. Applying the cuts (3.8) on the Michel spectrum, eq. (3.2), gives $a_{\text{geo}} = 5.3 \times 10^{-5}$ for the geometric acceptance of the background. Moreover, for a massless ALP, the number of background events in a $\pm 2\sigma$ band around $x_e = 1$ can be estimated by integrating the distribution provided in ref. [9], accounting for the quoted momentum resolution

$$\frac{\delta p_e}{p_e} = 0.13\%. \quad (3.10)$$

This leads to a background efficiency of $\epsilon_{\text{bkd}} = 1.7 \times 10^{-2}$, i.e., the fraction of background events satisfying (3.8) that are in the signal region $x_e \in [1 - 2\delta x_e, 1 + 2\delta x_e]$. The background acceptance in the signal region of the initial spin held sample is thus $a_{\text{bkd}} = a_{\text{geo}} \cdot \epsilon_{\text{bkd}} = 8.7 \times 10^{-7}$. Applying instead the 2σ band cut around $x_e \simeq 1$ directly on the Michel spectrum, eq. (3.2), gives a similar estimate for the background acceptance, $a_{\text{bkd}} \simeq 10^{-6}$. This shows that using the simple analytical Michel spectrum is good enough for the purposes of our estimates in the next section.

In order to obtain the correct limits we still need to take into account the effect of magnetic focusing. After magnetic focusing and the cuts in eq. (3.8) the total number of observed events in the spin held sample is $\simeq 7.4 \times 10^4$ (obtained by integrating the distribution published in ref. [9]). Assuming that these events are entirely due to the SM, and neglecting the focusing, would require $N_{\mu^+} \simeq 1.4 \times 10^9$ in the initial spin held sample, while in reality $N_{\mu^+} = 1.8 \times 10^7$. The magnetic focusing therefore leads to effectively larger geometric acceptance,⁷

$$F \equiv \frac{a_{\text{geo}+\text{focus}}}{a_{\text{geo}}} = 77.8. \quad (3.11)$$

The efficiency for the ALP signal in this experimental setup depends very much on the helicity structure of the ALP couplings to the SM current and on the ALP mass. For a massless isotropic ALP, i.e., $m_a \ll m_e$ and $C^V = 0$ or $C^A = 0$, we find $a_{\text{geo}}^{\text{ISO}} = 1.25 \times 10^{-2}$ after the angular cut in eq. (3.8). Assuming that the focusing lens act similarly on background and signal we find $a_{\text{geo}+\text{focus}}^{\text{ISO}} = 0.97$. The experimental analysis also assigns $\epsilon_{\text{sys}} = 0.9 \times 10^{-6}$ systematic uncertainty on the branching ratio for an isotropic ALP. Given the numbers above, eq. (3.5) then gives our recast bound of 2.8×10^{-6} , which agrees quite well with the result of the fit in ref. [9]. The agreement makes us confident that we understand the main features of the experimental setup. The above exercise also highlights that the systematic uncertainty in the experiment by Jodidio et al. was smaller than the statistical one, and thus the bound on $\text{BR}(\mu \rightarrow e a)$ would have benefitted from a factor of 10 bigger muon luminosity before hitting the bottleneck of systematics.

⁷We thank Angela Papa and Giovanni Signorelli for illuminating discussions about this point and refer to [63, 64] for a detailed derivation of the focusing power of a solenoid lens.

Next, we recast the result by Jodidio et al. for different chiral structures of ALP couplings and for higher ALP masses. We assume that there are no changes in systematic uncertainties and the solenoid focusing effect. The ratios of signal acceptances are then obtained by simply applying the angular cut in eq. (3.8) to the angular distribution predicted by the modified ALP coupling and mass, and compare it to the baseline massless isotropic ALP. We obtain $a_{\text{geo}}^{\text{RH}}/a_{\text{geo}}^{\text{ISO}} \simeq 2$ and $a_{\text{geo}}^{\text{LH}}/a_{\text{geo}}^{\text{ISO}} \simeq 0.012$ for the signal acceptances of right-handed (RH) ALP and left-handed (LH) ALP relative to the isotropic case, respectively. Rescaling the bound in eq. (3.9) gives for the massless right-handed ALP

$$C_{\mu e}^V = C_{\mu e}^A : \quad \text{BR}(\mu \rightarrow e a) < 2.5 \times 10^{-6} \quad (90\% \text{ CL}). \quad (3.12)$$

The improvement in the bound is marginal compared to the isotropic case because the magnetic focusing effect is already giving a signal acceptance very close to one. If focusing were a tunable parameter, the setup by Jodidio et al. would have benefitted from slight defocusing which would have decreased the SM background but kept the signal almost unchanged, giving the best bound on the $V + A$ ALP.

For massive ALP the measurement by Jodidio et al. still translates into a bound since the tail of the signal distribution leaks into the experimental signal region due to the finite experimental resolution on the positron momentum, eq. (3.10). For $p_e^{\text{line}} \simeq m_\mu/2$ the momentum spread is up to 70 keV. We recast the Jodidio et al. bound to $m_a \gtrsim m_e$ by taking the signal to be a Gaussian in p_e with the mean given by eq. (3.1) and the width set by the momentum resolution in eq. (3.10). The overlap between the Gaussian smearing of the signal and the flat bin centered around $p_e^{\text{line}} = m_\mu/2$ with a width of 140 keV sets the efficiency of the signal. The result of this extrapolation is shown in figure 1 and as expected the search strategy loses sensitivity very slowly for $m_a \gtrsim m_e$. The efficiency drops significantly only when $m_a \simeq m_\mu/10$, in which case the ALP mass starts to significantly shift the position of the positron line in eq. (3.1).

For pure left-handed ALP the systematic uncertainties related to the determination of the background endpoint are expected to grow significantly. The reason is that in this case the signal is suppressed more in the spin held dataset than in the spin precessed one. However, in the analysis [9] the latter was assumed to not be affected by the NP signal, in order to calibrate the instrument. Assuming that this is indeed still the case gives the green dashed line in figure 1, which should be viewed as only roughly indicative of what the correct bound for the left-handed ALP is.

The 2015 TWIST experiment [10] collected 5.8×10^8 muons after the selection cuts were applied. The experimental collaboration studied the $\mu^+ \rightarrow e^+ a$ decay, varying both the mass of the ALP and the chirality of its couplings to the SM. Their results are summarized with the blue lines in figure 1.

The experimental concept of TWIST is fundamentally different from the previously discussed experiment by Jodidio et al. TWIST detected positrons using a spectrometer with an approximate cylindrical symmetry surrounding the muon beam line, and momentum resolution at $x_e \simeq 1$ given by,

$$\frac{\delta p_e}{p_e} \simeq \frac{1.1\%}{|\sin \theta_e|}, \quad (3.13)$$

which correspond to $\delta p_e = 572 \text{ keV} / \sin \theta_e$ at $p_e = 52 \text{ MeV}$. Since the momentum resolution deteriorates in the forward direction (where $\sin \theta_e \rightarrow 0$), the experimental strategy of Jodidio et al. cannot be implemented, even though the TWIST muon beam is highly polarized. The momentum resolution in eq. (3.10) translates into an upper bound below which the ALP can be considered effectively massless

$$m_a \lesssim \sqrt{2m_\mu \cdot 1\text{MeV}} \simeq 14 \text{ MeV}. \quad (3.14)$$

The 1 MeV was obtained by setting $\cos \theta_e = 0.8$ in eq. (3.10), which is the positron angle for which TWIST had the widest momentum acceptance.

As shown in figure 1, for a massive ALP the TWIST search covers a region that was left unexplored by the 1986 Jodidio et al. result, and extends the coverage up to an ALP mass of 86.6 MeV. For higher masses the positron becomes too soft to be efficiently triggered on at TWIST. In this region older searches by Derenzo et al. [65] and Bilger et al. [66] and the recent PIENU result [67] complement the high mass coverage. We refer to the PIENU paper [67] for a summary plot focused on the high mass region.

TWIST becomes less sensitive for a massless ALP due to the systematic uncertainties related to the calibration of the $x_e = 1$ endpoint, which limited the sensitivity. Unlike Jodidio et al., TWIST collaboration had to rely on Monte Carlo modeling to calibrate the endpoint. The quoted systematic on the momentum edge represented an irreducible bottleneck to improving sensitivity of the massless ALP search. The 90% C.L. upper bounds for massless ALP by TWIST, extracted using the Feldman-Cousins method [68], are

$$C_{\mu e}^V = 0 \text{ or } C_{\mu e}^A = 0 : \quad \text{BR}(\mu \rightarrow e a) < 2.1 \times 10^{-5} \quad (90\% \text{ CL}), \quad (3.15)$$

$$C_{\mu e}^V = C_{\mu e}^A : \quad \text{BR}(\mu \rightarrow e a) < 1 \times 10^{-5} \quad (90\% \text{ CL}), \quad (3.16)$$

$$C_{\mu e}^V = -C_{\mu e}^A : \quad \text{BR}(\mu \rightarrow e a) < 5.8 \times 10^{-5} \quad (90\% \text{ CL}). \quad (3.17)$$

Only the last bound, on the left-handed ALP, is stronger than our recast of the previous bound from Jodidio et al. experiment, which, as discussed before, was less sensitive to left-handed ALPs. The gain of the TWIST experiment is due to the much larger muon luminosity compared to the one available to Jodidio et al., which compensates for the worse momentum resolution at TWIST.

3.2 New searches at PSI: MEGII-fwd and Mu3e-online

A proposal for MEGII-fwd. The MEG II experiment is expected to start its physics runs soon after completion of the engineering run that started in 2019. The goal of our proposal is to achieve at MEG II a configuration similar to the one used in the experiment by Jodidio et al. [9], so that the $\mu^+ \rightarrow e^+ a$ decays of polarized muons are detected in the forward region where the SM background is suppressed, see figure 4 and discussion in section 3.1. Such a set-up requires:

- A forward detector to collect energetic forward positrons. The final sensitivity to $\mu^+ \rightarrow e^+ a$ decays will depend on the energy resolution of this detector. We assume that a Lyso calorimeter (Lyso-ECAL) with a 10 cm diameter and high enough threshold for positron energies can be installed in the forward direction roughly 1.5 meters

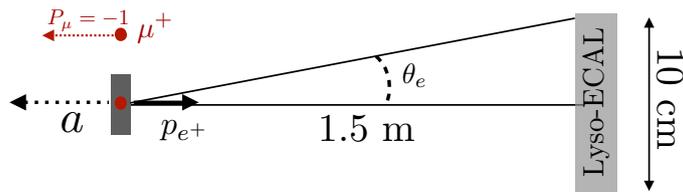


Figure 5. The proposed MEGII-fwd set-up. A Lyso-ECAL detector of 10 cm in diameter is placed along the muon beam line 1.5 m downstream from the stopping point. The muon polarization P_μ is in the opposite direction than the detected positron.

downstream from the muon stopping target, see figure 5. The resulting angular coverage is $\theta \simeq 1.91^\circ$. We vary the positron momentum resolution at $p_{e^+} = m_\mu/2$ in the range⁸

$$\delta x_e \equiv \delta p_{e^+}/m_e = 10^{-3} - 10^{-1}. \quad (3.18)$$

At a similar position the MEG II collaboration planned to put a Radiative Decay Counter (RDC), with the aim to reduce the accidental background for the search of $\mu^+ \rightarrow e^+ \gamma$. However, RDC is designed to detect low-momenta positrons and would not be useful for the $\mu^+ \rightarrow e^+ a$ search.

- A new configuration of the MEG II magnetic field in order to suppress depolarization effects and keep the μ^+ antiparallel to the outgoing positron. The two main sources of depolarization of the muon beam are the so called “halo muons”, emitted from pions decay in flight, and the angular divergence of the beam. How close the polarization can be kept to the maximal one, $\langle P_\mu \rangle = -1$ is crucial here, as this controls the suppression of the SM background, see (3.4), which directly affects the sensitivity to $\mu^+ \rightarrow e^+ a$. In what follows, we vary the depolarization in the range,

$$\Delta P_\mu \equiv \langle P_\mu \rangle + 1 = 10^{-3} - 10^{-1}. \quad (3.19)$$

In the predecessor experiment, the MEG experiment, the muon polarization at the stopping target was $\langle P_\mu \rangle = -0.86 \pm 0.06$ [70], which corresponds to the upper limit of the above range. The lower limit of the range assumes that the depolarization strategy similar to the one used in the 1998 experiment by Jodidio et al., ref. [9], can be put in place. Collecting in addition a less pure sample of polarized muons would help in calibrating the endpoint of the Michel spectrum, $x_e = 1$, exactly as it was done in ref. [9].

- A focusing lens to increase the positron luminosity in the forward direction. In the experiment by Jodidio et al., a solenoid lens was used to maximize the signal acceptance at the price of a higher SM background, cf. section 3.1. Focusing enlarges

⁸We are aware that the use of such a calorimeter for calibration purposes was discussed inside the MEG II collaboration. A momentum resolution $\delta p_{e^+}/p_{e^+} = 10^{-2}$ at $p_{e^+} = m_\mu/2$ is realistic for this ECAL [69]. We thank Angela Papa for private communications regarding this.

the effective size of the forward detector. The angular coverage of the Lyso-ECAL is very small so that without focusing the geometric acceptance of the signal is only,

$$a_{\text{geo}}^{\text{ISO}} = 5.6 \times 10^{-4} \quad , \quad a_{\text{geo}}^{\text{V+A}} = 1.2 \times 10^{-3} \quad , \quad (3.20)$$

for isotropic and $V + A$ ALP, respectively. The reach on the branching ratio scales with the focusing factor F as $\text{BR}(\mu^+ \rightarrow e^+a) \sim 1/\sqrt{F}$ as long as $F \times a_{\text{geo}} < 1$ and the systematics uncertainties are subdominant. In the projections we leave F as a free parameter, noting that $F \sim \mathcal{O}(10^2)$ was reached in the 1986 experiment by Jodidio et al., cf. eq. (3.11).

- Sufficient time devoted to the physics run in the $\mu^+ \rightarrow e^+a$ search configuration. As a reference point we will use

$$N_\mu = 1.2 \times 10^{14} \mu^+ \quad , \quad (3.21)$$

which corresponds to a 2 week run at an instantaneous luminosity of $10^8 \mu^+/\text{sec}$.⁹ This is seven orders of magnitude greater than the dataset collected by Jodidio et al. Crucially, the $\mu \rightarrow ea$ run could be performed at the very end of the $\mu \rightarrow e\gamma$ data taking at MEGII. This would allow to work on the necessary modifications and extend the physics purpose of MEG II before Mu3e starts.

In the left panel of figure 6 we show how the reach on the branching ratio $\text{BR}(\mu^+ \rightarrow e^+a)$ for a massless ALP depends on the average polarization and the angular resolution. Interestingly, from the shape of the contours one sees that augmenting the polarization purity of the muon beam should go together with an increase of the momentum resolution in order to lead to a better experimental reach. The orange star in the plot is one of our benchmark configurations, where we assume no focusing in the forward direction and set $\delta x_e = \langle P_\mu \rangle + 1 = 10^{-2}$. Already in these suboptimal conditions, MEGII-fwd could improve on the present bound from Jodidio et al.

In the right panel we set $\delta x_e = \langle P_\mu \rangle + 1$ and investigate the importance of the focusing. To be conservative we take as the benchmark $F = 100$, which corresponds to an effective diameter of the forward calorimeter roughly 10 times bigger than its actual size. This is of the same order of magnitude as the focusing achieved in ref. [9]. However, a larger focusing would always be beneficial until $F \times a_{\text{geo}}$ would be of order unity.

In table 1 and figure 1 we show the estimated reach of MEGII-fwd for the two benchmark configurations described above. The reach for higher m_a , shown in figure 1, is computed accounting for the amount of signal which will overlap with the massless hypothesis given the experimental resolution. Due to worse momentum resolution the drop in the f_a reach at high masses is slower at MEGII-fwd compared to the Jodidio et al. experiment.

⁹MEG II plans to run with the reduced instantaneous luminosity of $3 \times 10^7 \mu^+/\text{sec}$ in order to decrease accidental coincidences, which is the dominant background in the $\mu^+ \rightarrow e^+\gamma$ measurement. For the $\mu^+ \rightarrow e^+a$ search one should attempt to use the full instantaneous luminosity of $10^8 \mu^+/\text{sec}$ available at the $\pi E5$ beamline at PSI.

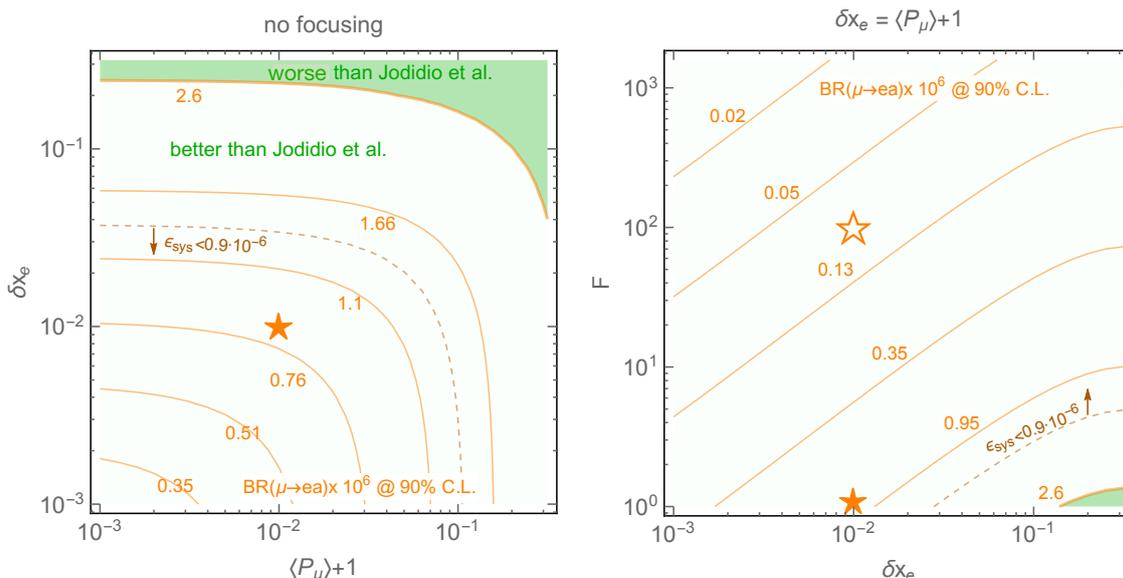


Figure 6. *Left:* The orange contours show the 90% C.L. expected sensitivity on $\text{BR}(\mu^+ \rightarrow e^+a)$ at MEGII-fwd for a massless isotropic ALP (either $C_{\mu e}^V \neq 0$ or $C_{\mu e}^A \neq 0$) as a function of the momentum resolution δx_e , and the deviation of average polarization $\langle P_\mu \rangle$ from -1 , assuming there is no magnetic focusing. *Right:* The expected sensitivity at MEGII-fwd (orange contours) as a function of the momentum resolution δx_e and focusing F , setting muon polarization to $\langle P_\mu \rangle = -1 + \delta x_e$. The angular resolution is assumed to be subdominant and systematic uncertainties below the statistical one. The darker green region are excluded by the 1986 Jodidio et al. experiment [9], cf. table 1. The dashed dark orange line indicates where our future projections assume that systematic uncertainties can be lowered compared to the ones in Jodidio et al. $\epsilon_{\text{sys}} = 0.9 \times 10^{-6}$.

Mu3e-online. The primary goal of the planned Mu3e experiment at PSI is to search for $\mu^+ \rightarrow e^+e^-e^+$ with the unprecedented sensitivity of 10^{-16} . The key feature of Mu3e is that it will operate without a hardware trigger. The full detector read-out will be streamed to the filter farm at 80 GB/sec, where the $\mu^+ \rightarrow e^+e^-e^+$ events will be identified and eventually stored on disk.

Recently, ref. [11] performed a preliminary study of the Mu3e sensitivity to $\text{BR}(\mu^+ \rightarrow e^+a)$, based on a more detailed simulation of the $\mu^+ \rightarrow e^+a$ channel at the phase I of the Mu3e experiment [42]. The dark thin red line in figure 1 shows the 95% C.L. limit on f_a for isotropic LFV ALP that Mu3e is projected to achieve with a physics run of 300 days. The $\sim 10^4$ improvement in f_a reach over the TWIST experiment is mostly driven by the seven orders of magnitude larger dataset, which, however, does not come without challenges.

For $\mu^+ \rightarrow e^+a$ search the Mu3e will be able to reconstruct online all the single positron events corresponding to “short tracks”, i.e., events with only four hits on the four detector layers. The online reconstruction of “short tracks” in the filter farm has been shown to reduce by a factor of 100 the data rate [71], making it possible to process all the short tracks events and store positron three-momenta, \vec{p}_e . A search for $\text{BR}(\mu^+ \rightarrow e^+a)$ positron line can be done as a bump hunt on the smooth SM $|\vec{p}_e|$ distribution, assuming that every positron event corresponds to a single track.

The tracks in Mu3e will gyrate around the magnetic field of roughly 1 T. The typical e^+ gyroradius will be much larger than the radius of the Mu3e instrumented region — a cylinder with radius of around 6 cm. Encountering the detector material the positron will lose momentum and eventually stop. The positron will typically stop after half a turn, i.e., after having encountered at least four detector layers: 2 central layers + 2 external layers. This justifies the assumption of one positron per one track. Positrons emitted perpendicularly to the muon beam will instead perform many turns in the central layers without being stopped in the detector material. Enforcing an angular cut will minimize the impact of these re-curling tracks. As a result, the angular acceptance of the Mu3e analysis is expected to be reduced to the region $\theta_e < \pi/2 - 0.1$.

The momentum resolution for short tracks will be in the $\delta p_{\text{short}} = 0.5 - 3$ MeV range, roughly comparable to TWIST. In principle, a momentum resolution down to $\delta p_{\text{long}} = 0.1 - 0.45$ MeV could be achieved, if an experimental upgrade allowed to also process online the “long tracks”, characterized by 6 or 8 hits in the detector layers. The Mu3e sensitivity on $\text{BR}(\mu^+ \rightarrow e^+ a)$ extracted from the long track analysis would then improve by a factor of $\sqrt{\delta p_{\text{long}}/\delta p_{\text{short}}} \simeq 0.4$ [42].

Finally, we comment on the challenges of calibrating the instrument. As already mentioned, the current concept foresees to use the endpoint of the positron momentum spectrum to calibrate the online reconstructed tracks. This method will not allow to efficiently search for $\mu^+ \rightarrow e^+ a$ decays when $m_a \lesssim 10 - 25$ MeV, depending on the precise positron momentum resolution. Given the large amount of single positron data collected at Mu3e, ref. [42] showed that, within a given mass assumption for the signal, one could in principle use the same momentum spectrum to simultaneously calibrate the apparatus and to perform the peak search. In the calibration fit one would remove from the calibration dataset the signal region, which is defined as a $2\delta p_{\text{short}}$ band below and above the expected momentum of the positron corresponding to the ALP mass. Given that for $m_a \lesssim 25$ MeV the signal region includes the Michel edge the determination of the scaling parameter $x_e = 1$ deteriorates, resulting in a limited sensitivity at low masses. Including this effect the expected sensitivities for a massless isotropic ALP are

$$C_{\mu e}^V = 0 \text{ or } C_{\mu e}^A = 0 : \quad \text{BR}(\mu \rightarrow e a) < 7.3 \times 10^{-8} \quad (90\% \text{ CL}). \quad (3.22)$$

Notice that the deterioration computed in ref. [42] includes only the broadening of the x_e statistical distribution estimated from a toy Monte Carlo sample of $10^9 \mu^+$ (the root mean square error grows from 7×10^{-6} for $m_a = 60$ MeV to 3.8×10^{-5} for $m_a = 10$ MeV). A similar effect is expected for both the left-handed and the right-handed ALPs so that the reach shown in figure 1 is a realistic estimate of the Mu3e reach for all chiral structures. It is possible, however, that alternative calibration strategies such as the one proposed in ref. [72] could improve the Mu3e reach at low masses.

3.3 The potential at $\mu^- \rightarrow e^-$ conversion experiments

The $\mu \rightarrow e$ conversion in nuclei experiments, COMET at J-PARC [4] and Mu2e at FNAL [5], are designed around very large muon fluxes, with over $10^9 \mu^-/s$ and $10^{10} \mu^-/s$

at the two respective experiments. One may hope that these could also be used for $\mu \rightarrow ea$ searches. However, in order to be able to deal with the vast data-streams the two experiments will not measure the full Michel spectrum but rather focus on the endpoint of the e^- energy, which for the $\mu^- + (A, Z) \rightarrow e^- + (A, Z)$ process is at $E_e^{\text{end}} = m_\mu - E_b - E_{\text{rec}}$, with E_b the binding energy of the muonic atom and E_{rec} the nuclear recoil energy. In Aluminum $E_e^{\text{end}} \simeq 105$ MeV, while in gold $E_e^{\text{end}} \simeq 95$ MeV.

The conversion with the emission of the ALP, $\mu^- + (A, Z) \rightarrow e^- + a + (A, Z)$, peaks instead at the electron energy of around $E_e \simeq m_\mu/2$, but with the tails of the electron energy distribution that go all the way up to E_e^{end} . Nominally, searches for the $\mu \rightarrow ea$ process at COMET and Mu2e would also rely on these tails of the distribution, which poses a challenge for obtaining a competitive reach. For instance, for $E_e > 100$ MeV only a small fraction of $\mu \rightarrow ea$ decays, about 2×10^{-10} , would be inside the signal region for $\mu \rightarrow e$ conversion on Al [73] (see also [74]). Even with 10^{18} muons a competitive search would thus likely require relaxing the lower bound on E_e and developing techniques to distinguish between the smooth shapes of the signal, $\mu \rightarrow ea$, and SM background, $\mu \rightarrow e\nu\bar{\nu}$, decays in orbit.

4 ALPs in $\mu \rightarrow e + \gamma + \text{invis. decays}$

The $\mu^+ \rightarrow e^+\gamma a$ decay offers a complementary probe of the LFV ALP which is less dependent on the chiral structure of the ALP couplings than the experimental searches for $\mu \rightarrow ea$. In section 4.1 we first discuss the searches performed at the Crystal Box experiment with a total number of 8.15×10^{11} stopped muons [47, 75]. In section 4.2 we discuss possible improvements in the reach were a similar search to be implemented at MEG II. A more detailed analysis of a dedicated trigger at MEG II for this channel is left for future work.

4.1 Past searches at Crystal Box

In order to lower the trigger rate the Crystal Box required at the trigger level a hard positron and a photon of similar energy [47, 75]. The search for the three-body $\mu^+ \rightarrow e^+\gamma a$ decay is then a search for a bump in the missing mass distribution in the collected data. The signal would be centered at $m_{\text{miss}} = m_a^2$ and spread by the photon and positron energy resolutions and the resolution on the angle between the two. The SM background has two main components: the four-body $\mu^+ \rightarrow e^+\gamma\nu\bar{\nu}$ decays, and the combinatorics background due to coincident $\mu^+ \rightarrow e^+\nu\bar{\nu}$ and $\mu^+ \rightarrow e^+\gamma\nu\bar{\nu}$ events. For the latter, a sufficiently hard positron from the $\mu^+ \rightarrow e^+\nu\bar{\nu}$ decay is detected within a 1.5 ns time window together with a hard photon from the $\mu^+ \rightarrow e^+\gamma\nu\bar{\nu}$ decay, while the soft positron is left undetected. The background and signal shapes at Crystal Box are shown in figure 7 (left).

The rate of the three-body decay $\Gamma(\mu \rightarrow ea\gamma)$ for an ALP of mass m_a is given by (in the limit $m_e \ll m_\mu$)

$$\Gamma(\mu \rightarrow ea\gamma) \approx \frac{\alpha_{\text{em}}}{32\pi^2} \frac{m_\mu^3}{F_{\mu e}^2} \mathcal{I}(x_{\text{min}}, y_{\text{min}}, \eta_a), \quad (4.1)$$

with the phase space integral given by

$$\mathcal{I}(x_{\min}, y_{\min}, \eta_a) = \int_{y_{\min}}^{1-\eta_a} dy \int_{\max(x_{\min}, 1-y-\eta_a)}^{\frac{1-y-\eta_a}{1-y}} dx \frac{y(1-x^2-\eta_a^2) - 2(1-\eta_a)(1-x-\eta_a)}{y^2(1-x-y-\eta_a)}, \quad (4.2)$$

where

$$x = 2E_e/m_\mu, \quad y = 2E_\gamma/m_\mu, \quad \eta_a = m_a^2/m_\mu^2. \quad (4.3)$$

These kinematic variables are related to the angle $\theta_{e\gamma}$ between the electron and photon momenta,

$$\cos \theta_{e\gamma} = 1 + \frac{2(1-x-y-\eta_a)}{xy}. \quad (4.4)$$

The branching ratio for the three-body decay, $\text{BR}(\mu \rightarrow e a \gamma)$, is related to the branching ratio for the 2-body decay, $\text{BR}(\mu \rightarrow e a)$,

$$\text{BR}(\mu \rightarrow e a \gamma) \approx \frac{\alpha_{\text{em}}}{2\pi(1-\eta_a)^2} \mathcal{I}(x_{\min}, y_{\min}, \eta_a) \text{BR}(\mu \rightarrow e a), \quad (4.5)$$

where in the expressions for both $\text{BR}(\mu \rightarrow e a \gamma)$ and $\text{BR}(\mu \rightarrow e a)$ the mass of electron should be neglected.

The infrared and collinear divergences (we are working in the limit $m_e \rightarrow 0$) are regulated by the experimental cuts on photon and positron energies x_{\min}, y_{\min} . The experimental cuts in the Cristal Box search were [47]

$$E_e > 38 - 43 \text{ MeV}, E_\gamma > 38 \text{ MeV} \quad \Rightarrow \quad x_{\min} = 0.72 - 0.81, \quad y_{\min} = 0.72. \quad (4.6)$$

$$\theta_{e\gamma} > 140^\circ \quad \Rightarrow \quad \cos \theta_{e\gamma} < -0.77. \quad (4.7)$$

Imposing the $\theta_{e\gamma}$ cut in the phase space integral, eq. (4.2), reduces $\mathcal{I}(x_{\min}, y_{\min}, \eta_a)$ by $\sim 10\%$ for a massless ALP, and has negligible effect for a massive ALP. The positron energy cut, $E_e > 38 \text{ MeV}$, reported by the Crystal Box collaboration, refers to the positron energy measured by the scintillation crystals. When translating this cut to x_{\min} in the phase space integral, eq. (4.2), one needs to account for the positron's energy loss during the propagation to the detector. The Crystal Box collaboration indicated that the positron could loose up to 5 MeV before reaching the detector. In our analysis we therefore vary the cut on the positron energy in the range from 38 to 43 MeV.

In the special case of vanishing axion mass, $\eta_a \rightarrow 0$, the above expressions, eq. (4.1)–(4.5), reduce to the ones obtained in ref. [76]. Notice that the expression used by the experimental collaboration in ref. [75] appears to have a typographical error (the roles of x and y in the phase space integral in eq. (4.2) were exchanged). However, this does not affect the analysis in ref. [47] because of symmetric cuts on photon and positron energies (up to the positron energy losses), cf. eq. (4.6).

The 90% C.L. bound on the $\mu \rightarrow e a \gamma$ branching ratio for $m_a \simeq 0$, obtained by the Crystal Box experiment, is

$$\text{BR}(\mu \rightarrow e a \gamma) < 1.1 \times 10^{-9} \quad (90\% \text{ CL}), \quad (4.8)$$

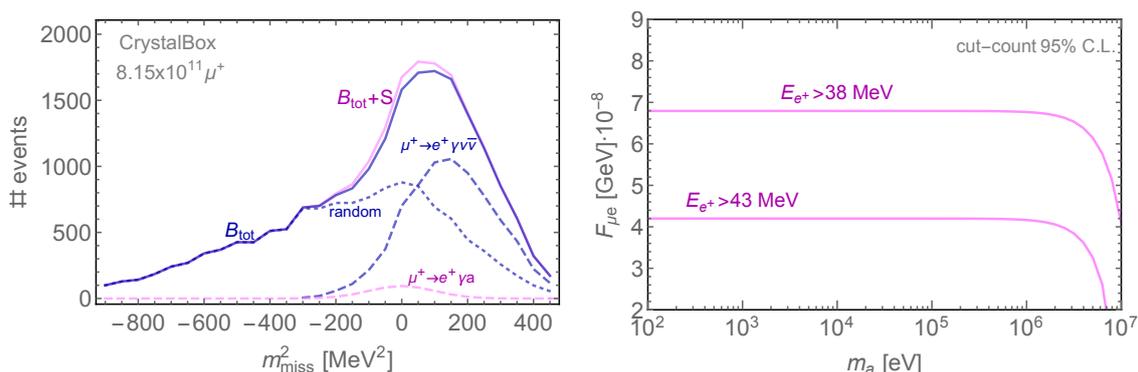


Figure 7. *Left:* Missing mass distribution of background and signal events in Crystal Box after the cuts in eqs. (4.6), (4.7) are implemented and experimental efficiencies taken into account. *Solid blue* is the total background composed by the *dashed blue* distribution of $\mu^+ \rightarrow e^+ \gamma \nu \bar{\nu}$ events and the *dotted blue* distribution of events where a hard photon from $\mu^+ \rightarrow e^+ \gamma \nu \bar{\nu}$ and a positron from $\mu^+ \rightarrow e^+ \nu \bar{\nu}$ are randomly coincident. *Solid magenta* shows the combined background plus signal distribution. The *dashed magenta* is the signal shape smeared by the experimental resolution for $m_a = 0$ and $\text{BR}(\mu^+ \rightarrow e^+ \gamma a) = 3 \times 10^{-9}$. *Right:* the 95% C.L. sensitivity of Crystal Box that we obtain using the simple cut and count scheme as described in the main text. The two lines indicate how the final bound on $F_{\mu e}$ depends on the positron energy loss before it reaches the detector. The actual experimental bound reaches a 30% higher value of $F_{\mu e}$ and is reported in eq. (4.9).

to be compared with the theory prediction,

$$\text{BR}(\mu \rightarrow e a \gamma) \approx (3.7 - 9.7) \times 10^{-10} \left(\frac{10^9 \text{ GeV}}{F_{\mu e}} \right)^2, \quad (4.9)$$

where the range is due to the variation in the phase space integral, $\mathcal{I}(x_{\min}, y_{\min}, \eta_a) \approx (0.004 - 0.011)$, obtained by varying the cut on the upstream positron energy between 38 MeV and 43 MeV. The 95% C.L. bound is then $F_{\mu e} \gtrsim (5.1 - 8.3) \times 10^8 \text{ GeV}$, and is always subdominant compared to the combination of the bounds obtained from $\mu^+ \rightarrow e^+ a$ by Jodidio et al. [9] and TWIST [10], cf. table 1.

For completeness, we discuss a simple procedure that we use to roughly reproduce the Crystal Box result and extend it to $m_a > 0$. The results of this recasting procedure are shown in the right panel of figure 7. We first extract the experimental efficiency for the $\mu^+ \rightarrow e^+ \gamma a$ signal from the bound on the branching ratio, $\text{BR}_{\text{excl.}}$, in eq. (4.8), which corresponds to $N_{\text{excl.}} = 165$ signal events [47]. Given that the total number of stopped muons is $N_{\mu^+} = 8.15 \times 10^{11}$ we get

$$\epsilon_S = \frac{N_{\text{excl.}}}{N_{\mu^+} \times \text{BR}_{\text{excl.}}} \simeq 0.18. \quad (4.10)$$

The signal shape can be extracted from the binned MC sample given in ref. [75], and is well described by a Gaussian centered at $m_{\text{miss}}^2 = 0$ with a variance of $\sigma_S = 100 \text{ MeV}^2$.

For a given mass m_a one can then estimate the sensitivity of Crystal Box on $\text{BR}(\mu \rightarrow e a \gamma)$ by considering the expected background and signal in a missing mass window, $(m_a^2 -$

$\sigma_S, m_a^2 + \sigma_S$) and using the asymptotic formula at 90% C.L. For $m_a = 0$, this procedure gives $\text{BR}(\mu \rightarrow ea\gamma) < 1.6 \times 10^{-9}$ which is weaker than the experimental bound in eq. (4.9) only by about 30%. This is not surprising since the full likelihood analysis has better discriminating power than the simple cut and count analysis we are performing. The bound on the branching ratio gets stronger for $m_{\text{miss}}^2 > 100 \text{ MeV}^2$ because the background is suppressed. However, this effect is compensated by the phase space suppression of the signal, eq. (4.1). All in all, the bound on $F_{\mu e}$ is constant up to $m_a \simeq 10 \text{ MeV}$ after which the signal phase space starts to shrink significantly with increasing ALP mass, given the strong energy cuts on the photon and positron energies.

4.2 Possibilities for future search at MEG II

We now briefly comment on the possibility for MEG II to improve on the Crystal Box result discussed above. Naively, the MEG II luminosity will exceed the one of Crystal Box by at least 3 orders of magnitude. In optimal conditions, this could lead to an increase of sensitivity on $F_{\mu e}$ of more than a factor of 5 with respect to the current bound in table. 1. With such an optimistic gain MEG II could start probing new parameter space beyond the current TWIST bound for the $V - A$ ALPs. An improvement on $\mu \rightarrow ea\gamma$ would complement the MEGII-fwd proposal, section 3.2, since this is only sensitive to ALPs with some amount of right handed couplings to the SM leptons. The combination of the two searches, $\mu \rightarrow ea\gamma$ and $\mu \rightarrow ea$ at MEGII-fwd would then fully cover the possible chiral structures of the ALP couplings.

Clearly, the above naive estimate for the improvement on the reach is far from guaranteed. A more realistic estimate of the reach would require a dedicated trigger study. The current energy and angular cuts of the MEG II trigger dedicated to the $\mu \rightarrow e\gamma$ search are designed to select a very energetic positron and a photon exactly back to back [77]. Keeping these cuts will greatly suppress the signal rate of $\mu \rightarrow e\gamma a$, and make it impossible to perform the search. An interesting possibility would be to relax the energy and angular cuts of the trigger down to similar values than the ones in Crystal Box. The final expected sensitivity at MEG II will strongly depend on the signal vs. background efficiencies whose detailed determination goes beyond the scope of this paper and is left for future investigations.

5 ALPs in τ decays

The search strategies for $\tau \rightarrow \mu a$ and $\tau \rightarrow ea$ decays are qualitatively different from the $\mu \rightarrow ea$ searches. The main differences can be traced to the fact that τ has a much shorter life-time ($3 \times 10^{-13} \text{ s}$ vs. $2 \times 10^{-6} \text{ s}$ for the muon), that it has many more decay channels, and that from the τ production it is not possible to unambiguously reconstruct the τ rest frame.

5.1 Past searches at ARGUS

The ARGUS experiment in 1995 derived bounds on tauonic LFV ALP couplings [43]. The tau data sample was produced from e^+e^- collisions in the DORIS II storage ring at

DESY at a center of mass energy varying between 9.4 and 10.6 GeV with a total integrated luminosity of 472 pb^{-1} .

The challenge in $\tau \rightarrow \ell_i a$ search is to disentangle the signal decay from the SM $\tau \rightarrow \ell_i \nu \bar{\nu}$ decays. The search would be easier in the tau rest frame, since then the lepton from $\tau \rightarrow \ell_i a$ is monochromatic and one can do a line search on top of the smooth $\tau \rightarrow \ell_i \nu \bar{\nu}$ background. Unfortunately, the tau is pair-produced in $e^+e^- \rightarrow \tau^+\tau^-$ collisions. Each of the tau's decays into a final state with at least one invisible particle, making exact reconstructions of the tau rest frames impossible. Instead, the ARGUS analysis used the “pseudo-rest frame” technique. The idea is to require one side of the $\tau^+\tau^-$ pair to decay into a three prong hadronic mode. The direction of the τ momentum is then approximated by the direction of the combined momentum of the three prong decay products. In the center of mass of e^+e^- collision the tau energy equals the beam energy, while the two taus are back to back. This gives enough constraints so that one can boost the leptonically decaying tau to its, approximate “pseudo-rest frame”. The crucial property of this frame is that the sensitivity on LFV two body tau decays into ALPs does not depend much on the ALP mass (see [43] for further details).

For a massless ALP ARGUS obtained [43]

$$\text{BR}(\tau \rightarrow e a) < 2.7 \times 10^{-3} \quad (95\% \text{ C.L.}) \quad \Rightarrow \quad F_{\tau e} \gtrsim 4.3 \times 10^6 \text{ GeV}, \quad (5.1)$$

$$\text{BR}(\tau \rightarrow \mu a) < 4.5 \times 10^{-3} \quad (95\% \text{ C.L.}) \quad \Rightarrow \quad F_{\tau \mu} \gtrsim 3.3 \times 10^6 \text{ GeV}. \quad (5.2)$$

The bound on $\text{BR}(\tau \rightarrow \mu a)$ is less stringent than $\text{BR}(\tau \rightarrow e a)$ at low masses while they become comparable for higher masses. The final bound on the ALP decay constant from $\text{BR}(\tau \rightarrow \mu a)$ is shown in figure 1. The mass dependence of the bound is predominantly due to the phase space suppression of the two-body decay for heavier ALPs.

5.2 Future searches at Belle-II

Belle and Belle-II. While Belle and Babar collected ≈ 2000 times larger datasets of τ 's than ARGUS, no experimental searches for $\tau \rightarrow \ell_i a$ were performed yet. However, a recent simulation of the expected limit at the Belle experiment with integrated luminosity of 1020 fb^{-1} was performed in ref. [48], and obtained for a massless ALP at 90% CL)

$$\text{Belle (1/ab) prospect: } \text{BR}(\tau \rightarrow \mu a) < 1.1 \times 10^{-4} \quad \Rightarrow \quad F_{\tau \mu} \gtrsim 2.1 \times 10^7 \text{ GeV}. \quad (5.3)$$

Notice that this bound is almost exactly a factor of $\sqrt{2000}$ more stringent than the present one from ARGUS, eq. (5.2). Using the same simple rescaling with the luminosity, we obtain for the expected 95% CL limit for a massless ALP for Belle II with 50 ab^{-1} ,

$$\text{Belle-II (50/ab) prospect: } \text{BR}(\tau \rightarrow \mu a) < 2.0 \times 10^{-5} \quad \Rightarrow \quad F_{\tau \mu} \gtrsim 4.9 \times 10^7 \text{ GeV}. \quad (5.4)$$

The limit as a function of ALP mass is shown with the purple line in figure 1. In the absence of MC analysis of Belle or Belle II reach for $\tau \rightarrow ea$, we estimate the Belle II sensitivity by performing the naive rescaling of the ARGUS result with luminosity, which gives

$$\text{Belle-II (50/ab) prospect: } \text{BR}(\tau \rightarrow e a) < 8.3 \times 10^{-6} \quad \Rightarrow \quad F_{\tau e} \gtrsim 7.7 \times 10^7 \text{ GeV}. \quad (5.5)$$

Belle II may improve the ARGUS searches for $\tau \rightarrow ea$ and $\tau \rightarrow \mu a$ transitions beyond mere increase in statistics. First of all, it could be interesting to explore the reach on $\text{BR}(\tau \rightarrow \mu a \gamma)$ and $\text{BR}(\tau \rightarrow e a \gamma)$, especially since for muons $\mu \rightarrow ea\gamma$ gives constraints that are not that far from the two-body $\mu \rightarrow ea$ decay (see [78] for similar comments in the context of a light Z'). Secondly, further improvements of $\tau \rightarrow \ell_i a$ searches may be possible. A possibly interesting direction, while still using the tau “pseudo-rest frame”, is to employ in addition variables that tag the tau polarization, such as the directions of pions in two prong tau decays. If successful, this could allow to further suppress the SM $\tau \rightarrow \ell_i \nu \bar{\nu}$ background, similarly to what was done for $\mu \rightarrow ea$ decays in the experiment by Jodidio et al. [9].

6 Bounds from astrophysics and cosmology

The bounds on ALP couplings from the astrophysical observations and from cosmology fall into two categories, depending on whether the ALP is assumed to constitute the observed DM relic abundance or not. In section 6.1 we first discuss the constraints from stellar cooling, which do not depend on whether or not ALP is the DM. In section 6.2 we then explore in which parts of the parameter space the LFV ALP could explain the observed DM abundance.

6.1 Bounds from stellar cooling

The emission of light particles inside stars can alter stellar evolution to an extent that is in conflict with observations. This leads to powerful constraints on the interactions of such light particles with matter and radiation [79]. Our primary interest here are the ALP couplings to leptons. In this context, ALP couplings to electrons can lead to efficient energy loss mechanisms in stars. For massless ALP (such as the QCD axion) the studies of red giants (RG) [80, 81] and white dwarfs (WD) [44], give roughly comparable bounds, which at 95% C.L. are¹⁰

$$F_{ee}^A \gtrsim 4.6 \times 10^9 \text{ GeV (WD)}, \quad F_{ee}^A \gtrsim 2.4 \times 10^9 \text{ GeV (RG)}. \quad (6.1)$$

In both cases the dominant cooling mechanism is ALP bremsstrahlung in electron–nucleus scattering, $e^- + N \rightarrow e^- + N + a$. This dominates over the Compton process, $\gamma + e^- \rightarrow e^- + a$, and electron-electron bremsstrahlung, $e^- + e^- \rightarrow e^- + e^- + a$, which are relevant only when electrons are non-degenerate [79].

For non-negligible ALP masses, the cooling rates are expected to be Boltzmann-suppressed. Following ref. [83], we estimate the resulting constraints on massive ALPs by rescaling the energy loss rates with the ratio $R(m_a, T)$ of ALP energy densities \mathcal{E}_a for the massive and massless case

$$R(m_a, T) \equiv \mathcal{E}_a(m_a, T) / \mathcal{E}_a(0, T). \quad (6.2)$$

¹⁰Interestingly, several stellar systems exhibit hints of non-standard energy losses. The global fit performed in ref. [82] finds that an axion/ALP solution to these anomalies with a coupling $5.4 \times 10^9 \text{ GeV} \lesssim F_{ee}^A \lesssim 8.1 \times 10^9 \text{ GeV}$ (1σ range) is preferred at the 3σ level over the case of only the standard energy loss through neutrinos.

The energy densities are given by

$$\mathcal{E}_a(m_a, T) = \frac{1}{2\pi^2} \int_{m_a}^{\infty} \frac{E^2 \sqrt{E^2 - m_a^2}}{e^{E/T} - 1} dE = \begin{cases} \frac{\pi^2}{30} T^4 & m_a \ll T, \\ \frac{1}{(2\pi)^{3/2}} T^4 \left(\frac{m_a}{T}\right)^{5/2} e^{-m_a/T} & m_a \gg T. \end{cases} \quad (6.3)$$

Since the cooling rates scale with $(F_{ee}^A)^{-2}$, the constraints on massive ALPs in figure 1 are obtained from the bounds on F_{ee}^A for the massless ALP by rescaling them with the factor $\sqrt{R(m_a, T)}$. Because of the Boltzmann suppression the star cooling bounds rapidly shut off for ALP masses above $m_a \approx 2T$, where $T_{\text{RG}} \approx 10^8 \text{K} \approx 8.6 \text{keV}$ and $T_{\text{WD}} \approx 10^7 \text{K} \approx 0.8 \text{keV}$.

For heavier ALPs the relevant astrophysics constraints are due to neutron star cooling in supernova (SN) explosions, since the nascent proto-neutron star (PNS) reaches a temperature of order 30 MeV a few seconds after the start of the supernova explosion [84, 85]. In order to estimate the SN bound on F_{aa}^A we use the expression for the energy loss rate per unit mass, ϵ , through electron-nucleon bremsstrahlung under highly degenerate conditions [79, 86],

$$\epsilon = \frac{\pi}{15} \alpha_{\text{em}}^2 \frac{T^4}{m_n (F_{ee}^A)^2} Y_p \mathcal{I}. \quad (6.4)$$

Here, $Y_p = n_p/n_B$ is the number density of protons relative to the baryon number density, while \mathcal{I} is the angular integral that includes plasma screening effects,

$$\mathcal{I} = \int \frac{d\Omega_2}{4\pi} \frac{d\Omega_a}{4\pi} \frac{(1 - \beta_F^2) [2(1 - c_{12}) - (c_{1a} - c_{2a})^2]}{(1 - c_{1a}\beta_F)(1 - c_{2a}\beta_F)(1 - c_{12})(1 - c_{12} + \kappa_{\text{DH}}^2)}. \quad (6.5)$$

It depends on the electron velocity at the Fermi surface, $\beta_F = p_F/E_F = p_F/\sqrt{p_F^2 + m_e^2} \approx 1$, while c_{12}, c_{1a}, c_{2a} are the cosines of the angles between the 3-momenta \mathbf{p}_1 (\mathbf{p}_2) of the incoming (outgoing) electron and the ALP 3-momentum \mathbf{p}_a , respectively. The screening effects enter through the parameter $\kappa_{\text{DH}}^2 = k_{\text{DH}}^2/2p_F^2$, where k_{DH} is the Debye screening scale. Note that the PNS with temperature $T_{\text{NS}} \approx 30 \text{MeV}$ and mass density $\rho \approx \rho_{\text{nuclear}}$ can be treated as composed of weakly coupled degenerate plasmas. The electron screening can then be neglected with respect to the proton screening, giving $k_{\text{DH}}^2 = 4\pi\alpha_{\text{em}}n_p/T_{\text{NS}}$. Since the main contribution to the angular integral (6.5) comes from the forward direction, $c_{12} = 1$ and $c_{1a} = c_{2a}$, the integral is well approximated by the simplified form obtained by setting $c_{2a} = c_{1a}$ in the denominator. In the ultra-relativistic case then [83]

$$\mathcal{I} = \frac{2 + \kappa_{\text{DH}}^2}{2} \log \frac{2 + \kappa_{\text{DH}}^2}{\kappa_{\text{DH}}^2} - 1. \quad (6.6)$$

For the numerical evaluation we use $n_B = 1.4 \times 10^6 \text{MeV}^3$, $Y_p = 0.2$, $p_F = 204 \text{MeV}$, so that $\kappa_{\text{DH}} = 0.10$ and $\mathcal{I} = 4.3$. This implies for the energy loss rate

$$\epsilon = 1.2 \times 10^{20} \frac{\text{erg}}{\text{g s}} \left(\frac{10^7 \text{GeV}}{F_{ee}^A} \right)^2. \quad (6.7)$$

Imposing the crude bound on the energy loss of $\epsilon \lesssim 10^{19} \text{ erg g}^{-1} \text{ s}^{-1}$ [83], leads in the case of a massless ALP to the bound

$$F_{ee}^A \gtrsim 3.4 \times 10^7 \text{ GeV (SN1987A)}. \quad (6.8)$$

For massive ALP we rescale this bound by the factor $R(m_a, T)$ in eq. (6.2) with $T \rightarrow T_{\text{NS}} \approx 30 \text{ MeV}$. Note that for heavy ALPs the $m_a^{5/2}$ dependence in the energy loss rate (6.3) appreciably counteracts the $\exp(-m_a/T)$ suppression, so that the SN bounds on F_{aa}^A are important up to $m_a \approx 200 \text{ MeV}$, cf. figure 1.

A qualitatively different regime is obtained for small values of F_{ee}^A . For small enough F_{ee}^A the ALP interactions become so strong that the ALP remains trapped within the stellar material, in which case there are no bounds on F_{ee}^A from the stellar cooling constraints. Following ref. [83], we estimate this limiting values of F_{ee}^A for RG, WD and NS systems by requiring that the mean free path λ of the ALP is smaller than the corresponding stellar effective radius R_0 , cf. table 2. The mean free path is calculated from ALP decay and absorption rates,

$$\lambda^{-1} = \beta^{-1} \Gamma_{\text{abs}} + (\beta\gamma)^{-1} \Gamma_{\text{decay}}. \quad (6.9)$$

The absorption rate is approximately given by the total energy loss rate per volume, $\rho\epsilon$, divided by the ALP energy density, \mathcal{E}_a (even if ALPs are not in thermal equilibrium)

$$\Gamma_{\text{abs}} = \rho\epsilon/\mathcal{E}_a. \quad (6.10)$$

As this estimate just follows from the principle of detailed balance, the absorption rate is independent of the ALP mass. We can therefore use for the energy loss per unit mass, ϵ , the result in eq. (6.4), while $\mathcal{E}_a = \mathcal{E}_a(0, T)$ in eq. (6.2). The dependence of λ on the ALP mass dominantly enters through the kinematical factors β and γ which are given by the integrals

$$\beta = \frac{\langle (E^2 - m_a^2) \rangle_T}{\langle E\sqrt{E^2 - m_a^2} \rangle_T}, \quad \beta\gamma = \frac{\langle E/m_a (E^2 - m_a^2) \rangle_T}{\langle E\sqrt{E^2 - m_a^2} \rangle_T}, \quad (6.11)$$

where $\langle \dots \rangle_T$ denotes a thermal average using the Bose-Einstein thermal distribution of ALPs. In the limit of large masses, $m_a \gg T$, one finds $\beta \approx \gamma \approx \sqrt{8/\pi} \sqrt{T/m_a}$.

For the decay rates, Γ_{decay} , we only consider decays to electrons and photons induced by F_{ee}^A . The resulting decay widths are given in section 2 (we set $E_{\text{UV}} = 1$ in these expressions, for definiteness). The contributions from $\gamma\gamma$ decay channel are always subdominant — for low ALP masses, where the decays to photons dominate Γ_{decay} , the absorption rate is more important in determining λ , while for higher ALP masses the decays to electrons or even to muons dominate over decays into photons.

We have used the inputs in table 2 (taken from ref. [79]) to derive the bounds in the trapping regime in figure 1. Note that for WDs the plasma is typically strongly coupled, so that the Debye screening given in eq. (6.6) is not an appropriate description. In these case we simply use $\mathcal{I} \approx 1$, which provides a good approximation [79, 87, 88].

	T	ρ	p_F	Y_p	R_0
RG	8.6 keV	4.3 MeV ⁴	409 keV	0.5	10 ⁴ km
WD	0.8 keV	7.7 MeV ⁴	495 keV	0.5	10 ³ km
PNS (SN)	30 MeV	1.3 × 10 ⁹ MeV ⁴	204 MeV	0.2	10 km

Table 2. Numerical inputs used in the evaluation of the stellar cooling bounds.

We have also included the recent bounds on ALP couplings to muons obtained in ref. [45] from SN1987A (see also ref. [46]). In contrast to our rough description of the PNS, these authors have used dedicated simulations which lead to the robust (and conservative) bound of $F_{\mu\mu}^A \geq 1.3 \times 10^8$ GeV for a massless axion. As for the case of the bounds on electron couplings from WDs and RGs, we simply rescale this bound by $\sqrt{R(m_a, T)}$ defined in eq. (6.2) (with $T \approx 30$ MeV) in order to account for non-zero axion masses. Note that the bound on muons is stronger than the SN bound on electron couplings in eq. (6.8), because energy loss through the Compton process for non-relativistic and non-degenerate muons is much more efficient than the same process for highly relativistic and degenerate electrons, see refs. [79, 89].

Finally, we comment on the bound on $\text{BR}(\mu \rightarrow ea)$ from SN1987A. This decay contributes to the cooling of the PNS, with a cooling rate that is given by (see ref. [23])

$$\epsilon = \frac{m_\mu^3 \Gamma(\mu \rightarrow ea)}{\pi^2 \rho (m_\mu^2 - m_n^2)} \int_0^\infty p_\mu dp_\mu \int_{p_e^{\min}}^{p_e^{\max}} p_e dp_e \frac{E_\mu - E_n}{E_\mu E_e} f_\mu(E_\mu) (1 - f_e(E_e)), \quad (6.12)$$

where p_e^{\max} (p_e^{\min}) is the maximal (minimal) electron momentum in the $\mu \rightarrow ea$ decay, if μ has momentum $p_\mu = |\mathbf{p}_\mu|$, and $E_i^2 = p_i^2 + m_i^2$ are the energies in the PNS's rest frame. Moreover, $f_i(E_i) = 1/(1 + \exp(\frac{E_i - \mu_i}{T}))$ are the Fermi distribution functions with μ_i the chemical potentials, which is the only non-trivial input in eq. (6.12). Using the numerical values $\mu_\mu \approx 41$ MeV and $\mu_e \approx 190$ MeV which we obtained as described below, the cooling rate becomes

$$\epsilon \approx 10^{19} \frac{\text{erg}}{\text{g s}} \left(\frac{\text{BR}(\mu \rightarrow ea)}{4 \times 10^{-3}} \right). \quad (6.13)$$

The resulting bound on $\text{BR}(\mu \rightarrow ea)$ is therefore about three orders of magnitude weaker than the constraints from laboratory experiments.

It is instructive to compare this result with the case where Pauli blocking in eq. (6.12) is neglected (i.e. $f_e(E_e) \rightarrow 1$), in which case the energy loss rate would be approximated by

$$\epsilon \simeq \frac{n_\mu}{\rho} (m_\mu - m_e) \Gamma(\mu \rightarrow ea) = \frac{Y_\mu}{m_n} (m_\mu - m_e) \Gamma(\mu \rightarrow ea). \quad (6.14)$$

Using $Y_\mu \approx 2.9 \times 10^{-3} \simeq Y_e \exp[-(m_\mu - m_e)/T]$ (see below), the resulting bound on $\text{BR}(\mu \rightarrow ea)$ is about a factor 50 larger than the exact result, indicating that Pauli blocking is an important effect. This can be understood by the large Fermi energy of the electrons, $E_F \approx p_F \approx 200$ MeV, which implies that muons at rest can hardly decay because almost all the relevant electron levels are filled.

The above values of μ_μ, μ_e and Y_μ were obtained by treating the PNS as a non-interacting Fermi gas, where all particles are in thermal and chemical equilibrium (see also ref. [23]). At given temperature the number density of a given particle then depends only on its chemical potential, which in turn can be written as a linear combination of the chemical potentials associated with the conserved quantum numbers, i.e., charge Q , baryon number B and the individual lepton numbers L_i . Furthermore, we assumed that neutrinos are trapped inside the PNS, which implies that the muon number density vanishes, $Y_{L_\mu} = Y_\mu + Y_{\nu_\mu} = 0$ [90]. For the electron number density (relative to the baryon number density n_B) we used $Y_{L_e} = 0.3$ [79]. With these three input values and the constraint of charge neutrality, $n_Q = 0$, one can numerically solve for the 4 unknown chemical potentials $\mu_Q, \mu_B, \mu_{L_e}, \mu_{L_\mu}$ and thus obtain number densities n_i and chemical potentials μ_i for all the involved particles. The tau leptons could be trivially included, taking $Y_{L_\tau} = 0$. However, unlike Y_μ , the tau number density is negligible, $Y_\tau \approx 0$. The electron's Fermi energy of about 200 MeV is just large enough to excite a relatively substantial muon population via electron scattering, while taus are simply too heavy. Note that this rather crude treatment of the PNS's thermodynamics appears to be consistent with the results from the dedicated simulations in ref. [45], since we obtain similar bounds on $F_{\mu\mu}^A$ using the expressions for the energy loss rate through Compton (as appropriate for muons) with the above input.

Finally, the derived SN bounds depend crucially on the SN explosion mechanism. If the SN1987A was not triggered by the canonical delayed neutrino mechanism but is rather due to the collapse-induced thermonuclear explosion, there would be no resulting bounds on the ALP couplings [91]. However, this interpretation would be disfavored, if the possible ALMA detection of a compact object in the remnant of SN 1987A [92], consistent with the neutron star [93], is confirmed.

6.2 ALP Dark Matter

Next, we explore under what circumstances the LFV ALP is a viable DM candidate, with the correct DM relic abundance. This will lead us to two important conclusions. Firstly, since the expected sensitivity of future LFV experiments is $f_a \gtrsim 10^{10}$ GeV, well above the present astrophysical bounds, the LFV experiments will explore a part of the parameter space of an ultralight ALP DM. Secondly, the LFV experiments cannot resolve ALP masses below 1 MeV (i.e. the typical experimental mass resolution). Below this mass range pinpointing the ALP mass will require experiments that search for ALP DM using other means, through its couplings to photons and/or electrons.

The minimal requirement for ALP to be the DM is that it is stable on timescales longer than the lifetime of the Universe. Assuming the $a \rightarrow \gamma\gamma$ decay dominates, this translates into the following constraint,

$$\frac{H_0}{\Gamma_{\text{tot}}} = H_0\tau_a > 1, \quad \text{where} \quad H_0\tau_a \simeq 5.4 \left(\frac{1}{E_{\text{eff}}^2} \right)^2 \left(\frac{10 \text{ keV}}{m_a} \right)^3 \left(\frac{f_a}{10^{10} \text{ GeV}} \right)^2. \quad (6.15)$$

Taking $f_a = 10^{10}$ GeV as a reference value, this means that the ALP DM probed in LFV experiments must have a mass below 10 keV. If other decay channels, such as $a \rightarrow \nu_i \bar{\nu}_j$,

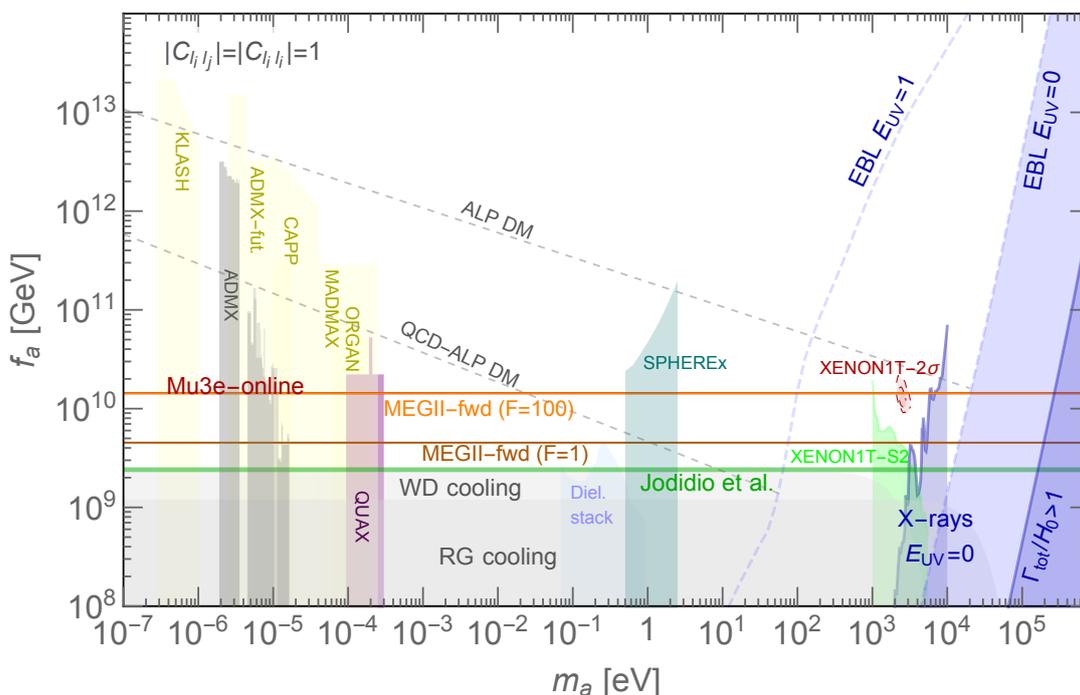


Figure 8. The impact of present and future $\mu \rightarrow ea$ searches compared to other light ALP DM searches, taking $E_{UV} = 1, 0$ as two representative examples. The *green solid* line shows the current best bound on the isotropic LFV ALP [9], the *(dark) orange thin* line gives our MEGII-fwd projection assuming $F = 100$ focusing enhancement (no focusing). The *dark red* line (overlapping with the orange thin line) shows the sensitivity of Mu3e-online analysis [42]. In the blue region enclosed by the *solid blue* line the ALP decays within the present Hubble time, while the region to the right of the *dashed blue* line is excluded by the extragalactic diffuse background light measurements for $E_{UV} = 0, 1$. We also show the X-rays constraints for $E_{UV} = 0$ [94, 95]. The red blob indicates where ALP DM could explain the XENON1T anomaly [96]. The *dashed gray lines* denote two scenarios where the observed DM relic abundance is due to ALPs produced through the misalignment mechanism. On the upper line the ALP mass is temperature independent, cf. eq. (6.17), while on the lower line the temperature dependence is parametrically similar to the one for the QCD axion, cf. eq. (6.18). The *gray shaded* regions are excluded by the star cooling bounds, and the ADMX data [97–99]. The light green region is excluded by the S2-only analysis of XENON1T [100] and Panda-X [101]. The *purple shaded* region shows the future reach of axion-magnon conversion experiments such as QUAX [102–104]. Regarding the coupling to photons, the *cyan band* shows the future sensitivity of SPHEREx estimated in ref. [105], assuming ALP decay exclusively to two photons, while the *yellow bands* show the future sensitivities of resonant microwave cavities such as ADMX [106], CAPP [107], KLASH [108], and ORGAN [109], dielectric haloscopes such as MADMAX [110] and the reach of the dielectric stack proposal [111] is shown with *light blue*.

are appreciable, then the above bound on the ALP mass is correspondingly lowered (this is for instance the case for the majoron, see section 7.4). For the rest of this section we will assume that the ALP decay channels apart from $a \rightarrow \gamma\gamma$ can be neglected.

The $a \rightarrow \gamma\gamma$ decays contribute to the extragalactic background light (EBL) and may be bounded by EBL measurements. The ALP decay results in a line at frequency $\nu_a = 1.2 \times 10^{14} m_a/\text{eV}$ Hz with intensity

$$\nu_a I_{\nu_a} = \frac{5 \times 10^{-3} \text{ Wm}^{-2}\text{sr}^{-1}}{\tau_{\gamma\gamma} H_0} \simeq 6 \times 10^{-13} \text{ Wm}^{-2}\text{sr}^{-1} \left(\frac{m_a}{1 \text{ eV}} \right)^3 \left(\frac{10^{10} E_{\text{eff}} \text{ GeV}}{f_a} \right)^2. \quad (6.16)$$

A conservative bound on the decay width $\Gamma(a \rightarrow \gamma\gamma)$ is obtained by requiring that the line intensity in eq. (6.16) is less than what is observed at that frequency [105]. An updated map of the EBL observations at different frequencies can be found in ref. [112]. For instance, the observed EBL intensity in the optical band is $10^{-8} \text{ Wm}^{-2}\text{sr}^{-1}$, constraining the axion width well below H_0 . Converting the EBL constraints to a bound on f_a gives the light blue exclusion region in figure 8 for $E_{\text{UV}} = 1$ and $E_{\text{UV}} = 0$. Stronger bounds can be obtained from the measurements of the X-ray microcalorimeters in the XQC rocket [94, 95]. These constraints will be further improved by future X-ray missions such as Athena [113] or by future line-intensity mapping campaigns [105]. Particularly relevant for the ALP parameter space is the SPHEREx project [114], a funded two-year mission by NASA with a planned launch in 2023, that will probe optical and near infrared frequencies corresponding to $m_a \sim \text{eV}$. The SPHEREx reach is denoted as green region in figure 8.

The ALP that satisfies the stability and decaying DM bounds could be a viable DM candidate. The main production mechanism in the allowed region of parameter space in figure 8 is the misalignment mechanism, first discussed in the context of the QCD axion in refs. [14–16] and then generalized to a generic ALP in ref. [115].¹¹ If inflation occurred below the scale of ALP global symmetry breaking, the initial misalignment of the ALP, a_0 , is frozen by the inflationary dynamics and acts as the initial condition. Conventionally, it is parametrized in terms of an angular variable, $a_0 = f_a \theta_0$, where $\theta_0 \in [0, \pi)$. As long as the Hubble expansion rate is large, $H > m_a$, the field is frozen at its initial value a_0 . At temperature T_{osc} , when $H(T_{\text{osc}}) \simeq m_a$, the field starts to oscillate and produces the ALP number density $n_a(T_{\text{osc}}) = \frac{1}{2} m_a^2 a_0^2$, which then expands adiabatically until the present time. For the case when ALP oscillations occur during the radiation dominated epoch the resulting ALP relic abundance is (see also [117])

$$\Omega_a^{T\text{-indep.}} h^2 = 0.12 \times 10^{-2} \sqrt{\frac{m_a}{\text{eV}}} \left(\frac{f_a}{10^{10} \text{ GeV}} \right)^2 \left(\frac{\theta_0}{\pi} \right)^2 \left(\frac{90}{g_*(T_{\text{osc}})} \right)^{1/4}. \quad (6.17)$$

Since the ALP mass is bounded from above by EBL constraints, the future reach of LFV searches (i.e. $f_a \approx 10^{10} \text{ GeV}$) will probe a region of parameter space where the production

¹¹The production of hot ALPs through freeze-in via lepton annihilation, $\ell^+ \ell^- \rightarrow a\gamma$, or lepton-photon collisions, $\ell^\pm \gamma \rightarrow e^\pm a$, could lead to a too large contribution to ΔN_{eff} . It is easy to check that this contribution is in fact negligible in the parameter space in figure 8, once the present bounds on the axion-electron couplings from stellar cooling are taken into account. (A systematic study of hot ALPs production through muon or tau couplings in the case where couplings to electrons are switched off has been performed in ref. [116].)

from misalignment does not suffice to obtain the total observed DM abundance ($\Omega_{\text{DM}}h^2 \simeq 0.12$ [118]) with an ALP mass independent on temperature. The relation between m_a and f_a that leads to the observed DM abundance for $\theta_0 \sim 1$ for a temperature independent ALP mass corresponds to the upper gray dashed line in figure 8. Below this line the ALP DM produced through the misalignment mechanism is under-abundant, while above this line a smaller value of θ_0 is needed in order to obtain $\Omega_{\text{DM}}h^2 = 0.12$.

An interesting alternative possibility that leads to enhanced misalignment production is if the ALP mass comes from a dynamical mechanism like the one of the QCD axion. At zero temperature the ALP mass is given by $m_a = \Lambda^2/f_a$, while at finite temperature the mass is suppressed, and is given by $m_a(T) = m_a(\Lambda/T)^b$, where $b = 4$ in QCD (an expression for b in a general gauge theory can be found in ref. [119]). The relic density from misalignment for this case has been studied in ref. [115], and more recently in ref. [117], and is given by

$$\Omega_a^{T\text{-dep}} = \Omega_a^{T\text{-indep.}} \left(\frac{2+b}{2}\right)^{\frac{3+b}{2+b}} \left(\frac{5}{8\pi} \frac{M_{\text{Pl}}}{f_a} \sqrt{\frac{90}{g_*(T_{\text{osc}})}}\right)^{\frac{b}{4+2b}}. \quad (6.18)$$

The two additional factors on the r.h.s. enhance the relic abundance with respect to the result in eq. (6.17). In figure 8 we show that for $b = 4$ this enhancement is large enough that the correct DM relic abundance is obtained in a large region of parameter space that will be tested by future LFV searches. The ALP abundance can be enhanced even further at small decay constants by tuning the ALP initial condition very close to $\theta_0 = \pi$, such that non-linearities dominate the production [19], or by mixing of the ALP with a dark photon in a presence of a magnetic field in the early Universe [18]. In summary, these different production mechanisms can make the ALP abundance match the current DM abundance in the whole parameter space shown in figure 8 which is not excluded by present constraints.

Parts of the LFV ALP parameter space will be probed by other means. In figure 8 we show two types of such probes, based either on axion couplings to electrons or to photon. All of these probes require the ALP to be the DM, and assume that the ALP is responsible for the full DM relic abundance. Note that in this case, for the range of masses shown in figure 8, the description of ALP in terms of classical background field is justified in the early Universe, since there are many ALPs inside a single de Broglie volume.

The ALP couplings to electrons can then lead to interesting constraints from electron recoil experiments where the ALP energy gets absorbed in the detector. The energy threshold of the DM experiment translate directly to the lower ALP mass that these experiment can probe. We show the current constraints from Panda-X [101] and from the S2-only analysis of XENON1T [100]. The fit of the XENON1T anomaly derived in [96] is also shown. Further improvements at lower ALP mass are expected from low threshold experiments like SENSEI [120]. The efficient axion-magnon conversion in an experiment such as QUAX [102, 103] could probe a portion of the ALP parameter space in the $m_a \sim 10\text{--}50 \mu\text{eV}$ window. The light purple region in figure 8 shows the future reach of QUAX derived in ref. [104] (the present sensitivity in ref. [121] is still outside the plotted range in f_a). At higher masses the low threshold DM absorption experiments based on existing technol-

ogy are generically weaker than the stellar cooling bounds, even if improvements can be foreseen with future technology under optimistic conditions [21, 22].

If the ALP coupling to photons is not suppressed, the standard searches for the QCD axion will cover significant parts of the parameter space, as seen in figure 8 for the case of $E_{UV} = 1$. The gray region around $0.2 \mu\text{eV}$ is excluded by current ADMX data [97–99]. The future axion haloscope campaigns will explore the ALP mass region between $0.2 - 20 \mu\text{eV}$. In figure 8 we show in yellow the estimated sensitivities of CAPP [107], KLASH [108], ORGAN [109], MADMAX [110] and the ADMX upgrade [106]. We also include the “dielectric stack” proposal which could have sensitivity beyond the current stellar cooling bounds between 0.1 and 1eV , depending on the value of E_{UV} [111]. The sensitivities of large-scale helioscopes such as IAXO [122, 123], and light-shining-through-wall experiments such as ALPS-II [124] lie below the current stellar cooling bounds for our choice of parameters and is not shown in figure 8.

7 LFV ALP models

So far we were concerned with the model independent bounds on LFV ALP couplings, eq. (2.1). In the remainder of this paper we focus instead on several representative models of ALPs with LFV couplings: the LFV QCD axion, the LFV axiflavor, the leptonic familon and the majoron. These examples are representative for broad classes of models, and illustrate how flavor-violating couplings to leptons can naturally arise for PNGBs of global symmetries addressing the strong CP problem, the SM flavor puzzle, or neutrino masses. In the first three models the LFV couplings are generated at tree level via non-universal charges of the global symmetry under which the ALP transform non-linearly while in the last one the LFV comes from loop-induced couplings. For each model above we also present the parameter space where the ALP can be a viable DM candidate as discussed in section 6.2.

The LFV QCD axion (section 7.1) and the LFV axiflavor (section 7.2) are two explicit realizations of the QCD axion, which elegantly solves the strong CP problem in the SM via the spontaneous breaking of a $U(1)$ Peccei-Quinn (PQ) symmetry that is anomalous under QCD. In section 7.1 we first show that in DFSZ-like models [12, 13] the QCD axion can naturally have LFV couplings while keeping the couplings to quarks flavor diagonal. In section 7.2 we go one step further and identify the PQ symmetry with a subgroup of the flavor symmetry that gives the hierarchical masses and mixings of the SM fermions. The LFV axiflavor is obtained for the $U(2)$ flavor group, since in this case the flavor violating couplings are parametrically larger in the leptonic than the quark sector.

The Leptonic familon (section 7.3) is the PNGB of a $U(1)$ flavor symmetry in the leptonic sector. The spontaneous breaking of this symmetry could explain the hierarchies among the charged leptons via the Froggatt-Nielsen mechanism [34, 125, 126] (for recent variations see [26, 127]). In the LFV familon setup the strengths of the LFV couplings depend on the texture of the PMNS matrix, as we will see in detail below.

Our final example, the majoron, is the PNGB associated with the spontaneous breaking of the lepton number [128, 129]. In section 7.4 we show that in a non-minimal class of

seesaw models the majoron has parametrically enhanced LFV couplings. In these theories an approximate generalized lepton number suppresses the neutrino masses [41, 130–138], without suppressing the majoron couplings to the SM.

7.1 The LFV QCD axion

The mass of the QCD axion is entirely due to the QCD anomaly, and is given by [139]

$$m_a = 5.691(51) \mu\text{eV} \left(\frac{10^{12} \text{ GeV}}{f_a} \right). \quad (7.1)$$

The value of the axion decay constraint f_a therefore completely determines the mass of the QCD axion, which for all the processes we consider is effectively massless.

Astrophysical constraints require the axion to be very weakly coupled, with a lifetime larger than the age of the universe and a mass below 3×10^{-2} eV. In this range the QCD axion is a perfectly viable cold DM candidate in large parts of the parameter space. One of the simplest scenarios for axion production is the misalignment mechanism described in section 6.2. In the QCD axion case the observed DM abundance is obtained for misalignment angles of order unity $\theta_0 \sim 1$ with an axion decay constants $f_a \sim 10^{(11 \div 13)}$ GeV. For smaller decay constants, within the reach of LFV experiments, the axion relic from the standard misalignment contribution is under-abundant unless non-trivial dynamics or tuning are invoked (see discussion in section 6.2).

The axion couplings to fermions in eq. (2.1) arise from rotating the PQ current to the fermion mass basis, with unitary rotations V^f defined by $V_L^{f\dagger} y_f V_R^f = y_f^{\text{diag}}$. Denoting the PQ charge matrices by X_f , one has

$$C_{f_i f_j}^{V,A} = -\frac{1}{2N} \left(V_R^{f\dagger} X_{f_R} V_R^f \pm V_L^{f\dagger} X_{f_L} V_L^f \right)_{ij}, \quad (7.2)$$

where $2N$ is the domain wall number. This implies that off-diagonal couplings arise when the PQ charges are not diagonal in the same basis as the Yukawa couplings, y_f . Their sizes depend on the misalignment between the two bases, which is parametrized by the unitary rotations $V_{R,L}^f$. We focus on the situation where the PQ charges in the quark sector are universal, so that the QCD axion only has flavor violating couplings in the lepton sector. (This is of course not the most general case. If PQ charges in the quark sector are not universal, the results from ref. [23] apply, with the bound from $K^+ \rightarrow \pi^+ a$ leading to tight constraints on f_a .)

In the following, we specify a DFSZ-like model of the QCD axion with LFV couplings. The field content of the theory consists of the SM fermions, two Higgs doublets, $H_{1,2}$, and a complex scalar S that is a gauge singlet. The model contains an anomalous global U(1) PQ symmetry under which the scalar fields carry charges $X_S = 1$, $X_{H_2} = 2 + X_{H_1}$. As a consequence, the scalar potential contains the couplings $H_2^\dagger H_1 S^2$ and $(S^\dagger S)^2$, but not, for instance, $H_1^\dagger H_2 S^2$ or S^4 . The fermionic U(1)_{PQ} charges are flavor universal in the quark sector, $X_{u_{Ri}} = -X_{H_1}$, $X_{d_{Ri}} = X_{H_2}$, $X_{q_{Li}} = 0$, $i = 1, 2, 3$, while they are generation dependent in the lepton sector, such that the Yukawa interaction Lagrangian takes the

form (here $\tilde{H}_i = i\sigma_2 H_i^*$ and $a, b = 2, 3$)

$$\mathcal{L} = y_{1a}^e \bar{\ell}_{L1} e_{Ra} \tilde{H}_1 + y_{a1}^e \bar{\ell}_{La} e_{R1} \tilde{H}_1 + y_{ab}^e \bar{\ell}_{La} e_{Rb} \tilde{H}_2 - y^u \bar{q}_L u_R H_1 + y^d \bar{q}_L d_R \tilde{H}_2 + \text{h.c.} \quad (7.3)$$

The first generation leptons carry charges $X_{e_{R1}} = X_{H_1}$, $X_{\ell_{L1}} = 2$ under $U(1)_{\text{PQ}}$, while the 2nd and 3rd generation leptons have $X_{e_{Ra}} = X_{H_2}$, $X_{\ell_{La}} = 0$, where $a = 2, 3$. In eq. (7.3) y_{1a}^e and y_{a1}^e are complex 2-vectors, y_{ab}^e is a complex 2×2 matrix, while for simplicity we do not display the flavor indices on 3×3 complex Yukawa matrices in the quark sector, y^u, y^d . The forms of $y_{1a}^e, y_{a1}^e, y_{ab}^e$ and $y^{u,d}$ are not fixed by the $U(1)_{\text{PQ}}$ symmetry and are thus external to the discussion (presumably there is additional flavor dynamics that gives their form and thus the required hierarchy of SM fermion masses).

The form of the scalar potential is assumed to be suitable to induce vacuum expectation values for all scalars, with the ratio of Higgs vevs given by $v_2/v_1 = \tan \beta \equiv t_\beta$ and $\langle S \rangle \equiv v_{\text{PQ}}/\sqrt{2} \gg v_i$, while $v_1^2 + v_2^2 = v^2$ with $v = 246 \text{ GeV}$ the electroweak vev. The axion a is then mainly contained in S , i.e., $S = \langle S \rangle \exp(ia/v_{\text{PQ}}) + \dots$, and partially in the two Higgs doublets, $H_i = \langle H_i \rangle \exp(iX_{H_i} a/v_{\text{PQ}}) + \dots$ (here we only show the leading dependence on a). Requiring that the axion is orthogonal to the Goldstone boson eaten by the Z completely fixes the embedding of $U(1)_{\text{PQ}}$, i.e., it fixes the PQ charge of the two Higgs doublets to be $X_{H_1} = -2s_\beta^2$, $X_{H_2} = 2c_\beta^2$, where we used the shortened notation $s_\beta \equiv \sin \beta$, $c_\beta \equiv \cos \beta$.

It is conventional to remove the axion from the Yukawa interactions (7.3) through phase redefinitions of the SM fermion fields. Working in this basis, the axion couples derivatively to the fermions as in eq. (2.1), and in addition has couplings to gluons and photons induced by the color and EM anomalies

$$\mathcal{L}_a = \frac{a}{f_a} \frac{\alpha_s}{8\pi} G_{\mu\nu} \tilde{G}^{\mu\nu} + \frac{E}{N} \frac{a}{f_a} \frac{\alpha_{\text{em}}}{8\pi} F_{\mu\nu} \tilde{F}^{\mu\nu} + \frac{\partial_\mu a}{2f_a} \bar{f}_i \gamma^\mu \left[C_{ij}^V + C_{ij}^A \gamma_5 \right] f_j. \quad (7.4)$$

The axion decay constant is related to the PQ breaking vev, $\langle S \rangle = \sqrt{2} N f_a$. The anomaly coefficients are given by $2N = -6$ and $E/N = 4/3$.

The couplings to fermions are given by eq. (7.2). In the quark sector the PQ charges are flavor universal, and so are the axion couplings,

$$C_{u_i u_j}^V = C_{u_i u_j}^A = \frac{s_\beta^2}{3} \delta_{ij}, \quad C_{d_i d_j}^V = C_{d_i d_j}^A = \frac{c_\beta^2}{3} \delta_{ij}. \quad (7.5)$$

In contrast, in the charged lepton sector the PQ charges are not universal and therefore the axion couplings to charged leptons depend on the unitary rotations that diagonalize the Yukawas as in eq. (7.2). It is useful to introduce the hermitian matrices

$$\epsilon_{ij}^{eL} \equiv (V_L^e)^*_{1i} (V_L^e)_{1j}, \quad \epsilon_{ij}^{eR} \equiv (V_R^e)^*_{1i} (V_R^e)_{1j}, \quad (7.6)$$

which satisfy

$$0 \leq \epsilon_{ii}^{eP} \leq 1, \quad \sum_i \epsilon_{ii}^{eP} = 1, \quad |\epsilon_{ij}^{eP}| = \sqrt{\epsilon_{ii}^{eP} \epsilon_{jj}^{eP}}, \quad (7.7)$$

	Bound on f_a (in GeV)		
	Benchmark “V”	Benchmark “V + A”	Benchmark “V - A”
SN1987A	9.4×10^7	9.4×10^7	9.4×10^7
WD cooling	1.3×10^9	9.3×10^8	9.3×10^8
$\mu \rightarrow e a$	1.1×10^8	8.0×10^8	1.2×10^8

Table 3. Bounds on the axion decay constant f_a (in GeV) for the three benchmarks of the LFV axion model choosing $\beta = 1$, cf. eqs. (7.11)–(7.13).

for $P = L, R$. In terms of these parameters one has

$$C_{\ell_i \ell_j}^V = \frac{1}{3} \left[c_\beta^2 \delta_{ij} - \epsilon_{ij}^{eR} + \epsilon_{ij}^{eL} \right], \quad C_{\ell_i \ell_j}^A = \frac{1}{3} \left[c_\beta^2 \delta_{ij} - \epsilon_{ij}^{eR} - \epsilon_{ij}^{eL} \right]. \quad (7.8)$$

Note that the flavor diagonal parts of the vectorial couplings, $C_{f_i f_j}^V$, eqs. (7.5), (7.8), can be set to zero through fermion field redefinitions (these would introduce couplings to EW boson field strengths as in eq. (7.4) that are, however, not relevant for our analyses).

To show the impact of the experimental searches for LFV processes with muons, we construct three benchmarks for off-diagonal matrices ϵ_{ij}^{eL} and ϵ_{ij}^{eR} . To do so we first choose a particular form of the leptonic Yukawa couplings, where we assume that in the basis in which y_{ab}^e is diagonal the Yukawa 2-vectors y_{1a}^e, y_{1b}^e have zero couplings between the 1st and 3rd generation. Within these assumptions the charged lepton mass matrix is completely fixed, apart from a single continuous parameter, η ,

$$m_{ij}^e = \begin{pmatrix} 0 & \eta m_\mu & 0 \\ -m_e/\eta & m_{\text{eff}} & 0 \\ 0 & 0 & m_\tau \end{pmatrix}, \quad m_{\text{eff}} \equiv \sqrt{m_\mu^2(1 - \eta^2) + m_e^2(1 - 1/\eta^2)}. \quad (7.9)$$

The parameter η controls the size of left- and right-handed rotations. We restrict its values to the range $m_e/m_\mu \leq \eta \leq 1$ such that there are no unnaturally large cancellations when diagonalizing the mass matrix. Choosing three representative values of η gives

$$(V_L^e)_{12} \approx \begin{cases} \frac{1}{\sqrt{2}} & \eta = \frac{1}{\sqrt{2}}, \\ \sqrt{\frac{m_e}{m_\mu}} & \eta = \sqrt{\frac{m_e}{m_\mu}}, \\ \frac{m_e}{m_\mu} & \eta = \sqrt{2} \frac{m_e}{m_\mu}, \end{cases} \quad (V_R^e)_{12} \approx \begin{cases} -\frac{m_e}{m_\mu} & \eta = \frac{1}{\sqrt{2}}, \\ -\sqrt{\frac{m_e}{m_\mu}} & \eta = \sqrt{\frac{m_e}{m_\mu}}, \\ -\frac{1}{\sqrt{2}} & \eta = \sqrt{2} \frac{m_e}{m_\mu}, \end{cases} \quad (7.10)$$

which we take as the three representative benchmarks: the “V - A”, “V” and “V + A” scenarios, respectively. As per our assumptions, the only flavor violating couplings are between the 1st and the 2nd generation leptons.

More explicitly, the axion couplings in the three scenarios are,

- **Benchmark “V”** ($\eta = \sqrt{m_e/m_\mu}$)

$$C_{\mu e}^V \approx 2/3 \sqrt{m_e/m_\mu}, \quad C_{\mu e}^A = 0, \quad C_{ee}^A \approx c_\beta^2/3 - 2/3, \quad (7.11)$$

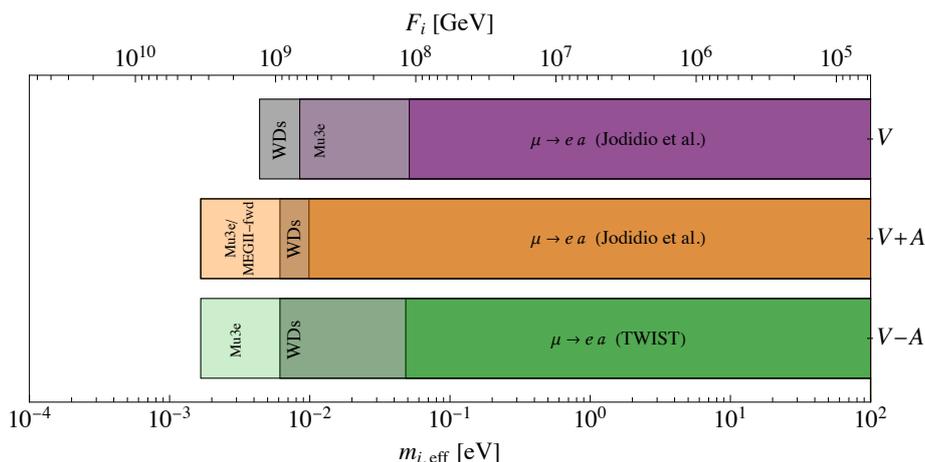


Figure 9. Present and expected future bounds on f_a and m_a for the LFV QCD axion in the three scenarios described in detail in the text, see also table 3 and eqs. (7.11)–(7.13). On the lower axis we indicate the corresponding values for the effective axion mass defined by $m_{i,\text{eff}} = 4.7 \text{ eV} \times 10^6 \text{ GeV}/F_i$.

- **Benchmark “V + A”** ($\eta = \sqrt{2}m_e/m_\mu$)

$$C_{\mu e}^V \approx C_{\mu e}^A \approx 1/6, \quad C_{ee}^A \approx c_\beta^2/3 - 1/2, \quad (7.12)$$

- **Benchmark “V – A”** ($\eta = 1/\sqrt{2}$)

$$C_{\mu e}^V \approx -C_{\mu e}^A \approx 1/6, \quad C_{ee}^A \approx c_\beta^2/3 - 1/2. \quad (7.13)$$

We can now reinterpret the model independent bounds on LFV ALPs, derived in sections 3–6, for the three LFV QCD axion benchmarks (choosing $\beta = 1$ as a representative value). The resulting bounds on f_a from $\mu \rightarrow ea$ and from WD cooling, obtained by rescaling respectively the bounds on $F_{\mu e}^{A,V}$, F_{ee}^A in table 1 by the appropriate values of $C_{\mu e}^{V,A}$, C_{ee}^A in the three benchmarks, are collected in table 3 and presented graphically in figure 9.

The SN1987 bound is modified with respect to the one discussed in section 6.1 due to the axion couplings to quarks and gluons that then result in the axion couplings to nucleons (due to smaller scattering cross sections the processes involving electrons lead only to subleading corrections). Adopting the treatment of ref. [140], the relevant bound is on the effective coupling to nucleons,

$$C_N \equiv \sqrt{C_n^2 + 0.29 C_p^2 + 0.27 C_p C_n}, \quad (7.14)$$

where $C_{p,n}$ are the axion couplings to protons and neutrons, respectively. Using the expressions in ref. [17] along with the values of couplings to quarks and gluons in eqs. (7.4), (7.5), we get for the LFV axion model, for all three benchmarks,

$$C_N = \sqrt{0.042 - 0.084 c_\beta^2 + 0.18 c_\beta^4}. \quad (7.15)$$

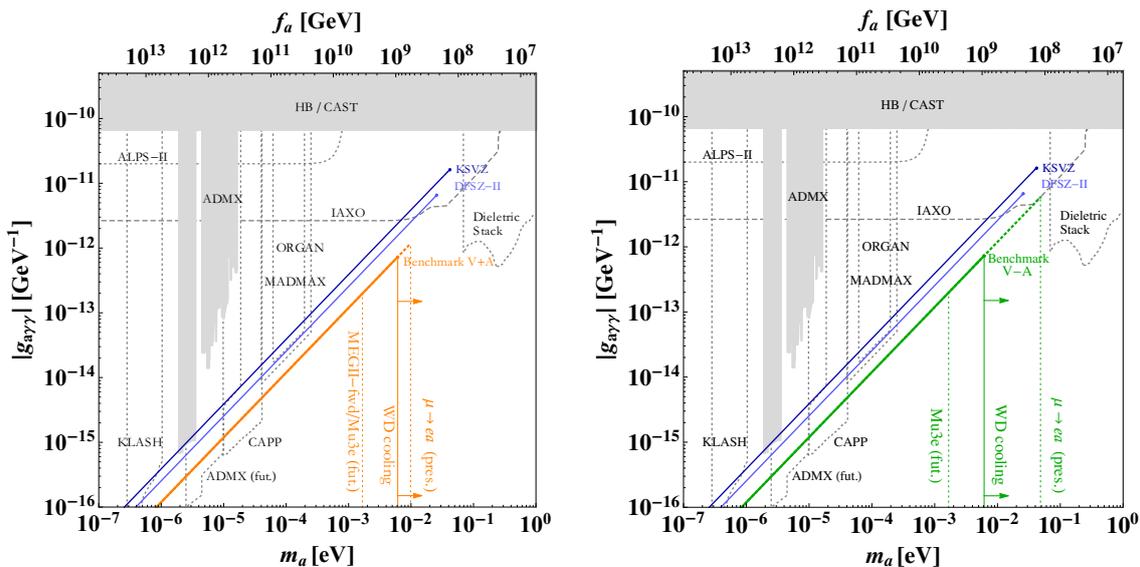


Figure 10. *Left:* the “ $V + A$ ” LFV axion corresponds to the *orange bold solid line*. *Right:* the “ $V - A$ ” LFV axion corresponds to the *green bold solid line*. The *grey shaded* regions show the present bounds on LFV QCD axion couplings to photons as a function of axion mass for the two benchmarks. The *grey dashed* lines denote future projected sensitivities on the photon coupling. The *solid orange/green* vertical lines show the present upper bound on the axion mass from WD cooling in the two models. The *dotted orange/green* vertical lines show the present bound and future reach on the axion mass from LFV experiments. For a comparison we also show lines corresponding to the standard KSVZ (*dark blue*) and DFSZ-II (*blue*) models. The former one is limited by the SN1987A bound, the latter one by WD cooling. For both the DFSZ-II and our LFV QCD axion models we set $\beta = 1$. See the main text for details.

For given c_β the bound on f_a follows from the bound on the effective decay constant $F_N \equiv 2f_a/C_N \geq 10^9$ GeV [140]. The resulting bound on f_a is of the order $f_a \gtrsim 10^8$ GeV, with mild dependence on c_β .

In figure 10 we show the constraints on the axion-photon couplings

$$g_{a\gamma\gamma} = \frac{1}{f_a} \frac{\alpha_{\text{em}}}{2\pi} (E/N - 1.92), \quad (7.16)$$

as a function of the axion mass, m_a . The orange (green) solid line in the left (right) panel shows $g_{a\gamma\gamma}$ as a function of m_a in the $V + A$ ($V - A$) benchmark up to the exclusion by the WD cooling constraints (the dotted continuation of the line is excluded by WD cooling but is not excluded by the direct $\mu \rightarrow e\gamma$ searches). Note that for the LFV QCD axion, $E/N = 4/3$, leading to smaller coupling to photons, $|g_{a\gamma\gamma}| \simeq 0.6 \times \alpha_{\text{em}}/(2\pi f_a)$, than for the flavor universal original KSVZ model, for which $E/N = 0$ and thus $|g_{a\gamma\gamma}| \simeq 1.9 \times \alpha_{\text{em}}/(2\pi f_a)$ [141, 142], and DFSZ-II model [12, 13], for which $E/N = 2/3$ and $|g_{a\gamma\gamma}| \simeq 1.3 \times \alpha_{\text{em}}/(2\pi f_a)$ (though large variations in this coupling are possible depending on the precise choices of the heavy fields and their charges [143–145]). The dotted orange (green) vertical lines show projected and present bounds from $\mu \rightarrow e\gamma$ searches, as denoted.

The gray regions are excluded by other axion experiments: CAST [146], cooling of horizontal branch stars (HB) and ADMX [97–99]. The gray dashed lines denote future projections from different axion searches already discussed in section 6.2. We show the reach of the future ADMX upgrade [106], of CAPP [107], KLASH [108], and ORGAN [109], MADMAX [110] and the “dielectric stack” proposal [111]. We also include the reach of large-scale helioscopes such as IAXO [122, 123], and light-shining-through-wall experiments such as ALPS-II [124].

Figure 10 demonstrates the complementarity between this diverse experimental program based on axion couplings to photons and electrons, and the reach of the LFV experiments MEG-fwd and Mu3e, in order to search for LFV axions. In particular, a signal in a future LFV search would be incompatible with an axion lighter than a few meV. Such an LFV axion line will be challenging to test in axion haloscopes because of the infamous DFSZ accidental suppression of the photon coupling, see eq. (7.16). In contrast, the future ADMX experimental campaign and CAPP will probe the LFV axion for masses between few μeV and few tens of μeV , which are inaccessible through LFV experiments.

7.2 The LFV axiflavor

We discuss next the possibility that the PQ symmetry that solves the strong CP problem is also responsible for explaining the smallness of the SM Yukawas, i.e., that the QCD axion is the *axiflavor*. This framework naturally results in a QCD axion with flavor-violating couplings.

The simplest scenarios of this kind arise when the $U(1)_{\text{PQ}}$ is responsible for explaining all the SM fermion mass hierarchies and mixings, along the lines of the Froggatt-Nielsen models [34, 125, 126], as in refs. [24, 33]. In these constructions the strongest bound on the axion decay constant always arises from $K^+ \rightarrow \pi^+ a$ constraints [24], since the axion coupling to sd is suppressed only by roughly the size of the Cabibbo angle, $V_{us} = \lambda \simeq 0.2$. Indeed, for the $U(1)$ axiflavor the LH quark charges, X_i^{qL} , are non-universal so that from eq. (7.2),

$$C_{sd}^V \sim V_{us} \frac{X_1^{qL} - X_2^{qL}}{2N} \sim V_{us}. \tag{7.17}$$

The sd axiflavor couplings, on the other hand, can be strongly suppressed in $U(2)_F$ flavor models [35, 147, 148]. In these classes of models the light generations form doublets of $U(2)_F$. The $U(1)_F$ factor acts as the PQ symmetry that gives rise to the QCD axion after spontaneous symmetry breaking. This scenario successfully explains the fermion mass hierarchies and mixings in terms of just two small parameters. In the model of ref. [35] both quark and lepton flavor violating couplings between the first two generations are equally suppressed because of the assumed structure that is compatible with $SU(5)$ unification. Here we present a variant of this model that leads instead to parametrically large μe couplings (their enhancement can be traced to large PMNS mixing angles). This model is a successful model of flavor and at the same time an example of a large class of flavored axion models where the LFV couplings are sizable while the FV couplings to quarks are suppressed.

	U_a^c	D_a^c	Q_a	U_3^c	D_3^c	Q_3	E_a^c	L_a	E_3^c	L_3	H	ϕ_a	χ
$SU(2)_F$	2	2	2	1	1	1	2	2	1	1	1	2	1
$U(1)_F$	1	1	1	0	1	0	1	1	0	-1	0	-1	-1

Table 4. The $U(2)_F$ quantum numbers of the SM fermions and the scalars H , ϕ_a and χ in the example LFV axiflavor model.

In the LFV axiflavor model the $U(2)_F = SU(2)_F \times U(1)_F$ quantum numbers of almost all the SM fermion are the same as in ref. [35]. In particular, all the fermions transform as $\mathbf{2} + \mathbf{1}$ under $SU(2)_F$. The only difference with respect to ref. [35] is that the $U(1)_F$ charge of the $SU(2)_F$ singlet left-handed lepton, L_3 , is -1 instead of 1. Table 4 summarizes the complete field content and the transformation properties under the $U(2)_F$ flavor group. In addition to the SM fermions, the electroweak doublets Q_i, L_i , and singlets U_i^c, D_i^c, E_i^c , $i = 1, 2, 3$ (in table 4, $a = 1, 2$), the model contains the SM Higgs doublet as well as two scalar spurions, ϕ_a and χ .

As in ref. [35], the breaking of the flavor symmetry is parametrized by two scalar spurions ϕ and χ , which transform under $U(2)_F$ as $\phi = \mathbf{2}_{-1}$ and $\chi = \mathbf{1}_{-1}$. These fields acquire the following flavor symmetry breaking vevs

$$\langle \phi \rangle = \begin{pmatrix} \epsilon_\phi \Lambda \\ 0 \end{pmatrix}, \quad \langle \chi \rangle = \epsilon_\chi \Lambda, \quad (7.18)$$

where Λ is, up to $\mathcal{O}(1)$ factors, the typical mass of the heavy states present in the full UV model (their exact structure is not important for our effective low energy discussion as they are integrated out). The values of the two small parameters in the two spurions are fixed by the fit to quark masses and mixings to be about $\epsilon_\phi \sim \lambda^2$ and $\epsilon_\chi \sim \lambda^3$.

Since the SM fermions are charged under $U(2)_F$, the Yukawa interactions between the SM fermions and the Higgs require insertions of the spurion fields in order to form invariants under $U(2)_F$. This leads to non-renormalizable interactions suppressed by appropriate powers of Λ . After replacing spurions with their vevs, eq. (7.18), the dependence on the heavy scale Λ drops out. The hierarchies in Yukawa matrices then arise from powers of the small parameters $\epsilon_{\phi, \chi}$, giving for the mass matrices

$$m_u \approx \frac{v}{\sqrt{2}} \begin{pmatrix} \lambda_{11}^u \epsilon_\phi^2 \epsilon_\chi^4 & \lambda_{12}^u \epsilon_\chi^2 & \lambda_{13}^u \epsilon_\phi \epsilon_\chi^2 \\ -\lambda_{12}^u \epsilon_\chi^2 & \lambda_{22}^u \epsilon_\phi^2 & \lambda_{23}^u \epsilon_\phi \\ \lambda_{31}^u \epsilon_\phi \epsilon_\chi^2 & \lambda_{32}^u \epsilon_\phi & \lambda_{33}^u \end{pmatrix}, \quad m_d \approx \frac{v}{\sqrt{2}} \begin{pmatrix} \lambda_{11}^d \epsilon_\phi^2 \epsilon_\chi^4 & \lambda_{12}^d \epsilon_\chi^2 & \lambda_{13}^d \epsilon_\phi \epsilon_\chi^3 \\ -\lambda_{12}^d \epsilon_\chi^2 & \lambda_{22}^d \epsilon_\phi^2 & \lambda_{23}^d \epsilon_\phi \epsilon_\chi \\ \lambda_{31}^d \epsilon_\phi \epsilon_\chi^2 & \lambda_{32}^d \epsilon_\phi & \lambda_{33}^d \epsilon_\chi \end{pmatrix}, \quad (7.19)$$

$$m_e \approx \frac{v}{\sqrt{2}} \begin{pmatrix} \lambda_{11}^e \epsilon_\phi^2 \epsilon_\chi^4 & \lambda_{12}^e \epsilon_\chi^2 & \lambda_{13}^e \epsilon_\phi \epsilon_\chi^2 \\ -\lambda_{12}^e \epsilon_\chi^2 & \lambda_{22}^e \epsilon_\phi^2 & \lambda_{23}^e \epsilon_\phi \\ \lambda_{31}^e \epsilon_\phi \epsilon_\chi & \lambda_{32}^e \epsilon_\phi \epsilon_\chi & \lambda_{33}^e \epsilon_\chi \end{pmatrix}, \quad m_\nu \approx \frac{v^2}{2M} \begin{pmatrix} \lambda_{11}^\nu \epsilon_\phi^2 \epsilon_\chi^4 & \lambda_{12}^\nu \epsilon_\phi^2 \epsilon_\chi^2 & \lambda_{13}^\nu \epsilon_\phi \epsilon_\chi \\ \lambda_{12}^\nu \epsilon_\phi^2 \epsilon_\chi^2 & \lambda_{22}^\nu \epsilon_\phi^2 & \lambda_{23}^\nu \epsilon_\phi \epsilon_\chi \\ \lambda_{13}^\nu \epsilon_\phi \epsilon_\chi & \lambda_{23}^\nu \epsilon_\phi \epsilon_\chi & \lambda_{33}^\nu \epsilon_\chi^2 \end{pmatrix}, \quad (7.20)$$

where we kept only the leading contributions in $\epsilon_{\phi, \chi}$. The parameters λ_{ij}^f are $\mathcal{O}(1)$ complex coefficients. The structure of the neutrino mass matrix m_ν was obtained under the assumption that the neutrinos are Majorana with the masses arising from the Weinberg operator with the UV suppression scale M . Notice that the 1-2 entry in m_ν of order ϵ_χ^2 vanishes due to anti-symmetrization, so that the first nonzero entry is of $\mathcal{O}(\epsilon_\phi^2 \epsilon_\chi^2)$.

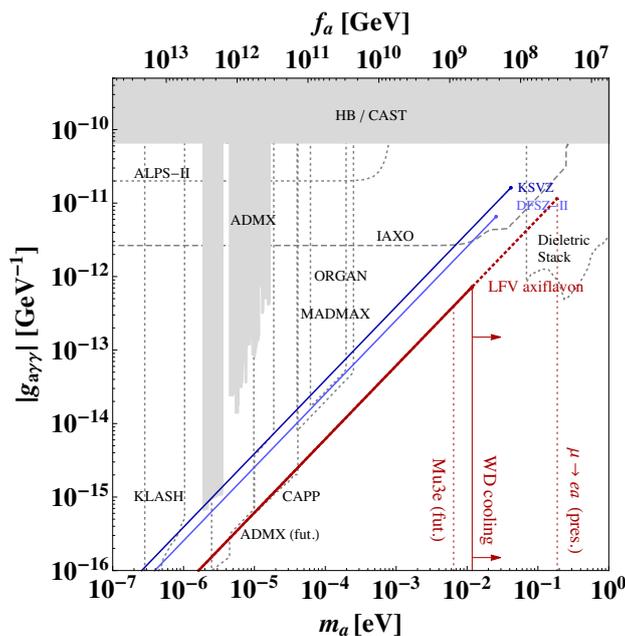


Figure 11. Present bounds and projected sensitivity for the LFV axiflavor model illustrated in section 7.2, with the LFV couplings given in eq. (7.30). The *red bold solid line* shows the predicted coupling to photons (as defined in eq. (7.16)) of the LFV axiflavor model. The *solid red vertical line* shows the upper bound on the axion mass from star cooling while the *dotted red vertical lines* the present bound and future sensitivity from LFV axion searches. The *grey shaded and grey dotted regions*, and the *blue and dark blue lines* are the same as in figure 10.

The quark sector of this model is identical to the one in ref. [35], and so are the unitary rotations that diagonalize the quark masses. Their parametric structure is given by,

$$V_L^u \sim V_R^u \sim \begin{pmatrix} 1 & \lambda & \lambda^7 \\ \lambda & 1 & \lambda^2 \\ \lambda^3 & \lambda^2 & 1 \end{pmatrix}, \quad V_L^d \sim \begin{pmatrix} 1 & \lambda & \lambda^3 \\ \lambda & 1 & \lambda^2 \\ \lambda^3 & \lambda^2 & 1 \end{pmatrix}, \quad V_R^d \sim \begin{pmatrix} 1 & \lambda & \lambda^5 \\ \lambda & 1 & 1 \\ \lambda & 1 & 1 \end{pmatrix}, \quad (7.21)$$

where $\lambda = V_{us} \sim 0.2$. The charged lepton mass matrix is parametrically the same as in ref. [35], and so are the parametric sizes of the unitary rotations that diagonalize it,

$$V_R^e \sim \begin{pmatrix} 1 & \lambda & \lambda^3 \\ \lambda & 1 & \lambda^2 \\ \lambda^3 & \lambda^2 & 1 \end{pmatrix}, \quad V_L^e \sim \begin{pmatrix} 1 & \lambda & \lambda^5 \\ \lambda & 1 & 1 \\ \lambda & 1 & 1 \end{pmatrix}. \quad (7.22)$$

It is easy to check that the neutrino sector can reproduce the PMNS matrix and neutrino mass differences for λ_{ij}^ν that are $\mathcal{O}(1)$ for normal neutrino mass ordering.

The PNGB corresponding to the $U(1)_F$ factor is the LFV axiflavor, i.e., it acts as the QCD axion that solves the strong CP problem (see ref. [35] for details). The anomaly coefficients controlling the couplings of the LFV axiflavor to gluons and to photons in eq. (7.4) are given by

$$N = 9/2, \quad E = 10. \quad (7.23)$$

The couplings of the LFV axiflavor to the SM fermions depend on the unitary rotations in eqs. (7.21), (7.22), as in eq. (7.2). Because of the $SU(2)_F$ structure of the LFV axiflavor model it is useful to introduce the matrices ($f = u, d, e; P = L, R$)¹²

$$\epsilon_{ij}^{fP} \equiv (V_P^f)^*_{3i} (V_P^f)_{3j}, \tag{7.24}$$

which satisfy

$$0 \leq \epsilon_{ii}^{fP} \leq 1, \quad \sum_i \epsilon_{ii}^{fP} = 1, \quad |\epsilon_{ij}^{fP}| = \sqrt{\epsilon_{ii}^{fP} \epsilon_{jj}^{fP}}. \tag{7.25}$$

They have the parametric structures

$$\epsilon_{ij}^{uL} \sim \epsilon_{ij}^{uR} \sim \epsilon_{ij}^{dL} \sim \epsilon_{ij}^{eR} \sim \begin{pmatrix} \lambda^6 & \lambda^5 & \lambda^3 \\ \lambda^5 & \lambda^4 & \lambda^2 \\ \lambda^3 & \lambda^2 & 1 \end{pmatrix}, \quad \epsilon_{ij}^{dR} \sim \epsilon_{ij}^{eL} \sim \begin{pmatrix} \lambda^2 & \lambda & \lambda \\ \lambda & 1 & 1 \\ \lambda & 1 & 1 \end{pmatrix}. \tag{7.26}$$

The axiflavor couplings to fermions are given in terms of charges and these parameters as

$$C_{f_i f_j}^V = \frac{X_{f_a^c} - X_{f_a}}{2N} \delta_{ij} + \frac{X_{f_3^c} - X_{f_a^c}}{2N} \epsilon_{ij}^{fR} - \frac{X_{f_3} - X_{f_a}}{2N} \epsilon_{ij}^{fL}, \tag{7.27}$$

$$C_{f_i f_j}^A = \frac{X_{f_a^c} + X_{f_a}}{2N} \delta_{ij} + \frac{X_{f_3^c} - X_{f_a^c}}{2N} \epsilon_{ij}^{fR} + \frac{X_{f_3} - X_{f_a}}{2N} \epsilon_{ij}^{fL}, \tag{7.28}$$

where $f = u, d, e$ denotes the fermion sector and $X_{f_a^c}, X_{f_a}, X_{f_3^c}, X_{f_3}$ are the $U(1)_F$ charges in table 4.

The sd couplings are strongly suppressed, $C_{sd}^{V,A} \sim \lambda^5/(2N)$, as a result of small rotations in the LH sector, $\epsilon_{L,12}^d \sim \lambda^5$, and the fact that RH rotation do not lead to off-diagonal terms because of universal charges in the RH sector, $X_{D_a^c} = X_{D_3^c}$. The RH contributions to the μe couplings are CKM-like suppressed, while the contribution from the LH rotations are large, since the corresponding charges are non-universal (in contrast to ref. [35]), $X_{L_a} \neq X_{L_3}$, giving $C_{\mu e}^V = -C_{\mu e}^A \sim \sqrt{m_e/m_\mu}/N$. The axiflavor couplings to nucleons and electrons are identical to ref. [35] and are to good approximation given by

$$C_{ee}^A = C_{uu}^A = C_{dd}^A = C_{cc}^A = C_{ss}^A \approx \frac{2}{9}, \quad C_{tt}^A \approx 0, \quad C_{bb}^A \approx \frac{1}{9}. \tag{7.29}$$

The $\mathcal{O}(1)$ coefficients in the rotation matrices can be fixed by performing an explicit fit to all the observables — the masses and mixings, including the neutrino sector. Using the same procedure and the SM inputs as in ref. [35], we find that a good fit is obtained by choosing $\epsilon_\phi = 0.023, \epsilon_\chi = 0.080, M = 4.8 \times 10^{11}$ GeV, with Yukawa couplings for the up quark sector $\lambda_{\{12,22,23,32,33\}}^u = \{-2.0, 1.0, -3.1, -1.1, -0.79\}$, and for the down quark sector $\lambda_{\{12,22,23,32,33\}}^d = \{1.3, 1.1, -0.76, 0.44, -0.85\}$, with the couplings $\lambda_{1,13,31}^u$ and $\lambda_{1,13,31}^d$ irrelevant because they give only subleading contributions to quark masses and mixings. The couplings for the lepton Yukawa matrices are

¹²Note the difference between ϵ_{ij}^{eP} defined here and in eq. (7.6). For this reason we restrict the use of ϵ_{ij}^{fP} symbols to the two respective sections.

$\lambda_{\{11,12,13,22,23,31,32,33\}}^e = \{1.0, 2.3, 1.0, 1.7, 0.40, -0.30, -1.3, 0.55\}$, while for the neutrinos they are $\lambda_{\{11,12,13,22,23,33\}}^V = \{-1.0, -1.0, 0.38, -1.2, 1.9, -0.69\}$. For this fit we find for the sd axiflavor coupling $C_{sd}^V = 1.6 \times 10^{-5}$, while the relevant couplings in the leptonic sector are

$$C_{\mu e}^V = -C_{\mu e}^A = 0.043, \quad C_e = 0.21, \quad C_{\tau e} = 0.029, \quad C_{\tau\mu} = 0.12, \quad (7.30)$$

where $C_{\ell_i\ell_j} \equiv \sqrt{|C_{\ell_i\ell_j}^V|^2 + |C_{\ell_i\ell_j}^A|^2}$. We use the above benchmark values for axiflavor coupling to derive the sensitivities of different observables to axiflavor in figure 11. Other phenomenologically viable choices of parameters that differ by “ $\mathcal{O}(1)$ ” factors can also give a good fit to the SM masses and mixings, so the constraints obtained in our benchmark should be viewed only as indicative, with $\mathcal{O}(1)$ variations, when this larger class of axiflavor parameters is considered.

The red line in figure 11 shows the predicted coupling to photons, $g_{a\gamma\gamma}$, which for the LFV axiflavor is given by eq. (7.16) with $E/N = 20/9$, as a function of axiflavor mass, eq. (7.1). The bound on f_a from WD cooling is denoted with a vertical solid red line. The next less stringent bound comes from the $\mu \rightarrow ea$ search by TWIST (dotted vertical red line). Figure 11 shows that the Mu3e future reach (dotted vertical red line) will exceed the WD cooling constraints. In order not to clutter figure 11, we do not show the less constraining present bound (future sensitivity) on LFV axiflavor from $K^+ \rightarrow \pi^+ a$ which is $f \gtrsim 5.4 \times 10^6$ GeV (1.6×10^7 GeV) [23]. The expected reach from Belle II is $f_a \gtrsim 3(1) \times 10^6$ GeV from $\tau \rightarrow \mu a$ ($\tau \rightarrow ea$) searches, which is outside the range plotted in figure 11. The other constraints, shown in grey with future sensitivities denoted with grey dashed lines, are as in figure 10. Clearly, there is significant parameter space where the LFV axiflavor can be discovered in LFV experiments, especially considering the potential astrophysical uncertainties in the WD cooling bounds.

7.3 The leptonic familon

Our next example of an LFV ALP is the familon, i.e., the PNCB arising from the spontaneous breaking of a global horizontal symmetry, which we take to be the Froggatt-Nielsen (FN) flavor symmetry, $U(1)_{\text{FN}}$ [34]. We consider the case where the $U(1)_{\text{FN}}$ only acts on the leptonic sector so that the LFV ALP is the *leptonic familon*. Unlike the previous two examples, the leptonic familon does not solve the strong CP problem and, as a consequence, its mass is not determined by the QCD anomaly. The mass of the leptonic familon is therefore taken to be a free parameter, yet still small enough that it can be produced in tau or muon decays. The predictive power of the model is limited to the parametric prediction of the LFV coupling, which are related to the neutrinos mass texture.

The couplings of the leptonic familon to the SM leptons are determined by the positive $U(1)_{\text{FN}}$ charges $[L]_i$ and $[e]_i$ carried by the lepton doublets L_i and singlets e_i^c , respectively. The $U(1)_{\text{FN}}$ symmetry is broken by the vev f_a of a scalar field Φ with charge $[\Phi] = -1$,

$$\Phi = \frac{f_a + \phi}{\sqrt{2}} e^{ia/f_a}, \quad (7.31)$$

where a is the familon, while ϕ is the radial mode with the mass $\mathcal{O}(f_a)$. Integrating out all the heavy fields the charged lepton Yukawa couplings and the Majorana neutrino mass matrix are given by (as determined through spurion analysis)

$$y_{ij}^e = a_{ij}^e \left(\frac{\langle \Phi^* \rangle}{M} \right)^{[L]_i + [e]_j}, \quad m_{ij}^\nu = \kappa_{ij}^\nu \frac{v^2}{\Lambda_N} \left(\frac{\langle \Phi^* \rangle}{M} \right)^{[L]_i + [L]_j}, \quad (7.32)$$

where a^e and κ^ν are assumed to be flavour-anarchical matrices of $\mathcal{O}(1)$ coefficients, M is a cut-off scale with

$$\epsilon \equiv \frac{\langle \Phi^* \rangle}{M} < 1,$$

controlling the hierarchies among charged lepton masses, while Λ_N is the lepton-number breaking scale suppressing the dimension 5 Weinberg operator.¹³ An anarchical PMNS matrix featuring $\mathcal{O}(1)$ mixing angles [149] can be achieved by taking equal charges for the lepton doublets,

$$([L]_1, [L]_2, [L]_3) = (L, L, L), \quad [\text{Pure Anarchy}]. \quad (7.33)$$

Good fits to the neutrino oscillation data can also be obtained for mildly hierarchical charges (at the price of somewhat larger values of $\epsilon \sim 0.3 - 0.4$) [150, 151],

$$([L]_1, [L]_2, [L]_3) = (L + 1, L, L), \quad [\mu\tau \text{ Anarchy}], \quad (7.34)$$

$$([L]_1, [L]_2, [L]_3) = (L + 2, L + 1, L), \quad [\text{Hierarchy}]. \quad (7.35)$$

The hierarchy of charged leptons is then reproduced by a suitable choice of the charges of the RH leptons. Up to some freedom due to uncertainties relative to the expansion parameter ϵ and more importantly to the $\mathcal{O}(1)$ coefficients a_{ij}^e , a successful charge assignment is the following:

$$([e]_1, [e]_2, [e]_3) = (5 - L, 3 - L, 2 - L), \quad \epsilon = 0.1 \quad [\text{Pure Anarchy}], \quad (7.36)$$

$$([e]_1, [e]_2, [e]_3) = (9 - L, 5 - L, 3 - L), \quad \epsilon = 0.3 \quad [\mu\tau \text{ Anarchy}], \quad (7.37)$$

$$([e]_1, [e]_2, [e]_3) = (11 - L, 6 - L, 4 - L), \quad \epsilon = 0.4 \quad [\text{Hierarchy}]. \quad (7.38)$$

The couplings of the familon to the leptons are given by,

$$\mathcal{L} \supset \frac{\partial^\mu a}{2f_a} \left(C_{\ell_i \ell_j}^V \bar{\ell}_i \gamma_\mu \ell_j + C_{\ell_i \ell_j}^A \bar{\ell}_i \gamma_\mu \gamma_5 \ell_j \right), \quad (7.39)$$

where

$$C^{V/A} = V_R^{e\dagger} X_R^e V_R^e \pm V_L^{e\dagger} X_L^e V_L^e. \quad (7.40)$$

The unitary rotation matrices are defined as $V_L^{e\dagger} y^e V_R^e \equiv y_{\text{diag}}^e$, while $(X_L^e)_{ij} = [L]_i \delta_{ij}$, $(X_R^e)_{ij} = -[e]_i \delta_{ij}$ are diagonal matrices of the FN charges. Up to $\mathcal{O}(1)$ coefficients the LH and RH rotations are given by $(V_{L/R}^e)_{ij} \approx \epsilon^{|[L/e]_i - [L/e]_j|}$.

¹³Introducing instead the RH neutrinos such that the neutrinos have Dirac masses would not change the following discussion and the resulting PMNS matrix.

The three charge assignments (7.33)–(7.35) lead to qualitatively different values of LFV familon couplings, eq. (7.40). In the purely anarchical model the $[L]_i$ are universal, cf. eq. (7.33), so that the off-diagonal couplings in eq. (7.40) are entirely due to the RH charges. The RH rotations are small, of the order of the ratios of lepton masses, $(V_R^e)_{ij} \approx (V_R^e)_{ji} \approx m_{\ell_i}/m_{\ell_j}$ ($i < j$), and thus the LFV couplings are suppressed,

$$C_{\ell_i \ell_j}^V = C_{\ell_i \ell_j}^A \approx ([e]_i - [e]_j) \frac{m_{\ell_i}}{m_{\ell_j}} \quad (i < j) \quad [\text{Pure Anarchy}]. \quad (7.41)$$

The purely anarchical leptonic familon couples to the $V + A$ current. It is thus subject to constraints from the old Jodidio et al. experiment and can be an important target for the proposed MEGII-fwd setup, despite the severe suppression of the LFV couplings, see figure 12 (left).

For the hierarchical model the $[L]_i$ are non-universal, eq. (7.35), and the LH rotations give the dominant contributions to the off-diagonal familon couplings in eq. (7.40). Even so, the $V + A$ couplings induced via RH rotations remain non-negligible, and are even enhanced compared to the anarchical case, $(V_R^e)_{ij} \approx (V_R^e)_{ji} \approx \epsilon^{[e]_i - [e]_j} \approx (m_{\ell_i}/m_{\ell_j})/\epsilon^{[L]_i - [L]_j}$ ($i < j$). The LFV couplings of the hierarchical leptonic familon are thus, from eq. (7.40),

$$\begin{aligned} C_{\ell_i \ell_j}^V &\approx ([e]_i - [e]_j) \frac{1}{\epsilon^{[L]_i - [L]_j}} \frac{m_{\ell_i}}{m_{\ell_j}} - ([L]_i - [L]_j) \epsilon^{[L]_i - [L]_j}, \\ C_{\ell_i \ell_j}^A &\approx ([e]_i - [e]_j) \frac{1}{\epsilon^{[L]_i - [L]_j}} \frac{m_{\ell_i}}{m_{\ell_j}} + ([L]_i - [L]_j) \epsilon^{[L]_i - [L]_j}, \end{aligned} \quad (i < j) \quad [\text{Hierarchy}]. \quad (7.42)$$

Because the hierarchical familon mostly couples to the $V - A$ current the dominant constraint comes from the bound on $\mu \rightarrow e a$ due to TWIST, eq. (3.17), even though this is otherwise the weakest bound on $\text{BR}(\mu \rightarrow e a)$. In the future, MEGII-fwd can improve the reach on the hierarchical familon beyond the present bounds, because of the non-vanishing $V + A$ contributions, cf. figure 12 (right).

The $\mu\tau$ Anarchy scenario is in between the above two cases. There is no LH mixing in the $2 - 3$ sector, while there is one in the case of $1 - 2$ and $1 - 3$ transitions. The RH mixings are suppressed, but with $1 - 2$ and $1 - 3$ mixing enhanced by $1/\epsilon$ over the pure anarchical case.

The familon coupling to electrons, relevant for assessing the star cooling bounds, is given by

$$C_{ee}^A \approx -([e]_1 + [L]_1). \quad (7.43)$$

Finally, the coupling to photons is controlled by the anomaly contribution E_{UV} , eq. (2.9), which can be calculated in terms of FN charges in eqs. (7.33)–(7.38) as

$$E_{UV} = \sum_i ([e]_i + [L]_i), \quad (7.44)$$

such that, for the three illustrative charge assignments (Anarchy, $\mu\tau$ Anarchy, Hierarchy), one gets respectively $E_{UV} = \{10, 18, 24\}$.

Figure 12 summarizes the present and expected future bounds on the FN breaking scale f_a for the purely anarchical (left panel) and for the hierarchical leptonic familon

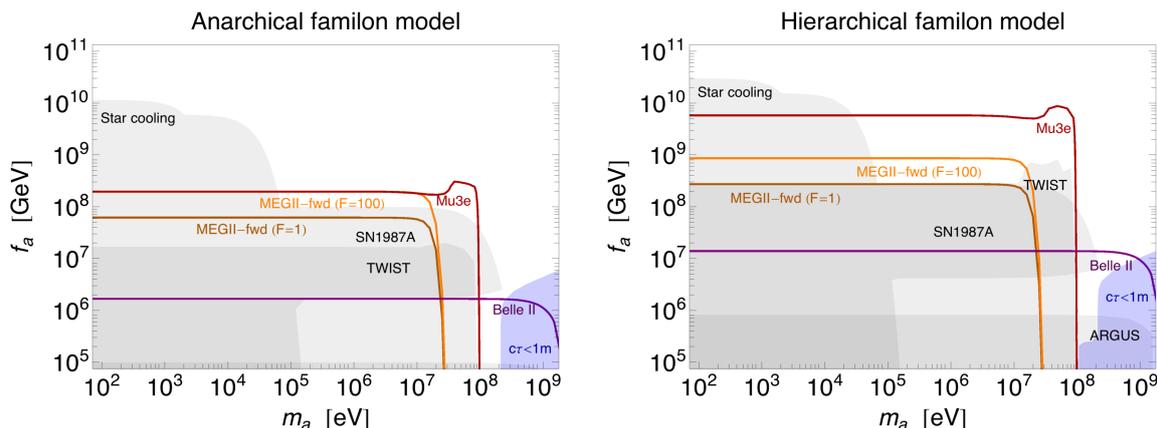


Figure 12. Present bounds and future sensitivities on the decay constant f_a (in GeV) vs. the mass m_a for two familon models. The *blue shaded* regions indicate where the familon decays promptly into SM states. The *grey shaded* regions are currently excluded by a combination of $\mu^+ \rightarrow e^+ a$ experiments and star cooling bounds on the ALP-electron coupling. We show projected sensitivities of MEGII-fwd in *orange*, Mu3e in *red*, Belle II in *purple*. *Left:* Anarchical model, cf. eq. (7.41), *Right:* Hierarchical model, cf. eq. (7.35).

(right panel). For the anarchical model, we set $[e]_1 - [e]_2 = 2$ and $[e]_2 - [e]_3 = 1$ in eq. (7.41), and $[e]_1 + [L]_1 = 8$ in eq. (7.43). For the hierarchical case we estimated the LH rotations appearing in eq. (7.42) employing $\epsilon = 0.4$ as in refs. [150, 151], while the RH couplings are obtained by setting $[e]_1 - [e]_2 = 5$ and $[e]_2 - [e]_3 = 2$. In this case the diagonal coupling to electrons in eq. (7.43) follows from $[e]_1 + [L]_1 = 13$. In all the cases the $\mathcal{O}(1)$ coefficients were set to 1 exactly. The blue area in figure 12 indicates where the familon proper decay length is shorter than 1 m (due to a sizeable width to leptons, $a \rightarrow \ell_i \ell_j$), so that searches for $\ell_i \rightarrow \ell_j +$ invisible lose sensitivity. Figure 12 shows that in both representative models all future searches for LFV processes performed by MEG II, Mu3e and Belle II can probe well into the yet unexplored parameter space.

As far as DM is concerned, the leptonic familon resembles very much the generic ALP DM discussion in section 8. We refer to that section for an extensive discussion of the implications of LFV searches on the ALP DM parameter space.

7.4 The majoron

The majoron [128, 129] is the PNGB due to spontaneous breaking of the lepton number. A natural context where this kind of ALP arises are the seesaw models, where the breaking scale of lepton number is associated with the mass scale of heavy new fields. In type-I seesaw models at least two singlet fermions, the right-handed neutrinos N_i , are added to the SM (for concreteness we will assume below two RH neutrinos). The RH neutrinos couple to LH leptons via the Yukawa coupling matrix, y_N , while their masses are described by the Majorana mass matrix M_N ,

$$\mathcal{L} = \mathcal{L}_{\text{SM}} + i\bar{N}\not{\partial}N - \left(y_N \bar{L} N \tilde{H} + \frac{1}{2} \bar{N}^c M_N N + \text{h.c.} \right). \quad (7.45)$$

In the seesaw limit, $M_N \gg m_D \equiv y_N v$, the RH neutrinos are heavy and can be integrated out, while the light neutrinos are predominantly part of the SM doublets, L_i , with the Majorana mass matrix given by

$$m_\nu = -m_D M_N^{-1} m_D^T. \quad (7.46)$$

The majoron J arises when the mass matrix M_N , which breaks the lepton number by two units, is generated dynamically by the vev of a new SM singlet scalar field σ . In this case the RH neutrino mass matrix in eq. (7.45) is replaced by the Yukawa couplings of RH neutrinos to the scalar σ , i.e., $M_N \rightarrow \lambda_N \sigma$ in eq. (7.45), where

$$\sigma = \frac{f_N + \hat{\sigma}}{\sqrt{2}} e^{iJ/f_N}, \quad (7.47)$$

so that the RH neutrino mass matrix is given in terms of the matrix λ_N as

$$M_N = \frac{\lambda_N f_N}{\sqrt{2}}. \quad (7.48)$$

The radial mode $\hat{\sigma}$ is heavy and can be integrated out, while the majoron J is a PNCB and is light (we take its mass to be a free parameter).

The majoron couples at tree level to neutrinos through the $\bar{N}^c \lambda_N \sigma N$ Yukawa terms. These Yukawa interactions then induce couplings of J to charged leptons and quarks at loop-level [128], see ref. [152] for complete expressions. Here we are interested only in the seesaw limit of these general expressions, which we match onto the effective Lagrangian (2.1) upon identifying $a \rightarrow J, f_a \rightarrow f_N$. Using the results of refs. [153, 154], we find for the majoron couplings to quarks and leptons, respectively,

$$C_{q_i q_j}^V = 0, \quad C_{q_i q_j}^A = -\frac{T_3^q}{16\pi^2} \delta_{ij} \text{Tr} \left(y_N y_N^\dagger \right), \quad (7.49)$$

$$C_{e_i e_j}^V = \frac{1}{16\pi^2} \left(y_N y_N^\dagger \right)_{ij}, \quad C_{e_i e_j}^A = \frac{1}{16\pi^2} \left[\frac{\delta_{ij}}{2} \text{Tr} \left(y_N y_N^\dagger \right) - \left(y_N y_N^\dagger \right)_{ij} \right], \quad (7.50)$$

where $T_3^{u,d} = \pm 1/2$. Note that $F_{\mu e}^V = -F_{\mu e}^A$, i.e., the LFV couplings of type-I seesaw majoron have the $V - A$ form (the couplings to the $V + A$ leptonic current are flavor conserving). The TWIST experiment is sensitive to such a majoron, while the more stringent bounds from the 1986 experiment by Jodidio et al. do not apply, see section 3.1 for details.

In contrast to the other ALP scenarios discussed above, in the case of a majoron, couplings to neutrinos are particularly relevant to assess the stability of the particle, hence whether it is a viable DM candidate. This is a consequence of the suppressed coupling to photons of the majoron, which in fact decays preferably to neutrinos when its mass is below the $2m_e$ threshold. In the seesaw limit, the coupling to light neutrinos is diagonal and the decay width reads [153, 154]

$$\Gamma(J \rightarrow \nu_i \bar{\nu}_i) = \frac{m_J}{16\pi f_N^2} m_i^2 \sqrt{1 - \frac{4m_i^2}{m_J^2}}, \quad (7.51)$$

where m_i are the light neutrino mass eigenvalues.

Another crucial issue is whether the experiments are able to probe scales that are interesting for the neutrino mass generation. We can distinguish two limits:

- In the *standard seesaw setup*, sizeable entries of the Yukawa matrix y_N are only compatible with the observed neutrino masses for an ultra-high seesaw scale. For instance, let us consider the case where elements of the Yukawa matrix are all of similar size, without any special structure, $|(y_N)_{ij}| \sim y$ (a hierarchical structure would not qualitatively change the argument). The light neutrino masses are thus of the size $m_\nu \sim y^2 v^2 / M_N$, where $v = 246$ GeV is the electroweak symmetry breaking scale. The effective scale suppressing the LFV $\ell_i \rightarrow \ell_j J$ decays is then given by $F \sim f_N / C \sim 16\pi^2 f_N / y^2 \sim 16\pi^2 v^2 / m_\nu \gtrsim 10^{16}$ GeV. The standard set-up therefore cannot be probed by the present nor any of the planned LFV experiments.
- In the *low-scale seesaw setup*, the neutrino masses are additionally suppressed, such that large couplings in y_N and a lower seesaw scale are compatible with the observed light neutrino masses. Indeed, since the neutrino masses $m_\nu \propto y_N M_N^{-1} y_N^T$ transform non-trivially under the lepton number (in contrast to $y_N y_N^\dagger$), it is possible that light neutrino masses are parametrically suppressed in the presence of an approximate generalized lepton number. Such scenarios, usually referred to as the “TeV scale seesaw mechanism”, have been extensively studied in the literature [41, 130–138]. In the following we use the results of ref. [41] to construct a concrete example of a majoron model with parametrically suppressed neutrino masses (and thus with enhanced majoron couplings to the SM leptons).

In the simplest low-energy seesaw model one considers two right-handed neutrinos $N_{1,2}$ with a 2×2 Dirac mass matrix, M_N , and 3×2 Dirac Yukawa couplings, y_N ,

$$M_N = \begin{pmatrix} 0 & M \\ M & 0 \end{pmatrix} = \frac{\lambda}{\sqrt{2}} \begin{pmatrix} 0 & f_N \\ f_N & 0 \end{pmatrix}, \quad y_N = \begin{pmatrix} y_{e1} & y_{e2} \\ y_{\mu 1} & y_{\mu 2} \\ y_{\tau 1} & y_{\tau 2} \end{pmatrix}, \quad (7.52)$$

where $\lambda \equiv \sqrt{2}M/f_N$ is a real free parameter. In the $y_{\ell 1} \rightarrow 0$ limit the model has a global U(1) symmetry, $M_N \rightarrow P M_N P$, $y_N \rightarrow e^{i\alpha} y_N P$ with $P = \text{diag}\{e^{i\alpha}, e^{-i\alpha}\}$. The majoron couplings, which are proportional to $y_N y_N^\dagger$, are invariant under this symmetry, while the neutrino masses are not, $m_\nu \rightarrow e^{2i\alpha} m_\nu$. This means that the neutrino masses are proportional to symmetry breaking parameters, $y_{\ell 1}$, which, if small, additionally suppress the neutrino masses compared to the majoron couplings.

Working in a basis where the charged lepton matrix is diagonal, we can adjust the input parameters, M and $y_{\ell i}$, such that all neutrino observables (2 mass differences + 3 mixing angles) are at the central experimental values. This leaves two free parameters, which we choose to be M , the mass scale of RH neutrinos, and the largest eigenvalue of the Dirac Yukawa matrix, $y = \max[\text{eig}(y_N y_N^\dagger)]$. Using the results of ref. [41], we obtain for the Normal Ordering (NO) in the seesaw limit,

$$y_N y_N^\dagger \approx y^2 \frac{m_3}{m_2 + m_3} A_i^* A_j, \quad \text{where } A_i = U_{i3} + iU_{i2} \sqrt{m_2/m_3}, \quad (7.53)$$

with $m_2 = (\Delta m_{21}^2)^{1/2}$ and $m_3 = (\Delta m_{31}^2)^{1/2}$ the light neutrino masses (the lightest neutrino mass is $m_1 = 0$ as we introduced only two RH neutrinos). The U_{ij} are the elements of the PMNS matrix, all of which are experimentally observable, while $y \approx (\text{Tr}(y_N y_N^\dagger))^{1/2}$ is a free parameter, only bounded by perturbativity, $y \lesssim 4$. The result for Inverted Ordering (IO) is obtained from (7.53) by replacing $m_3 \rightarrow m_2 = (-\Delta m_{32}^2)^{1/2}$, $m_2 \rightarrow m_1 = (-\Delta m_{21}^2 - \Delta m_{32}^2)^{1/2}$ and $U_{i3} \rightarrow U_{i2}, U_{i2} \rightarrow U_{i1}$.

For the numerical analysis we use the latest global neutrino oscillation fit results [155, 156], and set in eq. (7.53) the mass differences, the mixing angles and the Dirac CP phase in PMNS matrix to their central experimental values. This gives the 3×3 Hermitian matrix $y_N y_N^\dagger$ that still depends on y and one Majorana phases, α_m . For simplicity, we set the latter to zero, $\alpha_m = 0$. The effective suppression scales for the majoron couplings are in the NO case given by (similar results are obtained for IO)

$$F_{ee}^A = \frac{1.1 \times 10^{10} \text{ GeV}}{\lambda y^2} \left(\frac{M}{10^7 \text{ GeV}} \right), \quad F_{\mu e} = \frac{1.4 \times 10^{10} \text{ GeV}}{\lambda y^2} \left(\frac{M}{10^7 \text{ GeV}} \right), \quad (7.54)$$

$$F_{\tau e} = \frac{1.6 \times 10^{10} \text{ GeV}}{\lambda y^2} \left(\frac{M}{10^7 \text{ GeV}} \right), \quad F_{\tau\mu} = \frac{0.71 \times 10^{10} \text{ GeV}}{\lambda y^2} \left(\frac{M}{10^7 \text{ GeV}} \right), \quad (7.55)$$

where for the flavor-violating cases we quote the bound on combined A, V effective scale, as defined in eq. (2.2). For completeness we also show the results for the flavor diagonal couplings to muons and taus, although they are at the moment of little phenomenological relevance,

$$F_{\mu\mu}^A = -\frac{2.7 \times 10^{11} \text{ GeV}}{\lambda y^2} \left(\frac{M}{10^7 \text{ GeV}} \right), \quad F_{\tau\tau}^A = \frac{3.7 \times 10^{10} \text{ GeV}}{\lambda y^2} \left(\frac{M}{10^7 \text{ GeV}} \right). \quad (7.56)$$

The coupling to muons is comparatively suppressed due to an accidental cancellation between the two contributions to $C_{\mu\mu}^A$ in eq. (7.50).

The majoron also couples to nucleons via its couplings to quarks (7.49). These couplings do not depend on PMNS elements and are given by

$$F_N \approx \frac{0.88 \times 10^{10} \text{ GeV}}{\lambda y^2} \left(\frac{M}{10^7 \text{ GeV}} \right). \quad (7.57)$$

In figure 13, we summarize the current status and future prospect of the constraints on the parameter space of the low-energy seesaw majoron model described above. Besides the present and future bounds from LFV experiments and the astrophysical limits discussed in section 6.1 involving the coupling of the majoron to electrons, we display as a yellow-shaded area the region excluded by SN1987A, according to the study of ref. [157], due to the nucleon-majoron coupling of eq. (7.57).

We also show as a dashed blue line where the majoron lifetime equals the lifetime of the universe. Anywhere to the left of this line, the majoron is a viable DM candidate. In fact, there is no strong constraint from observations of the extragalactic background light, since coupling of the majoron to photons are suppressed. The two representative lines explained in section 6.2 show where the misalignment mechanism accounts for the whole DM abundance today with $\theta_0 \sim 1$. Below these lines the majoron is a sub-component

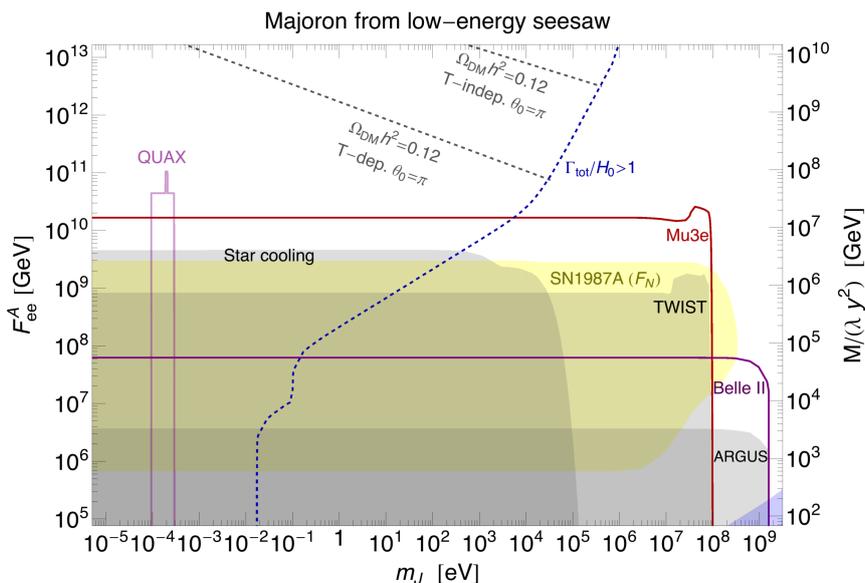


Figure 13. Present bounds and future sensitivities on the coupling to electrons F_{ee}^A (in GeV) and the RH neutrino mass scale M (normalized by the combination of free parameters λy^2) vs. the majoron mass m_J (in eV) for the low-energy seesaw majoron model. On the left of *dashed blue* line the majoron is stable on time scale of our Universe and could account for the DM abundance. The two *dashed gray lines* are examples of misalignment production discussed in section 6.2. The *gray shaded* region is excluded by star cooling bounds on the majoron-electron coupling and $\mu^+ \rightarrow e^+ a$ bounds. The *yellow shaded* is excluded by SN1987A. We show in *red* the future reach of Mu3e, in *purple* the reach of Belle II and in *light purple* the reach of QUAX which requires the majoron to be DM.

of the total DM abundance unless new dynamical mechanism or a tuning of the initial condition is put into place. Notice that here we ignore further mechanisms of production that could arise from Higgs portal-type of couplings (see ref. [158] for a discussion).

Finally, the light purple line denotes the expected reach of the QUAX experiment [102, 103, 121], which is directly sensitive to the coupling of the majoron to electrons. The shown sensitivity is as assessed in ref. [104], under the assumption that the majoron is responsible for 100% of the observed DM abundance.

Figure 13 shows the remarkable reach in terms of the effective scale F_{ee}^A and the RH mass neutrino mass scale M of the Mu3e experiment, which, by searching for the $\mu \rightarrow eJ$ decay, will be able to probe for the first time the uncharted territory beyond the star cooling and SN1987A bounds. In particular, Mu3e will be able to probe that part of the parameter space where the majoron can be a substantial component of DM.

Stringent constraints on low-energy seesaw models arise also from LFV decays, $\mu \rightarrow e\gamma$, eee , and $\mu \rightarrow e$ conversion, mediated by the heavy neutrinos at the scale M , if this is not much above the TeV scale. Following again ref. [41], one for instance finds for the branching ratio of $\mu \rightarrow e\gamma$

$$\text{BR}(\mu \rightarrow e\gamma) = \frac{3\alpha_{em}}{32\pi} \frac{|(m_D m_D^\dagger)_{12}|^2}{M^4} |H(X)|^2, \quad (7.58)$$

where $X = M^2/M_W^2$ and

$$H(x) = \frac{x(1 - 6x + 3x^2 + 2x^3 - 6x^2 \log x)}{(1 - x)^4} = \begin{cases} x & x \ll 1 \\ 2 + (11 - 6 \log x)/x & x \gg 1 \end{cases}. \quad (7.59)$$

This gives, in the $M \gg M_W$ limit,

$$\text{BR}(\mu \rightarrow e\gamma) = \frac{2.4 \times 10^{-13}}{\lambda^2} \left(\frac{10^3 \text{ GeV}}{M} \right)^2 \left(\frac{\text{BR}(\mu \rightarrow ea)}{5.8 \times 10^{-5}} \right). \quad (7.60)$$

This should be compared to the current 90% CL limit set by the MEG experiment, $\text{BR}(\mu \rightarrow e\gamma) < 4.2 \times 10^{-13}$ [77]. The LFV decays to majorons therefore typically provide stronger bounds on the seesaw scale M than the LFV processes, $\mu \rightarrow e\gamma$, $\mu \rightarrow 3e$, and $\mu N \rightarrow eN$, unless either $M \approx f_N \approx 1 \text{ TeV}$, or $\lambda \ll 1$, such that the mass scale of the RH neutrinos is much lower than the L -breaking scale $M \ll f_N$. Apart from these limits, $\mu \rightarrow eJ$ tends to provide the dominant constraint, as a consequence of different scalings with the heavy scale (the rate of $\mu \rightarrow eJ$ is $\propto M^{-2}$, while the rates of $\mu \rightarrow e\gamma$, $\mu \rightarrow 3e$, and $\mu N \rightarrow eN$, are $\propto M^{-4}$). See also the detailed discussion in ref. [154].

Finally, we comment on the chiral structure of the LFV majoron couplings which in our minimal implementation are entirely of $V - A$, given that LFV is mediated by the W -loops. This will not be the case in a two Higgs doublet model, where the majoron LFV couplings can be induced by loops of the charged Higgs. The $V + A$ LFV couplings are then going to be suppressed by $\sim y_H^2 v^2 / m_{H^+}^2$, where m_{H^+} is the mass of the charged Higgs and y_H its coupling to leptons. Given the strong bounds on H^+ coming from indirect searches in $b \rightarrow s\gamma$ [159] and direct collider searches [160] we expect the right-handed LFV couplings of the majoron to be generically suppressed with respect to the left-handed ones. It would be interesting to explore further the robustness of this statement in models where a light charged Higgs is still unconstrained [161].

8 Conclusions

Generically, axion like particles (ALPs) can have flavor violating couplings to the SM fermions. In this manuscript we explored the phenomenological consequences of such couplings; in the first part we took a model-independent perspective and summarized present constraints and future sensitivities on generic lepton flavor violating couplings of light ALPs, see figure 1. Here an important feature is our proposal for a new experimental set-up at MEG II, *MEGII-fwd*, which consists of a forward calorimeter to be installed in front of the MEG II beamline, see section 3.2 for details. The major benefit of such an experimental set-up is that the irreducible SM background from $\mu^+ \rightarrow e^+ \nu \bar{\nu}$ is reduced at the highest positron momentum in the forward region. Since the SM decay amplitude is controlled by left-handed couplings, it vanishes for an exactly forward positron of $p_{e^+} = m_\mu/2$ if produced from a muon that is completely negatively polarized. MEGII-fwd can be used to search for an effectively massless ALP produced in $\mu^+ \rightarrow e^+ a$ with some amount of right-handed LFV coupling to the SM leptons. The signal will appear as a sharp

line in positron energy distribution, in contrast to the smoothly falling SM background. The final reach of MEGII-fwd depends on how well depolarization effects can be controlled, on the positron momentum resolution, and on whether or not magnetic focusing is applied in order to increase the positron luminosity in the forward direction. Assuming realistic estimates for these parameters, we expect that a two week run at MEG II in MEGII-fwd configuration will allow to explore new ALP parameter space well before Mu3e and beyond the current astrophysical limits from star cooling.

Exploring this new region of parameter space will provide new insights on the couplings of an ALP DM produced non-thermally in the early Universe as discussed in section 6.2. A possible signal in an LFV experiment could be cross-correlated with the future experimental campaigns of axion haloscopes and future intensity mapping searches. This would give us more information about the light ALP mass beyond the expected resolution of MEGII-fwd.

In the second part of this paper we discussed several UV models where ALPs with flavor-violating couplings to leptons arise when addressing various shortcomings of the SM, such as the Strong CP Problem or the Flavor Puzzle, see section 7. These models explicitly demonstrate that the LFV couplings, which will be tested by future laboratory experiments, are correlated to the flavor-diagonal lepton couplings, which are strongly constrained by star cooling. Despite these constraints we find encouraging prospects for the expected sensitivity of MEGII-fwd to test these theoretically well-motivated models. We believe that this warrants a more systematic exploration of the feasibility of MEGII-fwd, including proper detector simulations. In particular, it would be interesting to investigate in more detail the required rearrangement of the MEGII magnetic field and its interplay with the achievable momentum resolution. We hope that this proposal will be explored further by the experimental collaboration. The MEGII-fwd proposal should also be carefully compared with other possibilities of improving the MEGII reach on LFV decays, for example the possibility of a dedicated trigger for $\mu^+ \rightarrow e^+ a \gamma$ decays as discussed in section 4.2. We hope to return to these issues in the future.

Beyond the new proposed experimental ideas, the paper includes several novel theoretical results. First of all, in section 6.1 we derived the astrophysical bounds from stellar cooling on leptonic couplings of an ALP with arbitrary mass. While the bounds on couplings of massless ALPs to electrons were readily available in the literature, this was not the case, to the best of our knowledge, for massive ALPs coupling to ee , as well as for the couplings to μe and $\mu\mu$ (recently, ref. [45] treated the case of massless ALP coupling to muons). Second, we showed how the chiral structure of the LFV couplings emerges in various explicit (and mostly novel) models, in which the presence of a light ALP is motivated by addressing the strong CP Problem (the LFV QCD Axion and the LFV Axiflavor), the hierarchical structure of SM fermion masses (the LFV Axiflavor and the Leptonic FAMILON) or the origin of neutrino masses (the Majoron). For the LFV Axiflavor and the Majoron the ALP couples mainly to left-handed leptons, and thus improvements in the reach on $\mu^+ \rightarrow e^+ \gamma a$ at MEGII (or ultimately the inclusive on-line bump hunt proposed at Mu3e) will be needed in order to test these scenarios. On the other hand, for the case of LFV QCD axion and the leptonic FAMILON the right-handed couplings are sufficiently large to allow MEGII-fwd to be the first experiment to discover the flavor violating ALP.

To conclude, we hope that this work may boost a renovated interest in the experimental possibilities at near future LFV experiments. The unprecedented luminosity of these facilities has been conceived to probe extremely rare lepton decays mediated by heavy new physics. The same data could also be used to probe decays with an (invisible) light new particle in the final state. We believe that we merely started to scratch the surface of the possible experimental improvements and theoretical motivations to search for light new physics in rare lepton flavor violating decays.

Acknowledgments

We thank Ryan Bayes, Marco Francesconi, Art Olin, Angela Papa, Ann-Kathrine Perrevoort, Francesco Renga, Andre Schöning and Giovanni Signorelli for many useful discussions and precious explanations on the different experimental setups. We thank Simon Knapen and Gustavo Marques-Tavares for discussions. We especially thank Angela Papa and Giovanni Signorelli for extensive feedback on our work and Matthias Linster for providing the numerical fit used in section 7.2. We thank Galileo Galilei Institute for kind hospitality and for providing the collaborative atmosphere that initiated this work. The work of RZ is supported by project C3b of the DFG-funded Collaborative Research Center TRR 257, “Particle Physics Phenomenology after the Higgs Discovery”. JZ acknowledges support in part by the DOE grant de-sc0011784.

Open Access. This article is distributed under the terms of the Creative Commons Attribution License ([CC-BY 4.0](https://creativecommons.org/licenses/by/4.0/)), which permits any use, distribution and reproduction in any medium, provided the original author(s) and source are credited.

References

- [1] MEG II collaboration, *The design of the MEG II experiment*, *Eur. Phys. J. C* **78** (2018) 380 [[arXiv:1801.04688](https://arxiv.org/abs/1801.04688)] [[INSPIRE](#)].
- [2] MU3E collaboration, *The Mu3e Experiment*, *Nucl. Phys. B Proc. Suppl.* **248-250** (2014) 35 [[INSPIRE](#)].
- [3] A. Blondel et al., *Research Proposal for an Experiment to Search for the Decay $\mu \rightarrow eee$* , [arXiv:1301.6113](https://arxiv.org/abs/1301.6113) [[INSPIRE](#)].
- [4] COMET collaboration, *COMET Phase-I Technical Design Report*, *PTEP* **2020** (2020) 033C01 [[arXiv:1812.09018](https://arxiv.org/abs/1812.09018)] [[INSPIRE](#)].
- [5] MU2E collaboration, *Mu2e Technical Design Report*, [arXiv:1501.05241](https://arxiv.org/abs/1501.05241) [[INSPIRE](#)].
- [6] BELLE-II collaboration, *Prospects for τ Lepton Physics at Belle II*, in *17th Conference on Flavor Physics and CP-violation*, (2019) [[arXiv:1906.08950](https://arxiv.org/abs/1906.08950)] [[INSPIRE](#)].
- [7] L. Calibbi and G. Signorelli, *Charged Lepton Flavour Violation: An Experimental and Theoretical Introduction*, *Riv. Nuovo Cim.* **41** (2018) 71 [[arXiv:1709.00294](https://arxiv.org/abs/1709.00294)] [[INSPIRE](#)].
- [8] J.L. Feng, T. Moroi, H. Murayama and E. Schnapka, *Third generation familons, b factories, and neutrino cosmology*, *Phys. Rev. D* **57** (1998) 5875 [[hep-ph/9709411](https://arxiv.org/abs/hep-ph/9709411)] [[INSPIRE](#)].

- [9] A. Jodidio et al., *Search for Right-Handed Currents in Muon Decay*, *Phys. Rev. D* **34** (1986) 1967 [Erratum *ibid.* **37** (1988) 237] [INSPIRE].
- [10] TWIST collaboration, *Search for two body muon decay signals*, *Phys. Rev. D* **91** (2015) 052020 [arXiv:1409.0638] [INSPIRE].
- [11] MU3E collaboration, *The Rare and Forbidden: Testing Physics Beyond the Standard Model with Mu3e*, *SciPost Phys. Proc.* **1** (2019) 052 [arXiv:1812.00741] [INSPIRE].
- [12] A.R. Zhitnitsky, *On Possible Suppression of the Axion Hadron Interactions* (in Russian), *Sov. J. Nucl. Phys.* **31** (1980) 260 [INSPIRE].
- [13] M. Dine, W. Fischler and M. Srednicki, *A Simple Solution to the Strong CP Problem with a Harmless Axion*, *Phys. Lett. B* **104** (1981) 199 [INSPIRE].
- [14] J. Preskill, M.B. Wise and F. Wilczek, *Cosmology of the Invisible Axion*, *Phys. Lett. B* **120** (1983) 127 [INSPIRE].
- [15] L.F. Abbott and P. Sikivie, *A Cosmological Bound on the Invisible Axion*, *Phys. Lett. B* **120** (1983) 133 [INSPIRE].
- [16] M. Dine and W. Fischler, *The Not So Harmless Axion*, *Phys. Lett. B* **120** (1983) 137 [INSPIRE].
- [17] G. Grilli di Cortona, E. Hardy, J. Pardo Vega and G. Villadoro, *The QCD axion, precisely*, *JHEP* **01** (2016) 034 [arXiv:1511.02867] [INSPIRE].
- [18] A. Hook, G. Marques-Tavares and Y. Tsai, *Scalars Gliding through an Expanding Universe*, *Phys. Rev. Lett.* **124** (2020) 211801 [arXiv:1912.08817] [INSPIRE].
- [19] A. Arvanitaki, S. Dimopoulos, M. Galanis, L. Lehner, J.O. Thompson and K. Van Tilburg, *Large-misalignment mechanism for the formation of compact axion structures: Signatures from the QCD axion to fuzzy dark matter*, *Phys. Rev. D* **101** (2020) 083014 [arXiv:1909.11665] [INSPIRE].
- [20] ARIADNE collaboration, *Progress on the ARIADNE axion experiment*, *Springer Proc. Phys.* **211** (2018) 151 [arXiv:1710.05413] [INSPIRE].
- [21] D.J.E. Marsh, K.-C. Fong, E.W. Lentz, L. Smejkal and M.N. Ali, *Proposal to Detect Dark Matter using Axionic Topological Antiferromagnets*, *Phys. Rev. Lett.* **123** (2019) 121601 [arXiv:1807.08810] [INSPIRE].
- [22] A. Mitridate, T. Trickle, Z. Zhang and K.M. Zurek, *Detectability of Axion Dark Matter with Phonon Polaritons and Magnons*, *Phys. Rev. D* **102** (2020) 095005 [arXiv:2005.10256] [INSPIRE].
- [23] J. Martin Camalich, M. Pospelov, P.N.H. Vuong, R. Ziegler and J. Zupan, *Quark Flavor Phenomenology of the QCD Axion*, *Phys. Rev. D* **102** (2020) 015023 [arXiv:2002.04623] [INSPIRE].
- [24] L. Calibbi, F. Goertz, D. Redigolo, R. Ziegler and J. Zupan, *Minimal axion model from flavor*, *Phys. Rev. D* **95** (2017) 095009 [arXiv:1612.08040] [INSPIRE].
- [25] J. Albrecht, E. Stamou, R. Ziegler and R. Zwicky, *Probing flavoured Axions in the Tail of $B_q \rightarrow \mu^+ \mu^-$* , arXiv:1911.05018 [INSPIRE].
- [26] Q. Bonnefoy, E. Dudas and S. Pokorski, *Chiral Froggatt-Nielsen models, gauge anomalies and flavourful axions*, *JHEP* **01** (2020) 191 [arXiv:1909.05336] [INSPIRE].

- [27] M.B. Gavela, R. Houtz, P. Quilez, R. Del Rey and O. Sumensari, *Flavor constraints on electroweak ALP couplings*, *Eur. Phys. J. C* **79** (2019) 369 [[arXiv:1901.02031](#)] [[INSPIRE](#)].
- [28] F. Björkeröth, L. Di Luzio, F. Mescia and E. Nardi, *U(1) flavour symmetries as Peccei-Quinn symmetries*, *JHEP* **02** (2019) 133 [[arXiv:1811.09637](#)] [[INSPIRE](#)].
- [29] T. Alanne, S. Blasi and F. Goertz, *Common source for scalars: Flavored axion-Higgs unification*, *Phys. Rev. D* **99** (2019) 015028 [[arXiv:1807.10156](#)] [[INSPIRE](#)].
- [30] Y. Ema, D. Hagihara, K. Hamaguchi, T. Moroi and K. Nakayama, *Supersymmetric Flaxion*, *JHEP* **04** (2018) 094 [[arXiv:1802.07739](#)] [[INSPIRE](#)].
- [31] F. Arias-Aragon and L. Merlo, *The Minimal Flavour Violating Axion*, *JHEP* **10** (2017) 168 [*Erratum ibid.* **11** (2019) 152] [[arXiv:1709.07039](#)] [[INSPIRE](#)].
- [32] K. Choi, S.H. Im, C.B. Park and S. Yun, *Minimal Flavor Violation with Axion-like Particles*, *JHEP* **11** (2017) 070 [[arXiv:1708.00021](#)] [[INSPIRE](#)].
- [33] Y. Ema, K. Hamaguchi, T. Moroi and K. Nakayama, *Flaxion: a minimal extension to solve puzzles in the standard model*, *JHEP* **01** (2017) 096 [[arXiv:1612.05492](#)] [[INSPIRE](#)].
- [34] C.D. Froggatt and H.B. Nielsen, *Hierarchy of Quark Masses, Cabibbo Angles and CP-violation*, *Nucl. Phys. B* **147** (1979) 277 [[INSPIRE](#)].
- [35] M. Linster and R. Ziegler, *A Realistic U(2) Model of Flavor*, *JHEP* **08** (2018) 058 [[arXiv:1805.07341](#)] [[INSPIRE](#)].
- [36] T. Higaki and J. Kawamura, *A low-scale flavon model with a Z_N symmetry*, *JHEP* **03** (2020) 129 [[arXiv:1911.09127](#)] [[INSPIRE](#)].
- [37] C.D. Carone and M. Merchand, *Flavor from the double tetrahedral group without supersymmetry: Flavorful axions and neutrinos*, *Phys. Rev. D* **100** (2019) 035006 [[arXiv:1904.11059](#)] [[INSPIRE](#)].
- [38] O. Davidi, R.S. Gupta, G. Perez, D. Redigolo and A. Shalit, *The hierarchion, a relaxion addressing the Standard Model's hierarchies*, *JHEP* **08** (2018) 153 [[arXiv:1806.08791](#)] [[INSPIRE](#)].
- [39] F. Björkeröth, E.J. Chun and S.F. King, *Flavourful Axion Phenomenology*, *JHEP* **08** (2018) 117 [[arXiv:1806.00660](#)] [[INSPIRE](#)].
- [40] M. Reig, J.W.F. Valle and F. Wilczek, *SO(3) family symmetry and axions*, *Phys. Rev. D* **98** (2018) 095008 [[arXiv:1805.08048](#)] [[INSPIRE](#)].
- [41] A. Ibarra, E. Molinaro and S.T. Petcov, *Low Energy Signatures of the TeV Scale See-Saw Mechanism*, *Phys. Rev. D* **84** (2011) 013005 [[arXiv:1103.6217](#)] [[INSPIRE](#)].
- [42] A.-K. Perrevoort, *Sensitivity Studies on New Physics in the Mu3e Experiment and Development of Firmware for the Front-End of the Mu3e Pixel Detector*, Ph.D. Thesis, University of Heidelberg, Germany (2018), [[DOI](#)].
- [43] ARGUS collaboration, *A search for lepton flavor violating decays $\tau \rightarrow e\alpha$, $\tau \rightarrow \mu\alpha$* , *Z. Phys. C* **68** (1995) 25 [[INSPIRE](#)].
- [44] M.M. Miller Bertolami, B.E. Melendez, L.G. Althaus and J. Isern, *Revisiting the axion bounds from the Galactic white dwarf luminosity function*, *JCAP* **10** (2014) 069 [[arXiv:1406.7712](#)] [[INSPIRE](#)].

- [45] R. Bollig, W. DeRocco, P.W. Graham and H.-T. Janka, *Muons in Supernovae: Implications for the Axion-Muon Coupling*, *Phys. Rev. Lett.* **125** (2020) 051104 [Erratum *ibid.* **126** (2021) 189901] [[arXiv:2005.07141](#)] [[INSPIRE](#)].
- [46] D. Croon, G. Elor, R.K. Leane and S.D. McDermott, *Supernova Muons: New Constraints on Z' Bosons, Axions and ALPs*, *JHEP* **01** (2021) 107 [[arXiv:2006.13942](#)] [[INSPIRE](#)].
- [47] R.D. Bolton et al., *Search for Rare Muon Decays with the Crystal Box Detector*, *Phys. Rev. D* **38** (1988) 2077 [[INSPIRE](#)].
- [48] BELLE collaboration, *MC study for the lepton flavor violating tau decay into a lepton and an undetectable particle*, *Nucl. Part. Phys. Proc.* **287-288** (2017) 218 [[INSPIRE](#)].
- [49] E. Izaguirre, T. Lin and B. Shuve, *Searching for Axionlike Particles in Flavor-Changing Neutral Current Processes*, *Phys. Rev. Lett.* **118** (2017) 111802 [[arXiv:1611.09355](#)] [[INSPIRE](#)].
- [50] Z. Chacko, A. Dev, P. Du, V. Poulin and Y. Tsai, *Cosmological Limits on the Neutrino Mass and Lifetime*, *JHEP* **04** (2020) 020 [[arXiv:1909.05275](#)] [[INSPIRE](#)].
- [51] Z. Chacko, A. Dev, P. Du, V. Poulin and Y. Tsai, *Determining the Neutrino Lifetime from Cosmology*, *Phys. Rev. D* **103** (2021) 043519 [[arXiv:2002.08401](#)] [[INSPIRE](#)].
- [52] M. Bauer, M. Neubert and A. Thamm, *Collider Probes of Axion-Like Particles*, *JHEP* **12** (2017) 044 [[arXiv:1708.00443](#)] [[INSPIRE](#)].
- [53] J. Heeck and W. Rodejohann, *Lepton flavor violation with displaced vertices*, *Phys. Lett. B* **776** (2018) 385 [[arXiv:1710.02062](#)] [[INSPIRE](#)].
- [54] M. Bauer, M. Neubert, S. Renner, M. Schnubel and A. Thamm, *Axionlike Particles, Lepton-Flavor Violation, and a New Explanation of a_μ and a_e* , *Phys. Rev. Lett.* **124** (2020) 211803 [[arXiv:1908.00008](#)] [[INSPIRE](#)].
- [55] C. Cornella, P. Paradisi and O. Sumensari, *Hunting for ALPs with Lepton Flavor Violation*, *JHEP* **01** (2020) 158 [[arXiv:1911.06279](#)] [[INSPIRE](#)].
- [56] MEG collaboration, *Search for lepton flavour violating muon decay mediated by a new light particle in the MEG experiment*, *Eur. Phys. J. C* **80** (2020) 858 [[arXiv:2005.00339](#)] [[INSPIRE](#)].
- [57] SINDRUM collaboration, *Limits for Shortlived Neutral Particles Emitted μ^+ or π^+ Decay*, *Phys. Lett. B* **175** (1986) 101 [[INSPIRE](#)].
- [58] MEG II and MU3E collaborations, *MEG II and Mu3e status and plan*, *EPJ Web Conf.* **212** (2019) 01004 [[INSPIRE](#)].
- [59] R. Conlin and A.A. Petrov, *Muonium-antimuonium oscillations in effective field theory*, *Phys. Rev. D* **102** (2020) 095001 [[arXiv:2005.10276](#)] [[INSPIRE](#)].
- [60] M. Endo, S. Iguro and T. Kitahara, *Probing $e\mu$ flavor-violating ALP at Belle II*, *JHEP* **06** (2020) 040 [[arXiv:2002.05948](#)] [[INSPIRE](#)].
- [61] F. Scheck, *Muon Physics*, *Phys. Rept.* **44** (1978) 187 [[INSPIRE](#)].
- [62] Y. Kuno and Y. Okada, *Muon decay and physics beyond the standard model*, *Rev. Mod. Phys.* **73** (2001) 151 [[hep-ph/9909265](#)] [[INSPIRE](#)].
- [63] M. Reiser, *Theory and design of charged particle beams*, Wiley (1995).

- [64] V. Kumar, *Understanding the focusing of charged particle beams in a solenoid magnetic field*, *Am. J. Phys.* **77** (2009) 737 [INSPIRE].
- [65] S.E. Derenzo, *Measurement of the low-energy end of the mu-plus decay spectrum*, *Phys. Rev.* **181** (1969) 1854 [INSPIRE].
- [66] R. Bilger et al., *Search for exotic muon decays*, *Phys. Lett. B* **446** (1999) 363 [hep-ph/9811333] [INSPIRE].
- [67] PIENU collaboration, *Improved search for two body muon decay $\mu^+ \rightarrow e^+ X_H$* , *Phys. Rev. D* **101** (2020) 052014 [arXiv:2002.09170] [INSPIRE].
- [68] G.J. Feldman and R.D. Cousins, *A unified approach to the classical statistical analysis of small signals*, *Phys. Rev. D* **57** (1998) 3873 [physics/9711021] [INSPIRE].
- [69] N. Atanov et al., *The Mu2e Calorimeter Final Technical Design Report*, arXiv:1802.06341 [INSPIRE].
- [70] MEG collaboration, *Muon polarization in the MEG experiment: predictions and measurements*, *Eur. Phys. J. C* **76** (2016) 223 [arXiv:1510.04743] [INSPIRE].
- [71] MU3E collaboration, D. vom Bruch, *Online data reduction using track and vertex reconstruction on gpus for the mu3e experiment*, *EPJ Web Conf.* **150** (2017) 00013.
- [72] G. Rutar et al., *A dedicated calibration tool for the MEG and MEG II positron spectrometer*, *Nucl. Instrum. Meth. A* **824** (2016) 575 [INSPIRE].
- [73] X. Garcia i Tormo, D. Bryman, A. Czarnecki and M. Dowling, *Bounds on majoron emission from muon to electron conversion experiments*, *Phys. Rev. D* **84** (2011) 113010 [arXiv:1110.2874] [INSPIRE].
- [74] Y. Uesaka, *Model identification in $\mu^- \rightarrow e^-$ conversion with invisible boson emission using muonic atoms*, *Phys. Rev. D* **102** (2020) 095007 [arXiv:2005.07894] [INSPIRE].
- [75] J.T. Goldman et al., *Light Boson Emission in the Decay of the μ^+* , *Phys. Rev. D* **36** (1987) 1543 [INSPIRE].
- [76] M. Hirsch, A. Vicente, J. Meyer and W. Porod, *Majoron emission in muon and tau decays revisited*, *Phys. Rev. D* **79** (2009) 055023 [Erratum *ibid.* **79** (2009) 079901] [arXiv:0902.0525] [INSPIRE].
- [77] MEG collaboration, *Search for the lepton flavour violating decay $\mu^+ \rightarrow e^+ \gamma$ with the full dataset of the MEG experiment*, *Eur. Phys. J. C* **76** (2016) 434 [arXiv:1605.05081] [INSPIRE].
- [78] J. Heeck, *Lepton flavor violation with light vector bosons*, *Phys. Lett. B* **758** (2016) 101 [arXiv:1602.03810] [INSPIRE].
- [79] G.G. Raffelt, *Stars as laboratories for fundamental physics*, University of Chicago Press (1996).
- [80] G. Raffelt and A. Weiss, *Red giant bound on the axion-electron coupling revisited*, *Phys. Rev. D* **51** (1995) 1495 [hep-ph/9410205] [INSPIRE].
- [81] N. Viaux et al., *Neutrino and axion bounds from the globular cluster M5 (NGC 5904)*, *Phys. Rev. Lett.* **111** (2013) 231301 [arXiv:1311.1669] [INSPIRE].
- [82] M. Giannotti, I.G. Irastorza, J. Redondo, A. Ringwald and K. Saikawa, *Stellar Recipes for Axion Hunters*, *JCAP* **10** (2017) 010 [arXiv:1708.02111] [INSPIRE].

- [83] G.G. Raffelt, *Astrophysical methods to constrain axions and other novel particle phenomena*, *Phys. Rept.* **198** (1990) 1 [INSPIRE].
- [84] A. Burrows and J.M. Lattimer, *The birth of neutron stars*, *Astrophys. J.* **307** (1986) 178 [INSPIRE].
- [85] H.A. Bethe, *Supernova mechanisms*, *Rev. Mod. Phys.* **62** (1990) 801 [INSPIRE].
- [86] G.G. Raffelt, *Axion bremsstrahlung in red giants*, *Phys. Rev. D* **41** (1990) 1324 [INSPIRE].
- [87] M. Nakagawa, Y. Kohyama and N. Itoh, *Axion Bremsstrahlung in Dense Stars*, *Astrophys. J.* **322** (1987) 291 [INSPIRE].
- [88] M. Nakagawa, T. Adachi, Y. Kohyama and N. Itoh, *Axion bremsstrahlung in dense stars. II — Phonon contributions*, *Astrophys. J.* **326** (1988) 241 [INSPIRE].
- [89] T. Hatsuda and M. Yoshimura, *Axion Flux From Nascent Neutron Stars*, *Phys. Lett. B* **203** (1988) 469 [INSPIRE].
- [90] M. Prakash, I. Bombaci, M. Prakash, P.J. Ellis, J.M. Lattimer and R. Knorren, *Composition and structure of protoneutron stars*, *Phys. Rept.* **280** (1997) 1 [nucl-th/9603042] [INSPIRE].
- [91] N. Bar, K. Blum and G. D’Amico, *Is there a supernova bound on axions?*, *Phys. Rev. D* **101** (2020) 123025 [arXiv:1907.05020] [INSPIRE].
- [92] P. Cigan et al., *High angular resolution ALMA images of dust and molecules in the SN 1987A ejecta*, *Astrophys. J.* **886** (2019) 51 [arXiv:1910.02960] [INSPIRE].
- [93] D. Page, M.V. Beznogov, I. Garibay, J.M. Lattimer, M. Prakash and H.-T. Janka, *NS 1987A in SN 1987A*, *Astrophys. J.* **898** (2020) 125 [arXiv:2004.06078] [INSPIRE].
- [94] A. Boyarsky, J.W. den Herder, A. Neronov and O. Ruchayskiy, *Search for the light dark matter with an X-ray spectrometer*, *Astropart. Phys.* **28** (2007) 303 [astro-ph/0612219] [INSPIRE].
- [95] XQC collaboration, *Searching for keV Sterile Neutrino Dark Matter with X-ray Microcalorimeter Sounding Rockets*, *Astrophys. J.* **814** (2015) 82 [arXiv:1506.05519] [INSPIRE].
- [96] I.M. Bloch, A. Caputo, R. Essig, D. Redigolo, M. Sholapurkar and T. Volansky, *Exploring new physics with $O(\text{keV})$ electron recoils in direct detection experiments*, *JHEP* **01** (2021) 178 [arXiv:2006.14521] [INSPIRE].
- [97] ADMX collaboration, *Extended Search for the Invisible Axion with the Axion Dark Matter Experiment*, *Phys. Rev. Lett.* **124** (2020) 101303 [arXiv:1910.08638] [INSPIRE].
- [98] ADMX collaboration, *Piezoelectrically Tuned Multimode Cavity Search for Axion Dark Matter*, *Phys. Rev. Lett.* **121** (2018) 261302 [arXiv:1901.00920] [INSPIRE].
- [99] ADMX collaboration, *A Search for Invisible Axion Dark Matter with the Axion Dark Matter Experiment*, *Phys. Rev. Lett.* **120** (2018) 151301 [arXiv:1804.05750] [INSPIRE].
- [100] XENON collaboration, *Light Dark Matter Search with Ionization Signals in XENON1T*, *Phys. Rev. Lett.* **123** (2019) 251801 [arXiv:1907.11485] [INSPIRE].
- [101] PANDAX collaboration, *Limits on Axion Couplings from the First 80 Days of Data of the PandaX-II Experiment*, *Phys. Rev. Lett.* **119** (2017) 181806 [arXiv:1707.07921] [INSPIRE].

- [102] R. Barbieri, M. Cerdonio, G. Fiorentini and S. Vitale, *Axion to magnon conversion: a scheme for the detection of galactic axions*, *Phys. Lett. B* **226** (1989) 357 [INSPIRE].
- [103] R. Barbieri et al., *Searching for galactic axions through magnetized media: the QUAX proposal*, *Phys. Dark Univ.* **15** (2017) 135 [arXiv:1606.02201] [INSPIRE].
- [104] S. Chigusa, T. Moroi and K. Nakayama, *Detecting light boson dark matter through conversion into a magnon*, *Phys. Rev. D* **101** (2020) 096013 [arXiv:2001.10666] [INSPIRE].
- [105] C. Creque-Sarbinowski and M. Kamionkowski, *Searching for Decaying and Annihilating Dark Matter with Line Intensity Mapping*, *Phys. Rev. D* **98** (2018) 063524 [arXiv:1806.11119] [INSPIRE].
- [106] T.M. Shokair et al., *Future Directions in the Microwave Cavity Search for Dark Matter Axions*, *Int. J. Mod. Phys. A* **29** (2014) 1443004 [arXiv:1405.3685] [INSPIRE].
- [107] CAPP/IBS collaboration, *Haloscope searches for dark matter axions at the Center for Axion and Precision Physics Research*, *EPJ Web Conf.* **164** (2017) 01012 [arXiv:1702.03664] [INSPIRE].
- [108] C. Gatti et al., *The Klash Proposal: Status and Perspectives*, in *14th Patras Workshop on Axions, WIMPs and WISPs*, (2018) [arXiv:1811.06754] [INSPIRE].
- [109] B.T. McAllister, G. Flower, E.N. Ivanov, M. Goryachev, J. Bourhill and M.E. Tobar, *The ORGAN Experiment: An axion haloscope above 15 GHz*, *Phys. Dark Univ.* **18** (2017) 67 [arXiv:1706.00209] [INSPIRE].
- [110] MADMAX WORKING GROUP collaboration, *Dielectric Haloscopes: A New Way to Detect Axion Dark Matter*, *Phys. Rev. Lett.* **118** (2017) 091801 [arXiv:1611.05865] [INSPIRE].
- [111] M. Baryakhtar, J. Huang and R. Lasenby, *Axion and hidden photon dark matter detection with multilayer optical haloscopes*, *Phys. Rev. D* **98** (2018) 035006 [arXiv:1803.11455] [INSPIRE].
- [112] R. Hill, K.W. Masui and D. Scott, *The Spectrum of the Universe*, *Appl. Spectrosc.* **72** (2018) 663 [arXiv:1802.03694] [INSPIRE].
- [113] D. Barret et al., *Athena+: The first Deep Universe X-ray Observatory*, (2013) [arXiv:1310.3814] [INSPIRE].
- [114] O. Doré et al., *Cosmology with the SPHEREX All-Sky Spectral Survey*, arXiv:1412.4872 [INSPIRE].
- [115] P. Arias, D. Cadamuro, M. Goodsell, J. Jaeckel, J. Redondo and A. Ringwald, *WISPy Cold Dark Matter*, *JCAP* **06** (2012) 013 [arXiv:1201.5902] [INSPIRE].
- [116] F. D’Eramo, R.Z. Ferreira, A. Notari and J.L. Bernal, *Hot Axions and the H_0 tension*, *JCAP* **11** (2018) 014 [arXiv:1808.07430] [INSPIRE].
- [117] N. Blinov, M.J. Dolan, P. Draper and J. Kozaczuk, *Dark matter targets for axionlike particle searches*, *Phys. Rev. D* **100** (2019) 015049 [arXiv:1905.06952] [INSPIRE].
- [118] PLANCK collaboration, *Planck 2015 results. XIII. Cosmological parameters*, *Astron. Astrophys.* **594** (2016) A13 [arXiv:1502.01589] [INSPIRE].
- [119] D.J. Gross, R.D. Pisarski and L.G. Yaffe, *QCD and Instantons at Finite Temperature*, *Rev. Mod. Phys.* **53** (1981) 43 [INSPIRE].
- [120] I.M. Bloch, R. Essig, K. Tobioka, T. Volansky and T.-T. Yu, *Searching for Dark Absorption with Direct Detection Experiments*, *JHEP* **06** (2017) 087 [arXiv:1608.02123] [INSPIRE].

- [121] QUAX collaboration, *Axion search with a quantum-limited ferromagnetic haloscope*, *Phys. Rev. Lett.* **124** (2020) 171801 [[arXiv:2001.08940](#)] [[INSPIRE](#)].
- [122] IAXO collaboration, *The International Axion Observatory IAXO. Letter of Intent to the CERN SPS committee*. CERN-SPSC-2013-022, [[INSPIRE](#)].
- [123] IAXO collaboration, *Physics potential of the International Axion Observatory (IAXO)*, *JCAP* **06** (2019) 047 [[arXiv:1904.09155](#)] [[INSPIRE](#)].
- [124] R. Bähre et al., *Any light particle search II — Technical Design Report, 2013 JINST* **8** T09001 [[arXiv:1302.5647](#)] [[INSPIRE](#)].
- [125] M. Leurer, Y. Nir and N. Seiberg, *Mass matrix models*, *Nucl. Phys. B* **398** (1993) 319 [[hep-ph/9212278](#)] [[INSPIRE](#)].
- [126] M. Leurer, Y. Nir and N. Seiberg, *Mass matrix models: The Sequel*, *Nucl. Phys. B* **420** (1994) 468 [[hep-ph/9310320](#)] [[INSPIRE](#)].
- [127] A. Smolković, M. Tammara and J. Zupan, *Anomaly free Froggatt-Nielsen models of flavor*, *JHEP* **10** (2019) 188 [[arXiv:1907.10063](#)] [[INSPIRE](#)].
- [128] Y. Chikashige, R.N. Mohapatra and R.D. Peccei, *Are There Real Goldstone Bosons Associated with Broken Lepton Number?*, *Phys. Lett. B* **98** (1981) 265 [[INSPIRE](#)].
- [129] J. Schechter and J.W.F. Valle, *Neutrino Decay and Spontaneous Violation of Lepton Number*, *Phys. Rev. D* **25** (1982) 774 [[INSPIRE](#)].
- [130] A. Broncano, M.B. Gavela and E.E. Jenkins, *The effective Lagrangian for the seesaw model of neutrino mass and leptogenesis*, *Phys. Lett. B* **552** (2003) 177 [Erratum *ibid.* **636** (2006) 332] [[hep-ph/0210271](#)] [[INSPIRE](#)].
- [131] M. Raidal, A. Strumia and K. Turzyski, *Low-scale standard supersymmetric leptogenesis*, *Phys. Lett. B* **609** (2005) 351 [Addendum *ibid.* **632** (2006) 752] [[hep-ph/0408015](#)] [[INSPIRE](#)].
- [132] J. Kersten and A.Y. Smirnov, *Right-Handed Neutrinos at CERN LHC and the Mechanism of Neutrino Mass Generation*, *Phys. Rev. D* **76** (2007) 073005 [[arXiv:0705.3221](#)] [[INSPIRE](#)].
- [133] A. Abada, C. Biggio, F. Bonnet, M.B. Gavela and T. Hambye, *Low energy effects of neutrino masses*, *JHEP* **12** (2007) 061 [[arXiv:0707.4058](#)] [[INSPIRE](#)].
- [134] M.B. Gavela, T. Hambye, D. Hernandez and P. Hernández, *Minimal Flavour Seesaw Models*, *JHEP* **09** (2009) 038 [[arXiv:0906.1461](#)] [[INSPIRE](#)].
- [135] A. Ibarra, E. Molinaro and S.T. Petcov, *TeV Scale See-Saw Mechanisms of Neutrino Mass Generation, the Majorana Nature of the Heavy Singlet Neutrinos and $(\beta\beta)_{0\nu}$ -Decay*, *JHEP* **09** (2010) 108 [[arXiv:1007.2378](#)] [[INSPIRE](#)].
- [136] D.N. Dinh, A. Ibarra, E. Molinaro and S.T. Petcov, *The $\mu - e$ Conversion in Nuclei, $\mu \rightarrow e\gamma$, $\mu \rightarrow 3e$ Decays and TeV Scale See-Saw Scenarios of Neutrino Mass Generation*, *JHEP* **08** (2012) 125 [Erratum *ibid.* **09** (2013) 023] [[arXiv:1205.4671](#)] [[INSPIRE](#)].
- [137] C.G. Cely, A. Ibarra, E. Molinaro and S.T. Petcov, *Higgs Decays in the Low Scale Type I See-Saw Model*, *Phys. Lett. B* **718** (2013) 957 [[arXiv:1208.3654](#)] [[INSPIRE](#)].
- [138] R. Alonso, M. Dhen, M.B. Gavela and T. Hambye, *Muon conversion to electron in nuclei in type-I seesaw models*, *JHEP* **01** (2013) 118 [[arXiv:1209.2679](#)] [[INSPIRE](#)].
- [139] M. Gorghetto and G. Villadoro, *Topological Susceptibility and QCD Axion Mass: QED and NNLO corrections*, *JHEP* **03** (2019) 033 [[arXiv:1812.01008](#)] [[INSPIRE](#)].

- [140] P. Carenza, T. Fischer, M. Giannotti, G. Guo, G. Martínez-Pinedo and A. Mirizzi, *Improved axion emissivity from a supernova via nucleon-nucleon bremsstrahlung*, *JCAP* **10** (2019) 016 [Erratum *ibid.* **05** (2020) E01] [[arXiv:1906.11844](#)] [[INSPIRE](#)].
- [141] J.E. Kim, *Weak Interaction Singlet and Strong CP Invariance*, *Phys. Rev. Lett.* **43** (1979) 103 [[INSPIRE](#)].
- [142] M.A. Shifman, A.I. Vainshtein and V.I. Zakharov, *Can Confinement Ensure Natural CP Invariance of Strong Interactions?*, *Nucl. Phys. B* **166** (1980) 493 [[INSPIRE](#)].
- [143] L. Di Luzio, M. Giannotti, E. Nardi and L. Visinelli, *The landscape of QCD axion models*, *Phys. Rept.* **870** (2020) 1 [[arXiv:2003.01100](#)] [[INSPIRE](#)].
- [144] L. Di Luzio, F. Mescia and E. Nardi, *Redefining the Axion Window*, *Phys. Rev. Lett.* **118** (2017) 031801 [[arXiv:1610.07593](#)] [[INSPIRE](#)].
- [145] L. Di Luzio, F. Mescia and E. Nardi, *Window for preferred axion models*, *Phys. Rev. D* **96** (2017) 075003 [[arXiv:1705.05370](#)] [[INSPIRE](#)].
- [146] CAST collaboration, *New CAST Limit on the Axion-Photon Interaction*, *Nature Phys.* **13** (2017) 584 [[arXiv:1705.02290](#)] [[INSPIRE](#)].
- [147] R. Barbieri, G.R. Dvali and L.J. Hall, *Predictions from a U(2) flavor symmetry in supersymmetric theories*, *Phys. Lett. B* **377** (1996) 76 [[hep-ph/9512388](#)] [[INSPIRE](#)].
- [148] R. Barbieri, L.J. Hall and A. Romanino, *Consequences of a U(2) flavor symmetry*, *Phys. Lett. B* **401** (1997) 47 [[hep-ph/9702315](#)] [[INSPIRE](#)].
- [149] L.J. Hall, H. Murayama and N. Weiner, *Neutrino mass anarchy*, *Phys. Rev. Lett.* **84** (2000) 2572 [[hep-ph/9911341](#)] [[INSPIRE](#)].
- [150] G. Altarelli, F. Feruglio, I. Masina and L. Merlo, *Repressing Anarchy in Neutrino Mass Textures*, *JHEP* **11** (2012) 139 [[arXiv:1207.0587](#)] [[INSPIRE](#)].
- [151] J. Bergstrom, D. Meloni and L. Merlo, *Bayesian comparison of U(1) lepton flavor models*, *Phys. Rev. D* **89** (2014) 093021 [[arXiv:1403.4528](#)] [[INSPIRE](#)].
- [152] A. Pilaftsis, *Astrophysical and terrestrial constraints on singlet Majoron models*, *Phys. Rev. D* **49** (1994) 2398 [[hep-ph/9308258](#)] [[INSPIRE](#)].
- [153] C. Garcia-Cely and J. Heeck, *Neutrino Lines from Majoron Dark Matter*, *JHEP* **05** (2017) 102 [[arXiv:1701.07209](#)] [[INSPIRE](#)].
- [154] J. Heeck and H.H. Patel, *Majoron at two loops*, *Phys. Rev. D* **100** (2019) 095015 [[arXiv:1909.02029](#)] [[INSPIRE](#)].
- [155] I. Esteban, M.C. Gonzalez-Garcia, A. Hernandez-Cabezudo, M. Maltoni and T. Schwetz, *Global analysis of three-flavour neutrino oscillations: synergies and tensions in the determination of θ_{23} , δ_{CP} , and the mass ordering*, *JHEP* **01** (2019) 106 [[arXiv:1811.05487](#)] [[INSPIRE](#)].
- [156] NuFIT 4.1 (2019), [www.nu-fit.org](#).
- [157] J.S. Lee, *Revisiting Supernova 1987A Limits on Axion-Like-Particles*, [arXiv:1808.10136](#) [[INSPIRE](#)].
- [158] M. Frigerio, T. Hambye and E. Masso, *Sub-GeV dark matter as pseudo-Goldstone from the seesaw scale*, *Phys. Rev. X* **1** (2011) 021026 [[arXiv:1107.4564](#)] [[INSPIRE](#)].

- [159] A. Katz, M. Reece and A. Sajjad, *Naturalness, $b \rightarrow s\gamma$, and SUSY heavy Higgses*, *JHEP* **10** (2014) 102 [[arXiv:1406.1172](#)] [[INSPIRE](#)].
- [160] N. Craig, F. D’Eramo, P. Draper, S. Thomas and H. Zhang, *The Hunt for the Rest of the Higgs Bosons*, *JHEP* **06** (2015) 137 [[arXiv:1504.04630](#)] [[INSPIRE](#)].
- [161] U. Nierste, M. Tabet and R. Ziegler, *Cornering Spontaneous CP Violation with Charged-Higgs-Boson Searches*, *Phys. Rev. Lett.* **125** (2020) 031801 [[arXiv:1912.11501](#)] [[INSPIRE](#)].