

ALPINIST: Axion-Like Particles In Numerous Interactions Simulated and Tabulated

Jan Jerhot,^a Babette Döbrich,^b Fatih Ertas,^c Felix Kahlhoefer^c
and Tommaso Spadaro^d

^a*Centre for Cosmology, Particle Physics and Phenomenology (CP3),
Université catholique de Louvain,
Chemin du Cyclotron 2, B-1348 Louvain-la-Neuve, Belgium*

^b*CERN,
Esplanade des Particules 1, 1211 Geneva 23, Switzerland*

^c*Institute for Theoretical Particle Physics and Cosmology (TTK), RWTH Aachen University,
D-52056 Aachen, Germany*

^d*Laboratori Nazionali di Frascati INFN,
Via E. Fermi 40, 00044 Frascati, Italy*

E-mail: jan.jerhot@cern.ch, babette.dobrich@cern.ch,
ertas@physik.rwth-aachen.de, kahlhoefer@physik.rwth-aachen.de,
tommaso.spadaro@cern.ch

ABSTRACT: Proton beam dump experiments are among the most promising strategies to search for light and feebly interacting states such as axion-like particles (ALPs). The interpretation of these experiments is however complicated by the wide range of ALP models and the multitude of different production and decay channels that can induce observable signals. Here we propose a new approach to this problem by separating the calculation of constraints and projected sensitivities into model-independent and model-dependent parts. The former rely on extensive Monte Carlo simulations of ALP production and decays, as well as estimates of the detection efficiencies based on simplified detector geometries. Once these simulations have been performed and tabulated, the latter parts only require simple analytical rescalings that can be performed using the public code ALPINIST released together with this work. We illustrate this approach by considering several ALP models with couplings to Standard Model gauge bosons. For the case of ALPs coupled to gluons we show that the sensitivity of proton beam dump experiments can be extended significantly by considering hadronic ALP decays into three-body final states.

KEYWORDS: Axions and ALPs, New Light Particles

ARXIV EPRINT: [2201.05170](https://arxiv.org/abs/2201.05170)

Contents

1	Introduction	1
2	General framework	2
3	ALP production	5
3.1	Photo-production	5
3.2	ALP production from meson mixing	6
3.3	ALP production from rare decays	7
4	ALP decay and detection of final states	9
4.1	General principles of the ALP Monte Carlo simulation	10
4.2	Decay channels	11
4.3	Propagation of decay products	12
5	Experimental framework	13
5.1	Results from past experiments	13
5.1.1	CHARM	13
5.1.2	NuCal	14
5.2	Acceptance of current experiments	14
5.2.1	NA62	15
5.3	Acceptance of proposed experiments	15
5.3.1	SHiP	16
5.3.2	DarkQuest	16
5.3.3	DUNE ND	17
5.3.4	SHADOWS	18
6	Results	18
6.1	Considering only di-photon decays	19
6.2	Including hadronic decays	21
7	Conclusions	22
A	Model-dependent calculations	24
A.1	FCNC transitions	24
A.2	ALP-meson mixing	25
B	Calculation of decay widths	26
C	Brief description of the ALPINIST code	27
D	Treatment of kinematics for the mixing production	28

1 Introduction

In recent years the search for physics beyond the Standard Model (SM) has begun to expand its focus from particles too heavy to be produced in laboratory experiments to light particles with extremely weak interactions [1]. As grows the range of models that predict such light degrees of freedom, so does the number of proposed and ongoing experiments that can probe the corresponding parameter spaces. A central challenge arising from this development is to identify those experiments that promise the greatest improvement in sensitivity to a wide variety of models. A popular way to address this challenge is to define a number of benchmark scenarios, for which exclusion bounds and sensitivity projections are to be calculated [2–4]. As the field matures, however, more flexible and model-independent methods are needed that allow for the reinterpretation of experimental bounds and sensitivities for a wide range of different scenarios.

This need is particularly great in the context of axion-like particles (ALPs), which are hypothetical elementary pseudoscalar particles arising as the pseudo-Nambu-Goldstone modes of an approximate global symmetry that is broken both spontaneously at some large energy scale Λ and explicitly by the non-zero ALP mass $m_a \ll \Lambda$. ALPs emerge naturally in many theories extending the SM symmetry group, which may feature new accidental global symmetries that are not experimentally observed and therefore assumed to be broken spontaneously [5]. A plethora of unobserved global symmetries also appears in theories with extra dimensions when these are compactified to four dimensions [6]. From a phenomenological point of view, the relation $m_a \ll \Lambda$ furthermore implies that ALPs can be viable dark matter (DM) candidates or mediate the interactions between SM and DM particles [7–10].

Since the details of the spontaneous symmetry breaking and the corresponding new physics at the energy scale Λ are typically unknown, ALPs can manifest themselves in a wide range of different ways at low energies. To approach this problem in a general way, one can express the interactions between ALPs and SM particles in terms of an effective field theory valid at energies $E \ll \Lambda$ [11–13]. This approach makes it possible to experimentally search for ALPs in a model-independent way, i.e. without referring to a specific symmetry breaking mechanism. While traditionally much energy has been focused on the effective ALP-photon coupling [14], recent studies have broadened the picture and considered also effective couplings to the other SM gauge bosons [15–19] as well as to SM Higgs bosons [20] and SM fermions [21, 22].

Ideally, one would like to study all of these effective couplings simultaneously, which requires the production of on-shell ALPs and the observation of their subsequent decays back into SM particles. For m_a in the MeV–GeV range, the most suitable laboratory for this purpose are particle accelerators. While colliders like the LHC have the advantage of larger ALP production cross sections [23], fixed-target experiments benefit from higher statistics and lower backgrounds and are therefore more sensitive to tiny couplings and small ALP masses [24]. To improve our understanding of ALPs and ultimately achieve their discovery, it is therefore essential to understand the signatures of ALPs in fixed-target experiments for general effective interactions.

In this work we propose a new approach for evaluating experimental sensitivities, which is based on the observation that the calculation of the predicted number of ALP events in a given experiment can be split into a number of *separate steps, each of which can be performed in a largely model-independent way*. Specifically, we

- write the differential ALP production cross section as the sum of several different production modes, each of which depends in a well-defined way on the effective ALP couplings;
- make use of the fact that the detection probability depends trivially on the various ALP branching ratios, with only the total ALP lifetime having a more complicated effect.

This makes it possible to perform all relevant calculations in advance and tabulate them in terms of a small number of model parameters. For a given combination of effective ALP couplings, these tables can then be evaluated and combined in a straight-forward way.

In addition to this conceptual innovation, our work also expands on the range of effects included in our calculations in order to extend existing approaches (using for example the automated simulations of beam-dump experiments in Maddump [25]) and to allow for a wide range of applications. On the ALP production side, we consider Primakoff production (both from on-shell [26] and off-shell [27] photons), ALP-meson mixing [20] and rare decays (both from B mesons [9] and D mesons [28]). On the ALP decay side, we include for the first time detailed experimental sensitivities for leptonic final states and three-body hadronic final states like $\pi^+\pi^-\pi^0$. For the latter we have implemented an automatic re-weighting of the Dalitz plot density based on detailed analytic calculations of hadronic decay modes [18, 29]. All these calculations are applied to a wide range of past, present and proposed experiments, for which we have implemented a modelling of detector geometries and efficiencies. To illustrate the flexibility of our approach, we apply it to several different ALP models featuring various effective couplings to SM gauge bosons.

Together with this work we release a public code named `ALPINIST` [30], which provides the relevant tables for ALP production and detection together with helpful routines for model-dependent calculations and a convenient user interface.

The remainder of this work is structured as follows. In section 2 we introduce the general framework used to obtain model-independent constraints on ALP models. The various ALP production and decay modes are discussed in sections 3 and 4, respectively. A detailed discussion of the experiments considered in this work is given in section 5. To illustrate our approach, we present experimental sensitivities for a range of ALP scenarios in section 6. The code `ALPINIST` is described in more detail in appendix C. The calculation of the various effective couplings and the resulting decay widths are discussed in the appendices A and B.

2 General framework

In this section we outline our general approach for simplifying the calculation of exclusion limits and sensitivity projections for ALPs in a model-independent way. For this purpose

we consider a pseudoscalar particle a with mass m_a that couples to the various particles of the SM via effective operators. Limiting ourselves to the case of dimension-5 operators, the strength of each interaction is parametrised by an effective coupling that can be expressed in terms of the new-physics scale Λ and the Wilson coefficients C of the various effective operators (see appendix A for details). For example, the coupling of the ALP to a pair of photons is given by

$$\mathcal{L}_{a\gamma\gamma} = e^2 \frac{C_{\gamma\gamma}}{\Lambda} a F^{\mu\nu} \tilde{F}_{\mu\nu}, \tag{2.1}$$

where $F^{\mu\nu}$ denotes the electromagnetic field strength tensor and $\tilde{F}^{\mu\nu}$ its dual. We denote the set of all couplings that define a specific parameter point of a given ALP model by \mathbf{C} . More details will be provided in section 6 and appendix A.

The focus of the present work is on proton beam fixed-target experiments, which aim to produce ALPs using a beam of highly energetic protons and detect the subsequent decays of ALPs into SM particles in a downstream detector. At such experiments ALPs can be produced via various mechanisms (production channels), for example through scattering of beam protons or secondary particles on the target material or through the decay of secondary particles. At first sight, the number N_{det} of expected ALP events in the detector depends in a complicated non-linear way on the couplings \mathbf{C} , which determine the production cross section, the kinematic distributions, the lifetime $\tau_a = 1/\Gamma_a$ and the branching ratios into the various final states. Upon closer inspection, it however becomes clear that this quantity can be factorized as follows:

$$N_{\text{det}} = \int d\theta_a dE_a \frac{d^2 N(m_a, \mathbf{C})}{d\theta_a dE_a} p_{\text{det}}(m_a, \Gamma_a, \theta_a, E_a, \mathbf{C}). \tag{2.2}$$

Here $d^2 N/d\theta_a dE_a$ denotes the spectrum of ALPs produced in the target as a function of the angle (relative to the incident beam) and the ALP energy in the laboratory frame, and p_{det} denotes the probability that an ALP with given angle and energy will lead to an observable signal in the detector, which depends on the position of the ALP decay as well as on the kinematic distribution of the decay products.

ALP production in the target can proceed via a number of different processes, which will be discussed in detail in section 3. The ALP spectrum is therefore simply a sum over all production channels:

$$\frac{d^2 N(m_a, \mathbf{C})}{d\theta_a dE_a} = \sum_i \frac{d^2 N_i(m_a, \mathbf{C})}{d\theta_a dE_a}. \tag{2.3}$$

Now the crucial observation is that for each individual production channel, the spectrum depends in a trivial way on the ALP coupling structure, in the sense that changing the couplings only affects the normalisation but not the shape of the spectrum. In other words, it is possible to calculate the spectrum for a given reference coupling \mathbf{C}_{ref} and then perform an appropriate rescaling:

$$\frac{d^2 N_i(m_a, \mathbf{C})}{d\theta_a dE_a} = f_i(\mathbf{C}, \mathbf{C}_{\text{ref}}) \frac{d^2 N_i(m_a, \mathbf{C}_{\text{ref}})}{d\theta_a dE_a}. \tag{2.4}$$

For example, if ALPs are produced via their coupling to photons, the production cross section is proportional to $C_{\gamma\gamma}^2$ and hence

$$f_{\gamma\gamma}(\mathbf{C}, \mathbf{C}_{\text{ref}}) = \frac{C_{\gamma\gamma}^2}{C_{\gamma\gamma,\text{ref}}^2}. \quad (2.5)$$

This simple formula holds even if the ALP effective photon coupling receives contributions from several different effective operators, see appendix A for details and ref. [20] for a comprehensive overview of such contributions.

We emphasize that the ALP mass, on the contrary, affects the ALP spectrum in a non-trivial way. It is therefore necessary to calculate the spectra $d^2N_i/d\theta_a dE_a$ separately for each ALP mass under consideration. For perturbative production processes such as photon fusion the ALP spectra can be derived from first principles. In many cases, however, production proceeds via non-perturbative processes and the spectra can be only obtained via simulations. These simulations however only need to be performed once (for the chosen values of \mathbf{C}_{ref}) and can then be applied to any ALP model.

Similarly, the detection probability can be split into different final states f :

$$p_{\text{det}}(m_a, \theta_a, E_a, \mathbf{C}) = \sum_f \text{BR}_{a \rightarrow f}(m_a, \mathbf{C}) p_{\text{det},f}(m_a, \Gamma_a, \theta_a, E_a), \quad (2.6)$$

where $\text{BR}_{a \rightarrow f} = \Gamma_{a \rightarrow f}/\Gamma_a$ denotes the branching ratio into the final state f . While these branching ratios may depend in a non-trivial way on the couplings \mathbf{C} , the detection probability $p_{\text{det},f}(m_a, \Gamma_a, \theta_a, E_a)$ for a given final state depends on the couplings only through the ALP lifetime, which determines the probability of the ALP to decay in a given part of the detector. Again, this makes it possible to calculate the detection probabilities $p_{\text{det},f}$ in advance and then apply an appropriate rescaling through the branching ratios $\text{BR}_{a \rightarrow f}$.

Putting everything together, we can therefore write

$$\begin{aligned} N_{\text{det}} &= \sum_{i,f} f_i(\mathbf{C}, \mathbf{C}_{\text{ref}}) \text{BR}_{a \rightarrow f}(m_a, \mathbf{C}) \times \int d\theta_a dE_a \frac{d^2N_i(m_a, \mathbf{C}_{\text{ref}})}{d\theta_a dE_a} p_{\text{det},f}(m_a, \Gamma_a, \theta_a, E_a) \\ &= \sum_{i,f} M_{if}(m_a, \mathbf{C}) \times E_{\text{det},if}(m_a, \Gamma_a), \end{aligned} \quad (2.7)$$

where M_{if} denotes the part that depends on the ALP model (i.e. on the couplings \mathbf{C}), while $E_{\text{det},if}$ encapsulates all the experimental details. Note that in the final line we have omitted the dependence on \mathbf{C}_{ref} , which will be kept implicit in the following.

As will become clear in the subsequent sections, the calculation of $E_{\text{det},if}$ for a given experiment can be quite challenging, typically requiring Monte Carlo (MC) simulations of both the production and the decay. However, these simulations only have to be performed once (for given m_a and Γ_a and a choice of \mathbf{C}_{ref}) and the (tabulated) functions $E_{\text{det},if}$ can easily be made publicly available. This makes it possible to perform fast analyses of any ALP model for which the model-dependent functions M_{if} have been calculated. Moreover, simulation upgrades including new production mechanisms or decay channels or more realistic experiment modelling can be quickly included in the framework, allowing straightforward estimates of the related phenomenological consequences.

3 ALP production

In this work we consider three types of ALP production channels, which are described in detail in the remainder of this section: production via the effective coupling of ALPs to photons (section 3.1), ALP mixing with neutral pseudoscalar mesons (section 3.2) and ALPs produced in heavy meson decays (section 3.3). Each of these production channels is efficient for production of ALPs in a different mass range and results in a different E_a - θ_a distribution, thus favouring a different experimental setup. For example, ALPs produced via photon-fusion tend to be forward-emitted and soft; those produced in heavy meson decays can be emitted at large angles and with higher energies.

As discussed in the previous section, the differential ALP yield for each of the production channels can be split into a model-dependent and a model-independent part. For production via the ALP-photon coupling, the model-dependent part is given by eq. (2.5). For production via meson mixing, the corresponding expression is

$$f_P(\mathbf{C}, \mathbf{C}_{\text{ref}}) = \left| \frac{\theta_{aP}}{\theta_{aP,\text{ref}}} \right|^2, \quad (3.1)$$

where $P = \pi^0, \eta, \eta'$ and θ_{aP} denotes the corresponding ALP-meson mixing angle (see below). Finally, for the production via B meson decay we have

$$f_{BK^{(*)}}(\mathbf{C}, \mathbf{C}_{\text{ref}}) = \frac{\text{BR}(B \rightarrow K^{(*)} + a)}{\text{BR}(B \rightarrow K^{(*)} + a)_{\text{ref}}}, \quad (3.2)$$

where

$$\text{BR}(B \rightarrow X) = \frac{\Gamma_{B \rightarrow X}}{\Gamma_B}, \quad (3.3)$$

denotes the B meson branching ratios, and analogous expressions for D meson decays.

3.1 Photo-production

The ALP photo-production, also called Primakoff process, can be realized in two ways at a beam-dump setup: the electro-magnetic (EM) scattering of a beam proton with a target nucleus and the interaction of on-shell photons from decays of secondary neutral pseudoscalar mesons with the nuclear EM field. The former process can be in turn either realized through the scattering of the whole proton on the target nuclei (elastic scattering) or through the scattering of their constituents. The elastic scattering, which dominates at beam energies of current beam dump facilities, has been studied in details in ref. [27]. In ref. [26] it has been shown that ALP photo-production via on-shell photons from meson decays constitutes an even a more important source of ALPs. In this work, we include both elastic scattering and Primakoff production from on-shell photons.

In both cases the ALP differential yield can be expressed as

$$\frac{d^2 N}{d\theta_a dE_a} = \frac{C_{\gamma\gamma}^2}{\Lambda^2} \frac{8e^4}{\sigma_N} \frac{\sin \theta_a}{\pi E_{\text{beam}}} \int dp_t^2 d\phi_a \frac{d\sigma_{\gamma N}}{d \cos \theta_a} \gamma(E_a, p_t^2), \quad (3.4)$$

where the differential cross section for the photon-nucleus interaction is taken from eq. (3.16) of ref. [27] and scales with Z_{target}^2 . This cross section has to be normalized to the total cross

section for the incident particle. In the case of elastic proton scattering, the total cross section for the incident particle is $\sigma_N \simeq 53 \text{ mb} \times A_{\text{target}}^{0.77}$. For the absorption of an on-shell photon, we use values from ref. [31].

The photon distribution γ can be calculated using the Weizsäcker-Williams approximation [32] in the case of the proton elastic scattering.¹ In order to obtain the γ distribution for on-shell photons from π^0, η, η' di-photon decays, we have to simulate the production of these secondary mesons. In this work the meson yields from p - p collisions at experimentally relevant energy values² were simulated using PYTHIA 8.2 [34] with `SoftQCD:all` flag and a pomeron flux parametrization `SigmaDiffractive:PomFlux(5)`, a setup which has been validated in ref. [26]. The corresponding photon distribution is then obtained from interpolation of the dataset resulting from the flat phase-space decay of these mesons into photon pairs. The decay products are boosted to the laboratory frame and the events are weighted with the π^0, η, η' di-photon branching fractions.

3.2 ALP production from meson mixing

If the ALP field is coupled to the SM quark and/or gluon fields, it naturally enters the kinetic and mass matrices of the effective Lagrangian of pseudoscalar mesons. Upon diagonalization of these matrices in order to obtain the mass eigenstates, the mixing angle between ALPs and mesons $P = \pi^0, \eta, \eta'$ emerges in the following form:

$$\theta_{aP} = \frac{f_\pi K_{aP} m_a^2 + m_{aP}^2}{f_a (m_a^2 - m_P^2)}, \quad (3.5)$$

where the ALP decay constant f_a depends on the underlying ALP couplings. We emphasize that this expression is valid for values m_a such that $|\theta_{aP}|^2 \ll 1$. For the explicit form of f_a , the kinetic mixing terms K_{aP} and the mass mixing terms m_{aP} we refer to appendix A and ref. [18].

In the presence of such a mixing it is possible to produce an ALP instead of the meson P via SM processes in a small fraction of cases proportional to $|\theta_{aP}|^2$. The corresponding differential yield can be written as

$$\frac{d^2 N}{d\theta_a dE_a} = |\mathcal{F}_{\text{VMD}}(m_a)|^2 \sum_P \frac{\hat{\alpha}_s^2(m_a)}{\hat{\alpha}_s^2(m_P)} \frac{d^2 N_P}{d\theta_P dE_P} \Big|_{\substack{E_P \rightarrow E_a \\ \theta_P \rightarrow \theta_a}} |\theta_{aP}|^2, \quad (3.6)$$

where $d^2 N_P/d\theta_P dE_P$ denotes the differential yield of the meson P . To consider ALP masses up to 3 GeV and capture the suppression of ALP production for $m_a \gg m_P$, we

¹It has been pointed out in ref. [33] that this approximation becomes inaccurate for small ALP masses as well as for low beam energies. Given that elastic scattering only gives a sub-dominant contribution compared to ALP production from photons produced in meson decays, we find that this approximation is nevertheless fully sufficient for our purposes.

²For 400, 120 and 70 GeV proton beams the respective p - p c.m. energies are approximately 27.4, 15 and 11.5 GeV.

follow ref. [18] and include the vector meson dominance (VMD) form factor

$$\mathcal{F}_{\text{VMD}}(m) = \begin{cases} 1, & \text{for } m \leq 1.4 \text{ GeV} \\ \sum_{i=0}^3 a_i m^i, & \text{for } 1.4 \text{ GeV} < m \leq 2 \text{ GeV} \\ (1.4 \text{ GeV}/m)^4, & \text{for } 2 \text{ GeV} < m \leq 3 \text{ GeV} \\ 0, & \text{for } m > 3 \text{ GeV}, \end{cases} \quad (3.7)$$

and the running strong coupling $\hat{\alpha}_s(m)$ in the form

$$\hat{\alpha}_s(m) = \begin{cases} 1, & \text{for } m \leq 1 \text{ GeV} \\ \sum_{i=0}^3 b_i m^i, & \text{for } 1 \text{ GeV} < m \leq 1.5 \text{ GeV} \\ 4\pi/7 \log(m/0.34 \text{ GeV})^2, & \text{for } m > 1.5 \text{ GeV}, \end{cases} \quad (3.8)$$

where parameters a_i and b_i are determined by requiring continuity of the resulting function and its first derivative. Note that compared to ref. [4] we assume a stronger suppression of the ALP yield for large ALP masses to ensure that our estimates are conservative.

For obtaining the distribution of $d^2 N_P/d\theta_P dE_P$, we can conveniently re-use the validated datasets from the simulation described in section 3.1. There is however an important subtlety: since the pseudoscalar meson masses are generally different from the ALP mass, it is not possible to simply replace a pseudoscalar by an ALP with the same energy and momentum. Even if the replacement is made in such a way that the three-momentum is conserved in a specific frame, it may not be conserved in a different frame. This introduces some degree of arbitrariness in the precise prescription used for the replacement. Here we require that the three-momentum of the ALP be equal to the three-momentum of the pseudoscalar meson in the p - p *cms frame*, calculate the corresponding ALP energy and then boost the result into the laboratory frame. For further details we refer to appendix D.

In figure 1 we show the predicted distribution of ALPs in the E_a - θ_a plane for the photon-from-meson production channel (left panel) and for production via ALP- η mixing (right panel). In both cases we find an anti-correlation between the two parameters in the sense that higher ALP energies imply smaller production angles and vice versa. For the meson mixing contribution, the differential yield peaks at larger ALP angles than for Primakoff production, but still gives a sizeable contribution in typical on-axis beam-dump experiments with an angular acceptance of a few mrad.

3.3 ALP production from rare decays

Flavour-changing ALP couplings arise not only from interactions of ALPs with quarks [10] and electroweak gauge bosons [15] but also in higher-order processes from ALP-gluon interactions [35]. In all of these cases ALPs can be produced in flavour-changing neutral current (FCNC) processes and can therefore be probed through rare meson decays in fixed-target experiments [19, 24]. However, in order for the produced ALPs to be focused in the forward direction it is essential that the parent meson decays before interacting with the target material. B and D mesons can thus be of particular importance since when

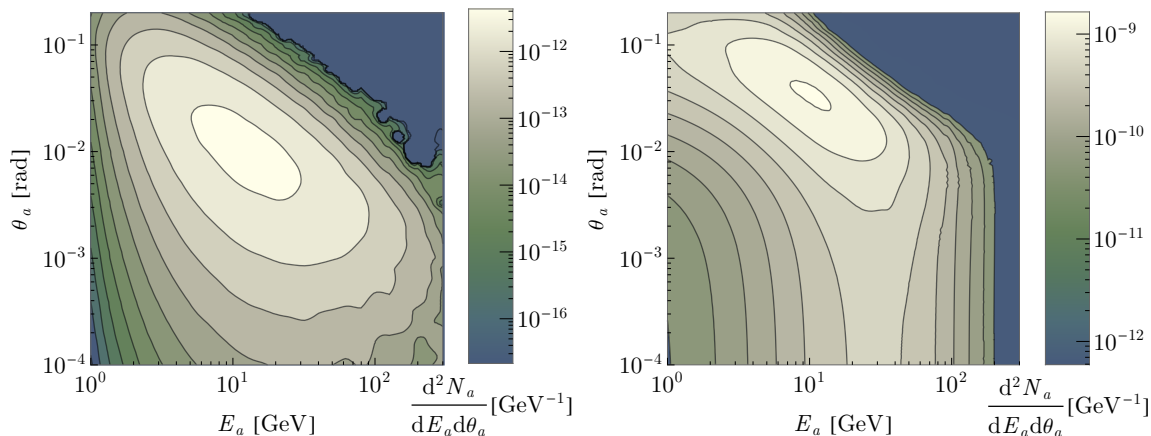


Figure 1. Production differential yields for an $m_a = 0.5$ GeV ALP produced at a copper target with a 400 GeV proton beam. For a *photon-from-meson* induced production (left) with a fixed $C_{\gamma\gamma}/\Lambda = 10^{-4} \text{ GeV}^{-1}$ and for an *a- η* mixing (right) assuming a fixed mixing angle $\theta_{a\eta} = 10^{-4}$.

produced at the beam energies of $\mathcal{O}(100)$ GeV their decay length is significantly smaller than the interaction length of the target material.³

Let us first consider the case of B mesons. In this work we include the ALP production from $B \rightarrow K^{(*)} + a$ decays (both charged and neutral). Since the decay is isotropic in the rest frame of the B meson, the differential ALP yield depends only on the distribution of B mesons and the branching ratios for the decays involving ALPs (see eq. (3.3)). The B meson distributions are again obtained using simulations of p - p collisions with PYTHIA 8.2. We follow the setup of ref. [24], i.e. we allow only the bottom quark production hard-QCD processes in order to have sufficient statistics and then reweight the final distribution according to the bottom production cross section at the given beam energies. The value reported by PYTHIA is $\sigma_{pp} = 39.85 \times 10^9$ pb and $\sigma_{bb} = 1.866 \times 10^3$ pb for 400 GeV beam, which has already been validated in ref. [24] using the values reported in ref. [37]. For 120 (and 70 GeV) the values are $\sigma_{pp} = 38.54 \times 10^9$ pb and $\sigma_{bb} = 3.8$ pb ($\sigma_{pp} = 38.38 \times 10^9$ pb and $\sigma_{bb} = 1.15 \times 10^{-7}$ pb respectively), for which no measurement has been found in the existing literature for validation. In any case, for experiments operating with such small beam energies the ALP production via B meson decays is found to be negligible (see section 6). A separate simulation involving the two-body decay kinematics of $B \rightarrow K^{(*)} + a$ is performed to obtain the resulting ALP distribution as a function of E_a and θ_a .

In an analogous way, we simulate the production of D mesons with

$$\sigma_{cc} = \begin{cases} 3.601 \times 10^6 \text{ pb,} & \text{for } E_{\text{beam}} = 400 \text{ GeV} \\ 0.518 \times 10^6 \text{ pb,} & \text{for } E_{\text{beam}} = 120 \text{ GeV} \\ 1.551 \times 10^5 \text{ pb,} & \text{for } E_{\text{beam}} = 70 \text{ GeV} \end{cases} \quad (3.9)$$

and a subsequent decay $D \rightarrow \pi + a$. Despite a larger production cross section of D mesons, for the ALP models considered in this work the branching ratio $\text{BR}_{D \rightarrow \pi+a}$ is many orders of

³The nuclear collision length for most of the target materials is ≈ 10 cm [36]. The decay length for a 100 GeV momentum B meson is < 1 cm and ≈ 1 cm for a D meson, while for K it is > 500 cm.




Figure 2. Production differential yield of an $m_a = 0.5$ GeV ALP produced in decays $B^{\pm,0} \rightarrow K^{\pm,0} + a$ (left) and $D^{\pm,0} \rightarrow \pi^{\pm,0} + a$ (right) of B and D produced at a copper target with 400 GeV proton beam assuming a referential value $\text{BR}_{B \rightarrow K+a} = \text{BR}_{D \rightarrow \pi+a} = 10^{-10}$.

magnitude smaller than $\text{BR}_{B \rightarrow K^{(*)}+a}$ as discussed in appendix A. Nevertheless, we enclose the D mesons data sets to allow also studies of specific ALP models enhancing the up-type quark transition (see e.g. ref. [28]).

The ALP distribution resulting from $B \rightarrow K + a$ and $D \rightarrow \pi + a$ decays are plotted in figure 2. Compared to the distributions shown in figure 1 we observe that the ALP distribution from rare meson decays peaks at even larger angles and energies. In particular ALPs produced from B mesons have very high energies, making this production mode particularly promising for off-axis experiments searching for relatively heavy ALPs with short lifetimes.

4 ALP decay and detection of final states

To determine the probability that an ALP induces an observable signal, we have modelled the various fixed-target experiments under consideration in a simplified⁴ and generic MC simulation of ALP propagation and decays. The layout consists of a decay volume, a spectrometer, analyzing magnets and a forward calorimeter (see schematic drawing in figure 3). Our set-up is described in detail in section 4.1.

The outcome of the simulation is the function $E_{\text{det},if}(m_a, \Gamma_a)$ defined in eq. (2.7). We provide tabulated data sets of $E_{\text{det},if}$ based on a simulation with $N = 10^6$ trials on a grid of Γ_a and m_a values for each of the production and decay channels i, f , and for each of the experiments described in section 5. We note that the function $E_{\text{det},if}$ includes the efficiency $\epsilon_{\text{det},f}$ of the experiment for the detection of a given final state f . For past experiments we use the efficiencies as determined by the experiment (see section 5), while for future experiments we simply set $\epsilon_{\text{det},f} = 1$. However, it is straight-forward in our approach to

⁴For example we do not include veto detectors in the simulation and we do not account for efficiencies of individual detectors and possible scattering or absorption of final state particles in the material of the experiment before they are detected.




Figure 3. Schematic view of the experimental setup for the case that the ALP decays via $a \rightarrow \pi^+ \pi^- \gamma$.

redo our analysis for different values of $\epsilon_{\text{det},f}$ making our framework a powerful tool to assess the impact of the experimental design on sensitivity projections.

4.1 General principles of the ALP Monte Carlo simulation

The probability $p_{\text{det},f}$ that an ALP with given m_a , Γ_a , θ_a and E_a decaying to channel f will be detected at the experiment is a potentially complicated function of the distance l that the particle travels before decaying and the momenta of the final-state particles, which can be characterised by a set of kinematical variables a_1, \dots, a_n . The conditions imposed on the decay products by experimental cuts can be summarized by a function $\Theta_{\text{det},f}(m_a, \theta_a, E_a, l, a_1, \dots, a_n)$, which is equal to 1 if all experimental conditions are met and 0 otherwise. Since it is typically impossible to perform the integration over all of these variables, it is useful to define $p_{\text{det},f}$ instead as a sum over randomly generated decay processes

$$p_{\text{det},f}(m_a, \Gamma_a, \theta_a, E_a) = \sum_{k=1}^N \epsilon_{\text{det},f} \frac{\Theta_{\text{det},f}(m_a, \Gamma_a, \theta_a, E_a; \{l, a_1, \dots, a_n\}_{\text{ran}}^k)}{N}, \quad (4.1)$$

where $\{l, a_1, \dots, a_n\}_{\text{ran}}^k$ denotes the k -th simulated event of the ALP decay.

To carry out these simulations we first calculate the ALP decay length in the laboratory frame $l_a = \gamma_a \beta_a / \Gamma_a$ with $\gamma_a = E_a / m_a$ and $\beta_a = p_a / m_a$ and use the distribution $p(l) = l_a^{-1} \exp(-l/l_a)$ to choose a random value of l . For the decay itself we generate the four-momenta of the second generation particles, as well as extra sets for the third generation particles if the former include unstable particles (the specific distributions used to simulate the ALP decays will be discussed in section 4.2). Given the magnetic fields of the experiment, these four-momenta enable us to calculate the trajectories of the final-state particles and check the experimental conditions in order to evaluate $\Theta_{\text{det},f}$.

In practice, in order to obtain the distribution $E_{\text{det},if}(m_a, \Gamma_a)$ for given values of m_a and Γ_a we first read the distribution $\Delta^2 N_i / \Delta \theta_a \Delta E_a$ for the production mode i from the previous section into a θ_a - E_a histogram using ROOT [38]. At each iteration we sample a random bin $\{\theta_a, E_a\}_{\text{ran}}^k$ according to the histogram and simulate the ALP decay for these values. By summing up a sufficiently large sample N , we then obtain a reliable estimate of $E_{\text{det},if}(m_a, \Gamma_a)$ in the form

$$E_{\text{det},if}(m_a, \Gamma_a) = \sum_{k=1}^N \epsilon_{\text{det},f} \frac{\Theta_{\text{det},f}(m_a; \{\theta_a, E_a, l, a_1, \dots, a_n\}_{\text{ran}}^k)}{N}. \quad (4.2)$$

	<i>neutral</i>	<i>charged</i>
2-body	2γ	e^+e^- $\mu^+\mu^-$
3-body	$3\pi^0$ $2\pi^0\eta$	$\pi^+\pi^-\gamma$ $\pi^+\pi^-\pi^0$ $\pi^+\pi^-\eta$

Table 1. Allowed decay channels for ALPs with masses up to about 1 GeV.

Note that the right-hand side depends implicitly on Γ_a through the distribution used to obtain l . Although the ALP detection probability depends sensitively on the ALP energy E_a and angle θ_a , the left-hand side of the equation above is independent of these quantities and only depends on the chosen production mechanism (represented by the index i). In other words, $E_{\text{det},if}$ can be thought of as the expectation value of the detection probability for the distribution of ALP energies and angles predicted by the chosen production mechanism.

4.2 Decay channels

Since beam dump experiments can probe ALPs of masses up to $\mathcal{O}(1)$ GeV, there are a number of kinematically allowed decay channels, listed in table 1. We restrict ourselves to two- or three-body decays, therefore accounting for the dominant contribution to the total decay width for the ALP masses and coupling scenarios under consideration (see section 6).

For the purposes of the MC simulation, we consider charged particles to be stable and π^0 and η to decay instantaneously into a γ - γ pair with the appropriate branching fraction. While two body decays of ALPs are treated with a flat phase space, for three-body decays a more careful treatment is necessary.

In a three-body decay, the invariant masses of first and second and second and third particles, m_{12} and m_{23} , can be used to define a Dalitz plot.⁵ If the transition amplitude square $|\mathcal{M}|^2$ of the process depends on the momenta of outgoing particles, the resulting Dalitz plot density is generally non-uniform. Such a momentum dependence arises in a wide range of models, for example if the decay proceeds via a resonance, as in the case of $a \rightarrow \pi\pi\gamma$, which receives a contribution from virtual ρ meson exchange.

In this work, for the simulation of hadronic decays, we employ the framework derived in ref. [18] using chiral perturbation theory. The results from ref. [18] are derived under the assumption that ALPs are coupled only to gluons, but they remain valid also for ALPs with flavour-universal couplings to light quarks. In such a case, the decay width $\Gamma_{a \rightarrow \text{hadrons}}$ changes but the ratio $|\mathcal{M}|^2_{a \rightarrow \text{hadrons}}/\Gamma_{a \rightarrow \text{hadrons}}$, which determines the distribution of the final-state hadrons, remains invariant. However, for models with non-universal couplings to light quarks, as considered for example in ref. [29], there may be significant differences in both the individual branching ratios and the corresponding distributions of the decay products.

⁵In all of the decays that we consider, there are at least two particles with the same mass. These are taken to calculate m_{12} .




Figure 4. Dalitz plot density for six values of ALP masses used for the $a \rightarrow \pi^+\pi^-\eta$ decay.

For a given ALP mass m_a , we use $|\mathcal{M}|_{a \rightarrow \text{hadrons}}^2 / \Gamma_{a \rightarrow \text{hadrons}}$ from ref. [18] to weight a flatly distributed Dalitz-plot density. The procedure is repeated for each m_a bin. As an example, the density for the decay $a \rightarrow \pi^+\pi^-\eta$ is shown in figure 4.

4.3 Propagation of decay products

The photons, either emitted as daughter particles of the ALP or as tertiary decay products, are propagated in straight line up to the upstream face of the forward calorimeter. When extrapolating any charged particle, we account for the bending induced by the spectrometer analyzing magnets, if present: we treat the field as that of an ideal solenoid and do not model effects due to detailed field maps, stray fields, etc.

Photons and e^\pm are assumed to release all their energy in the calorimeter. No resolution effects are simulated and any experimental condition is directly applied on the photon or e^\pm energy. A minimum separation between photon or e^\pm showers is required for “cluster counting”: a pair of energy releases is merged into a single “cluster” if the relative distance of the calorimeter impact points is below an experiment-dependent radius. In this case, the cluster energy is the energy sum of the two merged particles. For cluster counting, the particle impact point is required to be within the calorimeter sensitive region and an experiment-dependent minimum energy is required. Muons are extrapolated forward up to the front face of a muon detector, if present. For muon detection, the muon impact point at the front face of the muon detector is required to be in the sensitive region. The presence

Experiment	Status	E_{beam} [GeV]	N_{POT} [10^{18}]	Target	$l_{\text{n,DV}}$ [m]	$l_{\text{c,DV}}$ [m]	z_{DV} [m]	θ_{off} [mrad]	Ω_{cov} [μsr]
CHARM	completed	400	2.4	Cu	35	35	480	10	34
NuCal	completed	70	1.7	Fe	23	23	64	0	700
NA62	running	400	10	Cu	139	81	82	0	84
DUNE ND ⁶	proposed	120	1100	C	10	10	575	0	36
DarkQuest	proposed	120	1.44	Fe	13.5	1	5	0	12000
SHADOWS	proposed	400	10	Cu	26	23	10	75	9200
SHiP	proposed	400	200	Mo	50	28	45	0	4500

Table 2. Overview of the basic parameters for the experimental setups considered in this work. See text for details and definitions.

of a passive shield in front of a muon detector is only accounted for by applying a loose minimum-energy threshold.

It should be emphasized that we consider ALP decays into photons not only in the decay volume but also in the spectrometer area, effectively enlarging the decay volume up to the calorimeter location.

5 Experimental framework

In this section we briefly discuss how we have modelled the experimental setups considered in this work and the related acceptance conditions. Table 2 provides an overview of the basic parameters used for each experiment. A beam of protons with energy E_{beam} is made to interact with an absorber based on the element listed in the column “Target”. The total number of protons on target (POT) is denoted by N_{POT} . The experiment should be sensitive to ALPs produced in the target and reaching a decay volume (DV) located z_{DV} meters downstream. The longitudinal axis of the DV has an angle θ_{off} with respect to the beam axis, and Ω_{cov} denotes the solid angle covered by the calorimeter. Neutral (charged) ALP-daughter particles can be detected if the ALP decays within $l_{\text{n(c),DV}}$ meters from the DV entrance. Details are given in the following subsections.

5.1 Results from past experiments

To explore future opportunities, first of all we need to recast the results of the past CHARM and NuCal experiments in terms of their sensitivity for ALPs.

5.1.1 CHARM

For CHARM (using a copper target), an off-axis search for two photons in the final state was presented in ref. [39], with an average efficiency of 0.51. The search required at least

⁶Since the operational time is not yet decided, we evaluate the data sets for one year of data-taking at DUNE. Our data sets can be easily re-scaled and the resulting contours in 6 correspond to ten years of operation. In ref. [4], a ten year operation is considered with a larger number of POT integrated per year than that considered here. We assume the DUNE near detector to be operated on-axis. Changing the detector vertical position in our simulation is straightforward and can be performed on request.

one photon detected at the calorimeter. The sensitive area of the calorimeter has transverse dimensions of $3\text{ m} \times 3\text{ m}$. In the analysis quoted, the shower energy was required to be between 5 and 50 GeV and this condition has been implemented in our model. The same reference includes results for the detection of final states with two daughter muons.⁷ The detection efficiency is reported to be 0.85. In our simulation, the muons are required to have at least 1 GeV at the final detector plane in order to be seen as minimum ionizing particles throughout the detector material.

Regarding hadronic final states, we are only aware of on-axis searches, which go hand in hand with a larger background rate. Ref. [40] reports around 80 muon-less events based on a statistics of 7×10^{17} POT [41], after requiring a shower-like energy release of more than 20 GeV. According to the quoted analysis, 4 ± 21 events can be attributed to the interaction of other neutrinos or neutrino-like particles. Due to the limitation in statistics and due to the uncertainty induced by the subtraction of the background, the strongest exclusion by CHARM on hadronic decay modes is again given by a re-interpretation of the search for decays to final states with emitted photons (ref. [39]): final states such as $\pi^+\pi^-\gamma$, $\pi^+\pi^-\pi^0$, and $\pi^+\pi^-\eta$ are considered and at least one and at most two photons are required to be detected in the calorimeter, while all of the charged hadrons are required to escape the volume.

5.1.2 NuCal

The sensitivity of NuCal to final states with photons was evaluated in ref. [42], requiring the detection of one electromagnetic shower. While NuCal operated with a 70 GeV proton beam, it benefited from a relatively small distance between target and detector (64 m) and from a comparably large decay volume length (23 m). After rejecting hadronic-like showers, NuCal observed 1 event compared to a background expectation of 0.3 events. ALP signals corresponding to a prediction of more than 3.6 events are excluded with a confidence level of 90%. For the detection of a two-track final state, we follow the procedure from ref. [43] and require both tracks to reach a circle of 2.6 m diameter at the end of the decay volume with a combined energy of at least 10 GeV. As in ref. [43], we assume an overall detection efficiency of 0.7 for the photonic final state, and 0.8 for the di-muon final state.

By contrast, we are not aware of corresponding results for hadronic final states in NuCal. Thus, we proceed similarly as for CHARM: we reinterpret the search for final states with photons as appropriate (i.e. if the final state contains at least one γ , π^0 , or η) and require that the charged hadrons escape detection.

5.2 Acceptance of current experiments

The only operational proton beam experiment that we will consider in the following is NA62. Some first data taken in beam dump also exist from the KOTO experiment [44]. The modelling of KOTO might be easily added to our code. Doing so would however require us to re-run and validate the PYTHIA simulation for a 30 GeV beam energy, which is non-trivial because of limited available data.

⁷This is in contrast to the simulation performed in ref. [24], where it was mistakenly assumed that CHARM would be sensitive also to final states with just *one* muon.

5.2.1 NA62

The primary goal of NA62 is the precise measurement of the rare decay of $K^+ \rightarrow \pi^+ \nu \bar{\nu}$, but the experiment is also sensitive to a variety of exotics scenarios. To search for neutral final states, NA62 has to be run in beam-dump mode, otherwise the residual background can limit the sensitivity. For a dump-mode run, the standard T10 target of NA62 must be removed, and the movable TAX collimators located 22 m downstream of T10 must be closed and used to dump the proton beam. This operation has been validated during the 2016–2018 data-taking and about 10^{16} POT have been collected in beam-dump mode. Triggers sensitive to decay channels with neutral and charged daughter particles have been deployed. During the current 2021–2023 run $\sim 10^{18}$ POT will be collected in beam dump mode. Around 10^{17} POT have been collected in 2021 [45]. Here we discuss the implications of collecting data in beam dump mode in the data-taking period between the so-called long-shutdown 3 (LS3) and LS4 to obtain $\sim 10^{19}$ POT.⁸ We assume no background limitation, in line with the present knowledge based on 2021 data [45].

We model NA62 in a toy MC as follows: the beam-defining collimator for ‘regular’ data-taking is used to dump the beam and its distance to the start of the fiducial volume is 82 m. A decay region in vacuum contains a spectrometer with a first station 81 m downstream of the fiducial volume entrance. The spectrometer hosts a 0.7 T solenoidal magnet 95 m downstream of the entrance of the decay volume. The Liquid Krypton Calorimeter (LKr) is located about 139 m downstream of the entrance of the decay volume and it is modelled as an octagon, see e.g. ref. [47].

For the neutral final states, we require the following acceptance conditions: both photons from the ALP decay must be detected at a minimum mutual distance of 10 cm at the LKr plane. In addition, the energy of each photon must be above 1 GeV and each photon impact point at the LKr must be at least 15 cm away from the LKr central hole through which the beam-pipe passes. Finally, the energy sum of the two photons must exceed 3 GeV. The NA62 target material (i.e. the material of the upstream section of the final collimator) is copper.

For the di-muon decay, we require both muons to be in the acceptance of the NA62’s first and last tracking stations. The acceptance loss due to the central hole allowing the passage of the beam pipe is modelled, cf. also ref. [24]. Each of the muons should have at least 5 GeV in order to be tracked efficiently.

Finally, for hadronic final states that include photons, not only the π^+ and π^- must be tracked, but also the additional requirement is made that all photons reach the calorimeter and deposit at least 3 GeV of combined energy.

5.3 Acceptance of proposed experiments

To evaluate the sensitivity of future experiments, we base our analysis on the available proposals. Some of the endeavours we considered are more defined, while others might be subject

⁸Note that there are several projects in consideration for NA62 in the post-LS4 period which would allow collection of even larger statistics [46]. These are not considered in this work.

even to significant change in the future. Moreover, some of the proposed experiments may face non-negligible backgrounds, which we do not attempt to estimate in the present work.

5.3.1 SHiP

Since the envisaged geometry for SHiP has changed since the publication of the SHiP proposal [48], we follow the layout of ref. [49] for our estimates. The prospects for detection of photons from ALP decays in the SHiP calorimeter is modelled as follows. The fiducial region is taken to be 45 m downstream of the production point of the ALPs and contains a spectrometer with the first spectrometer station located 50.76 m after the entrance of the decay volume and a 0.15 T magnet. The calorimeter is situated after the fourth spectrometer station,⁹ which is 60.16 m downstream of the entrance of the decay volume.

We ask both photons to be in an acceptance area of $5 \text{ m} \times 10 \text{ m}$. The energy of each photon must be above 1 GeV. The sum of the photon energies should be above 3 GeV. The photon impact points at the calorimeter should be at least 10 cm apart. The target material is molybdenum and the POT are 10^{20} . Note that the proposed SHiP calorimeter has the potential of reconstructing the photon direction, which allows for an ALP mass reconstruction.

To model the detection of charged particles in SHiP we mimic the steps provided in ref. [50]. The decay vertex must lay in the decay volume and the two tracks must hit all spectrometer chambers in their sensitive volume, with a minimum distance of 5 cm away from the decay vessel walls. Each particle must have a minimum energy of 5 GeV. Finally, for hadronic final states that include also photons, all photons should reach the calorimeter and deposit at least 3 GeV of combined energy.

5.3.2 DarkQuest

To make projections for DarkQuest, we consider the proposed extension of the SeaQuest experiment [51], particularly the proposed phase-I parameters. Details of the DarkQuest setup can be found in ref. [52]. The target material is iron. Following ref. [53], we assume for the di-photon final state the need to detect 10 signal events to claim a signal beyond the background fluctuations.

It is assumed that the experiment is upgraded with a calorimeter placed between tracking stations 3 and 4 (at $\sim 18.5 \text{ m}$ downstream of the target). The fiducial volume entrance is 5 meters downstream of the target.¹⁰ The first spectrometer station is located about 6 m from the target and a geometric acceptance of $2 \text{ m} \times 2 \text{ m}$ is assumed in the transverse direction. We assume using the “phase-I statistics” of 1.44×10^{18} POT.

As with the other setups, we require a minimum energy of 1 GeV for each detected photon and a total energy of at least 3 GeV and — given the photon shower Moliere radius

⁹The exact distance between the calorimeter and the last spectrometer station is not stated in the used literature, so we use an approximate value of 50 cm inferred from the schematics.

¹⁰While this work was undergoing finalization, an independent study of the DarkQuest sensitivity to ALPs has appeared [54], using 7-8 m as fiducial volume (following [53], whereas our fiducial volume definition follows [52]). Apart from this major difference in the acceptance simulation, there are slight differences regarding the overall required energy and the photon separation.

— a minimum mutual distance of 10 cm at the calorimeter plane to avoid shower overlap. For the di-muon final state, we model tracking in the simulation using the KMAG magnet and impose that the particles be in the acceptance of the first and third tracking stations. In addition we require 1 GeV of minimum energy for each muon so that it is not stopped in the iron absorber.¹¹ Following ref. [51], muon track pairs should be detected with a $\sim 60\%$ efficiency for di-muon masses above 4.2 GeV. Here, we do not consider the possibility to also detect hadronic final states, as the K_L background expected for the phase-I set-up is expected to be a limiting factor [52].

5.3.3 DUNE ND

Motivated by the recent results of ref. [4] we also include the DUNE near detector [55] (ND) in our study. DUNE ND uses a neutrino beam produced at the Long Baseline Neutrino Facility (LBNF) at Fermilab by a 120 GeV proton beam impacting on a graphite target with integrated intensity of 1.1×10^{21} POT per year [56]. We note that the nominal intensity can increase with further updates and the target material can be changed to beryllium [56]. Since DUNE ND is located underground separated by approximately 0.5 km of earth from the target, it can also effectively serve as an on-axis¹² proton beam dump experiment.

The essential parts of DUNE ND used in our analysis are the Liquid Argon detector (ND-LAr) and a multipurpose detector based on gaseous argon (ND-GAr). The ND-LAr is a 5 m long time projection chamber with height of 3 m and width of 7 m located 574 m downstream of the target. The ND-GAr located right downstream of the ND-LAr is composed of a 5 m long cylindrical high pressure gaseous time projection chamber of 5.2 m diameter surrounded by an electromagnetic calorimeter, in a 0.5 T magnetic field and a muon system. In this analysis we treat the ND-LAr and ND-GAr as two independent decay volumes. Following the reconstruction in ref. [55], we allow the ALP decays to happen in $6 \text{ m} \times 3 \text{ m} \times 2 \text{ m}$ volume of ND-LAr, excluding 50 cm from the sides and upstream and 150 cm downstream corresponding to roughly $5 \text{ cm} \times 9 \text{ cm}$ Moliere radius to allow a full reconstruction of the showers. For ND-GAr we assume a 4 m long, 2 m high and 4.8 m wide block located 580 m from the target and exclude the outer layers following the exclusion in the active volume calculation of ref. [55]. The photons in the 2γ decays in the ND-GAr must have a minimum 0.8 deg (14 mrad) separation angle and a minimum 20 MeV energy each. Charged hadrons from hadronic ALP decays must have a minimum kinetic energy of 5 MeV, based on the measured resolution [55]. For decays taking place in the volume of ND-LAr we do not put any conditions on the minimum energy until the expected performance of the detector based on the ProtoDUNE [57] results is known.

¹¹Note that this requirement goes slightly beyond that made in ref. [52], where no minimum energy requirement is made. There, instead, only the transverse p_T kick induced by KMAG and the third tracking station are considered.

¹²The DUNE ND will also have the opportunity to take data up to 60 mrad off-axis as so-called DUNE-PRISM. This option is not considered in this work since the amount of time allocated to off-axis operation is unknown. We would like to stress, however, that since in certain scenarios ALPs in beam dumps are dominantly produced off-axis, DUNE-PRISM can have a large potential for this type of hidden sector physics searches even if operated for a short period of time.

For muons in DUNE, following figure 2.30 of [55], we assume muons with energies below 1 GeV are reconstructed with full efficiency. Above those energies, the angle w.r.t. to the beam axis must be smaller than 40 deg (700 mrad).

5.3.4 SHADOWS

SHADOWS is a proposed off-axis experiment in the NA62 experimental cavern, which can take data concurrently with NA62 operated in beam dump mode. It would use a 400 GeV proton beam dumped in the copper TAXes and can take up to 10^{19} POT in the data-taking period between LS3 and LS4. The layout of the detector we have employed is based on the 2021 PBC proposal [58] and can be subject to changes in the near future. The decay volume starts 10 m downstream of the target and accommodates a 2.5 m long spectrometer with four tracking stations. The spectrometer is located 33 m downstream of the target. The distance among the first two stations is 50 cm. A 1.5 m gap between the second and the third station is foreseen to accommodate a 1 T dipole magnet. The fourth station is 50 cm downstream of the third. A calorimeter is located about 50 cm behind the fourth spectrometer station. The calorimeter sensitive area has a rectangular shape of $2.5 \text{ m} \times 2.5 \text{ m}$ and in our implementation the detector is off axis by $\theta_{\text{off}} = 75 \text{ mrad}$, so that the calorimeter center is shifted by 2.25 m from the beam axis. The calorimeter is likely to be succeeded by a muon detector¹³ since muons generated in the TAXes present the main background component based on the preliminary studies in ref. [58]. In this analysis we optimistically assume negligible background, which may overestimate the sensitivity. The request on muons is similar to that of NA62 and SHiP: two spectrometer chambers should be hit as well as a minimum energy of 5 GeV for each particle.

6 Results

Having obtained the model-independent functions E_{if} by combining the differential ALP yields from section 3 with the various ALP decay models discussed in section 4 and the experimental set-ups considered in section 5, we are now in the position to calculate the predicted number of ALPs in a given experiment for any ALP model for which the model-dependent functions $M_{if} = f_i(\mathbf{C}, \mathbf{C}_{\text{ref}}) \text{BR}_{a \rightarrow f}(m_a, \mathbf{C})$ are known. The necessary formalism has been worked out in great detail for a broad class of ALP effective theories that include interactions of ALPs with SM quarks and leptons, as well as gauge and Higgs bosons. In the present work we will illustrate this procedure for the case that the ALP interacts dominantly with SM gauge bosons:

$$\mathcal{L}_{a,\text{int}} = g'^2 \frac{C_{BB}}{\Lambda} a B^{\mu\nu} \tilde{B}_{\mu\nu} + g^2 \frac{C_{WW}}{\Lambda} a W^{\mu\nu} \tilde{W}_{\mu\nu} + g_s^2 \frac{C_{gg}}{\Lambda} a G^{\mu\nu} \tilde{G}_{\mu\nu}. \quad (6.1)$$

Here g' , g and g_s denote the hypercharge, weak and strong gauge couplings respectively, $X^\mu = B^\mu, W^\mu, G^\mu$ denote the corresponding gauge fields (with the group generator index suppressed), $X^{\mu\nu}$ denotes the field strength tensor and $\tilde{X}^{\mu\nu} = \frac{1}{2} \epsilon^{\mu\nu\rho\sigma} X_{\rho\sigma}$ denotes its dual.

¹³We place the muon detector plane 1 m behind the calorimeter plane.

Rather than varying all three couplings independently, we will consider four different benchmark scenarios:

- (i) *B dominance*: $C_{BB} \neq 0$; $C_{WW} = C_{gg} = 0$
- (ii) *W dominance*: $C_{WW} \neq 0$; $C_{BB} = C_{gg} = 0$
- (iii) *Gluon dominance*: $C_{gg} \neq 0$; $C_{BB} = C_{WW} = 0$
- (iv) *Co-dominance*: $C_{BB} = C_{WW} = C_{gg} \neq 0$

The first three scenarios are chosen to highlight the different production channels. In particular, the first scenario is similar to the frequently studied case of *photon dominance* (see e.g. ref. [26]), except that it includes additional interactions between ALPs and Z bosons, which are however of no relevance for fixed-target experiments. The second scenario additionally predicts FCNC processes involving ALPs, effectively enhancing the ALP production. The third scenario features ALP-meson mixing and decays of ALPs into hadronic final states. Finally, the fourth scenario, first proposed in ref. [19] and explored further in ref. [4] investigates the potential interplay between the different couplings.

In order to predict ALP signals in beam dump experiments, one first needs to obtain explicit expressions for the functions $f_i(\mathbf{C}, \mathbf{C}_{\text{ref}})$, which describe how the contributions from the different production channels depend on the fundamental interactions in eq. (6.1). These functions are readily available in the literature and reproduced in appendix A. The second step is to calculate the ALP decay length and branching ratios as a function of the couplings and the ALP mass. In most cases the dominant decay mode will be the one into a pair of photons, with decays into leptons only appearing at the one-loop level and therefore giving a negligible contribution.¹⁴ For scenarios with $C_{gg} \neq 0$ and $m_a > 3m_\pi$ there will also be a relevant contribution from three-body decays into hadronic final states. Additional details are provided in appendix B.

Once the functions M_{if} have been obtained, it is straight-forward to calculate the predicted number of ALP events N_{det} for each experiment as a function of the fundamental couplings and the ALP mass. We provide these numbers in tabulated form together with this work for convenience. In the remainder of this section we show the exclusion limits (projected sensitivities) for past (future) experiments at 90% confidence level. We first consider the case that only the di-photon final state can be detected and subsequently explore the improvement in reach that can be achieved when also hadronic three-body final states are considered.

6.1 Considering only di-photon decays

For both scenarios (i) and (ii) one finds $\text{BR}(a \rightarrow \gamma\gamma) \approx 1$ since the one-loop processes $a \rightarrow \ell\ell$ are suppressed by many orders of magnitude. Moreover, since we assume $C_{gg} = 0$ in both cases, the production via meson mixing vanishes ($f_P = 0$), while $f_{\gamma\gamma} \neq 0$ thanks to the contribution from C_{BB} and C_{WW} to the effective ALP-photon coupling (see appendix A).

¹⁴This conclusion might change when considering direct couplings of the ALPs to leptons. We provide all necessary data sets as well as the implementation in the code to consider such a scenario. Note however that we do not currently consider ALP decays via the Bethe-Heitler process [59], which may be relevant in the case of ALPs coupled dominantly to leptons.




Figure 5. 90% CL exclusions limits for scenarios (i) (left) and (ii) (right). The exclusion from past experiments which accounts for signal and background studies is indicated by a full contour while the future experiments, assuming a zero background hypothesis, are indicated by a corresponding line.

The only difference between the two scenarios is that for (i) there are no FCNCs and hence $f_{BK^{(*)}} = 0$, whereas they give an important contribution for (ii). Likewise, the contribution of $f_{D\pi}$ vanishes for (i) and gives a marginal contribution for (ii) (see appendix A).

The existing exclusion limits and projected sensitivities for the two scenarios are compared in figure 5.¹⁵ As expected, all experiments under consideration follow a similar pattern, with the shape of the (projected) exclusions dictated by the ALP decay length (which decreases for large couplings and large ALP masses) and the production cross section (which decreases for small couplings and large ALP masses). Another notable feature in these plots is the low-mass cut-off in sensitivity for most future experiments. This cut-off results from the required separation distance between the two photons in the final state, which decreases for smaller ALP masses and correspondingly larger boost factors.

One finds that in order to extend the reach towards smaller couplings, the decisive quantity is the assumed POT, which is largest for DUNE and SHiP. To make improvements towards larger couplings and masses, the beam energy and detector geometry are decisive. In particular, we observe that the contribution from rare B meson decays in (ii) is especially relevant for experiments with $E_{\text{beam}} = 400$ GeV and clearly favours off-axis experiments like SHADOWS as well as on-axis experiments with large angular coverage like SHiP.

In scenarios (iii) and (iv) the gluon coupling is non-zero and therefore all production channels considered in this work become relevant.¹⁶ Moreover, one finds $\text{BR}(a \rightarrow \gamma\gamma) \ll 1$ for $m_a > 1$ GeV where hadronic decays dominate [18]. The effects of ALP-meson mixing in both production and decay is clearly visible in figure 6 and imprints a “pole structure”

¹⁵Note that the slight differences in reach compared to ref. [26] can be mainly attributed to the inclusion of η' decays and of a wider range of meson momenta.

¹⁶Note that we include the two-loop contribution to $B \rightarrow K^{(*)} + a$ from ALP-gluon couplings first calculated in ref. [35]. See appendix A for details.

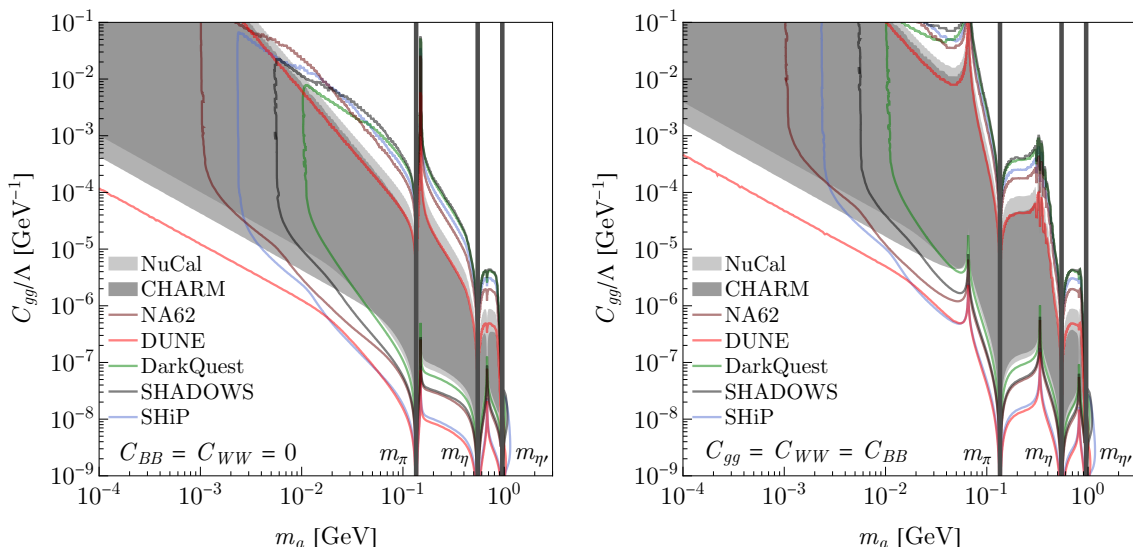


Figure 6. 90% CL exclusions limits for scenarios (iii) (left) and (iv) (right) assuming a di-gamma search only.

in the exclusion plots whenever the ALP mass approaches one of the pseudoscalar meson masses. Note that if the two masses become very close, the condition $|\theta_{aP}| \ll 1$ no longer holds. These regions are indicated by a dark shading and masked in our analysis.

In scenario (iv) there furthermore occurs a partial cancellation in the effective ALP-photon coupling, see eqs. (A.5) and (A.14), which has two effects: first of all, for small ALP masses ($m_a \lesssim 100$ MeV) it shifts all exclusion contours to somewhat larger couplings, which are necessary to achieve comparable ALP production yields. Second, it leads to additional “poles” outside of the masked regions whenever the different contributions to $C_{\gamma\gamma}$ cancel accidentally.

6.2 Including hadronic decays

For scenarios with ALP-gluon couplings the branching ratio into photons is strongly suppressed for heavy ALP masses and whenever there is a cancellation in the effective ALP-photon coupling. The incorporation of hadronic decays is therefore essential to probe these parameter regions and utilize the full potential of beam dump experiments for ALP searches. This can be seen in figure 7, which compares the exclusions and sensitivities for the di-photon channel only and the combination of all channels. We find that the improvement in sensitivity is most significant for SHiP, where the reach in terms of the ALP mass is extended by almost a factor of 2. The reason for this is simply that SHiP is the overall most sensitive experiment and therefore most dependent on the branching ratios of GeV-scale ALPs. Nevertheless, the effect of including hadronic final states is clearly visible also for all other experiments that we consider.

We emphasize that even experiments vetoing hadronic final states can partially cover regions with suppressed decays into photons thanks to hadronic decays with only photons in the final state and with charged hadrons escaping detection (provided that a partial reconstruction of the event is allowed). For experiments sensitive to a range of different

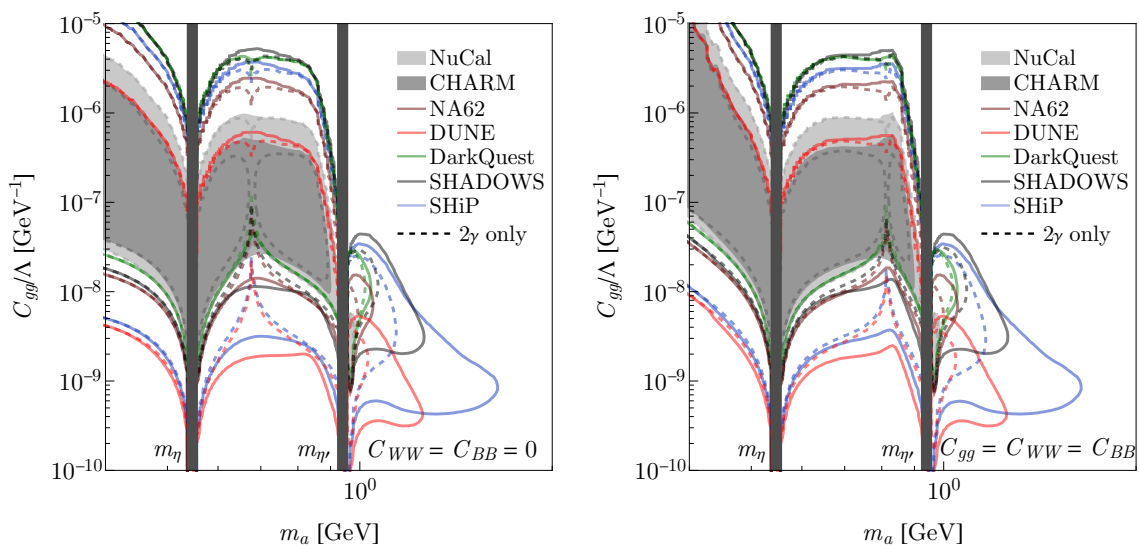


Figure 7. 90% CL exclusion limits for scenarios (iii) (left) and (iv) (right) given by a di-gamma decay only search (dashed line), compared with the exclusion when all mentioned hadronic channels are also taken into account (full line).

final states, on the other hand, our analysis highlights the exciting possibility that beam dump experiments may not only discover ALPs, but also infer their dominant branching ratios, giving us crucial clues regarding the underlying model.

7 Conclusions

The search for light feebly interacting particles at experiments with low centre-of-mass energy but high intensity is one of the most exciting frontiers of modern particle physics. A particularly attractive target are axion-like particles (ALPs), for which the smallness of the mass and the weakness of the interactions can simultaneously be motivated by their origin as Pseudo-Goldstone bosons of a spontaneously broken global symmetry. However, the wide range of possible ALP models raises the question how to systematically explore the parameter space and how to present experimental results in a model-independent way.

In this work we have addressed this challenge by developing a novel framework that allows to evaluate constraints and projected sensitivities from proton beam dump experiments for different ALP models without the need to re-run expensive Monte Carlo simulations. This is achieved by splitting the calculation into model-independent parts, which are computationally hard but can be performed and tabulated in advance, and model-dependent parts, which only require simple analytical rescaling functions.

Specifically, we have shown that ALP production can be written as a sum over a number of different production channels, each of which scales in a well-defined way with the effective ALP parameters at low energies. Most of these production channels require Monte Carlo simulations of the hadron shower created by the incident proton in the absorber and the subsequent decays of the resulting mesons. We have performed all necessary simulations using state-of-the-art tools and explicitly validated the resulting cross sections.

Regarding ALP propagation and decay, there are two main difficulties. The first one is to accurately simulate ALP three-body decays, which become important for GeV-scale ALPs with hadronic couplings. The second difficulty is to estimate the detection probability for a given final state in a wide range of experiments. To simulate three-body decays, we have developed a Monte Carlo technique based on the reweighting of Dalitz plots obtained from analytical calculations of hadronic ALP decays in chiral perturbation theory. Detector efficiencies are then estimated by propagating the decay particles through simplified detector geometries, which we have implemented for a range of past, present and future experiments.

To illustrate our approach we present exclusion limits and sensitivity projections for ALPs with couplings dominantly to Standard Model gauge bosons. To illustrate the impact of the detailed coupling structure, we consider four different scenarios, called B dominance, W dominance, gluon dominance and co-dominance. We find crucial differences between these four scenarios regarding the relative importance of different production and decay channels. In particular, for scenarios with gluon couplings, we find that experimental sensitivities can be significantly extended by including hadronic final states from three-body decays in the analysis.

The tabulated simulations and a Python script to perform the necessary rescaling for a user-defined combination of the various ALP interactions with SM particles are publicly available at <https://github.com/jjerhot/ALPINIST>. This code allows for the calculation of experimental constraints and sensitivities for a wide range of ALP models. Furthermore, it can be easily extended to include additional experiments, production and decay channels. Further details are provided in appendix C.

We emphasize that our implementation is still not fully model-independent. In particular the simulation of three-body decays presently assume the Dalitz plot densities predicted for ALPs with couplings to gluons and/or flavour-independent couplings to SM quarks. Different distributions are expected in more complicated new-physics scenarios, such as the one discussed in ref. [29]. We therefore aim to make also the code for Monte Carlo simulations publicly available, such that the user can easily modify decay distributions and experimental geometries. At the same time, we hope that independent Monte Carlo simulations will be developed by the experimental collaborations to produce similar tables to the ones used in our analysis based on more refined detector simulations. This will ultimately allow for a complete reinterpretation of experimental results in the general parameter space of ALP models.

Acknowledgments

We thank Yotam Soreq for useful discussions and Gaia Lanfranchi for comments on our implementation of SHiP and SHADOWS. This work is funded by the Deutsche Forschungsgemeinschaft (DFG) through the Collaborative Research Center TRR 257 “Particle Physics Phenomenology after the Higgs Discovery” under Grant 396021762 — TRR 257 and the Emmy Noether Grant No. KA 4662/1-1. This work is also supported through the European Research Council under grant ERC-2018-StG-802836 (AxScale project). JJ is supported through a FRIA grant by the F.R.S.-FNRS (Fonds de la Recherche Scientifique — FNRS), Belgium.

A Model-dependent calculations

Although our analysis in section 6 considers only couplings to B and W bosons and gluons (see eq. (6.1)), it is useful to start with a more general effective Lagrangian of ALP-SM interactions:

$$\begin{aligned} \mathcal{L}_a \supset & \frac{1}{2} \partial^\mu a \partial_\mu a - \frac{1}{2} m_a^2 a^2 + \sum_q \frac{C_{qq}}{2\Lambda} (\partial^\mu a) \bar{q} \gamma_\mu \gamma_5 q + \sum_\ell \frac{C_{\ell\ell}}{2\Lambda} (\partial^\mu a) \bar{\ell} \gamma_\mu \gamma_5 \ell \\ & + g^2 \frac{C_{WW}}{\Lambda} a W^{\mu\nu} \tilde{W}_{\mu\nu} + g'^2 \frac{C_{BB}}{\Lambda} a B^{\mu\nu} \tilde{B}_{\mu\nu} + g_s^2 \frac{C_{gg}}{\Lambda} a G^{\mu\nu} \tilde{G}_{\mu\nu}, \end{aligned} \quad (\text{A.1})$$

where Λ denotes the unknown scale of new physics (which may or may not be connected to the symmetry breaking scale [60]) and the parameters C denote the Wilson coefficients of the various effective operators. For the remainder of this discussion we will set $C_{qq} = 0$ at tree level. For a more detailed discussion of the effects of ALP-quark couplings, we refer to ref. [22].

At energy scales below EW symmetry breaking the most relevant interactions of ALPs with electroweak gauge bosons will then have the following form:

$$\mathcal{L}_{a,\text{EW}} \supset e^2 \frac{C_{\gamma\gamma}}{\Lambda} a F^{\mu\nu} \tilde{F}_{\mu\nu} + 2e^2 \frac{C_{\gamma Z}}{\Lambda} a F^{\mu\nu} \tilde{Z}_{\mu\nu} + 4g^2 \frac{C_{WW}}{\Lambda} a \varepsilon^{\mu\nu\alpha\beta} \partial_\mu W_\nu^+ \partial_\alpha W_\beta^-, \quad (\text{A.2})$$

where $e = g \sin \theta_W = g' \cos \theta_W$ and the Wilson coefficients are given by

$$C_{\gamma\gamma} = C_{WW} + C_{BB} \quad (\text{A.3})$$

$$C_{\gamma Z} = \frac{C_{WW} \cos^2 \theta_W - C_{BB} \sin^2 \theta_W}{\sin \theta_W \cos \theta_W}. \quad (\text{A.4})$$

At the one-loop level, $C_{\gamma\gamma}$ receives additional contributions from lepton and gauge boson loops, leading to

$$C_{\gamma\gamma} = C_{WW} + C_{BB} + \sum_\ell \frac{C_{\ell\ell}}{16\pi^2} B_1(4m_\ell/m_a) + \frac{2\alpha}{\pi} \frac{C_{WW}}{\sin^2 \theta_W} B_2(4m_W/m_a), \quad (\text{A.5})$$

where the functions $B_{1,2}(\tau)$ are defined in eq. (14) of ref. [20]. The analogous one-loop contributions to the effective leptonic coupling $C_{\ell\ell,\text{eff}}$ are given in eq. (26) of ref. [20], which is also implemented in our code.

A.1 FCNC transitions

Flavour-changing interactions are generated from the a - W interaction at the one-loop level [15] and from the a - g interaction at the two-loop level [35]. The corresponding effective interactions take the following form:

$$\begin{aligned} \mathcal{L}_{b \rightarrow sa} \supset & -\frac{C_{bs}}{\Lambda} (\partial^\mu a) \bar{s}_L \gamma_\mu b_L + \text{h.c.}, \\ \mathcal{L}_{c \rightarrow ua} \supset & -\frac{C_{cu}}{\Lambda} (\partial^\mu a) \bar{u}_L \gamma_\mu c_L + \text{h.c.} \end{aligned} \quad (\text{A.6})$$

The effective FCNC couplings can be written as $C_{qq'} = C_{WW}g_{qq',W} + C_{gg}g_{qq',g}$. The one-loop contribution from W bosons is given by

$$\begin{aligned} g_{bs,W} &= 3\alpha_W^2 \sum_{q=u,c,t} V_{qb}V_{qs}^* F(\xi_q), \\ g_{cu,W} &= 3\alpha_W^2 \sum_{q=d,s,b} V_{uq}V_{cq}^* F(\xi_q), \end{aligned} \quad (\text{A.7})$$

where $\alpha_W = g^2/(4\pi)$, $\xi_q = m_q^2/m_W^2$ and

$$F(x) = \frac{x[1+x(\ln x - 1)]}{(1-x)^2}. \quad (\text{A.8})$$

The two-loop contribution $g_{bs,g}$ has been calculated for the first time in ref. [35]:

$$g_{bs,g} = \frac{\alpha_s^2 \alpha_W}{4\pi} \left[f(\mu) + 2AC_F g(\mu) + 2BC_F \sum_{q=u,c,t} V_{qb}V_{qs}^* \xi_q \right] \quad (\text{A.9})$$

with functions $f(\mu)$ and $g(\mu)$ given in eqs. (B6) and (B7) of ref. [35]. The analogous expression for $g_{cu,g}$ is obtained by the obvious replacements. In our code, we associate μ with the new-physics scale Λ , which can be specified by the user as well as the model-dependent parameters A and B . For the results presented in section 6, we set $A = B = 0$, which was shown in ref. [35] to yield neither overly optimistic nor overly conservative exclusions.

We emphasize that both the contribution to the effective FCNC couplings from W bosons and the contribution from gluons are approximately proportional to ξ_Q , where Q is the heaviest quark in the loop. The resulting decay rates are then proportional to ξ_Q^2 . This leads to a strong suppression of the $c \rightarrow u + a$ transition (where $Q = b$) relative to the $b \rightarrow s + a$ transition (where $Q = t$), such that the process $B \rightarrow K^{(*)} + a$ turns out to be experimentally more important than $D \rightarrow \pi + a$ in spite of the greater D meson production cross section.

A.2 ALP-meson mixing

At energies below the QCD confinement scale, the ALP-gluon coupling gives rise to both kinetic and mass mixing of ALPs with the SM neutral pseudoscalars $P \in \{\pi^0, \eta, \eta'\}$. Removing this mixing requires a redefinition of the various fields:

$$P \rightarrow P + \theta_{aP} a \quad (\text{A.10})$$

with the mixing angle

$$\theta_{aP} = \frac{f_\pi}{f_a} \frac{K_{aP} m_a^2 + m_{aP}^2}{m_a^2 - m_P^2}, \quad (\text{A.11})$$

where $f_\pi = 93$ MeV and $f_a = \Lambda/32\pi^2 |C_{gg}|$ denote the pion and ALP decay constants.¹⁷

¹⁷We note that our definition of f_a and the ALP-meson mixing is not invariant under a chiral rotation of the SM quark fields, even though such a rotation should not affect physical observables. Rather than providing a more general treatment, we follow the common approach in the literature and assume that the chiral rotation is fixed in such a way that flavour-diagonal couplings to light quarks are absent.

The kinetic mixing terms are given by

$$\begin{aligned}
 K_{a\pi^0} &= -\frac{\kappa_u - \kappa_d}{2} \\
 K_{a\eta} &= -\frac{(\kappa_u + \kappa_d) \left(\frac{1}{\sqrt{3}} \cos \theta_{\eta\eta'} - \sqrt{\frac{2}{3}} \sin \theta_{\eta\eta'} \right) - \kappa_s \left(\frac{2}{\sqrt{3}} \cos \theta_{\eta\eta'} + \sqrt{\frac{2}{3}} \sin \theta_{\eta\eta'} \right)}{2} \\
 K_{a\eta'} &= -\frac{(\kappa_u + \kappa_d) \left(\frac{1}{\sqrt{3}} \sin \theta_{\eta\eta'} + \sqrt{\frac{2}{3}} \cos \theta_{\eta\eta'} \right) - \kappa_s \left(\frac{2}{\sqrt{3}} \sin \theta_{\eta\eta'} - \sqrt{\frac{2}{3}} \cos \theta_{\eta\eta'} \right)}{2}
 \end{aligned} \tag{A.12}$$

where $\theta_{\eta\eta'} \approx -13^\circ$ is the η - η' mixing angle and we make the convenient choice $\kappa_q = m_q^{-1} / \text{Tr}[\mathbf{m}^{-1}]$ with $\mathbf{m} = \text{diag}\{m_u, m_d, m_s\}$ to achieve $m_{a\pi^0} = m_{a\eta_8} = 0$ [11]. The remaining mass mixing terms are

$$\begin{aligned}
 m_{a\eta}^2 &= -\frac{\sqrt{6}B_0 \sin \theta_{\eta\eta'}}{\text{Tr}[\mathbf{m}^{-1}]} \\
 m_{a\eta'}^2 &= \frac{\sqrt{6}B_0 \cos \theta_{\eta\eta'}}{\text{Tr}[\mathbf{m}^{-1}]}
 \end{aligned} \tag{A.13}$$

where $B_0 = m_{\pi^0}^2 / (m_u + m_d)$.

The ALP-gluon coupling as well as ALP-meson mixing give rise to further contributions to the effective ALP-photon coupling, which is found to be [61]

$$C_{\gamma\gamma, \text{eff}} = C_{\gamma\gamma} - 1.92C_{gg} + \tilde{\alpha}_s(m_a^2) \mathcal{F}_{\text{VMD}} \sum_P \frac{\Lambda}{16\pi^2 f_P} \theta_{aP} \tag{A.14}$$

with $f_\eta \approx 93$ MeV and $f_{\eta'} \approx 73$ MeV. The functions \mathcal{F}_{VMD} and $\tilde{\alpha}_s$, defined in eqs. (3.7) and (3.8) respectively, allow us to extrapolate the mixing effects for $m_a > m_{\eta'}$.

B Calculation of decay widths

Based on the effective couplings derived in appendix A we can calculate the various decay widths relevant for this study.

We begin by noting that in all cases of interest the decaying particle has spin 0 and hence the squared matrix element does not depend on the phase space parameters, leading to an isotropic decay. Hence no information on the decay kinematics is lost when performing the phase space integration and it is sufficient to quote the corresponding decay width.

The width of the decay $a \rightarrow \gamma\gamma$ is given by

$$\Gamma_{a \rightarrow \gamma\gamma} = \frac{C_{\gamma\gamma, \text{eff}}^2 e^4 m_a^3}{\Lambda^2 4\pi} \tag{B.1}$$

with an effective ALP-photon coupling $C_{\gamma\gamma, \text{eff}}$. For the ALP leptonic decays $a \rightarrow \ell\ell$ ($\ell = e, \mu$) with decay width

$$\Gamma_{a \rightarrow \ell\ell} = \frac{C_{\ell\ell, \text{eff}}^2 m_\ell^2}{\Lambda^2 8\pi} \sqrt{m_a^2 - 4m_\ell^2} \tag{B.2}$$

where m_ℓ denotes the lepton mass and $C_{\ell\ell, \text{eff}}$ the effective leptonic coupling.

Other possible two-body decays for ALP masses up to 3 GeV are the hadronic decays $a \rightarrow \rho\rho$, $a \rightarrow \omega\omega$, $a \rightarrow K^*\bar{K}^*$ and $a \rightarrow \phi\phi$, which (in the absence of direct couplings to quarks) have a decay width proportional to C_{gg}^2 . Since these final states are not *detector-stable* for fixed-target experiments, we do not calculate them explicitly but we account for them in the decay widths for three-body decays. The corresponding decay widths used in this work employ amplitudes derived in ref. [18] using chiral perturbation theory and vector meson dominance.

Furthermore, we are also interested in the production of ALPs via flavour-changing two-body decays. The decay widths for the rare decays $B \rightarrow K^{(*)} + a$ can be calculated in terms of the *flavour-changing* effective couplings C_{bs} :

$$\begin{aligned}\Gamma_{B \rightarrow K+a} &= \frac{|C_{bs}|^2}{\Lambda^2} \frac{1}{64\pi m_B^3} (m_B^2 - m_K^2)^2 f_0^2(m_a) \lambda^{1/2}(m_B, m_K, m_a) \\ \Gamma_{B \rightarrow K^*+a} &= \frac{|C_{bs}|^2}{\Lambda^2} \frac{1}{64\pi m_B^3} A_0^2(m_a) \lambda^{3/2}(m_B, m_{K^*}, m_a),\end{aligned}\tag{B.3}$$

where $\lambda(x, y, z) = [x^2 - (x+z)^2][x^2 - (y-z)^2]$. Since we apply this result on a broad range of ALP masses, we use the following parametrization of the form factors [62, 63]:

$$f_0(q) = \frac{0.330}{1 - q^2/37.46},\tag{B.4}$$

$$A_0(q) = \frac{1.364}{1 - q^2/m_B^2} - \frac{0.990}{1 - q^2/36.78 \text{ GeV}^2},\tag{B.5}$$

which depend on the momentum transfer q .

In complete analogy, we find for the decay $D \rightarrow \pi + a$:

$$\Gamma_{D \rightarrow \pi+a} = \frac{|C_{cu}|^2}{\Lambda^2} \frac{1}{64\pi m_D^3} (m_D^2 - m_\pi^2)^2 f_0^2(m_a) \lambda^{1/2}(m_D, m_\pi, m_a),\tag{B.6}$$

where we use $f_0(0) = 0.612$ [64] (see also ref. [65]). To account for the q -dependence we use an analogous parametrization as above:

$$f_0(q) = \frac{f_0(0)}{1 - q^2/m_{\text{fit}}^2}\tag{B.7}$$

For fitted $m_{\text{fit}}^2 = 6.46 \text{ GeV}^2$ we reproduce the mean values of ref. [64] with a relative error $< 1\%$, which is sufficient for our purpose.

C Brief description of the ALPINIST code

Since the MC simulations of ALP production and decay are computing power demanding, our goal is to avoid as much as possible re-doing these simulations for each change in the model-dependent parameters. Thus we separate the code into several standalone modules, each of which processes the input tabulated data and passes the output data, again in the form of a table, to the next module in the sequence. This sequence is shown in the form of a diagram in figure 8, where the code-modules are represented by square boxes and the tabulated data by triangular boxes pointing in the direction of the data flow.

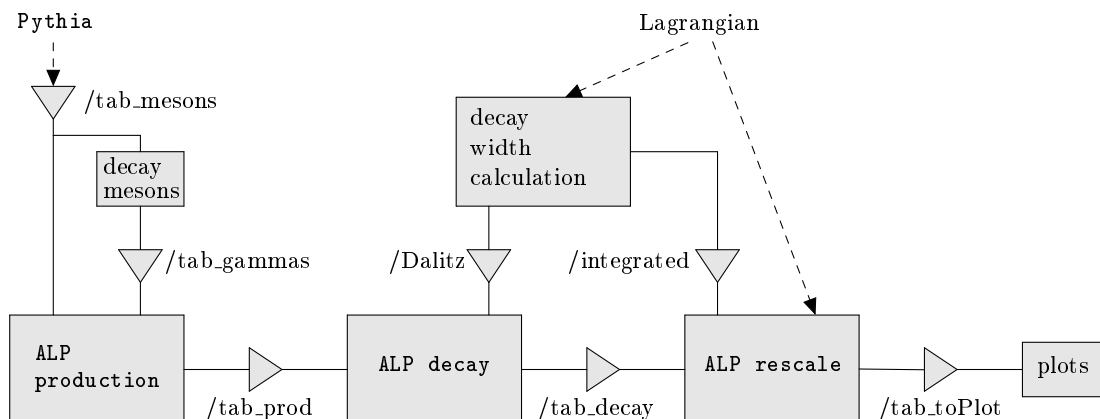


Figure 8. Diagram representing the internal structure of the ALPINIST framework.

- The **ALP production** module uses **Wolfram Mathematica** [66]. It corresponds to section 3, i.e. it generates the ALP distribution table for each production mode and for each experiment for a fixed value of the model-dependent parameters. For the ALP production from mesons it reads directly the output from the **PYTHIA** simulations. For the production from on-shell photons, it reads the photon distributions generated by a separate module, which simulates 2γ decays of π^0 , η and η' mesons. This module uses the **ROOT** framework [38].
- The **ALP decay** module also uses the **ROOT** framework libraries. It contains the main MC simulation described in sections 4 and 5, i.e. it handles the ALP decay kinematics for various decay modes in the conditions of the specific experiment. For the three-body decays it also requires the Dalitz plot densities, which are generated according to the specific Lagrangian (eq. (6.1) in this case) by a separate module.
- The **ALP rescale** module requires only **Python** [67] and the input tables to be run. For the set of model-dependent parameters given by the user it generates the table with the number of signal events in the common m_a - C_a plane used for plotting the exclusion limits. The equations used for the calculations can be found in appendices A, B as well as the references therein.

The **ALP rescale** module is the most relevant one for most users and thus is provided publicly at <https://github.com/jjerhot/ALPINIST> together with its input tables. Other modules may be needed for potential updates of experimental conditions and for including new experiments, production or decay modes to the simulation. Access to the full repository can be obtained upon request from the authors. The specific instructions for using of each module can be found in the respective **README.md** files.

D Treatment of kinematics for the mixing production

The mixing production serves as an approximation to the various processes in which an ALP can be produced. The total yield of produced ALPs can be derived from the yield

of the neutral pseudoscalar mesons and the respective mixing angles with the ALP. This approach, however, does not say anything about the change of the kinematics when the ALP is produced instead of the meson. The general minimal adjustment, which can be applied irrespective of the specific process which lead to the ALP production, is the change of the kinematics when the mass of the original pseudoscalar is changed. Since the center-of-mass frame of the processes that lead to the meson production is unknown, we choose a common frame for the mass adjustment to be the p - p collision *cm frame*. We emphasize that there is a certain level of arbitrariness in this choice, with the obvious alternative being the laboratory frame. However, we expect that the p - p *cm frame* leads to a more accurate description for highly-energetic ALPs, where only a smaller portion of the available energy is spent on the ALP production and the majority of the energy is carried in the ALP boost. These highly boosted ALPs are particularly important for the sensitivity of beam dump experiment searches.

The *cm* momentum of a particle with mass m is then simply found as

$$p_{\text{cm}}(p_{\text{lab}}, m) = \gamma_{\text{beam}} p_{\text{lab}} - \gamma_{\text{beam}} \beta_{\text{beam}} \sqrt{p_{\text{lab}}^2 + m^2} \tag{D.1}$$

and analogous relations can be used for the E_{cm} and, by inversion, for p_{lab} and E_{lab} . By adjusting the particle mass one necessarily violates the momentum or energy conservation (or both) in the original process. In our simulation we fix the momentum of the ALP in the p - p *cm frame* to be the same as the momentum of the original meson. The energy and angle of the ALP a in the laboratory frame then translates to the energy and angle of the original meson P in the laboratory frame as

$$\begin{aligned} E_P|_{p_{\text{fixed}}} &= \sqrt{p_{\text{lab}}(p_{\text{cm}}(p_a, m_a), m_P) + m_P^2}, \\ \theta_P|_{p_{\text{fixed}}} &= \arcsin \left[\sin(\theta_a) \frac{p_a}{p_{\text{lab}}(p_{\text{cm}}(p_a, m_a), m_P)} \right]. \end{aligned} \tag{D.2}$$

The plotted distributions for the case of $P = \eta'$ can be found in the upper row of figure 9.

If one chooses to fix the ALP energy to the meson energy in the p - p *cm frame*, one obtains in complete analogy the following relations:

$$\begin{aligned} E_P|_{E_{\text{fixed}}} &= E_{\text{lab}}(E_{\text{cm}}(p_a, m_a), m_P) \\ \theta_P|_{E_{\text{fixed}}} &= \arcsin \left[\sin(\theta_a) \frac{p_a}{\sqrt{E_{\text{lab}}(E_{\text{cm}}(E_a, m_a), m_P)^2 - m_P^2}} \right]. \end{aligned} \tag{D.3}$$

The resulting distributions are plotted in the bottom row of figure 9. By comparing the two distributions, we can see that even in the borderline case of choosing the p - p collision as a universal *cm* frame, we can get a significantly different E_a - θ_a distribution solely by choosing a different method of fixing the kinematics in the *cm frame*. This choice, however, impacts mostly ALPs with a small boost (large masses and low energies), which are of limited relevance for beam dump experiments.




Figure 9. Dependence of the energy (left) and the angle of the ALP (right) on the ALP mass in the laboratory frame for the case of the fixed momentum (upper panel) or fixed energy (lower panel) in the p - p cm frame when the original particle was the η' meson, produced by a 400 GeV proton beam. The angular distribution is evaluated for the case $E_a = 100$ GeV.

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. SCOAP³ supports the goals of the International Year of Basic Sciences for Sustainable Development.

References

- [1] P. Agrawal et al., *Feebly-interacting particles: FIPs 2020 workshop report*, *Eur. Phys. J. C* **81** (2021) 1015 [[arXiv:2102.12143](https://arxiv.org/abs/2102.12143)] [[INSPIRE](https://inspirehep.net/literature/1925000)].
- [2] J. Beacham et al., *Physics Beyond Colliders at CERN: Beyond the Standard Model Working Group Report*, *J. Phys. G* **47** (2020) 010501 [[arXiv:1901.09966](https://arxiv.org/abs/1901.09966)] [[INSPIRE](https://inspirehep.net/literature/1700000)].
- [3] G. Aielli et al., *Expression of interest for the CODEX-b detector*, *Eur. Phys. J. C* **80** (2020) 1177 [[arXiv:1911.00481](https://arxiv.org/abs/1911.00481)] [[INSPIRE](https://inspirehep.net/literature/1750000)].
- [4] K.J. Kelly, S. Kumar and Z. Liu, *Heavy axion opportunities at the DUNE near detector*, *Phys. Rev. D* **103** (2021) 095002 [[arXiv:2011.05995](https://arxiv.org/abs/2011.05995)] [[INSPIRE](https://inspirehep.net/literature/1850000)].

- [5] S. Andreas, O. Lebedev, S. Ramos-Sanchez and A. Ringwald, *Constraints on a very light CP-odd Higgs of the NMSSM and other axion-like particles*, *JHEP* **08** (2010) 003 [[arXiv:1005.3978](#)] [[INSPIRE](#)].
- [6] M. Cicoli, M. Goodsell and A. Ringwald, *The type IIB string axiverse and its low-energy phenomenology*, *JHEP* **10** (2012) 146 [[arXiv:1206.0819](#)] [[INSPIRE](#)].
- [7] Y. Nomura and J. Thaler, *Dark Matter through the Axion Portal*, *Phys. Rev. D* **79** (2009) 075008 [[arXiv:0810.5397](#)] [[INSPIRE](#)].
- [8] B. Batell, M. Pospelov and A. Ritz, *Exploring Portals to a Hidden Sector Through Fixed Targets*, *Phys. Rev. D* **80** (2009) 095024 [[arXiv:0906.5614](#)] [[INSPIRE](#)].
- [9] M. Freytsis, Z. Ligeti and J. Thaler, *Constraining the Axion Portal with $B \rightarrow Kl^+l^-$* , *Phys. Rev. D* **81** (2010) 034001 [[arXiv:0911.5355](#)] [[INSPIRE](#)].
- [10] M.J. Dolan, F. Kahlhoefer, C. McCabe and K. Schmidt-Hoberg, *A taste of dark matter: Flavour constraints on pseudoscalar mediators*, *JHEP* **03** (2015) 171 [Erratum *ibid.* **07** (2015) 103] [[arXiv:1412.5174](#)] [[INSPIRE](#)].
- [11] H. Georgi, D.B. Kaplan and L. Randall, *Manifesting the Invisible Axion at Low-energies*, *Phys. Lett. B* **169** (1986) 73 [[INSPIRE](#)].
- [12] M. Chala, G. Guedes, M. Ramos and J. Santiago, *Running in the ALPs*, *Eur. Phys. J. C* **81** (2021) 181 [[arXiv:2012.09017](#)] [[INSPIRE](#)].
- [13] M. Bauer, M. Neubert, S. Renner, M. Schnubel and A. Thamm, *The Low-Energy Effective Theory of Axions and ALPs*, *JHEP* **04** (2021) 063 [[arXiv:2012.12272](#)] [[INSPIRE](#)].
- [14] M.J. Dolan, T. Ferber, C. Hearty, F. Kahlhoefer and K. Schmidt-Hoberg, *Revised constraints and Belle II sensitivity for visible and invisible axion-like particles*, *JHEP* **12** (2017) 094 [Erratum *ibid.* **03** (2021) 190] [[arXiv:1709.00009](#)] [[INSPIRE](#)].
- [15] 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](#)].
- [16] I. Brivio et al., *ALPs Effective Field Theory and Collider Signatures*, *Eur. Phys. J. C* **77** (2017) 572 [[arXiv:1701.05379](#)] [[INSPIRE](#)].
- [17] G. Alonso-Álvarez, M.B. Gavela and P. Quilez, *Axion couplings to electroweak gauge bosons*, *Eur. Phys. J. C* **79** (2019) 223 [[arXiv:1811.05466](#)] [[INSPIRE](#)].
- [18] D. Aloni, Y. Soreq and M. Williams, *Coupling QCD-Scale Axionlike Particles to Gluons*, *Phys. Rev. Lett.* **123** (2019) 031803 [[arXiv:1811.03474](#)] [[INSPIRE](#)].
- [19] F. Ertas and F. Kahlhoefer, *On the interplay between astrophysical and laboratory probes of MeV-scale axion-like particles*, *JHEP* **07** (2020) 050 [[arXiv:2004.01193](#)] [[INSPIRE](#)].
- [20] M. Bauer, M. Neubert and A. Thamm, *Collider Probes of Axion-Like Particles*, *JHEP* **12** (2017) 044 [[arXiv:1708.00443](#)] [[INSPIRE](#)].
- [21] C. Cornella, P. Paradisi and O. Sumensari, *Hunting for ALPs with Lepton Flavor Violation*, *JHEP* **01** (2020) 158 [[arXiv:1911.06279](#)] [[INSPIRE](#)].
- [22] M. Bauer, M. Neubert, S. Renner, M. Schnubel and A. Thamm, *Flavor probes of axion-like particles*, [arXiv:2110.10698](#) [[INSPIRE](#)].
- [23] K. Mimasu and V. Sanz, *ALPs at Colliders*, *JHEP* **06** (2015) 173 [[arXiv:1409.4792](#)] [[INSPIRE](#)].

- [24] B. Döbrich, F. Ertas, F. Kahlhoefer and T. Spadaro, *Model-independent bounds on light pseudoscalars from rare B-meson decays*, *Phys. Lett. B* **790** (2019) 537 [[arXiv:1810.11336](#)] [[INSPIRE](#)].
- [25] L. Buonocore, C. Frugiuele, F. Maltoni, O. Mattelaer and F. Tramontano, *Event generation for beam dump experiments*, *JHEP* **05** (2019) 028 [[arXiv:1812.06771](#)] [[INSPIRE](#)].
- [26] B. Döbrich, J. Jaeckel and T. Spadaro, *Light in the beam dump — ALP production from decay photons in proton beam-dumps*, *JHEP* **05** (2019) 213 [*Erratum ibid.* **10** (2020) 046] [[arXiv:1904.02091](#)] [[INSPIRE](#)].
- [27] B. Döbrich, J. Jaeckel, F. Kahlhoefer, A. Ringwald and K. Schmidt-Hoberg, *ALPTraum: ALP production in proton beam dump experiments*, *JHEP* **02** (2016) 018 [[arXiv:1512.03069](#)] [[INSPIRE](#)].
- [28] A. Carmona, C. Scherb and P. Schwaller, *Charming ALPs*, *JHEP* **08** (2021) 121 [[arXiv:2101.07803](#)] [[INSPIRE](#)].
- [29] H.-C. Cheng, L. Li and E. Salvioni, *A theory of dark pions*, *JHEP* **01** (2022) 122 [[arXiv:2110.10691](#)] [[INSPIRE](#)].
- [30] J. Jerhot, B. Döbrich, E. Ertas, F. Kahlhoefer and T. Spadaro, *ALPINIST: v1.0.0*, [10.5281/zenodo.5844011](#) (2022).
- [31] M. Berger et al., *XCOM: Photon Cross Sections Database*, DOI (2010) [<http://physics.nist.gov/xcom>].
- [32] V.M. Budnev, I.F. Ginzburg, G.V. Meledin and V.G. Serbo, *The Two photon particle production mechanism. Physical problems. Applications. Equivalent photon approximation*, *Phys. Rept.* **15** (1975) 181 [[INSPIRE](#)].
- [33] L. Harland-Lang, J. Jaeckel and M. Spannowsky, *A fresh look at ALP searches in fixed target experiments*, *Phys. Lett. B* **793** (2019) 281 [[arXiv:1902.04878](#)] [[INSPIRE](#)].
- [34] T. Sjöstrand et al., *An introduction to PYTHIA 8.2*, *Comput. Phys. Commun.* **191** (2015) 159 [[arXiv:1410.3012](#)] [[INSPIRE](#)].
- [35] S. Chakraborty, M. Kraus, V. Loladze, T. Okui and K. Tobioka, *Heavy QCD axion in $b \rightarrow s$ transition: Enhanced limits and projections*, *Phys. Rev. D* **104** (2021) 055036 [[arXiv:2102.04474](#)] [[INSPIRE](#)].
- [36] PARTICLE DATA GROUP collaboration, *Review of Particle Physics*, *PTEP* **2020** (2020) 083C01 [[INSPIRE](#)].
- [37] C. Lourenco and H.K. Wohri, *Heavy flavour hadro-production from fixed-target to collider energies*, *Phys. Rept.* **433** (2006) 127 [[hep-ph/0609101](#)] [[INSPIRE](#)].
- [38] R. Brun and F. Rademakers, *ROOT: An object oriented data analysis framework*, *Nucl. Instrum. Meth. A* **389** (1997) 81 [[INSPIRE](#)].
- [39] CHARM collaboration, *Search for Axion Like Particle Production in 400-GeV Proton-Copper Interactions*, *Phys. Lett. B* **157** (1985) 458 [[INSPIRE](#)].
- [40] CHARM collaboration, *Bounds on Supersymmetric Particles From a Proton Beam Dump Experiment*, *Phys. Lett. B* **121** (1983) 429 [[INSPIRE](#)].
- [41] CHARM collaboration, *Experimental Study of Prompt Neutrino Production in 400-GeV Proton-Nucleus Collisions*, *Phys. Lett. B* **96** (1980) 435 [[INSPIRE](#)].

- [42] J. Blümlein et al., *Limits on neutral light scalar and pseudoscalar particles in a proton beam dump experiment*, *Z. Phys. C* **51** (1991) 341 [INSPIRE].
- [43] J. Blümlein and J. Brunner, *New Exclusion Limits on Dark Gauge Forces from Proton Bremsstrahlung in Beam-Dump Data*, *Phys. Lett. B* **731** (2014) 320 [arXiv:1311.3870] [INSPIRE].
- [44] H. Nanjo, *Dark Sector search at KOTO experiment*, contribution to *3rd meeting on Searches for Hidden Sectors at Kaon and Hyperon Factories*, PDF (2021).
- [45] K. Massri, *Update on the high-intensity Kaon programme at the SPS*, contribution to *Physics Beyond Colliders General WG meeting*, PDF (2021).
- [46] C. Lazzeroni, *NA62/KLEVER prospects for future high-intensity K^+ and K_L running, including beam dump*, contribution to *Physics Beyond Colliders Annual Workshop, BSM session*, PDF (2021).
- [47] NA62 collaboration, *The Beam and detector of the NA62 experiment at CERN*, *2017 JINST* **12** P05025 [arXiv:1703.08501] [INSPIRE].
- [48] SHiP collaboration, *A facility to Search for Hidden Particles (SHiP) at the CERN SPS*, arXiv:1504.04956 [INSPIRE].
- [49] SHiP collaboration, *SHiP Experiment — Progress Report*, CERN-SPSC-2019-010, SPSC-SR-248 (2019).
- [50] SHiP collaboration, *Sensitivity of the SHiP experiment to dark photons decaying to a pair of charged particles*, *Eur. Phys. J. C* **81** (2021) 451 [arXiv:2011.05115] [INSPIRE].
- [51] SEAQUEST collaboration, *The SeaQuest Spectrometer at Fermilab*, *Nucl. Instrum. Meth. A* **930** (2019) 49 [arXiv:1706.09990] [INSPIRE].
- [52] B. Batell, J.A. Evans, S. Gori and M. Rai, *Dark Scalars and Heavy Neutral Leptons at DarkQuest*, *JHEP* **05** (2021) 049 [arXiv:2008.08108] [INSPIRE].
- [53] A. Berlin, S. Gori, P. Schuster and N. Toro, *Dark Sectors at the Fermilab SeaQuest Experiment*, *Phys. Rev. D* **98** (2018) 035011 [arXiv:1804.00661] [INSPIRE].
- [54] N. Blinov, E. Kowalczyk and M. Wynne, *Axion-like particle searches at DarkQuest*, *JHEP* **02** (2022) 036 [arXiv:2112.09814] [INSPIRE].
- [55] DUNE collaboration, *Deep Underground Neutrino Experiment (DUNE) Near Detector Conceptual Design Report*, *Instruments* **5** (2021) 31 [arXiv:2103.13910] [INSPIRE].
- [56] DUNE collaboration, *Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE): Conceptual Design Report, Volume 1: The LBNF and DUNE Projects*, arXiv:1601.05471 [INSPIRE].
- [57] DUNE collaboration, *The Single-Phase ProtoDUNE Technical Design Report*, arXiv:1706.07081 [INSPIRE].
- [58] W. Baldini et al., *SHADOWS (Search for Hidden And Dark Objects With the SPS)*, arXiv:2110.08025 [INSPIRE].
- [59] J. Blümlein et al., *Limits on the mass of light (pseudo)scalar particles from Bethe-Heitler e^+e^- and $\mu^+\mu^-$ pair production in a proton-iron beam dump experiment*, *Int. J. Mod. Phys. A* **7** (1992) 3835 [INSPIRE].
- [60] G. Alonso-Álvarez, F. Ertas, J. Jaeckel, F. Kahlhoefer and L.J. Thormaehlen, *Leading logs in QCD axion effective field theory*, *JHEP* **07** (2021) 059 [arXiv:2101.03173] [INSPIRE].

- [61] G. Grilli di Cortona, E. Hardy, J. Pardo Vega and G. Villadoro, *The QCD axion, precisely*, *JHEP* **01** (2016) 034 [[arXiv:1511.02867](#)] [[INSPIRE](#)].
- [62] P. Ball and R. Zwicky, *New results on $B \rightarrow \pi, K, \eta$ decay formfactors from light-cone sum rules*, *Phys. Rev. D* **71** (2005) 014015 [[hep-ph/0406232](#)] [[INSPIRE](#)].
- [63] P. Ball and R. Zwicky, *$B_{d,s} \rightarrow \rho, \omega, K^*, \phi$ decay form-factors from light-cone sum rules revisited*, *Phys. Rev. D* **71** (2005) 014029 [[hep-ph/0412079](#)] [[INSPIRE](#)].
- [64] ETM collaboration, *Scalar and vector form factors of $D \rightarrow \pi(K)\ell\nu$ decays with $N_f = 2 + 1 + 1$ twisted fermions*, *Phys. Rev. D* **96** (2017) 054514 [*Erratum ibid.* **99** (2019) 099902] [*Erratum ibid.* **100** (2019) 079901] [[arXiv:1706.03017](#)] [[INSPIRE](#)].
- [65] 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](#)].
- [66] W.R. Inc., *Mathematica, Version 12.3.1*, Champaign, IL, U.S.A. (2021).
- [67] G. Van Rossum and F.L. Drake, *Python 3 Reference Manual*, CreateSpace, Scotts Valley, CA, U.S.A. (2009).