

Detector performance at SHiP for cascade-produced long-lived particles

Matei Climescu^{1,*}, Yehor Kyselov^{2,†} and Maksym Ovchinnikov^{3,‡}

¹*Ghent University, Ghent, Belgium*

²*Taras Shevchenko National University of Kyiv, Kyiv, Ukraine*

³*Theoretical Physics Department, CERN, 1211 Geneva 23, Switzerland*



(Received 12 June 2026; accepted 24 August 2026; published 14 September 2026)

Previous studies have shown that cascade production in the thick target of the Search for Hidden Particles (SHiP) experiment may substantially enhance the number of light long-lived particles (LLPs) decaying in the fiducial volume. However, cascade-produced LLPs are typically soft, so daughter-level acceptance and reconstruction effects can strongly suppress the observable event rate. We quantify this suppression for two representative cases: photophilic axionlike particles (ALPs) produced in electromagnetic cascades and heavy neutral leptons (HNLs) produced in decays of secondary kaons. We combine a semianalytic event-rate calculation with a detector-level study of ALP reconstruction in the electromagnetic calorimeter. For the nominal SHiP detector design, cascade ALPs give at most a moderate enhancement over primary production and only at the lightest masses; at higher masses, the cascade contribution becomes subdominant or negligible. For HNLs from secondary kaons, the cascade contribution is already subdominant after imposing daughter-level geometric acceptance. We also identify possible ways to recover part of the cascade event rate, including relaxed event-selection criteria and an active decay-volume subdetector.

DOI: [10.1103/h9gk-mjjh](https://doi.org/10.1103/h9gk-mjjh)

I. INTRODUCTION

Hypothetical particles with GeV-scale masses and tiny couplings to Standard Model (SM) fields may address several open problems in particle physics [1–7]. Despite their small couplings, such particles can be efficiently produced at accelerator experiments operating at much higher energy scales [6,8–14] and can be searched for using a wide range of techniques and signatures [15–26]. They are often referred to as long-lived particles (LLPs).

The upcoming Search for Hidden Particles (SHiP) experiment at Super Proton Synchrotron (SPS) [27,28], to be built at the SPS North Area, is one of the flagships in exploring the parameter space of GeV-mass LLPs via the signatures of displaced decays and scatterings, for a broad range of LLP scenarios [19,22,25]. This is achieved through a very high intensity of the SPS proton beam with the relatively large collision energy, together with

excellent background suppression and close-to-optimal detector geometry [29].

SHiP is currently in the technical design report (TDR) phase [30], making detector-performance studies increasingly important for realistic LLP signals and kinematic regimes. Most SHiP sensitivity studies [27,31,32] assume that LLPs are produced predominantly in the first interaction of the proton beam with the target. However, several recent works [33–38] have shown that secondary interactions inside the thick target can yield additional LLPs. These cascade-produced LLPs are typically much softer than those from primary production. It is therefore important to assess whether their decay products remain within the detector acceptance and can be reconstructed with sufficient efficiency.

In this work, we study the SHiP detector response to such low-energy cascade-induced signals. We extend previous estimates in two steps. First, using a semianalytic calculation, we evaluate the daughter-level geometric acceptance, namely, the fraction of LLP decays for which the visible decay products reach the relevant downstream detector. Second, for the two-photon final state, we perform a detector-level simulation and reconstruction study for events passing this geometric requirement.

We consider two representative scenarios: axionlike particles (ALPs) coupled to photons, which can be efficiently produced in electromagnetic cascades, and heavy neutral leptons (HNLs), which may be produced in decays

*Contact author: matei.climescu@cern.ch

†Contact author: kiselev883@gmail.com

‡Contact author: maksym.ovchinnikov@cern.ch

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Funded by SCOAP³.

of secondary kaons. The geometric-acceptance part of our analysis can be generalized to other low-energy signals, including vector particles such as dark photons or mediators coupled to the $L_\alpha - L_\beta$ current [36,37].

The paper is organized as follows. Section II describes the SHiP experiment and the baseline selection criteria. Section III reviews the main sources of cascade-produced LLPs in the thick SHiP target. Section IV summarizes existing estimates and their limitations for low-energy cascade signals. Section V presents our calculation of the LLP fluxes, geometric acceptance, and reconstruction efficiency. Section VI contains the main results and discussion: for ALPs, we compare the semianalytic acceptance estimate with the detector-level reconstruction study, quantifying the impact of the electromagnetic calorimeter (ECAL) acceptance, clustering, and impact-parameter requirements, and discuss possible ways to improve the sensitivity to low-energy two-photon signals; for HNLs, we show that the cascade contribution becomes subdominant once daughter-level geometric acceptance is imposed. We conclude in Sec. VII.

II. SHiP EXPERIMENT

The SHiP experiment is a proposed beam-dump facility at the CERN SPS designed to search for feebly interacting particles (FIPs) with macroscopic lifetimes. The experiment will use a slow-extracted proton beam with momentum 400 GeV and an annual intensity of about

$$N_{\text{PoT}} \simeq 4 \times 10^{19}, \quad (1)$$

directed onto a compact high- Z target. A representation of the experiment is shown in Fig. 1. The target serves as both the production point for hidden-sector particles and the primary beam dump. It has a transverse size of order 30×30 cm and a longitudinal depth of order 1.2–1.4 m, corresponding to approximately 12 nuclear interaction lengths. The baseline design uses a dense tungsten-based core. This choice maximizes the production of heavy mesons and photons, while suppressing pion and kaon decays in flight, thereby reducing the neutrino and muon backgrounds entering the detector region.

Immediately downstream of the target, the experimental layout contains a hadronic absorber and a magnetic muon shield. Their purpose is to absorb secondary hadrons and to




FIG. 1. Representation of the SHiP detector overlaid with a two-body FIP decay.

sweep away the intense flux of muons produced in the dump before these particles can reach the hidden-sector decay volume. This shielding system is essential for achieving the low-background environment required for searches for displaced visible decays.

The hidden-sector decay volume starts at $z \simeq 33$ m and ends at $z \simeq 83$ m, giving a fiducial length of approximately 50 m. Its geometry is approximately that of a truncated pyramid, aligned with the beam axis. The transverse opening is about 1×3 m at the entrance and expands to about 4×6 m at the downstream end. The volume is filled with a helium balloon to minimize secondary interactions in the material while maintaining a large geometrical acceptance for long-lived particle decays.

A magnetic spectrometer composed of four straw-tube-based tracking stations is located downstream of the decay volume, after a small $\simeq 1$ m gap. In the current layout, the spectrometer system occupies approximately the region

$$84 \lesssim z \lesssim 95 \text{ m}. \quad (2)$$

It consists of a large-aperture dipole magnet and four low-mass tracking stations, arranged with tracking stations before and after the magnetic-field region. The magnet has an aperture of approximately 4×6 m and provides a field integral of $\int B dl \simeq 0.6\text{--}0.8$ Tm, with a peak field of about $B \simeq 0.15$ T. The curvature of charged-particle trajectories in the magnetic field provides the momentum measurement, while the reconstructed tracks are used to determine the decay vertex, the impact parameter with respect to the production target, and the invariant mass under the relevant particle-identification hypotheses.

A timing detector is placed immediately downstream of the spectrometer, around $z \simeq 95$ m. It provides precise time information for charged tracks and helps reject accidental combinations of particles from uncorrelated backgrounds.

The ECAL begins at $z \simeq 95.5$ m and ends at $z \simeq 97.32$ m. Its transverse dimensions are approximately

$$A_{\text{ECAL}} \approx 4.32 \times 6.48 \text{ m}. \quad (3)$$

It uses the SplitCal detector concept [39,40], which provides electron and photon identification, measures electromagnetic shower energies, and reconstructs shower directionality. In combination with the spectrometer, it is crucial for reconstructing final states containing electrons or photons, including two-photon signatures and partially reconstructed hidden-sector decays.

The calorimeter is composed of scintillator bars and three high-precision layers based on scintillating fibers to improve angular resolution. The full detector comprises 46 layers, with a 1 m longitudinal gap separating the first five X_0 from the remainder of the detector, with the high-precision layers placed after 5 (right before the gap), 7, and 11 X_0 . An event display is shown in Fig. 2. The center of




FIG. 2. Event display of a reconstructed 10 GeV $\pi^0 \rightarrow \gamma\gamma$ event. The magenta and white lines correspond to the reconstructed photon directions; the pink dot on the left panel corresponds to the reconstructed vertex.

the calorimeter is placed 96.97 m away from the tungsten target, which is also the origin.

Finally, the hadronic calorimeter (HCAL) is located downstream of the ECAL. In the current layout, it starts at $z \simeq 97.42$ m and ends at $z \simeq 98.42$ m, and it is arranged as a λ_A iron sandwich calorimeter sampled every $5\lambda_A$ by scintillator layers similar to the ECAL's [40].

The HCAL provides muon and hadron identification and measures the energy deposited by hadronic showers, complementing the ECAL information in the downstream particle-identification system.

The entire calorimeter system provides particle identification in excess of 99% true reconstruction for electrons, charged pions, and muons.

A. Baseline selection criteria

Let us summarize the selection criteria on the events with decaying LLPs used in previous SHiP studies, following Refs. [27,28].

The central signature is a displaced-decay vertex reconstructed from the visible decay products. The baseline requirement is at least two reconstructible visible objects with total electric charge zero, such as $\gamma\gamma$, e^+e^- , $\mu^\pm\pi^\mp$, or other states.

Charged decay products are required to be reconstructed in the magnetic tracking spectrometer, which provides their momenta, the decay vertex, and the main kinematic information for charged final states. In the reconstruction considered here, charged tracks are required to lie within the acceptance of all four tracking stations of the spectrometer. The downstream particle-identification system, consisting primarily of the ECAL and HCAL, is then used to assign particle hypotheses and to reconstruct neutral electromagnetic final states such as photons. Thus, the spectrometer alone can reconstruct charged-track kinematics under assumed mass hypotheses, but calorimetry is needed to distinguish electrons, hadrons, and penetrating

muonlike tracks and to reconstruct final states containing photons or neutral hadrons. This particle-identification (PID) information, together with timing, vetoes, and topological requirements, is important for SHiP to operate in a nearly background-free regime. In particular, relaxing the PID requirements, for instance, by requiring electron decay products to be reconstructed only in the spectrometer without an ECAL match, would weaken background rejection. For this reason, when quantifying the geometric component of the decay-product acceptance, we require the visible decay products to remain within the relevant downstream detector acceptance: the ECAL for electromagnetic final states and the HCAL for hadronic or penetrating muonlike final states.

The SHiP Collaboration studies summarize baseline selection criteria that the decay products must pass to be considered as a reconstructed decay event:

- (i) Each visible decay product must have momentum, or energy in the case of photons,

$$p_{\text{daughter}} > 1 \text{ GeV}, \quad (4)$$

where for photons this condition should be understood as $E_\gamma > 1 \text{ GeV}$.

- (ii) The reconstructed tracks or directions must be compatible with a common displaced vertex, quantified by a distance of closest approach (DOCA),

$$\text{DOCA} < 1 \text{ cm}. \quad (5)$$

- (iii) The four-momentum of the LLP candidate, reconstructed from the decay-product kinematics, is extrapolated back toward the target; the corresponding transverse impact parameter (IP) must satisfy

$$\text{IP} < 2.5 \text{ m}. \quad (6)$$

For fully reconstructible final states made of charged particles (all decay products are reconstructed), this requirement can be tightened to $\text{IP} < 0.1 \text{ m}$.

After applying these selection criteria and requiring full event reconstruction, the search for displaced decays is expected to be effectively background-free. For the typical kinematics of LLPs X produced in primary, noncascade proton beam collisions, with energies $E_X \gtrsim 20 \text{ GeV}$ [41], the criteria (4)–(6) generically lead to an $\mathcal{O}(1)$ event reconstruction efficiency, assuming that the LLPs mostly decay into charged states [31,32]. See Sec. III for the distinction between primary and cascade production.

Performance of the SHiP detector to reconstruct the events with decays into a pair of photons has been studied in a recent work [40], considering photophilic ALPs produced in primary interactions as a concrete example. For decays into a pair of photons, an additional effective requirement is that the two photons be separated in the

transverse plane at the ECAL entrance $|\Delta \mathbf{r}_\perp|$ by approximately one Molière radius,

$$|\Delta \mathbf{r}_\perp| > r_M \approx 3.2 \text{ cm}. \quad (7)$$

The study confirmed good reconstruction efficiency for events satisfying the selection (4)–(7). Based on its framework, we will study how the reconstruction efficiency behaves when relaxing the energy criterion (4) and the impact-parameter criterion (6) (imposed in a tightened form in the detector-level study, see Sec. V D 2), which is important for secondary LLPs produced in cascade interactions (see Sec. V IC). As we will show there, relaxing the energy threshold is useful only in combination with a relaxed impact-parameter requirement: on its own it recovers very little signal, because the soft events it admits are in any case removed by the impact-parameter cut.

In this work, we use these cuts as an illustration of how the event rate depends on the selection criteria used by the collaboration (see Sec. VI A 1), while the detector-level study below also considers modified reconstruction requirements. In particular, we avoid using a hard cut on the photon’s energy. Instead, the detector-level study derives the efficiency from the simulated ECAL response, including cluster formation and reconstruction effects at low photon energies; the background implications of this choice require a dedicated study. Still, as we will see in Sec. VI, the calorimeter reconstruction efficiency typically decreases for lower-energy final states because the relative energy and angular resolutions worsen.

III. PARTICLES FROM LOW-ENERGY CASCADES

The thick target of beam-dump experiments like SHiP not only absorbs the primary proton beam but also initiates an extended hadronic and electromagnetic cascade. Throughout this work, we distinguish between the production of LLPs via primary and cascade particles. We call an LLP primary if it is produced in the first proton-target interaction or in the prompt interaction (decay or scattering) of particles produced in that interaction. We call an LLP cascade produced, or secondary, if the Standard Model particle responsible for its production is itself produced later inside the shower.

Cascade production is therefore not a single mechanism, but a collection of physically distinct source classes. One important class is electromagnetic cascades, composed mainly of photons, electrons, and positrons. These particles may produce LLPs through Primakoff scattering, photon fusion, electron or photon bremsstrahlung, or related electromagnetic processes. The relevant LLP examples include photophilic ALPs, dark photons, and vector mediators coupled to lepton currents [35–38]. The defining feature of this source is that the secondary particles are typically much softer than the prompt photons from neutral-meson decays, while still remaining relatively

aligned with the shower direction because electromagnetic interactions are dominated by small momentum transfers.

Another class is formed by secondary long-lived mesons, in particular, charged and neutral kaons [34]. These particles may decay in flight before reinteracting or being absorbed, producing LLPs through rare or weak decays. Examples include HNLs produced in leptonic and semi-leptonic kaon decays [11], as well as Higgs-like scalars and ALPs produced in rare kaon decays [12,14]. Compared to electromagnetic shower particles, secondary kaons have broader angular distributions, and the probability that they decay before absorption depends strongly on their energy. As a result, the corresponding LLP flux is typically softer and less forward than the flux produced by prompt charm or beauty hadrons.

A further class is provided by secondary short-lived mesons, such as π^0 , η , and η' . These cascade mesons are typically softer and more broadly distributed than those produced in the primary proton-target interaction. As a result, the LLPs originating from their decays inherit a low-energy, large-angle kinematic profile. Such contributions are particularly relevant for light mediators produced in meson decays, including dark photons, generic vector mediators, and light dark matter candidates [33,35–37].

Finally, secondary interactions may also increase the number of heavy mesons, such as D and B mesons. In this case, the effect is more moderate: simulations of the SHiP target indicate an enhancement of order $\simeq 1.5$ – 2 for heavy-meson production [42]. Moreover, the kinematic properties of the resulting LLPs are closer to those obtained from primary heavy-meson production than in the electromagnetic-cascade or secondary-kaon cases. LLP production from decays of secondary heavy mesons is therefore important for precision estimates of the total sensitivity, but it is not the main source of the low-energy reconstruction effects studied in this work.

A common feature of most cascade sources relevant here is that they can increase the number of produced LLPs while shifting the LLP spectrum toward lower energies. This may enhance the probability that an LLP decays inside the fiducial volume, since the decay probability in the long-lifetime regime scales approximately as $1/\gamma_X$. At the same time, the lower boost increases the opening angles of the visible decay products and makes them more vulnerable to geometric acceptance of the detector, energy threshold, deflection of charged particles by the spectrometer magnetic field, and reconstruction resolution. The observable impact of cascade production is therefore controlled by a competition between enhanced production and decay probability, on the one hand, and detector-level suppression on the other; we discuss this in more detail in Sec. IV C.

IV. EXISTING STUDIES

Several recent studies have investigated the impact of cascade production on the SHiP sensitivity to light LLPs.

These studies differ in the source particles considered, in the treatment of the target cascade, and in the level at which the detector response is modeled. In this work, we focus on two representative cases that probe the low-energy regimes where detector-level effects are expected to be particularly important. Photophilic ALPs produced in electromagnetic cascades test the case where the secondary source is relatively forward but very soft. HNLs produced in kaon decays test the case where the secondary source is both soft and broadly distributed. The discussion and treatment of other cascade-induced signals, such as vector mediators produced in electromagnetic cascades [36,37] and Higgs-like scalars or ALPs produced in kaon decays, are conceptually similar to the strategy developed below.

A. ALPs coupled to photons

We first consider an ALP coupled dominantly to photons, typically described by the effective interaction

$$\mathcal{L}_a = \frac{g_a}{4} a F_{\mu\nu} \tilde{F}^{\mu\nu}. \quad (8)$$

Here, a is the ALP field, g_a is the ALP-photon coupling, $F_{\mu\nu}$ is the electromagnetic field-strength tensor, and $\tilde{F}^{\mu\nu} = \frac{1}{2} \epsilon^{\mu\nu\alpha\beta} F_{\alpha\beta}$ is its dual.

At proton fixed-target experiments, such ALPs may be produced through several mechanisms. Primary production includes rare decays of neutral mesons

$$P^{\gamma+\gamma^*} \rightarrow \gamma + \gamma + a, \quad P = \pi^0, \eta, \eta', \quad (9)$$

photon fusion in proton-nucleus and proton-proton interactions,

$$p + p(Z)^{\gamma^*+\gamma^*} \rightarrow a + p + p(Z), \quad (10)$$

and the Primakoff process,

$$\gamma + Z \rightarrow a + Z, \quad \gamma + p \rightarrow a + p, \quad (11)$$

initiated by photons produced in the primary proton-target collision [41,43–45]. Various electromagnetic production mechanisms may be initiated by secondary photons, electrons, and positrons generated inside the target.

The dominant ALP decay is $a \rightarrow 2\gamma$. The probabilities of other channels, like $a \rightarrow \gamma + l^+ l^-$ and $a \rightarrow \gamma + \text{hadrons}$, are at the level of $\mathcal{O}(10^{-2})$ [46] and are therefore neglected.

Reference [38] studied the ALP flux induced by secondary photons in the electromagnetic cascade inside the SHiP target. The electromagnetic shower was generated using PETITE [35], a lightweight Monte Carlo tool designed to simulate electromagnetic cascades in thick targets and to provide fast interpolated shower fluxes for dark-sector production. In this framework, shower particles are

generated once and then used as input for different weakly coupled particle production channels, avoiding a full resimulation of the cascade for every model point. The ALP flux was obtained with ALPETITE: channels with a dark-photon analog were obtained by reweighting the corresponding PETITE samples, while Primakoff production was simulated with dedicated Monte Carlo samplers using the shower particles as input.

The resulting event rate in Ref. [38] was computed by assigning weights for ALP production and decay inside the fiducial volume. Additionally, as a proxy of the decay products acceptance, the trajectory of the ALP has been required to be within $\theta_a < \sqrt{S_{\text{deconvol}}}/(2 \times z_{\text{det}}) \approx 0.029$ rad, where $z_{\text{det}} \approx 83$ m is the distance from the target to the beginning of the SHiP spectrometer, and $S_{\text{deconvol}} \approx 4 \times 6$ m is the transverse area of the end of the decay volume. In addition, the energy-threshold requirement of $E_a > 200$ MeV has been applied to each generated ALP candidate.

The authors found that the decay event rate from secondary photons may exceed the rate from primary production modes by a factor of order 50 for light ALPs, while the enhancement decreases with increasing mass and becomes of order unity around $m_a \sim 1$ GeV. This study, therefore, provides a concrete example of the general cascade mechanism discussed above, in which an enhanced low-energy source population competes with the detector requirements needed to observe the LLP decay.

B. HNLs

We next consider HNLs coupled to the Standard Model through mixing with active neutrinos. In a minimal gauge-invariant description, these mixings arise from Yukawa interactions of the form

$$\mathcal{L}_N = Y_\alpha \bar{L}_\alpha \tilde{H} N + \text{H.c.} \quad (12)$$

Here, $L_\alpha, \alpha = e, \mu, \tau$ are the Standard Model lepton doublets, Y_α are the Yukawa couplings, $\tilde{H} = i\sigma_2 H^*$ is the conjugated Higgs doublet, and N is the HNL field. At the scales well below the electroweak transition, the HNL phenomenology at accelerator experiments is controlled by its mass m_N and by the active-sterile mixing angles $U_\alpha \simeq Y_\alpha v / (\sqrt{2} m_N)$, with v being the Higgs vacuum expectation value [11].

At SHiP, HNLs may be produced in decays of charm and beauty hadrons, for example,

$$D \rightarrow \ell_\alpha + N, \quad D \rightarrow h + \ell_\alpha + N, \quad (13)$$

and similarly,

$$B \rightarrow \ell_\alpha + N, \quad B \rightarrow h + \ell_\alpha + N, \quad (14)$$

where $h = K, \pi, \dots$ denotes a hadron. For the mass range $m_N \lesssim m_K - m_{l_a}$ and the HNLs mixing with $\nu_{e,\mu}$, kaon decays provide an additional source,

$$K^\pm \rightarrow \ell_a^\pm + N. \quad (15)$$

For the τ mixing case, the corresponding production mode is kinematically inaccessible.

The HNL decay modes relevant for displaced-decay searches include purely leptonic channels such as $N \rightarrow \ell_a^- \ell_b^+ \nu_\beta$, semileptonic charged channels such as $N \rightarrow \pi^\pm \ell_a^\mp$, and neutral channels such as $N \rightarrow \pi^0 \nu_\alpha$, together with the corresponding charge-conjugate modes [11].

Secondary kaons have been studied as a source of HNLs at SHiP in Ref. [34]. In that work, the kaon population inside the target was obtained using a detailed Geant4 simulation [47–49]. The authors found that roughly 0.29 (0.07) of K^+ (K^-) per proton on target decays in flight and may produce HNLs pointing to the SHiP decay volume, compared to produced yields of 8.01 (3.51). The rest of the kaons are either absorbed through scatterings or decay at rest, producing the flux of isotropic HNLs. As the decay volume of SHiP covers a tiny fraction of the solid angle, the at-rest decays contribute negligibly.

The estimates of the HNL decay event rate were organized as follows. The authors simulated the decays of the in-flight kaons into HNLs, propagated the HNLs to the decay volume, and then simulated their decays. The trajectories of the two visible decay products were then required to intersect the cross section of the end of the decay volume as a minimal decay-product acceptance requirement, and each of them was required to have $p > 1$ GeV.

The study found that, despite the small angular acceptance of HNLs from secondary kaons, the large kaon flux and the reduced HNL boost may provide an additional contribution to the SHiP sensitivity compared with the contribution from charm and beauty meson decays alone.

Production from decays of secondary heavy mesons gives another contribution to the total HNL sensitivity [41]. However, as discussed above, the enhancement of heavy-meson production in the cascade is moderate, and the resulting LLP kinematics are much closer to the primary production case. For this reason, the present work concentrates instead on lower-energy kaon sources.

Below, we consider the HNLs mixing with ν_e for numeric estimates. We comment on the mixing with ν_μ in Sec. VI.

C. Limitations of existing estimates

The observable number of cascade-induced LLP decays is controlled by several distinct factors. In the long-lifetime regime, a useful schematic comparison between the primary and cascade production is

$$\begin{aligned} \frac{N_{\text{ev,sec}}}{N_{\text{ev,prim}}} &\sim \frac{N_{X,\text{prod,sec}}}{N_{X,\text{prod,prim}}} \times \frac{\epsilon_{\text{geom,dec,sec}}}{\epsilon_{\text{geom,dec,prim}}} \times \frac{\langle \gamma_{X,\text{sec}}^{-1} \rangle}{\langle \gamma_{X,\text{prim}}^{-1} \rangle} \\ &\times \frac{\epsilon_{\text{geom,det,sec}}}{\epsilon_{\text{geom,det,prim}}} \times \frac{\epsilon_{\text{reco,sec}}}{\epsilon_{\text{reco,prim}}}. \end{aligned} \quad (16)$$

Here, $N_{X,\text{prod}}$ is the number of produced LLPs, $\epsilon_{\text{geom,dec}}$ is the fraction of LLPs intersecting the decay volume, $\langle \gamma_X^{-1} \rangle$ accounts for the decay probability in the long-lifetime regime, $\epsilon_{\text{geom,det}}$ is the probability that the visible decay products remain inside the relevant detector apertures, and ϵ_{reco} is the probability that the event can be reconstructed with the required vertexing, pointing, invariant mass, and particle-identification performance.

Existing cascade studies have addressed the first factors in this expression in considerable detail. In particular, they have estimated the secondary production rates, the angular and energy distributions of the produced LLPs, and the enhancement of the decay probability associated with the reduced boost. However, the last two factors, $\epsilon_{\text{geom,det}}$ and $\epsilon_{\text{reco,sec}}$, were treated with approximate prescriptions.

First, for the low-energy cascade LLPs, requiring the parent LLP to point toward the decay volume or a downstream plane does not by itself guarantee geometric detector acceptance for the decay products. The visible decay products must themselves remain inside the relevant detector apertures. For a two-body decay into light particles, the typical opening angle grows as the boost decreases, with $\Delta\theta \sim 2/\gamma_X$, with γ_X being the LLP's γ factor. As a result, the daughter-level acceptance can be strongly suppressed for cascade-produced LLPs. Qualitatively, the requirement for the decay products to be within the ECAL acceptance is

$$\Delta\theta \lesssim \frac{\sqrt{A_{\text{ECAL}}}}{\Delta z_{X\text{-ECAL}}}, \quad (17)$$

where $\Delta z_{X\text{-ECAL}} > 12$ m is the distance between the coordinate of LLP decay to the entrance of the ECAL. For the LLPs decaying at the beginning (end) of the decay volume, this requirement leads to the criterion $\gamma_X > 25$ ($\gamma_X > 5$), eliminating the low-energy part of the LLP flux.

Second, the reconstruction efficiency may introduce an additional suppression. Low-energy photons suffer from the worst reconstruction in the ECAL due to the degraded detector relative resolution, later showers, and worse achieved angular resolution. Low-momentum charged particles are more strongly deflected by the spectrometer magnet and may fail to leave enough usable tracking information for vertex and momentum reconstruction. This is particularly relevant for HNLs from secondary kaons, since many visible HNL decay modes contain charged particles. In exclusive channels such as $N \rightarrow \pi^\pm \ell^\mp$, the reconstruction also requires reliable

momentum measurement and particle identification, or at least a controlled assignment of mass hypotheses.

The purpose of the present work is to quantify the impact of these effects for representative low-energy signals at SHiP. We do not address an additional complication associated with low-energy signals, namely, the possible change in the background composition and rejection power. We regard this as the next step in evaluating the SHiP detector performance for cascade-induced signatures, to be carried out once the signal is shown to survive daughter-level acceptance and detector-level reconstruction requirements.

V. OUR APPROACH

Our approach to calculating the event rates with the LLPs is the following.

We use the fluxes of primary and secondary photons for ALPs and kaons for the HNLs as input to generate the angle-energy differential distribution of these LLPs. These are used to conduct two different cross-validated studies. Based on the semianalytic approach from Ref. [41], we estimate the impact of the geometric acceptance of the detector and the Monte Carlo truth selection criteria from Sec. II A on the number of events. Then, in the case of ALPs, where a more detailed investigation is needed (see Fig. 11), we perform a detector-level simulation to accurately calculate reconstruction efficiency, focusing on the domain of low-energy events.

A. Fluxes of primary and secondary SM particles

1. Photons

We have generated the flux of the primary photons originating from the decays of $\pi^0, \eta \rightarrow 2\gamma$ produced in the interactions of the proton beam with the target. For this purpose, we used calibrated Pythia8 on the data from collisions of a 400 GeV proton beam with a thin target

within 20% precision, see Ref. [44] for details. The primary photons were used as an input in the ALPETITE tool from Ref. [38] to generate the flux of secondary photons. The public ALPETITE/PETITE setup contains the molybdenum shower dictionaries needed for the validation run, whereas the current SHiP documents assume a tungsten target. We have therefore generated a dedicated tungsten PETITE dictionary. This is done by reusing the existing 400 GeV VEGAS adaptive phase-space maps and recomputing the integrated cross sections and rejection-sampling maxima for the tungsten material parameters. We keep photons with

$$E_\gamma > 100 \text{ MeV}, \quad (18)$$

following the ALPETITE benchmark setup. The normalization of the cascade photons is fixed by the primary photon yield,

$$w_{\text{prim}} \approx 2 \times \left(N_{\pi^0} \times \text{Br}(\pi^0 \rightarrow 2\gamma) + N_\eta \times \text{Br}(\eta \rightarrow 2\gamma) \right) \approx 8.8, \quad (19)$$

where N_{π^0}, N_η are the amounts of mother mesons per proton on target,

$$N_{\pi^0} \approx 4.2, \quad N_\eta \approx 0.5, \quad (20)$$

and a factor of 2 comes from the decay mode of these mesons into a pair of photons.

The density plots of the angle-energy distribution of the primary and secondary photons are shown in Fig. 3. We see that the secondary photons are much softer on average; namely, the energy of γ 's flying in the direction of the SHiP signal detectors (located at $\theta_\gamma < 0.044$ rad) is $\langle E_{\gamma, \text{casc}} \rangle \approx 1$ GeV versus $\langle E_{\gamma, \text{prim}} \rangle \approx 16.6$ GeV. For the Primakoff conversion, E_γ is an accurate proxy for the ALP energy, $E_a \approx E_\gamma$ [see relation (22)], and hence this property translates to the primary and secondary ALPs.




FIG. 3. Angle-energy distribution $dN_\gamma/d\theta_\gamma dE_\gamma$ for primary (left) and secondary (right) photons at SHiP, normalized per proton on target. The SHiP target is assumed to be made of tungsten. See text for details.




FIG. 4. Angle-energy distribution $dN_K/d\theta_K dE_K$ for charged kaons $K = K^+ + K^-$ at SHiP, normalized per proton on target, as obtained in Ref. [34] assuming a hybrid molybdenum-tungsten SHiP target. The constantlike behavior of the energy distribution with the polar angle θ_K is caused by the extrapolation in the domain of small $\theta_K < \theta_{K,\min} \approx 0.1$, for which the binning of Ref. [34] was not fine enough. Namely, $dN_K/d\theta_K dE_K \sim \frac{\theta_K}{\theta_{K,\min}} \left(\frac{dN_K}{d\theta_K dE_K} \right) \big|_{\theta_K \theta_{K,\min}}$. See text for details.

2. Kaons

For the decaying kaon flux and multiplicities, we will use the results of Ref. [34]. The authors prepared the flux of charged kaons assuming the molybdenum-tungsten target. The flux may be somewhat smaller for the current SHiP target design, which is made solely of tungsten. Nevertheless, this is enough for obtaining qualitative conclusions about the contribution of the cascade HNLs, as we will show in Figs. 7 and 11.

The primary kaons have a small probability of decaying before absorption compared to the secondary ones (due to the high energies of the primaries). For this reason, we effectively treat all the kaons decaying in flight (and hence HNLs from them) as cascade particles.

The angle-energy distribution of all the kaons decaying in flight is shown in Fig. 4. Compared to the photons, the kaons are much more isotropic and softer. This is caused by a strong kaon-energy dependence of the decay probability. Therefore, most of the HNLs would be produced close to isotropically.

B. Generating the LLP flux

Having the events with the primary and secondary photons, as well as kaons, we tabulate them and use them as an input to the SensCalc tool [41].¹ SensCalc is based on a semianalytic method to compute the number of events with different LLPs at various experiments, taking into account the geometry of the decay volume, detector, and the

selection criteria of the decay events based on the Monte-Carlo-truth-level cuts.

The differential rate of the produced ALPs is

$$\frac{dN_{a,\text{prod}}}{d\theta_\gamma dE_\gamma dq^2} \simeq \frac{1}{\sigma_{\gamma,\text{abs}}} \frac{dN_\gamma}{d\theta_\gamma dE_\gamma} \times \left[\frac{d\sigma_{\gamma Z \rightarrow aZ}}{dq^2} + Z \frac{d\sigma_{\gamma p \rightarrow ap}}{dq^2} \right]. \quad (21)$$

Here, the quantity $\sigma_{\gamma,\text{abs}}$ denotes the effective photon absorption cross section in the material, and $q^2 = -(p_a - p_\gamma)^2$ is the momentum transfer to the scattering target. Finally, the two terms in brackets of Eq. (21) describe coherent scattering on the nucleus and incoherent scattering on the individual protons. For the differential Primakoff cross sections, we use the expressions of Ref. [44], including the atomic screening factor for the nuclear target.

The distribution above must be connected to the ALP polar angle θ_a and energy E_a distributions. The previous implementation in SensCalc (Ref. [41]) used the collinear approximation $E_a \approx E_\gamma$, $\theta_a \approx \theta_\gamma$. It is a reliable description for photons with energy $E_\gamma \gg m_a$ but may break down for the soft cascade photons. This motivated an alternative strategy to obtain the ALP distribution: mapping the variables event by event for scattering processes on a target $T = Z, p$ of mass m_T .

Namely, for the given q^2 , we determine the ALP energy is

$$E_a = E_\gamma - \frac{q^2}{2m_T}. \quad (22)$$

In practice, the second term may often be neglected even for the light ALPs. The ALP angle ψ_T with respect to the incoming photon direction is given by

$$\cos(\psi_T) = \frac{2E_\gamma E_a - m_a^2 - q^2}{2E_\gamma \sqrt{E_a^2 - m_a^2}}, \quad (23)$$

with $|\cos(\psi_T)| < 1$. The azimuthal angle ϕ around the photon direction is taken to be uniformly distributed in $[0, 2\pi)$. Therefore, if the incoming photon has polar angle θ_γ with respect to the beam axis, the ALP polar angle is

$$\cos \theta_a = \cos \theta_\gamma \cos \psi_T + \sin \theta_\gamma \sin \psi_T \cos \phi. \quad (24)$$

Using this transformation, the weighted distribution in $(\theta_\gamma, E_\gamma, q^2)$ from Eq. (21) is converted into the final double-differential ALP distribution,

$$\frac{dN_{a,\text{prod}}}{d\theta_a dE_a} \equiv N_{a,\text{prod}} \times f_a(\theta_a, E_a), \quad (25)$$

where $N_{a,\text{prod}}$ is the total number of produced ALPs, and $f_a(\theta_a, E_a)$ is the ALP angle-energy distribution normalized to unity.

¹Available on Zenodo [50] and GitHub [51].




FIG. 5. Overall production probability (per proton on target) of HNLs mixing with ν_e (left) and photophilic ALPs (right). For the ALPs, we have considered the secondary photons with energies $E_\gamma > 100$ MeV. For the cascade HNLs, the shape of the mass dependence of the production probability is mainly caused by the mass dependence of the branching ratio of the process $K \rightarrow N + e$ (chirality suppression at small $m_N \ll m_K$), while for the ALPs it is driven by the steep rise of the photon flux toward lower E_γ (see Fig. 3) and approximate energy correspondence $E_a \approx E_\gamma$ in the Primakoff process [Eq. (22)].

For the HNLs, we straightforwardly simulate their production in kaon decays via the processes $K \rightarrow N + l_\alpha$, where α corresponds to the mixing angle U_α [11].

To demonstrate the potential importance of the cascade production, we show the integrated yields of the ALPs and HNLs per proton on target in Fig. 5. It is clearly seen that, before the geometric acceptance and event reconstruction are taken into account, the secondary production channels may dominate the flux of LLPs X with mass $m_X \lesssim 1$ GeV. In particular, recalling the discussion in Sec. IV C, even if geometric acceptance of the decay volume somehow suppresses the event rate with cascade particles (as may be the case for the HNLs from kaons, recall Fig. 4), this is partially compensated by the $1/\gamma_X$ scaling of the decay probability.

C. Semianalytic event-rate calculation

Once the cascade ALP and HNL fluxes are tabulated, we compute the event rate with their decays using SensCalc. Namely, for an LLP X with mass m_X and coupling to SM particles g_X , the number of decay events is [41]

$$N_{\text{ev}}(m_X, g_X) = N_{X,\text{prod}}(m_X, g_X) \times \int d\theta_X dE_X dz_X f_X(\theta_X, E_X) \frac{\Delta\phi(\theta_X, z_X)}{2\pi} \times \frac{dP_{\text{decay}}}{dz_X} \epsilon_{\text{dec}}(\theta_X, E_X, z_X). \quad (26)$$

Here, $N_{X,\text{prod}}$ is the total number of produced LLPs. θ_X, E_X, z_X are the polar angle, energy, and the longitudinal coordinate of the decaying X with respect to the target. $f_X(\theta_X, E_X)$ is the angle-energy distribution of the produced LLPs, normalized to unity. $\Delta\phi(\theta_X, z_X)$ is the fraction of the azimuthal angle covered by the decay volume. dP_{decay}/dz_X is the differential decay probability of the LLP,

$$\frac{dP_{\text{decay}}}{dz_X} = \frac{\exp\left[-\frac{z_X}{c\tau_X\sqrt{\gamma_X^2-1}\cos(\theta_X)}\right]}{c\tau_X\sqrt{\gamma_X^2-1}\cos(\theta_X)}, \quad (27)$$

with $\tau_X(m_X, g_X)$ being the X 's lifetime. Finally, ϵ_{dec} is the decay acceptance—the fraction of the events with a decaying LLP, where at least two decay products are within the geometric acceptance of the detector and also satisfy other cuts, such as energy, p_T , and transverse spatial separation.

Let us discuss the calculation of ϵ_{dec} in more detail. It is obtained for each point $(\theta_X, E_X, z_X, \{\phi_X\})$ by simulating decays of X into all allowed final states,

$$\epsilon_{\text{dec}}(\theta_X, E_X, z_X) \equiv \sum_f \sum_{\phi_{X,i} \in \{\phi_X\}} \frac{\epsilon_{\text{dec}}^{(f)}(\theta_X, E_X, z_X, \phi_{X,i})}{N_\phi} \times \text{Br}(X \rightarrow f). \quad (28)$$

Here, $\{\phi_X\}$ denotes the set of random N_ϕ azimuthal angles for which the decay point $\mathbf{r}_X(\theta_X, z_X, \phi_X)$ is inside the decay volume. The individual acceptance $\epsilon_{\text{dec}}^{(f)}(\theta_X, E_X, z_X, \phi_{X,i})$ is obtained by simulating a large number of decays of X (1000 by default) and may take a value

$$0 \leq \epsilon_{\text{dec}}^{(f)}(\theta_X, E_X, z_X, \phi_{X,i}) \leq 1. \quad (29)$$

Finally, f is the concrete final state, with $\text{Br}(X \rightarrow f)$ the corresponding branching ratio.

The trajectories of decay products are assumed to be unperturbed by medium interaction throughout the propagation, except for deflection caused by the spectrometer's magnetic field, which is implemented by a kick obtained through the Lorentz force at the end of the magnetized region.

TABLE I. Event selections used in the semianalytic estimate by Eq. (26) (for the detector-level simulation, see Sec. VD instead). The LLP is always required to decay inside the SHiP fiducial decay volume; the three cases differ by the requirements imposed on its visible decay products.

Selection	Decay-volume requirement	Daughter-level requirement	Meaning for $a \rightarrow \gamma\gamma$
$\epsilon_{\text{dec}} \equiv 1$	The LLP decay vertex lies inside the SHiP fiducial decay volume.	No requirement is imposed on the decay products.	All LLP decays inside the fiducial volume are counted, independent of the photon energies, directions, or separation.
Geom only	The LLP decay vertex lies inside the SHiP fiducial decay volume.	At least two visible decay products must reach the downstream detector aperture.	Both photons are required to lie within the aperture of the end of the ECAL.
Baseline	The LLP decay vertex lies inside the SHiP fiducial decay volume.	The baseline SHiP selection is imposed at MC-truth level, Eqs. (4)–(7).	Requirements on the photon energies and their transverse separation at the ECAL.

Additionally, the approach of Eq. (26) may accommodate the reconstruction efficiency map $\epsilon_{\text{rec}}(\theta_X, E_X, z_X)$, which is obtained by fitting the realistic detector response simulations, similar to how this is done in EvtGen [52].

For the calculation of the production and decay rates of HNLs and ALPs, we follow Refs. [8,11,41].

Within its purpose, the semianalytic approach is very accurate and has been tested against various methods, including the ALPINIST framework for ALPs [45] and the full SHiP simulation framework for dark photons and HNLs (see details in Ref. [41]).

We consider three types of decay event selection ϵ_{dec} as summarized in Table I.

D. Detector-level study

In this subsection, we focus on the ALP case. To study the impact of detector-level effects on signal reconstruction, we export the tabulated ALP angle-energy distribution and production probabilities to EventCalc-SHiP.² This tool samples the full decay-event kinematics of LLPs decaying inside the SHiP decay volume. Its inputs consist of the tabulated angle-energy distribution of the produced LLPs, together with the model-dependent decay information: the LLP lifetime as a function of mass and coupling, the decay branching ratios, and, for three-body decays, the squared matrix elements governing the decay-product distributions in the LLP rest frame.

Using tabulated LLP production distributions is useful because it separates the production calculation from the decay and detector-simulation steps. In this way, one does not need to implement every production channel directly inside the event generator. This is especially important for cascade-induced signals, where the production calculation may involve dedicated shower simulations, material-dependent cross sections, and model-specific reweighting

procedures. The tabulated interface, therefore, keeps the event-generation code modular and easier to maintain.

Predictions of EventCalc agree with SensCalc at the level of 5%, both in terms of the integrated quantities and the differential event distributions. However, given the sophisticated method of calculating the ALP flux, we additionally tested it against ALPETITE, namely, comparing the output of EventCalc with Fig. 3 of Ref. [38], finding an excellent agreement (see the Appendix).

The output—the information about the decay vertex coordinates, kinematics, and types of the decay products—may be passed to a detector simulation. Our focus here is on the ECAL performance, and below, we describe in detail the stand-alone Geant4 simulation to study it.

1. Geant4 simulation framework

The tabulated EventCalc-SHiP data are propagated into a custom simulation faithfully representing the SHiP ECAL and HCAL (although only the former is used in this work).

The detector geometry is described using GeoModel [54] and passed onto Geant4, which then handles propagation of particles in matter using the FTFP_BERT physics list. The raw Geant4 hits are digitized to most closely mimic detector response, converting energy deposition into parametrized photon distributions accounting for scintillation light, photon trapping, photon escape, and photodetection by the silicon photomultipliers. The obtained signals are then converted into Analog to Digital Converter counts, assuming usage of the Kanäle für Ladungsauslese von Silicon-photomultiplier ASIC as a front end based on previous experience operating prototypes of the detector [40,55].

The digitized detector response is then converted into reconstructed hits, clustered, and the information is used to reconstruct the event using an early version of the calorimeter reconstruction procedures, which will be documented in a future publication. Two clusters are expected per event when two photons are within detector acceptance.

²Available on GitHub [53].

2. Event vertex and invariant-mass reconstruction

The cluster energy is computed and calibrated in units of minimum-ionizing-particle depositions using a muon particle gun sample as a reference. The spatial distribution of hits in each cluster is then analyzed using a principal component analysis (PCA) [56]. The shower direction is estimated from the eigenvector of the hit-position covariance matrix corresponding to the largest eigenvalue. The cluster directions and positions are then used to compute a point of closest approach vertex. The energies and directions are also used to compute the two-photon invariant mass,

$$m_{\gamma\gamma}^2 = 2E_0E_1(1 - \cos \theta_{\text{open}}), \quad (30)$$

with θ_{open} the opening angle between the two clusters with energies E_0, E_1 . The photon cluster information is also used to derive an impact parameter of the vertex to the target, which may be used to reject background. Said background is nonetheless expected to be negligible and primarily emerging from deep-inelastic-scattering processes within the decay vessel walls. As a result, events are deemed to pass signal-level cuts with three levels of selection being defined:

- (i) The weak selection requires two reconstructed clusters found by the detector. This selection does not assume any vertex reconstruction. It effectively corresponds to the pure geometric requirement for the two photons to be within the ECAL and be separated enough to form nonoverlapping electromagnetic showers.
- (ii) The medium selection extends the weak selection by requiring the reconstructed vertex to lie inside a reduced fiducial decay volume. The reduced volume is obtained by shrinking the nominal decay volume by 1σ vertex resolution in each coordinate. As an example, if the vertex resolution for a given ALP scenario is $\sigma_{\text{vtx}} = \sigma_{\text{vty}} = 5$ cm and $\sigma_{\text{vtz}} = 5$ m in z , the fiducial volume will be defined as the decay volume's pyramidal frustum extending from the 0.9×2.9 m² plane centered around 0 in $z \sim 33$ m to 3.9×5.9 m² in $z \sim 83$ m from $z \sim 38$ to $z \sim 78$ m. No distance of closest approach cut was implemented as it was found to have little impact after the vertex selection.
- (iii) A tight selection, which further expands on the medium one, requires, in addition, that the reconstructed impact parameter to the target be smaller than 1 m to account for the relatively lower resolution in electromagnetic shower vertexing. This selection is only expected to be necessary in the real experiment if the background is found to be significant and is thus deemed very conservative.

The impact-parameter threshold of the tight selection is more stringent than the baseline criterion (6), $\text{IP} < 2.5$ m,

used in the semianalytic study of Sec. VC. We adopt the tighter value throughout the detector-level analysis as a conservative choice, so that the impact-parameter penalties reported in Sec. VIC should be understood as upper estimates.

VI. RESULTS AND DISCUSSION

A. Diagnostic effects of low-energy cascade kinematics

Cascade production may strongly increase the number of produced LLPs, but the resulting LLPs are typically much softer than those from primary production. This has two immediate consequences for displaced-decay searches. First, the visible decay products have larger opening angles, which may prevent them from simultaneously entering the downstream detector acceptance. Second, their lower energies make the reconstruction of individual objects and vertices more challenging. In this subsection, we isolate these effects before combining them into the total event yield.

1. Semianalytic estimates

We start with the semianalytic analysis by Eq. (26) and apply it to the $\epsilon_{\text{dec}} = 1$, Geom only, and baseline selections on decay products summarized in Table I.

To demonstrate the effects of adding the selection criteria, we consider the differential distribution of the decay events in the ALP energy E_a , which is obtained by integrating Eq. (26) in all quantities but E_a . We consider the primary and secondary ALPs for two reference masses, $m_a = 100$ and $m_a = 500$ MeV. We also consider the lifetime $c\tau_a = 300$ m, such that we are in the regime where the differential decay probability in Eq. (26) scales as $dP_{\text{decay}}/dz \approx 1/(c\tau_a \sqrt{\gamma_a^2 - 1} \cos(\theta_a))$, and the differential number of events has a simple scaling with the ALP coupling, $dN_{\text{ev}}/dE_a \propto g_a^4$. We focus on large lifetimes because the contribution of the cascade particles is most significant in this area of the parameter space: the small-lifetime regime prefers high-energy particles, which are naturally more populated by the primary production mechanisms.

The comparison is shown in Fig. 6. After including the detector geometric acceptance, the distributions show a significant suppression of the event rate in the low-energy ALP tail. This is the effect of the minimal opening angle between the decay products, which prevents both photons from simultaneously flying to the detector even if the mother ALP points to it, as discussed in Sec. IV C. The effect is much more prominent for cascade events, as they primarily populate this region of the ALP kinematics.

Adding the energy cut drops the entire energy range $E_a < 2$ GeV and somewhat suppresses the event fraction for the range $2 < E_a \lesssim 10$ GeV because there is a non-negligible fraction of events where one of the photons is below the energy threshold.




FIG. 6. The distribution of the ALP decay events in the ALP energy. Top: $m_a = 100$ MeV; bottom: $m_a = 500$ MeV. The ALP coupling g_a is fixed in a way such that $c\tau_a = 300$ m. The blue lines show the contribution from the primary ALP fluxes, including the Primakoff process, photon fusion, and decays of π^0, η mesons, whereas the red lines show the contribution from the cascade ALP flux originating from the Primakoff process. For each color, pale, dashed, and dark stylings show, correspondingly, the number of events calculated assuming “ $\epsilon_{\text{dec}} = 1$,” “Geom only,” and “baseline” selection criteria summarized in Table I.

The same daughter-level acceptance effect appears for HNLs produced in decays of secondary kaons. In this case, the relevant comparison is between primary HNLs, produced mostly by decays of D mesons, and secondary HNLs from kaon decays. For illustration, Fig. 7 shows the differential event rate for HNLs mixing with ν_e and having mass $m_N = 300$ MeV.

Compared to the ALP case, the suppression of cascade HNLs is even stronger. Secondary kaons have a broad angular distribution, and the HNLs produced in their decays inherit this feature. Therefore, even before considering a dedicated detector-level reconstruction, the purely geometric requirement on the visible decay products already removes a large fraction of the cascade contribution.

For the mixing with ν_μ , all the conclusions qualitatively hold, except for slightly better geometric acceptance. Indeed, decays of charged-current HNL mixing with ν_μ




FIG. 7. The distribution of the HNL decay events in the HNL energy for $m_N = 300$ MeV, assuming the mixing with electron neutrinos, for the lifetime $c\tau_N \approx 2.3 \times 10^7$ m. The meaning of the curves is the same as in Fig. 6, modulo the different sources of primary and cascade HNLs, as specified in the text.

contain μ instead of e , meaning that the decay products have less momentum in the rest frame and hence less angular spread.

2. Detector-level reconstruction

Let us now turn to the detector-level reconstruction of the ALP signal. A representative working point, $m_a = 1$ GeV and $c\tau_a = 1$ m, is shown in Fig. 8 to illustrate the absolute scale of the vertex, invariant-mass, and impact-parameter resolutions. Here, however, the main question is comparative: whether the detector reconstructs the soft cascade population as efficiently as the harder primary one.

This comparison is summarized in Fig. 9, which shows the cascade-to-primary ratios of the reconstructed resolutions across the $(m_a, c\tau_a)$ plane. Because cascade ALPs are systematically softer than primary ones, their decay photons are lower in energy and wider in opening angle at fixed mass. This is the single kinematic characteristic behind the behavior of all panels.

For the transverse vertex coordinates, the cascade resolution is poorer across essentially the whole plane: softer photons give broader, lower-amplitude showers, whose PCA directions are noisier. As a result, the $\sigma(x)$ and $\sigma(y)$ ratios sit slightly above unity. The longitudinal resolution is instead governed by the diphoton opening angle through the point-of-the-closest-approach vertex, which is the worst-constrained direction for a two-photon system. The ratio here reflects the competition between the larger cascade opening angle, which improves the conditioning of the z vertex, and the noisier per-cluster directions, and therefore stays closer to unity. Both primary and cascade resolutions degrade toward low ALP masses, where the minimum opening angle $\Delta\theta_{\min} \sim 2m_a/E_a$ collapses, and the photon pair becomes collinear. However, the degradation sets in earlier for the cascade population, which is why the lightest rows of Fig. 9 depart most strongly from




FIG. 8. Vertex, invariant-mass, and impact-parameter resolutions for a 1 GeV, $c\tau_a = 1$ m ALP. The efficiencies after fiducial-volume selection cuts are also provided for successive selections: decays occurring in the decay volume, and both photons in acceptance are taken at the truth level. Events with two reconstructed clusters, events with a vertex in fiducial acceptance, and events with an impact parameter below 1 m are taken at the reconstruction level. Each selection includes the previous conditions.




FIG. 9. Ratios of the resolution in various reconstructed quantities between the events with cascade and primary ALPs, for different masses m_a on the y axis and $c\tau_a$ on the x axis. From left to right, the panels show the invariant mass $m_{\gamma\gamma}$ and the x, y, z coordinates of the ALP decay vertex. The $m_a = 50$ MeV, $c\tau_a = 1$ m sample was interpolated between its same-mass neighbors.

unity. The dependence on $c\tau_a$ is the imprint of the surviving-energy spectrum: short lifetimes admit only the high-boost tail that reaches $z \gtrsim 33$ m, whose collimated photons inflate σ_z , while longer lifetimes admit a softer, wider-angle sample.

B. Event-rate impact in the nominal SHiP setup

We now combine the acceptance and reconstruction effects discussed above and evaluate the cascade contribution in the nominal SHiP setup as described in Sec. II.

For ALPs, the detector-level cutflow is shown in Fig. 10. It follows the event sample from the decay-in-volume stage

to the final reconstructed selection and, therefore, shows where the initial advantage of cascade production is lost throughout the reconstruction.

The efficiency comparison of Fig. 10 shows a clear cutflow pattern. At the decay-in-volume stage, the reduced cascade boost raises the in-volume decay probability through the $1/\gamma_X$ scaling of Eq. (27). Given this and also the discussion around Fig. 5, the cascade-to-primary ratio is largest toward large lifetimes and low masses. This advantage is then removed, step by step, by the requirements that make the event reconstructible.

The two-photon acceptance is the dominant suppressor. The wider opening angle of soft cascade decays makes it far




FIG. 10. Observed ratios of the number of events between cascade and primary ALPs, for different ALP masses m_a on the y axis and lifetimes $c\tau_a$ on the x axis, as obtained in the detector-level simulation. From left to right, the panels show the ratio of the number of events for ALPs decaying inside the decay volume; where both photons from the ALP decay are within the ECAL acceptance; where they form two separate clusters; where additionally the ALP decay vertex lies within the fiducial volume; and with the additional requirement on the transverse impact parameter, $IP < 1$ m. The first two panels correspond to the MC-truth results, whereas the rest are obtained using the detector simulation of the signal reconstruction. The $m_a = 50$ MeV, $c\tau_a = 1$ m sample was interpolated between its same-mass neighbors instead of using the simulation data, due to an artifact in the latter.

less likely that both photons reach the A_{ECAL} face. The two-cluster requirement removes a further low-mass fraction where the surviving pairs' increased collinearity degrades the clustering efficiency. Finally, the vertex and impact-parameter selections add an additional reconstruction-level suppression. The suppression coming from the impact-parameter cut is again stronger for cascade events than for primary events (we comment on this in Sec. VIC).

Let us now compare the integrated event rates from primary and cascade production. We use the simplified semianalytic analysis for both ALPs and HNLs and the full detector-level simulation for ALPs. The quantity shown in Fig. 11 is

$$N_{\text{ev,cascade}}/N_{\text{ev,primary}} \quad (31)$$

as a function of the LLP mass, in the regime where the decay length is large compared to the SHiP detector scale.

Let us start with the ALP result in the left panel of Fig. 11. The $\epsilon_{\text{dec}} = 1$ line shows a huge domination of

cascade events over the primary ones. However, already requiring both photons to be within the acceptance of the ECAL, shown by the Geom only curve, suppresses the secondary yield by the relative factor of $\simeq 4$ compared to the $\epsilon_{\text{dec}} = 1$ line. Finally, assuming the baseline selection criteria, the ratio (31) decreases further, especially in the domain of light ALPs with $m_a \lesssim 200$ MeV. Overall, for a 50-MeV ALP, the ratio decreases from $\simeq 70$ to $\simeq 3$ when going from $\epsilon_{\text{dec}} = 1$ to baseline.

On top of this, in the same panel, we show the results of the detector-level analysis. Given that it uses EVENTCALC for the signal generation, and EVENTCALC agrees well with SensCalc, the comparison between the ratios shows how realistic detector-level reconstruction differs from simple Monte Carlo truth cuts. The detector-level ratio sits close to the baseline curve, and its dependence on the ALP mass follows a similar trend except for the domain $m_a \lesssim 100$ MeV. There, the hard cut on the photon energy imposed in the semianalytic calculation causes the plateau-like behavior of the cascade-to-primary ratio, whereas the




FIG. 11. The ratio of the number of events with LLPs produced from cascade sources versus those having primary origin. Left: photophilic ALPs, for ALP lifetimes such that $c\tau_a \langle \gamma_a \rangle \gtrsim 100$ m. The continuous curves are obtained by the semianalytic approach with Eq. (26), and the meaning of the legends is the same as in Fig. 6. The points come from the detector-level simulation described in Secs. VD and VIA 2. Right: HNLs mixing with electron neutrinos. The meaning of the curves is the same, modulo the different sources of primary and cascade HNLs, as specified in the text. The nonmonotonic behavior of the HNL ratio is caused by opening up various HNL decay modes at $m_N \simeq m_\pi$, such as $N \rightarrow \pi^0 \nu$ and $N \rightarrow \pi + e$.

detector simulation shows a monotonic (yet slow) growth. This indicates that the semianalytic baseline selection captures the dominant suppression of the cascade ALP contribution, at least for the ratio of primary-to-cascade events.

The HNL result is shown in the right panel of Fig. 11. The comparison shows that, already after including the purely geometric part of the decay-product acceptance, the number of events from cascade HNLs becomes subdominant to the primary source. Moreover, the kaon flux used here was obtained for the hybrid molybdenum-tungsten target, which has a larger kaon interaction length and hence gives more chances for the kaons to decay in flight. If switching to the actual tungsten-only target, the flux of decaying kaons would become even more isotropic, suppressing the geometric acceptance for the HNLs and their decay products even further.

Given this, we do not perform a dedicated study of the reconstruction performance of low-energy HNL decay products. As discussed above, an additional complexity is that most of the visible HNL decay modes contain charged particles, and low-momentum species may be deflected out of the detector coverage by the spectrometer magnetic field. There is also the decay mode $N \rightarrow \pi^0 \nu$, where the detectable final state is a two-photon system. However, since the cascade contribution is already subdominant after imposing geometric acceptance, we leave a detector-level study of such channels for future work.

C. Sensitivity to the event-selection criteria

Let us now comment on the two selection requirements that most directly affect the detector-level cascade ALP yield: the threshold on the energy of the individual decay photons, Eq. (4), and the impact-parameter requirement, Eq. (6).

We begin with the energy threshold. Figure 12 quantifies the ECAL response to a threshold $E_\gamma^{\min}$ applied to the truth-level energy of the softer of the two decay photons; since Eq. (4) constrains each photon separately, an event survives the criterion if and only if its softer photon clears the threshold. Both panels are conditioned on the two photons being within the ECAL acceptance, so that the geometric suppression—quantified in the cutflow of Fig. 8 and, in cascade-to-primary ratio, in the second panel of Fig. 10—is removed. The left panel shows the conditional reconstruction efficiency, evaluated on the subsample passing the threshold: for the two-cluster requirement it is flat at $\simeq 0.78$ (primary) and $\simeq 0.93$ (cascade) below $\simeq 1$ GeV, and decreases above it for both origins. The decrease reflects the fact that, at this low mass, raising the threshold selects progressively more boosted decays, whose photons are increasingly collimated and degrade clustering efficiency. The right panel shows the same numerator normalized to the full in-acceptance sample and, therefore, gives the fraction of the signal retained by a cut at a given energy.




FIG. 12. Reconstruction performance for 100 000 $a \rightarrow \gamma\gamma$ decays as a function of a threshold $E_\gamma^{\min}$ imposed on the truth-level energy of the softer of the two decay photons, for the benchmark $m_a = 100$ MeV, $c\tau_a = 1$ m, a low mass where cascade enhancement is significant. The mass is the same as in Fig. 13, which shows several lifetimes at this mass including $c\tau_a = 1$ m, so that the two figures can be read together. Both panels are restricted to decays for which both photons lie within the ECAL acceptance, so that the geometric acceptance—whose impact is shown separately in Figs. 8 and 10—is removed and only the detector-level response is displayed. Left: conditional reconstruction efficiency; at each $E_\gamma^{\min}$ the numerator is the number of in-acceptance events above the threshold that are successfully reconstructed under the stated conditions (two reconstructed clusters, with or without the additional impact-parameter requirement), and the denominator is the number of in-acceptance events above the same threshold, so that the denominator moves with said threshold. Right: cumulative surviving fraction, i.e., the same numerator normalized to the fixed denominator of *all* in-acceptance events; here “cumulative” refers to the fraction of the signal surviving a cut at $E_\gamma^{\min}$, so this panel measures directly how much signal is discarded by the nominal criterion (4).

Comparing the low-threshold plateau of the right panel with its value at 1 GeV gives the signal recoverable by relaxing Eq. (4). At the two-cluster stage, the primary sample falls from ≈ 0.78 to ≈ 0.66 at 1 GeV, so the nominal threshold removes ≈ 0.12 of the in-acceptance signal: $\approx 15\%$ of the low-threshold plateau or, equivalently, a $\approx 18\%$ increase on the signal retained at 1 GeV if the criterion were fully relaxed. For cascade ALPs, the corresponding drop is from ≈ 0.93 to ≈ 0.64 , i.e., ≈ 0.29 , which is $\approx 31\%$ of the plateau or a gain of $\approx 45\%$ on the retained signal. Relaxing the threshold therefore recovers proportionally more cascade signal, as expected from their softer spectrum. Once the impact-parameter requirement is added, however, both curves are essentially flat across 1 GeV, so that lowering the threshold recovers almost nothing: the soft events admitted by a relaxed Eq. (4) are events that the impact-parameter requirement rejects in any case. The decrease of these curves at higher thresholds instead reflects the loss of genuinely reconstructible high-energy events and is not relevant to relaxing the criterion. Relaxing the energy threshold is therefore useful only in combination with a relaxed impact-parameter requirement, which we return to below.

Because the cascade spectrum is much softer ($\langle E_{\gamma, \text{casc}} \rangle \approx 1$ GeV versus $\langle E_{\gamma, \text{prim}} \rangle \approx 16.6$ GeV, recall Sec. VA), cascade events populate the low-threshold part of these curves, whereas primary events sit close to the plateau. This spectral effect is not, however, the whole story: at fixed softer-photon energy, the primary and cascade samples are not equivalent, and the comparison changes sign between the two selection stages. At the two-cluster level, the cascade curve lies above the primary one, by ≈ 0.15 in absolute efficiency at low threshold, the two converging as the threshold increases: at a given softer-photon energy, the cascade sample carries lower total ALP energy, so the decay opening angle is wider and the two showers are more readily resolved into separate clusters provided that both reach the calorimeter acceptance. Adding the requirement $\text{IP} < 1000$ mm reverses the ordering below ≈ 6 GeV and reduces the cascade efficiency to roughly half the primary one, since the decay direction is reconstructed from the cluster positions and degrades both with softer photons and with the shorter lever arm available for less boosted decays (recall Sec. IV C). The suppression of the cascade contribution after the full selection is therefore not of purely spectral origin: a residual detector-level penalty survives at fixed photon energy, and it enters through the directionality requirement rather than through cluster reconstruction. The nonmonotonic shape of the impact-parameter curves, which peak near 5–7 GeV, follows from two competing trends visible in the same panel: the fraction of reconstructed events passing the IP requirement grows with photon energy as the directional resolution improves, while the two-cluster efficiency itself falls above ≈ 1 GeV, as described above, because the increasingly collimated




FIG. 13. Relative signal efficiency as a function of the impact-parameter cut for the current reconstruction for several 100 MeV ALP samples. The signal efficiency is observed to fall faster for cascade ALPs compared to primary ones in the larger- $c\tau_a$ cases, where the cascade enhancement is most significant.

decays lead to degraded clustering efficiency. The peak marks the point at which the second effect overtakes the first.

Figure 13 shows the integrated counterpart of this effect, now as a function of the impact-parameter cut itself and for the same benchmark mass $m_a = 100$ MeV, so that the two figures can be read together. The requirement is especially detrimental to cascade events: because cascade ALPs are softer and their reconstructed directions are less precise, this cut removes relatively more cascade events than primary events.

If the impact-parameter cut can be softened, or removed altogether for sufficiently clean two-photon final states, the cascade enhancement may become significantly larger. For example, for 100 MeV ALPs with $c\tau_a > 30$ m, a simple two-cluster requirement may lead to an enhancement of up to a factor of ~ 8 . Whether this requirement can, in practice, be softened depends on the background composition and rejection power and, therefore, requires a dedicated background study.

To highlight the importance of the reconstruction of the low-energy events, we show in Fig. 14 the sensitivity of SHiP to ALPs as obtained using the semianalytic approach, under the requirements on decay products from Table I. The difference between the $\epsilon_{\text{dec}} = 1$ and baseline lines at the lower bound of the sensitivity may reach a factor ≈ 4 . Since the number of events scales as $N_{\text{ev}}(g_a) \propto g_a^4$, this would translate to a huge factor of $\lesssim 200$ enhancement in the event rates.

VII. CONCLUSIONS

Cascade interactions inside the thick target of the SHiP experiment may produce long-lived particles, whose decays can dominate the event rate inside the hidden-sector




FIG. 14. Sensitivity of the SHiP experiment to the ALPs obtained using the semianalytic approach via Eq. (26), as $N_{\text{ev}} \geq 2.3$, in the plane ALP mass m_a -ALP coupling g_a . The meaning of the lines is provided in Table I. The darker-gray and lighter-gray domains are constraints from laboratory searches and astrophysical observations, taken from Ref. [25]. In the low-mass region, SHiP as a laboratory probe would cover the range of couplings complementarily to the astrophysical bounds.

decay volume before detector-level requirements are imposed. In this paper, we have quantified the contribution of these cascade-induced signal sources relative to primary production mechanisms, focusing on the regime in which the produced LLPs are relatively soft, with typical energies $E_{\text{LLP}} \lesssim \text{few GeV}$.

We considered two representative case studies: photophilic ALPs, discussed in Sec. IV A, and HNLs, discussed in Sec. IV B. These particles can be produced from cascade photons and secondary kaons, respectively, as described in Sec. V A. Extending previous studies, we evaluated how the event rate is affected by two detector-level effects: first, the geometric acceptance of the detector, which requires the LLP decay products to remain within the SHiP detector system; and second, the reconstruction efficiency for low-energy decay products.

For the nominal SHiP setup, our main results are summarized in Figs. 10 and 11. For light LLPs, cascade production can significantly exceed the primary yield before imposing requirements on the decay products. However, the geometric acceptance of the detector already substantially suppresses this cascade-driven enhancement. For photophilic ALPs, the inclusion of realistic event reconstruction further reduces the effect, leaving only a moderate enhancement for the lightest masses. For HNLs from secondary kaons, the cascade contribution becomes subdominant already after imposing the geometric acceptance of the visible decay products.

At the same time, our results indicate that the SHiP search program could benefit from improving the detectability of such low-energy signals. One possibility is to enlarge the effective angular coverage of the electromagnetic calorimeter and to improve the background-rejection strategy, which may allow the transverse impact-parameter

requirement to be relaxed. As shown in Fig. 13, this requirement removes cascade ALP events more efficiently than primary ones, especially in the long-lifetime region where the cascade enhancement is most relevant.

Another possibility is to consider more usage of active decay-volume subdetectors, where the decay products could be reconstructed closer to the LLP decay vertex. A straightforward option such as the scattering detector at SHiP does not appear well suited for this purpose, given its small size and the intense muon background. An existing option would be the use of the SHiP surrounding background tagger [57] to attempt reconstruction of LLPs reaching the decay vessel but whose final states fail to (all) reach the downstream detectors. This would imply leveraging the detector's full coverage of the decay vessel and $\mathcal{O}(\text{cm})$ position resolution to reconstruct otherwise incomplete or fully missing events. The detector has poor efficiency for photon reconstruction but may allow one to recover part of the cascade enhancement in the lower-mass region. A different possibility is an LAr-based setup placed at the end of the SHiP detector, as discussed in Ref. [58]. A robust assessment of such an option would require a dedicated study of signal reconstruction and backgrounds, which we leave for future work.

This makes the present study particularly timely. By quantifying the event-rate suppression induced by realistic detector acceptance and reconstruction, we show that reliable sensitivity estimates for cascade-induced LLP signals require detector effects to be included up to the reconstruction level. At the same time, since SHiP is still in the TDR phase, there remains a concrete opportunity to optimize the setup and recover part of the cascade-induced event rate.

ACKNOWLEDGMENTS

The authors thank Yotam Soreq for useful discussions, Dmitry Gorbunov for discussions on the kaon production inside a thick SHiP target, and Nikita Blinov, Andrei Golovin, Anne-Marie Magnan, and Ryan Plestid for useful comments on the manuscript. M. C. has received support from the FWO, under Grant No. 12A4O26N (Belgium). Some of the authors (M. C., M. O.) are also members of the SHiP Collaboration. The present manuscript solely reflects the authors' views and not those of the SHiP Collaboration.

DATA AVAILABILITY

There are no publicly available research data or software supporting this manuscript. Requests for further information or data should be sent to the authors.

APPENDIX: COMPARISON WITH REF. [38]

To cross-check our calculation of the ALP flux, we have compared our results to Fig. 3 of Ref. [38]. For these purposes, within EVENTCALC, we have obtained the ALP




FIG. 15. The energy distribution of the number of events with ALPs having mass $m_a = 100$ MeV and lifetime $c\tau_a = 300$ m, assuming no cuts on decay products but requiring the ALP trajectory to satisfy $\theta_a < 0.029$ rad, as obtained by EVENTCALC (solid) and ALPETITE (dashed), for the primary and secondary event sources.

events energy distribution dN_{ev}/dE_a for mass $m_a = 100$ MeV and $c\tau_a = 300$ m, assuming no cuts on decay products, but imposing the cut $\theta_a < 0.029$ rad on the decaying ALP (this is the simplified analog of the decay-product acceptance used in Ref. [38]); the impact of this cut is quite low, because the ALPs are strongly collimated. Next, we have rescaled Fig. 3 of Ref. [38] by $N_{\text{PoT}} \times f_{c\tau_a=300\text{ m}}^{-4} \approx 3.6 \times 10^5 \text{ GeV}^{-4}$, where $f_{c\tau_a=300\text{ m}}^{-1} \approx 1.6 \times 10^{-4} \text{ GeV}^{-1}$. Finally, we have rescaled it by the factor $(1 - \exp(-1))^{-1}$ to account for a conservative assumption that only the first proton interaction length is

included within the target when obtaining the primary photon flux.

The resulting comparison is shown in Fig. 15. The results are in excellent agreement except for the domain $E_a \lesssim 1$ GeV. The small discrepancy there is likely due to a combination of a slightly different primary photon flux used in the calculations and also different target materials assumed in the calculation (molybdenum versus tungsten). The agreement confirms the mutual agreement in calculating the ALP number of events—from computing the photon flux to accounting for decay-volume geometry when sampling ALP decay events.

-
- [1] T. Asaka, S. Blanchet, and M. Shaposhnikov, The nuMSM, dark matter and neutrino masses, *Phys. Lett. B* **631**, 151 (2005).
 - [2] M. Shaposhnikov, A possible symmetry of the nuMSM, *Nucl. Phys. B* **763**, 49 (2007).
 - [3] M. Pospelov, A. Ritz, and M. B. Voloshin, Secluded WIMP dark matter, *Phys. Lett. B* **662**, 53 (2008).
 - [4] F. Bezrukov and D. Gorbunov, Light inflaton Hunter's guide, *J. High Energy Phys.* **05** (2010) 010.
 - [5] G. Krnjaic, Probing light thermal dark-matter with a Higgs portal mediator, *Phys. Rev. D* **94**, 073009 (2016).
 - [6] A. Berlin and F. Kling, Inelastic dark matter at the LHC lifetime frontier: ATLAS, CMS, LHCb, CODEX-b, FASER, and MATHUSLA, *Phys. Rev. D* **99**, 015021 (2019).
 - [7] A. Hook, S. Kumar, Z. Liu, and R. Sundrum, High quality QCD axion and the LHC, *Phys. Rev. Lett.* **124**, 221801 (2020).
 - [8] D. Gorbunov and M. Shaposhnikov, How to find neutral leptons of the νMSM ?, *J. High Energy Phys.* **10** (2007) 015; **11** (2013) 101(E).
 - [9] P. Ilten, Y. Soreq, M. Williams, and W. Xue, Serendipity in dark photon searches, *J. High Energy Phys.* **06** (2018) 004.
 - [10] D. Aloni, Y. Soreq, and M. Williams, Coupling QCD-scale axionlike particles to gluons, *Phys. Rev. Lett.* **123**, 031803 (2019).
 - [11] K. Bondarenko, A. Boyarsky, D. Gorbunov, and O. Ruchayskiy, Phenomenology of GeV-scale heavy neutral leptons, *J. High Energy Phys.* **11** (2018) 032.
 - [12] I. Boiarska, K. Bondarenko, A. Boyarsky, V. Gorkavenko, M. Ovchinnikov, and A. Sokolenko, Phenomenology of GeV-scale scalar portal, *J. High Energy Phys.* **11** (2019) 162.
 - [13] S. Foroughi-Abari and A. Ritz, Dark sector production via proton bremsstrahlung, *Phys. Rev. D* **105**, 095045 (2022).

- [14] M. Ovchinnikov and A. Zaporozhchenko, Advancing the phenomenology of GeV-scale axionlike particles, *Phys. Rev. D* **112**, 015001 (2025).
- [15] S. Alekhin *et al.*, A facility to search for hidden particles at the CERN SPS: The SHiP physics case, *Rep. Prog. Phys.* **79**, 124201 (2016).
- [16] P. Ilten, Y. Soreq, J. Thaler, M. Williams, and W. Xue, Proposed inclusive dark photon search at LHCb, *Phys. Rev. Lett.* **116**, 251803 (2016).
- [17] J. L. Feng, I. Galon, F. Kling, and S. Trojanowski, ForWard Search Experiment at the LHC, *Phys. Rev. D* **97**, 035001 (2018).
- [18] W. Altmannshofer *et al.* (Belle-II Collaboration), The Belle II physics book, *Prog. Theor. Exp. Phys.* **2019**, 123C01 (2019); **2020**, 029201(E) (2020).
- [19] J. Beacham *et al.*, Physics beyond colliders at CERN: Beyond the standard model working group report, *J. Phys. G* **47**, 010501 (2020).
- [20] D. Banerjee *et al.* (NA64 Collaboration), Search for Axion-like and Scalar Particles with the NA64 Experiment, *Phys. Rev. Lett.* **125**, 081801 (2020).
- [21] A. Blondel *et al.*, Searches for long-lived particles at the future FCC-ee, *Front. Phys.* **10**, 967881 (2022).
- [22] C. Antel *et al.*, Feebly interacting particles: FIPs 2022 workshop report, [arXiv:2305.01715](https://arxiv.org/abs/2305.01715).
- [23] V. Gorkavenko, B. K. Jashal, V. Kholoimov, Y. Kyselov, D. Mendoza, M. Ovchinnikov, A. Oyanguren, V. Svintozelskyi, and J. Zhuo, LHCb potential to discover long-lived new physics particles with lifetimes above 100 ps, *Eur. Phys. J. C* **84**, 608 (2024).
- [24] A. Abed Abud *et al.* (DUNE Collaboration), DUNE Phase II: Scientific opportunities, detector concepts, technological solutions, *J. Instrum.* **19**, P12005 (2024).
- [25] J. de Blas *et al.*, Physics briefing book: Input for the 2026 update of the European Strategy for Particle Physics, [arXiv:2511.03883](https://arxiv.org/abs/2511.03883).
- [26] R. Balkin, T. Coren, A. Jentsch, H. Liu, M. Ovchinnikov, Y. Soreq, and S. Trifinopoulos, Braking protons at the EIC: From invisible meson decay to new physics searches, [arXiv:2601.00068](https://arxiv.org/abs/2601.00068).
- [27] O. Aberle *et al.* (SHiP Collaboration), BDF/SHiP at the ECN3 high-intensity beam facility, Technical Report, CERN, Geneva, 2022, [http://cds.cern.ch/record/2839677](https://cds.cern.ch/record/2839677).
- [28] R. Albanese *et al.* (SHiP and HI-ECN3 Project Team Collaboration), SHiP experiment at the SPS beam dump facility, [arXiv:2504.06692](https://arxiv.org/abs/2504.06692).
- [29] K. Bondarenko, A. Boyarsky, R. Jacobsson, O. Mikulenko, and M. Ovchinnikov, Towards the optimal beam dump experiment to search for feebly interacting particles, *Eur. Phys. J. C* **83**, 1126 (2023).
- [30] R. Albanese (SHiP Collaboration), BDF/SHiP annual report 2025, Technical Report, CERN, Geneva, 2025, <https://cds.cern.ch/record/2948477>.
- [31] C. Ahdida *et al.* (SHiP Collaboration), Sensitivity of the SHiP experiment to heavy neutral leptons, *J. High Energy Phys.* **04** (2019) 077.
- [32] C. Ahdida *et al.* (SHiP Collaboration), Sensitivity of the SHiP experiment to dark photons decaying to a pair of charged particles, *Eur. Phys. J. C* **81**, 451 (2021).
- [33] C. Ahdida *et al.* (SHiP Collaboration), Sensitivity of the SHiP experiment to light dark matter, *J. High Energy Phys.* **04** (2021) 199.
- [34] D. Gorbunov, I. Krasnov, Y. Kudenko, and S. Suvorov, Heavy neutral leptons from kaon decays in the SHiP experiment, *Phys. Lett. B* **810**, 135817 (2020).
- [35] N. Blinov, P. J. Fox, K. J. Kelly, P. A. N. Machado, and R. Plestid, Dark fluxes from electromagnetic cascades, *J. High Energy Phys.* **07** (2024) 022.
- [36] T. Zhou, R. Plestid, K. J. Kelly, N. Blinov, and P. J. Fox, Long-lived vectors from electromagnetic cascades at SHiP, *J. High Energy Phys.* **02** (2025) 107.
- [37] N. Blinov, P. J. Fox, K. J. Kelly, R. Plestid, and T. Zhou, $L_\mu - L_\tau$ gauge bosons in beam dumps and supernovae, [arXiv:2511.09619](https://arxiv.org/abs/2511.09619).
- [38] S. Patrone, N. Blinov, and R. Plestid, Long-lived axionlike particles from electromagnetic cascades, *Phys. Rev. D* **113**, 075009 (2026).
- [39] W. M. Bonivento (SHiP Collaboration), Studies for the electro-magnetic calorimeter SplitCal for the SHiP experiment at CERN with shower direction reconstruction capability, *J. Instrum.* **13**, C02041 (2018).
- [40] M. A. V. Climescu, Design and development of calorimetry at SHiP and SND@LHC, Ph.D. thesis, Mainz U., Johannes Gutenberg University Mainz, 2025.
- [41] M. Ovchinnikov, J.-L. Tastet, O. Mikulenko, and K. Bondarenko, Sensitivities to feebly interacting particles: Public and unified calculations, *Phys. Rev. D* **108** (2023), 075028.
- [42] H. Dijkstra and T. Ruf (SHiP Collaboration), Heavy flavour cascade production in a beam dump, Online at <https://cds.cern.ch/record/2115534> (2015).
- [43] B. Döbrich, J. Jaeckel, F. Kahlhoefer, A. Ringwald, and K. Schmidt-Hoberg, ALPtraum: ALP production in proton beam dump experiments, *J. High Energy Phys.* **02** (2016) 018.
- [44] B. Döbrich, J. Jaeckel, and T. Spadaro, Light in the beam dump-ALP production from decay photons in proton beam-dumps, *J. High Energy Phys.* **05** (2019) 213; **10** (2020) 46.
- [45] J. Jerhot, B. Döbrich, F. Ertas, F. Kahlhoefer, and T. Spadaro, ALPINIST: Axion-like particles in numerous interactions simulated and tabulated, *J. High Energy Phys.* **07** (2022) 094.
- [46] M. Escudero Abenza, C. Garcia-Perez, and M. Ovchinnikov, Nucleosynthesis and CMB bounds on photophilic ALPs: A fresh look, *Eur. Phys. J. C* **86**, 500 (2026).
- [47] S. Agostinelli *et al.* (Geant4 Collaboration), Geant4—A simulation toolkit, *Nucl. Instrum. Methods Phys. Res., Sect. A* **506**, 250 (2003).
- [48] J. Allison *et al.*, Geant4 developments and applications, *IEEE Trans. Nucl. Sci.* **53**, 270 (2006).
- [49] J. Allison *et al.*, Recent developments in Geant4, *Nucl. Instrum. Methods Phys. Res., Sect. A* **835**, 186 (2016).
- [50] M. Ovchinnikov, SensCalc. Online at <https://doi.org/10.5281/zenodo.20030983> (2026).
- [51] M. Ovchinnikov, SensCalc. Online at <https://github.com/maismovchinnikov/SensCalc> (2026).
- [52] D. J. Lange, The EvtGen particle decay simulation package, *Nucl. Instrum. Methods Phys. Res., Sect. A* **462**, 152 (2001).

- [53] M. Ovchynnikov, EventCalc-SHiP. Online at <https://github.com/maksymovchynnikov/EventCalc-SHiP> (2026).
- [54] M. Bandieramonte, R.M. Bianchi, J. Boudreau, A. Dell'Acqua, and V. Tsulaia, The GeoModel tool suite for detector description, *Eur. Phys. J. Web Conf.* **251**, 03007 (2021).
- [55] K. Briggel, H. Chen, D. Schimansky, W. Shen, V. Stankova, and H.C. Schultz-Coulon, KLauS: A low power silicon photomultiplier charge readout ASIC in 0.18 UMC CMOS, *J. Instrum.* **11**, C03045 (2016).
- [56] H. Wind, Principal component analysis and its applications to track finding.
- [57] A. Brignoli *et al.*, Performance of prototypes with different reflector materials for the SHiP liquid scintillator surrounding background tagger, *J. Instrum.* **20**, P07023 (2025).
- [58] M. Ferrillo, M. Ovchynnikov, F. Resnati, and A. De Roeck, Improving the potential of BDF@SPS to search for new physics with liquid argon time projection chambers, *J. High Energy Phys.* **02** (2024) 196.