Displaced Vertex Searches for Electroproduced Strongly Interacting Massive Particles with the 2016 HPS Dataset

(HPS Collaboration)

(Dated: April 30, 2025)

The Heavy Photon Search (HPS) is a fixed target direct detection experiment looking for the production of thermal relic dark matter (DM) at Jefferson Laboratories' CEBAF facility. In a 2016 engineering run, it collected 10 608 nb^{-1} of cross section with electron beam energies of 3 GeV. This paper advances the models of DM that can be excluded using HPS data, namely to Strongly Interacting Massive Particles (SIMPs) whose thermal relic abundance is chiefly determined by a 3 to 2 DM decay. We will assume for our model of dark matter a $SU(3)$ dark matter field coupled via a dark photon to Standard Model; this yields 2 and 3 body interaction cross sections with which we will constrain 1–1000 MeV SIMP dark matter. We expand the known exclusion contour in a small region around ~ 70 MeV, and the methods we develop will establish a basis by which later physics runs can dramatically push the contour forward.

I. INTRODUCTION

In recent years, a number of extensions to the Standard Model (SM) including new gauge symmetries have been developed that allow for so-called dark sectors with indirect coupling to the SM. In the simplest of these, a new $U(1)$ gauge field is introduced which, via kinetic mixing with the SM photon, gives rise to a potentially massive "dark photon", denoted in the following as an A'. In these models, the dark photon kinetically mixes with the SM photon through a charged fermion loop, a process which is often simplified to an effective coupling of ϵ. This hypothesized coupling enables the electro-production of dark photons through a bremsstrahlung-like process on a nuclear target, as illustrated in Figure 1.

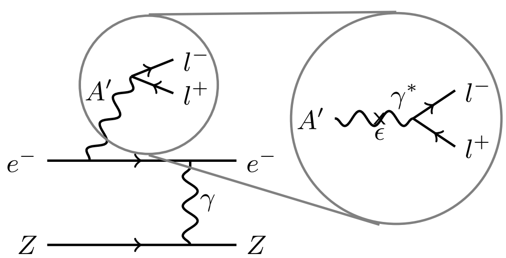

FIG. 1. Electro-production of A' through a bremsstrahlung-like process and subsequent visible leptonic decay. The inset highlights the conversion of dark photons to SM γ through kinetic mixing with strength ϵ.

This paper first introduces the investigated model of a dark sector with strongly interacting massive particles in Section II, highlighting both theoretical and experimental constraints. This is followed by brief descriptions of the setup of the Heavy Photon Search experiment (HPS), Section III, and the data collection and reconstruction, Section IV. Section V and Section VI, make up the main part of the paper, detailing the event selection and data analysis, respectively. Lastly, Section VII summarizes the findings and provides an outlook on future HPS analyses.

II. SIMP MODEL AND PARAMETER CONSTRAINTS

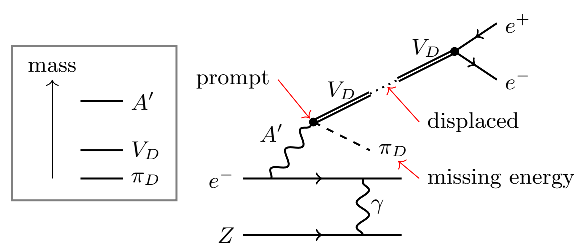

FIG. 2. Production of e^+e^- from the decay of a dark vector meson V_D via a virtual dark photon A'.

In contrast to the minimal dark photon scenario, where thermal freeze-out is achieved through $2 \to 2$ annihilation into SM particles, extended models permit alternative mechanisms. One such extension introduces a QCD-like $SU(3)_D$ gauge symmetry in the hidden sector, yielding strongly interacting massive particles (SIMPs). In this framework, dark pions (π_D) are the lightest states and serve as Dark Matter (DM) candidates.

Dark pion self-interactions allow for a $3\pi_D \to 2\pi_D$ annihilation process that depletes the DM relic density even after decoupling from the SM [1]. The inclusion of dark vector mesons (V_D) further enables a semi-annihilation channel, $\pi_D\pi_D \to \pi_D V_D$, followed by the decay $V_D \to SM$ through a virtual A'. This decay can produce a displaced e^+e^- pair – a signature well matched to the HPS detector's capabilities [2], as illustrated in Figure 2.

The SIMP model considered in this paper involves six key parameters: the dark photon, dark pion, and dark vector masses, $m_{A'}$, m_π, and m_V, respectively; the A' kinetic mixing strength ϵ with the SM photon; the hidden sector $U(1)_D$ gauge coupling constant α_D; and finally, the ratio of the dark pion mass to the dark pion decay constant m_π/f_π.

These parameters are constrained by both theoretical consistency and experimental requirements. Perturbativity demands $\alpha_D < 1$, and we fix $\alpha_D = 10^{-2}$ in this work. This implies $m_\pi/f_\pi \lesssim 4\pi$, since $m_\pi/f_\pi \sim g_D \sim 4\pi\alpha_D$.

The kinetic mixing parameter must fall within $10^{-6} < \epsilon < 10^{-2}$. Values of $\epsilon \gtrsim 10^{-2}$ suppress semi-annihilation, while $\epsilon \lesssim 10^{-6}$ fail to maintain kinetic equilibrium between the dark and visible sectors [2].

To ensure visible decays and a reconstructible signal in HPS, we apply further kinematic constraints:

- $m_{A'} > 2m_\pi$ to suppress $\pi\pi \to A'\pi$
- $m_{A'} > m_\pi + m_V$ to allow $A' \to \pi_D V_D$
- $m_{A'} < 2m_\mu$ and $m_{A'} < 2m_V$ to favor visible A' decays
- $m_V < 2m_\pi$ and $m_V < 2m_\mu$ to prevent $V_D \to \pi_D\pi_D$ and ensure visible decay

All constraints are summarized in Table I. To manage the complexity of the parameter space, we define two benchmark models with fixed mass ratios. The search is then performed as a function of $m_{A'}$ and ϵ, for two representative values of m_π/f_π.

theoretical requirements	**experimental constraints**
$\alpha_D < 1$	$m_{A'} < 2m_\mu$ and $m_{A'} < 2m_V$
$10^{-6} < \epsilon < 10^{-2}$	$m_{A'} > m_V + m_\pi$ and $m_{A'} > 2m_\pi$
$m_\pi/f_\pi < 4\pi$	$m_V < 2m_\pi$ and $m_V < 2m_\mu$

TABLE I. Summary of constraints on visibly decaying SIMP models. Coupling bounds follow from theoretical and cosmological arguments; mass requirements ensure reconstructible final states in HPS.

III. THE HPS EXPERIMENT

Although the HPS detector was originally designed to search for prompt and displaced nominal A's, its layout and capabilities also enable opportunistic sensitivity to SIMPs decays, which can produce similar e^+e^- final states but with different kinematics [3].

HPS uses the electron beam from the Continuous Electron Beam Accelerator Facility (CEBAF) at Jefferson Lab in Newport News, Virginia. CEBAF is an oval-shaped accelerator composed of two superconducting linacs connected by recirculating arcs. Electrons can make multiple passes through the linacs, gaining approximately 2.2 GeV per pass, for up to 5.5 passes, before being delivered to one of four experimental halls. Subharmonics of the 1.497 GHz RF frequency can be extracted simultaneously into different halls, enabling high-rate beam delivery – typically 499 MHz – to multiple experiments at once [4]. CEBAF's ability to provide a high-repetition-rate, multi-GeV electron beam with low per-bunch charge is essential to HPS, allowing for high-luminosity operation with minimal pile-up and manageable detector occupancies.

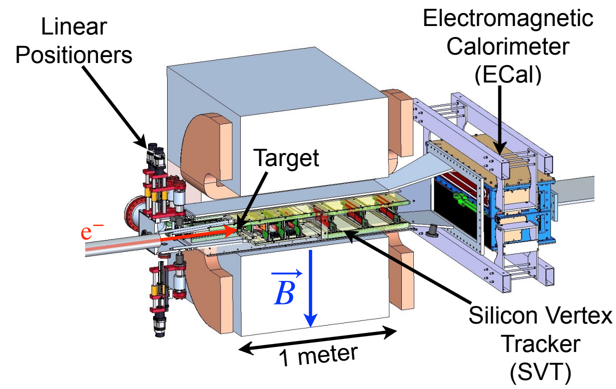

FIG. 3. A cutaway view of the HPS detector showing the SVT in a vacuum chamber inside the bore of the spectrometer magnet and the downstream ECal. The positions of the target and the front portions of the SVT are controlled by a set of linear positioning motors upstream of the detector.

HPS targets rare e^+e^- final states while rejecting large QED backgrounds. This requires precise measurement of invariant mass and decay vertex position. The detector's overall geometry is optimized for boosted, forward-going e^+e^- pairs, a feature shared by many potential signals, including both nominal A' and SIMP decays.

In the nominal A' scenario, the signal (and hence the e^+e^- pair) carries nearly all the beam energy, peaking at $x = E_{A'}/E_{\mathrm{beam}} \to 1$ [5]. In contrast, for the SIMP model the A' decays to dark-sector particles which may then decay to e^+e^-. This results in lower x for the pair and a less boosted decay with wider opening angles [2]. While HPS has limited acceptance for such events, it remains sensitive in regions where the SIMP decay products still fall within the detector's forward coverage.

To capture forward e^+e^- pairs, HPS places a magnetic spectrometer just downstream of the target. This includes a silicon vertex tracker (SVT) embedded in a vertical dipole field (0.24 T) to measure particle momenta and reconstruct displaced vertices. The magnetic field bends charged particles in the horizontal plane, separating signal tracks from beam-related backgrounds.

The SVT is split into upper and lower halves, positioned just above and below the beam plane to maximize acceptance near the beam. However, occupancy from beam electrons scattering in the target – up to $4\,\mathrm{MHz/mm^2}$ – limits how close the detector layers can be placed.

To manage these high rates, HPS takes advantage of the high repetition rate of the CEBAF beam (499 MHz), which spreads interactions over time. A fast e^+e^- trigger system filters for potential signal events, and both the SVT and the lead-tungstate electromagnetic calorimeter (ECal) are capable of selective, time-correlated readout to suppress backgrounds and maintain sensitivity.

The key components of the HPS apparatus are shown in Figure 3. More detailed motivations and detector specification are discussed in [3]. For reference, Table II lists all HPS data-taking runs and their respective energies and luminosities.

Run Period	Beam Energy [GeV]	Integrated Luminosity
2015 Engineering Run	1.04	$1.2\,\mathrm{pb}^{-1}$
2016 Engineering Run	2.40	$10.9\,\mathrm{pb}^{-1}$
2019 Physics Run	4.55	$110\,\mathrm{pb}^{-1}$
2021 Physics Run	3.74	$160\,\mathrm{pb}^{-1}$

TABLE II. Summary of HPS data-taking runs, beam energy, and delivered luminosity.

IV. DATA AND RECONSTRUCTION

The results presented here use data collected during the 2016 engineering run. All data used for analysis were collected at a beam energy of 2.30 GeV with a current of 200 nA on a tungsten foil target 4 µm ($\approx$0.125% X_0) thick. The total luminosity of this dataset is $10\,608\,\mathrm{nb}^{-1}$, comprising 7.2 billion triggered events from a total charge on target of 67.2 mC. In addition to physics runs, a number of special runs were taken, such as field-off runs and runs with a trigger dedicated to collecting scattered single electrons over a wide range of scattering angles. Data from these runs were used to calibrate and align the ECal and SVT.

In addition to experimental data, the analysis presented here makes use of Monte Carlo (MC) simulations to understand some attributes of the signal and background. MadGraph [6] is used to generate signal samples at a range of mass scales, as well as background samples, which include both Bethe-Heitler and radiative tridents (which are kinematically identical to signal) and their interference term, and converted wide-angle Bremsstrahlung (WAB) events. Simulation of Møller scattering events is also used to study the mass resolution. The beam backgrounds, predominantly scattered single electrons, are simulated using EGS5 [7] and overlaid on all MC samples, distributed according to the time structure of the beam. The simulation of generated samples uses GEANT4 [8] to model interactions with the detector, after which the detector response simulation and reconstruction are performed.

A. Event Reconstruction

The event reconstruction follows the procedure detailed in [3]. Briefly, energy deposits in the ECal are grouped into clusters, with per-crystal energy corrections applied using calibration tables.

In the SVT, tracks are reconstructed using a Kalman Filter (KF) for pattern recognition and the General Broken Lines (GBL) algorithm to fit trajectories, incorporating potential small-angle scatters within the detector material. Each track is then matched to an ECal cluster. A matched track-cluster pair is referred to as a reconstructed particle.

Pairs of oppositely charged tracks are combined to form vertex candidates. The vertex position is calculated using a global χ^2 minimization algorithm [9]. Only pairs with tracks in opposite halves of the detector volume are considered. To suppress out-of-time backgrounds, the two associated ECal clusters must be within 2.5 ns of each other, leveraging the 2 ns bunch spacing of the CEBAF beam.

Reconstructed electrons from elastic beam scattering are rejected by requiring the electron candidate momentum to be less than 2.15 GeV. In addition, vertices with total momentum above 2.8 GeV are excluded to remove clearly mis-reconstructed events.

All reconstruction-level requirements are summarized in Table III. Further analysis-level event selections are described in the next section.

Cut Description	Requirement
ECal Clusters in Opposite Volumes	$y_{\mathrm{clu},e^-} \times y_{\mathrm{clu},e^+} < 0$
Track-Cluster Time Difference (Data)	$\lvert t_{\mathrm{trk}} - t_{\mathrm{clu}} - 56\,\mathrm{ns}\rvert < 10\,\mathrm{ns}$
Track-Cluster Time Difference (Sim)	$\lvert t_{\mathrm{trk}} - t_{\mathrm{clu}} - 43\,\mathrm{ns}\rvert < 10\,\mathrm{ns}$
Track-Cluster X Position Difference	$\lvert x_{\mathrm{trk\ at\ Ecal}} - x_{\mathrm{clu}}\rvert < 20.0\,\mathrm{mm}$
Track-Cluster Y Position Difference	$\lvert y_{\mathrm{trk\ at\ Ecal}} - y_{\mathrm{clu}}\rvert < 20.0\,\mathrm{mm}$
Cluster Time Difference	$\Delta(t_{\mathrm{clu},e^-}, t_{\mathrm{clu},e^+}) < 2.5\,\mathrm{ns}$
Beam Electron Cut	$p_{e^-} < 2.15\,\mathrm{GeV}$
Vertex Momentum	$p_{\mathrm{vtx}} < 2.8\,\mathrm{GeV}$

TABLE III. Reconstruction level requirements for vertex candidates. Track-Cluster time difference in simulation and data is corrected using offsets calibrated in [3]. The track positions are found by extrapolating the track from the last layer hit to the face of the ECal.

V. EVENT SELECTION

After the data samples go through reconstruction, further event selection is required to remove background SM processes and isolate potential signal events. This additional event selection was performed in two stages.

A. Preselection

The preselection cuts are designed to remove poorly reconstructed tracks and vertices as well as accidental e^+e^- pairs from the data sample leaving pairs from trident and WAB events. The preselected sample, in MC, is also used to calculated the fraction of radiative events in the data sample as a function of e^+e^- invariant mass. The preselection cuts are summarized in Table IV.

Each reconstructed event is then required to have exactly one of these preselected vertices. This requirement mostly removes events where no quality vertex was able to be reconstructed; however, this selection also prevents side effects of pileup and statistical overlap of the two hit-content categories studied below.

Cut Description	Requirement
Trigger	Pair1
Track Time	$\lvert t_{\mathrm{trk}}\rvert \leq 6\,\mathrm{ns}$
Cluster Time Difference	$\Delta(t_{\mathrm{clu},e^-}, t_{\mathrm{clu},e^+}) \leq 1.45\,\mathrm{ns}$
Track-Cluster Time Difference	$\Delta(t_{\mathrm{trk}}, t_{\mathrm{clu}}) \leq 4.0\,\mathrm{ns}$
Track Quality	$\chi^2_{\mathrm{trk}}/\mathrm{n.d.f.} \leq 20.0$
Beam Electron Cut	$p_{e^-} \leq 1.75\,\mathrm{GeV}$
Minimum Hits on Track	$N_{\mathrm{2d\ hits}} \geq 7$
Unconstrained Vertex Quality	$\chi^2_{\mathrm{vtx}} \leq 20.0$
Vertex Momentum	$p_{e^-+e^+} \leq 2.4\,\mathrm{GeV}$

TABLE IV. Preselection requirements for e^+e^- vertex candidates.

B. Tight Selection

The "tight selection" is the final selection stage. It splits the analysis into mutually exclusive categories based on the track hit content and applies a few additional cuts tuned to perform best within these categories to help eliminate the falsely displaced background.

1. Hit-Content Categories

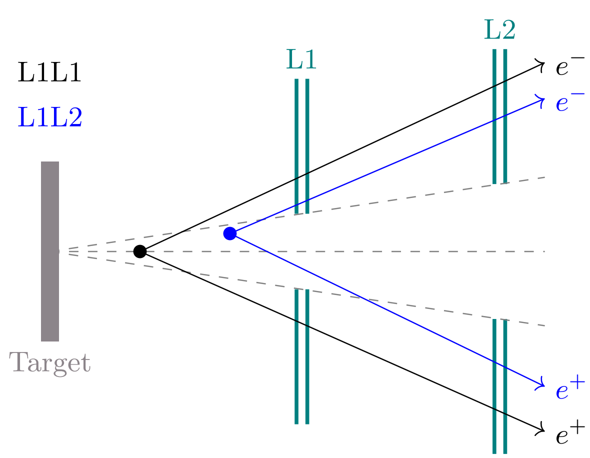

FIG. 4. Diagrams for the two mutually exclusive categories based on the track hit content within a vertex "L1L1" (black) and "L1L2" (blue).

The analysis depends on the track hit content because the mass and vertex resolution, as well as the nature of the falsely displaced background, depend on the SVT layers used to reconstruct the tracks of a vertex candidate.

The first analysis category is called "L1L1", which consists of vertices where both tracks leave hits in both sensors in the first two tracking layers (L1 and L2). These events have the best vertex resolution, though the events in the acceptance are limited to decay lengths much less than the position of L1, as depicted in Figure 4. A hit in L2 is required because the presence of a hit in L2 improves the tracking algorithm's ability to extrapolate the track backward towards L1 and pick up the correct L1 hit.

The second analysis category is called "L1L2" and picks up events where one track misses L1 due to a hit inefficiency or reduced acceptance from longer lifetimes. The track that misses L1 is required to have a hit in L3 again to improve the tracking algorithm's ability to extrapolate the track backward. The L1L2 category introduces more complicated backgrounds, such as an increased rate of WAB conversions coming from the L1 material, and also has a degraded vertex resolution, requiring a slightly different approach to the analysis.

2. Additional Constraints

An observed SIMP signal vertex has lost some energy to the light dark meson π_D carried away by the unobserved light dark meson π_D This shifts our signal region in total momentum from near the beam energy, for the nominal A' analysis, to significantly lower than it; thus, a selection on the sum of the momentum magnitudes is applied.

$$P_{\mathrm{sum}} = |\vec{p}_{e^-}| + |\vec{p}_{e^+}| \tag{1}$$

P_{sum} is chosen to cover the same range for both the L1L1 and the L1L2 categories. Specifically, the signal region (SR) used for the SIMP search requires $1.0\,\mathrm{GeV} < P_{\mathrm{sum}} < 1.9\,\mathrm{GeV}$ and the control region (CR) used for determining the trident differential production rate is $1.9\,\mathrm{GeV} < P_{\mathrm{sum}} < 2.4\,\mathrm{GeV}$.

Since we are searching for the dark vector boson V_D via its 2-body decay into e^+e^-, we expect the invariant mass of the vertex m_{reco} to be within the resolution of the detector σ_m of the mass m_{V_D} we are searching for.

$$p_m = \frac{|m_{\mathrm{reco}} - m_{V_D}|}{\sigma_m} \tag{2}$$

Applying an upper limit on p_m is often refered to as a "mass window" since it results in m_{reco} residing within a small range around m_{V_D}.

C. Displaced Vertex Selection

Real reconstructed vertices of interest should be consistent with originating from the beamspot on the target. This is verified by projecting a vertex candidate back towards the target at z_{target}, using the reconstructed vertex momentum. The target-projected vertex has new coordinates x_{target} and y_{target} can then be used to calculate a significance using the beamspot mean and standard deviations. The shape, size, and position of the beamspot on the target depend on the beam conditions for a given run and a therefore characterized on a run-by-run basis. Analogously, the beamspot is characterized for MC, though without the run dependence. The Vertex Projection Significance (VPS) is then required to be below some threshold in order to keep the vertex candidate.

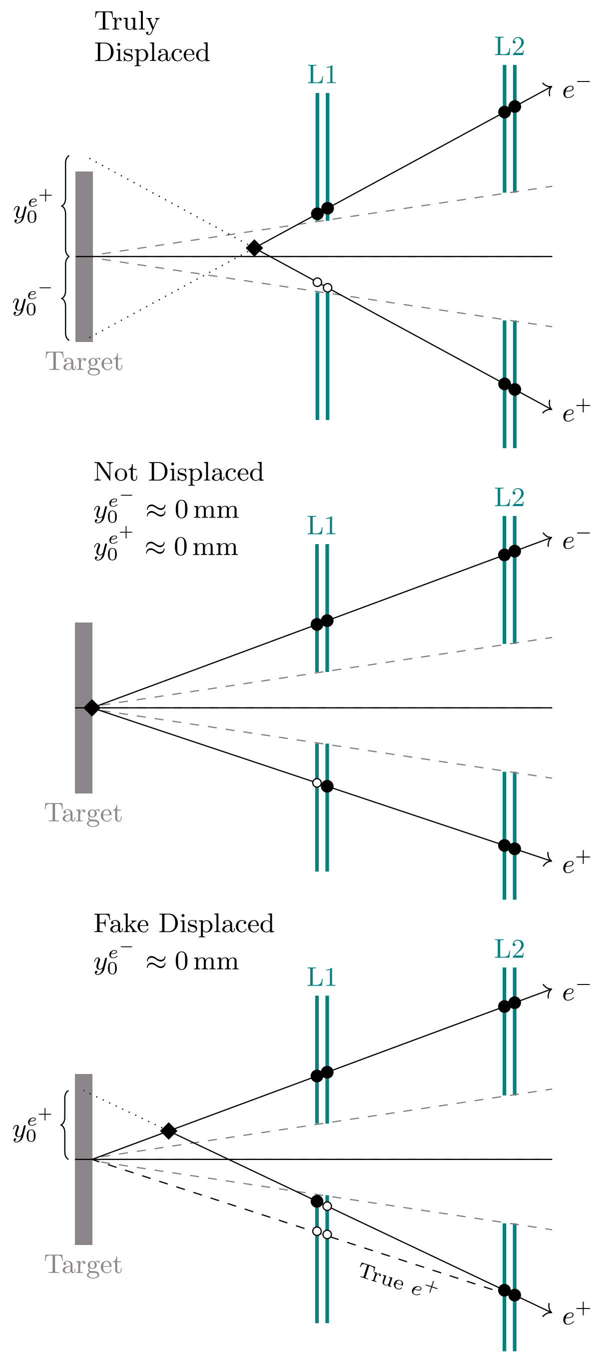

FIG. 5. Illustrations of the vertical track impact parameters y_0 at the target for truly-displaced events (top), not-displaced events (middle), and fake-displaced events (bottom) due to scattering or reconstruction errors.

Since the detector's tracking modules are oriented to be most sensitive in the vertical direction, the vertical impact parameter y_0 has higher precision compared to the horizontal impact parameter. For truly-displaced signal vertices, both tracks creating the vertex would have y_0 far from zero while background vertices would have at least one track with y_0 near zero (undisplaced vertices would have both, but mis-reconstructed fake-displaced vertices could have one far from zero). These scenarios are depicted in Figure 5. This motivates selecting vertices based on requiring the minimum of the two absolute value y_0 to be above a certain threshold.

$$y_{0,\mathrm{min}} = \min(|y_{0,e^-}|, |y_{0,e^+}|) \tag{3}$$

which more sharply distinguishes truly-displaced vertices compared to the vertex z often muddled by fake-displaced vertices where one track is mis-reconstructed at high $|y_0|$.

The uncertainty of the vertical impact parameter σ_{y_0} is a helpful quality parameter measuring how confident the track fit is in the y_0 value. Placing an upper limit on this value for both tracks within a vertex effectively requires both tracks to have good vertical resolution, helping remove some highly-displaced vertices presumably arising from mis-reconstructed tracks.

$$\sigma_{y_0,\mathrm{max}} = \max(\sigma_{y_0,e^-}, \sigma_{y_0,e^+}) \tag{4}$$

Vertex z is left for late-stage statistical analysis of the results and – being highly correlated with $y_{0,\mathrm{min}}$– is redundant with this variable.

1. *Selection Optimization*

The selections for both L1L1 and L1L2 categories were optimized independently on a 10 % subsample of the collected data representing the population of background events and simulated signal samples. First, all of the selections except $y_{0,\mathrm{min}}$ were optimized by keeping the signal efficiency high (at least 80 %) while removing background events with relatively high values for $y_{0,\mathrm{min}}$. While the $\sigma_{y_0,\mathrm{max}}$ parameter was not found to be powerful for the L1L1 category, it was helpful in removing highly-displaced background events within the L1L2 category. Afterward these additional quality selections, the $y_{0,\mathrm{min}}$ parameter was optimized by maximizing the binomial significance of the signal yield above the leftover background. The signal yield was calculated as described in Section VI B 1 and then scaled up by a factor of $0.1/\epsilon$ in order to put it on the same level as the background yield within this subsample. In order to avoid biasing the selection to the specific 10% of the dataset chosen, the selections chosen from this maximization were then fit with a second (first) order polynomial for the L1L1 (L1L2) category.

The final values of these cuts are summarized in Table VI where

$$y_{0,\mathrm{min}}^{\mathrm{cut}}(m) = A + Bm + Cm^2 \tag{5}$$

for L1L1 and

$$y_{0,\min}^{\text{cut}}(m) = \begin{cases} A & m \leq 40\,\text{MeV} \\ B + Cm & 40\,\text{MeV} < m < 120\,\text{MeV} \\ D & m \geq 120\,\text{MeV} \end{cases} \quad (6)$$

for L1L2 and the parameters of these functions are given in Table V.

Figure 6 shows the distributions for the L1L1 and L1L2 hit-content categories after all of these selections on the data sample.

Parameter	L1L1	L1L2
A	$1.176\,\text{mm}$	$1.66\,\text{mm}$
B	$-7.44 \times 10^{-3}\,\text{mm/MeV}$	$1.86\,\text{mm}$
C	$1.59 \times 10^{-5}\,\text{mm/MeV}^2$	$-5.1 \times 10^{-3}\,\text{mm/MeV}$
D	–	$1.25\,\text{mm}$

TABLE V. Parameters for Equation (5) and Equation (6).

Selection	L1L1	L1L2
Missing Energy	$1.0\,\text{GeV} < P_{\text{sum}} < 1.9\,\text{GeV}$	
Mass Resonance	$p_m < 1.5$	
From Beamspot	VPS < 2	VPS < 4
Lower y_0 Error	–	$\sigma_{y_0,\max} < 0.4\,\text{mm}$
Highly Displaced	$y_{0,\min} > y_{0,\min}^{\text{cut}}(m_{\text{reco}})$	

TABLE VI. Summary of the final tight selection depending on hit-content category.

VI. DATA ANALYSIS

This analysis is searching for an excess of highly-displaced vertices reconstructed at a particular mass resonance. Both dimensions of this search are necessary in order to separate the signal process from known SM backgrounds, since trident or WAB processes that occur within the target in combination with reconstruction effects (e.g. hit inefficiencies) are able to mimic signal behavior.

A. Search Procedure

Before applying the final selection on $y_{0,\min}$, we perform an ABCD-like background estimation technique in the $y_{0,\min}$–m_{reco} space and compare this estimate to the observation to check for a signal-like excesses.

We separate our search space into signal regions and side bands in m_{reco} and $y_{0,\min}$. Along the m_{reco} axis, there are two sidebands – one below and one above the signal region – while there is one lower sideband along the $y_{0,\min}$ axis. Table VII gives the definition of these regions and Figure 7 shows an example of these regions along with the calcuation described below for the L1L1 channel.

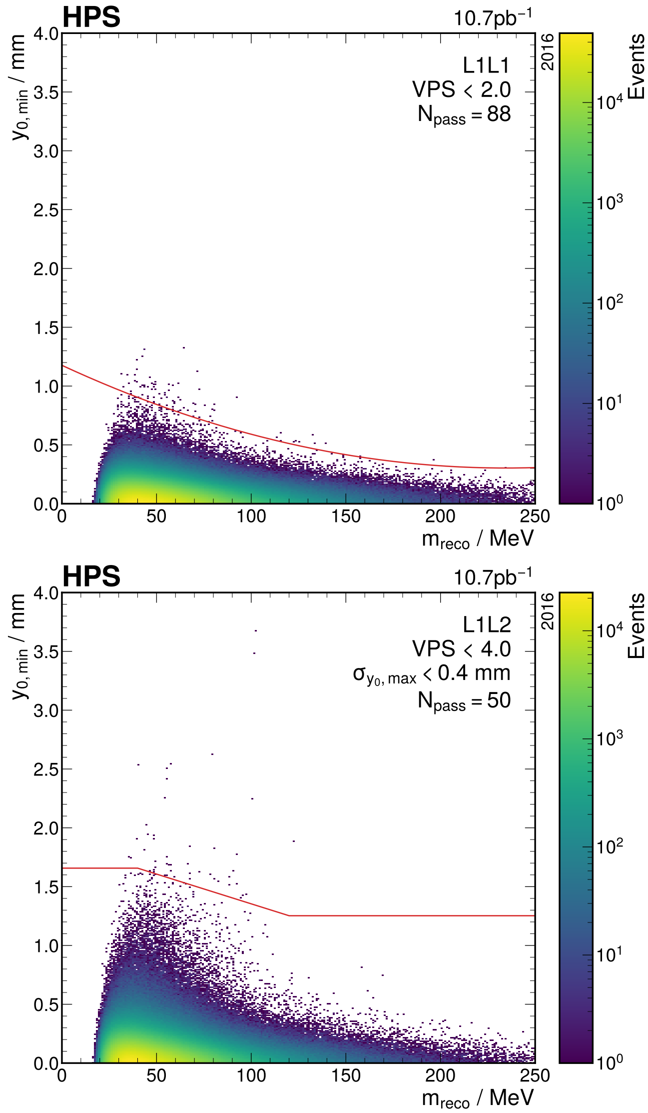

FIG. 6. $y_{0,\min}$ distribution as a function of reconstructed invariant mass m_{reco} with the final selection $y_{0,\min}^{\text{cut}}$ drawn in red for the L1L1 (L1L2) hit-content category in red on top (bottom).

We cast the sidebands into region F to obtain the expected number of events F_{exp} with

$$\text{F}_{\text{exp}} = \text{C} \times \frac{\max(\text{A} + \text{E}, 0.4)}{\text{B} + \text{D}} \quad (7)$$

where A stands for the number of events within region A, B for number of events in region B, etc. The limiting value of 0.4 was chosen to keep this prediction well behaved when the events are sparse in the higher mass region. A Poisson mean of 0.4 is the highest Poisson mean with zero observed counts being the most probable outcome.

The statistical test for excess is performed using 10 000 toy counting experiments. We construct the null distribution by sampling C and B+D from normal distribu-

Region	$m_{\rm reco}$ Range	$y_{0,\rm min}$ Range
A	$(m_{V_D} - 4.5\sigma_m, m_{V_D} - 1.5\sigma_m)$	$(y_{0,\rm min}^{\rm cut}, \infty)$
B	$(m_{V_D} - 4.5\sigma_m, m_{V_D} - 1.5\sigma_m)$	$(y_{0,\rm min}^{\rm floor}, y_{0,\rm min}^{\rm cut})$
C	$(m_{V_D} - 1.5\sigma_m, m_{V_D} + 1.5\sigma_m)$	$(y_{0,\rm min}^{\rm floor}, y_{0,\rm min}^{\rm cut})$
D	$(m_{V_D} + 1.5\sigma_m, m_{V_D} + 4.5\sigma_m)$	$(y_{0,\rm min}^{\rm floor}, y_{0,\rm min}^{\rm cut})$
E	$(m_{V_D} + 1.5\sigma_m, m_{V_D} + 4.5\sigma_m)$	$(y_{0,\rm min}^{\rm cut}, \infty)$
F	$(m_{V_D} - 1.5\sigma_m, m_{V_D} + 1.5\sigma_m)$	$(y_{0,\rm min}^{\rm cut}, \infty)$

TABLE VII. Region definitions for use in background estimation via sidebands. Region F is the signal region in which we are searching for an excess. m_{V_D} is the mass point we are searching for, σ_m is the detector mass resolution evaluated at m_{V_D}, $y_{0,\rm min}^{\rm cut}$ is the optimized cut value evaluated at m_{V_D}, and $y_{0,\rm min}^{\rm floor}$ is the maximum value of $y_{0,\rm min}$ such that region C has at least one thousand events in it.

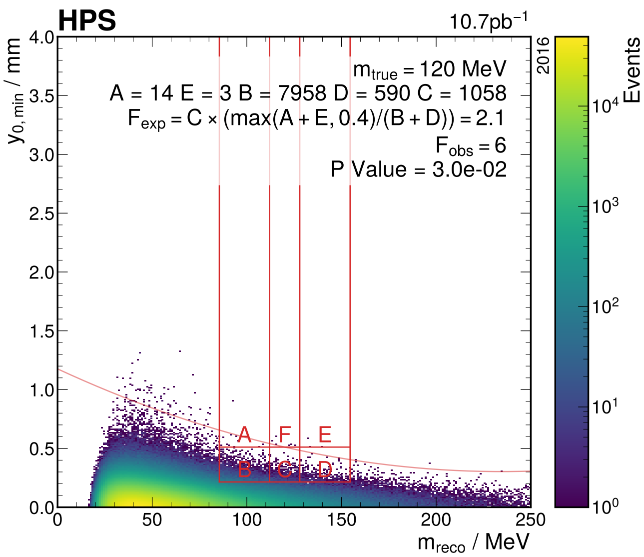

FIG. 7. Example search calculation within the L1L1 channel showing the six regions and how the calculation is performed.

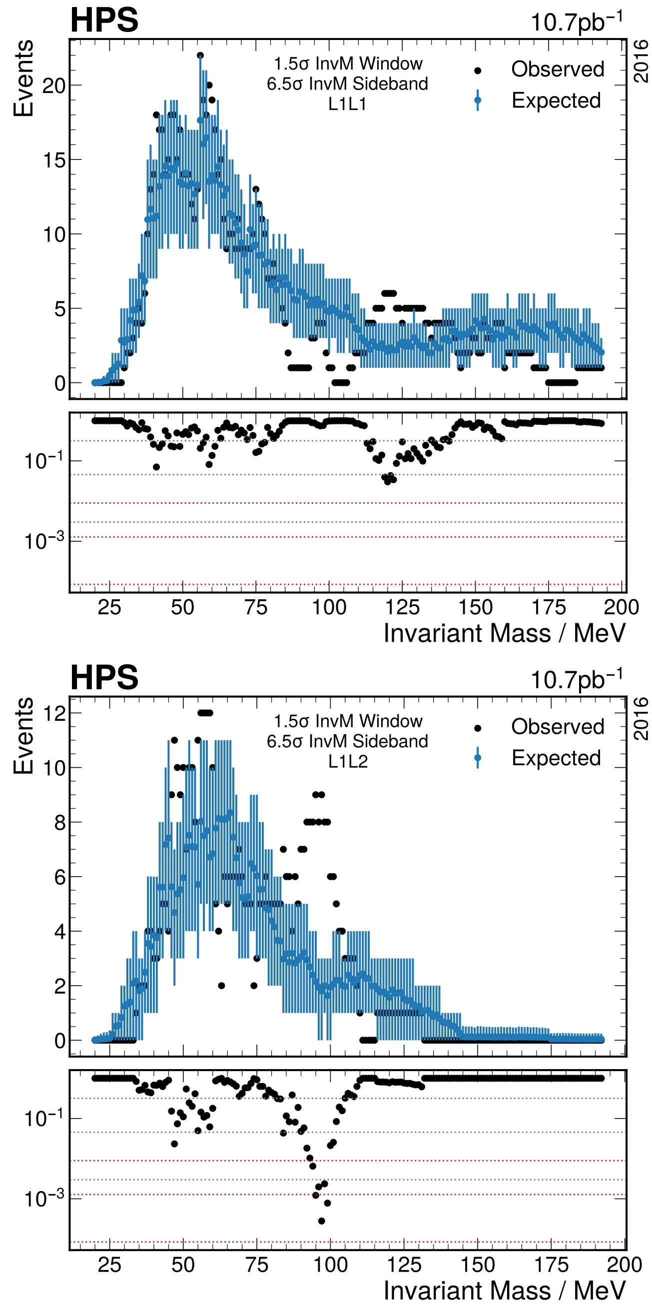

FIG. 8. Search results for the L1L1 (L1L2) hit-content category on top (bottom). The gray (red) dotted lines in the lower panels are 1σ, 2σ, and 3σ local (global) significance lines. The global significance is estimated by dividing the local significance by an approximate number of independent mass bins in which the search was performed.

tions and sampling A+E from a Poisson distribution and then using Equation (7) with these samples to calculate the sampled F. This null distribution is then integrated from the observed number of events in region F up to infinity to obtain an approximate probability that the observed number aligns with the background prediction, which we use as the local p-value.

This procedure is repeated for each mass m_{V_D} in our search range, producing Figure 8 showing the comparison between expected and observed event yields in region F and their corresponding p-values derived from these toy experiments. The lowest observed p-value at $m_{\rm inv} = 97\,\rm MeV$ achieves less than 3σ global significance. Furthermore, the excess only exists within the L1L2 category, leading to the conclusion that it is a normal (although rare) statistical fluctuation.

B. Exclusion Procedure

Without statistically-sound evidence for a SIMP-like signal excess, the question can be inverted instead to what SIMP parameters can be excluded given the lack of excess. This exclusion calculation is done by estimating the sensitivity of this analysis which is defined as the ratio of the expected signal yield to the maximum allowed signal yield. The maximum allowed signal yield is calculated using the Optimum Interval Method (OIM) [10]. We describe the expected signal yield calculation in greater detail below since it is more specific to HPS.

1. Expected Signal Yield

The expected signal yield for a given mass and kinetic mixing strength, $N_{\mathrm{sig}}(m_{A'}, \epsilon)$, is calculated with all other SIMP parameters fixed ($m_{A'} : m_{V_D}$, $m_{A'} : m_{\pi_D}$, α_D, $m_{\pi_D} : f_\pi$). First, the total expected A' production rate in a given dataset, $N_{A'}(m_{A'}, \epsilon)$, is calculated using the simulation-derived terms called radiative fraction and radiative acceptance. The A' can visibly decay into two different neutral dark vectors, ρ_D and ϕ_D. Therefore, the total expected signal can be calculated by measuring the combined acceptance × efficiency for both $\rho_D \to e^+e^-$ and $\phi \to e^+e^-$ using simulated signal. The ρ_D and ϕ_D vectors each have their own production branching ratio and lifetime, which is a function of ϵ, so the acceptance × efficiency for each vector is calculated as a function of ϵ.

The A' production cross-section for dark photons of mass $m_{A'}$ is related to the radiative trident production cross-section by [5].

$$\sigma_{A'} = \frac{3\pi m_{A'}\epsilon^2}{2N_{\mathrm{eff}=1}\alpha}\frac{\mathrm{d}\sigma_{\gamma^*}}{\mathrm{d}m_{l^+l^-}}\Bigg|_{m_{l^+l^-}=m_{A'}} \tag{8}$$

where the differential cross-section is evaluated at the particular mass $m_{A'}$. Multiplying both sides of Equation (8) by the integrated luminosity gives the A' production yield given the differential radiative trident rate,

$$N_{A'}(m_{A'}, \epsilon) = \frac{3\pi m_{A'}\epsilon^2}{2N_{\mathrm{eff}=1}\alpha}\frac{\mathrm{d}N_{\gamma^*}}{\mathrm{d}m_{A'}} \tag{9}$$

The differential radiative trident rate in Equation (9) is broken into three components as

$$\frac{\mathrm{d}N_{\gamma^*}}{\mathrm{d}m_{A'}} = \left(\frac{\mathrm{d}N_{\gamma^*,\mathrm{CR}}}{\mathrm{d}m_{A'}}\Big/\frac{\mathrm{d}N_{\mathrm{CR}}}{\mathrm{d}m_{\mathrm{reco}}}\right)\left(\frac{\mathrm{d}N_{\gamma^*}}{\mathrm{d}m_{A'}}\Big/\frac{\mathrm{d}N_{\gamma^*,\mathrm{CR}}}{\mathrm{d}m_{A'}}\right)\frac{\mathrm{d}N_{\mathrm{CR}}}{\mathrm{d}m_{\mathrm{reco}}} \tag{10}$$

The first term in Equation (10) is the radiative fraction ($f_{\mathrm{rad}}(m_{A'})$), which measures the expected contribution of radiative tridents to the reconstructed and selected background in the CR. The second term in Equation (10) is the inverse of the radiative trident acceptance × efficiency in the CR, referred to as the radiative acceptance ($A_{\mathrm{rad}}(m_{A'})$). The last term in Equation (10) is simply the reconstructed and selected background rate in the CR, and provides a means to scale the production rate to a given dataset, whether in simulation or data.

In this signal hypothesis, we do not observe the dark photons production or decay, instead, the dark photon decays to an unobservable dark pion and the neutral dark vector meson V_D then decays back into a e^+e^- pair. Thus, in order to calculate the expected signal rate, we need to account for these two processes. The first process has a branching ratio $BR(A' \to \pi_D V_D)$ which is complicated by the fact that there are two different dark neutral vector mesons that fit our requirements. The second process is embedded in the decay rate $\Gamma(V_D \to e^+e^-)$. Let $E(z)$ be the signal efficiency of the analysis as a function of the z where the V_D decayed into the e^+e^- pair. Then we can sum over the possible V_D and estimate the fraction of $N_{A'}$ that produce a V_D which decays and passes the analysis requirements.

$$N_{\mathrm{sig}} = N_{A'}\int_{z_{\mathrm{target}}}^{\infty}\sum_{V_D\in\{\rho_D,\phi_D\}} D_{V_D}(z)E(z)dz \tag{11}$$

where

$$D_{V_D}(z) = BR(A' \to \pi_D V_D)\frac{e^{-(z-z_{\mathrm{target}})/(\gamma c\tau_{V_D})}}{\gamma c\tau_{V_D}} \tag{12}$$

The branching ratio $BR(A' \to \pi_D V_D)$ and lifetime τ_{V_D} are taken from [2] where the lifetime explicitly depends on $m_{A'}$ and ϵ^2. The V_D energy (and thus the relativistic γ) used in $D_{V_D}(z)$ is only distributed over a small range (within $\mathcal{O}(100\,\mathrm{MeV})$) so we replace it with the mean $\langle\gamma\rangle$ in order to make the calculation more practical.

2. Systematic Errors

All systematic errors arising from the experiment and this analysis have been quantified individually for the two hit-content categories. The considered systematic errors were found to be within $\sim 1\,\%$ of each other for both hit-content categories and thus the largest error of the two is used for the final exclusion contour in either category and their combination. Table VIII summarizes the systematics considered which are described in this section.

The systematic error of the radiative fraction of $7\,\%$ is estimated from the uncertainty on the total cross sections of the different trident processes. A detailed description of this is given in [3].

Both preselection and final cuts have systematic errors that are found to be negligible. The difference in efficiency between data and simulated background samples is less than a few percent for the selection variables used and the simulated background is lower than the data efficiency; thus, accounting for this difference is neglected. We find that the radiative acceptance is influenced most by the smearing of the preselection cut variables and appears to be underestimated by $\sim 12\,\%$. However, we don't include this systematic as this would artificially improve the sensitivity since the signal yield (and therefore the sensitivity) is inversely proportional to the radiative acceptance.

Systematic	Value
radiative fraction	7 %
preselection cuts	neglected
final selection cuts	neglected
radiative acceptance	
from pre-selection	neglected
from target uncertainty	$\sim 5\,\%$
signal yield	
from target uncertainty	2 %
from mass resolution	0.5 %
beamspot	neglected
P_{sum} shape	$\sim 3\%$
total	$\sim 10\,\%$

TABLE VIII. Summary of systematic errors considered and the values determined. Values marked preceded by $\sim$ are mass-dependent and the maximum value within the most-sensitive mass range is what is listed.

The uncertainty on the target position affects both the radiative acceptance and the signal yield. To determine the resulting systematics, two simulated samples with the target position offset by ± 5 mm were created. This value corresponds to the known uncertainty of the position of the target . From these samples, the radiative acceptance was found to be overestimated by $\sim 5\,\%$ and the signal yield was found to be overestimated by 2 % due to selections on target position-dependent variables.

The width of the beamspot and the mass resolution of the detector are underestimated within the simulation relative to the data. In order to account for this underestimate the resulting analysis variables were smeared accordingly which was found to only have a small impact on the results. Due to a higher efficiency of events passing the VPS cut, the beamspot smearing improves the signal yield, so we choose to neglect it in order to keep this exclusion estimate conservative. The mass smearing, however, was found to decrease the signal yield by 0.5 % which is included in the total systematic uncertainty.

Finally, the shape of the P_{sum} distribution is different between data and simulated background. The effect of this systematic was determined by re-weighting events according to the ratio of the data and simulation P_{sum} distributions and then re-estimating the signal yield with these new weights. This led to a decrease in signal yield of $\sim 3\,\%$ for the most sensitive mass range, rising to $\sim 15\,\%$ in the lower masses.

These systematics were summed in quadrature leading to a total of $< 10\,\%$ for all but the lowest mass points evaluated (rising up to $\sim 18\,\%$).

3. Combined Exclusion Contour

Figure 9 shows the sensitivity for both hit-content categories. The contours are drawn where the sensitivity equals one after being suppressed by potential systematic errors described in the previous section. To improve on our reach, we consolidate the L1L1 and L1L2 results. The combined sensitivity of the two categories is calculated by adding the two expected yields together and estimating the maximum allowed using the "Minimum Limit" combination technique for OIM results [11]. Figure 10 shows the resulting sensitivity along with the combined exclusion contour, including systematic errors. Compared to the individual sensitivities of the two hit-content categories, the combined result continuously covers a broader range in invariant mass and extends to $\epsilon^2 < 10^{-6}$ which neither category reaches in their own.

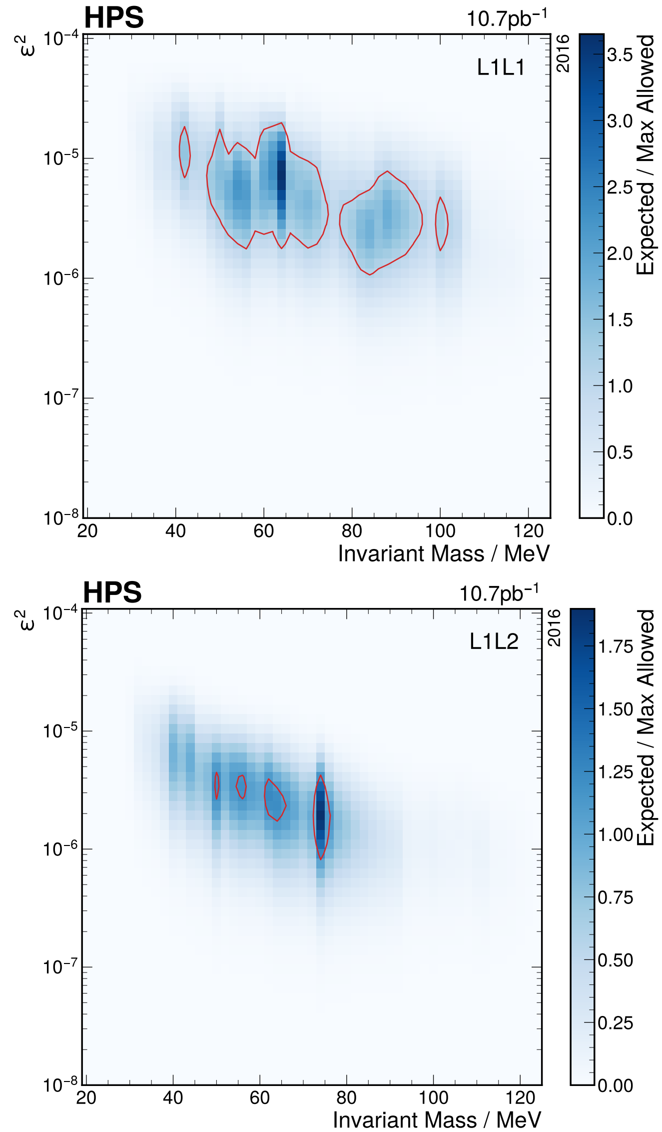

FIG. 9. Sensitivity and resulting exclusion contours for both L1L1 (top) and L1L2 (bottom) hit-content categories.

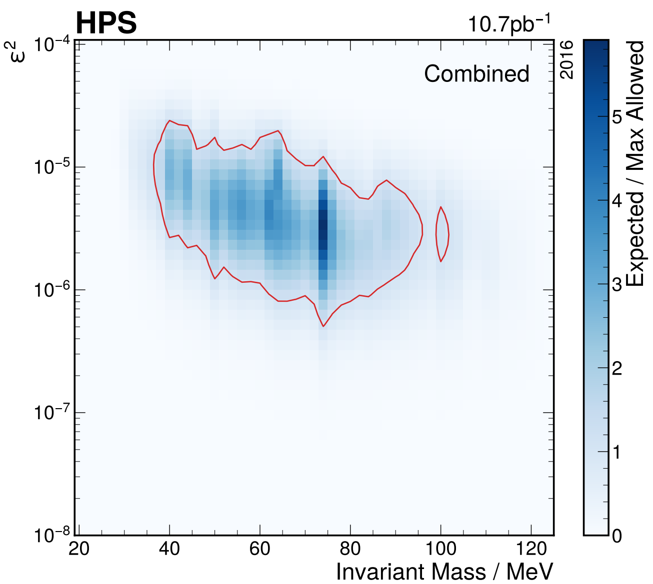

FIG. 10. Sensitivity and resulting exclusion contour for the combination of the two hit-content categories.

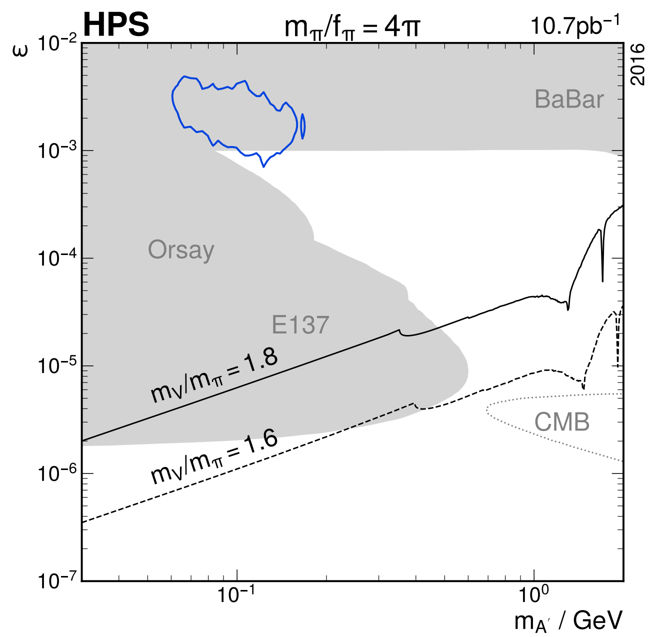

FIG. 11. Exclusion contour from this analysis with comparisons to other experiments (gray) and theoretical predictions for this model (black).

VII. CONCLUSION

In the investigated region of the SIMP parameter space, couplings above $\epsilon^2 = 10^{-6}$ have been excluded by BaBar [12]. Our result contributes to this effort by confirming the BaBar results and probing a small portion of previously unexplored SIMP parameter space. Since this analysis was developed as the first of its kind (displaced-vertex search in the low-P_{sum} region) for HPS, it opens the door to further refinement and investigation with later and larger datasets. For example, a possible extension to the analysis is given by a third hit category "L2L2" where both tracks miss the first tracking layer. This category also suffers from complex backgrounds and significantly reduced vertex resolution, but it does have acceptance to even longer lifetimes where both tracks decay without hitting L1. The L2L2 category is particularly interesting in the context of the SIMP search because there is greater acceptance for longer decay lengths. Future analyses based on the ~ 10 times more luminous 2021 data sample could include this new hit category.

[1] Y. Hochberg, E. Kuflik, T. Volansky, and J. G. Wacker, Mechanism for thermal relic dark matter of strongly interacting massive particles, Phys. Rev. Lett. **113**, 171301 (2014).

[2] A. Berlin, N. Blinov, S. Gori, P. Schuster, and N. Toro, Cosmology and accelerator tests of strongly interacting dark matter, Physical Review D **97**, 10.1103/physrevd.97.055033 (2018).

[3] P. H. Adrian *et al.* (HPS), Searching for prompt and long-lived dark photons in electroproduced electron-positron pairs with the heavy photon search experiment at jlab, Physical Review D **108**, 10.1103/physrevd.108.012015 (2023).

[4] C. Leemann, D. Douglas, and G. Krafft, The Continuous Electron Beam Accelerator Facility: CEBAF at the Jefferson Laboratory, Ann. Rev. Nucl. Part. Sci. **51**, 413 (2001).

[5] J. D. Bjorken, R. Essig, P. Schuster, and N. Toro, New fixed-target experiments to search for dark gauge forces, Phys. Rev. D **80**, 075018 (2009).

[6] J. Alwall *et al.*, The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations, Journal of High Energy Physics **2014**, 79 (2014).

[7] H. Hirayama, Y. Namito, A. F. Bielajew, S. J. Wilderman, and W. R. Nelson, The EGS5 code system, SLAC-R-730, KEK-2005-8, KEK-REPORT-2005-8 (2005).

[8] S. Agostinelli *et al.* (GEANT4), GEANT4–a simulation toolkit, Nucl. Instrum. Meth. A **506**, 250 (2003).

[9] P. Billoir and S. Qian, Fast vertex fitting with a local parametrization of tracks, Nucl. Inst. Methods A **311**, 139 (1992).

[10] S. Yellin, Finding an upper limit in the presence of an unknown background, Phys. Rev. D **66**, 032005 (2002).

[11] S. Yellin, Some ways of combining optimum interval upper limits (2011), arXiv:1105.2928 [physics.data-an].

[12] J. P. Lees *et al.* (BaBar), Search for Invisible Decays of a Dark Photon Produced in e+e- Collisions at BaBar, Physical Review Letters **119**, 10.1103/physrevlett.119.131804 (2017).