

SRC-CT Analysis Note for Axion-Like Particle Search

J.R. Pybus¹

¹MIT, Cambridge, MA

December 19, 2023

Abstract

This analysis note describes the search for photon-coupling Axion-Like Particles (ALPs) in the mass region of 200 to 450 MeV using Primakoff production in the Hall D SRC-CT photonuclear data from Carbon-12. The purpose of this note is to receive group approval for unblinding our data, which has thus far been limited to 10% of the total. The analysis procedure and statistical calculation will be described.

Contents

1	Introduction	3
2	Event Selection	5
2.1	Reaction and Measurement	5
2.2	Trigger	6
2.3	Skimming	6
2.4	Post Selection	11
2.5	Veto and Cut Optimization	13
2.6	Good Run Selection	15
3	Monte Carlo	18
3.1	Generator	18
3.2	GEANT	19
3.3	Random Trigger Event Mixing	19
3.4	Efficiency and Resolution Determination	19
4	Flux/Luminosity	22
5	Results	24
5.1	Blinding Procedures	24
5.2	Statistical Analysis	24
5.3	Tests of Discovery	26

5.3.1	Look-Elsewhere Effect	26
5.3.2	Signal Injection Tests	28
5.3.3	η Discovery Test	28
5.3.4	Polynomial Order Examination	29
5.4	Tests of Exclusion	30
5.4.1	Normalization	30
5.4.2	Expected Exclusion Range	33
6	Reference Channel Normalization Check	37
A	Comments from SRC-CT Group	41

1 Introduction

This note describes a search for Axion-Like Particles (ALPs) using the Carbon dataset of the Hall D SRC-CT experiment. The method used here was explored in Ref. [1], and similarly searches for QCD-scale ALPs in the mass region of $200 < m_{ALP} < 450$ MeV with dominant coupling to photons. Fig. 1 shows the production mechanism for these ALP candidates, which is Primakoff production via t -channel photons. Fig. 2 shows the existing limits for this model of ALP candidates. A similar method may be obtained by performing a bump hunt on the invariant mass spectrum for 2-photon final states in our dataset.




Figure 1: Diagram of the reaction of interest in this study. The incoming beam photon interacts with the nucleus coherently and produces two final-state photons through the mediation of an intermediate particle.

In this note we describe the process used to obtain similar coupling limits using this dataset, which has thus far been blinded to 10% of the full data. We describe the event selection criteria used for obtaining the final invariant mass spectrum, the simulation methods used to estimate ALP signal properties, and the statistical methods used to determine signal discovery and exclusion. Finally, we show the prediction for the coupling limits the full dataset will afford us.

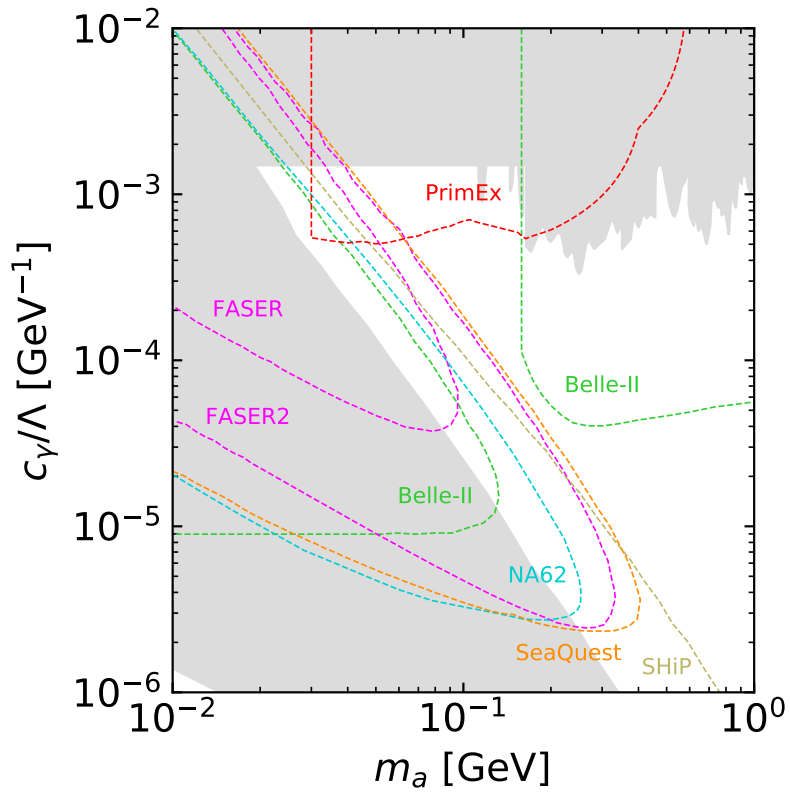


Figure 2: Existing limits on ALP coupling as a function of mass (grey shaded region) [1, 2, 3, 4, 5, 6], as well as predicted limits for other experiments (dashed lines). [1, 7, 8, 9, 10]

2 Event Selection

2.1 Reaction and Measurement

The interaction we intend to study is the Primakoff production of scalar or pseudoscalar particles with subsequent decay into a pair of photons. We describe the sequence of the reaction and its measurement in the following paragraphs.

The beam photon is produced via coherent bremsstrahlung of the 11 GeV CEBAF beam electron from the diamond radiator in the Hall D tagger hall. The scattered electron from the bremsstrahlung event can be measured in the Tagger Hodoscope (TAGH) or the Tagger Microscope (TAGM), which use a magnetic field and measurements of the electron deflection to determine the energy of the scattered electron. This allows inference of the energy and 4-momentum p_{beam} of the corresponding beam photon incident on the target. Details of the Hall D tagger setup may be found in Ref. [11].

In the reaction of interest, the beam photon interacts with the nucleus of 4-momentum p_A coherently, producing two final-state photons via the intermediate process which may only be inferred. The final-state photons produce showers in the Forward and Barrel Calorimeters (primarily the forward calorimeter), denoted FCAL and BCAL respectively. The calorimeters allow us to measure the energy of the photons via their energy deposition in the showers. Because the final state is measured only by two showers, no vertex information is measured, and must therefore be assumed to be at the center of the ^{12}C target at $(0, 0, 65)$ cm. This assumed vertex, together with the shower energies and positions in the calorimeters, allow for the reconstruction of the final-state photon momenta $p_{\gamma 1}$ and $p_{\gamma 2}$. Fig. 3 shows an example of the event display for a simulated $\gamma^{12}\text{C} \rightarrow X^{12}\text{C} \rightarrow \gamma\gamma^{12}\text{C}$ event, with $m_X = 300$ MeV. Fig. 6 shows the distribution of shower position in the FCAL.

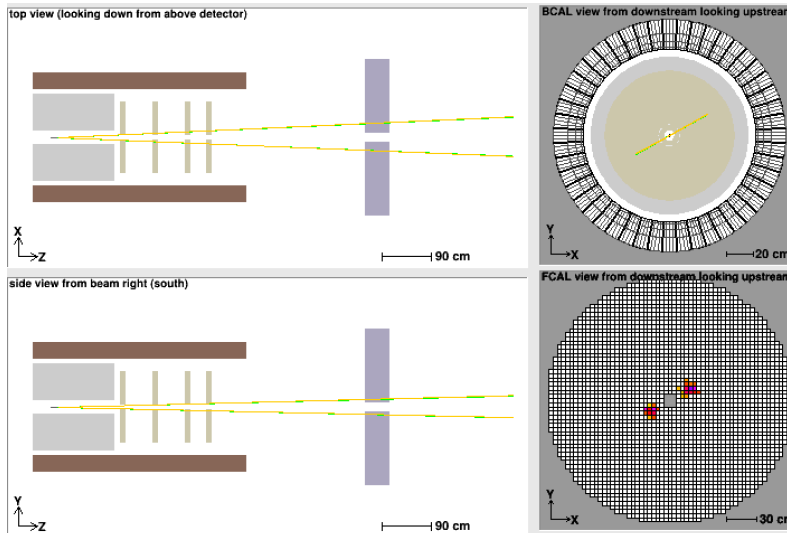


Figure 3: Event display for simulated event with $m_X = 300$ MeV. Larger-mass events result in increased opening angle.

From this measured two-photon final state, the relevant kinematic variables of the reaction can be inferred assuming that Fig. 1 accurately describes the reaction process. The

4-momentum of the intermediate particle, denoted X , can be calculated as $p_X \equiv p_{\gamma 1} + p_{\gamma 2}$, with an invariant mass calculated $m_X^2 = p_X^2 = 2E_1 E_2 (1 - \cos \theta_{12})$. The 4-momentum transfer to the nucleus from the beam photon may be calculated $q \equiv p_{beam} - p_X$. The squared 4-momentum transfer $t = q^2$ is also an important variable for separating Primakoff production mechanisms from competing reactions, and is always negative.

As the opening angle of the photon pair correlates strongly with its invariant mass, pairs with low mass are detected almost exclusively in the FCAL; for the mass region examined in this study ($m_X < m_\eta$), only showers in the FCAL were considered.

2.2 Trigger

Collection of data for a given event during the experiment was contingent on satisfying one of several trigger conditions:

- **Start Counter Trigger:** A trigger combining shower energy deposition and Start Counter hits was used to allow detection of minimally-ionizing final states. This trigger required satisfaction of the criteria $4E_{FCAL} + 5E_{BCAL} > 1$ GeV and at least one hit in the Start Counter.
- **Shower-Only Trigger:** An additional trigger relying only on shower energy deposition was also used, with the requirement $E_{FCAL} + 3E_{BCAL} > 4.2$ GeV.
- **PS Trigger:** A trigger was also used to select events with a hit in the Pair Spectrometer; these events were used to measure the photon flux and luminosity.
- **Random Trigger:** In addition to the above physics triggers depending on detector information, each run included a random trigger at a rate of 100 Hz. This provides a random sample of background during the run which helps in assessing pileup, as described in subsection 3.3.

Events in this search were required to satisfy the shower-energy-only trigger, as the energy deposition of such high-energy photons is large and hits in the Start Counter are indicative of charged-particle background.

2.3 Skimming

Event selection was performed by running a custom plugin over the cooked REST files using the Hall D reconstruction software. Initial event selection criteria required an event with exactly two neutral showers (not associated with any charged tracks) satisfying the following three criteria.

- First, the showers require a “propagated timing” within 3 nanoseconds of the RF time. The propagated timing is given by the time of the shower in the calorimeter minus the propagation time, which is calculated assuming propagation at a speed of light, c , from the center of the target. Figure 4 shows the distribution of propagated timing relative to the RF timing in all measured neutral showers. This helps to remove background resulting from slow-moving neutral particles.

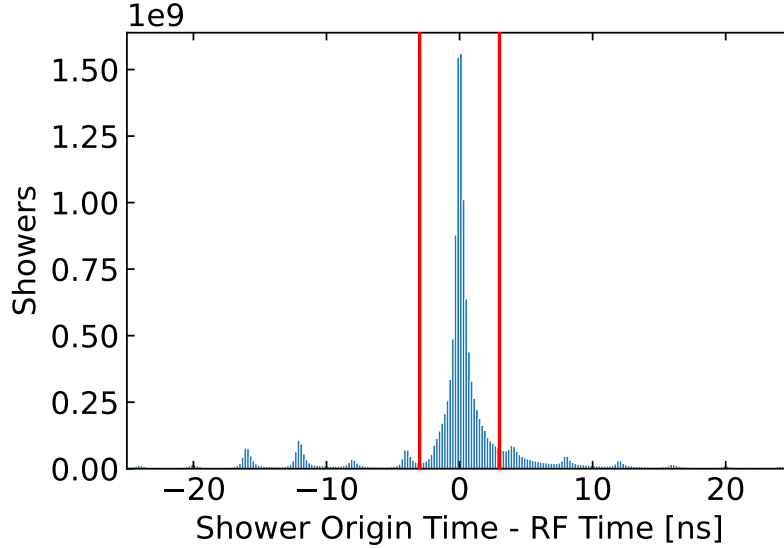


Figure 4: The timing of all neutral showers in a subset of data relative to the RF time, after taking into account propagation time d/c . Valid showers require timing coincidence within 3 ns.

- Second, the showers are required to have a measured energy greater than 0.1 GeV, consistent with the lower energy limits of the FCAL as specified in Rec. [11] and used in Ref. [12]. Figure 5 shows the distribution of measured energy in all neutral showers, as well as showers in selected events. The exact placement of this cut can be seen to have little impact on the selected event sample.




Figure 5: *Left:* The measured energy of all neutral showers in a subset of data. Valid showers require energy greater than 0.1 GeV. *Right:* The measured energy of photon showers in events following all selection cuts. The location of the lower energy cut may be seen to have little effect on the final events after all selection cuts have been applied.

- Third, the center of the shower was required to be located outside the innermost layer of the FCAL to ensure accurately reconstructed energies, as showers near the beamline may leak energy out the detector and therefore have misreconstructed energy. This excludes the layer of scintillators in the FCAL located closest to the beamline, which also record high rates of background. Figure 6 shows the distribution of FCAL neutral shower hit positions, with the location of the inner layer outlined in red.

- Finally, the center of the shower was required to be within 105.5 cm of the center of the beamline. This serves to fiducialize the outer edge of the FCAL, as showers near the edge of the detector may also be subject to misreconstructed energy due to leakage out of the detector.

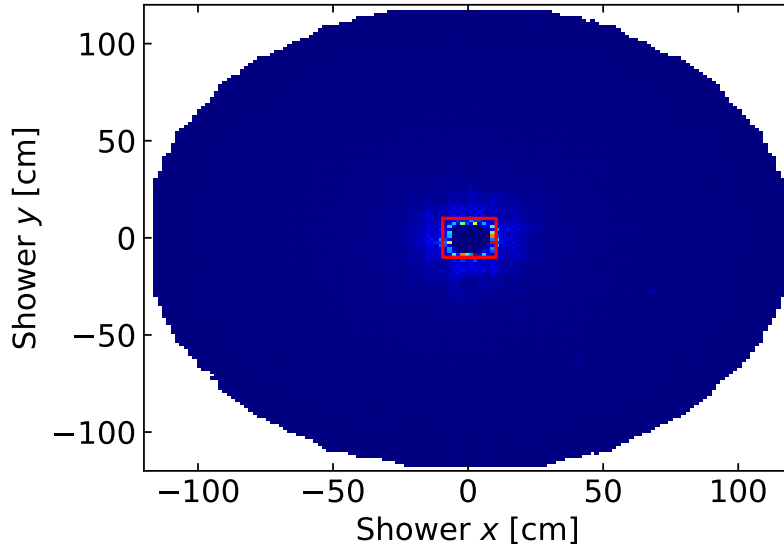


Figure 6: The measured hit position of all neutral showers in the FCAL in a subset of data. The red line outlines the innermost layer of the FCAL. Valid showers require a hit position outside the innermost layer.

These latter three criteria are standard preselection cuts used in Hall D for the identification of neutral showers.

Neutral showers satisfying each of these four criteria are labelled “valid” showers. Additional showers that did not satisfy these criteria were initially ignored in the event selection (though are included in later veto cuts). Figure 7 shows the total number of neutral showers in a given event, as well as the total number of valid neutral showers per event, where invalid showers are initially ignored pending event vetoing.

Following these initial cuts, several subsequent vetoes were applied within the selection plugin to remove background processes using additional detector information. Details of veto/cut optimization are given in subsection 2.5. Signal event generation and simulation is detailed in subsection 3. Comparisons with simulation will be shown in the following description of cuts in order to demonstrate their efficacy. Background processes include charged-particle production and off-vertex production, the former of which is largely removed through the described vetos.

- Events with a hit in the FTOF within 6.5 ns of the RF time (after subtracting time of propagation to the TOF hit) and within 6 cm of one neutral shower were vetoed in order to remove charged-particle-induced showers. Time of propagation was again calculated assuming speed c from the center of the target. Figure 8 shows the distribution of the nearest time-correlated (within 6.5 ns) TOF hits for showers in our candidate

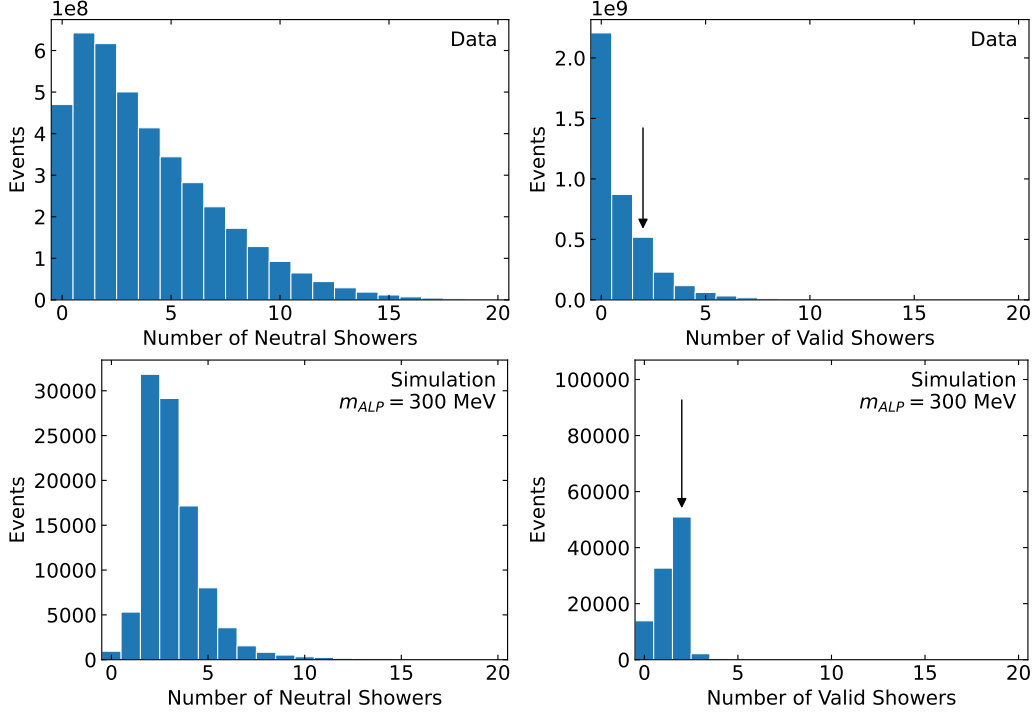


Figure 7: *Top left:* The total number of neutral showers per event in a subset of data. *Top right:* The number of neutral showers satisfying the listed requirements per event in a subset of data. Events required exactly two such showers, with extra showers initially ignored. *Bottom:* Same as top figures in simulated signal with $m_{\chi} = 300$ MeV. Shower validity conditions may be seen to be generally applicable for signal events.

events, in both data and simulation. Signal simulation still showed correlated TOF hits despite the absence of charged particles, likely indicating conversion of the photons in the detector. The value of this cut was optimized using simulation as described in subsection 2.5.




Figure 8: The location of the nearest time-correlated TOF hits to a neutral shower in data (*left*) and simulation (*right*). The location of the veto is shown in the red circle; all events with a TOF hit closer to a candidate shower were discarded.

- Events with an additional FCAL shower within 4 ns of the RF time (after subtracting

for propagation) were vetoed. Figure 10 shows the distribution of extra FCAL showers in time relative to RF time for our candidate events, in both data and simulation. Signal simulation shows a clear peak at 0 ns, likely from mis-reconstructing a shower as multiple showers. Data shows a broader distribution in time, likely resulting from extra photons, neutrons, or charged particles.




Figure 9: The timing of extra FCAL showers with respect to the RF time in data (*left*) and simulation (*right*), after subtracting the propagation time d/c . The location of the vetoes are shown in the red line; events with an FCAL shower in this region were discarded.

- Events with an additional BCAL shower within 6 ns of the RF time (after subtracting for propagation) were vetoed. Figure 10 shows the distribution of extra BCAL showers in time relative to RF time for our candidate events, in both data and simulation. Signal simulation shows no correlation with the RF time, while data shows a clear 0 ns peak from background processes.




Figure 10: The timing of extra BCAL showers with respect to the RF time in data (*left*) and simulation (*right*), after subtracting the propagation time d/c . The location of the vetoes are shown in the red line; events with a BCAL shower in this region were discarded.

For each remaining event a number of candidate tagged beam photons were selected. As the beam photon candidates are independent of the initial event selection this may be done in parallel.

Fig. 11 shows the spectrum of tagged photon energies for all events in run 90271. As the Pair Spectrometer allows flux measurement only for beam photons above 6 GeV, photons below this energy were discarded in order to allow comparison to simulation using the measured photon flux. The structure of the photon tagger may be observed in the energy distribution; efficiency is higher in the tagging microscope for photons with energies between 7.2 and 8.3 GeV, and the hodoscope has only partial coverage at lower photon energies, resulting in gaps in the measured distribution.

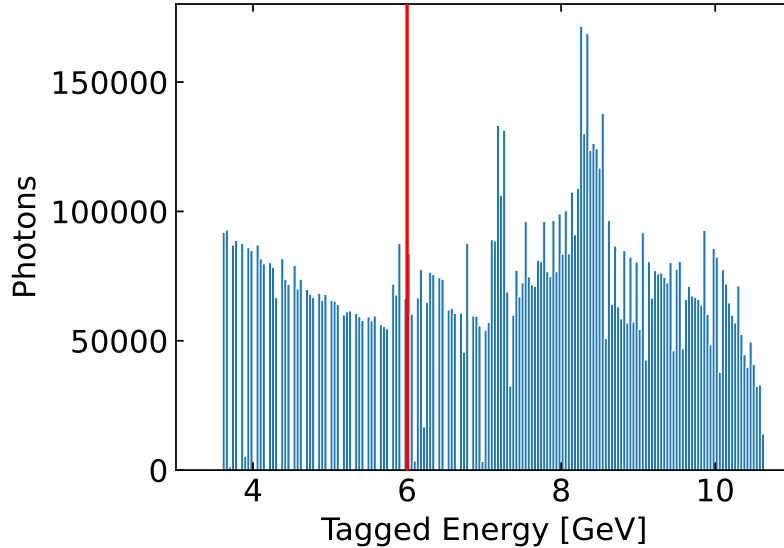


Figure 11: The spectrum of tagged beam photon energies in (subset of) data. Photons below 6 GeV were discarded, as flux can only be measured for higher-energy photons.

2.4 Post Selection

As tagger pileup in GlueX is very common, in many events multiple beam photons fall within the $|\Delta t| < 2.004$ ns timing window of a single bunch with respect to the beam time. In many analyses it is convenient to account for the systematic impact of these “accidental” beam photons via subtraction. However, in the case of this analysis the impact of accidental photons is not worth subtracting. Accidental subtraction mitigates the impact of accidental background on the shape and scale of measured distributions, but the distributions observed in this analysis are already background-dominated; the primary analysis challenge is to distinguish any signal peak from the background, for which the uncertainties in the spectrum are more important than the absolute scale. This subtraction actually has the impact of *increasing* statistical uncertainties (and introducing ambiguity in the distribution of counts in each bin) while reducing the background impact on the shape. It was therefore determined that accidental subtraction would provide certain drawbacks without the usual benefits of such a procedure, and this was therefore not used. Instead, cuts on beam-dependent quantities such as elasticity $\frac{E_X}{E_{beam}}$ are considered satisfied if *at least one* on-time beam photon satisfied such a cut.

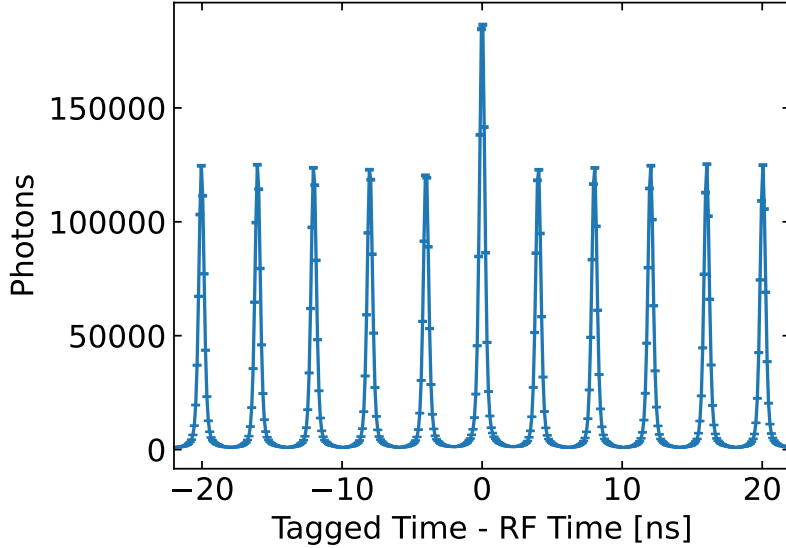


Figure 12: Beam photon timing with respect to the RF time.

A cut was placed on the “elasticity” of the event, which is defined as the fraction of the beam photon energy carried by the final-state photon pair. Simulations of the detector response indicated a elasticity cut of $0.95 < \frac{E_X}{E_{beam}} < 1.05$ would account for the effects of resolution on this measured quantity.




Figure 13: Accidental-subtracted elasticity distribution in data (*left*) and simulation with $m_a = 300$ MeV (*right*).

In order to limit the events to the region dominated by Primakoff production mechanisms, a cut was placed on the scattering angle θ_X of the sum of the two final-state photon momenta $p_X = p_{\gamma 1} + p_{\gamma 2}$. This cut was chosen to be $\theta_X < 0.5^\circ$ in comparison with Monte-Carlo predictions; see Fig. 14. This cut is also consistent with that proposed in Ref. [1].

Fig. 15 shows the effect of the various vetos and cuts on the invariant mass spectrum. The selection may be seen to dramatically reduce the number of events, particularly at low masses, and to increase the proportional number of $(\eta \rightarrow \gamma\gamma)$ and $(\omega \rightarrow \pi^0\gamma)$ events.




Figure 14: Accidental-subtracted diphoton angular distribution in data (*left*) and simulation with $m_a = 300$ MeV (*right*).

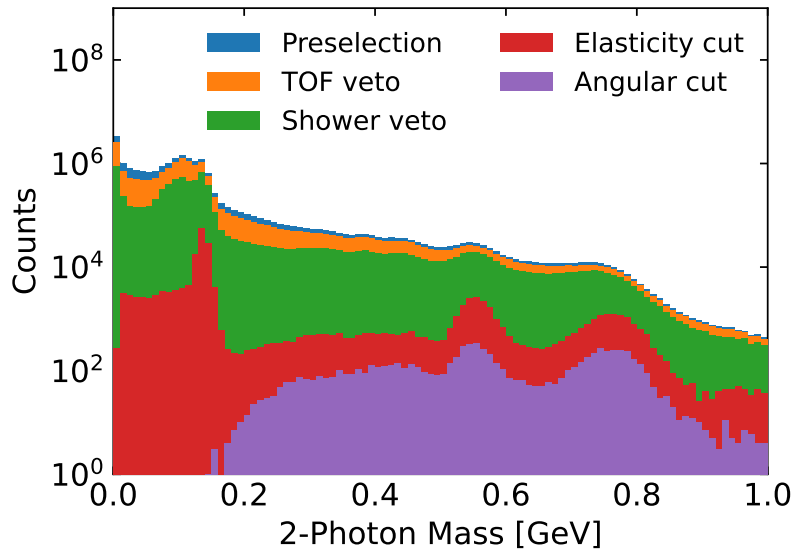


Figure 15: The invariant 2-photon mass spectrum at each level of cut and veto, with cuts applied sequentially.

2.5 Veto and Cut Optimization

We require some optimization of the event selection cuts and vetoes in order to maximize our sensitivity to signal. This is done by comparing the impact of these on the data as well as on simulated signal at a given ALP mass. In order to properly assess the impact of vetoes on signal, simulation of detector response to signal events (generated as described in Section 3.1) was superimposed with “random trigger” data. This allows us to simulate the impact of pileup on signal when applying vetoes.

We define the “signal significance” as our figure-of-merit for optimizing our event selection. A standard figure-of-merit for optimizing signal S with respect to background B is the product of the efficiency and the purity of the cut:

$$F \propto \epsilon \times p \quad (1)$$

The efficiency of the cut is proportional to the number of signal events in simulation which

successfully pass the cut, K :

$$\epsilon \propto S \propto K \quad (2)$$

The purity of the cut cannot be known in absolute terms without knowledge of the normalization of the signal, which is itself an unconstrained parameter. However, we find that the purity should in all cases be *proportional* to the ratio of simulation to number of data events passing the cut, N :

$$p = \frac{S}{S+B} \propto \frac{K}{N} \quad (3)$$

We may therefore take a figure-of-merit which is proportional to our desired quantity and may therefore be optimized in its place:

$$F = \frac{K^2}{N} \quad (4)$$

The absolute scale of this quantity is arbitrary but its maximum location will nonetheless allow optimization of our cuts.

For the background vetoes and physics cuts applied, this figure-of-merit was calculated for various values of the veto/cut parameters in order to optimize the signal significance. As this figure-of-merit differs for different signal hypotheses, different values of m_a were tested to verify that the cuts were optimally selected. To ensure the signal significance was calculated over the relevant region for a given mass hypothesis, a cut on the diphoton mass $0.75 < \frac{m_{\gamma\gamma}}{m_a} < 1.25$ was applied for each calculation of the figure-of-merit. Finally, the figure-of-merit was normalized to a maximum of 1 in each case to ensure the results could be easily interpreted.

This procedure was performed iteratively. The veto parameters were first sequentially optimized for the TOF veto, and the extra shower vetoes. The physics cuts were then optimized on the data following the vetoes, in the order of the elasticity and then the θ_X cut. Then the vetoes were removed, and again sequentially optimized on the data *following* the selected physics cuts. This process was iterated until a stable set of vetoes and cuts were determined:

- *TOF* hit veto radius of 6 cm
- Extra FCAL shower veto window of 4 ns
- Extra BCAL shower veto window of 6 ns
- Elasticity cut of $|1 - \frac{E_X}{E_{beam}}| < 0.05$
- Diphoton deflection angle cut of $\theta_X < 0.5^\circ$

Figures 16, 17 and 18 show the figure-of-merit curves for the *TOF* hit veto radius, the extra FCAL shower timing window, and the extra BCAL shower veto timing window, respectively, given data with the elasticity cut $0.95 < \frac{E_X}{E_{beam}} < 1.05$ and the deflection angle cut $\theta_X < 0.5^\circ$ already applied.

Figures 19 and 20 show the figure-of-merit curves for the elasticity cut and deflection angle θ_X cut, respectively, following the application of the background vetoes. In the case of

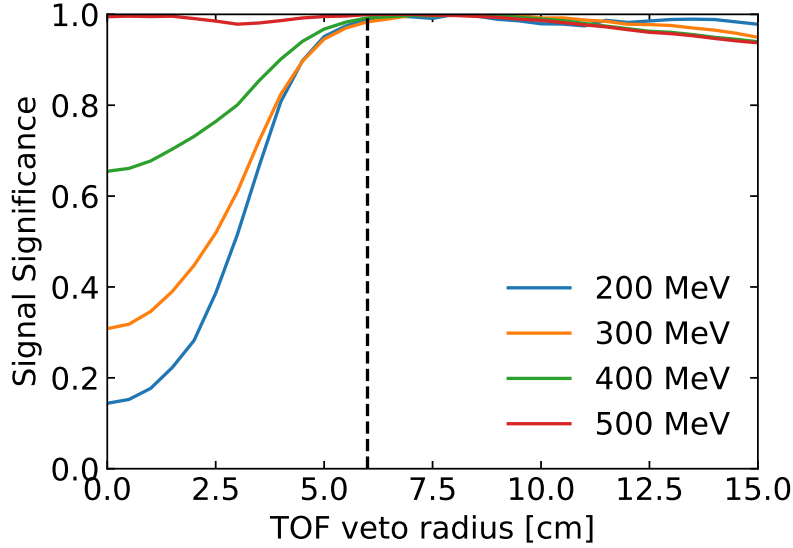


Figure 16: Figure-of-merit for the TOF veto hit radius after application of the physics cuts.




Figure 17: Figure-of-merit for the extra FCAL shower veto timing window after application of the physics cuts and TOF veto.

the θ_X cut the figure-of-merit different significantly between different m_a values in the signal simulation; a cut of $\theta_X < 0.5^\circ$ was chosen to compromise between the different regions of parameter space.

2.6 Good Run Selection

A total number of 662 runs were collected during the experiment. Their run conditions are stored in the run condition database (RCDB) of the Hall D software package. Production runs have to satisfy a set of requirements of RCDB parameters. Their calibration and data

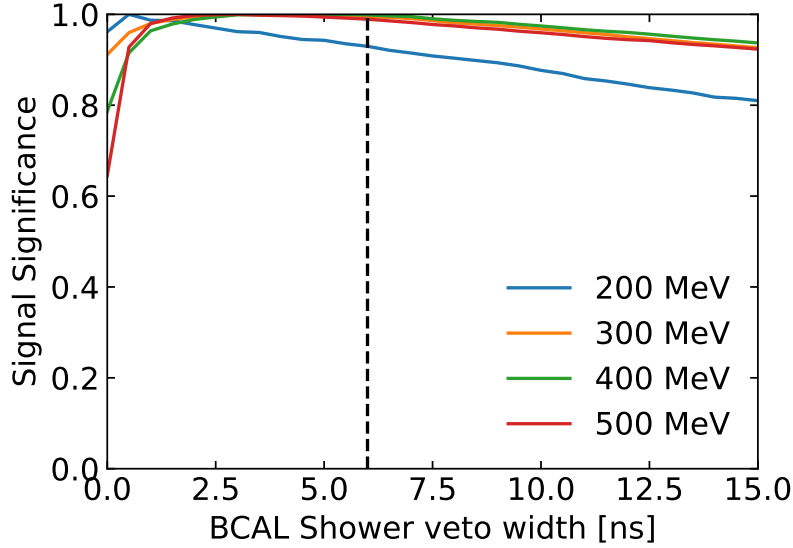


Figure 18: Figure-of-merit for the extra BCAL shower veto timing window after application of the physics cuts and TOF veto. The cut value was chosen to be somewhat conservatively to reduce non-Primakoff η yield.

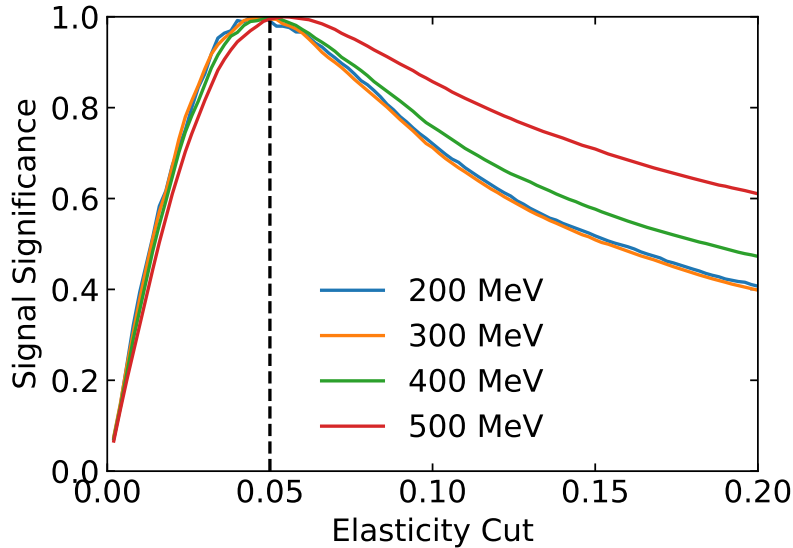


Figure 19: Figure-of-merit for the elasticity cut after application of the vetoes.

quality are also checked to be approved for analysis. The other runs are marked as junk and discarded.

After selection, 480 runs are kept (93.43B triggers in 55988 REST files). They include 78 runs on deuterium (16.44B triggers in 9890 REST files), 164 runs on helium-4 (29.78 B triggers in 17680 REST files), 232 runs on carbon-12 (47.01B triggers in 28313 REST files) and 6 runs on empty target (0.21B triggers in 105 REST files).

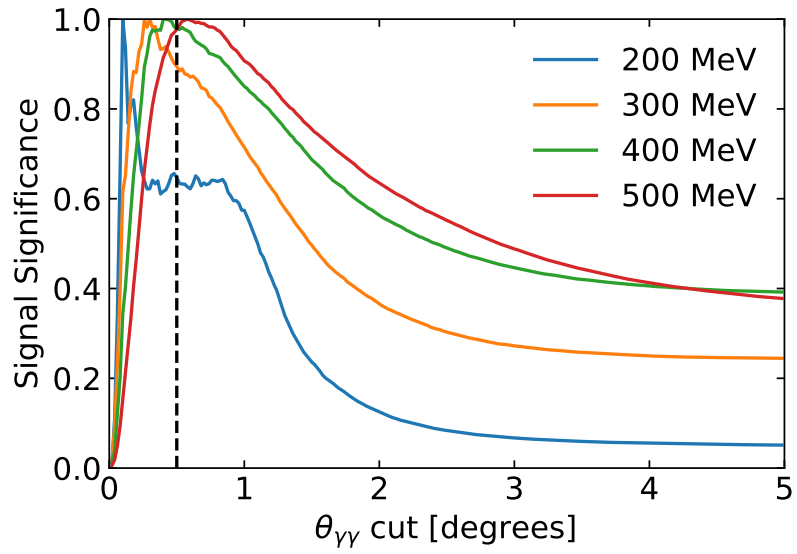


Figure 20: Figure-of-merit for the deflection angle cut after application of the vetoes and elasticity cut.

3 Monte Carlo

3.1 Generator

The cross section for this Primakoff process was calculated as described in [1]:

$$\frac{d\sigma_{\gamma A \rightarrow aA}}{dt} = \alpha Z^2 F_A^2(t) \Gamma_{a \rightarrow \gamma\gamma} \mathcal{H}(m_A, m_a, s, t) \quad (5)$$

where

$$\mathcal{H}(m_A, m_a, s, t) \equiv 128\pi \frac{m_A^4 m_a^2 t (m_A^2 + s) - m_a^4 m_A^2 - t \left((s - m_A^2)^2 + st \right)}{m_a^3 t^2 (s - m_A^2)^2 (t - 4m_A^2)^2} \quad (6)$$

and

$$\Gamma_{a \rightarrow \gamma\gamma} = \left(\frac{c_\gamma}{\Lambda} \right)^2 \frac{m_a^3}{64\pi} \quad (7)$$

Here $F_A(t)$ is the nuclear electric form factor, calculated as a Fourier transform of the nuclear charge density:

$$F_A(t) = \frac{1}{Z} \int d^3r \frac{\sin(qr)}{qr} \rho_A(r) \quad (8)$$

Here the 3-momentum transfer is calculated as $q = \sqrt{t(t/4m_A^2 - 1)}$ and the charge density $\rho_A(r)$ for carbon was modelled using the 3-parameter Fermi model:

$$\rho_A(r) = C \frac{1 + wr^2/c^2}{1 + \exp(r - c)/z} \quad (9)$$

where $c = 2.355$ fm, $z = 0.5224$ fm, and $w = -0.149$ are taken from [13], and C is chosen such that $\int d^3x \rho_A(x) = Z$.




Figure 21: *Left:* Charge density parameterization for Carbon-12. *Right:* Electric form factor calculation for Carbon-12.

This cross section model was implemented into an event generator to produce pseudo-data for a given ALP mass hypothesis. The event generator randomly selected a beam photon energy according to the distribution of tagged photon energy as measured using the Pair Spectrometer, as described in Section 4.

Upon selection of the beam photon energy, the 4-momentum transfer t and the azimuthal ALP production angle ϕ_X were randomly selected from a phase-space distribution of choice $\frac{1}{2\pi}P(t)$; ϕ_X is the azimuthal angle of the ALP momentum p_X around the beamline, and the value of t is defined earlier. The generated event was then assigned a weight according to the ratio of the cross section and the sampling distribution:

$$w_i = \frac{2\pi}{P(t)} \frac{d\sigma_{\gamma A \rightarrow aA}}{dt}(E_\gamma, t, m_a) \quad (10)$$

This process was used to produce a large number of weighted events distributed according to the selected sampling distribution, with varying weights. To limit the number of events simulated in the detector model, these events were reservoir sampled using the A-ExpJ algorithm as described in Ref. [14] to produce a smaller, but unweighted, set of events distributed according to the cross section model. This maximized the statistical power of events passed through the detector model.

3.2 GEANT

In order to assess the experimental effects of the GlueX detector on signal diphoton events, simulated events were run through a full GEANT4 model of Hall D and the GlueX detector. Details of the simulated geometry and detector response are given in Section 14 of Ref. [11]. The output of this simulation and smearing are formatted in the same manner as raw data, and therefore are subjected to the same reconstruction and event selection steps as described in Section 2.

3.3 Random Trigger Event Mixing

It is necessary to assess the additional impact of pileup on signal events. Accidental coincidence of additional tracks, showers, and other detector hits can impact the efficiency of the event selection criteria for signal events, and this effect must be quantified.

As described in subsection 2.2, a random-trigger event sample was taken during each run in order to assess pileup impact. This random data sample was superimposed onto the simulation, with the detector response for a given simulated event being mixed with the detector response from a random trigger event. This simulates the pileup impact due to random coincidence, allowing for a realistic calculation of cut and veto efficiencies for signal events.

3.4 Efficiency and Resolution Determination

The output of the signal simulation allows us to examine the impact of the detector and our cuts on the desired signal events. The yield of signal events was taken at various stages of analysis to determine the impact of various analysis steps on the acceptance. Figure 22 shows this impact for different ALP mass hypothesis, showing the fraction of events accepted in initial preselection criteria, satisfying trigger conditions, passing the selection vetoes and finally satisfying the selection cuts. The final line, which shows the compounded impact of

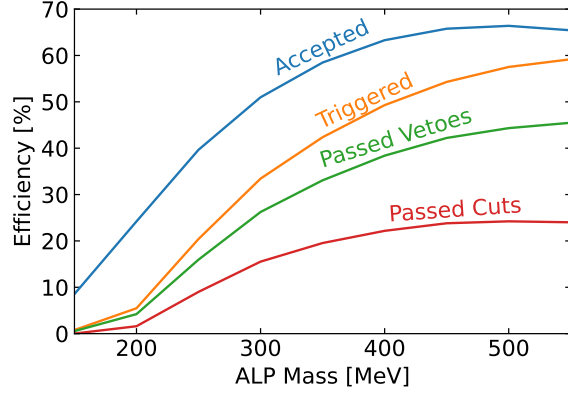


Figure 22: Cut efficiencies.

all selection criteria on the final acceptance, is a necessary factor in converting measured yield limits to coupling limits.

The shape of the signal for ALP events was examined in order to determine our resolution on the reconstructed diphoton mass. The signal (assuming no measurable decay width) was found to be well-described by a Gaussian distribution. The extraction of the measured resolution on the diphoton mass is shown in Figure 23, both for simulation of a $m_a = 300$ MeV ALP and for the data at the η resonance (where the latter assumed a constant background). Figure 24 shows the diphoton resolution as a function of the diphoton mass,

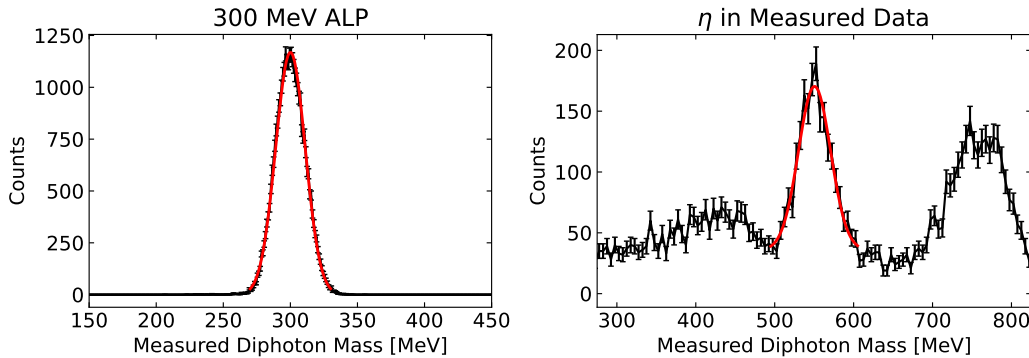


Figure 23: *Left*: The distribution of the measured diphoton mass in simulation with an ALP mass of $m_a = 300$ MeV, together with a Gaussian fit used to extract the mass resolution. *Right*: The same fit for the η in measured data.

including simulation for different values of m_a and the measured η in data. The resolution is found to be 3 – 4% of the true diphoton mass for our signal after all selection criteria were applied. This was found to agree with the data for the η within statistical uncertainties.

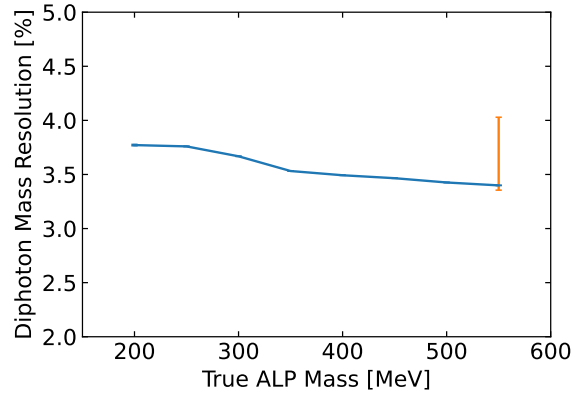


Figure 24: *Left:* The measured resolution of the diphoton as a function of the diphoton mass for the simulation of different ALP masses (blue) and the η in data (orange), where the latter has significant statistical uncertainty.

4 Flux/Luminosity

One of the important tasks of the Hall D pair spectrometer (PS) is to determine a flux of collimated beam photons incident on the GlueX target, which is needed to measure cross sections of various physics processes. The photon flux is obtained by reconstructing electron-positron pairs produced by beam photons during the physics run. The PS was integrated into the GlueX trigger system and allowed to record e^+e^- candidates in parallel with taking experimental data. The number of beam photons (N_γ) is related to the number of electron-positron pairs, ($N_{e^+e^-}$), detected by the pair spectrometer according to the following expression:

$$N_\gamma = \frac{N_{e^+e^-}}{N_{\text{conv}} \sigma_{e^+e^-} \epsilon A}, \quad (11)$$

where N_{conv} is the number of atoms in the pair spectrometer converter, $\sigma_{e^+e^-}$ is the pair production cross section, ϵ is the efficiency of detecting leptons in the PS counters, and A is the PS acceptance. The denominator in Eq. 11, $K = N_{\text{conv}} \sigma_{e^+e^-} \epsilon A$, was obtained during PS calibration runs, where we simultaneously measured the number of electromagnetic pairs and the number of photons in the beam. For the calibration, we used a small electromagnetic calorimeter, which was inserted into the photon beam and allowed us to directly count the number of beam photons.

The number of beam photons is determined for each tagger counter by requiring a hit coincidence between the PS and tagging detectors. Each tagger counter corresponds to the specific beam energy, so a hit in the tagging detectors is required to define the energy of a beam photon. An example of the tagged beam photon energy distribution is presented in Fig. 25. Holes in the distribution correspond to gaps between tagging counters. The tagger hodoscope counters are sampled below about 7.2 GeV. An enhancement in the distribution around 8.5 GeV corresponds to the coherent peak in the photon beam energy spectrum. The shape of the energy distribution depends on the type of the beam radiator (diamond or amorphous) and orientation of the diamond radiator, and varied from run to run.

The tagged photon energy distribution is used in the event generator by randomly selecting a beam photon energy according to this distribution. The tagged energy spectra are determined from the PS data for every run in the SRC experiment and are stored in the calibration database. The Hall D simulation framework allows to generate MC samples according to realistic run-by-run dependent distributions of tagged photon energy spectra and electron beam energies.

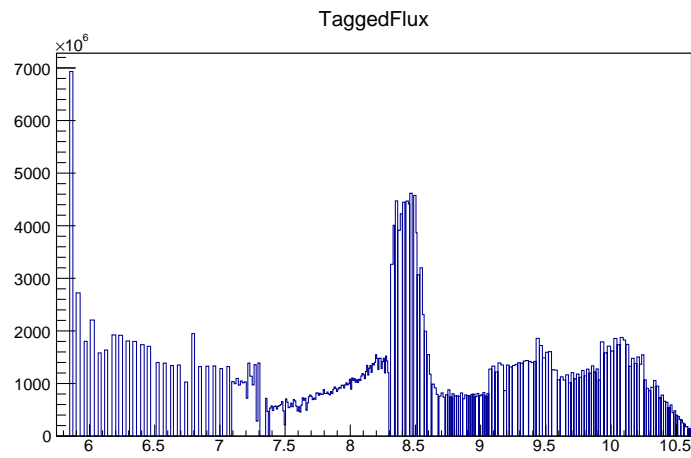


Figure 25: Tagged beam photon energy distribution for run 90271 as measured using the Pair Spectrometer.

5 Results

5.1 Blinding Procedures

We performed a blinded analysis on the data by initially only analyzing 10% of the data. We compared our 10% sample with simulation and tested our bump hunt procedure by inserting a fake signal into our data (see results in Section 5.3.2). We also analyzed the η channel resolution in this sample with the expected mass resolution from simulation in order to validate our understanding of the mass resolution in our data. We made predictions using the 10% sample to estimate the reach that would be obtained with the same results in the full 100% data sample. We changed no data cuts in event selection, vetoes, and analysis for the full data sample. It is the objective of this document to obtain approval for the full unblinding of our data.

The diphoton mass spectrum of the blinded subset of data is shown in Figure 40. Peaks resulting from $\eta \rightarrow \gamma\gamma$ and $\omega \rightarrow \pi^0\gamma$ decay may be clearly seen, and limit searches in the higher-mass region. Below the η mass is a smooth background; this remaining background is presumed to result from η and ω production in air downstream from the target. The assumed vertex results in a lower reconstructed mass due to an underestimated opening angle, and the lack of vertexing in a neutral-only final-state limits our ability to remove such backgrounds.

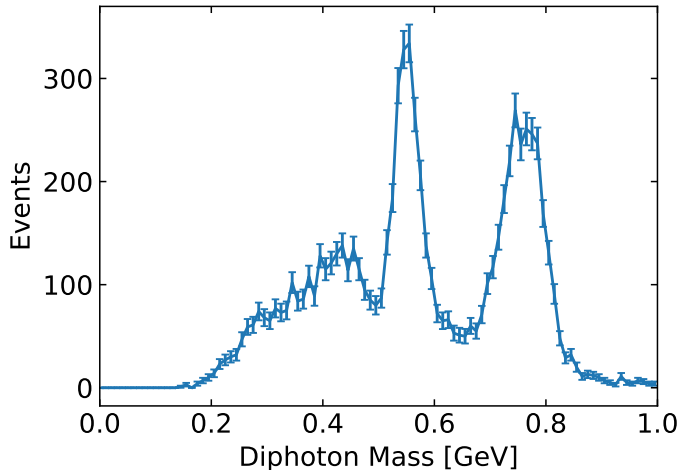


Figure 26: 2-photon mass spectrum after event selection in blinded subset of data.

5.2 Statistical Analysis

We performed a bump-hunt on the 2-photon mass spectrum using the statistical procedure outlined in Ref. [15]. The distribution of 2-photon resonance signal was seen in simulation to follow a Gaussian shape, and the resolution of this Gaussian $\sigma_m(m_X)$ was taken from simulation for a given m_X hypothesis (see Fig. 24). This signal distribution function shall be denoted

$$S_{m_X}(m) = \frac{1}{\sqrt{2\pi\sigma_m^2}} \exp\left(-\frac{(m - m_X)^2}{2\sigma_m^2}\right)$$

clearly normalized to $\int S_{m_X}(m) dm = 1$, where m is the *measured* 2-photon mass and m_X is the truth-level mass of the hypothesized resonance. (The normalization of this signal level will be later introduced as a parameter.) The smooth, non-resonant 2-photon background was modeled as a polynomial distribution, denoted

$$B_{m_X}(m|\mathbf{a}) = \sum_{i=0}^l a_i (m - m_X)^i$$

where $\mathbf{a} = \{a_i\}$ is the set of coefficients defining the polynomial. (The overall normalization of this background is ultimately a nuisance parameter and therefore given no special attention.) Ref. [16] shows that a relationship exists between the size of the fitting window relative to the resolution of the signal and the appropriate order of polynomial fit to the background to avoid signal bias. We selected a fit window width of $\Delta m = 20\sigma_m$; that is, for a given mass hypothesis m_X the 2-photon mass spectrum was considered in the window of $m_X - 10\sigma_m(m_X) < m < m_X + 10\sigma_m(m_X)$. This was somewhat altered in the upper mass region, where the presence of the $\eta \rightarrow \gamma\gamma$ peak presented a non-smooth background; in this region, a boundary at 500 MeV was implemented, and the window was instead $500 \text{ MeV} - 20\sigma_m(m_X) < m < 500 \text{ MeV}$. Similarly, 180 MeV was taken as a low-mass limit due to the lack of data below that. In each window the number of bins was selected as 400, giving a bin width of $\sigma_m/20$. (The effect of different choice of binning is shown later.) To pair with this a polynomial order of $l = 4$ was considered.

Combining these distributions for signal and background, the overall distribution for the 2-photon mass spectrum is parameterized as $B_{m_X}(m|\mathbf{a}) + \mu S_{m_X}(m)$, where we have introduced the signal strength parameter μ , equivalent to the *number of signal counts* in the distribution. In a given bin j , we may integrate over the mass range to determine the normalized number of signal counts

$$s_j = \int_{m \in \text{bin } j} S_{m_X}(m) dm$$

and the number of background counts

$$b_j(\mathbf{a}) = \int_{m \in \text{bin } j} B_{m_X}(m|\mathbf{a}) dm$$

where we have suppressed the m_X index, but each mass hypothesis is considered separately. We denote the *observed* number of counts in the bin as n_j , which will naturally be a sum of signal and background events. The likelihood of observing that number of counts assuming the underlying distribution defined by μ and \mathbf{a} is given by a Poisson distribution:

$$L_j(\mu, \mathbf{a}) = \frac{(b_j(\mathbf{a}) + \mu s_j)^{n_j}}{n_j!} e^{-(b_j(\mathbf{a}) + \mu s_j)}$$

Taking all bins into consideration, the overall likelihood function is the product of each bin's individual Poisson likelihood:

$$L(\mu, \mathbf{a}) = \prod_{j \in \text{bins}} \frac{(b_j(\mathbf{a}) + \mu s_j)^{n_j}}{n_j!} e^{-(b_j(\mathbf{a}) + \mu s_j)}$$

For a frequentist statistical analysis, we must consider maximization of this likelihood function of the fit parameters μ, \mathbf{a} . We define the parameters that *globally* maximize this likelihood over the full parameter space $\hat{\mu}, \hat{\mathbf{a}}$.

As signal strength μ is our parameter of interest, we can also consider maximizing this likelihood at a fixed value of μ . The background parameters which maximize L at a fixed value of μ are defined $\hat{\mathbf{a}} = \hat{\mathbf{a}}(\mu)$. This may be used to quantify the goodness-of-fit of a signal strength parameter μ through the profile likelihood ratio:

$$\lambda(\mu) = \frac{L(\mu, \hat{\mathbf{a}})}{L(\hat{\mu}, \hat{\mathbf{a}})}, \quad (12)$$

where we compare our best estimator $\hat{\mu}$ against the alternatives. We will use the profile likelihood ratio both in test of discovery as well as test of exclusion.

5.3 Tests of Discovery

We first wish to determine if these data allow us to identify any statistically significant signal at a given value of m_X . To this end we define our first test statistic q_0

$$q_0 = -2 \ln \lambda(0) = -2 \ln \frac{L(0, \hat{\mathbf{a}})}{L(\hat{\mu}, \hat{\mathbf{a}})} \quad (13)$$

This statistic quantifies the difference in goodness-of-fit between the null (background-only) hypothesis, and the global best fit to the data, in which μ may be nonzero. In the Wald approximation [17], this quantity is equivalent to the squared Z-score of the null hypothesis:

$$q_0 = Z_0^2 \quad (14)$$

This means that we may directly calculate the p-value of the null hypothesis using this quantity:

$$p_0 = \begin{cases} 1 - \Phi(+\sqrt{q_0}) & \hat{\mu} > 0 \\ 1 - \Phi(-\sqrt{q_0}) & \hat{\mu} < 0 \end{cases} \quad (15)$$

Here we have defined Φ as the cumulative distribution function for the standard Gaussian with a mean of 0 and a standard deviation of 1.

By performing the global minimization for $L(\hat{\mu}, \hat{\mathbf{a}})$ and the background-only minimization for $L(0, \hat{\mathbf{a}})$, we may calculate the value of q_0 and therefore p_0 for each considered hypothesis m_X . Figure 27 shows the local p-value for signal discovery over the examined mass range between 200 MeV and 450 MeV. The largest observed excess is seen at $m_a \approx 210$ MeV, showing a 2σ significance for signal over the background hypothesis.

5.3.1 Look-Elsewhere Effect

The *global* p-value for signal discovery must be calculated accounting for the Look-Elsewhere Effect, which states that the probability of observing a statistically significant signal *some-where* in our search range is larger than the *local* probability of any specific signal observation. The probability of a signal observation must therefore be modified using a “trial




Figure 27: Local p-value for signal discovery over the examined mass range. The largest excess is observed at $m_a \approx 440$ MeV, with 2σ significance.

factor”, which accounts for this probability by quantifying the size of the search region and accounting for the fact that such a search is effectively performing multiple quasi-independent searches.

We use the asymptotic formula given in Ref. [18] to calculate the trial factor at a given local significance Z_{local} :

$$trial\# \equiv \frac{p_{global}}{p_{local}} \approx 1 + \sqrt{\frac{\pi}{2}} \mathcal{N} Z_{local} \quad (16)$$

Here \mathcal{N} denotes the “effective number” of independent search regions within our spectrum. This value can be estimated using the “upcrossing” approximation of Ref. [18], which relates it to the number of upcrossings $N(q_{ref})$ of $q_0 = Z_0^2$ as a function of mass above some threshold value q_{ref} :

$$\mathcal{N} = \langle N(q_{ref}) \rangle e^{q_{ref}/2} \quad (17)$$

This value may be taken from data. Setting $q_{ref} = 0$ allows us to examine the number of upcrossings of the p-value above the 0σ line; looking at Fig. 27 we may see that we have $\mathcal{N} = \langle N(q_{ref} = 0) \rangle = 3 \pm 1.7$. For a maximum value of $Z_{local} = 2.08$ and a minimum value of $p_{local} = 1.88 \times 10^{-2}$, we calculate a trial factor of

$$trial\# \approx 8.8 \pm 4.5$$

a global p-value of

$$p_{global} \approx 1.7 \pm 0.84 \times 10^{-1}$$

and a global significance of

$$Z_{global} = 0.97 \pm 0.34$$

We therefore do not observe any excess with a larger global significance than 1σ .

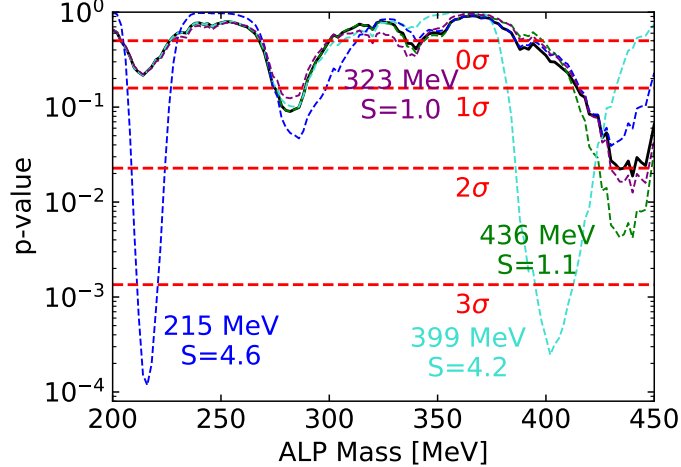


Figure 28: Local p-value for signal discovery over the examined mass range. The black curve shows the p-value for the existing dataset, while the colored curves show the p-value after random signal injection.

5.3.2 Signal Injection Tests

In order to confirm the ability of our method to successfully identify signals in data, we performed tests by injecting randomly generated signal event samples in the data and again using the test of discover. The mass of the signal was randomly generated between 200 and 450 MeV, with the resolution being taken from simulation. The significance of the signal was randomly generated between 0 and 5, where the signal significance was defined as the ratio between the inserted signal and the uncertainty on the data integral within 2σ of the signal peak:

$$S = \frac{N_{injected}}{\sqrt{\int_{\pm 2\sigma} N_{data}(m) dm}} \quad (18)$$

Fig. 28 shows the results in the case of 4 randomly generated signals compared with that for the existing dataset. The measured p-value was smaller than the significance of the injected signal, largely due to the ability of the background model to absorb some portion of the signal strength, but the signals clearly resulted in decreases at the appropriate masses of the p-value.

5.3.3 η Discovery Test

The ability of the algorithm to positively identify the signal of the $\eta \rightarrow \gamma\gamma$ decays in data was examined to verify the discovery test. Fig. 28 shows the local p-value evaluated near the mass of the η meson $m_\eta = 548$ MeV. The p-value shows an extreme dip in the vicinity of the known η mass, clearly identifying the peak seen in Fig. 40 as highly statistically significant.

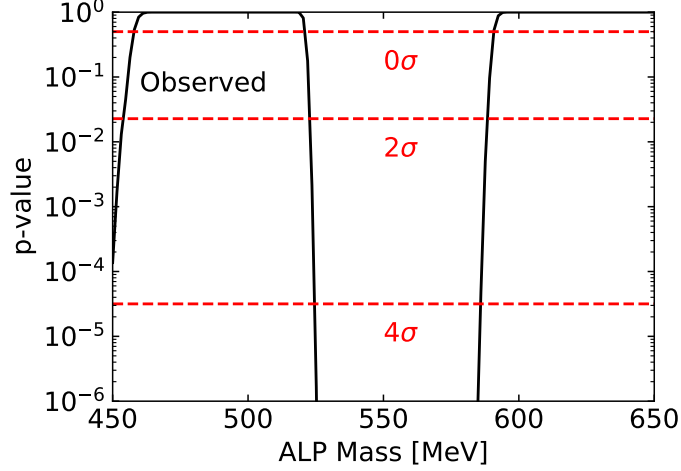


Figure 29: Local p-value for signal discovery over the mass range near the mass of the η meson. The signal of the η production and decay may be clearly seen to be highly significant.

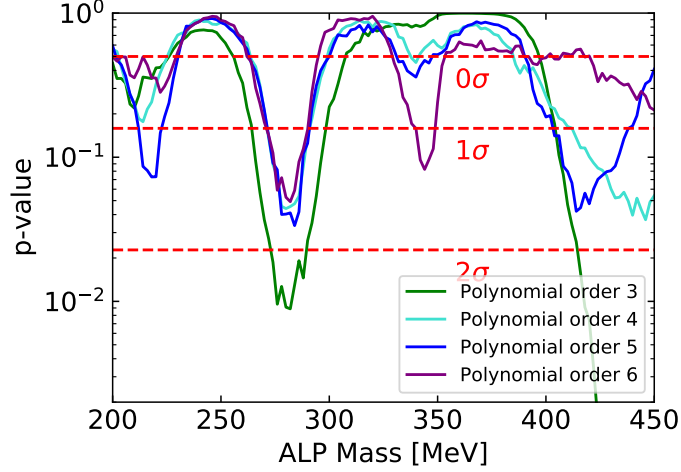


Figure 30: Local p-value for signal discovery over the examined mass range. Different colors correspond to different background polynomial orders.

5.3.4 Polynomial Order Examination

The effect of the order of polynomial used to fit the background was also examined. Fig. 30 shows the local-pvalue for different background models using polynomial order $l = 3, 4, 5, 6$. The specifics of the obtained p-value differ as a function of the polynomial order, but the qualitative features of the local p-value as a function of the mass remain the same. With the exception of the $l = 3$ polynomial, which is clearly insufficiently flexible to describe the mass spectrum shoulder near 450 MeV, no background model suggests a significant signal in the data.

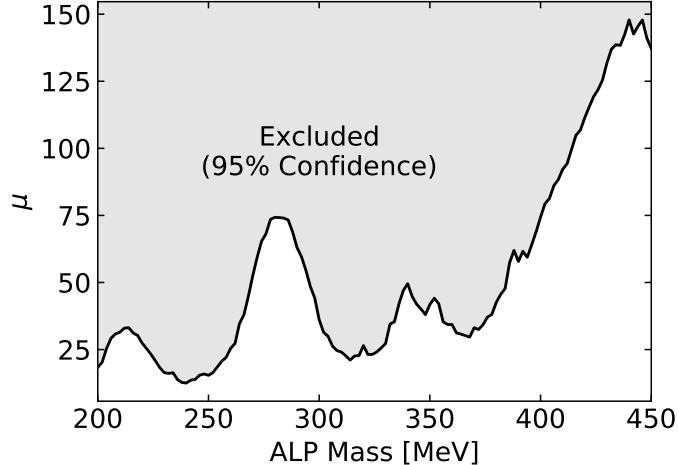


Figure 31: 95% exclusion range on μ as a function of m_a .

5.4 Tests of Exclusion

We also wish to use the data to place upper limits on the signal strength μ , in order to translate these to limits on the coupling c_γ/Λ . This requires us to compare fits at fixed μ with the best fit hypothesis to determine a specific shift in the likelihood. However, we should account for the fact that fits with $\mu < 0$ are nonphysical when performing this comparison. Assuming the background-only hypothesis of the underlying data distribution is accurate, we expect roughly half the time to observe $\hat{\mu} < 0$ due to statistical fluctuations. In these cases we wish to instead compare to the background-only null hypothesis, to avoid the possibility of having a negative upper limit. We therefore define a truncated version of the profile likelihood ratio:

$$\lambda(\mu) = \begin{cases} \frac{L(\mu, \hat{\mathbf{a}})}{L(\hat{\mu}, \hat{\mathbf{a}})} & \hat{\mu} > 0 \\ \frac{L(\mu, \hat{\mathbf{a}})}{L(0, \hat{\mathbf{a}}(0))} & \hat{\mu} < 0 \end{cases} \quad (19)$$

Using this we may define our second test statistic \tilde{q}_μ :

$$\tilde{q}_\mu = \begin{cases} -2 \ln \tilde{\lambda}(\mu) & \mu > \hat{\mu} \\ 0 & \mu < \hat{\mu} \end{cases} \quad (20)$$

This quantity is equivalent to the squared Z-score of the hypothesis at fixed μ as compared to the best-fit hypothesis for a given mass. To 95% confidence, we find that hypotheses with $\tilde{q}_\mu > (1.64)^2$ are excluded, giving us our range of excluded values of μ for a given mass. Fig. 31 shows the limits extracted on the value of μ over our search range to 95% confidence.

5.4.1 Normalization

Conversion of the limit on μ , equivalent to the observed signal yield, requires that we take into account the overall normalization of the measured data. A precise method for doing so was proposed in Ref. [1]; in normalizing against the primakoff production of known resonances

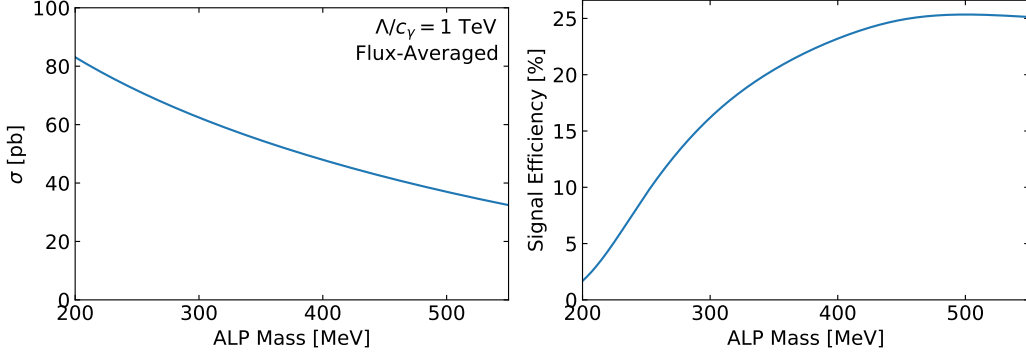


Figure 32: *Left*: Primakoff photoproduction cross section as a function of diphoton resonance mass m_a at fixed coupling $c_\gamma/\Lambda = 1 \text{ TeV}^{-1}$, averaged over the tagged beam photon flux with $E_\gamma > 6 \text{ GeV}$. *Right*: Efficiency for Primakoff diphoton signal detection with $E_\gamma > 6 \text{ GeV}$.

such as the π^0 or the η , we may reduce the systematics involved with the experiment and with the modelling of nuclear structure effects, and achieve higher precision than the use of luminosity measurements would allow.

The yield Y for some signal process is equal to the product of the luminosity \mathcal{L} , the cross section σ , and the efficiency of detection ϵ for the process (as well as the branching ratio in the case of the η):

$$Y = \mathcal{L} \times \sigma \times \epsilon \times \mathcal{B}(X \rightarrow \gamma\gamma) \quad (21)$$

Here the luminosity is integrated above a threshold of $E_\gamma > 6 \text{ GeV}$, the cross section is averaged over the tagged photon flux, and the efficiency is averaged over final-state kinematics as described in Section 3.4.

The equal luminosity between the ALP and η Primakoff production measurements leads to the following equality:

$$\frac{Y_a}{\sigma_a \epsilon_a} = \frac{Y_\eta}{\sigma_\eta \epsilon_\eta \mathcal{B}(\eta \rightarrow \gamma\gamma)} \quad (22)$$

As the yield Y is equal to the signal strength μ , the ratio between the ALP and η production cross sections is proportional to the ratio of the signal strengths:

$$\frac{\sigma_a}{\sigma_\eta} = \frac{\epsilon_\eta \mu_a}{\epsilon_a \mu_\eta} \mathcal{B}(\eta \rightarrow \gamma\gamma) \quad (23)$$

As shown in Eqs. 5 and 7, the cross section is proportional to the square of the coupling ratio c_γ/Λ ; this may be factorized out, with both cross sections being calculable at a fixed reference coupling ratio $(c_\gamma/\Lambda)_{ref}$:

$$\frac{(c_\gamma/\Lambda)_a^2 \sigma_{ref}(m = m_a)}{(c_\gamma/\Lambda)_\eta^2 \sigma_{ref}(m = m_\eta)} = \frac{\epsilon_\eta \mu_a}{\epsilon_a \mu_\eta} \mathcal{B}(\eta \rightarrow \gamma\gamma) \quad (24)$$

The ratio of the squares of the coupling constants may be given by a correction factor \mathcal{F} to the ratio of the signal strengths:

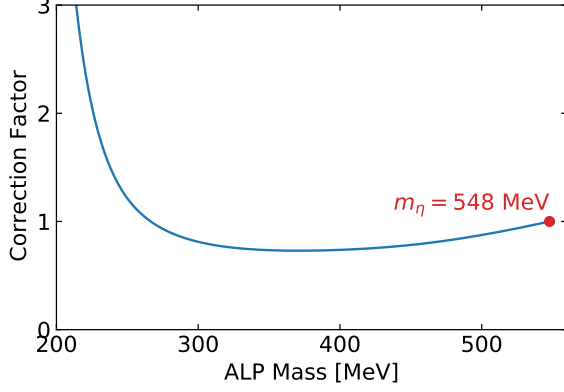


Figure 33: Correction factor \mathcal{F} as a function of m_a to the yield ratio to calculate the squared coupling ratio (relative to the η meson).

$$\frac{(c_\gamma/\Lambda)_a^2}{(c_\gamma/\Lambda)_\eta^2} = \frac{\sigma_{ref}(m = m_\eta)\epsilon_\eta}{\sigma_{ref}(m = m_a)\epsilon_a} \times \frac{\mu_a}{\mu_\eta} \times \mathcal{B}(\eta \rightarrow \gamma\gamma) = \mathcal{F}(m_a) \times \frac{\mu_a}{\mu_\eta} \times \mathcal{B}(\eta \rightarrow \gamma\gamma) \quad (25)$$

Fig. 32 shows the value of the fixed-coupling cross section and the signal efficiency, each as a function of resonance mass. Fig. 33 shows the value of the correction factor of Eq. 25.

The coupling of the ALP may from here be calculated relative to the $\eta - \gamma$ coupling using the measured signal strength η . Specifically, we are interest in the 95%-confidence limit on the ALP coupling, which is similarly related to the 95% limit on μ :

$$\left(\frac{c_\gamma}{\Lambda}\right)_{95\%} = \sqrt{\mathcal{F}(m_a) \times \frac{\mu_{95\%}}{\mu_\eta} \times \mathcal{B}(\eta \rightarrow \gamma\gamma)} \left(\frac{c_\gamma}{\Lambda}\right)_\eta \quad (26)$$

We may calculate the absolute value of this quantity using the known properties of the η meson. Using Eq. 7, we may determine the $\eta - \gamma$ coupling from its partial decay width to $\gamma\gamma$:

$$\left(\frac{c_\gamma}{\Lambda}\right)_\eta = \sqrt{\frac{64\pi}{m_\eta^3} \Gamma_{\eta \rightarrow \gamma\gamma}} = \sqrt{\frac{64\pi}{m_\eta^3} \Gamma_\eta \mathcal{B}(\eta \rightarrow \gamma\gamma)} \quad (27)$$

Using the values of Ref. [19], we calculate this coupling to be $(c_\gamma/\Lambda)_\eta = 25.1 \pm 0.48 \text{ TeV}^{-1}$.

We use the same fitting procedure previously described to extract the signal strength of the η peak in our mass spectrum. In this case we use a $10\sigma_m$ mass window to fit the spectrum, due to the complexity of the background shape around the η mass, along with a 4th-order polynomial background, a Gaussian signal of known mean m_η and width $\sigma_m(m_\eta)$, and 400 bins. This fit is shown in Fig. 34, with a reduced number of bins in aid in viewing the data. This fit, as well as examination of the test statistic \tilde{q}_μ to quantify uncertainties, provides us with a measured signal strength of $\mu_\eta = 1200 \pm 100$.

These values together allow us to use the limits extracted in Fig. 31 along with Eq. 26 to place exclusion limits on the ALP coupling. Fig. 35 shows the exclusion as a function of m_a




Figure 34: Fit to the mass spectrum around the mass of the η meson. The dashed red line shows the full best-fit to the data, while the solid magenta line shows the background component of this fit.

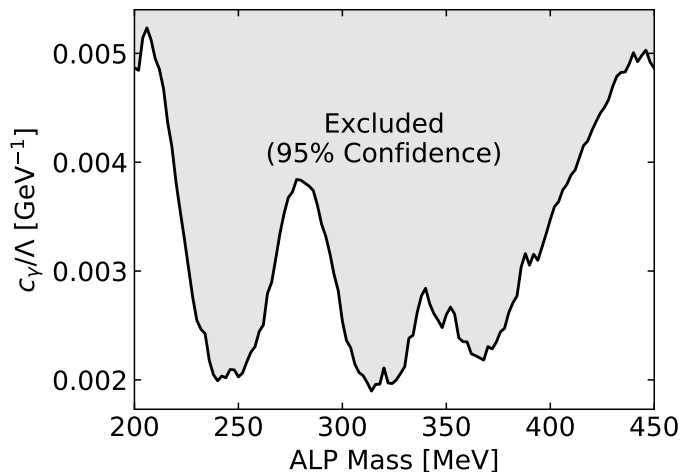


Figure 35: 95% exclusion range on c_γ/Λ as a function of m_a .

over the examined range. We observed that the current sample of data places upper limits on the coupling around $c_\gamma/\Lambda \lesssim 5 \text{ TeV}^{-1}$ for most of the range.

5.4.2 Expected Exclusion Range

As an extension of this, we would like to determine the *expected* exclusion limit on μ given the background-only hypothesis. That is, in the event that the best-fit background-only hypothesis for the underlying distribution is accurate, we would like to calculate what level of exclusion we would expect given our number of statistics, as well as the expected level of deviation in this. The benefits of this calculations are twofold. First, this allows us to verify that the fluctuations observed in the limits set by the data are within expectations for the quantity of data used. Second, this allows us to project the limits we expect to set when using the full set of data, allowing us to verify that the unblinded results are consistent with




Figure 36: *Left*: The best-fit background-only (blue, solid) and signal-allowed (red, dashed) fits to the data when testing the $m_a = 300$ MeV hypothesis. *Right*: The Asimov dataset for this mass range under the null hypothesis.

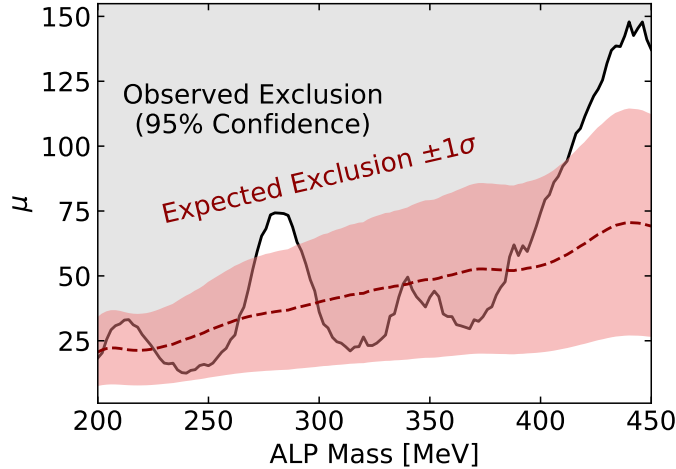


Figure 37: 95% exclusion range on μ as a function of m_a . The black line and shaded region show the exclusion limits set by the current subset of the data. The red band shows the range of expected exclusion of the current dataset given the background-only hypothesis.

the blinded predictions.

This may be calculated using what we call the “Asimov dataset”, as described in Ref. [15]. The Asimov dataset defines the set of data in which each bin takes on precisely its expectation value as its contents (under some hypothesis). This has the effect of smoothing out local statistical fluctuations. In our case, we may take the best background-only fit (the null hypothesis) as the basis for our Asimov dataset at each mass point, with one such example shown in Fig. 36.

This pseudo-dataset may be used to perform the same exclusion test as described previously in order to calculate the 95% limit on μ as set by $\tilde{q}_\mu = (1.64)^2$. We may additionally calculate the expected variation in this bound by changing this equality; the $\pm 1\sigma$ range on the exclusion is calculated using $\sqrt{\tilde{q}_\mu} = 1.64 \pm 1$, giving us a band of expected exclusion limits, as shown in Fig. 37.

Repeating the same normalization procedure as described in Section 5.4.1, we may convert this exclusion range on μ to the expected range of exclusion on the coupling. The red band in

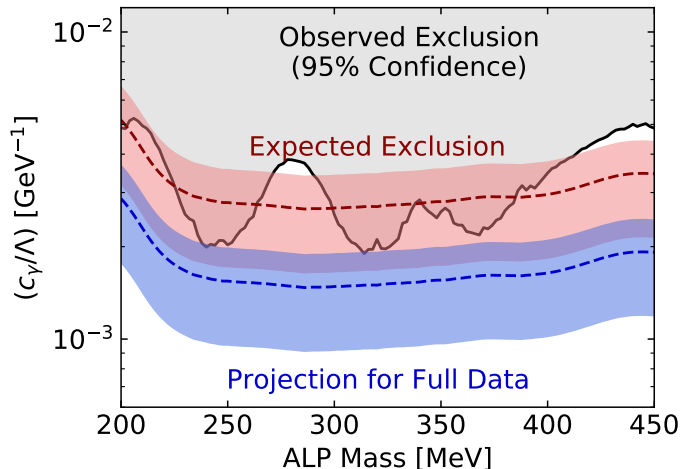


Figure 38: 95% exclusion range on the ALP-photon coupling as a function of m_a . The black line and shaded region show the exclusion limits set by the current subset of the data. The red band shows the range of expected exclusion of the current dataset given the background-only hypothesis. The blue band shows the projections for the limits set by the full dataset.

Fig. 38 shows the 1σ range of expected exclusion on the ALP-photon coupling as a function of the mass.

Finally, we may use this exclusion expectation to estimate the limits set by the full (unblinded) dataset. As the limits on the cross section will scale as $1/\sqrt{N}$ and the cross section scales as a square of the coupling $(c_\gamma/\Lambda)^2$, this projection may be performed by multiplying the expected exclusion band by a factor of $f^{-1/4}$, where $f \approx 10\%$ is the fraction of data used in this blinded analysis. The blue band in Fig. 38 shows the projected exclusion for the full dataset.

Fig. 39 shows the limits set by the current data and the projected full data limits compared with existing limits set by data, as well as projections for future studies. The limits set by the full data are unlikely to be world-leading over most of the range except in the case of a favorable fluctuation in the spectrum. In most cases the limits set by the Belle and BESIII experiments supersede those set by this data. A higher-statistics set of data with a target such as 208-Pb would clearly provide more stringent limits in our mass range.




Figure 39: 95% exclusion range on the ALP-photon coupling as a function of m_a . The black line and shaded blue region show the exclusion limits set by the current subset of the data and the projected dataset, respectively. The grey region shows limits set by existing analysis. The dashed lines show the limits projected for other experiments.

6 Reference Channel Normalization Check

In order to verify that the absolute normalization of the limits set by this work are accurate, a check was performed of the absolute yield of measured Primakoff $\eta \rightarrow \gamma\gamma$ events.

The total number of photon-tagged events produced per unit photon energy is given by the product of the cross section (evaluated at a given E_γ) and the differential luminosity:

$$\frac{dN_{produced}}{dE_\gamma} = \frac{d\mathcal{L}}{dE_\gamma} \times \sigma(E_\gamma) \quad (28)$$

For beam photons with an energy greater than our threshold of 6 GeV, the differential luminosity may be expressed as the product of the total luminosity L and the normalized tagged photon energy distribution $P(E_\gamma)$:

$$\frac{d\mathcal{L}}{dE_\gamma} = L \times P(E_\gamma) \quad (29)$$

The total number of produced events may be determined by integrating over the beam photon energy above our threshold:

$$N_{produced} = \int dE_\gamma \frac{dN_{produced}}{dE_\gamma} = \mathcal{L} \times \int dE_\gamma P(E_\gamma) \sigma(E_\gamma) = \mathcal{L} \times \langle \sigma \rangle \quad (30)$$

Here we have defined $\langle \sigma \rangle$ as the flux-averaged total cross section. The number of events resulting a 2-photon final-state requires us to factor in the branching fraction to $\gamma\gamma$

$$N_{\gamma\gamma} = \mathcal{L} \times \langle \sigma \rangle \times \mathcal{B}(\eta \rightarrow \gamma\gamma) \quad (31)$$

Finally, the efficiency ϵ for detection of Primakoff 2-photon final-states must be factored in to calculate the total measured yield:

$$N_{meas} = \mathcal{L} \times \langle \sigma \rangle \times \mathcal{B}(\eta \rightarrow \gamma\gamma) \times \epsilon \quad (32)$$

The total tagged luminosity for photons with energy above 6 GeV was measured to be 8.21 pb⁻¹·nucleus. The cross section model and event generator described in Section 3.1 was used to calculate a total flux-averaged cross section of 20.58 nb/nucleus. The branching fraction for ($\eta \rightarrow \gamma\gamma$) was taken from Ref. [19] to be 39.41%. Finally, the efficiency was calculated using the event generator output as described in Section 3.4 to be 25%.

Combining these factors and the data fraction 10% currently being used, we estimate:

- 169,000 Primakoff η events in data
- 16,900 Primakoff η events in the blinded subset of data
- 6,660 Primakoff ($\eta \rightarrow \gamma\gamma$) events in the blinded subset of data
- 1,660 detected Primakoff ($\eta \rightarrow \gamma\gamma$) events in the blinded subset of data

As described in Section 5 and shown in Fig. 34, the measured number of ($\eta \rightarrow \gamma\gamma$) events in our mass spectrum was found to be 1200 ± 100 . Given the difficulty in modelling absolute detector efficiency, as well as the model-dependency of Primakoff calculations on the nuclear form factors, this level of agreement is assessed to be sufficient to validate the use of the η meson yield to normalization. The listed factors of detector response and nuclear effects in large part cancel in the ratio of Primakoff yields.

References

- [1] Daniel Aloni et al. “Photoproduction of Axionlike Particles”. In: *Phys. Rev. Lett.* 123 (7 Aug. 2019), p. 071801. DOI: 10.1103/PhysRevLett.123.071801. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.123.071801>.
- [2] G. Abbiendi et al. and The OPAL Collaboration. “Multi-photon production in ee collisions at $\sqrt{s}=181\text{--}209$ GeV”. In: *The European Physical Journal C - Particles and Fields* 26.3 (2003), pp. 331–344. DOI: 10.1140/epjc/s2002-01074-5. URL: <https://doi.org/10.1140/epjc/s2002-01074-5>.
- [3] Simon Knapen et al. “Searching for Axionlike Particles with Ultraperipheral Heavy-Ion Collisions”. In: *Phys. Rev. Lett.* 118 (17 Apr. 2017), p. 171801. DOI: 10.1103/PhysRevLett.118.171801. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.118.171801>.
- [4] J. D. Bjorken et al. “Search for neutral metastable penetrating particles produced in the SLAC beam dump”. In: *Phys. Rev. D* 38 (11 Dec. 1988), pp. 3375–3386. DOI: 10.1103/PhysRevD.38.3375. URL: <https://link.aps.org/doi/10.1103/PhysRevD.38.3375>.
- [5] J. Blümlein et al. “Limits on neutral light scalar and pseudoscalar particles in a proton beam dump experiment”. In: *Zeitschrift für Physik C Particles and Fields* 51.3 (1991), pp. 341–350. DOI: 10.1007/BF01548556. URL: <https://doi.org/10.1007/BF01548556>.
- [6] Medina Ablikim et al. “Search for an axion-like particle in J/ψ radiative decays”. In: (Nov. 2022). arXiv: 2211.12699 [hep-ex].
- [7] Babette Döbrich et al. “ALPtraum: ALP production in proton beam dump experiments”. In: *Journal of High Energy Physics* 2016.2 (2016), p. 18. DOI: 10.1007/JHEP02(2016)018. URL: [https://doi.org/10.1007/JHEP02\(2016\)018](https://doi.org/10.1007/JHEP02(2016)018).
- [8] Jonathan L. Feng et al. “Axionlike particles at FASER: The LHC as a photon beam dump”. In: *Phys. Rev. D* 98 (5 Sept. 2018), p. 055021. DOI: 10.1103/PhysRevD.98.055021. URL: <https://link.aps.org/doi/10.1103/PhysRevD.98.055021>.
- [9] Asher Berlin et al. “Dark sectors at the Fermilab SeaQuest experiment”. In: *Phys. Rev. D* 98 (3 Aug. 2018), p. 035011. DOI: 10.1103/PhysRevD.98.035011. URL: <https://link.aps.org/doi/10.1103/PhysRevD.98.035011>.
- [10] Matthew J. Dolan et al. “Revised constraints and Belle II sensitivity for visible and invisible axion-like particles”. In: *Journal of High Energy Physics* 2017.12 (2017), p. 94. DOI: 10.1007/JHEP12(2017)094. URL: [https://doi.org/10.1007/JHEP12\(2017\)094](https://doi.org/10.1007/JHEP12(2017)094).
- [11] S. Adhikari et al. “The GlueX beamline and detector”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 987 (2021), p. 164807. ISSN: 0168-9002. DOI: <https://doi.org/10.1016/j.nima.2020.164807>. URL: <https://www.sciencedirect.com/science/article/pii/S0168900220312043>.

- [12] K. Moriya et al. “A measurement of the energy and timing resolution of the GlueX Forward Calorimeter using an electron beam”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 726 (2013), pp. 60–66. ISSN: 0168-9002. DOI: <https://doi.org/10.1016/j.nima.2013.05.109>. URL: <https://www.sciencedirect.com/science/article/pii/S0168900213007201>.
- [13] C. W. De Jager, H. De Vries, and C. De Vries. “Nuclear charge and magnetization density distribution parameters from elastic electron scattering”. In: *Atom. Data Nucl. Data Tabl.* 14 (1974). [Erratum: *Atom. Data Nucl. Data Tabl.* 16, 580–580 (1975)], pp. 479–508. DOI: 10.1016/S0092-640X(74)80002-1.
- [14] Pavlos S. Efraimidis and Paul G. Spirakis. “Weighted random sampling with a reservoir”. In: *Information Processing Letters* 97.5 (2006), pp. 181–185. ISSN: 0020-0190. DOI: <https://doi.org/10.1016/j.ipl.2005.11.003>. URL: <https://www.sciencedirect.com/science/article/pii/S002001900500298X>.
- [15] Glen Cowan et al. “Asymptotic formulae for likelihood-based tests of new physics”. In: *The European Physical Journal C* 71.2 (2011), p. 1554. DOI: 10.1140/epjc/s10052-011-1554-0. URL: <https://doi.org/10.1140/epjc/s10052-011-1554-0>.
- [16] M. Williams. “A novel approach to the bias-variance problem in bump hunting”. In: *Journal of Instrumentation* 12.09 (Sept. 2017), P09034–P09034. DOI: 10.1088/1748-0221/12/09/p09034. URL: <https://doi.org/10.1088/1748-0221/12/09/p09034>.
- [17] Abraham Wald. “Tests of Statistical Hypotheses Concerning Several Parameters When the Number of Observations is Large”. In: *Transactions of the American Mathematical Society* 54.3 (1943), pp. 426–482. ISSN: 00029947. URL: <http://www.jstor.org/stable/1990256> (visited on 09/26/2022).
- [18] Eilam Gross and Ofer Vitells. “Trial factors for the look elsewhere effect in high energy physics”. In: *The European Physical Journal C* 70.1 (2010), pp. 525–530. DOI: 10.1140/epjc/s10052-010-1470-8. URL: <https://doi.org/10.1140/epjc/s10052-010-1470-8>.
- [19] R. L. Workman et al. “Review of Particle Physics”. In: *PTEP* 2022 (2022), p. 083C01. DOI: 10.1093/ptep/ptac097.

A Comments from SRC-CT Group

1. Introduction needs more materials justifying this search and this particle candidate.

Further details on the motivation and theoretical basis will be added in the final paper in consultation with theorists.

2. Impact of selection criteria on background should be shown individually.

Fig. 15 has been added to show the effects of the cuts.

3. TOF-RF veto timing window should be reduced from 6.5 ns to 0.5 ns.

As this cut does not impact the accuracy of the normalization of the η yield and only results in a somewhat lower efficiency for signal events, the more conservative timing window has been kept.

4. RF-FCAL and RF-BCAL timing vetoes should be different, and particularly the BCAL veto should be wide enough to remove potential incoherent contributions to the η peak.

The RF-BCAL cut was widened to 6 ns at Igal's suggestion to avoid overestimating the η Primakoff yield.

5. Show elasticity vs. diphoton mass for the mass range and polar angle considered for the search.

See below:




Figure 40: Elasticity vs Diphoton Mass.

6. Elasticity may be used to improve 2-photon mass resolution by correcting for photon energy smearing.

It is true that the assumption of perfect elasticity enables an improvement in the invariant mass resolution for the 2-photon combinations, compensating for energy resolution effects in the FCAL. However, this introduces some complications in proper handling of the statistics of beam-event combinations, particularly in calculation of the uncertainties of low-count bins, and the improvement in the reach is assessed to be insufficient to change the non-world-leading result. As such, the traditional calculation of the invariant mass is sufficient for this analysis.

7. A validation channel should be used to confirm that the normalization of this method is accurate.

A comparison of the η yield to theory calculations based on the measured luminosity has been added; see Section 6. The agreement between the luminosity-based estimates and the measured yield was found to be good.

8. Mass regions closer to the η peak may be reached by including the η in the background fit.

Expansion of the current mass region is not within the current scope of the work, and including the eta peak results in a more complex background fit, introducing convergence issues.

9. Check the sensitivity of the result to the width of the mass windows around the candidate mass.

Following-up with Yotam about this.

10. The definition of μ should be positive definite, but this is not taken into account in the calculation of the p-value.

μ was allowed to be below 0 for the purposed of calculating the p-value (in order to observe symmetry of the fluctuations around $\mu = 0$) but was defined as positive definite for the purposes of calculating the limits.

11. Do you have a plot of μ_0 /(the expected 1 sigma bound) as a function of ma?

Still need clarification on what this is asking for

12. The relation of Eq. 21 is true for fixed photon energy, but the photon energy in the experiment varies.

Clarifying text was added after Eq. 21 that the cross section is averaged over the tagged flux, at which point the relation is exactly correct.

13. Could you describe how the detection efficiency is determined in Fig. 32?

Clarifying text was added after Eq. 21 pointing to Sec. 3.4, which describes the calculation of efficiency for signal events.

14. Check if the sharp features in the Asimov bands are accurate.

Recalculation of the band removes the feature, which was a result of non-convergence.

15. How sensitive is the mass to the assumed vertex? Does the finite length of the target impact the analysis?

The impact of the finite target length is much smaller than energy resolution effects, and may be safely neglected. (This was resolved in Dec. 12 meeting discussion.)

16. The possible production of η mesons from dense material downstream of the target may introduce false resonances when statistics are increased. The impact of this should be estimated.

This has been looked into and is a valid concern. Any possible “discovery” must be carefully examined and simulation may be used to check consistency with resulting from downstream material at a particular location.

17. Is there a possibility that the “number of upcrossings” may severely fluctuate after the unblinding of the data, causing the Look-Elsewhere Effect to have a large impact on the final p-value?

The number of upcrossings should not vary largely as a function of our statistics, but only as a function of our resolution. While the total impact of the LEE may differ such a difference should be small.