Searching for axionlike particles with data scouting at ATLAS and CMS

Simon Knapen,[1,2,3,*] Soubhik Kumar,[2,3,†] and Diego Redigolo[1,4,‡]

[1]*CERN, Theoretical Physics Department, Geneva 1211, Switzerland*
[2]*Berkeley Center for Theoretical Physics, Department of Physics, University of California, Berkeley, California 94720, USA*
[3]*Theoretical Physics Group, Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA*
[4]*INFN, Sezione di Firenze Via G. Sansone 1, 50019 Sesto Fiorentino, Italy*

(Received 16 March 2022; accepted 17 May 2022; published 10 June 2022)

We investigate the physics case for a dedicated trigger-level analysis for very low mass diphoton resonances at ATLAS and CMS and introduce a new photon isolation criterion. This greatly increases the acceptance for light particles at the expense of a somewhat larger Standard Model background. We show how such an analysis would likely yield new experimental coverage for axionlike particles for masses below 15 GeV, bridging the gap with the region covered by flavor factories.

DOI: 10.1103/PhysRevD.105.115012

I. INTRODUCTION

Diphoton resonances have long been recognized as a fertile ground for potential new discoveries, most notably by the discovery of the Standard Model (SM) Higgs boson in the $\gamma\gamma$ channel [1,2]. The latter was, of course, anticipated before the construction of the LHC, and as a result both ATLAS and CMS have excellent photon identification and reconstruction capabilities. On the theory front, it is well known that any new scalar or pseudoscalar particle with coupling charged fermions will typically generate a coupling to two photons at the one loop level. As was the case for the SM Higgs, this naively subleading channel can be a primary discovery mode if the decay widths to the SM fermions are suppressed or background limited. These considerations have resulted in a number of powerful searches, so far with sensitivity in the mass range between 65 GeV and 3 TeV [3–6]. Extending the reach to lower invariant masses is feasible in the near future with present data, and proposals have already been put forward for both ATLAS and CMS [7,8] as well as LHCb [9].[1] In view of future LHC runs, it is important to investigate how much further new trigger techniques could extend the exploration of the diphoton spectrum at low invariant masses.

The diphoton invariant mass can be written as

$$m_{\gamma\gamma} = \sqrt{2p^{\gamma}_{T_1}p^{\gamma}_{T_2}(\cosh\Delta\eta_{\gamma\gamma} - \cos\Delta\phi_{\gamma\gamma})} \simeq \sqrt{p^{\gamma}_{T_1}p^{\gamma}_{T_2}}\Delta R_{\gamma\gamma}, \tag{1}$$

where $p^{\gamma}_{T_1}$ ($p^{\gamma}_{T_2}$) is the transverse momentum of the leading (subleading) photon and $\Delta R_{\gamma\gamma} = \sqrt{\Delta\eta^2_{\gamma\gamma} + \Delta\phi^2_{\gamma\gamma}}$ is their angular separation in the limit where both pseudorapidity separation $\Delta\eta_{\gamma\gamma} \ll 1$ and azimuthal separation $\Delta\phi_{\gamma\gamma} \ll 1$. This simple formula is instrumental to understand why searches for low invariant mass resonances face severe challenges related to the trigger system. The primary problem is the huge rate of low p_T photons produced in parton showers and in meson decays, π^0 decays in particular. To keep the rate manageable, minimal p_T thresholds and photon isolation criteria are needed. Given Eq. (1), these result in a lower bound on the attainable diphoton invariant mass in a given search.

For example, in [5] CMS relied on a trigger which required a $p^{\gamma}_{T_1} > 30$ GeV ($p^{\gamma}_{T_2} > 18$ GeV) threshold on the leading (subleading) photon, tight isolation, and an invariant mass requirement of $m_{\gamma\gamma} > 55$ GeV. The latter limited the accessible mass range in the analysis to $m_{\gamma\gamma} > 70$ GeV, once the turn-on of the trigger was taken into account. To address this issue, ATLAS has implemented a special trigger stream geared toward low mass diphoton resonances [13]. At the hardware/level 1 trigger (L1 trigger), it requires two photons with a symmetric $p^{\gamma}_{T_{1,2}} > 15$ GeV requirement and loose isolation. This L1 trigger seeds a software/high-level trigger (HLT) with $p^{\gamma}_{T_{1,2}} > 22$ GeV. This trigger

*smknapen@lbl.gov
†soubhik@berkeley.edu
‡d.redigolo@gmail.com

[1]Both ATLAS and CMS have carried out a search for lighter diphoton resonances in ultraperipheral heavy ion collisions [10,11]. These searches are, however, only sensitive if the branching ratio of the new state to the $\gamma\gamma$ channel is $\mathcal{O}(1)$ [12].

combination collected already 138.5 fb^{-1} data during the past LHC run, and it currently represents the best way of exploring the low invariant mass region with ATLAS. CMS has a similar, but asymmetric trigger with $p_{T_1}^{\gamma} > 25$ GeV ($p_{T_2}^{\gamma} > 12$ GeV) L1 thresholds for the leading (subleading) photon, with and without a loose isolation criterion [14].

Here we explore an alternative, forward-looking strategy to probe low mass resonances, which we benchmark against the existing ATLAS trigger [13]. The minimal photon p_T and the diphoton angular separation are interconnected by simple kinematics if we look at diphoton resonances, as in Eq. (1). For this reason, once the trigger thresholds on the photons are accounted for, at low masses the resonances will be sufficiently boosted such that the two photons can spoil each other's isolation criteria. Our proposal tries to address this problem in two ways.

First, we assess whether a strategy of directly using photons reconstructed at the trigger level (called turbo stream [15], data scouting [16], or trigger-level analysis [17]) could yield a further improvement relative to the existing ATLAS trigger [13]. In such a strategy, only a small fraction of the full event record is written to tape, which allows for a much higher output rate. The ATLAS, CMS, and LHCb Collaborations have successfully deployed this technique for low-mass resonance searches, in particular in the context of dijet [16,17] and dimuon [18–20] resonances. For the diphoton case, applying this technique could theoretically allow for lower p_T thresholds and possibly less stringent isolation requirements, as compared to the normal trigger stream.

In addition, we investigate whether the tension between the p_T thresholds and isolation cut can be further alleviated by defining a dedicated isolation variable for a diphoton pair: when summing the p_T of the activity within the isolation cone, we will explicitly exclude the hardest, other photon candidate. We also attempt to assess how the background from jets faking photons behaves when subjected to this different isolation requirement. However, we do not attempt to modify the shower shape variables, and our study can thus be seen as a step toward a broader class of "photon-jet" objects at trigger level [21–25], where the isolation requirements and identification criteria on the shower shape are tailored to the features of the targeted signal.

The experimental implementation of an analysis such as what we propose here is both a very time- and resource-consuming endeavor, and we therefore seek to first estimate its possible gain with a phenomenological study. Our results will hopefully aid the experimental collaborations in determining whether a diphoton analysis directly at the trigger level is feasible and desirable, given the resources at their disposal.

To quantify the possible gain in sensitivity, we use a well-motivated, Kim, Shifman, Vainshtein, and Zakharov (KSVZ) style [26,27] axionlike particle (ALP) as an example benchmark model, defined by

$$\mathcal{L}_a \supset -\frac{1}{2}m_a^2 a^2 - \frac{\alpha_s}{8\pi}\frac{a}{f_a}G\tilde{G} + \frac{E}{N}\frac{\alpha_{\rm em}}{8\pi}\frac{a}{f_a}F\tilde{F}. \quad (2)$$

Here N and E are the anomaly coefficients for the gluon and photon couplings, respectively, and the $a \to \gamma\gamma$ branching ratio is specified by the ratio of E/N. In KSVZ models which are compatible with grand unification one typically finds $E/N \sim \mathcal{O}(1)$, resulting in $\mathrm{Br}(a \to \gamma\gamma) \sim 10^{-3}$ and $\mathrm{Br}(a \to jj) \sim 1$. Despite this seemingly small branching ratio, the $\gamma\gamma$ channel is still a superior signature due to the huge dijet background. The gluon coupling, on the other hand, can yield very large cross sections for m_a below 100 GeV, partially offsetting the small branching ratio to photons.

The remainder of our paper is organized as follows: in Sec. II we describe our simulation framework and analysis strategy, followed by a brief summary of the theory motivation for low mass diphoton resonances in Sec. III. We conclude in Sec. IV. In the Appendix, we provide further details on our simulation framework.

II. ANALYSIS

In this section we describe our simulation framework and analysis strategy, with a special emphasis on how our isolation criteria are defined. We take the ATLAS calorimeter as a primary example, but most aspects of our study should carry over to CMS.

A. Photon identification and isolation

In ATLAS Run-2 data taking the photon candidate has to satisfy reconstruction, identification, and isolation requirements which are described in [28–30]. Here we briefly review the main steps, with a focus on how these requirements can limit the signal efficiency for low mass diphoton resonances. First of all, the photon selection at L1 is based on a reduced Electromagnetic Calorimeter (ECAL) granularity. At present, a trigger ECAL tower is taken to be a rectangle of $\Delta\eta \times \Delta\phi = 0.075 \times 0.098$ [31].[2]

Off-line, a photon candidate is classified as a (loosely) reconstructed photon if a number of criteria are satisfied: first, the energy deposited in a 0.25×0.25 window of the hadronic calorimeter (HCAL) behind the ECAL energy cluster has to be very small compared to the energy deposited in the ECAL. (The ratio between HCAL and ECAL energy deposits is called hadronic leakage in the experimental literature.) Second, ATLAS imposes requirements on the shape of the electromagnetic shower in the middle ECAL layer, such as an upper bound on the lateral width of the shower. In addition, they define the variables

[2]The actual granularity of a single ECAL cell is $\Delta\eta \times \Delta\phi = 0.025 \times 0.0245$ in the middle layer and goes down to $\Delta\eta = 0.0031$ if we consider the strip towers in the first layer of the ATLAS ECAL. This more refined information is, however, only available at the HLT.

$$R_\eta \equiv \frac{E_{3\times7}}{E_{7\times7}} \quad \text{and} \quad R_\phi \equiv \frac{E_{3\times3}}{E_{3\times7}}. \tag{3}$$

Here, e.g., $E_{3\times7}$ is the energy deposited in a 3×7 block of ECAL cells, centered around the barycenter of the photon candidate, where the cells are organized in an $\eta\times\phi$ grid. Both R_η and R_ϕ are bounded from above, to ensure the shower of the photon candidate is sufficiently contained. The above variables are used for loose reconstruction and their primary role is to reject photons from collimated π^0 decays [28,32,33].

Loose (tight) photon isolation furthermore imposes an upper bound on the amount of extra energy deposit within a cone of $\Delta R < 0.2$ ($\Delta R < 0.4$) around the photon candidate. In particular, the ATLAS loose photon isolation variable is defined as

$$r_\gamma \equiv \frac{1}{p_{T_\gamma}}\left(\sum_{\substack{\Delta R<0.2 \\ p_{T_i}>0.5\ \text{GeV}}} p_{T_i}\right), \tag{4}$$

and similarly for tight isolation. A photon is considered isolated if $r_\gamma < 0.05$.[3] The sum runs over all tracks and calorimeter deposits with $p_T > 0.5$ GeV and a distance $\Delta R < 0.2$ from the photon candidate, excluding the candidate itself. This isolation criterion is meant to veto photon candidates that originated from inside a jet, but it can also reject genuine, light diphoton resonances when combined with the trigger requirements on the photon p_T.

To address this problem, we apply a trigger-level photon isolation criterion where the photons are not counted toward each other's isolation variable, as will be explained in Eq. (5). We do not modify the other reconstruction and identification criteria. As such, we still require that the minimal separation between a diphoton pair is such that two L1 trigger ECAL towers do not overlap, which roughly corresponds to $\Delta R_{\gamma\gamma} \gtrsim 0.1$. Of course, a more detailed study by the experimental collaborations is required to properly assess the behavior of the reconstruction and identification algorithms in combination with the modified isolation defined here. For instance, the use of the ATLAS R_η and R_ϕ variables as defined now may not be compatible with our modified isolation.

For our simulations, we rely on the DELPHES3 framework [35], which propagates the particles of the event through a simplified detector volume and assigns a fraction of the particle's energy to a particular collection of calorimeter cells. The candidate photons after this process are identified with those ECAL cells for which there is no significant energy deposit in the corresponding tower in the hadronic calorimeter and for which there is no track pointing toward the ECAL cell.

To determine which photon candidates qualify as reconstructed photons, DELPHES3 uses the same isolation criterion as in Eq. (4). Candidates that fail isolation are included among the inputs for the jet clustering algorithm [36]. Since we wish to preserve photon candidates that have another hard photon candidate within their isolation cone, we modify the standard isolation procedure by explicitly subtracting the p_T of the hardest photon candidate in the isolation cone from the sum. Concretely, we modified the DELPHES3 isolation module to define a new isolation variable $r_{\gamma\gamma}$, which is more suitable for photons that are expected to come in pairs,

$$r_{\gamma\gamma} \equiv \frac{1}{p_{T_\gamma}}\left(\sum_{\substack{\Delta R<0.2 \\ p_{T_i}>0.5\ \text{GeV}}} p_{T_i} - p_{T_{\gamma,1}}\right), \tag{5}$$

where p_{T_γ} is the transverse momentum of the primary photon candidate and $p_{T_{\gamma,1}}$ is the transverse momentum of the leading photon candidate in the isolation cone, excluding the primary photon candidate itself. The sum is defined as in Eq. (4). A photon candidate is then classified as an isolated photon if it satisfies $r_{\gamma\gamma} < 0.05$. A similar modification of the isolation variable was already used in an offline analysis [37].

The isolation criterion described in Eq. (4) relies on tracking information. However, tracking is a computing-intensive step, and it may therefore not be available at high enough rates in an analysis that includes photons reconstructed directly at the trigger level, due to the limited resources available in the trigger farm. In the context of diphoton resonances, tracking primarily improves the efficacy of pileup subtraction methods and isolation cuts, leading to a lower number of fake photons.[4] We therefore expect larger backgrounds for an analysis which cannot rely on tracking. To estimate how much a search would degrade in the absence of tracking when computing the isolation, we also repeat the procedure outlined above while setting the tracking efficiency for all charged particles to zero.

B. Analysis cuts

As the implementation of our signal and the background simulations are informed by our fiducial cuts, we first discuss those here. We must make a number of assumptions, in particular in relation to which p_T cuts may be viable for a trigger-level analysis. For guidance, we look at the ATLAS L1 thresholds for two electromagnetic objects [13], which were set to $p_T > 10$ GeV and $p_T > 15$ GeV in

[3] In other studies, such as Ref. [34], isolation is defined as an upper bound on the total transverse energy inside the isolation cone, $E_T^{\text{MAX}} = r_\gamma p_{T_\gamma}$.

[4] Tracking is also important to extract data driven estimates for the photon fake rate [38]. If tracking is not available for the main trigger stream, a separate prescaled stream may be needed to obtain the required control samples.

the 2015 and 2016 datasets, respectively (L1_2EM10VH and L1_2EM15VH). As of the 2017 dataset, a loose isolation requirement was added (L1_2EM15VHI). We will here optimistically consider $p_T > 15$ GeV but assume no isolation requirement at L1, effectively mimicking the cuts of the L1_2EM15VH trigger item. We assume that the modified isolation criteria described in the previous subsection would be implemented online after the L1 trigger. The rate for the L1_2EM15VH selection was measured to be roughly 5 kHz [13], which is similar to what was recorded for an existing trigger-level analysis [19].

We estimate the expected sensitivity for the following three cases:

(I) Trigger level with tracking information, modified isolation [Eq. (5)], and requiring $p_{T\gamma} > 15$ GeV for each photon, denoted by "Tracking."
(II) Same as above but without tracking, denoted by "No tracking."
(III) As a stand-in for an off-line analysis with the existing ATLAS trigger, we require $p_{T\gamma} > 25$ GeV for each photon [13].[5] For this case we assume standard isolation [Eq. (4)] with tracking information. Going forward, we will refer to this as the "Off-line" case.

Each event is required to contain two isolated photons with pseudorapidity cut $|\eta_\gamma| < 2.5$, each above the p_T threshold as mentioned above.

C. Signal

The ALP benchmark model was implemented in FeynRules [39], to facilitate the generation of a Monte Carlo sample with MadGraph 5 [40] at leading order. The showering and hadronization was done with PYTHIA8 [41,42]. We assumed a KSVZ ALP with vectorlike matter in a set of 5-$\bar{5}$ multiplets, compatible with $SU(5)$ unification. This choice corresponds to $E/N = 8/3$ in relation to Eq. (2) and is only relevant for the $a \to \gamma\gamma$ branching ratio. The $a \to \gamma\gamma$ branching ratio is sensitive to Next-to-next-to Leading Order (NNLO) corrections to gluon width [43] and has therefore a mild dependence on m_a. For 5 GeV $< m_a <$ 100 GeV, it is well approximated by the phenomenological formula

$$\mathrm{Br}(a \to \gamma\gamma) \approx 7.8 \times 10^{-4} \times \left(\frac{E/N}{8/3}\right)^2 \times \left[1 + 0.67 \log\left(\frac{m_a}{10\ \mathrm{GeV}}\right) + 0.12 \log^2\left(\frac{m_a}{10\ \mathrm{GeV}}\right)\right]. \quad (6)$$

[5]We choose $p_{T\gamma} > 25$ GeV, as compared to the $p_{T\gamma} > 22$ GeV in the ATLAS HLT trigger strategy [13], to approximately take into account the turn-on of the trigger.

The detector response was simulated with DELPHES3, with the standard ATLAS card, which we modified to more accurately represent the angular resolution of the ATLAS ECAL (see the Appendix A 2). We assumed an average of 40 pileup collisions per event and employed basic pileup mitigation strategies, as built into DELPHES3. Concretely, all tracks with a longitudinal separation of more than 0.1 mm from the primary vertex are removed and a pile-up density correction is applied to the photon isolation variable [35].

At low m_a, the most relevant process is an ALP produced in gluon fusion together with an associated jet. The latter is needed to ensure that the photons from the ALP decay can pass the trigger p_T thresholds. The a + jet production process was modeled at the level of the hard matrix element, using MadGraph 5 with a minimal cut on the jet p_T: $p_{T,j} > 30$ GeV as in [44]. The fiducial rate does not increase if a fully inclusive (matched) sample is used instead, as long as m_a is less than twice the photon p_T threshold. We also verified that matching up to two jets does not significantly change the relevant differential distributions. Following [45], we then compute the $a + 1j$ cross section at the NLO with MadGraph@NLO [46,47]. The fiducial cross sections subject to the cuts specified in Sec. II B are shown in Fig. 1 (solid lines).

For m_a larger than twice the photon p_T threshold no recoiling jet is needed to satisfy the trigger requirements. An inclusive signal sample without demanding an additional jet is therefore more suitable, as indicated by the dot-dashed lines in Fig. 1. The sample was rescaled to match the NNLO cross section in [7], which was computed with the ggHiggs v3.5 package [48–50]. A smooth interpolation

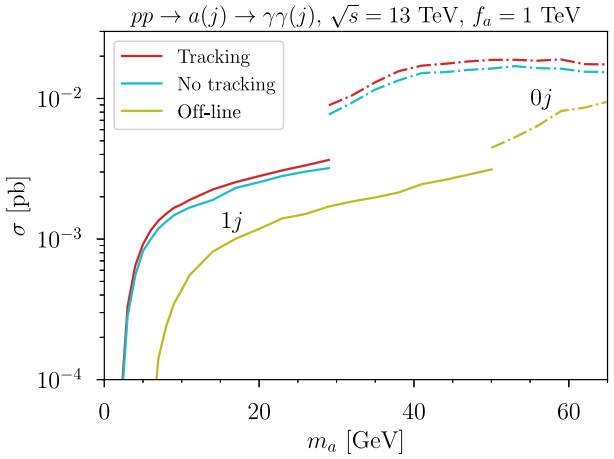

FIG. 1. Fiducial cross sections for the a + jet (solid lines) and a (dot-dashed lines) inclusive cross sections, for the selections described in the text. We switch between the two samples when m_a is twice the photon p_T threshold. The discontinuity can be resolved as a smooth step transition by matching the samples at next-to-leading order, which we did not attempt here. For $m_a \lesssim 10$ GeV, the photons are too collimated to pass the isolation criterion for the "off-line" selection.

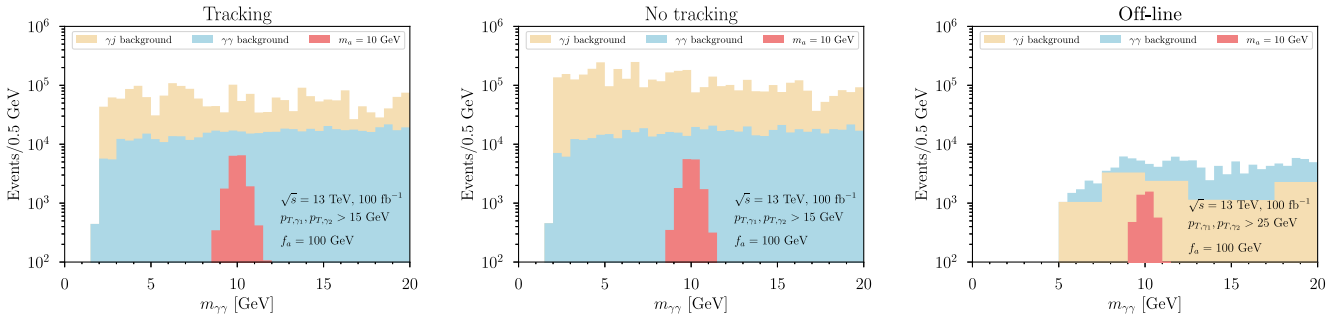

FIG. 2. Stacked invariant mass distribution for the signal and the background for the three scenarios described in Sec. II B. A representative signal resonance is shown for $m_a = 10$ GeV with the total number of events corresponding to $f_a = 100$ GeV. For the same plot with a larger mass range, see Fig. 6 in the Appendix A. The "γj background" in the panel labeled "Off-line" uses larger bin widths due to more limited simulation statistics. See Sec. II D for details.

between both regimes can be achieved by matching the zero jet and one jet sample at NLO, which we did not attempt here.

The main take-away from Fig. 1 is that the off-line selection starts to lose acceptance for $m_a \lesssim 15$ GeV, whereas both the "tracking" and "no tracking" selections maintain reasonable acceptance as low as $m_a \approx 5$ GeV. We also see from Fig. 1 that the signal efficiency does not degrade much in the absence of tracking, as expected.

D. Background

One of the main backgrounds is an event with pairs of relatively collimated photons, which recoil against a jet with moderate p_T. In the SM, photons arise from the hard scattering matrix element, from emission of a fermion line in the parton shower, or in hadron decays, notably π^0 decays. Isolation criteria are designed to single out the former contribution, which give us the most insight in the short distance physics. The shower and hadronization contributions are captured in the perturbative and nonperturbative components of the photon fragmentation functions (e.g., see [51]), and are ideally counted toward their corresponding jet. This is achieved to a very high degree with the use of suitable isolation criteria. However, even very small fake probabilities (probability that an isolated photon arises through showering or hadronization) can be important due to the comparatively high dijet and photon plus jet cross sections. This means that backgrounds from jets faking photons cannot *a priori* be neglected.

We opt to simulate all components with a combination of MadGraph 5 and PYTHIA8, as described below. We validate our simulation against the NNLO calculation by Catani *et al.* [34], as well as experimental data [38] (see also [52] for a recent update at $\sqrt{s} = 13$ TeV). This is described in the Appendix A 1.

Photons originating from the hard collision were modeled with an inclusive $\gamma\gamma$ sample, matched up to two jets. The events were then passed through DELPHES3 as described in the previous sections. The collinear photon-quark singularity was regulated with an angular cut of $\Delta R_{q\gamma} > 0.05$ at the generator level. The overall cross section was rescaled to the NLO inclusive $\gamma\gamma$ cross section as computed with MadGraph@NLO. The resulting invariant mass distributions are shown in Fig. 2, labeled as "$\gamma\gamma$ background," for our three different benchmark strategies, as described in Sec. II B.

A second background component arises from photon plus jet production, where the jet fakes an additional photon. The most dominant contribution is the $qg \to q\gamma$ process, whenever an additional hard photon is emitted during the fragmentation of the outgoing quark. Fragmentation photons may be identified as isolated if they are radiated at a wide angle from their corresponding quark line. Such fake photons are relatively rare; however, it is essential to estimate their rate, because the γ plus jet cross section is much larger than the $\gamma\gamma$ cross section. This is partially due to the additional $\alpha_{\rm em}$ suppression of the latter, but also due to luminosity functions: At leading order, the $\gamma\gamma$ process requires a $q\bar{q}$ initial state, whose luminosity function is much smaller than the qg luminosity.

To characterize this contribution it is helpful to look at the kinematics of the $q \to q\gamma$ splitting function: In the collinear limit the angle between the photon and the outgoing quark is given by

$$\theta_{\gamma q} \approx \frac{Q}{E_q}\sqrt{\frac{1}{z(1-z)}}, \tag{7}$$

where Q and E_q are the virtuality and energy of the incoming quark, respectively, while $z = E_\gamma/E_q$ is the energy fraction carried away by the photon. The PYTHIA8 parton shower is ordered according to virtuality, which through Eq. (7) implies that wide angle radiation occurs primarily in the initial stages of the parton shower. However, emissions with $\Delta R_{q\gamma} > 0.05$ populate the same region of phase space as the hard process simulation described in the previous paragraph. We therefore veto them to avoid double counting. The effect of this veto is shown in Fig. 3, where we show

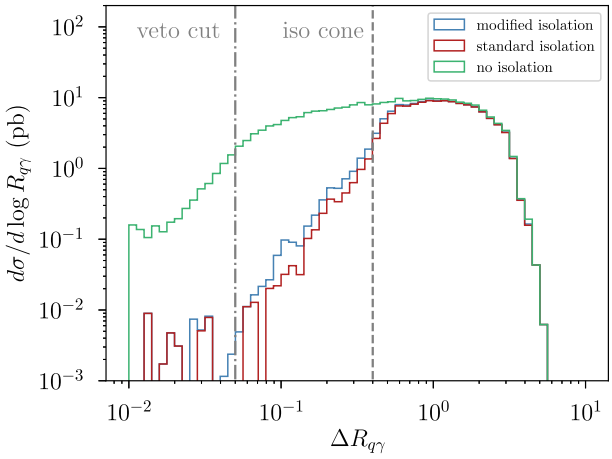

FIG. 3. Angular separation of hard ($p_T > 15$ GeV) fragmentation photons from their sister quark in a leading order $gq \to q\gamma$ sample, simulated with PYTHIA8 and rescaled to match the NLO, inclusive photon plus jet cross section. Only photons from the perturbative parton shower are included. The veto cut and isolation cone are indicated for reference. (See text for details.)

the $\Delta R_{q\gamma}$ distribution for fragmentation photons, subjected normal isolation, modified isolation, and no isolation. Photons from the nonperturbative parts of the fragmentation process, e.g., hadron decays, are accounted for separately and not included in this plot. As expected, all three distributions align for $\Delta R_{q\gamma}$ larger than the isolation cone, while the two isolated distributions drop rapidly for $\Delta R_{q\gamma} < 0.4$. Since all photons with $\Delta R_{q\gamma} > 0.05$ are vetoed, we find that this contribution to the isolated photon distribution can be neglected.

Fake photons can also be produced in the nonperturbative part of the quark fragmentation function, primarily from π^0 decays. To estimate this part, we simulate a $pp \to \gamma j(j)$ sample with MadGraph 5, matched up to two jets. The cross section was rescaled to match the inclusive, NLO γj cross section as computed with MadGraph@NLO. The events were subsequently showered and hadronized with PYTHIA8. In this step all electromagnetic radiation within the jet was vetoed, as we separately accounted for this component in the previous paragraph. The resulting distribution is shown in Fig. 2, labeled as "γj background." It is comparable to or larger than the true photon background for our modified isolation criterion, and smaller than the true photon background for standard isolation. The fake photon background can be further reduced by imposing tighter isolation cuts; we comment on this in the Appendix A 1.

In our discussion above we have always assumed that at least one photon originated from the hard matrix element. *A priori* it is possible that both the photons are supplied by the fragmentation of quarks in multijet processes. This was found to be negligible by Catani *et al.* [34], although in a slightly different kinematical regime. We verified that this conclusion extends to our kinematical cuts, as long as the fragmentation probability for both photons are assumed to be uncorrelated. Whether this assumption holds for the low $m_{\gamma\gamma}$ regime of interest must be verified with data.

Finally, in a realistic experimental environment fake photons also arise from electrons for which the track was not reconstructed or from hadrons which stop in the first part of the ECAL. Both are difficult to model with a theorist-level simulation, and they are neglected in this study.

E. Results

As expected, Fig. 2 makes it clear that both signal and background go up substantially for a trigger-level analysis, as compared to the off-line scenario. A trigger-level approach is most clearly useful for $m_a \lesssim 10$ GeV, as here the signal efficiency of the off-line selection dies off rapidly, as seen from Fig. 1. For $m_a \gtrsim 10$ GeV the relative power of both approaches will depend on a number of important experimental subtleties, such as the precise thresholds and turn-on of the various triggers as well as various sources of systematic uncertainties.

To produce an estimate for the reach of both strategies, a number of additional assumptions are therefore needed. In particular, we will assume that the dominant uncertainty on the background is statistical rather than systematic. This is a reasonable assumption for a bump hunt over a smooth background; however, any features in the invariant mass distribution of the diphoton background below the trigger threshold might introduce further complications. We also do not attempt a full fledged likelihood analysis and instead simply compare the number of signal and background events in bins with width set to the expected invariant mass resolution (see the Appendix A 2). This can be justified because of the large values of expected signal and background events. For example, as can be seen in Fig. 2, the number of background and signal events (for $m_a < 100$ GeV and $f_a < 1$ TeV) in a single invariant mass bin containing the signal is $\sim 10^4$ and $\sim 10^2$, respectively. One can first construct a test statistic [53]

$$t = -2 \ln\left(\frac{L(n|b)}{L(n|s+b)}\right) \tag{8}$$

to obtain the significance. Here $L(n|m)$ denotes the probability of observing n events originating from a Poisson distribution with mean m, and s and b denote the expected number of signal and background events, respectively. By assuming $n = s + b \gg 1$ and expanding in $s \ll b$, t can be written as

$$t = 2\left(n \ln\left(\frac{s+b}{b}\right) - s\right) \approx \frac{s^2}{b}. \tag{9}$$

In the limit of a large number of samples, t is χ^2-distributed with 1 degree of freedom under the null hypothesis.

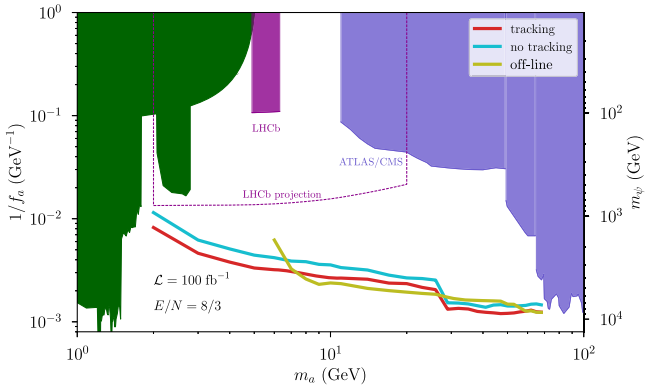

FIG. 4. Estimated 2σ exclusion reach for 100 fb^{-1}, along with existing constraints from ATLAS [3,7], CMS [7,54–56], LHCb [9], and exotic B to a decays [57]. The smooth green contour represents the bound on the inclusive $B \to sa$, while the remaining green contours represent several exclusive channels. These bounds are somewhat model dependent; see [57] for details. We also show an aggressive projection for LHCb, assuming 300 fb^{-1} [9]. The right-hand y axis indicates the mass of the colored fermions (ψ) for the example UV completion described in Sec. III A.

Therefore, for each value of m_a, we consider a single invariant mass bin which contains 90% of the signal and determine f_a such that $\sqrt{t} = 2$, to obtain a 2σ confidence limit.

With this simplified recipe we find similar sensitivity for all approaches for $m_a \gtrsim 10$ GeV but see that a trigger-level analysis can potentially uniquely probe the $m_a \lesssim 10$ GeV region, as shown in Fig. 4. While the overall normalization of our estimate is sensitive to all the subtleties listed above, we emphasize that this feature should be robust, as it is driven by the behavior of the signal efficiency in Fig. 1. All strategies would, however, be a significant improvement over the existing bounds.

Finally, we comment on the possibility to observe the Standard Model $\eta_b \to \gamma\gamma$ decay.[6] The η_b is the lightest state in the bottomonium system with mass $m_{\eta_b} \approx 9.4$ GeV. For a slightly boosted η_b with $p_T \approx 20$ GeV, its production cross section at the LHC is expected to be between 1 and 10 nb [58]. The η_b branching ratio to photons has not been measured yet but is thought to be roughly 5×10^{-5} [59]. This rate might be high enough to be observable in a trigger-level analysis, depending on the efficiency of our isolation requirements. We leave this interesting question for a future study.

III. THEORY MOTIVATION

We first comment on the generic features of the UV completions of the effective theory in Eq. (2). Then we briefly review some of the top-down motivation for low mass diphoton resonances.

A. Example UV completion

For completeness, we should investigate whether there exist healthy UV completions of the effective action in Eq. (2), given the values of f_a that are probed in Fig. 4 and the complementary searches at the LHC probing new physics at this scale.

As a minimal example we take the model by KSVZ [26,27], where the ALP is the phase of a complex scalar singlet Φ. The model also includes a set of heavy, vector-like colored and electrically charged fermions that are charged under the Peccei Quinn (PQ) symmetry. This specifies the action

$$\mathcal{L}_{\mathrm{KSVZ}} \supset g_* \Phi \bar{\psi} \psi, \qquad \Phi = \frac{v_a + \varphi}{\sqrt{2}} e^{ia/v_a}. \tag{10}$$

In our conventions, the vacuum expectation value of Φ (v_a) is related to the ALP decay constant by $v_a = 2Nf_a$. The anomaly coefficients N and E are specified by the multiplicity and the representation of the fermions. Taking the ψ to fill out a 5-$\bar{5}$ representation of $SU(5)$, we have $N = N_\psi/2$ and $E = 4/3N_\psi$ with N_ψ the number of ψ flavors. Finally, the g_* coupling determines the mass of the fermions $m_\psi = N_\psi g_* f_a/\sqrt{2}$. Taking, for example, $N_\psi = 5$ and $g_* = 3$ ensures that the ψ states can be heavier than the reach of the many LHC searches for new colored particles. (See right-hand y axis in Fig. 4.) This example indicates that the discovery of a low mass diphoton resonance would likely be the hallmark of an elaborate composite sector near the TeV scale.

Alternatively, it is not difficult to increase the branching ratio to photons in Eq. (6) by choosing exotic representations that yield a larger E/N ratio. One can also consider low values of f_a but lift the masses of the colored particles in the multiplet by means of an additional mass term in Eq. (10). This does not reduce the overall signal rate, since the reduced production cross section is exactly compensated by an increase in the branching ratio to photons, as long as the decay width to photons remains small compared to the decay width to gluons. This remains true as long as the mass of the colored particles satisfies $M \lesssim (\alpha_s/\alpha_{\mathrm{em}}) g_* f_a$. For the parameters in Fig. 4, this easily pushes the colored states outside the energy reach of the LHC.

B. More complete models

The ALP benchmark model defined in Eq. (2) can be the hallmark of more complete models addressing the hierarchy, the strong CP, or the dark matter problems. These three motivations were reviewed comprehensively in [8,45], and we therefore only summarize the main points.

[6]We thank Michelangelo Mangano for pointing this out to us.

The main virtue of very low scale supersymmetry breaking is that a viable gravitino abundance can be achieved without the need for a nonstandard cosmological history [60]. In such scenarios, one also expects its associated $U(1)_R$ symmetry to be broken at a low scale, leading to a parametrically light pseudo Goldstone boson, the R-axion [61–63]. The R-axion couples to the gauginos, which induce the operators in Eq. (2) in the low energy theory.

Axions famously provide a solution to the strong CP problem by promoting the QCD θ angle to a dynamical field [64,65]. This mechanism, however, comes with a *PQ quality problem*, which states that any sources of PQ breaking other than the QCD dynamics must be very small. The suppression of potentially dangerous Planck-suppressed operators can be achieved through extending the UV structure of the model [66–69] or by choosing the axion decay constant sufficiently low. The latter option is already excluded in its most minimal form [70–75], though the constraints can be evaded in models which raise the axion mass without affecting the alignment of its potential with the CP preserving vacuum [44,76–88]. The collider phenomenology of such models maps directly onto the simplified model in Eq. (2) if the axion is of the KSVZ class [26,27]. In the case of Dine-Fischler-Srednicki-Zhitnitsky axions [89,90] coupling to SM leptons is also present, even though the best sensitivity at these masses might still come from the diphoton final state [91–95].

Finally, the ALP in Eq. (2) could couple to the dark matter particle and be responsible for its freeze-out into SM states. A natural realization of this model with $\mathcal{O}(1)$ coupling constants predicts $f_a \sim$ TeV [8].

In the above, we have always assumed that the diphoton resonance is directly produced in the partonic collision, typically associated with a gluon. This need not be the case, however, and such a new state could also be produced in the decay of a heavier state, such as the Higgs. Specifically, if we consider a scalar field S with couplings

$$\mathcal{L} \supset \frac{\lambda_S}{2} S^2 H^\dagger H + \frac{\alpha_{\rm em}}{4\pi}\frac{1}{\Lambda_{\gamma\gamma}} S F^{\mu\nu} F_{\mu\nu} + \frac{\alpha_s}{4\pi}\frac{1}{\Lambda_{gg}} S G^{\mu\nu} G_{\mu\nu}, \tag{11}$$

we can neglect possible couplings to the WW, ZZ, and $Z\gamma$ operators as long as $m_S \ll m_W$. This simplified model reproduces exotic Higgs decays of the form $h \to 4\gamma$, $h \to 2\gamma 2g$, and $h \to 4g$, whose relative importance depends on $\Lambda_{\gamma\gamma}/\Lambda_{gg}$. The four gluon channel is likely hopeless, but ATLAS and CMS have already performed searches for the $h \to 2\gamma 2g$ [96] and $h \to 4\gamma$ [37,97] channels. These searches place the following constraints for $m_S \approx 20$ GeV:

$$\begin{aligned} &\mathrm{Br}(h \to SS) \times \mathrm{Br}(S \to \gamma\gamma) \times \mathrm{Br}(S \to gg) \lesssim 0.1, \\ &\mathrm{Br}(h \to SS) \times \mathrm{Br}(S \to \gamma\gamma)^2 \lesssim 2 \times 10^{-4}. \end{aligned} \tag{12}$$

In the context of the simplified model in Eq. (11), the $h \to 4\gamma$ channel is therefore always more powerful, and no trigger-level analysis in the diphoton topology is needed.

The reason is the $\mathrm{Br}(S \to gg) + \mathrm{Br}(S \to \gamma\gamma) \leq 1$ constraint, which is, however, easily evaded by adding an additional scalar

$$\mathcal{L} \supset \lambda_S S_1 S_2 H^\dagger H + \frac{\alpha}{4\pi}\frac{1}{\Lambda_{\gamma\gamma}} S_1 F^{\mu\nu} F_{\mu\nu} + \frac{\alpha_s}{4\pi}\frac{1}{\Lambda_{gg}} S_2 \mathrm{Tr}[G^{\mu\nu} G_{\mu\nu}]. \tag{13}$$

This model only induces the $h \to S_1 S_2 \to gg\gamma\gamma$ topology. While we are not aware of a direct, top-down theory motivation for this specific topology, Eq. (13) is a perfectly plausible and fairly economical extension of the SM Higgs sector. As such, it is important that it be covered as well as possible, since nature may very well be more clever (or devious) than modern day model builders. The ATLAS search [96] in particular makes use of the diphoton trigger path discussed in the Introduction, which limits their reach to $m_{\gamma\gamma} \geq 20$ GeV. If a trigger-level analysis for lower invariant masses proves to be feasible, the $h \to 2\gamma 2g$ topology should therefore benefit as well.

In addition to the above simple, two step cascade models more complicated dynamics can be realized in hidden valley models [98,99], which can yield one or more low mass diphoton resonances [21,100].

IV. DISCUSSION

Low mass diphoton resonances could hide in plain sight, being rejected by the experimental trigger selections. We showed that the uncharted mass window below the current reach of diphoton searches and above the existing constraints from flavor physics can be probed by ATLAS and CMS. While probing masses down to roughly 10 GeV is probably feasible with the current triggers, we find that a trigger-level analysis with modified isolation requirements may cover the whole mass region down to masses of order 1 GeV.

We considered a simple, modified isolation requirement which was already used by ATLAS in a prior analysis [37] and argue that it could significantly increase the experimental acceptance for light resonances if implemented in a trigger-level analysis. Though the increased signal acceptance comes at the cost of a larger background rate, we nevertheless find otherwise inaccessible parameter space. That said, there are likely major technical challenges associated with a high rate, trigger-level analysis which can only be assessed to a satisfactory level by the collaborations themselves. The size of the fake photon background in particular should be verified in data before a definitive claim can be made.

We remark that our analysis strategy is deliberately simplistic, and perhaps even naive, given the expected computational constraints on a trigger-level analysis. A number of more sophisticated analysis strategies have already been

proposed for scenarios where the trigger is not the main challenge, most notably the topology $h \to 2a \to 4\gamma$ [101,102] but also cascade decays from heavier, exotic resonances [103]. Jet substructure and machine learning techniques applied to the energy depositions in the photon isolation cone appear to be promising [21–25]. Should it be possible to implement some of these techniques in a trigger-level analysis, they should yield substantially better sensitivity than our comparatively simplistic isolation requirement.

A more ambitious possibility would be to push to even lower $m_{\gamma\gamma}$, to the extent that both photons merge and one must rely on the single photon L1 seed. This effectively involves modifying the photon identification requirements themselves. This was already accomplished in an (off-line) ATLAS [104] search for a pair of photon jets. However, the payoff could be very substantial if this can also be done for a single photon jet as a trigger-level analysis, given that the single photon trigger jumps from 22 GeV to 120 GeV between L1 and HLT [13]. While the technical challenges should not be underestimated, we conclude that trigger-level searches for low mass diphoton resonances and photon jets could open qualitatively new parameter space for the discovery of low mass diphoton resonances.

ACKNOWLEDGMENTS

We thank Liron Barak, Matt Low, Alberto Mariotti, Marco Montella, José Ocariz, Michele Papucci, Luis Pascual, Dean Robinson, Filippo Sala, and Kohsaku Tobioka for many discussions, comments on the draft, and earlier work on similar topics. We thank Christian Bauer, Zoltan Ligeti, Michelangelo Mangano, and Nicholas Rodd in particular for discussions on the details of the photon fragmentation functions. We are especially grateful to Caterina Doglioni for the discussions that initiated this project and for detailed comments on the manuscript. S. Kn. was supported by the Office of High Energy Physics of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. S. Ku. was supported in part by the National Science Foundation (NSF) Grant No. PHY-1915314 and the U.S. Department of Energy Contract No. DE-AC02-05CH11231. S. Ku. thanks Aspen Center of Physics, supported by NSF Grant No. PHY-1607611, for its hospitality while this work was in progress. All authors thank the Galileo Galilei Institute for Theoretical Physics for its hospitality and the INFN for partial support during the initial and final phases of this work.

APPENDIX A: SIMULATION DETAILS

1. Background simulation

To validate our simulation framework, we compare with the NNLO calculations performed by Catani *et al.* [34,105], which have shown good agreement with ATLAS data at 7 TeV [106]. This comparison is shown in Fig. 5. Note that these results are obtained with a fixed isolation threshold $E_T^{\mathrm{MAX}} = 2$ GeV, rather than with the sliding threshold we

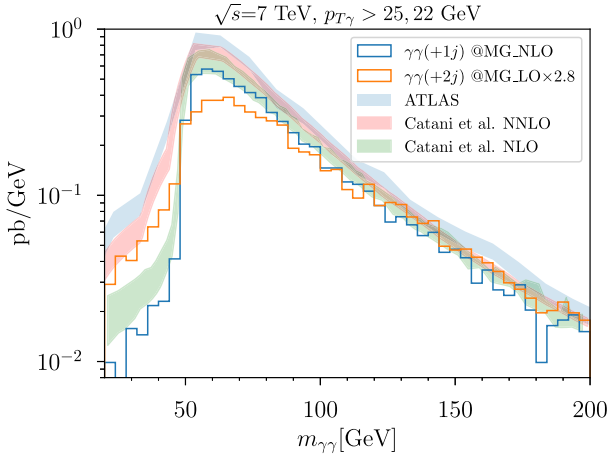

FIG. 5. Diphoton invariant mass plot from Catani *et al.* [34], overlaid with our simplified procedure (orange line, rescaled by an NLO K-factor of 2.8) and MadGraph@NLO (blue line). See text for details.

used in Eq. (4) and Eq. (5). The leading and subleading photon must satisfy $p_T > 25$ GeV and $p_T > 22$ GeV, respectively, as well as a rapidity cut $|y_\gamma| < 2.37$. At leading order, the process is just the $q\bar{q} \to \gamma\gamma$, for which there is no phase space for $m_{\gamma\gamma} < 50$ GeV. At NLO, new processes contribute such as $qg \to \gamma\gamma q$ and $qg \to \gamma q$, where a second photon is supplied by the quark fragmentation function. This contribution is enhanced by the parton distribution functions (PDFs), despite being formally higher order than the $q\bar{q} \to \gamma\gamma$ process. The NLO calculation also contains virtual corrections to $q\bar{q} \to \gamma\gamma$. The additional jet at NLO opens up the phase space for $m_{\gamma\gamma} < 50$ GeV. Finally, the NNLO calculation allows for an additional soft jet and includes virtual corrections via the box diagrams. The latter are important because they include the $gg \to \gamma\gamma$ process, which is further PDF enhanced compared to the LO and NLO contributions. The blue curve, diphoton events matched up to one jet, in Fig. 5 is obtained with MadGraph@NLO and is in good agreement with the Catani *et al.* NLO calculation.

For computational reasons, we used a leading order matched sample, rescaled to the inclusive NLO cross section, as described in Sec. II D. To estimate the error this simplification introduces, we apply it to the kinematical selections used by Catani *et al.*, as is shown by the orange curve in Fig. 5. Here we neglected photons from the nonperturbative part of the fragmentation, which is a subleading component for standard isolation (see the rightmost panel of Fig. 2). Our simplified procedure somewhat underestimates the peak but is in $\mathcal{O}(1)$ agreement with the low $m_{\gamma\gamma}$ tail. While this may seem far from ideal, it worth noting that the sensitivity projection in Fig. 4 only scales as the fourth root of the background. Our qualitative conclusions are therefore not very sensitive to this level of background mismodeling. For completeness

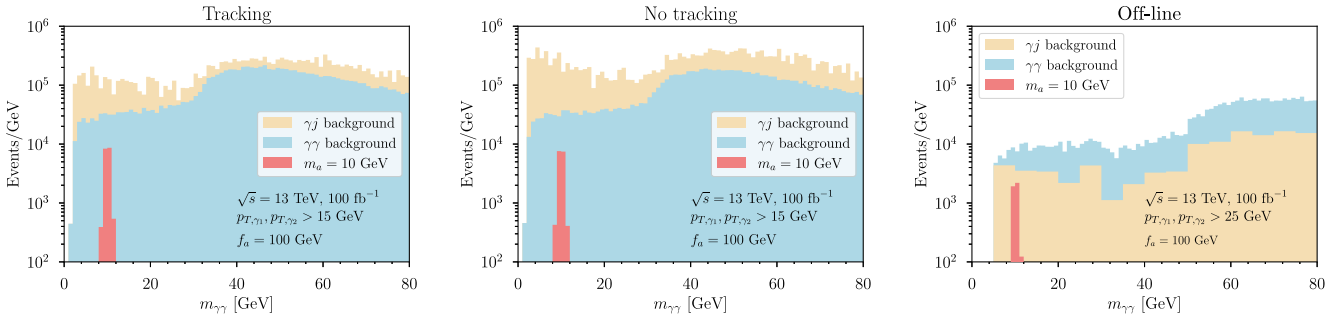

FIG. 6. Stacked invariant mass distribution for the signal and the background for the three scenarios described in Sec. II B, for a larger mass range. A representative signal resonance is shown for $m_a = 10$ GeV with the total number of events corresponding to $f_a = 100$ GeV.

we extend the invariant mass distributions in Fig. 2 to the full mass range in Fig. 6. We clearly see the feature around twice the p_T threshold, consistent with the discussion in the previous paragraph.

Finally, we also comment on the dependence of the hadronic contribution to the photon background on the choice of isolation criterion. For this purpose, we simulated a large $q\bar{q}$ sample with PYTHIA8 and quantify the probability of a jet producing an isolated photon from a hadronic decay. To cleanly separate this contribution from the perturbative part of the fragmentation function we veto all electromagnetic emissions during the parton shower. (These contributions were discussed separately in Sec. II D, specifically in Fig. 3.) The result is shown in Fig. 7.

We find that with the modified isolation criterion, the fake rate increases by a bit more than an order of magnitude. Using tight isolation does reduce the background somewhat, but does not fully compensate for the aforementioned increase. That said, given the rather special corner of phase space that is being probed here, a data-driven approach will be essential to validate the Monte Carlo predictions.

2. Mass resolution

Given the limited number of handles at our disposal, the main background discrimination comes from hunting for a bump on top of a large background. The resolution on the invariant mass is therefore very important. This is handled automatically by DELPHES3, as shown in Fig. 8. Given its importance, we seek to verify the build in DELPHES3 parametrization explicitly. Concretely, the energy resolution of the ATLAS ECAL is parametrized by [107]

$$\frac{\delta E_\gamma}{E_\gamma} = 10\% \left(\frac{\text{GeV}}{E_\gamma}\right)^{1/2} \oplus 0.7\%, \tag{A1}$$

where the $\oplus$ notation indicates that the errors are added in quadrature. The mass resolution is given by

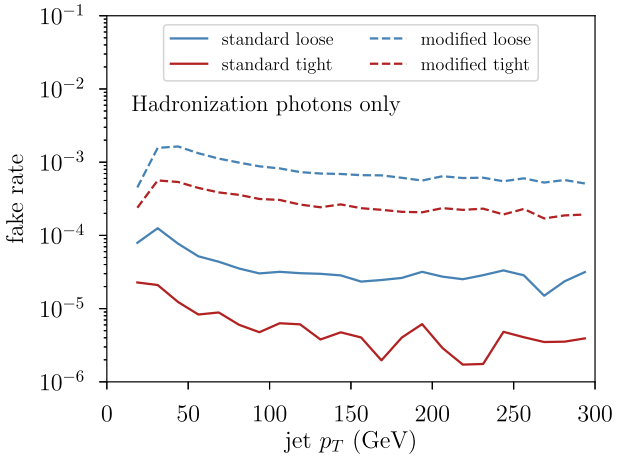

FIG. 7. Probability of a quark jet fragmenting into an isolated photon with $p_T > 15$ GeV, for standard and modified isolation, with loose and tight selections (see Sec. II A). Only photons from hadronic decays are considered. For photons from the perturbative part of the parton shower we refer to Fig. 3.

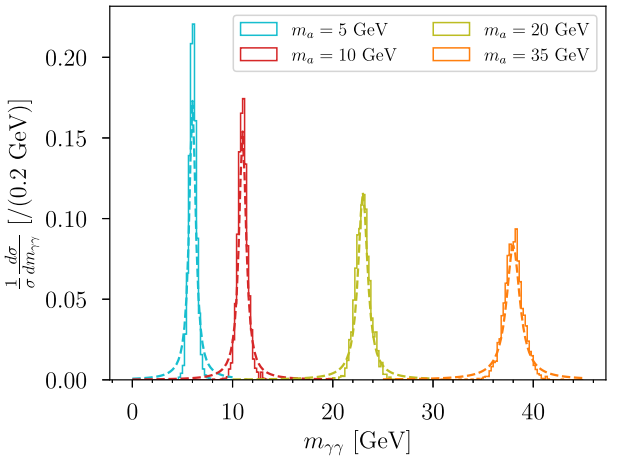

FIG. 8. Mass resolution for a diphoton resonance, as reproduced by DELPHES3 with the standard ATLAS configuration card (solid line). To determine the width of the overlaid Lorentzian line shape we take a p_T threshold of 15 GeV (dashed line), as described in the text.

$$\frac{\delta m_{\gamma\gamma}}{m_{\gamma\gamma}} = \frac{1}{2}\left(\frac{\delta E_{\gamma,1}}{E_{\gamma,1}} \oplus \frac{\delta E_{\gamma,2}}{E_{\gamma,2}}\right) \oplus \frac{0.0247}{\Delta R_{\gamma\gamma}}, \tag{A2}$$

where for the angular term we have taken the off-line granularity of the calorimeter. This is slightly worse than the true off-line resolution of the ECAL but better than the L1 angular resolution. The assumption here is that the fine grained information for the ECAL tower can be written out for a small region of interest around the candidate resonance. To compare with Fig. 8, we estimate $\delta m_{\gamma\gamma}$ by taking $\eta \approx 0$ for both photons and saturating $E_{\gamma,1}$ and $E_{\gamma,2}$ to the p_T thresholds. This determines the energy resolution via Eq. (A1). We then solve for $\Delta R_{\gamma\gamma}$ in Eq. (1) by using the injected invariant mass. The resulting numbers are plugged into Eq. (A2) to obtain the estimated mass resolution. We overlaid a Lorentzian line shape with width $\delta m_{\gamma\gamma}$ to visualize the agreement with the full DELPHES3 simulation.

[1] S. Chatrchyan *et al.* (CMS Collaboration), Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC, Phys. Lett. B **716**, 30 (2012).

[2] G. Aad *et al.* (ATLAS Collaboration), Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC, Phys. Lett. B **716**, 1 (2012).

[3] G. Aad *et al.* (ATLAS Collaboration), Search for Scalar Diphoton Resonances in the Mass Range 65–600 GeV with the ATLAS Detector in pp Collision Data at $\sqrt{s} = 8$ TeV, Phys. Rev. Lett. **113**, 171801 (2014).

[4] A. M. Sirunyan *et al.* (CMS Collaboration), Search for physics beyond the standard model in high-mass diphoton events from proton-proton collisions at $\sqrt{s} = 13$ TeV, Phys. Rev. D **98**, 092001 (2018).

[5] A. M. Sirunyan *et al.* (CMS Collaboration), Search for a standard model-like Higgs boson in the mass range between 70 and 110 GeV in the diphoton final state in proton-proton collisions at $\sqrt{s} = 8$ and 13 TeV, Phys. Lett. B **793**, 320 (2019).

[6] G. Aad *et al.* (ATLAS Collaboration), Search for resonances decaying into photon pairs in 139 fb^{-1} of pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector, Phys. Lett. B **822**, 136651 (2021).

[7] A. Mariotti, D. Redigolo, F. Sala, and K. Tobioka, New LHC bound on low-mass diphoton resonances, Phys. Lett. B **783**, 13 (2018).

[8] X. Cid Vidal *et al.*, Report from Working Group 3: Beyond the standard model physics at the HL-LHC and HE-LHC, CERN Yellow Rep. Monogr. **7**, 585 (2019).

[9] X. Cid Vidal, A. Mariotti, D. Redigolo, F. Sala, and K. Tobioka, New axion searches at flavor factories, J. High Energy Phys. 01 (2019) 113; 06 (2020) 141(E).

[10] A. M. Sirunyan *et al.* (CMS Collaboration), Evidence for light-by-light scattering and searches for axion-like particles in ultraperipheral PbPb collisions at $\sqrt{s_{\mathrm{NN}}} = 5.02$ TeV, Phys. Lett. B **797**, 134826 (2019).

[11] G. Aad *et al.* (ATLAS Collaboration), Measurement of light-by-light scattering and search for axion-like particles with 2.2 nb^{-1} of Pb + Pb data with the ATLAS detector, J. High Energy Phys. 03 (2021) 243.

[12] S. Knapen, T. Lin, H. K. Lou, and T. Melia, Searching for Axionlike Particles with Ultraperipheral Heavy-Ion Collisions, Phys. Rev. Lett. **118**, 171801 (2017).

[13] G. Aad *et al.* (ATLAS Collaboration), Performance of electron and photon triggers in ATLAS during LHC Run 2, Eur. Phys. J. C **80**, 47 (2020).

[14] A. M. Sirunyan *et al.* (CMS Collaboration), Performance of the CMS Level-1 trigger in proton-proton collisions at $\sqrt{s} = 13$ TeV, J. Instrum. **15**, P10017 (2020).

[15] R. Aaij *et al.*, A comprehensive real-time analysis model at the LHCb experiment, J. Instrum. **14**, P04006 (2019).

[16] V. Khachatryan *et al.* (CMS Collaboration), Search for Narrow Resonances in Dijet Final States at $\sqrt{s} = 8$ TeV with the Novel CMS Technique of Data Scouting, Phys. Rev. Lett. **117**, 031802 (2016).

[17] M. Aaboud *et al.* (ATLAS Collaboration), Search for Low-Mass Dijet Resonances using Trigger-Level Jets with the ATLAS Detector in pp Collisions at $\sqrt{s} = 13$ TeV, Phys. Rev. Lett. **121**, 081801 (2018).

[18] A. M. Sirunyan *et al.* (CMS Collaboration), Search for a Narrow Resonance Lighter than 200 gev Decaying to a Pair of Muons in Proton-Proton Collisions at $\sqrt{s} = 13$ TeV, Phys. Rev. Lett. **124**, 131802 (2020).

[19] CMS Collaboration, Search for long-lived particles decaying into two muons in proton-proton collisions at $\sqrt{s} = 13$ TeV using data collected with high rate triggers, CERN Technical Report No. CMS-PAS-EXO-20-014, 2021.

[20] R. Aaij *et al.* (LHCb Collaboration), Search for $A' \to \mu^+\mu^-$ Decays, Phys. Rev. Lett. **124**, 041801 (2020).

[21] S. D. Ellis, T. S. Roy, and J. Scholtz, Phenomenology of photon-jets, Phys. Rev. D **87**, 014015 (2013).

[22] B. C. Allanach, D. Bhatia, and A. M. Iyer, Dissecting multi-photon resonances at the large hadron collider, Eur. Phys. J. C **77**, 595 (2017).

[23] B. Sheff, N. Steinberg, and J. D. Wells, Higgs boson decays into narrow diphoton jets and their search strategies at the Large Hadron Collider, Phys. Rev. D **104**, 036009 (2021).

[24] D. Wang, L. Wu, J. M. Yang, and M. Zhang, Photon-jet as a probe of axion-like particles at the LHC, Phys. Rev. D **104**, 095016 (2021).

[25] J. Ren, D. Wang, L. Wu, J. M. Yang, and M. Zhang, Detecting an axion-like particle with machine learning at the LHC, J. High Energy Phys. 11 (2021) 138.

[26] J. E. Kim, Weak Interaction Singlet and Strong *CP* Invariance, Phys. Rev. Lett. **43**, 103 (1979).

[27] M. A. Shifman, A. Vainshtein, and V. I. Zakharov, Can confinement ensure natural CP invariance of strong interactions?, Nucl. Phys. **B166**, 493 (1980).
[28] The ATLAS collaboration, Expected photon performance in the ATLAS experiment, CERN Technical Report No. ATL-PHYS-PUB-2011-007, 2011.
[29] G. Aad *et al.* (ATLAS Collaboration), Electron and photon performance measurements with the ATLAS detector using the 2015–2017 LHC proton-proton collision data, J. Instrum. **14**, P12006 (2019).
[30] G. Aad *et al.* (ATLAS Collaboration), Expected performance of the ATLAS experiment—detector, trigger and physics, arXiv:0901.0512.
[31] M. Aaboud *et al.* (ATLAS Collaboration), Measurement of the photon identification efficiencies with the ATLAS detector using LHC Run-1 data, Eur. Phys. J. C **76**, 666 (2016).
[32] The ATLAS collaboration, Expected electron performance in the ATLAS experiment, CERN Technical Report No. ATL-PHYS-PUB-2011-006, 2011.
[33] M. Aaboud *et al.* (ATLAS Collaboration), Measurement of the photon identification efficiencies with the ATLAS detector using LHC Run 2 data collected in 2015 and 2016, Eur. Phys. J. C **79**, 205 (2019).
[34] S. Catani, L. Cieri, D. de Florian, G. Ferrera, and M. Grazzini, Diphoton production at the LHC: A QCD study up to NNLO, J. High Energy Phys. 04 (2018) 142.
[35] J. de Favereau, C. Delaere, P. Demin, A. Giammanco, V. Lemaître, A. Mertens, and M. Selvaggi (DELPHES 3 Collaboration), DELPHES3, A modular framework for fast simulation of a generic collider experiment, J. High Energy Phys. 02 (2014) 057.
[36] M. Cacciari, G. P. Salam, and G. Soyez, FastJet user manual, Eur. Phys. J. C **72**, 1896 (2012).
[37] G. Aad *et al.* (ATLAS Collaboration), Search for new phenomena in events with at least three photons collected in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector, Eur. Phys. J. C **76**, 210 (2016).
[38] G. Aad *et al.* (ATLAS Collaboration), Measurement of the isolated di-photon cross-section in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector, Phys. Rev. D **85**, 012003 (2012).
[39] A. Alloul, N. D. Christensen, C. Degrande, C. Duhr, and B. Fuks, FeynRules 2.0—A complete toolbox for tree-level phenomenology, Comput. Phys. Commun. **185**, 2250 (2014).
[40] J. Alwall, M. Herquet, F. Maltoni, O. Mattelaer, and T. Stelzer, MadGraph 5: Going Beyond, J. High Energy Phys. 06 (2011) 128.
[41] T. Sjostrand, S. Mrenna, and P. Z. Skands, PYTHIA6.4 physics and manual, J. High Energy Phys. 05 (2006) 026.
[42] T. Sjostrand, S. Mrenna, and P. Z. Skands, A brief introduction to PYTHIA8.1, Comput. Phys. Commun. **178**, 852 (2008).
[43] K. G. Chetyrkin, B. A. Kniehl, M. Steinhauser, and W. A. Bardeen, Effective QCD interactions of CP odd Higgs bosons at three loops, Nucl. Phys. **B535**, 3 (1998).
[44] A. Hook, S. Kumar, Z. Liu, and R. Sundrum, High Quality QCD Axion and the LHC, Phys. Rev. Lett. **124**, 221801 (2020).
[45] Y. Gershtein, S. Knapen, and D. Redigolo, Probing naturally light singlets with a displaced vertex trigger, Phys. Lett. B **823**, 136758 (2021).
[46] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. S. Shao, T. Stelzer, P. Torrielli, and M. Zaro, The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations, J. High Energy Phys. 07 (2014) 079.
[47] P. Artoisenet *et al.*, A framework for Higgs characterisation, J. High Energy Phys. 11 (2013) 043.
[48] R. D. Ball, M. Bonvini, S. Forte, S. Marzani, and G. Ridolfi, Higgs production in gluon fusion beyond NNLO, Nucl. Phys. **B874**, 746 (2013).
[49] M. Bonvini, R. D. Ball, S. Forte, S. Marzani, and G. Ridolfi, Updated Higgs cross section at approximate N^3LO, J. Phys. G **41**, 095002 (2014).
[50] M. Bonvini, S. Marzani, C. Muselli, and L. Rottoli, On the Higgs cross section at $N^3LO + N^3LL$ and its uncertainty, J. High Energy Phys. 08 (2016) 105.
[51] L. Bourhis, M. Fontannaz, and J. P. Guillet, Quarks and gluon fragmentation functions into photons, Eur. Phys. J. C **2**, 529 (1998).
[52] G. Aad *et al.* (ATLAS Collaboration), Measurement of the production cross section of pairs of isolated photons in pp collisions at 13 TeV with the ATLAS detector, J. High Energy Phys. 11 (2021) 169.
[53] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, Asymptotic formulae for likelihood-based tests of new physics, Eur. Phys. J. C **71**, 1554 (2011); **73**, 2501(E) (2013).
[54] A. M. Sirunyan *et al.* (CMS Collaboration), Search for low mass vector resonances decaying into quark-antiquark pairs in proton-proton collisions at $\sqrt{s} = 13$ TeV, J. High Energy Phys. 01 (2018) 097.
[55] CMS Collaboration, Search for new resonances in the diphoton final state in the mass range between 80 and 110 GeV in pp collisions at $\sqrt{s} = 8$ TeV, CERN Technical Report No. CMS-PAS-HIG-14-037, 2015.
[56] CMS Collaboration, Search for new resonances in the diphoton final state in the mass range between 70 and 110 GeV in pp collisions at $\sqrt{s} = 8$ and 13 TeV, CERN Technical Report No. CMS-PAS-HIG-17-013, 2017.
[57] S. Chakraborty, M. Kraus, V. Loladze, T. Okui, and K. Tobioka, Heavy QCD axion in b → s transition: Enhanced limits and projections, Phys. Rev. D **104**, 055036 (2021).
[58] A. K. Likhoded, A. V. Luchinsky, and S. V. Poslavsky, Production of heavy quarkonia in hadronic experiments, JETP Lett. **105**, 739 (2017).
[59] N. Fabiano, Theoretical determination of eta(b)'s electromagnetic decay width, Nucl. Phys. B, Proc. Suppl. **126**, 255 (2004).
[60] K. Osato, T. Sekiguchi, M. Shirasaki, A. Kamada, and N. Yoshida, Cosmological constraint on the light gravitino mass from CMB lensing and cosmic shear, J. Cosmol. Astropart. Phys. 06 (2016) 004.
[61] A. E. Nelson and N. Seiberg, R symmetry breaking versus supersymmetry breaking, Nucl. Phys. **B416**, 46 (1994).
[62] K. A. Intriligator, N. Seiberg, and D. Shih, Supersymmetry breaking, R-symmetry breaking and metastable vacua, J. High Energy Phys. 07 (2007) 017.

[63] B. Bellazzini, A. Mariotti, D. Redigolo, F. Sala, and J. Serra, *R*-Axion at Colliders, Phys. Rev. Lett. **119**, 141804 (2017).

[64] R. D. Peccei and H. R. Quinn, *CP* Conservation in the Presence of Instantons, Phys. Rev. Lett. **38**, 1440 (1977).

[65] R. D. Peccei and H. R. Quinn, Constraints imposed by *CP* conservation in the presence of instantons, Phys. Rev. D **16**, 1791 (1977).

[66] K. I. Izawa, T. Watari, and T. Yanagida, Higher dimensional QCD without the strong *CP* problem, Phys. Lett. B **534**, 93 (2002).

[67] A. Fukunaga and K. I. Izawa, Warped QCD without the strong *CP* problem, Phys. Lett. B **562**, 251 (2003).

[68] H.-C. Cheng and D. E. Kaplan, Axions and a gauged Peccei-Quinn symmetry, arXiv:hep-ph/0103346.

[69] M. Redi and R. Sato, Composite accidental axions, J. High Energy Phys. 05 (2016) 104.

[70] G. G. Raffelt, Astrophysical axion bounds diminished by screening effects, Phys. Rev. D **33**, 897 (1986).

[71] G. G. Raffelt and D. S. Dearborn, Bounds on hadronic axions from stellar evolution, Phys. Rev. D **36**, 2211 (1987).

[72] A. Lai *et al.* (NA48 Collaboration), Precise measurement of the decay $K_L \to \pi^0\gamma\gamma$, Phys. Lett. B **536**, 229 (2002).

[73] A. Artamonov *et al.* (E949 Collaboration), Search for the decay K^+ to π^+ gamma gamma in the π^+ momentum region $P > 213$ MeV/c, Phys. Lett. B **623**, 192 (2005).

[74] E. Abouzaid *et al.* (KTeV Collaboration), Final results from the KTeV experiment on the decay $K_L \to \pi^0\gamma\gamma$, Phys. Rev. D **77**, 112004 (2008).

[75] C. Lazzeroni *et al.* (NA62 Collaboration), Study of the $K^\pm \to \pi^\pm\gamma\gamma$ decay by the NA62 experiment, Phys. Lett. B **732**, 65 (2014).

[76] V. Rubakov, Grand unification, and heavy axion, JETP Lett. **65**, 621 (1997).

[77] Z. Berezhiani, L. Gianfagna, and M. Giannotti, Strong *CP* problem and mirror world: The Weinberg-Wilczek axion revisited, Phys. Lett. B **500**, 286 (2001).

[78] A. Hook, Anomalous Solutions to the Strong *CP* Problem, Phys. Rev. Lett. **114**, 141801 (2015).

[79] H. Fukuda, K. Harigaya, M. Ibe, and T. T. Yanagida, Model of visible QCD axion, Phys. Rev. D **92**, 015021 (2015).

[80] S. Dimopoulos, A. Hook, J. Huang, and G. Marques-Tavares, A collider observable QCD axion, J. High Energy Phys. 11 (2016) 052.

[81] B. Holdom and M. E. Peskin, Raising the axion mass, Nucl. Phys. **B208**, 397 (1982).

[82] K. Choi, C. W. Kim, and W. K. Sze, Mass Renormalization by Instantons and the Strong *CP* Problem, Phys. Rev. Lett. **61**, 794 (1988).

[83] B. Holdom, Strong QCD at high-energies, and a heavy axion, Phys. Lett. **154B**, 316 (1985); **156B**, 452(E) (1985).

[84] M. Dine and N. Seiberg, String theory and the strong *CP* problem, Nucl. Phys. **B273**, 109 (1986).

[85] J. M. Flynn and L. Randall, A computation of the small instanton contribution to the axion potential, Nucl. Phys. **B293**, 731 (1987).

[86] K. Choi and H. D. Kim, Small instanton contribution to the axion potential in supersymmetric models, Phys. Rev. D **59**, 072001 (1999).

[87] P. Agrawal and K. Howe, Factoring the strong *CP* problem, J. High Energy Phys. 12 (2018) 029.

[88] P. Agrawal and K. Howe, A flavorful factoring of the strong *CP* problem, J. High Energy Phys. 12 (2018) 035.

[89] M. Dine, W. Fischler, and M. Srednicki, A simple solution to the strong *CP* problem with a harmless axion, Phys. Lett. **104B**, 199 (1981).

[90] A. Zhitnitsky, On Possible suppression of the axion hadron interactions. (In Russian), Sov. J. Nucl. Phys. **31**, 260 (1980).

[91] S. Chatrchyan *et al.* (CMS Collaboration), Search for a Light Pseudoscalar Higgs Boson in the Dimuon Decay Channel in pp Collisions at $\sqrt{s} = 7$ TeV, Phys. Rev. Lett. **109**, 121801 (2012).

[92] A. M. Sirunyan *et al.* (CMS Collaboration), Search for a Narrow Resonance Lighter than 200 GeV Decaying to a Pair of Muons in Proton-Proton Collisions at $\sqrt{s} =$ TeV, Phys. Rev. Lett. **124**, 131802 (2020).

[93] R. Aaij *et al.* (LHCb Collaboration), Search for a dimuon resonance in the Υ mass region, J. High Energy Phys. 09 (2018) 147.

[94] G. Cacciapaglia, G. Ferretti, T. Flacke, and H. Serodio, Revealing timid pseudo-scalars with taus at the LHC, Eur. Phys. J. C **78**, 724 (2018).

[95] D. Buarque Franzosi, G. Cacciapaglia, X. Cid Vidal, G. Ferretti, T. Flacke, and C. Vázquez Sierra, Exploring new possibilities to discover a light pseudo-scalar at LHCb, Eur. Phys. J. C **82**, 3 (2022).

[96] M. Aaboud *et al.* (ATLAS Collaboration), Search for Higgs boson decays into pairs of light (pseudo)scalar particles in the $\gamma\gamma jj$ final state in pp collisions at $\sqrt{s} =$ 13 TeV with the ATLAS detector, Phys. Lett. B **782**, 750 (2018).

[97] CMS Collaboration, Search for exotic decay of the Higgs boson into two light pseudoscalars with four photons in the final state at $\sqrt{s} = 13$ TeV, CERN Technical Report No. CMS-PAS-HIG-21-003, 2021.

[98] M. J. Strassler and K. M. Zurek, Echoes of a hidden valley at hadron colliders, Phys. Lett. B **651**, 374 (2007).

[99] T. Han, Z. Si, K. M. Zurek, and M. J. Strassler, Phenomenology of hidden valleys at hadron colliders, J. High Energy Phys. 07 (2008) 008.

[100] S. Knapen, J. Shelton, and D. Xu, Perturbative benchmark models for a dark shower search program, Phys. Rev. D **103**, 115013 (2021).

[101] B. A. Dobrescu, G. L. Landsberg, and K. T. Matchev, Higgs boson decays to *CP* odd scalars at the Tevatron and beyond, Phys. Rev. D **63**, 075003 (2001).

[102] P. Draper and D. McKeen, Diphotons from tetraphotons in the decay of a 125 GeV Higgs at the LHC, Phys. Rev. D **85**, 115023 (2012).

[103] N. Toro and I. Yavin, Multiphotons and photon jets from new heavy vector bosons, Phys. Rev. D **86**, 055005 (2012).

[104] M. Aaboud *et al.* (ATLAS Collaboration), A search for pairs of highly collimated photon-jets in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector, Phys. Rev. D **99**, 012008 (2019).

[105] S. Catani, L. Cieri, D. de Florian, G. Ferrera, and M. Grazzini, Diphoton Production at Hadron Colliders: A Fully-Differential QCD Calculation at NNLO, Phys. Rev. Lett. **108**, 072001 (2012); **117**, 089901(E) (2016).

[106] G. Aad *et al.* (ATLAS Collaboration), Measurement of isolated-photon pair production in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector, J. High Energy Phys. 01 (2013) 086.

[107] M. Aharrouche *et al.* (ATLAS Electromagnetic Barrel Calorimeter Collaboration), Energy linearity and resolution of the ATLAS electromagnetic barrel calorimeter in an electron test-beam, Nucl. Instrum. Methods Phys. Res., Sect. A **568**, 601 (2006).