

Electroweak ALP searches at a muon collider

Yunjia Bao,^a JiJi Fan^{a,b} and Lingfeng Li^a

^a*Department of Physics, Brown University,
Providence, RI, 02912, U.S.A.*

^b*Brown Theoretical Physics Center, Brown University,
Providence, RI, 02912, U.S.A.*

E-mail: yunjia_bao@brown.edu, jiji_fan@brown.edu,
lingfeng_li@brown.edu

ABSTRACT: A high-energy muon collider with center-of-mass energy around and above 10 TeV is also a vector boson fusion (VBF) machine, due to the significant virtual electroweak (EW) gauge boson content of high-energy muon beams. This feature, together with the clean environment, makes it an ideal collider to search for TeV-scale axion-like particles (ALP) coupling to Standard Model EW gauge bosons, which current and other future colliders have limited sensitivities to. We present detailed analyses of heavy ALP searches in both the VBF and associated production channels at a muon collider with different running benchmarks. We also show projected constraints on the ALP couplings in the effective field theory, including an operator with its coefficient not determined by the mixed Peccei-Quinn anomaly. We demonstrate that a muon collider could probe new ALP parameter space and push the sensitivities of the couplings between the ALP and EW gauge bosons by one order of magnitude compared to HL-LHC. The projected limits and search strategies for ALPs could also be applied to other types of resonances coupling to EW gauge bosons.

KEYWORDS: Axions and ALPs, Specific BSM Phenomenology

ARXIV EPRINT: [2203.04328](https://arxiv.org/abs/2203.04328)

Contents

1	Introduction	1
2	The effective theory of ALP coupling to EW gauge bosons	4
3	ALP collider signals and (HL)-LHC constraints	5
3.1	Vector boson fusion	5
3.2	Associated production	7
3.3	Non-resonant (indirect) signals	9
4	Muon collider phenomenology: simulation setups	9
4.1	Simulating ALP signals	10
4.2	Simulating major SM backgrounds	11
5	Muon collider phenomenology: VBF channels	13
5.1	Event reconstruction and analyses	14
5.2	Model independence of the limits	19
5.3	Analysis including forward muons	20
6	Muon collider phenomenology: associated production channels	21
6.1	Analysis of $a(\rightarrow \gamma\gamma)\gamma$	21
7	Constraining the EFT	23
8	Summary and outlook	26

1 Introduction

An axion-like particle (ALP) refers to a general periodic pseudo-scalar, $a \cong a + 2\pi f_a$ with f_a the decay constant. ALPs appear ubiquitously in extensions of the Standard Model (SM), usually as pseudo Nambu-Goldstone bosons arising from spontaneous breaking of some global symmetries [1, 2]. ALPs enjoy a wide range of phenomenological applications¹ and serve as one of the most motivated new physics scenarios. While most of the studies focus on very light ALPs, there has been growing interests to search for a heavy ALP with a mass m_a around or above the weak scale at collider-type experiments. A heavy ALP could acquire its mass from a strongly-interacting hidden sector with a confining scale above the weak scale.² Experimentally, the Large Hadron Collider (LHC) could probe m_a up to the

¹The earliest examples include the QCD axion from breaking of a global Peccei-Quinn symmetry, which could solve the strong CP problem and is also a cold dark matter candidate [1–11].

²Ref. [12] proposes to use an enlarged but unified color sector to bring the QCD axion mass up to the TeV scale.

scale of $\sim (2 - 3)$ TeV, around and above which the parameter space is still widely open. In this article, we will focus on searches of heavy ALPs with m_a of TeV scale, which couple to SM electroweak (EW) gauge bosons, at a possible future muon collider.

Searches for heavy ALP signals could be challenging at LHC and future colliders. A heavy ALP with $m_a \sim \mathcal{O}(1)$ TeV requires highly energetic beams for direct production. In addition, its one-loop-suppressed couplings and a large f_a further restrain the signal rate. Consequently, combining a high integrated luminosity and low background level is necessary for detecting it. A variety of direct heavy ALP searches are available (e.g., [13–26]), constraining visible ALPs at LEP and LHC. Future circular e^+e^- colliders [27, 28] are not optimal for the detection of TeV-scale ALPs as their achievable energies are limited by the synchrotron radiation energy loss. Future linear e^+e^- colliders [29, 30], although almost free from radiation energy loss, suffer from smaller luminosities as each bunch is dumped instead of stored after the first crossing. Heavy ALPs could be efficiently produced at future hadron colliders [31]. However, the signal could easily be overwhelmed by the enormous QCD backgrounds or vulnerable to large systematic uncertainties.

A future muon collider provides the potential solution to all the issues regarding energy, luminosity, and background cleanliness [32–37]. As $m_\mu \approx 207m_e$, the synchrotron radiation energy loss is under control since the energy loss rate is proportional to m_ℓ^{-4} . The circular collider design with a high luminosity could be achievable even for multi-TeV muon beams. From the conservative projection [38, 39], the expected integrated luminosity reaches $\mathcal{O}(1)$ ab^{-1} for TeV-scale muon colliders and could potentially increase further. Muon also carries a significant portion of beam energy as an elementary particle. In contrast, multiple quarks and gluons share the beam energy at a future hadron collider, with the high energy tip of parton distribution functions (PDFs) suppressed. When searching for a heavy ALP that couples to EW vector bosons, the advantage of a muon collider is even more explicit since the beam muon, unlike a beam proton, radiates a significant fraction of its energy to all EW vector bosons. For example, about 5% of the total beam energy is carried by γ , W , and Z in a TeV-scale muon beam [40], while this fraction drops to below 1% in a proton beam [41]. In addition, a muon collider can directly produce an ALP almost as heavy as the center-of-mass energy \sqrt{s} from $\mu^+\mu^-$ annihilations in association with an EW vector boson, which is unlikely at hadron colliders. Last but not least, the inclusive event rate of muon collisions is much lower than that of hadronic collisions. It is well known that the proton-proton inclusive cross section scales as the proton size ($\sim \Lambda_{\text{QCD}}^{-2} \sim \mathcal{O}(\text{mb})$). On the other hand, muons are elementary particles, and their inclusive cross section remains small at large \sqrt{s} , dominated by vector boson fusion (VBF) processes. This feature greatly reduces the overall background level and background-induced systematic uncertainties.

The main difficulty of the muon collider concept comes from the short muon lifetime $\sim 2 \mu\text{s}$. It is highly challenging to produce low emittance muon beams and store them in the collider ring before they decay. Thanks to the progress of accelerator physics, the particle physics community’s interest in a high-energy muon collider grows again. Several leading approaches have been developed quickly over the past years [42, 43]. The US Muon Accelerator Program (MAP) [44–47] focuses on proton-driven muon sources where secondary pions decay to muons. The program’s target scenario covers a wide range of

beam energy, from $\mathcal{O}(1-10)$ GeV neutrino factories to Higgs factories and multi-TeV colliders. Such proton-driven muon beams occupy a large phase space volume, which need to be cooled significantly before muons decay. The issue is addressed by the development of Muon Ionization Cooling Experiment (MICE) [48–50], which demonstrates the potential of cooling muon beams over short time scales. More recently, the novel approach of Low Emittance Muon Accelerator (LEMMA) allows high luminosity muon beams to be produced with low transverse emittance [47, 51]. The design uses a 45-GeV positron beam to produce muon pairs at the threshold, resulting in a small beam transverse emittance without extra beam cooling. The technology may allow a muon collider operating at even higher energies up to $\sqrt{s} = \mathcal{O}(100)$ TeV but with limited luminosities.

The abundant VBF interaction rates and low SM backgrounds make a high energy muon collider optimal for precision SM tests and beyond-the-Standard-Model (BSM) searches. With high energy and high luminosity muon beams, a future muon collider enables a plethora of possible discoveries. Latest phenomenological studies demonstrate the advantages of a muon collider in many different contexts. They already cover quite a few interesting aspects, such as measuring various Higgs properties [52–57], probing dark matter models [58–60], investigating models motivated by flavor physics [61–64] or the $g_\mu - 2$ anomaly [65–72], higher-dimensional operator analyses [73–76], and other BSM studies [77–86]. General prospects of muon collider phenomenology can be found in [87–89]. However, the panorama of future muon collider phenomenology remains to be fully explored.

Given all the potential advantages outlined above, we implement and report in this paper a detailed analysis of TeV-scale ALPs coupling to EW gauge bosons at a future muon collider. We show that the search of electroweakly coupled ALPs provides another showcase for the great physics potential of a high-energy muon collider. The search strategies and projected sensitivities on the cross sections could be applied to other types of EW resonances.

The paper is organized as follows. In section 2, we present the effective field theory (EFT) of an ALP that couples to SM EW gauge bosons. In section 3, we present the common production mechanisms of heavy ALPs at a collider. The HL-LHC projections on ALP couplings are deduced from current LHC limits, as HL-LHC is the inevitable step en route to any future colliders. In section 4, details of simulations for ALP searches at a muon collider are presented. Section 5 enumerates seven analyses on ALPs produced in the VBF channel, aiming at seven final states. Meanwhile, we show that the analysis results could apply to general ALP couplings and other new physics resonance models. The gain in sensitivity from forward-region information is also discussed. Another important ALP production mechanism, associated production with the tri-photon final state, is analyzed in section 6. In section 7, the projected constraints on the EFT couplings are deduced from limits derived in the previous sections. We conclude in section 8.

2 The effective theory of ALP coupling to EW gauge bosons

We will focus on the EFT of ALP coupling to the EW gauge bosons in the SM. Here, the heavy axion of interest has a mass of $\mathcal{O}(\text{TeV})$. Before electroweak symmetry breaking (EWSB), the Lagrangian reads:

$$\begin{aligned} \mathcal{L} = & \frac{1}{2} [(\partial a)^2 - m_a^2 a^2] + \left(\frac{g_1}{4\pi}\right)^2 C_{BB} \frac{a}{f_a} B_{\mu\nu} \tilde{B}^{\mu\nu} \\ & + \left(\frac{g_2}{4\pi}\right)^2 C_{WW} \frac{a}{f_a} W_{\mu\nu}^i \tilde{W}^{i;\mu\nu} + \left(\frac{g_1}{4\pi}\right) \left(\frac{g_2}{4\pi}\right) C_{BW} \frac{a}{f_a} B_{\mu\nu} \tilde{W}^{3;\mu\nu}, \end{aligned} \quad (2.1)$$

where $W(B)_{\mu\nu}$ are the field strength tensors of SM $SU(2)_W(U(1)_Y)$ gauge groups before EWSB, $\tilde{W}(\tilde{B})_{\mu\nu} \equiv \epsilon_{\rho\sigma\mu\nu} W(B)^{\rho\sigma}/2$ are the corresponding dual field strengths, $i = 1, 2, 3$ are the $SU(2)_W$ indices, and $g_{1(2)}$ are the corresponding coupling constants of $U(1)_Y(SU(2)_W)$. These effective couplings could be generated by integrating out heavy fermions carrying EW charges at one loop. The coefficients are thus suppressed by the loop factor, $1/(4\pi)^2$.³ Note that other conventions for the EFT Lagrangian exist. In particular, another widely used convention is to rescale the coefficients $\{C_{WW}, C_{BB}, C_{BW}\} \rightarrow \{(g_2/4\pi)^2 C_{WW}, (g_1/4\pi)^2 C_{BB}, g_1 g_2 / (4\pi)^2 C_{BW}\}$ to absorb the gauge couplings and loop factors [90]. We will follow the convention in eq. (2.1) throughout the analysis, while the other convention will also be presented in some plots for the readers' convenience. We will ignore other possible ALP couplings like the ones to SM fermions and the Higgs in our study since these couplings could lead to different production channels and detection strategies beyond the scope of this paper, e.g., s -channel muon annihilation $\mu^+ \mu^- \rightarrow a(+\gamma) \rightarrow \mu^+ \mu^- (+\gamma)$ if an axion has a sizable coupling to muon.

It is noteworthy that a non-zero C_{BW} coupling in eq. (2.1) is not always included in the literature. An ALP coupling is usually proportional to associated global anomaly coefficient, which is true for an ALP coupling to massless SM gauge bosons, such as photons or gluons. Yet as discussed in refs. [91–95], this coupling could arise in several ALP models and shall be considered in a generic analysis. Although there is no non-vanishing $U(1)_{\text{global}} \times U(1)_Y \times SU(2)_W$ anomaly coefficient for C_{BW} , the coupling term turns out to arise from dimension-7 operators such as $a H^\dagger \tau^i H W_{\mu\nu}^i \tilde{B}^{\mu\nu} / \Lambda^3$, where H is the SM Higgs doublet, τ^i 's are $SU(2)_W$ generators, and Λ is the UV cutoff scale. After EWSB, this results in an operator of the form $(a/\Lambda)(v_{\text{EW}}/\Lambda)^2 \tilde{B}^{\mu\nu} W_{\mu\nu}^3$, in which v_{EW} is the Higgs VEV. This term could be numerically important when Λ is not far beyond the EW scale. There have been efforts to construct models in which such terms appear. In the standard DFSZ axion model [7, 8], the light axion is a linear combination of angular modes of the two Higgs doublets and the PQ scalar [96, 97]. The top loop generates a non-zero C_{BW} coupling [95]. Another example of a possible UV completion is constructed in ref. [93] via a DFSZ-type extension in which heavy EW charged chiral fermions are integrated out. In this scenario, the new fermions receive their masses mainly from EWSB and are thus within reach of current and near-future collider searches. One may also note that other operators exist

³The loop factors here are expected in general so that we could reproduce the axion coupling to photons with the right form, i.e., $\alpha/(4\pi)$ times a quantized coefficient with α the fine structure constant, after EWSB.

beyond dimension 7; however, at and below dimension 7, the operators listed in eq. (2.1) are equivalent to the complete set of gauge-invariant operators that involve one axion field and EW gauge bosons. In this paper, we remain agnostic about the UV completion of the C_{BW} term in the ALP EFT.

After EWSB, the ALP couplings to the gauge boson mass eigenstates can be written as:

$$\mathcal{L} \supset \frac{1}{2}(\partial a)^2 - \frac{1}{2}m_a^2 a^2 + \frac{\alpha}{4\pi} \frac{a}{f_a} \left(C_{\gamma\gamma} F\tilde{F} + 2C_{\gamma Z} F\tilde{Z} + C_{ZZ} Z\tilde{Z} + \frac{2}{s_W^2} C_{WW} W^+\tilde{W}^- \right), \quad (2.2)$$

with

$$F\tilde{Z} \equiv \frac{1}{2}\epsilon^{\mu\nu\rho\sigma} F_{\mu\nu} Z_{\rho\sigma}, \quad s_W \equiv \sin\theta_W, \quad c_W \equiv \cos\theta_W, \quad (2.3)$$

where F , Z , W are EW field strengths, α is the fine structure constant, and θ_W is the Weinberg angle. The magnitude C_{WW} remains the same as in eq. (2.1). The other three effective couplings after EWSB are related to the linear combinations of the original ones as

$$\begin{aligned} C_{\gamma\gamma} &\equiv C_{BB} + C_{WW} + C_{BW}, & C_{\gamma Z} &\equiv C_{WW} \frac{c_W}{s_W} - C_{BB} \frac{s_W}{c_W} + \frac{1}{2} C_{BW} \left(\frac{c_W}{s_W} - \frac{s_W}{c_W} \right), \\ C_{ZZ} &\equiv C_{WW} \frac{c_W^2}{s_W^2} + C_{BB} \frac{s_W^2}{c_W^2} - C_{BW}. \end{aligned} \quad (2.4)$$

With the couplings above, the partial decay widths of axion to vector bosons are given by

$$\begin{aligned} \Gamma(a \rightarrow \gamma\gamma) &= \frac{1}{2} \left(\frac{C_{\gamma\gamma}}{f_a} \right)^2 \frac{\alpha^2 m_a^3}{32\pi^3}, \\ \Gamma(a \rightarrow \gamma Z) &= \left(\frac{C_{\gamma Z}}{f_a} \right)^2 \frac{\alpha^2 m_a^3}{32\pi^3} \left[1 - \left(\frac{m_Z}{m_a} \right)^2 \right]^3, \\ \Gamma(a \rightarrow ZZ) &= \frac{1}{2} \left(\frac{C_{ZZ}}{f_a} \right)^2 \frac{\alpha^2 m_a^3}{32\pi^3} \left[1 - \left(\frac{2m_Z}{m_a} \right)^2 \right]^{3/2}, \\ \Gamma(a \rightarrow W^+W^-) &= \left(\frac{C_{WW}}{f_a} \right)^2 \frac{\alpha^2 m_a^3}{32s_W^4\pi^3} \left[1 - \left(\frac{2m_W}{m_a} \right)^2 \right]^{3/2}, \end{aligned} \quad (2.5)$$

where m_Z and m_W are the masses of the Z and W gauge bosons.

3 ALP collider signals and (HL)-LHC constraints

In this section, we will introduce two production channels of heavy ALPs based on the EFT in eq. (2.1) and discuss projected constraints on the ALP couplings to EW gauge bosons at HL-LHC. We will also comment on the indirect probes where the heavy ALP is a mediator.

3.1 Vector boson fusion

VBF channels refer to both WW fusion and neutral boson fusion (ZZ , $Z\gamma$, and $\gamma\gamma$ fusion) processes, as shown in figure 1. They could be schematically written as

$$f\bar{f}' \rightarrow VV' + X\text{'s} \rightarrow a + X\text{'s}, \quad (3.1)$$

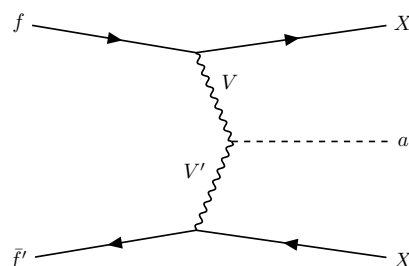


Figure 1. Feynman diagram for the VBF production of ALP at a collider.

where f and f' are beam fermions, V and V' are SM vector bosons, and X 's denotes recoiling beam remnants. The representation above applies to both hadron and lepton colliders with different choices of $f^{(i)}$ and corresponding X 's. Due to the enhanced amplitude when X 's are collinear with the beams, X 's are mainly present in the forward regions with large $|\eta|$'s, which could be helpful to veto non-VBF backgrounds. However, the detector performance is usually limited in these regions. The VBF-produced ALPs will also have low p_T 's on average and generate a sharp invariant mass peak $\approx m_a$ if their final states are fully visible. Since the width $\Gamma_a \ll m_a$ for realistic couplings, it is justified to take the narrow-width approximation and make the ALP production and decay independent to each other.

At hadron colliders like the LHC, the initiating fermions $f^{(i)}$ are quarks and the recoiling remnant X 's contain two (or even more) jets in the high- $|\eta|$ region. LHC has already implemented several relevant searches with the $\sqrt{s} = 13$ TeV dataset, e.g., searches for diphoton resonances [98, 99], inclusive WW or ZZ resonances [100–103], and $Z\gamma$ resonances [104, 105]. The large SM backgrounds from hadronic interactions significantly hinder signal reconstruction and affect the discovery potential. For example, soft particles from pileup could pose serious challenges to reconstruct mother particles in the decay chains, especially for hadronic W or Z fat jets as their large jet areas contain more pileup particles [106]. Most exclusion limits at 95% C.L. on $\sigma(pp \rightarrow a + X\text{'s}) \times \text{BR}(a \rightarrow VV')$ are of $\mathcal{O}(1-10)$ fb when $m_a \approx 1$ TeV. These limits strengthen by about one order of magnitude when $m_a \approx 2$ TeV since the SM backgrounds drop rapidly as the hard scattering energy scale increases.

We adopt the four VBF rates of ALP productions from ref. [107], using the LUXqed photon PDF [108] and matrix element (ME) method. Notice that the interference between initial state γ and transverse Z bosons is ignored as the PDF set including W and Z bosons is not yet available. In addition, such an interference is unlikely to change the ALP production rate significantly [109, 110]. Hence, we leave a more precise determination of ALP production rates at hadron colliders to future work. In the left panel of figure 2, we show the corresponding production cross section $\sigma(pp \rightarrow a + X\text{'s})_{\text{VBF}}$ as a function of axion mass m_a , normalizing each $C_{VV'}/f_a = 1 \text{ TeV}^{-1}$ after EWSB. Hierarchies between processes originate from vector boson properties and definitions in eq. (2.2). For example, the WW fusion rate is enhanced by the large s_W^{-2} factor in eq. (2.2) and becomes much larger than

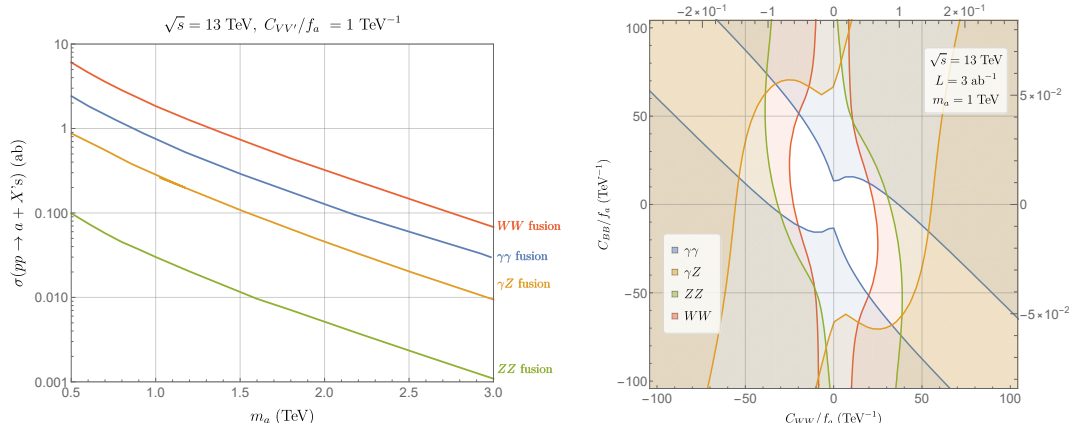


Figure 2. *Left:* Different VBF production rates as functions of m_a at the 13-TeV LHC, fixing $C_{VV'}/f_a = 1 \text{ TeV}^{-1}$. It should be noted that it is impossible to obtain all $C_{VV'}/f_a = 1 \text{ TeV}^{-1}$ simultaneously since there are only three Wilson coefficients, $(C_{BB}/f_a, C_{BW}/f_a, C_{WW}/f_a)$, before EWSB. Nonetheless, treating γ , Z , and W as partons of protons allows us to show four fusion channels individually as benchmarks. *Right:* Projected HL-LHC 95% C.L. constraints of C_{WW}/f_a and C_{BB}/f_a , fixing $C_{BW} = 0$ and $m_a = 1 \text{ TeV}$. The bottom and left axes indicate values of the axion couplings defined in eq. (2.1), while the top and right axes indicate another popular EFT convention in [90].

the other three. It is followed by the $\gamma\gamma$ fusion rate as it is easy for the energetic proton beam to radiate a massless photon. For $m_a, f_a \gtrsim 1 \text{ TeV}$, and assuming $\mathcal{O}(1)$ EFT couplings, the typical production rates are $\lesssim 10^{-2} \text{ fb}$, which are at least two orders of magnitudes below the current LHC limits.

To estimate the HL-LHC sensitivities on couplings between ALP and EW gauge bosons, we take the current limits and scale them with the integrated luminosity. We assume that systematic uncertainties will scale as statistical ones.⁴ The constraints on C_{WW}/f_a and C_{BB}/f_a are plotted in the right panel of figure 2 for a 1-TeV ALP, fixing $C_{BW} = 0$. We take into account of $a \rightarrow W^+W^-$ and $a \rightarrow ZZ$ constraints from [103], $a \rightarrow Z\gamma$ constraint from the scalar production benchmark in ref. [100], and $a \rightarrow \gamma\gamma$ constraint from [99].⁵ For $f_a \gtrsim m_a \sim \text{TeV}$, HL-LHC could only probe $C_{VV'}$ of $\mathcal{O}(10)$ or higher. The situation still persists when $C_{BW} \neq 0$. In another convention of ALP EFT [90], where the loop factors and gauge couplings in eq. (2.1) are absorbed in the definition of couplings, the projected sensitivities will appear stronger, as shown by the top and right axes of the right panel in figure 2.

3.2 Associated production

Another important ALP production mechanism is the associated production of an ALP with an EW vector boson, sometimes also called ALP-strahlung:

$$f\bar{f}' \rightarrow Va, \quad V = \gamma, W^\pm, Z. \quad (3.2)$$

⁴We also assume difficulties induced by the high pileup level in the high luminosity era [111] can be mitigated by techniques such as time-of-flight measurements [112] and machine learning [113, 114].

⁵We ignore the signal efficiency differences between ALP and the benchmark models (e.g., KK graviton or radion) used in the LHC analyses.

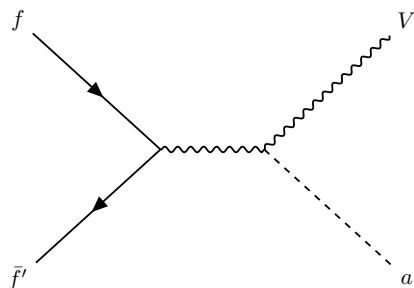


Figure 3. Feynman diagram for ALP production associated with a SM vector boson.

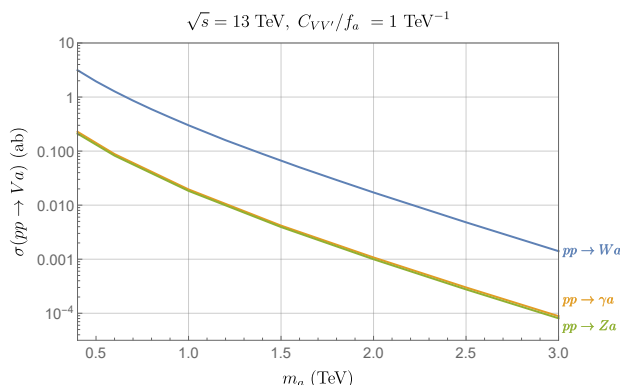


Figure 4. Different associated production rates as a function of m_a at the 13-TeV LHC, fixing $C_{VV'}/f_a = 1 \text{ TeV}^{-1}$.

As the ALP further decays to two bosons, the typical collider signal is a tri-boson final state, $VV'V''$, with $V'V''$ forming a narrow resonance. The powers of loop factors and f_a suppressions in these processes are the same as the VBF production channels. Since we ignore the ALP couplings to SM fermions, the associated production at colliders is dominated by the s -channel diagram through an off-shell vector boson, as shown in figure 3.

At LHC, several searches for tri-boson signals are implemented with the 13-TeV data. There are limits on cross sections of $W\gamma\gamma$ and $Z\gamma\gamma$ processes [115]. Measurements of WWW , WWZ , WZZ and ZZZ final states are also presented in [116–118]. Nonetheless, for $WZ\gamma$ and $WW\gamma$ final states, the 13-TeV analysis is not available yet, leaving the latest constraints from the 8-TeV data [119]. Most analyses above are non-resonant measurements of tri-boson rates, which all turn out to be compatible with the SM predictions. The typical uncertainties on $\sigma(VV'V'')$ are of $\mathcal{O}(1 - 100)$ fb depending on the final states. The resonance $Wa(\rightarrow WW)$ limit is also deduced in ref. [118], covering $m_a \in [200, 600]$ GeV range.⁶ Figure 4 shows the typical $pp \rightarrow Va$ cross sections at $\sqrt{s} = 13 \text{ TeV}$, with each relevant $C_{VV'}/f_a = 1 \text{ TeV}^{-1}$ after EWSB. The production rates drop faster than the VBF ones as m_a increases, due to the power suppression of energy from the s -channel propagator.

Because of the lower rates and essentially non-resonant limits, the projected HL-LHC bounds from Va associated production are even weaker than the VBF ones in the right

⁶The signal regions, however, do not target narrow WW resonances; thus, the limits are still qualitatively similar to non-resonant ones.

panel of figure 2. For a 1-TeV ALP (with f_a of the same order), typical constraints on $C_{VV'}$ are $\gtrsim \mathcal{O}(10^2)$, dominated by the $W\gamma\gamma$ and $Z\gamma\gamma$ measurements with lower experimental uncertainties. Even for future resonance searches associated with an extra boson, the small Va production rates still make this channel less relevant at HL-LHC.

3.3 Non-resonant (indirect) signals

Inevitably, heavy ALPs coupled to the SM could also modify event rates as a mediator. In this case, there will be no narrow ALP resonances in the final states. Since the most significant correction stems from the interference terms between the SM diagrams and ALP-exchanging diagrams, the modifications have the same loop factor and f_a suppressions as the VBF processes. However, as discussed in the last section, non-resonant searches have large SM backgrounds and are subject to significant systematic uncertainties. The current LHC data from this channel are only relevant when $m_a \lesssim m_Z$ [120].

For $m_a \gg m_Z$, the impact of virtual ALPs can be described by several dimension-8 anomalous quartic gauge coupling (aQGC) operators. They are suppressed by powers of m_a^{-1} , in addition to f_a^{-1} and loop factors from the EFT in eq. (2.1). The natural scaling of aQGC terms from ALP mediation follows

$$\mathcal{L}_{\text{eff}} \supset \frac{C_{VV'}^2 \alpha^2}{16\pi^2 f_a^2 m_a^2} \mathcal{O}_{\text{SM}}^8, \tag{3.3}$$

where $\mathcal{O}_{\text{SM}}^8$ represents dimension-8 aQGC operators. The limits can be set by both multi-boson [121] and vector boson scattering (VBS) [110] measurements. The extra suppression factor $\sim \alpha^2/(16\pi^2) \lesssim 10^{-6}$ strongly hinders the discovery potential of this non-resonance method. According to the HL-LHC performance study [122], the coefficient of the dimension-8 operators above can be measured down to $\sim 0.5 \text{ TeV}^{-4}$ in the high luminosity phase. To make the signal detectable for $m_a, f_a \simeq 1 \text{ TeV}$, $C_{VV'}$'s need to be as large as $\mathcal{O}(10^3)$. Such bounds are much weaker than those given by the direct searches. Though corrections to the discussions above may arise from finite $m_a \lesssim \sqrt{s}$, it is safe to conclude that the non-resonant/indirect approach will be unlikely to play an important role in heavy ALP searches at the HL-LHC.

4 Muon collider phenomenology: simulation setups

In the forthcoming sections, we will study the potential of searching a heavy ALP, in particular, through its coupling to EW gauge bosons at a future muon collider.

Throughout this work, we use MADGRAPH 5 [123] to generate parton-level hard processes for both ALP signals and SM backgrounds. We then use PYTHIA 8 [124] to handle hadronic and electromagnetic shower effects. The detector effects are simulated by DELPHES 3 [125] with its built-in muon collider detector template.

For simplicity, we follow the detector template's default parameters for reconstructing elementary objects, e.g., photon γ and light leptons ℓ . A photon with transverse momentum $p_T > 0.5 \text{ GeV}$ is considered isolated if the scalar sum of p_T over all the other particles inside a cone with a radius $\Delta R = 0.1$ around the photon is smaller than 20% of the photon's p_T .

For simplicity, we adopt the approximation that the detector efficiencies and resolutions of electrons and muons are the same. In practice, only processes with muonic final states are simulated, with their event rates rescaled to the inclusive ones of both electrons and muons accordingly. The p_T and isolation requirements for isolated muons are the same as the photon ones. Both isolated muons and photons must satisfy $|\eta| < 2.5$. Besides, muons in VBF beam remnants encode critical information like the identities of initial-state vector bosons. The ability to detect beam muons, usually in the forward region, enables a better signal background discrimination [58]. These recoiling muons are then measured with the default forward muon spectrometer in the muon collider template, covering the $2.5 < |\eta| < 8$ region. We use the Valencia algorithm to cluster jets for hadronic ALP decay final states, which was reported to provide better performances at high-energy lepton colliders [126]. The W and Z bosons from ALP decays are highly boosted on average. Therefore, we cluster them into fat jets with cone sizes $R = 1.0$ or 1.2 , depending on the detection channels. All fat jets must have $p_T > 200$ GeV, above which W/Z bosons are sufficiently boosted. We also assume the tracker's resolution on the impact parameter is much smaller than τ 's lifetime $\sim 87 \mu\text{m}$. Charged particles from τ decays are thus long-lived enough to have significant impact parameters. Such features are distinctive from prompt ALP decays. Therefore, we safely ignore any backgrounds having $Z \rightarrow \tau\tau$ decays.

We follow the muon collider's beam energy and integrated luminosity benchmarks in [38, 39, 88]. For each beam energy, the analysis results are projected according to the conservative (L_{con}) or the optimistic (L_{opt}) scenario of the integrated luminosity. In particular, we choose four different energy benchmarks and corresponding luminosities:

$$\begin{aligned} \sqrt{s} &= 10, 14, 30, 50 \text{ TeV}, \\ L_{\text{con}} &= 10, 10, 10, 10 \text{ ab}^{-1}, \\ L_{\text{opt}} &= 10, 20, 90, 250 \text{ ab}^{-1}. \end{aligned} \tag{4.1}$$

4.1 Simulating ALP signals

To evaluate different ALP production mechanisms at a muon collider, their rates at $\sqrt{s}=14$ TeV are calculated by the full ME method without any parton-level cuts. For every m_a , we take $\{C_{WW}, C_{BB}, C_{BW}\}/f_a = \{1, 1, 0\} \text{ TeV}^{-1}$. From the results in figure 5, the VBF mechanism (corresponding to the blue curve) is the most efficient way to produce heavy ALPs at a muon collider, with the largest rate for any $m_a < \sqrt{s}$. Meanwhile, the associated Va production rate remains small in the entire mass range considered. However, the SM tri-boson backgrounds for the Va channel are also small. Therefore, the associated production mechanism could still be a relevant search channel. For the indirect search, the difference in the VBS rate induced by non-resonant ALP diagrams, namely $\sigma(\mu^+\mu^- \rightarrow W^+W^- + X's)|_{\text{off-shell+SM}} - \sigma(\mu^+\mu^- \rightarrow W^+W^- + X's)|_{\text{SM}}$, is also plotted. Here $\sigma(\mu^+\mu^- \rightarrow W^+W^- + X's)|_{\text{off-shell+SM}}$ is the cross section of the same process including t -channel and off-shell s -channel ALP diagrams. Similar to the discussion in section 3.3 for the LHC, the indirect search channels are not as useful as the direct ones due to the absence of narrow ALP resonances, though their rates could be larger than those of Va channels.

Although most of the beam energy is carried by the muon, a fraction of it could be transferred to collinear vector bosons and neutrinos. Such initial-state-radiation (ISR)

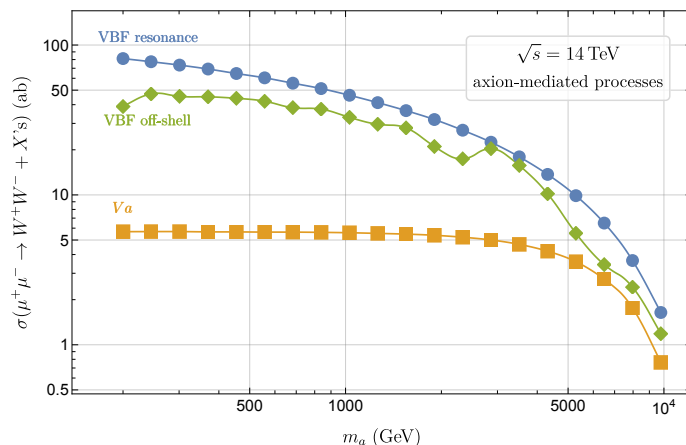


Figure 5. Benchmark production rates of the W^+W^- final state via the heavy axion as a function of its mass at a future muon collider with $\sqrt{s} = 14$ TeV. The Va channel includes both γa and Za processes with $a \rightarrow W^+W^-$ decay. The indirect effect of an off-shell exchange of an axion is estimated by $\sigma_{\text{off-shell}} \equiv \sigma(\mu^+\mu^- \rightarrow W^+W^- + X's)|_{\text{off-shell+SM}} - \sigma(\mu^+\mu^- \rightarrow W^+W^- + X's)|_{\text{SM}}$, in which the $\sigma(\mu^+\mu^- \rightarrow W^+W^- + X's)|_{\text{SM}}$ is the SM VBS cross section, and $\sigma(\mu^+\mu^- \rightarrow W^+W^- + X's)|_{\text{off-shell+SM}}$ is the cross section of the same process including t -channel and off-shell s -channel ALP diagrams. Similarly, the corresponding VBF resonance cross section is also for W^+W^- final state. Note that the fluctuation of $\sigma_{\text{off-shell}}(m_a)$ lies within the Monte Carlo uncertainties and is likely due to numerical artifacts.

effect introduces corrections to $\mu^+\mu^-$ annihilations. With ISR, the effective center-of-mass energy of $\mu^+\mu^-$ annihilation drops below \sqrt{s} , leading to a mild increase of event rates. Since the built-in muon PDF for muon beam is not yet available for our simulation framework, we approximate the ISR effect by inclusive collinear photon emissions. In particular, events with at most one extra photon γ_{col} are simulated for both signals and backgrounds, i.e., $\mu^+\mu^- \rightarrow a\gamma + (\gamma_{\text{col}})$ and $\mu^+\mu^- \rightarrow 3\gamma + (\gamma_{\text{col}})$. As we are interested in cases where the ISR photon is unidentified, the extra ISR photon must satisfy $p_T < 0.1$ GeV and $|\eta| > 2.5$. For other photons, only a $p_T(\gamma) > 10$ GeV cut is imposed at the parton level.

4.2 Simulating major SM backgrounds

SM backgrounds are more involved to simulate compared to ALP signals. At a muon collider, the irreducible backgrounds for VBF ALP channels are the SM diboson processes, especially the VBS [110] due to large EW boson PDFs and sizable SM gauge couplings. The situation is analogous to the $gg \rightarrow jj$ backgrounds for the dijet resonance search at hadron colliders. The diboson invariant mass distributions of VBS backgrounds are continuous, on top of which we search for the sharp peak of an ALP resonance. The overall background level can thus be modeled by a background fit. Moreover, vector bosons from ALP decays give characteristic final states, e.g., a hard photon, a lepton pair at the Z pole, or a fat jet with mass $\approx m_{W/Z}$, which further suppress reducible backgrounds like $\mu^+\mu^- \rightarrow q\bar{q}$.

In contrast to signal simulations, the soft and collinear singularities in VBS amplitudes require parton-level cuts, which render the Monte Carlo integrations convergent. However,

Background Processes	Parton-level Cuts	Methods	σ_{BKGD} (fb)
$\mu^+\mu^- \rightarrow \gamma\gamma + X$'s			1.05×10^2
$\mu^+\mu^- \rightarrow ZZ(2\mu^+2\mu^-) + X$'s	$p_{T,\gamma} \geq 10 \text{ GeV}$ $\& p_{T,\ell} \geq 0.5 \text{ GeV}$ $\& \eta_\ell \leq 8$	ME	0.310
$\mu^+\mu^- \rightarrow ZZ(4j) + X$'s			1.27×10^2
$\mu^+\mu^- \rightarrow \gamma Z + X$'s			3.18×10^2
$\mu^+\mu^- \rightarrow \gamma W^\pm + X$'s			15.8
$\mu^+\mu^- \rightarrow W^+W^- + X$'s			6.26×10^3
$\mu^+\mu^- \rightarrow Z(\mu^+\mu^-)Z(jj) + X$'s		ME + m_μ	12.5
$\mu^\pm\gamma^* \rightarrow Z(\mu^+\mu^-)W^\pm(jj) + X$'s	no cut	ME + EVA	26.6
$\mu^\pm\gamma^* \rightarrow Z(jj)W^\pm(jj) + X$'s			5.33×10^2

Table 1. Summary of cuts and cross sections for SM backgrounds simulated by MADGRAPH: γ^* denotes a partonic photon from the muon beam, and X 's denote the beam remnants, including muon, muon neutrino, and their antiparticles. The background cross section σ_{BKGD} is understood to include branching fractions to all secondary decays indicated in the parentheses. Decay products from Z 's and W 's are not subject to parton-level cuts.

stringent parton-level cuts on beam remnants could lead to background underestimations. It is necessary to choose proper cuts to achieve realistic background distributions and rates without oversampling in the divergent part of the phase space. Details of each background channel simulated are summarized in table 1.

Depending on the beam remnants, X 's, VBS backgrounds could be classified into three types: 1) X 's contain only muon neutrinos, with W^+W^- as the corresponding initial-state vector bosons; 2) X 's contain only muons, with the initial states being $\gamma\gamma$, $Z\gamma$ or ZZ ; 3) X 's contain one muon and one muon neutrino, with the initial states WZ or $W\gamma$. For type 1), the background rates are not sensitive to cuts on neutrino beam remnants. For type 2), the final state could only be W^+W^- due to the SM gauge symmetry. Equivalently, ZZ , $Z\gamma$, $\gamma\gamma$ final states could only be from type 1). The process initiated by $\gamma\gamma$ are susceptible to variations of parton-level cuts due to a collinear divergence. Its cross section is about one order of magnitude above the other processes in this category. Fortunately, at the cost of a long integration time, the ME method still provides plausible distributions of $\gamma\gamma$ -initiated processes. In practice, we use the same parton-level cuts as in type 1).

For type 3), i.e., VBS ZW and γW events, the final state dibosons have a net electric charge. They only contribute to the reducible backgrounds when the W decays hadronically and is misidentified as a Z . For $\mu^+\mu^- \rightarrow \gamma W + X$'s background, the ME method with soft parton-level cuts can give converging results. However, the ME method leads to numerical instability when generating $\mu^+\mu^- \rightarrow ZW + X$'s samples. In this case, we use the effective vector boson approximation (EVA) based on the improved Weizsäcker-Williams formula [127, 128], incorporated in MADGRAPH [129] to generate $\mu^\pm\gamma^* \rightarrow ZW + X$'s samples, where γ^* denotes the partonic photon from the muon beam. The vertex of $\mu \rightarrow W\nu$ emission is still handled by the ME calculation in this case. For simplicity, the Z boson component in the muon beam is ignored, as its contribution to the overall VBS ZW rate is negligible. For our background simulation, we set the factorization scale $Q = \sqrt{s}/2$. We

Final States	p_T (GeV)	$m_{\ell\ell}$	R_J	m_J (GeV)	Resonance Window
$\gamma\gamma$	$\gamma_1 \geq 350$	—	—	—	$[m_a - 3\sigma_m, m_a + 3\sigma_m]$
$\gamma Z(\ell\ell)$	$\gamma_1 \geq 350$ $\ell_1 \geq 100$	$m_Z \pm 5\Gamma_Z$	—	—	$[m_a - 3\sigma_m, m_a + 3\sigma_m]$
$\gamma Z(jj)$	$\gamma_1 \geq 350$ $J_1^* \geq 200$	—	1.2	$m_Z^{+2.5}_{-1.5}\Gamma_Z$	$[m_a - 1.5\sigma_m, m_a + 0.5\sigma_m]$
$ZZ(4\ell)$	$\ell_1 \geq 300$ $\ell_2 \geq 200$	$m_Z \pm 5\Gamma_Z$	—	—	$[m_a - 3\sigma_m, m_a + 3\sigma_m]$
$ZZ(2\ell 2j)$	$\ell_1 \geq 300$ $J_1 \geq 200$	$m_Z \pm 5\Gamma_Z$	1.2	$m_Z^{+2.5}_{-1.5}\Gamma_Z$	$[m_a - 1.5\sigma_m, m_a + 0.5\sigma_m]$
$ZZ(4j)$	$J_{1,2} \geq 300$	—	1.0	$m_Z^{+5}_{-2}\Gamma_Z$	$[m_a - 1.5\sigma_m, m_a + 0.5\sigma_m]$
$WW(4j)$	$J_{1,2} \geq 300$	—	1.0	$m_W^{+3}_{-5}\Gamma_W$	$[m_a - 1.5\sigma_m, m_a + 0.5\sigma_m]$

Table 2. Detector-level cuts imposed on various analyses: — means that no value is applicable; J represents a detector-level fat jet, while j corresponds to a parton-level jet. The subscript of a particle name denotes its detector-level p_T ordering. The jet radius $R = 1$ is used for final states with two J 's (corresponding to $4j$'s), and $R = 1.2$ is used for final states with one J ($2j$'s). The superscript $*$ indicates that the requirement is only for fat jets that do not contain photons. The width of each resonance peak is characterized by σ_m , see the main text for more details.

also calculate the ZW background rates with $Q = \sqrt{s}$ and $Q = \sqrt{s}/4$ and estimate the factorization scale induced uncertainty to be $\sim 6\%$.

The discussions above focus on the VBS backgrounds, the leading ones for ALP searches in the VBF channels. For the search in the associated production channel, we need to consider the SM tri-boson background. In particular, we will consider the 3γ final state, the simplest one in the associated production channel. Most VBS 3γ background can be removed by requiring the invariant mass of 3γ to be $\approx \sqrt{s}$. Simulation shows that the VBS 3γ background rate after applying the parton-level cut $m_{3\gamma} \in [0.9, 1.0]\sqrt{s}$ is four orders of magnitude below that of $\mu^+\mu^-$ annihilations and can be ignored safely. The ISR effects are also included by the same method described in section 4.1, giving a cross section of $\mu^+\mu^- \rightarrow 3\gamma + (\gamma_{\text{col}}) = 0.44 \text{ fb}$.

5 Muon collider phenomenology: VBF channels

As discussed in section 3 and 4, VBF production of ALPs provides a heavy diboson resonance with a narrow width. Therefore, we focus on the diboson final states which are fully visible, i.e., no $W \rightarrow \ell\nu$ or $Z \rightarrow \nu\bar{\nu}$ decays. All final states containing τ 's will also be excluded from our analysis since they always contain neutrinos. Here we propose several analyses targeting all four ALP decay modes and their final states. Details of each analysis are described in section 5.1.

The general procedures of our analyses are straightforward. We first select events containing the same number of photons, leptons, and/or jets as the target ALP final state. Hard cuts on p_T of $\mathcal{O}(m_a)$ are often applied to these objects at this stage to suppress SM

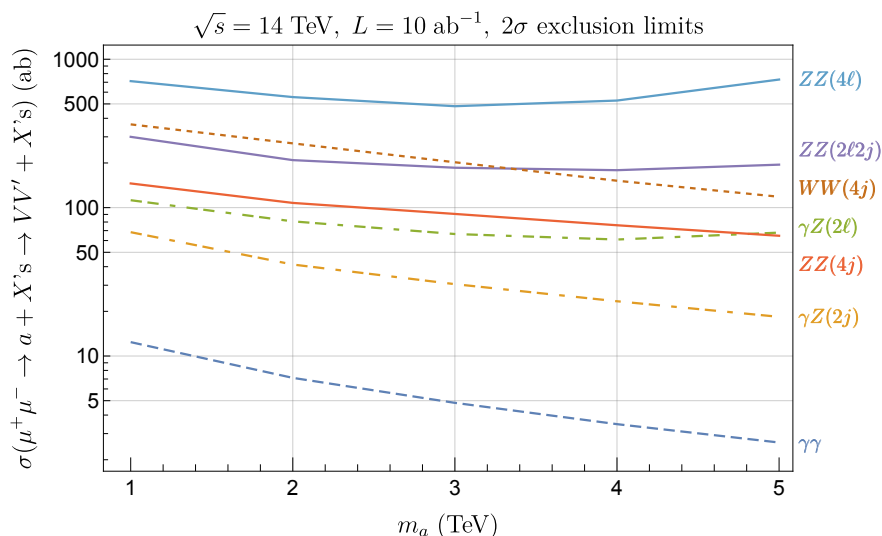


Figure 6. Projected 2σ exclusion limits on the VBF production cross section times the branching fraction of the ALP to a specific diboson final state, at a 14-TeV muon collider with $L = 10 \text{ ab}^{-1}$. Solid lines are exclusion limits from the analyses of $a \rightarrow ZZ$ with different secondary decays of Z 's, short dashed line is the exclusion limit from the analysis of $a \rightarrow WW$ with W 's decaying hadronically, long dashed line is the exclusion limit from the $\gamma\gamma$ channel, and exclusion limits from $a \rightarrow \gamma Z$ analyses are given by the dot-dashed lines. The W and Z decay products are indicated in the parentheses.

VBS backgrounds. Z/W reconstruction cuts are applied if the final state involves Z/W decays. Sensitivities on $\sigma(\mu^+\mu^- \rightarrow a+X's) \times \text{BR}(a \rightarrow ZZ, Z\gamma, \gamma\gamma, WW)$ are deduced from the signal and background yields within a narrow invariant mass window around m_a . The width of each invariant window is determined by the final state, the detector resolution, and the ALP mass. As to be discussed in section 5.2, different initial states of VBF production only cause mild variations in signal efficiencies; hence, results provided in this section are approximately model independent. On the other hand, forward-region observables, e.g., number and energies of forward muons, are highly correlated with the initial boson states. Including forward-region information into the analysis may further enhance the sensitivity; see section 5.3. Since the forward detector at a very-high-energy muon collider is still under development, we present limits utilizing forward-region information only as suggestive values.

5.1 Event reconstruction and analyses

To show concrete kinematic distributions and efficiencies, in the remainder of this section we will mainly present distributions and results for the $\sqrt{s} = 14 \text{ TeV}$ benchmark. Analyses on seven final states are demonstrated in the following paragraphs, covering all four ALP diboson decay modes. Table 2 summarizes the key cut parameters for all the analyses. For convenience, we use J to represent detector-level fat jets, while j corresponds to parton-level jets. Also, the subscript of a particle name denotes its p_T ordering at the detector level. The projected 2σ exclusion limits on $\sigma(\mu^+\mu^- \rightarrow a+X's \rightarrow VV'+X's)$ are plotted in

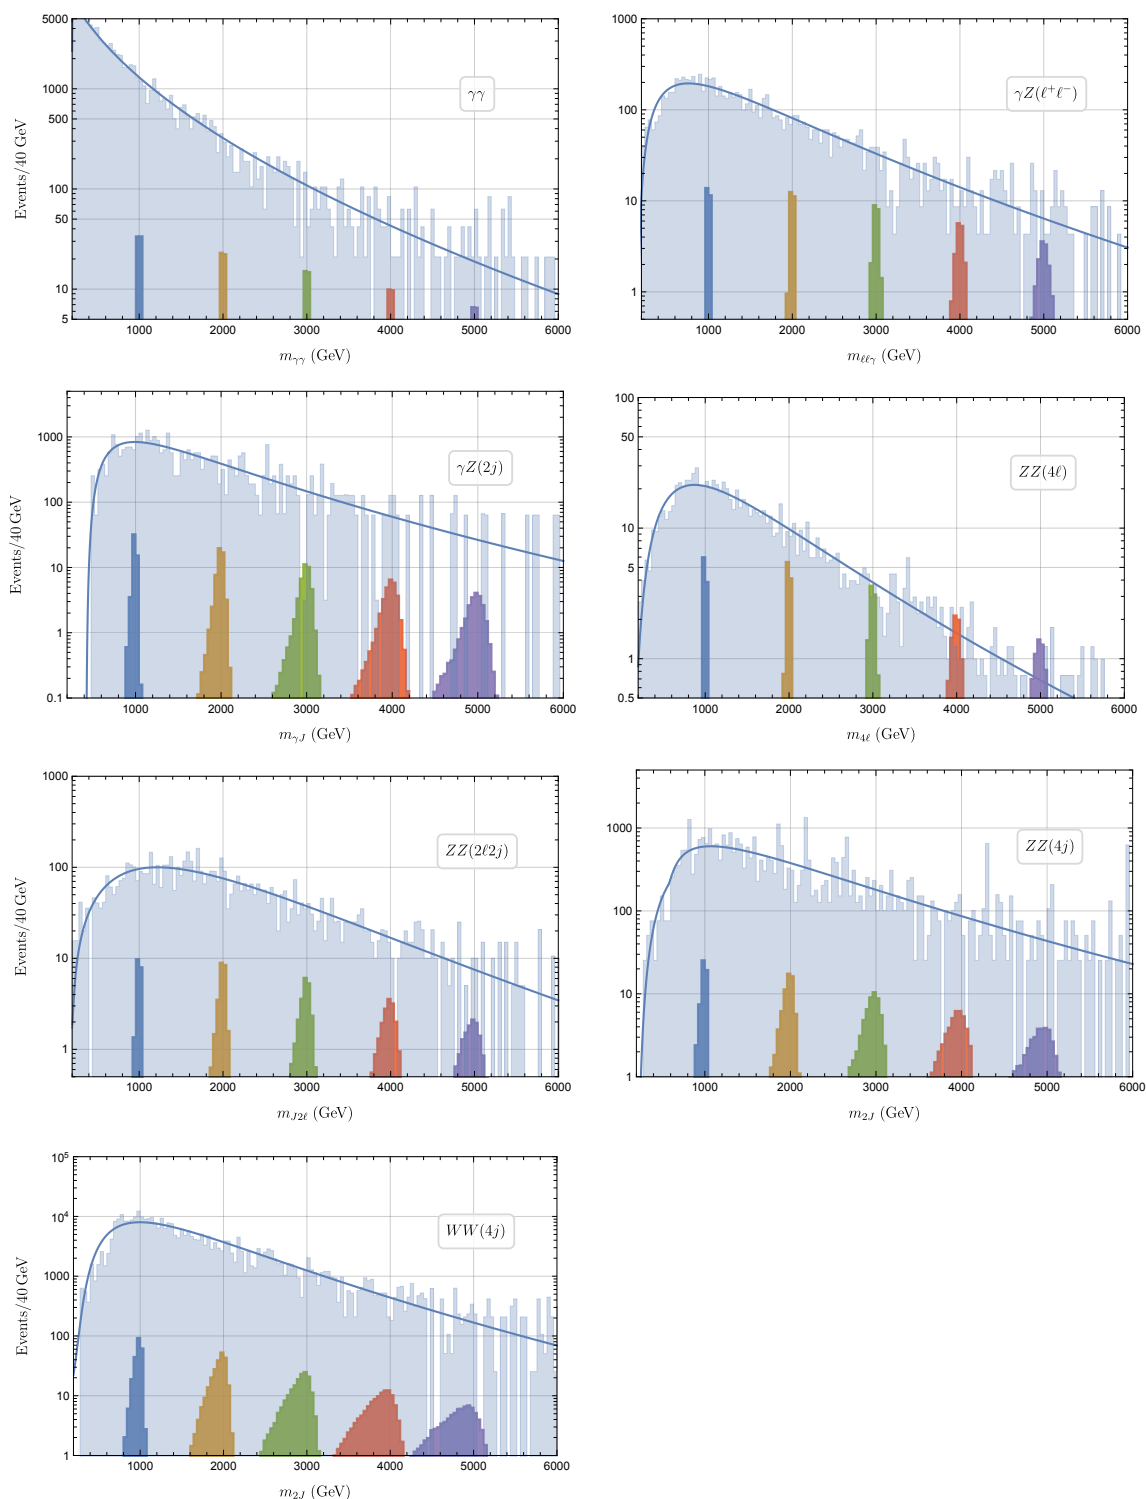


Figure 7. SM backgrounds *vs.* ALP resonance signals for various ALP masses ($m_a = 1, 2, 3, 4, 5$ TeV) at a 14-TeV muon collider with $L = 10 \text{ ab}^{-1}$. Simulated data are shown in histograms with a bin width of 40 GeV while fitted background distributions are shown as solid curves. For all signal peaks, the cross sections are normalized to their 2σ exclusion limits.

figure 6, assuming Poisson statistics [130]. Since the typical signal and background yields are sufficiently larger than one, signal significance is well approximated by the Gaussian limit $S/\sqrt{S+B}$, in which S (B) denotes the signal (background) counts in the signal region. Depending on ALP reconstruction efficiencies, EW boson decay branching ratios, and background levels, the limits from different analyses can vary by two orders of magnitude. Combining possible final states, the $a \rightarrow \gamma\gamma$ mode has the highest sensitivity, followed by $a \rightarrow Z\gamma$, $a \rightarrow ZZ$, and $a \rightarrow WW$ with the least sensitivity.

Analysis of $a \rightarrow \gamma\gamma$. As the preselection criteria, only events containing exactly two isolated photons with the transverse momentum of the leading photon satisfying $p_T(\gamma_1) > 350$ GeV are kept. Also, no more than one charged track other than forward muons with $p_T > 10$ GeV is allowed. The signal efficiency for the $m_a = 1(5)$ TeV benchmark is $\sim 57(74)\%$ after the preselection. The same cuts keep only 8.5% of VBS diphoton backgrounds since most of them have low $p_T(\gamma)$ and are vetoed by the hard $p_T(\gamma_1)$ cut.

The diphoton invariant mass ($m_{\gamma\gamma}$) distributions for the VBS SM background and ALP benchmarks are shown in the upper left panel of figure 7. The typical widths of the $a \rightarrow \gamma\gamma$ signal peaks are $\lesssim 20$ GeV due to the high resolution on the photons' momenta. The signal region for each ALP benchmark is then defined as $|m_{\gamma\gamma} - m_a| < 3\sigma_{m_{\gamma\gamma}}$. Here the characteristic width $\sigma_{m_{\gamma\gamma}}$ is the half width at half maximum (HWHM) of each signal peak fitted by the Breit-Wigner distribution. Background yields in each signal region are calculated from a background fit using a generalized gamma distribution,⁷ rather than the event counts to avoid large fluctuations.⁸ Thanks to the high reconstruction efficiency and low backgrounds, the 2σ exclusion limit on $\sigma(\mu^+\mu^- \rightarrow a + X\text{'s} \rightarrow \gamma\gamma + X\text{'s})$ reaches a few to 10 ab for TeV-scale axion at a $\sqrt{s} = 14$ TeV muon collider with $L = 10$ ab⁻¹.

A similar search strategy applies to other benchmark beam energies. Figure 8 shows the 2σ limits on $\sigma(\mu^+\mu^- \rightarrow a + X\text{'s} \rightarrow \gamma\gamma + X\text{'s})$ for various muon collider energy benchmarks. Limits for optimistic luminosity scenarios are lower, corresponding to higher sensitivities, due to the aggressive increase of the integrated luminosities. However, in a conservative scenario where the luminosity is fixed, the exclusion limit becomes weaker for a higher \sqrt{s} . The reason is that as m_a becomes much smaller than \sqrt{s} , the ALP produced picks up a greater boost in the beam direction on average. Consequently, the ALP decay products have larger $|\eta|$'s and are harder to reconstruct, reducing the signal efficiency. Meanwhile, SM backgrounds also increase moderately.

Analysis of $a \rightarrow \gamma Z(\rightarrow \ell\ell)$. We first select events containing exactly one isolated photon with $p_T(\gamma) > 350$ GeV, and two isolated leptons with the leading $p_T(\ell_1) > 100$ GeV. They must have opposite charges and same flavor. In addition, the dilepton invariant mass must satisfy $|m_{\ell\ell} - m_Z| < 5\Gamma_Z$ in which Γ_Z denotes the full Z width. All events containing

⁷We have also fitted the backgrounds with other common distributions, e.g., exponential, log-normal, and Pareto. It was found that generalized gamma distribution provides a similar if not better background fit in general.

⁸To improve the quality of the background fit, any event with $m_{\gamma\gamma} < 200$ GeV is dropped. Similar requirements are applied to all other VBF channel analyses.

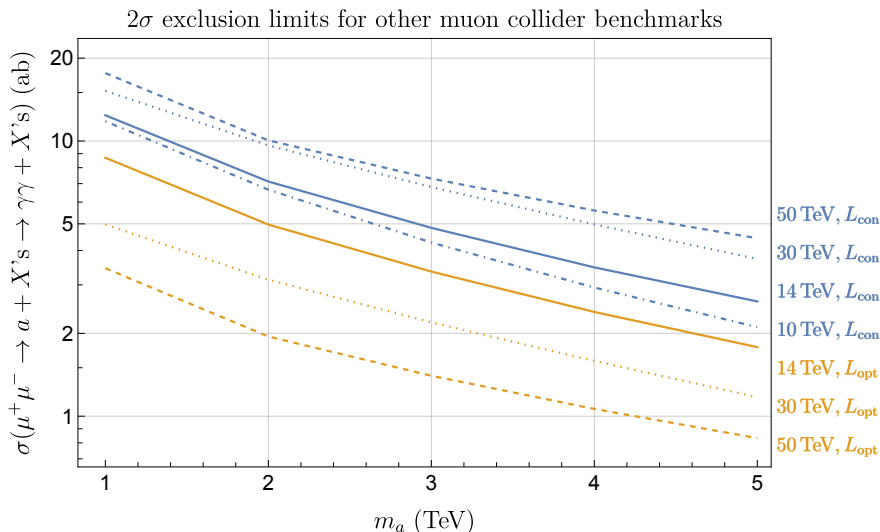


Figure 8. 2σ limits on $\sigma(\mu^+\mu^- \rightarrow a + X's \rightarrow \gamma\gamma + X's)$ at various muon collider benchmarks.

more than one extra charged track with $p_T > 10$ GeV are vetoed, excluding the lepton pairs discussed above and forward muons.

The $m_{\ell\ell\gamma}$ distributions of signal and background samples are shown in the upper right panel of figure 7. The signal peaks are also narrow, thanks to the high lepton and photon momentum resolutions. The signal region is defined as $|m_{\ell\ell\gamma} - m_a| < 3\sigma_{m\ell\ell\gamma}$, where $\sigma_{m\ell\ell\gamma}$ is the HWHM of the fitted signal peak. Although the signal reconstruction is straightforward and backgrounds are small, the 2σ upper limits on $\sigma(\mu^+\mu^- \rightarrow a + X's \rightarrow \gamma Z + X's)$ from this analysis are only of $\mathcal{O}(100)$ ab because of the low $\text{BR}(Z \rightarrow \ell\ell) \approx 7\%$.

Analysis of $a \rightarrow \gamma Z (\rightarrow jj)$. As the preselection, events with a hard isolated photon with $p_T(\gamma) > 350$ GeV are chosen. Note that the inclusive Valencia jet algorithm may also recognize the isolated photon as a hard fat jet. Thus, the requirement on the number of fat jets with $R = 1.2$ is relaxed to $N_J \in [1, 2]$. For those events with two J 's, it is necessary to ensure that one J comes from the hard photon. More specifically, we require one of them to have $m_J < 1$ GeV and contain at most one charged track. In addition, the leading isolated photon must be inside this J . This light J will be identified and treated as the photon instead of a hadronic J . The p_T of the remaining J (the true hadronic fat jet) must be greater than 200 GeV. To suppress $W\gamma$ backgrounds, its m_J is further required to be $\in [m_Z - 1.5\Gamma_Z, m_Z + 2.5\Gamma_Z]$. The m_J window is asymmetric around m_Z to suppress the $W \rightarrow jj$ background and provide adequate signal efficiencies. Numerically, the m_J selection above suppresses the $W\gamma$ backgrounds by a factor of $\sim 3\%$ at the price of keeping $\sim 40\%$ of the signal events. Similar W jet suppressions are also present in other analyses with hadronic Z decays.

The middle left panel of figure 7 shows the $m_{\gamma J}$ distributions of SM backgrounds and several signal benchmarks. As expected, the finite jet momentum resolution makes signal peaks much wider than previous cases. Since signal peaks around m_a are wide

and asymmetric, the optimized signal region is defined as $m_{\gamma J} \in [m_a - 1.5\sigma_{m_{\gamma J}}, m_a + 0.5\sigma_{m_{\gamma J}}]$ for each m_a . Here $\sigma_{m_{\gamma J}}$ is simply the standard deviation of each signal peak, as the asymmetric peaks make the fitted HWHM less meaningful. Benefited from the large $\text{BR}(Z \rightarrow jj)$, the 2σ limit on $\sigma(\mu^+\mu^- \rightarrow a + X\text{'s} \rightarrow Z\gamma + X\text{'s})$ from the $a \rightarrow \gamma Z(\rightarrow jj)$ analysis is at least a factor of two better than its $a \rightarrow \gamma Z(\rightarrow \ell\ell)$ counterpart, in the entire axion mass range we have studied.

Analysis of $a \rightarrow ZZ(\rightarrow 4\ell)$. For preselection, events with exactly four isolated leptons (excluding forward muons) with net electric charge zero are selected. Only those containing two same-flavor, opposite-sign lepton pairs both satisfying $|m_{\ell+\ell^-} - m_Z| \leq 5\Gamma_Z$ are kept. To suppress the soft SM $ZZ \rightarrow 4\ell$ backgrounds, $p_T(\ell_{1,2})$ must be greater than 300 and 200 GeV, respectively.

The resulting distributions of four lepton invariant mass, $m_{4\ell}$, are shown in the middle right panel of figure 7. The narrow signal peaks leave a simple signal region definition as $|m_{4\ell} - m_a| < 3\sigma_{m_{4\ell}}$, where $\sigma_{m_{4\ell}}$ is the HWHM of each fitted signal peak. The projected 2σ upper bounds on $\sigma(\mu^+\mu^- \rightarrow a + X\text{'s} \rightarrow ZZ + X\text{'s})$ are of $\mathcal{O}(1)$ fb, considerably weaker than other channels due to the small $\text{BR}(Z \rightarrow \ell\ell)^2 \sim 5 \times 10^{-3}$.

Analysis of $a \rightarrow Z(\rightarrow \ell\ell)Z(\rightarrow jj)$. We first preselect events containing exactly two opposite-sign same-flavor leptons (excluding forward muons) with $p_T(\ell_1) > 300$ GeV. Their invariant mass must satisfy $|m_{\ell+\ell^-} - m_Z| \leq 5\Gamma_Z$. A fat jet with $R = 1.2$ and $m_J \in [m_Z - 1.5\Gamma_Z, m_Z + 2.5\Gamma_Z]$ is also required to suppress the reducible VBS $W(\rightarrow jj)Z(\rightarrow \ell\ell)$ background. The signal and background $m_{J2\ell}$ distributions are shown in the lower left panel in figure 7.

The signal regions are defined by $m_{J2\ell} \in [m_a - 1.5\sigma_{m_{J2\ell}}, m_a + 0.5\sigma_{m_{J2\ell}}]$, where $\sigma_{m_{J2\ell}}$ is the standard deviation of each asymmetric signal peak. Benefited from the high $\text{BR}(Z \rightarrow jj)$, the sensitivity provided by this analysis is stronger than the corresponding $a \rightarrow ZZ(4\ell)$ one.

Analysis of $a \rightarrow ZZ(\rightarrow 4j)$. Events with exactly two fat jets ($R = 1.0$) with both $p_T(J_{1,2}) > 300$ GeV are selected. Both J 's must satisfy $m_J \in [m_Z - 2\Gamma_Z, m_Z + 5\Gamma_Z]$, where a wider m_J window is to ensure sufficient signal efficiencies. In addition, events containing more than one isolated leptons or photons with $p_T(\gamma/\ell) > 10$ GeV are vetoed.

The resulting distributions of two jet invariant mass are shown in the lower right panel of figure 7. Finally, the signal region is defined by $m_{2J} \in [m_a - 1.5\sigma_{m_{2J}}, m_a + 0.5\sigma_{m_{2J}}]$ for each m_a benchmark, where $\sigma_{m_{2J}}$ is the standard deviation of the fitted signal peak. Benefited from the sizeable $\text{BR}(Z \rightarrow jj)$, the analysis provides the leading constraint on ZZ detection channel with $\sigma(\mu^+\mu^- \rightarrow a + X\text{'s} \rightarrow ZZ + X\text{'s}) \lesssim \mathcal{O}(100)$ ab.

Analysis of $a \rightarrow WW(\rightarrow 4j)$. Similar to the previous analysis, events having two $R = 1.0$ fat jets with both $p_T(J_{1,2}) > 300$ GeV are chosen. To veto VBS $ZZ(\rightarrow 4j)$ and $ZW(\rightarrow 4j)$ backgrounds, each fat jet must satisfy $m_J \in [m_W - 5\Gamma_W, m_W + 3\Gamma_W]$. Also, no more than one isolated lepton or photon with $p_T(\gamma/\ell) > 10$ GeV is allowed.

The m_{JJ} distributions of signal and backgrounds of this analysis are shown in the last panel of figure 7. The signal region is defined by $m_{2J} \in [m_a - 1.5\sigma_{m_{2J}}, m_a + 0.5\sigma_{m_{2J}}]$

Model	$C_{WW} : C_{BB} : C_{BW}$	$\text{BR}(\gamma\gamma) : \text{BR}(\gamma Z) : \text{BR}(ZZ) : \text{BR}(WW)$
Default	1 : 1 : 0	7 : 6 : 22 : 65
$\gamma\gamma$ -phobic WW -phobic	0 : -1 : 1	0 : 63 : 37 : 0
$\gamma\gamma$ -phobic WW -philic 1	-2 : 1 : 1	0 : 11 : 23 : 65
$\gamma\gamma$ -phobic WW -philic 2	-1 : 1 : 0	0 : 20 : 16 : 65
$\gamma\gamma$ -phobic WW -philic 3	-1 : 0 : 1	0 : 5 : 31 : 64
$\gamma\gamma$ -philic WW -phobic	0 : 1 : 1	89 : 0 : 10 : 0
$\gamma\gamma$ -philic γZ -phobic	1 : 9.33 : 5.16	87 : 0 : 0 : 13
$\gamma\gamma$ -philic ZZ -phobic	1 : 3.33 : 4.33	60 : 12 : 0 : 29
Random model 1	1 : 1 : 1	15 : 12 : 11 : 61
Random model 2	2.34 : 1 : -2.26	0 : 3 : 33 : 64
Random model 3	1 : 2.69 : 4	53 : 15 : 0 : 32
Random model 4	1 : 1.54 : -7.04	11 : 13 : 58 : 19

Table 3. Details of ALP models in the model independence test. The branching fractions are normalized to percentages for readers’ convenience.

for each m_a benchmark, where σ_{m2J} is the standard deviation of the signal peak. This analysis is subject to the highest background rate, stemming from the large VBS WW cross section and $\text{BR}(W \rightarrow jj) \sim 70\%$. Meanwhile, the reconstructed signal peaks are also the widest among all analyses due to the broad m_J window. As a result, the constraint on $\sigma(\mu^+\mu^- \rightarrow a + X\text{'s} \rightarrow WW + X\text{'s})$ is only about (100 - 400) ab.

5.2 Model independence of the limits

The signal samples used to derive the limits summarized in figure 6 are simulated assuming $C_{BB} = C_{WW}$ and $C_{BW} = 0$. It is then necessary to check whether collider limits depend on a particular choice of ALP couplings. The concern is that with different combinations of couplings, the signal cut efficiency can be affected as the relative strengths in various production channels change. Even though multiple sets of EFT couplings allow ALPs to decay to the same final states, they can still affect the kinematics of produced ALPs as the importance of different initial states varies. Different initial states could lead to distinctive ALP momentum distributions, which may lead to different signal efficiencies in multiple analyses. For example, the ALPs produced in a $\gamma\gamma$ -philic model (i.e., the ALP coupling to $\gamma\gamma$ is significantly larger than the others after EWSB) will have a greater average longitudinal boost than the ALPs produced in a WW -philic model. Then, in the $\gamma\gamma$ -philic model, the average $|\eta|$ of the ALP decay products will also be larger, causing more difficult event reconstructions and reduced signal efficiencies.

To examine the generality of limits in section 5.1, we survey twelve models parameterized by $\{C_{BB}, C_{WW}, C_{BW}\}$ and compare their signal efficiencies. The descriptions of tested models are provided in table 3. The relative standard deviations of the signal efficiencies, i.e. the standard deviations of signal efficiencies divided by the mean signal efficiencies of

m_a (TeV)	$\gamma\gamma$	$\gamma Z(2\ell)$	$\gamma Z(2j)$	$ZZ(4\ell)$	$ZZ(2\ell 2j)$	$ZZ(4j)$	$WW(4j)$
1	5.0	10.5	16.4	11.0	12.8	3.2	5.5
2	1.7	10.7	4.7	10.3	11.6	2.1	4.1
3	1.3	12.4	1.4	11.5	10.8	2.4	2.3
4	1.0	12.7	1.8	13.3	12.5	2.3	3.3
5	1.3	14.5	2.7	13.3	14.7	2.7	1.2

Table 4. Percentage variations of the signal efficiencies for different search channels, based on models listed in table 3. The percentage variation is defined as the standard deviation of signal efficiencies divided by the mean signal efficiency of all the models we checked.

models considered, represent the analyses’ model dependence⁹ and are collected in table 4. One could see that the signal efficiencies only vary by $\lesssim 20\%$ for all the analyses in a wide range of m_a . Therefore, it is reasonable to take the results in figure 6 as model-independent limits for VBF produced pseudo-scalars at a 14-TeV muon collider. Although minor model-dependent variations remain, the exclusion limits in figure 6 provide a solid order-of-magnitude estimate for further studies.

We also investigate some different new physics models with similar final state topologies. Specifically, we simulate models of heavy scalars with couplings to the square of EW field strengths [131] and heavy vector boson with Chern-Simons couplings to EW gauge bosons [132, 133]. It turns out that signal efficiencies of these scenarios match with the ALP models within the 1σ range shown in table 4. Hence, it is very likely the upper limits in figure 6 apply to generic VBF-produced resonances decaying to EW boson pairs.

5.3 Analysis including forward muons

One possible way to suppress the SM backgrounds $\mu^+\mu^- \rightarrow (ZZ, Z\gamma, \gamma\gamma) + X$ ’s is to use the beam remnants’ information in the forward region. As discussed in section 4.2, these processes are always associated with W^+W^- initial state due to the SM gauge symmetry; therefore, X ’s in these processes are only muon neutrinos. They will be strongly suppressed if we require one or two energetic forward muons in the analysis. Nevertheless, additional forward-region cuts also reduce signal efficiencies. Typically, only $\mathcal{O}(20\%)$ of signal events will remain if two forward muons are required. In addition, limits obtained with forward-region cuts are more sensitive to the model variation since different $C_{VV'}$ couplings in the EFT significantly affect the identity and kinematics of the beam remnants.

Exclusion limits using only events with two forward muons are shown in figure 9, together with corresponding ones without requiring forward muons from figure 6. In the high- m_a region with small SM backgrounds already, results from analyses with no forward-muon cut take the lead as their signal efficiencies are higher. For some analyses, the forward-muon cuts could help improve the sensitivities slightly in the low- m_a region. For example, requiring two forward muons improves the $a \rightarrow \gamma Z(\rightarrow 2\ell)$ and $a \rightarrow \gamma\gamma$ analyses when $m_a \lesssim 3$ TeV. However, due to the small SM background and decreased signal efficiency

⁹Note that all VV' -phobic models are dropped when evaluating the model dependence of a particular VV' detection channel since there will be no such a decay.

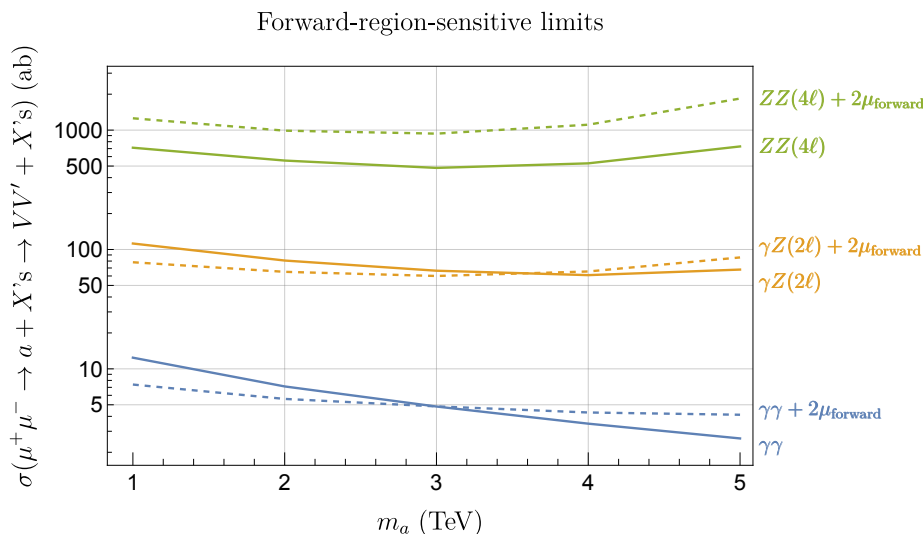


Figure 9. Comparison between 2σ exclusion limits with (dashed lines) and without the forward-muon number cut (solid lines, the same as figure 6) for $\sqrt{s} = 14$ TeV and $L = 10$ ab^{-1} . For cleaner reconstruction, we consider only channels with γ or $Z \rightarrow \ell^+\ell^-$ final states.

from the forward-muon cut, the limits from forward-muon-specific analyses are always weaker in the cleanest channel, i.e., $a \rightarrow ZZ(\rightarrow 4\ell)$.

6 Muon collider phenomenology: associated production channels

From figure 5, the associated production $\mu^+\mu^- \rightarrow Va$ appears to be a subdominant production channel for ALP at a muon collider. However, this channel has several interesting features that could potentially benefit the detection and calls for further study. Firstly, as is clear from figure 5, $\sigma(\mu\mu \rightarrow Va)$ is approximately a constant for a wide range of m_a up to the threshold $m_a \sim \sqrt{s}$, where VBF production is less efficient. Secondly, the major SM backgrounds for this channel are much smaller than the VBS backgrounds for the VBF channel [53, 87, 88]. Last but not least, the Va production kinematics is almost independent of the ALP model choices when $\sqrt{s} \gg m_Z$. The limits derived are thus approximately model independent.

6.1 Analysis of $a(\rightarrow \gamma\gamma)\gamma$

Inspired by the strong limit set by the VBF $a \rightarrow \gamma\gamma$ analysis, we focus on the $\mu^+\mu^- \rightarrow \gamma a(\gamma\gamma)$ process as the benchmark of associated production channels. The top panel of figure 10 displays the three photon invariant mass ($m_{3\gamma}$) distributions of the SM background and four representative m_a benchmarks, simulated according to the method discussed in section 4. Note that $m_{3\gamma}$ can be regarded as a proxy for the energy scale of the hard process, sometimes denoted as $\sqrt{\hat{s}}$ in ISR-related studies. Aside from the detector smearing effects, the ISR-induced low- $m_{3\gamma}$ tails are also obvious. The $m_{3\gamma}$ distribution of the $e^+e^- \rightarrow 3\gamma$ process in SM with ISR effects (using the electron ISR PDF provided by MADGRAPH5) is

m_a (TeV)	0.5	1	5	10
$\sigma_{\text{Exc, no ISR}}$ (ab)	28.6	33.5	44.4	74.8
$\sigma_{\text{Exc, with ISR}}$ (ab)	31.6	37.7	48.6	80.7

Table 5. Exclusion limit at 2σ level for the tri-photon search at a 14-TeV muon collider with 10 ab^{-1} data.

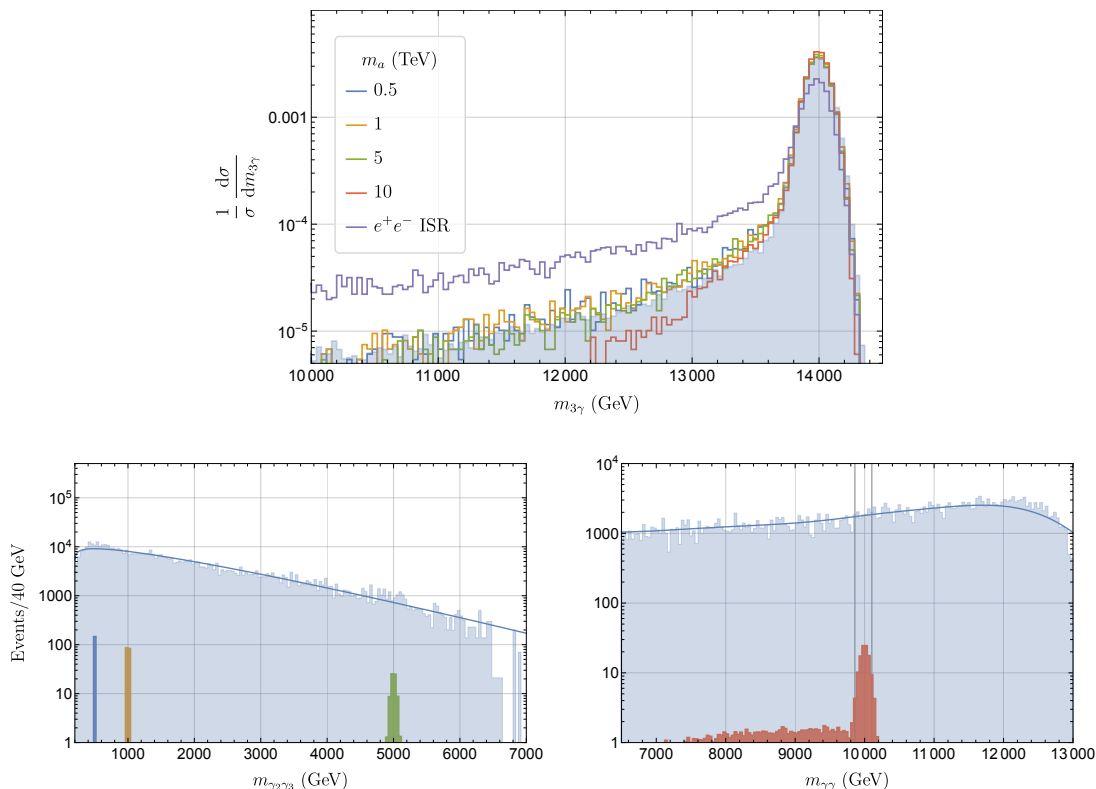


Figure 10. Distributions for $\gamma_a(\gamma\gamma)$ associated production at a 14-TeV muon collider with $L = 10 \text{ ab}^{-1}$ data. In all panels, the light blue histograms represent the SM backgrounds. *Top:* Normalized distributions of reconstructed $m_{3\gamma}$. The smooth $m_{3\gamma}$ tail below \sqrt{s} is the direct consequence of the undetected ISR photon. As a reference, $e^+e^- \rightarrow 3\gamma$ in SM with ISR at $\sqrt{s} = 14 \text{ TeV}$ is also plotted. *Bottom left:* Reconstructed ALP peaks (for $m_a = 0.5, 1, 5 \text{ TeV}$) and the SM background. Signals are scaled to the expected event counts when the production cross sections saturate the 2σ exclusion limits. *Bottom right:* reconstructed ALP peaks (for $m_a = 10 \text{ TeV}$) over the SM background. The two vertical lines stand for the *ad hoc* $m_{\gamma\gamma} \in [9.86, 10.1] \text{ TeV}$ window.

also shown for comparison. Its ISR-induced tail is even more significant as expected since electrons radiate more than muons.

Due to different kinematics, the analysis procedures for “light ALP” ($m_a \leq \sqrt{s}/2$) and “heavy ALP” ($m_a > \sqrt{s}/2$) benchmarks differ. For the light ALP case, we expect the associatively produced photon to be harder than the ALP decay products. We hence identify the two softer photons, i.e., γ_2 and γ_3 , as ALP decay products. The detector-level cuts $p_T(\gamma_1) \geq 4.5 \text{ TeV}$ and $\Delta R(\gamma_2\gamma_3) \leq 2$ are imposed to further suppress backgrounds with

small $m_{\gamma_2\gamma_3}$. To suppress the potential VBF background, we also impose a detector-level cut on the hard scattering scale $m_{3\gamma} \geq 0.9\sqrt{s}$. The lower left panel of figure 10 shows the signal and background distributions of $m_{\gamma_2\gamma_3}$. Only events with a resonance peak and $m_{\gamma_2\gamma_3} \in [m_a - 3\sigma_{m_{\gamma_2\gamma_3}}, m_a + 3\sigma_{m_{\gamma_2\gamma_3}}]$ are selected, where $\sigma_{m_{\gamma_2\gamma_3}}$ is the signal HWHM. The corresponding 2σ limits are shown in table 5. Notice that the continuous SM background level increases slowly when $m_{\gamma_2\gamma_3}$ decreases. Nevertheless, the larger background is compensated by the improved photon resolution (hence narrower $m_{\gamma_2\gamma_3}$ window) and higher signal efficiencies at low $m_{\gamma_2\gamma_3}$, rendering stronger constraints for smaller m_a . The limits obtained from samples without ISR are also shown for comparison. As expected, the ISR effect enhances the SM background rate and thus weakens the sensitivity slightly.

For the last benchmark ($m_a > \sqrt{s}/2$), we use an angular-separation-ordered strategy instead. This is inspired by the fact that the associated photon is less energetic than the heavy ALP, and the heavy resonance has a small boost inside the detector. Therefore, the photon pair from the ALP decay shall have a larger angular separation, and the signal peak shall be reconstructed by the photon pair with the largest ΔR . Our simulation also suggests that imposing $\min\{p_{T,\gamma_i}\} \geq 1 \text{ TeV}$, $\max_{i,j}\{\Delta R_{\gamma_i\gamma_j}\} \leq 5$, and $m_{3\gamma} \geq 0.9\sqrt{s}$ helps separate the signal and SM backgrounds. The $m_{\gamma\gamma}$ range of the signal region is determined *ad hoc* as $m_{\gamma\gamma} \in [9.86, 10.1] \text{ TeV}$ since the reconstructed signal peak shape is highly non-trivial. Nonetheless, this choice of the signal region carves out most of the resonance peak around m_a . Relevant $m_{\gamma\gamma}$ distributions are shown in the lower right panel of figure 10, and the 2σ exclusion limit for the $m_a = 10 \text{ TeV}$ benchmark is listed in table 5.

To conclude, limits from the $a(\rightarrow \gamma\gamma)\gamma$ analysis are weaker than those of the VBF diphoton channel by a factor of $\mathcal{O}(1 - 10)$ at a 14-TeV muon collider. A more detailed study on the ALP associated production channel with different final states and collider energies will be left to future work.

7 Constraining the EFT

With ALP production rates, branching ratios, and corresponding collider sensitivities known, one can put constraints on the ALP EFT in eq. (2.1). ALP production rates, computed by MADGRAPH5, are quadratic functions of the EFT couplings, and the decay rates in terms of the couplings are given in eq. (2.5). For collider sensitivities, we take the model-averaged limits discussed in section 5 and 6.

We first focus on the $C_{WW}-C_{BB}$ subspace with $C_{BW} = 0$, which is the most common scenario. The projected 2σ constraints for a 1-TeV ALP at a 14-TeV muon collider with $L = 10 \text{ ab}^{-1}$ is shown in figure 11. We show the parameter space covered by the four most sensitive VBF analyses and the 3γ analysis with associated production. While diphoton channel provides the strongest constraint on the signal rate, as demonstrated in figure 7, it is insensitive to the $C_{BB} \simeq -C_{WW}$ direction in the coupling plane, as shown in figure 11. This could be understood from eq. (2.4): along this direction, the ALP coupling to photon vanishes. Thus, we need other analyses to complement the diphoton resonance search. In particular, the $\gamma Z(2j)$ analysis provides the strongest sensitivity along the $C_{BB} \simeq -C_{WW}$ direction. The combined 2σ contour takes a butterfly shape, as visible in figure 11. In brief,

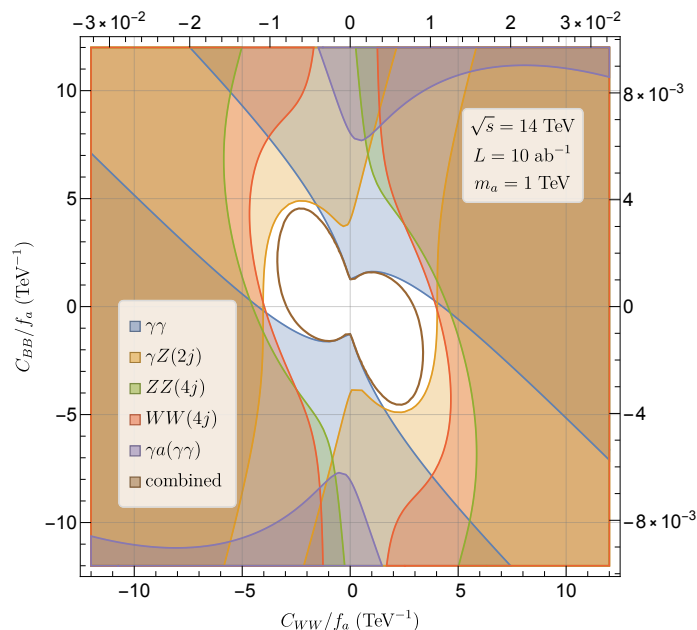


Figure 11. Projected 2σ constraint contours in the $C_{BB}/f_a - C_{WW}/f_a$ plane, for a 1-TeV ALP at a 14-TeV muon collider with $L = 10 \text{ ab}^{-1}$, assuming $C_{BW} = 0$. Detectable regions of different search channels are shaded with various colors. The solid brown butterfly-shaped curve around the center is the combined 2σ exclusion contour. Note that the bottom and left axes indicate values of the ALP couplings defined in eq. (2.1), while the top and right axes indicate values of these couplings absorbing the gauge couplings and loop factors following the convention in ref. [90].

a 14-TeV muon collider with $L = 10 \text{ ab}^{-1}$ data could probe C_{BB}/f_a and C_{WW}/f_a down to $\mathcal{O}(1) \text{ TeV}^{-1}$, which is at least one order of magnitude more sensitive than HL-LHC!

Other muon collider energy/luminosity benchmarks are considered in figure 12. The VBF $a \rightarrow \gamma\gamma$ exclusion limits for other running benchmarks are adopted from section 5.1 directly. For limits on other final states, they are approximated by rescaling $\sigma(\mu^+\mu^- \rightarrow a + X\text{'s} \rightarrow \gamma\gamma + X\text{'s})$ limits in figure 8. In particular, we assume that the ratios between the limits at any (\sqrt{s}, L) and $(14 \text{ TeV}, 10 \text{ ab}^{-1})$ in every search channel stay the same as those from di-photon limits. We checked this assumption by completing the analysis with full simulations in the next-to-most-contributing detection channel, $\gamma Z(jj)$ channel, with a variety of values of m_a and \sqrt{s} . The percentage variations between the rescaled exclusion limits and the exclusion limits with full simulations are $\lesssim 20\%$ for $\sqrt{s} = 10, 30 \text{ TeV}$ benchmarks and are at most $\sim 50\%$ for the $\sqrt{s} = 50 \text{ TeV}$ benchmark. When \sqrt{s} increases, the VBF ALP production rates increase, and consequently, the projected limits improve for both conservative and optimistic luminosity scenarios. The left panel of figure 12 shows the limits in the conservative luminosity scenario, while the right panel shows the ones in the optimistic scenario. By comparing both panels, the significant benefit from high integrated luminosities is evident.

Now we turn to the case in which $C_{BW} \neq 0$. We treat C_{BW} on an equal footing as C_{WW} and C_{BB} . To better understand the effects of a non-zero C_{BW} , we show exclusion

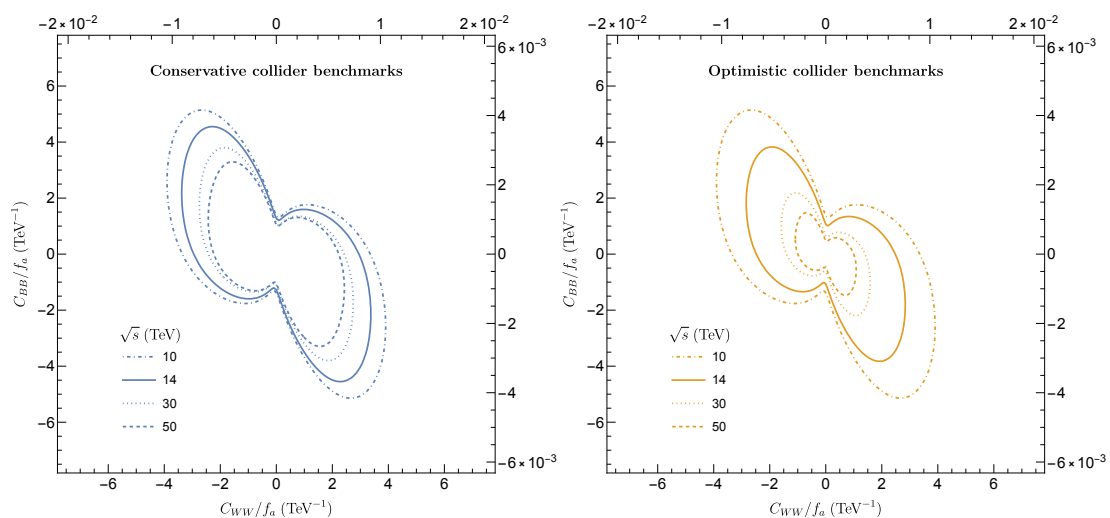


Figure 12. Projected 2σ constraint contours in the $C_{BB}/f_a - C_{WW}/f_a$ plane for different muon collider running benchmarks, assuming $C_{BW} = 0$. *Left:* L_{con} scenario. *Right:* L_{opt} scenario.

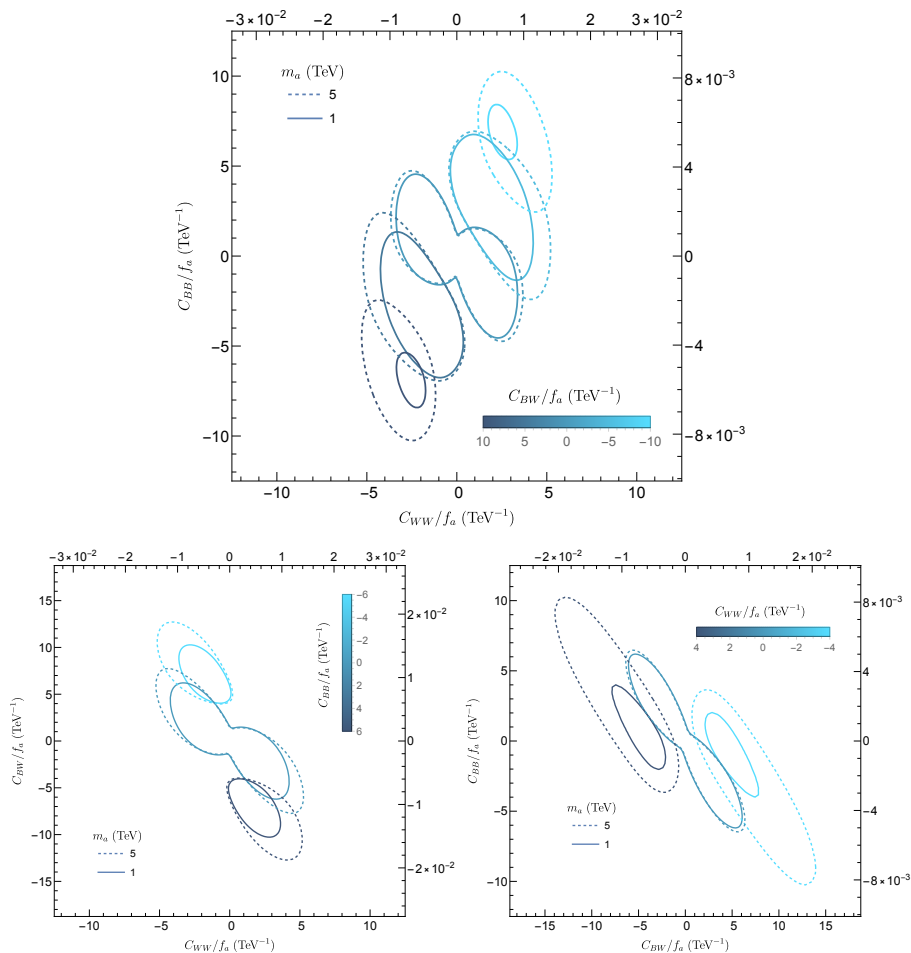


Figure 13. Combined projected 2σ constraints in various parameter subspaces at a 14-TeV muon collider with $L = 10 \text{ ab}^{-1}$. Solid contours: $m_a = 1 \text{ TeV}$. Dashed contours: $m_a = 5 \text{ TeV}$.

contours in the $C_{WW}-C_{BB}$ plane with $C_{BW}/f_a = \{-10, -5, 0, 5, 10\}$ TeV^{-1} , in the top panel of figure 13, with darker color indicating a higher value of C_{BW} . Near the origin of the plane, we obtain the familiar butterfly-shaped contour at the center of figure 11. As C_{BW} becomes non-zero, the constraints shift away from the origin due to the interference between operators. A larger $|C_{BW}|$ (so as the other couplings) also gives rise to a higher ALP production rate in general, leading to shrinking contours.

Similar analyses are implemented in the $C_{WW}-C_{BW}$ plane with $C_{BB}/f_a = \{-6, 0, 6\}$ TeV^{-1} , and in the $C_{BW}-C_{BB}$ plane with $C_{WW}/f_a = \{-4, 0, 4\}$ TeV^{-1} . It is noteworthy that the overall constraints are the weakest in a direction $C_{WW} : C_{BB} : C_{BW} \simeq -3 : -5 : 8$ that simultaneously renders a small ALP production rate and elusive ALP decays with $\text{BR}(a \rightarrow \gamma Z) \sim 5\%$ and $\text{BR}(a \rightarrow \gamma\gamma) < 1\%$.

8 Summary and outlook

In this article, we present a detailed analysis of TeV-scale ALP searches at a future high-energy muon collider. In particular, we focus on the searches that probe axion couplings to EW gauge bosons. The dominant ALP production channel at a muon collider is the VBF channel due to the high virtual EW gauge boson content of high-energy muon beams, analogous to the virtual gluon content of high-energy proton beams. We show that the $a \rightarrow \gamma\gamma$ final state enjoys the highest sensitivity, followed by the $Z\gamma$ mode. We also analyze the associated production channel with a lower rate but a smaller background, which turns out to be less sensitive compared to the VBF diphoton channel. Meanwhile, we show the projected constraints on various subspaces of ALP couplings in the EFT. A muon collider with energy $\gtrsim 10$ TeV could improve the sensitivity to ALP couplings by at least one order of magnitude compared to the HL-LHC, as well as expand the mass range of ALPs that could be searched for. We also demonstrate that the model-independent limits on ALP VBF production and decays are also applicable to generic BSM resonances coupling to EW gauge bosons, which could benefit future related studies. This study serves as another example of the great physics potentials of a high-energy muon collider.

There are several directions to expand the work. First, we base our study on the EFT of an ALP coupling to EW gauge bosons. While the EFT usually serves as a useful model-independent theoretical framework for experimental searches, the UV completions could predict (model-dependent) degrees of freedom and signals that could also be within reach. In particular, the couplings of heavy ALPs that current and future colliders are sensitive to are pretty large. They could be induced by heavy fermions carrying EW charges, which could be searched for as well. It will be useful to survey and classify possible UV completions of the ALP EFTs, check whether there are some generic predictions for collider phenomenologies, and compare the sensitivities to the associated UV degrees of freedom and to ALPs. Secondly, the beam remnants in the forward region in the VBF channels encode additional information. We carry out a crude estimate in this paper and find that they could lead to minor improvements of sensitivities in some mass range of ALPs. It will be useful to examine how to employ the forward-region information in a more sophisticated way, which could be valuable inputs for the design of a future muon

collider. Thirdly, despite of the weaker constraints on the EFT couplings from $\gamma a(\gamma\gamma)$ associated production channel, a generic Va channel is robust against model variations and systematic uncertainties. In this study, we use the most straightforward observables and analysis strategy on the benchmark tri-photon channel. Dedicated designs of observables and analyses in other Va channels can potentially complement the VBF search at a future muon collider and are worth further investigations.

Note added. After our preprint draft appeared on arxiv, we became aware of another work [134] on the same subject. While both papers share some components in analyses, they differ and complement each other in several aspects.

Acknowledgments

We thank Keping Xie for useful discussions. JF and LL are supported by the DOE grant DE-SC-0010010 and the NASA grant 80NSSC18K1010.

Open Access. This article is distributed under the terms of the Creative Commons Attribution License ([CC-BY 4.0](https://creativecommons.org/licenses/by/4.0/)), which permits any use, distribution and reproduction in any medium, provided the original author(s) and source are credited. SCOAP³ supports the goals of the International Year of Basic Sciences for Sustainable Development.

References

- [1] R.D. Peccei and H.R. Quinn, *Constraints imposed by CP conservation in the presence of instantons*, *Phys. Rev. D* **16** (1977) 1791 [[INSPIRE](#)].
- [2] R.D. Peccei and H.R. Quinn, *CP conservation in the presence of instantons*, *Phys. Rev. Lett.* **38** (1977) 1440 [[INSPIRE](#)].
- [3] S. Weinberg, *A new light boson?*, *Phys. Rev. Lett.* **40** (1978) 223 [[INSPIRE](#)].
- [4] F. Wilczek, *Problem of strong P and T invariance in the presence of instantons*, *Phys. Rev. Lett.* **40** (1978) 279 [[INSPIRE](#)].
- [5] J.E. Kim, *Weak interaction singlet and strong CP invariance*, *Phys. Rev. Lett.* **43** (1979) 103 [[INSPIRE](#)].
- [6] M.A. Shifman, A.I. Vainshtein and V.I. Zakharov, *Can confinement ensure natural CP invariance of strong interactions?*, *Nucl. Phys. B* **166** (1980) 493 [[INSPIRE](#)].
- [7] A.R. Zhitnitsky, *On possible suppression of the axion hadron interactions* (in Russian), *Sov. J. Nucl. Phys.* **31** (1980) 260 [*Yad. Fiz.* **31** (1980) 497] [[INSPIRE](#)].
- [8] M. Dine, W. Fischler and M. Srednicki, *A simple solution to the strong CP problem with a harmless axion*, *Phys. Lett. B* **104** (1981) 199 [[INSPIRE](#)].
- [9] J. Preskill, M.B. Wise and F. Wilczek, *Cosmology of the invisible axion*, *Phys. Lett. B* **120** (1983) 127 [[INSPIRE](#)].
- [10] M. Dine and W. Fischler, *The not so harmless axion*, *Phys. Lett. B* **120** (1983) 137 [[INSPIRE](#)].

- [11] L.F. Abbott and P. Sikivie, *A cosmological bound on the invisible axion*, *Phys. Lett. B* **120** (1983) 133 [INSPIRE].
- [12] M.K. Gaillard, M.B. Gavela, R. Houtz, P. Quilez and R. Del Rey, *Color unified dynamical axion*, *Eur. Phys. J. C* **78** (2018) 972 [arXiv:1805.06465] [INSPIRE].
- [13] J. Jaeckel, M. Jankowiak and M. Spannowsky, *LHC probes the hidden sector*, *Phys. Dark Univ.* **2** (2013) 111 [arXiv:1212.3620] [INSPIRE].
- [14] K. Mimasu and V. Sanz, *ALPs at colliders*, *JHEP* **06** (2015) 173 [arXiv:1409.4792] [INSPIRE].
- [15] J. Jaeckel and M. Spannowsky, *Probing MeV to 90 GeV axion-like particles with LEP and LHC*, *Phys. Lett. B* **753** (2016) 482 [arXiv:1509.00476] [INSPIRE].
- [16] S. Knapen, T. Lin, H.K. Lou and T. Melia, *Searching for axionlike particles with ultraperipheral heavy-ion collisions*, *Phys. Rev. Lett.* **118** (2017) 171801 [arXiv:1607.06083] [INSPIRE].
- [17] M. Bauer, M. Neubert and A. Thamm, *Collider probes of axion-like particles*, *JHEP* **12** (2017) 044 [arXiv:1708.00443] [INSPIRE].
- [18] N. Craig, A. Hook and S. Kasko, *The photophobic ALP*, *JHEP* **09** (2018) 028 [arXiv:1805.06538] [INSPIRE].
- [19] L. Lee, C. Ohm, A. Soffer and T.-T. Yu, *Collider searches for long-lived particles beyond the standard model*, *Prog. Part. Nucl. Phys.* **106** (2019) 210 [Erratum *ibid.* **122** (2022) 103912] [arXiv:1810.12602] [INSPIRE].
- [20] M. Bauer, M. Heiles, M. Neubert and A. Thamm, *Axion-like particles at future colliders*, *Eur. Phys. J. C* **79** (2019) 74 [arXiv:1808.10323] [INSPIRE].
- [21] A. Hook, S. Kumar, Z. Liu and R. Sundrum, *High quality QCD axion and the LHC*, *Phys. Rev. Lett.* **124** (2020) 221801 [arXiv:1911.12364] [INSPIRE].
- [22] A. Carmona, C. Scherb and P. Schwaller, *Charming ALPs*, *JHEP* **08** (2021) 121 [arXiv:2101.07803] [INSPIRE].
- [23] D. Wang, L. Wu, J.M. Yang and M. Zhang, *Photon-jet events as a probe of axionlike particles at the LHC*, *Phys. Rev. D* **104** (2021) 095016 [arXiv:2102.01532] [INSPIRE].
- [24] P. Agrawal et al., *Feebly-interacting particles: FIPs 2020 workshop report*, *Eur. Phys. J. C* **81** (2021) 1015 [arXiv:2102.12143] [INSPIRE].
- [25] A. Flórez et al., *Probing axionlike particles with $\gamma\gamma$ final states from vector boson fusion processes at the LHC*, *Phys. Rev. D* **103** (2021) 095001 [arXiv:2101.11119] [INSPIRE].
- [26] Y. Liu and B. Yan, *Searching for the axion-like particle at the EIC*, arXiv:2112.02477 [INSPIRE].
- [27] FCC collaboration, *FCC-ee: The Lepton Collider: Future Circular Collider Conceptual Design Report Volume 2*, *Eur. Phys. J. ST* **228** (2019) 261 [INSPIRE].
- [28] CEPC STUDY GROUP collaboration, *CEPC conceptual design report: volume 2. Physics & detector*, arXiv:1811.10545 [INSPIRE].
- [29] P. Bambade et al., *The International Linear Collider: a global project*, arXiv:1903.01629 [INSPIRE].

- [30] CLICDP, CLIC collaboration, *The compact linear collider (CLIC) — 2018 Summary report*, [arXiv:1812.06018](#) [INSPIRE].
- [31] FCC collaboration, *FCC-hh: the Hadron Collider: Future Circular Collider Conceptual Design Report Volume 3*, *Eur. Phys. J. ST* **228** (2019) 755 [INSPIRE].
- [32] A.N. Skrinsky and V.V. Parkhomchuk, *Cooling methods for beams of charged particles* (in Russian), *Sov. J. Part. Nucl.* **12** (1981) 223.
- [33] D. Neuffer, *Principles and applications of muon cooling*, *Conf. Proc. C* **830811** (1983) 481 [INSPIRE].
- [34] D. Neuffer, *Multi-TeV muon colliders*, *AIP Conf. Proc.* **156** (1987) 201 [INSPIRE].
- [35] V.D. Barger, M.S. Berger, J.F. Gunion and T. Han, *s-channel Higgs boson production at a muon muon collider*, *Phys. Rev. Lett.* **75** (1995) 1462 [[hep-ph/9504330](#)] [INSPIRE].
- [36] V.D. Barger, M.S. Berger, J.F. Gunion and T. Han, *Higgs boson physics in the s channel at $\mu^+\mu^-$ colliders*, *Phys. Rept.* **286** (1997) 1 [[hep-ph/9602415](#)] [INSPIRE].
- [37] C.M. Ankenbrandt et al., *Status of muon collider research and development and future plans*, *Phys. Rev. ST Accel. Beams* **2** (1999) 081001 [[physics/9901022](#)] [INSPIRE].
- [38] Snowmass, *Discussion forum on muon colliders*, (2021).
- [39] C. Aime et al., *Muon collider physics summary*, [arXiv:2203.07256](#) [INSPIRE].
- [40] T. Han, Y. Ma and K. Xie, *High energy leptonic collisions and electroweak parton distribution functions*, *Phys. Rev. D* **103** (2021) L031301 [[arXiv:2007.14300](#)] [INSPIRE].
- [41] NNPDF collaboration, *Illuminating the photon content of the proton within a global PDF analysis*, *SciPost Phys.* **5** (2018) 008 [[arXiv:1712.07053](#)] [INSPIRE].
- [42] R. B. Palmer, *Muon Colliders*, *Rev. Accel. Sci. Tech.* **07** (2014) 137 [INSPIRE].
- [43] J.P. Delahaye et al., *Muon colliders*, [arXiv:1901.06150](#) [INSPIRE].
- [44] J.-P. Delahaye et al., *Enabling intensity and energy frontier science with a muon accelerator facility in the U.S.: a white paper submitted to the 2013 U.S. community summer study of the division of particles and fields of the American Physical Society*, [arXiv:1308.0494](#) [INSPIRE].
- [45] J.-P. Delahaye et al., *A staged muon accelerator facility for neutrino and collider physics*, [arXiv:1502.01647](#) [INSPIRE].
- [46] R. Ryne et al., *Design Concepts for Muon-Based Accelerators*, in the proceedings of the 6th International Particle Accelerator Conference (IPAC'15), May 3–8, Richmond, U.S.A. (2015), DOI [INSPIRE].
- [47] K. Long, D. Lucchesi, M. Palmer, N. Pastrone, D. Schulte and V. Shiltsev, *Muon colliders to expand frontiers of particle physics*, *Nature Phys.* **17** (2021) 289 [[arXiv:2007.15684](#)] [INSPIRE].
- [48] MICE collaboration, *First demonstration of ionization cooling in MICE*, [arXiv:1806.01807](#) [INSPIRE].
- [49] MICE collaboration, *Recent results from the study of emittance evolution in MICE*, [arXiv:1806.04409](#)[INSPIRE].
- [50] MICE collaboration, *Demonstration of cooling by the muon ionization cooling experiment*, *Nature* **578** (2020) 53 [[arXiv:1907.08562](#)] [INSPIRE].

- [51] M. Antonelli, M. Boscolo, R. Di Nardo and P. Raimondi, *Novel proposal for a low emittance muon beam using positron beam on target*, *Nucl. Instrum. Meth. A* **807** (2016) 101 [[arXiv:1509.04454](#)] [[INSPIRE](#)].
- [52] M. Chiesa, F. Maltoni, L. Mantani, B. Mele, F. Piccinini and X. Zhao, *Measuring the quartic Higgs self-coupling at a multi-TeV muon collider*, *JHEP* **09** (2020) 098 [[arXiv:2003.13628](#)] [[INSPIRE](#)].
- [53] T. Han, D. Liu, I. Low and X. Wang, *Electroweak couplings of the Higgs boson at a multi-TeV muon collider*, *Phys. Rev. D* **103** (2021) 013002 [[arXiv:2008.12204](#)] [[INSPIRE](#)].
- [54] M. Chiesa, B. Mele and F. Piccinini, *Multi Higgs production via photon fusion at future multi-TeV muon colliders*, [arXiv:2109.10109](#) [[INSPIRE](#)].
- [55] M. Cepeda, S. Gori, V.M. Outchoorn and J. Shelton, *Exotic Higgs decays*, [arXiv:2111.12751](#) [[INSPIRE](#)].
- [56] J. Chen, T. Li, C.-T. Lu, Y. Wu and C.-Y. Yao, *Measurement of Higgs boson self-couplings through $2 \rightarrow 3$ vector bosons scattering in future muon colliders*, *Phys. Rev. D* **105** (2022) 053009 [[arXiv:2112.12507](#)] [[INSPIRE](#)].
- [57] MUON COLLIDER PHYSICS AND DETECTOR WORKING GROUP collaboration, *Higgs boson couplings at muon collider*, [PoS\(EPS-HEP2021\)619](#) [[INSPIRE](#)].
- [58] T. Han, Z. Liu, L.-T. Wang and X. Wang, *WIMPs at high energy muon colliders*, *Phys. Rev. D* **103** (2021) 075004 [[arXiv:2009.11287](#)] [[INSPIRE](#)].
- [59] R. Capdevilla, F. Meloni, R. Simoniello and J. Zurita, *Hunting wino and higgsino dark matter at the muon collider with disappearing tracks*, *JHEP* **06** (2021) 133 [[arXiv:2102.11292](#)] [[INSPIRE](#)].
- [60] A.D. Medina, N.I. Mileo, A. Szykman and S.A. Tanco, *The elusive muonic WIMP*, [arXiv:2112.09103](#) [[INSPIRE](#)].
- [61] G.-y. Huang, S. Jana, F.S. Queiroz and W. Rodejohann, *Probing the $RK(*)$ anomaly at a muon collider*, *Phys. Rev. D* **105** (2022) 015013 [[arXiv:2103.01617](#)] [[INSPIRE](#)].
- [62] P. Asadi, R. Capdevilla, C. Cesarotti and S. Homiller, *Searching for leptoquarks at future muon colliders*, *JHEP* **10** (2021) 182 [[arXiv:2104.05720](#)] [[INSPIRE](#)].
- [63] G. Haghightat and M. Mohammadi Najafabadi, *Search for lepton-flavor-violating ALPs at a future muon collider and utilization of polarization-induced effects*, *Nucl. Phys. B* **980** (2022) 115827 [[arXiv:2106.00505](#)] [[INSPIRE](#)].
- [64] P. Bandyopadhyay, A. Karan and R. Mandal, *Distinguishing signatures of scalar leptoquarks at hadron and muon colliders*, [arXiv:2108.06506](#) [[INSPIRE](#)].
- [65] R. Capdevilla, D. Curtin, Y. Kahn and G. Krnjaic, *Discovering the physics of $(g-2)_\mu$ at future muon colliders*, *Phys. Rev. D* **103** (2021) 075028 [[arXiv:2006.16277](#)] [[INSPIRE](#)].
- [66] D. Buttazzo and P. Paradisi, *Probing the muon $g-2$ anomaly with the Higgs boson at a muon collider*, *Phys. Rev. D* **104** (2021) 075021 [[arXiv:2012.02769](#)] [[INSPIRE](#)].
- [67] W. Yin and M. Yamaguchi, *Muon $g-2$ at multi-TeV muon collider*, [arXiv:2012.03928](#) [[INSPIRE](#)].
- [68] R. Capdevilla, D. Curtin, Y. Kahn and G. Krnjaic, *No-lose theorem for discovering the new physics of $(g-2)_\mu$ at muon colliders*, *Phys. Rev. D* **105** (2022) 015028 [[arXiv:2101.10334](#)] [[INSPIRE](#)].

- [69] N. Chen, B. Wang and C.-Y. Yao, *The collider tests of a leptophilic scalar for the anomalous magnetic moments*, [arXiv:2102.05619](#) [INSPIRE].
- [70] T. Li, M.A. Schmidt, C.-Y. Yao and M. Yuan, *Charged lepton flavor violation in light of the muon magnetic moment anomaly and colliders*, *Eur. Phys. J. C* **81** (2021) 811 [[arXiv:2104.04494](#)] [INSPIRE].
- [71] R. Dermisek, K. Hermanek and N. McGinnis, *Di-Higgs and tri-Higgs boson signals of muon $g - 2$ at a muon collider*, *Phys. Rev. D* **104** (2021) L091301 [[arXiv:2108.10950](#)] [INSPIRE].
- [72] R. Capdevilla, D. Curtin, Y. Kahn and G. Krnjaic, *Systematically testing singlet models for $(g - 2)_\mu$* , *JHEP* **04** (2022) 129 [[arXiv:2112.08377](#)] [INSPIRE].
- [73] L. Di Luzio, R. Gröber and G. Panico, *Probing new electroweak states via precision measurements at the LHC and future colliders*, *JHEP* **01** (2019) 011 [[arXiv:1810.10993](#)] [INSPIRE].
- [74] D. Buttazzo, R. Franceschini and A. Wulzer, *Two paths towards precision at a very high energy lepton collider*, *JHEP* **05** (2021) 219 [[arXiv:2012.11555](#)] [INSPIRE].
- [75] S. Spor and M. Köksal, *Investigation of anomalous triple gauge couplings in $\mu\gamma$ collision at multi-TeV muon colliders*, [arXiv:2201.00787](#) [INSPIRE].
- [76] S. Chen, A. Glioti, R. Rattazzi, L. Ricci and A. Wulzer, *Learning from radiation at a very high energy lepton collider*, *JHEP* **05** (2022) 180 [[arXiv:2202.10509](#)] [INSPIRE].
- [77] E. Eichten and A. Martin, *The muon collider as a H/A factory*, *Phys. Lett. B* **728** (2014) 125 [[arXiv:1306.2609](#)] [INSPIRE].
- [78] N. Chakrabarty, T. Han, Z. Liu and B. Mukhopadhyaya, *Radiative return for heavy Higgs boson at a muon collider*, *Phys. Rev. D* **91** (2015) 015008 [[arXiv:1408.5912](#)] [INSPIRE].
- [79] D. Buttazzo, D. Redigolo, F. Sala and A. Tesi, *Fusing vectors into scalars at high energy lepton colliders*, *JHEP* **11** (2018) 144 [[arXiv:1807.04743](#)] [INSPIRE].
- [80] P. Bandyopadhyay and A. Costantini, *Obscure Higgs boson at colliders*, *Phys. Rev. D* **103** (2021) 015025 [[arXiv:2010.02597](#)] [INSPIRE].
- [81] W. Liu and K.-P. Xie, *Probing electroweak phase transition with multi-TeV muon colliders and gravitational waves*, *JHEP* **04** (2021) 015 [[arXiv:2101.10469](#)] [INSPIRE].
- [82] T. Han, S. Li, S. Su, W. Su and Y. Wu, *Heavy Higgs bosons in 2HDM at a muon collider*, *Phys. Rev. D* **104** (2021) 055029 [[arXiv:2102.08386](#)] [INSPIRE].
- [83] C. Sen, P. Bandyopadhyay, S. Dutta and A. KT, *Displaced Higgs production in Type-III seesaw at the LHC/FCC, MATHUSLA and muon collider*, *Eur. Phys. J. C* **82** (2022) 230 [[arXiv:2107.12442](#)] [INSPIRE].
- [84] W. Liu, K.-P. Xie and Z. Yi, *Testing leptogenesis at the LHC and future muon colliders: a Z' scenario*, *Phys. Rev. D* **105** (2022) 095034 [[arXiv:2109.15087](#)] [INSPIRE].
- [85] C. Cesarotti, S. Homiller, R.K. Mishra and M. Reece, *Probing new gauge forces with a high-energy muon beam dump*, [arXiv:2202.12302](#) [INSPIRE].
- [86] P. Bandyopadhyay, A. Karan and C. Sen, *Discerning signatures of seesaw models and complementarity of leptonic colliders*, [arXiv:2011.04191](#) [INSPIRE].
- [87] A. Costantini et al., *Vector boson fusion at multi-TeV muon colliders*, *JHEP* **09** (2020) 080 [[arXiv:2005.10289](#)] [INSPIRE].

- [88] H. Al Ali et al., *The muon Smasher's guide*, *Rept. Prog. Phys.* **85** (2022) 084201 [[arXiv:2103.14043](#)] [[INSPIRE](#)].
- [89] R. Franceschini and M. Greco, *Higgs and BSM physics at the future muon collider*, *Symmetry* **13** (2021) 851 [[arXiv:2104.05770](#)] [[INSPIRE](#)].
- [90] I. Brivio et al., *ALPs effective field theory and collider signatures*, *Eur. Phys. J. C* **77** (2017) 572 [[arXiv:1701.05379](#)] [[INSPIRE](#)].
- [91] G. Alonso-Álvarez, M.B. Gavela and P. Quilez, *Axion couplings to electroweak gauge bosons*, *Eur. Phys. J. C* **79** (2019) 223 [[arXiv:1811.05466](#)] [[INSPIRE](#)].
- [92] J. Quevillon and C. Smith, *Axions are blind to anomalies*, *Eur. Phys. J. C* **79** (2019) 822 [[arXiv:1903.12559](#)] [[INSPIRE](#)].
- [93] Q. Bonnefoy, L. Di Luzio, C. Grojean, A. Paul and A.N. Rossia, *The anomalous case of axion EFTs and massive chiral gauge fields*, *JHEP* **07** (2021) 189 [[arXiv:2011.10025](#)] [[INSPIRE](#)].
- [94] J. Quevillon, C. Smith and P.N.H. Vuong, *Axion effective action*, *JHEP* **08** (2022) 137 [[arXiv:2112.00553](#)] [[INSPIRE](#)].
- [95] J. Bonilla, I. Brivio, M.B. Gavela and V. Sanz, *One-loop corrections to ALP couplings*, *JHEP* **11** (2021) 168 [[arXiv:2107.11392](#)] [[INSPIRE](#)].
- [96] M. Srednicki, *Axion couplings to matter*, *Nucl. Phys. B Proc. Suppl.* **260** (1985) 689 [[INSPIRE](#)].
- [97] M.A. Buen-Abad, J. Fan, M. Reece and C. Sun, *Challenges for an axion explanation of the muon $g - 2$ measurement*, *JHEP* **09** (2021) 101 [[arXiv:2104.03267](#)] [[INSPIRE](#)].
- [98] CMS collaboration, *Search for high-mass diphoton resonances in proton-proton collisions at 13 TeV and combination with 8 TeV search*, *Phys. Lett. B* **767** (2017) 147 [[arXiv:1609.02507](#)] [[INSPIRE](#)].
- [99] 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** (2021) 136651 [[arXiv:2102.13405](#)] [[INSPIRE](#)].
- [100] ATLAS collaboration, *Search for heavy resonances decaying to a photon and a hadronically decaying Z/W/H boson in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Phys. Rev. D* **98** (2018) 032015 [[arXiv:1805.01908](#)] [[INSPIRE](#)].
- [101] ATLAS collaboration, *Combination of searches for heavy resonances decaying into bosonic and leptonic final states using 36 fb⁻¹ of proton-proton collision data at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Phys. Rev. D* **98** (2018) 052008 [[arXiv:1808.02380](#)] [[INSPIRE](#)].
- [102] CMS collaboration, *A multi-dimensional search for new heavy resonances decaying to boosted WW, WZ, or ZZ boson pairs in the dijet final state at 13 TeV*, *Eur. Phys. J. C* **80** (2020) 237 [[arXiv:1906.05977](#)] [[INSPIRE](#)].
- [103] ATLAS collaboration, *Search for heavy diboson resonances in semileptonic final states in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Eur. Phys. J. C* **80** (2020) 1165 [[arXiv:2004.14636](#)] [[INSPIRE](#)].
- [104] ATLAS collaboration, *Search for heavy resonances decaying to a Z boson and a photon in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Phys. Lett. B* **764** (2017) 11 [[arXiv:1607.06363](#)] [[INSPIRE](#)].

- [105] CMS collaboration, *Search for $Z\gamma$ resonances using leptonic and hadronic final states in proton-proton collisions at $\sqrt{s} = 13$ TeV*, *JHEP* **09** (2018) 148 [[arXiv:1712.03143](#)] [[INSPIRE](#)].
- [106] M. Cacciari, G.P. Salam and G. Soyez, *The catchment area of jets*, *JHEP* **04** (2008) 005 [[arXiv:0802.1188](#)] [[INSPIRE](#)].
- [107] E. Molinaro and N. Vignaroli, *Diphoton resonances at the LHC*, *Mod. Phys. Lett. A* **32** (2017) 1730024 [[arXiv:1707.00926](#)] [[INSPIRE](#)].
- [108] A. Manohar, P. Nason, G.P. Salam and G. Zanderighi, *How bright is the proton? A precise determination of the photon parton distribution function*, *Phys. Rev. Lett.* **117** (2016) 242002 [[arXiv:1607.04266](#)] [[INSPIRE](#)].
- [109] B. Fornal, A.V. Manohar and W.J. Waalewijn, *Electroweak gauge boson parton distribution functions*, *JHEP* **05** (2018) 106 [[arXiv:1803.06347](#)] [[INSPIRE](#)].
- [110] D. Buarque Franzosi et al., *Vector boson scattering processes: status and prospects*, *Rev. Phys.* **8** (2022) 100071 [[arXiv:2106.01393](#)] [[INSPIRE](#)].
- [111] CMS collaboration, *Projected Performance of an Upgraded CMS Detector at the LHC and HL-LHC: Contribution to the Snowmass Process*, in *Community Summer Study 2013: Snowmass on the Mississippi*, 7, 2013 [[arXiv:1307.7135](#)] [[INSPIRE](#)].
- [112] CMS collaboration, *A MIP timing detector for the CMS Phase-2 upgrade*, CERN-LHCC-2019-003 (2019) [[INSPIRE](#)].
- [113] P.T. Komiske, E.M. Metodiev, B. Nachman and M.D. Schwartz, *Pileup Mitigation with Machine Learning (PUMML)*, *JHEP* **12** (2017) 051 [[arXiv:1707.08600](#)] [[INSPIRE](#)].
- [114] J. Arjona Martínez, O. Cerri, M. Pierini, M. Spiropulu and J.-R. Vlimant, *Pileup mitigation at the Large Hadron Collider with graph neural networks*, *Eur. Phys. J. Plus* **134** (2019) 333 [[arXiv:1810.07988](#)] [[INSPIRE](#)].
- [115] CMS collaboration, *Measurements of the $pp \rightarrow W^\pm\gamma\gamma$ and $pp \rightarrow Z\gamma\gamma$ cross sections at $\sqrt{s} = 13$ TeV and limits on anomalous quartic gauge couplings*, *JHEP* **10** (2021) 174 [[arXiv:2105.12780](#)] [[INSPIRE](#)].
- [116] ATLAS collaboration, *Evidence for the production of three massive vector bosons with the ATLAS detector*, *Phys. Lett. B* **798** (2019) 134913 [[arXiv:1903.10415](#)] [[INSPIRE](#)].
- [117] CMS collaboration, *Observation of the Production of Three Massive Gauge Bosons at $\sqrt{s} = 13$ TeV*, *Phys. Rev. Lett.* **125** (2020) 151802 [[arXiv:2006.11191](#)] [[INSPIRE](#)].
- [118] CMS collaboration, *Search for the production of $W^\pm W^\pm W^\mp$ events at $\sqrt{s} = 13$ TeV*, *Phys. Rev. D* **100** (2019) 012004 [[arXiv:1905.04246](#)] [[INSPIRE](#)].
- [119] ATLAS collaboration, *Study of $WW\gamma$ and $WZ\gamma$ production in pp collisions at $\sqrt{s} = 8$ TeV and search for anomalous quartic gauge couplings with the ATLAS experiment*, *Eur. Phys. J. C* **77** (2017) 646 [[arXiv:1707.05597](#)] [[INSPIRE](#)].
- [120] J. Bonilla, I. Brivio, J. Machado-Rodríguez and J.F. de Trocóniz, *Nonresonant searches for axion-like particles in vector boson scattering processes at the LHC*, *JHEP* **06** (2022) 113 [[arXiv:2202.03450](#)] [[INSPIRE](#)].
- [121] D.R. Green, P. Meade and M.-A. Pleier, *Multiboson interactions at the LHC*, *Rev. Mod. Phys.* **89** (2017) 035008 [[arXiv:1610.07572](#)] [[INSPIRE](#)].

- [122] P. Azzi et al., *Report from Working Group 1: Standard Model Physics at the HL-LHC and HE-LHC*, *CERN Yellow Rep. Monogr.* **7** (2019) 1 [[arXiv:1902.04070](#)] [[INSPIRE](#)].
- [123] J. Alwall et al., *The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations*, *JHEP* **07** (2014) 079 [[arXiv:1405.0301](#)] [[INSPIRE](#)].
- [124] T. Sjöstrand, S. Mrenna and P.Z. Skands, *A Brief Introduction to PYTHIA 8.1*, *Comput. Phys. Commun.* **178** (2008) 852 [[arXiv:0710.3820](#)] [[INSPIRE](#)].
- [125] DELPHES 3 collaboration, *DELPHES 3, a modular framework for fast simulation of a generic collider experiment*, *JHEP* **02** (2014) 057 [[arXiv:1307.6346](#)] [[INSPIRE](#)].
- [126] M. Boronat, J. Fuster, I. Garcia, E. Ros and M. Vos, *A robust jet reconstruction algorithm for high-energy lepton colliders*, *Phys. Lett. B* **750** (2015) 95 [[arXiv:1404.4294](#)] [[INSPIRE](#)].
- [127] C.F. von Weizsacker, *Radiation emitted in collisions of very fast electrons*, *Z. Phys.* **88** (1934) 612 [[INSPIRE](#)].
- [128] E.J. Williams, *Nature of the high-energy particles of penetrating radiation and status of ionization and radiation formulae*, *Phys. Rev.* **45** (1934) 729 [[INSPIRE](#)].
- [129] R. Ruiz, A. Costantini, F. Maltoni and O. Mattelaer, *The Effective Vector Boson Approximation in high-energy muon collisions*, *JHEP* **06** (2022) 114 [[arXiv:2111.02442](#)] [[INSPIRE](#)].
- [130] G. Cowan, K. Cranmer, E. Gross and O. Vitells, *Asymptotic formulae for likelihood-based tests of new physics*, *Eur. Phys. J. C* **71** (2011) 1554 [*Erratum ibid.* **73** (2013) 2501] [[arXiv:1007.1727](#)] [[INSPIRE](#)].
- [131] H.-Y. Ren, L.-H. Xia and Y.-P. Kuang, *Model-independent Probe of anomalous heavy neutral Higgs bosons at the LHC*, *Phys. Rev. D* **90** (2014) 115002 [[arXiv:1404.6367](#)] [[INSPIRE](#)].
- [132] I. Antoniadis, A. Boyarsky, S. Espahbodi, O. Ruchayskiy and J.D. Wells, *Anomaly driven signatures of new invisible physics at the Large Hadron Collider*, *Nucl. Phys. B* **824** (2010) 296 [[arXiv:0901.0639](#)] [[INSPIRE](#)].
- [133] K. Bondarenko, A. Boyarsky, M. Ovchinnikov, O. Ruchayskiy and L. Shchutska, *Probing new physics with displaced vertices: muon tracker at CMS*, *Phys. Rev. D* **100** (2019) 075015 [[arXiv:1903.11918](#)] [[INSPIRE](#)].
- [134] T. Han, T. Li and X. Wang, *Axion-like particles at high energy muon colliders — A White paper for Snowmass 2021*, [arXiv:2203.05484](#) [[INSPIRE](#)].