

Global analysis of ALP-mediated multiboson production at the LHC

Fabian Esser^{1,*}, Alexandre Salas-Bernárdez^{2,†}, Veronica Sanz^{3,‡} and Maria Ubiali^{4,§}

¹*Institute of Particle and Nuclear Physics (IPNP), Faculty of Mathematics and Physics, Charles University Prague, V Holešovičkách 2, 180 00 Praha 8, Czech Republic*

²*Departamento Análisis Matemático y Matemática Aplicada and IPARCOS, Universidad Complutense de Madrid, Plaza de las Ciencias 3, 28040 Madrid, Spain*

³*Instituto de Física Corpuscular (IFIC), Universidad de Valencia-CSIC, E-46980 Valencia, Spain*

⁴*DAMTP, University of Cambridge, Wilberforce Road, Cambridge CB3 0WA, United Kingdom*



(Received 12 November 2025; accepted 16 March 2026; published 26 May 2026)

Axionlike particles (ALPs) provide a well-motivated framework for physics beyond the Standard Model, coupling to gauge bosons through dimension-five operators protected by an approximate shift symmetry. At the LHC, such interactions lead to distinctive signatures in multiboson production, where the ALP appears as an off-shell mediator rather than a narrow resonance. In this work, we present the first global analysis of ALP-mediated multiboson processes, combining measurements of diphoton, ZZ , W^+W^- , dijet, and vector-boson-fusion (VBF) final states. On the theory side, motivated from a UV perspective, we assume that the ALP couples only to the gauge sector of the SM and classify the ALP-multiboson vertices that directly modify collider observables. Using Run-2 LHC measurements, we extract bounds on the Wilson coefficients ($c_{\tilde{G}}$, $c_{\tilde{W}}$, $c_{\tilde{B}}$) that parametrize gluonic and electroweak ALP interactions. Our results show that the dijet channel dominates the sensitivity to ALP couplings and determines the limits on $c_{\tilde{G}}$, while diboson and VBF processes provide complementary constraints on the electroweak couplings. We further assess the validity of the effective field theory expansion at the multi-TeV scales probed by the data. This global study provides the most comprehensive picture to date of ALP-gauge interactions from multiboson production at the LHC and highlights the opportunities for significant improvements with future high-luminosity measurements.

DOI: 10.1103/x7f6-rwgy

I. INTRODUCTION

Axionlike particles (ALPs) are among the most compelling candidates for physics beyond the Standard Model (SM). While the original Peccei-Quinn axion was proposed as a solution to the strong CP problem [1–3], generic ALPs arise naturally in a broad class of ultraviolet completions, including string compactifications [4,5] and composite Higgs scenarios [6–8]. These pseudo-Nambu-Goldstone bosons inherit shift-symmetry protection and couple to the SM gauge bosons through dimension-five operators suppressed by a high scale f_a . An up-to-date overview can be found in Ref. [9], while a comprehensive public repository of constraints on ALP parameters—spanning high-energy,

nuclear, and atomic physics, as well as cosmology and astrophysics—is provided in Ref. [10].

At high-energy colliders, such couplings open up novel production mechanisms. In particular, ALP-mediated processes can lead to striking signatures involving multiple gauge bosons in the final state. Multiboson production is especially interesting for several reasons. First, these final states are rare within the SM and therefore constitute sensitive probes of new physics. Second, the structure of ALP-gauge couplings leads to irreducible contact interactions that directly generate multi-gauge-boson vertices. Third, measurements of vector boson scattering and multi-boson production have become increasingly precise at the Large Hadron Collider (LHC), providing an ideal testing ground for this framework.

Previous work has analyzed individual ALP off-shell signatures at the LHC, focusing on diphoton, massive diboson, or vector-boson-scattering processes [11–14], as well as ALP couplings to the top quark [15–19], gluons [13,16], and di-Higgs final states [20]. However, a comprehensive treatment of ALP-mediated multiboson production combining different channels has not yet been undertaken, or only partially as in Ref. [21], where

*Contact author: fabian.esser@matfyz.cuni.cz

†Contact author: alexsala@ucm.es

‡Contact author: veronica.sanz@uv.es

§Contact author: m.ubiali@damtp.cam.ac.uk

vector-boson-fusion channels were explored. Moreover, in much of the existing literature, assumptions are often made about the gluon coupling, since gluon-induced production is typically dominant whenever it is present. Yet such assumptions can obscure the interplay between different effective operators. A global approach is therefore essential: only by combining complementary observables can one disentangle the underlying parameter space and truly probe the range of scenarios that are consistent with data.

In this paper, we provide a systematic study of axion-mediated multiboson production at the LHC. Working in the linear ALP effective field theory, we classify the irreducible ALP-multiboson vertices, identify the relevant collider signatures, and confront them with existing LHC measurements. We present a global fit including diphoton, ZZ , W^+W^- , dijet, and vector-boson-fusion final states and reinterpret the bounds in terms of the linear ALP effective field theory (EFT) operators. Our analysis delivers the most complete picture to date of the constraints on ALP-gauge interactions from multiboson processes and highlights the prospects for future improvements at the high-luminosity LHC. Throughout this work, we focus on a gauge-ALP benchmark scenario, which allows us to isolate the irreducible multiboson signatures of ALP exchange without committing to a specific ultraviolet completion.

The paper is organized as follows. In Sec. II, we introduce the effective Lagrangian for ALPs, describe their couplings to the SM gauge bosons, and discuss possible UV origins of such interactions. Section III presents our collider analysis: we first study the individual production channels—diphoton, ZZ , W^+W^- , dijet, and vector-boson-fusion final states—and then combine them into a global fit. Section IV reports the limits obtained under different

statistical prescriptions (projection, profiling, marginalization) and discusses EFT validity. We summarize our findings and outline future directions in Sec. V. Finally, in Appendix A we discuss interference effects and we list the selection cuts for our signal simulation in Appendix B.

II. THEORETICAL SETUP

We work within the framework of the linear ALP effective field theory, in which the axionlike particle a is introduced as a pseudo-Nambu-Goldstone boson coupling to the SM gauge bosons through dimension-five operators. We consider an ALP coupled exclusively at leading order to the gauge sector of the SM, a scenario that we will motivate from an explicit UV completion in Sec. II A. In this scenario and at leading order, the relevant Lagrangian reads

$$\mathcal{L}_{\text{ALP}} = \frac{a}{f_a} [c_{\tilde{G}} G_{\mu\nu}^a \tilde{G}^{a\mu\nu} + c_{\tilde{W}} W_{\mu\nu}^i \tilde{W}^{i\mu\nu} + c_{\tilde{B}} B_{\mu\nu} \tilde{B}^{\mu\nu}], \quad (2.1)$$

where f_a denotes the ALP decay constant, $c_{\tilde{X}}$ the dimensionless Wilson coefficients, and $\tilde{X}_{\mu\nu} = \frac{1}{2} \epsilon_{\mu\nu\rho\sigma} X^{\rho\sigma}$ the dual field-strength tensor for $X = G, W, B$.

After electroweak symmetry breaking, these interactions translate into couplings of the ALP to gluons, photons, and massive electroweak bosons,

$$\mathcal{L}_{\text{ALP}} = -\frac{a}{f_a} (g_{agg} G_{\mu\nu} \tilde{G}^{\mu\nu} + g_{aWW} W_{\mu\nu} \tilde{W}^{\mu\nu} + g_{aZZ} Z_{\mu\nu} \tilde{Z}^{\mu\nu} + g_{a\gamma Z} F_{\mu\nu} \tilde{Z}^{\mu\nu} + g_{a\gamma\gamma} F_{\mu\nu} \tilde{F}^{\mu\nu}). \quad (2.2)$$

Explicitly, one obtains

$$\begin{aligned} g_{a\gamma\gamma} &= \frac{4}{f_a} (c_{\tilde{B}} \cos^2 \theta_W + c_{\tilde{W}} \sin^2 \theta_W), \\ g_{aZZ} &= \frac{4}{f_a} (c_{\tilde{W}} \cos^2 \theta_W + c_{\tilde{B}} \sin^2 \theta_W), \\ g_{aZ\gamma} &= \frac{4}{f_a} \sin \theta_W \cos \theta_W (c_{\tilde{W}} - c_{\tilde{B}}), \\ g_{aWW} &= \frac{4}{f_a} c_{\tilde{W}}, \quad g_{agg} = \frac{4}{f_a} c_{\tilde{G}}. \end{aligned} \quad (2.3)$$

In this work, we restrict ourselves to bosonic ALP couplings and neglect direct couplings to fermions. For most SM fermions, these operators are phenomenologically uninteresting at the LHC, as their effects are either suppressed or redundant. The exception is the top quark, whose large Yukawa coupling and role in loop-induced processes make it especially relevant. We have studied ALP-top interactions in previous works [15,16], and here we neglect them at tree level, but note that they can contribute at loop level. In the next section, we describe a UV scenario in which no direct couplings to fermions arise, consistent with the possibility that the ALP does not couple to the electroweak symmetry-breaking sector.

The interactions above generate contact interaction terms involving both two and three gauge bosons, as listed in

Table I. In our analysis, we focus on the two-boson vertices, which currently provide the strongest experimental sensitivity. Channels with three gauge bosons—akin to those targeted in anomalous quartic gauge coupling (aQGC) searches—are not yet competitive, but they are expected to gain relevance with the increased luminosity of Run 3 and beyond and would represent a natural extension of our present study.

A. Model building a gauge-ALP coupling

We consider an ALP a as a light degree of freedom emerging from the spontaneous breaking of an approximate global $U(1)_X$ at a scale f_a . The associated shift symmetry $a \rightarrow a + c$ (with c constant) protects the mass and

TABLE I. Irreducible tree-level ALP-multiboson vertices in the linear EFT. For the three-boson final states, we define $p_{WW\gamma} = p_{W^+} + p_{W^-} + p_\gamma$, $p_{WWZ} = p_{W^+} + p_{W^-} + p_Z$, and $p_{ggg} = p_{g1} + p_{g2} + p_{g3}$.

Vertex	Feynman rule (linear ALP EFT)
$a\gamma\gamma$	$-\frac{4i}{f_a}(\cos^2\theta_W c_{\tilde{B}} + \sin^2\theta_W c_{\tilde{W}})\epsilon^{\mu\nu\rho\sigma}(p_{\gamma 1})_\rho(p_{\gamma 2})_\sigma$
$a\gamma Z$	$\frac{2i}{f_a}\sin(2\theta_W)(c_{\tilde{B}} - c_{\tilde{W}})\epsilon^{\mu\nu\rho\sigma}(p_Z)_\rho(p_\gamma)_\sigma$
aZZ	$-\frac{4i}{f_a}(\sin^2\theta_W c_{\tilde{B}} + \cos^2\theta_W c_{\tilde{W}})\epsilon^{\mu\nu\rho\sigma}(p_{Z1})_\rho(p_{Z2})_\sigma$
aW^+W^-	$-\frac{4i}{f_a}c_{\tilde{W}}\epsilon^{\mu\nu\rho\sigma}(p_{W^+})_\rho(p_{W^-})_\sigma$
agg	$-\frac{4i}{f_a}c_{\tilde{G}}\delta_{ab}\epsilon^{\mu\nu\rho\sigma}(p_{g1})_\rho(p_{g2})_\sigma$
$a\gamma W^+W^-$	$\frac{4ie}{f_a}c_{\tilde{W}}\epsilon^{\mu\nu\rho\sigma}(p_{WW\gamma})_\sigma$
aZW^+W^-	$\frac{4ie}{f_a}\frac{\cos\theta_W}{\sin\theta_W}c_{\tilde{W}}\epsilon^{\mu\nu\rho\sigma}(p_{WWZ})_\sigma$
$aggg$	$\frac{4}{f_a}g_s c_{\tilde{G}} f_{abc}\epsilon^{\mu\nu\rho\sigma}(p_{ggg})_\sigma$

interactions of a , and anomalous couplings to the SM gauge fields arise when the $U(1)_X$ current has mixed anomalies with $SU(3)_C \times SU(2)_L \times U(1)_Y$. After integrating out heavy states responsible for the anomaly, the low-energy effective interaction takes the form given in Eq. (2.1) with coefficients determined by anomaly matching, which schematically can be written as $c_{\tilde{X}} \propto \text{Tr}[Q_X T_X^2]$ for heavy fermions carrying the global charge Q_X and gauge generators T_X . This structure is the hallmark of UV completions where the ALP couples to the SM only through gauge anomalies.

Specifically, let $U(1)_X$ be an approximate global symmetry spontaneously broken below the scale f_a , with a the corresponding pseudo-Nambu-Goldstone boson. Take a set of heavy vectorlike fermions Ψ carrying $U(1)_X$ charge q_X^Ψ and transforming under $SU(3)_C \times SU(2)_L \times U(1)_Y$ in representations R_3^Ψ , R_2^Ψ and hypercharge Y_Ψ . In the UV, the ALP couples derivatively to the $U(1)_X$ current,

$$\mathcal{L} \supset \frac{\partial_\mu a}{f_a} J_X^\mu, \quad J_X^\mu = \sum_\Psi q_X^\Psi \bar{\Psi} \gamma^\mu \Psi. \quad (2.4)$$

The $U(1)_X$ current is anomalous with the SM gauge groups,

$$\partial_\mu J_X^\mu = \sum_\Psi \left[\frac{g_s^2}{16\pi^2} 2q_X^\Psi T(R_3^\Psi) G_{\mu\nu}^a \tilde{G}^{a\mu\nu} + \frac{g^2}{16\pi^2} 2q_X^\Psi T(R_2^\Psi) W_{\mu\nu}^i \tilde{W}^{i\mu\nu} + \frac{g^2}{16\pi^2} 2q_X^\Psi \mathcal{Y}_\Psi^2 B_{\mu\nu} \tilde{B}^{\mu\nu} \right], \quad (2.5)$$

where $T(R)$ is the Dynkin index [$T(\mathbf{3}) = \frac{1}{2}$, $T(\mathbf{2}) = \frac{1}{2}$, $T(\mathbf{1}) = 0$] and $\mathcal{Y}_\Psi$ denotes the hypercharge with the convention $Q = T_3 + Y$.

Integrating out the heavy fermions as in Fig. 1 or, equivalently, performing a chiral rotation that removes




FIG. 1. Representative anomaly triangle generating a $aF\tilde{F}$ term in the effective Lagrangian with heavy fermions Ψ .

the derivative coupling $(\partial_\mu a/f_a)\bar{\Psi}\gamma^\mu\Psi$,¹ the Fujikawa Jacobian generates the Wess-Zumino term

$$\mathcal{L}_{\text{WZ}} = \frac{a}{f_a} \left[\frac{g_s^2}{16\pi^2} \mathcal{A}_3 G_{\mu\nu}^a \tilde{G}^{a\mu\nu} + \frac{g^2}{16\pi^2} \mathcal{A}_2 W_{\mu\nu}^i \tilde{W}^{i\mu\nu} + \frac{g^2}{16\pi^2} \mathcal{A}_1 B_{\mu\nu} \tilde{B}^{\mu\nu} \right], \quad \mathcal{A}_X \equiv 2 \sum_\Psi q_X^\Psi \mathcal{I}_X^\Psi, \quad (2.6)$$

with $\mathcal{I}_3^\Psi = T(R_3^\Psi)$, $\mathcal{I}_2^\Psi = T(R_2^\Psi)$, and $\mathcal{I}_1^\Psi = \mathcal{Y}_\Psi^2$. Comparing Eq. (2.6) with our low-energy parametrization in Eq. (2.1),

$$\mathcal{L}_{\text{ALP}} = \frac{a}{f_a} [c_{\tilde{G}} G\tilde{G} + c_{\tilde{W}} W\tilde{W} + c_{\tilde{B}} B\tilde{B}],$$

we identify

$$c_{\tilde{G}} = \frac{g_s^2}{16\pi^2} \mathcal{A}_3, \quad c_{\tilde{W}} = \frac{g^2}{16\pi^2} \mathcal{A}_2, \quad c_{\tilde{B}} = \frac{g^2}{16\pi^2} \mathcal{A}_1. \quad (2.7)$$

Equivalently, one may write the schematic relation stated in the text, $c_{\tilde{X}} \propto \text{Tr}[Q_X T_X^2]$, with the constants of proportionality fixed by Eq. (2.7).

For example, for a single Dirac fermion Ψ in $(R_3, R_2)_Y$ with charge q_X , we get

$$c_{\tilde{G}} = \frac{g_s^2}{16\pi^2} 2q_X T(R_3), \quad c_{\tilde{W}} = \frac{g^2}{16\pi^2} 2q_X T(R_2), \quad c_{\tilde{B}} = \frac{g^2}{16\pi^2} 2q_X Y^2. \quad (2.8)$$

If several vectorlike fermions contribute, the $\mathcal{A}_i$ simply add. Color-neutral heavy spectra can yield $c_{\tilde{G}} = 0$ while keeping $c_{\tilde{W}}, c_{\tilde{B}} \neq 0$; conversely, colored heavy states naturally

¹Start from $\mathcal{L} \supset \frac{\partial_\mu a}{f_a} \bar{\Psi} \gamma^\mu (q_X) \Psi$. Perform the field redefinition $\Psi \rightarrow \exp[i(q_X a/f_a)\gamma_5] \Psi$. This removes the derivative term but generates, via the anomalous Jacobian,

$$\Delta\mathcal{L} = \frac{a}{f_a} \sum_X \frac{g_X^2}{16\pi^2} 2q_X \mathcal{I}_X X_{\mu\nu} \tilde{X}^{\mu\nu},$$

which is precisely Eq. (2.6).

generate $c_{\tilde{G}} \neq 0$. Alignment among charges/representations can suppress particular linear combinations, leading to photophobiclike patterns where $g_{a\gamma\gamma} \propto c_{\tilde{B}} \cos^2 \theta_W + c_{\tilde{W}} \sin^2 \theta_W$ is small at leading order. See, e.g., Refs. [22,23] for discussions connecting UV charge assignments to the pattern of electroweak ALP couplings and string/composite realizations with Wess-Zumino-Witten origins [24,25]. Systematic classifications of anomaly-induced electroweak ALP couplings and their dependence on the heavy fermion spectrum can be found in Ref. [26].

In particular, a simple and economical realization employs vectorlike fermions Ψ that are charged under the SM gauge group and carry $U(1)_X$ charge, while all SM fermions and the Higgs boson are neutral under $U(1)_X$. When the heavy Ψ are integrated out, Eq. (2.1) is generated with calculable $c_{\tilde{X}}$, but there are no tree-level derivative couplings $(\partial_\mu a) \bar{f} \gamma^\mu \gamma_5 f$ to SM fermions and no portal to the electroweak symmetry-breaking (EWSB) sector. In particular, there is no a -Higgs boson mixing and no direct $a \bar{f} f$ or $a H^\dagger H$ operators at leading order. This “gauge-ALP” setup is the UV rationale for our working assumption in the analysis: we focus on bosonic couplings and neglect direct fermionic ones.

Different UV choices for the heavy spectrum and charge assignments select distinctive patterns among $(c_{\tilde{G}}, c_{\tilde{W}}, c_{\tilde{B}})$. For instance, if the heavy fermions are colored, one naturally obtains $c_{\tilde{G}} \neq 0$ together with electroweak couplings; if they are color neutral, gluonic interactions can be absent, while $c_{\tilde{W}}$ and/or $c_{\tilde{B}}$ remain sizable, leading to electroweak-dominated phenomenology. Alignment symmetries can further suppress particular combinations, such as the “photophobic” limit in which the $a\gamma\gamma$ coupling cancels at leading order while $aZ\gamma$, aZZ , and aW^+W^- remain nonzero, a construction explored in collider contexts [27]. These scenarios emphasize the importance of combining several multiboson channels rather than assuming a fixed hierarchy among the couplings.

The effective couplings to mass eigenstates after EWSB follow from Eq. (2.1) and electroweak mixing, Eq. (2.3), which directly control the irreducible contact vertices and the off-shell ALP exchange entering our multiboson analysis.

We emphasize that the gauge-ALP scenario adopted here should be understood as a well-defined benchmark rather than as the most general ALP effective theory. In generic ultraviolet completions, additional operators (such as ALP-fermion couplings or Higgs-portal interactions) may be generated at loop level. In the class of constructions discussed above, however, these effects are parametrically suppressed and do not play a leading role in the multiboson observables considered in this work, which are dominantly controlled by the anomaly-induced gauge couplings. They arise only radiatively and are typically suppressed by loop factors and/or additional heavy-mass insertions. As a result, their impact on the high-energy multiboson observables analyzed in this work is expected to be subleading.

From a phenomenological perspective, these operators primarily affect processes involving fermionic final states, Higgs production, or low-energy precision observables, rather than the purely bosonic channels that dominate the sensitivity of our global analysis. While they may play an important role in complementary probes of ALP physics, they do not qualitatively modify the interpretation of the multiboson constraints derived here, which are driven by off-shell ALP exchange through gauge interactions.

A concrete illustration of the phenomenological impact of such loop-induced operators arises at low energies. Even if direct ALP-fermion couplings are absent at tree level, gauge-ALP interactions generate them radiatively under renormalization-group evolution. At scales below the electroweak scale, these induced couplings can mediate flavor-changing processes such as $K \rightarrow \pi a$ or $B \rightarrow Ka$, contribute to anomalous magnetic moments and electric dipole moments, or lead to signals in beam-dump and fixed-target experiments. Those effects were discussed in a previous work [15] and are only relevant in a small mass parameter space.

These observables provide some of the most stringent constraints on light ALPs in complementary regions of parameter space. At the same time, they probe a different set of effective interactions than the high-energy multiboson processes considered in this work. As a result, while loop-induced ALP-fermion and Higgs-portal operators can be crucial for low-energy phenomenology, they do not qualitatively affect the interpretation of LHC multiboson data, which are dominated by off-shell ALP exchange through gauge couplings.

III. GLOBAL ANALYSIS OF ALP-MEDIATED MULTIBOSON PRODUCTION

An important feature of the multiboson signatures considered here is that the ALP is exchanged off shell. In contrast to standard resonant searches, where one looks for a narrow peak in an invariant-mass distribution, our analysis does not rely on kinematic cuts around the ALP mass m_a . Instead, the ALP enters as a virtual mediator that modifies the tails of SM distributions, as depicted in Fig. 2.

This distinction is crucial, as it allows us to probe light ALPs that would otherwise evade conventional resonance searches due to selection cuts. In particular, for $m_a \ll \sqrt{\hat{s}}$ with $\hat{s}$ the squared center-of-mass energy of the parton-level process, the ALP exchange leads to boosted gauge-boson final states, whose kinematic features can still be exploited experimentally even though no resonant structure is present.

Because of this off-shell character, the relevant observables are inclusive multiboson production rates and differential distributions, rather than resonance bumps. In practice, the absence of resonance cuts means that the sensitivity depends primarily on the effective couplings $c_{\tilde{X}}$ and only weakly on the ALP mass, provided $m_a \ll \sqrt{\hat{s}}$.




FIG. 2. Schematic parton-level topologies for multiboson production. (a) SM continuum (box/topology example). (b) Virtual ALP exchange with no resonant cut. Our analysis targets this nonresonant regime. Here $V, V' \in \{g, \gamma, Z, W\}$.

For example, at the parton level, the s -channel amplitude for $V_1 V_2 \rightarrow V_3 V_4$ through an intermediate ALP takes the schematic form

$$\mathcal{M}_a \sim \frac{1}{\hat{s} - m_a^2 + im_a \Gamma_a} \left(\frac{4}{f_a} c_{V_1 V_2} \epsilon_{\mu\nu\rho\sigma} p_1^\rho p_2^\sigma \epsilon_1^\mu \epsilon_2^\nu \right) \times \left(\frac{4}{f_a} c_{V_3 V_4} \epsilon_{\alpha\beta\lambda\kappa} k_1^\lambda k_2^\kappa \epsilon_3^\alpha \epsilon_4^\beta \right), \quad (3.1)$$

where Γ_a is the ALP decay width. Squaring and averaging over polarizations, one finds, up to angular factors,

$$|\mathcal{M}_a|^2 \propto \frac{(c_{V_1 V_2} c_{V_3 V_4})^2}{f_a^4} \frac{\hat{s}^4}{(\hat{s} - m_a^2)^2 + m_a^2 \Gamma_a^2}. \quad (3.2)$$

In the regime of interest, $m_a^2 \ll \hat{s}$, the propagator reduces to $1/\hat{s}^2$, and the explicit m_a and Γ_a dependence drop out. The scaling with $\hat{s}$ in the numerator of the amplitude can be traced back to the momentum factors in the $\epsilon_{\mu\nu\rho\sigma}$ tensors of the ALP-gauge vertices. This behavior is generic as long as $m_{V_i}^2, m_a^2 \ll \hat{s}$ and follows directly from the Feynman rules in Table I. The cross section therefore scales as

$$\hat{\sigma}(V_1 V_2 \rightarrow V_3 V_4) \propto \frac{(c_{V_1 V_2} c_{V_3 V_4})^2}{f_a^4} \hat{s}, \quad (3.3)$$

with the growth with $\hat{s}$ already discussed in Ref. [13]. We stress that this growth with $\hat{s}$ is a generic feature of ALP-mediated processes in the linear EFT and reflects the derivative nature of the aVV couplings. As a consequence, the sensitivity of the LHC measurements considered in this work is driven by the high-energy tails of kinematic distributions. The resulting bounds should therefore be interpreted as constraints on the effective couplings $c_{\tilde{X}}/f_a$ within the domain of validity of the EFT, rather than as exclusions of specific ultraviolet completions.

In the following subsections, we explore in detail the individual production channels, including $\gamma\gamma$, ZZ , W^+W^- , γZ , and multiboson final states and discuss the current LHC measurements that are most sensitive to each of them. Each channel probes a different linear combination of the underlying couplings $c_{\tilde{G}}$, $c_{\tilde{W}}$, $c_{\tilde{B}}$, and their complementarity

is essential for a consistent interpretation. In the next section, we then combine these ingredients and present the results of our global fit, which constrains the full parameter space of ALP-gauge interactions. Note that our global analysis combines different channels without requiring assumptions on a specific resonance region, making it complementary to direct ALP resonance searches.

In Table II, we list a compilation of experimental measurements in multiboson final states, the provided distributions, and their key features. We highlight in bold the analyses that we reinterpreted in terms of the signal processes detailed below.

In all the presented analyses, we set $m_a = 1$ MeV and $f_a = 1$ TeV to generate events. As discussed before, for $m_a \ll \sqrt{\hat{s}}$ the signal cross section is roughly independent of the ALP mass. Note that we are only sensitive to the ratios $\frac{c_{\tilde{G}}}{f_a}$, $\frac{c_{\tilde{W}}}{f_a}$, and $\frac{c_{\tilde{B}}}{f_a}$, cf. Eq. (3.2). Limits on $c_{\tilde{X}}$ in this article are reported for $f_a = 1$ TeV, but any other value can be obtained by using the simple rescaling $\frac{c_{\tilde{X}}}{f_a}$.

Throughout this work, ALP-mediated signal events are generated at leading order and consistently normalized across all channels. Since the ALP contribution enters as a deformation of the high-energy tails of kinematic distributions rather than as an absolute rate measurement, the sensitivity of the analysis is dominated by shape information and by experimental uncertainties, rather than by higher-order corrections to the overall signal normalization.

A. Diphoton final state

The diphoton channel provides one of the cleanest probes of ALP-gauge interactions at the LHC. We base our study on the ATLAS measurement of inclusive diphoton production at $\sqrt{s} = 13$ TeV with 139 fb^{-1} [30], for which the WiKi for physics analysis, the HepData entry, and the corresponding Rivet routine are listed in Table II. The most sensitive distribution is provided by the invariant mass of the photon pair $m_{\gamma\gamma}$ [30]. The differential cross section, the experimental uncertainty, and different SM background estimations are given in bins ranging from 13.5 up to 1020 GeV and one overflow bin up to 13 TeV in HepData.

TABLE II. List of experimental analyses used in this work. The searches and distributions highlighted in *italics* are included in our global analysis of ALP-mediated multiboson final states.

Final state	Experiment	$\sqrt{s}$ (TeV)	Luminosity (fb ⁻¹)	Observables	References and data
$\gamma\gamma$	CMS	7	5	$d\sigma_{\gamma\gamma}/dm_{\gamma\gamma}$	[28] HepData [29]
	ATLAS	13	139	$d\sigma_{\gamma\gamma}/dm_{\gamma\gamma}$	[30] Analysis HepData Rivet [31]
ZZ	CMS	13	137	$d\sigma/m_{ZZ}, d\sigma/p_T^l$	[32] Analysis HepData [33]
	ATLAS	13.6	29	$d\sigma/dm_{4l}$	[34] HepData [35]
W^+W^-	CMS	13	34.8	total cross section	[36] HepData [37]
	ATLAS	13	36.1	$d\sigma_{WW}/dm_{e\mu}$	[38] HepData Rivet [39]
Dijet (gg)	ATLAS	13	3.2	$d^2\sigma_{jj}/dm_{jj}dy^*$	[40] HepData Rivet [41]
	CMS	13	2.3	$d^2\sigma_{jj}/dm_{\text{jet}}dp_T$	[42] HepData Rivet [43]
	CMS	13	2.6	$1/\sigma d^2\sigma_{jj}/d\chi dm_{jj}$	[44] HepData Rivet [45]
	CMS	13	35.9	$1/\sigma d^2\sigma_{jj}/d\chi dm_{jj}$	[46] HepData [47]
$\gamma Z jj$ (VBF)	CMS	13	139	$d\sigma/dp_T^{\gamma}, d\sigma/dp_T^{\text{lead},l}, d\sigma/p_T^{\text{lead},j}, d^2\sigma/dy_{jj}dm_{jj}$	[48] HepData [49]
$ZZjj, W^+W^-jj$ (VBF)	ATLAS	13	140	Events/200 GeV in m_{VV} bins	[50] Additional material [51]

In order to compare the data and background with an ALP-mediated signal, we generate signal events for the process

$$pp \rightarrow a^* \rightarrow \gamma\gamma$$

with MadGraph5_aMC@NLO [52] at LO, using NNPDF4.0 [53] in the four-flavor scheme and the ALP_linear_UFO [54]. We choose a benchmark scenario with $c_{\tilde{G}} = 1$, $c_{\tilde{W}} = 1$, $c_{\tilde{B}} = 0$, and note that the distributions for any other choices of $(c_{\tilde{G}}, c_{\tilde{W}}, c_{\tilde{B}})$ can be easily obtained from the Feynman rules in Table I and Eq. (3.3). The event selection follows the fiducial cuts specified in Ref. [30]. We then construct the $m_{\gamma\gamma}$ distribution using the same binning and fiducial selection cuts as in the ATLAS analysis, normalizing the signal events to 139 fb⁻¹ with the cross section from the Les Houches Event (LHE) header. The SM background prediction from Sherpa, together with the measured data points and their statistical uncertainties, is taken from Ref. [30], allowing for a direct, bin-by-bin comparison. Figure 3 shows the signal contribution superimposed on the SM background and data, expressed in terms of the cross section per bin normalized by the bin width. The impact of the ALP is to modify the high-mass tail of the distribution, consistent with the off-shell character of the exchange discussed above.

From Eq. (3.2), the diphoton cross section depends on the product $g_{agg}g_{a\gamma\gamma}$. In terms of the underlying Wilson coefficients, this corresponds to

$$g_{agg}g_{a\gamma\gamma} \propto c_{\tilde{G}}(c_{\tilde{B}}\cos^2\theta_W + c_{\tilde{W}}\sin^2\theta_W). \quad (3.4)$$

The diphoton channel therefore constrains only this specific combination of $c_{\tilde{G}}, c_{\tilde{W}}, c_{\tilde{B}}$ and cannot disentangle the individual coefficients. Independent information from other channels is required to lift this degeneracy, as it will be shown in the global analysis. This single channel cannot

provide independent bounds on either coupling, but only on their product. This limitation highlights the importance of combining complementary channels.

B. ZZ final state

The four-lepton final state provides a clean probe of ZZ production, with excellent mass resolution and relatively low backgrounds. We base our study on the ATLAS measurement of $pp \rightarrow ZZ \rightarrow 4\ell$ at $\sqrt{s} = 13$ TeV with 29 fb⁻¹ [34]. The measured differential cross section in m_{4l} , the uncertainties, and the background simulation are taken from the HepData entry listed in Table II, with bins extending from 180 to 1200 GeV. The analysis requires both lepton pairs to fall in the Z-mass window $50 < m_{\ell\ell} < 150$ GeV and the four-lepton invariant mass




FIG. 3. Diphoton invariant-mass distribution. The signal is shown for $c_{\tilde{G}} = 1$, $c_{\tilde{W}} = 1$, and $c_{\tilde{B}} = 0$, compared to ATLAS data [30] and the SM background (BG) prediction from Sherpa.




FIG. 4. Differential cross section for the four-lepton invariant mass $m_{4\ell}$ in ZZ production. The signal is shown for $c_{\tilde{G}} = 1$, $c_{\tilde{W}} = 1$, $c_{\tilde{B}} = 0$, compared to ATLAS data [34] and the SM prediction.

in the range $180 < m_{4\ell} < 1200$ GeV, with additional fiducial cuts on lepton rapidity and separation.

Signal events for the process

$$pp \rightarrow a^* \rightarrow ZZ \rightarrow 4\ell$$

and the benchmark setting $c_{\tilde{G}} = 1$, $c_{\tilde{W}} = 1$, $c_{\tilde{B}} = 0$ are generated with MadGraph5_aMC@NLO, using the ALP_linear_UFO model and normalized to the integrated luminosity of 29 fb^{-1} with the cross section and acceptance derived from the LHE files. We reconstruct the distribution in the four-lepton invariant mass $m_{4\ell}$ using the same binning as the ATLAS measurement. Figure 4 shows the signal prediction compared to the ATLAS data and the SM background simulation. As expected from the off-shell ALP exchange, the main effect is a distortion of the high-mass tail of the $m_{4\ell}$ spectrum rather than a resonant structure.

C. W^+W^- final state

The W^+W^- final state provides another clean probe of electroweak ALP couplings. We base our analysis on the ATLAS measurement of $pp \rightarrow W^+W^- \rightarrow \ell^+\nu\ell^-\bar{\nu}$ with different flavors for ℓ^- and ℓ^+ at $\sqrt{s} = 13$ TeV with an integrated luminosity of 36.1 fb^{-1} [38]. The HepData entry linked in Table II gives the differential cross section in the invariant mass of the final state leptons $m_{\ell\mu}$, the experimental uncertainty, and the SM background simulation in bins from 55 to 1500 GeV. For comparison, we also note the more recent CMS analysis [55], which we have not included in our study due to insufficient public information to reliably add the ALP signal on top of the SM results.

Signal events for $c_{\tilde{G}} = 1$, $c_{\tilde{W}} = 1$, $c_{\tilde{B}} = 0$ are generated with MadGraph5_aMC@NLO, using the ALP_linear_UFO model for the process




FIG. 5. Differential cross section in the WW final state as a function of the diboson invariant mass. The signal is shown for $c_{\tilde{G}} = 1$, $c_{\tilde{W}} = 1$, and $c_{\tilde{B}} = 0$, compared to the ATLAS data and the SM background simulation in [38].

$$pp \rightarrow a^* \rightarrow W^+W^- \rightarrow \ell^+\nu\ell^-\bar{\nu}$$

for $\ell^+\ell^- \in \{e^+\mu^-, \mu^+e^-\}$. The events are analyzed with the same fiducial selection and m_{WW} binning as the ATLAS measurement, with normalization to 36.1 fb^{-1} using the cross section from the LHE file. The observable of interest, the invariant mass of the W -boson pair, is reconstructed from the dilepton system.

As in the ZZ case, the effect of an off-shell ALP is to distort the high-mass tail of the m_{WW} spectrum. Figure 5 shows the predicted signal superimposed on the ATLAS data and the SM background. This channel depends on the combination of gluon fusion into an ALP off-shell state (g_{agg}) to produce a WW final state, via

$$g_{aWW} = \frac{4}{f_a} c_{\tilde{W}}$$

and therefore provides a direct and complementary constraint on the previous channels probing the electroweak sector of the ALP EFT.

D. Dijet final state

The dijet channel probes the coupling of the ALP to gluons and provides a direct constraint on $c_{\tilde{G}}$. We base our analysis on the ATLAS measurement of the double-differential dijet cross section

$$\frac{d^2\sigma}{dm_{jj} dy^*}$$

in the invariant mass of the jet pair m_{jj} and the half-rapidity separation $y^* = \frac{1}{2}|y_{j1} - y_{j2}|$ at $\sqrt{s} = 13$ TeV with 3.2 fb^{-1} [40]. The ATLAS analysis divides the phase space into six regions of y^* as follows:

$$\begin{aligned}
R1: 0 < y^* < 0.5, & \quad R2: 0.5 < y^* < 1.0, & \quad R3: 1.0 < y^* < 1.5, \\
R4: 1.5 < y^* < 2.0, & \quad R5: 2.0 < y^* < 2.5, & \quad R6: 2.5 < y^* < 3.0.
\end{aligned}$$

The m_{jj} bins depend on the y^* region and extend up to several TeV in the most central region. The individual data points of the double-differential distribution are extracted from the HepData entry linked in Table II. Note that the ATLAS analysis does not provide any background estimation. We take the simulation for the SM background in the same binning from next-to-next-to-leading-order (NNLO) QCD predictions in the leading-color approximation, computed with the NNLOjet code [56,57]. The central factorization and renormalization scales were chosen as $\mu_F = \mu_R = m_{jj}$. These predictions were released [58] as interpolation grids in the APPLfast format through the Ploughshare website [59], converted to the PineAPPL format [60] and released along with the NNPDF public code [53,61]. We do not account for next-to-leading-order (NLO) electroweak corrections or photon-initiated contributions [62].

Because of the double-differential character of the provided data, the analysis comprises 136 bins in total, and taking into account only the diagonal uncertainties is not enough, as will become clear below. Thus, we utilize the full experimental covariance matrix including experimental correlated systematic uncertainties, which is imported from the NNPDF public code [53,61]. Moreover, given that the NNLO QCD theory predictions in the large m_{jj} region that most strongly constrain our bounds are dominated by parton distribution function (PDF) uncertainties [63,64], we explicitly checked the impact of the inclusion of the main theory uncertainty of the SM background on the bounds we obtain. To do so, we added the PDF uncertainty (including full correlations among the various bins due to the PDFs) via the addition of a PDF theory covariance matrix to the likelihood [65,66].

Signal events for $pp \rightarrow a^* \rightarrow gg$ and $c_{\tilde{G}} = 1$, $c_{\tilde{W}} = 0$, $c_{\tilde{B}} = 0$ are generated with MadGraph5_aMC@NLO and the ALP_linear_UFO model. Note that, since this channel depends only on $c_{\tilde{G}}$, the choice of $c_{\tilde{W}}$ and $c_{\tilde{B}}$ does not matter here. 10^6 events are produced, which populate the y^* regions unevenly: about half fall into the most central bin ($0 < y^* < 0.5$), while the most forward bin ($2.5 < y^* < 3$) contains only a few hundred events. This feature mirrors the expected kinematics of gluon fusion. We reconstruct the double-differential cross section in the same binning as the experimental analysis (see Appendix B).

To compare signal, data, and background, we rescale the SM background simulation to the data by applying a k factor in each y^* region such that the total number of background events matches the observed yield. This procedure compensates for missing higher-order corrections in the generator and is applied consistently to both signal and background.

Figure 6 shows the distributions in the six different y^* regions. As for the diphoton, ZZ , and W^-W^+ final states, the effect of the ALP-mediated signal becomes apparent in the tail of the invariant-mass distribution. It is more visible in the more central regions with smaller y^* simply because there are more signal events generated in these regions. For illustration purposes, we show the distributions with error bars corresponding to the diagonal entries of the covariance matrix. For the statistical analysis, however, we employ the full covariance.

The dependence of the χ^2 on $c_{\tilde{G}}$ is shown in Fig. 7. Here we define the test statistic as

$$\begin{aligned}
\chi^2(c_{\tilde{G}}, c_{\tilde{B}}, c_{\tilde{W}}) = & (N^{\text{data}} - N^{\text{th}}(c_{\tilde{G}}, c_{\tilde{B}}, c_{\tilde{W}}))^T \\
& \times \Sigma^{-1} (N^{\text{data}} - N^{\text{th}}(c_{\tilde{G}}, c_{\tilde{B}}, c_{\tilde{W}})), \quad (3.5)
\end{aligned}$$

where N^{data} denotes the vector of observed event counts in each bin, N^{th} the corresponding theory prediction, and Σ the covariance matrix of experimental uncertainties. The theory prediction is given by

$$N^{\text{th}} = N^{BG} + c_{\tilde{G}}^4 N^{\text{ALP}}(c_{\tilde{G}} = 1), \quad (3.6)$$

where $N^{\text{ALP}}(c_{\tilde{G}} = 1)$ denotes the number of signal events for the benchmark scenario. The scaling of the total cross section and thus the number of events with $c_{\tilde{G}}$ follows directly from Eq. (3.2). In the plots, the error bars shown correspond to the diagonal entries of Σ , while the fits are performed using the full covariance, which captures correlations across bins and can substantially affect the resulting limits.

Using the full covariance matrix, we find that values $c_{\tilde{G}} \gtrsim 0.4$ are excluded at the 2σ level. By contrast, an analysis based only on the diagonal entries would prefer a nonzero best-fit value around $c_{\tilde{G}} \simeq 0.4$ – 0.5 , illustrating the importance of accounting for correlated uncertainties. This highlights the importance of including bin-by-bin experimental correlations in the likelihood as soon as those are provided by experimentalists. On the other hand, by comparing the orange and the red curves in Fig. 7, we see that the inclusion of the fully correlated PDF uncertainty does not modify the bounds in any significant way. Rather counterintuitively, the inclusion of PDF uncertainties and PDF-induced correlations in the likelihood makes the bounds slightly tighter. This is due to the correlations among bins due to PDFs that have a negative sign. Given that the inclusion of the PDF uncertainties does not change our conclusions in any significant way, we do not include it systematically at this stage. Note that most of the other




FIG. 6. Double-differential distribution in m_{jj} and y^* for $pp \rightarrow a^* \rightarrow gg$, compared to ATLAS data [40] and the SM background. For better visibility, the signal is shown for $c_{\tilde{G}} = 3$, $c_{\tilde{B}} = c_{\tilde{W}} = 0$.




FIG. 7. χ^2 dependence on $c_{\tilde{G}}$ from the dijet double-differential spectrum. The red curve includes the full experimental covariance matrix, while the blue curve assumes uncorrelated bin uncertainties. The orange curve adds the PDF theory covariance matrix to the experimental covariance matrix.

measurements that we include in this analysis do not provide yet the experimental correlations that are much more relevant than any correlated (or uncorrelated) theory uncertainties. Hence an inclusion of further correlated theory uncertainties without the inclusion of all experimental correlations would not be consistent. We leave the systematic inclusion of correlations and full theory uncertainties to future work, as soon as experimentalists release tables with data and full breakdown of systematics.

E. Vector-boson-fusion $Z\gamma jj$ final state

Vector boson fusion provides sensitivity to ALP couplings in the electroweak sector and enables us to constrain




FIG. 8. Diagrams contributing to $Z\gamma jj$: Vector boson fusion with an s -channel intermediate ALP for $V_1 V_2 \in \{W^+ W^-, ZZ, Z\gamma, \gamma\gamma\}$ (left) and vector boson scattering with a t -channel intermediate ALP for $V_1, V_2 \in \{\gamma, Z\}$ (right).

the couplings in the $(c_{\tilde{W}}, c_{\tilde{B}})$ plane, independent of $c_{\tilde{G}}$. We consider the CMS measurement of $pp \rightarrow Z\gamma jj$ at $\sqrt{s} = 13$ TeV with 137 fb^{-1} [48], where the Z boson is reconstructed in the dilepton channel. The corresponding HepData entry is listed in Table II. Signal events are generated with MadGraph5_aMC@NLO for the process

$$pp \xrightarrow{\text{NP}=2} Z\gamma jj, \quad (3.7)$$

using the ALP_linear_UFO model and setting $c_{\tilde{G}} = 0$, $c_{\tilde{W}} = 1$, and $c_{\tilde{B}} = 0$. Note that it is important to set c_{Φ} , which governs the ALP-fermion couplings, to zero in MadGraph5_aMC@NLO. We use the syntax in Eq. (3.7), setting the new physics (NP) order $\text{NP} = 2$, because not all diagrams that contribute to the signal process contain an s -channel ALP, which distinguishes this process significantly from the final states discussed before. In addition to the typical vector-boson-fusion diagram with an s -channel ALP, there are contributions from vector boson scattering (VBS) with a t -channel ALP, cf. Fig. 8.

Combining all possible diagrams, this process consists of eight subprocesses. Since we do not know their relative frequency, we cannot rescale the cross section for different values of $c_{\tilde{W}}$ and $c_{\tilde{B}}$ as easily as for the previous analyses. Hence, we calculate the cross section in MadGraph5_aMC@NLO

looping over various combinations for values of $c_{\tilde{W}}$ and $c_{\tilde{B}}$ and fit a quartic polynomial of the form

$$\sigma_{Z\gamma jj}(c_{\tilde{W}}, c_{\tilde{B}}) = \sum_{m,n=0}^{4,n+m \leq 4} \alpha_{nm} c_{\tilde{W}}^n c_{\tilde{B}}^m, \quad (3.8)$$

which covers all possible contributions from vector-boson-fusion and vector-boson-scattering diagrams to the observed cross sections. With the fit at hand, the number of events for arbitrary $c_{\tilde{W}}$ and $c_{\tilde{B}}$ can be obtained from

$$N^{\text{ALP}}(c_{\tilde{W}}, c_{\tilde{B}}) = N^{\text{ALP}}(c_{\tilde{W}} = 1, c_{\tilde{B}} = 0) \times \frac{\sigma_{Z\gamma jj}(c_{\tilde{W}}, c_{\tilde{B}})}{\sigma_{Z\gamma jj}(c_{\tilde{W}} = 1, c_{\tilde{B}} = 0)} \quad (3.9)$$

with $N^{\text{ALP}}(c_{\tilde{W}} = 1, c_{\tilde{B}} = 0)$ and $\sigma_{Z\gamma jj}(c_{\tilde{W}} = 1, c_{\tilde{B}} = 0)$ the distribution and cross section, respectively, for the benchmark scenario.

The fiducial event selection follows the CMS analysis, requiring two isolated leptons consistent with a Z boson, one isolated photon, and at least two jets in the forward region. The signal events for $c_{\tilde{W}} = 1$ and $c_{\tilde{B}} = 0$ are normalized to 137 fb^{-1} using the cross section from MadGraph5_aMC@NLO, the acceptance estimated from the LHE files for this benchmark scenario and multiplying with the branching ratio of Z into a lepton pair, $\text{BR}(Z \rightarrow l^+ l^-)$. We further assume that the acceptance to phase space cuts does not change when rescaling the total cross section.

We analyze two differential distributions: the photon transverse momentum p_T^γ and the transverse momentum of the leading jet $p_T^{j,\text{lead}}$. The binning matches the CMS measurement, with p_T^γ spanning 40–1000 GeV and $p_T^{j,\text{lead}}$ spanning 30–800 GeV. Figure 9 shows the signal predictions compared to the CMS data and the SM background. The characteristic effect of the ALP is an enhancement in the high- p_T tails of both distributions, arising from the momentum dependence of the aVV vertices.




FIG. 9. Differential distributions in p_T^γ (left) and $p_T^{j,\text{lead}}$ (right) for the $Z\gamma jj$ final state. Signal is shown for $c_{\tilde{G}} = 1$, $c_{\tilde{W}} = 1$ and $c_{\tilde{B}} = 0$, compared to CMS data [48] and the SM background.




FIG. 10. χ^2 and exclusion limits from the γZjj final state in p_T^γ (left) and $p_T^{j,\text{lead}}$ (right).

Performing a χ^2 analysis for both distributions leads to the 1σ and 2σ exclusion plots shown in Fig. 10. The distribution in p_T^γ gives slightly stronger limits and thus we take this distribution into account for the combined χ^2 fit.

F. Vector-boson-fusion $ZZjj$ and W^+W^-jj final states

Vector-boson fusion (VBF) provides direct sensitivity to the electroweak ALP couplings $c_{\tilde{W}}$ and $c_{\tilde{B}}$ via off-shell ALP exchange in $VV'jj$ topologies. Different VV' states will probe a different combination of the two Wilson coefficients. Adding analyses for various combinations of VV' therefore constrains complementary directions in the $(c_{\tilde{B}}, c_{\tilde{W}})$ plane. The analysis in [50] measures the distributions of $VV'jj$ with $V, V' \in \{Z, W\}$.

The $VV'jj$ amplitudes receive contributions from both VBF and VBS with several ALP vertices, cf. Fig. 11, and an analytic closed form of the fiducial cross section in terms of $(c_{\tilde{W}}, c_{\tilde{B}})$ is therefore not available. Similar to $Z\gamma jj$, we proceed by scanning a grid in $(c_{\tilde{W}}, c_{\tilde{B}})$ with $c_{\tilde{G}} = 0$ and fitting for each channel a quartic polynomial

$$\begin{aligned}\sigma_{ZZjj}^{\text{fid}}(c_{\tilde{W}}, c_{\tilde{B}}) &= \sum_{m,n=0}^{4,n+m \leq 4} \alpha_{nm}^{(ZZ)} c_{\tilde{W}}^n c_{\tilde{B}}^m, \\ \sigma_{WWjj}^{\text{fid}}(c_{\tilde{W}}, c_{\tilde{B}}) &= \sum_{m,n=0}^{4,n+m \leq 4} \alpha_{nm}^{(WW)} c_{\tilde{W}}^n c_{\tilde{B}}^m,\end{aligned}\quad (3.10)$$

which captures the SM-ALP interference (quadratic terms) and the pure ALP contribution (quartic terms) from the process depicted in Fig. 11.

We consider three fiducial final states in which one boson (V) is reconstructed hadronically (two jets) and the other (V') leptonically:

- (i) $ZZjj$ with $Z \rightarrow \nu\bar{\nu}$ (0ℓ),
- (ii) $ZZjj$ with $Z \rightarrow \ell^+\ell^-$ (2ℓ), and
- (iii) W^+W^-jj with $W^\pm \rightarrow \ell^\pm\nu$ (1ℓ).

To avoid ambiguities in assigning the hadronic boson jets at generator level, we simulate the corresponding leptonic decay modes and then rescale the event yields with the appropriate branching ratios to obtain the correct fiducial cross sections. The events are subsequently normalized to 140 fb^{-1} using the MadGraph5_aMC@NLO cross section and the acceptance after fiducial selections.

We split events into “resolved” if $\Delta R_{jj} > 0.4$ and “merged” otherwise to account for boosted hadronic bosons reconstructed as either two resolved jets or a single fat jet. Note that the analyses in [50] do not provide any HepData and we have to digitize the measured distribution, the experimental uncertainty, and the SM background simulation from the additional figures linked in Table II. For each channel and each region (resolved and merged), we analyze the differential distribution in the diboson invariant mass m_{VV} , using the binning of the experimental plots we digitize (ten bins up to $m_{VV} \sim 2 \text{ TeV}$ in the resolved category, slightly fewer in the merged category depending on the analysis’s figure).




FIG. 11. Diagrams contributing to $ZZ\gamma jj$ and W^+W^-jj : Vector boson fusion with an s -channel intermediate ALP for $V_1 V_2 \in \{W^+W^-, ZZ, Z\gamma, \gamma\gamma\}$ (left) and vector boson scattering with a t -channel intermediate ALP for $V_1, V_2 \in \{\gamma, Z\}$ for $ZZjj$ and $V_1 V_2 \in \{W^+W^-\}$ for W^+W^-jj (right).

Background templates and the data points (with upper/lower uncertainty envelopes) are read from the digitized CSV files from the public experimental results. Error bars in the figures correspond to the diagonal uncertainties, while the fit uses those same bin-by-bin uncertainties (no published covariance matrix is available for these plots).

1. Data treatment and uncertainties

Since only figure overlays are publicly available for these VBF selections, we digitize the published spectra for the 0ℓ , 2ℓ ($ZZjj$) and 1ℓ ($WWjj$) final states and propagate the quoted bin uncertainties. The plots in Fig. 12

show the signal overlaid on the digitized data and background for both resolved and merged categories; error bars correspond to the diagonal uncertainties used later in the fit. As expected for off-shell ALP exchange, the sensitivity resides in the high- m_{VV} bins, with the merged 2ℓ category providing the single strongest constraint among the six subchannels.

2. Constraints in the $(c_{\tilde{B}}, c_{\tilde{W}})$ plane

For each subchannel and category, we build a binned $\chi^2(c_{\tilde{B}}, c_{\tilde{W}})$ using the digitized data, SM background, and the polynomial surrogate for the signal. Figure 13 displays




FIG. 12. Differential distributions in m_{VV} for $VVjj$ final states. One boson decays hadronically, reconstructed either as two resolved jets (left column) or one merged fat jet (right column), while the other boson decays leptonically: $Z \rightarrow \nu\bar{\nu}$ (0ℓ), $Z \rightarrow \ell^+\ell^-$ (2ℓ), $W^\pm \rightarrow \ell^\pm\nu$ (1ℓ). For illustration, we show the signal for $c_{\tilde{G}} = 0$, $c_{\tilde{W}} = 10$, and $c_{\tilde{B}} = 0$, with vanishing fermionic ALP couplings.




FIG. 13. χ^2 in the $(c_{\tilde{B}}, c_{\tilde{W}})$ plane for the six $VVjj$ categories: $Z \rightarrow 0\ell$ (top), $Z \rightarrow 2\ell$ (middle), $W \rightarrow 1\ell$ (bottom); resolved (left) vs merged (right).

the resulting contours; the 2ℓ merged selection dominates the sensitivity, with the 0ℓ resolved and 1ℓ merged providing complementary directions through their different acceptance and background compositions. Combining all six subchannels yields the global constraint shown in Fig. 14.

IV. GLOBAL FIT RESULTS

The sensitivity of each experimental channel to ALP effects depends not only on the statistical power of the measurement, but also on the number of free parameters varied in the fit. A limit obtained by varying a single coupling while fixing all others to zero is necessarily

stronger than one obtained in a simultaneous multiparameter fit, where flat directions and correlations can dilute the constraints. To quantify these effects consistently, we perform combined fits to the three Wilson coefficients $(c_{\tilde{G}}, c_{\tilde{W}}, c_{\tilde{B}})$, which control the ALP couplings to gluons and electroweak bosons. For definiteness, we fix the ALP decay constant to $f_a = 1$ TeV, noting that all results trivially rescale as $1/f_a$.

The global likelihood is then explored under different statistical prescriptions (projection, profiling, and marginalization) that correspond to different ways of treating the parameter space when extracting one-dimensional limits. The resulting exclusion contours can therefore differ significantly, and the comparison between them provides




FIG. 14. Combined χ^2 in the $(c_{\tilde{B}}, c_{\tilde{W}})$ plane from all $VVjj$ categories. The contour is dominated by the $ZZjj$ 2ℓ merged selection, with complementary information from the other channels.

insight into the robustness of the constraints and the complementarity among different final states.

Before presenting the results of the global fit, we emphasize that all constraints derived in this work should be interpreted as bounds on the effective couplings $c_{\tilde{X}}/f_a$. The choice $f_a = 1$ TeV is made for convenience and does not carry physical significance. Owing to the exact factorization of the ALP contribution, cf. Eq. (3.2), all limits can be trivially rescaled to any other value of f_a .

Note that a global on-shell ALP analysis was performed in Ref. [67]. Previous work based on the off-shell contributions from ALP mediation considered global analyses restricted to vector-boson-scattering observables [21]. Here, by including also diphoton, diboson, dijet, and VBF signatures, we provide, to the best of our knowledge, the first global treatment of ALP-mediated multiboson production at the LHC.

A. Global fits in two and three dimensions

Figure 15 shows the full three-dimensional (3D) fit including all channels. The dijet final state provides the dominant sensitivity to $c_{\tilde{G}}$ and tightly constrains this direction. For $c_{\tilde{G}} > 0$, the combination of diphoton, ZZ , and WW channels further restricts the electroweak coefficients $c_{\tilde{B}}$ and $c_{\tilde{W}}$. In contrast, when $c_{\tilde{G}} \simeq 0$ these channels lose sensitivity, and the inclusion of vector-boson-fusion topologies, such as γZjj and $VVjj$, becomes essential to close the parameter space and obtain finite limits.

To visualize the interplay between parameters, we present in Fig. 16 several two-dimensional slices of the full 3D likelihood. The top row shows projections onto the $(c_{\tilde{B}}, c_{\tilde{W}})$ plane for $c_{\tilde{G}} = 0$ (left) and $c_{\tilde{G}} = 0.3$ (right). The




FIG. 15. Three-dimensional global fit including all channels. The shaded regions indicate the 1σ and 2σ confidence levels. The dijet analysis dominates the constraint on $c_{\tilde{G}}$, while the electroweak boson channels constrain $c_{\tilde{B}}$ and $c_{\tilde{W}}$ once $c_{\tilde{G}}$ is fixed.

bottom row shows projections in the $(c_{\tilde{G}}, c_{\tilde{B}} = c_{\tilde{W}})$ and $(c_{\tilde{G}}, c_{\tilde{B}} = -\tan^2(\theta_W)c_{\tilde{W}})$ planes, respectively. The first case with $c_{\tilde{B}} = c_{\tilde{W}}$ corresponds to the scenario in which the ALP coupling to one photon and one Z boson, $g_{a\gamma Z}$, vanishes, meaning that the ALP can couple only to either two photons or two Z bosons in the neutral electroweak sector. The second case with $c_{\tilde{B}} = -\tan^2(\theta_W)c_{\tilde{W}}$ mimics the photophobic scenario, in which the ALP does not couple to a pair of photons, $g_{a\gamma\gamma} = 0$, cf. Table I. Note that $c_{\tilde{G}} \gtrsim 0.41$ is excluded at the 2σ level by the dijet channel.

Figure 17 shows the contours that one obtains from subsets of the full set of final states shown for the previously discussed projections in the $c_{\tilde{B}} = c_{\tilde{W}}$ (top) and $c_{\tilde{B}} = -\tan^2(\theta_W)c_{\tilde{W}}$ (bottom) planes. As expected, diboson final states provide the tightest constraints for $c_{\tilde{G}} \neq 0$ and $c_{\tilde{B}} = c_{\tilde{W}} \neq 0$ in both cases. Dijet final states depend only on $c_{\tilde{G}}$ and close the contour for $c_{\tilde{B}} = c_{\tilde{W}} = 0$. VBF final states, on the other hand, depend on $c_{\tilde{B}}$ and $c_{\tilde{W}}$ but not $c_{\tilde{G}}$ and close the limits for $c_{\tilde{G}} \rightarrow 0$.

B. Projected, profiled, and marginalized limits

Table III summarizes the numerical limits on $c_{\tilde{G}}$, $c_{\tilde{B}}$, and $c_{\tilde{W}}$ under three different statistical prescriptions:

- (i) Projected limits: obtained by slicing the 3D likelihood along the axis of a single coefficient while fixing the others to zero. These limits correspond to the 2σ threshold $\Delta\chi^2 = 8.02$ in three dimensions.
- (ii) Profiled limits: obtained by minimizing χ^2 with respect to the other two coefficients. This produces an effective 1D likelihood $\chi^2_{\text{prof}}(c_i)$ with 2σ corresponding to $\Delta\chi^2 = 4$. As expected, the profiled limits are stronger than the projected ones.




FIG. 16. Two-dimensional projections of the global fit. Top: $(c_{\tilde{B}}, c_{\tilde{W}})$ plane for $c_{\tilde{G}} = 0$ (left) and $c_{\tilde{G}} = 0.3$ (right). Bottom: projections onto the $(c_{\tilde{G}}, c_{\tilde{B}} = c_{\tilde{W}})$ and $(c_{\tilde{G}}, c_{\tilde{B}} = -\tan^2(\theta_W)c_{\tilde{W}})$ planes.

- (iii) Marginalized limits: defined as the extremal values of each coefficient on the 1σ and 2σ surfaces of the 3D fit, with the 2σ limit corresponding to $\Delta\chi^2 = 8.02$. They have a direct geometric interpretation and facilitate comparison with the projected case, although they are somewhat weaker due to the full volume of parameter space being taken into account.

As expected, the profiled bounds are tighter than the projected ones, since correlations with the other couplings are taken into account. The marginalized limits are somewhat weaker than the projected case, reflecting the volume of parameter space that remains allowed once all three coefficients vary simultaneously.

Among the three operators, $c_{\tilde{G}}$ is most strongly constrained, with $c_{\tilde{G}} \lesssim 0.4$ at the 2σ level, consistent with the dijet analysis that dominates the $c_{\tilde{G}}$ direction in Fig. 15. The WW and ZZ channels provide the leading sensitivity to $c_{\tilde{W}}$, while the bounds on $c_{\tilde{B}}$ remain weaker, consistent with the partial cancellations between $c_{\tilde{B}}$ and $c_{\tilde{W}}$ in the $a\gamma Z$ coupling, as seen in the 2D projections of Fig. 16.

Numerically, profiling improves the limits relative to projection by about 15% for $c_{\tilde{G}}$, 10% for $c_{\tilde{B}}$, and nearly 25% for $c_{\tilde{W}}$. This illustrates the importance of treating correlations correctly in a multiparameter fit. All results assume $f_a = 1$ TeV and rescale as $1/f_a$. The numbers quoted in the table therefore provide a compact summary of the exclusion regions shown in Figs. 15 and 16.

C. Leading new physics contributions and relative constraining power

Focusing on off-shell intermediate ALPs, all processes contributing to an observable $\mathcal{O}^a$ necessarily contain two insertions of new physics vertices, meaning that we can parametrize the amplitude by

$$\mathcal{A}^a = \mathcal{A}^{a,SM} + \sum_{ij} \mathcal{A}_{ij}^{a,ALP} c_i c_j \quad (4.1)$$

for $c_i, c_j \in \{c_{\tilde{G}}, c_{\tilde{B}}, c_{\tilde{W}}\}$. Then, the observable, e.g., the differential cross section, is calculated from




FIG. 17. Individual contributions from diboson (blue), dijet (red), and VBF (green) channels and combined limits (black), in the plane $c_{\tilde{B}} = c_{\tilde{W}}$ (top) and in the photophobic plane $c_{\tilde{B}} = -\tan^2(\theta_W) c_{\tilde{W}}$ (bottom).

$$\mathcal{O}^a = |\mathcal{A}^{a,SM}|^2 + 2\text{Re}\left[\sum_{ij} (\mathcal{A}^{a,SM})^* \mathcal{A}_{ij}^{a,ALP} c_i c_j\right] + \text{Re}\left[\sum_{ijkl} (\mathcal{A}_{ij}^{a,ALP})^* \mathcal{A}_{kl}^{a,ALP} c_i c_j c_k c_l\right] + \dots \quad (4.2)$$

In this article, we focus on the quadratic new physics contribution and do not consider any interference terms

TABLE III. Comparison of projected, profiled, and marginalized 1σ and 2σ limits for $c_{\tilde{G}}$, $c_{\tilde{B}}$, and $c_{\tilde{W}}$. All numerical values quoted for the Wilson coefficients $c_{\tilde{X}}$ correspond to the reference choice $f_a = 1$ TeV and should be interpreted as bounds on the ratios $c_{\tilde{X}}/f_a$.

	Projected		Profiled		3D marginalized	
	1σ	2σ	1σ	2σ	1σ	2σ
$c_{\tilde{G}}$	0.40	0.44	0.35	0.41	0.46	0.48
$c_{\tilde{B}}$	1.77	1.89	1.60	1.78	1.92	2.00
$c_{\tilde{W}}$	1.07	1.16	0.85	1.09	1.26	1.33

with the SM, a choice that we comment on in Appendix A. Under this assumption, the quartic term in c , i.e., the last term in Eq. (4.2), provides the dominant ALP contribution to the observables. In particular, the linear term in c_i and the Fisher matrix, which is commonly used to quantify the constraining power of different observables on each Wilson coefficient, are identical zero. Physically speaking, the minimum of the χ^2 distribution is found for all $c_i = 0$, but the slope and curvature around the minimum are zero, and the deviation from the minimum is a quartic polynomial in c_i . This also implies that the fit cannot conclusively fix the sign of any Wilson coefficient, a fact that is reflected by the exclusion plots being symmetric under $c_i \rightarrow -c_i$ in quadratic terms or $c_i \rightarrow -c_i, c_j \rightarrow -c_j$ in mixed terms. With

$$S_{ijkl}^a \equiv \text{Re}[(\mathcal{A}_{ij}^{a,ALP})^* \mathcal{A}_{kl}^{a,ALP}], \quad (4.3)$$

we define for each observable

$$I_{ijkl}^a = \sum_b S_{ijkl}^a \Sigma_{ab}^{-1} S_{ijkl}^b, \quad (4.4)$$

with the experimental covariance matrix Σ_{ab} between different observables, which is relevant, e.g., for a distribution in various bins when the correlation between the bins is reported. For uncorrelated observables, this reduces to

$$I_{ijkl}^a = \left(\frac{S_{ijkl}^a}{\sigma_a}\right)^2 \quad (4.5)$$

for the experimental uncertainty σ_a for observable $\mathcal{O}_a$. All observables considered in this article belong to one of the sectors $\gamma\gamma$, ZZ , WW , dijet, or VBF. For a given sector A , we define

$$I_{ijkl}^A = \sum_{a \in A} I_{ijkl}^a, \quad (4.6)$$

which holds for uncorrelated observables in each sector. This allows us to calculate the relative constraining power of a dataset A on a set of Wilson coefficients $\{c_i, c_j, c_k, c_l\}$ as

$$f_{ijkl} = \frac{I_{ijkl}^A}{\sum_B I_{ijkl}^B}. \quad (4.7)$$

These relative constraining powers for various combinations of Wilson coefficients are shown in Table IV. Since the new physics enters via quartic terms in the Wilson coefficients, the one-coefficient-at-a-time assumption does not reveal much new information: Diboson production processes, i.e., $\gamma\gamma$, ZZ , and WW , depend on both $c_{\tilde{G}}$ and a combination of $c_{\tilde{B}}$ and $c_{\tilde{W}}$. Hence, they cannot constrain any Wilson coefficient on its own and the individual constraints on $c_{\tilde{G}}$ exclusively come from dijet measurements and the individual constraints on $c_{\tilde{B}}$ and $c_{\tilde{W}}$ from

TABLE IV. The relative constraining power of the five datasets on combinations of Wilson coefficients according to Eq. (4.7). The first three rows show scenarios with individual Wilson coefficients, the next three rows the constraining power on pairs of Wilson coefficients, and the last row all three Wilson coefficients at a time.

$(c_{\tilde{G}}, c_{\tilde{B}}, c_{\tilde{W}})$	$\gamma\gamma$	ZZ	WW	Dijet	VBF
(1,0,0)	0	0	0	1	0
(0,1,0)	0	0	0	0	1
(0,0,1)	0	0	0	0	1
(1,1,0)	0.303	0.103	0	0.594	8.83×10^{-6}
(1,0,1)	9.02×10^{-6}	0.0662	0.931	0.003	1.25×10^{-6}
(0,1,1)	0	0	0	0	1
(1,1,1)	0.003	0.162	0.832	0.00233	1.09×10^{-6}

VBF. The next three rows constrain pairs of Wilson coefficients. $c_{\tilde{G}} = c_{\tilde{B}} = 1$ shows an interplay between $\gamma\gamma$ and dijet while $c_{\tilde{G}} = c_{\tilde{W}} = 1$ is dominated by WW measurements. Only for $c_{\tilde{B}} = c_{\tilde{W}} = 1$ VBF becomes crucial to break the degeneracy in $c_{\tilde{G}}$. The last row shows the importance of each dataset for $c_{\tilde{G}} = c_{\tilde{B}} = c_{\tilde{W}} = 1$, with the largest contribution coming from WW. This analysis puts the relative constraints into context and also highlights the interplay of the different channels to break degenerate directions.

D. EFT validity

Our analysis relies on an ALP EFT with dimension-five operators $(4/f_a)c_i a X \tilde{X}$. For off-shell exchange, the relevant hard scale is $\sqrt{\hat{s}} \simeq m_{VV}$ (or m_{jj} in dijets). A conservative way to assess EFT validity is to compare $\sqrt{\hat{s}}$ with the EFT cutoff estimated by naive dimensional analysis (NDA), $\Lambda_{\text{NDA}} \sim 4\pi f_a / \max |c_i|$. Requiring $\sqrt{\hat{s}}_{\text{max}} \ll \Lambda_{\text{NDA}}$ yields

$$|c_i| \simeq \frac{4\pi f_a}{\sqrt{\hat{s}}_{\text{max}}}.$$

With $f_a = 1$ TeV and the kinematic reaches used here (dibosons $m_{VV} \lesssim 0.5$ – 1.5 TeV, VBF $m_{VV} \lesssim 1$ TeV, dijets $m_{jj} \lesssim 8$ TeV), this translates into indicative EFT-consistency windows,

$$\begin{aligned} \text{dibosons } (\sim 1.5 \text{ TeV}): |c_i| &\simeq \frac{4\pi}{1.5} \approx 8.4, \\ \text{VBF } (\sim 1 \text{ TeV}): |c_i| &\simeq 4\pi \approx 12.6, \\ \text{dijets } (\sim 8 \text{ TeV}): |c_i| &\simeq \frac{4\pi}{8} \approx 1.6. \end{aligned}$$

Our 95% confidence-level bounds (Table III) lie comfortably within these NDA ranges: for example, $c_{\tilde{G}} \lesssim 0.4$ from dijets is well below the ~ 1.6 EFT ceiling implied by

$m_{jj} \sim 8$ TeV, and the electroweak coefficients remain $\mathcal{O}(1)$, far below their diboson/VBF ceilings. For an ultraconservative criterion, one may take the cutoff as $\Lambda \sim f_a / \max |c_i|$, which would instead require $|c_i| \lesssim f_a / \sqrt{\hat{s}}_{\text{max}}$ (giving $|c_i| \lesssim 0.67, 1.0, 0.125$ for the three channels above); our limits then remain safely within the diboson/VBF ranges, while dijets become closer to the edge. In practice, one can enforce EFT robustness by truncating the fit to bins with $\sqrt{\hat{s}} < \kappa f_a$ (e.g., $\kappa = 1$ – 2); we have checked that our conclusions are stable under this replacement of the high-energy tail by an explicit truncation.

Note that, rather than imposing a hard cut on the partonic center-of-mass energy $\hat{s}$, we retain the full experimental information and accompany our results with an explicit discussion of the energy scales probed by the most sensitive bins. Introducing a sharp truncation would be highly process dependent and would require additional assumptions on the structure of the ultraviolet completion. Our approach follows standard practice in global EFT analyses, where high-energy data are included but the resulting constraints are interpreted as bounds on effective operators, whose UV completion must restore perturbative unitarity above the probed energies. In the most sensitive channels, the dominant constraints arise from bins with typical invariant masses of order a few TeV, which we indicate explicitly in the discussion above.

In parallel to these EFT-consistency estimates based on the size of the expansion parameter, recent work [68] has derived substantially stronger constraints by applying partial-wave unitarity and positivity bounds to ALP-gauge amplitudes. These bounds rely on different theoretical assumptions—most notably the on-shell formalism and coupled-channel unitarity—and therefore are not directly comparable to our conservative NDA criteria; however, they provide a valuable and highly complementary perspective that we now incorporate and cite in the updated version.

V. CONCLUSIONS

We have presented a global analysis of ALP-mediated multiboson production at the LHC in an EFT with bosonic, dimension-five couplings. A key feature of our approach is that the ALP is exchanged off shell, so sensitivity arises from distortions of high-energy tails rather than resonant bumps. This makes the analysis largely insensitive to the ALP mass provided $m_a \ll \sqrt{s}$ and allows us to combine information across channels with different kinematics and backgrounds.

On the theory side, we assumed ALP couplings exclusive to the SM gauge sector and motivated this pattern from a UV perspective, then classified the irreducible ALP-multiboson vertices and mapped the setup onto the low-energy couplings $(c_{\tilde{G}}, c_{\tilde{W}}, c_{\tilde{B}})$. On the phenomenology

side, we confronted the EFT with a broad set of LHC measurements: inclusive $\gamma\gamma$, $ZZ \rightarrow 4\ell$, $W^+W^- \rightarrow \ell\ell\nu\nu$, double-differential dijets, and VBF topologies ($Z\gamma jj$ and $VVjj$ with resolved/merged categories). Each channel probes a different linear combination of Wilson coefficients: dijets fix $c_{\tilde{G}}$ through g_{agg} ; WW and ZZ , in turn, constrain $c_{\tilde{W}}$ via g_{aWW} and g_{aZZ} ; and $Z\gamma jj/VVjj$ access both $c_{\tilde{W}}$ and the difference $(c_{\tilde{W}} - c_{\tilde{B}})$ through $aZ\gamma$ and aZZ interference. Their complementarity is essential to break flat directions.

We then performed a simultaneous three-parameter fit to $(c_{\tilde{G}}, c_{\tilde{B}}, c_{\tilde{W}})$ at $f_a = 1$ TeV (with all limits rescaling as $1/f_a$), reporting projected, profiled, and marginalized constraints. Numerically (Table III), the dijet spectrum yields the strongest bound on the gluonic coefficient, $c_{\tilde{G}} \lesssim 0.4$ at 2σ , while WW/ZZ lead the constraints on $c_{\tilde{W}}$, and the VBF channels close the $(c_{\tilde{W}}, c_{\tilde{B}})$ plane, particularly in the photophobiclike direction $c_{\tilde{W}} \simeq c_{\tilde{B}}$. Profiling improves the limits by $\sim 15\%$ ($c_{\tilde{G}}$), $\sim 10\%$ ($c_{\tilde{B}}$), and nearly $\sim 25\%$ ($c_{\tilde{W}}$) relative to simple projection, underscoring the need to treat correlations consistently in a multiparameter fit. We also checked EFT robustness using the relevant hard scales (m_{VV} , m_{jj}): our bounds lie comfortably within the NDA windows for diboson and VBF, and remain below the conservative ceiling implied by the $m_{jj} \sim 8$ TeV reach of the dijet analysis.

Several improvements are straightforward. On the experimental side, fuller use of unfolded covariance matrices (where available), inclusion of the highest-luminosity WW/ZZ measurements, and additional electroweak channels ($Z\gamma$, $W\gamma$, $\gamma\gamma$ +jet, tribosons, and polarized VBS) will sharpen the $(c_{\tilde{W}}, c_{\tilde{B}})$ constraints. On the theory side, incorporating NLO QCD electroweak corrections and matched predictions for the most sensitive tails would reduce theory systematics; extending the fit to include loop-induced fermionic effects (notably the top) will turn our discussion of c_t into a quantitative four-parameter analysis. Finally, repeating the fit with an explicit energy truncation $\sqrt{s} < \kappa f_a$ provides a conservative cross-check that our conclusions are stable against the highest-energy bins.

In summary, off-shell multiboson production offers a coherent and competitive program to test ALP-gauge interactions. The present global fit delivers stringent bounds on $(c_{\tilde{G}}, c_{\tilde{W}}, c_{\tilde{B}})$ from existing LHC data and clarifies the unique role of VBF and dijet channels in closing the parameter space. With high-luminosity LHC (HL-LHC) statistics and improved measurements, this strategy can decisively probe $\mathcal{O}(1)$ electroweak ALP couplings and push the sensitivity to gluonic interactions well into the EFT-safe regime. Finally, looking further ahead, channels with three gauge bosons are not yet competitive, but as triboson and aQGC measurements improve with HL-LHC, they will provide a powerful complementary handle on ALP interactions.

ACKNOWLEDGMENTS

The work of V. S. is supported by the Spanish Grants No. PID2023-148162NB-C21, No. CNS-2022-135688, and No. CEX2023-001292-S. The work of M. U. is supported by the European Research Council under the European Union's Horizon 2020 research and innovation program (Grant Agreement No. 950246) and in part by STFC consolidated Grant No. ST/X000664/1. F. E. is supported from Charles University through Project No. PRIMUS/24/SCI/013. A. S. B. acknowledges support from Grants No. PID2023-148162NB-C21 and No. PID2022-137003NB-I00 from Spanish MCIN/AEI/10.13039/501100011033/ and EU FEDER.

DATA AVAILABILITY

The data that support the findings of this article are not publicly available upon publication because it is not technically feasible and/or the cost of preparing, depositing, and hosting the data would be prohibitive within the terms of this research project. The data are available from the authors upon reasonable request.

APPENDIX A: ON THE INTERFERENCE WITH SM AMPLITUDES

As commented in Sec. IV C, for all limits given in this article, we focus on the pure quadratic ALP contributions to the total cross section and do not take into account interference terms with the SM. Since we can analytically trace the scaling of the quadratic ALP contribution to the cross section in all channels, i.e., the last term in Eq. (4.2), it suffices to generate signal events for one finite choice of c_i . Therefore, we can model the resulting χ^2 distribution as an




FIG. 18. Differential cross section in the WW final state as a function of the diboson invariant mass. The signal is shown for $c_{\tilde{G}} = 1$, $c_{\tilde{W}} = 1$, and $c_{\tilde{B}} = 0$, compared to the ATLAS data and the SM background simulation in [38]. We add the contours of the full differential cross section including interference terms with the SM for $c_{\tilde{W}} = -1$ (yellow) and $c_{\tilde{W}} = 1$ (green).

analytic function in $c_{\tilde{G}}$, $c_{\tilde{B}}$, and $c_{\tilde{W}}$ and determine precise contours. If we included interference effects, we could not determine the scaling of the total cross section with the Wilson coefficients and would instead have to sample events for many different choices and interpolate the signal on this grid simulations.

Here, we justify the choice to focus on the quadratic ALP contribution in hindsight by estimating the importance of the interference terms for the two channels that, according to, e.g., Table IV, have the largest relative constraining power, namely, WW and dijet.

In the WW channel, and similarly in the other diboson channels $\gamma\gamma$ and ZZ, the new physics amplitudes do not

interfere with the SM model at the partonic level, as there is no SM contribution to $gg \rightarrow W^+W^-$. Hence, the interference with the SM arises only at the hadronic level after convoluting the partonic amplitudes with the parton distribution functions in the proton and are accordingly small. The new physics amplitude, which is independent of $c_{\tilde{B}}$, scales with $c_{\tilde{G}}c_{\tilde{W}}$ and the interference term depends on the sign of $c_{\tilde{G}}$ and $c_{\tilde{W}} = 1$. In Fig. 18, we show the differential cross section for the pure quadratic ALP signal (red) and the SM background (blue) in the diboson invariant mass, similar to Fig. 5, and add contours for the full cross section including the interference terms for $c_{\tilde{G}} = 1$ and the two choices $c_{\tilde{W}} = -1$ (yellow) and $c_{\tilde{W}} = 1$ (green). It is striking




FIG. 19. Double-differential distribution in m_{jj} and y^* for $pp \rightarrow a^* \rightarrow gg$ for pure ALP signal (red) and including interference terms (green), compared to ATLAS data [40] and the SM background. For better visibility, the signal is shown for $c_{\tilde{G}} = 10$, $c_{\tilde{B}} = c_{\tilde{W}} = 0$.

that the full cross section including the interference agrees well with the signal without interference.

To quantify this statement at the level of the total cross section, for $c_{\tilde{G}} = 1$ and $c_{\tilde{W}} = 1$, we find

$$\sigma_{WW}^{\text{SM}} = 76.74 \text{ pb}, \quad \sigma_{WW}^{\text{ALP}} = 181.13 \text{ pb}, \quad \sigma_{WW}^{\text{Int}} = -0.02 \text{ pb}, \quad (\text{A1})$$

meaning that, for values of the Wilson coefficients in the range of the limits found in this article, the interference is indeed negligible with respect to the pure quadratic ALP signal.

Next, we consider the dijet final, for which the new physics amplitude depends quadratically on $c_{\tilde{G}}$, which implies that the interference does not depend on the sign of $c_{\tilde{G}}$. The process $gg \rightarrow gg$ can be mediated in the SM via s -, t - and u -channel exchange of a gluon. We find that the SM cross section is 1 order of magnitude larger than the pure ALP cross section and the interference effects are of the same order as the pure ALP cross section. In Fig. 19, we show the double-differential cross section in the invariant dijet mass and y^* , as in Fig. 6, for the background distribution and the signal, assuming $c_{\tilde{G}} = 3$ for better visibility. In green we add the contour for the full cross section with the interference term for $c_{\tilde{G}} = 3$.

For most bins, in particular, in the signal regions with small y^* , in which many events are observed, the new physics amplitudes interfere positively with the SM cross section and including the interference terms predicts a higher signal. In principle, one would expect that including the interference effects we could get better limits on $c_{\tilde{G}}$. However, apart from the observation that we cannot analytically trace the scaling of the interference term with $c_{\tilde{G}}$, there is another caveat: In Sec. III D, we use the SM background estimation provided by ATLAS, which includes higher-order calculation including QCD effects to an accuracy that we cannot match with our MG5 simulation. As long as we consider only the pure ALP signal, we can estimate the total number of events by adding our signal estimation to the provided background. Once we include interference effects, the MG5 simulation will inevitably also contain the SM background estimation that can no longer be disentangled, and we cannot use the ATLAS SM background. Therefore, the calculated χ^2 difference will be huge due to the worse background estimation that overshadows the actual contribution of the interference term, and all limits derived from this distribution are worse than the limits given for the pure ALP signal in this article.

We conclude that, in the dijet channel, interference effects can be important and could, in principle, be used to derive stronger limits on $c_{\tilde{G}}$, but to match the sensitivity of our analysis, an improved background estimation in MG5 would be necessary. The limits assuming a pure

quadratic ALP signal that are derived in this article should therefore be seen as conservative bounds on the couplings.

APPENDIX B: FIDUCIAL CUTS

1. $\gamma\gamma$ final state

The fiducial cuts used for our signal estimation are taken from [30], which require the presence of at least two clusters of energy depositions with $|\eta| < 2.47$, with transverse momentum $p_T > 30$ GeV, and requiring that at least one photon has $p_T > 40$ GeV. The angular separation of the photons is required to be $\Delta R_{\gamma\gamma} > 0.4$ and we require that invariant mass of the photon to be $m_{\gamma\gamma} > 200$ GeV. The integrated luminosity is 139 fb^{-1} .

The binning used for the differential cross section with respect to the invariant mass of the photon pair is shown only up to 500 GeV with bin limits at 13.5, 20, 22.6, 25, 27, 29.1, 31, 33.2, 35.4, 37.8, 40.4, 43, 46, 49, 52.4, 55.8, 59.6, 63.4, 67.6, 72.2, 77, 82, 87.5, 93, 99, 105, 112, 118.5, 126, 133.5, 141, 149.5, 158, 166.5, 175.5, 185, 194.5, 205, 215, 226, 238, 24, 262, 274, 287, 301, 315, 330, 345, 361, 377, 394, 411, 429, 448, 467, 487, 508 GeV.

2. ZZ final state

The fiducial cuts based on the ATLAS measurement of $pp \rightarrow ZZ \rightarrow 4\ell$ at $\sqrt{s} = 13$ TeV with 29 fb^{-1} integrated luminosity [34] are as follows: The cuts for leptons in the final state are $p_T > 5$ GeV and $|\eta| < 2.47$ with a minimal angular separation of $\Delta R > 0.05$. In addition, the two same flavor opposite charge lepton pair invariant masses are required to satisfy $m_{ll} \in (66, 116)$ GeV. Finally, the invariant mass of the four leptons in the two pairs is required to be $m_{4l} > 180$ GeV.

The binning used for the differential cross section with respect to the invariant mass of the Z -boson pair is defined by the bin limits: 180, 200, 220, 240, 265, 290, 320, 350, 390, 440, 520, 620, 1200 GeV.

3. W^+W^- final state

Fiducial cuts for the ALP-mediated $pp \rightarrow W^+W^- \rightarrow \ell^+\nu\ell^-\bar{\nu}$ are based on the ATLAS measurement at $\sqrt{s} = 13$ TeV with an integrated luminosity of 36.1 fb^{-1} [38]. Leptons are required to have $p_T > 27$ GeV and $|\eta| < 2.5$. The minimum missing transverse energy corresponding to the neutrinos is $E_T^{\text{miss}} > 20$ GeV. The minimum angular separation for leptons is $\Delta R > 0.4$. The minimum invariant mass and transverse momentum for all leptons and neutrinos is $m_{T,4l} > 55$ and $p_{T,4l} > 30$ GeV.

The binning used for the differential cross section with respect to the invariant mass of the W -boson pair is defined by the bin limits 55, 75, 85, 95, 110, 125, 140, 160, 185, 220, 160, 185, 220, 280, 380, 600, 1500 GeV and all signal events fall in the binning region.

4. Dijet final state

The fiducial cuts for simulating our ALP-mediated $pp \rightarrow jj$ signal are based on the ATLAS measurement [40] with integrated luminosity of 3.2 fb^{-1} . These are as follows: The minimum transverse momentum of individual jets is $p_T > 75 \text{ GeV}$ with angular separation $\Delta R > 0.4$ and maximum rapidity of $|\eta| < 3$. The total invariant mass for the two jets satisfies $m_{T,2j} \in (260, 7000) \text{ GeV}$. Finally, the sum of magnitudes of transverse momenta of the two jets fulfills $H_{T,2j} > 200 \text{ GeV}$.

The binning used for the double-differential cross section with respect to the invariant mass of the jet pair and the half-rapidity separation $\frac{d^2\sigma}{dm_{jj}dy^*}$ is as follows:

- (i) For the half-rapidity region R1, $0 < y^* < 0.5$, there are 29 bins from 260 to 7000 GeV.
- (ii) For the region R2, $0.5 < y^* < 1$, there are 29 bins from 310 to 7630 GeV.
- (iii) For the region R3, $1 < y^* < 1.5$, there are 28 bins from 440 to 8317 GeV.
- (iv) For the region R4 $1.5 < y^* < 2$, there are 25 bins from 760 to 7630 GeV.
- (v) For the region R5, $2 < y^* < 2.5$, there are 22 bins from 1060 to 9066 GeV.
- (vi) For the region R6, $2.5 < y^* < 3$, there are 9 bins from 1600 to 9066 GeV.

For the different bin widths, we refer to the ATLAS data repository of [40]. Of our generated signal events, 93.0781% fall into the given ATLAS bins.

5. Vector-boson-fusion events

The fiducial cuts for our VBF ALP-mediated signal are as follows:

- (1) $Z\gamma$ final state: For this final state, the selection cuts for our signal estimation are the same as in the CMS measurement of $pp \rightarrow Z\gamma jj$ at $\sqrt{s} = 13 \text{ TeV}$ with 137 fb^{-1} [48], where the Z boson is reconstructed in the dilepton channel. The cuts in the transverse momentum are $p_T > 30 \text{ GeV}$ for the jets and $p_T > 20 \text{ GeV}$ for the photon; in the pseudorapidity $|\eta| < 4.7$ for the jets and $|\eta| < 2.5$ for the photons. The minimum angular distance is $\Delta R > 0.5$ and imposed both between two jets and between the photon and a jet. The minimum invariant mass for the jet pair is $m_{jj} > 500 \text{ GeV}$.

The binning has bin limits at 20, 80, 120, 200, 400 GeV for p_T^γ and 30, 150, 250, 350, 800 GeV for $p_T^{j,\text{lead}}$.

- (2) ZZ and W^+W^- final states: In accordance with [50], we simulate our events for ALP-mediated VBF di-weak-boson production at $\sqrt{s} = 13 \text{ TeV}$ with the following cuts:

- (a) Zero-lepton channel: The cut in transverse momentum is $p_T > 30 \text{ GeV}$ for the jets and $E_T^{\text{miss}} > 200 \text{ GeV}$ for the missing energy coming from neutrinos. The rapidity for jets is $|\eta| < 5$ and the angular distance among jets fulfills $\Delta R > 0.4$ with a jet pair having invariant mass $m_{jj} > 400 \text{ GeV}$.

The binning has bin limits at 0, 200, 400, 600, 800, 1000, 1200, 1400, 1600 GeV for the invariant mass of the ZZ system in the merged and resolved regions.

- (b) One-lepton channel: The cut in transverse momentum is $p_T > 30 \text{ GeV}$ for the jets and $p_T > 20 \text{ GeV}$ for the lepton ($p_T > 27 \text{ GeV}$ for the leading lepton). $E_T^{\text{miss}} > 80 \text{ GeV}$ for the missing energy coming from the neutrino. The rapidity for jets is $|\eta| < 5$ and the angular distance among jets fulfills $\Delta R_{jj} > 0.4$ with a jet pair having invariant mass $m_{jj} > 400 \text{ GeV}$. Finally, the angular distance among jets and the leptons is cut to be $\Delta R_{\ell j} > 0.4$ and the invariant mass for all leptons fulfills $m_{\ell\ell} > 64 \text{ GeV}$.

The binning has bin limits at 0, 200, 400, 600, 800, 1000, 1200, 1400, 1600, 1800, 2000 GeV for the invariant mass of the W^+W^- system in the merged and resolved regions.

- (c) Two-lepton channel: The cut in transverse momentum is $p_T > 30 \text{ GeV}$ for the jets and $p_T > 27 \text{ GeV}$ for the leptons. The rapidity for jets is $|\eta| < 5$ and $|\eta| < 4$ for leptons. The angular distance among all particles fulfills $\Delta R > 0.4$ with a jet pair having invariant mass $m_{jj} > 400 \text{ GeV}$.

The binning has bin limits at 0, 200, 400, 600, 800, 1000, 1200, 1400, 1600, 1800, 2000 GeV for the invariant mass of the ZZ system in the merged and resolved regions.

[1] R.D. Peccei and H.R. Quinn, *CP conservation in the presence of instantons*, *Phys. Rev. Lett.* **38**, 1440 (1977).
[2] S. Weinberg, *A new light boson?*, *Phys. Rev. Lett.* **40**, 223 (1978).

[3] F. Wilczek, *Problem of strong P and T invariance in the presence of instantons*, *Phys. Rev. Lett.* **40**, 279 (1978).
[4] P. Svrcek and E. Witten, *Axions in string theory*, *J. High Energy Phys.* **06** (2006) 051.

- [5] A. Arvanitaki, S. Dimopoulos, S. Dubovsky, N. Kaloper, and J. March-Russell, String axiverse, *Phys. Rev. D* **81**, 123530 (2010).
- [6] R. Contino, D. Marzocca, D. Pappadopulo, and R. Rattazzi, On the effect of resonances in composite Higgs phenomenology, *J. High Energy Phys.* **10** (2011) 081.
- [7] G. Ferretti and D. Karateev, Fermionic UV completions of composite Higgs models, *J. High Energy Phys.* **03** (2014) 077.
- [8] V. Sanz and J. Setford, Composite Higgses with seesaw EWSB, *J. High Energy Phys.* **12** (2015) 154.
- [9] K. Choi, S. H. Im, and C. Sub Shin, Recent progress in the physics of axions and axion-like particles, *Annu. Rev. Nucl. Part. Sci.* **71**, 225 (2021).
- [10] C. O'Hare, cajohare/axionlimits: Axionlimits, <https://cajohare.github.io/AxionLimits/> (2020).
- [11] I. Brivio, M. B. Gavela, L. Merlo, K. Mimasu, J. M. No, R. del Rey, and V. Sanz, ALPs effective field theory and collider signatures, *Eur. Phys. J. C* **77**, 572 (2017).
- [12] M. Bauer, M. Neubert, and A. Thamm, Collider probes of axion-like particles, *J. High Energy Phys.* **12** (2017) 044.
- [13] M. B. Gavela, J. M. No, V. Sanz, and J. F. de Trocóniz, Nonresonant searches for axionlike particles at the LHC, *Phys. Rev. Lett.* **124**, 051802 (2020).
- [14] J. Feng, Y.-n. Mao, and K. Wang, Discovery prospects for photophobic axion-like particles in the $WWjj$ final state at the High-Luminosity LHC, [arXiv:2511.21003](https://arxiv.org/abs/2511.21003).
- [15] F. Esser, M. Madigan, V. Sanz, and M. Ubiali, On the coupling of axion-like particles to the top quark, *J. High Energy Phys.* **09** (2023) 063.
- [16] J. Butterworth, M. Cullingworth, J. Egan, F. Esser, V. Sanz, and M. Ubiali, Probing the coupling of axions to tops and gluons with LHC measurements, *J. High Energy Phys.* **02** (2026) 073.
- [17] Y. Hosseini and M. Mohammadi Najafabadi, Exploring axionlike particle couplings through single top tW -channel and top pair production at the LHC, *Phys. Rev. D* **110**, 055026 (2024).
- [18] S. Bisal, Constraining ALP-top interaction from the chromoelectric dipole moment of the top quark, *Phys. Rev. D* **112**, 055027 (2025).
- [19] S. Barbosa, M. Coelho, S. Fichet, G. G. da Silveira, and M. Machado, LHC as an axion-photon collider, *Phys. Rev. Lett.* **135**, 181801 (2025).
- [20] F. Esser, M. Madigan, A. Salas-Bernardez, V. Sanz, and M. Ubiali, Di-Higgs production via axion-like particles, *J. High Energy Phys.* **10** (2024) 164.
- [21] 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, *J. High Energy Phys.* **06** (2022) 113.
- [22] G. Alonso-Álvarez, M. B. Gavela, and P. Quílez, Axion couplings to electroweak gauge bosons, *Eur. Phys. J. C* **79**, 223 (2019).
- [23] A. Biekötter and K. Mimasu, Axions and axion-like particles: Collider searches, [arXiv:2508.19358](https://arxiv.org/abs/2508.19358).
- [24] P. Agrawal, M. Nee, and M. Reig, Axion couplings in grand unified theories, *J. High Energy Phys.* **10** (2022) 141.
- [25] P. Agrawal, M. Nee, and M. Reig, Axion couplings in heterotic string theory, *J. High Energy Phys.* **02** (2025) 188.
- [26] G. Alonso-Álvarez, M. B. Gavela, and P. Quílez, Axion couplings to electroweak gauge bosons, *Eur. Phys. J. C* **79**, 223 (2019).
- [27] N. Craig, A. Hook, and S. Kasko, The photophobic ALP, *J. High Energy Phys.* **09** (2018) 028.
- [28] S. Chatrchyan *et al.* (CMS Collaboration), Measurement of differential cross sections for the production of a pair of isolated photons in pp collisions at $\sqrt{s} = 7$ TeV, *Eur. Phys. J. C* **74**, 3129 (2014).
- [29] <https://www.hepdata.net/record/ins1298393>.
- [30] G. Aad *et al.* (ATLAS Collaboration), Measurement of the production cross section of pairs of isolated photons in pp collisions at 13 TeV with the ATLAS detector, *J. High Energy Phys.* **11** (2021) 169.
- [31] <https://www.hepdata.net/record/104925>.
- [32] A. M. Sirunyan *et al.* (CMS Collaboration), Measurements of $pp \rightarrow ZZ$ production cross sections and constraints on anomalous triple gauge couplings at $\sqrt{s} = 13$ TeV, *Eur. Phys. J. C* **81**, 200 (2021).
- [33] <https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PAPERS/STDm-2022-17/>.
- [34] G. Aad *et al.* (ATLAS Collaboration), Measurement of ZZ production cross-sections in the four-lepton final state in pp collisions at $s = 13.6$ TeV with the ATLAS experiment, *Phys. Lett. B* **855**, 138764 (2024).
- [35] <https://www.hepdata.net/record/144768>.
- [36] A. Hayrapetyan *et al.* (CMS Collaboration), Measurement of inclusive and differential cross sections for W^+W^- production in proton-proton collisions at $\sqrt{s} = 13.6$ TeV, *Phys. Lett. B* **861**, 139231 (2025).
- [37] <https://www.hepdata.net/record/ins2796231>.
- [38] M. Aaboud *et al.* (ATLAS Collaboration), Measurement of fiducial and differential W^+W^- production cross-sections at $\sqrt{s} = 13$ TeV with the ATLAS detector, *Eur. Phys. J. C* **79**, 884 (2019).
- [39] <https://www.hepdata.net/record/ins1734263?version=1>.
- [40] M. Aaboud *et al.* (ATLAS Collaboration), Measurement of inclusive jet and dijet cross-sections in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector, *J. High Energy Phys.* **05** (2018) 195.
- [41] https://rivet.hepforge.org/analyses/ATLAS_2018_I1634970.
- [42] A. M. Sirunyan *et al.* (CMS Collaboration), Measurements of the differential jet cross section as a function of the jet mass in dijet events from proton-proton collisions at $\sqrt{s} = 13$ TeV, *J. High Energy Phys.* **11** (2018) 113.
- [43] <https://www.hepdata.net/record/ins1682495>.
- [44] A. M. Sirunyan *et al.* (CMS Collaboration), Search for new physics with dijet angular distributions in proton-proton collisions at $\sqrt{s} = 13$ TeV, *J. High Energy Phys.* **07** (2017) 013.
- [45] <https://www.hepdata.net/record/ins1519995>.
- [46] A. M. Sirunyan *et al.* (CMS Collaboration), Search for new physics in dijet angular distributions using proton-proton collisions at $\sqrt{s} = 13$ TeV and constraints on dark matter and other models, *Eur. Phys. J. C* **78**, 789 (2018); *Eur. Phys. J. C* **82**, 379(E) (2022).
- [47] <https://www.hepdata.net/record/ins1663452>. <https://www.hepdata.net/record/ins1869513>
- [48] A. Tumasyan *et al.* (CMS Collaboration), Measurement of the electroweak production of $Z\gamma$ and two jets in

- proton-proton collisions at $\sqrt{s} = 13$ TeV and constraints on anomalous quartic gauge couplings, *Phys. Rev. D* **104**, 072001 (2021).
- [49] <https://www.hepdata.net/record/ins1869513>.
- [50] G. Aad *et al.* (ATLAS Collaboration), Electroweak diboson production in association with a high-mass dijet system in semileptonic final states from pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector, [arXiv:2503.17461](https://arxiv.org/abs/2503.17461).
- [51] <https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PAPERS/STDM-2018-27/>.
- [52] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. S. Shao, T. Stelzer, P. Torrielli, and M. Zaro, The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations, *J. High Energy Phys.* **07** (2014) 079.
- [53] R. D. Ball *et al.* (NNPDF Collaboration), The path to proton structure at 1% accuracy, *Eur. Phys. J. C* **82**, 428 (2022).
- [54] I. Brivio, ALP_linear_UFO, http://feynrules.irmp.ucl.ac.be/attachment/wiki/ALPsEFT/ALP_linear_UFO.tar.gz, Accessed: November 2022, 30.
- [55] A. M. Sirunyan *et al.* (CMS Collaboration), Evidence for WW production from double-parton interactions in proton-proton collisions at $\sqrt{s} = 13$ TeV, *Eur. Phys. J. C* **80**, 41 (2020).
- [56] A. Gehrmann-De Ridder, T. Gehrmann, E. W. N. Glover, A. Huss, and J. Pires, Triple differential dijet cross section at the LHC, *Phys. Rev. Lett.* **123**, 102001 (2019).
- [57] A. Huss *et al.* (NNLOJET Collaboration), NNLOjet: A parton-level event generator for jet cross sections at NNLO QCD accuracy, [arXiv:2503.22804](https://arxiv.org/abs/2503.22804).
- [58] D. Britzger *et al.*, NNLO interpolation grids for jet production at the LHC, *Eur. Phys. J. C* **82**, 930 (2022).
- [59] Ploughshare website, <https://ploughshare.web.cern.ch/ploughshare/>.
- [60] S. Carrazza, E. R. Nocera, C. Schwan, and M. Zaro, PineAPPL: Combining EW and QCD corrections for fast evaluation of LHC processes, *J. High Energy Phys.* **12** (2020) 108.
- [61] R. D. Ball *et al.* (NNPDF Collaboration), An open-source machine learning framework for global analyses of parton distributions, *Eur. Phys. J. C* **81**, 958 (2021).
- [62] C. Schwan, PineAPPL: NLO EW corrections for PDF processes, *SciPost Phys. Proc.* **8**, 079 (2022).
- [63] X. Chen, T. Gehrmann, E. W. N. Glover, A. Huss, and J. Mo, NNLO QCD corrections in full colour for jet production observables at the LHC, *J. High Energy Phys.* **09** (2022) 025.
- [64] A. Chiefa, M. N. Costantini, J. Cruz-Martinez, E. R. Nocera, T. R. Rabemananjara, J. Rojo, T. Sharma, R. Stegeman, and M. Ubiali, Parton distributions confront LHC Run II data: A quantitative appraisal, *J. High Energy Phys.* **07** (2025) 067.
- [65] R. Abdul Khalek *et al.* (NNPDF Collaboration), Parton distributions with theory uncertainties: General formalism and first phenomenological studies, *Eur. Phys. J. C* **79**, 931 (2019).
- [66] R. D. Ball and R. L. Pearson, Correlation of theoretical uncertainties in PDF fits and theoretical uncertainties in predictions, *Eur. Phys. J. C* **81**, 830 (2021).
- [67] S. Bruggisser, L. Grabitz, and S. Westhoff, Global analysis of the ALP effective theory, *J. High Energy Phys.* **01** (2024) 092.
- [68] L. C. Bresciani, G. Levati, and P. Paradisi, Positivity and partial wave unitarity bounds on ALP theories via amplitude methods, [arXiv:2510.13953](https://arxiv.org/abs/2510.13953).