



Exploring new possibilities to discover a light pseudo-scalar at LHCb

Diogo Buarque Franzosi^{1,2,a}, Giacomo Cacciapaglia^{3,4,b}, Xabier Cid Vidal^{5,c}, Gabriele Ferretti^{1,d}, Thomas Flacke^{6,e}, Carlos Vázquez Sierra^{7,f} 

¹ Department of Physics, Chalmers University of Technology, Fysikgården, 41296 Gothenburg, Sweden

² Physics Department, University of Gothenburg, 41296 Gothenburg, Sweden

³ Institut de Physique des 2 Infinis (IP2I), CNRS/IN2P3, UMR5822, 69622 Villeurbanne, France

⁴ Université de Lyon, Université Claude Bernard Lyon 1, 69001 Lyon, France

⁵ Instituto Galego de Física de Altas Enerxías (IGFAE), Universidade de Santiago de Compostela, 15782 Santiago de Compostela, Spain

⁶ Center for Theoretical Physics of the Universe, Institute for Basic Science (IBS), Daejeon 34126, Korea

⁷ European Organization for Nuclear Research (CERN), Geneva, Switzerland

Received: 23 June 2021 / Accepted: 10 December 2021

© The Author(s) 2021

Abstract We study the possibility of observing a light pseudo-scalar a at LHCb. We target the mass region $2.5 \text{ GeV} \lesssim m_a \lesssim 60 \text{ GeV}$ and various decay channels, some of which have never been considered before: muon pairs, tau pairs, D meson pairs, and di-photon. We interpret the results in the context of models of 4D Composite Higgs and Partial Compositeness in particular.

1 Introduction

The search for resonantly produced particles Beyond the Standard Model (BSM) is a high-priority goal at the LHC. Many searches focus on resonant production and decays into di-bosons (see, e.g., [1,2]) or di-fermions (see, e.g., [3–5]) in the high-mass region. This has led the LHC experiments to push the exclusion to mass scales, in most cases, well above the TeV. Yet, new states with small masses may still be allowed and lie in unconstrained oases of the BSM parameter space. One example is provided by electrically neutral, colorless, scalars with a mass below the Z pole and above the heavy meson mass scales [6]. Rare meson decays, in fact, provide an additional class of strong bounds [7–10]. At the LHC, low masses are mainly constrained by di-muon [11–15] and di-photon searches [16–21].

A light spin-0 state could emerge in many BSM scenarios like supersymmetry, Higgs sector extensions, and models based on composite dynamics. Specifically, a pseudo-scalar a with a mass much lighter than the BSM scale is naturally realized as pseudo-Nambu–Goldstone boson (pNGB) associated with the spontaneous breaking of an approximate global symmetry. A time-honored example is provided by axions emerging from the Peccei–Quinn solution to the strong CP problem of QCD [22–25]. Other models of different nature featuring a light pNGB fall under the generic class of models of axion-like particles (ALPs) [26,27]. ALPs can be found in supersymmetry [28], models of composite electroweak symmetry breaking [29–31], and models with extended scalar sectors, like multiple Higgs doublet models [32,33] (including 2HDMs), type-II see-saw models for neutrino masses [34], and models with custodial triplets [35,36]. Among them, we focus specifically on composite Higgs models with an underlying fermionic UV description [37]. In particular, a light ALP is ubiquitous [6,38,39] in models with two species of confining fermions, needed to implement top partial compositeness [40,41].

The physics of ALPs can be encoded in a generic effective Lagrangian at low energy. The pNGB nature of the pseudo-scalar bears additional information on its coupling to Standard Model (SM) fermions (via derivative interactions that yield couplings proportional to the fermion masses) and gauge bosons (via Wess–Zumino–Witten terms). Moreover, in composite models, the coefficients are related to each other [38], as they emerge from the same underlying dynamics, and thus the branching ratios of a into SM particles can be correlated and predicted. Hence, con-

^a e-mail: bruarque@chalmers.se

^b e-mail: g.cacciapaglia@ipnl.in2p3.fr

^c e-mail: xabier.cid.vidal@cern.ch

^d e-mail: gabriele.ferretti@chalmers.se

^e e-mail: tom.flacke@gmail.com

^f e-mail: carlos.vazquez@cern.ch (corresponding author)

trary to a generic ALP scenario, detection prospects in different decay channels can be compared. A pseudo-scalar pNGB with a mass below the Z pole is expected to dominantly decay into the heaviest accessible fermion pairs ($b\bar{b}$, $c\bar{c}$, $\tau^+\tau^-$) or gluons gg (i.e. light hadrons), for which no low-mass LHC searches are available, so far (see Ref. [39] for a proposal of a low-mass di-tau search). The branching ratios into the experimentally tested $\mu^+\mu^-$ and $\gamma\gamma$ channels are typically small. In this mass range, the production mode at LHC is completely dominated by gluon fusion and this is the only production mode considered in this work. For studies of ALP production at lepton colliders see e.g. Refs. [42–45].

In this article, we present the first study of LHCb prospects to observe a composite ALP a in the $c\bar{c}$ and $\tau^+\tau^-$ channels. For comparison with existing bounds, we also re-interpret searches in the $\mu^+\mu^-$ and projections for the $\gamma\gamma$ channel. As benchmarks, we consider the 12 composite Higgs models (M1–M12) defined in Refs. [6,38,41]. However, results are also presented model-independently and can be applied to any other light pseudo-scalar model. At the LHC, the LHCb detector [46] is a forward spectrometer, whose special features make it appropriate for the types of signatures described in this work [47]. This includes the capability of triggering on soft objects, excellent vertex reconstruction that is useful to distinguish shorter lifetime objects, such as τ leptons, and very good invariant mass resolution that provides advantages for the discrimination against large continuous backgrounds. This last feature is crucial in the case of $c\bar{c}$. LHCb is currently undergoing an upgrade [48–50], after which it is expected to collect 15 fb^{-1} over the next three years. Overall, the experiment will collect 300 fb^{-1} in its whole lifetime [51,52].

The paper is organized as follows. In Sect. 2 we briefly introduce the effective Lagrangian and the benchmark models used in this article. In Sect. 3 we present the recast and projections for the existing di-muon searches. In Sect. 4 we compare to the reach obtainable in the di-photon final state. In the following two sections, 5 and 6, we describe in detail new search proposals for final states containing taus and D mesons. Finally, we offer our conclusion and summary plots in Sect. 7.

2 Model and simulation

The phenomenology of a light ALP a can be generically described by the following effective Lagrangian

$$\mathcal{L}_{\text{eff}} \supset \frac{1}{2}(\partial_\mu a)(\partial^\mu a) - \frac{1}{2}m_a^2 a^2 - i \sum_\psi \frac{C_\psi m_\psi}{f} a \bar{\psi} \gamma^5 \psi$$

Table 1 The numerical values of the coefficients in the lagrangian (1) for the 12 models considered. Only the sum $K_\gamma = K_W + K_B$ is relevant for the mass range considered in this work. The coupling to the top quark C_t can take a discrete set of values depending on the spuriotic charge assignments. In this work we chose the value that leads to the largest constructive interference to the gluon coupling and thus the largest production cross section via gluon fusion

	$C_{\psi \neq t}$	C_t	K_g	$K_\gamma = K_W + K_B$
M1	2.17	5.79	−7.24	10.4
M2	2.61	4.79	−8.70	17.7
M3	2.17	2.54	−6.34	0.483
M4	1.46	2.43	−10.9	−5.82
M5	1.46	6.31	−4.85	4.04
M6	1.46	6.31	−4.85	5.50
M7	2.61	4.79	−8.70	20.3
M8	1.90	3.16	−1.58	−0.422
M9	0.702	1.87	−10.3	−16.2
M10	0.702	1.87	−9.36	−13.7
M11	1.66	2.22	−3.33	−2.22
M12	1.83	2.84	−4.06	−1.69

$$+ \frac{a}{16\pi^2 f} \left(g_s^2 K_g G_{\mu\nu}^a \tilde{G}^{a\mu\nu} + g^2 K_W W_{\mu\nu}^i \tilde{W}^{i\mu\nu} + g'^2 K_B B_{\mu\nu} \tilde{B}^{\mu\nu} \right), \tag{1}$$

where ψ are the SM fermions, $F_{\mu\nu}$ the SM gauge field strengths ($F = G^a, W^i, B$), $\tilde{F}_{\mu\nu} = \frac{1}{2}\epsilon_{\mu\nu\rho\sigma} F^{\rho\sigma}$, and f the ALP decay constant. In models addressing the hierarchy problem, like composite Higgs ones, the scale f is typically assumed to be in the TeV range. For recent systematic studies of the dynamics of the Lagrangian above see Refs. [53–57]. Note that all couplings are considered of order 1 for generic ALP scenarios.

The presence of an explicit mass for the ALP evades the usual constraints that exclude the Peccei–Quinn–Weinberg–Wilczek (PQWW) axion [22–25] for a TeV scale f but also precludes its use to solve the strong CP problem. Conversely, this scenario arises naturally in models of composite Higgs with top partial compositeness, emerging from an underlying gauge theory with fermions [41]. An ALP state, potentially light, is an unavoidable byproduct of the global symmetries broken by the condensates [6,38], which generate both a composite Higgs and composite top partners.

A main feature of this class of composite ALPs is that the coefficients in the effective Lagrangian can be computed in terms of the underlying theory, hence rendering the theory highly predictive. In this article, we follow the set of benchmark models M1–M12 that are defined in Ref. [6]. They yield predictions for the couplings C_ψ, K_g, K_W, K_B summarized in Table 1. We emphasize that the couplings are computed from the underlying theory and not arbitrarily chosen. The

couplings to gauge bosons $K_{g,W,B}$ are generated by anomalies, and do not include the effects of loops of SM fermions, which are computed separately. In particular, the coupling to the photon $K_\gamma = K_W + K_B$ can take values between -13.7 and 20.3 yielding very different branching ratios in this channel. For a detailed discussion of the models and the derivation of the coupling constants, we refer to Ref. [6]. In the following, we will only summarize the aspects relevant for the LHCb study presented here.

Contrary to the generic scenario with all order 1 couplings, the composite ALP scenarios M1–M12 offer cases where specific couplings could be substantially enhanced or suppressed. This opens up the possibility of discovery in what would be generically considered sub-leading channels or vice versa. As a concrete example, one finds models like M3 and M8, where the photon channel is suppressed. This feature arises from the specific electroweak couplings of the confining fermions. In fact, the main difference of Eq. (1) from similar Lagrangians arising in the context of 2HDMs is that, for the models at hand, the $K_{g,W,B}$ constants denote the anomalous contributions from the (confined) hyperquarks of the UV theory and are thus non-zero even before integrating out the heavy SM quarks (mainly the top and bottom, for the mass ranges we consider in this study). Note that such terms could also be present in 2HDMs and other extensions of the scalar sector of the SM if heavy non-SM fermions are included.

In the remainder of the paper, we focus on the four most promising signatures for the search of composite ALPs at LHCb, with special focus on the role it can play in the coming runs. We consider, in turn, the decay modes:

$$a \rightarrow \mu^+\mu^-, \gamma\gamma, \tau^+\tau^-, c\bar{c}. \tag{2}$$

The most studied decay channel, and the one likely to be dominant under generic assumptions on the couplings, is the di-muon channel. Here, we offer a straightforward recasting of the existing searches targeting 2HDMs [11–15] and convert the bounds on the mixing angle among Higgses to those of the tuning parameter v/f , where $v = 246$ GeV is the Higgs vacuum expectation value.

In the di-photon channel, we use the projections in Ref. [16], (obtained from inclusive diphoton cross-section measurements imposing that the signal events are less than the total measured events plus twice their uncertainty), which can be treated similarly to the di-muon case. In our models, this channel is expected to give relatively weak exclusion bounds.

The di-tau and charm channels have not been considered before in the LHCb context. The di-tau channel can potentially cover a significant mass range (from 14 to 40 GeV). It benefits from a branching ratio enhanced by a factor $(m_\tau/m_\mu)^2 \approx 283$ compared to the di-muon (for $C_\tau \approx C_\mu$) but suffers from the presence of neutrinos and hadrons in the

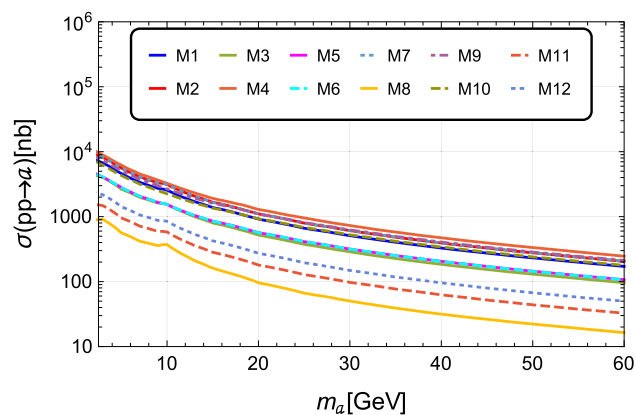


Fig. 1 Cross section at $\sqrt{s} = 14$ TeV pp collision for the 12 benchmark models defined by Eq. (1) with the couplings given in Table 1 for $f = v$. The cross sections scale as $(v/f)^2$ and are computed at NNLO in QCD using the HIGLU program [62,63]. The mass range starts at 2.5 GeV to stay away from the non-perturbative region (see Ref. [64])

final states. The $a \rightarrow c\bar{c}$ decay mode is relevant in the mass range $3.8 \text{ GeV} \lesssim m_a \lesssim 6 \text{ GeV}$. (For $m_a < m_{J/\psi}$ the non-observation of the $J/\psi \rightarrow a\gamma$ process puts strong bounds on f [58].) Besides the recast of existing bounds, we will offer projected LHCb reaches for all four channels for integrated luminosities of 15 fb^{-1} and 300 fb^{-1} .

For the signal simulations, we compute the ALP total production cross sections $\sigma(pp \rightarrow a)$, dominated by the gluon-fusion channel, with the HIGLU program [59] at NNLO in QCD, using the NNPDF 3.1 containing LHCb data [60] for 14 TeV in the pp center-of-mass. The numerical results are shown in Fig. 1 for each of the 12 models for $f = v$. The cross sections always scale as $(v/f)^2$. In the plot we use renormalization and factorization scales $\mu_F = \mu_R = m_a$. It should be noted that at low masses (few GeVs) a large scale dependence is present and the predictions start to become unreliable. In Fig. 1 we only show the central value obtained for each model and refer to Ref. [61] for a discussion of the expected error estimated by varying μ_F and μ_R . We use these values only for limit setting. The branching ratios of a into the four decay channels (2) are shown in Fig. 2. We use analytical expressions for the partial widths into $\mu^+\mu^-, \tau^+\tau^-, c\bar{c}, b\bar{b}, \gamma\gamma$ [6]. For the $a \rightarrow gg$ decay channel we use instead the HIGLU program [62,63].

3 Recast and projections for $a \rightarrow \mu^+\mu^-$

If the ALP couples to the muon with a typical strength $C_\mu \approx 1$, we expect the muon channel to give strong bounds for a wide range of masses. Searches in this channel have already been performed by various collaborations [11–15] and we start by presenting a recast of these results. We use the summary plot in fig. 10 of Ref. [11], presenting upper limits at

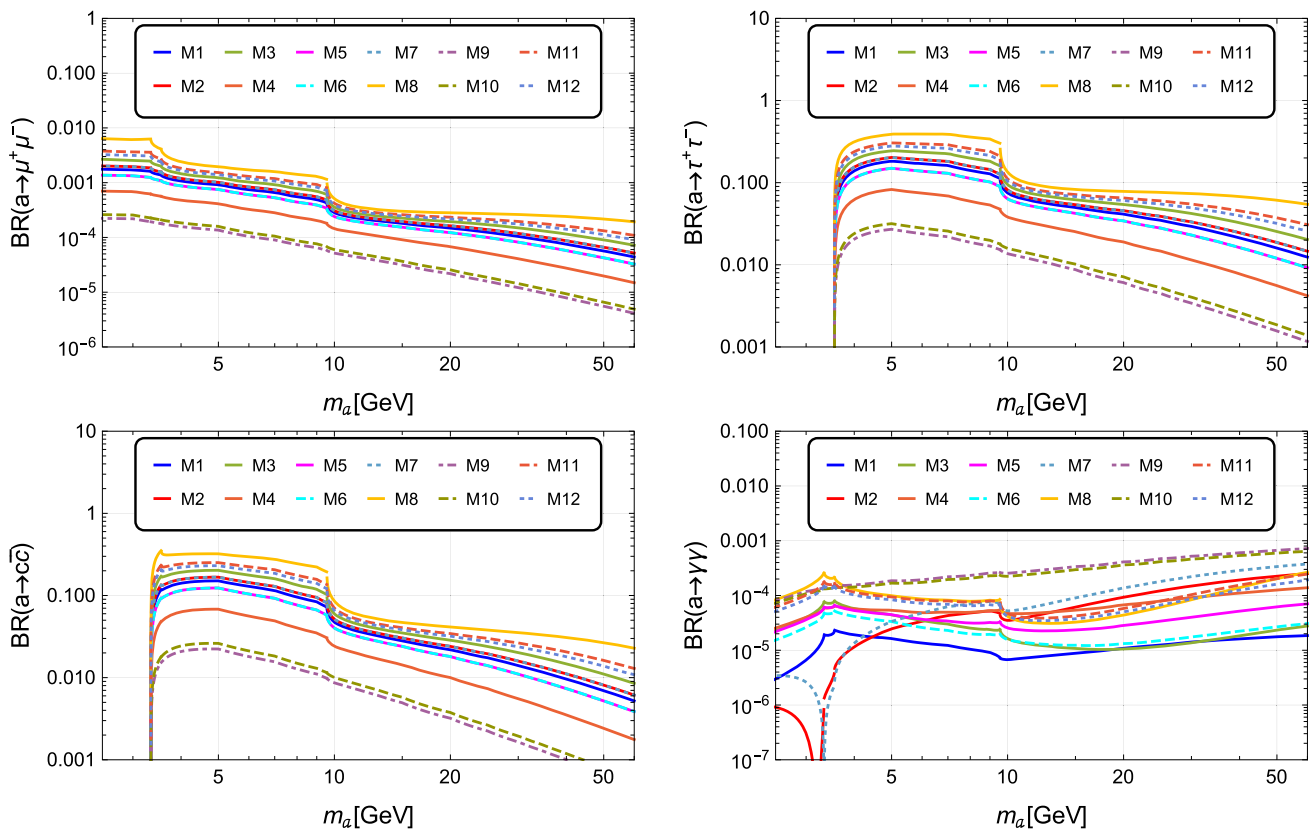


Fig. 2 Branching ratios of a into the four channels considered in (2) for the 12 models. Note that in these models the ALP is always promptly decaying with a narrow width, ranging, in this mass region, from few

keV to few MeV for $f = 1$ TeV. The HIGLU program has been used for the computation of the $a \rightarrow gg$ decay channel [62,63]

90% confidence level on the mixing angle $\sin \theta_H$ between the pseudo-scalar of a 2HDM and the imaginary component of a complex singlet. The plot is obtained from the total di-muon cross-section of the previous searches [11–15] and tests the type IV 2HDM model of [61,65] at $\tan \beta = 0.5$. The bounds are set on the mixing angle $\sin \theta_H$ as a function of the ALP mass, ranging between 1–60 GeV.

The recast is easily implemented thanks to the following two observations. The first one is that, for all models, the narrow width approximation holds very well and thus $\sigma(pp \rightarrow a \rightarrow \mu^+\mu^-) = \sigma(pp \rightarrow a) \times \mathcal{B}(a \rightarrow \mu^+\mu^-)$. Note however that in these models the ALP is always promptly decaying since the width ranges from few keV to few MeV for $f = 1$ TeV.

The second observation is that all $\sigma(pp \rightarrow a)$ are proportional to $\sin^2 \theta_H$ in the 2HDM and proportional to v^2/f^2 in the composite ALP models of Eq. (1), while all the branching ratios are independent on these quantities. Thus, denoting the cross-section of the 2HDM model with $\sin \theta_H = 1$ by $\bar{\sigma}^{2\text{HDM}}(pp \rightarrow a)$ and, similarly, the cross-section of the models of Table 1 with $f = v$ by $\bar{\sigma}^{\text{Mi}}(pp \rightarrow a)$, the bounds on v/f are obtained from the bounds on $\sin \theta_H$ in Ref. [11] by the simple rescaling

$$\frac{v}{f} \Big|_{\text{Mi}} = \sqrt{\frac{\bar{\sigma}^{2\text{HDM}}(pp \rightarrow a) \mathcal{B}^{2\text{HDM}}(a \rightarrow \mu^+\mu^-)}{\bar{\sigma}^{\text{Mi}}(pp \rightarrow a) \mathcal{B}^{\text{Mi}}(a \rightarrow \mu^+\mu^-)}} \sin \theta_H \tag{3}$$

Using this procedure one could also recast the LHCb search [11] to obtain bounds on any of the four types of 2HDMs (I, II, III, and IV) for any value of $\tan \beta$. If one were content with working at leading order, employing the detailed balance $\sigma(pp \rightarrow a) \propto \Gamma(a \rightarrow gg)$ one could replace the cross-sections in Eq. (3) with the partial width and obtain an analytic formula valid to within few % in the mass region $m_a > 15$ GeV.

The relevant widths are listed in Ref. [65] for the 2HDM and in Ref. [6] for the ALP models. For the total width we sum over the channels $\mu^+\mu^-$, $\tau^+\tau^-$, $c\bar{c}$, $b\bar{b}$, gg , $\gamma\gamma$. For the $a \rightarrow gg$ decay, giving the full hadronic decay at lower masses, we used HIGLU and compared the results with analytic ones including one-loop renormalization of the quark masses and the gauge couplings as well as the finite part of the QCD corrections [62,63]. We find good agreement between the two. More precisely, the mean deviation between the two estimates, averaging over the 12 models, is

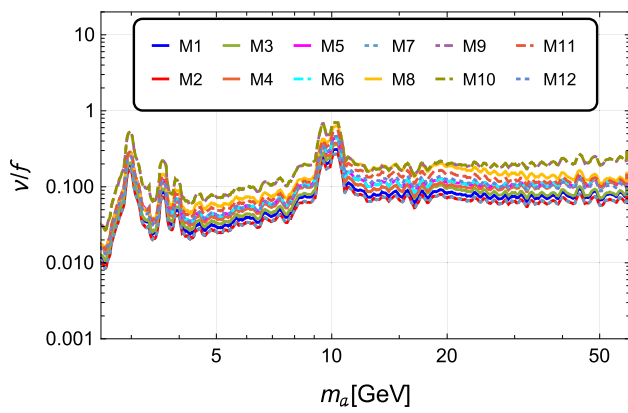


Fig. 3 Bounds on v/f as a function of m_a for the di-muon channel at 90% C.L. This is a recasting of Fig. 10 in Ref. [11] using the envelope of the exclusion curves from the searches [11–15]. In order to smooth out some of the variability of the original exclusion curve we performed a moving average over the 10 nearby mass point for each point on the axis

20%, 4.1%, 1.5%, and 0.56% for ALP masses of 2.5, 5, 30, and 60 GeV respectively. The non-perturbative aspects of ALP decay into light hadrons become crucial for lighter ALPs and are discussed in Ref. [64]. The cross-sections are also computed numerically at NNLO with HIGLU as described in Sect. 2.

We reiterate that the important difference between 2HDMs and Eq. (1) is the presence in the latter of contact terms to gluons and photons coming from the anomaly of the hyperfermions, which are absent in 2HDMs and lead to an enhancement of the gg coupling strength. Just like for the value of $\tan \beta$ in the 2HDM, we need to fix the coupling strengths C_ψ to the fermions in our models. These are given in Table 1. The coefficients for all fermions other than the top quark are fixed in the underlying theory as explained in Refs. [6,38]. The variability comes from different discrete choices for the coupling of the top quark due to different spurion charge assignments. Throughout the paper we present the bounds arising from the choice of C_t giving the largest effective coupling between the ALP and the gluon. This leads to the strongest bounds and sensitivities for three of the four channels in Eq. (2) due to the constructive interference between the hyper-fermion anomaly and the top quark coupling. The exception is the di-photon channel, discussed in more detail in Sect. 4.

Figure 3 shows the bounds on v/f for the 12 benchmark models that arise from the rescaling of the di-muon bounds of Ref. [11] as described in Eq. (3). For each ALP mass hypothesis we use the strongest bound of the four experimental analyses: BABAR [12], CMS Run 1 [13], CMS Run 2 [14], LHCb Run 1 [15], and LHCb Run 2 [11]. These bounds are the strongest constraints on v/f for these models in the considered m_a mass range to-date, showing that this channel is the most sensitive one under the assumption that

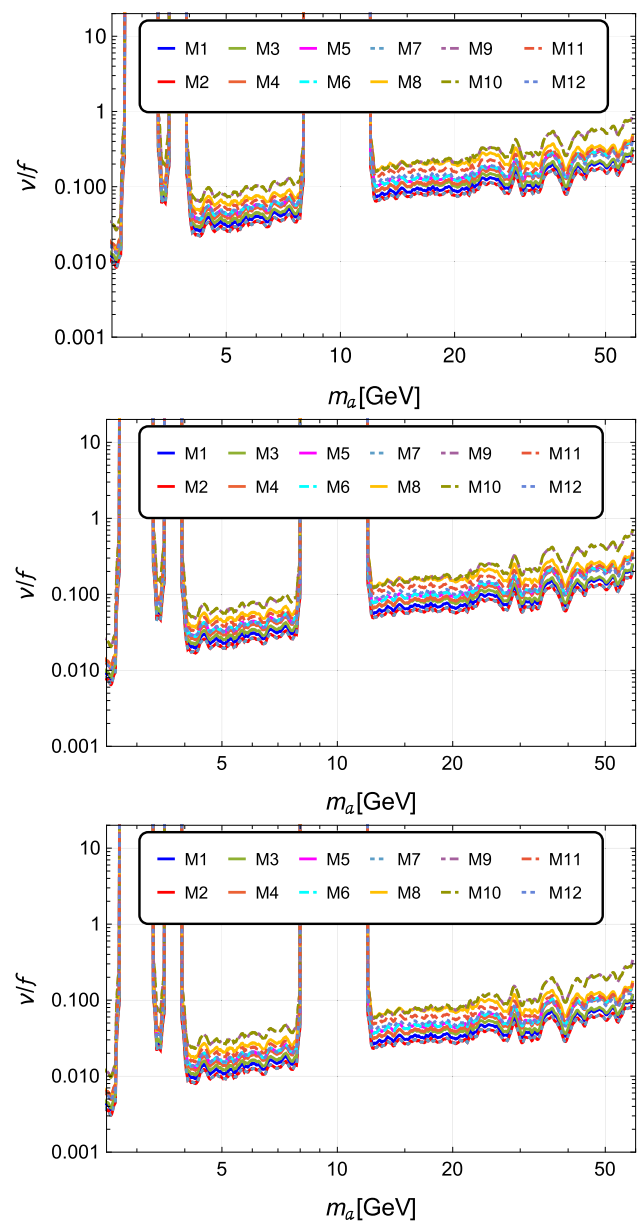


Fig. 4 Bounds on v/f as a function of m_a for the di-muon channel. Top: Same as Fig. 3 but using LHCb data Run 2 only ($\mathcal{L} = 5.1 \text{ fb}^{-1}$). Middle: Projections at 90% C.L. for LHCb at $\mathcal{L} = 15 \text{ fb}^{-1}$. Bottom: Projections at 90% C.L. for LHCb at $\mathcal{L} = 300 \text{ fb}^{-1}$

the ALP couples to the muon with standard strength. However, for a “muon-phobic” ALP the other channels become relevant and should be considered in order to broaden the reach.

We conclude this section by presenting in Fig. 4 the projections for the di-muon channel obtained by rescaling the LHCb results [11] (Run 2 only, top panel) from 5.1 fb^{-1} to 15 fb^{-1} (middle panel), and 300 fb^{-1} (bottom panel). We expect the results to be dominated by statistics. Hence, noting that $S/\sqrt{B} \propto \sqrt{\mathcal{L}}$ and that $S \propto (v/f)^2$, the exclusion in v/f scales like $\mathcal{L}^{-1/4}$. Note that the projections do not

account for the removal of the first trigger level at LHCb [48], based on hardware, which is expected to happen from Run 3 onward. While doing this would yield more stringent exclusions, we choose to extrapolate from the existing LHCb results, based on data, which provides more realistic estimates of the background.

4 Projections for $a \rightarrow \gamma\gamma$

We continue by examining the di-gamma decay channel. We restrict ourselves to the low mass region $2.5 \text{ GeV} < m_a < 20 \text{ GeV}$ where the strongest bounds derive from the analysis in Ref. [16] of the data from ATLAS, CMS, LHCb, and BABAR [17–21]. The exclusions [16] translate into fairly weak current bounds for our models and we thus simply present the projected LHCb bounds obtained from Ref. [66,67] in the same spirit as for the di-muon channel.

Reference [66] investigated the $\gamma\gamma$ channel for a fully fermiophobic model ($C_\psi = 0$) whose anomaly coefficients can be written in our notation (1) as $K_g = 10$ and $K_W + K_B = 80/3$. The same reasoning that yields Eq. (3) now gives (“fph” stands for fermiophobic)

$$\frac{v}{f} \Big|_{\text{Mi}} = \sqrt{\frac{\bar{\sigma}^{\text{fph}}(pp \rightarrow a) \mathcal{B}^{\text{fph}}(a \rightarrow \mu^+\mu^-)}{\bar{\sigma}^{\text{Mi}}(pp \rightarrow a) \mathcal{B}^{\text{Mi}}(a \rightarrow \mu^+\mu^-)}} \frac{v}{f} \Big|_{\text{fph}}, \quad (4)$$

the barred cross-sections denoting the values with $f = v$.

In Fig. 5 we show the projected sensitivity for the 12 models for the LHCb run with integrated luminosities of 15 fb^{-1} (top) and 300 fb^{-1} (bottom) respectively. For the calculation of the partial width of $a \rightarrow \gamma\gamma$, we take into account the one-loop corrections from $t, b, c,$ and τ loops, which yield sizable corrections, in particular for low m_a . The remaining calculation is performed exactly as explained in the di-muon section.

In Fig. 5 we still choose the value of C_t leading to the largest effective coupling of the ALP to the gluon and thus to the largest production cross section (within each benchmark model). However, contrary to all other channels, the C_t with largest ALP production cross section does not necessarily lead to the strongest exclusion bound in the di-photon channel. This is because, contrary to the fermionic decay channels, the decay width $\Gamma(a \rightarrow \gamma\gamma)$ depends strongly on the value of C_t , as explained in Ref. [38].

Specifically, the top loop correction to the effective coupling of the gluon and the photon are, respectively

$$K_{g,\text{eff.}} = K_g - \frac{1}{2}C_t p(4m_t^2/m_a^2) + \dots \quad (5)$$

$$K_{\gamma,\text{eff.}} = K_W + K_B - \frac{4}{3}C_t p(4m_t^2/m_a^2) + \dots \quad (6)$$

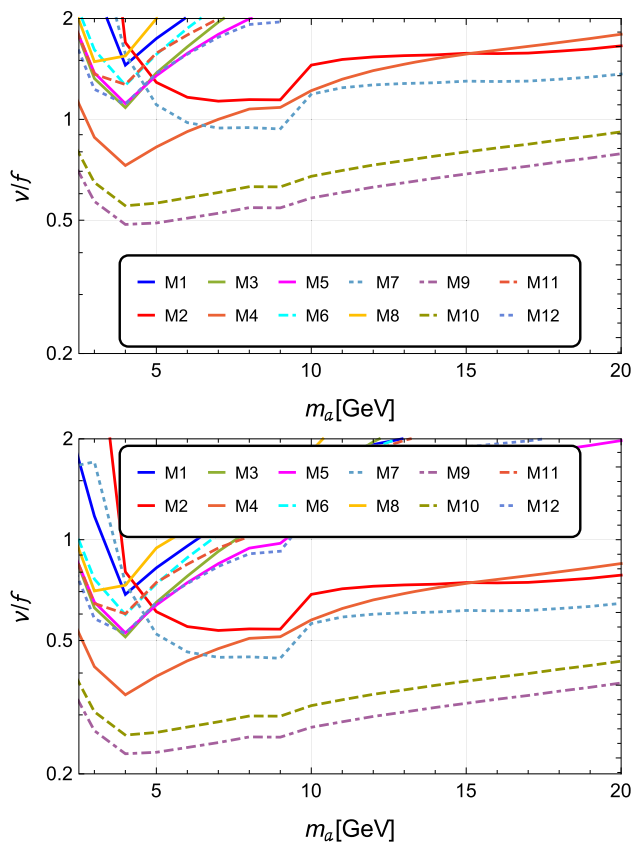


Fig. 5 Di-photon projected sensitivity on v/f as a function of m_a for the LHCb run with integrated luminosities of 15 fb^{-1} (top) and 300 fb^{-1} (bottom). The results are obtained by converting the sensitivity reported in Fig. 1 of Ref. [66] using (4) with values for the couplings in Table 1. Note that for $\mathcal{L} = 15 \text{ fb}^{-1}$ many models are basically unconstrained

where K_g, K_W, K_B are the anomaly coefficients in Table 1 from the hyper-fermions, $p(\tau) = \tau \arctan^2(\sqrt{1/(\tau-1)})$ and the dots represent contributions from the lighter SM fermions. For all models one always has $K_g < 0$, and C_t ranges in an interval containing positive and negative values. The largest $|K_{g,\text{eff.}}|^2$ is thus attained by picking the spurion charges giving the largest (positive) C_t for each model. On the contrary, $K_W + K_B$ can have positive or negative values and there can be destructive interference if one picks the largest C_t .

For a consistent comparison with (projected) bounds in the di-muon, di-tau, and di-charm channel, we use the same values of C_t for the analysis of the di-photon search, but it should be noted that other choices of C_t can lead to altered sensitivity in the di-photon channel, up to an order of magnitude.

5 New search in $a \rightarrow \tau^+\tau^-$

In this section we show the potential of the LHCb experiment to search for $a \rightarrow \tau^+\tau^-$ decays. Despite the presence of

neutrinos and the fact of being a non-hermetic spectrometer, LHCb has already shown its potential to search for decay modes including τ leptons in the final state. Notable examples include τ leptons produced either from a displaced low-mass vertex (semileptonic decays of B mesons [68] or pairs from the leptonic decay of a B_s^0 meson [69], where τ leptons are reconstructed using three charged pions in the final state), or from a prompt high-mass vertex (decay of a Z boson into $\tau^+\tau^-$ [70], where several combinations of the decay modes of both τ leptons – hadronic, semileptonic, and fully leptonic – are explored).

The signal topology described in this paper is challenging for LHCb, due to the fact that both τ leptons are produced from the decay of a prompt object, which has a relatively low invariant-mass. However, the excellent capabilities of the detector to reconstruct soft objects in the final state help suppressing most of the dominant background components that would pollute our signal. A previous proposal for a $a \rightarrow \tau^+\tau^-$ search at CMS and ATLAS can be found in Ref. [39].

5.1 Simulation and analysis strategy

The signal process is defined by the production of a prompt light pseudo-scalar in proton-proton collisions at a center-of-mass energy of 14 TeV, decaying into a pair of tau leptons with opposite charges. It is simulated with the PYTHIA 8.305 program [71] with fully spin-correlated tau decays [72].

A signal fiducial region is defined in order to account for differences in the production mechanism between this simplified PYTHIA model and the NLO model described in more details in Sect. 6. This region is defined by selecting pseudo-scalar PYTHIA objects with a pseudo-rapidity between 2 and 4.5 and a transverse momentum p_T between 15 and 150 GeV. We also impose requirements on the two τ leptons: a pseudo-rapidity between 1.5 and 5, at least one τ with $p_T > 7.5$ GeV and the second tau with $p_T > 5$ GeV. These requirements are also part of the selection imposed to the reconstructed objects, as described in the following paragraphs. We checked that any leak in selected data from events not in the fiducial region is negligible.

The decay modes of the τ leptons lead to very distinct signatures and background contributions. The main possibilities are:

- Fully leptonic ($e\mu$): $\tau^+ \rightarrow e^+ \bar{\nu}_\tau \nu_e$ and $\tau^- \rightarrow \mu^- \nu_\tau \bar{\nu}_\mu$.
- Semileptonic with an electron (h_3e): $\tau^+ \rightarrow \pi^+ \pi^- \pi^+ \bar{\nu}_\tau$ and $\tau^- \rightarrow e^- \nu_\tau \bar{\nu}_e$.
- Semileptonic with a muon ($h_3\mu$): $\tau^+ \rightarrow \pi^+ \pi^- \pi^+ \bar{\nu}_\tau$ and $\tau^- \rightarrow \mu^- \nu_\tau \bar{\nu}_\mu$.
- Fully hadronic (h_3h_3): $\tau^+ \rightarrow \pi^+ \pi^- \pi^+ \bar{\nu}_\tau$ and $\tau^- \rightarrow \pi^- \pi^- \pi^+ \nu_\tau$.

Charge-conjugate final states are left understood.

Three main SM background processes are expected to contaminate the signal selection:

- QCD multijet production, $pp \rightarrow jj$ with j standing for light quarks, gluons, and charm quarks.
- QCD heavy-flavor $pp \rightarrow b\bar{b}$ production. Most of the background leptons are expected to originate from b -hadron decays.
- Drell–Yan $pp \rightarrow \tau^+\tau^-$ production.

Other background components (such as other Drell–Yan productions and Z/W boson pair production with leptonic and semileptonic decays) were found to be negligible. All background processes have been simulated with PYTHIA8.

The QCD multijet background is challenging to estimate and is expected to be dominant for the fully hadronic and semileptonic channels. On the other hand, its contribution to the fully leptonic mode is estimated to be at most 10% of the dominant $b\bar{b}$ background. We also expect the fully leptonic channel to be the most sensitive one, see Appendix A for more details, hence we will neglect the other channels in computing the limits. We therefore only consider the two dominant backgrounds sources in the analysis of the fully leptonic mode, namely, heavy-flavor $b\bar{b}$ and Drell–Yan productions.

We limit our study to the mass region above 14 GeV and below 40 GeV. Below 14 GeV, the QCD background becomes unacceptably large and the Υ resonances are present, severely limiting the sensitivity at LHCb. Conversely, above 40 GeV the signal efficiency becomes compromised due to acceptance limitations.

5.2 Analysis of the $e\mu$ mode

All charged stable tracks under consideration (muons, electrons, pions, kaons, and protons) are required to be inside the LHCb pseudo-rapidity acceptance, $2 < \eta < 5$. A requirement of having the particles produced within the Vertex Locator (VELO) region is also imposed, that is, a minimum p_T of 0.5 GeV, with $V_r < 30$ mm and $V_z < 200$ mm, where (V_z, V_r) denotes the spatial position of their production vertices, expressed in cylindrical coordinates. Kaons, pions, and protons are later used only to define the quantities used to define the electron and muon track isolation.

Electrons and muons are required to have a p_T greater than 3.5 GeV, a minimum energy of 10 GeV, a minimum impact parameter (IP, defined below) of 0.01 mm for muons (and 0.03 mm for electrons), and to be well isolated from other tracks in the event. For this purpose, a quantity I to measure the *isolation* of a track is defined as the fraction of its p_T over the sum of the p_T of all stable charged tracks (muons, electrons, pions, kaons, and protons) inside a cone of a certain $\Delta R^2 = \Delta\phi^2 + \Delta\eta^2$ value, built around the track

of interest. We require at least one of the muons or electrons to be well isolated, imposing a tight cut to one of them, this is, $\max(I_\mu, I_e) > 0.99$, for $\Delta R^2 = 0.05$. Note that the discrimination provided by isolation tends to be overestimated in simulation with respect to data. Since this is an effect that is hard to estimate, we use the discrimination provided by our simulation, but we acknowledge this is a limitation of our study.

ALP a candidates are reconstructed by summing the 4-momenta of the selected (e, μ) pair. The reconstructed a is required to be prompt with a maximum distance of flight of 1 mm, to have an IP smaller than 0.2 mm, a pseudo-rapidity between 2 and 4.5, and a transverse momentum between 15 and 150 GeV (otherwise, the signal would be polluted by low p_T QCD background contributions). The muon-electron pairs from the decay of a are required to have a maximum distance of closest approach (DOCA) of 0.4 mm. Each lepton must have a minimum $p_T > 5$ GeV and at least one of them must have $p_T > 7.5$ GeV. The DOCA is the minimum distance between two different trajectories, defined as $|\mathbf{V} \times \rho|/|\rho|$ where $\mathbf{V} = \mathbf{V}_1 - \mathbf{V}_2$ and $\rho = \mathbf{p}_1 \times \mathbf{p}_2$. Here, \mathbf{V}_j and \mathbf{p}_j are a vertex 3D position and the tri-momentum associated to a track j ($j = 1, 2$), respectively. With the same vertex and tri-momentum definitions per track, the IP of a track j is defined as $|\frac{\mathbf{p}_j}{|\mathbf{p}_j|} \times \mathbf{V}_j|$.

Moreover, for each signal mass hypothesis, the invariant mass of the $e\mu$ system is required to be in a range that enhances the signal over background ratio. The ranges are shown in Table 2 for each ALP mass hypothesis.

5.3 Computation of efficiencies and bounds

Signal, $b\bar{b}$, and DY background efficiencies, ϵ^{ALP} , ϵ^{bb} , and ϵ^{DY} respectively, are obtained for several a mass hypotheses, and shown in Tables 2 and 3. The signal efficiency ϵ^{ALP} is the product of two contributions: the first one, computed with the NLO model, is the ratio of the cross section in the signal fiducial region and the inclusive cross section $pp \rightarrow a \rightarrow \tau^+\tau^-$; the second one, computed with the PYTHIA8 LO model, is the efficiency of the fully leptonic analysis in the fiducial region.

For each mass hypothesis, the number of signal and background events are

$$S = \mathcal{L} \times \sigma(pp \rightarrow a) \times \mathcal{B}(a \rightarrow \tau^+\tau^-) \times \epsilon^{ALP} \times \mathcal{B}(\tau^+ \rightarrow e^+\bar{\nu}_\tau \nu_e) \times \mathcal{B}(\tau^- \rightarrow \mu^-\nu_\tau \bar{\nu}_\mu) \times 2 \quad (7)$$

$$B = \mathcal{L} \times \left(\sigma(pp \rightarrow \tau^+\tau^-) \times \epsilon^{DY} \times \mathcal{B}(\tau^+ \rightarrow e^+\bar{\nu}_\tau \nu_e) \times \mathcal{B}(\tau^- \rightarrow \mu^-\nu_\tau \bar{\nu}_\mu) \times 2 + \sigma(pp \rightarrow b\bar{b}) \times \epsilon^{b\bar{b}} \right), \quad (8)$$

Table 2 Background efficiencies for the $e\mu$ reconstruction mode, for Drell–Yan and $b\bar{b}$ components. A mass window requirement is defined per signal mass hypothesis to be as efficient as possible for the signal, while helping to suppress a large fraction of background. These mass requirements are presented in the “Mass range” column, together with the corresponding value in GeV of the signal mass hypothesis used to obtain them, separated by a vertical line. These background efficiencies are provided in the full acceptance, from PYTHIA8 LO simulations

Mass range (GeV)	ϵ^{DY} (%)	$\epsilon^{b\bar{b}}$ (%)
(2.5, 10.0) 14	0.00676	1.75×10^{-7}
(4.0, 16.0) 20	0.0120	2.65×10^{-7}
(4.0, 16.0) 22	0.0120	2.65×10^{-7}
(4.5, 20.0) 25	0.0164	3.31×10^{-7}
(5.0, 24.0) 30	0.0289	3.92×10^{-7}
(7.0, 30.0) 40	0.0745	4.46×10^{-7}

Table 3 Signal efficiencies for the $\tau^+\tau^-$ channel in the $e\mu$ reconstruction mode, considering the NLO production model in the full acceptance. Mass window requirements are imposed on top of the selection, as described in Table 2 caption

Mass (GeV)	ϵ^{ALP} (%)
14	0.0523
20	0.108
22	0.109
25	0.139
30	0.186
40	0.206

where σ is the production cross-sections, \mathcal{B} the branching fractions of a and τ , and \mathcal{L} the integrated luminosity.

The $b\bar{b}$ production cross-section, $\sigma(pp \rightarrow b\bar{b}) = (562 \pm 82) \times 10^9$ fb, is taken from Ref. [73]. The Drell–Yan cross-section for the different mass windows are taken from Ref. [74] to be $\sigma(pp \rightarrow \tau^+\tau^-) = (4.494 \pm 0.237) \times 10^6$ fb, and the values of the branching fractions of the different τ decay modes are taken from Ref. [75].

Apart from discriminating against the background, our tight selection is defined to maximize the reconstruction and trigger efficiencies at LHCb. The removal of the first trigger level at LHCb [48], mentioned above, makes this assumption more realistic. Therefore, we assume this efficiency to be 100%. Moreover, we neglect experimental resolution effects that would affect the computation of the invariant masses, since the inaccuracy on this is dominated by the presence of neutrinos. The signal and background di-lepton invariant mass $m(e\mu)$ distributions are shown in Fig. 6, for $\mathcal{L} = 300 \text{ fb}^{-1}$, and model M1 with $v/f = 0.1$ for the signal.

We also neglect systematic uncertainties and provide expected limits at 90% C.L based purely on statistical uncer-

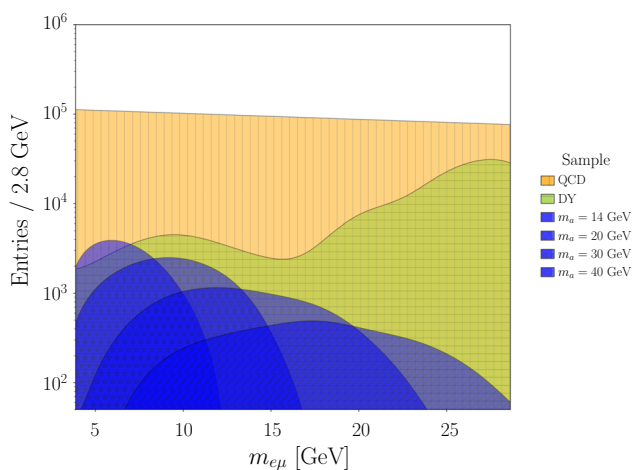


Fig. 6 Signal and background distribution of $m(\tau^+\tau^-)$ for the $e\mu$ reconstruction channel. The yields correspond to 300 fb^{-1} . For signal, the cross-sections are those predicted by model M1 and $v/f = 0.1$

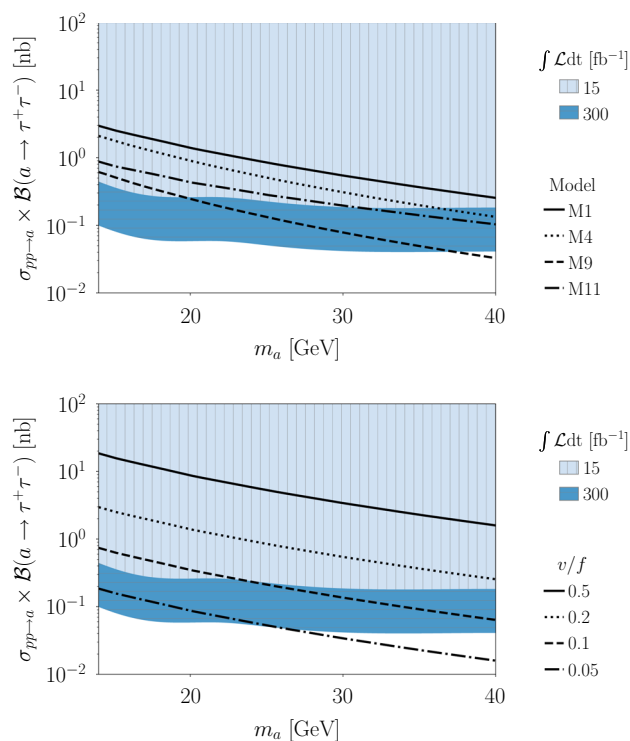


Fig. 7 Model independent projected bounds at 90% C.L. on $\sigma(pp \rightarrow a) \times \mathcal{B}(a \rightarrow \tau^+\tau^-)$ for the reference integrated luminosities. Black lines are central predictions for (top) models M1, M4, M9, and M11 with $v/f = 0.2$ and (bottom) model M1 with different values of v/f

tainties using the CL_s method [76,77]. An actual experimental search would be needed to control systematic effects.

In Fig. 7 we show the projected bounds on the signal cross section $\sigma(pp \rightarrow a) \times \mathcal{B}(a \rightarrow \tau^+\tau^-)$ as a function of the ALP mass, assuming an integrated luminosity of $\mathcal{L} = 15 \text{ fb}^{-1}$ (light blue) and $\mathcal{L} = 300 \text{ fb}^{-1}$ (dark blue). Signal

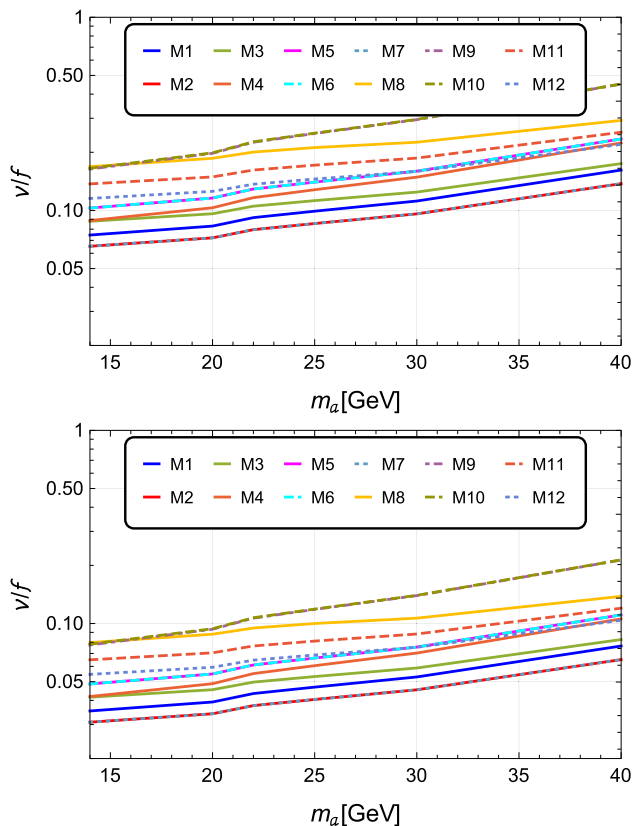


Fig. 8 Projections for the bounds at 90% C.L. on v/f as a function of m_a for the di-tau channel for the 12 models at $\mathcal{L} = 15 \text{ fb}^{-1}$ (top) and $\mathcal{L} = 300 \text{ fb}^{-1}$ (bottom)

predictions for benchmark models M1, M4, M9, and M11 with $v/f = 0.2$ are shown on the upper panel while model M1 is shown with different values of v/f on the lower panel.

In Fig. 8 we show the projected bounds on v/f for the 12 benchmark models, assuming an integrated luminosity of $\mathcal{L} = 15 \text{ fb}^{-1}$ (top) and $\mathcal{L} = 300 \text{ fb}^{-1}$ (bottom).

6 New search with D mesons targeting $a \rightarrow c\bar{c}$

The $c\bar{c}$ channel is only relevant for a small range of low ALP masses, $3.8 \text{ GeV} \lesssim m_a \lesssim 6 \text{ GeV}$. It is especially motivated in scenarios where the muon obtains its mass from a different mechanism and a couples only to quarks of the up type, similarly to the scenarios in Ref. [58], although we will not consider flavor violating couplings (which have also been studied recently in Refs. [78,79]) nor the mass range relevant for J/ψ decay.

To estimate the LHCb sensitivity in this channel we perform a novel dedicated analysis. We generate signal events using MG5_AMC@NLO [80] with the Higgs Characterization model [81] and pass them through PYTHIA8 [71] for showering, hadronization, and decays. For the purpose of

computing efficiencies, the signal event generation is performed at NLO in QCD.

The background is expected to be fully dominated by the QCD production of $c\bar{c}$. We use the total background cross section $\sigma^B(c\bar{c}) = 7.1 \pm 3.4$ mb [82] and simulate $c\bar{c}$ events with PYTHIA8 at LO. The idea is that slightly above threshold ($m_a \geq 3.8$ GeV) the ALP decay leads to a fully reconstructable D^+D^- pair whose invariant mass fits into a very narrow bin (≈ 40 MeV) allowing to overcome the huge background.

The LHCb capabilities in terms of particle identification allow identifying D^\pm with nearly 100% efficiency if they decay into $K^-2\pi^+$ ($K^+2\pi^-$) [83], so we select events with at least one D^+ and one D^- each decaying into this mode. The rate of events with at least one D^+D^- pair compared to the total $c\bar{c}$ production is denoted by $f_{c\bar{c} \rightarrow D^+D^-}$. We select events in which all the six decay products K^\pm and $2 \times \pi^\pm$ are within LHCb coverage $2 < \eta < 5$ and $p_T > 0.25$ GeV. The corresponding acceptance is denoted by f_{Acc} . The values obtained for the background QCD $c\bar{c}$ productions are

$$f_{c\bar{c} \rightarrow D^+D^-}^B = 9.86\% \quad f_{Acc}^B = 2.70\% . \tag{9}$$

The corresponding values for the signal for different values of m_a are shown in Table 4.

Lastly, we require the events to fulfill $m(D^+D^-) = m_a \pm 20$ MeV, given the LHCb high invariant mass resolution reconstruction of the D^+D^- system. This last cut, denoted by f_{mass} is almost fully efficient for the signal in the low mass region and allows to dramatically reduce the background rates. This resolution corresponds to approximately $\pm 2\sigma$, where $\sigma \sim 9$ MeV is the D^+D^- invariant mass resolution, taken from the LHCb measurement of the $B^0 \rightarrow D^+D^-$ decay [84]. The invariant masses of the D^+D^- system with a Gaussian smearing according to this resolution are shown, for model M1 with $v/f = 1$ and with $\mathcal{L} = 300 \text{ fb}^{-1}$, in Fig. 9. The cumulative efficiencies of this final selection,

$$\epsilon_{S,B} = f_{c\bar{c} \rightarrow D^+D^-}^{S,B} \times f_{Acc}^{S,B} \times f_{mass}^{S,B}, \tag{10}$$

are shown in Table 4. The signal efficiency drops rapidly for $m_a \gtrsim 5.5$ GeV due to the opening of other decay channels.

Figure 10 shows the resulting bound, obtained by the CLs method at 90% C.L., on the signal cross section $\sigma(pp \rightarrow a) \times \mathcal{B}(a \rightarrow c\bar{c})$ as a function of the ALP mass, assuming an integrated luminosity of $\mathcal{L} = 15 \text{ fb}^{-1}$ (light blue) and $\mathcal{L} = 300 \text{ fb}^{-1}$ (dark blue). Following Ref. [85], systematic uncertainties are expected to be below 0.01%, since the background yields can be also extrapolated here from the invariant mass side bands, and therefore are neglected. For reference, in Fig. 10 (top) we show the central prediction for the signal cross section $\sigma(pp \rightarrow a) \times \mathcal{B}(a \rightarrow c\bar{c})$ of the benchmark models M1, M4, M9, and M11 with $v/f = 1$

Table 4 Cumulative efficiencies (in %) for both signal and background. The background fragmentation fraction $f_{c\bar{c} \rightarrow D^+D^-}^B$ and acceptance f_{Acc}^B are independent on the mass points and are given in (9). $\epsilon_{S,B} = f_{c\bar{c} \rightarrow D^+D^-}^{S,B} \times f_{Acc}^{S,B} \times f_{mass}^{S,B}$. The total number of signal events are given by $S = \sigma(pp \rightarrow a) \times \mathcal{B}(a \rightarrow c\bar{c}) \times \mathcal{B}(D \rightarrow K\pi\pi)^2 \times \epsilon_S \times \mathcal{L}$ and similarly for the background. $\mathcal{B}(D \rightarrow K\pi\pi) = 9.38\%$ is the branching ratio $D^+ \rightarrow K^-2\pi^+$ (and $D^- \rightarrow K^+2\pi^-$), and \mathcal{L} the integrated total luminosity

$m_a[\text{GeV}]$	Signal (in %)			Background (in %)	
	$f_{c\bar{c} \rightarrow D^+D^-}^S$	$f_{c\bar{c} \rightarrow D^+D^-}^S \times f_{Acc}^S$	ϵ_S	ϵ_B	
3.8	22.0	1.71	1.62	0.000390	
4.0	17.7	1.27	1.16	0.000768	
4.2	14.9	1.12	1.04	0.00101	
4.4	14.1	1.02	0.891	0.00122	
4.6	14.1	0.962	0.814	0.00138	
4.8	13.5	0.897	0.691	0.000390	
5.0	12.5	0.818	0.560	0.00152	
5.2	11.8	0.768	0.483	0.00164	
5.4	10.8	0.673	0.307	0.00166	
5.6	10.1	0.636	0.185	0.00167	
5.8	8.89	0.491	0.0109	0.00163	
6.0	9.06	0.475	0.00110	0.00164	

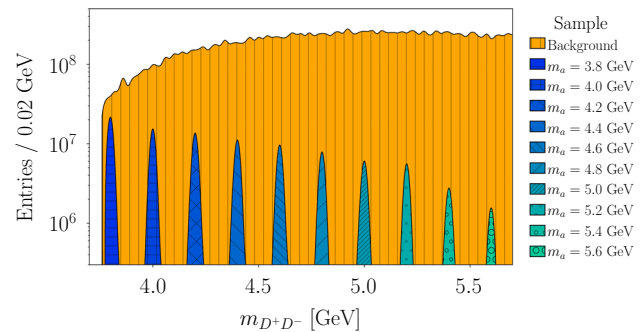


Fig. 9 Signal and background distribution of $m(D^+D^-)$. The yields correspond to 300 fb^{-1} . For signal, the cross-sections are those predicted by model M1 and $v/f = 1$

(see Figs. 1 and 2 for production cross section and branching ratios of other benchmark models). In Fig. 10 (bottom), we show the central prediction for the signal cross section for model M1 with different values of v/f . Figure 11 shows the projected bounds on v/f in the 12 benchmark models, assuming an integrated luminosity of $\mathcal{L} = 15 \text{ fb}^{-1}$ (top) and $\mathcal{L} = 300 \text{ fb}^{-1}$ (bottom), which result from simple scaling of the projected bounds in Fig. 10.

Having discussed the decay into charm pairs as a possible discovery channel, one may wonder why not considering bottom pairs as well. Unfortunately applying the same strategy to the $a \rightarrow b\bar{b}$ channel is not viable, in spite of the larger ALP branching ratio $\mathcal{B}(a \rightarrow b\bar{b}) \approx 21 \times \mathcal{B}(a \rightarrow c\bar{c})$, due

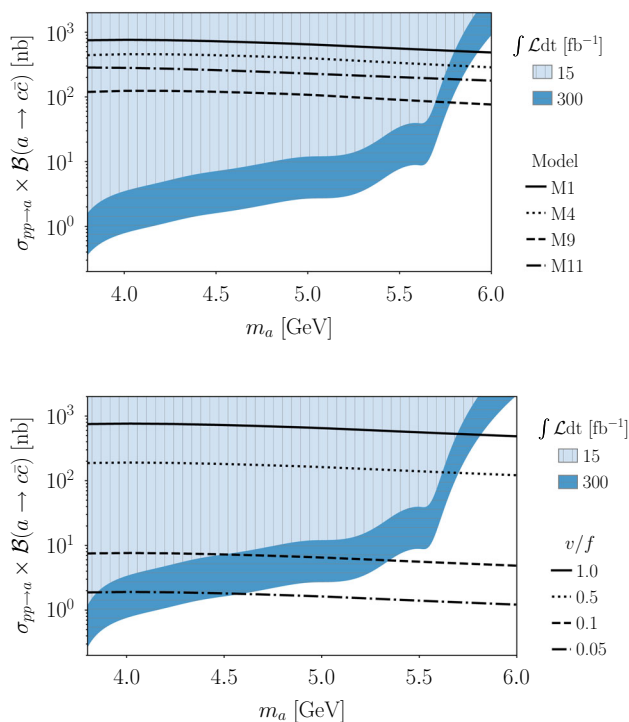


Fig. 10 Model independent projected bounds on $\sigma(pp \rightarrow a) \times \mathcal{B}(a \rightarrow c\bar{c})$ for the reference integrated luminosities. Black lines are central predictions for (top) models M1, M4, M9, and M11 with $v/f = 1$ and (bottom) model M1 with different values of v/f

to the lack of fully reconstructable hadronic decay modes of the B meson with a large branching ratio.

7 Conclusion

Light pseudo-scalar particles are present in many Standard Model extensions. In particular pseudo-Nambu–Goldstone bosons resulting from a spontaneously broken underlying global symmetry may successfully evade current searches even if they are as light as a few GeV. Their couplings to gauge bosons are suppressed as they arise at loop level or through anomaly contributions, while their couplings to fermions are proportional to the fermion masses.

In this article, we provided a first study for the prospects of detecting a light pseudo-scalar at LHCb in $\tau^+\tau^-$ and in $c\bar{c}$ (D^+D^-) channels, with $\sqrt{s} = 14$ TeV and luminosities of 15 fb^{-1} or 300 fb^{-1} . We also compared these new channels to the projected reach of existing searches in $\mu^+\mu^-$ and prospects in di-photons (the main production mode in all these channels is gluon fusion). As benchmarks, we use 12 models of composite Higgs with top partial compositeness, which lead to calculable couplings of the ALP to the SM fermions and gauge bosons. The results can, however, be

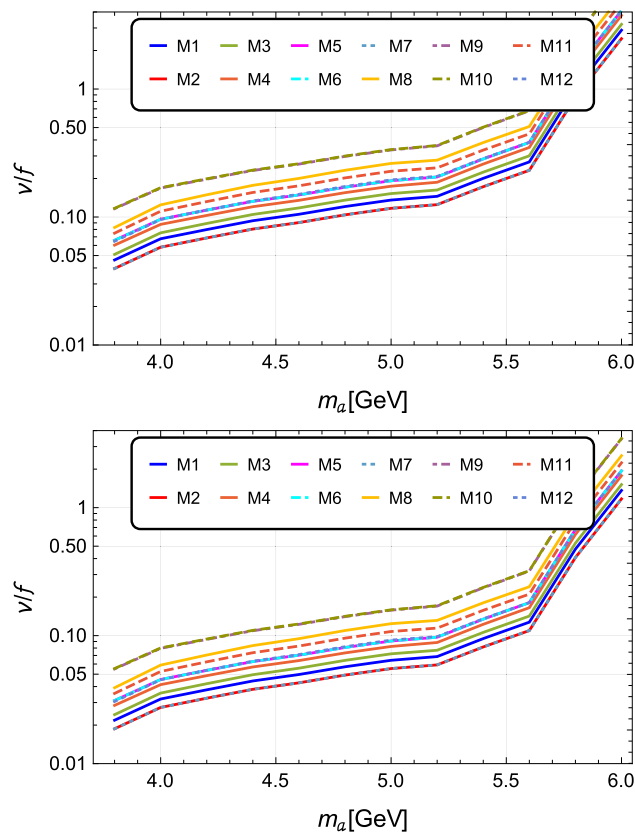


Fig. 11 Bounds on v/f at 90% C.L. using the $c\bar{c}$ analysis for all models with $\mathcal{L} = 15 \text{ fb}^{-1}$ (top) and $\mathcal{L} = 300 \text{ fb}^{-1}$ (bottom)

applied to generic ALP scenarios as well, and model independent projected bounds are shown in Figs. 7 and 10 for $\tau^+\tau^-$ and $c\bar{c}$, respectively.

For the $a \rightarrow \tau^+\tau^-$ channel discussed in Sect. 5 we designed an analysis strategy targeting prompt $a \rightarrow \tau^+\tau^-$ with subsequent di-tau decays in four categories: fully hadronic, semileptonic with a muon, semileptonic with an electron, and fully leptonic ($e\mu$). However, for the computation of the limits we only considered the fully leptonic mode, which is found to be highly dominant over the other three decay categories. We find the exclusion reach on $\sigma(pp \rightarrow a) \times \mathcal{B}(a \rightarrow \tau^+\tau^-)$ given in Fig. 7 for a luminosity of 15 fb^{-1} and 300 fb^{-1} . A dedicated study on the signal efficiencies of the hadronic and semileptonic modes is discussed in Appendix A.

For the $a \rightarrow c\bar{c}$ channel discussed in Sect. 6, we focused on the exclusive final state D^+D^- , which is fully reconstructable at LHCb. This final state is only relevant for masses right above the threshold of 3.8 GeV. The limits quickly deteriorate above 5.5 GeV, yielding limits on $\sigma(pp \rightarrow a) \times \mathcal{B}(a \rightarrow c\bar{c})$ in Fig. 10.

Comparing prospects to find a pseudo-scalar a in the newly proposed $a \rightarrow \tau^+\tau^-$ or $a \rightarrow c\bar{c}$ channels to the already studied $a \rightarrow \mu^+\mu^-$ or $a \rightarrow \gamma\gamma$ channels is inher-

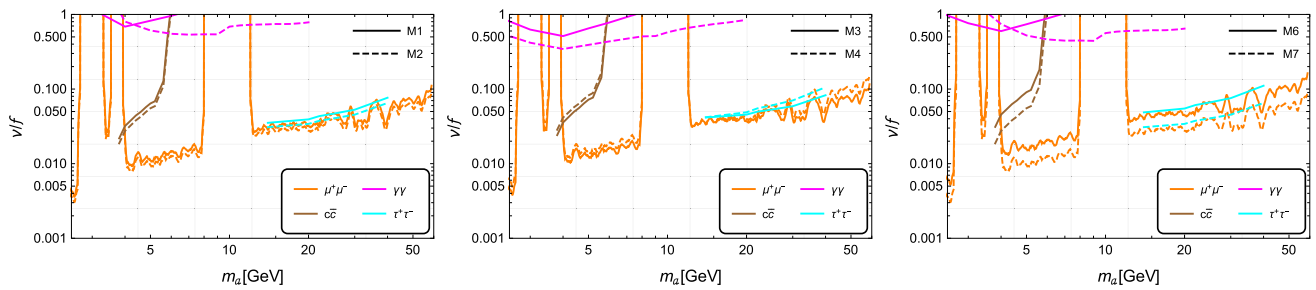


Fig. 12 Summary plots for $\mathcal{L} = 300 \text{ fb}^{-1}$ showing 90% C.L. exclusion bounds on v/f . Solid (dashed) lines refer to M1 (M2), M3 (M4), M6 (M7). Note that the muon bounds are a recast from an actual LHCb

result (or prospects built upon this), while the rest are estimates not based on the full LHCb simulation, so they are not expected to be as accurate

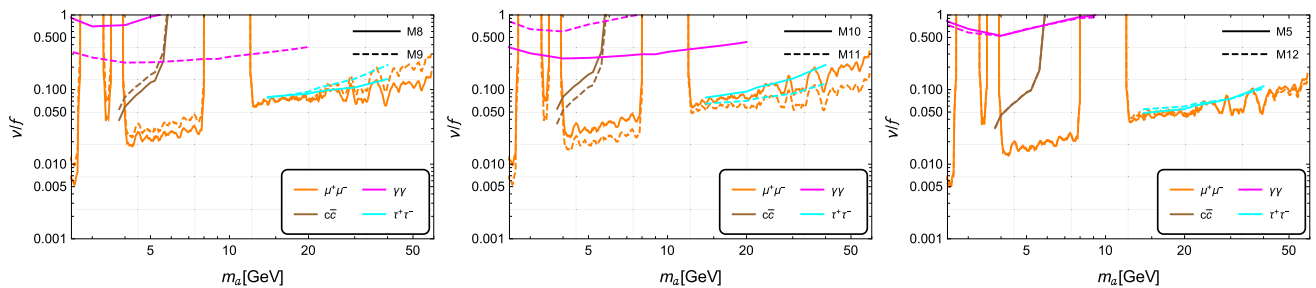


Fig. 13 Summary plots for $\mathcal{L} = 300 \text{ fb}^{-1}$ showing 90% C.L. exclusion bounds on v/f . Solid (dashed) lines refer to M8 (M9), M10 (M11), M5 (M12). Note that the muon bounds are a recast from an actual LHCb

result (or prospects built upon this), while the rest are estimates not based on the full LHCb simulation, so they are not expected to be as accurate

ently model-dependent, as the comparison depends on the branching ratios of a . For the 12 models under consideration we focused on the top couplings that maximise the production cross section in gluon fusion and the overall sensitivity in the fermionic channels at the price of reducing the sensitivity to di-photon, for some models.

As summary, in Figs. 12 and 13 we show the projected limits for a luminosity of 300 fb^{-1} for the 12 models, where models with the same global symmetries at low energy are shown in the same panel (with the exception of M5 and M12 in the last one. See Ref. [38] for details). The plots reveal that the muon channel remains the dominant one, but the new ones, particularly the di-tau, provide comparable and complementary information albeit in a more limited mass range. The photon channel is always minor, even though it provides unique limits in the mass ranges not covered by the muons due to the presence of large background from J/ψ and Υ resonances.

We conclude that LHCb has excellent prospects to investigate light di-tau resonances. For models with an identical coupling to μ and τ ($C_\mu = C_\tau$ in Eq. (1)), the first feasibility study presented here promises an exclusion range which is comparable to the well-established and highly optimized di-muon resonance searches. The di-charm resonance search is applicable only in a small mass range near the $D^+ D^-$ thresh-

old, and yields weaker bounds than the di-muon search (under the assumption $C_c = C_\mu$), but offers bounds which can partially cover the gap left in the di-muon search near the J/ψ resonance.

Finally, we emphasize that searches in the di-tau and di-charm channel are of course interesting in their own right. The assumption of uniform pseudo-scalar fermion coupling is theoretically well motivated in these models but not guaranteed in general (see e.g. Refs. [58, 78] for counter examples).

Acknowledgements We thank A. Mariotti, D. Redigolo, F. Sala, and K. Tobioka for providing the data used for the recast in Fig. 5. We thank M. Spira for explaining details about the ALP decay into gluons. We thank P. Ilten for important suggestions about the simulation of the backgrounds. We are grateful to the Mainz Institute for Theoretical Physics (MITP) for its hospitality and its partial support during the initial stages of this project. DBF and GF are supported by the Knut and Alice Wallenberg foundation under the grant KAW 2017.0100 (SHIFT project). TF’s work is supported by IBS under the project code IBS-R018-D1. The work of XCV is supported by MINECO (Spain) through the Ramón y Cajal program RYC-2016-20073 and by XuntaGAL under the ED431F 2018/01 project. He has also received financial support from XuntaGAL (Centro singular de investigación de Galicia accreditation 2019-2022), the European Union ERDF, the “María de Maeztu” Units of Excellence program MDM-2016-0692, and the Spanish Research State Agency. GC is grateful to the LABEX Lyon Institute of Origins (ANR-10-LABX-0066) of the Université de Lyon for its financial support within the pro-

gram “Investissements d’Avenir” (ANR-11-IDEX-0007) of the French government operated by the National Research Agency (ANR).

Data Availability Statement This manuscript has no associated data or the data will not be deposited. [Authors’ comment: This sensitivity study is solely based on Pythia and MadGraph simulations, which are easily reproducible with the information in the paper.]

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Funded by SCOAP³.

Appendix A: Hadronic and semileptonic τ modes

As mentioned in Sect. 5, we have also conducted a study of the hadronic and semileptonic modes of τ decays.

The QCD component of the background for these modes, dominated by $c\bar{c}$ pairs and jets produced from gluons and light quarks, is overwhelming. Given the fact that our simulation and reconstruction framework becomes substantially slower when the τ 3-prong decay modes are involved (due to the additional selections and due to the vertex reconstruction of the three pions), it is unfeasible to simulate enough QCD background events that pass our full selection. Instead, a proper study of the h_3h_3 , $h_3\mu$, and h_3e categories should be done by using a minimum-bias LHCb dataset of proton-proton collisions, to which we have no access for this study.

Furthermore, our limits are fully dominated by the $e\mu$ decay mode. We have tested this by computing the bound in $\sigma(pp \rightarrow a) \times \mathcal{B}(a \rightarrow \tau^+\tau^-)$ with the CL_s method, using the four categories h_3h_3 , $h_3\mu$, h_3e , and $e\mu$. As a conservative check, we have scaled the $b\bar{b}$ contribution by a large factor to account for the missing sources of background in the QCD component as previously mentioned, leading to almost negligible changes in the combination.

For all the above reasons, the bounds in Sect. 5 are obtained using only the $e\mu$ channel. Nevertheless, we believe it is worth presenting in this appendix the selection and reconstruction procedure, as well as the signal efficiencies, of the hadronic and semileptonic categories to provide a useful input for potential future studies for these modes using LHCb data.

In discussing the semileptonic channel, pions are required to have a minimum p_T of 1 GeV (just as for the electrons and muons), while for the fully hadronic channel this requirement

Table 5 Signal efficiencies for the hadronic and semileptonic modes of $\tau^+\tau^-$, considering the NLO production model in the full acceptance. Mass window requirements are imposed on top of the selection, as described in the caption of Table 6

Mass (GeV)	$\epsilon_{h_3h_3}$ (%)	$\epsilon_{h_3\mu}$ (%)	ϵ_{h_3e} (%)
14	0.130	0.0817	0.0465
20	0.102	0.173	0.109
22	0.271	0.177	0.111
25	0.315	0.221	0.142
30	0.425	0.277	0.187
40	0.433	0.306	0.204

is loosened to 0.5 GeV. In all cases a minimum IP of 0.01 mm, and a minimum momentum of 2 GeV are required. The pions are then combined for all these channels considering all possible three-body combinations of appropriate charge.

The reconstruction procedure goes as follows: the three-pion combination is required to have an invariant mass below 1.7 GeV, a minimum p_T of 10 GeV (2.5 GeV for the h_3h_3 category), an IP smaller than 0.2 mm (0.1 mm for the h_3h_3 category), and a *corrected mass* between 1.2 GeV and 2.5 GeV. In order to determine these quantities, a τ decay vertex is defined as the point in space that minimizes the sum of the distances to the three daughter pions. In addition, a maximum DOCA of 0.05 mm for all two-body combinations of these sets of tracks is required. Finally, for the h_3e and $h_3\mu$ modes we also require $I > 0.99$ for $\Delta R^2 = 0.05$.

The *corrected mass* [86] is defined as $\sqrt{m^2(\pi\pi\pi) + p_T^2(\pi\pi\pi) + p_T(\pi\pi\pi)}$, where the p_T is computed with respect to the τ direction of flight. This quantity, which necessarily requires to know the τ decay vertex and its direction of flight, serves as a good proxy to the *real* invariant mass of the tau, accounting for the presence of an invisible massless particle, and has been widely used in LHCb analyses involving a neutrino in the final state.

We then combine pairs of daughters to reconstruct candidates, computing (pseudo-)decay vertices of the a candidate for the h_3h_3 , $h_3\mu$, and h_3e categories, in a similar way as for the $e\mu$ case. The candidates are reconstructed using two τ leptons for the h_3h_3 category, or one τ lepton and a selected $e(\mu)$ for the $h_3e(h_3\mu)$ category. The cuts are very similar to those of the $e\mu$ candidates. The only exception is the h_3h_3 category, with tighter cuts, being the maximum distance of flight is 0.25 mm and the IP smaller than 0.1 mm. As for the a daughters, the cuts are again the same except for the h_3h_3 . For these, the DOCA should not exceed 0.1 mm and both τ leptons are required to be produced promptly, that is, $0.1\text{mm} < V_r < 5\text{mm}$, being V_r the radial position of their production vertices.

Signal efficiencies and mass windows are reported in Tables 5 and 6.

Table 6 Mass windows for the $\tau^+\tau^-$ hadronic and semileptonic modes

Mass range (GeV)	Category
(6.6, 13.6) 14	h_3h_3
(4.0, 12.2) 14	$h_3\mu$
(4.0, 12.2) 14	h_3e
(9.5, 14.0) 20	h_3h_3
(5.5, 17.5) 20	$h_3\mu$
(6.0, 18.0) 20	h_3e
(10.0, 21.5) 22	h_3h_3
(6.0, 20.0) 22	$h_3\mu$
(6.0, 20.0) 22	h_3e
(12.0, 24.5) 25	h_3h_3
(7.0, 24.0) 25	$h_3\mu$
(7.0, 22.0) 25	h_3e
(13.0, 29.0) 30	h_3h_3
(8.0, 28.0) 30	$h_3\mu$
(7.0, 26.5) 30	h_3e
(18.0, 40.0) 40	h_3h_3
(10.0, 36.0) 40	$h_3\mu$
(10.0, 36.0) 40	h_3e

References

- G. Aad et al., Eur. Phys. J. C **80**(12), 1165 (2020). <https://doi.org/10.1140/epjc/s10052-020-08554-y>
- A. Tumasyan et al. (2021). [arXiv:2109.06055](https://arxiv.org/abs/2109.06055)
- G. Aad et al., Phys. Lett. B **796**, 68 (2019). <https://doi.org/10.1016/j.physletb.2019.07.016>
- A.M. Sirunyan et al., JHEP **07**, 208 (2021). [https://doi.org/10.1007/JHEP07\(2021\)208](https://doi.org/10.1007/JHEP07(2021)208)
- G. Aad et al., JHEP **10**, 061 (2020). [https://doi.org/10.1007/JHEP10\(2020\)061](https://doi.org/10.1007/JHEP10(2020)061)
- G. Cacciapaglia, G. Ferretti, T. Flacke, H. Serôdio, Front. Phys. **7**, 22 (2019). <https://doi.org/10.3389/fphy.2019.00022>
- W. Altmannshofer, S. Gori, D.J. Robinson, Phys. Rev. D **101**(7), 075002 (2020). <https://doi.org/10.1103/PhysRevD.101.075002>
- S. Gori, G. Perez, K. Tobioka, JHEP **08**, 110 (2020). [https://doi.org/10.1007/JHEP08\(2020\)110](https://doi.org/10.1007/JHEP08(2020)110)
- C. Cai, H.H. Zhang, M.T. Frandsen, M. Rosenlyst, G. Cacciapaglia, Phys. Rev. D **102**(7), 075018 (2020). <https://doi.org/10.1103/PhysRevD.102.075018>
- M. Bauer, M. Neubert, S. Renner, M. Schnubel, A. Thamm (2021). [arXiv:2110.10698](https://arxiv.org/abs/2110.10698)
- R. Aaij et al., JHEP **10**, 156 (2020). [https://doi.org/10.1007/JHEP10\(2020\)156](https://doi.org/10.1007/JHEP10(2020)156)
- J.P. Lees et al., Phys. Rev. D **87**(3), 031102 (2013). <https://doi.org/10.1103/PhysRevD.87.031102> [Erratum: Phys. Rev. D **87**, 059903 (2013)]
- S. Chatrchyan et al., Phys. Rev. Lett. **109**, 121801 (2012). <https://doi.org/10.1103/PhysRevLett.109.121801>
- A.M. Sirunyan et al., Phys. Rev. Lett. **124**(13), 131802 (2020). <https://doi.org/10.1103/PhysRevLett.124.131802>
- R. Aaij et al., JHEP **09**, 147 (2018). [https://doi.org/10.1007/JHEP09\(2018\)147](https://doi.org/10.1007/JHEP09(2018)147)
- A. Mariotti, D. Redigolo, F. Sala, K. Tobioka, Phys. Lett. B **783**, 13 (2018). <https://doi.org/10.1016/j.physletb.2018.06.039>
- G. Aad et al., JHEP **01**, 086 (2013). [https://doi.org/10.1007/JHEP01\(2013\)086](https://doi.org/10.1007/JHEP01(2013)086)
- M. Aaboud et al., Phys. Rev. D **95**(11), 112005 (2017). <https://doi.org/10.1103/PhysRevD.95.112005>
- S. Chatrchyan et al., Eur. Phys. J. C **74**(11), 3129 (2014). <https://doi.org/10.1140/epjc/s10052-014-3129-3>
- S. Benson, A. Puig Navarro, Tech. Rep., CERN, Geneva (2018). <https://cds.cern.ch/record/2314368>
- J.P. Lees et al., Phys. Rev. Lett. **107**, 221803 (2011). <https://doi.org/10.1103/PhysRevLett.107.221803>
- R.D. Peccei, H.R. Quinn, Phys. Rev. Lett. **38**, 1440 (1977). <https://doi.org/10.1103/PhysRevLett.38.1440>
- R.D. Peccei, H.R. Quinn, Phys. Rev. D **16**, 1791 (1977). <https://doi.org/10.1103/PhysRevD.16.1791>
- S. Weinberg, Phys. Rev. Lett. **40**, 223 (1978). <https://doi.org/10.1103/PhysRevLett.40.223>
- F. Wilczek, Phys. Rev. Lett. **40**, 279 (1978). <https://doi.org/10.1103/PhysRevLett.40.279>
- J. Jaeckel, A. Ringwald, Annu. Rev. Nucl. Part. Sci. **60**, 405 (2010). <https://doi.org/10.1146/annurev.nucl.012809.104433>
- P. Arias, D. Cadamuro, M. Goodsell, J. Jaeckel, J. Redondo, A. Ringwald, JCAP **06**, 013 (2012). <https://doi.org/10.1088/1475-7516/2012/06/013>
- H.E. Haber, G.L. Kane, Phys. Rep. **117**, 75 (1985). [https://doi.org/10.1016/0370-1573\(85\)90051-1](https://doi.org/10.1016/0370-1573(85)90051-1)
- D.B. Kaplan, H. Georgi, S. Dimopoulos, Phys. Lett. B **136**, 187 (1984). [https://doi.org/10.1016/0370-2693\(84\)91178-X](https://doi.org/10.1016/0370-2693(84)91178-X)
- M.J. Dugan, H. Georgi, D.B. Kaplan, Nucl. Phys. B **254**, 299 (1985). [https://doi.org/10.1016/0550-3213\(85\)90221-4](https://doi.org/10.1016/0550-3213(85)90221-4)
- H.M. Georgi, In: Workshop on Electroweak Symmetry Breaking (1986)
- T.D. Lee, Phys. Rev. D **8**, 1226 (1973). <https://doi.org/10.1103/PhysRevD.8.1226>
- G. Branco, P. Ferreira, L. Lavoura, M. Rebelo, M. Sher, J.P. Silva, Phys. Rep. **516**(1–2), 1–102 (2012). <https://doi.org/10.1016/j.physrep.2012.02.002>
- R.N. Mohapatra, G. Senjanovic, Phys. Rev. D **23**, 165 (1981). <https://doi.org/10.1103/PhysRevD.23.165>
- H. Georgi, M. Machacek, Nucl. Phys. B **262**, 463 (1985). [https://doi.org/10.1016/0550-3213\(85\)90325-6](https://doi.org/10.1016/0550-3213(85)90325-6)
- M.S. Chanowitz, M. Golden, Phys. Lett. B **165**, 105 (1985). [https://doi.org/10.1016/0370-2693\(85\)90700-2](https://doi.org/10.1016/0370-2693(85)90700-2)
- G. Cacciapaglia, C. Pica, F. Sannino, Phys. Rep. **877**, 1 (2020). <https://doi.org/10.1016/j.physrep.2020.07.002>
- A. Belyaev, G. Cacciapaglia, H. Cai, G. Ferretti, T. Flacke, A. Parolini, H. Serodio, JHEP **01**, 094 (2017). [https://doi.org/10.1007/JHEP01\(2017\)094](https://doi.org/10.1007/JHEP01(2017)094) [Erratum: JHEP **12**, 088 (2017)]
- G. Cacciapaglia, G. Ferretti, T. Flacke, H. Serodio, Eur. Phys. J. C **78**(9), 724 (2018). <https://doi.org/10.1140/epjc/s10052-018-6183-4>
- J. Barnard, T. Gherghetta, T.S. Ray, JHEP **02**, 002 (2014). [https://doi.org/10.1007/JHEP02\(2014\)002](https://doi.org/10.1007/JHEP02(2014)002)
- G. Ferretti, D. Karateev, JHEP **03**, 077 (2014). [https://doi.org/10.1007/JHEP03\(2014\)077](https://doi.org/10.1007/JHEP03(2014)077)
- C. Fruguele, E. Fuchs, G. Perez, M. Schlaffer, JHEP **10**, 151 (2018). [https://doi.org/10.1007/JHEP10\(2018\)151](https://doi.org/10.1007/JHEP10(2018)151)
- A.S. Cornell, A. Deandrea, B. Fuks, L. Mason, Phys. Rev. D **102**(3), 035030 (2020). <https://doi.org/10.1103/PhysRevD.102.035030>
- N. Steinberg, J.D. Wells (2021). [arXiv:2101.00520](https://arxiv.org/abs/2101.00520)
- D. d'Enterria (2021). [arXiv:2102.08971](https://arxiv.org/abs/2102.08971)
- I. Bediaga et al. (2012). CERN-LHCC-2012-007
- X. Cid Vidal, C. Vázquez Sierra et al. (2021). [arXiv:2105.12668](https://arxiv.org/abs/2105.12668)
- LHCb Trigger and Online Upgrade Technical Design Report. Tech. rep. (2014). <https://cds.cern.ch/record/1701361>
- LHCb Tracker Upgrade Technical Design Report. Tech. rep. (2014). <https://cds.cern.ch/record/1647400>

50. LHCb VELO Upgrade Technical Design Report. Tech. rep. (2013). <https://cds.cern.ch/record/1624070>
51. J. Albrecht, M.J. Charles, L. Dufour, M.D. Needham, C. Parkes, G. Passaleva, A. Schopper, E. Thomas, V. Vagnoni, M.R.J. Williams, G. Wilkinson, Tech. rep., CERN, Geneva (2019). <https://cds.cern.ch/record/2653011>
52. R. Aaij, et al., Tech. rep. (2018). 10347/15157. <https://cds.cern.ch/record/2636441>. ISBN 978-92-9083-494-6
53. I. Brivio, M.B. Gavela, L. Merlo, K. Mimasu, J.M. No, R. del Rey, V. Sanz, Eur. Phys. J. C **77**(8), 572 (2017). <https://doi.org/10.1140/epjc/s10052-017-5111-3>
54. M. Bauer, M. Neubert, A. Thamm, JHEP **12**, 044 (2017). [https://doi.org/10.1007/JHEP12\(2017\)044](https://doi.org/10.1007/JHEP12(2017)044)
55. M. Bauer, M. Heiles, M. Neubert, A. Thamm, Eur. Phys. J. C **79**(1), 74 (2019). <https://doi.org/10.1140/epjc/s10052-019-6587-9>
56. M. Chala, G. Guedes, M. Ramos, J. Santiago, Eur. Phys. J. C **81**(2), 181 (2021). <https://doi.org/10.1140/epjc/s10052-021-08968-2>
57. M. Bauer, M. Neubert, S. Renner, M. Schnubel, A. Thamm, JHEP **04**, 063 (2021). [https://doi.org/10.1007/JHEP04\(2021\)063](https://doi.org/10.1007/JHEP04(2021)063)
58. A. Carmona, C. Scherb, P. Schwaller (2021). [arXiv:2101.07803](https://arxiv.org/abs/2101.07803)
59. M. Spira, Nucl. Instrum. Meth. A **389**, 357 (1997). [https://doi.org/10.1016/S0168-9002\(97\)00129-0](https://doi.org/10.1016/S0168-9002(97)00129-0)
60. V. Bertone, R. Gauld, J. Rojo, JHEP **01**, 217 (2019). [https://doi.org/10.1007/JHEP01\(2019\)217](https://doi.org/10.1007/JHEP01(2019)217)
61. U. Haisch, J.F. Kamenik, Phys. Rev. D **93**(5), 055047 (2016). <https://doi.org/10.1103/PhysRevD.93.055047>
62. M. Spira, A. Djouadi, D. Graudenz, P.M. Zerwas, Nucl. Phys. B **453**, 17 (1995). [https://doi.org/10.1016/0550-3213\(95\)00379-7](https://doi.org/10.1016/0550-3213(95)00379-7)
63. M. Spira, Fortsch. Phys. **46**, 203 (1998). [https://doi.org/10.1002/\(SICI\)1521-3978\(199804\)46:3<203::AID-PROP203>3.0.CO;2-4](https://doi.org/10.1002/(SICI)1521-3978(199804)46:3<203::AID-PROP203>3.0.CO;2-4)
64. D. Aloni, Y. Soreq, M. Williams, Phys. Rev. Lett. **123**(3), 031803 (2019). <https://doi.org/10.1103/PhysRevLett.123.031803>
65. U. Haisch, J.F. Kamenik, A. Malinauskas, M. Spira, JHEP **03**, 178 (2018). [https://doi.org/10.1007/JHEP03\(2018\)178](https://doi.org/10.1007/JHEP03(2018)178)
66. X. Cid Vidal, A. Mariotti, D. Redigolo, F. Sala, K. Tobioka, JHEP **01**, 113 (2019). [https://doi.org/10.1007/JHEP01\(2019\)113](https://doi.org/10.1007/JHEP01(2019)113) [Erratum: JHEP **06**, 141 (2020)]
67. X. Cid Vidal et al., CERN Yellow Rep. Monogr. **7**, 585 (2019). <https://doi.org/10.23731/CYRM-2019-007.585>
68. R. Aaij et al., Phys. Rev. D **97**(7), 072013 (2018). <https://doi.org/10.1103/PhysRevD.97.072013>
69. R. Aaij et al., Phys. Rev. Lett. **118**(25), 251802 (2017). <https://doi.org/10.1103/PhysRevLett.118.251802>
70. R. Aaij et al., JHEP **09**, 159 (2018). [https://doi.org/10.1007/JHEP09\(2018\)159](https://doi.org/10.1007/JHEP09(2018)159)
71. T. Sjöstrand, S. Ask, J.R. Christiansen, R. Corke, N. Desai, P. Ilten, S. Mrenna, S. Prestel, C.O. Rasmussen, P.Z. Skands, Comput. Phys. Commun. **191**, 159 (2015). <https://doi.org/10.1016/j.cpc.2015.01.024>
72. P. Ilten, Nucl. Phys. B Proc. Suppl. **253–255**, 77 (2014). <https://doi.org/10.1016/j.nuclphysbps.2014.09.019>
73. R. Aaij et al., Phys. Rev. Lett. **118**(5), 052002 (2017). <https://doi.org/10.1103/PhysRevLett.118.052002> [Erratum: Phys. Rev. Lett. **119**, 169901 (2017)]
74. A.M. Sirunyan et al., JHEP **12**, 059 (2019). [https://doi.org/10.1007/JHEP12\(2019\)059](https://doi.org/10.1007/JHEP12(2019)059)
75. P.A. Zyla et al., PTEP **2020**(8), 083C01 (2020). <https://doi.org/10.1093/ptep/ptaa104>
76. T. Junk, Nucl. Instrum. Meth. A **434**, 435 (1999). [https://doi.org/10.1016/S0168-9002\(99\)00498-2](https://doi.org/10.1016/S0168-9002(99)00498-2)
77. A.L. Read, J. Phys. **G28**, 2693 (2002). <https://doi.org/10.1088/0954-3899/28/10/313>
78. H. Davoudiasl, R. Marcarelli, N. Miesch, E.T. Neil (2021). [arXiv:2105.05866](https://arxiv.org/abs/2105.05866)
79. G. Haghighat, M. Mohammadi Najafabadi (2021). [arXiv:2106.00505](https://arxiv.org/abs/2106.00505)
80. J. Alwall, M. Herquet, F. Maltoni, O. Mattelaer, T. Stelzer, JHEP **06**, 128 (2011). [https://doi.org/10.1007/JHEP06\(2011\)128](https://doi.org/10.1007/JHEP06(2011)128)
81. P. Artoisenet et al., JHEP **11**, 043 (2013). [https://doi.org/10.1007/JHEP11\(2013\)043](https://doi.org/10.1007/JHEP11(2013)043)
82. D. d'Enterria, A.M. Snigirev, Phys. Rev. Lett. **118**(12), 122001 (2017). <https://doi.org/10.1103/PhysRevLett.118.122001>
83. L. Collaboration, LHCb PID Upgrade Technical Design Report. Tech. rep. (2013). <https://cds.cern.ch/record/1624074>
84. R. Aaij et al., Phys. Rev. Lett. **117**(26), 261801 (2016). <https://doi.org/10.1103/PhysRevLett.117.261801>
85. R. Aaij et al., Phys. Rev. Lett. **124**(4), 041801 (2020). <https://doi.org/10.1103/PhysRevLett.124.041801>
86. K. Abe et al., Phys. Rev. Lett. **80**, 660 (1998). <https://doi.org/10.1103/PhysRevLett.80.660>