

Setting limits on Effective Field Theories: the case of Dark Matter

Federico Pobbe,^{a,b} Andrea Wulzer^{a,b,c,d} and Marco Zanetti^{a,b}

^a*Dipartimento di Fisica e Astronomia, Università di Padova,
via Marzolo 8, 35131 Padova, Italy*

^b*INFN — Sezione di Padova,
via Marzolo 8, 35131 Padova, Italy*

^c*Theoretical Physics Department, CERN,
Geneva, Switzerland*

^d*Theoretical Particle Physics Laboratory, Institute of Physics,
EPFL, Lausanne, Switzerland*

E-mail: federico.pobbe@pd.infn.it, andrea.wulzer@pd.infn.it,
marco.zanetti@cern.ch

ABSTRACT: The usage of Effective Field Theories (EFT) for LHC new physics searches is receiving increasing attention. It is thus important to clarify all the aspects related with the applicability of the EFT formalism in the LHC environment, where the large available energy can produce reactions that overcome the maximal range of validity, i.e. the cutoff, of the theory. We show that this does not forbid to set rigorous limits on the EFT parameter space through a modified version of the ordinary binned likelihood hypothesis test, which we design and validate. Our limit-setting strategy can be carried on in its full-fledged form by the LHC experimental collaborations, or performed externally to the collaborations, through the Simplified Likelihood approach, by relying on certain approximations. We apply it to the recent CMS mono-jet analysis and derive limits on a Dark Matter (DM) EFT model. DM is selected as a case study because the limited reach on the DM production EFT Wilson coefficient and the structure of the theory suggests that the cutoff might be dangerously low, well within the LHC reach. However our strategy can also be applied, if needed, to EFT's parametrising the indirect effects of heavy new physics in the Electroweak and Higgs sectors.

KEYWORDS: Dark matter, Hadron-Hadron scattering (experiments)

ARXIV EPRINT: [1704.00736](https://arxiv.org/abs/1704.00736)

Contents

1	Introduction	1
2	Hypothesis test with unknown additional signal	5
2.1	The problem	5
2.2	Possible solutions	6
2.2.1	The full story	7
2.2.2	A first simplification	12
2.2.3	Simplified likelihood	14
2.2.4	Cross-section measurements	15
2.2.5	Dealing with interference	16
2.2.6	Why not a signal-strength-based test ?	17
3	CMS mono-jet search reinterpretation	18
3.1	Signal simulation	19
3.2	Setting limits	21
4	Conclusions	23
A	Toy example	26
B	Derivation of the χ^2 formula	29

1 Introduction

The preliminary LHC results made clear that new physics does not assume the “vanilla” form we had imagined for it, leaving us uncertain on how to proceed in the search for new phenomena. The problem is that no single explicit new physics model, through which complete signal predictions could be made, singles out at present as compelling or particularly motivated. Not having to our disposal a sharp signal hypothesis to be compared with the data is obviously a complication. The problem has been addressed, and solved in most cases, by developing comprehensive phenomenological descriptions aimed at parametrising a wide set of new physics models or scenarios in terms of few free parameters. The transition from the constrained MSSM to phenomenological parametrisations of the signal topologies broadly expected in SUSY, advocated in ref. [1], is the paradigmatic example of this approach.

Phenomenological models are useful, but they must be employed with care. It must be kept in mind that they owe their simplicity and generality to the fact that they are designed to describe only a specific set of physical processes. Their few free parameters indeed

correspond to those combinations of the underlying theory parameters that control the processes we decided to focus on. If used outside their domain of validity, their predictions might be incompatible with the new physics scenario they aim at describing, or they correspond to a very specific model and lose their generality. For instance, it is obvious that one should not try to study Dark Matter (DM) in SUSY through a model designed to describe the gluino pair-production topology. Or, one should not use a model for vector resonances production [2] to describe the shape of the signal outside the resonant peak. Making a proper usage of phenomenological models is simple in most cases: it suffices to select for the analysis only those events that are within the domain of validity of the model. However this might not be possible, for instance because the events to which the model applies are not experimentally distinguishable from the others. When this occurs, the situation is more complicated and a special treatment is needed. Addressing this issue for DM searches based on Effective Field Theories (EFT) is the purpose of the present paper. The strategy we will develop can be also applied, if needed, to other EFT-based analyses.

The formalism of low-energy EFT's allows us to produce phenomenological models in a rigorous and systematic way. EFT's provide a complete description of the low-energy dynamics of the lightest particles of a given microscopic (UV) theory, under the sole assumption that the rest of the spectrum sits at an energy scale M_{cut} which is much above the light particles' mass. The validity domain of EFT-based phenomenological models is thus not restricted to a limited set of topologies, but rather to a limited range of energy. All reactions with a center of mass energy E_{cm} below M_{cut} , called the EFT cutoff, are accurately described by the EFT in terms of a few leading EFT interaction operators. The coefficients of these operators, called Wilson coefficients, are free parameters of the EFT. Any UV theory with appropriate low-energy spectrum corresponds to a given value of the Wilson coefficients. By varying the latter ones can thus effectively explore all the UV models with the light particles (and symmetries, which restrict the allowed interaction operators) content we assumed in the construction of the EFT. One can of course also reverse the logic and, starting from experimental results (e.g., from exclusion limits) given in terms of the Wilson coefficients, translate them into constraints on any specific model, in which the Wilson coefficients can be calculated. All this however only works if the EFT is properly employed, i.e. if the comparison of the theory with data relies in no way on the EFT predictions for reactions with E_{cm} at or above the cutoff M_{cut} . As E_{cm} approaches M_{cut} , indeed, the EFT description obtained with the leading operators becomes less and less accurate and more and more sub-leading interactions should be added. At $E_{\text{cm}} \geq M_{\text{cut}}$, infinitely many operators become relevant and the leading-order EFT prediction loses any resemblance to the one of the original theory. In the high energy region the EFT predictions should be replaced with those of the underlying UV theory, which however is precisely what we want to be agnostic about.

EFT's have been employed for new physics searches in a multitude of different contexts. Applications range from flavour and electroweak (EW) precision physics to DM direct and indirect searches and, more recently, to LHC studies of DM [3–10] and of the EW-plus-Higgs sector [11–17]. The limited range of validity of the EFT is never an issue for non-LHC applications, because the experimental conditions automatically forbid E_{cm} to exceed M_{cut} .

For instance in flavour measurements, $E_{\text{cm}} < \text{few GeV}$, while the new physics scale one probes by the EFT (i.e., M_{cut}) is easily above or much above the TeV. A lot of energy is instead available at the LHC and the EFT validity problem must be carefully addressed. This has been studied in ref.s [18–26] in the context of DM and in ref.s [27, 28] for EW-plus-Higgs sector EFT’s. The cases discussed in ref. [27], and in ref. [28] for the neutral Drell-Yan analysis, are very easy to deal with, as E_{cm} can be measured experimentally. By an $E_{\text{cm}} < M_{\text{cut}}$ selection cut one can thus restrict the data to the region where the EFT predictions are reliable. With this cut in place, data can be compared with the EFT prediction with the standard statistical tools (see e.g. [29–32]) producing measurements of the Wilson coefficients or, more often, exclusion limits.¹ Obviously E_{cm} cannot be measured in DM searches because the DM particles escape from the detector. Still we will show that a fully rigorous limit-setting strategy can be set up and used to interpret LHC “mono- X ” DM searches in the EFT context. Our approach is instead not suited to characterise a discovery. We will comment on this point in the Conclusions.

Mono- X DM searches consist in measuring the detector-level Missing Transverse Energy (\cancel{E}_T) distribution in events characterised by the presence of one visible object X . Concretely, X can be a jet [33, 34], a photon [35], a vector or Higgs boson [36–38] or even a top pair [39]. The production of DM contributes to the \cancel{E}_T distribution, or better to its binned version measured in N bins, by an amount

$$n_i^{\text{DM}} = \int_{\text{bin}_i} d\cancel{E}_T \frac{dn_i^{\text{DM}}}{d\cancel{E}_T} = \int_{\text{bin}_i} d\cancel{E}_T \int^{\sqrt{S}} dE_{\text{cm}} \frac{d^2n_i^{\text{DM}}}{dE_{\text{cm}}d\cancel{E}_T}, \quad (1.1)$$

where $d^2n_i^{\text{DM}}$ is the detector-level DM+ X production cross-section times the integrated luminosity, doubly differential in \cancel{E}_T and in the center of mass energy.² The equation has been written in a format that outlines how the DM production signal in each \cancel{E}_T bin receives contributions from all the E_{cm} energy scales up to the total collider energy $\sqrt{S} = 13 \text{ TeV}$ at the LHC run-2. Notice instead that a lower bound (depending on \cancel{E}_T , see e.g. eq. (3.5)) can be set on the E_{cm} integration domain. This is not relevant for the discussion, that focuses on high-energy contributions, and thus it has not been reported in the equation.

The problem with the EFT is that it does not allow us to compute eq. (1.1) fully, but only part of it. Namely, splitting the E_{cm} integral in the regions below and above the cutoff, we have

$$n_i^{\text{DM}} = \int_{\text{bin}_i} d\cancel{E}_T \int^{M_{\text{cut}}} dE_{\text{cm}} \frac{d^2n_i^{\text{DM}}}{dE_{\text{cm}}d\cancel{E}_T} + \int_{\text{bin}_i} d\cancel{E}_T \int_{M_{\text{cut}}}^{\sqrt{S}} dE_{\text{cm}} \frac{d^2n_i^{\text{DM}}}{dE_{\text{cm}}d\cancel{E}_T} \equiv n_i^{\text{EFT}} + \Delta n_i. \quad (1.2)$$

The first term accounts from DM production occurring at $E_{\text{cm}} < M_{\text{cut}}$, i.e. within the domain of validity of the EFT. We can thus compute it by using the EFT, as a function

¹Notice that the limits depend on M_{cut} , which is a free parameter to be scanned over. Limits are thus effectively set in the enlarged parameter space defined by the Wilson coefficients plus the EFT cutoff M_{cut} . This is perfectly correct because the cutoff is indeed one of the parameters of the EFT, as extensively discussed in [24].

²An unambiguous definition of E_{cm} is postponed to section 3.

of its free parameters (including the cutoff), obtaining the prediction n_i^{EFT} . The second contribution, Δn_i , is instead not known from the viewpoint of the EFT because it accounts for “hard” DM production, originating from energy scales where the EFT description no longer applies. All what we know for sure is that it is necessarily positive, $\Delta n_i \geq 0 \forall i$. This being so is obvious from the definition, and physically due to the fact that n_i^{EFT} and Δn_i emerge from two kinematically different (and thus quantum-mechanically distinguishable) processes. As such they cannot interfere, producing two positive contributions that add up to each other. In what follows, we will denote n_i^{EFT} as “the signal”, in spite of the fact that it is only one of the two contributions to the total DM production rate. We will refer to the second contribution, Δn_i , as “unknown additional signal”.

The statistical problem of setting limits in the presence of an unknown additional signal has never been seriously addressed in the literature. A simple but unsatisfactory solution (adopted for instance in ref. [24]) is to turn the experimental analysis into a cut-and-count experiment, i.e. to measure the distribution in a single \cancel{E}_T bin. Together with the estimate of the SM background, this single measurement produces an upper limit, call it $\bar{n}_{\text{exc.}}^{\text{DM}}$, on the total DM production signal (or, equivalently, on the DM production cross-section) one can tolerate in the search region with a given (typically, 95%) Confidence Level (CL). The limit is fully model-independent and it is readily interpreted in the EFT by writing

$$\bar{n}^{\text{EFT}} \leq \bar{n}^{\text{DM}} \leq \bar{n}_{\text{exc.}}^{\text{DM}}. \quad (1.3)$$

This method is rigorous, but largely unsatisfactory. First because the choice of the search region needs to be optimised not only for each given EFT model, but also for each point in the EFT parameter space. This was shown in ref. [24], where the 4 search regions considered by ATLAS in ref. [40] were found to have different limit-setting performances in different regions of the EFT parameter space. Moreover, the cut-and-count approach is never fully optimal as it does not make use of the whole available experimental information. For that, one would need to employ the entire measured distribution, but this is not straightforward in the presence of additional signal. The standard statistical methods to compare a binned histogram with the data crucially rely on the ability of the model to provide a definite prediction for (at least) the *shape* of the distribution. The EFT model instead gives us only a *lower limit* on the distribution and tells us nothing about the shape, which is sensitive to the one of the additional signal component. A slight variation of the standard hypothesis test methodology is needed in order to deal with this situation.

The rest of the paper is organised as follows. In section 2 we discuss several strategies to address the problem of limit-setting in the presence of additional signal. We start from a complete approach, to be adopted by the LHC experimental collaboration that have access to all the details of the analysis, and simplify it down to a level that allows it to be performed outside the collaborations, relying on the minimal amount of experimental information. The usage of simplified likelihoods [41] (see also [42]) emerges as a valid strategy, but alternatives could also be considered. In section 3 we return to DM EFT and set limits on a specific $d = 6$ EFT operator by re-interpreting the CMS mono-jet analysis in ref. [34], for which the simplified likelihood is available. On top of serving as an illustration of the limit-setting method, this example will give us the opportunity to

introduce and to validate the operative definition of the center of mass energy E_{cm} to be employed for the $E_{\text{cm}} < M_{\text{cut}}$ restriction in the EFT signal simulation. In section 4 we report our conclusions and comment on the applicability of our study outside the domain of DM searches. A validation of our limit-setting strategy and some algebraic derivations are reported in appendices A and B, respectively.

2 Hypothesis test with unknown additional signal

2.1 The problem

Consider measuring N observables $\vec{O} = \{O_i\}, i = 1, \dots, N$, whose joint p.d.f. (the probability model) depends on N “known” quantities S_i (the signal), plus N unknown additional signal components $\Delta_i \geq 0$. Namely, the likelihood function is

$$L = L(\vec{S}, \vec{\Delta}; \vec{O}) = \mathcal{P}(\vec{O} | \vec{S} + \vec{\Delta}). \tag{2.1}$$

In what follows we will mostly be interested in interpreting \vec{O} as event countings in N bins, \vec{S} as the expected yields as predicted by the EFT and $\vec{\Delta}$ as the additional signal yield from reactions above the EFT cutoff, in direct correspondence with eq. (1.2). However \vec{O} might also represent parton-level differential cross-section measurements, in which case \vec{S} and $\vec{\Delta}$ are the cross-section predictions from low and high energy processes, respectively. Other parameters $\vec{\nu}$, among which the backgrounds, also influence the probability model. These are routinely treated as nuisance parameters (see e.g., [31, 32]) and constrained by means of auxiliary measurements. Nuisance parameters play no role for the conceptual point we are making here, we thus momentarily ignore their presence.

In order to construct an hypothesis test one defines in some way (typically out of the likelihood function) a test statistic

$$t(\vec{S}, \vec{\Delta}; \vec{O}), \tag{2.2}$$

which is such that a large value of t signals tension of the model with observations. The test statistic is itself a random variable, because of its dependence on \vec{O} . Its distribution follows from the probability model for \vec{O} and therefore it depends, a priori, on \vec{S} and $\vec{\Delta}$ in a highly non trivial manner. Namely

$$t \sim f(t | \vec{S}, \vec{\Delta}). \tag{2.3}$$

The probability for the model to produce an experimental result which is equally or more incompatible than the one observed is quantified by the p -value

$$p(\vec{S}, \vec{\Delta}; \vec{O}) = \int_{t(\vec{S}, \vec{\Delta}; \vec{O})}^{\infty} dt' f(t' | \vec{S}, \vec{\Delta}). \tag{2.4}$$

The model is said to be excluded at a given CL α if $p < 1 - \alpha$, it is not excluded otherwise.

The hypothesis test defined by eq. (2.4) is of course not what we want. It is a test for the probability model of \vec{O} , which requires both \vec{S} and $\vec{\Delta}$ to be specified. Our goal is instead to test a physics model that only predicts \vec{S} and leaves complete freedom on what

the additional signal $\vec{\Delta}$ could be. We thus want our test to be such that the model is excluded only if its prediction are incompatible with the data for any value of $\vec{\Delta}$, provided of course $\Delta_i \geq 0 \forall i$.³ In practice, we want the model to be excluded *only if* it fails the “ordinary” statistical test (i.e., $p < 1 - \alpha$ in eq. (2.4)) *for any* possible choice of the Δ_i ’s. We would also like our test to exclude as much as possible, namely we want the model *not* to be excluded *only if* there exists *at least one* point in the $\vec{\Delta}$ -space that would pass the ordinary test. Both these requirements are fulfilled by defining the p -value as

$$p_{\max}(\vec{S}; \vec{O}) = \sup_{\Delta_i \geq 0} \left\{ p(\vec{S}, \vec{\Delta}; \vec{O}) \right\} = \sup_{\Delta_i \geq 0} \left\{ \int_{t(\vec{S}, \vec{\Delta}; \vec{O})}^{\infty} dt' f(t' | \vec{S}, \vec{\Delta}) \right\}, \quad (2.5)$$

and use it to set a limit in the \vec{S} -space depending on whether $p_{\max} \leq 1 - \alpha$.

Eq. (2.5) is in principle all what we need to set limits in the presence of unknown additional signal, however it is not applicable in concrete. The main obstruction is the need of determining the p.d.f. of the test statistic, which in general cannot be computed analytically and must be obtained numerically by running a toy Monte Carlo. The p.d.f. depends on $\vec{\Delta}$, therefore one toy Monte Carlo would be needed at each point in the $\vec{\Delta}$ -space, resulting in a too demanding numerical procedure. Fortunately this problem is easy to address because the standard test statistic variables employed in LHC analyses are defined in such a way that their p.d.f.s assume a specific form (typically, a χ^2 distribution) under suitable conditions, namely in the so-called Asymptotic Limit (AL). The AL approximation relies on the assumption that a large set of data is employed in the analysis. Further qualifications and checks of its validity are reported in the following section and in appendix A. It suffices here to notice that in the AL the p.d.f. of t becomes independent of $\vec{\Delta}$ (and of \vec{S}) so that eq. (2.5) becomes

$$p_{\max} \stackrel{\text{AL}}{=} \sup_{\Delta_i \geq 0} \left\{ \int_{t(\vec{S} + \vec{\Delta}; \vec{O})}^{\infty} dt' f(t') \right\} = \int_{t_{\min}}^{\infty} dt' f(t'), \quad \text{where } t_{\min}(\vec{S}; \vec{O}) = \inf_{\Delta_i \geq 0} \{t(\vec{S}, \vec{\Delta}; \vec{O})\}. \quad (2.6)$$

The problem thus reduces to the one of minimising t in the $\vec{\Delta}$ space. We will show in the next section how to treat it by a combination of analytical and numerical techniques.

Before concluding this section it is worth noticing that it would be legitimate to regard the $\vec{\Delta}$ ’s as additional nuisance parameters. Exactly like the \vec{v} ’s, indeed, they are parameters whose values we are not interested in testing that however enter in the probability model for the observables. We decided not to adopt this interpretation and we introduced the terminology of “unknown additional signal” because ordinary nuisance parameters are constrained by auxiliary measurements (or by theoretical considerations, for theory-driven errors) while there is no way to constrain the $\vec{\Delta}$ ’s.

2.2 Possible solutions

We discuss now how to deal with eq. (2.5) (actually with its much simpler version (2.6)) in concrete. We begin with the full-fledged treatment, to be carried on directly by the

³Each given UV-completion of the EFT model corresponds to a given value of $\vec{\Delta}$. Asking for the EFT to be excluded for any $\vec{\Delta}$ thus ensures that all its possible UV-completions are excluded.

experimental collaborations starting from the actual data and from the detailed knowledge of the nuisance parameters, and we progressively simplify it.

2.2.1 The full story

In an LHC experimental analysis, the observations \vec{O} are event countings in N disjoint bins. Their p.d.f. is thus the product of N independent Poisson distributions and it is fully specified by the mean values $\vec{M} = \{M_i\}$. The \vec{M} 's consist of three terms

$$\vec{M}(\vec{s}, \vec{\Delta}, \vec{\nu}) = \vec{S}(\vec{s}, \vec{\nu}) + \vec{B}(\vec{\nu}) + \vec{\Delta}, \tag{2.7}$$

where \vec{S} is the signal, \vec{B} is the background, $\vec{\Delta}$ is the additional signal and $\vec{\nu} = \{\nu^{\hat{a}}\}$ ($\hat{a} = 1, \dots, \kappa$) represents the nuisance parameters. The signal is the component of \vec{M} that emerges from the new physics model we aim at testing, i.e. from the EFT with the $E_{\text{cm}} < M_{\text{cut}}$ restriction as in eq. (1.2).⁴ It depends of course on nuisance parameters such as the uncertainties on the luminosity, on the efficiency for trigger and final state reconstruction, etc. Notice that the signal \vec{S} is taken here to depend on N parameters $\vec{s} = \{s_i\}$, one for each histogram bin. Those are meant to be N quantities of purely theoretical nature, calculable within our model with no reference to experimental effects that would induce a dependence on the nuisance parameters. For definiteness, we will think to \vec{s} as the parton-level EFT cross-section in each bin, however we will never need to specify its exact definition because \vec{S} will be directly obtained from simulations performed within the model of interest, with no need of computing it as a function of \vec{s} as an intermediate step. Notice that our approach is different from the standard one [30–32]. In the latter, one takes into account from the very beginning that \vec{S} is actually function of a smaller set of parameters, namely of the free parameters of the BSM model under consideration. One actually goes even further than that, and treats all the BSM parameters as fixed, aside from the total cross-section normalisation μ (the so-called signal-strength). What we are doing here is to first construct an hypothesis test for a generic model in which the cross-section in each bin is a free parameter, and next restrict it to the model of interest in which \vec{s} (and in turn \vec{S}) is function of μ only. The concrete implications of this different approach are described below and the reason why we did not adopt the standard strategy is explained in section 2.2.6.

The other terms in eq. (2.7) are easier to understand. \vec{B} is the “ideally reducible” SM background (namely, the one that does not interfere with the signal, see section 2.2.5), some of whose components are obtained from simulations while others are measured in control regions. The background is affected by nuisance parameters which are not necessarily the same ones appearing in the signal prediction. For instance, the background components estimated from data are themselves nuisance parameters, which obviously do not affect the signal prediction. Our notation, in which all the nuisance are collected in a single vector $\vec{\nu}$, is general enough to account for all possibilities. The last term in eq. (2.7) is the additional

⁴The $E_{\text{cm}} < M_{\text{cut}}$ restriction must be applied to the complete physical signal event yield calculation. Specifically this means that the restricted prediction must include all the SM diagrams, if any, that interfere with the EFT vertex. More details in section 2.2.5.

signal component $\vec{\Delta}$ and it is independent of the nuisance. One might find this confusing if thinking to the additional signal as emerging from an additional term in the parton-level cross-section. A detector simulation would be needed to relate this additional cross-section to the additional signal and a dependence on the nuisance would be introduced. However the additional cross-section is completely arbitrary (but positive) and therefore it produces a completely arbitrary (but positive, $\Delta_i \geq 0$) additional signal. We can thus just take the latter, instead of the additional cross-section, to specify our probability model.

The complete likelihood function reads

$$L[\vec{M}(\vec{s}, \vec{\Delta}, \vec{\nu}), \vec{\nu}; \vec{O}] = \prod_{i=1}^N \frac{M_i^{O_i}}{O_i!} e^{-M_i} \cdot L_{\vec{\nu}}, \quad (2.8)$$

where $L_{\vec{\nu}}$ is the likelihood for the nuisance parameters obtained by a set of auxiliary measurements, such as those performed in control region or in detector calibration studies. Auxiliary measurements being by assumption insensitive to the putative signal, $L_{\vec{\nu}}$ is exclusively a function of $\vec{\nu}$. Out of the complete likelihood in eq. (2.8), the test statistic is defined as the Profile Log Likelihood Ratio

$$t(\vec{s}, \vec{\Delta}; \vec{O}) = -2 \log \left[\sup_{\vec{\nu}} \{L(\vec{M}, \vec{\nu}; \vec{O})\} / \sup_{\vec{\nu}, \vec{s}} \{L(\vec{M}, \vec{\nu}; \vec{O})\} \right]. \quad (2.9)$$

In the numerator, the likelihood is maximised with respect to the nuisance parameters only, while keeping \vec{s} fixed. In the denominator, one maximises with respect to both \vec{s} and $\vec{\nu}$, leading to the absolute maximum of the likelihood in the entire parameter space of the \vec{s} 's and of the $\vec{\nu}$'s. No maximisation is instead performed over the $\vec{\Delta}$'s that are treated, at this stage, as fixed constants. Our first goal is indeed to define a test for each individual $\vec{\Delta}$ hypothesis, to be eventually turned into a $\vec{\Delta}$ -independent test with the strategy outlined in the previous section.

The test statistic in eq. (2.9) perfectly complies with the general definition of Profile Log Likelihood Ratio given in [30, 31, 43, 44], however it differs from the standard definition employed in LHC analyses [31, 32]. In the standard case, as previously explained, the signal is restricted from the very beginning to be the one predicted by the model at hand and the signal-strength μ is taken as the only “parameter of interest”. This means that the supremum in the denominator of the likelihood ratio formula is taken over μ , rather than over the N parameters s_i as in our case. This difference has two implications. The first is that the denominator can be computed analytically in our case. Given that we now have N free parameters to maximise over, we can use them to set $M_i = O_i$ in each bin, reaching in this way the absolute maximum of the Poisson likelihood.⁵ At the same time one of

⁵This might require setting some of the s_i 's to a negative value. Doing so means that we are constructing a test for an extremely general model in which the signal cross-sections s_i 's can even be negative, possibly due to some destructive interference. After constructing the test for such a general model we will apply it to the EFT, in which of course the s_i 's are positive.

course sets $\vec{\nu}$ to whatever value (call it $\vec{\nu}_0$) maximises $L_{\vec{\nu}}$. We can thus rewrite eq. (2.9) as

$$\begin{aligned}
 t &= -2 \log \left[\sup_{\vec{\nu}} \{L(\vec{M}, \vec{\nu}; \vec{O})\} / L(\vec{O}, \vec{\nu}_0; \vec{O}) \right] \\
 &= \inf_{\vec{\nu}} \left\{ 2 \sum_{i=1}^N \left(M_i - O_i - O_i \log \frac{M_i}{O_i} \right) - 2 \log \frac{L_{\vec{\nu}}}{L_{\vec{\nu}_0}} \right\}. \tag{2.10}
 \end{aligned}$$

This formula is the generalisation, including nuisance, of the maximum likelihood goodness-of-fit test for Poisson countings (see e.g. [29]).

The second difference with the ordinary test is that t is now distributed, in the AL, as a χ^2 with N degrees of freedom, unlike the ordinary “signal-strength-based” test statistic that is distributed as a χ^2 with 1 degree of freedom independently of the number of bins. This result straightforwardly follows from the Wilks’ theorem [43, 44], according to which the Profile Log Likelihood Ratio follows a χ^2 distribution with a number of degrees of freedom equal to the number of parameters of interest. The parameters of interest are the vector \vec{s} , of length N , therefore

$$t \stackrel{\text{AL}}{\sim} f_{\chi^2}(t; N) = \frac{1}{2^{N/2} \Gamma(N/2)} t^{N/2-1} e^{-t/2}. \tag{2.11}$$

The similarity of eq. (2.10) with the maximum likelihood goodness-of-fit, which is also (in the AL limit) a χ^2 with N degrees of freedom, is apparent also in this respect.

The AL formula (2.11) is, as previously explained, essential for our construction. We must thus discuss in details the conditions for its validity. The AL is the one in which the data sample is large, a condition that we interpret in a restrictive manner by asking for $O_i \gg 1$ in each bin.⁶ Of course one cannot check the validity of the latter condition before the experiment is performed. Therefore for the design of the hypothesis test, prior to the experiment, we interpret the AL condition as $M_i \gg 1$ for the expected countings in each bin. The idea is that if O_i will eventually turn out to be much different from M_i , for instance if $M_i \gg O_i \not\gg 1$, the hypothesis will be excluded anyhow and the possible unaccuracy of the AL formula in this regime will not be an issue. Of course M_i depends on $\vec{\nu}$, however the nuisance parameters are normally well constrained by the auxiliary measurements so that their variation within few sigmas around the central value $\vec{\nu}_0$ does not change M_i radically. The $M_i \gg 1$ condition can thus be enforced at $\vec{\nu} = \vec{\nu}_0$. The M_i ’s also depend on the unknown additional signal components Δ_i , and clearly we need the AL formula to hold for any Δ_i . Since $\Delta_i \geq 0$, asking for $M_i \gg 1$ at $\Delta_i = 0$, i.e.

$$S_i(\vec{s}, \vec{\nu}_0) + B_i(\vec{\nu}_0) \gg 1, \tag{2.12}$$

is the necessary and sufficient condition for the validity of the AL formula in the entire $\vec{\Delta}$ -space. One might ask how large $S_i + B_i$ concretely needs to be for the AL to hold with good accuracy. In appendix A we address this question quantitatively in a toy example and we find that values as small as 3 are sufficient. Thousands of countings are expected

⁶A similar condition must hold for the auxiliary measurements that constrain $\vec{\nu}$.

in the mono-jet DM search we study in the next section, therefore the number of events is never an issue. The situation might be different for other applications of our method.

In order to construct a $\vec{\Delta}$ -independent hypothesis test we define p_{\max} like in eq. (2.5) and we use the AL formula (2.11) to express it, similarly to eq. (2.6), as

$$p_{\max} \stackrel{\text{AL}}{=} \int_{t_{\min}}^{\infty} dt' f_{\chi^2}(t'; N), \text{ where } t_{\min}(\vec{s}; \vec{O}) = \inf_{\Delta_i \geq 0} \{t(\vec{s}, \vec{\Delta}; \vec{O})\}, \quad (2.13)$$

with t given in eq. (2.10). Notice that t is itself the result of a minimisation, with respect to $\vec{\nu}$. We must further minimise, with respect to $\vec{\Delta}$, in order to obtain t_{\min} . However it is convenient to invert the order of the two minimisations, because the one with respect to $\vec{\Delta}$, if performed first, can be done analytically. The point is that the terms in the round brackets of eq. (2.10) have a unique minimum (equal to zero) at $M_i = O_i$ and monotonically increase (decrease) for M_i larger (smaller) than O_i . The Δ_i 's, one for each bin, are positive shifts of M_i with respect to their minimal values $S_i + B_i$ (see eq. (2.7)), therefore two possibilities are given. If $S_i + B_i \leq O_i$, i.e. if the observed counting over-fluctuates with respect to the prediction, a positive value of Δ_i exists that makes $M_i = O_i$ and the absolute minimum is reached in the corresponding bin. If instead $S_i + B_i > O_i$, i.e. in an under-fluctuating bin, there is no way to reach $M_i = O_i$ with a positive Δ_i . The minimum is thus at $\Delta_i = 0$, i.e. $M_i = S_i + B_i$. The position of the minimum in the $\vec{\Delta}$ -space and the corresponding formula for t_{\min} are

$$\begin{aligned} [\Delta_i(\vec{s}, \vec{\nu})]_{\min} &= \begin{cases} O_i - S_i(\vec{\nu}, \vec{s}) - B_i(\vec{\nu}), & \text{if } S_i(\vec{\nu}, \vec{s}) + B_i(\vec{\nu}) \leq O_i \\ 0, & \text{if } S_i(\vec{\nu}, \vec{s}) + B_i(\vec{\nu}) > O_i \end{cases}, \\ &\Downarrow \\ t_{\min} &= \inf_{\vec{\nu}} \left[\sum_{\{S_i+B_i > O_i\}} 2 \left(S_i(\vec{s}, \vec{\nu}) + B_i(\vec{\nu}) - O_i - O_i \log \frac{S_i(\vec{s}, \vec{\nu}) + B_i(\vec{\nu})}{O_i} \right) - 2 \log \frac{L_{\vec{\nu}}}{L_{\vec{\nu}_0}} \right]. \end{aligned} \quad (2.14)$$

The result is trivial for its simplicity. Over-fluctuating bins do not contribute to t_{\min} because for them we set $M_i = O_i$. The very simple reason for this is that $O_i \geq S_i + B_i$ is perfectly compatible with the hypothesis, the over-fluctuation of the observed being possibly due to some amount of additional signal. Under-fluctuating bins do instead contribute to t_{\min} , and their contribution is conservatively evaluated at $\Delta_i = 0$. Notice that in spite of the fact that they do not contribute to t_{\min} , over-fluctuating bins are implicitly taken into account in our procedure. All the bins are indeed counted in the number of degrees of freedom N of the χ^2 distribution by which compatibility or incompatibility is established. An over-fluctuating (and thus perfectly compatible with the hypothesis, as explained) bin does not contribute to t_{\min} but contributes to N and thus it increases the p -value making the model more compatible with data as it should.

Having defined the test for a generic model, with arbitrary parton-level cross-section in each bin, we can now straightforwardly restrict it to the model of interest, in which \vec{s}

depends on a single signal-strength parameter μ as $s_i(\mu) = \mu \bar{s}_i$.⁷ We just have to set, in eq. (2.14)

$$S_i(\vec{s}, \vec{\nu}) = S_i(\vec{s}(\mu), \vec{\nu}) = \mu \bar{S}_i(\vec{\nu}). \quad (2.15)$$

We stress once again that this step does not require us to first compute \bar{S} as a function of \vec{s} and next compute and plug-in \vec{s} as a function of μ . Everything can be done in one step by simulating the model with the benchmark value for its parameters, that corresponds to $\mu = 1$, compute \bar{S}_i and eventually rescale by μ . Once this is done, the substitution in eq. (2.15) makes t_{\min} become a function of μ (and of the observations) only and the exclusion limit on μ , call it μ_{exc} , is set by solving the equation

$$p_{\text{exc}}(\mu_{\text{exc}}; \vec{O}) = \int_{t_{\min}(\mu_{\text{exc}}; \vec{O})}^{\infty} dt' f_{\chi^2}(t'; N) = 1 - \alpha, \quad (2.16)$$

like for an ordinary hypothesis test.

We described the steps needed to apply our method in a rather pedantic manner, with the purpose of outlining that it does not require more work than the standard LHC limit-setting strategy. It requires signal simulations at different points of the $\vec{\nu}$ -space, in order to assess the impact of nuisance on the signal, the determination of the backgrounds in control regions and the numerical minimisation over the nuisance parameters. All these steps are equally necessary for a standard analysis, which on the other hand cannot deal with additional signal. One additional technical difficulty might actually emerge in our case, due to the fact that whether or not one bin over-fluctuates, and should thus be included in the sum or not, might depend on the value of the nuisance parameters. This occurs when one of the bins has O_i nearly equal to $S_i + B_i$ at the central value $\vec{\nu} = \vec{\nu}_0$ of the nuisance likelihood. In this case nuisance variations around the central value can turn the bin from under-fluctuating to over-fluctuating. The condition of summing over under-fluctuating bins in eq. (2.14) is practically implemented by a step function for each bin in front of the term in the round bracket. These step functions produce discontinuities where their arguments change sign. Actually we don't have discontinuities in the function or in its first derivative, because the term in the round bracket vanishes, together with its derivative, at the singular point $S_i + B_i = O_i$. Discontinuities appear in the second and higher derivatives. Such discontinuities might be impossible to treat for certain minimisation algorithms, in which case they should be regularised and this would result in a more demanding numerical procedure. We are not in the position to assess how severe the problem would actually be in a concrete LHC analysis (we easily dealt with it in the toy example in appendix A), however we are confident that it can be circumvented. Nevertheless in the next section, in the process of simplifying our test, we will describe an approximate strategy that does not require minimising a non doubly-differentiable function.

⁷We are assuming here that there is no interference between the signal and the SM, otherwise the dependence of s on μ is not just a rescaling. We will see in section 2.2.5 how to include interference in the discussion.

2.2.2 A first simplification

A simple approximate version of our test is obtained by relying on three assumptions. The first one is that nuisance parameters $(\nu^{\hat{a}}, \hat{a} = 1, \dots, \kappa)$ have Gaussian likelihood, i.e.

$$L_{\vec{\nu}} = \frac{1}{(\sqrt{2\pi})^\kappa |V|} e^{-\frac{1}{2}(\vec{\nu}-\vec{\nu}_0)^t V^{-1}(\vec{\nu}-\vec{\nu}_0)}, \quad (2.17)$$

with constant ($\vec{\nu}$ -independent) covariance matrix V . The second assumption is that nuisance affects \vec{S} and \vec{B} mildly, so that it is legitimate to Taylor-expand them around $\vec{\nu} = \vec{\nu}_0$

$$M_i = S_i + B_i + \Delta_i \simeq S_i^0 + B_i^0 + \Delta_i + \delta M_i \equiv M_i^0 + \delta M_i, \quad (2.18)$$

$$\delta M_i = (\nu^{\hat{a}} - \nu_0^{\hat{a}}) \partial_{\hat{a}} M_i^0 = (\nu^{\hat{a}} - \nu_0^{\hat{a}}) \left[\frac{\partial S_i}{\partial \nu^{\hat{a}}} + \frac{\partial B_i}{\partial \nu^{\hat{a}}} \right]_{\vec{\nu}=\vec{\nu}_0} \ll M_i^0.$$

This is of course necessarily true if $\vec{\nu}$ is very close to $\vec{\nu}_0$. The assumption we make is that the linear expansion holds when $\vec{\nu}$ is few sigmas (in terms of the covariance matrix V) away from $\vec{\nu}_0$, which is the relevant region for the minimisation of eq. (2.10). Eq. (2.18) thus expresses the condition that the errors on the nuisance are small, in such a way that the few sigma region around $\vec{\nu}_0$ is small and the linear expansion is justified. Notice that few sigma variations of $\vec{\nu}$ resulting in a mild change of M , i.e. $\delta M \ll M$, does not imply that the uncertainties on the nuisance are going to have a mild impact on the limit. Whether or not this is the case depends on how large or small M is in absolute value, which in turn determines the magnitude of the statistical error. In appendix A we will come back to this point and to a quantitative assessment of how small the relative error due to nuisance needs to be for the validity of the approximation.

The third assumption is that the countings are approximately Gaussian distributed, with mean and variance M_i , so that the $\vec{\Delta}$ -dependent test statistic in eq. (2.10) becomes

$$t(\vec{s}, \vec{\Delta}; \vec{O}) = \inf_{\vec{\nu}} \left\{ \sum_{i=1}^N \frac{(M_i - O_i)^2}{M_i} + (\vec{\nu} - \vec{\nu}_0)^t V^{-1} (\vec{\nu} - \vec{\nu}_0) \right\}. \quad (2.19)$$

Poisson countings automatically become Gaussian in the large- M_i limit we already assumed in eq. (2.12) in order to obtain a χ^2 -distributed test statistic as in eq. (2.11). However the Gaussian approximation for the Poisson distribution becomes accurate enough only for values of M_i which are larger (above around few tenths, see appendix A) than those needed for the validity of eq. (2.11).

Few straightforward manipulations, reported for completeness in appendix B, allow us to compute eq. (2.19) analytically in the $\delta M_i \ll M_i$ limit of eq. (2.18). The result is the familiar χ^2 formula

$$t \simeq \chi^2 = \sum_{i,j=1}^N (M_i^0 - O_i) (\Sigma_{\text{tot}}^{-1})_{ij} (M_j^0 - O_j), \quad \Sigma_{\text{tot}} = \Sigma_M + \Sigma_\nu, \quad (2.20)$$

where the total covariance matrix Σ_{tot} is the sum of statistical (Σ_M) and systematical (Σ_ν) errors. Let us discuss them in turn. The statistical error is the variance of the

Poisson distribution in each bin, evaluated at the central value of the nuisance parameters $\vec{\nu} = \vec{\nu}_0$, i.e.

$$(\Sigma_M)_{ij} = M_i^0 \delta_{ij}. \tag{2.21}$$

Obviously it is diagonal, since statistical errors are uncorrelated. The systematical component depends on the nuisance covariance matrix V and on how sensitive the M_i 's are to nuisance departures from the central value. It is given by a formula

$$(\Sigma_\nu)_{ij} = \sum_{\hat{a}, \hat{b}=1}^{\kappa} \partial_{\hat{a}} M_i^0 V_{\hat{a}\hat{b}} \partial_{\hat{b}} M_j^0, \tag{2.22}$$

which seems, at first sight, complicated to apply in a concrete analysis where often neither V nor the functional dependence of \vec{M} on $\vec{\nu}$ are known explicitly. However by the standard error propagation formula we can rewrite it as

$$(\Sigma_\nu)_{ij} = \int d^{\kappa} \nu \mathcal{L}_{\vec{\nu}} \cdot (M_i(\vec{\nu}) - M_i^0)(M_j(\vec{\nu}) - M_j^0) = E [(M_i - M_i^0)(M_j - M_j^0)], \tag{2.23}$$

where “ E ” denotes the expectation value over the nuisance parameters, treated here as random variables with p.d.f. $\mathcal{L}_{\vec{\nu}}$. By sampling the $\vec{\nu}$ -space with weight $\mathcal{L}_{\vec{\nu}}$, computing \vec{M} by means of simulations and taking averages, Σ_ν is easily obtained with no need of determining V and $\partial_{\hat{a}} M_i$ as an intermediate step.⁸

Having expressed $t(\vec{s}, \vec{\Delta}; \vec{O})$ analytically by the χ^2 formula, we should now minimise it over the additional signal components $\Delta_i \geq 0$, as prescribed by eq. (2.13). While this minimisation will in general be performed numerically, it is important to get an analytical understanding of the dependence of t on $\vec{\Delta}$, the way it emerges from the various elements that compose the χ^2 formula (2.20). The first element is M_i^0 , which is given by

$$M_i^0 = \mu \bar{S}_i(\vec{\nu}_0) + B_i(\vec{\nu}_0) + \Delta_i, \tag{2.24}$$

having made use of eq. (2.15) to restrict the signal to the one predicted by our model in terms of the unique signal-strength parameter μ . The M_i 's depend on $\vec{\Delta}$, but in a trivial (additive) manner. The terms \bar{S}_i^0 and B_i^0 are of course independent of $\vec{\Delta}$ and thus they can be computed once and for all by setting the nuisance parameters to their central values. For the nuisance that correspond to instrumental effects (e.g., trigger or reconstruction), this amounts to perform one single signal and background simulation with the nuisance set to their nominal values. For nuisance associated with backgrounds, it amounts to compute the background at the central value obtained from the fit in the control regions. The dependence on μ of the signal is included by rescaling the benchmark simulation. The second element we need in the χ^2 formula is Σ_ν which, importantly enough, is independent of $\vec{\Delta}$. This being the case is obvious from its definition, but also from rewriting eq. (2.23) explicitly as

$$\begin{aligned} (\Sigma_\nu)_{ij} = & \mu^2 E[(\bar{S}_i - \bar{S}_i^0)(\bar{S}_j - \bar{S}_j^0)] + \mu E[(\bar{S}_i - \bar{S}_i^0)(B_j - B_j^0) + i \leftrightarrow j] \\ & + E[(B_i - B_i^0)(B_j - B_j^0)]. \end{aligned} \tag{2.25}$$

⁸A similar logic can be applied to M_i^0 , which is equal to $E[M_i]$. However the approach of first computing $\vec{\nu}_0$ as the maximum of $\mathcal{L}_{\vec{\nu}}$ and next using it to obtain the M_i^0 's is probably more convenient.

This equation allows us to compute Σ_ν as a function of the signal strength parameter by sampling the nuisance space as previously explained. Together with eq. (2.24), it provides us with t as function of μ and $\vec{\Delta}$. After minimising over $\vec{\Delta}$ we obtain $t_{\min}(\mu; \vec{O})$ and in turn μ_{exc} by solving eq. (2.16).

The minimisation of t will have to be performed numerically, but this does not pose any conceptual difficulty as the χ^2 is an infinitely differential function of $\vec{\Delta}$. The procedure can however be rather slow because each function call requires the inversion of Σ_{tot} , which depends on $\vec{\Delta}$ through Σ_M . A considerable simplification is obtained by replacing eq. (2.20) with the “modified” χ^2 formula

$$t \simeq \chi_{\text{mod}}^2 = \sum_{i,j=1}^N (M_i^0 - O_i)(\Sigma_{\text{tot}}^{-1})_{ij}(M_j^0 - O_j), \quad \Sigma_{\text{tot}} = \Sigma_O + \Sigma_\nu, \quad (2.26)$$

where Σ_O is the statistical error matrix computed with the observed countings rather than with the expected ones. The logic behind this simplification is that M_i cannot be much different from O_i , in the limit of high statistic, for configurations that are not trivially excluded. The modified χ^2 is a simple quadratic function of $\vec{\Delta}$, which can be minimised either analytically (though with some complication due to the condition $\Delta_i \geq 0$) or numerically with a fast procedure since no $\vec{\Delta}$ -dependent matrix inversion is involved in the evaluation of the function. The modified χ^2 is known to be a poor approximation of the exact one, that requires large statistics to become accurate. We verify this fact explicitly with our toy example in appendix A. However eq. (2.26) is useful also when the statistics is not large in order to get a first estimate of the location of the minimum in the $\vec{\Delta}$ -space, to be used as a convenient starting point for the minimisation of the exact χ^2 formula.

2.2.3 Simplified likelihood

The limit-setting strategies outlined above require complete knowledge of the data and full control of the sources of systematic uncertainty. Thus they must be carried on internally by the experimental collaborations. Alternative approaches, in which the limit is set externally to the collaborations, are highly desirable because they avoid the presentation of the experimental results being committed to a specific (though quite general, like an EFT) new physics model. The way to proceed is, ideally, rather obvious. The experiments should select a set of quantities calculable with ease by theorists that effectively parametrise a generic signal hypothesis, and report the likelihood as a function of these quantities. If the likelihood is Gaussian, it is of course sufficient to report central values and covariance. In the case of DM, valid quantities to be considered are certainly the parton-level DM plus X production cross-sections in each \cancel{E}_T bin, call them σ_i .⁹ If $L(\vec{\sigma})$ was given, one could straightforwardly apply the logic we outlined in section 2.1 and carry on limit-setting with no need of further experimental information. One would do so by adopting for eq. (2.1) a slightly different interpretation than the one of sections 2.2.1 and 2.2.2. The signal \vec{S} would

⁹These should not be confused with the s_i 's we defined in section 2.2.1. The latter are the EFT contribution to the cross-section, while σ_i is the total cross-section emerging from the sum of the EFT and of the unknown UV contribution.

be the contribution to $\vec{\sigma}$ (call it $\vec{\sigma}_{\text{EFT}}$) coming from the EFT truncated by the $E_{\text{cm}} < M_{\text{cut}}$ restriction. $\vec{\Delta}$ would be the additional parton-level DM production cross-section rather than the additional detector-level signal as in sections 2.2.1 and 2.2.2. The test statistic would be defined as the Log Likelihood Ratio and would be distributed as a χ^2 with N degrees of freedom in the AL. After computing $\vec{\sigma}_{\text{EFT}} = \mu \vec{\sigma}_{\text{EFT}}$ at parton level, plugging in t and maximising over $\vec{\Delta}$ we would end up with eq. (2.16) and use it to set the limit.

The above procedure has limitations, the first one being that it requires quite a bit of additional experimental work. Computing $L(\vec{\sigma})$ requires taking the benchmark model simulation, rescaling it with one signal-strength parameter for each histogram bin and studying the dependence of the likelihood on these parameters. The second difficulty is that it exposes us to errors that is difficult to quantify, associated with the fact that other kinematical distributions (other than \cancel{E}_T) of the benchmark simulation might be different from those of the EFT signal to which the analysis will eventually be applied. Such difference might affect the experimental efficiencies and produce a non-accurate result. Even more worrisome is the fact that the kinematical distributions of the additional signal are completely unknown. Notice that this is not an issue with the approaches of sections 2.2.1 and 2.2.2 because $\vec{\Delta}$ there represents the additional signal yield, whose connection with the additional cross-section needs not to be specified.

Another approach to limit reinterpretation, that does not suffer of the issues above, is based on a ‘‘Simplified Likelihood’’ [41, 42]. The way it works is nicely described by looking at eq. (2.25). Assume the approximations we made in section 2.2.2 are justified, so that t becomes the χ^2 in eq. (2.20), with Σ_ν as in eq. (2.25). Suppose also, and this needs to be checked case-by-case, that none of the relevant sources of nuisance affect the signal and the background at the same time, or that nuisance effects on the signal are negligible. In the mono-jet DM search we discuss in section 3 this is quite a reasonable approximation since data-driven background estimate is the major source of uncertainty. In the hypothesis that signal and background nuisance parameters are disjoint, the mixed terms in eq. (2.25) drop and one obtains

$$(\Sigma_\nu)_{ij} = \mu^2 E[(\bar{S}_i - \bar{S}_i^0)(\bar{S}_j - \bar{S}_j^0)] + E[(B_i - B_i^0)(B_j - B_j^0)] \equiv \mu^2(\Sigma_S)_{ij} + (\Sigma_B)_{ij}. \quad (2.27)$$

The simplified likelihood proposal [41, 42] is that the experimental collaborations report Σ_B , leaving to theorist the determination of Σ_S , if needed. The central value of the background in each bin, B_i^0 , will also be reported, while the evaluation of \bar{S}_i^0 is again left to theory estimate. The potential limitation of this method is that it relies on an accurate simulation of the detector-level signal \bar{S}_i^0 and, even worse, of Σ_S , if not negligible. In section 3 we will apply the simplified likelihood method to the CMS mono-jet search, and we will argue that it should be reasonably accurate for our purpose. Validation from a full-fledged EFT experimental analysis would however be welcome.

2.2.4 Cross-section measurements

The last case which is worth discussing is the one in which the experimental collaborations report, as the result of the analysis, measurements of unfolded parton-level differential

cross-sections. When accurate theoretical predictions of the SM backgrounds are available, this makes setting limits on the EFT a very simple task. Clearly this approach is not applicable to DM searches, where the background cannot be predicted.¹⁰ However it could be useful for other EFT studies such as the ones proposed in refs. [28, 45]. The logic outlined in section 2.1 straightforwardly applies to cross-section measurements. The signal \vec{S} in eq. (2.1) represents now the parton-level EFT cross-sections $\vec{\sigma}^{\text{EFT}}$, in N bins, computed within the EFT with the habitual $E_{\text{cm}} < M_{\text{cut}}$ restriction. $\vec{\Delta}$ is the additional signal cross-section from reactions above the EFT cutoff and the observations \vec{O} are the measured cross-sections $\vec{\sigma}^{\text{m}}$. The test statistic is just the χ^2

$$t(\vec{S}, \vec{\Delta}; \vec{O}) = \chi^2 = \sum_{i,j=1}^N (\vec{\sigma}^{\text{EFT}} + \vec{\Delta} + \vec{b} - \vec{\sigma}^{\text{m}})_i \Sigma_{ij}^{-1} (\vec{\sigma}^{\text{EFT}} + \vec{\Delta} + \vec{b} - \vec{\sigma}^{\text{m}})_j, \quad (2.28)$$

where Σ is the covariance matrix of the measurements. The SM background \vec{b} is taken here to be predicted with infinite accuracy, but of course it would not be hard to include the theory uncertainties on the background prediction in the χ^2 formula by proceeding like in the previous section. The test statistic is distributed as a χ^2 with N d.o.f., and it is very easy to maximise over $\Delta_i \geq 0$, being just a quadratic function. The limit is thus set as a trivial application of eq. (2.6).

2.2.5 Dealing with interference

Until now we described our limit-setting strategy having implicitly (or even explicitly, see eq. (2.15)) in mind the case in which the EFT produces a phenomenon (e.g., DM production) that does not occur in the SM. If this is the case, no quantum mechanical interference is present between the SM and the EFT Feynman diagrams and a clear distinction can be made between the “background” and the “signal”, to be further split into the proper EFT signal and the unknown additional signal from high-energy reactions. Specifically, in DM mono- X searches the background is the complete SM contribution to the \cancel{E}_T distribution, emerging from SM processes with arbitrarily high E_{cm} . The signal is instead DM production from $E_{\text{cm}} < M_{\text{cut}}$, to be computed starting from the EFT diagrams evaluated with the $E_{\text{cm}} < M_{\text{cut}}$ restriction on the phase space. The additional signal is the UV contribution to DM production.

Other interesting EFT’s are those that describe BSM effects in the EW plus Higgs sector. These EFT’s normally produce BSM corrections to processes that do occur also in the SM, and as such they do interfere with the SM diagrams. In the presence of interference our method should be applied as follows. The background must be interpreted as the sum of all SM processes that do not interfere with the EFT because they are characterised by a different parton-level final or initial state. We can call it the “ideally reducible” background, as one might get rid of it by an ideal detector capable to reconstruct all the particles (including, say, the neutrinos) with infinite accuracy. The “signal” is all the rest.

¹⁰In principle, one might consider providing both a measurement of the \cancel{E}_T distribution and a measurement of the SM background from control regions. Addressing the feasibility of this strategy goes beyond the purpose of the present article.

It contains the pure EFT contribution, the SM-EFT interference and the SM prediction for the relevant final state. This is further split into “proper” low-energy EFT signal with the $E_{\text{cm}} < M_{\text{cut}}$ restriction, plus “additional signal” from the UV. Specifically this means that the signal also contains SM terms (namely, the square of SM diagrams), which are also truncated by $E_{\text{cm}} < M_{\text{cut}}$. Similarly, the unknown additional signal is the physical contribution to the distribution that comes from $E_{\text{cm}} > M_{\text{cut}}$ processes and as such it also includes high energy SM contributions. Notice that with this definition the additional signal Δ_i is positive (and the signal as well) because the distinction between the EFT signal and the additional one is made on a physical (kinematical) basis. This would not be the case if we had left the SM out of the truncation and regarded it as part of the background like was done in ref. [46].

With this interpretation in mind one can easily go through the previous sections and check that all the consideration we made remain valid. Our method can thus straightforwardly deal with interference, eq. (2.15) being the only formula that concretely needs to be modified. It assumes the signal being proportional to the signal-strength parameter μ , which is interpreted as the total signal cross-section. In the presence of interference μ should be better viewed as the coefficient of the EFT operator and eq. (2.15) takes the form

$$S_i(\vec{s}, \vec{v}) = S_i(\vec{s}(\mu), \vec{v}) = \mu^2 \bar{S}_i^{(2)}(\vec{v}) + \mu \bar{S}_i^{(1)}(\vec{v}) + \bar{S}_i^{(0)}(\vec{v}). \quad (2.29)$$

The three terms correspond respectively to the square of the EFT diagram, to the interference and to the SM contribution. All the previous formulas that rely on eq. (2.15) can be easily modified according to eq. (2.29).

2.2.6 Why not a signal-strength-based test ?

We extensively discussed in section 2.2.1 that our test is slightly “unusual” because it is first constructed for a generic model in which the signal cross-section in each bin is a free parameter, and later restricted to the model of interest where the signal-strength μ is the only parameter. Ordinary LHC tests are instead “signal-strength-based” from the very beginning, namely they are constructed having in mind that μ is the only parameter of the probability model. Focussing for simplicity on the case in which the interference is absent, the restricted signal is the one in eq. (2.15). In the construction of signal-strength-based tests, this restriction is applied already in eq. (2.9), both in the numerator and the denominator of the likelihood ratio. This makes no difference for the numerator, where we also make use of eq. (2.15), eventually, but it is a big change for the denominator, where now the maximisation is performed over μ only rather than on \vec{s} . The signal-strength-based test has typically a stronger expected limit than the one we used, one might thus wonder why we did not take this direction in our construction. The point is that what the signal-strength-based test actually does is comparing the hypothesis we want to test with the most favourable hypothesis (i.e., the maximal μ) that is present in the restricted set defined by eq. (2.15). If all the hypotheses in the set are far from the data, it returns an artificially high p -value. This is precisely the situation we encounter if we consider very large (positive) values for the additional signal Δ_i , much above the $\mu \bar{S}_i$ contribution

from the EFT and the observed O_i . In this configuration, the total expected is much larger than the observed and the model is in tension with observations for any value of μ . Moreover the likelihood is nearly constant in μ because $\mu \bar{S}_i \ll \Delta_i$, therefore the likelihood divided by its maximum is nearly 1, i.e. $t \simeq 0$, and the signal-strength-based test returns perfect compatibility. Thus a signal-strength-based test cannot be applied in the presence of additional signal because the minimisation of t over $\vec{\Delta}$ would always return $t_{\min} = 0$, with the minimum reached at $\vec{\Delta} \rightarrow \infty$.

3 CMS mono-jet search reinterpretation

As a simple example of DM EFT, we consider DM being a Majorana particle χ in the singlet of the EW group. If DM (with mass m_{DM}) is the lightest particle of its sector, and for energies below the mass “ M_{med} ” of the other new particles (where $M_{\text{med}} \gg m_{\text{DM}}$), its interactions with the quarks are universally described by a set of $d = 6$ effective operators

$$\mathcal{L}_{\text{int}} = \frac{1}{M_*^2} \sum_i c_i O_i, \tag{3.1}$$

classified in ref. [10]. The effective operator Wilson coefficients are conveniently parametrised in terms of dimensionless parameters c_i and of an overall interaction scale M_* . The latter scale should not be confused with the EFT cutoff M_{cut} or with the mass M_{med} of the heavy particles mediating the quark-DM interaction. In order to avoid confusion it would have been convenient to trade it for a parameter $G_* \equiv 1/M_*^2$, analog to the Fermi constant. We will nevertheless adopt the standard notation and use M_* instead. A comprehensive study of all the operators in eq. (3.1) would be interesting, and straightforward with our methodology. However for the sake of simplicity we restrict here to a single axial-axial operator

$$\mathcal{L}_{\text{int}} = -\frac{1}{M_*^2} (\bar{X} \gamma^\mu \gamma^5 X) \left(\sum_q \bar{q} \gamma_\mu \gamma^5 q \right), \tag{3.2}$$

where q runs over the six species of SM quarks. This specific operator is not particularly motivated from a BSM perspective, still it has been extensively studied in the literature and several mediator models have been proposed for its microscopic origin. References can be found in [24].

The EFT we will study is thus endowed with a 3-dimensional parameter space, the parameters being the DM mass m_{DM} , the interaction scale M_* and the cutoff of the EFT M_{cut} . The latter should be regarded and treated, for all practical purposes, as one of the free parameters of the EFT [24]. Notice that there is no way to define rigorously M_{cut} , nor to get a hint of its value, before the UV-completion is specified. This is why it is important to treat it as a free parameter and show how the limits change as a function of M_{cut} . Qualitatively, M_{cut} is of the order of the mass of the mediator particles, M_{med} , but it does not necessarily coincides with that. M_{cut} is the maximal energy at which the EFT predictions resemble those of the UV model, and thus it will typically have to be taken slightly below M_{med} . Deciding how much below, and thus choosing the relevant

m_{DM} [GeV]	1	100	200	400	600	800	1000
$\bar{\sigma}$ [pb]	25	16	11	5.1	2.6	1.3	0.68

Table 1. Total DM pair-production cross-section, for $M_* = 1$ TeV at the 13 TeV LHC.

M_{cut} exclusion contour among the ones we will draw, is left to model-builders aimed at using our results to set limits on their specific UV model. On top of displaying limits in the $m_{\text{DM}}-M_*$ plane at fixed M_{cut} , one can also visualise exclusions at fixed g_* , with g_* defined as

$$G_* = \frac{1}{M_*^2} = \frac{g_*^2}{M_{\text{cut}}^2}. \tag{3.3}$$

The advantage is that g_* ranges in a finite domain because basic perturbativity considerations require $g_* \lesssim 4\pi$ and $g_* \ll 1$ is expected for WIMP-like DM [24].

In what follows, limits are derived on this specific EFT by re-interpreting the CMS 13 TeV mono-jet search [34], for which the Simplified Likelihood is available, as an application of the methodology developed in the previous section.

3.1 Signal simulation

We simulated the DM production signal with MADGRAPH 5 [47], interfaced with PYTHIA 6 [48] for showering and hadronization and with DELPHES [49] for the simulation of the detector response. DM pair-production $pp \rightarrow \chi\chi$ is simulated with up to two parton-level jets in the final state and the resulting event samples are combined by the MLM showering/parton-level matching implemented in MADGRAPH. The effective operator scale has been set to $M_* = 1$ TeV and seven simulations have been performed at the m_{DM} points listed in table 1. The values reported in the table represent the total production cross-section, inclusive in the number of jets, with no cuts.

In a conventional BSM search, the only cuts to be imposed on the generated event sample would be those related with trigger, acceptance and selection that define the search region employed in the experimental analysis.¹¹ For an EFT-based signal instead, the $E_{\text{cm}} < M_{\text{cut}}$ restriction must also be put in place and the question arises of how the center of mass energy E_{cm} should be concretely defined. If the signal was uniquely associated to a single parton-level hard process we would naturally define E_{cm} as the center of mass energy of that partonic process. If for instance the signal was entirely produced by $pp \rightarrow \chi\chi j$, E_{cm} would be the total DM-pair plus parton-level jet invariant mass. Extra low- p_T jets emitted by parton showering must be excluded from the calculation of E_{cm} because those emissions are low virtuality QCD processes, whose occurrence does not invalidate the accuracy of the EFT prediction. While $pp \rightarrow \chi\chi j$ is indeed the dominant process in our sample, after restricting it to the mono-jet search region, the contribution from other partonic configurations is not completely negligible and a more refined definition of E_{cm} is needed. This is constructed by noticing that the purpose of matching algorithms

¹¹In our case [34], the cut that define the signal region are $E_T > 200$ GeV, $p_T > 100$ GeV and $|\eta| < 2.5$ for the leading jet in p_T .

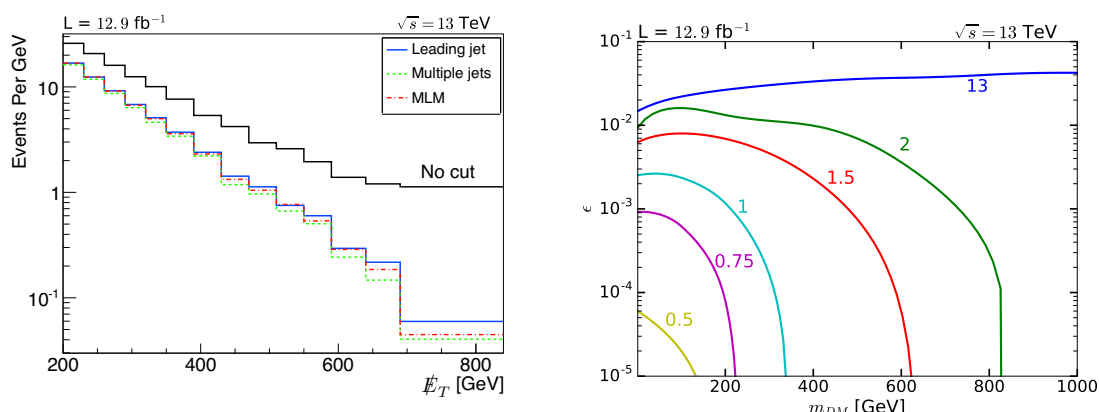


Figure 1. Left panel: comparison among the three definitions of E_{cm} described in the text. The distribution without E_{cm} cut is also shown. Right panel: the efficiency ϵ as a function of m_{DM} at different M_{cut} values.

is precisely to distinguish hard jet emissions, that do contribute to E_{cm} , from soft ones that do not. The definition of E_{cm} is particularly straightforward to construct within the MLM matching algorithm [50], because this algorithm removes all the events in which soft QCD emissions are generated at parton level, so that the soft emissions are exclusively generated by showering in all the events that compose the final sample. The parton-level configuration that produced each event, stored in the output file, is thus the proper hard process to be associated with the event, and E_{cm} is computed out of that. Once the detailed implementation of the algorithm (MLM Kt-jet as implemented in MADGRAPH 5, in our case) and its parameters (“xqcut = 30 GeV”, “ptj=xqcut”, “etaj = 7”, “maxjetflavor = 5” and “QCUT = 100”) are specified, this definition of E_{cm} is unambiguous and fully reproducible.

The E_{cm} definition given above, based on the MLM method, is theoretically accurate and easy to implement, it is thus the one we will use in what follows. However it is interesting to compare it with alternative reasonable definitions, given below

- *Leading jet*: the total invariant mass of the system composed of the two DM particles and of the p_T -leading jet after the jet reconstruction.
- *Multiple jets*: the jets of an event are selected and ordered in p_T . Then the leading jet’s 4-momentum is summed to the one of the DM particles. If the transverse component of the total 4-momentum calculated this way is at least 90% of the event’s \cancel{E}_T , E_{cm} is calculated as the total 4-momentum invariant mass. If this does not happen, the second p_T -ordered jet is considered, its 4-momentum is summed to the previous one and then the transverse component is evaluated again. The procedure goes on until the \cancel{E}_T value is balanced at 90% percent by the jets.

On the left panel of figure 1, the \cancel{E}_T distribution is shown for $m_{\text{DM}} = 200$ GeV, after the cuts specified in footnote 11 and the $E_{\text{cm}} < M_{\text{cut}} = 2$ TeV restriction, with the three different definitions of E_{cm} given above. The distribution without any E_{cm} cut is also shown for

comparison. The three definitions give quite similar results, showing however appreciable differences. The “Leading jet” definition tends to overestimate the EFT signal because it does not count hard extra jets emission in the calculation of E_{cm} , thus producing a lower E_{cm} and in turn a higher signal. The “Multiple jets” instead underestimates the EFT signal, showing that showering emissions give a significant contribution to the transverse momentum unbalance of the event.

After the $E_{\text{cm}} < M_{\text{cut}}$ restriction and with the cuts in footnote 11, the total EFT signal cross-section takes the form

$$\sigma_{\text{EFT}}(M_*, m_{\text{DM}}, M_{\text{cut}}) = \left[\frac{1 \text{ TeV}}{M_*} \right]^4 \cdot \bar{\sigma}(m_{\text{DM}}) \cdot \epsilon(m_{\text{DM}}, M_{\text{cut}}), \quad (3.4)$$

with $\bar{\sigma}$ as in table 1. The scaling with the fourth power of $1/M_*$ is dictated by the one of the effective operator coefficient. The efficiency ϵ is evaluated on the seven m_{DM} points in table 1 and for seven values of M_{cut} :

M_{cut} [GeV]						
500	750	1000	1500	2000	4000	13000

The highest value is of course equivalent to $M_{\text{cut}} = \infty$, namely it assumes that the EFT is valid at all the energy scale the LHC can probe.

It is interesting to inspect the dependence of ϵ on m_{DM} and M_{cut} , shown on the right panel of figure 1. In the Naive EFT, the efficiency monotonically increases with m_{DM} because heavier DM is produced in more energetic reactions, for which it is easier to produce large enough \cancel{E}_T and jet p_T to contribute to the signal region. The efficiency instead decreases with m_{DM} when the E_{cm} cut is in place, because asking for $E_{\text{cm}} < M_{\text{cut}}$ forbids hard reactions and reduces the allowed phase space. The efficiency goes exactly to zero above the kinematical threshold

$$(m_{\text{DM}}^{\text{max}})^2 = \frac{M_{\text{cut}}^2}{4} \left(1 - 2 \frac{\cancel{E}_T}{M_{\text{cut}}} \right). \quad (3.5)$$

3.2 Setting limits

The CMS mono-jet analysis [34] contains all the information we need to set limits on the EFT as a straightforward application of section 2.2.3. Namely it reports (see table 1 and figure 14 of [34]) the observed countings O_i , the expected backgrounds B_i^0 with their covariance matrix Σ_B (see eq. (2.27)), in 22 \cancel{E}_T bins. However we considered the number of countings in the high- \cancel{E}_T bins slightly too small for a safe application of the χ^2 formula (see appendix A), therefore we aggregated bins number 14 to 16 and 17 to 22 in two single bins and we employed an $N = 15$ bins histogram to set the limit. The covariance matrix for the aggregated bins is obtained as in ref. [41]. The simplified likelihood approach prevents a full-fledged treatment of nuisance, but still it allows us to introduce, through Σ_S , nuisance parameters that affect only the signal such as the luminosity and the trigger uncertainties. We verified that those effects do not change our limit appreciably, therefore we set $\Sigma_S = 0$ in what follows.




Figure 2. Limits in the $m_{\text{DM}}-M_*$ plane at fixed M_{cut} (left panel) and g_* (right panel), obtained from the reinterpretation of ref. [34]. The dashed lines on the right plot represent the boundaries of the theoretically forbidden region $2m_{\text{DM}} > M_{\text{cut}} = g_*M_*$ at fixed g_* . In this region, DM production never occurs within the range of validity of the EFT and no limit can be set. The meaning of the “Pure EFT” line on the left plot is explained in the text.

All what is missing to compute the χ^2 in eq. (2.20) are the signal central values S_i^0 , which we obtain by simulations and parametrise by the obvious generalisation of eq. (3.4)

$$S_i^0(M_*, m_{\text{DM}}, M_{\text{cut}}) = \mathcal{L} \cdot \left[\frac{1 \text{ TeV}}{M_*} \right]^4 \cdot \bar{\sigma}(m_{\text{DM}}) \cdot \epsilon_i(m_{\text{DM}}, M_{\text{cut}}), \quad (3.6)$$

where $\mathcal{L} = 12.9 \text{ fb}^{-1}$ is the nominal integrated luminosity and ϵ_i is the efficiency in each bin. Efficiencies are calculated as interpolating functions in the $m_{\text{DM}}-M_{\text{cut}}$ plane, using the 7×7 grid of simulations described above. This allows a fast exploration of the parameter space.

The last step is the minimisation of the χ^2 over the unknown additional signal components $\Delta_i \geq 0$. Out of $\chi_{\text{min}}^2 = t_{\text{min}}$ one obtains the $\alpha = 95\%$ limit by solving eq. (2.16). What plays the role of the signal strength μ here is $1/M_*^4$, therefore the upper exclusion limits on μ will be reported as lower limits, M_*^{exc} , on the effective interaction scale M_* . The minimisation of χ^2 is performed numerically, using as starting point of the minimisation algorithm the minimum of the modified χ^2 in eq. (2.26). Being a quadratic function, the location of the minimum of χ_{mod}^2 can be obtained analytically and it is found to be a very good approximation of the true one. The result of this procedure is shown in figure 2 in the $m_{\text{DM}}-M_*$ plane at fixed M_{cut} (left panel) and g_* (right panel). On the left plot, all the points below the curves are excluded, while in the one on the right what is excluded is the interior of the curves. This peculiar behaviour [24] is due to the fact that at fixed g_* (see eq. (3.3)) low values of M_* correspond to small M_{cut} and the analysis is not sensitive to very small M_{cut} because of the $\cancel{E}_T > 200 \text{ GeV}$ cut (see eq. (3.5)) on the signal region.

The meaning of the black curve on the right panel of figure 2, labeled as “Pure EFT” needs to be explained. The statistical treatment of the unknown additional signal, which eventually led us to our limit-setting strategy, aims at taking rigorously into account the possible contamination of the EFT signal due to high-energy DM production, occurring and $E_{\text{cm}} > M_{\text{cut}}$. However it legitimate to assume that high-energy contributions are small or

absent, for instance because the mass of the mediator particles, and in turn the EFT cutoff, is above the total LHC energy of 13 TeV.¹² Under this restrictive assumption, no unknown additional signal is present and the ordinary limit-setting strategy can be carried on. The result, obtained with the “ \tilde{q}_μ ” test statistic of ref. [32] and with the simplified likelihood, is dubbed “Pure EFT” limit in figure 2, since it assumes no contamination in the signal region from non-EFT reactions. The Pure EFT limit is stronger than the one obtained with $M_{\text{cut}} = 13$ TeV using our procedure. This was expected because the \tilde{q}_μ test is signal-strength based while ours it is not, for the reasons explained in section 2.2.6. Physically this is due to the fact that the Pure EFT test employs the full theoretical information about the shape of the signal distribution, which is partially lost in our analysis because of the additional signal. The difference between the Pure EFT and the 13 TeV lines in figure 2 gives a measure of this effect. Notice that the Pure EFT limit is in some sense more correct than the $M_{\text{cut}} = 13$ TeV one, because assuming $M_{\text{cut}} = 13$ TeV guarantees that no additional signal is present. The limit on the $M_{\text{cut}} = 13$ TeV configuration should thus be set with a signal-strength-based approach that has a stronger expected limit. For lower M_{cut} instead additional signal must be taken into account and our limit-setting strategy is the only viable one.

Validation

Our results rely on two main approximations: the Simplified Likelihood approach and the usage of DELPHES for detector simulation. Both these tools have been validated in the literature, however it is worth cross-checking. We do so by reproducing, using the same tools, the limits on a conventional DM benchmark model (see [34] for details) where DM is a Dirac fermion coupled to a spin-1 Z' -like mediator. The results, obtained with the “ \tilde{q}_μ ” test statistic [32], are reported as grey triangles in figure 3, showing perfect agreement.

4 Conclusions

The absence so far of direct discoveries makes the indirect exploration of heavy new physics by means of EFT’s one of the priorities of the LHC experimental program. Significant progresses in this direction are expected from run-2 and run-3 data, and even more from the High-Luminosity LHC upgrade. It is thus important to clarify all the aspects related with the usage of EFT’s in the LHC environment. In this paper we addressed the issue related with the limited range of validity of EFT’s and we found a rigorous and simple procedure to set limits on the EFT parameter space, duly enlarged to include the EFT cutoff parameter. A full-fledged treatment, to be carried on by the experimental collaborations, is possible and not more complicated than ordinary hypothesis tests. A Simplified Likelihood procedure can also be constructed, relying on approximations that we analysed in great detail.

¹²Actually it is legitimate to assume $M_{\text{cut}} > 13$ TeV only if $M_* \gtrsim 13/4\pi = 1$ TeV, since $M_{\text{cut}} \lesssim 4\pi M_*$ by perturbativity. Given that $M_* = 1$ TeV is right at the boundary of the region that can be excluded by the analysis, $M_{\text{cut}} > 13$ TeV is a rather marginal configuration. However the considerations that follow also apply to mediators lighter than 13 TeV but still in the multi-TeV range such that they are too heavy to be produced.




Figure 3. The grey triangles, superimposed on the CMS limit, show 95% C.L. lower limits on the mediator mass m_{med} for some values of m_{DM} . The figure is obtained within the benchmark model employed by CMS in ref. [34] and for benchmark values, reported on the plot, of its parameters.

The procedure has been applied to the CMS mono-jet search [34], which we reinterpreted as limits on axial-axial DM EFT. The resulting plots, shown in figure 2, can be employed in two ways. First, they provide a semi-quantitative assessment of how heavy the mediator particles (i.e., how large M_{cut}) need to be for the “Pure EFT” limit to hold to good approximation. For instance we conclude that the Pure EFT constraint approximately applies for mediator masses above around 2 TeV. Second, the plot can be quantitatively reinterpreted in specific UV DM models, where the appropriate value of M_{cut} can be worked out. Concrete examples of this reinterpretation are given in ref. [24]. Notice that the bounds one obtains in this way might not be the strongest experimental limits on the UV model at hand. Especially at low M_{cut} , i.e., low mediator mass, mediators are likely to be efficiently constrained by direct searches. UV models will be constrained by combining patches, the EFT limit in figure 2 being one of those patches. The right panel of figure 2 gives another important semi-quantitative information. Namely that WIMP-like DM models, where $g_* \simeq 1$, are poorly constrained by mono-jet DM searches. Direct mediator searches are likely to be more effective in that regime.

The limit-setting strategy we developed can be straightforwardly applied to EFT’s aimed at describing the effects of heavy new physics in the EW-plus-Higgs sector. However whether our approach is really needed or not is question that needs to be addressed case-by-case. For a concreteness, consider the EFT study proposed by one of us in ref. [28], namely the search, in neutral (l^+l^-) and charged ($l\nu$) Drell-Yan at high mass, of two $d = 6$ operators that induce “oblique” corrections to the SM vector bosons propagators. For the neutral Drell-Yan analysis there is obviously no need for our statistical procedure because the center of mass energy of the events can be measured experimentally. When this is the case, the limit on the EFT can be set by using only events with $E_{\text{cm}} < M_{\text{cut}}$, employing the standard statistical tools, as explained in the Introduction. This is not the case for the charged Drell-Yan analysis, where the relevant distribution is the one in transverse mass,

which can receive contributions from arbitrarily high E_{cm} . In principle, charged Drell-Yan should be treated with our strategy, in full analogy with the DM mono- X searches. There is however a substantial quantitative difference. The charged Drell-Yan analysis has such a powerful reach on the EFT Wilson coefficients that it is conceivable to assume that the mediator particles are as heavy as 10 TeV (in “plausible” Universal models, they could be even much higher if other microscopic origins were considered for the EFT operators). It is thus expected (and could be checked) that the mediator particles cannot be produced at a sufficient rate to contribute significantly to the signal region. The “Pure EFT” limit-setting procedure, in which one assumes vanishing additional signal, is thus justified in this case. On the other hand, one might want to study what happens if the mediators are light (but still heavy enough not to be directly seen). This would indeed require applying our limit-setting strategy.

Having in mind limit-setting, in this paper we completely ignored the possibility that a significant excess is observed with respect to the SM prediction. If this occurs, our procedure will obviously return a much larger p -value than expected, and the result might be used to exclude the SM by computing the probability of such a high p -value in the SM hypothesis. However, differently from what happens for a “normal” signal, the dependence of the p -value on the signal-strength μ will not tell us much about its true value. The behaviour of the p -value in the case of a discovery is easily understood by looking at eq. (2.14) and assuming that the observed countings are those predicted by the EFT with signal-strength $\bar{\mu} \neq 0$, plus a certain amount of positive additional signal $\bar{\Delta}_i$ in each bin. If the hypothesis $\mu = 0$ is considered in eq. (2.14), and if the observed are sufficiently above the SM backgrounds B_i , a *small* value of t_{min} will be found, and thus a large p -value, because most of the bins are over-fluctuating and do not contribute to the signal. The $\mu = 0$ hypothesis would thus be found to be fully consistent with the data, which is correct because the over-fluctuation could very well be due to the additional signal rather than to the EFT contribution. The p -value will stay constant and large as μ increases, until $\mu = \bar{\mu}$. At that point the minimisation over Δ_i will of course return $\Delta_i = \bar{\Delta}_i$ and perfect compatibility will be found. The p -value will start decreasing for $\mu > \bar{\mu}$, with a slope that depends on the O_i 's, and thus in turn on true values of the additional signal components $\bar{\Delta}_i$. It is only when μ is large enough for $S_i + B_i$ to overcome O_i in several bins that large t_{min} , and thus small p -value, will be obtain obtained. Unlike a “normal” p -value, ours does not show a peak at $\mu = \bar{\mu}$, but rather an (not necessarily sharp) end point. It can thus be used to set an upper limit on μ , not to measure its value. Therefore, our statistical analysis is not suited to characterise a discovery in terms of the measurement of the EFT Wilson coefficients. Still the EFT can be used to characterise the discovery under the “Pure EFT” assumption $\Delta_i = 0$. If this assumption is not self-consistent, or if it is not consistent with data, explicit UV models should be employed to characterise the discovered signal.

Acknowledgments

We thank P. Checchia and D. Racco for useful discussions. A. W. acknowledges the MIUR-FIRB grant RBF12H1MW and the PRAT grant “Fundamental Physics at the TeV scale”.




Figure 4. The expected p -value, computed at the median over 100 toy Monte Carlos, for $B = 3$ (left) and $B = 10$ (right). The crosses represent the ideal p -value, continuous, dashed and dotted lines represent the full, the χ^2 and the modified χ^2 tests, respectively. The levels of p -value corresponding to 1, 2 and 3 σ are also displayed.

A Toy example

In order to illustrate the validity of the limit-setting strategies developed in sections 2.1, 2.2.1 and 2.2.2, we consider an histogram with $N = 3$ equally spaced bins, constructed for a variable $x \in [0, 3]$. The signal contribution to the expected countings, i.e. the signal from the EFT, is taken to be exponentially distributed with total cross-section μ . The total SM background is called B and the background distribution is also an exponential, but with slightly different slope. Namely

$$\frac{dS}{dx} = \mu \frac{e^{-\lambda_s x}}{1 - e^{-3\lambda_s}}, \quad \frac{dB}{dx} = B \frac{e^{-\lambda_b x}}{1 - e^{-3\lambda_b}}, \quad (\text{A.1})$$

with $\lambda_s = 1/5$ and $\lambda_b = 2/5$. The slopes are chosen in such a way that the signal-over-background ratio is comparable in the three bins. This avoids the limit being dominated by a single bin, in which case our limit-setting strategy would effectively reduce to the one in eq. (1.3). We momentarily assume that the signal and the background are exactly known. We will study the impact of nuisance parameters later.

We consider background values $B = 3$ and $B = 10$, which by integrating eq. (A.1) in the bins lead to $\vec{B} \simeq \{1.4, 0.95, 0.64\}$ and $\vec{B} \simeq \{4.7, 3.2, 2.1\}$, respectively. In each case we generated 100 toy Monte Carlos, i.e. a set of observed \vec{O} vectors distributed according to the background-only hypothesis, and studied the performances of our tests on each of them. The ‘‘ideal’’ test defined by eq. (2.5) is obviously the most difficult one to implement, because it requires the determination of the distribution of the test statistic point-by-point in the $\vec{\Delta}$ -space. Rather than by the Monte Carlo method, we found more convenient to compute the p -value directly by summing the probabilities of the possible outcomes of the experiment that are more incompatible than the actual observation \vec{O} . Namely, we compute, for each μ , $\vec{\Delta}$ and \vec{O}

$$p(\mu, \vec{\Delta}; \vec{O}) = \sum_{o_{1,2,3}=1}^{\infty} \Theta[t(\mu, \vec{\Delta}; \vec{o}) - t(\mu, \vec{\Delta}; \vec{O})] \prod_{i=1}^3 \frac{(\mu \bar{S}_i + B_i + \Delta_i)^{o_i}}{o_i!} e^{-(\mu \bar{S}_i + B_i + \Delta_i)}, \quad (\text{A.2})$$




Figure 5. The expected p -value, computed at the median over 1000 toy Monte Carlos, for $B = 10$, $\delta = 0.1$ (left) and $\delta = 0.01$ (right). Continuous, dashed and dotted lines represent the full, the χ^2 and the modified χ^2 tests, respectively. The levels of p -value corresponding to 1, 2 and 3 σ are also displayed.

where Θ is the step function and the test statistic t is the standard Log Likelihood Ratio, i.e. the argument of the curly bracket in eq. (2.10). The normalised signal $\bar{S}_i \simeq \{0.40, 0.33, 0.27\}$ is defined in eq. (2.15) and is obtained by integrating eq. (A.1) in the bins. The sum over o_i is of course truncated at a maximal value in the calculation, and this maximum is increased in suitably designed step through an iterative procedure. The procedure stops when the relative change of the p -value with respect to the previous step is below 10^{-3} . Afterwards, $p_{\max}(\mu; \vec{O})$ needs to be computed by maximising p over $\Delta_i \geq 0$. This step is greatly facilitated by taking as starting point of the algorithm the maximum of t , which is provided by eq. (2.14). The result of this procedure is reported in figure 4 as a function of μ , together with the “full” p -value (continuous line) defined in section 2.2.1, the one obtained by the χ^2 formula in eq. (2.20) (dashed line) and the “modified” χ^2 in eq. (2.26) (dotted line). The full p -value and the χ^2 are obtained as explained in the main text and don’t pose any computational issue. Notice in particular that the full p -value can be computed analytically in the present example, where no nuisance parameters are included, thanks again to eq. (2.14). The modified χ^2 formula (2.26) is not applicable for $B = 3$, because the expected countings are so low that the observed ones are very likely to have one vanishing entry.

Figure 4 shows the expected p -value, obtained as the median over the possible outcomes, as a function of μ . The agreement of the full version of the test with the ideal one is remarkable, in spite of the fact that the expected background countings are of order one for $B = 3$, while the full test is supposed to hold in the AL where the number of countings is large. This confirms the common lore according to which the statistical distribution of the Log Likelihood Ratio is well-described by the AL formula in eq. (2.11) even if the number of counting is not large. Other AL formulas are much less accurate. For instance the Gaussian approximation for the Poisson distribution is known to become accurate for a number of countings well above 5. Indeed we see in the figure that the χ^2 approximation of the p -value is completely inaccurate for $B = 3$ and that it is not yet satisfactory (though the accuracy improves) for $B = 10$. The modified χ^2 is found to be completely off, compatibly once again with generic expectations. Notice that the agreement of the expected p -value




Figure 6. The expected p -value for $B = 100$, $\delta = 0.1$ (left) and $\delta = 0.01$ (right).

is not strictly speaking sufficient to establish the accuracy of our approximations, which we need to hold for each individual possible outcome of the experiment. We inspected the sample of 100 toy Monte Carlos and verified that for each of them (including those that display large under- and over-fluctuations with respect to the background-only hypothesis) the level of agreement of the different p -value calculation is similar to the one we found for the expected one.

Next, we complicate our probability model in eq. (A.1) by adding five sources of nuisance $\vec{\nu} = \{\hat{\mathcal{L}}, \beta_1, \beta_2, \beta_3, R\}$, corresponding respectively to the integrated luminosity normalised to the nominal one, to the determination of the background in the 3 bins from the control region and to the transfer factor needed to relate the latter measurements to the expected background in the signal region. Namely, the signal and background expected countings are taken to be

$$S_i(\vec{\nu}) = \mu \hat{\mathcal{L}} \bar{s}_i, \quad B_i(\vec{\nu}) = R \beta_i b_i, \quad (\text{A.3})$$

where \bar{s}_i (previously denoted as \bar{S}_i) and b_i (previously, B_i) are the signal and background true values obtained by integrating eq. (A.1) in the bins. The true value of each nuisance parameter is equal to 1 and its “measurement”, i.e. the central value of its likelihood, ν_i^0 , as obtained by auxiliary measurements, is taken to be Gaussian-distributed with standard deviation δ . In order to mimic the possible outcome of repeated measurements we generated 1000 points in the $\vec{\nu}^0$ space, with the previously described distribution, and 1000 toy Monte Carlos for \vec{O} distributed around the background true values b_i . The covariance matrix V that defines the likelihood of the nuisance parameter in eq. (2.17) is $1/\delta^2$ times the identity and it is taken not to fluctuate in the repeated experiments. For each point in the $\vec{\nu}^0$ and \vec{O} space we evaluate the full test in eq. (2.14) and the ones based on the χ^2 and the modified χ^2 in eq.s (2.20) and (2.26), as a straightforward application of the formulas in sections 2.1, 2.2.1 and 2.2.2. The presence of step functions in eq. (2.14) does not pose any computational issue to the automatic minimisation algorithm implemented in MATHEMATICA. We didn’t try to implement the calculation of the ideal p -value, which would probably be too demanding already in our toy example. We consider the results obtained in the no-nuisance case sufficient to conclude that the full p -value is an accurate estimate of the ideal one, for $B \geq 10$, therefore in what follows we will consider the former as reference to assess the validity of the other methods.




Figure 7. The expected p -value for $B = 1000$, $\delta = 0.1$ (left) and $\delta = 0.01$ (right).

Results are obtained for different values of B and δ , starting from $B = 10$ and $\delta = 0.1$ and $\delta = 0.01$, shown in figure 5. The level of agreement is comparable to the one in the right panel of figure 4 and it does not improve substantially lowering δ . This shows that the un-accuracy of the χ^2 and the modified χ^2 approximations is dominated by the low statistics, while the expansion around $\vec{\nu} = \vec{\nu}^0$ we performed in section 2.2.2 to obtain the χ^2 formula does not introduce much additional error. The agreement improves for $B = 100$ (see figure 6), because with larger statistics the Gaussian approximation for the Poisson countings becomes more accurate. The un-accuracy of the χ^2 formula due to the $\vec{\nu} = \vec{\nu}^0$ expansion starts becoming visible in this case, resulting in an improvement of the χ^2 approximation (dashed line) when δ is lowered from 0.1 to 0.01. The behaviour of the modified χ^2 p -value (dotted line) requires further explanations. For $\delta = 0.1$ it provides a rather accurate approximation, but this is an accident due to the fact that for low B (see $B = 10$, $\delta = 0.1$ in figure 5) the modified χ^2 underestimate the true p -value while for large B (see $B = 1000$, $\delta = 0.1$ in figure 7) it overestimates it. The point $B = 100$ happens to be close to the transition between the two regimes. Notice also that for $\delta = 0.01$ (both at $B = 100$ and $B = 1000$) the performances of the modified χ^2 are inferior to those of the χ^2 as expected. We thus recommend the latter for a more accurate (and conservative, in all cases we studied) estimate of the p -value. Finally, the value $B = 1000$ is considered in figure 7. In this case the statistic is large enough for the Gaussian approximation to work extremely well and the expansion around the nuisance central value is the leading source of un-accuracy. The χ^2 approximation becomes indeed substantially exact when δ is lowered from 0.1 to 0.01.

B Derivation of the χ^2 formula

The derivation basically consists in expanding the argument of the “inf” eq. (2.19) up to quadratic order in the nuisance parameters around $\vec{\nu} = \vec{\nu}_0$ as prescribed by eq. (2.18). Afterwards the minimisation over $\vec{\nu}$ can be performed analytically. The only subtlety is that we expand to the second order only the $(M_i - O_i)^2$ numerator, while treating the denominator at the zeroth order, i.e. setting $1/M_i = 1/M_i^0$. This is justified by the fact that the numerator is much more sensitive to the nuisance fluctuations than the denominator. Technically, the contributions to the Taylor series from the expansion of the denominator

are proportional to $(M_i - O_i)^2$, which is small because $M_i^0 - O_i \ll M_i^0$ for configurations that are not trivially excluded. By proceeding in this way we obtain

$$t = \inf_{\vec{\nu}} \left\{ c + 2\vec{b}^t(\vec{\nu} - \vec{\nu}_0) + (\vec{\nu} - \vec{\nu}_0)^t A(\vec{\nu} - \vec{\nu}_0) \right\} = c - \vec{b}^t A^{-1} \vec{b}, \quad (\text{B.1})$$

where \vec{b} and A are respectively a vector and a matrix in the κ -dimensional nuisance space, while c is a constant. They read

$$\begin{aligned} A &= V^{-1} + \sum_{i=1}^N \frac{1}{M_i^0} \vec{\nabla} M_i^0 (\vec{\nabla} M_i^0)^t, \\ \vec{b} &= \sum_{i=1}^N \frac{M_i^0 - O_i}{M_i^0} \vec{\nabla} M_i^0, \\ c &= \sum_{i=1}^N \frac{(M_i^0 - O_i)^2}{M_i^0}. \end{aligned} \quad (\text{B.2})$$

In order to evaluate eq. (B.1) we need to compute

$$\vec{b}^t A^{-1} \vec{b} = \sum_{i,j} \frac{M_i^0 - O_i}{M_i^0} (\vec{\nabla} M_i^0)^t A^{-1} \vec{\nabla} M_j^0 \frac{M_j^0 - O_j}{M_j^0}, \quad (\text{B.3})$$

which in turn is obtained as follows. Start from the identity

$$\mathbb{K} = A^{-1} A = A^{-1} V^{-1} + \sum_k \frac{1}{M_k^0} A^{-1} \vec{\nabla} M_k^0 (\vec{\nabla} M_k^0)^t, \quad (\text{B.4})$$

and multiply it by $V \vec{\nabla} M_j$ on the right and by $(\vec{\nabla} M_i^0)^t$ on the left. We obtain a set of scalar equations in the nuisance space, labeled by i and j that run over the N bins. Those can be expressed as a matrix equation in the bin space, namely

$$\Sigma_\nu = \Sigma_M a (\Sigma_M + \Sigma_\nu) \Rightarrow a = \Sigma_M^{-1} - (\Sigma_M + \Sigma_\nu)^{-1}, \quad (\text{B.5})$$

having defined

$$(\Sigma_M)_{ij} = M_i^0 \delta_{ij}, \quad (\Sigma_\nu)_{ij} = (\vec{\nabla} M_i^0)^t V \vec{\nabla} M_j^0, \quad (a)_{ij} = \frac{1}{M_i^0} (\vec{\nabla} M_i^0)^t A^{-1} \vec{\nabla} M_j^0 \frac{1}{M_j^0}. \quad (\text{B.6})$$

By using eq. (B.5) in eq. (B.3) we obtain

$$\vec{b}^t A^{-1} \vec{b} = \sum_{i,j} (M_i^0 - O_i) [\Sigma_M^{-1} - (\Sigma_M + \Sigma_\nu)^{-1}]_{ij} (M_j^0 - O_j), \quad (\text{B.7})$$

from which eq. (2.20) follows.

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

References

- [1] N. Arkani-Hamed et al., *MARMOSET: the path from LHC data to the new standard model via on-shell effective theories*, [hep-ph/0703088](#) [[INSPIRE](#)].
- [2] D. Pappadopulo, A. Thamm, R. Torre and A. Wulzer, *Heavy vector triplets: bridging theory and data*, *JHEP* **09** (2014) 060 [[arXiv:1402.4431](#)] [[INSPIRE](#)].
- [3] M.W. Goodman and E. Witten, *Detectability of certain dark matter candidates*, *Phys. Rev. D* **31** (1985) 3059 [[INSPIRE](#)].
- [4] M. Beltrán, D. Hooper, E.W. Kolb and Z.C. Krusberg, *Deducing the nature of dark matter from direct and indirect detection experiments in the absence of collider signatures of new physics*, *Phys. Rev. D* **80** (2009) 043509 [[arXiv:0808.3384](#)] [[INSPIRE](#)].
- [5] Q.-H. Cao, C.-R. Chen, C.S. Li and H. Zhang, *Effective dark matter model: relic density, CDMS II, Fermi LAT and LHC*, *JHEP* **08** (2011) 018 [[arXiv:0912.4511](#)] [[INSPIRE](#)].
- [6] M. Beltrán, D. Hooper, E.W. Kolb, Z.A.C. Krusberg and T.M.P. Tait, *Maverick dark matter at colliders*, *JHEP* **09** (2010) 037 [[arXiv:1002.4137](#)] [[INSPIRE](#)].
- [7] J. Goodman, M. Ibe, A. Rajaraman, W. Shepherd, T.M.P. Tait and H.-B. Yu, *Constraints on light Majorana dark matter from colliders*, *Phys. Lett. B* **695** (2011) 185 [[arXiv:1005.1286](#)] [[INSPIRE](#)].
- [8] Y. Bai, P.J. Fox and R. Harnik, *The Tevatron at the frontier of dark matter direct detection*, *JHEP* **12** (2010) 048 [[arXiv:1005.3797](#)] [[INSPIRE](#)].
- [9] J. Fan, M. Reece and L.-T. Wang, *Non-relativistic effective theory of dark matter direct detection*, *JCAP* **11** (2010) 042 [[arXiv:1008.1591](#)] [[INSPIRE](#)].
- [10] J. Goodman, M. Ibe, A. Rajaraman, W. Shepherd, T.M.P. Tait and H.-B. Yu, *Constraints on dark matter from colliders*, *Phys. Rev. D* **82** (2010) 116010 [[arXiv:1008.1783](#)] [[INSPIRE](#)].
- [11] A. Pomarol and F. Riva, *Towards the ultimate SM fit to close in on Higgs physics*, *JHEP* **01** (2014) 151 [[arXiv:1308.2803](#)] [[INSPIRE](#)].
- [12] G. Buchalla, O. Catà and C. Krause, *Complete electroweak chiral lagrangian with a light Higgs at NLO*, *Nucl. Phys. B* **880** (2014) 552 [*Erratum ibid.* **B 913** (2016) 475] [[arXiv:1307.5017](#)] [[INSPIRE](#)].
- [13] J.D. Wells and Z. Zhang, *Precision electroweak analysis after the Higgs boson discovery*, *Phys. Rev. D* **90** (2014) 033006 [[arXiv:1406.6070](#)] [[INSPIRE](#)].
- [14] B. Batell, S. Gori and L.-T. Wang, *Higgs couplings and precision electroweak data*, *JHEP* **01** (2013) 139 [[arXiv:1209.6382](#)] [[INSPIRE](#)].
- [15] R. Alonso, M.B. Gavela, L. Merlo, S. Rigolin and J. Yepes, *The effective chiral lagrangian for a light dynamical “Higgs Particle”*, *Phys. Lett. B* **722** (2013) 330 [*Erratum ibid.* **B 726** (2013) 926] [[arXiv:1212.3305](#)] [[INSPIRE](#)].
- [16] R. Contino, M. Ghezzi, C. Grojean, M. Muhlleitner and M. Spira, *Effective Lagrangian for a light Higgs-like scalar*, *JHEP* **07** (2013) 035 [[arXiv:1303.3876](#)] [[INSPIRE](#)].
- [17] J. Ellis, V. Sanz and T. You, *The effective standard model after LHC Run I*, *JHEP* **03** (2015) 157 [[arXiv:1410.7703](#)] [[INSPIRE](#)].
- [18] J. Goodman and W. Shepherd, *LHC bounds on UV-complete models of dark matter*, [arXiv:1111.2359](#) [[INSPIRE](#)].

- [19] I.M. Shoemaker and L. Vecchi, *Unitarity and Monojet Bounds on Models for DAMA, CoGeNT and CRESST-II*, *Phys. Rev. D* **86** (2012) 015023 [[arXiv:1112.5457](#)] [[INSPIRE](#)].
- [20] S. Profumo, W. Shepherd and T. Tait, *Pitfalls of dark matter crossing symmetries*, *Phys. Rev. D* **88** (2013) 056018 [[arXiv:1307.6277](#)] [[INSPIRE](#)].
- [21] G. Busoni, A. De Simone, E. Morgante and A. Riotto, *On the validity of the effective field theory for dark matter searches at the LHC*, *Phys. Lett. B* **728** (2014) 412 [[arXiv:1307.2253](#)] [[INSPIRE](#)].
- [22] G. Busoni, A. De Simone, J. Gramling, E. Morgante and A. Riotto, *On the validity of the effective field theory for dark matter searches at the LHC, Part II: complete analysis for the s-channel*, *JCAP* **06** (2014) 060 [[arXiv:1402.1275](#)] [[INSPIRE](#)].
- [23] G. Busoni, A. De Simone, T. Jacques, E. Morgante and A. Riotto, *On the validity of the effective field theory for dark matter searches at the LHC part III: analysis for the t-channel*, *JCAP* **09** (2014) 022 [[arXiv:1405.3101](#)] [[INSPIRE](#)].
- [24] D. Racco, A. Wulzer and F. Zwirner, *Robust collider limits on heavy-mediator dark matter*, *JHEP* **05** (2015) 009 [[arXiv:1502.04701](#)] [[INSPIRE](#)].
- [25] S. Matsumoto, S. Mukhopadhyay and Y.-L.S. Tsai, *Singlet Majorana fermion dark matter: a comprehensive analysis in effective field theory*, *JHEP* **10** (2014) 155 [[arXiv:1407.1859](#)] [[INSPIRE](#)].
- [26] S. Matsumoto, S. Mukhopadhyay and Y.-L.S. Tsai, *Effective theory of WIMP dark matter supplemented by simplified models: singlet-like Majorana fermion case*, *Phys. Rev. D* **94** (2016) 065034 [[arXiv:1604.02230](#)] [[INSPIRE](#)].
- [27] R. Contino, A. Falkowski, F. Goertz, C. Grojean and F. Riva, *On the validity of the effective field theory approach to SM precision tests*, *JHEP* **07** (2016) 144 [[arXiv:1604.06444](#)] [[INSPIRE](#)].
- [28] M. Farina, G. Panico, D. Pappadopulo, J.T. Ruderman, R. Torre and A. Wulzer, *Energy helps accuracy: electroweak precision tests at hadron colliders*, *Phys. Lett. B* **772** (2017) 210 [[arXiv:1609.08157](#)] [[INSPIRE](#)].
- [29] G. Cowan, *Statistical data analysis*, Oxford Science Publications, Oxford U.K. (1998).
- [30] PARTICLE DATA GROUP collaboration, K.A. Olive et al., *Review of particle physics*, *Chin. Phys. C* **38** (2014) 090001 [[INSPIRE](#)].
- [31] W.A. Rolke, A.M. Lopez and J. Conrad, *Limits and confidence intervals in the presence of nuisance parameters*, *Nucl. Instrum. Meth. A* **551** (2005) 493 [[physics/0403059](#)] [[INSPIRE](#)].
- [32] G. Cowan, K. Cranmer, E. Gross and O. Vitells, *Asymptotic formulae for likelihood-based tests of new physics*, *Eur. Phys. J. C* **71** (2011) 1554 [*Erratum ibid.* **C 73** (2013) 2501] [[arXiv:1007.1727](#)] [[INSPIRE](#)].
- [33] ATLAS collaboration, *Search for new phenomena in final states with an energetic jet and large missing transverse momentum in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector*, *Phys. Rev. D* **94** (2016) 032005 [[arXiv:1604.07773](#)] [[INSPIRE](#)].
- [34] CMS collaboration, *Search for dark matter produced with an energetic jet or a hadronically decaying W or Z boson at $\sqrt{s} = 13$ TeV*, *JHEP* **07** (2017) 014 [[arXiv:1703.01651](#)] [[INSPIRE](#)].

- [35] ATLAS collaboration, *Search for new phenomena in events with a photon and missing transverse momentum in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *JHEP* **06** (2016) 059 [[arXiv:1604.01306](#)] [[INSPIRE](#)].
- [36] ATLAS collaboration, *Search for dark matter in association with a Higgs boson decaying to b-quarks in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Phys. Lett. B* **765** (2017) 11 [[arXiv:1609.04572](#)] [[INSPIRE](#)].
- [37] CMS collaboration, *Search for associated production of dark matter with a Higgs boson decaying to $b\bar{b}$ or $\gamma\gamma$ at $\sqrt{s} = 13$ TeV*, [arXiv:1703.05236](#) [[INSPIRE](#)].
- [38] CMS collaboration, *Search for dark matter and unparticles in events with a Z boson and missing transverse momentum in proton-proton collisions at $\sqrt{s} = 13$ TeV*, *JHEP* **03** (2017) 061 [[arXiv:1701.02042](#)] [[INSPIRE](#)].
- [39] CMS collaboration, *Search for the production of dark matter in association with top-quark pairs in the single-lepton final state in proton-proton collisions at $\sqrt{s} = 8$ TeV*, *JHEP* **06** (2015) 121 [[arXiv:1504.03198](#)] [[INSPIRE](#)].
- [40] ATLAS collaboration, *Search for new phenomena in monojet plus missing transverse momentum final states using 10 fb^{-1} of pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector at the LHC*, *ATLAS-CONF-2012-147* (2012).
- [41] CMS collaboration, *Simplified likelihood for the re-interpretation of public CMS results*, *CMS-NOTE-2017-001* (2017).
- [42] S. Fichtel, *Taming systematic uncertainties at the LHC with the central limit theorem*, *Nucl. Phys. B* **911** (2016) 623 [[arXiv:1603.03061](#)] [[INSPIRE](#)].
- [43] S.S. Wilks, *The large-sample distribution of the likelihood ratio for testing composite hypotheses*, *Annals Math. Statist.* **9** (1938) 60.
- [44] A. Wald, *Tests of statistical hypotheses concerning several parameters when the number of observations is large*, *Trans. Amer. Math. Soc.* **54** (1943) 426.
- [45] J. de Blas, M. Chala and J. Santiago, *Global constraints on lepton-quark contact interactions*, *Phys. Rev. D* **88** (2013) 095011 [[arXiv:1307.5068](#)] [[INSPIRE](#)].
- [46] A. Falkowski, M. Gonzalez-Alonso, A. Greljo, D. Marzocca and M. Son, *Anomalous triple gauge couplings in the effective field theory approach at the LHC*, *JHEP* **02** (2017) 115 [[arXiv:1609.06312](#)] [[INSPIRE](#)].
- [47] J. Alwall, M. Herquet, F. Maltoni, O. Mattelaer and T. Stelzer, *MadGraph 5: going beyond*, *JHEP* **06** (2011) 128 [[arXiv:1106.0522](#)] [[INSPIRE](#)].
- [48] T. Sjöstrand, S. Mrenna and P.Z. Skands, *PYTHIA 6.4 physics and manual*, *JHEP* **05** (2006) 026 [[hep-ph/0603175](#)] [[INSPIRE](#)].
- [49] DELPHES 3 collaboration, J. de Favereau et al., *DELPHES 3, a modular framework for fast simulation of a generic collider experiment*, *JHEP* **02** (2014) 057 [[arXiv:1307.6346](#)] [[INSPIRE](#)].
- [50] J. Alwall et al., *Comparative study of various algorithms for the merging of parton showers and matrix elements in hadronic collisions*, *Eur. Phys. J. C* **53** (2008) 473 [[arXiv:0706.2569](#)] [[INSPIRE](#)].