

On the origin of long-lived particles

Jared Barron and David Curtin

*Department of Physics, University of Toronto,
60 St. George St., Toronto, ON, Canada*

E-mail: jared.barron@mail.utoronto.ca, dcurtin@physics.utoronto.ca

ABSTRACT: MATHUSLA is a proposed large-volume displaced vertex (DV) detector, situated on the surface above CMS and designed to search for long-lived particles (LLPs) produced at the HL-LHC. We show that a discovery of LLPs at MATHUSLA would not only prove the existence of BSM physics, it would also uncover the theoretical origin of the LLPs, despite the fact that MATHUSLA gathers no energy or momentum information on the LLP decay products. Our analysis is simple and robust, making it easily generalizable to include more complex LLP scenarios, and our methods are applicable to LLP decays discovered in ATLAS, CMS, LHCb, or other external detectors. In the event of an LLP detection, MATHUSLA can act as a Level-1 trigger for the main detector, guaranteeing that the LLP production event is read out at CMS. We perform an LLP simplified model analysis to show that combining information from the MATHUSLA and CMS detectors would allow the LLP production mode topology to be determined with as few as ~ 100 observed LLP decays. Underlying theory parameters, like the LLP and parent particle masses, can also be measured with $\lesssim 10\%$ precision. Together with information on the LLP decay mode from the geometric properties of the observed DV, it is clear that MATHUSLA and CMS together will be able to characterize any newly discovered physics in great detail.

KEYWORDS: Beyond Standard Model, Hadron-Hadron scattering (experiments), Lifetime

ARXIV EPRINT: [2007.05538](https://arxiv.org/abs/2007.05538)

Contents

1	Introduction	1
2	The MATHUSLA detector	3
2.1	Detector overview	4
2.2	Physics reach	4
2.3	Measuring LLP boost	5
3	LLP simplified models	6
3.1	LLP production modes	7
3.2	Event generation	9
3.3	LLP decay in MATHUSLA	9
4	Determining production mode	10
4.1	Hierarchical classification algorithm	11
4.2	Performance	17
5	Determining model parameters	20
5.1	Exotic B -meson decay	22
5.2	Charged current	23
5.3	Heavy parent	25
5.4	Exotic Higgs decay	27
5.5	Direct pair production	28
5.6	Heavy resonance	28
6	Uncertainty in matching LLP decay to production vertex	31
6.1	Impact on production mode classification	31
6.2	Impact on parameter estimation	32
7	Conclusions	33
A	Simulation details	34

1 Introduction

Since the Higgs boson was discovered by the Large Hadron Collider (LHC) experiments ATLAS and CMS in 2012 [1, 2], the Standard Model (SM) of particle physics has been complete, with every fundamental particle predicted having been found. Nonetheless, many questions remain unresolved, including the nature of the dark matter, neutrino masses, and baryogenesis, to name a few [3–6]. All of these require the existence of physics beyond the

Standard Model (BSM). In particular, symmetry-based solutions to the hierarchy problem of the Higgs boson require the existence of TeV-scale particles to avoid fine-tuning of the Higgs boson mass. Motivated by theories like supersymmetry (SUSY) [7] as well as composite or extra-dimensional models that can naturally give rise to such TeV-scale particles and solve the hierarchy problem [8], the efforts of the experimental particle physics community at the LHC have been largely devoted to searches for heavy, promptly decaying particles as the most likely avenue for discovery. The LHC main detectors are explicitly designed for such a search program. Unfortunately, despite strong theoretical motivations and a decade of searching, no BSM particles have been discovered to date [9–12].

As many canonical models to solve issues like the hierarchy problem are becoming highly constrained, it is important to consider other possible signatures of BSM physics that may have escaped notice so far. One of these possibilities is that BSM particles may exist at masses accessible by current colliders, but are long-lived rather than promptly decaying [13]. Particles with macroscopic decay lengths can elude detection at large collider experiments because of triggers designed for prompt decays, complicated backgrounds, and low acceptance due to a high probability of escaping the detector before decaying if the lifetime is long. Furthermore, LLPs are highly theoretically motivated in their own right, both from a bottom-up point of view, since various mechanisms resulting in meta-stable states in the SM can equally well operate in BSM theories, as well as from a top-down point of view, since LLPs arise and are instrumental in many BSM scenarios postulated to explain the hierarchy problem, dark matter, baryogenesis, and neutrino masses [14].

Recently, several new auxiliary detectors [15–17] and CERN experiments [18] have been proposed to extend the LLP reach of the existing LHC detectors. One of these is MATHUSLA (MAssive Timing Hodoscope for Ultra-Stable neutral pArticles) [14, 15, 19]. MATHUSLA is a dedicated large-volume displaced vertex detector for the HL-LHC, to be positioned $\mathcal{O}(100)$ meters away from the LHC interaction point (IP). The detector will be composed of horizontal layers of trackers over an empty, air-filled decay volume. Its goal is to discover ultra-long lived particles produced at the LHC by reconstructing their decays inside its volume as displaced vertices (DVs). In the long lifetime regime, MATHUSLA can achieve cross-section sensitivity to weak-scale LLPs up to three orders of magnitude beyond ATLAS, depending on the decay and production mode [14]. This gain in sensitivity is largely due to MATHUSLA’s much lower backgrounds for LLP signals than main detector experiments, made possible by copious shielding from the LHC collision and the imposition of stringent signal reconstruction requirements.

MATHUSLA is fundamentally a giant tracker without any energy or momentum measurement. This is well-suited to reconstructing DVs in a low-background environment, but it might not be obvious how to characterize any new particles that are discovered. Fortunately, MATHUSLA is not only an LLP-discovery machine, and its physics utility goes beyond initial observation of LLP signals. Despite its simplicity, MATHUSLA can use geometric information to extract the likely velocity and decay mode of detected LLPs [20]. This also allows the LHC bunch crossing that produced the LLP to be identified. Furthermore, MATHUSLA can act as a Level-1 trigger for CMS, ensuring information about the LLP production event is written to tape regardless of main detector trigger thresholds.

Together, these abilities suggest that it should be possible to combine information from both MATHUSLA and CMS in order to study the properties of detected LLPs and details of the underlying theory.

In this work we analyze the prospect of using correlated MATHUSLA and main detector information to diagnose the production mode of LLPs, and estimate underlying model parameters under a given production mode hypothesis. One initial obstruction to this effort is the vast space of possible models that give rise to LLPs. Fortunately, recent progress has been made to organize much of this space into a set of just a few ‘simplified models’ that describe particular experimental signatures [13]. Complicated scenarios exist that do not fit easily into any of these categories, but the simplified model framework provides a good starting point for analyses that cover a wide range of commonly studied BSM scenarios in a relatively model-independent fashion.

Assuming that MATHUSLA observes a number of LLP decays that all originate from a single production mode, we show that with simple cuts using only a few event observables, observation of $\sim \mathcal{O}(100)$ events allows the production mode to be determined with $\gtrsim 90\%$ confidence. Even for $\mathcal{O}(10)$ observed events, some simplified models can be correctly identified with high probability, while higher event yield obviously results in even greater classification confidence. Assuming correct model classification, the LLP and parent particle masses can be determined to $\sim 10\%$ precision with similarly modest statistics.

Our simple analysis can be generalized to include more production modes, and more complicated scenarios with several LLP production and decay modes considered simultaneously. Our methods should also be applicable to other proposed external LLP detectors [16, 17], as well as to LLPs discovered in the LHC main detectors themselves (where additional information on the LLP decay products may be available). Together with information on the LLP decay mode from the geometric properties of the DVs [20], a discovery at MATHUSLA would therefore not only prove the existence of BSM physics, it would allow the theoretical origin of LLPs to be uncovered.

The paper is structured as follows. First we will review the MATHUSLA detector and its potential to be used as a trigger for CMS in section 2. In section 3 we will summarize the simulation of LLP production under several simplified models and across a wide range of parameters. In section 4 we will develop a robust, physically motivated classifier to diagnose the LLP production mode using both CMS and MATHUSLA observables, and evaluate its performance. We describe how simple analyses can extract the LLP and parent particle masses for the considered LLP production modes in section 5, and perform maximum likelihood estimation using pseudo-data compared to simulation-derived distributions to estimate the achievable precision. Section 6 contains a brief discussion of complications due to multiple possible LLP production vertices, and we make concluding remarks in section 7. Appendix A contains various Monte Carlo details for LLP event generation in our analysis.

2 The MATHUSLA detector

We first review the principles and geometry of the MATHUSLA detector in section 2.1 to provide context for the rest of the analysis. Section 2.2 contains a brief overview of the

physics reach of MATHUSLA. In section 2.3 we review how MATHUSLA can measure LLP boost, and the implications for coordinating between MATHUSLA and CMS.

2.1 Detector overview

The proposed MATHUSLA detector is a large-volume surface detector, to be placed near CMS [19, 21], see figure 1. The goal is to instrument a large volume to search for displaced vertex (DV) decays of ultra long-lived particles produced at the LHC, focusing on decay lengths $c\tau \gtrsim 100\text{m}$. It has been shown [19] that MATHUSLA can operate in a low-background environment by making use of strong restrictions on the final state multiplicity, timing, and geometry of candidate events. The detector principle is to place several (~ 5) layers of trackers, likely to be plastic scintillator, above a large decay volume of air. LLPs decaying to SM charged particles inside MATHUSLA will give rise to upwards going tracks. By requiring reconstructed tracks to travel upwards relativistically and originate both spatially and temporally at a point inside the empty detector volume, backgrounds such as LHC muons, atmospheric neutrinos and cosmic rays can be rejected.

MATHUSLA will be equipped with an internal trigger system for LLPs, which relies on real-time tagging of upwards-traveling track-candidates within groups of neighboring detector modules [19]. The trigger rate will be low enough that the MATHUSLA trigger can also act as a Level-1 burst-trigger for CMS: if the upwards tracks originate from the decay of an LLP, there is a range of < 10 candidate LHC bunch crossings that are very likely to include the production event at CMS. The MATHUSLA Level-1 trigger can therefore request a range of CMS events to be written to tape, which would enable the kind of off-line analysis we study in this work. If only CMS triggers were available, then the post-discovery analysis of many scenarios we consider here would be significantly more difficult or even impossible, due to much smaller trigger efficiencies in the main detector [14].

In this work, we assume the original MATHUSLA benchmark geometry [14, 15], which is a $200\text{m} \times 200\text{m}$ square detector with a 20m high decay volume, placed 100m displaced from the interaction point both horizontally and vertically. More recently, the collaboration has considered more realistic designs with a $100\text{m} \times 100\text{m}$ detector volume at an available site near CMS, with significantly smaller horizontal and vertical displacement compensating for the smaller size. This results in practically identical sensitivities to LLP decays [22]. Our results apply almost verbatim to this more realistic design as well.

2.2 Physics reach

Due to MATHUSLA's near-zero backgrounds, ~ 10 observed events are likely enough for a discovery in most of the targeted LLP scenarios, especially if they decay to hadrons. This low threshold allows MATHUSLA to probe much lower cross-sections than ATLAS and CMS in the long-lifetime regime. The number of LLPs MATHUSLA will observe depends primarily on the production cross-section and decay length of the LLPs, and has only mild model dependence. In the long-lifetime regime, the lowest cross-section that MATHUSLA can exclude scales linearly with $\bar{b}c\tau$, where \bar{b} is the mean boost of the LLPs. The minimum




Figure 1. Schematic illustration of MATHUSLA, using the original 200 m benchmark geometry [14, 15]. Our results will apply to the more realistic geometries currently considered by the MATHUSLA collaboration, with $\sim (100 \text{ m})^2$ area situated closer to the collision point, as well. The tracking planes in the roof detect charged particles, allowing for the reconstruction of displaced vertices in the air-filled decay volume. The floor protector provides vetoing capability against charged particles entering the detector from below.

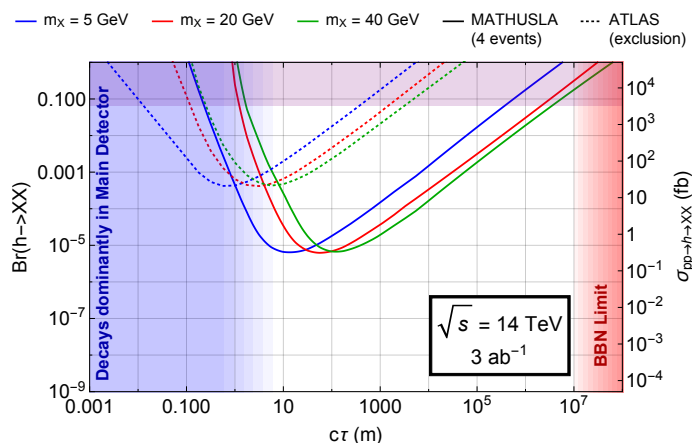


Figure 2. MATHUSLA sensitivity to new LLPs X pair-produced in exotic Higgs decays, where X decays hadronically. Exclusion curves correspond to 4 events in MATHUSLA. Plot from [14].

signal cross-section required for discovery is roughly [14]

$$\sigma_{\text{sig, min}}^{\text{LHC}} \sim (1 \text{ fb}) \left(\frac{\bar{b}c\tau}{10^3 \text{ m}} \right). \quad (2.1)$$

To give a concrete example of the detector’s physics reach, figure 2 shows that MATHUSLA can probe exotic Higgs decay branching ratios to LLPs down to 10^{-5} . For this study we will work in terms of the number of observed LLP decays, N_{obs} , ranging from 10 to 1000 without specifying the corresponding cross-sections for each model under consideration. Clearly this is a plausible event yield for many BSM scenarios.

2.3 Measuring LLP boost

There are two important reasons to measure LLP boost at MATHUSLA [20]. First, the boost is highly correlated with LLP mass for a given production mode. Second, knowledge

of LLP boost allows for the LHC production event to be identified, or at least narrowed down to a few possible bunch crossings, which allows us to use information from both MATHUSLA and CMS to learn about the nature of newly detected LLPs.

Ref. [20] showed that MATHUSLA can determine the velocity of detected LLPs using only geometric information of a displaced vertex and the knowledge of LLP's direction of travel, from the LHC IP to the reconstructed DV. For a two-body decay, the velocity of the LLP decaying to two particles at angles θ_1, θ_2 is

$$\beta_X = \frac{\beta_1 \beta_2 \sin(\theta_1 + \theta_2)}{\beta_1 \sin(\theta_1) + \beta_2 \sin(\theta_2)} \tag{2.2}$$

Since in the vast majority of cases the LLP decay products are ultra-relativistic, β_1 and β_2 can be set to 1 while incurring only negligible error. Regardless, the timing resolution of the proposed tracking detectors should allow these velocities to be determined within $\sim 5\%$. The uncertainty in θ_1 and θ_2 is dominated by detector effects and is estimated at roughly 0.2% for cm spatial resolution.

In the case of an LLP decay mode whose final states have high multiplicity, i.e. hadronic decays, there is an alternative method to estimate its velocity, relying on the approximate sphericity of charged decay products in the LLP rest frame. Labelling the final-state momenta p_i , we solve for β_X in the constraint

$$\hat{p}_X \cdot \sum_i \hat{p}_i(\beta_X) = 0 \tag{2.3}$$

Essentially, the measured track momenta are boosted backward along the direction toward the LHC IP until they are approximately spherically distributed, which approximates the LLP rest frame. The accuracy of this estimate is negatively affected by the fact that the decay product momenta may not be lightlike as well as the fact that some decay products might be downgoing in the lab frame and therefore not included in the calculation. Thus the measured LLP boost distribution suffers some systematic bias (which can be accounted for) and a small increase in spread. Nonetheless, the LLP boost can be determined sufficiently well event-by-event to narrow down the candidate bunch crossings in which it was produced to ~ 5 . There may be multiple hard-scattering vertices in the list of candidate events, but as discussed in 6, this does not jeopardize our analysis strategy.

3 LLP simplified models

The space of all physical models that could produce long-lived particles is immense. To identify, consider, and discard them one by one would be enormously impractical. Simplified models break down BSM theories into smaller effective models that reflect experimental signatures and can be characterized by a small number of parameters that correspond to experimental observables, like cross sections, branching ratios, and masses. In the past decade, the use of simplified models in searches for prompt particles at the LHC has become a standard practice [23]. Recently, considerable work has been done to extend the simplified model framework to LLPs [13]. This includes, in particular, separate simplified

models for LLP production and decay, since they do not a priori have to be related. This is very suitable for our analysis, since the decay mode can be reasonably well characterized by MATHUSLA alone [20], and we focus on the production mode. In this section, we review and slightly expand upon the simplified model basis for LLP production modes introduced in [13], and briefly explain how we generate Monte Carlo event samples for our study.

3.1 LLP production modes

In this paper we will employ the LLP simplified models recently defined and presented in the LHC LLP white paper [13], with the addition of LLP production in exotic B -meson decays. This set of simplified models, each representing an LLP production event topology, is meant to cover a wide range of well-motivated LLP production scenarios while remaining agnostic of the underlying mechanism generating the particles' long lifetimes. They are as follows, with schematic Feynman diagrams for each model in figure 3.

1. The Exotic B -meson Decay channel is abbreviated as **BB** and has a the SM B -meson decay to an LLP plus one or more SM products. This scenario is not included in the simplified model library constructed in [23], but we include this mode due to its importance for low-mass LLP benchmark models [24] below ~ 5 GeV. Two distinct scenarios for this production channel are right-handed neutrino models [25, 26] and light dark scalars [27]. Both are explored in this paper.
2. In the Charged Current (**CC**) scenario, a BSM particle possessing a SM electric charge is produced in the s -channel and decays to a charged SM particle and an LLP. In this case, assuming the dark sector particle is of integer charge, the charged SM decay product can be assumed to be a charged lepton. Models containing W' bosons [28–30] decaying to long-lived right-handed neutrinos [31, 32] are especially representative of this topology.
3. In the Heavy Parent (**HP**) scenario some BSM particle is pair produced before each decays to an LLP and one or more standard model particles. The visible SM particles in this case could be quarks or leptons, leading to a signature in CMS of hard leptons or jets accompanying the observation of an LLP in MATHUSLA.¹ Supersymmetry [7] is full of examples that fall into this topology, since a dominant signature is superpartner pair production with subsequent decays to the lightest detector-stable supersymmetric state. HP scenarios with hard leptons or photons in the final state are relatively conspicuous and easy to diagnose, so we focus on jetty final states. Squark (gluino) pair production and subsequent decay (via an off-shell squark) to a long-lived neutralino would produce one (two) jet(s) per decay chain. We consider both possibilities separately, and comment on how our methods could be easily extended to include leptonic or multi-step decay chains.

¹It is also possible, though less commonly considered, that the SM objects produced in association with the LLP are neutrinos. This challenging case of LLPs produced in association with exclusively invisible particles deserves further study.




Figure 3. Schematic Feynman diagrams for the models we consider. (a): exotic B -meson decay (**BB**), with heavy neutral lepton and scalar LLPs. (b): charged Current (**CC**). (c): heavy Parent (**HP**). (d): exotic Higgs Decay (**HIG**) with gluon fusion and vector boson fusion production channels. (e): direct Pair Production (**DPP**). (f): heavy Resonance (**RES**).

4. In the Exotic Higgs Decay (**HIG**) scenario, the Standard Model Higgs boson couples to the dark sector and decays to two long-lived particles. Hidden sector models with a Higgs portal fall into this category [24, 33]. More complex Higgs decay final states are of course a possibility, especially with decays into confining dark sectors [34–37]. The two-particle final state serves as a starting point when considering this space of models.

5. In the Direct Pair Production (**DPP**) scenario, two long-lived particles are produced non-resonantly, through some effective operator, such as that generated by a very heavy squark leading to neutralino pair production, or pair-production of a hidden sector LLP via an off-shell s -channel mediator.
6. The Heavy Resonance (**RES**) model has some BSM particle, such as a Z' boson (see e.g. [38–40]), produced on-shell in the s -channel that decays to two long-lived particles. In the high-mass limit, the Heavy Resonance model and Direct Pair Production model should coincide.

We will use these simplified models as our ‘basis set’ of LLP production topologies. Many searches at the LHC constrain the available parameter space for these models. However, in general, production cross-sections and branching ratios can be dialed down while still ensuring events in MATHUSLA, so we do not concern ourselves with current bounds on these signatures, except to note that they are not excluded for weak-scale masses.

3.2 Event generation

Event generation was performed with MadGraph5 version 2.6.6 [41], parton showering and hadronization with Pythia8 [42, 43], and CMS detector simulation with DELPHES 3.4.2 [44]. For each of the simplified models, a corresponding simulation model was chosen from the LLP Simplified Model Library of [13] or existing standard MadGraph models [7, 45–52], and used to generate events across a range of mass-parameters relevant for HL-LHC production.

All of these simplified models have either one or two parameters that are varied in simulation. The only parameter that can be freely chosen in the BB, DPP, and HIG models is the LLP mass m_{LLP} . The CC, HP, and RES models also have a parent particle that decays to the LLP, so the two free parameters are m_{parent} and m_{LLP} . The width of BSM parent particles can in principle be varied as well. For simplicity, Γ_{parent} is set to $0.01 \times m_{\text{parent}}$ for each of these models, since the exact width does not matter as long as it is small enough. A high-width parent particle case is also simulated for the RES model, with $\Gamma_{\text{parent}} = 0.3 \times m_{\text{parent}}$, to illustrate the effect of parent particle width on classification accuracy.

The range of LLP and parent masses for which simulation was generated in each model is summarized in table 1. In single-parameter models, the chosen m_{LLP} values are equally spaced within the indicated ranges (on a linear scale). In models with a variable parent mass, m_{parent} and $m_{\text{LLP}}/m_{\text{parent}}$ values are equally spaced within the indicated ranges. This is relevant in interpreting our probabilistic results regarding the achievable accuracy in classifying datasets of observed LLP decays. Further details on the simulation of each simplified model are in appendix A.

3.3 LLP decay in MATHUSLA

For each point in parameter space, 10^6 events were initially simulated. Only events with LLPs in the right rapidity range $\eta \in (0.8, 2)$ to potentially intersect MATHUSLA were considered. To increase the efficiency of our simulations for LLP decays in MATHUSLA

Model	MadGraph model	m_{parent} range [GeV]	m_{LLP} range
B decay	SM	~ 5.3 GeV	0.1 – 4 GeV
Charged Current	ChargeWN [13]	50 – 1850 GeV	$0.05 - 0.9 \times m_{\text{parent}}$
Heavy Parent	MSSM [7, 45–50]	50 – 1850 GeV	$0.05 - 0.9 \times m_{\text{parent}}$
Exotic Higgs Decay	heft [51, 52]	125 GeV	20 – 55 GeV
Direct Pair Production	MSSM	—	50 – 2000 GeV
Heavy Resonance	Zp_{2LLP} [13]	50 – 1850 GeV	$0.05 - 0.4 \times m_{\text{parent}}$

Table 1. Summary of signal simulation parameters. The values of m_{LLP} or $(m_{\text{parent}}, m_{\text{LLP}}/m_{\text{parent}})$ used for simulation samples are equally spaced within the above indicated ranges on a linear scale.

we took advantage of the fact that the orientation of each event around the beam axis is random. Therefore, we rotated and appropriately reweighed events to assign LLPs a random azimuthal angle that guaranteed them flying through some part of MATHUSLA. Each event was also weighted to account for the probability of the LLP decaying inside MATHUSLA by a factor $L/bc\tau$, where L is the distance each particular LLP travels through MATHUSLA, b is its boost, and $c\tau$ is its decay length. In the long lifetime regime, this is proportional to the probability that the LLP decays within MATHUSLA. This event selection and weighting defines our high-statistics weighted samples. Since our analysis assumes some number of observed LLP decays N_{obs} , there is no weighting applied for a particular cross-section, and our results are valid for arbitrary $\bar{b}c\tau \gg 10^2$ m, where \bar{b} is average LLP boost. (Repeating our analysis for shorter lifetimes would be straightforward, except it would allow for additional information on the LLP lifetime to be extracted from the geometric distribution of decays in the detector.) For each of these sets of weighted events, an unweighted event sample of size $\sim \mathcal{O}(10^4)$ was drawn. These are used to construct representative signal samples of $N_{\text{obs}} = 10, 100, 1000$ observed events.

We make two assumptions that simplify our analysis and simulations. The first is that however the LLP decays, MATHUSLA will accurately measure its boost using the methods outlined in section 2.3. This allows us to omit detailed MATHUSLA detector simulation of the LLP reconstruction, and makes our analysis independent of details of the final detector design. The second assumption is that each LLP decay can be uniquely matched to a hard scattering and hence a single primary vertex in an event at CMS. This is to simplify our first pilot study of analyzing LLP events in MATHUSLA and CMS, and a more realistic analysis can be generalized to account for the possibility that each LLP decay is associated with several possible hard-scattering candidate events due to high pile-up at the HL-LHC. As we discuss in section 6, this complication should not qualitatively impact our conclusions on production mode classification and parameter estimation.

4 Determining production mode

In this section we will describe the strategy for identifying the LLP production mechanism at the LHC using combined data from CMS and MATHUSLA.

Model	Variable	Cut for model classification
B decay (BB)	Fraction of events with no jets of $p_T > 20$ GeV	>0.9
Charged Current (CC)	Fraction of events with one lepton of $p_T > 10$ GeV	>0.2
Heavy Parent (LLP +1 jet)	Fraction of events with at least one jet of $p_T > 20$ GeV	>0.85
Heavy Parent (LLP +2 jet)	Fraction of events with at least six jets with $p_T > 20$ GeV	>0.3
Exotic Higgs Decay	Fraction of events y with highest jet $p_T > 90$ GeV	$0.04 < y < 0.22$
	Fraction of events x with two jets with $\Delta\eta_{jj} > 2.5$	$y < 3.5x - 0.2$
Direct Pair Production	Standard deviation σ and mean μ of LLP $\log_{10}(b)$ distribution	$\sigma > 2\mu + .22$
Heavy Resonance	—	—

Table 2. Cuts defining the LLP production model classification algorithm, applied in sequence from top to bottom. An observed event sample that fails all cuts is classified as Heavy Resonance production mode by a process of elimination.

We assume that MATHUSLA observes $N_{\text{obs}} = 10, 100$ or 1000 LLP decays, all resulting from the same single production topology. *Sample-level variables* describing characteristics of the entire observed LLP dataset, like fraction of events with some number of jets above some p_T in CMS, are used to classify the production mode. The algorithm is summarized in figure 4 and table 2.

Using characteristic features of each production mode, we find that simple cuts in sample-level observables can be used to achieve $\gtrsim 90\%$ probability of correct model classification for all but small corners of BSM particle parameter space with $\mathcal{O}(100)$ observed events, and $\gtrsim 98\%$ with $\mathcal{O}(1000)$ observed events. For the BB, CC, and HP models, $> 90\%$ probabilities of correct classification can be achieved with only $N_{\text{obs}} = 10$ events.

4.1 Hierarchical classification algorithm

Our goal is to find robust, physically motivated observables that allow us to classify samples of LLP production events into the categories defined by the different simplified models in figure 3. We therefore construct a hierarchical classification algorithm that addresses each production mode hypothesis in order of increasing difficulty. It is summarized in figure 4 and table 2. Below, each step is discussed in detail.

The two easiest tasks are to identify data arising from the Exotic B -meson Decay and Charged Current models. We start with BB. Because the vast majority of B -mesons are produced in relatively soft $b\bar{b}$ events at the LHC with $p_T < 20$ GeV, we can expect very few jets and even fewer b-tags in events from this model. Only $\approx 0.5\%$ of events have a b-tagged jet with $p_T > 20$ GeV in the Delphes simulation. The presence of a b-tag could thus only be

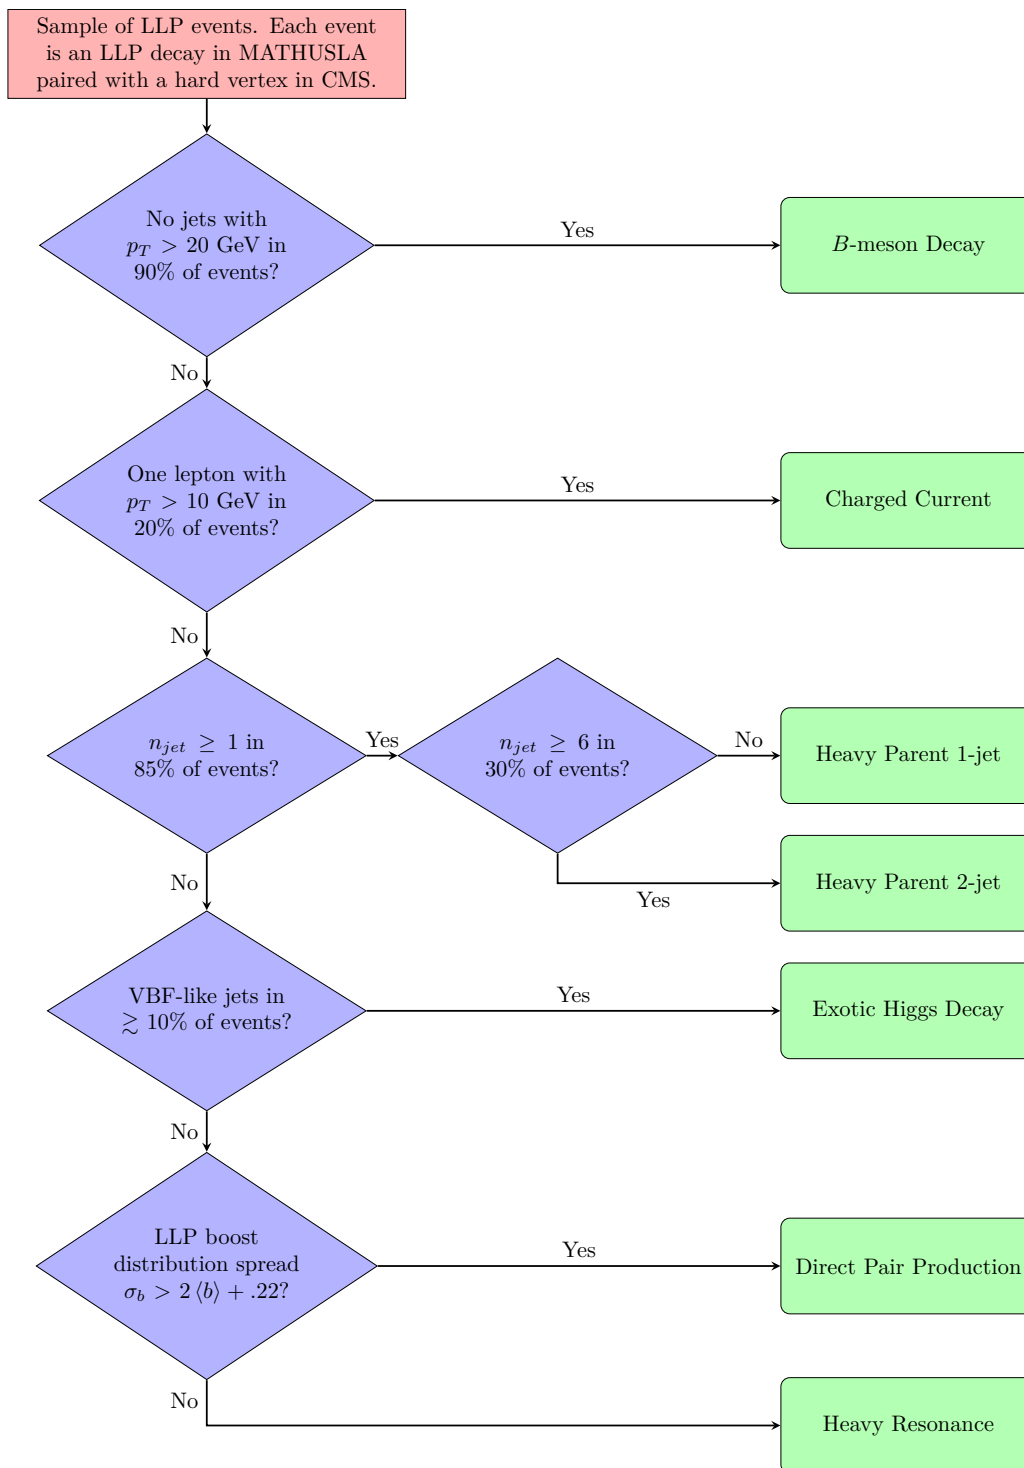


Figure 4. Summary of our hierarchical LLP production mode classification algorithm.

useful when $\mathcal{O}(1000)$ LLPs have been detected. However, there is an upside to the low characteristic energy scale of these events - a complete absence of hard jets in $> 90\%$ of events. All of the other models under consideration have a hard process at some high scale, producing ISR jets at the least in a sizeable fraction of events. The cut used to classify a sample as BB-like is essentially a jet veto, demanding zero jets above a certain p_T^{\min} cutoff in a high fraction of events. If the event sample fails this jet veto, we move on to the next possibility.

Next, we check whether the sample could arise from the Charged Current model. In this model, the LLP produced is accompanied by a SM charged lepton. For most values of m_{LLP} and m_{parent} , this lepton is generically produced at high transverse momentum, and in the opposite azimuthal direction from MATHUSLA for events where the LLP is detected at MATHUSLA. None of the other simplified models have a mechanism to produce a *single* hard SM lepton in such a high fraction of events. The cut for this model will demand an isolated lepton above some minimum p_T^{\min} in a high fraction of events.

The next step in the classification is to check for the Heavy Parent model. In the other models featuring LLPs produced in the decay of a heavy particle, the only jets in an LLP production event come from initial state radiation (with the exception of VBF jets in Higgs production). The spectrum of this radiation is dependent on the characteristic energy scale (and therefore BSM particle masses) of each model, but in general the number of jets per event and their p_T is lower than for the HP model, where we assume the SM products of the heavy parent decay give rise to jets. To be classified as HP, some minimum fraction of events in a sample will have to contain one or more jets above some p_T^{\min} . Within the HP model, we can further distinguish between the decay of the heavy parent particle to one SM particle plus an LLP, and the decay to two SM particles plus an LLP. This secondary classification asks for an even higher jet multiplicity in some minimum fraction of events, and may take advantage of the fact that 4 hard quarks or gluons at parton-level may be quite likely to split into even more jets after showering.

After these easier classification decisions, we must turn to more difficult cases requiring slightly more sophisticated checks. The Heavy Resonance and Direct Pair Production modes produce no associated SM particles in the hard process that creates the LLP, and the same is true in the dominant gluon fusion channel for Higgs production. However, in the vector boson fusion (VBF) and associated vector boson (VH) production channels for Higgs production, there are additional particles created in the hard process. In particular, we will look for the signature of VBF events to distinguish the HIG model. These events are characterized by a pair of hard jets well-separated in pseudorapidity ($\Delta\eta_{jj}$), originating from the two quarks surviving from the VBF process. Since VBF accounts for approximately 10% of the Higgs production cross-section, we expect to see roughly the same fraction of Higgs events with this distinctive signature. The DPP and RES models are much less likely to produce events of this nature. Apart from this, the jets present in HIG events have a weaker p_T spectrum than for RES or DPP in most regions of parameter space, due to the heavy resonance having mass $> m_H$ or the produced LLPs having $2m_{\text{LLP}} > m_H$, controlling the scale of the interaction and therefore the momentum of produced ISR jets. With both of these considerations in mind, we can construct a combined cut in the fraction of events with highest jet p_T larger than some cutoff, along with frac-

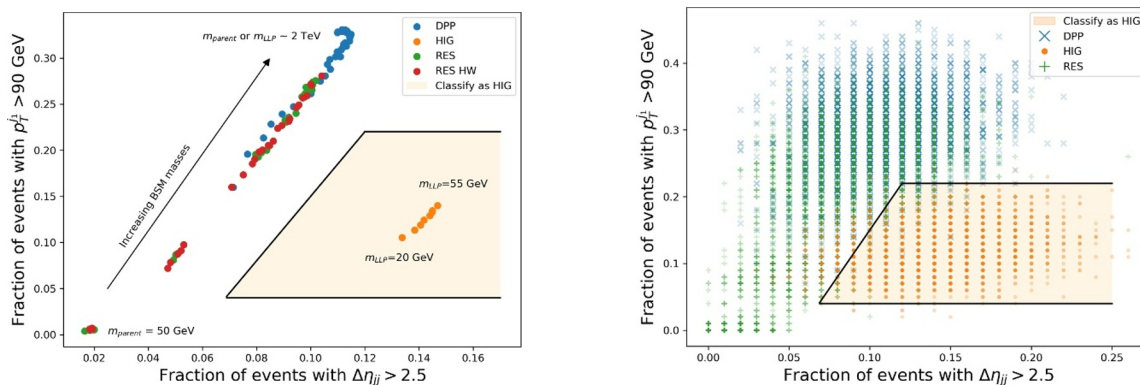


Figure 5. For Higgs Decay, Heavy Resonance and Direct Pair Production simulation samples, for the range of masses in table 1, the left plot shows for each sample the fraction of events with $\Delta n_{jj} > 2.5$ vs the fraction of events with hardest jet p_T above 90 GeV. Clearly, there is very good separation between HIG and RES/DPP in the high-statistics limit. The right plot shows the distribution of these two variables for samples of size $N_{\text{obs}} = 100$ for the various production modes and masses. There is now significantly more spread, but HIG events are still well separated. The orange shaded region defines our cuts the sample must pass to be classified as Higgs production.

tion of events with Δn_{jj} greater than some cutoff. See figure 5 for the shape of this cut, along with the high-statistics prediction of these two variables for the HIG, RES and DPP samples with various masses, and the spread for observed samples of size $N_{\text{obs}} = 100$.

Finally, we must devise a strategy to distinguish between the two most similar models, Heavy Resonance and Direct Pair Production. Indeed, in the high m_{parent} limit, the RES and DPP models coincide. For the first time, we will use the information gathered by MATHUSLA for production mode classification, not just the fact that it triggers CMS. In the RES model, the boost of the detected LLPs is determined almost exclusively by $m_{\text{LLP}}/m_{\text{parent}}$, since the heavy resonant particles will be produced dominantly on-threshold, giving the LLPs a boost $b \sim m_{\text{parent}}/2m_{\text{LLP}}$. This produces a fairly narrow boost distribution for the LLPs if the heavy resonance has a narrow width. Contrastingly, in the DPP model the LLP boost is correlated directly with m_{LLP} , with higher masses corresponding to lower boosts. Because there is no resonance decaying to the LLPs in the DPP model, the energy distribution of the LLPs is wider, leading to a wider boost distribution compared to even a high-width heavy resonance.

The difference in correlations of the LLP boost with particle masses between the two models will allow us to define a cut that separates the two models. Figure 6 shows the LLP boost mean and standard deviation of the low- and high-width Heavy Resonance samples for the various simulated masses, compared to the Direct Pair Production samples, in both the high-statistics limit and for observed samples of size $N_{\text{obs}} = 100$. A diagonal cut in this plane can separate DPP and RES. This cut only loses some of its distinguishing power in the high-width limit for the RES model.²

²For the initial classification we will make no effort to distinguish between the low- and high-width limits of the RES model. Since the width of the heavy resonant particle is a continuous parameter, the task of determining the width is better framed as part of parameter estimation. Section 5.6 sketches how this could be done.

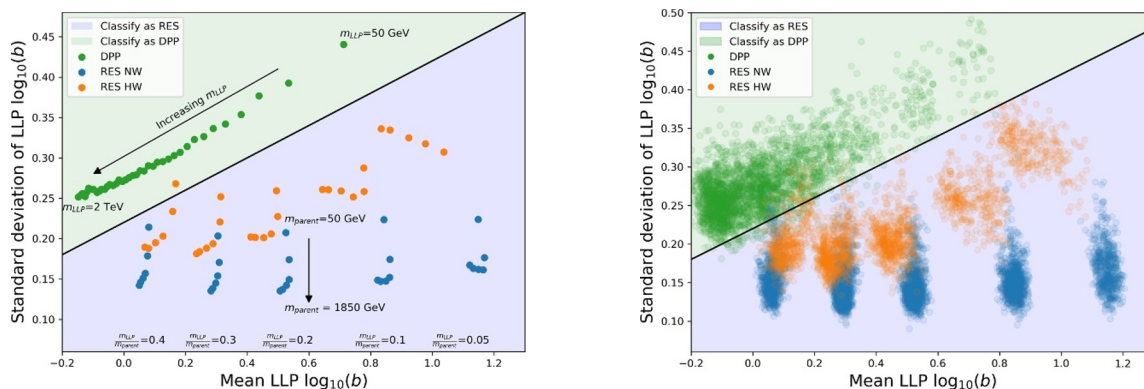


Figure 6. LLP boost mean and standard deviation of the low- and high-width Heavy Resonance samples for the various simulated masses, compared to the Direct Pair Production samples. Left: high-statistics limit. Right: for observed samples of size $N_{\text{obs}} = 100$. The indicated diagonal cut in this plane can separate RES from DPP samples. We also indicate the parent and LLP masses of each high statistics sample. Gaps in the RES distribution are due to the coarse grid of simulated masses.

For all of the above cuts we optimize the precise thresholds to maximize classification performance for event samples of size $N_{\text{obs}} = 100$. The resulting cuts are summarized in table 2, and visualized as a flowchart in figure 4. Note that these cuts are applied in order, such that only samples failing to meet the criteria of a given step are tested in the next. While more sophisticated simulations or data-driven optimizations might result in slightly different optimal values of these cuts, their physical basis nevertheless makes them robust.

Our classification framework is easily extendable to include other LLP scenarios. For example, leptonic decay in the Heavy Parent model can be accommodated by testing for a minimum fraction of events with two or more hard leptons. A modification of the two-jet Heavy Parent scenario where the heavy parent undergoes a decay chain rather than three-body decay could be distinguished using jet kinematic variables like M_{T2} [53]. Finally, a generalized Higgs decay scenario with indeterminate mass for the Higgs-like scalar could still be identified by the presence of VBF production events,³ with the task of distinguishing the scalar resonance from the SM Higgs becoming a problem of parameter estimation similar to the cases we discuss in section 5. We have omitted these scenarios from our first pilot analysis for simplicity, or because they are easier to classify than the cases we study (in particular the leptonic heavy-parent models). Even so, we can anticipate how our classification algorithm would be modified to include these possibilities in figure 7. We hope this will be useful for future extensions of our analysis.

It is important to note that nothing about our methods requires the LLP to be detected in an external detector like MATHUSLA. If LLPs are discovered at the LHC main detectors, the same kind of analysis can be used to characterize them as well. In fact, more information

³A singlet scalar without electroweak gauge charge would be difficult to distinguish from the Z' -like Heavy Resonance model if only LLP production events are considered, but if such a state existed, the pattern of its visible decays in the main detector could distinguish gauge-ordered from Yukawa-ordered couplings and hence favor a vector or scalar interpretation.

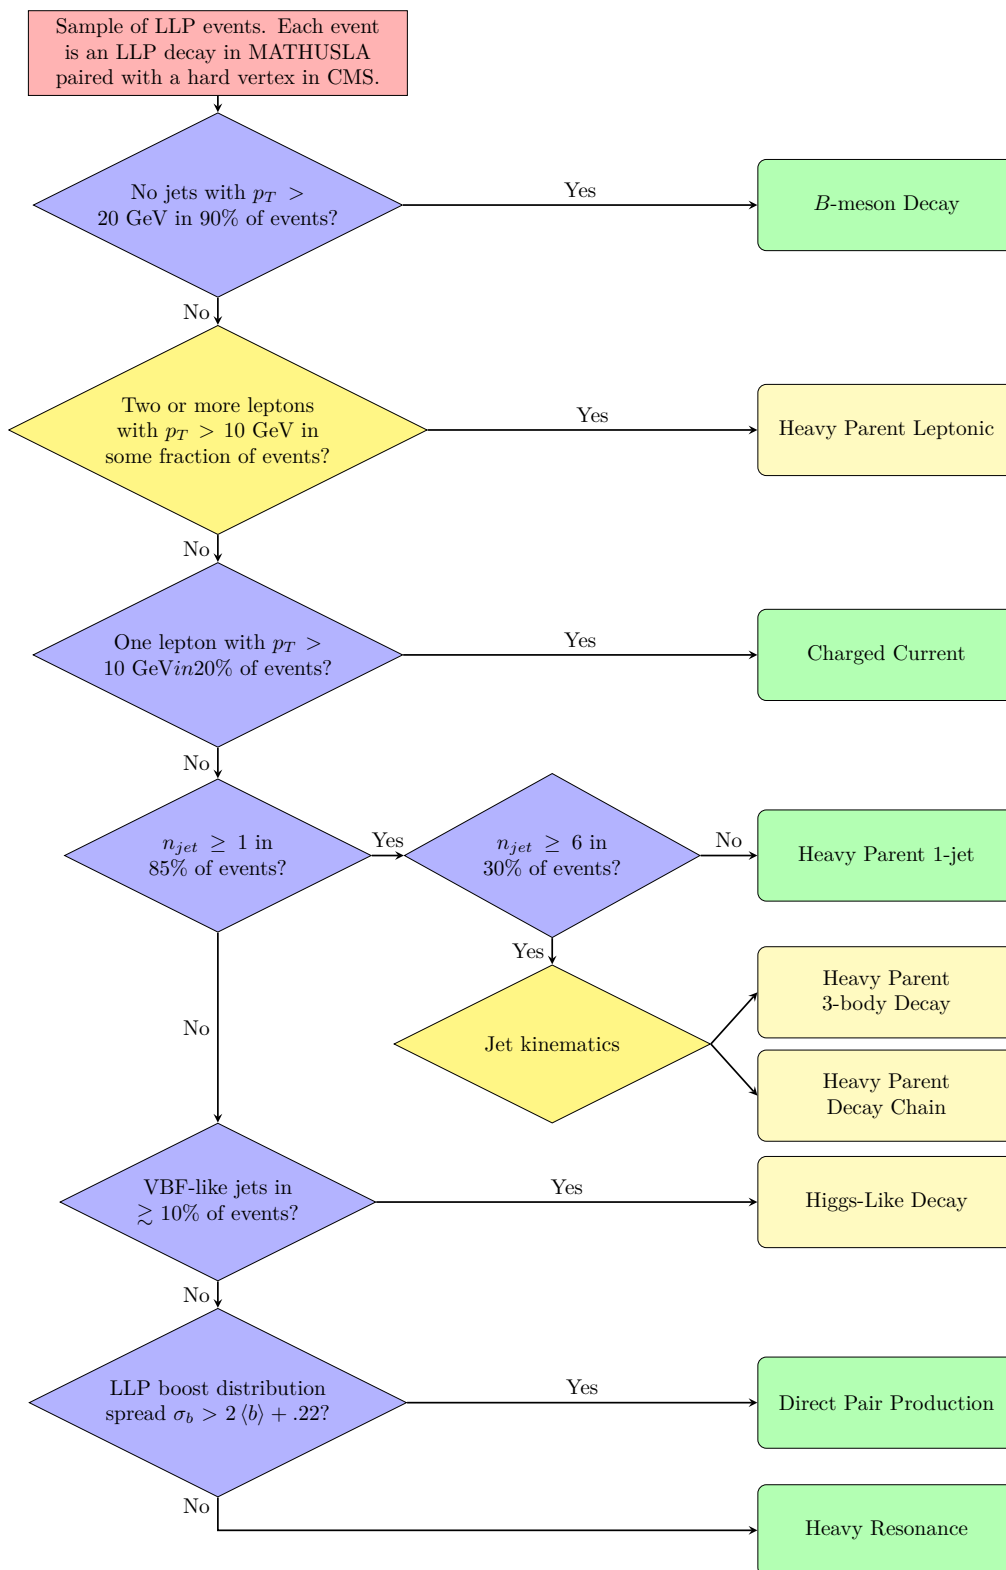


Figure 7. Flow chart summary of our hierarchical LLP production mode classification algorithm, showing possible additions to include other LLP production scenarios not considered in our detailed analysis.

N_{obs}	10	100	1000
B -decay (BB)	98	100	100
Charged Current (CC)	94	98	98
Heavy Parent (HP)	91	92	93
Higgs Decay (HIG)	36	91	100
Direct Pair Production (DPP)	48	98	100
Resonance, narrow-width (RES NW)	71	93	98
Resonance, high-width (RES HW)	63	86	94

Table 3. Classification accuracies (%) for each model at varying N_{obs} , averaged over all simulated LLP and parent particle masses for each model. For more details, see table 4.

will be available, since the main detectors might be able to extract energy and momentum data on the LLP decay products, depending on the position of the DV within the detector.

4.2 Performance

Unweighted sets of pseudo-data with $N_{\text{obs}} = 10, 100, 1000$ were generated for the models and mass ranges shown in table 1. For most models, 30-40 parameter-points were generated, forming a rough grid in the $m_{\text{parent}}, m_{\text{LLP}}/m_{\text{parent}}$ plane. For the one-parameter models, between 8 (HIG) and 40 (DPP) values of m_{LLP} were generated.

The classification algorithm was run on all of these pseudo-data samples, and the rates of correct and incorrect classifications were recorded. At each of the three benchmark values N_{obs} , we report an average classification accuracy for each model. This average of classification accuracies is taken first over all samples at a given point in parameter space, then over all points in parameter space, for each model. This average is particular to the prior imposed by our choice of simulated BSM particle masses, but generally reflects the ability of the classifier to distinguish a given model from the others across parameter space. Corners of parameter space for which the classification accuracy differs significantly from the average are discussed on a case-by-case basis, as are trends of variation in the accuracy as a function of BSM particle masses and widths.

Table 3 summarizes the performance of the classifier for each model, while table 4 shows the full classification matrices for each benchmark N_{obs} .

1. All pseudo-experiments of BB data with $N_{\text{obs}} = 100$ and 1000 were correctly classified, and no pseudo-experiments were incorrectly classified as BB. There were 29 samples of 100 events per m_{LLP} point on average, and 3 samples of 1000 events on average. The limited statistics available is not concerning, since the classification accuracy reaches 100% already at $N_{\text{obs}} = 100$.
2. In the CC model, all pseudo-experiments with $N_{\text{obs}} = 100$ were correctly classified except for those with $m_{\text{parent}} = 50$ GeV and $m_{\text{LLP}} = 45$ GeV. No pseudo-experiments from other models were incorrectly classified as CC. That performance persisted for

$N_{\text{obs}} = 10$:

Result \ Truth	BB	CC	HP 1-jet	HP 2-jet	HIG	DPP	RES
BB (4570)	$98.4_{-0.4}^{+0.3}$	0	0	0	0	$1.4_{-0.3}^{+0.4}$	$0.2_{-0.1}^{+0.2}$
CC (20920)	$4.0_{-0.3}^{+0.3}$	$93.5_{-0.3}^{+0.3}$	$0.004_{-0.004}^{+0.017}$	0	$0.03_{-0.02}^{+0.03}$	$2.2_{-0.2}^{+0.2}$	$0.2_{-0.1}^{+0.1}$
HP 1-jet (26510)	$0.1_{-0.03}^{+0.04}$	0	$81.8_{-0.5}^{+0.5}$	$9.2_{-0.3}^{+0.4}$	$1.3_{-0.1}^{+0.1}$	$5.7_{-0.3}^{+0.3}$	$2.0_{-0.2}^{+0.2}$
HP 2-jet (12360)	0	0	$19.9_{-0.7}^{+0.7}$	$78.8_{-0.7}^{+0.7}$	$0.5_{-0.1}^{+0.1}$	$0.8_{-0.1}^{+0.2}$	$0.1_{-0.0}^{+0.1}$
HIG (28550)	$0.02_{-0.01}^{+0.02}$	$0.04_{-0.02}^{+0.03}$	$9.5_{-0.3}^{+0.3}$	0	$36.1_{-0.6}^{+0.6}$	$29.4_{-0.5}^{+0.5}$	$24.9_{-0.5}^{+0.5}$
DPP (29120)	$0.01_{-0.01}^{+0.02}$	0	$17.5_{-0.4}^{+0.4}$	0	$12.3_{-0.4}^{+0.4}$	$48.0_{-0.6}^{+0.6}$	$22.2_{-0.5}^{+0.5}$
RES NW (40160)	$4.7_{-0.2}^{+0.2}$	0	$7.6_{-0.3}^{+0.3}$	0	$14.0_{-0.3}^{+0.3}$	$3.0_{-0.2}^{+0.2}$	$70.7_{-0.4}^{+0.4}$
RES HW (23100)	$3.9_{-0.2}^{+0.3}$	0	$6.6_{-0.3}^{+0.3}$	0	$14.9_{-0.5}^{+0.5}$	$11.4_{-0.4}^{+0.4}$	$63.1_{-0.6}^{+0.6}$

 $N_{\text{obs}} = 100$:

Result \ Truth	BB	CC	HP 1-jet	HP 2-jet	HIG	DPP	RES
BB (457)	$100.0_{-0.5}^{+0.0}$	0	0	0	0	0	0
CC (2092)	0	$97.6_{-0.7}^{+0.6}$	0	0	0	$2.4_{-0.6}^{+0.7}$	0
HP 1-jet (2651)	0	0	$90.5_{-1.2}^{+1.1}$	$0.9_{-0.3}^{+0.4}$	$1.0_{-0.3}^{+0.4}$	$7.3_{-0.9}^{+1.0}$	$0.3_{-0.2}^{+0.3}$
HP 2-jet (1236)	0	0	$16.0_{-2.0}^{+2.1}$	$83.5_{-2.1}^{+2.0}$	$0.4_{-0.2}^{+0.5}$	$0.1_{-0.1}^{+0.3}$	0
HIG (2855)	0	0	0	0	$90.4_{-1.1}^{+1.0}$	$7.4_{-0.9}^{+1.0}$	$2.2_{-0.5}^{+0.6}$
DPP (2912)	0	0	$0.1_{-0.1}^{+0.2}$	0	$1.6_{-0.4}^{+0.5}$	$97.9_{-0.6}^{+0.5}$	$0.4_{-0.2}^{+0.3}$
RES NW (4016)	0	0	0	0	$4.8_{-0.6}^{+0.7}$	$1.9_{-0.4}^{+0.5}$	$93.3_{-0.8}^{+0.7}$
RES HW (2310)	0	0	0	0	$6.3_{-0.9}^{+1.0}$	$8.0_{-1.1}^{+1.2}$	$85.7_{-1.5}^{+1.4}$

 $N_{\text{obs}} = 1000$:

Result \ Truth	BB	CC	HP 1-jet	HP 2-jet	HIG	DPP	RES
BB (40)	$100.0_{-6.0}^{+0.0}$	0	0	0	0	0	0
CC (191)	0	$97.6_{-3.0}^{+1.5}$	0	0	0	$2.4_{-1.5}^{+3.0}$	0
HP 1-jet (246)	0	0	$92.7_{-3.8}^{+2.8}$	0	0	$7.3_{-2.8}^{+3.8}$	0
HP 2-jet (112)	0	0	$15.0_{-5.7}^{+7.5}$	$85.0_{-7.5}^{+5.7}$	0	0	0
HIG (282)	0	0	0	0	$100.0_{-0.9}^{+0.0}$	0	0
DPP (274)	0	0	0	0	0	$100.0_{-0.9}^{+0.0}$	0
RES NW (390)	0	0	0	0	0	$1.7_{-0.9}^{+1.7}$	$98.3_{-1.7}^{+0.9}$
RES HW (219)	0	0	0	0	$0.3_{-0.2}^{+1.5}$	$5.7_{-2.5}^{+3.7}$	$94.0_{-3.7}^{+2.6}$

Table 4. Breakdown of the Production Mode Classifier output, for pseudo-data samples with 10, 100 or 1000 events, averaged over all LLP and parent particle masses simulated for each model. Entries in row i , column j show the percentage of samples from model i classified as model j . 95% confidence intervals are shown for non-zero classification accuracies, accounting only for statistical uncertainty due to the limited number of samples tested. The number of samples tested for each model is listed in brackets in the first column.

$N_{\text{obs}} = 1000$. For $N_{\text{obs}} = 10$, performance degraded for $250 \text{ GeV} \leq m_{\text{parent}} < 1 \text{ TeV}$ down to $\approx 97 - 99\%$ classification accuracy, and for $m_{\text{parent}} = 50 \text{ GeV}$, classification accuracies varied between 0-70% for m_{LLP} ranging from 45 GeV down to 5 GeV. Still, no samples from other models were incorrectly classified as CC.

3. For $m_{\text{parent}} \geq 250 \text{ GeV}$ in the HP model, all pseudo-experiments with $N_{\text{obs}} \geq 100$ were classified as HP, whether in the one or two jet decay modes. There were $\mathcal{O}(50)$ pseudo-data samples tested per point in parameter space. For $m_{\text{parent}} = 50 \text{ GeV}$, the lowest simulated parent particle mass, classification accuracy ranged from 88% at $m_{\text{LLP}}/m_{\text{parent}} = 0.1$ down to 0% for $m_{\text{LLP}}/m_{\text{parent}} > 0.5$. With such low available energy for SM products in these events, this low classification accuracy reflects the physical similarity of the experimental signature of the HP model to the signatures of the other simplified models in the low-mass corner of parameter space. This pattern of performance persisted from $N_{\text{obs}} = 10$ to $N_{\text{obs}} = 1000$. Above $m_{\text{parent}} \approx 500 \text{ GeV}$, the one- and two-jet decay modes were completely distinguishable. At $m_{\text{parent}} = 200 \text{ GeV}$, a significant fraction of HP 2-jet samples were classified as HP 1-jet, ranging from 100% incorrectly classified at $m_{\text{LLP}} = 180 \text{ GeV}$ down to 18% at $m_{\text{LLP}} = 20 \text{ GeV}$.
4. For the HIG model, 91% of pseudo-experiments of $N_{\text{obs}} = 100$ events were correctly classified. This performance is independent of m_{LLP} .
5. For the DPP model, correct classification rates of pseudo-experiments with $N_{\text{obs}} = 100$ varied from 90% at $m_{\text{LLP}} = 50 \text{ GeV}$ up to approximately $\approx 98 - 100\%$ for $m_{\text{LLP}} \geq 700 \text{ GeV}$. The number of unweighted samples of 100 events available varied from 20 at low $m_{\text{LLP}} \sim 50 \text{ GeV}$ to 80 at high $m_{\text{LLP}} \gtrsim 500 \text{ GeV}$, with a total of 2275 samples tested. For both the HP and CC models with $m_{\text{parent}} \sim 50 \text{ GeV}$, $m_{\text{LLP}} \sim 45 \text{ GeV}$, most samples were incorrectly classified as DPP.
6. For the RES model with a narrow-width resonance, the average classification accuracy of 100 event pseudo-data samples across parameter space was 93%. Accuracies varied from $\approx 85 - 100\%$ across parameter space, with the exception of $m_{\text{parent}} = 50 \text{ GeV}$, $m_{\text{LLP}}/m_{\text{parent}} = 0.4$, where 52% of samples were mis-classified as DPP. Generally, classification accuracy rose with m_{parent} .
7. For the RES model with a high-width resonance ($\Gamma_{Z'} = 0.3 \times m_{Z'}$), classification accuracy was lower than for narrow-width, averaging 86% across the simulated parameter space. The variation of classification accuracy across parameter space was similar to the narrow-width case.

Clearly, our classification algorithm performs well in most regions of BSM parameter space. A small minority of scenarios is misclassified, but in each such case this reflects either a genuine physical degeneracy of observables, like in various low-mass limits, requiring much more detailed analysis, or simply a ‘special case’, where an extension of our simple classifier could greatly improve accuracy. Having demonstrated that classification is generally robust, we leave such further details to future studies.

5 Determining model parameters

The second task for which we would like to estimate the prospective capabilities of MATHUSLA and CMS is the measurement of the properties of the newly discovered BSM particles. In ref. [20] it was shown that the measured LLP boost distribution gives information about the LLP mass under the assumption of production by exotic Higgs decay. We will attempt to demonstrate similar relations for the rest of the simplified models under our consideration. Without the knowledge of an intermediate parent particle whose mass is itself known, this task becomes more difficult. Each subsection below will describe the estimators of BSM particle masses and precision achievable for each of the simplified models. One important assumption that we make here is that the production mode of the LLPs in the sample has already been accurately identified, so that we know which analysis to perform.⁴ The main conclusion of our analysis is that MATHUSLA can determine BSM particle masses with useful $\mathcal{O}(10\%)$ precision assuming only $\mathcal{O}(100)$ observed events.

For each production mode, we use binned maximum-likelihood estimation to compare distributions of event-level observables from pseudo-data with $N_{\text{obs}} = 10, 100$ or 1000 events against distributions generated using our full high-statistics simulations. Best-fit masses are computed for many pseudo-data samples, and the standard deviation of the distribution of best-fit masses is taken as an estimate of the precision MATHUSLA can achieve when measuring LLP model parameters. The log-likelihood (LLH) function to be minimized as a function of the theory parameters is

$$-\ln L(m_{\text{parent}}, m_{\text{LLP}}) = \sum_i n_i \ln \left(\frac{n_i}{\mu_i(m_{\text{parent}}, m_{\text{LLP}})} \right) \quad (5.1)$$

where i indexes bins in a histogram of the chosen observables, n_i is the observed count in bin i for a given sample of pseudo-data, and $\mu_i(m_{\text{parent}}, m_{\text{LLP}})$ is the expected count in bin i . Bins with 0 observed count are ignored. The likelihood function is evaluated over a grid of points in $(m_{\text{parent}}, m_{\text{LLP}}/m_{\text{parent}})$ with available simulation, then interpolated and minimized. For models without varying parent particle masses, simulation is generated for m_{LLP} at regular intervals, and the likelihood depends only on m_{LLP} .

The spacing and size of the grid were chosen to accurately demonstrate the precision achievable at $N_{\text{obs}} = 100$. This means that for $N_{\text{obs}} = 10$, some outlier pseudo-experiments do not have best-fit mass values within the grid. In such cases, the reported precision should be regarded as an approximate lower bound on the uncertainty in BSM particle masses MATHUSLA can achieve. For $N_{\text{obs}} = 1000$ the grid spacing is much larger in some cases than the variation in best-fit masses, meaning that the actual distribution of best-fit masses is not accurately portrayed. However, systematic uncertainties driven by the final experimental design are likely to dominate in this regime, so the exact values obtained for 1000 observed events are less important than the qualitative lesson that the BSM parameter measurement might become systematics dominated.

⁴This assumption could be loosened if we extended our methods to some kind of global template-fit over all BSM models, but this would have to contend with detailed mis-modeling effects which our simple cut-based classifier algorithm mostly avoids.

Model	x_1	x_2	N_{obs}	$\frac{m_{\text{LLP}}}{m_{\text{parent}}}$ or m_{LLP} precision	m_{parent} precision	
B decay	$\log_{10}(b_{\text{LLP}})$	-	10	0.3 – 0.7	—	
			100	0.1 – 0.2		
			1000	$\lesssim 0.05$		
Charged Current		p_T^ℓ		10	0.1	0.1
				100	0.05	0.02
				1000	0.01	0.01
Heavy Parent		H_T		10	0.2	0.2
				100	0.05	0.05
				1000	0.01	0.01
Exotic Higgs decay	-		10	0.15	—	
			100	0.05		
			1000	0.01		
Direct Pair Production	-		10	0.3 – 0.5	—	
			100	0.1 – 0.2		
			1000	.03 – .07		
Heavy Resonance (narrow)	n_{jet}		10	0.07	0.15	
			100	0.02		
			1000	0.01		
Heavy Resonance (wide)	n_{jet}		10	0.12	0.15	
			100	0.05		
			1000	0.02		

Table 5. Summary of parameter estimation performance for all of the simplified models we consider. The variables x_1, x_2 chosen to estimate BSM particle masses are listed. The precisions shown are the characteristic standard deviation/mean of best-fit masses for benchmark BSM particle masses that are approximately representative for each model.

For single-parameter models without a resonant LLP parent particle (Direct Pair Production), or where the parent mass is known (B -decay, Higgs decay), we only need a single variable $x_1(m_{\text{LLP}}) = b = |\vec{p}_{\text{LLP}}|/m_{\text{LLP}}$, the LLP boost that MATHUSLA can measure from the DV geometry. The binned distribution $\mu_i(m_{\text{LLP}})$ in this variable is sufficient to measure LLP mass, since $b \sim m_{\text{parent}}/m_{\text{LLP}}$.

For models with unknown parent particle mass (Charged Current, Heavy Parent and Heavy Resonance), we also make use of the LLP boost $x_1(m_{\text{LLP}}, m_{\text{parent}}) = b$. However, we need to define an additional variable $x_2(m_{\text{LLP}}, m_{\text{parent}})$ that can be constructed from

information at MATHUSLA and CMS, and which supplies information about the LLP or parent particle masses independent of the boost. We do this case-by-case below, and the variables are summarized in table 5, along with the approximate precision MATHUSLA can achieve for each model.

Due to the high computational cost of generating simulated datasets on the finely-spaced grid of LLP and parent particle masses necessary to accurately maximize the likelihood over parameter space, we demonstrate parameter estimation performance for one set of BSM particle masses in the two-parameter models. We choose $m_{\text{parent}} = 1$ TeV, $m_{\text{LLP}} = 300$ GeV for all models as a benchmark that is fairly representative of MATHUSLA’s abilities across a large range of parameter space.

5.1 Exotic B -meson decay

The first model we consider is Exotic B -meson Decay, in the two cases of a heavy neutral lepton (HNL) LLP or a scalar LLP. We assume these can be distinguished by examining the decay products of observed LLPs. In the HNL case, the LLP decay includes an invisible neutrino, so the visible decay products in MATHUSLA do not point back at the LHC interaction point. The details of how this classification could be accomplished would require careful further study as well as detailed knowledge on the final experimental design of MATHUSLA, which is not currently available. Furthermore, the boost measurement requires special treatment when the LLP mass is not much greater than that of its decay products. These issues are important and require their own dedicated investigation, which is beyond the scope of this work. For our purposes, we will assume that the LLP decay mode can be determined and its boost measured with useful accuracy.

With the mass of SM mesons known, the only free parameter in this model is the LLP mass. If $m_{\text{LLP}} \ll m_B$, the LLPs produced will be highly boosted, and their boost is strongly correlated with their mass. Figure 8 shows this relationship for both scalar and RHN LLPs, using the mean boost to demonstrate the dependence on m_{LLP} . It should be noted that as m_{LLP} increases, the boost rapidly loses its dependence on m_{LLP} .

The precision of the m_{LLP} estimation was evaluated at three benchmark LLP masses: 0.5, 1.0, 3.0 GeV, for both HNL and scalar LLPs. Maximum likelihood estimation was performed on many samples of $N_{\text{obs}} = 10, 100, 1000$, comparing each sample’s boost histogram against expected histograms for $\log_{10}(b_{\text{LLP}})$ at each mass point created from high-statistics simulated samples. Tables 6 and 7 show the standard deviation of best-fit masses for each mass point and pseudo-experiment sample size under the two different models. In cases where no local minimum in $-LLH$ was found within the range of masses with available simulation, typically occurring for high m_{LLP} , the best-fit LLP mass was set to 4 GeV. Consequently, the spread of best-fit masses is generally underestimated for the $m_{\text{LLP}} = 3$ GeV case.

Comparing the two kinds of LLP under consideration, we find that LLP mass can be estimated with lower standard error in the scalar case than the HNL case. At or below ≈ 1 GeV, the HNL LLP mass can be determined with relative precision $\approx 0.15 - 0.2$ with $N_{\text{obs}} = 100$, while the scalar LLP mass can be determined with relative precision ≈ 0.1 . At $m_{\text{LLP}} = 3$ GeV, the spread in best-fit masses is large enough for both types of LLP




Figure 8. Average LLP boost as a function of m_{LLP} for production in B -meson decays, assuming the LLP is either a HNL or a light neutral scalar.

True m_{LLP} [GeV]	0.5			1.0			3.0		
N_{obs}	10	100	1000	10	100	1000	10	100	1000
N_{samples}	1083	110	10	1574	155	15	2144	213	21
Relative precision of best-fit m_{LLP}	0.68	0.17	0.051	0.71	0.19	0.047	> 0.43	\approx 0.28	\approx 0.17

Table 6. Exotic B decay \rightarrow Heavy neutral lepton LLP + SM mass estimation results. N_{samples} is the number of pseudodata sets used to estimate the spread of best-fit values. Relative spread of best-fit masses is reported for three choices of m_{LLP} and $N_{\text{obs}} = 10, 100, 1000$.

True m_{LLP} [GeV]	0.5			1.0			3.0		
N_{obs}	10	100	1000	10	100	1000	10	100	1000
N_{samples}	504	50	5	946	95	9	1993	200	20
Relative precision of best-fit m_{LLP}	0.29	0.092	0.020	0.34	0.091	0.032	> 0.28	\approx 0.16	\approx 0.052

Table 7. Exotic B decay \rightarrow Scalar LLP + SM mass estimation results. N_{samples} is the number of pseudodata sets used to estimate the spread of best-fit values. Relative spread of best-fit masses is reported for three choices of m_{LLP} and $N_{\text{obs}} = 10, 100, 1000$.

that a large fraction of samples have no local $-LLH$ minimum within the range of masses simulated, indicating a large uncertainty in the LLP mass.

5.2 Charged current

Second, let us consider the Charged Current production model, where two new particles participate in the LLP production process. We need to define a variable that is complementary to the LLP boost $b = x_1$, in order to determine both parent and LLP masses.

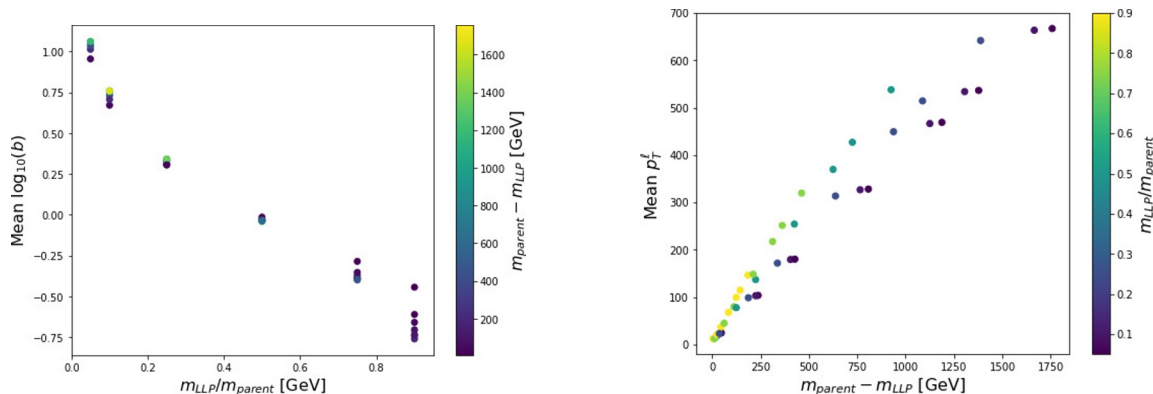


Figure 9. For the Charged Current (CC) production mode, we show on the left and right plots how the high-statistics distributions of our estimator variables $x_1 = \log_{10}(b)$ and $x_2 = p_T^\ell$ depend on $m_{\text{LLP}}/m_{\text{parent}}$ and m_{parent} in complementary ways. The color legend shows the subdominant variation with the complementary combination of LLP and parent mass.

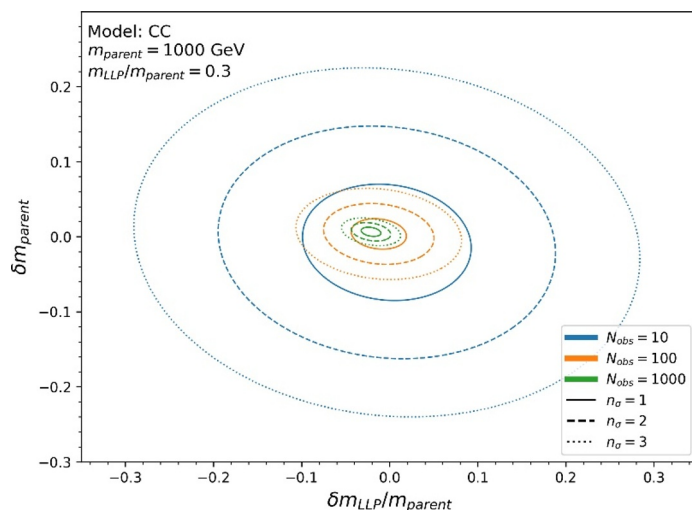


Figure 10. Parameter estimation for a benchmark point $(m_{\text{parent}}, m_{\text{LLP}}/m_{\text{parent}}) = (1000 \text{ GeV}, 0.3)$ in the Charged Current model, showing 1, 2, and 3 σ ellipses for $N_{\text{obs}} = 10, 100, 1000$ observed events. The axes are in units of $\delta x \equiv \frac{x_{\text{fit}} - x_{\text{true}}}{x_{\text{true}}}$, where x_{true} and x_{fit} are the actual values of the parameters and their measured value extracted from the fit.

In the CC model, the LLP is produced in a decay whose other product is a SM lepton ℓ , which is most likely observed and whose momentum is measured in CMS. Its transverse momentum p_T^ℓ is highly correlated with $m_{\text{parent}} - m_{\text{LLP}}$, which is an orthogonal combination of masses to m_{LLP} . We therefore set $x_2 = p_T^\ell$.⁵

The strong dependences of the high-statistics distribution average of $x_1 = \log_{10}(b)$ on $m_{\text{LLP}}/m_{\text{parent}}$, and the peak value of the $x_2 = p_T^\ell$ distribution on $m_{\text{parent}} - m_{\text{LLP}}$ are illustrated in figure 9.

⁵We also checked the performance of an estimator that uses the full lepton momentum measured at CMS and the LLP velocity measured at MATHUSLA to directly reconstruct the parent mass, but surprisingly we found that using just p_T^ℓ gave better precision, likely because it depends more directly on $m_{\text{parent}} - m_{\text{LLP}}$.

N_{obs}	10	100	1000
N_{samples}	838	84	8
Relative precision of $m_{\text{LLP}}/m_{\text{parent}}$	0.096	0.032	0.011
Relative precision of m_{parent}	0.078	0.020	5.7×10^{-3}
Pearson correlation coefficient	-0.09	-0.14	-0.33

Table 8. Charged Current model mass estimation results with $m_{\text{LLP}} = 300$ GeV, $m_{\text{parent}} = 1$ TeV. N_{samples} is the number of pseudo-data sets tested for the indicated number of events N_{obs} . The low correlation coefficient indicates that the uncertainties in $m_{\text{LLP}}/m_{\text{parent}}$ and m_{parent} separately are not highly correlated.

Figure 10 shows the precision attained for m_{parent} and $m_{\text{LLP}}/m_{\text{parent}}$ in the cases $N_{\text{obs}} = 10, 100, 1000$ for our benchmark parameter point of $(m_{\text{parent}}, m_{\text{LLP}}) = (1 \text{ TeV}, 300 \text{ GeV})$. Table 8 gives the relative precisions for parent and LLP masses, defined as the standard deviation divided by the mean of best-fit masses. The most important result is that under the CC production model, with 100 LLP observations in MATHUSLA successfully matched with their production events in CMS, the parent particle mass and the $m_{\text{LLP}}/m_{\text{parent}}$ mass ratio can both be determined with better than 5% resolution. Even with only 10 events, these resolutions degrade to $\approx 8\%$ and 10% respectively.

5.3 Heavy parent

Next, consider the Heavy Parent model. Again, one of our variables will be the LLP boost $\log_{10}(b) = x_1$, and we need to find a complementary variable x_2 . We select $x_2 = H_T$, the scalar sum of jet transverse momenta in CMS. Its distribution is highly correlated with $m_{\text{parent}} - m_{\text{LLP}}$, similarly to p_T^ℓ in the Charged Current case, and works well for both one- and two-jet parent decay scenarios.⁶ Figure 11 shows the dependence of the high-statistics average value of both estimators on the BSM particle masses.

Ellipses showing the spread of best-fit masses for pseudo-data samples of size $N_{\text{obs}} = 10, 100, 1000$ using H_T are displayed in figure 12 for the one- and two-jet decay modes of our benchmark parameter point $(m_{\text{parent}}, m_{\text{LLP}}) = (1 \text{ TeV}, 300 \text{ GeV})$. Tables 9 and 10 summarize the distributions of best-fit particle masses in the one- and two-jet Heavy Parent decay models respectively. For the one-jet case, 5% resolution in m_{parent} and $m_{\text{LLP}}/m_{\text{parent}}$ is possible with 100 observed LLP events in MATHUSLA, and 20% resolution with just 10 events. In the two-jet case, resolution of 5-10% in m_{parent} and $m_{\text{LLP}}/m_{\text{parent}}$ is possible with 100 observed LLP events in MATHUSLA, and 20-30% resolution with just 10 events.

Since the H_T variable can be used to estimate masses in the one-jet and two-jet Heavy Parent decay scenarios, it is interesting to consider whether the parameter estimation even requires an accurate determination of the production mode (within the Heavy Parent class of scenarios). In other words, could one measure the parent particle mass in the

⁶Again, as for the Charged Current mode, we investigated variables that directly reconstruct the parent mass using MATHUSLA and CMS information, which in this case have to also contend with combinatorics background. In analogy to the CC analysis, the simple H_T again delivered superior performance.

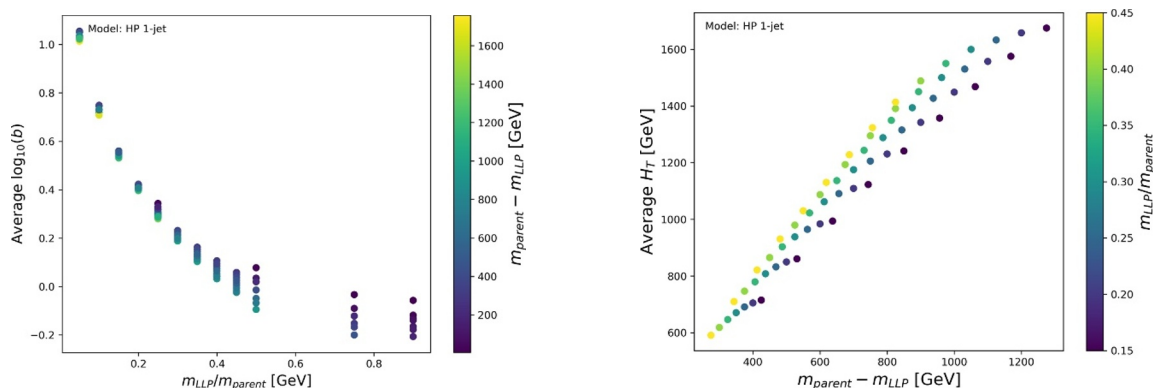


Figure 11. For the Heavy Parent (HP) production mode, we show on the left and right plots how the high-statistics distributions of our estimator variables $x_1 = \log_{10}(b)$ and $x_2 = H_T$ depend on $m_{\text{LLP}}/m_{\text{parent}}$ and $m_{\text{parent}} - m_{\text{LLP}}$ in complementary ways. The color legend shows the subdominant variation with the complementary combination of LLP and parent mass.

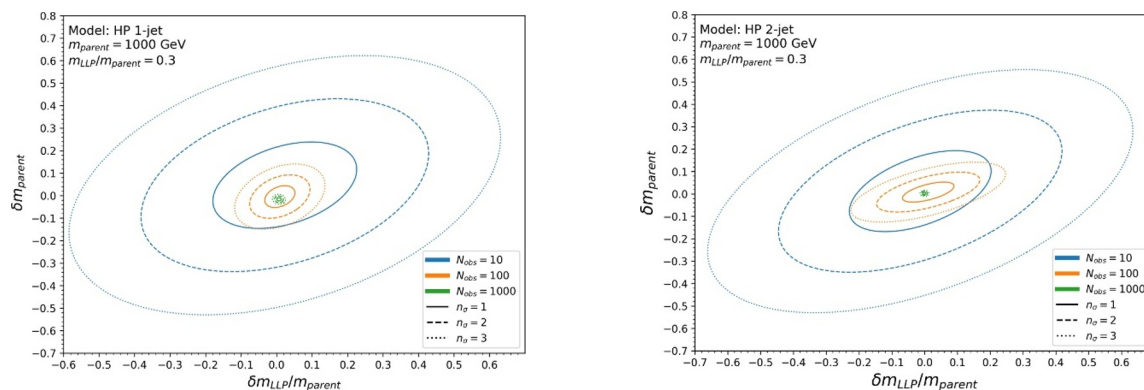


Figure 12. Parameter estimation results for a benchmark point $(m_{\text{parent}}, m_{\text{LLP}}/m_{\text{parent}}) = (1000 \text{ GeV}, 0.3)$ in the Heavy Parent model, with 1-jet decay on the left and 2-jet decay on the right, showing 1, 2, and 3σ ellipses for $N_{\text{obs}} = 10, 100, 1000$ observed events. The axes are in units of $\delta x \equiv \frac{x_{\text{fit}} - x_{\text{true}}}{x_{\text{true}}}$, where x_{true} and x_{fit} are the actual values of the parameters and their measured value extracted from the fit.

N_{obs}	10	100	1000
N_{samples}	1154	116	11
Relative precision of $m_{\text{LLP}}/m_{\text{parent}}$	0.20	0.042	5.8×10^{-3}
Relative precision of m_{parent} [GeV]	0.18	0.048	7.2×10^{-3}
Pearson correlation coefficient	0.37	0.34	-0.093

Table 9. Heavy Parent 1-jet decay model mass estimation results for $m_{\text{LLP}} = 300 \text{ GeV}$, $m_{\text{parent}} = 1 \text{ TeV}$, using estimators $x_1 = \log_{10}(b_{\text{LLP}})$, $x_2 = H_T$. N_{samples} is the number of pseudo-data sets tested for the indicated number of events N_{obs} .

N_{obs}	10	100	1000
N_{samples}	2450	244	23
Relative precision of $m_{\text{LLP}}/m_{\text{parent}}$	$\gtrsim 0.22$	0.078	4.3×10^{-3}
Relative precision of m_{parent} [GeV]	0.18	0.044	4.5×10^{-3}
Pearson correlation coefficient	0.50	0.61	0.054

Table 10. Heavy Parent 2-jet decay model mass estimation results for $m_{\text{LLP}} = 300$ GeV, $m_{\text{parent}} = 1$ TeV, using estimators $x_1 = \log_{10}(b_{\text{LLP}})$, $x_2 = H_T$. N_{samples} is the number of pseudo-data sets tested for the indicated number of events N_{obs} .

N_{obs}	10	100	1000
N_{samples}	3229	321	32
Relative precision of m_{LLP}	0.15	0.037	0.011

Table 11. Exotic Higgs decay model mass estimation results with estimator $\log_{10}(b_{\text{LLP}})$ and true $m_{\text{LLP}} = 35$ GeV. N_{samples} is the number of pseudo-data sets tested for the indicated number of events N_{obs} .

HP scenario, without knowing the exact hadronic decay mode of the parent? To test this, we performed a parameter estimation on HP two-jet data with $m_{\text{parent}} = 1000$ GeV, $m_{\text{LLP}} = 300$ GeV, but using the one-jet hypothesis as our fitting templates. The best-fit masses had a large positive bias, in many cases beyond the range of simulated masses. This is not unexpected, since the final states of a 2-body decay have lower combined H_T than the final states of a 3-body decay. This bias should therefore persist for lower masses, indicating that precise measurements of m_{parent} and m_{LLP} requires fairly precise knowledge of the parent decay mode.⁷

5.4 Exotic Higgs decay

In the Exotic Higgs Decay model, the average boost $b \sim m_h/2m_{\text{LLP}}$ of LLPs is strongly correlated with LLP mass, see figure 13 (left). This scenario was already analyzed in ref. [20], but we include it here for completeness. Maximum likelihood estimation was performed on the boost distributions of many pseudo-data samples at $m_{\text{LLP}} = 35$ GeV, similarly to the procedure for exotic B -meson decays to evaluate the mass estimation precision. Because the precision does not vary significantly over the range of m_{LLP} considered for this model, we only show results for this one mass point. Table 11 summarizes the distributions of best-fit masses for $N_{\text{obs}} = 10, 100, 1000$.

⁷It might be possible to construct some H_T -like variable that takes final state jet multiplicity into account event-by-event, such that a measurement of m_{parent} becomes possible even without a detailed hypothesis on the precise decay mode. We leave this for future investigations.




Figure 13. Left: average LLP boost as a function of LLP mass in the Exotic Higgs Decay model. Right: for the Direct Pair Production model. Note the dependence becomes weaker at higher masses for DPP.

True m_{LLP} [GeV]	100			1000			1700		
N_{obs}	10	100	1000	10	100	1000	10	100	1000
N_{samples}	779	78	8	2523	244	25	2372	237	23
Relative precision of m_{LLP}	0.50	0.13	0.033	.28	0.085	0.049	> 0.25	\approx 0.14	0.068

Table 12. Direct Pair Production model mass estimation results with estimator $\log_{10}(b_{\text{LLP}})$. N_{samples} is the number of pseudo-data sets tested for the indicated number of events N_{obs} .

5.5 Direct pair production

Our strategy for the Direct Pair Production model is the same as for Higgs decay. The correlation between boost and LLP mass is strong for this model, degrading in the high-mass limit, see figure 13 (right). Maximum likelihood estimation was performed for ensembles of pseudo-experiments with $N_{\text{obs}} = 10, 100, 1000$ GeV, for benchmark masses 100, 1000, 1700 GeV. Table 12 gives the mean and relative precision of the best-fit masses obtained for each choice. Note that since the highest mass point simulated was 2 TeV, this is the upper bound on best-fit masses reported. Therefore, the spread in best-fit masses may be underestimated for $m_{\text{LLP}} = 1700$ GeV, $N_{\text{obs}} = 10, 100$ (although the prior on detecting much heavier states at the LHC would be suppressed due to kinematic suppression of the production rate). For all three benchmark masses, resolution of 15% or better can be achieved with 100 observed LLPs.

5.6 Heavy resonance

The final simplified model to consider is the intermediate s -channel Heavy Resonance. Like the Heavy Parent scenario, there are two sub-categories. In this case they are the narrow- and high-width limits for the resonant particle. Just as for HP, we will first consider parameter estimation in each case separately, then investigate the impact of using the incorrect width hypothesis to estimate parameters. Finally, we will demonstrate that the shape of the boost distribution provides useful information on the resonance width.




Figure 14. For the Heavy Resonance (RES) production mode, we show on the left and right plots how the high-statistics distributions of our estimator variables $x_1 = \log_{10}(b)$ and $x_2 = n_{jet}$ depend on m_{LLP}/m_{parent} and m_{parent} in complementary ways. The color legend shows the subdominant variation with the complementary combination of LLP and parent mass.

In the RES simplified model, the LLP boost distribution is almost entirely dependent on m_{LLP}/m_{parent} , with very little separate dependence on m_{parent} . This is demonstrated in figure 14 (left). This not only makes $x_1 = b$ the obvious choice for one of our fit variables, it also means the LLP to parent mass ratio can be extracted independently.

To find the second complementary variable that independently measures m_{parent} , the only handle we have available is the fact that the energy scale of ISR jets in the detector is determined by the overall energy scale of the hard process, which is m_{parent} . Several variables were considered as candidates for x_2 , including the highest main-detector jet transverse momentum, p_T^{j1} , the scalar sum of jet transverse momenta H_T , and the number of jets with $p_T > 20$ GeV, n_{jet} . It was found that n_{jet} gave the best sensitivity, and its dependence on m_{parent} is shown in figure 14 (right). The weakening dependence of $\langle n_{jet} \rangle$ at high m_{parent} is consistent with the logarithmic dependence of ISR kinematics on the hard scale of the event.

These two variables were used for the two-dimensional binned likelihood fit to measure m_{parent} and m_{LLP}/m_{parent} in the same fashion as the Charged Current and Heavy Parent analyses, using pseudo-data samples with $m_{parent} = 1000$ GeV, $m_{LLP} = 300$ GeV. Unfortunately, the variance in n_{jet} is so high that only for $N_{obs} = 1000$ was m_{parent} estimated with precision better than the range of heavy resonance masses simulated. We therefore do not report a precision for measurement of m_{parent} for $N_{obs} \lesssim 100$ for our benchmark point, but given the steeper m_{parent} dependence of n_{jet} for lighter parent masses, we expect that observation of fewer events might be sufficient to measure parent masses in the $\mathcal{O}(100$ GeV) range.

Parameter estimation in the Heavy Resonance model is likely to benefit from other measurements at the LHC main detectors, since the s -channel mediator must have a visible SM decay channel, at minimum through the same process by which it was produced. This is similar to dark matter simplified models involving a heavy vector mediator, where some of the strongest constraints come from LHC searches for di-jet resonances [54–57]. Current

N_{obs}	10	100	1000
N_{samples}	6565 (5104)	2279 (508)	228 (51)
Relative precision of $m_{\text{LLP}}/m_{\text{parent}}$	0.066 (0.12)	0.022 (0.044)	7.2×10^{-3} (0.017)
Relative precision of m_{parent} [GeV]	—	—	0.12 (0.12)
Pearson correlation coefficient	—	—	-0.57 (-0.84)

Table 13. Heavy Resonance model mass estimation results for $m_{\text{LLP}} = 300$ GeV, $m_{\text{parent}} = 1$ TeV, using estimators $x_1 = \log_{10}(b_{\text{LLP}})$, $x_2 = n_{\text{jet}}$. N_{samples} is the number of pseudo-data sets tested for the indicated number of events N_{obs} . Results for high-width are shown in brackets next to those for narrow-width. Note that a measurement of m_{parent} using LLP events is only possible for ~ 1000 observed decays for our benchmark point.

and future colliders have the potential to detect the heavy mediator and measure its mass with $\mathcal{O}(10\%)$ or better precision with standard resonance search techniques [58, 59]. If LLPs are discovered at MATHUSLA and classified as being produced via an s -channel resonance, the observed events in MATHUSLA could be used to independently measure $m_{\text{LLP}}/m_{\text{parent}}$. Table 13 demonstrates that this ratio can be measured with $\sim 5\%$ precision or better for 10-100 observed events. Resonance searches at the main detectors would then provide an independent measurement of m_{parent} , allowing all mass parameters of the production mode to be determined.

To evaluate the robustness of parameter estimation for this model under variation of the width of the heavy resonance, likelihood fits were performed for samples of narrow width events using template distributions from the high-width case, and vice versa. Using this method, $m_{\text{LLP}}/m_{\text{parent}}$ can be estimated with similar precision to that achieved using the correct template distributions, but with a small bias. The average $m_{\text{LLP}}/m_{\text{parent}}$ for narrow width samples under the high-width assumption is higher than the true value by approximately 1%, but this is smaller than the spread of the distribution of best-fit masses up to $N_{\text{obs}} = 1000$. The average mass ratio for high-width samples under the narrow width assumption is biased slightly lower than the true value, by approximately 3%. The LLP to parent mass ratio can therefore be reliably extracted from the LLP boost distribution measured at MATHUSLA alone, once the production mode has been classified as Heavy Resonance, regardless of the resonance width.

On the other hand, estimation of m_{parent} using the incorrect width fails completely. In both cases, the best-fit masses were at the edge of the range of masses simulated, at least 50% off of the true value, indicating that it is not possible to estimate m_{parent} without an estimate or assumption of the heavy resonant particle’s width, regardless of N_{obs} . This emphasizes the importance of main detector resonance searches to determine all parameters of this LLP simplified production model.

In the high statistics limit, the width and mass of the resonance could be measured independently of a separate resonance search by performing a full three-dimensional likelihood fit for the LLP mass, parent mass, and parent width. This is beyond our scope, but we demonstrate that this is possible in principle by classifying pseudo-data sets as low-

N_{obs}	10	100	1000
N_{samples}	22812 (5104)	2279 (508)	228 (51)
Classification accuracy	0.81 (0.57)	0.96 (0.96)	1.0 (1.0)

Table 14. Heavy Resonance model width classification results from likelihood ratios of $\log_{10}(b_{\text{LLP}})$ distribution with true $m_{\text{parent}} = 1$ TeV, $m_{\text{LLP}} = 300$ GeV, and $\Gamma_{\text{parent}} = 10$ GeV, 300 GeV. Results for high-width are shown in brackets next to those for narrow width. N_{samples} is the number of pseudo-data sets tested for the indicated number of events N_{obs} .

or high-width based on the shape of the sample’s boost distribution, using the ratio of maximum likelihoods under the two width hypotheses as a discriminant. The results are summarized in table 14 for true $m_{\text{parent}} = 1$ TeV, $m_{\text{LLP}} = 300$ GeV. It is clear that at least the two extremes of possible heavy resonance widths can be reliably distinguished, and that with enough observed events the heavy resonance’s width could be estimated simultaneously with BSM particle masses.

6 Uncertainty in matching LLP decay to production vertex

The analyses above demonstrate that combining information from MATHUSLA and the main detector in the event of an LLP discovery would allow both the production mode and the mass parameters of the underlying simplified model to be determined. For simplicity, we have so far assumed that each LLP decay at MATHUSLA could be associated with a single hard primary vertex at the HL-LHC. However, this is unlikely to be the case, partially due to some uncertainty in the LLP boost measurement at MATHUSLA, which narrows down the bunch crossings that produced the LLP down to a few, but more importantly because of the high pile-up multiplicity of 140-200 during the HL-LHC running conditions [60, 61],

Even given the fact that ‘interesting’ hard events are in absolute terms fairly rare in pp collisions, and that the primary vertex with the highest associated total energy in the main detector is likely to be the one that corresponds to the LLP production event for most production modes, it seems apparent that in a realistic analysis, each LLP observation at MATHUSLA will correspond to a list of candidate primary vertices, instead of just one. We must therefore ask how our classification and parameter estimation analysis would be affected by this complication. We find that our methods are robust and should be adaptable to account for an ambiguity in the exact production vertex.

6.1 Impact on production mode classification

If there is an ‘ensemble’ of multiple candidate hard primary vertices per LLP, then instead of cuts on the fraction of events meeting certain criteria, the classifier would impose cuts on the fraction of *ensembles* of candidate primary vertices meeting criteria. This could be as simple as demanding that at least one of the candidate vertices corresponding to an LLP observation satisfy certain conditions. We explore the impact of such a change for each of the simplified models.

1. Since the requirement to classify a sample as Exotic B -meson Decay is a lack of hard scattering events, the additional candidate hard scatters will not lead to other models being classified as BB. However, it is likely that BB samples will be more likely to fail this criterion. The maximum fraction of events across all ensembles with jets of $p_T > 20$ GeV can be decreased to accommodate this, or we could impose a cut on not the fraction of events with a hard jet, but the fraction of ensembles where a certain fraction of events have a hard jet. This can be optimized in a realistic analysis.
2. The correct production vertex in the Charged Current model can still be identified by a hard lepton, which does not occur often in the detector. Therefore, classification for this model should be relatively unaffected.
3. Classification of the Heavy Parent model relies on jet multiplicity, which is vulnerable to contamination from spurious vertices. This contamination can be reduced by raising the minimum p_T and/or jet multiplicity for a candidate vertex to pass selection criteria. There may be a small penalty to classification accuracy for this model, but it should still perform similarly to what we demonstrated in this work.
4. The main signature used to identify Exotic Higgs Decay samples is the presence of well-separated jets from VBF. While correct identification of the production vertex may be more difficult for the GF production channel, the presence of VBF-like jets in roughly 10% of the sets of candidate vertices should be roughly unaltered, as should the classifier's performance for this model.
5. The classification of the DPP model uses only LLP boost information, and is therefore unaffected by uncertainty in the production vertex identification.

With only small modifications, the classification procedure outlined in section 4 should therefore be remarkably robust under the introduction of multiple possible LLP production vertices. We turn now to a discussion of how parameter estimation would proceed in these circumstances.

6.2 Impact on parameter estimation

For the Exotic B -meson Decay, Exotic Higgs Decay, and Direct Pair Production models, parameter estimation only uses LLP boost, so it is unaffected. For the Charged Current and Heavy Parent models, the correct production vertex can likely be identified in a high fraction of events due to associated hard objects or leptons, so parameter estimation can proceed as normal. Finally, for the Heavy Resonance model, estimation of $m_{\text{LLP}}/m_{\text{parent}}$ uses only LLP boost and is unaffected. Estimation of m_{parent} using the ISR spectrum is unreliable even when the correct vertex is identified, and is likely to rely on separate resonance searches performed with the main detectors. Therefore, the ability of MATHUSLA to estimate BSM particle masses is essentially unaffected by the complication of multiple possible LLP production vertices.

7 Conclusions

Many well-motivated scenarios for BSM physics include long-lived particles, and it is important that the experimental search program for long-lived particles continue to be developed and expanded. The proposed MATHUSLA detector represents a vital part of that program, able to probe parameter space inaccessible by any other experiment. Allowing MATHUSLA to act as a Level-1 Trigger for CMS is required to take full advantage of both detectors' potential. This would enable an analysis of LLP observations at MATHUSLA to be augmented by observation of their production events in the main detector. With information from both experiments, both the production and decay topology of long-lived particles, as well as the BSM particle masses, can be accurately determined.

We have shown that for heavy LLPs, an accurate diagnosis at the simplified model level can be achieved $> 90\%$ of the time with only 100 observed events, and $\approx 98\%$ of the time with 1000 observed events. Similar performance is possible for lighter LLPs, except for the Heavy Parent and Charged Current models with squeezed spectra $m_{\text{LLP}} \sim m_{\text{parent}} < 100$ GeV, where classification fails because the associated objects being used to identify the production mode become too soft. With similar statistics, the underlying parameters of the simplified model, like LLP and parent particle mass, can be measured with $\sim 10\%$ precision or better in most cases. The exception is the Heavy Resonance production mode, where the ratio of LLP to parent particle mass can be reliably extracted from the LLP boost distribution, but measurement of the parent particle mass would likely benefit from separate resonance searches conducted with the main detectors, in analogy to mediator resonance searches in Dark Matter simplified models [54–57].

This performance is achieved with extremely simple cuts and analyses using robust, physically motivated features of LLP production events. Further work is sure to improve on our demonstrated classification accuracy and measurement precision. One case that will require further study is LLP production in exotic B -meson decays, where the masses of LLP decay products cannot be neglected, and the measurement of LLP boost at MATHUSLA is consequently more difficult. We have also discussed how to extend our work to include other LLP production scenarios like Heavy Parent decay to leptons + LLP, Heavy Parent with multi-step decay chains, or the consideration of a Higgs-like intermediate resonance of indeterminate mass. It is also imaginable that post-discovery, we suffer an embarrassment of riches, with many different LLP production and decay modes being observed. In that case, our analysis would have to be extended to allow for more than one dominant channel, but this work still demonstrates in principle how the information in each channel could be extracted. Finally, if the LLP lifetime is in the range $\bar{b}c\tau \sim 1 - 100$ m, the geometric distribution of DVs in MATHUSLA's decay volume could be used to extract the lifetime directly. Our methods are applicable not just to MATHUSLA and CMS, but also to other external LLP detector proposals, or even LLPs discovered using LHC main detectors alone. This emphasizes not only the great discovery potential of new LHC detectors like MATHUSLA [15], FASER [17] or CODEX-b [16], but also shows that in the event of a discovery, however it took place, the origin of Long-Lived Particles can be uncovered in great detail.

Acknowledgments

We thank Erez Etzion, Henry Lubatti, Jessie Shelton, David Wandler, Steven Winter, and Charlie Young for helpful discussions. The research of JB and DC was supported in part by a Discovery Grant from the Natural Sciences and Engineering Research Council of Canada, and by the Canada Research Chair program. The research of JB was supported in part by a Canada Graduate Scholarship from the Natural Sciences and Engineering Research Council of Canada.

A Simulation details

This section contains a detailed description of how LLP production events were simulated for each of the simplified models considered in this work.

1. For the B -meson decay model, $b\bar{b}$ production was simulated with the MadGraph Standard Model. The B -mesons produced were identified in the simulation at hadronization level. If two were present, one was randomly chosen and manually decayed to two- or three-body final states consisting of SM electrons and/or mesons as well as the LLP in accordance with the relative branching ratios in [25, 27], depending on the LLP type. Heavier B -baryons were present in approximately 10% of events, but were ignored for simplicity. (A more realistic analysis would include them separately, but this would not affect our conclusions.) If a jet with $p_T > 20$ GeV corresponding to the B -meson chosen for LLP decay was present in the simulated event at detector level, it was removed from the event. A jet was defined as corresponding to a B -meson if the $\Delta R = \sqrt{(\Delta\phi)^2 + (\Delta\eta)^2}$ between the jet and B -meson was less than 0.5. This value was chosen by considering the distribution of ΔR to find the characteristic distance between B -mesons and their corresponding jets at detector level. The SM products of the exotic B decay were not added back into the event. Since most $b\bar{b}$ production at the LHC is near threshold energy, and the B -mesons characteristically have momentum ≈ 5 GeV, only $\approx 2\%$ of events contained jets with p_T above the 20 GeV threshold at detector level. Therefore, the effect of completely discarding the SM products of the exotic decay is negligible. A reweighting of these events was performed by comparing the B -meson p_T spectrum to one obtained from FONLL for B -mesons with $0 < \eta < 3$, roughly corresponding to those heading towards MATHUSLA [62].
2. For the Charged Current model, the *ChargeWN* MadGraph model from the LLP Simplified Model Library was used [13]. The simulated process was production of an (anti-)muon plus the lightest neutralino, mediated by an s-channel W'^{\pm} boson. Because of the distinctive hard lepton present in these events, rendering analysis of jet variables irrelevant for model classification, jet matching was not performed for these samples. The choice was made to restrict to the case of decay of the W' to muons plus LLPs. Decay to other leptons or lepton-universal decays would not present any additional challenges for LLP production mode identification.

3. For the Heavy Parent model, the minimal supersymmetric standard model (*MSSM*) model was used [7, 45–49]. In the one-jet decay mode the simulated process was squark anti-squark production, with each squark (anti-squark) decaying to either one quark or anti-quark as well as the lightest neutralino, which serves as the LLP. In the two-jet decay mode the simulated process was gluino pair production, with each gluino decaying to two quarks via an off-shell 10 TeV squark. Because the most important jets in these events originate from the hard process, jet matching was not performed on these samples. The choice was made to restrict the decay of the heavy parent particle to jets only, because this presents the most difficult case to distinguish from other simplified models. Decay to leptons plus LLPs would be identifiable by a large proportion of events containing multiple hard leptons.
4. For the Exotic Higgs Decay model, Higgs production was simulated using the *heft* model in Madgraph [50–52]. The gluon fusion, vector boson fusion, and associated vector boson production modes were included. Jet matching was performed up to one jet for the gluon fusion channel, with the jet matching variable *xqcut* set to 15. Jet matching was performed to increase the reliability of simulated jet-related variables, since the only jets in these events are due to initial state radiation, and none are produced in the hard process. The Higgs boson was not permitted to decay. For each event, a pair of LLP 4-vectors with energy $m_H/2$ were manually generated pointing in a random direction and its opposite, as one would expect in Higgs decay to two scalars in the Higgs rest frame. These 4-vectors were then boosted into the lab frame using the Higgs momentum from simulation.
5. For the Direct Pair Production model, the same *MSSM* package was used, and pair production of the lightest neutralino as the LLP was the simulated process. The intermediate, t-/u-channel squark was set to 10 TeV to remove it as a dynamical degree of freedom. Jet matching was performed up to one jet, and the variable *xqcut* was set to $0.4 \times m_{LLP}$. This choice of *xqcut* was made to produce smooth distributions of jet kinematics variables that were stable under small variation of *xqcut*.
6. For the Heavy Resonance model, the simplified model Zp_{2LLP} was used [13]. The process simulated was dark scalar pair production with a Z' in the s-channel, with the dark scalar's width set to 1×10^{-20} GeV so that they did not decay in simulation. Jet matching was performed up to one jet, and the variable *xqcut* was set to $10 + 0.1 \times m_{\text{parent}}$ after a similar optimization procedure to that used for Direct Pair Production.

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] ATLAS collaboration, *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC*, *Phys. Lett. B* **716** (2012) 1 [[arXiv:1207.7214](#)] [[INSPIRE](#)].
- [2] CMS collaboration, *Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC*, *Phys. Lett. B* **716** (2012) 30 [[arXiv:1207.7235](#)] [[INSPIRE](#)].
- [3] PLANCK collaboration, *Planck 2015 results. XIII. Cosmological parameters*, *Astron. Astrophys.* **594** (2016) A13 [[arXiv:1502.01589](#)] [[INSPIRE](#)].
- [4] V.C. Rubin, N. Thonnard and J. Ford, W. K., *Rotational properties of 21 Sc galaxies with a large range of luminosities and radii, from NGC 4605 ($R = 4$ kpc) to UGC 2885 ($R = 122$ kpc)*, *Astrophys. J.* **238** (1980) 471 [[INSPIRE](#)].
- [5] SUPER-KAMIOKANDE collaboration, *Evidence for oscillation of atmospheric neutrinos*, *Phys. Rev. Lett.* **81** (1998) 1562 [[hep-ex/9807003](#)] [[INSPIRE](#)].
- [6] SNO collaboration, *Measurement of the rate of $\nu_e + d \rightarrow p + p + e^-$ interactions produced by ^8B solar neutrinos at the Sudbury Neutrino Observatory*, *Phys. Rev. Lett.* **87** (2001) 071301 [[nucl-ex/0106015](#)] [[INSPIRE](#)].
- [7] S.P. Martin, *A supersymmetry primer*, *Adv. Ser. Direct. High Energy Phys.* **18** (1998) 1 [*Adv. Ser. Direct. High Energy Phys.* **21** (2010) 1] [[hep-ph/9709356](#)] [[INSPIRE](#)].
- [8] C. Csáki, S. Lombardo and O. Telem, *TASI lectures on non-supersymmetric BSM models*, in *Proceedings, Theoretical Advanced Study Institute in Elementary Particle Physics: anticipating the next discoveries in particle physics (TASI 2016)*, Boulder, CO, U.S.A., 6 June–1 July 2016, R. Essig and I. Low eds., *World Scientific*, Singapore (2018), pg. 501 [[arXiv:1811.04279](#)] [[INSPIRE](#)].
- [9] ATLAS collaboration, *Search for a scalar partner of the top quark in the jets plus missing transverse momentum final state at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *JHEP* **12** (2017) 085 [[arXiv:1709.04183](#)] [[INSPIRE](#)].
- [10] ATLAS collaboration, *Search for R-parity-violating supersymmetric particles in multi-jet final states produced in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC*, *Phys. Lett. B* **785** (2018) 136 [[arXiv:1804.03568](#)] [[INSPIRE](#)].
- [11] CMS collaboration, *Search for direct top squark pair production in events with one lepton, jets, and missing transverse momentum at 13 TeV with the CMS experiment*, *JHEP* **05** (2020) 032 [[arXiv:1912.08887](#)] [[INSPIRE](#)].
- [12] ATLAS collaboration, *Search for new resonances in mass distributions of jet pairs using 139 fb^{-1} of pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *JHEP* **03** (2020) 145 [[arXiv:1910.08447](#)] [[INSPIRE](#)].
- [13] J. Alimena et al., *Searching for long-lived particles beyond the Standard Model at the Large Hadron Collider*, *J. Phys. G* **47** (2020) 090501 [[arXiv:1903.04497](#)] [[INSPIRE](#)].
- [14] D. Curtin et al., *Long-lived particles at the energy frontier: the MATHUSLA physics case*, *Rept. Prog. Phys.* **82** (2019) 116201 [[arXiv:1806.07396](#)] [[INSPIRE](#)].
- [15] J.P. Chou, D. Curtin and H.J. Lubatti, *New detectors to explore the lifetime frontier*, *Phys. Lett. B* **767** (2017) 29 [[arXiv:1606.06298](#)] [[INSPIRE](#)].

- [16] V.V. Gligorov, S. Knapen, M. Papucci and D.J. Robinson, *Searching for long-lived particles: a compact detector for exotics at LHCb*, *Phys. Rev. D* **97** (2018) 015023 [[arXiv:1708.09395](#)] [[INSPIRE](#)].
- [17] J.L. Feng, I. Galon, F. Kling and S. Trojanowski, *ForwArd Search ExpeRiment at the LHC*, *Phys. Rev. D* **97** (2018) 035001 [[arXiv:1708.09389](#)] [[INSPIRE](#)].
- [18] SHiP collaboration, *A facility to Search for Hidden Particles (SHiP) at the CERN SPS*, [arXiv:1504.04956](#) [[INSPIRE](#)].
- [19] MATHUSLA collaboration, *A Letter of Intent for MATHUSLA: a dedicated displaced vertex detector above ATLAS or CMS*, [arXiv:1811.00927](#) [[INSPIRE](#)].
- [20] D. Curtin and M.E. Peskin, *Analysis of long lived particle decays with the MATHUSLA detector*, *Phys. Rev. D* **97** (2018) 015006 [[arXiv:1705.06327](#)] [[INSPIRE](#)].
- [21] MATHUSLA collaboration, *Explore the lifetime frontier with MATHUSLA*, *2020 JINST* **15** C06026 [[arXiv:1901.04040](#)] [[INSPIRE](#)].
- [22] I. Alkhatib, *Geometric optimization of the MATHUSLA detector*, [arXiv:1909.05896](#) [[INSPIRE](#)].
- [23] LHC NEW PHYSICS WORKING GROUP collaboration, *Simplified models for LHC new physics searches*, *J. Phys. G* **39** (2012) 105005 [[arXiv:1105.2838](#)] [[INSPIRE](#)].
- [24] J. Beacham et al., *Physics beyond colliders at CERN: beyond the Standard Model working group report*, *J. Phys. G* **47** (2020) 010501 [[arXiv:1901.09966](#)] [[INSPIRE](#)].
- [25] K. Bondarenko, A. Boyarsky, D. Gorbunov and O. Ruchayskiy, *Phenomenology of GeV-scale heavy neutral leptons*, *JHEP* **11** (2018) 032 [[arXiv:1805.08567](#)] [[INSPIRE](#)].
- [26] E.J. Chun, A. Das, S. Mandal, M. Mitra and N. Sinha, *Sensitivity of lepton number violating meson decays in different experiments*, *Phys. Rev. D* **100** (2019) 095022 [[arXiv:1908.09562](#)] [[INSPIRE](#)].
- [27] J.A. Evans, *Detecting hidden particles with MATHUSLA*, *Phys. Rev. D* **97** (2018) 055046 [[arXiv:1708.08503](#)] [[INSPIRE](#)].
- [28] J.C. Pati and A. Salam, *Lepton number as the fourth color*, *Phys. Rev. D* **10** (1974) 275 [*Erratum ibid.* **11** (1975) 703] [[INSPIRE](#)].
- [29] R.N. Mohapatra and J.C. Pati, *A natural left-right symmetry*, *Phys. Rev. D* **11** (1975) 2558 [[INSPIRE](#)].
- [30] G. Senjanović and R.N. Mohapatra, *Exact left-right symmetry and spontaneous violation of parity*, *Phys. Rev. D* **12** (1975) 1502 [[INSPIRE](#)].
- [31] W.-Y. Keung and G. Senjanović, *Majorana neutrinos and the production of the right-handed charged gauge boson*, *Phys. Rev. Lett.* **50** (1983) 1427 [[INSPIRE](#)].
- [32] F.F. Deppisch, P.S. Bhupal Dev and A. Pilaftsis, *Neutrinos and collider physics*, *New J. Phys.* **17** (2015) 075019 [[arXiv:1502.06541](#)] [[INSPIRE](#)].
- [33] D. Curtin et al., *Exotic decays of the 125 GeV Higgs boson*, *Phys. Rev. D* **90** (2014) 075004 [[arXiv:1312.4992](#)] [[INSPIRE](#)].
- [34] M.J. Strassler and K.M. Zurek, *Echoes of a hidden valley at hadron colliders*, *Phys. Lett. B* **651** (2007) 374 [[hep-ph/0604261](#)] [[INSPIRE](#)].

- [35] M.J. Strassler and K.M. Zurek, *Discovering the Higgs through highly-displaced vertices*, *Phys. Lett. B* **661** (2008) 263 [[hep-ph/0605193](#)] [[INSPIRE](#)].
- [36] N. Craig, S. Knapen and P. Longhi, *Neutral naturalness from orbifold Higgs models*, *Phys. Rev. Lett.* **114** (2015) 061803 [[arXiv:1410.6808](#)] [[INSPIRE](#)].
- [37] N. Craig, A. Katz, M. Strassler and R. Sundrum, *Naturalness in the dark at the LHC*, *JHEP* **07** (2015) 105 [[arXiv:1501.05310](#)] [[INSPIRE](#)].
- [38] F. Yu, *Di-jet resonances at future hadron colliders: a Snowmass whitepaper*, [arXiv:1308.1077](#) [[INSPIRE](#)].
- [39] A. Das, P.S.B. Dev and N. Okada, *Long-lived TeV-scale right-handed neutrino production at the LHC in gauged $U(1)_X$ model*, *Phys. Lett. B* **799** (2019) 135052 [[arXiv:1906.04132](#)] [[INSPIRE](#)].
- [40] C.-W. Chiang, G. Cottin, A. Das and S. Mandal, *Displaced heavy neutrinos from Z' decays at the LHC*, *JHEP* **12** (2019) 070 [[arXiv:1908.09838](#)] [[INSPIRE](#)].
- [41] J. Alwall et al., *The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations*, *JHEP* **07** (2014) 079 [[arXiv:1405.0301](#)] [[INSPIRE](#)].
- [42] T. Sjöstrand, S. Mrenna and P.Z. Skands, *PYTHIA 6.4 physics and manual*, *JHEP* **05** (2006) 026 [[hep-ph/0603175](#)] [[INSPIRE](#)].
- [43] T. Sjöstrand, S. Mrenna and P.Z. Skands, *A brief introduction to PYTHIA 8.1*, *Comput. Phys. Commun.* **178** (2008) 852 [[arXiv:0710.3820](#)] [[INSPIRE](#)].
- [44] DELPHES 3 collaboration, *DELPHES 3, a modular framework for fast simulation of a generic collider experiment*, *JHEP* **02** (2014) 057 [[arXiv:1307.6346](#)] [[INSPIRE](#)].
- [45] H.P. Nilles, *Supersymmetry, supergravity and particle physics*, *Phys. Rept.* **110** (1984) 1 [[INSPIRE](#)].
- [46] H.E. Haber and G.L. Kane, *The search for supersymmetry: probing physics beyond the Standard Model*, *Phys. Rept.* **117** (1985) 75 [[INSPIRE](#)].
- [47] J. Rosiek, *Complete set of Feynman rules for the minimal supersymmetric extension of the Standard Model*, *Phys. Rev. D* **41** (1990) 3464 [[INSPIRE](#)].
- [48] P.Z. Skands et al., *SUSY Les Houches accord: interfacing SUSY spectrum calculators, decay packages, and event generators*, *JHEP* **07** (2004) 036 [[hep-ph/0311123](#)] [[INSPIRE](#)].
- [49] B.C. Allanach et al., *SUSY Les Houches accord 2*, *Comput. Phys. Commun.* **180** (2009) 8 [[arXiv:0801.0045](#)] [[INSPIRE](#)].
- [50] R.P. Kauffman and S.V. Desai, *Production of a Higgs pseudoscalar plus two jets in hadronic collisions*, *Phys. Rev. D* **59** (1999) 057504 [[hep-ph/9808286](#)] [[INSPIRE](#)].
- [51] T.G. Rizzo, *Gluon final states in Higgs boson decay*, *Phys. Rev. D* **22** (1980) 178 [*Addendum ibid.* **22** (1980) 1824] [[INSPIRE](#)].
- [52] R.P. Kauffman, S.V. Desai and D. Risal, *Amplitudes for Higgs bosons plus four partons*, in *American Physical Society (APS) meeting of the Division of Particles and Fields (DPF 99)*, (1999) [[hep-ph/9903330](#)] [[INSPIRE](#)].
- [53] M. Burns, K. Kong, K.T. Matchev and M. Park, *Using subsystem MT2 for complete mass determinations in decay chains with missing energy at hadron colliders*, *JHEP* **03** (2009) 143 [[arXiv:0810.5576](#)] [[INSPIRE](#)].

- [54] M. Fairbairn, J. Heal, F. Kahlhoefer and P. Tunney, *Constraints on Z' models from LHC dijet searches and implications for dark matter*, *JHEP* **09** (2016) 018 [[arXiv:1605.07940](#)] [[INSPIRE](#)].
- [55] M. Chala, F. Kahlhoefer, M. McCullough, G. Nardini and K. Schmidt-Hoberg, *Constraining dark sectors with monojets and dijets*, *JHEP* **07** (2015) 089 [[arXiv:1503.05916](#)] [[INSPIRE](#)].
- [56] J. Alexander et al., *Dark sectors 2016 workshop: community report*, [arXiv:1608.08632](#) [[INSPIRE](#)].
- [57] E. Bagnaschi et al., *Global analysis of dark matter simplified models with leptophobic spin-one mediators using MasterCode*, *Eur. Phys. J. C* **79** (2019) 895 [[arXiv:1905.00892](#)] [[INSPIRE](#)].
- [58] CMS collaboration, *Search for dijet resonances in 7 TeV pp collisions at CMS*, *Phys. Rev. Lett.* **105** (2010) 211801 [[arXiv:1010.0203](#)] [[INSPIRE](#)].
- [59] CMS collaboration, *Search for low mass vector resonances decaying into quark-antiquark pairs in proton-proton collisions at $\sqrt{s} = 13$ TeV*, *Phys. Rev. D* **100** (2019) 112007 [[arXiv:1909.04114](#)] [[INSPIRE](#)].
- [60] ATLAS collaboration, *Expected pileup values at the HL-LHC*, Tech. Rep. [ATL-UPGRADE-PUB-2013-014](#), CERN, Geneva, Switzerland (2013).
- [61] G. Soyez, *Pileup mitigation at the LHC: a theorist's view*, *Phys. Rept.* **803** (2019) 1 [[arXiv:1801.09721](#)] [[INSPIRE](#)].
- [62] M. Cacciari, M. Greco and P. Nason, *The p_T spectrum in heavy flavor hadroproduction*, *JHEP* **05** (1998) 007 [[hep-ph/9803400](#)] [[INSPIRE](#)].