Published for SISSA by Springer

Received: *September 23, 2021*
Accepted: *November 29, 2021*
Published: *December 20, 2021*

Unsupervised hadronic SUEP at the LHC

Jared Barron,[a] **David Curtin,**[a] **Gregor Kasieczka,**[b] **Tilman Plehn**[c] **and Aris Spourdalakis**[a]

[a] *Department of Physics, University of Toronto, 60 St. George St., Toronto, Canada*
[b] *Institut für Experimentalphysik, Universität Hamburg, Luruper Chaussee 149 D-22761, Hamburg, Germany*
[c] *Institut für Theoretische Physik, Universität Heidelberg, Philosophenweg 16, Heidelberg, Germany*

E-mail: jared.barron@mail.utoronto.ca, dcurtin@physics.utoronto.ca, gregor.kasieczka@cern.ch, plehn@uni-heidelberg.de, aspourda@physics.utoronto.ca

Abstract: Confining dark sectors with pseudo-conformal dynamics produce SUEPs, or Soft Unclustered Energy Patterns, at colliders: isotropic dark hadrons with soft and democratic energies. We target the experimental nightmare scenario, SUEPs in exotic Higgs decays, where all dark hadrons decay promptly to SM hadrons. First, we identify three promising observables: the charged particle multiplicity, the event ring isotropy, and the matrix of geometric distances between charged tracks. Their patterns can be exploited through a cut-and-count search, supervised machine learning, or an unsupervised autoencoder. We find that the HL-LHC will probe exotic Higgs branching ratios at the per-cent level, even without a detailed knowledge of the signal features. Our techniques can be applied to other SUEP searches, especially the unsupervised strategy, which is independent of overly specific model assumptions and the corresponding precision simulations.

Keywords: Beyond Standard Model, Hadron-Hadron scattering (experiments), Jets, Rare decay

ArXiv ePrint: 2107.12379

Open Access, © The Authors.
Article funded by SCOAP[3].

https://doi.org/10.1007/JHEP12(2021)129

JHEP12(2021)129

Contents

1 Introduction **1**

2 Theory of SUEPs **3**
- 2.1 Dark hadron production in exotic Higgs decays 3
- 2.2 Dark hadron decay 5
- 2.3 Prompt hadronic benchmark scenario 5

3 Simulation **7**
- 3.1 Signal 7
- 3.2 Background 8

4 Analysis **9**
- 4.1 SUEP observables 9
- 4.2 Supervised ML-classifier 11
- 4.3 Unsupervised autoencoder 12
- 4.4 Data-driven background estimation 14

5 Results **14**

6 Conclusion **18**

A Additional results **20**

JHEP12(2021)129

1 Introduction

Hidden sectors are one of the most interesting and generic paradigms for physics beyond the Standard Model (BSM). These kinds of new particles and forces are not only plausible from a bottom-up point of view, but also arise in many top-down BSM theories [1–24]. In most scenarios of interest, hidden sector particles couple to SM fields via feeble interactions or heavy messengers, and the nature of these portal couplings determines their collider phenomenology. A case of special interest are hidden valleys [1], referring to hidden sectors with a confining gauge group which gives rise to rich infrared (IR) dynamics from very simple ultraviolet (UV) theory structures. The production of dark quarks at the LHC leads to a dark shower and high-multiplicity production of dark hadrons, in analogy to QCD jets. Depending on the portal by which the dark hadrons are produced and decay, these dark showers produce a wide variety of LHC signatures, which have been the subject of intense theoretical and experimental study in the last decade [2, 8, 25–35].

We consider one of the most challenging varieties of dark showers, Soft Unclustered Energy Patterns (SUEPs). If a hidden valley possesses a large gauge coupling that is

pseudo-conformal above its confinement scale, then large-angle emission is unsuppressed for most of the parton shower evolution. This means the dark hadrons are not arranged in narrow QCD-like jets, but emitted approximately isotropically in the shower centre-of-mass frame [36, 37]. This defines the SUEP final state as a high-multiplicity spherically-symmetric shower of hidden sector states. While existing searches can be sensitive to SUEPs produced at high energy scales, or with dark hadrons that decay to sufficiently conspicuous final states like leptons or Long-Lived Particles (LLPs) [8, 37], no dedicated SUEP searches exist to date. Furthermore, SUEPs *without* conspicuous final states are not captured by any existing search and represent an unusually cruel signal, since their soft, isotropic distributions can mimic the ubiquitous pile-up produced by simultaneous LHC collisions.

To ensure that all types of SUEP signals can be discovered at the LHC, we focus on a well-motivated SUEP nightmare scenario, where SUEP is produced in exotic Higgs decays and the dark hadrons decay promptly and exclusively to SM hadrons. The modest energy scale of exotic Higgs decays and the lack of conspicuous final states forces us to rely on the kinematics of the resulting SM hadrons to extract the signal from the overwhelming QCD background. This production mode also allows us to side-step the problem of how to trigger on SUEPs [37] by using leptons in associated Vh production. The analysis techniques we develop will not only allow for the detection of this SUEP nightmare scenario, but should also increase the LHC experiments' sensitivity to all other SUEP possibilities.

An acute obstacle for SUEP searches is the lack of rigorous predictions and simulations for the strongly coupled pseudo-conformal dark sectors. Rather than hoping for conspicuous final states, we utilize the kinematics of the SM hadrons without relying on fine details of the signal beyond the robust SUEP characteristics of isotropic, soft, and democratic dark hadron energies. We therefore simulate SUEP production using a simple QCD fireball model of thermal dark hadron emission [20], and find robust observables and event representations. An important observable for SUEP searches is the inter-particle distance matrix ΔR_{ij} between charged hadrons. This matrix encodes the essential geometric differences between QCD-like and SUEP-like hadron production. It forms the backbone of all our analysis strategies, together with known variables like event isotropy [38] and the charged particle multiplicity.

To demonstrate the distinguishing power of our observables, as well as the drastic improvements from more sophisticated techniques, we examine three strategies for systematic SUEP searches at the HL-LHC. First, we simulate a simple cut-and-count analysis, which will turn out to allow for impressive sensitivities to $\mathrm{Br}(h \to \mathrm{SUEP}) \sim 1\%$ at the HL-LHC. We anticipate that a realistic analysis with data-driven background estimation will perform even better, since our study is limited by background simulation statistics. To improve on this, we utilize supervised as well as unsupervised machine learning (ML). Unsupervised analysis concepts along the lines of autoencoder neural networks [39] have the potential to transform LHC analyses [40–56], including searches for dark showers [45]. We point out how unsupervised methods are especially appealing for difficult-to-simulate signals like SUEP since they only rely on the known QCD background for training, while yielding at least several times greater SUEP sensitivity than the cut-and-count analysis.

JHEP12(2021)129

Our investigation establishes that SUEP searches need not rely on conspicuous SM final states for excellent sensitivity at the HL-LHC. The unique dark hadron kinematics, which robustly follows from their origin in pseudo-conformal strongly coupled dynamics, allows for the SUEP final state to be distinguished from its overwhelming QCD background. Our techniques can be applied to all SUEP searches to dramatically enhance their sensitivity, regardless of energy scale or SM final state.

This paper is structured as follows. In section 2, we briefly review the SUEP theory and define the benchmark scenario for our study. Signal and background simulation is discussed in section 3. In section 4 we define the relevant observables to distinguish SUEP from QCD and discuss supervised and unsupervised machine learning techniques. In section 5 we present our results, including projections for the Br($h \to$ SUEP) sensitivity at the HL-LHC, and we conclude in section 6.

2 Theory of SUEPs

Hidden valley models are a large class of BSM theories in which the SM is extended by additional gauge groups under which SM particles are neutral. New particles that are charged under non-Abelian extensions can give rise to a wide range of interesting hidden sector dynamics [1, 2, 25] and various challenging SM signatures. These models appear in the context of many top-down constructions [57], including string theory, of course [21], and are compatible with various potential resolutions to the hierachy problem such as supersymmetry [58], little Higgs models, TeV extra dimensions and Randall-Sundrum scenarios [1, 59–61]. In Neutral Naturalness scenarios, hidden valleys actually solve the little hierarchy problem directly [3]. Models containing a hidden valley have also been studied in the context of dark matter [22], matter-antimatter asymmetry [23] and the origin of neutrino masses [24].

In most hidden sector scenarios with a confining gauge force, the dark parton shower is qualitatively similar to QCD: the asymptotic freedom of the running gauge coupling enhances soft and collinear emission, resulting in the production of hidden sector states in collimated jets. If the hidden 't Hooft coupling is large ($\lambda \equiv g^2 N_c \gg 1$) and approximately constant over a significant energy range, the distribution of the produced dark hadron momenta will be much more democratic than the hierarchical jet-like behaviour we see in QCD. This is the SUEP class of signals, characterized by relatively soft, isotropic emission of dark hadrons [36, 37]. It is worth keeping in mind that strongly coupled hidden sector dynamics is not the only scenario that can lead to SUEPs. Similar final states can be produced in many-step cascade decays in the hidden sector [62], or in theories with extra spatial dimensions (see e.g. [63–65]). We now describe the general features of production, evolution and decay that constitute the SUEP signal, and define a benchmark scenario for our study.

2.1 Dark hadron production in exotic Higgs decays

Production of SUEP can occur through various portals coupling SM particles to SM-singlet states charged under a dark gauge group. Higgs and vector boson portals are commonly

JHEP12(2021)129

studied examples. Alternatively, a new particle charged under both SM and the dark gauge group could be produced [20]. We assume that hidden sector states ψ_D can be produced via the Higgs portal:

$$\mathcal{O}_{\text{production}} \sim |H|^2 \psi_D \bar{\psi}_D \ . \tag{2.1}$$

ψ_D could represent a fermion or scalar dark quark charged under the hidden gauge group. In the fermion case, the above operator is actually dimension 5 with a coupling $\sim 1/\Lambda$ for some UV scale Λ. This operator will give rise to the production of hidden states in exotic Higgs decays, as long as the dark hadrons are kinematically accessible. The dark quarks hadronize into a large multiplicity of dark hadrons.

For a confining gauge theory, the evolution of the system from the hard scale Q at which the first partons charged under the gauge group are generated to the IR confinement scale Λ governs the hadron multiplicity generated during showering. Q is of the order m_h in our case. For QCD jets, this evolution can be reliably simulated as a parton shower, but for SUEPs different strategies are required.

The average hadron multiplicity is given by

$$\langle n(Q) \rangle = \int_0^1 F(x, Q) dx \tag{2.2}$$

where F is the fragmentation function describing the final state distribution of momenta [66], and x is the momentum fraction of a given splitting. Calculating its evolution from the scale Q down to an IR scale Λ involves resumming the divergent contributions from the anomalous dimensions, with the leading contribution obtained from just the first ($j = 1$) Mellin transform of the anomalous dimension. If the theory is conformal between Q and Λ, the running of the coupling can (by definition) be neglected. In the limit where the 't Hooft coupling is large $\lambda \gg (1 - x)$ it was first shown in ref. [67] that

$$\langle n \rangle \sim \left(\frac{Q}{\Lambda} \right)^{2\gamma_T(1)} \ , \tag{2.3}$$

where $\gamma_T(1)$ is the first Mellin-Moment of the time-like anomalous dimension of the fragmentation function. Further expanding to zeroth order in the coupling yields

$$\langle n \rangle \sim \left(\frac{Q}{\Lambda} \right)^{1+\mathcal{O}(1/\sqrt{\lambda})} \tag{2.4}$$

Note that the small momentum fraction x carried by each individual splitting follows from $\lambda \gg (1 - x)$. In this strong regime branching is expected to yield emissions with $x \sim \Lambda/Q$ that are relatively isotropic in direction and democratic in momentum [37, 67]. Thus, with a large enough scale separation $Q \gg \Lambda$, low-x, high-multiplicity final states are generated. Branching terminates after $N_{\text{final}} \sim \log\langle n \rangle$ splittings at $Q_{N_{\text{final}}} \sim Q/2^{N_{\text{final}}} \sim \Lambda$, at which point hadronization takes over.

Hadron production in QCD at high temperatures is a close analogue to our situation, and statistical models have consistently shown that hadron multiplicities follow a thermal

JHEP12(2021)129

distribution [67–77]. We use this picture as a toy model of dark hadron production in SUEP, modelling the distribution of dark meson momenta as a relativistic Boltzmann distribution

$$\frac{dN}{d^3\mathbf{p}} \sim \exp\left(-\frac{\sqrt{\mathbf{p}^2 + m_D^2}}{T_D}\right) , \tag{2.5}$$

where m_D is the mass of the final dark states and T_D acts as the Hagedorn temperature of the hidden confining gauge force, with $T_D \sim \Lambda$ [37, 71]. This temperature controls the kinetic energy of the dark hadron distribution.

2.2 Dark hadron decay

Generically the decay of the dark hadrons ϕ_D into the SM will occur through some effective coupling of the form

$$\mathcal{O}_{\text{decay}} = \phi_D \mathcal{O}_{\text{SM}} \tag{2.6}$$

where $\mathcal{O}_{\text{SM}}$ contains fields charged under the SM gauge group. The phenomenology of the SUEP signal is determined to a large extent by the dominant decay portal. For example, if the dark hadrons decay to massive dark photons that mix with the SM photon, then the SM final state contains both hadrons and leptons in roughly gauge-ordered proportions, though for various dark photon masses in the GeV-regime, hadrons can dominate [37]. Alternatively, if the dark hadrons decay through the Higgs portal, the final state will contain hadrons and leptons in roughly Yukawa-ordered proportions. For sufficiently small portal coupling, the decay length of the dark hadrons may also be macroscopic, resulting in LLP signatures [8].

It is also possible for the hidden sector states to decay purely hadronically, which is the most experimentally challenging case at the LHC. A very simple example is the gluon portal as described in [20]:

$$\mathcal{L} \supset -\frac{1}{2} m_a^2 a^2 - \frac{\alpha_2}{8\pi} \frac{1}{f_a} a G_{\mu\nu} \tilde{G}^{\mu\nu} - i y_{\psi_D} a \psi_D \psi_D^* \, . \tag{2.7}$$

Here, a is a heavy elementary pseudo-scalar in the dark sector and ψ_D is the dark quark. Dark hadrons, which are bound states of ψ_D, could then decay to SM hadrons via an effective operator $\phi_D G\tilde{G}$. Another example is the hadrophilic (or leptophobic) Z' portal [78], where a new heavy gauge boson couples to SM quarks but not leptons, allowing dark hadrons to decay via an effective operator like $\phi_D q\bar{q}$.

2.3 Prompt hadronic benchmark scenario

The production of hidden sector states via the Higgs portal generally and exotic Higgs decays in particular is one of the most motivated and plausible discovery scenarios for new physics [79]. It is therefore vital that our experimental search strategies cover all possibilities for a signal at the LHC. This is especially urgent since for final states that are not covered by existing searches, branching fractions of ∼10% are easily allowed by current measurements of Higgs couplings and invisible decays [80–84].

JHEP12(2021)129

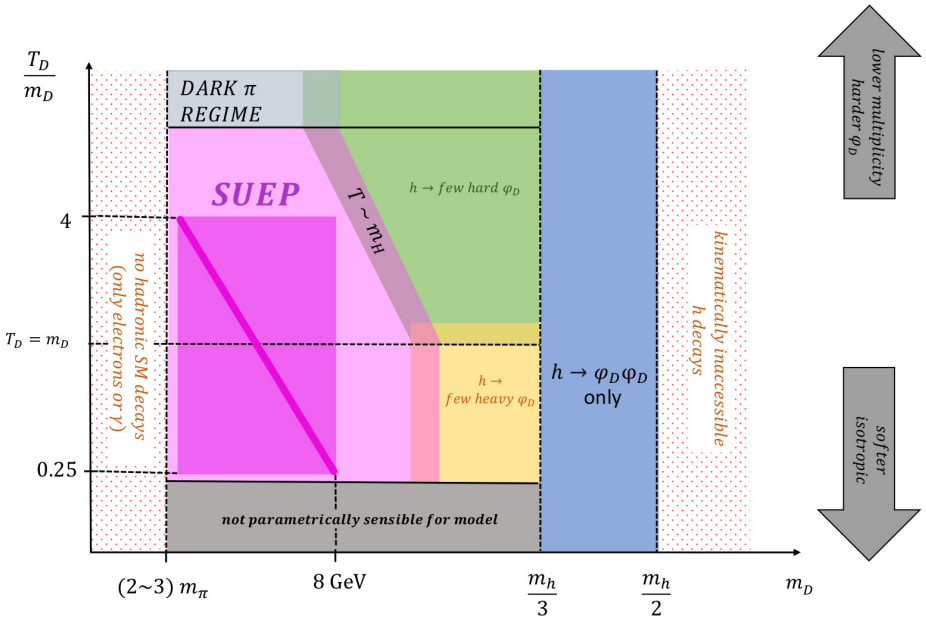


Figure 1. Cartoon of the our benchmark scenario for SUEP produced in Higgs decays with prompt decay of dark hadrons into SM hadrons. The SUEP description applies in the purple area, for T within a factor of a few of the dark hadron mass. In the green region, dark hadron production is not thermal, but described by processes more analogous to chiral QCD. In the orange and green regions, searches for Higgs decays to a few resonances would be sensitive to this dark sector. The parameter space region indicated with the darker purple rectangle is the focus of our analysis. Our cuts are optimized for the most archetypically SUEP-like final states, schematically indicated by the lower-left corner of this rectangle, demarcated with the diagonal line.

We therefore focus on the experimental worst-case scenario for SUEP produced in exotic Higgs decays: purely hadronic and prompt decays, with a particular interest in low dark hadron masses that make resonance searches [85, 86] or applications of jet substructure techniques [87] challenging. While the simplest gluon portal scenarios suggest that dark hadrons lighter than ∼10 GeV have macroscopic decay lengths [20] (which could allow for the use of long-lived particle search techniques [8, 33]), other possibilities can easily realize prompt, purely hadronic decays over a much wider range of dark hadron masses. For example, the hadrophilic (leptophobic) vector portal [78] with a hypothetical confining sector where the lightest dark hadron is a dark-rho-like vector ρ_D would allow dark hadrons lighter than a GeV to decay promptly into SM hadrons. In focusing on the prompt case we can develop techniques that allow SUEP production to be identified using only the geometrical and momentum distribution of its SM final states. These techniques will enhance our sensitivity for any kind of general SUEP, in addition to whatever other features of the final state, like displaced decays, leptons, or photons, can also be exploited.

The parameter space of our benchmark scenario is shown schematically in figure 1. The Higgs production portal sets the high scale for the event to $Q = 125$ GeV. With this scale fixed, the distribution of final states is in principle determined by the dark hadron mass m_D and the dark Hagedorn temperature T_D of eq. (2.5), shown on the vertical and horizontal

JHEP12(2021)129

axes of figure 1. In reality, there may be different dark hadrons with different spins, decays, and distributions, but this simplified description is sufficient for our purposes. On the left and right, the relevant parameter space is kinematically bounded by the red-hatched areas. Dark hadron production in Higgs decays requires $m_D < m_h/2$. As we explain below, we focus on dark hadrons that decay to SM hadrons, which requires $m_D > (2-3)m_\pi$, depending on the exact decay portal.

Even within that mass range, hidden sectors with pseudo-conformal dynamics do not always manifest as SUEP signatures in exotic Higgs decays. If $m_D > m_h/3$ (blue area in figure 1), then dark hadrons are only pair-produced in Higgs decays, making this scenario equivalent to standard exotic Higgs decays to pairs of various new particles (see e.g. [79]). If the dark hadron mass is fairly large, $m_D \sim m_h/(\text{few})$ (yellow area), or the dark Hagedorn temperature is comparable to or above the Higgs mass $T \gtrsim m_h$ (green area), then exotic Higgs decay would produce only a small multiplicity of dark hadrons that are either fairly hard or fairly heavy. In both cases, dark hadron production is not thermal but is described by processes more akin to chiral QCD. These regions are likely accessible through modified searches for resonances in exotic Higgs decays [85, 86], and we do not focus on them here. The gray area where $T/m \ll 1$ is not expected to be realized by any pseudo-conformal hidden sector, since the dark hadron mass and temperature are both related to the strong coupling scale Λ. On the other hand, $T/m \gg 1$ is possible in what we call the "dark pion regime", where the dark hadrons are pseudo-Goldstone bosons of an approximate symmetry, meaning their mass can be much smaller than Λ. We do not focus on this region, but it would be an interesting target for future investigations.

This leaves us with the actual SUEP regime for dark hadron production in exotic Higgs decays, indicated by the light purple area. In this investigation, we will focus on dark hadron masses below 8 GeV and T/m in the reasonable range of ~ 0.25 to 4. This target SUEP parameter space is marked out as the darker purple rectangle. Our cuts will be particularly optimized for the lower-left region of the rectangle, demarcated with the thick diagonal line. This is the region of low dark hadron mass and/or temperature, corresponding to the softest SM final states that are most difficult to search for using existing techniques.

3 Simulation

We briefly outline how we simulate our SUEP signal and the most important QCD backgrounds, where the latter is necessary to develop our analysis techniques, even though a realistic experimental analysis would rely on data-driven background estimation.

3.1 Signal

We generate event samples for exotic Higgs decay into SUEP using the `SUEP_Generator` plugin [20] for `Pythia 8.243` [88], which models the dark shower as a spherical distribution of dark pseudo-scalar mesons with momenta drawn from the relativistic Boltzmann distribution eq. (2.5). As in ref. [37], we make the simplifying assumption that there is only one flavor of dark meson produced in the exotic Higgs decay. The free parameters of

JHEP12(2021)129

the SUEP shower are the dark hadron mass m_D and the effective temperature T_D, setting the energy scale at which dark hadrons are produced dominantly.

We simulate associated Higgs production at the 14 TeV HL-LHC, $pp \to Vh, V \to \ell\ell/\ell\nu$, in `Pythia 8`. The Higgs is decayed to the SUEP final state of dark mesons, which then decay directly to a $u\bar{u}$ quark pair that in turn undergoes SM hadronization. The exact choice of hadronic decay mode does not significantly affect our analysis, so we use this single channel as a stand-in for other purely hadronic portals. The events are then passed through the simplified detector simulation code `Delphes 3` [89] with CMS detector settings. The simulated detector-level objects output by Delphes are used for our analysis. Our SUEP search will only use charged-track information. Since charged tracks can be traced back to the primary vertex, they are very robust with respect to pile-up contamination. We can therefore neglect the effects of pile-up in the remainder of our study.

In our event samples we cover m_D from 400 MeV to 8 GeV, and T_D/m_D from 0.25 to 4. The lower bound on the dark hadron mass ensures that the dark mesons are kinematically allowed to decay to two pions. The upper bound is chosen to show where our search loses sensitivity. The range of temperatures is chosen to satisfy $T_D \sim m_D$, which is the regime where the thermal picture of SUEP production is valid. The signal cross section is [90]

$$\begin{aligned} \sigma(pp \to e\nu/\mu\nu + \text{ SUEP}) &= (0.34 \text{ pb}) \cdot \text{Br}(h \to \text{ SUEP}) \\ \sigma(pp \to ee/\mu\mu + \text{ SUEP}) &= (0.066 \text{ pb}) \cdot \text{Br}(h \to \text{ SUEP}) \end{aligned} \tag{3.1}$$

for SUEP production in exotic Higgs decays in association with leptonic W/Z-bosons that decay into electrons or muons. In total, we generate 2.0×10^5 Zh and Wh signal events, proportional to their respective cross section, for each set of signal parameters $(m_D, T_D/m_D)$.

3.2 Background

The dominant background to our signal is production of one or two leptons in association with any number of QCD jets. It is highly challenging to model reliably, and in a realistic study, data-driven background estimation would be employed, see section 4.4. However, for the purpose of developing our analysis techniques, we simulate QCD+leptons background samples using `MadGraph5_aMC@NLO 2.6.6` and `Pythia 8.243`.

Ideally, one should simulate fully matched multi-jet + $\ell\ell/\ell\nu$ samples to capture the background distribution as closely as possible. However, due to the large statistics needed for our analysis, and the fact that such a simulation is anyway unlikely to be a perfect representation of the detailed hadronic distributions at the relevant high multiplicities and relatively low energy scale of the Higgs mass, this approach is not practical. Instead, we simulate $nj + \ell\ell/\ell\nu$, where $n = 2, 3, 4$ without jet matching and $p_T > 15$ GeV at generator level, to determine the effect of jet multiplicity at the hard event level on our analysis. We find that $n > 2$ leads to lower cross section while being *more distinguishable* from the SUEP signal using the analysis techniques we develop here. Therefore, to be conservative, we simulate $2j + \ell\ell/\ell\nu$ as our background samples for Zh and Wh production and decay into SUEP, respectively. In total, we use 10^8 background events to represent the

$$\sigma(\text{QCD } + \ \ell\ell/\ell\nu) \approx 3.7 \times 10^3 \text{ pb} \tag{3.2}$$

JHEP12(2021)129

lowest-order MadGraph5 cross section for this background. While this is sufficient to develop our analysis techniques, the Monte Carlo background sample has $\sim 1/100$ the statistics of the full HL-LHC dataset. This is important for the interpretation of our results in section 5.

4 Analysis

The goal of this paper is to devise strategies for extracting SUEP signals from a large background events without relying on the details of the simulated signal. We first describe our trigger assumptions and baseline cuts, and define a cut-based classifier to establish how sensitive such a simple approach can be. section 4.2 introduces a supervised neural network classifier, to demonstrate both the advantages and limitations of the supervised approach for our physics problem. We then introduce our primary tool in section 4.3 — an unsupervised neural network that we employ as an anomaly detector for SUEP.

4.1 SUEP observables

We define our trigger pre-selection by requiring that all events have at least one charged electron or muon with $p_T \geq 40\,\text{GeV}$, or two opposite sign charged leptons with $p_T \geq 30(20)\,\text{GeV}$. We also require that the scalar p_T-sum of hadronic charged tracks from Delphes is above 30 GeV. Both signal and background have a trigger efficiency of ≈40%, relative to the cross sections in eq. (3.1) and eq. (3.2).

We focus on the most challenging region of SUEP parameter space, with either low dark hadron masses or low dark shower temperatures. This gives rise to the most archetypically SUEP-like final states with a high multiplicity of isotropically distributed, relatively soft SM hadrons. The three observables that best capture the characteristics of this signal are the charged particle multiplicity N_{charged}, the event isotropy $\mathcal{I}$ [38], and the interparticle distance. In all steps of the analysis that follow, we only use charged particle tracks with $p_T \geq 300\,\text{MeV}$ from the primary vertex, excluding the one or two hard leptons associated with the decaying gauge boson.

The event isotropy observable $\mathcal{I} \in (0, 1)$ [38] quantifies the energy mover's distance between a collider event and an idealized isotropic event with uniform energy distribution, so $\mathcal{I} = 0$ indicates a fully isotropic event. Originally, three different definitions of the event isotropy are laid out, utilizing different geometries — spherical, cylindrical, or in a two-dimensional ring. We compute the *ring isotropy* of the set of charged hadronic tracks of each event, since at a pp collider we have no way to know the longitudinal boost of the Higgs that decays to SUEP. Since the SUEP is isotropic in the Higgs rest frame, we boost the hadronic charged track system of each event into its transverse rest-frame before computing the ring isotropy. To do this we assume that all hadronic charged tracks belong to particles with the pion mass, but this is sufficient to significantly separate signal and background events.

The variable that we introduce for the specific purpose of studying SUEP events is the interparticle distance matrix ΔR_{ij} for charged hadron tracks in the lab frame. It captures the unique topology of SUEP events while being very suitable for machine-learning

JHEP12(2021)129

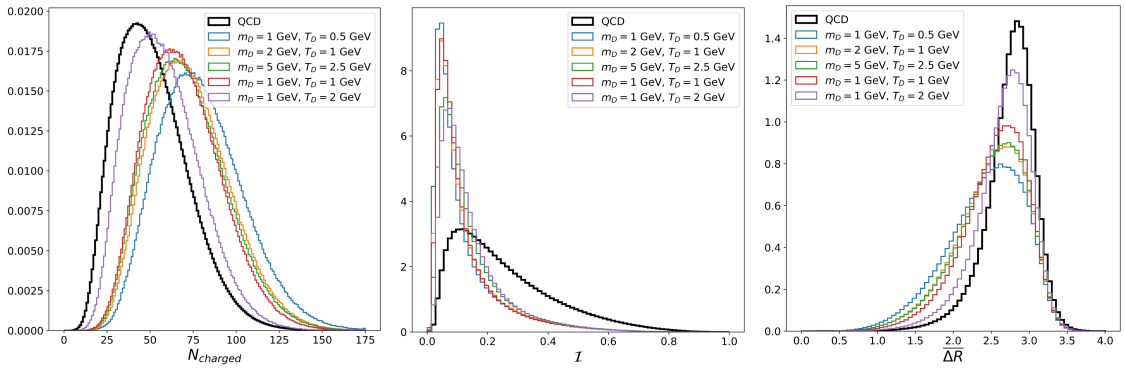


Figure 2. Comparison of QCD and SUEP distributions of selected observables.

applications, since it is invariant under re-definitions of the azimuthal angle around the beam axis. It is also useful to define the mean $\overline{\Delta R}$ of all matrix entries.

Figure 2 shows distributions of N_{charged}, $\mathcal{I}$, and $\overline{\Delta R}$ for the QCD background and a variety of SUEP benchmark points after trigger selection. The separation between signal and background is clear, with SUEP having higher multiplicity, more isotropic distribution of tracks, and a significantly wider spread of inter-particle distances. To understand how these observables change across SUEP parameter space, we show their average values (across the whole sample after trigger selection) as a function of m_D and T_D in figure 3. The pairwise correlations between each of these observables are included in the appendix in figure 8, demonstrating that each of these three variables encode distinct information about each event.

In addition to the trigger requirements, we therefore impose the following pre-selection cuts on all events:

$$N_{\text{charged}} \geq 70 \quad , \quad \mathcal{I} < 0.07 \quad , \quad \overline{\Delta R} < 3 \, . \tag{4.1}$$

This cut targets the most SUEP-like parts of signal parameter space with low dark Hagedorn temperature and/or low dark hadron mass (see figure 1). All but 2.2% of the post-trigger background is eliminated, while leaving 31.8% of signal for $m_D = 0.4\,\text{GeV}$ and $T_D = 0.4\,\text{GeV}$. These requirements are less optimal for larger dark hadron masses or temperatures — for example, the signal efficiency is only 1.1% for $m_D = 5\,\text{GeV}$, $T_D = 20\,\text{GeV}$. However, larger temperatures and masses generally lead to higher-energy final states or separable resonances, and are therefore not the focus of our present analysis.

We now consider three options for SUEP searches at the HL-LHC, using events which pass the baseline pre-selection as a starting point:

1. The simplest strategy is a cut-and-count analysis using high-level observables. It will serve as a baseline for more sophisticated machine-learning techniques, and can be implemented very easily and effectively with a stricter cut on $\overline{\Delta R}$ compared to eq. (4.1). Varying the cut threshold yields a significance improvement curve (SIC) of the signal efficiency vs the background efficiency for each point in signal parameter space, which we will then be able to compare to results using machine learning methods. As we will see, this already yields very promising SUEP sensitivities.

JHEP12(2021)129

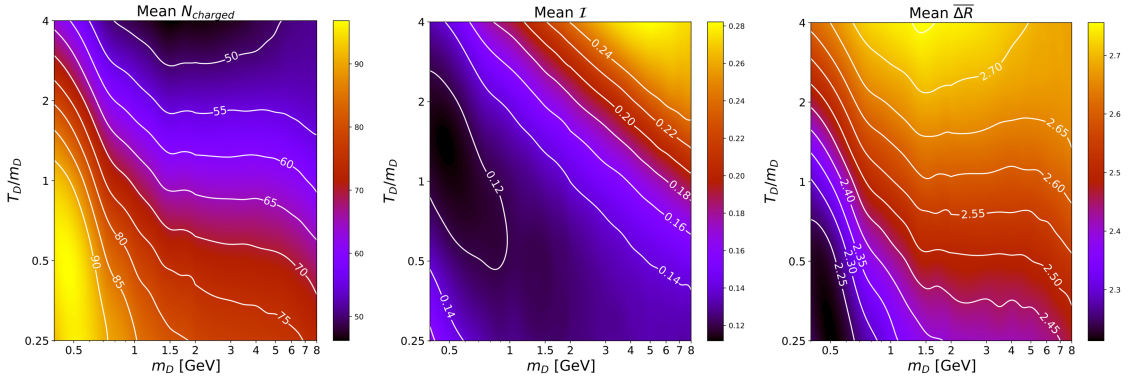


Figure 3. Average values of charged particle multiplicity, event isotropy and mean-interparticle-distance as a function of m_D and T_D for SUEP. QCD average values: $\langle N_{\text{charged}} \rangle = 51$, $\langle \mathcal{I} \rangle = 0.25$, $\langle \overline{\Delta R} \rangle = 2.7$.

2. A supervised ML-classifier requires detailed knowledge of the signal, since it is trained on signal and background. This makes supervised techniques unlikely to be a realistic analysis candidate for broad SUEP searches. However, we perform a simple supervised study in section 4.2 to demonstrate the best-case scenario for SUEP sensitivity if the signal was very well understood.

3. In section 4.3, we use an unsupervised autoencoder trained only on the background as an anomaly detector to improve on the sensitivity of the cut-and-count analysis. This is likely to be a realistic analysis candidate since it can be performed using data-driven background estimation techniques without precise knowledge of the signal.

4.2 Supervised ML-classifier

While a supervised classifier is explicitly dependent on the characteristics of the signal (and background) simulation used in its training, which we cannot trust in detail, it can still provide a useful comparison to an unsupervised network, and give an indication of how sensitive a search based on similar methods could be with improved modelling of the SUEP signal. An additional limitation is that even with reliable signal simulation, the parameters of the real SUEP are unknown, and a supervised network trained to recognize SUEP with one set of parameters may fail if the true parameters change. As we will see, this is indeed the case.

The supervised network architecture we choose is a dynamical graph convolutional neural network [91, 92]. We implement the network in PyTorch. The input feature representation for both the supervised and unsupervised networks is the interparticle distance matrix, with the redundant left-lower half set to zero and track p_T information added to the diagonal:

$$\Delta\tilde{R}_{ij} \equiv \begin{cases} \Delta R_{ij} & i > j \\ p_{T,i}/\text{GeV} \quad \text{for} & i = j \\ 0 & i < j \end{cases} \tag{4.2}$$

JHEP12(2021)129

All events are required to have at least $N_{\text{charged}} \geq 70$ tracks, and all events are truncated to keep only the 70 highest-p_T tracks to ensure each event has the same dimensionality in the analysis.

The input matrix ΔR_{ij} is used to generate graph edges between each particle (node) and its $k = 7$ nearest neighbours in ΔR space. The node features for each particle are the 70-dimensional vector of ΔR distances to all other particles in the event. The graph network has two EdgeConv blocks, each comprising a three-layer perceptron with leaky ReLU activation [93]. The EdgeConv operation updates the node features x_i of each particle as

$$\vec{x}_i^{\,\prime} = \frac{1}{k} \sum_{j=1}^{k} \vec{h}_\theta(\vec{x}_i, \vec{x}_i - \vec{x}_j) \ , \tag{4.3}$$

where k is the number of neighbours assigned to each node, and h_θ is a non-linear function of learned parameters θ, implemented as a three-layer perceptron. The graph edges are re-computed between the first and second blocks using the Euclidean distance between the feature vectors of each node to determine its 14 nearest neighbours. The first block has feature dimension 64, while the second block has feature dimension 32. Batch normalization follows each layer. The output of the graph layers is averaged, then passed through two fully connected layers, first expanding to dimension 128, then down to output dimension 2. The loss function is the cross-entropy loss between the output and the true class label of each event.

$$\mathcal{L}(x, \text{class}) = -\log \frac{x[\text{class}]}{\sum_j x[j]} \tag{4.4}$$

To test how the performance of the model depends on the choice of training parameters, we train twelve neural networks on twelve different choices of $(m_D, T_D/m_D)$, with $m_D = 0.5, 1, 2$ GeV and $T_D/m_D = 0.5, 1, 2, 4$, and evaluate their effectiveness over the whole signal parameter space. We also evaluate the efficacy of a 'cocktail approach' [92, 94–96] by training on a mixed sample including signal events from each of the twelve parameter choices. A conditional training [97, 98] on the signal parameters would be a similar approach. Each network is trained for 10 epochs with a decaying learning rate. Longer training periods were tested and found to be unnecessary. Loss values for each test sample event are obtained for the model realizations of the last 5 training epochs before being averaged.

4.3 Unsupervised autoencoder

Supervised ML-classification is an extremely powerful analysis tool, but it only works for signals with well-defined and universal features and corresponding precision simulations. At the LHC, this is not always the case, and SUEP with its toy shower is a perfect example for a more broadly defined signal. Here we prefer not to train a classifier on signal simulations. Instead, we employ anomaly detection methods, where we train a network only on the well-understood background dataset, so that it can flag events that are anomalous in comparison. We use an autoencoder [39, 45, 46] and train it on data without class labels, with a loss function that incentivizes its output to be as close as possible to the input. The intermediate network layers have a restricted number of nodes compared to the

JHEP12(2021)129

dimension of the input and output, forcing the network to compress the information in the input, and then decompress it to recover the output. The principle of the autoencoder's use as an anomaly detector is that it should fail to accurately reconstruct events that are anomalous compared to the dataset it was trained on. A high reconstruction loss flags an event as being potentially anomalous, or in collider physics parlance, a signal (SUEP) candidate event.

The autoencoder uses the same modified $\Delta\tilde{R}_{ij}$ matrix input as used by the supervised network, see eq. (4.2). Other representations were tested, including the matrices of both k_T and anti-k_T distances $d_{ij} = \min(k_{ti}^{\pm 2}, k_{tj}^{\pm 2})(\Delta\tilde{R}_{ij})^2/R$ between particles [99], the high-level observables used in the pre-selection cuts, as well as the raw p_T, ϕ, η values for each particle; all yielded results that were much less useful than the analysis we present here. This emphasizes the importance of choosing the correct input representation over sophisticated network architecture for SUEP searches.

The matrix $\Delta\tilde{R}_{ij}$ is flattened into a vector of length N^2_{charged} and fed into the autoencoder. The neural network comprises five fully connected layers, with the number of nodes decreasing to the bottleneck size in the third layer, then increasing back to $N^2_{\text{charged}} = 4900$. An alternative 3-layer network performs only slightly worse. For the bottleneck size we find that larger bottlenecks consistently lead to better performance than smaller ones, so we use $N_{\text{bottleneck}} = 1000$. Each layer of the network has a leaky ReLU activation with slope -0.2 for negative values of x, except the final layer which has a ReLU activation. The loss function of the network measures the difference between the network's input and output as

$$\mathcal{L}(y^{\text{in}}, y^{\text{out}}) = \frac{1}{N_{\text{charged}}} \sum_i \frac{|y_i^{\text{out}} - \sigma(y_i^{\text{in}})|^m}{|\sigma(y_i^{\text{in}})|^n} \tag{4.5}$$

where $\sigma(x) = 1/(1 + e^{-x})$. Among different values of m and n, starting with the usual mean squared error $m = 2, n = 0$, the best-performing is $m = 3, n = 0$. The sigmoid normalization of the input is essential to the network's success. Without it, the autoencoder encodes SUEP events with slightly *lower* loss than the QCD background on which it was trained, and completely fails to identify anomalous events. We hypothesize that this is because, unlike many experimental signatures, SUEP is less complex than its QCD background, and it has smaller values of ΔR than QCD. This complexity bias has been noted before [45, 52, 53]. The sigmoid function reduces sensitivity to large values of ΔR and p_T by mapping them to values very close to 1, while remaining approximately linear for small input values, but offset to a minimum value of 0.5. These effects make it easier to accurately reconstruct QCD events while enhancing the network's sensitivity to deviations from the input on the SUEP events with characteristically smaller absolute values of the input features.

Out of 8.8×10^5 background Monte Carlo events that pass the pre-selection cuts, 2.4×10^5 are used for training, 5×10^4 for validation when tuning network hyperparameters, and 5.9×10^5 for testing. The number of signal events that pass the cuts varies with m_D and T_D, but generally a few $\times 10^4$ events remain at each parameter point to be used for testing. SUEP events with $m_D = 1\,\text{GeV}$, $T_D = 0.5\,\text{GeV}$ are used for validation purposes.

JHEP12(2021)129

The network is trained for 15 epochs with a decaying learning rate. Longer training periods were tested and found to be unnecessary. Loss values for each test sample event are obtained for the model realizations of the last 5 training epochs before being averaged.

Other architectures were investigated, including variational autoencoders [100] and a graph convolutional autoencoder utilizing the same EdgeConv operations as the supervised network, which itself led to the use of the $\Delta\tilde{R}_{ij}$ event representation. Interestingly, the simpler, fully connected architecture consistently delivered much better background rejection than any of the more sophisticated graph networks in the unsupervised approach.

4.4 Data-driven background estimation

Following the logic of this section further, we briefly describe how a data-driven background sample could be derived for use in a realistic experimental analysis based on our study. The total background cross-section can be measured directly (and compared against simulation), since the leptons+QCD signal region is completely background dominated for SUEP production in exotic Higgs decays. The background efficiency of the classifier, whether it is based on cuts, an unsupervised neural network, or a supervised one, can be estimated by evaluating it on a control region.

A variety of control regions are possible, but perhaps the most promising is defined by replacing the lepton criterion by a mono-photon criterion. Such a sample should be free of signal contamination, and its hadronic content is extremely similar to the hadronic background to our SUEP search, since it is unconstrained in its detailed production channel except that it recoils off a hard electroweak boson. The QCD distributions in this mono-photon-plus-jets sample should therefore closely match those in the Z+jets signal region, especially if the control region sample is reweighted to match the shape of the photon p_T spectrum to the dilepton p_T spectrum in the Z+jets signal sample. A variant of this approach can likely be adapted to estimate the background in the W+jets channel as well, by using simulation to compute a transfer function from $p_{T,W}$ to $p_{T,\ell}$ and applying it to $p_{T,\gamma}$ in the control region before applying the $p_{T,\ell}$ reweighing.

Based on the assumption that such a strategy can be implemented, we therefore will use benchmark estimates of 1% and 10% for the systematic background uncertainty in our analysis to estimate the final physics reach.

5 Results

Figure 4 shows significance improvement curves for the cut on $\overline{\Delta R}$, the autoencoder, and a selection of supervised networks trained on signal samples with a variety of different dark hadron masses and temperatures m_{train} and T_{train}. The horizontal axis shows either signal efficiency (after triggering) or number of remaining signal events at the HL-LHC as the threshold is varied. The vertical axis shows signal significance improvement relative to the trigger sample. From these specific examples, a few general features are apparent:

- The autoencoder consistently outperforms the simple $\overline{\Delta R}$ cut significantly;
- The supervised networks outperform the autoencoder for signal parameters close to their training values, but can perform much worse for different parameters;

JHEP12(2021)129

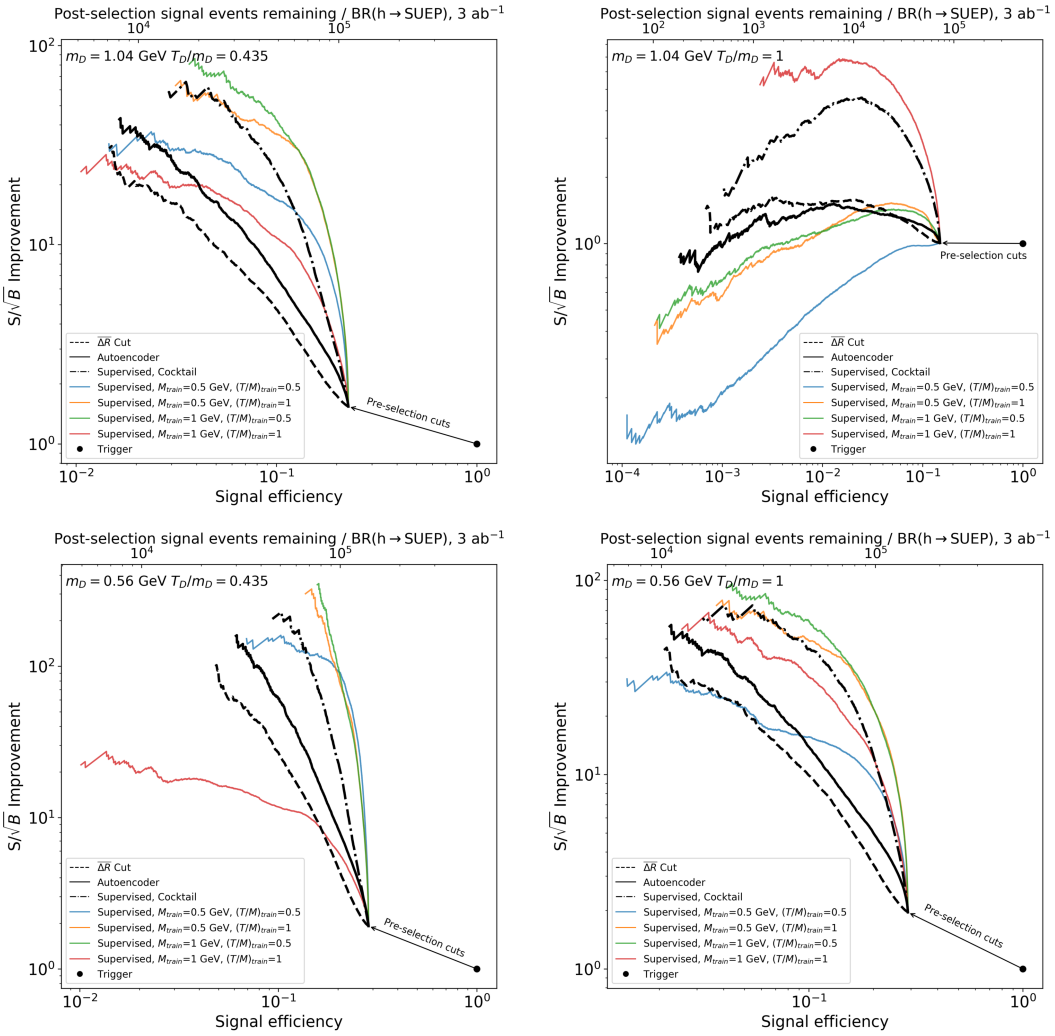


Figure 4. Some examples of Significance Improvement Curves (SIC), relative to trigger selection, of the autoencoder (black solid), cut on $\overline{\Delta R}$ (black dashed), and supervised graph networks (black dash-dotted and colored) for SUEP with test parameters of $(m_D, T_D/m_D) = (1.04, 0.435)$, (1.04, 1), (0.56, 0.435), (0.56, 1). For the autoencoder, the peak sensitivity improvement we are able to probe reliably with the statistics of our Monte Carlo samples (see text) corresponds to signal and background efficiencies of $(1.6 \times 10^{-2}, 1 \times 10^{-7})$, $(3.7 \times 10^{-4}, 1.5 \times 10^{-7})$, $(6.1 \times 10^{-2}, 1 \times 10^{-7})$ and $(2.2 \times 10^{-2}, 1 \times 10^{-7})$ relative to trigger selection, respectively.

- Our analysis is not optimized for larger dark hadron masses and temperatures;
- For lower dark hadron masses or temperatures, both the cut-and-count and autoencoder analysis strategies are very powerful, yielding orders of magnitude improvement in signal significance compared to the baseline preselection cuts of eq. (4.1).

In fact, the statistics of our QCD background sample is insufficient to capture the true power of our analysis techniques. We would expect the significance improvement to increase as the signal efficiency decreases, but only until the SI curve turns around when the cut

JHEP12(2021)129

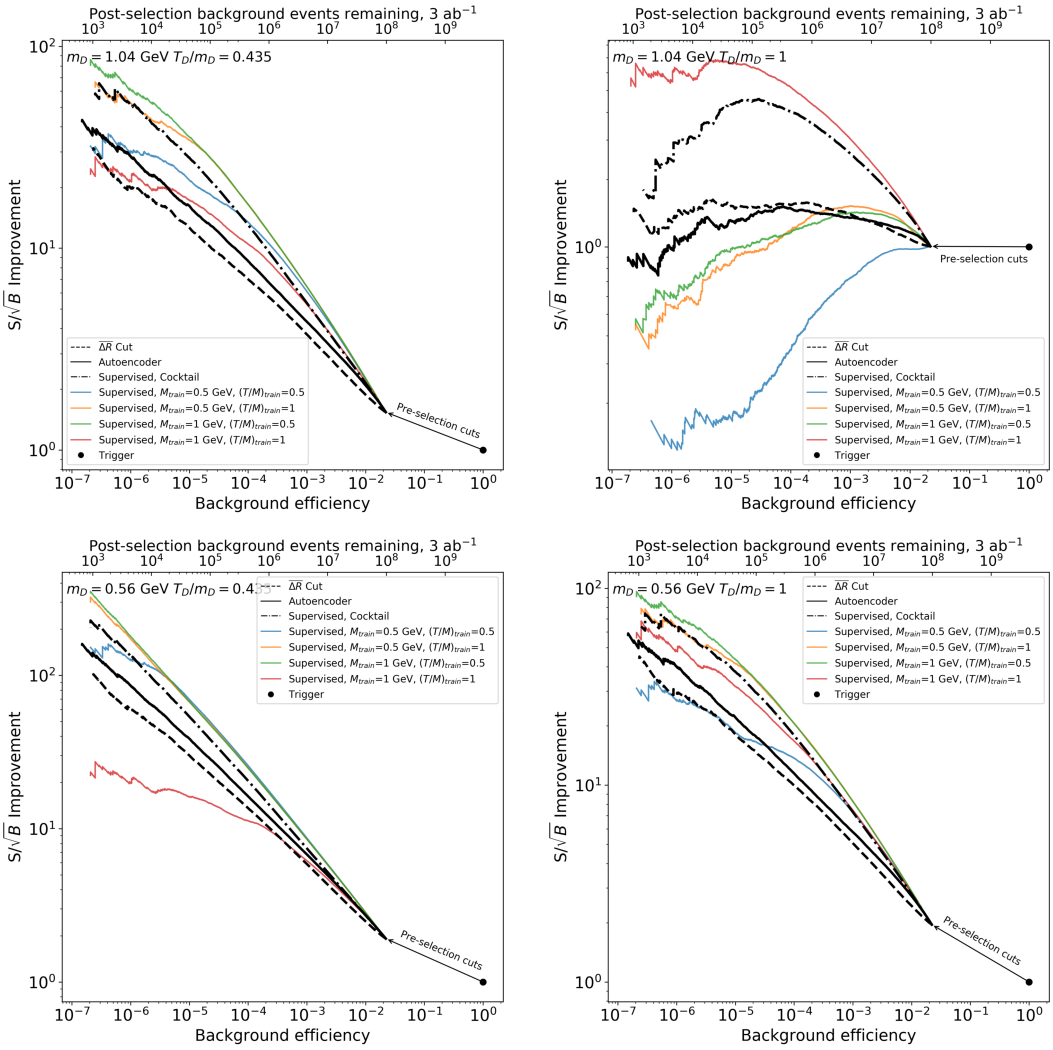


Figure 5. Same SIC curves as figure 5, but with background efficiency on the horizontal axis. This demonstrates that with the full HL-LHC dataset, significantly harder cuts could increase sensitivity beyond the limitations of what we can demonstrate with our simulated background dataset.

becomes too harsh. This is clearly the case in the top-right panel of figure 4, for signal parameters to which our analysis is not optimized. However, for the three other examples, we never reach this turn-over point before running out of simulated background events. It is not even clear if this turn-over is reached with the full statistics of the HL-LHC (100 times greater than our simulated background sample).

To understand how much harder we might be able to cut on the background, figure 5 shows the same SIC curves but with background efficiency on the vertical axis. Naively extrapolating, we can anticipate that a realistic autoencoder search with a fully data-driven background sample at the HL-LHC might be able to reach the very-low background regime while retaining enough signal to probe Br($h \to$ SUEP) as much as an order of magnitude smaller than the sensitivity estimates we can rigorously derive here.

JHEP12(2021)129

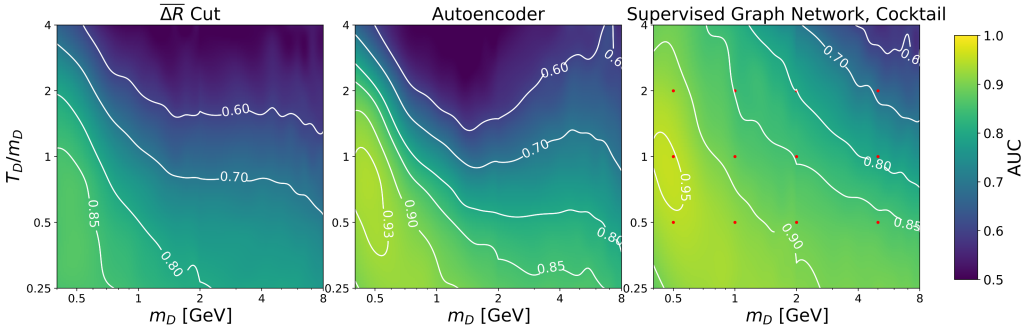


Figure 6. AUC for (a) cut on $\overline{\Delta R}$, (b) fully connected autoencoder, (c) supervised graph network trained using cocktail of signal parameter choices (training parameters indicated with red dots). These plots illustrate the significant performance improvements of the autoencoder relative to the simple $\overline{\Delta R}$ cut, and of the supervised cocktail approach relative to the autoencoder.

As a result, the reach projections we present in this paper will be very conservative. Furthermore, the limited statistics of our background sample means that reach projections will be very similar for the three analysis methods despite their obvious differences in performance in figures 4 and 5. It is therefore important to additionally evaluate our classifiers using a somewhat orthogonal metric.

The area under the curve (AUC) of the signal efficiency as a function of background rejection can be computed as a function of the number of events kept after the baseline preselection cuts, see eq. (4.1). This metric is standard to measure the performance classifiers independently of the choice of threshold. figure 6 shows the AUC achieved by the cut-based classifier, the fully connected autoencoder, and the supervised graph network trained on a cocktail of signal events, where for the latter the signal parameters are indicated with red dots in the $(m_D, T_D/m_D)$-plane). Unsurprisingly, the highest AUC values are achieved at low m_D or T_D/m_D, with the supervised graph network modestly outperforming the autoencoder, which outperforms the simple cut, across the SUEP parameter space. The performance of supervised networks trained on single signal parameter points are presented in the appendix.

A key physics result is the actual sensitivity of HL-LHC searches to the SUEP final state. We therefore extract the smallest Br($h \rightarrow$ SUEP) branching ratio for which $S/\sqrt{B + u_{sys}B^2} > 2$, where u_{sys} gives the systematic uncertainty on the background. Because of the high degree of background rejection required to be sensitive to relevant Higgs branching ratios, and the limited size of our background dataset, very few or no simulated background events remain when the classifier's threshold is set to maximize sensitivity to low branching ratios. This introduces a significant statistical uncertainty to the estimated LHC reach. We can decouple our estimate somewhat from these limitations by demanding the statistical uncertainty from the limited size of the background dataset to remain below 50%. This is conservative, since a realistic analysis will employ even harsher cuts with better significance improvement. For the same reason, the difference in reach between

JHEP12(2021)129

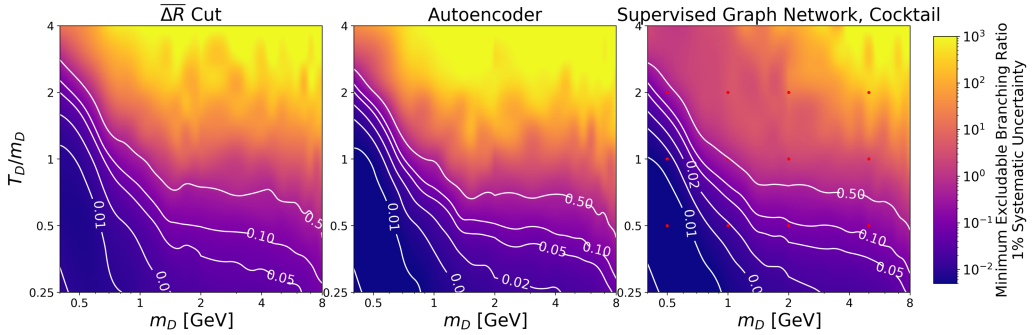


Figure 7. Minimum excludable Br($h \to$ SUEP) at the HL-LHC, assuming 1% systematic uncertainty on QCD background for (a) cut on $\overline{\Delta R}$, (b) fully connected autoencoder, (c) supervised graph network trained using cocktail of signal parameter choices (indicated with red dots). Note that the limited statistics of our QCD background sample leads these projections to be very conservative while also de-emphasizing the performance differences between the three methods.

our three methods is likely underestimated. The performance gaps between the $\overline{\Delta R}$ cut, autoencoder, and supervised network are more clearly shown by the individual significance improvement curves and the AUC differences.

Finally, figure 7 shows the SUEP sensitivity achievable by the $\overline{\Delta R}$ cut, the autoencoder, and the cocktail-trained supervised network, all assuming 1% systematic uncertainty on the background. Both the cut and the autoencoder can probe %-level branching ratios. In the appendix we show sensitivity projections assuming a much larger 10% systematic background uncertainty, with only very modest degradation in reach. This shows that the overwhelming QCD background has been reduced to low enough levels to make the search statistics limited, and speaks to the robustness of our results.

6 Conclusion

SUEPs represent a highly plausible but extremely challenging experimental signature of confining hidden sectors, which typically results in a high multiplicity of soft SM final states. To date, there are no targeted LHC searches for SUEP, and most existing searches have limited or no sensitivity. Furthermore, since SUEP is produced by hidden sectors featuring fairly strongly-coupled, approximately conformal dynamics with a wide variety of possible dark hadronization scenarios, modelling the detailed production of SUEP is fraught with uncertainties. Existing proposals for SUEP searches [8, 37] target conspicuous, qualitative features of the final state, such as displaced vertices or leptons. Their existence is a prediction for some decay portals of the dark hadrons, and targeting them with searches is fairly robust with respect to modelling uncertainties.

An important but previously unaddressed question is how well we can look for SUEP using only its basic kinematic features, without any conspicuous SM final states. This is not only an experimental challenge, the design of such a search also has to be very mindful of uncertainties in the signal modelling. We investigated this worst-case scenario by studying

JHEP12(2021)129

SUEPs from dark hadrons produced in exotic Higgs decays, which decay promptly and purely hadronically. Our choice of model is motivated by the Higgs portal being one of the most plausible production modes for new physics. The modest energy scale of the Higgs decay also eliminates high-energy or high-mass observables as discriminants. Finally, Higgs production in association with a W/Z circumvents the problem of triggering on the SUEP final state directly [37].

Our first results are observables which capture the essential SUEP features from the track momenta. We focused on the charged particle multiplicity N_{charged}, the event ring-isotropy [38] $\mathcal{I}$, and the interparticle distance ΔR_{ij}. All three are robust with respect to modelling uncertainties of the SUEP final state, since they capture the essential model features: high multiplicity of final states, and dark hadron production that is more isotropic and democratic in momentum than the QCD background. The charged track interparticle distance matrix ΔR_{ij} is particularly suitable for ML applications.

Based on these observables, we devised three strategies for $h \to$ SUEP searches. The first is based on a simple cut on the interparticle distance matrix. The second assumes that our naive signal simulation tools can be trusted, and uses supervised ML techniques. The third is an unsupervised ML approach using a fully connected autoencoder trained only on simulated QCD events as an anomaly detector. Both ML approaches used the slightly modified interparticle distance matrix of eq. (4.2) as the event representation. The cut-and-count approach serves as a simple baseline, over which the unsupervised ML approach represents a significant improvement, even without detailed knowledge of the signal. This is to be compared to the higher sensitivity of the supervised machine learning approach, which is unlikely to be robust with respect to modeling uncertainties or choice of training parameters.

All three approaches will probe exotic Higgs decays to prompt, hadronic SUEP at the HL-LHC for branching fractions at the percent level, with both the autoencoder and supervised approaches probing rates as low as 1%. We assumed a systematic background uncertainty of one percent, but increasing this to ten percent only modestly decreases sensitivity, signaling that our analyses have sufficient differentiation power to reduce the enormous QCD background to a statistics-dominated level. These estimates are conservative, since our simulated background samples were still too small to fully cover the exceedingly large background rejection. A realistic search combining the autoencoder with data-driven background estimates will achieve significantly higher sensitivities.

Our results show that even without a detailed theoretical description of the SUEP showering process, an analysis using an unsupervised neural network can be highly sensitive to exotic Higgs decays to SUEP. Our observables and techniques can equally well be used in SUEP searches using leptons, photons or displaced vertices, to significantly enhance their sensitivity based on the inherently SUEP-y kinematics. Generalizations of our methods, for instance including searches for explicit dark hadron resonances, should allow for sensitivity to SUEPs with higher dark hadron masses or dark Hagedorn temperatures than those targeted by our analysis. Our new, unsupervised search strategy can be applied to a wide range of LHC scenarios to discover new physics, even if the true BSM model should differ from our exact theory expectations.

JHEP12(2021)129

Acknowledgments

The authors would like to thank Anja Butter, Simon Knapen, Jessie Shelton, and Jennifer Thompson for helpful conversations. The research of JB and DC is supported in part by a Discovery Grant from the Natural Sciences and Engineering Research Council of Canada, and by the Canada Research Chair program. JB also acknowledges funding from a Postgraduate Doctoral Scholarship (PGS D) provided by the Natural Sciences and Engineering Research Council of Canada. GK acknowledges the support of the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy — EXC 2121 "Quantum Universe" — 390833306. The research of TP is supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under grant 396021762 — TRR 257 *Particle Physics Phenomenology after the Higgs Discovery*.

Computations were performed on the Niagara supercomputer at the SciNet HPC Consortium [101, 102]. SciNet is funded by: the Canada Foundation for Innovation; the Government of Ontario; Ontario Research Fund — Research Excellence; and the University of Toronto. This research was enabled in part by support provided by Compute Canada (www.computecanada.ca).

A Additional results

Figure 8 shows the two-dimensional correlations between each pair of the observables N_{charged}, $\mathcal{I}$, and $\overline{\Delta R}$. While the average values of N_{charged} and $\overline{\Delta R}$ vary in a similar fashion with m_D and T_D, their distributions for a fixed choice of signal parameters are not highly correlated.

Figure 9 shows the AUC for each of the twelve supervised graph networks trained on a single sample of signal events, with parameters indicated by the red dot in each plot. This vividly demonstrates how the parameter choice of the signal dataset on which the supervised network is trained has a large impact on the regime of parameter space where it can achieve high AUC, providing a strong argument for using the cocktail approach.

The sensitivity projection for these supervised networks with one percent background systematic is shown in figure 10. The peak sensitivity for each network is also at the 1%-level, similar to the other approaches, but for potentially a different range of SUEP parameters, depending greatly on the training parameters of each network.

Figures 11 and 12 show sensitivity projections for 10% systematic background uncertainty.

JHEP12(2021)129

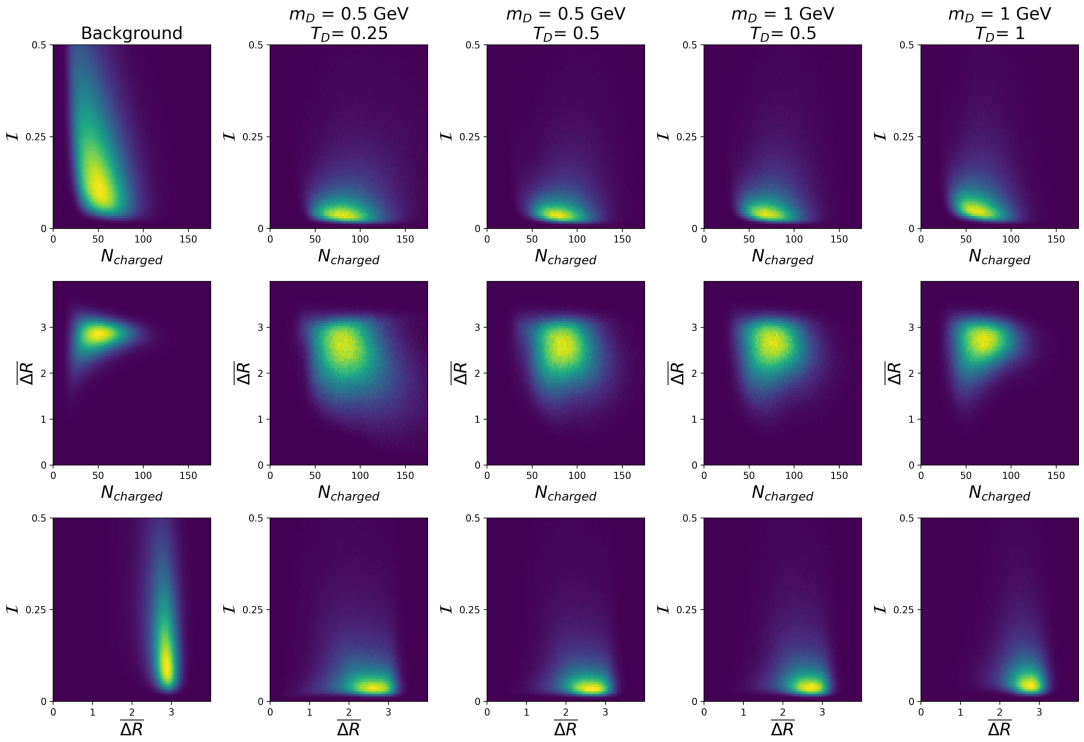


Figure 8. Two-dimensional unit-normalized histograms showing pairwise correlations between N_{charged}, $\mathcal{I}$, and $\overline{\Delta R}$ for the background sample and the same four choices of signal parameters as in figure 4.

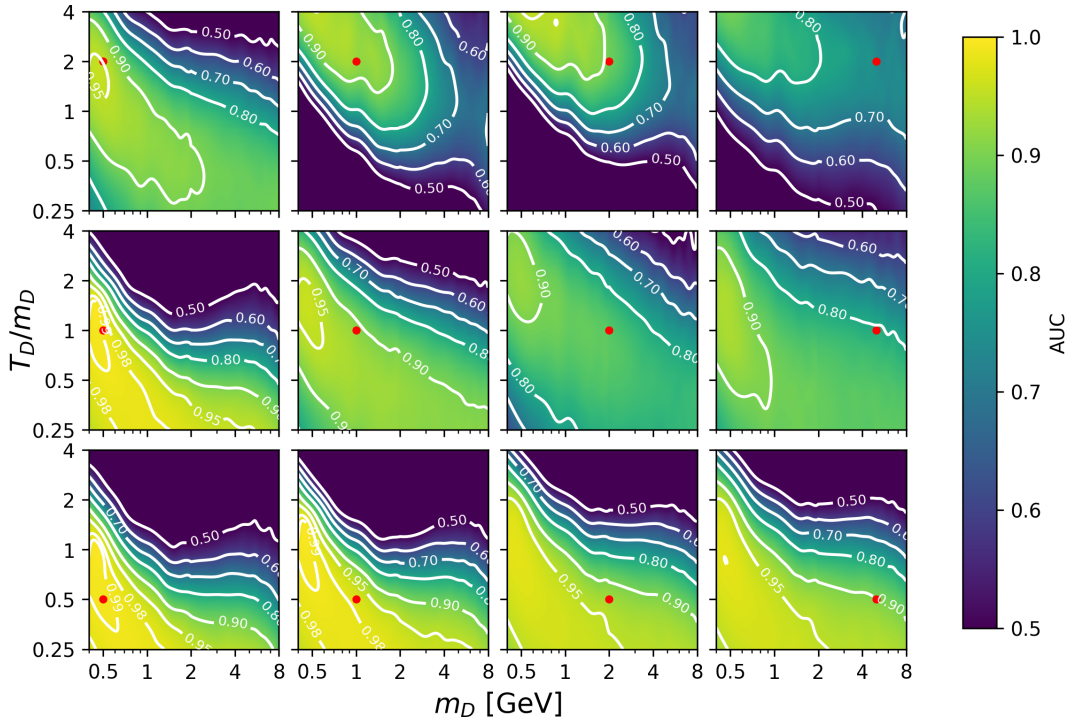


Figure 9. AUC for supervised graph networks trained using different signal parameter choices. Red dots indicate the training parameters in each plot.

JHEP12(2021)129

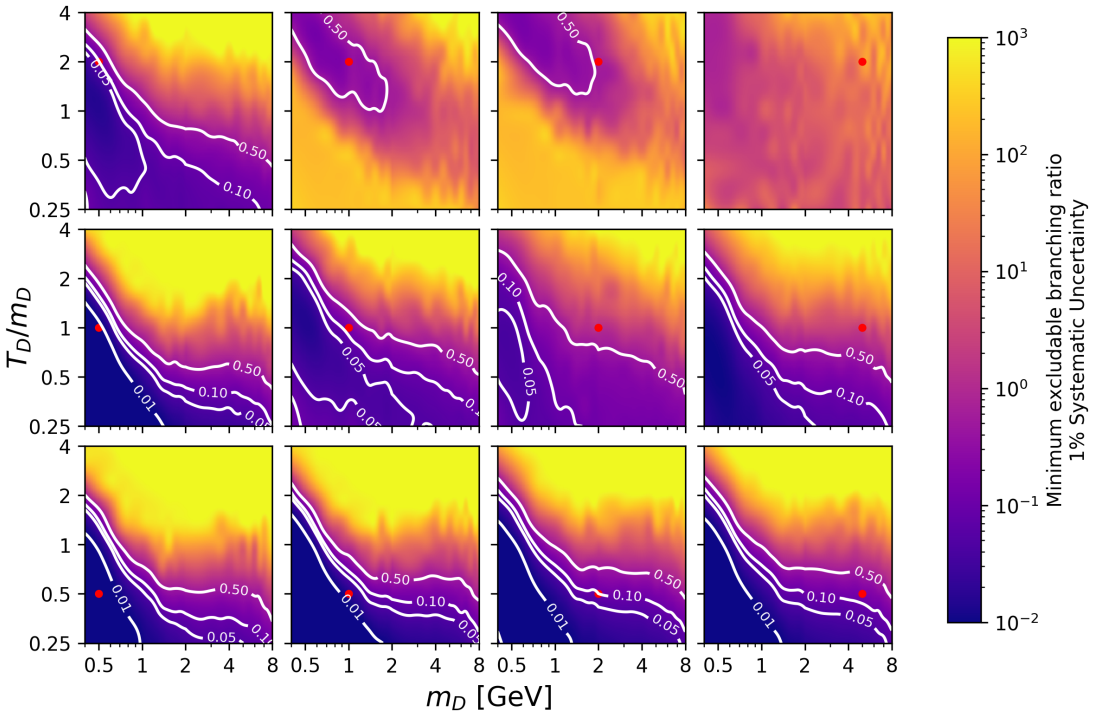
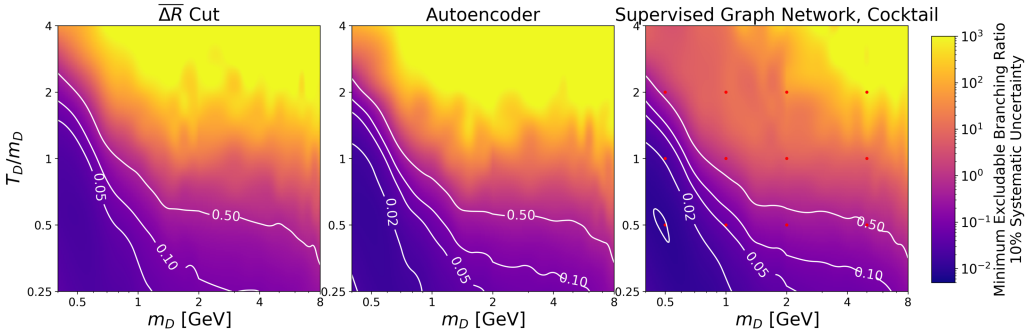


Figure 10. Minimum branching ratio excludable by supervised graph networks trained on different choices of (m_D,T_D). Red dots indicate the training parameters in each plot.



(a) (b) (c)

Figure 11. Same as figure 7 but for 10% systematic background uncertainty.

JHEP12(2021)129

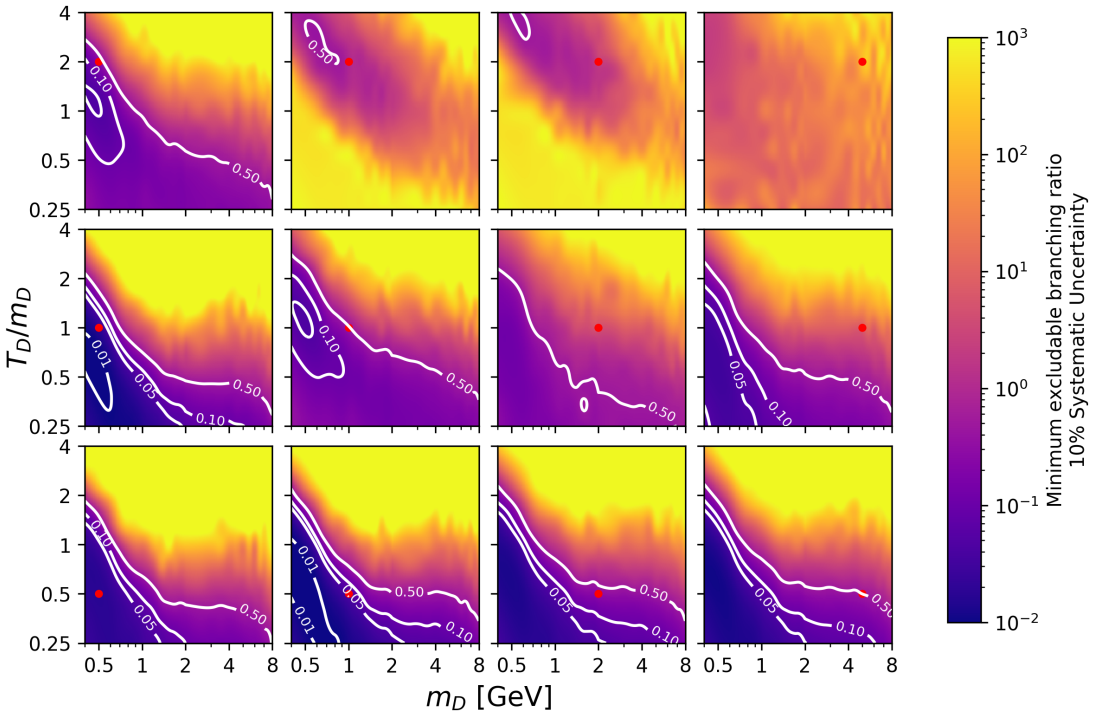


Figure 12. Same as figure 10 but for 10% systematic background uncertainty.

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

References

[1] 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].

[2] M.J. Strassler, *On the phenomenology of hidden valleys with heavy flavor*, arXiv:0806.2385 [INSPIRE].

[3] Z. Chacko, H.-S. Goh and R. Harnik, *The twin Higgs: natural electroweak breaking from mirror symmetry*, *Phys. Rev. Lett.* **96** (2006) 231802 [hep-ph/0506256] [INSPIRE].

[4] R.M. Schabinger and J.D. Wells, *A Minimal spontaneously broken hidden sector and its impact on Higgs boson physics at the large hadron collider*, *Phys. Rev. D* **72** (2005) 093007 [hep-ph/0509209] [INSPIRE].

[5] B. Patt and F. Wilczek, *Higgs-field portal into hidden sectors*, hep-ph/0605188 [INSPIRE].

[6] J.R. Espinosa and M. Quirós, *Novel effects in electroweak breaking from a hidden sector*, *Phys. Rev. D* **76** (2007) 076004 [hep-ph/0701145] [INSPIRE].

[7] J. March-Russell, S.M. West, D. Cumberbatch and D. Hooper, *Heavy dark matter through the Higgs portal*, *JHEP* **07** (2008) 058 [arXiv:0801.3440] [INSPIRE].

[8] 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].

JHEP12(2021)129

[9] D. Curtin and S. Gryba, *Twin Higgs portal dark matter*, *JHEP* **08** (2021) 009 [arXiv:2101.11019] [INSPIRE].

[10] B. Holdom, *Two* U(1)*'s and epsilon charge shifts*, *Phys. Lett. B* **166** (1986) 196 [INSPIRE].

[11] S.A. Abel, M.D. Goodsell, J. Jaeckel, V.V. Khoze and A. Ringwald, *Kinetic mixing of the photon with hidden* U(1)*s in string phenomenology*, *JHEP* **07** (2008) 124 [arXiv:0803.1449] [INSPIRE].

[12] B. Batell, M. Pospelov and A. Ritz, *Probing a secluded* U(1) *at B-factories*, *Phys. Rev. D* **79** (2009) 115008 [arXiv:0903.0363] [INSPIRE].

[13] J. Jaeckel and A. Ringwald, *The low-energy frontier of particle physics*, *Ann. Rev. Nucl. Part. Sci.* **60** (2010) 405 [arXiv:1002.0329] [INSPIRE].

[14] R. Foot, *Mirror dark matter: cosmology, galaxy structure and direct detection*, *Int. J. Mod. Phys. A* **29** (2014) 1430013 [arXiv:1401.3965] [INSPIRE].

[15] D. Feldman, Z. Liu and P. Nath, *The Stueckelberg Z-prime extension with kinetic mixing and milli-charged dark matter from the hidden sector*, *Phys. Rev. D* **75** (2007) 115001 [hep-ph/0702123] [INSPIRE].

[16] M. Pospelov, A. Ritz and M.B. Voloshin, *Secluded WIMP dark matter*, *Phys. Lett. B* **662** (2008) 53 [arXiv:0711.4866] [INSPIRE].

[17] E. Dudas, L. Heurtier, Y. Mambrini and B. Zaldivar, *Extra* U(1)*, effective operators, anomalies and dark matter*, *JHEP* **11** (2013) 083 [arXiv:1307.0005] [INSPIRE].

[18] H. An, X. Ji and L.-T. Wang, *Light dark matter and* Z' *dark force at colliders*, *JHEP* **07** (2012) 182 [arXiv:1202.2894] [INSPIRE].

[19] G.D. Kribs, A. Martin, B. Ostdiek and T. Tong, *Dark mesons at the LHC*, *JHEP* **07** (2019) 133 [arXiv:1809.10184] [INSPIRE].

[20] S. Knapen, J. Shelton and D. Xu, *Perturbative benchmark models for a dark shower search program*, *Phys. Rev. D* **103** (2021) 115013 [arXiv:2103.01238] [INSPIRE].

[21] M. Cvetič, P. Langacker and G. Shiu, *Phenomenology of a three family standard like string model*, *Phys. Rev. D* **66** (2002) 066004 [hep-ph/0205252] [INSPIRE].

[22] T. Hur, D.-W. Jung, P. Ko and J.Y. Lee, *Electroweak symmetry breaking and cold dark matter from strongly interacting hidden sector*, *Phys. Lett. B* **696** (2011) 262 [arXiv:0709.1218] [INSPIRE].

[23] Y. Bai and P. Schwaller, *Scale of dark QCD*, *Phys. Rev. D* **89** (2014) 063522 [arXiv:1306.4676] [INSPIRE].

[24] Y. Grossman and D.J. Robinson, *Composite Dirac neutrinos*, *JHEP* **01** (2011) 132 [arXiv:1009.2781] [INSPIRE].

[25] T. Han, Z. Si, K.M. Zurek and M.J. Strassler, *Phenomenology of hidden valleys at hadron colliders*, *JHEP* **07** (2008) 008 [arXiv:0712.2041] [INSPIRE].

[26] M. Buschmann, J. Kopp, J. Liu and P.A.N. Machado, *Lepton jets from radiating dark matter*, *JHEP* **07** (2015) 045 [arXiv:1505.07459] [INSPIRE].

[27] N. Arkani-Hamed and N. Weiner, *LHC signals for a SuperUnified Theory of dark matter*, *JHEP* **12** (2008) 104 [arXiv:0810.0714] [INSPIRE].

[28] S.D. Ellis, T.S. Roy and J. Scholtz, *Phenomenology of photon-jets*, *Phys. Rev. D* **87** (2013) 014015 [arXiv:1210.3657] [INSPIRE].

JHEP12(2021)129

[29] N. Toro and I. Yavin, *Multiphotons and photon jets from new heavy vector bosons*, *Phys. Rev. D* **86** (2012) 055005 [arXiv:1202.6377] [INSPIRE].

[30] ATLAS collaboration, *A search for pairs of highly collimated photon-jets in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Phys. Rev. D* **99** (2019) 012008 [arXiv:1808.10515] [INSPIRE].

[31] T. Cohen, M. Lisanti and H.K. Lou, *Semivisible Jets: dark matter undercover at the LHC*, *Phys. Rev. Lett.* **115** (2015) 171804 [arXiv:1503.00009] [INSPIRE].

[32] G. Burdman and G. Lichtenstein, *Displaced vertices from hidden glue*, *JHEP* **08** (2018) 146 [arXiv:1807.03801] [INSPIRE].

[33] P. Schwaller, D. Stolarski and A. Weiler, *Emerging jets*, *JHEP* **05** (2015) 059 [arXiv:1502.05409] [INSPIRE].

[34] T. Cohen, M. Lisanti, H.K. Lou and S. Mishra-Sharma, *LHC searches for dark sector showers*, *JHEP* **11** (2017) 196 [arXiv:1707.05326] [INSPIRE].

[35] T. Cohen, J. Doss and M. Freytsis, *Jet substructure from dark sector showers*, *JHEP* **09** (2020) 118 [arXiv:2004.00631] [INSPIRE].

[36] M.J. Strassler, *Why unparticle models with mass gaps are examples of hidden valleys*, arXiv:0801.0629 [INSPIRE].

[37] S. Knapen, S. Pagan Griso, M. Papucci and D.J. Robinson, *Triggering soft bombs at the LHC*, *JHEP* **08** (2017) 076 [arXiv:1612.00850] [INSPIRE].

[38] C. Cesarotti and J. Thaler, *A robust measure of event isotropy at colliders*, *JHEP* **08** (2020) 084 [arXiv:2004.06125] [INSPIRE].

[39] D. Rumelhart, G. Hinton and R. Williams, *Parallel distributed processing*, volume 1, MIT Press, U.S.A. (1986), chapter 8.

[40] P. Asadi, M.R. Buckley, A. DiFranzo, A. Monteux and D. Shih, *Digging deeper for new physics in the LHC data*, *JHEP* **11** (2017) 194 [arXiv:1707.05783] [INSPIRE].

[41] E.M. Metodiev, B. Nachman and J. Thaler, *Classification without labels: learning from mixed samples in high energy physics*, *JHEP* **10** (2017) 174 [arXiv:1708.02949] [INSPIRE].

[42] A. Andreassen, I. Feige, C. Frye and M.D. Schwartz, *JUNIPR: a framework for unsupervised machine learning in particle physics*, *Eur. Phys. J. C* **79** (2019) 102 [arXiv:1804.09720] [INSPIRE].

[43] J.H. Collins, K. Howe and B. Nachman, *Anomaly detection for resonant new physics with machine learning*, *Phys. Rev. Lett.* **121** (2018) 241803 [arXiv:1805.02664] [INSPIRE].

[44] A. De Simone and T. Jacques, *Guiding new physics searches with unsupervised learning*, *Eur. Phys. J. C* **79** (2019) 289 [arXiv:1807.06038] [INSPIRE].

[45] T. Heimel, G. Kasieczka, T. Plehn and J.M. Thompson, *QCD or what?*, *SciPost Phys.* **6** (2019) 030 [arXiv:1808.08979] [INSPIRE].

[46] M. Farina, Y. Nakai and D. Shih, *Searching for new physics with deep autoencoders*, *Phys. Rev. D* **101** (2020) 075021 [arXiv:1808.08992] [INSPIRE].

[47] T.S. Roy and A.H. Vijay, *A robust anomaly finder based on autoencoders*, arXiv:1903.02032 [INSPIRE].

[48] T. Cheng, J.-F. Arguin, J. Leissner-Martin, J. Pilette and T. Golling, *Variational autoencoders for anomalous jet tagging*, arXiv:2007.01850 [INSPIRE].

JHEP12(2021)129

[49] B. Nachman and D. Shih, *Anomaly detection with density estimation*, *Phys. Rev. D* **101** (2020) 075042 [arXiv:2001.04990] [INSPIRE].

[50] M.A. Md Ali, N. Badrud'din, H. Abdullah and F. Kemi, *Alternate methods for anomaly detection in high-energy physics via semi-supervised learning*, *Int. J. Mod. Phys. A* **35** (2020) 2050131 [INSPIRE].

[51] B. Bortolato, B.M. Dillon, J.F. Kamenik and A. Smolkovič, *Bump Hunting in Latent Space*, arXiv:2103.06595 [INSPIRE].

[52] B.M. Dillon, T. Plehn, C. Sauer and P. Sorrenson, *Better latent spaces for better autoencoders*, *SciPost Phys.* **11** (2021) 061 [arXiv:2104.08291] [INSPIRE].

[53] T. Finke, M. Krämer, A. Morandini, A. Mück and I. Oleksiyuk, *Autoencoders for unsupervised anomaly detection in high energy physics*, *JHEP* **06** (2021) 161 [arXiv:2104.09051] [INSPIRE].

[54] G. Kasieczka et al., *The LHC olympics 2020: a community challenge for anomaly detection in high energy physics*, arXiv:2101.08320 [INSPIRE].

[55] T. Aarrestad et al., *The dark machines anomaly score challenge: benchmark data and model independent event classification for the Large Hadron Collider*, arXiv:2105.14027 [INSPIRE].

[56] O. Cerri, T.Q. Nguyen, M. Pierini, M. Spiropulu and J.-R. Vlimant, *Variational autoencoders for new physics mining at the Large Hadron Collider*, *JHEP* **05** (2019) 036 [arXiv:1811.10276] [INSPIRE].

[57] D.E. Morrissey, T. Plehn and T.M.P. Tait, *Physics searches at the LHC*, *Phys. Rept.* **515** (2012) 1 [arXiv:0912.3259] [INSPIRE].

[58] N. Arkani-Hamed, S. Dimopoulos and S. Kachru, *Predictive landscapes and new physics at a TeV*, hep-th/0501082 [INSPIRE].

[59] N. Arkani-Hamed, A.G. Cohen and H. Georgi, *Electroweak symmetry breaking from dimensional deconstruction*, *Phys. Lett. B* **513** (2001) 232 [hep-ph/0105239] [INSPIRE].

[60] L. Randall and R. Sundrum, *A large mass hierarchy from a small extra dimension*, *Phys. Rev. Lett.* **83** (1999) 3370 [hep-ph/9905221] [INSPIRE].

[61] L. Randall and R. Sundrum, *An Alternative to compactification*, *Phys. Rev. Lett.* **83** (1999) 4690 [hep-th/9906064] [INSPIRE].

[62] G. Elor, N.L. Rodd and T.R. Slatyer, *Multistep cascade annihilations of dark matter and the Galactic Center excess*, *Phys. Rev. D* **91** (2015) 103531 [arXiv:1503.01773] [INSPIRE].

[63] C. Cesarotti, M. Reece and M.J. Strassler, *Spheres to jets: tuning event shapes with* 5*d simplified models*, *JHEP* **05** (2021) 096 [arXiv:2009.08981] [INSPIRE].

[64] A. Costantino, S. Fichet and P. Tanedo, *Effective field theory in AdS: continuum regime, soft bombs, and IR emergence*, *Phys. Rev. D* **102** (2020) 115038 [arXiv:2002.12335] [INSPIRE].

[65] S.C. Park, *Black holes and the LHC: a review*, *Prog. Part. Nucl. Phys.* **67** (2012) 617 [arXiv:1203.4683] [INSPIRE].

[66] R.K. Ellis, W.J. Stirling and B.R. Webber, *QCD and collider physics*, volume 8, Cambridge University Press, Cambridge U.K. (2011).

[67] Y. Hatta and T. Matsuo, *Jet fragmentation and gauge/string duality*, *Phys. Lett. B* **670** (2008) 150 [arXiv:0804.4733] [INSPIRE].

JHEP12(2021)129

[68] E. Fermi, *High-energy nuclear events*, *Prog. Theor. Phys.* **5** (1950) 570 [INSPIRE].

[69] R. Hagedorn, *Statistical thermodynamics of strong interactions at high-energies*, *Nuovo Cim. Suppl.* **3** (1965) 147 [INSPIRE].

[70] J.D. Bjorken and S.J. Brodsky, *Statistical model for electron-positron annihilation into hadrons*, *Phys. Rev. D* **1** (1970) 1416 [INSPIRE].

[71] P. Blanchard, S. Fortunato and H. Satz, *The hagedorn temperature and partition thermodynamics*, *Eur. Phys. J. C* **34** (2004) 361 [hep-ph/0401103] [INSPIRE].

[72] F. Becattini and G. Passaleva, *Statistical hadronization model and transverse momentum spectra of hadrons in high-energy collisions*, *Eur. Phys. J. C* **23** (2002) 551 [hep-ph/0110312] [INSPIRE].

[73] J. Cleymans, *The thermal model at the Large Hadron Collider*, *Acta Phys. Polon. B* **43** (2012) 563 [arXiv:1203.5640] [INSPIRE].

[74] F. Becattini, P. Castorina, J. Manninen and H. Satz, *The thermal production of strange and non-strange hadrons in e^+e^- collisions*, *Eur. Phys. J. C* **56** (2008) 493 [arXiv:0805.0964] [INSPIRE].

[75] F. Becattini, P. Castorina, A. Milov and H. Satz, *A comparative analysis of statistical hadron production*, *Eur. Phys. J. C* **66** (2010) 377 [arXiv:0911.3026] [INSPIRE].

[76] F. Becattini, P. Castorina, A. Milov and H. Satz, *Predictions of hadron abundances in pp collisions at the LHC*, *J. Phys. G* **38** (2011) 025002 [arXiv:0912.2855] [INSPIRE].

[77] L. Ferroni and F. Becattini, *Statistical hadronization with exclusive channels in e^+e^- annihilation*, *Eur. Phys. J. C* **71** (2011) 1824 [arXiv:1109.5185] [INSPIRE].

[78] E. Bernreuther, F. Kahlhoefer, M. Krämer and P. Tunney, *Strongly interacting dark sectors in the early Universe and at the LHC through a simplified portal*, *JHEP* **01** (2020) 162 [arXiv:1907.04346] [INSPIRE].

[79] D. Curtin et al., *Exotic decays of the 125 GeV Higgs boson*, *Phys. Rev. D* **90** (2014) 075004 [arXiv:1312.4992] [INSPIRE].

[80] ATLAS collaboration, *Search for invisible Higgs boson decays in vector boson fusion at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Phys. Lett. B* **793** (2019) 499 [arXiv:1809.06682] [INSPIRE].

[81] ATLAS collaboration, *Combination of searches for invisible Higgs boson decays with the ATLAS experiment*, *Phys. Rev. Lett.* **122** (2019) 231801 [arXiv:1904.05105] [INSPIRE].

[82] ATLAS collaboration, *Combined measurements of Higgs boson production and decay using up to 80 fb^{-1} of proton-proton collision data at $\sqrt{s} = 13$ TeV collected with the ATLAS experiment*, *Phys. Rev. D* **101** (2020) 012002 [arXiv:1909.02845] [INSPIRE].

[83] CMS collaboration, *Search for invisible decays of a Higgs boson produced through vector boson fusion in proton-proton collisions at $\sqrt{s} = 13$, TeV*, *Phys. Lett. B* **793** (2019) 520 [arXiv:1809.05937] [INSPIRE].

[84] A. Biekoetter, T. Corbett and T. Plehn, *The Gauge-Higgs Legacy of the LHC Run II*, *SciPost Phys.* **6** (2019) 064 [arXiv:1812.07587] [INSPIRE].

[85] A. Pierce, B. Shakya, Y. Tsai and Y. Zhao, *Searching for confining hidden valleys at LHCb, ATLAS, and CMS*, *Phys. Rev. D* **97** (2018) 095033 [arXiv:1708.05389] [INSPIRE].

[86] ATLAS collaboration, *Search for Higgs boson decays into two new low-mass spin-0 particles in the 4b channel with the ATLAS detector using pp collisions at $\sqrt{s} = 13$ TeV*, *Phys. Rev. D* **102** (2020) 112006 [arXiv:2005.12236] [INSPIRE].

JHEP12(2021)129

[87] M. Park and M. Zhang, *Tagging a jet from a dark sector with Jet-substructures at colliders*, *Phys. Rev. D* **100** (2019) 115009 [arXiv:1712.09279] [INSPIRE].

[88] T. Sjöstrand et al., *An introduction to PYTHIA 8.2*, *Comput. Phys. Commun.* **191** (2015) 159 [arXiv:1410.3012] [INSPIRE].

[89] DELPHES 3 collaboration, *DELPHES 3, a modular framework for fast simulation of a generic collider experiment*, *JHEP* **02** (2014) 057 [arXiv:1307.6346] [INSPIRE].

[90] M. Cepeda et al., *Report from working group 2: Higgs physics at the HL-LHC and HE-LHC*, *CERN Yellow Rep. Monogr.* **7** (2019) 221 [arXiv:1902.00134] [INSPIRE].

[91] Y. Wang, Y. Sun, Z. Liu, S.E. Sarma, M.M. Bronstein and J.M. Solomon, *Dynamic graph CNN for learning on point clouds*, arXiv:1801.07829 [INSPIRE].

[92] E. Bernreuther, T. Finke, F. Kahlhoefer, M. Krämer and A. Mück, *Casting a graph net to catch dark showers*, *SciPost Phys.* **10** (2021).

[93] H. Qu and L. Gouskos, *ParticleNet: jet tagging via particle clouds*, *Phys. Rev. D* **101** (2020) 056019 [arXiv:1902.08570] [INSPIRE].

[94] J.A. Aguilar-Saavedra, J.H. Collins and R.K. Mishra, *A generic anti-QCD jet tagger*, *JHEP* **11** (2017) 163 [arXiv:1709.01087] [INSPIRE].

[95] O. Knapp, O. Cerri, G. Dissertori, T.Q. Nguyen, M. Pierini and J.-R. Vlimant, *Adversarially learned anomaly detection on CMS open data: re-discovering the top quark*, *Eur. Phys. J. Plus* **136** (2021) 236 [arXiv:2005.01598] [INSPIRE].

[96] P. Baldi, K. Cranmer, T. Faucett, P. Sadowski and D. Whiteson, *Parameterized neural networks for high-energy physics*, *Eur. Phys. J. C* **76** (2016) 235 [arXiv:1601.07913] [INSPIRE].

[97] G. Louppe, M. Kagan and K. Cranmer, *Learning to pivot with adversarial networks*, arXiv:1611.01046 [INSPIRE].

[98] CMS collaboration, *A deep neural network to search for new long-lived particles decaying to jets*, *Mach. Learn. Sci. Tech.* **1** (2020) 035012 [arXiv:1912.12238] [INSPIRE].

[99] M. Cacciari, G.P. Salam and G. Soyez, *The anti-k_t jet clustering algorithm*, *JHEP* **04** (2008) 063 [arXiv:0802.1189] [INSPIRE].

[100] D.P. Kingma and M. Welling, *Auto-encoding variational Bayes*, arXiv:1312.6114 [INSPIRE].

[101] M. Ponce et al., *Deploying a top-100 supercomputer for large parallel workloads: the Niagara supercomputer*, in the proceedings of the *Practice and Experience in Advanced Research Computing on Rise of the Machines (Learning) (PEARC'19)*, July 28–August 1, Chicago, U.S.A. (2019).

[102] C. Loken et al., *SciNet: lessons learned from building a power-efficient top-20 system and data centre*. *J. Phys.: Conf. Ser.* **256** (2010) 012026.