

Soft-unclustered-energy patterns from quirks

David Curtin^{1,2,*} Sascha Dreyer^{3,†} Max Fusté Costa^{3,4,‡} Sarah Heim^{3,4,§} Gregor Kasieczka^{4,||}
 Louis Moureaux^{4,¶} David Rouso^{3,**} David Shih^{5,††} and Manuel Sommerhalder⁴

¹Department of Physics, *University of Toronto*, Toronto, Ontario, Canada M5S 1A7

²Theoretical Physics Department, *CERN*, 1211 Geneva 23, Switzerland

³Deutsches Elektronen-Synchrotron DESY, 22603 Hamburg, Germany

⁴Institut für Experimentalphysik, *Universität Hamburg*, 22761 Hamburg, Germany

⁵NHETC, Department of Physics and Astronomy, *Rutgers University*, Piscataway, New Jersey 08854, USA



(Received 19 June 2025; accepted 25 November 2025; published 9 January 2026)

We propose searching for physics beyond the Standard Model in the low-transverse-momentum tracks accompanying hard-scatter events at the LHC. TeV-scale resonances connected to a dark QCD sector could be enhanced by selecting events with anomalies in the track distributions. As a benchmark, a quirk model with microscopic string lengths is developed, including a setup for event simulation. For this model, strategies are presented to enhance the sensitivity compared to inclusive resonance searches: a simple cut-based selection, a supervised search, and a model-agnostic weakly supervised anomaly search with the CATHODE method. Expected discovery potentials and exclusion limits are shown for 140 fb^{-1} of 13 TeV proton-proton collisions at the LHC.

DOI: [10.1103/g7c7-6qh2](https://doi.org/10.1103/g7c7-6qh2)

I. INTRODUCTION

In the quest of finding physics beyond the Standard Model (SM) of particle physics at the LHC, it is crucial that no signatures are missed, in particular those involving soft-unclustered-energy patterns (SUEPs) from low-transverse-momentum (low- p_T) particles. Such signatures are challenging, because they often hide under large non-perturbative backgrounds which are hard to model and which make the online trigger selection difficult.

In the spirit of resonant anomaly detection (for reviews and original references, see [1–3]), we propose to search for phenomena beyond the SM (BSM) in the low- p_T tracks ($\lesssim \text{GeV}$) accompanying TeV-scale resonances. The presence of a resonance allows the online trigger to select high- p_T detector objects, and the sidebands of the resonance can be used to estimate backgrounds without having to rely on

simulated events. The proposed signature is motivated by a number of theories beyond the SM, some of which predict resonances that are, in fact, deexcited bound states. A typical signature would be a heavy particle-antiparticle pair produced in an excited state, which radiates hundreds of low- p_T particles as it approaches its ground state, where it finally annihilates, producing a resonant pair of high- p_T SM particles.

“Hidden valley” models are an entire class of models beyond the SM with so-far undetected particles that are charged under a new, usually QCD-like force [4,5], here referred to as “dark QCD.” The showering and hadronization according to the new force and the subsequent decays can produce a large number of low- p_T SM particles. Hidden valley models have been discussed in the context of various solutions to the SM problems, such as the dark-matter question [6], and the hierarchy problem. The latter is addressed in theories of neutral naturalness through a twin Higgs scenario [7,8], as well as folded supersymmetry [9–12], and the quirky little Higgs scenarios [13]. These models can give rise to so-called *quirks* [11,14], which are dark quarks with a mass that is much larger than the dark QCD confinement scale and would be produced in colliders as quirk-antiquark pairs. In the absence of low-mass dark quarks, the dark color string connecting two pair-produced heavy quirks cannot break but instead supplies a linear restoring force, creating a bound state with the two heavy particles oscillating back and forth on the end of the color string, until they deexcite by various radiative processes and eventually annihilate. Here, we study a microscopic quirk scenario where quirk-antiquark pairs that carry SM

* Contact author: dcurtin@physics.utoronto.ca

† Contact author: sascha.dreyer@desy.de

‡ Contact author: max.fuste.costa@desy.de

§ Contact author: sarah.heim@desy.de

|| Contact author: gregor.kasieczka@uni-hamburg.de

¶ Contact author: louis.moureaux@cern.ch

** Contact author: david.rousso@desy.de

†† Contact author: shih@physics.rutgers.edu

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Funded by SCOAP³.

QCD charge are produced in proton (pp) collisions at the LHC and then rapidly deexcite via the isotropic emission of low- p_T SM pions. Once deexcited, the quirks annihilate into a pair of high- p_T jets, as described in Sec. II.

We evaluate existing constraints on the developed quirk model from inclusive dijet searches performed by the ATLAS collaboration, and present strategies to enhance their sensitivity: A simple cut-based selection, a supervised classifier trained on the background estimate and simulated signal events, and a weakly supervised search for anomalies with the CATHODE framework [15]. Discovery potentials and exclusion ranges are also presented.

Our work demonstrates that the deexciting pion emission gives rise to a large number of detectable soft tracks that are isotropic in the quirk rest frame, which can serve as a striking discovery signature if it is utilized to extend the reach of a standard dijet search. We also provide a comparative analysis of various classical and machine learning approaches for the specific class of models under study, i.e., resonances accompanied by track anomalies.

II. SIGNAL MODEL

A very generic possibility for BSM physics is the existence of a dark QCD sector [4], meaning new states charged under a new asymptotically free gauge force with confinement scale Λ' . If the dark QCD sector has no light dark-quark flavors, which means all dark-QCD-charged matter is much heavier than the dark confinement scale Λ' , then the lightest dark hadrons are dark glueballs [16,17] with a mass $m_{GB} \approx 6\Lambda'$ [18] for a dark $SU(N)$ confining gauge force. This has important implications for the behavior of any heavy dark-sector states charged under dark QCD.

In any confining gauge theory, if a heavy quark-antiquark pair ($q'\bar{q}'$, with $m_{q'} \gg \Lambda'$) is produced, it is initially connected by a flux tube of confined force lines that classically would provide a linear restoring force. In SM QCD, this semiclassical behavior does not manifest, since it is energetically more favorable for this flux tube to immediately break by the creation of light quark-antiquark pairs with mass below the confinement scale Λ , corresponding to the production of light mesons in hadronization. However, in a dark QCD sector with no light quark flavors, the flux tube cannot break. The semiclassical linear restoring force of a confining gauge theory in the long distance ($r > 1/\Lambda'$) regime is therefore manifest, and the dark quark-antiquark behaves as though it is connected by a spring, oscillating back and forth with maximal linear separation $l \sim m_{q'}/\Lambda'^2$ (assuming that the dark quarks are created with kinetic energy $\sim m_{q'}$ in their rest frame). Such dark quarks are called *quirks*.

Upon creation of the “quirkonium” state, it will start to oscillate. The quirks can in principle annihilate each time they pass close to each other, but like any accelerating charges, they also immediately start deexciting by the

emission of radiation. This imparts orbital angular momentum onto the quirk pair, which greatly suppresses their annihilation rate, until the quirkonium reaches its ground state. In most scenarios, the quirkonium lifetime is therefore determined by the *deexcitation time scale*.

The type of radiation emitted depends on the quirk quantum numbers. If it is charged under electromagnetism [or any unbroken gauge force, like a dark $U(1)_D$], it will emit (dark) photons via bremsstrahlung. If it carries SM QCD charge, the quirks will emit SM pions. The one guaranteed deexcitation channel is emission of dark glueballs, since the quirk has to be charged under dark QCD. Quirk deexcitation dynamics is famously complicated. Photon emission is calculable, while the emission of pions can be crudely modeled to be isotropic and approximately thermal with a Hagedorn momentum spectrum [14], provided the string length is much smaller than Λ^{-1} . On the other hand, no rigorous treatment of quirk deexcitation via glueball emission exists. A similar ansatz to pion emission could be applied, but only as long as the quirk kinetic energy remains much larger than the glueball mass. Even so, the absolute glueball emission rate is very difficult to predict.

The above illustrates that quirks can exhibit highly exotic behavior [8,10,11,13,19–37]. We will focus on the regime where the quirk string is microscopic compared to the SM hadronic scale. Our signal model consists of quirks charged under both QCD and dark QCD, with masses in the phenomenologically interesting $m_{q'} \sim \mathcal{O}(\text{TeV})$ range for LHC resonance searches, while $\Lambda' \gtrsim \mathcal{O}(10 \text{ GeV})$ to ensure the quirk string satisfies $l \sim E_{\text{kin}}/\Lambda'^2 \sim m_{q'}/\Lambda'^2 < \Lambda^{-1}$ [14]. Without going into a detailed calculation of pion emission, this also ensures prompt (within a ps) deexcitation and annihilation, even for an implausibly small pion emission probability per oscillation. We are agnostic as to whether these dark quarks carry electromagnetic charge, since their phenomenology will be determined by their strong interactions. For concreteness, we take the dark QCD to be $SU(3)_D$, and the quirks to be fundamentals under $SU(3)_c \otimes SU(3)_D$. For their LHC production cross sections (see Fig. 1) we can use the next-to-leading order (NLO) results for vector-like quark (VLQ) production in Ref. [38], multiplied by a dark color factor of 3. The dark color contribution to the K factor will be a subdominant effect. After deexcitation, the $q'\bar{q}'$ state can annihilate into QCD and dark-QCD gluons, with a branching ratio to SM jets given by $\text{Br}(q'\bar{q}' \rightarrow jj) = \alpha_S^2/(\alpha_S^2 + \alpha_D^2)$.¹ Both couplings are evaluated at the scale $m_{q'}$, and $\alpha_D(m_{q'})$ is

¹This assumes that the quirk annihilates in the SM color-singlet state, which is justified for microscopic quirk strings with $l < \Lambda^{-1}$. In analogy to SM heavy quarkonium, annihilation of the deexcited quirk pair will be dominated by the s -wave color-singlet state, with the color-octet annihilation being suppressed by v^2 , where v is the nonrelativistic speed of the quirks in the deexcited bound-state rest frame [39].




FIG. 1. Cross section of quirk-antiquirk production in pp collisions at 13 TeV as a function of the quirk mass (top). Branching ratio for the diquark system to decay into two SM jets, depending on the quirk mass and Λ' (bottom). The dashed red line indicates where quirk annihilation to two dark glueballs is kinematically forbidden, invalidating the presented simplified estimate. Close to the dashed red line, the quirks dominantly or entirely annihilate to SM jets.

determined, assuming that q' is the lightest dark-QCD charged state in the theory, by one-loop pure Yang-Mills RG evolution from $\alpha_D^{-1}(\Lambda') = 0$. The resulting annihilation branching fraction to QCD jets is shown in Fig. 1 (bottom).

The estimate of the branching ratio to SM jets given above does not take into account the phase space suppression from the dark glueball mass, which can be significant at large Λ' . We therefore also show in Fig. 1 (bottom) a dashed red line indicating when the lightest dark glueball mass $m_{GB} \approx 6\Lambda'$ exceeds $m_{q'}$. Above this line, annihilation to two dark glueballs is kinematically forbidden, and the branching ratio to SM jets will be close to 1 (though a subdominant mixed decay to one dark glueball and a SM jet may occur). Computing the precise decay width including glueball mass effects is beyond the scope of this paper, but around this threshold, the SM jet branching ratio of quirk annihilation will be close to 1.

Following $q'\bar{q}'$ pair creation, we model deexcitation of the quirkonium state by converting the quirk kinetic energy

E_{kin} in the quirkonium rest frame into isotropic pion emission with a thermal energy spectrum $dN/d^3\mathbf{p} \propto \exp(-E/T_H)$, where $T_H = 1.2\Lambda$ is the Hagedorn temperature taken to be 150 MeV. This results in the emissions of $N_\pi \sim E_{\text{kin}}/T_H \sim m_{q'}/T_H \sim \mathcal{O}(10^3)$ pions per quirk pair, with relative probabilities of pion charges given by isospin (1/3 each for positive, negative, and neutral). The last emitted pion charge is dictated by overall charge conservation. See also Ref. [11] for an early discussion of this soft-track quirk signal.

The fraction of quirkonium internal energy that is radiated away as dark glueballs is unknown. In general, the quirks may also have other dark sector interactions that open additional deexcitation channels, such as dark photon emission. We account for this uncertainty by introducing a parameter $0.2 \leq E_{\text{frac}} \leq 1$ into our model, setting the fraction of quirkonium internal energy that is converted into thermal isotropic pions.

The final model therefore has three parameters: the mass of the quirks $m_{q'}$, the energy scale Λ' , and the fraction of energy going to pions E_{frac} . In practice, for the range of dark confinement scales we study, the exact value of Λ' only changes the cross section normalization of the dijet annihilation signal by determining $\text{Br}(q'\bar{q}' \rightarrow jj)$. In order to test the sensitivity, we also introduce the signal strength μ , which scales the cross section of the model with the chosen parameters. For example, a sensitivity to $\mu = 10$ means that we can find the signature, but only if its cross section is 10 times as large as predicted by the nominal model.

III. SIGNAL AND BACKGROUND SIMULATION

Assuming pp collisions at 13 TeV center-of-mass energy, the MG5_aMCatNLO 3.4.2 [40] MC generator was used to generate the signal and background events, together with PYTHIA 8.2 [41] for parton shower, hadronization, and the underlying event from multiple parton interaction (MPI). The NNPDF23 set [42] was used. No pileup events were added, as the simplifying assumption is made that the track-vertex association reduces the contribution from pileup tracks significantly (to 10% or less, see Ref. [43]) and the remaining contribution can be estimated [44]. Obviously, the analysis strategy will have to be revisited for the high-luminosity LHC with its more challenging pileup conditions but improved tracking detectors.

A. Signal simulation

The generation of simulated signal samples was performed in a semianalytical hybrid approach. The simulation consists of three components: the simulation of the quirk-antiquirk production, the radiation of pions during the oscillation/deexcitation process, and the decay of the quirkonium. Quirk-antiquirk production is approximated using top-quark-like heavy-vector-like quark triplets [45,46]

generated with the MG5_aMCatNLO MC generator and PYTHIA. The nonperturbative emissions of the pions during the deexcitation are simplified by treating them as isotropic radiation, thereby bypassing detailed simulation of each radiated particle by assuming an isotropic thermal distribution. Depending on the relative kinetic energy of the two quirks, and the chosen parameter value of E_{frac} , charged and neutral pion four-vectors are produced isotropically with an energy sampled from a relativistic Maxwell-Boltzmann distribution with a temperature that corresponds to T_H . This thermal approach is inspired by earlier SUEP studies [47]. The momentum of the pion system is adjusted to that of the diquirk system. Finally, the quirk-antiquirk annihilation to two hadronic jets is produced separately with MG5_aMCatNLO from the decay of a heavy scalar particle, whose mass corresponds to twice the quirk mass, also with a total momentum adjusted to that of the diquirk system. Events were simulated for masses ranging between 750 and 1500 GeV, and E_{frac} values between 0.2 and 1.0.

Note that in general we neglect the momentum imparted on the quirk bound state during deexcitation, and treat total pion emission as exactly isotropic. This is a good approximation when deexcitation is dominated by pion emission, i.e., E_{frac} is close to 1. Given most quirk deexcitations emit thousands of pions, with an average pion momentum of $\langle p_\pi \rangle \sim T_H$, the total momentum of the emitted pion radiation is roughly $|\vec{p}| \sim N_\pi^{1/2} \langle p_\pi \rangle \sim (E_{\text{kin}} T_H)^{1/2} \sim \mathcal{O}(10 \text{ GeV})$, which is a very small correction to our event kinematics given the quirk mass scales we study. Similar arguments apply when $E_{\text{frac}} < 1$ and the additional dark radiation is dominated by light states like dark photons. The case where $E_{\text{frac}} < 1$ and the dark radiation is dominated by dark glueballs deserves special mention. The dynamics of glueball emission during quirk deexcitation are unknown, but we can write the average number of emitted glueballs to be $N_{GB} \sim \alpha E_{\text{kin}} / m_{GB} \sim \alpha E_{\text{kin}} / (6\Lambda')$, with $\alpha < 1$, while each glueball's average momentum is likely $\langle p_{GB} \rangle \sim \beta m_{GB} = 6\beta\Lambda'$, with $\beta < 1$ if emission is Hagedorn-like with an energy set by $T'_H \sim \Lambda'$. For $E_{\text{frac}} \ll 1$, this yields an estimate for the total momentum of the emitted glueball radiation of $|\vec{p}| \sim (N_{GB})^{1/2} \langle p_{GB} \rangle \sim \beta\alpha^{1/2} (E_{\text{kin}} \Lambda')^{1/2}$. Given that we study Λ' up to $\mathcal{O}(100 \text{ GeV})$, the resulting momentum shift on the quirk system may not be entirely negligible, and may appreciably shift the kinematics of the resulting dijet signal. However, we neglect this effect for two reasons. First, it is still expected to be subdominant and not greatly impact the reach of the dijet search once the additional pion track criteria are imposed. Second, given the unknown details of deexcitation via glueball emission, we choose to keep the analysis agnostic of those details (thereby letting the results apply to other sources of dark radiation). It would certainly be interesting to repeat our studies for $E_{\text{frac}} < 1$ at some future time, when the dynamics of deexcitation via glueball emission are better understood.

B. Background simulation

The dijet background, including MPI, was simulated at leading order with the MG5_aMCatNLO generator and PYTHIA 8. A trijet sample was also generated at leading order to confirm that the contribution of events with a third high- p_T jet is negligible even in the tails of the studied distributions. Events were simulated for each of the following dijet invariant mass regions: 1.1–1.5 TeV (400M), 1.5–2.2 TeV (80M), 2.2–4 TeV (10M), 4–6 TeV (10M), 6–9 TeV (1M).

C. Preselection and distributions

The generated events were processed through a simplified LHC detector simulation [48] based on DELPHES 3.5 [49].

The applied selection criteria approximate the ATLAS dijet search [50], except that the radius parameter of the selected jets is increased from anti- k_t $R = 0.4$ to 1.0, in order to improve the dijet invariant mass resolution through better including final state radiation (FSR). Of course, this also increases the number of pion tracks that are wrongly included in the resonance. The resulting signal width is about 20% of the resonance mass if approximated by a Gaussian distribution. Without the radius increase, the signal, even with zero additional pion radiation, would be significantly wider [50], causing a reduction in significance and challenges for the CATHODE algorithm. A dedicated FSR recovery algorithm would, of course, be even more powerful, and more robust against initial state radiation (ISR) and the inclusion of pion radiation.

Events are required to have at least two jets in the detector acceptance with transverse momentum larger than 150 GeV. Tracks must be in the detector acceptance and have transverse momenta larger than 0.5 GeV, the current threshold for track reconstruction in ATLAS [51]. A selection is applied on half of the rapidity separation between the two leading jets, $y^* = (y_1 - y_2)/2$, where $y_1(y_2)$ is the rapidity of the (sub)leading jet. To favor s-channel processes and suppress QCD, only events with $|y^*| < 0.6$ are selected. The distribution of the dijet invariant mass m_{jj} is only considered above 1.3 TeV, as for lower values the required number of simulated background events would be too high. The m_{jj} signal region (SR) depends on the probed $m_{q'}$. Its definition aims to contain as much signal as possible while keeping the window narrow enough for a good signal over background ratio and for the CATHODE algorithm to perform well. Further considerations are the exponentially falling background and the need to accommodate different signal shapes for different values of E_{frac} . The SR is defined such that 68% of the signal events for $E_{\text{frac}} = 1.0$ are required to fall below its upper bound in m_{jj} . The lower bound is defined such that the difference between it and $2m_{q'}$ is half the difference between $2m_{q'}$ and the upper

TABLE I. Generated (Gen.) and expected (Exp.) number of events after event selection for dijet background and different signal hypotheses ($\Lambda' = 10$ GeV) in the chosen mass window, based on 140 fb^{-1} of 13 TeV pp data.

E_{frac}	$m_{q'} = 0.75 \text{ TeV}$		$m_{q'} = 1.0 \text{ TeV}$		$m_{q'} = 1.5 \text{ TeV}$	
	0.2	1.0	0.2	1.0	0.2	1.0
Gen. signal	1.4×10^5	1.5×10^5	2.1×10^4	5700	2.3×10^4	2.3×10^4
Exp. signal	1.3×10^4	1.3×10^4	1800	1900	78	79
Gen. background	5.4×10^7		7.9×10^6		6.5×10^5	
Exp. background	1.2×10^7		2.3×10^6		1.4×10^5	
Exp. $S/\sqrt{B}$	3.6	3.8	1.2	1.3	0.21	0.21

bound. For example, for $m_{q'} = 1.5 \text{ TeV}$, the window is 2.82–3.36 TeV to ensure good coverage of the invariant mass peak at approximately $2m_{q'} = 3 \text{ TeV}$.

Together with the invariant mass requirement, the y^* selection translates into an implicit selection cut on the transverse momentum of the jets, which is significantly higher than the initial 150 GeV, such that they would pass the corresponding trigger requirements applied by the experiments. The expected numbers of signal and background events for different mass points and values of E_{frac} , for 140 fb^{-1} of pp data, corresponding to the LHC run-2 dataset collected by ATLAS or CMS, are shown in Table I. The difference in selected events between the different E_{frac} values can be attributed to slightly higher reconstructed dijet masses with more pion radiation. It can be seen that the dijet event selection leads to an $S/\sqrt{B}$ of 0.21 to 3.8 depending on the quirk mass, where S is the number of signal and B the number of background events in the SR.

Figure 2 shows a number of reconstruction-level distributions, comparing the simulated dijet background with the quirk signal, where $m_{q'}$ is chosen to be 1.5 TeV and E_{frac} is varied between 0.2 and 1.0. The top figure shows the invariant mass of the two leading jets. The resonant structure can clearly be seen, as well as the shift to higher invariant masses when significant pion radiation is included ($E_{\text{frac}} = 1.0$). The remaining figures show track properties, where the tracks include both the jet tracks and the tracks from the pion radiation to avoid issues related to infrared safety: the number of tracks N_{tracks} , which increases dramatically with larger E_{frac} , the absolute value of the polar angle θ shifted by $\pi/2$, which encapsulates the higher centrality of the pion radiation in signal events compared to the QCD background, the transverse sphericity [52], which reflects the isotropy of the radiated pions (a value of 0 corresponds to a pencil-like event, and a value of 1 to a spherical event), and $\langle(\Delta R)^2\rangle$, which describes the average distance between each track pair in the event. The higher density signal events feature tracks that are closer together than what is observed in the QCD background. Decreasing the quirk mass reduces the energy scale of the signal events and with that the track multiplicity and therefore the discrimination power of the track-based observables.

D. Constraints from existing searches

So far, to our knowledge, no search for resonances combined with anomalies in the low- p_T track distributions has been performed at any collider. However, thanks to the presence of the resonance, “inclusive searches,” i.e., bump hunts over a falling background spectrum after a dijet selection, should provide some sensitivity.

An example is the search by the ATLAS collaboration for dijet resonances, using approximately 140 fb^{-1} of 13 TeV pp data produced by the LHC [50]. With our simulated dijet background events and Gaussian signals, we can reproduce the expected ATLAS cross-section limits for narrow-width signals to within 40%. The main difference comes from our imperfectly modeled dijet background. In Ref. [50], the observed fiducial cross-section limits are given for Gaussian signal shapes with widths up to 15%. Comparing the corresponding excluded fiducial cross sections to the cross sections and acceptances of our model, quirk masses between 750 and 1000 GeV (depending on the probed Λ') can be excluded. This mass exclusion range is about 10% stronger than what we present as “inclusive” in our results in Sec. V, due to multiple reasons: the difference between our simulated background distribution and the ATLAS dijet spectrum, our simplified signal extraction (cut-and-count instead of a fit, see Sec. IV), and differences between the Gaussian and our actual signal shape.

General SUEP searches (see, e.g., Ref. [53]) do not apply to the presented models with deexciting QCD-charged quirks for two reasons. First, they usually target signals arising from a dark QCD sector with quasiconformal running, which produces a SUEP of dark hadrons with some mass and effective temperature. Since these dark hadrons decay to SM final states, the associated mass scale is $\mathcal{O}(\text{GeV})$ or larger. The signature discussed here arises from the direct isotropic quasithermal emission of SM pions, such that the resulting “SM-QCD SUEP” covers a different parameter space, $\mathcal{O}(100 \text{ MeV})$, than the one covered by general SUEP searches. Second, even if these searches were sensitive to the discussed signal, they do not include the high mass resonance requirement, which greatly increases the sensitivity.




FIG. 2. Reconstruction-level distributions, comparing the shape of the simulated dijet background (black) with the quirk signal (blue), where $m_{q'}$ is chosen to be 1.5 TeV and E_{frac} is varied between 0.2 and 1.0. Also shown are the dijet background distributions estimated by a CATHODE run (pink). The invariant mass of the dijet system is shown on the top, while the other figures show track distributions in the m_{jj} signal window.

IV. SELECTION STRATEGIES

This study evaluates the power of the different selection strategies by approximating the discovery sensitivity by $S/\sqrt{B}$ based on simulated signal and background events and their associated truth labels in the chosen dijet invariant mass window. Of course, in an actual data analysis, a bump hunt would be performed that involves fitting analytical functions or templates for the signal and background components to the resulting dijet invariant mass spectrum and systematic uncertainties would be included.

A. Rectangular selection cut

The simplest way to improve the sensitivity over a resonance-only search is to apply a selection cut to the most powerful discriminant observable, which is the number of tracks (found via ablation studies using the CATHODE approach). This observable also has the advantage of being very general—many models with track anomalies will exhibit an increased number of tracks, be it from dark shower decays or oscillation radiation.

The best choice of cut depends on the model. In the following, we show results for a relatively loose background efficiency of 2%, which works better for models with less radiation, and for a background efficiency of 0.05%, which is tighter and benefits models with many tracks in the final state, as it improves the signal-background separation.

B. Supervised classifier

With the primary goal to understand the maximally achievable sensitivity to the benchmark model, a supervised classifier was trained on the signal and background distributions.

The considered observables are the number of tracks, the absolute value of θ shifted by $\pi/2$, the transverse sphericity, and $\langle(\Delta R)^2\rangle$, as described in Sec. III C and shown in Fig. 2.

A neural network is implemented in PYTORCH [54] containing three hidden layers with 64 nodes each and a binary cross-entropy loss. It is trained for 100 epochs with a batch size of 128, using the Adam [55] optimizer with a learning rate of 10^{-3} . Half of the generated signal and background events in the signal region are used for training, 1/6 for validation, and 1/3 for testing and determination of the results. Ten trials are made, with the events shuffled between the training, validation, and testing samples. The results presented in Sec. V are based on the average of these ten trials.

To understand the dependence on the variations within the benchmark model, the supervised classifier was also applied in a mismatched way—trained on a model with a chosen E_{frac} value and evaluated for a model with another E_{frac} value.

C. Anomaly detection with CATHODE

Anomaly detection strategies are employed to be as model-agnostic as possible. They allow finding regions of phase space where a selected set of observables show differences between data (background plus possible signal) and the expected background, the latter ideally estimated purely from data control regions.

The anomaly detection strategy chosen here is a weakly supervised search with CATHODE [15]. CATHODE is based on weak supervision [56], a training paradigm in which a classifier is trained to distinguish between a background estimate of probability density $p_{\text{bkg}}^{\text{est}}$ and data, which may contain a small fraction ϵ of signal events and has the probability density $p_{\text{data}} = (1 - \epsilon)p_{\text{bkg}} + \epsilon p_{\text{signal}}$. If $p_{\text{bkg}}^{\text{est}} \approx p_{\text{bkg}}$, the classifier learns the likelihood ratio $p_{\text{data}}/p_{\text{bkg}} = 1 - \epsilon + p_{\text{sig}}/p_{\text{bkg}}$, which is monotonically related to the optimal discriminant $p_{\text{sig}}/p_{\text{bkg}}$ between signal and background.

To obtain the background estimate, CATHODE starts by splitting the sample into a SR and sidebands based on a known feature, here chosen to be the dijet invariant mass m_{jj} . The SRs are constructed as described in Sec. III using the m_{jj} distributions for the $E_{\text{frac}} = 1$ models. The conditional probability density $p_{\text{bkg}}^{\text{est}}(x|m_{jj})$ of the background in the sidebands is learned using a normalizing flow [57] conditioned on m_{jj} , trained using events from the sidebands. The flow architecture and training uses SK_CATHODE [58] with the Pyro [59] back-end and the default settings. In a second step, the flow is sampled for values of m_{jj} obtained from a kernel density estimate [60] of the data in the signal region. These background samples are then compared to the data using a weakly supervised classifier.

The considered features are the same as in the supervised case. Different classifiers and ensembling strategies were tested. The presented results are based on a neural network with the same architecture and hyperparameters as listed above for the supervised search. Unlike in the supervised case, which is trained based on large simulated samples, the training in anomaly searches relies on the limited number of events available in the respective dataset. To understand the statistical spread, for every tested signal mass and every probed μ , ten trials are made. The results in Sec. V are the average of these trials, while the uncertainty bands correspond to the 1σ ranges. Each trial involves sampling the expected number of events for 140 fb^{-1} of data from the generated signal and background events. In each trial, two nonoverlapping samples are drawn. Half of the first sample is used for training, 1/6 for validation, and 1/3 for testing. The second independent sample is used in its entirety to extract the results.

In Sec. V, results for two versions of the above are presented: Results using a so-called idealized anomaly detector (IAD), where it is assumed that the background estimate is perfect (here, the background estimate is then

simply taken from a dedicated independent background sample), and the full CATHODE chain, which includes the normalizing flow estimate of the background in the signal region, as shown in Fig. 2.

V. RESULTS

As stated above, the results in this section are based on simulated samples, assuming 140 fb^{-1} of pp collisions at 13 TeV, which approximates the parameters of the LHC run-2 dataset. Setting the stage, for the inclusive resonance search, the sensitivity $S/\sqrt{B}$ versus the signal strength is shown in Fig. 3 for different quirk masses and two values of E_{frac} . The sensitivity depends linearly on the signal strength through the number of signal events S , with varying slopes due to the expected signal and background cross sections. The effect of increasing the value of E_{frac} is very small and can be attributed to a small shift in the dijet peak to higher values, leading to a slightly increased $S/\sqrt{B}$ (see also Table I). It can be seen that discovery sensitivity ($S/\sqrt{B} = 5$) is almost reached for the nominal benchmark model ($\mu = 1$) at a mass of 750 GeV, while the other masses require higher signal strengths.

A. Significance improvement curves

A selection of significance improvement curves with respect to the inclusive resonance search is shown in Fig. 4. The curves are calculated as a function of the false positive rate (that is, the background efficiency) for a quirk mass of 1.5 TeV and two choices of E_{frac} (0.2 and 1.0), as well as the different additional selection strategies described in Sec. IV. The uncertainty bands are derived by sampling the background and signal samples 10 times. The signal strengths chosen for the anomaly detection curves correspond to an $S/\sqrt{B} = 1$ and 2 after the preselection; the other classifiers are independent of this value.




FIG. 3. Sensitivity $S/\sqrt{B}$ for different signal strengths μ after the preselection, presented for $\Lambda' = 10 \text{ GeV}$, different quirk masses and two values of E_{frac} .




FIG. 4. Significance improvement curves for a quirk mass of 1.5 TeV as a function of the false positive rate, for the different selection strategies and two values of E_{frac} . The significance improvement for CATHODE and IAD depends on the signal strength; results for $S/\sqrt{B} = 1$ and 2 at preselection are presented. The uncertainty bands are derived from ten different samplings of the background and signal samples. The vertical line indicates the chosen false positive rate.

$E_{\text{frac}} = 1$ is a striking signature with a lot of emitted pions modifying the track multiplicity distribution in particular. All selection strategies perform very well, with a slight decrease in the simple cut-based selection compared to the multivariate methods. For both $E_{\text{frac}} = 1$ and $E_{\text{frac}} = 0.2$, the significance improvement in CATHODE is worse than in IAD, due to the additional challenge of estimating the background through the normalizing flow, especially in the tail of the track multiplicity. We leave the improvement of the normalizing flow in the tails of observables to future work and point out that the performance difference is less significant in the discovery potential and exclusion limits probed below. Tail-based anomaly detection has also been shown to work for the models probed in Ref. [61].

$E_{\text{frac}} = 0.2$ is a relatively challenging signature, adding fewer additional tracks. As expected, the supervised

classifier achieves the best improvement, as it is multivariate and tailored to the specific signal model. Interestingly, a mismatch of E_{frac} in the training compared to the test signal does not lead to a significant performance decrease, making the supervised classifier a very attractive option for this class of models. Both anomaly detection variants lose power if the signal strength becomes too small, as the classifier is trained to discern $S + B$ from B , which becomes challenging for small signals. The rectangular selection cut on the track multiplicity shows a significance improvement of 2–3, and, like the supervised classifier, has the advantage to perform independently of the signal strength, but the disadvantage of a certain model dependence.

B. Discovery potential

In the end, the relevant quantities are not the improvements in significance, but the detection potential, that is, the required signal strength for a 5σ excess, and, in the absence of a signal, the exclusion limits that can be set. For these evaluations, an operating point needs to be chosen, which is usually done by picking a background efficiency. Since all selection strategies, except for the track multiplicity cut, exhibit a rising improvement with a tighter cut (see Fig. 4), we show the results for a relatively tight background efficiency of 0.05%. For the track multiplicity cut, we also show the results for a looser background efficiency of 2%.

Figure 5 shows the required signal strengths for a 5σ discovery for different quirk masses and different values of E_{frac} , employing the different selection strategies discussed in Sec. IV. For the anomaly detection methods, the ranges are again the 1σ bands from the ten samplings. All strategies show significant improvements over the inclusive resonance search, except in small regions of phase space around $E_{\text{frac}} \approx 0.2$. The nominal cross sections of the benchmark model ($\mu = 1$) can be probed for all tested masses, if E_{frac} is high enough. CATHODE performs worse than the inclusive search for $m_q = 750$ GeV and $E_{\text{frac}} \approx 0.2$, as the challenging signature reduces the classifier to a random selection, decreasing the $S/\sqrt{B}$ value at a given operating point. As expected, the supervised classifier requires the smallest signal strength for discovery. The performance decreases only minimally when the classifier is trained with a mismatched E_{frac} . The rectangular cut on the track multiplicity with a background efficiency of 0.05% performs very well, with required discovery signal strengths that are higher by $\lesssim 2$ compared to the supervised classifier. The looser track multiplicity cut improves the performance only minimally at low values of E_{frac} . For the chosen benchmark model, which should be representative of a whole class of models featuring a resonance and a higher track multiplicity compared to the background, the IAD and CATHODE performance is very similar to the track multiplicity cut. This is consistent with the SIC curves in Fig. 4: For example, for a quirk mass of 1.5 TeV and




FIG. 5. Required signal strengths for a 5σ discovery for $\Lambda' = 10$ GeV, different quirk masses and different values of E_{frac} , for different selection strategies (lower is better). For CATHODE and IAD, the uncertainty bands are dominated by the 1σ spread of ten trial runs. For the other curves, they correspond to the expected statistical uncertainties on S and B (as the trial spread is negligible). The dashed horizontal line indicates a signal strength of 1.

$E_{\text{frac}} = 0.2$, the signal strength needed for a discovery with IAD or CATHODE corresponds to significances after preselection of $S/\sqrt{B} \approx 1$ to 1.5, which can be determined by finding the $S/\sqrt{B}$ values in Fig. 3 that correspond to the signal strength values read off the curves in Fig. 5. These are values for which the anomaly detection methods exhibit similar performance as the track multiplicity cut, also in

the SIC curves. It was verified that the anomaly detection methods indeed correctly identify the track multiplicity as the most dominant observable and rely almost exclusively on it for low amounts of signal, which do not allow the classifiers to benefit from the more subtle observables, like the angular distributions. Models with additional powerful and more model-dependent observables benefit more from the multivariateness and improved model dependence of anomaly detection methods. The performance difference between CATHODE and IAD for the lowest mass point is due to the less discriminant features at low quirk masses, which make the classifier more sensitive to mismodelings of the normalizing flow.

C. Exclusion limits

In the absence of signal, limits at the 95% confidence limit are set as a function of the quirk mass, and different values of Λ' . The strategy followed for the limit setting with CATHODE and IAD is described in the Appendix of Ref. [62].

Figure 6 shows the exclusion limit contours for $E_{\text{frac}} = 0.2$. It can be seen that quirk masses up to $m_{q'} \approx 1450$ GeV can be excluded with the supervised classifier. The performance of the mismatched supervised classifier is not much worse. The loose selection cut on the track multiplicity follows, whereas IAD and CATHODE do not




FIG. 6. Expected 95% CL exclusion limits as a function of the quirk mass and Λ' for $E_{\text{frac}} = 0.2$ (higher is better). The area to the left under the curves is excluded by the different search strategies. For CATHODE and IAD, the uncertainty bands are dominated by the 1σ spread of ten trial runs. For the other curves, they correspond to the expected statistical uncertainties on S and B (as the trial spread is negligible). The dashed red line indicates when quirk annihilation to two dark glueballs is kinematically forbidden, invalidating the naive estimate of the branching fraction to SM jets in Fig. 1. Close to the dashed red line, the quirks dominantly or entirely annihilate to SM jets, and the actual experimental limit therefore slightly exceeds the $m_{q'}$ mass reach found at very low Λ' values. Note that the entire parameter plane is excluded for $E_{\text{frac}} = 1$ by the supervised classifiers and the tight track multiplicity cut.

improve over the inclusive selection. The worse performance of the anomaly detection compared to the discovery potential is due to the smaller values of $S/\sqrt{B}$ probed by the exclusion results. These small values make it very difficult for the classifier to find a difference between $S+B$ and B even in the track multiplicity observable, as the $S+B$ distribution can be lower than B due to statistical fluctuations.

As for the discovery sensitivities, $E_{\text{frac}} = 1$ is a striking signature and applying the supervised classifier or the tight track multiplicity cut allows to exclude the whole probed phase space of $m_{q'} < 1500$ GeV and $\Lambda' < 200$ GeV. The performance of IAD and CATHODE is very similar.

VI. CONCLUSION

In conclusion, we propose to search for a TeV-scale resonance accompanied by anomalies in the low-transverse momentum ($\lesssim 1$ GeV) tracks at the LHC, thereby enhancing the discovery potential for new phenomena like deexcited bound states or dark showers. A benchmark quirk model and a corresponding simulation strategy were developed, representing phenomena leading to final states with a TeV-scale dijet resonance and increased track multiplicity from radiated pions. Different classifiers relying on track information were tested. These can enhance the discovery potential of a subsequent bump hunt in the dijet invariant mass spectrum. Estimated discovery sensitivities and exclusion ranges are presented for 140 fb^{-1} of 13 TeV LHC data.

For the chosen benchmark model, the track multiplicity is by far the dominant discriminator, and a rectangular selection based on this observable is a simple, relatively model-independent way to enhance the sensitivity, as it does not make any assumptions on other characteristics of the new tracks. A supervised classifier optimized for the specific model increases the sensitivity further by making use of other, less discriminant, track properties, like angular distributions. Interestingly, the supervised classifier also performs well when an incorrect track multiplicity is assumed in the training. However, its performance could degrade if other track properties differ or are severely mismodeled. CATHODE, the anomaly detection method

tested in this work, yields a discovery sensitivity similar to the simple track selection, due to the dominance of this observable and the small amounts of signal needed for a discovery, which prevents the method from using more granular information. The use of anomaly detection methods could, however, be beneficial to maintain a low model dependence also for more complex models with large signal-background differences in other observables.

The suggested phenomenology covers a phase space that has not been probed at the LHC so far. It is rich and warrants further studies both by varying the type of resonance, e.g., enhancing dilepton or diphoton resonance searches, and by varying the type of particles and emission patterns of the radiation spectra.

ACKNOWLEDGMENTS

The work of D. C. was supported in part by Discovery Grants from the Natural Sciences and Engineering Research Council of Canada (NSERC), the Canada Research Chair program, the Alfred P. Sloan Foundation, the Ontario Early Researcher Award, and the University of Toronto McLean Award. M. F. C., S. H., G. K., L. M., and M. S. acknowledge support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy—EXC2121 “Quantum Universe”—390833306 as well as by the Bundesministerium für Bildung und Forschung under Project No. 05H24GUA. S. H. thanks the Helmholtz Association for the support through the EBP initiative. The work of D. S. was supported by DOE Grant No. DE-SC0010008. This work has benefited from computing services provided by the German National Analysis Facility (NAF).

DATA AVAILABILITY

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

-
- [1] G. Kasieczka *et al.*, The LHC Olympics 2020: A community challenge for anomaly detection in high energy physics, *Rep. Prog. Phys.* **84**, 124201 (2021).
 - [2] V. Belis, P. Odagiu, and T. K. Aarrestad, Machine learning for anomaly detection in particle physics, *Rev. Phys.* **12**, 100091 (2024).
 - [3] G. Karagiorgi, G. Kasieczka, S. Kravitz, B. Nachman, and D. Shih, Machine learning in the search for new fundamental physics, *Nat. Rev. Phys.* **4**, 399 (2022).
 - [4] M. J. Strassler and K. M. Zurek, Echoes of a hidden valley at hadron colliders, *Phys. Lett. B* **651**, 374 (2007).

- [5] M. J. Strassler, On the phenomenology of hidden valleys with heavy flavor, [arXiv:0806.2385](#).
- [6] 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**, 262 (2011).
- [7] Z. Chacko, H.-S. Goh, and R. Harnik, The twin Higgs: Natural electroweak breaking from mirror symmetry, *Phys. Rev. Lett.* **96**, 231802 (2006).
- [8] N. Craig, A. Katz, M. Strassler, and R. Sundrum, Naturalness in the dark at the LHC, *J. High Energy Phys.* **07** (2015) 105.
- [9] G. Burdman, Z. Chacko, H.-S. Goh, and R. Harnik, Folded supersymmetry and the LEP paradox, *J. High Energy Phys.* **02** (2007) 009.
- [10] G. Burdman, Z. Chacko, H.-S. Goh, R. Harnik, and C. A. Krenke, The quirky collider signals of folded supersymmetry, *Phys. Rev. D* **78**, 075028 (2008).
- [11] R. Harnik and T. Wizansky, Signals of new physics in the underlying event, *Phys. Rev. D* **80**, 075015 (2009).
- [12] D. Curtin and C. B. Verhaaren, Discovering uncolored naturalness in exotic Higgs decays, *J. High Energy Phys.* **12** (2015) 072.
- [13] H. Cai, H.-C. Cheng, and J. Terning, A quirky little Higgs model, *J. High Energy Phys.* **05** (2009) 045.
- [14] J. Kang and M. A. Luty, Macroscopic strings and “quirks” at colliders, *J. High Energy Phys.* **11** (2009) 065.
- [15] A. Hallin, J. Isaacson, G. Kasieczka, C. Krause, B. Nachman, T. Quadfasel, M. Schlaffer, D. Shih, and M. Sommerhalder, Classifying anomalies through outer density estimation, *Phys. Rev. D* **106**, 055006 (2022).
- [16] J. E. Juknevich, D. Melnikov, and M. J. Strassler, A pure-glue hidden valley I. States and decays, *J. High Energy Phys.* **07** (2009) 055.
- [17] J. E. Juknevich, Pure-glue hidden valleys through the Higgs portal, *J. High Energy Phys.* **08** (2010) 121.
- [18] A. Athenodorou and M. Teper, SU(N) gauge theories in $3+1$ dimensions: Glueball spectrum, string tensions and topology, *J. High Energy Phys.* **12** (2021) 082.
- [19] K. Cheung, W.-Y. Keung, and T.-C. Yuan, Phenomenology of iquarkonium, *Nucl. Phys.* **B811**, 274 (2009).
- [20] G. D. Kribs, T. S. Roy, J. Terning, and K. M. Zurek, Quirky composite dark matter, *Phys. Rev. D* **81**, 095001 (2010).
- [21] S. Nussinov, Quirks and strings attached as the ultimate communication and acceleration devices, [arXiv:1410.6409](#).
- [22] R. Harnik, G. D. Kribs, and A. Martin, Quirks at the tevatron and beyond, *Phys. Rev. D* **84**, 035029 (2011).
- [23] R. Fok and G. D. Kribs, Chiral quirkonium decays, *Phys. Rev. D* **84**, 035001 (2011).
- [24] S. Knapen, H. K. Lou, M. Papucci, and J. Setford, Tracking down quirks at the large hadron collider, *Phys. Rev. D* **96**, 115015 (2017).
- [25] M. Farina and M. Low, Constraining quirky tracks with conventional searches, *Phys. Rev. Lett.* **119**, 111801 (2017).
- [26] S. Knapen, S. Pagan Griso, M. Papucci, and D. J. Robinson, Triggering soft bombs at the LHC, *J. High Energy Phys.* **08** (2017) 076.
- [27] Z. Chacko, D. Curtin, and C. B. Verhaaren, A quirky probe of neutral naturalness, *Phys. Rev. D* **94**, 011504 (2016).
- [28] H.-C. Cheng, S. Jung, E. Salvioni, and Y. Tsai, Exotic quarks in twin Higgs models, *J. High Energy Phys.* **03** (2016) 074.
- [29] J. A. Evans and M. A. Luty, Stopping quirks at the LHC, *J. High Energy Phys.* **06** (2019) 090.
- [30] J. Li, T. Li, J. Pei, and W. Zhang, The quirk trajectory, *Eur. Phys. J. C* **80**, 651 (2020).
- [31] J. Li, T. Li, J. Pei, and W. Zhang, Uncovering quirk signal via energy loss inside tracker, *Phys. Rev. D* **102**, 056006 (2020).
- [32] J. Li, J. Pei, L. Ran, and W. Zhang, The quirk signal at FASER and FASER 2, *J. High Energy Phys.* **12** (2021) 109.
- [33] M. Borsato *et al.*, Unleashing the full power of LHCb to probe stealth new physics, *Rep. Prog. Phys.* **85**, 024201 (2022).
- [34] J. Forsyth, M. Low, C. Tenney, and C. B. Verhaaren, Visible collider signals of natural quirks, [arXiv:2504.02940](#).
- [35] Q. Sha, D. Murnane, M. Fieg, S. Tong, M. Zakharyan, Y. Fang, and D. Whiteson, Learning to reconstruct quirky tracks, [arXiv:2410.00269](#).
- [36] J. L. Feng, J. Li, X. Liao, J. Ni, and J. Pei, Discovering quirks through timing at FASER and future forward experiments at the LHC, *J. High Energy Phys.* **06** (2024) 197.
- [37] J. Li, X. Liao, J. Ni, and J. Pei, Detection prospects of long-lived quirk pairs at the LHC far detectors, *Phys. Rev. D* **109**, 095005 (2024).
- [38] B. Fuks and H.-S. Shao, QCD next-to-leading-order predictions matched to parton showers for vector-like quark models, *Eur. Phys. J. C* **77**, 135 (2017).
- [39] G. T. Bodwin, E. Braaten, and G. P. Lepage, Rigorous QCD analysis of inclusive annihilation and production of heavy quarkonium, *Phys. Rev. D* **51**, 1125 (1995); **55**, 5853(E) (1997).
- [40] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. S. Shao, T. Stelzer, P. Torrielli, and M. Zaro, The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations, *J. High Energy Phys.* **07** (2014) 079.
- [41] T. Sjöstrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten, S. Mrenna, S. Prestel, C. O. Rasmussen, and P. Z. Skands, An introduction to PYTHIA 8.2, *Comput. Phys. Commun.* **191**, 159 (2015).
- [42] R. D. Ball *et al.*, Parton distributions with LHC data, *Nucl. Phys.* **B867**, 244 (2013).
- [43] ATLAS Collaboration, Performance of the ATLAS inner detector track and vertex reconstruction in the high pileup LHC environment, CERN, Report No. ATLAS-CONF-2012-042, 2012.
- [44] ATLAS Collaboration, Measurements of observables sensitive to colour reconnection in $t\bar{t}$ events with the ATLAS detector at $\sqrt{s} = 13$ TeV, *Eur. Phys. J. C* **83**, 518 (2023).
- [45] M. Buchkremer, G. Cacciapaglia, A. Deandrea, and L. Panizzi, Model independent framework for searches of top partners, *Nucl. Phys.* **B876**, 376 (2013).
- [46] M. Buchkremer, G. Cacciapaglia, A. Deandrea, and L. Panizzi, VLQ UFO, <https://feynrules.irmp.ucl.ac.be/wiki/VLQ> (accessed: 2023-06-22).

- [47] J. Barron, D. Curtin, G. Kasieczka, T. Plehn, and A. Spourdalakis, Unsupervised hadronic SUEP at the LHC, *J. High Energy Phys.* **12** (2021) 129.
- [48] G. Kasieczka, B. Nachman, and D. Shih, R&D dataset for LHC Olympics 2020 anomaly detection challenge, (2019), [10.5281/zenodo.4536377](https://zenodo.org/record/4536377).
- [49] J. de Favereau, C. Delaere, P. Demin, A. Giammanco, V. Lemaître, A. Mertens, and M. Selvaggi, DELPHES 3, A modular framework for fast simulation of a generic collider experiment, *J. High Energy Phys.* **02** (2014) 057.
- [50] ATLAS Collaboration, Search for new resonances in mass distributions of jet pairs using 139 fb^{-1} of pp collisions at $\sqrt{s} = 13\text{ TeV}$ with the ATLAS detector, *J. High Energy Phys.* **03** (2020) 145.
- [51] ATLAS Collaboration, Software performance of the ATLAS track reconstruction for LHC run 3, *Comput. Software Big Sci.* **8**, 9 (2024).
- [52] M. Sas and J. Schoppink, Event shapes and jets in e^+e^- and pp collisions, *Nucl. Phys.* **A1011**, 122195 (2021).
- [53] CMS Collaboration, Search for soft unclustered energy patterns in proton-proton collisions at 13 TeV, *Phys. Rev. Lett.* **133**, 191902 (2024).
- [54] A. Paszke *et al.*, PYTORCH: An imperative style, high-performance deep learning library, *Advances in Neural Information Processing Systems* 32, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Curran Associates, Inc., Red Hook, NY, 2019), pp. 8024–8035.
- [55] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [56] E. M. Metodiev, B. Nachman, and J. Thaler, Classification without labels: Learning from mixed samples in high energy physics, *J. High Energy Phys.* **10** (2017) 174.
- [57] M. Germain, K. Gregor, I. Murray, and H. Larochelle, MADE: Masked autoencoder for distribution estimation, in *Proceedings of the 32nd International Conference on Machine Learning Research, Lille, France—Vol. 37* (PMLR, 2015), pp. 881–889, <https://proceedings.mlr.press/v37/germain15.html>.
- [58] M. Sommerhalder, J. Birk, L. Moureaux, O. Amram, and C. Ewen, SK_CATHODE, https://github.com/uhh-pd-ml/sk_cathode (accessed: 2025-05-15).
- [59] E. Bingham, J. P. Chen, M. Jankowiak, F. Obermeyer, N. Pradhan, T. Karaletsos, R. Singh, P. Szerlip, P. Horsfall, and N. D. Goodman, Pyro: Deep universal probabilistic programming, *J. Mach. Learn. Res.* **20**, 1 (2018), <https://jmlr.org/papers/v20/18-403.html>.
- [60] E. Parzen, On estimation of a probability density function and mode, *Ann. Math. Stat.* **33**, 1065 (1962).
- [61] G. Bickendorf, M. Drees, G. Kasieczka, C. Krause, and D. Shih, Combining resonant and tail-based anomaly detection, *Phys. Rev. D* **109**, 096031 (2024).
- [62] CMS Collaboration, Model-agnostic search for dijet resonances with anomalous jet substructure in proton–proton collisions at $\sqrt{s} = 13\text{ TeV}$, *Rep. Prog. Phys.* **88**, 067802 (2025).