

The Intrinsic Dimension of Collider Events and Model-Independent Searches in 100 Dimensions

Raffaele Tito D’Agnolo,^{1,2} Alfredo Glioti,³ Gabriele Rigo,¹ Alessandro Valenti⁴

¹*Université Paris-Saclay, CNRS, CEA, Institut de Physique Théorique, F-91191 Gif-sur-Yvette, France*

²*Laboratoire de Physique de l’École Normale Supérieure, ENS, Université PSL, CNRS, Sorbonne Université, Université Paris Cité, F-75005 Paris, France*

³*INFN Sezione di Roma, Piazzale Aldo Moro 2, I-00185 Rome, Italy*

⁴*Department of Physics, University of Basel, Klingelbergstrasse 82, CH-4056 Basel, Switzerland*

ABSTRACT: The phase space of hadron collider events spans hundreds of dimensions, generating an intricate geometry that we are just starting to explore. The number of possible new physics signals is exponential in the number of dimensions and detecting all of them is currently impossible for any human or artificial intelligence. In this work we introduce a method to search for new physics model-independently in this high-dimensional space. It is based on the measurement of the most basic property of the manifold of collider events, its dimensionality. Our proposed technique does not suffer from a look-elsewhere effect that grows exponentially with the number of dimensions of the dataset, and by construction is insensitive to energy scale uncertainties. We illustrate its potential by finding new physics in simulated events with hundreds of phase space dimensions, taking as input single particles rather than jets. This study sets the stage for new model-independent search strategies based on global properties of collider data manifolds.

Contents

1	Introduction	1
2	Intrinsic Dimension from Nearest Neighbors	4
2.1	Properties of NNID Estimators at finite N	7
2.2	Comparison with Existing Work	8
2.3	The Swiss-Roll as a Toy Example	10
3	A Model-Independent New Physics Test	12
3.1	Statistical Test	13
4	Flexibility, Robustness, Systematics and No Curse of Dimensionality	15
4.1	Leptonic Decays of Charginos and Neutralinos	16
4.2	Leptonic Z'	24
4.3	Semileptonic Decays of Charginos and Neutralinos	28
4.4	Summary of Results	30
4.4.1	Qualitative Comparisons with Other Search Strategies	31
5	Conclusion and Outlook	33
A	Empirical T^2 Distribution, Validation and PCA	35
B	More Details on the MSSM Analyses	38

1 Introduction

Proton–proton collisions at the Large Hadron Collider (LHC) fill the detector with hundreds of particles, populating a highly complex manifold¹ in their phase space. Identifying all potential new physics signals within this manifold is extremely challenging and goes beyond what traditional analysis strategies or novel machine-learning approaches can reliably achieve. In this work we propose a methodology that allows to take a first step in the search for new physics with hundreds or even thousands of observables, when the signal model is unknown.

Our approach is based on the measurement of the number of intrinsic dimensions of phase space. To derive this quantity from data we need to define a metric on this space. In this work we use the energy mover’s distance (EMD) proposed in [1], which connects to several known collider observables [2] and allows a transparent interpretation of our results

¹Without any infrared regulator this is not a real manifold. Its dimension is unbounded due to the emission of infinitely many soft photons. However the detector itself regulates the divergence. We thank J. Thaler for pointing this out to us.

in terms of measurable energy scales. We compute the intrinsic dimension of the data manifold using a nearest-neighbors estimator [3–6] described in Section 2, and we refer to the resulting observable as NNID (Nearest Neighbor Intrinsic Dimension) to distinguish it from other existing ways to measure the intrinsic dimension of a dataset (see [7, 8] for a review).

Our proposal has been developed in the context of a growing interest in the geometry of collider events. New observables were presented in [2, 9, 10], a series of optimal transport distances between energy flow distributions were introduced in [11–13] and a number of applications have followed. In particular, a different notion of intrinsic dimension, compared to the one used in this paper, was applied to jets in the CMS open data in [14, 15] and other strategies for new physics searches were discussed in [16–18]. We comment in detail in Section 2 on the differences between our notion of intrinsic dimension, the correlation dimension used in [1, 14, 15] and other estimators for the dimension of a data manifold.

Additionally, the EMD has been used to speed-up collider computations [11, 12, 19, 20] and classify jets [21]. A covariant description of N -body phase space has been introduced in [22] and has led to the definition of a Riemannian metric on this manifold in [23, 24].

From this discussion it is clear that we are not the first ones to take a “global look” at collider events, but we still want to focus on the most basic property of the phase space manifold, its dimensionality. There are multiple reasons to make this choice. First of all, to the best of our knowledge, we are the first to use the intrinsic dimension of the data to search for new physics. This approach has several qualitative advantages, that we detail below, over existing model-independent searches in high-dimensional spaces.

Besides, the NNID that we discuss in this paper, which is well-known in computer science [3, 5, 7, 25–31] and statistical physics [32, 33], can be computed from a maximum-likelihood procedure founded on solid statistical principles, with well-understood convergence properties [34].

Finally, in its simplest form the NNID is dominated by the shortest distances in the dataset as opposed to the correlation dimension used in [1, 14, 15] which is dominated by the largest distances. Furthermore, one can generalize its simplest definition to explore the kinematics of the phase space manifold at different energy scales.²

After introducing the intrinsic dimension and its properties in Section 2, in Section 3 we build a test for new physics based on it, which includes also information on the average distance between neighboring points in phase space. In Section 4 we find that our test is sensitive to a wide variety of signals, to the point that it can be considered a model-independent search for new physics. Any signal that we have tested has a characteristic shape for its intrinsic dimension as a function of energy which differs from the Standard Model (SM), even when the kinematics of the events are similar.

A large number of model-independent new physics search strategies exists in the literature [16, 35–142]. Most of them have been developed in the last decade in response to the tension between our theory priors and LHC results. Theoretical considerations (as the

²This can be done also with correlation dimensions [1, 14, 15], but we find in toy models that the NNID better captures the short-distance physics.

electroweak hierarchy problem) led us to expect new physics at the LHC, but hundreds of experimental searches have measured only the Standard Model (SM).

There are several qualitative differences between our approach and other existing ideas. To better understand the main difference, it is useful to divide existing techniques into two broad categories. The first one, developed for instance in [37, 45, 64], consists in choosing a fixed analysis strategy which is independent of the dimensionality of the dataset. A concrete example could be: “Take the two leading muons in the event and search for a bump in their invariant mass”. This example greatly oversimplifies the ideas in [37, 45, 64], but captures their two basic features: they do not suffer from a look-elsewhere effect that increases exponentially with the number of dimensions, but they are insensitive to large classes of signals. The second category, which includes [38, 50, 99], lives at the opposite end of the spectrum. It is sensitive to essentially any signal in the dataset, but suffers from a correspondingly large look-elsewhere effect. Techniques based on autoencoders, such as [41, 42, 78, 139] or a pre-training dimensionality reduction step [141–145] can live in either one of the two categories, depending on their implementation. If the number of features extracted in the dimensionality reduction step is increased with the input dimension³, then we have a strategy in the second category, with an exponential increase of the look-elsewhere effect. If instead this number is fixed, the analysis strategy falls in the first category, i.e. the look-elsewhere effect is kept under control, but at the price of reducing the number of signals that one is sensitive to. We explain this point in more detail in Section 4.4.

Our methodology is conceptually similar to the first category of techniques, since it is a fixed analysis strategy, independent of the dimensionality of the dataset, and does not suffer from a large look-elsewhere effect. The strategy is essentially the comparison of the NNID, and of distances in phase space, between different datasets as a function of energy. We are choosing *ab initio* a handful of fixed observables and using them to distinguish signal from background, so by construction we are not sensitive to any possible signal in the dataset. However, this methodology ends up being much more flexible than most existing techniques, because the NNID captures any subtle difference in kinematics between the SM and new physics, well beyond “bumps”, “tails” or other features that one typically looks for, as discussed in Section 4. The approaches in [141, 142] are quite similar in their final goal, but quite different in terms of implementation. The observables used there are selected using contrastive learning from a higher-dimensional set of input features defined on a event-by-event basis, while we rely on global observables that are only defined on the whole dataset.

A second unique feature of our new physics test is its robustness against an important set of systematic errors. As we discuss in Section 2 the NNID is measured from a set of ratios of dimensionful quantities. This makes it completely insensitive to energy scale uncertainties. Additionally, it involves an average over the dataset and this strongly reduces the impact of energy and angular resolution effects. We demonstrate these properties in Section 4.

It is important to point out that our search strategy is not designed for precision

³This is done to reliably describe the null hypothesis i.e. the SM.

searches in the sense of high energy physics, since we need at least one dimension of the dataset to be observably distorted compared to the SM prediction. We are not looking for a part in 10^5 effect on top of a SM distribution that has already been studied in great detail, but rather for signals with large significances. Nonetheless our technique is designed to solve the small signal-over-background problems in high-dimensional spaces that are relevant to high energy physics, where the signal can be rare, weakly expressed and embedded in the data manifold.

We want a framework to analyze datasets with hundreds or thousands of phase space dimensions, which can go beyond clustering jets and, if possible, even reconstructing tracks. Large signals (in terms of statistical significance) could still be hiding in these spaces, since at the LHC we essentially reconstruct what we know, i.e. well-behaved tracks⁴ of particles with charge $\sim \pm 1$ without long-range correlations that are then clustered into well-behaved anti- k_T jets. Any light new physics creating correlations between tracks, jets or constituents of different jets, or deforming them could still be hiding in the data⁵. We consider the intrinsic dimension the natural first step to detect these signals. In this work, as a first test and to develop some analytical intuition, we study the performance of this technique on more traditional signals, but its real goal is to find new physics that we might not even have imagined yet.

2 Intrinsic Dimension from Nearest Neighbors

Consider a sample of D -dimensional vectors $\{\mathbf{x}_l\}_{l=1}^N$, where each vector is drawn from a probability distribution having support on a submanifold of dimension $d \leq D$. Our goal is to estimate d from the data without prior knowledge of their probability distribution. Here, we focus on nearest-neighbor based estimation methods [3–6], although other approaches are possible (see [7, 8] for a review). We choose nearest-neighbor estimators because they are grounded on a firm statistical foundation, allowing for exact asymptotic results and a simple interpretation.

We assume that *locally* the probability density function (pdf) of our dataset on the support of dimension d is constant. In practice we are taking a small neighborhood of a point in the dataset and keeping only the constant term in a Taylor expansion of the pdf. We then imagine to have enough points in the dataset that this constant value does not fluctuate appreciably when we choose different starting points. Notice that this assumption is needed if we want to measure d , but it is not needed for our new physics search strategy in Section 3. In that case we want to identify an observable that distinguishes signal from background, but does not necessarily have to be equal to d , even in expectation value. In practice, when searching for new physics, we compute a property of the manifold that converges to d when $N \rightarrow \infty$, but in some cases we purposefully stay away from this asymptotic limit. The $N \rightarrow \infty$ limit sees any contribution to d including spurious dimensions introduced by the finite momentum and angular resolution of the detector. In a realistic setting, with finite detector resolution and systematic uncertainties, the “true”

⁴For an alternative approach to the problem which seeks to reconstruct non-helical tracks see [146].

⁵Although a number of efforts in this direction has been made in the context of specific models [147–153].

dimension d of the manifold also includes these effects. As a simple illustration, consider $2 \rightarrow 2$ events with on-shell final-state particles. When the initial energies are fixed, as at a lepton collider, the kinematics span only two dimensions in phase space so $d = 2$. At a hadron collider, by contrast, the initial parton energies are not fixed, and the same final state spans four dimensions so⁶ $d = 4$. However a finite momentum resolution of our detector makes $d \rightarrow 6$. If the particles have momenta much larger than the detector resolution, these noise directions are significantly smaller than the physical ones, and one may avoid the asymptotic small-scale limit of the estimator to reduce detector effects on the measurement of d . More details on this point are given in Section 2.1.

For now we set aside this discussion and focus on the measurement of d in the limit where the pdf ρ is constant. In this limit the number of points falling within a sphere of radius r with center $\mathbf{x}_k$ can be thought of as a realization of a homogeneous Poisson process with constant intensity function ρ , i.e. the probability of finding m points within the sphere of volume $V(r)$ is

$$P(m; r, \rho) = (\rho V(r))^m \frac{\exp(-\rho V(r))}{m!},$$

$$V(r) = \frac{\pi^{d/2} r^d}{\Gamma(d/2 + 1)}.$$
(2.1)

The number of points in the hyperspherical shells between any consecutive neighbors are independent and identically distributed exponential random variables, with constant parameter ρ [5, 154].

To remove any dependence on the unknown dimensionful variable ρ , one can consider the dimensionless ratio between the distance of point $\mathbf{x}_k$ to its j -th and i -th nearest neighbors ($j > i$),

$$\mu_{k,ij} \equiv \frac{r_{k,j}}{r_{k,i}}.$$
(2.2)

Expressing $\mu_{k,ij}$ in terms of the volume of the hyperspherical shells, it is then straightforward to obtain its probability distribution [155],

$$f(\mu_{k,ij}; d) = \frac{d \mu_{k,ij}^{-id-1} (1 - \mu_{k,ij}^{-d})^{j-i-1}}{B(j-i, i)},$$
(2.3)

where B denotes the Euler Beta function. As promised, Eq. (2.3) does not depend on the unknown variable ρ , but it has an explicit dependence on the parameter d that we want to estimate. It is therefore natural to construct the negative log-likelihood,

$$L = - \sum_{k=1}^N \log f(\mu_{k,ij}; d) = -N \log d + (1 + id) \sum_{k=1}^N \log \mu_{k,ij} - (j - i - 1) \sum_{k=1}^N \log (1 - \mu_{k,ij}^{-d}),$$
(2.4)

⁶At a lepton collider one may parametrize the configuration by the two angles θ and ϕ of one of the final-state particles, while at a hadron collider one must additionally specify the total energy and rapidity of the two-particle center of mass.

where the sum extends over all points in the dataset. The likelihood in Eq. (2.4) can be minimized to obtain the Maximum Likelihood Estimator (MLE) $\hat{d}_{i,j}$ for the intrinsic dimension d . In principle any choice of i, j is measuring the same value of d and should give the same result, but this is true only asymptotically ($N \rightarrow \infty$), when ρ is truly constant across the dataset. For the specific case $j = i + 1$ the minimization can be carried out analytically, resulting in

$$\begin{aligned}\hat{d}_{i,i+1} &= \frac{N/i}{\sum_k \log \mu_{k,i(i+1)}}, \\ E[\hat{d}_{i,i+1}] &= d \frac{N}{N-1}.\end{aligned}\tag{2.5}$$

In general Eq. (2.4) should be minimized numerically, but the particular choice of j in Eq. (2.5) allows to compute $\hat{d}_{i,i+1}$ from first principles in QCD as was done in [15] for a different notion of intrinsic dimension. For $i = 1$, $\hat{d}_{i,i+1}$ in Eq. (2.5) is the same as the two nearest neighbors estimator originally introduced in [5], although here we obtain it from a ML principle rather than a fit to the empirical CDF of μ .

To better understand our estimator for d for collider events, it is important to recall the definition and the properties of the metric that we have chosen. The EMD introduced in [1] has dimensions of energy and measures the distance between two events $\mathcal{E}_{A,B}$. Here by events we mean collections of three momenta of massless particles, and in practice we use for each particle its p_T , y and ϕ coordinates, since we are mainly considering applications to hadron collider physics. The EMD is the minimum amount of “work” needed to move all the particles in event $\mathcal{E}_A$ so that the event is rearranged into event $\mathcal{E}_B$. Formally it can be obtained via the minimization process

$$\text{EMD}(\mathcal{E}_A, \mathcal{E}_B) = \min_{\{\xi_{km}\}} \sum_{k,m} \xi_{km} \frac{\theta_{km}}{R} + \left| \sum_{k \in \mathcal{E}_A} E_k - \sum_{m \in \mathcal{E}_B} E_m \right|,\tag{2.6}$$

subject to the constraints

$$\xi_{km} \geq 0, \quad \sum_{k \in \mathcal{E}_A} \xi_{km} \leq E_m, \quad \sum_{m \in \mathcal{E}_B} \xi_{km} \leq E_k, \quad \sum_{k,m} \xi_{km} \leq E_{\min}.\tag{2.7}$$

In the above equations we have introduced a dimensionless parameter R which sets the relative importance of the two terms in the EMD. We have also used the angular distance between two particles

$$\theta_{km} \equiv \sqrt{(y_k - y_m)^2 + (\phi_k - \phi_m)^2}\tag{2.8}$$

and the minimal energy $E_{\min} = \min(E_{\mathcal{E}_A}, E_{\mathcal{E}_B})$, where $E_{\mathcal{E}_{A,B}}$ are the total energies in the two events.

One can show [1] that the EMD is a true metric that satisfies the triangle inequality as long as θ_{km} is a metric and $R \geq \theta_{\max}/2$, where $\theta_{\max}$ is the largest angular distance in the dataset. Additionally the EMD is an infrared and collinear (IRC) safe observable [1].

Note that the EMD is not the only possible definition of distance between collider events and more recent proposals, such as [156, 157], reduce the minimization in Eq. (2.6) which is a two-dimensional optimal transport problem, to a one-dimensional problem that can be solved analytically. However in this work we use Eq. (2.6) since it allows a simple physical interpretation of our results which is needed to construct the new physics search strategy in Sections 3 and 4. We leave to future work the exploration of different definitions of distance for the estimation of the NNID.

For us the most relevant feature of Eq. (2.6) is that small values of the EMD correspond to small energy differences and small spatial separations on the collider “celestial sphere”. Similarly, large EMDs correspond to large energy differences and spatial separations.

We are now ready to develop some intuition on our intrinsic dimension estimator $\hat{d}_{i,j}$. From now on we focus on the estimator evaluated for $j = 2i$, that we call $\hat{d}_i$,

$$\hat{d}_i \equiv \hat{d}_{i,2i}. \quad (2.9)$$

This choice gives a small statistical error (i.e. a stable estimate for d) over the largest range in i (and hence energy scales), as verified empirically in [6] and confirmed by our studies of the models in Section 4.

2.1 Properties of NNID Estimators at finite N

Consider the estimator $\hat{d}_i$ in (2.9). As we increase i at fixed N we are looking at the phase space manifold from further and further away, progressively becoming insensitive to small energy differences and spatial separations, so there is some optimal value of i that cuts out all detector effects. For example, a finite resolution in energy can introduce a fictitious small dimension that we might want to get rid of.

If we take i large enough the manifold will appear to us as a point-like object. Therefore, even if $\hat{d}_{i,j}$ is independent of i, j asymptotically (that is for $N \rightarrow \infty$), at fixed N we expect $\hat{d}_{i,j} \rightarrow 0$ as i grows. This can be useful to filter out noise, but also to explore the kinematical features of our events at different energy scales. At the qualitative level, we expect $\hat{d}_{i,j} \rightarrow d$, independent of i, j , when $N \rightarrow \infty$, because in that limit more and more points are added at short distances and any finite choice of i, j is probing infinitesimal distances in the dataset. Therefore in the asymptotic limit $\hat{d}_{i,j}$ converges to the result that a “perfect short-distance observer” would measure.

If we want to retain the dependence of $\hat{d}_i$ on the energy scales in the dataset also in the asymptotic limit, we can compute it at fixed i/N rather than i . We can justify quantitatively the choice of varying i/N rather than i as follows. The pdf of the distance r_l between a point and its l -th neighbor is [155]

$$f(r_l) = \frac{d}{r_l(l-1)!} (\rho \omega_d r_l^d)^l \exp(-\rho \omega_d r_l^d), \quad (2.10)$$

with ω_d the volume of the d -dimensional sphere of unit radius. From Eq. (2.10) one has $E[r_l] \simeq (l/\rho \omega_d)^{1/d}$, so that on average $r_l/r_N \approx (l/N)^{1/d}$. If we imagine that r_N is bounded by the kinematics and varies slowly with N , eventually saturating in the asymptotic limit,

by fixing the *proportions* $i/N, j/N$ we are probing fixed distances in the dataset. This should lead to curves which are compatible between datasets with different number of points and retains the dependence of $\hat{d}_i$ on the energy scales in the dataset also in the $N \rightarrow \infty$ limit. We verify this intuition on the models studied in Section 4.

The last property of $\hat{d}_i$ that we want to highlight and we are going to use in Section 3 is its distribution in the asymptotic limit. In the $N \rightarrow \infty$ limit, the estimators $\hat{d}_{i,j}$ minimizing Eq. (2.4) are expected to follow a multivariate Gaussian distribution under standard regularity conditions. For the case $j = i + 1$, this is easily seen and follows directly from the Central Limit Theorem and Delta method applied to (2.5) [158]. For the vector of estimators $\hat{\mathbf{d}} = (\hat{d}_{i_1, j_1}, \hat{d}_{i_2, j_2}, \dots)^T$ across different (i, j) pairs, joint asymptotic normality follows from multivariate M-estimator theory [159, 160]. Asymptotically $\sqrt{N}(\hat{\mathbf{d}} - \mathbf{d}_0) \rightarrow N(\mathbf{0}, \mathbf{\Sigma})$, where the covariance matrix $\mathbf{\Sigma}$ accounts for correlations arising from shared data points and nearest-neighbor relationships across different (i, j) pairs. This is true under “physically sensible” regularity conditions, including finite second moments of $\log \mu_{k,ij}$, existence of the Jacobian and non-singularity of the Hessian matrix of (2.4) at the true solution $\mathbf{d}_0$.

In all the cases discussed later, we empirically verify the normality of our estimators before using this property in any new physics search.

To summarize, when performing the new physics search described in Section 3 we are going to compute our estimators $\hat{d}_i$ as a function of i/N for two reasons: 1) to make them comparable across different datasets, 2) to retain the dependence on i (and thus on energy scale) also in the $N \rightarrow \infty$ limit. Additionally we are going to verify that the entries of the vector $\{\hat{d}_i\}_{i=1, \dots, N}$ follow a joint normal distribution that we then use to compute significances.

2.2 Comparison with Existing Work

We can compare the NNID to the correlation dimension computed in [1, 14, 15] and other strategies commonly used to measure the dimension of a dataset. First of all we recall the broader context of dimensionality reduction. A number of methodologies [161–166] define a loss function which is essentially a projection error and then try to minimize it. The dimensionality of the sample is inferred from discontinuities in the eigenvalues of the loss function. These ideas face serious difficulties when trying to define the notion of discontinuity for a function (the spectrum of eigenvalues) defined over discrete values.

We take the alternative point of view that the data are sampled from a high-dimensional “intrinsic manifold” and we estimate their dimensionality from that of its tangent space at each point. Fractal methods, as the one used in [1, 14, 15] to measure the *correlation dimension* of jets, are another application of the same logic. Fractal methods rely on the intuition that the number of points in a sphere of radius r should grow as r^d for constant density. Their main drawback is that an accurate estimate of the dimension requires a number of data points that grows exponentially with d , as discussed in [167, 168].

In practice, correlation dimensions have a hard time reproducing the dimension that a “UV observer” with access to the shortest distances in the dataset would measure. This would be the number that we are interested in if we wanted to know how many neurons




Figure 1: Sketch of a Swiss-Roll manifold.

we need in the latent space of an autoencoder, or how many additional degrees of freedom one needs to describe a jet compared to the 3 components of the momentum of a massless parton. The correlation dimension in [1, 14, 15] captures a different and physically interesting property of a dataset. It was shown to measure the logarithmic increase in the number of jet constituents as one decreases the jet energy scale [14]. In the regime considered in [1, 14, 15] it is essentially measuring the entropy in the dataset. The NNID is instead insensitive to the density of points, and thus the entropy of the dataset, as long as the density is constant on the scale defined by the typical distance to the j -th neighbor.

A second comparison worth exploring is about the dependence of the dimension on the typical energy scales in the problem. The correlation dimension of [1, 14, 15] was defined as

$$d_C(r_{\max}) = \frac{d}{d \log r_{\max}} \sum_{1 \leq A \leq B \leq N} \Theta [\text{EMD}(\mathcal{E}_A, \mathcal{E}_B) < r_{\max}] , \quad (2.11)$$

where N is the total number of events and $r_{\max}$ an energy scale. The value of d_C strongly depends on the choice of $r_{\max}$, i.e. the largest distances (energy scales) that we include in its calculation. Instead, the NNID estimator based on the two nearest neighbors ($j = i+1 = 2$) is insensitive to $r_{\max}$. As discussed before, we can introduce a dependence on the energy scale in the NNID by considering different values of i, j , i.e. using farther neighbors of each point to estimate the dimension of the dataset. If we do this, we should recover the intuition discussed above, i.e. at large i, j the dimension goes to 0, as if we were looking at the manifold from further and further away, while for the closest neighbors ($i, j = 1, 2$) we should recover the result that we would expect based on first principles, when one has a perfect knowledge of the dataset. We illustrate these points in the toy example below.




Figure 2: A comparison of the correlation dimension d_C (solid lines) and the NNID (dotted lines) for a Swiss-Roll. The NNID is computed at $i = 1$ and $j = 2$ and by definition is insensitive to cutting the largest distances in the dataset. The parameter δ , defined in the text, quantifies the thickness of the Swiss Roll. The shaded areas give the 95% C.L. interval for the two estimators.

2.3 The Swiss-Roll as a Toy Example

We consider here a toy example consisting of a *Swiss-Roll* manifold to make all our statements more concrete and easy to visualize. The Swiss-Roll is a two-dimensional manifold embedded in a three-dimensional space. It is parameterized by

$$\vec{x}_{\text{SR}} = (t \cos t, t \sin t, z), \quad t \in \left[\frac{3}{2}\pi, \frac{9}{2}\pi \right], \quad z \in [0, 21]. \quad (2.12)$$

We show a sketch in Fig. 1. The (Euclidean) distance between consecutive spiral windings in the $x - y$ plane is 2π , that together with the extension along z sets the typical length scales of the manifold.

We add some noise to this manifold in the form of a random vector extracted from a Normal distribution with mean zero and unit variance, rescaled by a factor δ ,

$$\begin{aligned} \vec{x} &= \vec{x}_{\text{SR}} + \delta \cdot \vec{x}_G, \\ \vec{x}_G &\sim \mathcal{N}(\vec{0}, \text{diag}(1, 1, 1)). \end{aligned} \quad (2.13)$$

In this way, when $\delta = 0$ the manifold is genuinely 2-dimensional, while $\delta \neq 0$ “fills” the spiral and makes it 3-dimensional. Later we will choose $\delta = 0.5, 1$, corresponding to a $\approx 10\%, 20\%$ noise with respect to the typical length scales of the spiral.

We first compare the NNID for the closest nearest neighbors $\hat{d}_1$ to the correlation dimension by computing both under the constraint on the distances $r \leq r_{\text{max}}$. We employ




Figure 3: The intrinsic dimension of a Swiss-Roll for different choices of the noise δ , i.e. its thickness. On the x -axis we show the average distance of the i -th neighbor, on the y -axis the NNID calculated at $j = 2i$. The shaded areas give the 95% C.L. interval for the estimator. We truncate our lines at the value of $\bar{r}_i$ below which there are no more neighbors.

the Euclidean distance, and to obtain confidence bands we use 100 datasets each comprised of 10^4 points. The result, reported in Fig. 2, shows that the NNID (dotted bands) is independent of the largest distances in the dataset and always sees either a 2-dimensional manifold (when $\delta = 0$) or a 3-dimensional manifold (when $\delta > 0$). The correlator dimension d_C (solid bands) at large $r_{\max}$ sees a point-like object, then as $r_{\max}$ decreases it starts resolving more features of the Swiss-Roll and therefore increases. It peaks around $r_{\max} \simeq 6 - 7$, the length scale of the spiral, where it can not yet resolve the substructure of the manifold and sees it as a filled cylinder. Then, as we further lower $r_{\max}$, it starts resolving the spiral structure and slowly converges to the true result.

At very small values of $r_{\max}$ both measurements of the dimension break down, because we do not have enough points in the dataset for a reliable estimate. For example, in the case of the NNID we are starting to cut out of the dataset the nearest neighbors used for the estimate. It is interesting to notice that the NNID knows that it is failing (i.e. the error bands get larger and larger as we decrease $r_{\max}$), while the same is not true for the correlation dimension d_C .

To illustrate the dependence of $\hat{d}_i$ on i that we introduced earlier, we plot in Fig. 3 the NNID estimator against the average distance between a point and the nearest neighbor used in the estimate ($\bar{r}_i$), in turn averaged over our 100 datasets. In practice, when varying $\bar{r}_i$ on the x axis of the plot, we are changing the nearest-neighbor number i in the NNID

estimation (with fixed $j = 2i$). We truncate our lines at the value of $\bar{r}_i$ below which there are no more neighbors.

We observe a dependence of the NNID on distance that resembles the behavior of d_C as a function of $r_{\max}$ in Fig. 2, and the qualitative description of the curve is the same. Interestingly, however, our estimator remains stable even when probing the smallest distances in the dataset, unlike d_C which diverges in that regime. In fact, our estimator correctly converges and stops at $d_{i,2i} = 2$ in the case of no noise and at $d_{i,2i} = 3$ when noise is present.

3 A Model-Independent New Physics Test

The NNID of events with a fixed number of particles is well defined and can in principle be computed on a piece of paper, simply from kinematics. Consider, for example, a generic $2 \rightarrow n$ scattering process with no on-shell mediators and in which the masses of all the particles are known, as well as the 3-momenta of the incoming ones. If we can measure the momenta of $m \leq n$ of the outgoing particles, simple kinematic considerations set $d = \min(3n - 4, 3m)$.

Suppose, however, that some of the final momenta are correlated; for example, because two of the m particles in the final state come from a new intermediate on-shell particle X that decays promptly. This reduces the dimension of the support manifold by one, since the momenta of the decay products satisfy $(p_1 + p_2)^2 = m_X^2$. Hence, the measured NNID in the data will be lower than the naïve SM expectation, potentially revealing the presence of new physics.

In reality, things are much more complicated. Even if we just consider the simple example above, without systematics or detector effects, our new particle X has some finite width Γ_X , and the kinematical constraint $(p_1 + p_2)^2 \simeq m_X^2$ does not eliminate one dimension, it just shrinks it to a size of order Γ_X . If we know Γ_X in advance we can choose the appropriate i, j in the estimator $d_{i,j}$ so that we are insensitive to energies smaller than Γ_X . However, if we want to design a model-independent search strategy, we won't know Γ_X in advance, or even the nature of the kinematical constraints in the Beyond the Standard Model theory which is hiding in the data.

Additionally, every measurement has an associated experimental error. This source of noise introduces fictitious dimensions that make the empirical estimate of the NNID deviate from the kinematical one and introduce new energy scales in the problem. Furthermore, hadron collider events coming from the same high-energy scattering process do not have a fixed number of particles due to pile-up, particles from the underlying event, ISR, FSR, and the stochastic nature of the quantum mechanical processes in jet showers. Even if we could isolate two quarks with exactly the same momentum in two different events, the two reconstructed jets would not contain the same number of particles.

In practice the best strategy to find new physics does not seem to analytically predict the NNID of a given class of events or its dependence on i, j and compare it with the data. We should rather leverage the differences in NNID between different processes as a function of energy scale (or in practice of $i/N, j/N$) to detect new physics.

To do it we first construct a reference sample of events that describe our null hypothesis (i.e. the SM). This can be done in several ways, depending on our final goal:

1. We take a data sample and split it into subsamples that were collected by the detector at different times. One is the reference and by testing its compatibility with the others we can perform data quality monitoring tasks.
2. We exploit the approximate symmetries of the SM, as proposed in [169]. For example we take a data sample that contains at least one high-quality electron and compare it with another sample containing at least one high-quality muon, after correcting for detector effects.
3. We generate the reference sample by simulating the SM using Monte Carlo (MC) methods.

The first strategy is straightforward to implement, but its applicability to new physics searches is limited. The second strategy is purely data-driven and can be used to compare complex datasets without clustering jets. It does not allow to probe every possible new physics signal, but it is still quite flexible, as one can repeat the exercise with all sorts of different datasets (two electrons vs two muons, three electrons vs three muons, and so on). Additionally, we remind the reader that we are not doing a precision search in the sense of high energy physics, but rather looking for signals with small S/B , but potentially large significances, hiding in the large dimensionality of the space, so the systematic error associated with replacing a muon with an electron is a small effect for us.

The last strategy is the most flexible, i.e. the one which is sensitive to the largest number of new physics signals, but requires a reliable MC simulation. We do not view this as a real drawback, for two reasons. On the one hand MC simulation has progressed a lot since the early days of the LHC, to the point that it is routinely used in precision measurements [170, 171]. Secondly, we do not need to simulate the SM with high accuracy, as the dimension of the dataset is a robust quantity, as we demonstrate in Section 4. Its sensitivity to new physics degrades only slightly going from parton-level to detector-level events, it is not affected by momentum scale uncertainties and only mildly by momentum resolution effects, even if at detector level we use hundreds of particles to compute the NNID as opposed to the few partons coming out of the hard collision.

3.1 Statistical Test

Equipped with the information contained in the NNID estimators $\hat{d}_i$ and in the neighbor distances $\bar{r}_i$, one can perform model-independent searches for new physics. Here we consider the simplest possible implementation of such a test, which has the virtue of being easy to interpret and computationally very efficient in the calculation of significances, since it exploits analytic results in the asymptotic limit. In the previous Sections we have described at length the NNID and its properties, but in the test we include also information on the average distance between neighbors $\bar{r}_i$. Its properties have already been studied in [1]. For our purposes it is sufficient to recall that since $\bar{r}_i$ is an average we also expect it to follow a Gaussian distribution asymptotically.

Our test compares NNID and $\bar{r}_i$, as a function of i/N , between collider data and a reference hypothesis (the SM). We consider a single collider dataset, while the reference ensemble is split into smaller subsamples and used to estimate the mean and covariance under the null hypothesis of the asymptotic Gaussian distribution of NNID and $\bar{r}_i$. The test then quantifies the compatibility of the experimental data with the reference distribution. We leave to future work the study of more elaborate approaches that can further optimize this procedure, for instance by feeding the information on $\hat{d}_i$, $\bar{r}_i$, or related quantities into machine-learning algorithms.

For concreteness, let us focus on the NNID. We denote by $\{\mathbf{d}_k\}_{k=1}^n$ a reference set of m -dimensional NNID vectors computed from n datasets, where $n > m$, sampled from the same reference distribution under identical conditions, and by $\mathbf{d}'$ an m -dimensional vector computed from the experimental data, whose compatibility with the reference set we want to test. The m entries of the vectors are populated by varying i/N in the NNID computation.

In Section 2 we assumed the probability density of points to be locally constant and then derived from this assumption a series of properties for the NNID estimator. This was useful to set the stage for the new physics search presented here. However it is important to keep in mind that even when this initial assumption breaks down, our new physics test can still be performed as described in this Section.

Even if ρ is not constant across the dataset, the random vector $\mathbf{d}$ can still be distributed as a multivariate normal random variable with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$, $\mathbf{d} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$. Call $\hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\Sigma}}$ the parameters empirically estimated employing the n datasets associated to the reference distribution, and call $\hat{T}^2$ the would-be χ^2 ,

$$\hat{T}^2 = \frac{n}{n+1} (\mathbf{d}' - \hat{\boldsymbol{\mu}})^T \hat{\boldsymbol{\Sigma}}^{-1} (\mathbf{d}' - \hat{\boldsymbol{\mu}}). \quad (3.1)$$

Under the null hypothesis, namely $\mathbf{d}' \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, the variable $F \equiv T^2(n-m)/(m(n-1))$ follows an F -distribution with $(m, n-m)$ degrees of freedom, $F \sim F_{m, n-m}$ [172].⁷ This can be regarded as a generalization of the familiar Student's t -distribution. Note that in the limit $n \rightarrow \infty$, $T_{n,m}^2 \rightarrow \chi_m^2$.

A p -value for the compatibility between the reference model and the new data can be computed simply as

$$p = P\left(F \geq \hat{F}\right) \quad F \sim F_{m, n-m}. \quad (3.2)$$

and then converted in the more familiar significance $Z = \Phi^{-1}(1-p)$.

This procedure can be seen as a model-independent test of the agreement between the reference distribution and the experimental data. The biggest numerical burden is the computation of the distance matrix between events, which for a two-dimensional optimal transport problem, such as the EMD computation, is of $\mathcal{O}(N^2 D^3)$ for a dataset with N points and D features.

In Section 4 we also consider for the test the averages of the distances $\bar{r}_i$, which by the Central Limit Theorem are approximately normally distributed. In practice, this amounts

⁷Formally, $\hat{T}^2$ follows a *Hotelling* distribution [173], $\hat{T}^2 \sim T_{n-1, m}^2$.

to repeating the same procedure above but for the $\{\bar{\mathbf{r}}_k\}_{k=1}^n$ vectors, where $\bar{\mathbf{r}}$ denotes the vector of the various $\bar{r}_i$ for a given dataset. We choose to carry out the two tests, on NNID and $\bar{r}_i$, separately, to discuss their different physical properties, but to maximize performances one should concatenate the two vectors, for instance as outlined in Section A.

We stress that the reliability of the resulting p -value depends on the validity of the assumptions underlying the test; most crucially, that Eq. (3.1) follows an $F_{m,n-m}$ distribution under the null hypothesis. In practice, this approximation can break down when the NNID vector components become highly correlated, or when their number m approaches the number of reference samples n . In these regimes, the empirical covariance matrix $\hat{\Sigma}$ acquires small eigenvalues, making its inversion numerically unstable and causing the test statistic to fluctuate more than expected from the F -distribution.

A standard approach to mitigate this problem is *Principal Component Analysis* (PCA). We project the NNID vectors onto the subspace spanned by the first M principal components, i.e. the directions with largest eigenvalues, and discard components associated with small eigenvalues that are most prone to numerical noise. The optimal number of retained components M is determined by performing a Cramér–von Mises (CvM) test on the empirical distribution of T^2 obtained via a leave-one-out procedure under the null hypothesis: we choose the largest M for which the CvM statistic indicates acceptable agreement with the theoretical F distribution. As an additional qualitative cross-check, we inspect quantile–quantile (Q–Q) plots of the resulting T^2 values. Full details of this procedure are given in Section A, where we show its application to the examples of Section 4.

In the next Section we perform this novel new physics test on a series of benchmark models and verify its main properties:

1. **Flexibility:** Sensitivity to diverse Beyond the Standard Model scenarios with different kinematics.
2. **Robustness** (to showering and detector effects): The performances do not change appreciably when going from the parton-level to the detector-level simulation.
3. **Robustness to Systematics:** Systematic errors can be incorporated in a statistically well-defined way and the test is insensitive to energy scale errors.
4. **Protection from the Curse of Dimensionality:** This feature is already present by construction in our test, but we demonstrate it explicitly in two examples, going up to ~ 100 dimensions.

4 Flexibility, Robustness, Systematics and No Curse of Dimensionality

In this Section we show the performances of our new physics search strategy on six benchmark models. The results are a first proof-of-principle test of our methodology and therefore we only consider the main background in each search. In Section 4.4, we summarize the significance of these results for each of the four properties listed at the end of the previous Section (Flexibility, Robustness to Detector Effects, Robustness to Systematics, No Curse of Dimensionality).

To test the performances of our methodology in a controlled setting, we apply our new physics test following a procedure that is not optimal as far as finding new physics goes, but it is the most transparent for testing. More concretely, we start by choosing a final state, for example two muons and one electron or a W and a Z boson. Then we identify the irreducible SM backgrounds and design a set of cuts to eliminate all other backgrounds. Only then we select a signal that can produce the chosen final state. The cuts are designed on the SM backgrounds, assuming no knowledge of the signal. Our cuts are chosen to simplify the analysis, i.e. to select a single SM background process. For this reason they often penalize our signal. An example of this is our main benchmark (called MSSM-3 in what follows), that leads to very boosted particles in the final state. The efficiency of our cuts on MSSM-3 is not optimal, as they were designed on WZ production in the SM. Additionally, the cuts are designed to be exclusive, since their goal is to reduce the list of backgrounds relevant to the analysis, but the real purpose of our techniques is to find unexpected new signals and it should be applied on final states that are as inclusive as possible. We leave this to future work in collaboration with LHC experiments. Our choice simplifies the task of interpreting and validating the NNID and $\bar{r}_i$ curves and develop some physics intuition on them.

4.1 Leptonic Decays of Charginos and Neutralinos

As a first example, we focus on a purely leptonic final state arising from the production of new weakly coupled particles, namely a set of charginos and neutralinos. The model consists of an electrically charged fermion $\chi_1^\pm$ and two neutral fermions χ_2^0, χ_1^0 , with R -parity conserving couplings to the SM. We take χ_1^0 , to be the lightest fermion in the spectrum, so R -parity makes χ_1^0 stable.

We analyze the process $pp \rightarrow \chi_1^+ \chi_2^0$, with on-shell χ_1^+ and χ_2^0 decaying as $\chi_1^+ \rightarrow W^+ \chi_1^0$ and $\chi_2^0 \rightarrow Z \chi_1^0$, respectively. The final state we consider contains a pair of χ_1^0 's, a positron-neutrino pair from $W^+ \rightarrow e^+ \nu_e$, and a muon-antimuon pair from $Z \rightarrow \mu^+ \mu^-$; the full diagram is depicted in Fig. 4. The leading SM background consists of $pp \rightarrow W^+ Z$ events, with the same W^+ and Z decays as above.

This is a MSSM-inspired simplified model, in which, however, we let the total cross section vary freely to study the behavior of the NNID with respect to the number of BSM events. Therefore the detailed admixture of Wino, Higgsino and Bino that makes up our particles is not relevant to the search.

Simulation and Event Selection. We simulate the process with $\sqrt{s} = 13.6$ TeV center of mass energy and choose a reference luminosity of 200 fb^{-1} . The simulation of the signal and background is carried out at parton level using MADGRAPH5_AMC@NLO [174] with the MSSM_SLHA2 model [175]. Hadronization and showering are performed with PYTHIA 8 [176], and detector effects are simulated with DELPHES [177].

The data is split into several samples (see Results paragraph below), and for each the NNID is obtained as described in Section 2. Specifically, we first compute the EMDs among events and then obtain the NNID by minimizing Eq. (2.4). The EMDs are computed with




Figure 4: Leading diagram for the leptonic MSSM signal in Section 4.1.

the Python package **EnergyFlow** [1, 2, 14, 178–181], using $\beta = 1$ and $R = 3$, and taking as input the (p_T, η, ϕ) of each particle (all treated as massless).

Since our new physics test will be applied at both parton and detector level to assess the difference in performance, we report here the details of the event simulation, reconstruction, and selection in both cases:

- i) Parton level:* For each event, the muon, antimuon, and positron from the hard collision simulated in MADGRAPH are directly used in the computation of the NNID. The leptons are required to satisfy $p_T > 10$ GeV and $|\eta| < 2.5$. Furthermore, we select isolated leptons by requiring the total p_T of other particles within a cone of radius $R = 0.3$ around each lepton to be smaller than 0.12 of the lepton’s p_T for electrons and 0.25 for muons. After all cuts we find an overall efficiency $\epsilon \simeq 0.44$ for the SM. In terms of the number of events, taking the results of [182, 183] for the total cross section and $W^\pm$ polarization, and the W^+, Z branching ratios from [184], about 12,000 SM events are expected;
- ii) Detector level:* We employ the default CMS card [185] in DELPHES, with the only modification being an adjusted lepton isolation requirement. Specifically, we apply the same isolation criteria as at parton level. Jets are clustered with the anti- k_T algorithm with radius $R = 0.5$. Events selected for the analysis must contain exactly a muon, an antimuon, and a positron with $p_T > 10$ GeV and $|\eta| < 2.5$. Additional jets (from initial and final state radiation) are allowed if $p_{T,j} < 40$ GeV. For the EMD computation only the three leptons are used. These criteria lead to approximately 4800 selected SM events; the efficiency relative to the parton level selection is about 43% for the SM.

Results. To get some intuition on the physics of the $\bar{r}_i$ and NNID search strategy, we first inspect the NNID curve as a function of i/N and $\bar{r}_i$ in pure SM and BSM datasets at parton level. For the BSM model, here and in the following we fix $m_{\tilde{\chi}_1^0} = 150$ GeV and $m_{\tilde{\chi}_2^0} = m_{\tilde{\chi}_1^\pm}$, with $m_{\tilde{\chi}_2^0, \tilde{\chi}_1^\pm} = 250, 500, 1000$ GeV, which we refer to as MSSM-1,2,3, respectively.




Figure 5: Parton level NNID curves as a function of $2i/N$ (left panel) and the average distance to the i -th neighbor (right panel) in the SM (black) and in three MSSM-like benchmark models described in the text. The curves are obtained from 100 samples. In the left panel the central line represents the median, while the band is the 68% confidence interval. In the right panel we show the 68% confidence interval for both axes as error bars.

The NNID curve is computed taking $j = 2i$ and $m = 200$ points logarithmically spaced between $i = 1$ and $i = \lfloor N/2 \rfloor$, with $N = 10^4$ points.⁸ This operation is repeated for 100 signal and background samples, to obtain confidence level intervals for the NNID. We also compute the average distances $\bar{r}_i$ for each i and plot the NNID against it. The results are given in Fig. 5.

In the SM we expect the NNID estimate to asymptote to either 8 or 9 for small values of i/N , depending on whether we can resolve the Z boson width or not. The theoretical counting of the independent degrees of freedom that justifies our expectation goes as follows. We are only including the two muons and the positron in the calculation of the NNID, and we treat them as massless, so we have at our disposal 9 variables. The 2 kinematical constraints from momentum conservation in the plane transverse to the collision do not hold because we are not including the neutrino in the calculation. The only additional kinematical constraint that we can impose comes from the Z being on-shell. The W does not count because again we are not including the neutrino in our calculation.

The result in Fig. 5 (left) respects perfectly this estimate, falling between 8 and 9, but closer to 8, meaning that we can barely resolve the Z width with the number of events at our disposal.

For the MSSM models, we also ignore missing energy and only consider particles to which we can associate a three-component momentum. With this assumption the counting proceeds in exactly the same manner as in the SM. Hence, also here we expect a saturation of the NNID between 8 and 9. In Fig. 5, we show that this counting is approximately respected for the MSSM-1 benchmark. The MSSM-2 and MSSM-3 benchmarks give a

⁸At this stage we are only interested in the qualitative behavior of the NNID and use a large number of events for illustrative purposes.

very different result; the latter, in particular, shows the biggest deviation compared to our theoretical counting. The muons and positron-neutrino pairs arising from the Z and W decays are very boosted, since $m_{\chi_2^0} \gg m_{\chi_1^0}, m_{Z,W}$. This leads to a strong collimation of their momenta, which may render it difficult for the estimator to disentangle them, thereby reducing the measured dimension of the manifold. This is evident from the right panel of Fig. 5, in which we can appreciate the boost in the form of much bigger typical distances $\bar{r}_i$ for the MSSM-2,3 cases.

Overall, the NNID shows a remarkable sensitivity to small kinematical changes in the models. For example MSSM-1 and the SM are quite close kinematically, but have different NNID curves. Interestingly, as the proportion i/N increases, the difference persists or increases. This qualitative behavior shows that we can use the NNID as a test for new physics, at least for these specific models.

Fig. 5 also highlights the topological nature of the NNID (i.e. its insensitivity to absolute scales). The distances $\bar{r}_i$ are sensitive to the typical energy scale of the events and show clearly that MSSM-2 and MSSM-3 live at much higher scales compared to the SM. This is only partially reflected in the NNID which registers the collapse of some of the dimensions due to the MSSM events' boost, but is also losing some of the kinematical information contained in the $\bar{r}_i$, as one can see by comparing the two panels of the Figure. We see this both as an issue and a positive feature. It is an issue because of the loss of information, but it is a feature because the NNID is completely orthogonal to traditional strategies that target high-energy tails or bumps. Additionally the topological nature of the NNID makes it much less sensitive to a series of effects that we describe below (for example errors on p_T scale and resolution, extra input dimensions identical for signal and background).

To solve the issue we enlarge our new physics test to include also the *average* distance $\bar{r}_i$ between neighbors, as discussed in the previous Section. We choose the average over the median because of asymptotic normality and to give a bigger weight to the small number of highly energetic new physics events as opposed to the large number of low-energy events in the SM. We build a vector of $\bar{r}_i$'s and treat it as the vector of NNID 's in the test. We also performed a combined test concatenating the two vectors, but here we only show the results of the two separate tests to better understand their strengths and weaknesses.

To proceed with the new physics test outlined in Section 3.1, we simulate 200 SM NNID curves, now with 20 logarithmically spaced i/N points, to speed up computations. We collect in Section A the comparison of our test statistic under the null hypothesis with the theoretical T^2 distribution.

Each SM sample has a variable number of events extracted from a Poisson distribution whose mean μ_{SM} is itself extracted from a Gaussian distribution with variance $\delta\mu_{\text{SM}}$. Here $\delta\mu_{\text{SM}}$ accounts for the systematic error on the total number of SM events, given by the sum in quadrature of the error on the luminosity [186] and that on the cross section of our main background [182]. The error on the cross section $\delta\sigma_{WZ}/\sigma_{WZ} \simeq 4\%$ dominates over the error on the luminosity $\delta\mathcal{L}/\mathcal{L} \simeq 1.6\%$ and the Poisson statistical error, that for our choice of luminosity is around 1%.

Similarly, we compute 50 SM+BSM curves, with the number of SM events determined

as above. On top of the SM events, we add $\text{Poisson}(\mu_{\text{BSM}})$ BSM events whose “nominal” rate is fixed by the total cross-section ratio $\sigma_{\text{BSM}}/\sigma_{\text{SM}}$. Note that in the new physics test the effective rate used is the nominal one rescaled according to the SM and BSM efficiencies; in particular, the efficiency of our parton-level cuts on these three benchmarks is $\simeq 0.7, 0.82, 0.73$ for MSSM-1,2,3, and roughly 33% pass also the detector-level cuts for MSSM-3. This means that, for example, the effective S/B at parton level for MSSM-3 is about 1.7 times larger than the nominal rate. We then compare the NNID and $\bar{r}_i$ curves using the statistical test described in Section 3, including the selection of the best number of points using PCA (see Section A). This set of SM+BSM samples allows us to obtain confidence level intervals for the significance of our statistical test.

In Fig. 6 we show the result of the new physics test on the benchmark models that we have just described. Both of our signal-agnostic tests outperform a simple counting experiment by orders of magnitude for MSSM-3, while they cannot meaningfully distinguish MSSM-1 from the SM. In this setup, the test based on the $\bar{r}_i$ variables achieves significantly better sensitivity than the one using the intrinsic dimension: the strong boost of the MSSM events amplifies the $\bar{r}_i$ values, leading to a much clearer separation between SM and MSSM. However this is not always the case as we are going to see in the next Section for a different type of signal.

Detector effects reduce the performance of the tests, as shown in the right panel of Fig. 6. This reduction is mainly due to the lower event selection efficiency at detector level, which effectively halves the available events in the reference samples and consequently increases the statistical fluctuations of the points along the NNID and $\bar{r}_i$ curves. As expected from general significance tests, this loss of statistics weakens the overall discriminating power.

For comparison, the dashed line in the figures shows the result of the “trivial” counting test, in which the p -value is obtained solely from the total number of events N and the corresponding significance is computed as $Z = \mu_{\text{BSM}}/\sqrt{\mu_{\text{SM}} + \delta\mu_{\text{SM}}^2}$, where $\delta\mu_{\text{SM}}$ takes into account the uncertainty in the SM cross section and luminosity knowledge mentioned in the previous paragraphs.

No Curse of Dimensionality. Our methodology is inherently free from a look-elsewhere effect that grows exponentially with the number of dimensions. This is because we apply to each dataset a fixed analysis strategy, independent of its phase-space dimensionality. Consequently, we expect the performance of our method to remain stable when adding “spectator” particles that are distributed identically in the SM and in the new physics signal.

We demonstrate this in Fig. 7, where we extend the MSSM-3 benchmark by including up to 60 additional dimensions corresponding to spectator particles. For each such particle, the values of (p_T, η, ϕ) are randomly and independently sampled from the distributions of SM muons. The results show that the performance of the NNID test on detector-level events remains unchanged when up to 15 spectator dimensions are added. By contrast, for other model-independent searches the p -value would degrade exponentially with dimensionality [50].




Figure 6: Left Panel: Parton level Significance of the NNID and $\bar{r}_i$ new physics search versus the fraction of BSM events in the sample, for the three MSSM leptonic benchmarks described in Section 4.1. The solid and dashed lines are the median significances and are shown with the corresponding 68% CL band. The black dotted line corresponds to the simple counting significance, $Z \sim S/\delta B$. Right Panel: Comparison of the parton level and detector level significances for the MSSM-3 benchmark model.

A performance decrease becomes visible only when adding 30 ($m = 10$) or 60 ($m = 20$) extra dimensions, though still at the level of about one sigma. This behavior is expected: in high dimensions, the distance between nearest neighbors grows with d (the growth scaling as $\sqrt{d}$ for an Euclidean distance, for example), and once the typical inter-event distance exceeds the smallest intrinsic scale of the dataset, the search can no longer resolve it. For example, in comparing SM and MSSM-3 events, if d is increased sufficiently, even the two SM muons may appear overlapping, effectively reducing the number of informative dimensions in the NNID curves below six. At that point, distinguishing the more boosted MSSM-3 events becomes impossible unless the total number of events N is increased.

In contrast, the test based on $\bar{r}_i$ suffers a much stronger degradation when spectator particles are added. The addition of new dimensions distorts the $\bar{r}_i$ distribution, and since $\bar{r}_i$ enters the test directly, rather than through a topological quantity as the NNID, these distortions immediately wash out the distinctive MSSM contribution. As a result, it becomes significantly harder to identify a new effect hidden in the background.

Robustness to Systematics. As a first assessment of the robustness of our method against systematic errors, we concentrate on scale and resolution uncertainties in the charged particles' p_T . Once again, we take as a benchmark the detector-level MSSM-3 model.

Any global rescaling of p_T uniformly rescales the distances $\bar{r}_i$, while leaving the NNID completely unchanged. Indeed, from the definition of the EMD in Eq. (2.6), it follows directly that all distances scale linearly under dilations, as expected from their scaling dimension

$$p_T \rightarrow \lambda p_T \quad \Rightarrow \quad r_{ki} \rightarrow \lambda r_{ki}. \quad (4.1)$$

Since the NNID depends only on ratios of these distances, it is exactly invariant under




Figure 7: Top row: NNID curves at detector level in the SM and MSSM-3 benchmark models described in Section 4.1 with a variable number m of spectator particles (i.e. fictitious particles that are identical in the MSSM-3 and SM samples). Bottom row: the significance of the NNID new physics test for the MSSM-3 benchmark model remains stable if the number m of spectator particles is increased, while the $\bar{r}_i$ test suffers more.

p_T rescalings. We demonstrate this in Fig. 8, where we apply a constant rescaling that increases all particles' p_T values, mimicking the procedure used to estimate detector energy-scale uncertainties. We set $\lambda = (1+\alpha)$ in Eq. (4.1) with $\alpha = \{5\%, 10\%, 20\%, 40\%\}$. Both the NNID curve (zoomed in here to highlight the 68% confidence band) and the corresponding significance remain unchanged compared to the unrescaled case. The same conclusion holds, of course, for rescalings that decrease p_T . For the $\bar{r}_i$, while the rescaling seems to alter the distribution of distances, the uniform transformation in Eq. (4.1) ensures that the test statistic in Eq. (3.1) is invariant. Consequently, the significance is unaffected, as is evident from the curves shown in Fig. 8.

The impact of the finite p_T resolution of the detector on the value of the NNID is suppressed by the sample size ($N \approx 5000$ in the Figure). While the CMS DELPHES card that we used already incorporates a few-percent p_T smearing for leptons, we deliberately amplify this effect by applying an additional 5%, 10%, 20% or 40% smearing to the selected final state events.

By standard M-estimator theory, the variance of the NNID, σ_d^2 , scales as $1/N$. An additional per-particle smearing of relative size σ_{p_T} inflates the variance of the distances, and hence of each log-ratio $\mu_{k,ij} = r_{k,j}/r_{k,i}$, by an $\mathcal{O}(\sigma_{p_T}^2)$ term, but its contribution to




Figure 8: A constant, uniform rescaling of all particles’ p_T leaves the NNID invariant, and thus the significance remains unchanged (left panels). The distances $\bar{r}_i$ get uniformly rescaled, but the significance is still unaltered (right panels). We rescale the p_T by up to 40% in the Figure, but we could have further increased this factor with no visible effect. The Figure demonstrates the insensitivity of our new physics search to energy scale uncertainties.

σ_d^2 is suppressed by $1/N$. Consequently, even a $\mathcal{O}(10\%)$ smearing in p_T produces only a few-percent increase in the standard error of d , which is already at the percent level. From the top-left panel of Fig. 9, we can indeed appreciate how the 68% confidence level bands for the NNID maintain the same width. The most notable differences are the shifts in the median, particularly for $\sigma_{p_T} = 20\%, 40\%$, associated to the hard cut that we impose to keep $p_T > 0$ which distorts the distribution of the distances, as shown in the top-right panel of Fig. 9. The same argument applies to the variance of $\bar{r}_i$, which is also already at percent level and hence is essentially unaffected even for big additional smearing.

Since the significance depends on the relative distance between the NNID and $\bar{r}_i$ curves in units of their standard deviation, it should not be appreciably affected by the smearing, given its small impact on the standard deviation of their distribution. Indeed this is what we can see from the bottom panels of Fig. 9, in which the significances remain essentially unchanged, except for the biggest smearing values which induce a small decrease of Z .

Overall, this demonstrates that our $\bar{r}_i$ and NNID-based significance is both exactly insensitive to global scale shifts and highly resilient to large random p_T variations, making the method very robust to detector effects on particles’ energies and momentum. We




Figure 9: Robustness of the NNID and the associated statistical test to a Gaussian smearing of all particles' p_T . The Figure demonstrates the insensitivity of our new physics search to energy resolution uncertainties.

expect similar results to hold for errors on the particles' positions (η and ϕ), given that their contribution to the EMD is similar to that of the particle p_T .

In principle, such a transformation would also affect which events pass the p_T selection cuts, since their absolute thresholds are not invariant under rescaling. In our analysis, the rescaling (and later the smearing) is applied only after event selection, so this effect is not included. In Appendix B we rescale the p_T of all particles before event selection and we find a non-zero effect, as expected, but still a very modest impact on the significance as shown in Fig. 18 and explained in the Appendix.

4.2 Leptonic Z'

To test our approach as a model-independent search, we now consider a completely different signal. Specifically, we focus on a light Z' model arising from gauging the $U(1)_{L_\mu-L_\tau}$ symmetry of the SM, where $L_{\mu,\tau}$ are the muon and tau lepton numbers. The process we study is $pp \rightarrow 2\mu^+2\mu^-$, where a Z' boson is radiated by a final-state muon and subsequently decays into a $\mu^+\mu^-$ pair (the characteristic topology is depicted in Fig. 10 a). This setup is quite different compared to the previous example, where the final-state particles originating from the decay of heavy new physics were highly energetic. Here we do not have boosted particles in the final state and, additionally, we have a relatively low mass resonance (in what follows we take $m_{Z'} < m_Z$).




Figure 10: Dominant production mechanism for the Z' signal in Section 4.2 (a) and leading SM backgrounds (b), (c).

The leading irreducible Standard Model background consists of diagrams similar to that of the Z' , i.e. with a radiatively emitted Z boson, along with other contributions, that include Z bosons and virtual photons as intermediate states, as shown in Fig. 10. We include all relevant tree-level processes in our background simulation (but not the 1-loop box diagram with virtual quarks, considered in [187]).

Simulation and Event Selection. Events are generated at parton level with a custom MADGRAPH implementation of the Z' model [188]. We set $\sqrt{s} = 13.6$ TeV and a reference luminosity of 400 fb^{-1} , doubled with respect to Section 4.1 to provide a sizeable amount of events. We then run the data through PYTHIA 8 and DELPHES to simulate hadronization and detector effects. The procedure for computing $\bar{r}_i$ and the NNID is the same as in Section 4.1, and it is carried out again for two sets of events:

- i) Parton level:* for each event, the 2 muons and 2 antimuons from MADGRAPH are directly employed in the computation of distances, subject to the same cuts as in Section 4.1 regarding rapidity, p_T and muon isolation. This leads to datasets of ≈ 2600 SM events with the reference luminosity of 400 fb^{-1} ;
- ii) Detector level:* we employ again the custom CMS card, and select only events with at least two isolated muons and two isolated antimuons with $p_T > 10 \text{ GeV}$, $|\eta| < 2.5$. As before, jets are allowed if $p_{T,j} < 40 \text{ GeV}$, although they are not employed for the computation of distances. The efficiency of the selection after the detector simulation retains only about 50% of the events compared to the parton case.

Results. We generate 500 samples of the SM background, and compute the NNID and $\bar{r}_i$ with 50 logarithmically spaced values of the normalized index i/N . The granularity is improved compared to the model in Section 4.1 to compensate the smaller number of events per sample in this benchmark. We collect in Section A the comparison of our test statistic under the null hypothesis with the theoretical T^2 distribution.

For the Z' signal, we fix the mass to $M_{Z'} = 50 \text{ GeV}$ and vary the gauge coupling to muons and taus, denoted as $g_{Z'}$, in the range $[0.01, 0.1]$ in steps of 0.01. As a reference, the number of Z' events passing the cuts and selection at detector level ranges from a few




Figure 11: NNID curves at parton level as a function of $2i/N$ and $\bar{r}_i$ in the pure SM and pure Z' samples (left). In the right panels we show a comparison between the SM and the SM+ Z' samples. The mass of the Z' is fixed to $M_{Z'} = 50$ GeV with coupling $g_{Z'} = 0.1$, and the curves are obtained with a fixed number of SM events $N = 2613$ for illustrative purposes. The bottom row shows the invariant mass distribution of the four muons in the event. We highlight with a dashed black box the region of the NNID curves corresponding to the two shoulders in the invariant mass distribution of the bottom panels.

($g_{Z'} = 0.01$) to about 100 ($g_{Z'} = 0.1$). To obtain confidence level bands, we generate 50 samples for each value of the coupling.

The NNID curves for both the pure and the combined background+signal samples are shown in Fig. 11. Interestingly, both the pure SM and SM+ Z' cases exhibit a distinct “plateau” at large values of $2i/N$, corresponding to scales of $\mathcal{O}(100)$ GeV. A comparison with the invariant mass distributions of the four leptons $2\mu^+2\mu^-$, shown in the bottom




Figure 12: Significance bands at 68% CL of the new physics test at parton and detector level, as a function of the Z' coupling to leptons. The curves are obtained using the test described in Section 3.1, either using information from the NNID (blue curves) or from $\bar{\tau}_i$ (orange curves).

panels of Fig. 11, suggests a possible connection with the two shoulders highlighted by the dashed boxes. This is confirmed by the absence of such a bump in the NNID curves and invariant mass distributions in the leptonic MSSM case (see Fig. 5) and in the pure Z' samples in Fig. 11.

The significances obtained from the NNID and $\bar{\tau}_i$ tests described in Section 3.1, at both parton and detector level, are shown in Fig. 12 as a function of $g_{Z'}$. The sensitivities at parton and detector level are compatible within one sigma.

In this case, the test based on the NNID is more powerful than that based on $\bar{\tau}_i$, in contrast to the findings of Section 4.1. The reduced sensitivity of the $\bar{\tau}_i$ test can be traced to the kinematics of the process: the events associated with Z' production are not significantly more boosted than the corresponding SM background ones, since $m_{Z'}$ is close to the Z -boson mass. As a result, the distribution of $\bar{\tau}_i$ alone provides limited discrimination. The NNID, on the other hand, appears to enhance the separation between signal and background by effectively exploiting a non-linear function of $\bar{\tau}_i$, thereby leading to a more powerful test.

As in previous cases, the simple counting experiment (dashed black line in the figure) is dominated by systematic uncertainties (4.4% on the SM cross section, combined with a 1.6% uncertainty on the luminosity [189]).

The main message of this Section is that our NNID test is sensitive also to the Z' signal. The kinematics explored in this Section and the previous one are completely different from each other and support our statements on the flexibility of our new physics search strategy. Additionally, the performances at parton level and detector level remain

comparable, further confirming the robustness of our test.

4.3 Semileptonic Decays of Charginos and Neutralinos

In this section we show that our techniques can find new physics without clustering jets in a ~ 100 -dimensional phase space. To this end, we employ the MSSM-3 model introduced in Section 4.1, but consider a semileptonic final state where the W boson decays to jets: $W \rightarrow jj$.

Simulation and Event Selection. As in the previous sections, parton-level events, hadronization, and detector effects are simulated using MADGRAPH, PYTHIA, and DELPHES, respectively, with the same configuration cards as in the leptonic study. We take $\sqrt{s} = 13.6$ TeV and generate batches of 10,000 clean parton level events, formally fixing the effective luminosity through the corresponding cross section (see below). This choice is motivated by the fact that the cross section of the relevant background process would otherwise produce datasets far too large for us to process with our available resources, if we were to match the luminosity used in Section 4.1. We perform the analysis on two sets of events:

- i) *Parton level:* the four charged partons from the hard collision generated by MADGRAPH are used directly. The selection cuts are $p_T > 10$ GeV for muons and $p_T > 30$ GeV for jets, both within $|\eta| < 2.5$. The same muon isolation as in Section 4.1 is imposed, and we further require $\Delta R_{jj} \geq 0.4$ to mimic the detector-level jet separation (since jets are clustered with anti- k_T and radius $R = 0.4$, see below). We also impose $70 \text{ GeV} \leq m_{jj} \leq 90 \text{ GeV}$ and $76 \text{ GeV} \leq m_{\mu\mu} \leq 106 \text{ GeV}$;
- ii) *Detector level:* we use the custom CMS card of Sections 4.1 and 4.2, this time with a jet radius $R = 0.4$. Events are required to contain an isolated muon and an isolated antimuon with $p_T > 10$ GeV and $76 \text{ GeV} \leq m_{\mu\mu} \leq 106 \text{ GeV}$, as well as at least two jets with $p_T > 30$ GeV and $70 \text{ GeV} \leq m_{jj} \leq 90 \text{ GeV}$.

With these cuts, the leading SM background consists of $Z+2$ jets events; the cross section for WZ is orders of magnitude smaller. The dimuon invariant mass window suppresses the $t\bar{t}$ background. We stress that additional cuts, for instance on the missing transverse momentum as in dedicated MSSM searches for this channel [190], could further improve the signal isolation. However, since our goal is to identify potential traces of new physics within a fixed final-state configuration (here $2\mu 2j$), without prior knowledge of the signal, imposing such cuts would artificially enhance the significance. With the cuts above, the parton-level cross section for the SM background is ≈ 3 pb; the relative efficiency from parton to detector level is about 45% for both signal and background.⁹

To test the robustness of our search against the dimensionality of the input space, we compute distances and NNID at detector level in two setups: (a) using muons and

⁹This cross section and our events are obtained simulating $Z + 2j$ at LO with MADGRAPH. We do not generate an inclusive matched sample because we choose to neglect the corrections to the p_T distribution of our jets with $p_T > 30$ GeV (that are of $\mathcal{O}(\alpha_s \log(m_Z/p_T))$) in this first proof-of-principle study of the methodology.

clustered jets as inputs, and (b) using muons and the *subcomponents* of the selected jets. The latter case results in an average of about 30 particles per event, corresponding to a search performed in an $O(100)$ -dimensional space.

Results. We compute the NNID on 200 SM samples. We collect in Section A the comparison of our test statistic under the null hypothesis with the theoretical T^2 distribution for the various sets of events described below (parton level, detector level clustered and unclustered). Then, we proceed as in the leptonic case: we include a variable fraction of MSSM events in the SM samples, and perform the new physics test. The fraction is set in a fiducial region defined by loose cuts:

$$p_{T,j} > 25 \text{ GeV}, \quad p_{T,\mu^\pm} > 10 \text{ GeV}, \quad |\eta_{j,\mu^\pm}| < 2.75.$$

Since the efficiency for SM and MSSM with respect to our event selection is different (about 10% for the SM and 60% for the MSSM at parton level), the effective S/B ratio in the actual computation is almost a factor of ten larger. This explains why in this case a smaller percentage of signal is needed to get sizeable significances compared to Sections 4.1 and 4.2.

In Fig. 13 we show the NNID as a function of the proportionality vector $2i/N$ and of the average distance to the i -th neighbor, $\bar{r}_i$, for pure SM samples and SM + MSSM-3 events. The input to the NNID computation varies between 12 kinematical variables in the parton case, to more than 100 for detector-level unclustered events. Nonetheless, NNID curves do not vary dramatically between these two sets of events. The results in the left panel of Fig. 13 show that the transition between partons and jets encodes information, as the NNID grows at small i/N between the two sets of events from roughly 9 to around 11. Additionally, the ID correctly recognizes that the 100 or so particles in the unclustered events are in reality encoding information about a much smaller set of degrees of freedom. Indeed we can see in the left panel of Fig. 13 that the NNID computed at detector level starting from jets (so overall 12 input variables) gives roughly the same result as the NNID computed from jet constituents (so around 100 input variables). In the right panel of the figure we show that also the typical distances between neighbors do not change appreciably between clustered and unclustered calculations, further demonstrating the robustness of these observables. The shorter distances probed in the parton-level dataset arise from the larger number of available events, which allows the nearest-neighbors estimator to explore smaller scales. At fixed $\bar{r}_i$, the intrinsic dimension extracted from detector-level data is indeed larger, as expected, because hadronization and detector smearing introduces additional noisy directions.

In Fig. 14 we display the usual significance plots as a function of a variable BSM cross section. In this case, for the reference “naive” counting experiment we take a total error of $\approx 9\%$ on the theory SM prediction for $Z + 2j$ [191]. Quite remarkably clustered and unclustered events at detector level give the same median significance with also the error bar on the significance remaining stable. We emphasize that, in the unclustered case, this corresponds to a model-independent new physics search in a space of order $\mathcal{O}(100)$




Figure 13: NNID curves at parton, clustered detector and unclustered detector level as a function of $2i/N$ (left panel) and the average distance to the i -th neighbor (right panel). The central line represents the median, while the band is the 68% confidence interval.

dimensions. This demonstrates the robustness of our NNID and $\bar{r}_i$ new physics test and supports its application to real data.

The parton-level performances are instead higher than the detector-level ones. We can trace this difference to the different efficiency of our event selection cuts. The larger number of events at parton level reduces the statistical error on NNID and $\bar{r}_i$.

In Fig. 19 (Section B) we provide the same significance plot in terms of S/B , since the efficiency of the cuts going from the region of phase space where we define our fiducial cross sections to the signal region that we use to compute the NNID makes S/B larger than the fiducial cross section ratio.

4.4 Summary of Results

This work represents an initial proof-of-principle study in which we have included the Standard Model background most closely resembling the signal. Nevertheless, the results obtained here are already highly encouraging and clearly illustrate the capabilities and distinctive advantages of our new physics search strategy.

1. **Flexibility:** sensitivity to diverse Beyond the Standard Model scenarios with different kinematics. This is demonstrated by the variety of signals explored in the three previous Sections. We had boosted leptons from the MSSM-3 signal in Section 4.1, low-mass resonances from the Z' in Section 4.2, boosted jets in Section 4.3 and heavy non-resonant new physics from the MSSM-2 signal in Section 4.1.
2. **Robustness** (to showering and detector effects): the performances are degraded when going from the parton-level to the detector-level simulation. We trace this effect to the lower efficiency of the detector-level event selection and the increased statistical error on NNID and $\bar{r}_i$. However, the performances remain unaltered going from clustered jets to jets' constituents and remain comparable between leptonic and hadronic decays of the W . We take this as a promising sign of the robustness of our methodology, leaving to future work the study of different showering algorithms.




Figure 14: Significance bands at 68% CL of the new physics test in the semileptonic MSSM-3 model at parton level and detector level (with clustered and unclustered jets), as a function of the ratio between the BSM and SM cross section in the fiducial region.

This is shown in Figs 6 and 12 for the MSSM-3 leptonic and Z' signals, respectively and in Fig. 14 for the MSSM-3 semileptonic signal. A comparison between leptonic and semileptonic MSSM is shown in Appendix B.

3. **Robustness to Systematics:** systematic errors can be incorporated in a statistically well-defined way. We included p_T scale and resolution uncertainties in the leptonic MSSM-3 benchmark in Figs 8 and 9 and found no sensitivity to momentum scale uncertainties and a mild sensitivity to momentum resolution uncertainties. We expect similar results for angular resolution uncertainties.
4. **Protection from the Curse of Dimensionality:** This feature is already present by construction in our definition of the NNID, but we demonstrate it explicitly in a toy example in Fig. 7, and by detecting new physics in 100 dimensions in Section 4.3.

4.4.1 Qualitative Comparisons with Other Search Strategies

A quantitative comparison with other model-independent search strategies goes beyond the scope of this work. However, at least conceptually, it is easy to show that our technique is complementary to the two broad categories of model-independent searches identified in the introduction. Our methodology does not require a localized resonance in some kinematical variable and in this sense is complementary to the techniques put forward in [37, 45, 64] that require one. The more flexible strategies in [38, 50, 99] can in principle detect any of our signals, but only at much larger cross sections. Their global p -value decreases exponentially with the dimension of the input space, placing the $d \simeq 100$ and $d \simeq 60$ examples that we have discussed well beyond their reach with any future LHC

luminosity. Complementing [38, 50, 99] with a pre-training step of dimensionality reduction as in [141, 142] can solve the problem, but for now it has been done with event-by-event observables for localized signals. So our “global” approach is conceptually quite different and it would be interesting to compare the performances of the two ideas in more detail.

It is, of course, possible to find examples where other techniques would shine and ours would underperform. It is the case, for instance, of the LHC Olympics dataset [192], which contains only resonant signals. Those signals are easily detectable with [37, 45, 64], but require a larger cross section with our methodology. Similarly, any intricate low-dimensional signal that requires combining multiple final states each with kinematics close to the SM and a small cross section is better searched for using [38, 50, 99] rather than what we proposed in this work.

To conclude this brief and qualitative discussion, we would like to emphasize two aspects of model-independent searches that are obvious to the practitioners, but often underappreciated by the broader community. We chose to compare our sensitivities to that of a counting experiment because this is the only option when one knows nothing about the signal and other model-independent search strategies fail by construction. Asking if a model-dependent search targeted to our Z' signal, for instance, would do better, is not a honest question. It forces a comparison between what we do when we know nothing about the signal and what we do when we know everything about it. Obviously, there is always a better way to construct a search compared to what we did, if we know what we are looking for. So all the hundred (or so) model-independent search strategies developed so far are truly useful only in two situations. The first is when we suspect that signals lying outside our current theory priors might be hidden in the data. The second is when the range of reasonable possibilities is so large that we do not have the manpower to search for all of them individually. An example of the first case is an object that does not trace a helix in the magnetic field of the LHC detectors. An example of the second is a Higgs-like particle with sizeable branching ratios in many final states but a cross section small enough that we must combine all channels with the right weights in order to observe it.

The second aspect of model-independent searches that is obvious, but worth to state explicitly is that anomaly detection is only the first step in a new physics search. We can establish if an anomalous event is coming from a new physics phenomenon only if we can define a reference hypothesis to compare it with and output a p -value. After identifying a set of anomalous events, we still have to set up a statistical test to determine whether they are just rare SM events or truly new physics. If we consider for instance anomaly detection with autoencoders, the most natural way to do it (but not the only way) would be to use the autoencoder output as a test statistics and then perform a standard frequentist goodness-of-fit test.

This short discussion completes the analogy between the most common use of autoencoders in high energy physics and the second category of model-independent search strategies that we mentioned in the introduction. The ideas in those two sets of papers are just two different ways of obtaining a test statistic needed to perform a goodness-of-fit test. A good test statistic necessarily describes correctly the probability distribution of the reference data (i.e. the SM) against which we want to compare the experimental data,

so an autoencoder needs a larger and larger number of parameters in the hidden layer as we increase the dimensionality d of the training data. The increase in internal parameters must be comparable to that in [38, 50, 99] since both methodologies aim at fitting the probability distribution of the reference data with a neural network. Hence we expect a look-elsewhere effect that increases exponentially in d also for autoencoders.

The analogy between these two methodologies becomes an equality in the case of normalized autoencoders [139, 193] where the autoencoder is built explicitly to learn the probability distribution of the training data, just as was done in [38, 50, 99] with different loss functions. Note that normalized autoencoders are built precisely to correct what makes traditional autoencoders unsuitable for anomaly detection, i.e. that they are not learning correctly the probability distribution of their training data for outliers.

However, as stated in the introduction, this is not the only possible use of autoencoders or dimensionality reduction in general. If we keep the number of extracted features fixed and do not increase it with the input dimension of the dataset, we end up with an analysis strategy in the first category, with a smaller look-elsewhere effect, but a correspondingly smaller expressivity on the signal. In that case we aim at well-describing the SM only on a lower-dimensional subspace of the input space.

5 Conclusion and Outlook

We have discussed how to perform a data-driven measurement of the dimension of the phase space of collider events. In Sections 3 and 4 we have demonstrated that it can be used as a model-independent new physics search strategy. It allows to capture new physics precisely where it could still be hiding at the LHC, i.e. in subtle correlations between hundreds or thousands of observables that in existing searches are coarse-grained and reduced to a few high-level quantities. Our methodology, as is, already allows to search for new physics without clustering jets (see Section 4.3) and we see it as a first step towards performing a search directly on raw data. There are of course a number of open questions on how to analyze raw data directly, so many of them that we will not even list them explicitly, but this is a first conceptual step in this direction.

The results presented here aim at illustrating the salient conceptual features of the methodology. We see this as a first paper that will be followed by more detailed studies which will include more background processes, comparisons between different showering algorithms and MC tunes, the full set of systematic uncertainties for each analysis and studies of even higher-dimensional spaces.

In addition to new physics searches, there is a number of other applications of the NNID that we have barely mentioned in the rest of the paper, but we would like to explore in future work. We highlight here those that we find most interesting:

1. *What is the intrinsic dimension of a jet?* The answer depends on the energy scale (as noted also in [1, 14, 15]) and encodes interesting information on the transition between the perturbative and non-perturbative regimes. Compared to existing studies [1, 14, 15] the NNID is more robust experimentally at small EMDs for finite datasets, as

illustrated by the toy model of Section 2.3. This leaves open the question of its calculability in perturbative QCD, which is itself an interesting problem for future work.

2. *What is the right dimension for the latent space of an autoencoder or other dimensionality reduction techniques?* This question is usually answered heuristically in high energy physics, but the NNID provides a (rather standard) data-driven methodology to find the optimal size of these spaces.
3. *Is this a quark or a gluon jet?* From preliminary studies we found that the NNID of quark and gluon jets, computed from their constituents, is appreciably different and could be used as an observable to improve their classification, as was done with the correlation dimension in [15].
4. *Can the NNID improve model-dependent searches?* The application of neural networks (and older techniques such as boosted decision trees) to model-dependent analyses seems to have left very little margin of improvement to human ingenuity. However model-dependent searches are optimized using event-by-event observables and not global properties of the dataset. We would like to explore the possibility that the NNID curves computed in this work can be used to further increase the sensitivity of model-dependent searches.
5. *More “global” observables for new physics searches.* As mentioned just above, new physics is often searched for using observables that can be defined on an event-by-event basis, but here we have discussed a way to do it using only the measurement of a property of the entire dataset. The NNID has no meaning for a single event. In this sense our strategy is orthogonal to anomaly detection and could be the first of many other applications of global properties of a dataset to detect new physics. We already took a first step in this direction by using the average distances $\bar{r}_i$ to detect new physics.

In this work we have explored the intrinsic dimension of a dataset as a tool for model-independent new physics searches. As such, it allows to study very high-dimensional spaces, and could reveal signals still hiding in LHC data. We see this technique as something qualitatively new in the already vast landscape of model-independent searches.

Furthermore, the NNID of a dataset can have a long list of other applications to the analysis of collider data, including the study of non-perturbative aspects of QCD and the optimization of model-dependent searches.

Acknowledgements

We would like to thank L. Wang, J. Thaler, A. Larkoski and G. Grosso for useful discussions. RTD would like to thank P. Urbani for introducing him to the NNID. RTD and GR acknowledge ANR grant ANR-23-CE31-0024 EUHiggs for partial support. This research was supported in part by grant NSF PHY-2309135 to the Kavli Institute for Theoretical

Physics (KITP). The work of AG was supported in part by the European Union - Next Generation EU under Italian MUR grant PRIN-2022-RXEZCJ. The work of AV has received funding from the Swiss National Science Foundation (SNF) through the Eccellenza Professorial Fellowship “Flavor Physics at the High Energy Frontier”, project number 186866. Some of the calculations were performed at sciCORE (<http://scicore.unibas.ch/>) scientific computing center at University of Basel.

A Empirical T^2 Distribution, Validation and PCA

In this Appendix, we expand upon and illustrate the procedure outlined in Section 3.1 for validating the T^2 statistic under the null hypothesis. The primary source of discrepancy between the theoretical and empirical distributions arises when the NNID vector components are highly correlated or when the feature count m becomes comparable to the number of reference samples n . In these regimes, the empirical covariance matrix $\hat{\Sigma}$ develops small eigenvalues, its inversion is numerically unstable and amplifies statistical noise, causing large fluctuations in the test statistic.

To diagnose this instability, we construct the empirical T^2 distribution via a leave-one-out approach under the null hypothesis: for each $j = 1, \dots, n$, we remove the j -th reference vector, compute $\hat{\mu}^{(-j)}$ and $\hat{\Sigma}^{(-j)}$ from the remaining $n - 1$ samples, then evaluate

$$\hat{T}_j^2 = \frac{n-1}{n} (\mathbf{d}_j - \hat{\mu}^{(-j)})^T [\hat{\Sigma}^{(-j)}]^{-1} (\mathbf{d}_j - \hat{\mu}^{(-j)}). \quad (\text{A.1})$$

This yields n (approximately independent, up to $\mathcal{O}(1/n)$ correlations) realizations of $\hat{T}^2$ under the null hypothesis.

A qualitative check is provided by a Q-Q plot of the empirical quantiles of $\{\hat{T}_j^2\}$ against the theoretical quantiles of the $T_{n-1,m}^2$ distribution (equivalently proportional to $F_{m,(n-1)-m}$). Deviations from the diagonal, beyond expected finite-sample fluctuations in the upper tail, signal a breakdown of the null hypothesis assumptions.

For a quantitative assessment, one may employ nonparametric goodness-of-fit tests such as Kolmogorov–Smirnov [194] or Cramér–von Mises (CvM) [195]. We prefer the CvM statistic because it weights deviations in the tails and near the median equally, yielding a more balanced measure of agreement.

Since the root cause of the discrepancy is covariance-matrix instability, a natural remedy is dimensionality reduction. We implement a PCA-based selection of principal directions as follows.

PCA Selection Algorithm:

1. For each reference vector $\mathbf{d}_j$, remove it from the dataset and perform PCA on the remaining $n - 1$ vectors. Apply the same rotation induced by PCA to $\mathbf{d}_j$.
2. For each $k = 1, \dots, m$, compute $\hat{T}_j^2(k)$ using only the first k principal components (ordered by descending covariance matrix eigenvalue).

This produces an $n \times m$ matrix of $\hat{T}_j^2(k)$ values. For each column k , we apply the CvM test to compare its empirical distribution to the theoretical $T_{n-1,k}^2$ distribution, yielding a p -value $p(k)$. Plotting $p(k)$ versus k reveals the largest k for which the empirical and theoretical distributions remain in statistical agreement. Finally, we inspect the Q–Q plot at the chosen k to confirm visually the adequacy of the fit. This final step also helps identify and potentially remove remaining outliers, which are still visible in some plots.

The same procedure can be repeated for the $\{\bar{\mathbf{r}}\}$ vectors and for the joint $\{(\mathbf{d}, \bar{\mathbf{r}})\}$ set. When combined, PCA should be performed globally on the concatenated and standardized vector. Standardizing the concatenated vector before applying PCA is important, as otherwise observables with big relative but small absolute variations may be obscured by those with larger scales.




Figure 15: CvM p -value as a function of the number of principal components retained for the SM background of the Z' signal in Section 4.2. The right panel shows the qualitative agreement at the optimal number of components via a Q–Q plot for the NNID vector.

A representative example is shown in Fig. 15, applied to the Standard-Model background of the Z' signal in Section 4.2. The left panel displays the CvM p -value curves for the NNID vectors and the $\bar{r}_i$ vectors. For small k , the CvM tests shows good compatibility between the theoretical and empirical reference distributions. As k increases, small eigenvalues introduce instability and the p -values fall. Setting a 5% cutoff selects, for instance, $k = 21$ components for the NNID vector; the corresponding Q–Q plot in the right panel confirms the adequacy of this choice.




Figure 16: CvM p -values as a function of the number of principal components retained for the SM background of the MSSM leptonic signal in Section 4.1 at parton and detector level.




Figure 17: Q-Q plot corresponding to the optimal number of PCA components for the NNID vectors at parton level (top right), and CvM p -value as a function of the number of principal components retained for the SM background of the MSSM semileptonic signal in Section 4.3.

For completeness, we include in the following pages the diagnostic plots for the other benchmark models analyzed in Section 4, following the same procedure used in Fig. 15. In Figs. 16 and 17 we present the corresponding results for the leptonic and semileptonic MSSM backgrounds discussed in Section 4.1, 4.3.

Interestingly, at very large k the p -value curves sometimes recover, particularly in the

case of $\bar{r}_i$ for the MSSM background (Figs 16, 17), wrongly suggesting a good fit. We attribute this to the diminishing discriminating power of the $F_{k,(n-1)-k}$ distribution when k is large: noise amplifications fail to produce statistically significant deviations, giving an illusion of agreement. This regime, however, is unsuitable for new-physics discovery and should be avoided.

B More Details on the MSSM Analyses

In Section 4.1 we showed the impact on the significance test based on the NNID and $\bar{r}_i$ when uniformly rescaling the particle transverse momenta (p_T) *after* the event selection had been carried out. Here we examine how the results change if the rescaling is applied *before* the event selection. In this case, we downscale p_T by 5, 10, and 20% with respect to the true values.




Figure 18: Impact of a global uniform downscaling of p_T *before* the event selection on the significance tests with NNID and $\bar{r}_i$ in the leptonic MSSM case.




Figure 19: Significance bands at 68% CL of the new physics test in the leptonic (left) and semileptonic (right) MSSM-3 model as a function of the signal over background rate after all cuts.

In the leptonic case, the event selection required $p_T > 10$ GeV for all leptons, and additionally $p_T < 40$ GeV for all extra jets present in the event. Since the leptons origi-

nate from on-shell Z and W bosons, the downscaling has a negligible effect on the lepton p_T cut. The dominant effect instead arises from the jet veto: as p_T is reduced, more events are accepted because they satisfy the veto condition. Quantitatively, the overall cut efficiency increases by about 5% between the no-rescaling and the 20% downscaling cases. Consequently, we expect a small improvement on the significance test, unlike the post-cut rescaling analyzed in Section 4.1, where the test remained exactly invariant. This expectation is confirmed by our results, shown in Fig. 18. Interestingly, the effect is more pronounced for the $\bar{r}_i$ test than for the NNID one.

To conclude we show the sensitivity curve for the leptonic and semileptonic analysis of the MSSM-3 benchmark as a function of S/B after all cuts in Fig. 19.

References

- [1] P.T. Komiske, E.M. Metodiev and J. Thaler, *Metric Space of Collider Events*, *Phys. Rev. Lett.* **123** (2019) 041801 [[1902.02346](#)].
- [2] P.T. Komiske, E.M. Metodiev and J. Thaler, *The Hidden Geometry of Particle Collisions*, *JHEP* **07** (2020) 006 [[2004.04159](#)].
- [3] E. Levina and P.J. Bickel, *Maximum likelihood estimation of intrinsic dimension*, in *Neural Information Processing Systems*, 2004, <https://api.semanticscholar.org/CorpusID:14865278>.
- [4] C. Ceruti, S. Bassis, A. Rozza, G. Lombardi, E. Casiraghi and P. Campadelli, *Danco: dimensionality from angle and norm concentration*, *arXiv preprint arXiv:1206.3881* (2012) .
- [5] E. Facco, M. d’Errico, A. Rodriguez and A. Laio, *Estimating the intrinsic dimension of datasets by a minimal neighborhood information*, *Scientific reports* **7** (2017) 12140.
- [6] F. Denti, D. Doimo, A. Laio and A. Mira, *The generalized ratios intrinsic dimension estimator*, *Scientific Reports* **12** (2022) 20005.
- [7] P. Campadelli, E. Casiraghi, C. Ceruti and A. Rozza, *Intrinsic dimension estimation: Relevant techniques and a benchmark framework*, *Mathematical Problems in Engineering* **2015** (2015) 759567.
- [8] F. Camastra and A. Staiano, *Intrinsic dimension estimation: Advances and open problems*, *Information Sciences* **328** (2016) 26.
- [9] C. Cesarotti and J. Thaler, *A Robust Measure of Event Isotropy at Colliders*, *JHEP* **08** (2020) 084 [[2004.06125](#)].
- [10] D. Ba, A.S. Dogra, R. Gambhir, A. Tasissa and J. Thaler, *SHAPER: can you hear the shape of a jet?*, *JHEP* **06** (2023) 195 [[2302.12266](#)].
- [11] T. Cai, J. Cheng, N. Craig and K. Craig, *Linearized optimal transport for collider events*, *Phys. Rev. D* **102** (2020) 116019 [[2008.08604](#)].
- [12] T. Cai, J. Cheng, K. Craig and N. Craig, *Which metric on the space of collider events?*, *Phys. Rev. D* **105** (2022) 076003 [[2111.03670](#)].
- [13] T. Cai, J. Cheng, B. Schmitzer and M. Thorpe, *The linearized hellinger–kantorovich distance*, *SIAM Journal on Imaging Sciences* **15** (2022) 45 [<https://doi.org/10.1137/21M1400080>].

- [14] P.T. Komiske, R. Mastandrea, E.M. Metodiev, P. Naik and J. Thaler, *Exploring the Space of Jets with CMS Open Data*, *Phys. Rev. D* **101** (2020) 034009 [[1908.08542](#)].
- [15] P.T. Komiske, S. Kryhin and J. Thaler, *Disentangling quarks and gluons in CMS open data*, *Phys. Rev. D* **106** (2022) 094021 [[2205.04459](#)].
- [16] A. Mullin, S. Nicholls, H. Pacey, M. Parker, M. White and S. Williams, *Does SUSY have friends? A new approach for LHC event analysis*, *JHEP* **02** (2021) 160 [[1912.10625](#)].
- [17] M. Crispim Romão, N.F. Castro, J.G. Milhano, R. Pedro and T. Vale, *Use of a generalized energy Mover’s distance in the search for rare phenomena at colliders*, *Eur. Phys. J. C* **81** (2021) 192 [[2004.09360](#)].
- [18] A. Davis, T. Menzo, A. Youssef and J. Zupan, *Earth mover’s distance as a measure of CP violation*, *JHEP* **06** (2023) 098 [[2301.13211](#)].
- [19] S. Tsan, R. Kansal, A. Aportela, D. Diaz, J. Duarte, S. Krishna et al., *Particle Graph Autoencoders and Differentiable, Learned Energy Mover’s Distance*, in *35th Conference on Neural Information Processing Systems*, 11, 2021 [[2111.12849](#)].
- [20] O. Kitouni, N. Nolte and M. Williams, *Finding NEEMo: Geometric Fitting using Neural Estimation of the Energy Mover’s Distance*, [2209.15624](#).
- [21] S. Alipour-Fard, P.T. Komiske, E.M. Metodiev and J. Thaler, *Pileup and Infrared Radiation Annihilation (PIRANHA): a paradigm for continuous jet grooming*, *JHEP* **09** (2023) 157 [[2305.00989](#)].
- [22] A.J. Larkoski and T. Melia, *Covariantizing phase space*, *Phys. Rev. D* **102** (2020) 094014 [[2008.06508](#)].
- [23] T. Cai, J. Cheng, N. Craig, G. Koszegi and A.J. Larkoski, *The phase space distance between collider events*, *JHEP* **09** (2024) 054 [[2405.16698](#)].
- [24] T. Cai, N. Craig, K. Craig and X. Lin, *Multi-scale Optimal Transport for Complete Collider Events*, [2501.10681](#).
- [25] K.W. Pettis, T.A. Bailey, A.K. Jain and R.C. Dubes, *An intrinsic dimensionality estimator from near-neighbor information*, *IEEE Transactions on pattern analysis and machine intelligence* (1979) 25.
- [26] L. van der Maaten, E.O. Postma and J. van den Herik, *Dimensionality reduction: A comparative review*, 2008, <https://api.semanticscholar.org/CorpusID:12051918>.
- [27] A. Rozza, G. Lombardi, C. Ceruti, E. Casiraghi and P. Campadelli, *Novel high intrinsic dimensionality estimators*, *Machine Learning* **89** (2012) 37 .
- [28] A. Ansuini, A. Laio, J.H. Macke and D. Zoccolan, *Intrinsic dimension of data representations in deep neural networks*, *Advances in Neural Information Processing Systems* **32** (2019) .
- [29] P. Pope, C. Zhu, A. Abdelkader, M. Goldblum and T. Goldstein, *The intrinsic dimension of images and its impact on learning*, *arXiv preprint arXiv:2104.08894* (2021) .
- [30] L. Valeriani, D. Doimo, F. Cuturello, A. Laio, A. Ansuini and A. Cazzaniga, *The geometry of hidden representations of large transformer models*, *ArXiv* **abs/2302.00294** (2023) .
- [31] J. Gliozzo, V. Guarino, A. Bonometti, A. Cabri, E. Cavalleri, M. Soto-Gomez et al., *Intrinsic-dimension analysis for guiding dimensionality reduction and data-fusion in multi-omics data processing*, *bioRxiv* (2024) .

- [32] P. Grassberger and I. Procaccia, *Characterization of Strange Attractors*, *Phys. Rev. Lett.* **50** (1983) 346.
- [33] H.G.E. Hentschel and I. Procaccia, *The infinite number of generalized dimensions of fractals and strange attractors*, *Physica* **8** (1983) 435.
- [34] E. Levina and P. Bickel, *Maximum likelihood estimation of intrinsic dimension*, in *Advances in Neural Information Processing Systems*, L. Saul, Y. Weiss and L. Bottou, eds., vol. 17, MIT Press, 2004, https://proceedings.neurips.cc/paper_files/paper/2004/file/74934548253bcab8490ebd74afed7031-Paper.pdf.
- [35] CMS collaboration, *MUSiC: a model-unspecific search for new physics in proton–proton collisions at $\sqrt{s} = 13$ TeV*, *Eur. Phys. J. C* **81** (2021) 629 [2010.02984].
- [36] J.A. Aguilar-Saavedra, J.H. Collins and R.K. Mishra, *A generic anti-QCD jet tagger*, *JHEP* **11** (2017) 163 [1709.01087].
- [37] J.H. Collins, K. Howe and B. Nachman, *Anomaly Detection for Resonant New Physics with Machine Learning*, *Phys. Rev. Lett.* **121** (2018) 241803 [1805.02664].
- [38] R.T. D’Agnolo and A. Wulzer, *Learning New Physics from a Machine*, *Phys. Rev. D* **99** (2019) 015014 [1806.02350].
- [39] A. De Simone and T. Jacques, *Guiding New Physics Searches with Unsupervised Learning*, *Eur. Phys. J. C* **79** (2019) 289 [1807.06038].
- [40] J. Hajer, Y.-Y. Li, T. Liu and H. Wang, *Novelty Detection Meets Collider Physics*, *Phys. Rev. D* **101** (2020) 076015 [1807.10261].
- [41] M. Farina, Y. Nakai and D. Shih, *Searching for New Physics with Deep Autoencoders*, *Phys. Rev. D* **101** (2020) 075021 [1808.08992].
- [42] T. Heimel, G. Kasieczka, T. Plehn and J.M. Thompson, *QCD or What?*, *SciPost Phys.* **6** (2019) 030 [1808.08979].
- [43] A. Casa and G. Menardi, *Nonparametric semisupervised classification for signal detection in high energy physics*, 1809.02977.
- [44] 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 [1811.10276].
- [45] J.H. Collins, K. Howe and B. Nachman, *Extending the search for new resonances with machine learning*, *Phys. Rev. D* **99** (2019) 014038 [1902.02634].
- [46] T.S. Roy and A.H. Vijay, *A robust anomaly finder based on autoencoders*, 1903.02032.
- [47] B.M. Dillon, D.A. Faroughy and J.F. Kamenik, *Uncovering latent jet substructure*, *Phys. Rev. D* **100** (2019) 056002 [1904.04200].
- [48] A. Blance, M. Spannowsky and P. Waite, *Adversarially-trained autoencoders for robust unsupervised new physics searches*, *JHEP* **10** (2019) 047 [1905.10384].
- [49] M. Romão Crispim, N.F. Castro, R. Pedro and T. Vale, *Transferability of Deep Learning Models in Searches for New Physics at Colliders*, *Phys. Rev. D* **101** (2020) 035042 [1912.04220].
- [50] R.T. D’Agnolo, G. Grosso, M. Pierini, A. Wulzer and M. Zanetti, *Learning multivariate new physics*, *Eur. Phys. J. C* **81** (2021) 89 [1912.12155].

- [51] B. Nachman and D. Shih, *Anomaly Detection with Density Estimation*, *Phys. Rev. D* **101** (2020) 075042 [[2001.04990](#)].
- [52] A. Andreassen, B. Nachman and D. Shih, *Simulation Assisted Likelihood-free Anomaly Detection*, *Phys. Rev. D* **101** (2020) 095004 [[2001.05001](#)].
- [53] O. Amram and C.M. Suarez, *Tag N' Train: a technique to train improved classifiers on unlabeled data*, *JHEP* **01** (2021) 153 [[2002.12376](#)].
- [54] M. Crispim Romão, N.F. Castro, J.G. Milhano, R. Pedro and T. Vale, *Use of a generalized energy Mover's distance in the search for rare phenomena at colliders*, *Eur. Phys. J. C* **81** (2021) 192 [[2004.09360](#)].
- [55] 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 [[2005.01598](#)].
- [56] ATLAS collaboration, *Dijet resonance search with weak supervision using $\sqrt{s} = 13$ TeV pp collisions in the ATLAS detector*, *Phys. Rev. Lett.* **125** (2020) 131801 [[2005.02983](#)].
- [57] B.M. Dillon, D.A. Faroughy, J.F. Kamenik and M. Szewc, *Learning the latent structure of collider events*, *JHEP* **10** (2020) 206 [[2005.12319](#)].
- [58] M. Crispim Romão, N.F. Castro and R. Pedro, *Finding New Physics without learning about it: Anomaly Detection as a tool for Searches at Colliders*, *Eur. Phys. J. C* **81** (2021) 27 [[2006.05432](#)].
- [59] T. Cheng, J.-F. Arguin, J. Leissner-Martin, J. Pilette and T. Golling, *Variational autoencoders for anomalous jet tagging*, *Phys. Rev. D* **107** (2023) 016002 [[2007.01850](#)].
- [60] C.K. Khosa and V. Sanz, *Anomaly Awareness*, *SciPost Phys.* **15** (2023) 053 [[2007.14462](#)].
- [61] P. Thaprasop, K. Zhou, J. Steinheimer and C. Herold, *Unsupervised Outlier Detection in Heavy-Ion Collisions*, *Phys. Scripta* **96** (2021) 064003 [[2007.15830](#)].
- [62] J.A. Aguilar-Saavedra, F.R. Joaquim and J.F. Seabra, *Mass Unspecific Supervised Tagging (MUST) for boosted jets*, *JHEP* **03** (2021) 012 [[2008.12792](#)].
- [63] S. Alexander, S. Gleyzer, H. Parul, P. Reddy, M.W. Toomey, E. Usai et al., *Decoding Dark Matter Substructure without Supervision*, [2008.12731](#).
- [64] K. Benkendorfer, L.L. Pottier and B. Nachman, *Simulation-assisted decorrelation for resonant anomaly detection*, *Phys. Rev. D* **104** (2021) 035003 [[2009.02205](#)].
- [65] A.A. Pol, V. Berger, G. Cerminara, C. Germain and M. Pierini, *Anomaly Detection With Conditional Variational Autoencoders*, in *Eighteenth International Conference on Machine Learning and Applications*, 10, 2020 [[2010.05531](#)].
- [66] V. Mikuni and F. Canelli, *Unsupervised clustering for collider physics*, *Phys. Rev. D* **103** (2021) 092007 [[2010.07106](#)].
- [67] M. van Beekveld, S. Caron, L. Hendriks, P. Jackson, A. Leinweber, S. Otten et al., *Combining outlier analysis algorithms to identify new physics at the LHC*, *JHEP* **09** (2021) 024 [[2010.07940](#)].
- [68] S.E. Park, D. Rankin, S.-M. Udrescu, M. Yunus and P. Harris, *Quasi Anomalous Knowledge: Searching for new physics with embedded knowledge*, *JHEP* **21** (2020) 030 [[2011.03550](#)].

- [69] D.A. Faroughy, *Uncovering hidden new physics patterns in collider events using Bayesian probabilistic models*, *PoS ICHEP2020* (2021) 238 [[2012.08579](#)].
- [70] G. Stein, U. Seljak and B. Dai, *Unsupervised in-distribution anomaly detection of new physics through conditional density estimation*, in *34th Conference on Neural Information Processing Systems*, 12, 2020 [[2012.11638](#)].
- [71] G. Kasieczka et al., *The LHC Olympics 2020 a community challenge for anomaly detection in high energy physics*, *Rept. Prog. Phys.* **84** (2021) 124201 [[2101.08320](#)].
- [72] P. Chakravarti, M. Kuusela, J. Lei and L. Wasserman, *Model-Independent Detection of New Physics Signals Using Interpretable Semi-Supervised Classifier Tests*, [2102.07679](#).
- [73] J. Batson, C.G. Haaf, Y. Kahn and D.A. Roberts, *Topological Obstructions to Autoencoding*, *JHEP* **04** (2021) 280 [[2102.08380](#)].
- [74] A. Blance and M. Spannowsky, *Unsupervised event classification with graphs on classical and photonic quantum computers*, *JHEP* **21** (2020) 170 [[2103.03897](#)].
- [75] B. Bortolato, A. Smolkovič, B.M. Dillon and J.F. Kamenik, *Bump hunting in latent space*, *Phys. Rev. D* **105** (2022) 115009 [[2103.06595](#)].
- [76] J.H. Collins, P. Martín-Ramiro, B. Nachman and D. Shih, *Comparing weak- and unsupervised methods for resonant anomaly detection*, *Eur. Phys. J. C* **81** (2021) 617 [[2104.02092](#)].
- [77] B.M. Dillon, T. Plehn, C. Sauer and P. Sorrenson, *Better Latent Spaces for Better Autoencoders*, *SciPost Phys.* **11** (2021) 061 [[2104.08291](#)].
- [78] T. Finke, M. Krämer, A. Morandini, A. Mück and I. Oleksiyuk, *Autoencoders for unsupervised anomaly detection in high energy physics*, *JHEP* **06** (2021) 161 [[2104.09051](#)].
- [79] D. Shih, M.R. Buckley, L. Necib and J. Tamanas, *via machinae: Searching for stellar streams using unsupervised machine learning*, *Mon. Not. Roy. Astron. Soc.* **509** (2021) 5992 [[2104.12789](#)].
- [80] O. Atkinson, A. Bhardwaj, C. Englert, V.S. Ngairangbam and M. Spannowsky, *Anomaly detection with convolutional Graph Neural Networks*, *JHEP* **08** (2021) 080 [[2105.07988](#)].
- [81] A. Kahn, J. Gonski, I. Ochoa, D. Williams and G. Brooijmans, *Anomalous jet identification via sequence modeling*, *JINST* **16** (2021) P08012 [[2105.09274](#)].
- [82] T. Aarrestad et al., *The Dark Machines Anomaly Score Challenge: Benchmark Data and Model Independent Event Classification for the Large Hadron Collider*, *SciPost Phys.* **12** (2022) 043 [[2105.14027](#)].
- [83] T. Dorigo, M. Fumanelli, C. Maccani, M. Mojsovska, G.C. Strong and B. Scarpa, *RanBox: anomaly detection in the copula space*, *JHEP* **01** (2023) 008 [[2106.05747](#)].
- [84] S. Caron, L. Hendriks and R. Verheyen, *Rare and Different: Anomaly Scores from a combination of likelihood and out-of-distribution models to detect new physics at the LHC*, *SciPost Phys.* **12** (2022) 077 [[2106.10164](#)].
- [85] E. Govorkova, E. Puljak, T. Aarrestad, M. Pierini, K.A. Woźniak and J. Ngadiuba, *LHC physics dataset for unsupervised New Physics detection at 40 MHz*, *Sci. Data* **9** (2022) 118 [[2107.02157](#)].
- [86] G. Kasieczka, B. Nachman and D. Shih, *New Methods and Datasets for Group Anomaly*

Detection From Fundamental Physics, in *Conference on Knowledge Discovery and Data Mining*, 7, 2021 [2107.02821].

- [87] S. Volkovich, F. De Vito Halevy and S. Bressler, *A data-directed paradigm for BSM searches: the bump-hunting example*, *Eur. Phys. J. C* **82** (2022) 265 [2107.11573].
- [88] E. Govorkova et al., *Autoencoders on field-programmable gate arrays for real-time, unsupervised new physics detection at 40 MHz at the Large Hadron Collider*, *Nature Mach. Intell.* **4** (2022) 154 [2108.03986].
- [89] A. Hallin, J. Isaacson, G. Kasieczka, C. Krause, B. Nachman, T. Quadfasel et al., *Classifying anomalies through outer density estimation*, *Phys. Rev. D* **106** (2022) 055006 [2109.00546].
- [90] B. Ostdiek, *Deep Set Auto Encoders for Anomaly Detection in Particle Physics*, *SciPost Phys.* **12** (2022) 045 [2109.01695].
- [91] K. Fraser, S. Homiller, R.K. Mishra, B. Ostdiek and M.D. Schwartz, *Challenges for unsupervised anomaly detection in particle physics*, *JHEP* **03** (2022) 066 [2110.06948].
- [92] P. Jawahar, T. Aarrestad, N. Chernyavskaya, M. Pierini, K.A. Wozniak, J. Ngadiuba et al., *Improving Variational Autoencoders for New Physics Detection at the LHC With Normalizing Flows*, *Front. Big Data* **5** (2022) 803685 [2110.08508].
- [93] J. Herrero-Garcia, R. Patrick and A. Scaffidi, *A semi-supervised approach to dark matter searches in direct detection data with machine learning*, *JCAP* **02** (2022) 039 [2110.12248].
- [94] C.G. Lester and R. Tombs, *Using unsupervised learning to detect broken symmetries, with relevance to searches for parity violation in nature. (Previously: "Stressed GANs snag desserts")*, 2111.00616.
- [95] J.A. Aguilar-Saavedra, *Anomaly detection from mass unspecific jet tagging*, *Eur. Phys. J. C* **82** (2022) 130 [2111.02647].
- [96] R. Tombs and C.G. Lester, *A method to challenge symmetries in data with self-supervised learning*, *JINST* **17** (2022) P08024 [2111.05442].
- [97] V. Mikuni, B. Nachman and D. Shih, *Online-compatible unsupervised nonresonant anomaly detection*, *Phys. Rev. D* **105** (2022) 055006 [2111.06417].
- [98] S. Chekanov and W. Hopkins, *Event-Based Anomaly Detection for Searches for New Physics*, *Universe* **8** (2022) 494 [2111.12119].
- [99] R.T. d'Agnolo, G. Grosso, M. Pierini, A. Wulzer and M. Zanetti, *Learning new physics from an imperfect machine*, *Eur. Phys. J. C* **82** (2022) 275 [2111.13633].
- [100] F. Canelli, A. de Cosa, L.L. Pottier, J. Niedziela, K. Pedro and M. Pierini, *Autoencoders for semivisible jet detection*, *JHEP* **02** (2022) 074 [2112.02864].
- [101] V.S. Ngairangbam, M. Spannowsky and M. Takeuchi, *Anomaly detection in high-energy physics using a quantum autoencoder*, *Phys. Rev. D* **105** (2022) 095004 [2112.04958].
- [102] J.A. Aguilar-Saavedra, *Taming modeling uncertainties with mass unspecific supervised tagging*, *Eur. Phys. J. C* **82** (2022) 270 [2201.11143].
- [103] T. Buss, B.M. Dillon, T. Finke, M. Krämer, A. Morandini, A. Mück et al., *What's anomalous in LHC jets?*, *SciPost Phys.* **15** (2023) 168 [2202.00686].
- [104] L. Bradshaw, S. Chang and B. Ostdiek, *Creating simple, interpretable anomaly detectors for new physics in jet substructure*, *Phys. Rev. D* **106** (2022) 035014 [2203.01343].

- [105] M. Birman, B. Nachman, R. Sebbah, G. Sela, O. Turetz and S. Bressler, *Data-directed search for new physics based on symmetries of the SM*, *Eur. Phys. J. C* **82** (2022) 508 [[2203.07529](#)].
- [106] J.A. Raine, S. Klein, D. Sengupta and T. Golling, *CURTAINS for your sliding window: Constructing unobserved regions by transforming adjacent intervals*, *Front. Big Data* **6** (2023) 899345 [[2203.09470](#)].
- [107] M. Letizia, G. Losapio, M. Rando, G. Grosso, A. Wulzer, M. Pierini et al., *Learning new physics efficiently with nonparametric methods*, *Eur. Phys. J. C* **82** (2022) 879 [[2204.02317](#)].
- [108] C. Fanelli, J. Giroux and Z. Papandreou, ‘*Flux+Mutability*’: a conditional generative approach to one-class classification and anomaly detection, *Mach. Learn. Sci. Tech.* **3** (2022) 045012 [[2204.08609](#)].
- [109] T. Finke, M. Krämer, M. Lipp and A. Mück, *Boosting mono-jet searches with model-agnostic machine learning*, *JHEP* **08** (2022) 015 [[2204.11889](#)].
- [110] R. Verheyen, *Event Generation and Density Estimation with Surjective Normalizing Flows*, *SciPost Phys.* **13** (2022) 047 [[2205.01697](#)].
- [111] B.M. Dillon, R. Mastandrea and B. Nachman, *Self-supervised anomaly detection for new physics*, *Phys. Rev. D* **106** (2022) 056005 [[2205.10380](#)].
- [112] S. Alvi, C.W. Bauer and B. Nachman, *Quantum anomaly detection for collider physics*, *JHEP* **02** (2023) 220 [[2206.08391](#)].
- [113] B.M. Dillon, L. Favaro, T. Plehn, P. Sorrenson and M. Krämer, *A normalized autoencoder for LHC triggers*, *SciPost Phys. Core* **6** (2023) 074 [[2206.14225](#)].
- [114] S. Caron, R.R. de Austri and Z. Zhang, *Mixture-of-Theories training: can we find new physics and anomalies better by mixing physical theories?*, *JHEP* **03** (2023) 004 [[2207.07631](#)].
- [115] S.E. Park, P. Harris and B. Ostdiek, *Neural embedding: learning the embedding of the manifold of physics data*, *JHEP* **07** (2023) 108 [[2208.05484](#)].
- [116] G. Kasieczka, R. Mastandrea, V. Mikuni, B. Nachman, M. Pettee and D. Shih, *Anomaly detection under coordinate transformations*, *Phys. Rev. D* **107** (2023) 015009 [[2209.06225](#)].
- [117] J.F. Kamenik and M. Szewc, *Null hypothesis test for anomaly detection*, *Phys. Lett. B* **840** (2023) 137836 [[2210.02226](#)].
- [118] A. Hallin, G. Kasieczka, T. Quadfasel, D. Shih and M. Sommerhalder, *Resonant anomaly detection without background sculpting*, *Phys. Rev. D* **107** (2023) 114012 [[2210.14924](#)].
- [119] J.Y. Araz and M. Spannowsky, *Quantum-probabilistic Hamiltonian learning for generative modeling and anomaly detection*, *Phys. Rev. A* **108** (2023) 062422 [[2211.03803](#)].
- [120] R. Mastandrea and B. Nachman, *Efficiently Moving Instead of Reweighting Collider Events with Machine Learning*, in *36th Conference on Neural Information Processing Systems: Workshop on Machine Learning and the Physical Sciences*, 12, 2022 [[2212.06155](#)].
- [121] J. Schuhmacher, L. Boggia, V. Belis, E. Puljak, M. Grossi, M. Pierini et al., *Unravelling physics beyond the standard model with classical and quantum anomaly detection*, *Mach. Learn. Sci. Tech.* **4** (2023) 045031 [[2301.10787](#)].

- [122] T. Golling et al., *The Mass-ive Issue: Anomaly Detection in Jet Physics*, in *34th Conference on Neural Information Processing Systems*, 3, 2023 [[2303.14134](#)].
- [123] S. Roche, Q. Bayer, B. Carlson, W. Ouligian, P. Serhiayenka, J. Stelzer et al., *Nanosecond anomaly detection with decision trees and real-time application to exotic Higgs decays*, *Nature Commun.* **15** (2024) 3527 [[2304.03836](#)].
- [124] D. Sengupta, S. Klein, J.A. Raine and T. Golling, *CURTAINS Flows For Flows: Constructing Unobserved Regions with Maximum Likelihood Estimation*, [2305.04646](#).
- [125] J.A. Aguilar-Saavedra, E. Arganda, F.R. Joaquim, R.M. Sandá Seoane and J.F. Seabra, *Gradient Boosting MUST taggers for highly-boosted jets*, [2305.04957](#).
- [126] L. Vaslin, V. Barra and J. Donini, *GAN-AE: an anomaly detection algorithm for New Physics search in LHC data*, *Eur. Phys. J. C* **83** (2023) 1008 [[2305.15179](#)].
- [127] ATLAS collaboration, *Anomaly detection search for new resonances decaying into a Higgs boson and a generic new particle X in hadronic final states using $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector*, *Phys. Rev. D* **108** (2023) 052009 [[2306.03637](#)].
- [128] V. Mikuni and B. Nachman, *High-dimensional and Permutation Invariant Anomaly Detection*, *SciPost Phys.* **16** (2024) 062 [[2306.03933](#)].
- [129] T. Golling, G. Kasieczka, C. Krause, R. Mastandrea, B. Nachman, J.A. Raine et al., *The interplay of machine learning-based resonant anomaly detection methods*, *Eur. Phys. J. C* **84** (2024) 241 [[2307.11157](#)].
- [130] S.V. Chekanov and R. Zhang, *Enhancing the hunt for new phenomena in dijet final states using anomaly detection filters at the high-luminosity large Hadron Collider*, *Eur. Phys. J. Plus* **139** (2024) 237 [[2308.02671](#)].
- [131] CMS ECAL collaboration, *Autoencoder-Based Anomaly Detection System for Online Data Quality Monitoring of the CMS Electromagnetic Calorimeter*, *Comput. Softw. Big Sci.* **8** (2024) 11 [[2309.10157](#)].
- [132] G. Bickendorf, M. Drees, G. Kasieczka, C. Krause and D. Shih, *Combining resonant and tail-based anomaly detection*, *Phys. Rev. D* **109** (2024) 096031 [[2309.12918](#)].
- [133] T. Finke, M. Hein, G. Kasieczka, M. Krämer, A. Mück, P. Prangchaikul et al., *Tree-based algorithms for weakly supervised anomaly detection*, *Phys. Rev. D* **109** (2024) 034033 [[2309.13111](#)].
- [134] E. Buhmann, C. Ewen, G. Kasieczka, V. Mikuni, B. Nachman and D. Shih, *Full phase space resonant anomaly detection*, *Phys. Rev. D* **109** (2024) 055015 [[2310.06897](#)].
- [135] M. Freytsis, M. Perelstein and Y.C. San, *Anomaly detection in the presence of irrelevant features*, *JHEP* **02** (2024) 220 [[2310.13057](#)].
- [136] E.M. Metodiev, J. Thaler and R. Wynne, *Anomaly detection in collider physics via factorized observables*, *Phys. Rev. D* **110** (2024) 055012 [[2312.00119](#)].
- [137] A. Banda, C.K. Khosa and V. Sanz, *Strengthening Anomaly Awareness*, [2504.11520](#).
- [138] E. Puljak, M. Pierini and A. Garcia-Saez, *Tensor Network for Anomaly Detection in the Latent Space of Proton Collision Events at the LHC*, [2506.00102](#).
- [139] CMS collaboration, *Wasserstein normalized autoencoder for anomaly detection*, [2510.02168](#).

- [140] CMS collaboration, *Model-agnostic search for dijet resonances with anomalous jet substructure in proton–proton collisions at $\sqrt{s} = 13$ TeV*, *Rept. Prog. Phys.* **88** (2025) 067802 [[2412.03747](#)].
- [141] S. Bright-Thonney, C. Reissel, G. Grosso, N. Woodward, K. Govorkova, A. Novak et al., *Autoscidact: Automated scientific discovery through contrastive embedding and hypothesis testing*, *arXiv preprint arXiv:2510.21935* (2025) .
- [142] G. Grosso, S.S.R. Hindupur, T. Fel, S. Bright-Thonney, P. Harris and D. Ba, *Sparse, self-organizing ensembles of local kernels detect rare statistical anomalies*, 2025.
- [143] A. van den Oord, Y. Li and O. Vinyals, *Representation learning with contrastive predictive coding*, *arXiv preprint arXiv:1807.03748* (2018) .
- [144] Y. Wang and M.I. Jordan, *Desiderata for representation learning: A causal perspective*, *arXiv preprint arXiv:2109.03795* (2021) .
- [145] M.M. Bronstein, J. Bruna, T. Cohen and P. Veličković, *Geometric deep learning: Grids, groups, graphs, geodesics, and gauges*, *arXiv preprint arXiv:2104.13478* (2021) .
- [146] L. Condren and D. Whiteson, *Finding Unexpected Non-Helical Tracks*, [2509.08878](#).
- [147] J. Forsyth, M. Low, C. Tenney and C.B. Verhaaren, *Visible Collider Signals of Natural Quirks*, [2504.02940](#).
- [148] M. Farina and M. Low, *Constraining Quirky Tracks with Conventional Searches*, *Phys. Rev. Lett.* **119** (2017) 111801 [[1703.00912](#)].
- [149] Q. Sha, D. Murnane, M. Fieg, S. Tong, M. Zakharyan, Y. Fang et al., *Learning to Reconstruct Quirky Tracks*, [2410.00269](#).
- [150] S. Knapen, H.K. Lou, M. Papucci and J. Setford, *Tracking down Quirks at the Large Hadron Collider*, *Phys. Rev. D* **96** (2017) 115015 [[1708.02243](#)].
- [151] J. Li, T. Li, J. Pei and W. Zhang, *Uncovering quirk signal via energy loss inside tracker*, *Phys. Rev. D* **102** (2020) 056006 [[1911.02223](#)].
- [152] J. Li, T. Li, J. Pei and W. Zhang, *The quirk trajectory*, *Eur. Phys. J. C* **80** (2020) 651 [[2002.07503](#)].
- [153] S. Knapen, S. Pagan Griso, M. Papucci and D.J. Robinson, *Triggering Soft Bombs at the LHC*, *JHEP* **08** (2017) 076 [[1612.00850](#)].
- [154] J. Kingman, *Poisson Processes*, Oxford science publications, Clarendon Press (1993).
- [155] F. Denti, D. Doimo, A. Laio and A. Mira, *Distributional results for model-based intrinsic dimension estimators*, *arXiv preprint arXiv:2104.13832* (2021) .
- [156] A.J. Larkoski and J. Thaler, *A spectral metric for collider geometry*, *JHEP* **08** (2023) 107 [[2305.03751](#)].
- [157] R. Gambhir, A.J. Larkoski and J. Thaler, *SPECTER: efficient evaluation of the spectral EMD*, *JHEP* **12** (2025) 219 [[2410.05379](#)].
- [158] G. Casella and R. Berger, *Statistical inference*, Duxbury Press (1990).
- [159] P.J. Huber, *Robust estimation of a location parameter*, *The Annals of Mathematical Statistics* (1964) 73.
- [160] P.J. Huber, *Robust regression: asymptotics, conjectures and monte carlo*, *The annals of statistics* (1973) 799.

- [161] I. Jolliffe, *Principal Component Analysis*, Springer (2002), <https://doi.org/10.1007/b98835>.
- [162] M.A.A. Cox and T.F. Cox, *Multidimensional scaling*, in *Handbook of Data Visualization*, (Berlin, Heidelberg), pp. 315–347, Springer Berlin Heidelberg (2008), DOI.
- [163] S.T. Roweis and L.K. Saul, *Nonlinear dimensionality reduction by locally linear embedding*, *Science* **290** (2000) 2323 [<https://www.science.org/doi/pdf/10.1126/science.290.5500.2323>].
- [164] R.R. Coifman and S. Lafon, *Diffusion maps*, *Applied and Computational Harmonic Analysis* **21** (2006) 5.
- [165] J.B. Tenenbaum, V. de Silva and J.C. Langford, *A global geometric framework for nonlinear dimensionality reduction*, *Science* **290** (2000) 2319 [<https://www.science.org/doi/pdf/10.1126/science.290.5500.2319>].
- [166] M. Belkin and P. Niyogi, *Laplacian eigenmaps for dimensionality reduction and data representation*, *Neural Computation* **15** (2003) 1373.
- [167] E. Facco, M. d’Errico, A. Rodriguez and A. Laio, *Estimating the intrinsic dimension of datasets by a minimal neighborhood information*, *Scientific Reports* **7** (2017) .
- [168] D. Granata and V. Carnevale, *Accurate estimation of the intrinsic dimension using graph distances: Unraveling the geometric complexity of datasets*, *Scientific Reports* **6** (2016) 31377.
- [169] S. Bressler, I. Savoray and Y. Zurgil, *Learning new physics from data: A symmetrized approach*, *Phys. Rev. D* **110** (2024) 095004 [2401.09530].
- [170] CMS collaboration, *High-precision measurement of the W boson mass with the CMS experiment at the LHC*, **2412.13872**.
- [171] ATLAS collaboration, *Measurement of the W -boson mass and width with the ATLAS detector using proton–proton collisions at $\sqrt{s} = 7$ TeV*, *Eur. Phys. J. C* **84** (2024) 1309 [2403.15085].
- [172] T. Anderson, *An Introduction to Multivariate Statistical Analysis*, Wiley Series in Probability and Statistics, Wiley (2003).
- [173] H. Hotelling, *The generalization of student’s ratio*, *Annals of Mathematical Statistics* **2** (1931) 54.
- [174] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer et al., *The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations*, *JHEP* **07** (2014) 079 [1405.0301].
- [175] B.C. Allanach et al., *SUSY Les Houches Accord 2*, *Comput. Phys. Commun.* **180** (2009) 8 [0801.0045].
- [176] T. Sjöstrand, S. Ask, J.R. Christiansen, R. Corke, N. Desai, P. Ilten et al., *An introduction to PYTHIA 8.2*, *Comput. Phys. Commun.* **191** (2015) 159 [1410.3012].
- [177] DELPHES 3 collaboration, *DELPHES 3, A modular framework for fast simulation of a generic collider experiment*, *JHEP* **02** (2014) 057 [1307.6346].
- [178] P.T. Komiske, E.M. Metodiev and J. Thaler, *Energy flow polynomials: A complete linear basis for jet substructure*, *JHEP* **04** (2018) 013 [1712.07124].
- [179] P.T. Komiske, E.M. Metodiev and J. Thaler, *Energy Flow Networks: Deep Sets for Particle Jets*, *JHEP* **01** (2019) 121 [1810.05165].

- [180] P.T. Komiske, E.M. Metodiev and J. Thaler, *Cutting Multiparticle Correlators Down to Size*, *Phys. Rev. D* **101** (2020) 036019 [[1911.04491](#)].
- [181] A. Andreassen, P.T. Komiske, E.M. Metodiev, B. Nachman and J. Thaler, *OmniFold: A Method to Simultaneously Unfold All Observables*, *Phys. Rev. Lett.* **124** (2020) 182001 [[1911.09107](#)].
- [182] CMS collaboration, *Measurement of the inclusive and differential WZ production cross sections, polarization angles, and triple gauge couplings in pp collisions at $\sqrt{s} = 13$ TeV*, *JHEP* **07** (2022) 032 [[2110.11231](#)].
- [183] CMS collaboration, *Measurement of the inclusive WZ production cross section in pp collisions at $\sqrt{s} = 13.6$ TeV*, *JHEP* **04** (2025) 115 [[2412.02477](#)].
- [184] PARTICLE DATA GROUP collaboration, *Review of particle physics*, *Phys. Rev. D* **110** (2024) 030001.
- [185] Delphes default CMS card.
https://github.com/delphes/delphes/blob/master/cards/delphes_card_CMS.tcl.
- [186] CMS collaboration, *Precision luminosity measurement in proton-proton collisions at $\sqrt{s} = 13$ TeV in 2015 and 2016 at CMS*, *Eur. Phys. J. C* **81** (2021) 800 [[2104.01927](#)].
- [187] ATLAS collaboration, *Search for a new Z' gauge boson in 4μ events with the ATLAS experiment*, *JHEP* **07** (2023) 090 [[2301.09342](#)].
- [188] F. del Aguila, M. Chala, J. Santiago and Y. Yamamoto, *Collider limits on leptophilic interactions*, *JHEP* **03** (2015) 059 [[1411.7394](#)].
- [189] CMS collaboration, *Measurements of $pp \rightarrow ZZ$ production cross sections and constraints on anomalous triple gauge couplings at $\sqrt{s} = 13$ TeV*, *Eur. Phys. J. C* **81** (2021) 200 [[2009.01186](#)].
- [190] ATLAS collaboration, *Searches for new phenomena in events with two leptons, jets, and missing transverse momentum in 139 fb^{-1} of $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector*, *Eur. Phys. J. C* **83** (2023) 515 [[2204.13072](#)].
- [191] CMS collaboration, *Measurement of differential cross sections for Z boson production in association with jets in proton-proton collisions at $\sqrt{s} = 13$ TeV*, *Eur. Phys. J. C* **78** (2018) 965 [[1804.05252](#)].
- [192] G. Kasieczka, B. Nachman, D. Shih, O. Amram, A. Andreassen, K. Benkendorfer et al., *The lhc olympics 2020 a community challenge for anomaly detection in high energy physics*, *Reports on Progress in Physics* **84** (2021) 124201.
- [193] S. Yoon, Y.-K. Noh and F.C. Park, *Autoencoding under normalization constraints*, 2023.
- [194] K. An, *Sulla determinazione empirica di una legge di distribuzione*, *Giorn Dell'inst Ital Degli Att* **4** (1933) 89.
- [195] T.W. Anderson, *On the distribution of the two-sample cramer-von mises criterion*, *The Annals of Mathematical Statistics* (1962) 1148.