

Multiscale optimal transport for complete collider events

Tianji Cai^{*}

*School of Physical Science and Engineering, [Tongji University](#), Shanghai 200092, China
and State Key Laboratory of Autonomous Intelligent Unmanned Systems,
MOE Frontiers Science Center for Intelligent Autonomous Systems,
Tongji University, Shanghai 200092, China*

Nathaniel Craig[†]

*Department of Physics, [University of California](#), Santa Barbara, California 93106, USA
and [Kavli Institute for Theoretical Physics](#), Santa Barbara, California 93106, USA*

Katy Craig[‡]

Department of Mathematics, [University of California](#), Santa Barbara, California 93106, USA

Xinyuan Lin[§]

Department of Physics, [University of California](#), San Diego, California 92093, USA



(Received 9 February 2025; accepted 8 July 2025; published 29 August 2025)

Building upon the success of optimal transport metrics defined for single collinear jets, we develop a multiscale framework that models entire collider events as distributions on the manifold of their constituent jets, which are themselves distributions on the ground space of the calorimeter. This hierarchical structure of optimal transport effectively captures relevant physics at different scales. We demonstrate the versatility of our method in two event classification tasks, which respectively emphasize intrajet substructure and interjet spatial correlations. Our results highlight the importance of a nested structure of manifolds in the treatment of full collider events, broadening the applicability of optimal transport methods in collider analyses.

DOI: [10.1103/7nhv-hndt](#)

I. INTRODUCTION

Distance-based representations of collider data have been shown to provide powerful inputs to modern machine learning methods (see, e.g., [1–9] for broad overviews). Treating collider events as distributions on the calorimeter naturally suggests the use of optimal transport (OT) distances, which quantify the “work” required to rearrange one distribution into another. In recent years, optimal transport methods have opened promising new avenues for analyzing LHC data [10–24].¹

In the usual application of OT to collider physics, both jets and entire events are considered as transverse momentum distributions in the rapidity-azimuth plane of the calorimeter. Optimal transport distances between these distributions are typically quite effective whenever the underlying task hinges on a single intrinsic scale in the data. However, entire collider events often feature multiple intrinsic scales, such as the scales associated with jet structure and substructure. This raises the natural question: What is the distance between multiscale collider events?

In this paper, we propose a distance between complete collider events that models them as distributions on the manifold of jets. Inspired by previous work, we endow the manifold of jets with a type of 2-Wasserstein (W_2) distance, induced by the Euclidean ground distance in the rapidity-azimuth plane of the collider. However, unlike previous work, we introduce additional weights which allow us to adjust the relative importance of jet substructure and the locations of the jets within the full collider event. Then, we endow the manifold of events with another 2-Wasserstein distance, induced by the above ground distance on the manifold of jets. In this way, our approach can be thought of as “ W_2 -on- W_2 .” This hierarchical structure allows us to

^{*}Contact author: tianjiresearch@gmail.com

[†]Contact author: ncraig@ucsb.edu

[‡]Contact author: kcraig@math.ucsb.edu

[§]Contact author: xil244@ucsd.edu

¹For a broad overview, see [25]. Other suitable distances between collider events have recently been proposed in [26,27]. Optimal transport has also been employed to great effect within machine learning architectures for collider physics, see, e.g., [28–33].

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

effectively distinguish between differences between full collider events, at various scales of the data.

Our paper is organized as follows. Section II gives an overview of our methodology. First, the usual single-scale optimal transport method is presented, with discussions on suitable representations that account for relevant symmetries acting at the jet and event level. We then introduce the new *multiscale optimal transport* framework which incorporates both intrajet substructure and interjet spatial correlations for quantifying the dissimilarity between two complete events. Section III applies the new method to two event classification tasks, i.e., $t\bar{t}$ vs QCD classification and beyond the Standard Model (BSM) paired dijet vs QCD classification, underlining the broad applicability of our method. The paper concludes with Sec. IV, which highlights several promising directions for future work using the multiscale OT framework.

II. METHODOLOGY

A. The 2-Wasserstein distance

We begin with a brief overview of OT theory relevant to collider physics, focusing on the p -Wasserstein distance and the linearization of the $p = 2$ case.

Consider an event represented by its energy flow, i.e., a discrete distribution $\mathcal{E} = \sum_i E_i \delta_{x_i}$ where the i th constituent particle, assumed to be massless, has an energy E_i and is located at position x_i in some predefined ground space Ω . Here an “event” could denote either a jet or the calorimeter deposit of an entire collision event.²

Given events $\mathcal{E}$ and $\mathcal{E}'$ with the same total energy, i.e., $\sum_i E_i = \sum_{j'} E'_{j'}$, the p -Wasserstein distance between them is defined in terms of the minimization problem

$$W_p(\mathcal{E}, \mathcal{E}') = \min_{g_{ij'} \in \Gamma(\mathcal{E}, \mathcal{E}')} \left(\sum_{ij'} g_{ij'} (d_{\text{ground}}(x_i, x'_{j'}))^p \right)^{1/p},$$

$$\Gamma(\mathcal{E}, \mathcal{E}') = \left\{ g_{ij'} : g_{ij'} \geq 0, \sum_{j'} g_{ij'} = E_i, \sum_i g_{ij'} = E'_{j'} \right\},$$

(1)

where the *transport plan* $g_{ij'} \in \Gamma(\mathcal{E}, \mathcal{E}')$ represents how energy is transferred from particle i in event $\mathcal{E}$ to particle j' in event $\mathcal{E}'$. The W_p distance measures the difference between two events based on the *minimal effort* required to transfer all the energy from $\mathcal{E}$ to $\mathcal{E}'$, where the *effort* of a transport plan $g_{ij'}$ is proportional to $d_{\text{ground}}(x_i, x'_{j'})$, penalizing the distance over which energy is moved in the

ground space. The $g_{ij'}$ that achieves the minimum in Eq. (1) is known as the *optimal transport plan*.

It is clear from the definition that the choice of ground metric encodes the essential physics of the problem at hand. In the context of hadron colliders, a natural ground space is the physical rapidity vs azimuthal-angle plane of the cylindrical collider, i.e., $\Omega = (y, \phi) \in \mathbb{R}^2$, with the ground metric being the Euclidean distance $d_{\text{ground}}(x_i, x'_{j'}) = \|(y_i, \phi_i) - (y'_{j'}, \phi'_{j'})\|$ in this y - ϕ plane. Furthermore, the “energy” E_i is often replaced by the transverse momentum p_T of the particle. We refer to Eq. (1) defined on this ground space as the *single-scale optimal transport* problem.

Different choices of p penalize the ground distance at different length scales. Two common choices are the W_1 metric, also called the Earth Mover’s distance [10,11,14,15], and the 2-Wasserstein (W_2) distance [12,13,34]. The W_2 distance is unique due to its formal Riemannian structure [35–37], which lends itself conveniently to efficient linearization, as discussed below. Therefore, the default choice in our study is $p = 2$.

In general, optimal transport distances are more expensive to compute than classical distances between distributions, such as the l^2 norm. For events containing n constituent particles, the Bertsekas’ auction algorithm for the W_p distance demands $\mathcal{O}(n^3)$ operations, while entropic regularization and the Sinkhorn algorithm requires $\mathcal{O}(n^2 \log(n))$ [38–40]. In comparison, after binning two distributions, the l^2 norm between distributions represented on a grid of size n only requires $\mathcal{O}(n)$ steps. Furthermore, to use OT to compare the mutual differences in a typical collider dataset containing $\mathcal{O}(10^5)$ events, one would need to calculate $\mathcal{O}(10^{10})$ pairwise OT distances, exacerbating the computational demand.

One way to mitigate this problem is to leverage the Riemannian structure of W_2 to *linearize* the distance [12,41]. In this approach, one projects all events onto the tangent plane of the event manifold at a specific reference event $\mathcal{R} = \sum_i R_i \delta_{y_i}$.³ If r_{ij} denotes the optimal transport plan between a given event $\mathcal{E}$ and the reference $\mathcal{R}$, the projection is

$$\mathcal{E} \mapsto z_i := \frac{1}{R_i} \sum_j r_{ij} x_j. \quad (2)$$

While the numerical cost of this projection is comparable to the cost of computing a W_2 distance, the projection only needs to be computed once per event. Then, all pairwise comparisons between events can be done via cheap l^2 norm computations in the tangent plane, i.e., the difference

²In the usual application of OT to collider physics, both jets and events are considered as *single-scaled* transverse momentum distributions in the rapidity-azimuth plane, with no distinction made among possible intrinsic scales in the data. In this subsection, we will refer to both as “events,” but will distinguish them beginning in Sec. II B.

³While the rigorous geometric theory only holds for reference events that do not give mass to small sets, for simplicity, we explain the approach in the fully discrete case and refer to [12] for more detail.

between events $\mathcal{E}$ and $\mathcal{E}'$ is given by

$$\text{LinW}_2(\mathcal{E}, \mathcal{E}') = \left(\sum_i R_i \|z_i - z'_i\|^2 \right)^{1/2}. \quad (3)$$

While LinW_2 itself is not a metric, it converges to a true metric if the reference event satisfies certain properties (see Proposition 1 in [12]), which are explicitly enforced in the construction of the reference in our study.

B. Representation of jets and events

In order to use OT to quantify the dissimilarity between energy flows (whether single jets or entire events), one must represent the objects in a manner that takes into account inherent symmetries, so that the OT distance between two distributions that are identical, up to symmetry, is zero.

We consider jets as equivalence classes of distributions $\mathcal{J} = \sum_i p_T^i \delta_{(y_i, \phi_i)}$ in the y - ϕ plane, where we identify two jets that are identical up to rotational symmetry.⁴ To remove the rotational symmetry on the y - ϕ plane, we translate each jet so that its axis is at $(y, \phi) = (0, 0)$, align the principal component of each jet, and then translate each jet back to its original location. The principal component of each jet [42,43] is given by the principal eigenvector of its moment tensor,

$$\mathcal{T} = \sum_{i=1}^n p_T^i \begin{pmatrix} y_i^2 & y_i \phi_i \\ y_i \phi_i & \phi_i^2 \end{pmatrix}. \quad (4)$$

As the principal component is only defined up to a sign, we rotate the jet so that its principal component is along the ϕ axis and reflect the jet to ensure that more total p_T resides in the upper part of the y - ϕ plane, i.e., $\phi > 0$. Note that in the case that $p_T^i = \frac{1}{n}$ for all i , this coincides with the usual principal component analysis (PCA).

In a similar spirit, we consider entire collider events (henceforth “events”) as equivalence classes of distributions $\mathcal{E} = \sum_s J_s \delta_{\mathcal{J}_s}$ on the manifold of jets, where we identify events that are equal up to symmetry. To build this representation, we must both remove symmetries and translate the complete collision event from a distribution on the plane of the detector into a distribution over its composite jets.

To accomplish this, we first boost an event from the lab frame back to its own c.m. frame, so that the event’s c.m. is at $y = 0$. Next, due to the fragmentation of high-energy quarks and gluons, the final-state hadrons of a complete

collider event are usually clustered around a few collimated jets. Using the anti- k_T jet clustering algorithm [44], in which a radius R determines the angular extent of a jet, we partition the final-state hadrons of an event $\mathcal{E}$ into distinct jets $\mathcal{J}_s$. The entire collider event is then represented by $\mathcal{E} = \sum_s J_s \delta_{\mathcal{J}_s}$, where there are a variety of choices for how to define the weight J_s of the s th jet in the distribution; we devote particular attention to the cases of weighting jets uniformly and in terms of their total p_T (see Sec. II C). Lastly, we translate all jets $\mathcal{J}_s$ simultaneously so that the axis of the jet with the highest p_T is at $\phi = 0$. This preprocessing approach guarantees that events differing only by a boost and/or rotation around the beam axis are represented identically.⁵

Finally, as described above, while we typically represent the jets in an event as equivalence classes of distributions, up to rotational symmetry, in our comparisons of our multiscale approach with classical single-scale optimal transport, we also consider the case in which the jets are not rotated to align the principle access. We distinguish between these two representations of events in the single-scale optimal transport case as $X_{\text{c.m.}+\text{jet}}$, when the jets are rotated, and $X_{\text{c.m.}}$, when they are not.

C. Multiscale optimal transport

The partition of complete collider events into jets gives rise to two significant length scales on the y - ϕ plane. Within a single jet, the relevant length scale is $\mathcal{O}(\frac{R}{N})$, where $R \sim \mathcal{O}(1)$ is the jet radius and $N \sim \mathcal{O}(100)$ is particle multiplicity, so that intrajet substructure is manifested on a length scale of $\mathcal{O}(1/100)$. On the other hand, the jets that compose a given event are distributed on the cylindrical calorimeter in a highly nonuniform way, with small clusters separated by large regions mostly devoid of hard particles; see Fig. 1. The separation between jets (red arrow) has a length scale of $\mathcal{O}(\frac{C}{M})$, where $C \sim \mathcal{O}(5)$ is the diameter of the detector and $M \sim \mathcal{O}(5)$, leading to an $\mathcal{O}(1)$ length scale for the interjet separation. Since this is a hundred times larger than the intrajet scale, a single-scale W_2 distance between entire events as distributions in the y - ϕ plane would be dominated by the spatial separation between jets,

⁵Comparing to the jet preprocessing scheme, here we do not rotate an entire event on the y - ϕ plane, as such a transformation is not physically meaningful outside the collinear limit. In principle, if the event enjoys a perfect $O(3)$ symmetry, then an extra rotational degree of freedom around the $y, \phi = 0$ axis needs to be fixed. One could, for example, further require the axis of the second highest p_T jet to be in the $\phi = 0$ plane. However, this symmetry is broken at hadronic colliders as the beam axis (i.e., the $+z$ axis) is unique, due to the fact that the actual four-momenta of the underlying colliding partons are unknown and that a large fraction of final-state particles down the beam line (i.e., close to the $+z$ axis) are not measured by detectors. We therefore do not include this rotational symmetry in our event preprocessing schemes.

⁴For many events, the distribution of jets around their jet axes are correlated, for example by color flow or spin correlations. For simplicity, in the spirit of simplified models, we proceed with our current preprocessing scheme, though such correlation could naturally be included by a mild modification of the ground metric on the jet manifold.




FIG. 1. A sample $t\bar{t}$ event composed of four jets (shown in different colors). Constituent particles are displayed on the y - ϕ plane, with sizes proportional to their p_T .

with little hope to capture the finer distinctions at the substructure level.

Instead, by representing events as distributions on the manifold of jets, i.e., $\mathcal{E} = \sum_s J_s \delta_{\mathcal{J}_s}$, we are able to circumvent this issue and include both small-scale substructure within jets and large-scale information about the jets' location on the detector. In particular, the difference between two events $\mathcal{E}$ and $\mathcal{E}'$ can be quantified in terms of a W_2 metric, where the ground distance itself is given by a modified W_2 distance on the jet manifold. Given $c_{\text{intra}}, c_{\text{inter}} > 0$, this is defined as follows:

$$W_{2,\text{event}}(\mathcal{E}, \mathcal{E}') = \min_{\pi_{s't'} \in \Gamma(\mathcal{E}, \mathcal{E}')} \left(\sum_{s't'} \pi_{s't'} d_{\text{jet}}^2(\mathcal{J}_s, \mathcal{J}'_{t'}) \right)^{1/2}, \quad (5)$$

where

$$d_{\text{jet}}(\mathcal{J}, \mathcal{J}') = ((c_{\text{intra}} W_{2,\text{jet}}(\mathcal{J}^*, \mathcal{J}'^*))^2 + (c_{\text{intra}} \|(\bar{y}, \bar{\phi}) - (\bar{y}', \bar{\phi}')\|)^2)^{1/2}, \quad (6)$$

with $W_{2,\text{jet}}(\mathcal{J}^*, \mathcal{J}'^*)$ as defined in Eq. (1) ($p = 2$). Here $(\bar{y}, \bar{\phi})$ represents the center of mass of a jet, $\|\cdot\|$ denotes the Euclidean distance on the collider cylinder, and $\mathcal{J}^*$ and $\mathcal{J}'^*$ represent the translations of jets $\mathcal{J}$ and $\mathcal{J}'$ to have zero center of mass. Due to the fact that two events could, in principle, be composed of a large number of jets, in practice, we use the Lin W_2 approximation for $W_{2,\text{jet}}$; see Eq. (3). Furthermore, we use the jet axis from the anti- k_T algorithm as a proxy for the c.m. of the jet. With these choices, Eq. (6) defines the *multiscale optimal transport* problem.

For any $c_{\text{intra}}, c_{\text{inter}} > 0$, d_{jet} is a metric on the manifold of jets, so $W_{2,\text{event}}$ is a metric on the manifold of complete collider events, where each event is represented as a distribution on the manifold of jets [45]. It is a classical

result that, when $c_{\text{intra}} = c_{\text{inter}} = c > 0$, $d_{\text{jet}}(\mathcal{J}, \mathcal{J}') = c W_{2,\text{jet}}(\mathcal{J}, \mathcal{J}')$ (see Proposition 2.5 in [46]), so that the ground metric on the jet manifold is exactly the 2-Wasserstein distance between the original untranslated jets.

We allow $c_{\text{intra}} \neq c_{\text{inter}}$ in order to capture the effect of the two distinct length scales in the composition of an event. A larger ratio of $c_{\text{intra}}/c_{\text{inter}}$ places a greater emphasis on the substructure (shape) of the jets composing an event, while a smaller ratio of $c_{\text{intra}}/c_{\text{inter}}$ highlights the relative locations of the jets in the collider plane. As will be seen later, for our specific datasets, $\|(\bar{y}, \bar{\phi}) - (\bar{y}', \bar{\phi}')\|$ are $\sim O(1)$ whereas $W_{2,\text{jet}}(\mathcal{J}^*, \mathcal{J}'^*)$ are $\sim O(0.1)$. Thus, the choice $(c_{\text{intra}}, c_{\text{inter}}) = (10, 1)$ roughly gives equal weighting on both terms.

As described in Sec. II B, when representing an event as a distribution on the manifold of jets, $\mathcal{E} = \sum_s J_s \delta_{\mathcal{J}_s}$, there are several potential choices one could make for the weights J_s of each jet $\mathcal{J}_s$. A first choice would be to define

$$J_s = \begin{cases} 1 & \text{if } \mathcal{J}_s \text{ is the hardest jet,} \\ 0 & \text{otherwise,} \end{cases}$$

where by *hardest jet*, we indicate the jet with the largest total p_T . In this case, the 2-Wasserstein distance between events simplifies to

$$W_{2,\text{event}}(\mathcal{E}, \mathcal{E}') = ((c_{\text{intra}} W_{2,\text{jet}}(\mathcal{J}_1^*, \mathcal{J}'_1{}^*))^2 + (c_{\text{inter}} \|(\bar{y}_1, \bar{\phi}_1) - (\bar{y}'_1, \bar{\phi}'_1)\|)^2)^{1/2}, \quad (7)$$

where $\mathcal{J}_1$ is the jet with the highest p_T in event $\mathcal{E}$ and $\mathcal{J}'_1$ for event $\mathcal{E}'$. While this choice of weighting is sensitive to both the substructure difference between the two hardest jets and their spatial separation on the y - ϕ plane, it does not take into account any of the features of softer jets. This simple method will be referred to as “Multiscale OT: Jet 1.”

More sophisticated choices of weighting are to either choose J_s uniformly, letting $J_s \equiv 1/M$, where M is the number of jets in event $\mathcal{E}$, or to choose J_s proportional to the transverse momentum of the s th jet, i.e.,

$$J_s = \frac{p_T^{\text{s}th \text{ jet}}}{\sum_{s'} p_T^{\text{s}th \text{ jet}}}.$$

For this second approach, jets with higher p_T are more influential in determining the overall OT distance between events. In what follows we will refer to these approaches as “Multiscale OT: uniform weighted OT” and “Multiscale OT: p_T weighted OT,” respectively.

D. A toy example

We close our discussion of the method with an illustration of its performance on a toy example of two artificial “events.” Here each event is composed of three jets, and we want to quantify the dissimilarity of the events in terms of




FIG. 2. (a) A series of multiscale OT distances between the two events under different $(c_{\text{intra}}, c_{\text{inter}})$ pairs, where only the event-level transport plans are shown, denoting the matching among jets as a whole; (b) Single-scale OT between the two events, showing the transport plan among the particles between the two events (represented by the gray lines).

both small-scale (shape of jets) and large-scale (physical location of jets) features. All particles that make up each jet are equally weighted, and the event is considered as a uniformly weighted distribution on its constituent jets. We choose a simple Euclidean x - y plane as the ground space. We compare the performance of a classical single-scale W_2 distance between the distributions with our new multiscale OT distance. In both cases, the distance is determined by the corresponding optimal transport plan. Consequently, we illustrate the differences by comparing the structure of the optimal transport plans in each case. Further details of the setup are given in Appendix A.

We first compute the single-scale W_2 distance between $\mathcal{E}$ and $\mathcal{E}'$, using Eq. (1) with $p = 2$. Figure 2(b) shows the optimal transport plan among all particles, which only considers the vicinity of the particles on the y - ϕ plane. To also capture the difference in the shape of the jets, we need the multiscale OT distance. In Fig. 2(a), we consider five pairs of coefficients,

$$(c_{\text{intra}}, c_{\text{inter}}) = (1, 0), (42, 1), (30, 1), (12, 1), (0, 1),$$

which are chosen to illustrate the role of small scale vs large scale effects. As expected, when $c_{\text{intra}}/c_{\text{inter}}$ is small, the event level optimal transport plan prefers to match the jets between the events based on the location of the jet in physical space. Conversely, when $c_{\text{intra}}/c_{\text{inter}}$ is large, the event level optimal transport plan prefers to match the jets between the events based on the similarity of their shape.

III. APPLICATION TO EVENT CLASSIFICATION

We now apply both single-scale and multiscale OT frameworks to two different event classification tasks to illustrate the utility of the latter. The first task aims to distinguish boosted $t\bar{t}$ events from QCD events, while the

second task aims to distinguish a pair-produced BSM particle decaying into dijets from QCD. At a typical jet radius of $R = 0.5$, the two tasks respectively benefit from intrajet and interjet information.

Although we use $R = 0.5$ jets for our primary benchmark, it is also illuminating to see how intrajet and interjet contributions change as the jet radius is varied. As R decreases, interjet information becomes more important as decays of parent particles are increasingly resolved and the radiation pattern within each jet exhibits less structure. Conversely, intrajet information becomes more important as R increases and the decay products of parent particles are clustered together.

Once a distance between events is obtained using either of the two OT frameworks, a classifier is needed to take the distance as input and output the class label for each event in a supervised manner. Given that the focus of our study is on the role of the metrics between events, not on the classifier, we simply use the k -nearest neighbor (kNN) algorithm [47], as it depends solely on the notion of pairwise distances. kNN classifies a new data point based on the majority of its k nearest neighbors, with neighbors defined by the user-specified distance. Other more sophisticated classifiers, including neural networks, can be coupled to both single-scale and multiscale OT frameworks; we leave this for future study.

Our dataset for each classification task always consists of 5000 background events and 5000 signal events, further split into 50% for training the classifier, 25% for validation, and 25% for testing. The optimal value of the model hyperparameter k is selected using the validation dataset, with k ranging from 10 to 200 with an increment of 10. The performance of kNN is evaluated on the held-out test set, which outputs the final area under the curve (AUC) values.

For comparison with non-OT methods, we choose the classical event shape observable *transverse thrust* [48–50],

which is widely used in pp collisions and served as a benchmark for OT-based event isotropy studies [15]. The transverse thrust is defined as

$$T_{\perp} = \max_{\hat{n}_{\perp}} \frac{\sum_i |\vec{p}_{T,i} \cdot \hat{n}_{\perp}|}{\sum_i |\vec{p}_{T,i}|}, \quad (8)$$

where $\vec{p}_{T,i}$ is the transverse momentum of the i th particle in the event, and $\hat{n}_{\perp}$ is a unit vector on the transverse plane. $T_{\perp}$ ranges from $\frac{2}{\pi}$ for a fully isotropic event, to 1 for a completely collinear event. For convenience and following the convention in [15], we use a rescaled version of $T_{\perp}$ so that the range becomes $[0, 1]$. That is, each event corresponds to a single scalar,

$$\tilde{T}_{\perp} = \frac{\pi}{\pi - 2} \left(T_{\perp} - \frac{2}{\pi} \right). \quad (9)$$

We apply a cut on $\tilde{T}_{\perp}$ to separate signal events from background events. A sweep over the cut value then outputs the AUC value for the classification using this $\tilde{T}_{\perp}$ observable, which may be compared with the two OT methods.

A. $t\bar{t}$ vs QCD event classification

All events are generated in MadGraph 2.9.16 [51] using pp collisions with a c.m. energy $\sqrt{s} = 13$ TeV. For QCD events, we consider $pp \rightarrow jjjj$, where a jet j is originated by either a gluon or u, d, c, s quarks (and the corresponding antiquarks). For $t\bar{t}$ pair production, we consider $pp \rightarrow t\bar{t} \rightarrow W^+ b W^- \bar{b} \rightarrow (jj)(jj)bb$ events, where the boosted W boson decays hadronically. To increase event yield, a minimum p_T cut of 150 GeV is applied to each j product of the QCD events, whereas a $p_T > 450$ GeV cut is applied to the top quarks themselves. Hadronization and showering are done in Pythia 8.310 [52] with default parameters. Events are then clustered into jets with FastJet 3.4 [53], using the anti- k_T algorithm [44] with a jet radius of $R = 0.5$. Only events with at least four jets are retained. A rapidity cutoff of $y \leq 4$ is applied to each jet. The two highest p_T jets (labeled Jet 1 and Jet 2) are required to have $p_T > 300$ GeV, whereas the next two jets (labeled Jet 3 and Jet 4) have a threshold of $p_T > 150$ GeV. In cases where more than four jets are identified in an event, we only keep the first four jets with the largest p_T and discard the rest. These four jets define an event in our dataset.

In order to study the effect of varying jet radius, we further create two more datasets based on the above dataset. For each event in the above $R = 0.5$ dataset containing exactly four jets, we recluster all its constituent particles using the anti- k_T algorithm with two new radii, respectively $R = 0.3$ and $R = 1.0$. We emphasize that this procedure




FIG. 3. AUC scores for the $t\bar{t}$ vs QCD event classification task, when events are clustered into jets with radius $R = 0.5$ (upper), 0.3 (lower left), and 1.0 (lower right). The colored curves denote the respective results for the three multiscale OT methods, i.e., “ p_T weighted OT,” “uniform weighted OT,” and “Jet 1,” as we vary the values for the coefficient pair $(c_{\text{intra}}, c_{\text{inter}})$ (expressed in terms of “ $c_{\text{intra}} : c_{\text{inter}}$ ” on the x -axis). The performance of the single-scale OT framework are shown in gray dashed lines, with the two different preprocessing schemes, i.e., “c.m.” and “c.m. + Jet.” Also included as benchmark is the performance of the transverse thrust (gray solid line), where events are preprocessed according to the “c.m.” scheme.

only includes particles that had been clustered in the leading four $R = 0.5$ jets of each original event.⁶ Although artificial, this provides the clearest illustration of the interplay between the jet radius R and the contributions from intrajet and interjet distributions.

Figure 3 shows the classification performance (measured in AUC) on the test set for both the single-scale and multiscale OT frameworks coupled to a kNN classifier, where jets are clustered with different radii $R = 0.3, 0.5, 1.0$. Also displayed for comparison is the performance of transverse thrust, which stays the same for all three datasets as the underlying events are the same.

When using the single-scale OT method, events are preprocessed according to either the “c.m.” or the “c.m. + jet” schemes introduced in Sec. II B. As can be seen, the two preprocessing schemes make little difference to the

⁶The jet mass distribution of $R = 0.5$ jets and the impact of reclustering on jet mass and multiplicity distributions is illustrated in Appendix B.

classification performance. This suggests that the additional jet-level preprocessing step generates no meaningful improvement under the single-scale OT framework. Such behavior is to be expected, as the existence of multiple scales intrinsic to a full collision event is being ignored.

The multiscale OT framework shows improved performance in almost all cases considered. One interesting coefficient pair is $(c_{\text{intra}}, c_{\text{inter}}) = (1, 0)$, where no interjet spatial correlation is used and the information given to both single-scale and multiscale OT frameworks is the same. In spite of the same inputs, for $R = 0.5$ and 1.0 , the multiscale OT method still outputs an AUC about 25% higher than that of the two single-scale benchmarks. Such a large improvement underlines the importance of the ability to highlight substructure variation via the ground metric of Eq. (6).

For more general choices of $(c_{\text{intra}}, c_{\text{inter}})$, all three combination methods under the multiscale framework (i.e., “ p_T weighted OT,” “uniform weighted OT,” and “Jet 1”) give relatively comparable performance across different R values, with slight preference to “ p_T weighted OT.” That “uniform weighted OT” matches the performance of “ p_T weighted OT” is to be expected, as all included jets in a given event have more or less similar p_T , especially for larger R . On the other hand, the “Jet 1” method performs well overall in part because of the nature of this particular classification task, as most of the hardest jets are W or t jets for $R = 0.5$ and 1.0 whose substructure alone contains enough information to identify signal events from backgrounds. In Sec. III B, we will see that this is sensitive to the classification task.

The three multiscale methods perform best when more weight is put on intrajet substructure, which contains key discriminative information for this task. Indeed, for $R = 1.0$, the best AUC of 0.782 is obtained with the coefficient pair $(c_{\text{intra}}, c_{\text{inter}}) = (1, 0)$, and the performance drops when more emphasis is given to interjet spatial correlation. For smaller $R = 0.3$ or 0.5 , the highest AUC is achieved with the “ p_T weighted OT” method around $(c_{\text{intra}}, c_{\text{inter}}) = (10, 1)$, which is roughly where intrajet and interjet parts contribute equally. Particularly for $R = 0.3$, the decay products are expected to become more resolved and the corresponding jets have less substructure with which to distinguish them from QCD jets. In this case, interjet spatial correlation becomes more important, as verified by the differing trend of the curves with respect to the $(c_{\text{intra}}, c_{\text{inter}})$ ratio, comparing to larger R values.

Additionally, for $R = 0.5$ and 1.0 , the performance of the “ p_T weighted OT” and “uniform weighted OT” methods decline sharply around $(c_{\text{intra}}, c_{\text{inter}}) = (1, 1)$, whereas for the “Jet 1” method, the AUC displays a noticeable drop only when there is no contribution from intrajet substructure. In these cases, when the jet radii are larger, the highest p_T jet is more likely to be a complete boosted jet, i.e., a top jet, and its substructure carries more discriminative

information. Since the event preprocessing procedure aligns the axis of its hardest jet, the interjet spatial correlation contributes less to the event W_2 distance for the “Jet 1” method than the other two multiscale methods. Consequently, while the other multiscale schemes deteriorate as too much emphasis is placed on the (less informative) interjet differences, “Jet 1” doesn’t deteriorate until ratio $(0, 1)$, when events are compared exclusively based on interjet location differences (ignoring substructure). Furthermore, at ratio $(0, 1)$, “Jet 1” is worse than other multiscale approaches, since, in this case, the method only classifies based on the interjet location differences of a single jet in each event, as compared to the other multiscale methods which consider multiple jets in each event. The same phenomenon also occurs in the BSM vs QCD event classification with $R = 1.0$ in Sec. III B, for similar reasons.

Finally, even when $(c_{\text{intra}}, c_{\text{inter}}) = (0, 1)$, multiscale OT (with the exception of the “Jet 1” method as explained above) still gives noticeably better performance than single-scale OT, again affirming the importance of a multiscale treatment.

B. BSM vs QCD event classification

The second classification task involves discrimination between a paired dijet signal from BSM physics and the corresponding four-jet Standard Model QCD background. Whereas jet substructure played a significant role in the first classification task, here the jet substructure plays a less decisive role compared to the jet spatial distributions themselves.

We consider a new BSM scalar particle ϕ , which carries electromagnetic charge $+\frac{1}{3}$ and has mass $m_\phi = 100$ GeV with width $\Gamma_\phi = 2$ GeV. It is pair produced via $pp \rightarrow \phi\bar{\phi}$ and subsequently decays into two QCD jets, with equal coupling to all quark pairs that conserve charge. The Feynman rules for the new particle are computed using FeynRules [51], and passed to MadGraph for the computation of the matrix elements. Both the BSM signal events and QCD background events involve a $4j$ final state, and a minimum jet p_T of 150 GeV is imposed at simulation level. Hadronization and showering are again conducted with Pythia and jet clustering is performed in FastJet using the anti- k_T algorithm with $R = 0.5$. A rapidity cutoff of $y \leq 4$ and a p_T threshold of $p_T \geq 150$ GeV for each jet are applied. As before, we only select events with at least four jets and retain only the four largest- p_T jets for each chosen event. As with the first classification task, we create two additional datasets by reclustering the particles in the $R = 0.5$ dataset using $R = 0.3$ and $R = 1.0$. The jet mass distribution of $R = 0.5$ jets and the impact of reclustering on jet mass and multiplicity distributions is illustrated in Appendix C.

Figure 4 shows the performance of transverse thrust, single-scale OT, and multiscale OT methods on the BSM vs




FIG. 4. AUC scores for the BSM vs QCD event classification task, when events are clustered into jets with radius $R = 0.5$ (upper), 0.3 (lower left), and 1.0 (lower right). The colored curves denote the respective results of the three multiscale OT methods, i.e., “ p_T weighted OT,” “uniform weighted OT,” and “Jet 1,” as we vary the values of the coefficient pair $(c_{\text{intra}}, c_{\text{inter}})$ (expressed in terms of “ $c_{\text{intra}} : c_{\text{inter}}$ ” on the x -axis). The performance of the single-scale OT framework are shown in gray dashed lines, with the two different preprocessing schemes, i.e., “c.m.” and “c.m. + Jet.” Also included as benchmark is the performance of the transverse thrust (gray solid line), where events are preprocessed according to the “c.m.” scheme.

QCD classification task. Here, the benchmark observable transverse thrust already performs well, with AUC around 0.8 for all three R choices. It is on the same par with (and even slightly better than) the single-scale OT framework.

The p_T -weighted and uniform-weighted multiscale OT methods exhibit the best performance for optimal values of $c_{\text{intra}} : c_{\text{inter}}$, and the interplay between jet radius and $c_{\text{intra}} : c_{\text{inter}}$ is as expected for these methods. For $R = 0.5$ the decay products of each ϕ are typically resolved, favoring larger c_{inter} ; this persists for smaller values of R . Conversely, by $R = 1.0$ the decay products of each ϕ are typically clustered into fat jets, favoring larger c_{intra} . That being said, the optimal performance for $R = 0.5$ is achieved around $(c_{\text{intra}}, c_{\text{inter}}) = (10, 1)$ and $(5, 1)$, where both interjet and intrajet contributions are comparable; additional information is captured by the intrajet component, given that the decay products from some ϕ decays are clustered together.

In contrast, the “Multiscale OT: Jet 1” method performs poorly at and below $R = 0.5$. This is to be expected, as at $R = 0.5$ the hardest jets in both signal and background events originate from individual quarks or gluons. For $R = 1.0$, however, the decay products of the ϕ are typically collected into fat jets, and the “Jet 1” method becomes comparable to the other two multiscale OT methods.

IV. CONCLUSIONS AND FUTURE DIRECTIONS

Optimal transport has proven to provide useful metrics between individual jets at hadron colliders [10–16, 18–21]. In this work, we have introduced a new multiscale OT framework that extends the original single-scale optimal transport methods to complete collision events, effectively leveraging information encoded in both intrajet and interjet distributions. We have achieved this by first representing jets as distributions on the calorimeter, and entire events as distributions on the manifold of jets.

We demonstrated the versatility of the multiscale OT distance by considering two event classification tasks of different natures. For almost all cases considered, the multiscale OT framework is consistently better than the original single-scale OT distance and the non-OT benchmark of transverse thrust. For standard $R = 0.5$ jets, the best performance is obtained by either p_T -weighted or uniformly-weighted multiscale methods, with numerically comparable contributions from interjet and intrajet structure. The broad efficacy of these methods suggests they may be useful in general search and classification tasks at colliders.

There are various avenues to further development. For our current work, each collision event is first normalized to have the same total “mass” of one unit, before computing the balanced 2-Wasserstein distance with Eq. (6). Such normalization ignores the potentially relevant information stored in the total “mass” difference between two events, with “mass” being either the total scalar p_T of the event in the “ p_T weighted OT” scheme, or the jet multiplicity in the “uniform weighted OT” scheme. Analogous to [34, 54], which used the Hellinger-Kantorovich extension of the 2-Wasserstein distance to compare individual jets with different total p_T , the Hellinger-Kantorovich approach could also be used in the present setting to capture the additional information of the total “mass” difference between events, and this may improve the performance of our current multiscale OT framework.

In collider experiments, one typically has access to more information than just calorimetry, such as particle charge and tracks. Such additional information can in principle also be incorporated into the definition of an optimal transport distance, and may significantly improve the sensitivity of the distance to the underlying hard process. Another future direction would be to combine our novel notion of distance between complete collider events with the optimal transport approach for background modeling

developed by Manole, *et al.* [55], which is similar in spirit to our approach, due to its nested use of optimal transport.

More generally, the multiscale OT framework developed here can be applied to many scenarios within and beyond collider physics. Any data that features multiple length scales and can be described in terms of distributions on similarly nested manifolds could benefit from a comparable approach.

ACKNOWLEDGMENTS

The authors thank the referee for their valuable suggestions. T. C. would like to further thank Rikab Gambhir for suggesting the inclusion of missing references. The work of K. C. has been supported by the National Science Foundation under Grant No. NSF DMS-2145900. The work of T. C. and N. C. was supported in part by the U.S. Department of Energy under Grant No. DE-SC0011702 and performed in part at the Kavli Institute for Theoretical Physics, supported by the National Science Foundation under Grant No. NSF PHY-1748958. T. C. was additionally supported by the U.S. Department of Energy Awards No. DE-FOA-0002705, KA/OR55/22 (AIHEP), and No. DE-AC02-76SF00515.

DATA AVAILABILITY

The data are not publicly available. The data are available from the authors upon reasonable request.

APPENDIX A: TOY EXAMPLE

Here we give more details on the two artificial events in the toy example in Sec. IID. Both events consist of a circular “jet,” a triangular “jet,” and a rectangular “jet,” mimicking the morphological difference of real jets; see Fig. 5. The total “mass” of each of the three jets is set to be $\frac{1}{3}$, giving unit mass to both events. The particles within each jet all have the same mass $\frac{1}{3n}$, where n is the number of particles inside each jet.

In particular, the particles in the first event $\mathcal{E}$ are:

- (i) Jet 1 (circular): 30 points uniformly sampled within a circle centered at $(-1, 0.5)$ with a radius of 0.3;
- (ii) Jet 2 (triangular): 40 points uniformly sampled within an equilateral triangle centered at $(1, 0.5)$ with a height of 0.6; and
- (iii) Jet 3 (rectangular): 50 points uniformly sampled within a square centered at $(0, -1)$ with a side of 0.6.

The locations of the three jets are chosen so that the center of mass of the whole event is approximately at $(0, 0)$. The second event $\mathcal{E}'$ is generated with slight modifications to $\mathcal{E}$:

- (i) Jet 1' (rectangular): 50 points uniformly sampled within a rectangle centered at $(-0.6, -0.8)$ with sides being 1.0 and 0.4;




FIG. 5. Two artificially generated events. Each event contains three jets, highlighted in different colors and with their overall shape delineated for illustrative purposes.

- (ii) Jet 2' (circular): 30 points uniformly sampled within an ellipse centered at $(0.6, 0.8)$ with axes being $a = 0.48$ and $b = 0.21$; and
- (iii) Jet 3' (triangular): 40 points uniformly sampled within an upside down equilateral triangle centered at $(0, -1.6)$ with height again being 0.6.

These numerical choices for the jet locations are made so that the center of mass of the complete event stays approximately the same for both events.

APPENDIX B: $t\bar{t}$ VS QCD EVENTS

In this and the next appendix, we provide supplemental information about the events used in the $t\bar{t}$ vs QCD and BSM vs QCD classification tasks, respectively.

The jet mass distributions of the four p_T -ordered $R = 0.5$ jets passing selection cuts in both the $t\bar{t}$ and QCD event samples are shown in Fig. 6. A noticeable fraction of the highest- p_T Jet 1 in the signal events are top jets, peaking around a mass of ~ 170 GeV, together with a majority of W jets with a mass peak around 80 GeV. Many of Jet 2 are also W jets, while both Jet 3 and Jet 4 are mostly QCD jets. We therefore expect intrajet substructure to play a significant role in the classification task.

The effects of reclustering the $R = 0.5$ jets with $R = 0.3, 1.0$ are shown in Fig. 7. Figures 7(a) and 7(b) show the number of jets per event when the radius is 0.3 and 1.0,




FIG. 6. Mass of the first four highest p_T jets (clustered with $R = 0.5$) in the signal $t\bar{t}$ events (top) and the background QCD events (bottom).




FIG. 7. Jet multiplicity in each event, where jets are clustered with a radius of (a) $R = 0.3$, or (b) $R = 1.0$. Note that the label on the left tick of the x -axis shows the number of jets in each bin, i.e., for example, in (b) roughly 700 $t\bar{t}$ events have 2 jets with $R = 1.0$. (c) Jet mass for the first largest p_T jet in the signal $t\bar{t}$ events (Jet 1), where jets are clustered with different jet radii $R = 0.3, 0.5, 1.0$.

respectively, while Fig. 7(c) shows the jet mass distribution of the highest p_T jet for $R = 0.3, 0.5, 1.0$. These distributions are consistent with the fact that intrajet contributions play a significant role in the classification task for $R = 0.5$ and $R = 1.0$, and less of a role for $R = 0.3$.

APPENDIX C: BSM VS QCD EVENTS

Figure 8 again plots the jet mass distribution of the first four highest p_T jets in the BSM signal events and QCD background events. Both Jet 1 and Jet 2 show a noticeable peak around $m_\phi = 100$ GeV, but are dominated by pure QCD jets. In comparison, the first two jets in the $t\bar{t}$ events have additional information coming from the easily identifiable W boson jets, and include far fewer pure QCD jets. This is consistent with the observation that both intrajet and interjet distributions play a role in the BSM vs QCD classification task with $R = 0.5$ jets.

The effects of reclustering the $R = 0.5$ jets with $R = 0.3, 1.0$ are shown in Fig. 9. For $R = 1.0$, the decay products of ϕ are all clustered together, and the first jet in the signal




FIG. 8. Mass for the first four highest p_T jets (clustered with $R = 0.5$) in the signal BSM events (top) and the background QCD events (bottom).




FIG. 9. Jet multiplicity in each event, where jets are clustered with a radius of (a) $R = 0.3$, or (b) $R = 1.0$. (c) Jet mass for the first largest p_T jet in the signal BSM events (Jet 1), where jets are clustered with different jet radii $R = 0.3, 0.5, 1.0$.

events is almost always a BSM jet. This is consistent with the improved performance of the “Jet 1” multiscale method for $R = 1.0$ jets. On the other hand, for both $R = 0.3$ and 0.5 , the vast majority of Jet 1 are QCD jets, consistent with the greater role of interjet contributions in the classification task.

[1] A. J. Larkoski, I. Moult, and B. Nachman, Jet substructure at the large hadron collider: A review of recent advances in theory and machine learning, *Phys. Rep.* **841**, 1 (2020).

[2] D. Guest, K. Cranmer, and D. Whiteson, Deep learning and its application to LHC physics, *Annu. Rev. Nucl. Part. Sci.* **68**, 161 (2018).

- [3] K. Albertsson *et al.*, Machine learning in high energy physics community white paper, *J. Phys. Conf. Ser.* **1085**, 022008 (2018).
- [4] A. Radovic, M. Williams, D. Rousseau, M. Kagan, D. Bonacorsi, A. Himmel, A. Aurisano, K. Terao, and T. Wongjirad, Machine learning at the energy and intensity frontiers of particle physics, *Nature (London)* **560**, 41 (2018).
- [5] G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, and L. Zdeborová, Machine learning and the physical sciences, *Rev. Mod. Phys.* **91**, 045002 (2019).
- [6] D. Bourilkov, Machine and deep learning applications in particle physics, *Int. J. Mod. Phys. A* **34**, 1930019 (2019).
- [7] M.D. Schwartz, Modern machine learning and particle physics, *Harvard Data Sci. Rev.* **3** (2021).
- [8] G. Karagiorgi, G. Kasieczka, S. Kravitz, B. Nachman, and D. Shih, Machine learning in the search for new fundamental physics, [arXiv:2112.03769](https://arxiv.org/abs/2112.03769).
- [9] P. Shanahan, K. Terao, and D. Whiteson, Snowmass 2021 computational frontier compf03 topical group report: Machine learning, [arXiv:2209.07559](https://arxiv.org/abs/2209.07559).
- [10] P. T. Komiske, E. M. Metodiev, and J. Thaler, Metric space of collider events, *Phys. Rev. Lett.* **123**, 041801 (2019).
- [11] P. T. Komiske, E. M. Metodiev, and J. Thaler, The hidden geometry of particle collisions, *J. High Energy Phys.* **07** (2020) 006.
- [12] T. Cai, J. Cheng, N. Craig, and K. Craig, Linearized optimal transport for collider events, *Phys. Rev. D* **102**, 116019 (2020).
- [13] T. Cai, J. Cheng, K. Craig, and N. Craig, Which metric on the space of collider events?, *Phys. Rev. D* **105**, 076003 (2022).
- [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**, 034009 (2020).
- [15] C. Cesarotti and J. Thaler, A robust measure of event isotropy at colliders, *J. High Energy Phys.* **08** (2020) 084.
- [16] C. Cesarotti, M. Reece, and M. J. Strassler, The efficacy of event isotropy as an event shape observable, *J. High Energy Phys.* **07** (2021) 215.
- [17] T. Cheng, J.-F. Arguin, J. Leissner-Martin, J. Pilette, and T. Golling, Variational autoencoders for anomalous jet tagging, *Phys. Rev. D* **107**, 016002 (2023).
- [18] 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**, 192 (2021).
- [19] G. Aad, B. Abbott *et al.* (The ATLAS Collaboration), Measurements of multijet event isotropies using optimal transport with the ATLAS detector, *J. High Energy Phys.* **10** (2023) 060.
- [20] N. Craig, J. N. Howard, and H. Li, Exploring optimal transport for event-level anomaly detection at the large hadron collider, [arXiv:2401.15542](https://arxiv.org/abs/2401.15542).
- [21] L. Gouskos, F. Iemmi, S. Liechti, B. Maier, V. Mikuni, and H. Qu, Optimal transport for a novel event description at hadron colliders, *Phys. Rev. D* **108**, 096003 (2023).
- [22] C. Cesarotti and M. LeBlanc, A field guide to event-shape observables using optimal transport, [arXiv:2409.13150](https://arxiv.org/abs/2409.13150).
- [23] R. Gambhir, A. J. Larkoski, and J. Thaler, SPECTER: Efficient evaluation of the spectral EMD, *J. High Energy Phys.* **12** (2025) 219.
- [24] D. Ba, A. S. Dogra, R. Gambhir, A. Tasissa, and J. Thaler, SHAPER: Can you hear the shape of a jet?, *J. High Energy Phys.* **06** (2023) 195.
- [25] T. Cai, Optimal transport for high energy physics, Ph.D. thesis, UC, Santa Barbara (main) (2023), <https://escholarship.org/uc/item/93n2s2ff>.
- [26] A. Mullin, S. Nicholls, H. Pacey, M. Parker, M. White, and S. Williams, Does SUSY have friends? A new approach for LHC event analysis, *J. High Energy Phys.* **02** (2021) 160.
- [27] T. Cai, J. Cheng, N. Craig, G. Koszegi, and A. J. Larkoski, The phase space distance between collider events, *J. High Energy Phys.* **09** (2024) 054.
- [28] G. Stein, U. Seljak, and B. Dai, Unsupervised in-distribution anomaly detection of new physics through conditional density estimation, [arXiv:2012.11638](https://arxiv.org/abs/2012.11638).
- [29] G. Kasieczka *et al.*, The LHC olympics 2020 a community challenge for anomaly detection in high energy physics, *Rep. Prog. Phys.* **84**, 124201 (2021).
- [30] J. N. Howard, S. Mandt, D. Whiteson, and Y. Yang, Learning to simulate high energy particle collisions from unlabeled data, *Sci. Rep.* **12**, 7567 (2022).
- [31] G. D. Guglielmo, F. Fahim, C. Herwig, M. B. Valentin, J. Duarte, C. Gingu, P. Harris, J. Hirschauer, M. Kwok, V. Loncar, Y. Luo, L. Miranda, J. Ngadiuba, D. Noonan, S. Ogren-ci-Memik, M. Pierini, S. Summers, and N. Tran, A reconfigurable neural network ASIC for detector front-end data compression at the HL-LHC, *IEEE Trans. Nucl. Sci.* **68**, 2179 (2021).
- [32] R. Kansal, J. Duarte, H. Su, B. Orzari, T. Tomei, M. Pierini, M. Touranakou, J.-R. Vlimant, and D. Gunopulos, Particle cloud generation with message passing generative adversarial networks, [arXiv:2106.11535](https://arxiv.org/abs/2106.11535).
- [33] J. H. Collins, An exploration of learnt representations of jets, [arXiv:2109.10919](https://arxiv.org/abs/2109.10919).
- [34] T. Cai, J. Cheng, B. Schmitzer, and M. Thorpe, The linearized Hellinger-Kantorovich distance, *SIAM J. Imag. Sci.* **15**, 45 (2022).
- [35] J. Lott, Some geometric calculations on wasserstein space, [arXiv:math/0612562](https://arxiv.org/abs/math/0612562).
- [36] L. Ambrosio, N. Gigli, and G. Savaré, *Gradient Flows: In Metric Spaces and in the Space of Probability Measures* (Springer Science & Business Media, New York, 2008).
- [37] F. Otto, The geometry of dissipative evolution equations: The porous medium equation, *Commun. Partial Differ. Equations* **26**, 101 (2001).
- [38] D. Bertsekas, The auction algorithm: A distributed relaxation method for the assignment problem, *Ann. Oper. Res.* **14**, 105 (1988).
- [39] J. Altschuler, J. Weed, and P. Rigollet, Near-linear time approximation algorithms for optimal transport via Sinkhorn iteration, *Adv. Neural Inf. Process. Syst.* **30**, 1961 (2018).
- [40] M. Cuturi, Sinkhorn distances: Lightspeed computation of optimal transportation distances, *Adv. Neural Inf. Process. Syst.* **26**, 2292 (2013).
- [41] W. Wang, D. Slepčev, S. Basu, J. A. Ozolek, and G. K. Rohde, A linear optimal transportation framework for

- quantifying and visualizing variations in sets of images, *Int. J. Comput. Vis.* **101**, 254 (2013).
- [42] G. Gur-Ari, M. Papucci, and G. Perez, Classification of energy flow observables in narrow jets, [arXiv:1101.2905](#).
- [43] J. Gallicchio and M. D. Schwartz, Quark and gluon jet substructure, *J. High Energy Phys.* **04** (2013) 090.
- [44] M. Cacciari, G. P. Salam, and G. Soyez, The anti- k_t jet clustering algorithm, *J. High Energy Phys.* **04** (2008) 063.
- [45] L. Ambrosio and N. Gigli, A user's guide to optimal transport, in *Modelling and Optimisation of Flows on Networks: Cetraro, Italy 2009*, Editors: Benedetto Piccoli, Michel Rascle (Springer, Berlin, Heidelberg, 2013), pp. 1–155.
- [46] J. Cuesta-Albertos, C. Matrán, and J. Rodríguez-Rodríguez, Approximation to probabilities through uniform laws on convex sets, *J. Theor. Probab.* **16**, 363 (2003).
- [47] T. M. Cover and P. E. Hart, Nearest neighbor pattern classification, *IEEE Trans. Inf. Theory* **13**, 21 (1967).
- [48] I. A. Bertram (D0 Collaboration), Jet results at the D0 experiment, *Acta Phys. Pol. B* **33**, 3141 (2002).
- [49] Z. Nagy, Next-to-leading order calculation of three-jet observables in hadron-hadron collisions, *Phys. Rev. D* **68**, 094002 (2003).
- [50] A. Banfi, G. P. Salam, and G. Zanderighi, Resummed event shapes at hadron-hadron colliders, *J. High Energy Phys.* **08** (2004) 062.
- [51] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H.-S. Shao, T. Stelzer, P. Torrielli, and M. Zaro, The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations, *J. High Energy Phys.* **07** (2014) 079.
- [52] C. Bierlich, S. Chakraborty, N. Desai, L. Gellersen, I. Helenius, P. Ilten, L. Lönnblad, S. Mrenna, S. Prestel, C. T. Preuss, T. Sjöstrand, P. Skands, M. Uthmeim, and R. Verheyen, A comprehensive guide to the physics and usage of Pythia 8.3, *SciPost Phys. Codebases* **2022**, 8 (2022).
- [53] M. Cacciari, G. P. Salam, and G. Soyez, FastJet user manual, *Eur. Phys. J. C* **72**, 1896 (2012).
- [54] T. Cai, N. Craig, K. Craig, and X. Lin, Multi-scale optimal transport for complete collider events, [arXiv:2501.10681](#).
- [55] T. Manole, P. Bryant, J. Alison, M. Kuusela, and L. Wasserman, Background modeling for double Higgs boson production: Density ratios and optimal transport, [arXiv:2208.02807](#).