

MIT Open Access Articles

SHAPER: can you hear the shape of a jet?

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Journal of High Energy Physics. 2023 Jun 29;2023(6):195

As Published: [https://doi.org/10.1007/JHEP06\(2023\)195](https://doi.org/10.1007/JHEP06(2023)195)

Publisher: Springer Berlin Heidelberg

Persistent URL: <https://hdl.handle.net/1721.1/151063>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of use: Creative Commons Attribution



SHAPER: can you hear the shape of a jet?

**Demba Ba,^{a,b} Akshunna S. Dogra,^{a,b,c,d} Rikab Gambhir,^{b,e} Abiy Tasissa^f
and Jesse Thaler^{b,e}**

^a*John A. Paulson School of Engineering and Applied Sciences, Harvard University,
Western Ave, Cambridge, MA, U.S.A.*

^b*The NSF AI Institute for Artificial Intelligence and Fundamental Interactions, U.S.A.*

^c*Department of Mathematics, Imperial College London,
Queen's Gate, London, U.K.*

^d*EPSRC CDT in Mathematics of Random Systems: Analysis, Modelling and Simulation,
Oxford, U.K.*

^e*Center for Theoretical Physics, Massachusetts Institute of Technology,
Massachusetts Ave, Cambridge, MA 02139, U.S.A.*

^f*Department of Mathematics, Tufts University,
College Ave, Medford, MA 02155, U.S.A.*

E-mail: demba@seas.harvard.edu, adogra@nyu.edu, rikab@mit.edu,
Abiy.Tasissa@tufts.edu, jthaler@mit.edu

ABSTRACT: The identification of interesting substructures within jets is an important tool for searching for new physics and probing the Standard Model at colliders. Many of these substructure tools have previously been shown to take the form of optimal transport problems, in particular the Energy Mover's Distance (EMD). In this work, we show that the EMD is in fact *the* natural structure for comparing collider events, which accounts for its recent success in understanding event and jet substructure. We then present a Shape Hunting Algorithm using Parameterized Energy Reconstruction (SHAPER), which is a general framework for defining and computing shape-based observables. SHAPER generalizes N -jettiness from point clusters to any extended, parametrizable shape. This is accomplished by efficiently minimizing the EMD between events and parameterized manifolds of energy flows representing idealized shapes, implemented using the dual-potential Sinkhorn approximation of the Wasserstein metric. We show how the geometric language of observables as manifolds can be used to define novel observables with built-in infrared-and-collinear safety. We demonstrate the efficacy of the SHAPER framework by performing empirical jet substructure studies using several examples of new shape-based observables.

KEYWORDS: Jets and Jet Substructure, Automation, Specific QCD Phenomenology

ARXIV EPRINT: [2302.12266](https://arxiv.org/abs/2302.12266)

Contents

1	Introduction	2
2	The unreasonable effectiveness of Wasserstein in collider physics	4
2.1	Event shapes, jet shapes, and geometric shapes	4
2.2	Measure-ing the energy flow	6
2.3	Geometrizing IRC safety	8
2.4	The importance of being faithful	10
3	Hearing shapes	11
3.1	Shape observables and shape parameters	12
3.2	Composing shapes	14
3.3	Examples of novel shapes	15
3.3.1	N -ringiness	15
3.3.2	N -diskiness	17
3.3.3	N -ellipsiness	17
3.3.4	... plus pointiness	18
3.3.5	... plus pileup	19
3.3.6	... and more!	19
4	The Shaper framework	20
4.1	The SHAPER algorithm	20
4.2	The dual formulation of Wasserstein	21
4.3	Reviewing the Sinkhorn divergence	22
4.4	Implementation details	23
5	Empirical studies with jets	25
5.1	Dataset	25
5.2	Benchmarking Sinkhorn: jet isotropy	26
5.3	Benchmarking optimization: N -subjettiness	28
5.4	Hearing gradients	29
5.5	Hearing jets ring (and disk, and ellipse)	32
5.6	Shapiness and shape parameters	32
5.7	Pileup-mitigating shapes	39
6	Conclusions and outlook	41
A	Energy flows as measures	43

B Constructing the Wasserstein metric	45
B.1 Shaping up the loss	46
B.2 Why Wasserstein?	47

1 Introduction

Collisions at the Large Hadron Collider (LHC) produce events with hundreds of particles in the final state, which must be carefully analyzed to extract information about the underlying physics. In order to make sense of these high-dimensional data, increasingly sophisticated observables are required that are well understood at both the theoretical and experimental levels. Event shape [1–4] and jet shape [5, 6] observables have played an important role in refining our understanding of the structure of high energy collisions, by relating hadronic final states to perturbatively accessible partonic degrees of freedom. Many shape observables, such as event thrust [1, 7, 8] and jet angularities [9, 10], have been computed to next-to-next-to-next-to leading log (N³LL) accuracy [11] and next-to-next-to leading log accuracy N²LL [12] in e^+e^- collisions, respectively. Shape observables have been extensively measured and used to search for new physics signatures [13–24].

It was shown in ref. [25] that many of these event shapes and jet shapes can be cast as optimal transport problems, using the Energy Mover’s Distance (EMD). The EMD was introduced in ref. [26] in order to provide a quantitative measure of the “distance” between two collider events, \mathcal{E} and \mathcal{E}' . The EMD is based off the “earth mover’s distance” from computer vision [27–31], which itself is a special case of the Wasserstein metric [32, 33]. The EMD has since seen many uses in collider physics applications, such as in building a metrized latent space of events [26, 34–36] and in event/jet tagging and classification [37–39]. The EMD has also been used to define a novel shape observable, the event isotropy [40–42], which probes how “uniform” an event \mathcal{E} looks by comparing it to the idealized isotropic event \mathcal{U} .

In this paper, we seek to explain the effectiveness of the Wasserstein metric, by showing that it is the unique metric on collider events that both is continuous and respects the detector geometry *faithfully*. As shown in ref. [26], continuity encodes the collider physics concept of infrared-and-collinear (IRC) safety. Geometric faithfulness is, to our knowledge, a new concept for the collider community, which allows statements to be made about spatial distributions of energy within events. After advocating for Wasserstein geometry, we then generalize the notion of event shapes and jet shapes, motivated by the EMD. Our framework — called SHAPER — not only allows observables to be defined that probe *any* geometric substructure of events and jets in an IRC-safe way, analogous to the event isotropy probing the uniform structure of events, but also allows those observables to be numerically estimated.

In particular, we:

1. **Motivate the Wasserstein Metric:** in section 2, we show that the Wasserstein metric is *the* natural structure for building shape-based observables for collider physics, justifying its success in refs. [25, 26, 34–42] and beyond. By adopting a measure-theoretic language for energy flows, we show that the EMD is *not* an ad-hoc structure to impose on the space of collider events, but rather the only structure that faithfully respects the detector geometry and continuity on the space of events. Further details of this argument are provided in appendices A and B.
2. **Use Optimal Transport to Define Shapes:** in section 3, we build off the work in ref. [25], where it was shown that several well-known shape observables can be described in the form:

$$\mathcal{O}_{\mathcal{M}}(\mathcal{E}) \equiv \min_{\mathcal{E}_{\theta} \in \mathcal{M}} \text{EMD}^{(\beta, R)}(\mathcal{E}, \mathcal{E}_{\theta}), \quad (1.1)$$

$$\theta_{\mathcal{M}}(\mathcal{E}) \equiv \text{argmin}_{\mathcal{E}_{\theta} \in \mathcal{M}} \text{EMD}^{(\beta, R)}(\mathcal{E}, \mathcal{E}_{\theta}), \quad (1.2)$$

where \mathcal{M} is a parameterized manifold of energy flows that define the shape, R sets a length scale for the shape, and β is a distance weighting exponent. Importantly, both the observable $\mathcal{O}_{\mathcal{M}}$ and the optimal shape parameters $\theta_{\mathcal{M}}$ can be separately extracted from the EMD. We extend this construction to define many new shape observables, by greatly expanding the class of manifolds \mathcal{M} considered, which can be constructed as explicit geometric shapes. We develop a prescription for defining new custom shape observables by parameterizing probability distributions, and even prescriptions for composing old shape observables together to form new ones.

3. **Introduce SHAPER:** in section 4, we introduce SHAPER, or **S**hape **H**unting **A**lgorithm using **P**arameterized **E**nergy **R**econstruction. This is a computational framework for defining shapes and evaluating eqs. (1.1) and (1.2) on data. SHAPER leverages the Sinkhorn approximation of the Wasserstein metric [43–46], which enables fast numerical calculation and even gradient estimation with respect to entire events, enabling easy and efficient optimization. See NEEMO [47] for an alternative gradient-based Wasserstein estimator.
4. **Evaluate Empirical Examples:** to demonstrate the potential of SHAPER, in section 5 we define and evaluate several observables on a top and QCD jet benchmark dataset [48, 49]. These shape observables can be used to extract dynamic jet radii and jet energies, and even non-radially-symmetric structures, such as jet eccentricity. In particular, we show empirical examples of our new observables for jet substructure analysis and automated pileup removal.

Generalized shape observables defined using SHAPER can be used to probe interesting collider signatures. For example, SHAPER can be used to build specialized jet algorithms with dynamic radii¹ and even dynamic pileup mitigation. This can be viewed as a generalization of k -means type clustering algorithms, such as XCONE [54]: rather than finding

¹See e.g. [50–53] for other examples of dynamic jet radii.

k points that best approximate an event, shape observables can be used to find the k geometric structures that best describe the event. This means that it is possible to design specialized jet algorithms that select for e.g. elliptical or non-isotropic jets, or that even probe the soft and collinear structure of jets separately. This can prove especially useful, for example, in boosted top or heavy vector boson decays that produce “fat jets” with multi-pronged substructure, which may not be described well by circular or isotropic patterns. We comment on further phenomenological studies in section 6.

2 The unreasonable effectiveness of Wasserstein in collider physics

In this section, we aim to answer the question, “If I had never heard of event or jet shapes before, how could I have come up with them myself?” Our discussion builds off the work of refs. [25, 26], wherein the EMD was introduced as a new language for event and jet shape observables. Here, we show that the EMD is *the* natural language for event and jet shapes. We do this by showing that the EMD is the unique metric between a geometric shape \mathcal{E}' and an event \mathcal{E} that encodes IRC safety (through its topological features) and faithfully respects the geometry of the detector. This section is largely self-contained, and readers primarily interested in the construction of new shape observables can skip to section 3.

We begin with a review of event shapes and jet shapes, noting that they all share a general form — they can all be written as a minimization of a universal loss function between the event and a parameterized set of “idealized” events, which can be interpreted as geometric shapes. We show that if the universal loss function is both IRC-safe and reduces to the ground metric on the detector geometry (that is, it *faithfully* lifts the ground metric, without distorting extended objects), then the universal loss function must indeed be the Wasserstein metric. More details of this construction can be found in appendices A and B.

2.1 Event shapes, jet shapes, and geometric shapes

We begin with a review of event shapes and jet shapes. *Event shapes* are observables that probe the geometric distribution of energy in events. Many different event shapes, such as thrust [1, 7, 8], sphericity [55], broadening [56], and N -jettiness [54, 57], have been defined and extensively studied over the years in the context of e^+e^- collisions [3], with analogues studied in the context of pp collisions [58, 59]. We may also include jet algorithms, such as XCONE [54] and sequential recombination algorithms (k_T [60, 61], Cambridge-Aachen [62, 63], and anti- k_T [64]), in this list. Similarly, *jet shapes* probe the geometric distribution of energy within individual jets rather than the global event. Examples of commonly studied jet shapes include the *integrated* jet shape² [65, 66], angularities [9, 10], and N -subjettiness [68]. While there are significant theoretical complications when considering the difference between event and jet shapes, such as the introduction of non-

²A note about nomenclature: the original “jet shape” refers to the observable $\Psi(r/R)$, the radial jet energy fraction [65, 66]. However, the word has been hijacked by ref. [67] to refer to observables consisting of weighted sums of particle momenta. We later justify the name “shapes” by showing how this corresponds to fitting actual geometric shapes to event data.

Event Shape	Description	Expression
Thrust [1, 7, 8]	How Pencil-Like?	$t(\mathcal{E}) = 2 \min_{\hat{n}} (\sum_i E_i (1 - \hat{n}_i \cdot \hat{n}))$
Sphericity [55]	How Transverse-Planar?	$s(\mathcal{E}) = \min_{\hat{n}} (\sum_i E_i \hat{n}_i \times \hat{n} ^2)$
Broadening [56]	How 2-Pronged?	$b(\mathcal{E}) = \min_{\hat{n}_1, \hat{n}_2} (\sum_i E_i \min(d_{i1}, d_{i2}))$
N -jettiness [54, 57]	How N -particle like?	$\mathcal{T}_N^{(\beta)}(\mathcal{E}) = \min_{\hat{n}_1, \dots, \hat{n}_N} (\sum_i E_i \min(R^\beta, d_{i1}^\beta, \dots, d_{iN}^\beta))$
Isotropy [40]	How Uniform?	$\mathcal{I}^{(\beta)}(\mathcal{E}) = \min_{\mathcal{U} \in \mathcal{M}} (\text{EMD}^{(\beta, R)}(\mathcal{E}, \mathcal{U}))$
XCONE [54]	Which N -particles?	$\hat{n}_i(\mathcal{E}) = \text{argmin}_{\hat{n}_1, \dots, \hat{n}_N} (\sum_i E_i \min(R^\beta, d_{i1}^\beta, \dots, d_{iN}^\beta))$
S. Recomb. [60–64]	Clustering History?	$d_{ij}^N(\mathcal{E}) = \min(E_i^{2p}, E_j^{2p}) \frac{d_{ij}^2}{R^{2p}}; d_{iR}^N(\mathcal{E}) = E_i^{2p}$

Table 1. Common event shapes and jet algorithms studied in collider physics. Note that most of these observables take the general form of eq. (2.1). Here, we do not necessarily normalize energies.

global logarithms [69, 70], for our purposes, we will treat event shapes and jet shapes interchangeably.

Following the definition in ref. [67], an event shape (jet shape) is an IRC-safe weighted sum over the four-momenta of the particles in an event (jet). These observables probe the geometric distribution of energy in an event (jet), and typically depend on the *detector ground metric*, $d(x, y)$, which defines distances between points on the detector. Expressions for common event and jet shapes can be found in tables 1 and 2, respectively. We note that all the above listed event shapes can be written in the generic form:

$$\mathcal{O}(p_1, \dots, p_M) = \min_{\theta \in \mathcal{M}} F \left(\sum_{i=1}^M E_i \phi_\theta(x_i) \right), \quad (2.1)$$

where F and ϕ_θ are generic functions, and the observable may involve a minimization (or maximization) over auxiliary parameters θ living in some constrained manifold \mathcal{M} . The choice of F , ϕ_θ , and \mathcal{M} define the event shape. Not every shape requires an optimization — for instance, while recoil-free jet angularities require an optimization over possible jet axes, it is also possible to define angularities with respect to a fixed axis [25]. In this case, the minimization may be written over the trivial manifold isomorphic to $\mathcal{M} = \{0\}$. It is also common to divide by the total energy scale E_{tot} , or some other hard scale, as this reduces the sensitivity of the event (jet) shape on experimental jet energy scale uncertainties [59]. Unless otherwise stated, we will normalize our events such that $E_{\text{tot}} = 1$ without loss of generality.

We propose to write eq. (2.1) in a universal form, such that the event shape is instead specified solely by the choice of \mathcal{M} :

$$\mathcal{O}_{\mathcal{M}}(\mathcal{E}) = \min_{\theta \in \mathcal{M}} [\mathcal{L}(\mathcal{E}; \theta)], \quad (2.2)$$

where \mathcal{L} is a universal loss function. All geometrical information about the event shape is then contained in the construction of \mathcal{M} . To emphasize this, we will adopt the notation $\mathcal{O}_{\mathcal{M}}$ for these observables, to remind us that the observable is defined through the choice of \mathcal{M} .

Jet Shape	Description	Expression
Angularities [9, 10]	Angular Moments? ... Recoil Free?	$\lambda_\beta(\mathcal{J}) = \sum_i E_i d_{iJ}^\beta$ $\lambda_\beta(\mathcal{J}) = \min_{\hat{n}} \left(\sum_i E_i d_{i\hat{n}}^\beta \right)$
N -subjettiness [68]	How N -Particle Like?	$\mathcal{T}_N^{(\beta)}(\mathcal{J}) = \min_{\hat{n}_1, \dots, \hat{n}_N} \left(\sum_i E_i \min(d_{i1}^\beta, \dots, d_{iN}^\beta) \right)$
Int. Shape [65, 66]	Radial Energy CDF?	$\psi_{\mathcal{J}}(r/R) = (\sum_i E_i \Theta(r - d_{iJ})) / (\sum_i E_i \Theta(R - d_{iJ}))$

Table 2. Common jet shapes studied in collider physics. Note that most of these observables take the general form of eq. (2.1). The notation d_{iJ} refers to the distance from particle i to the jet axis. Here, we do not necessarily normalize energies.

The task is now to determine what universal \mathcal{L} reproduces all event and jet shape observables — we will argue in section 2.4 that \mathcal{L} must be the Wasserstein metric. To begin, we may rewrite eq. (2.2) in a more suggestive form. We note that for all of the event and jet shapes in tables 1 and 2, there is always some optimal \mathcal{E}^* , not necessarily unique, such that $\mathcal{O}_{\mathcal{M}}(\mathcal{E}^*) = 0$. For example, the \mathcal{E}^* for thrust is a perfectly back-to-back event, the \mathcal{E}^* for N -subjettiness is an event with exactly N particles, and so on. Thus, it is convenient to rewrite eq. (2.2), such that the minimization is over a space of events \mathcal{E}_θ , and that $\mathcal{L} = 0$ is achieved when $\mathcal{E} = \mathcal{E}_\theta$, where θ parameterizes the space of all \mathcal{E}^* 's:

$$\mathcal{O}_{\mathcal{M}}(\mathcal{E}) = \min_{\mathcal{E}_\theta \in \mathcal{M}} [\mathcal{L}(\mathcal{E}; \mathcal{E}_\theta)]. \quad (2.3)$$

Eq. (2.3) provides a nice geometric intuition for event and jet shapes. We can interpret $\mathcal{O}_{\mathcal{M}}(\mathcal{E})$ as the answer to “How close, in event space, is my event to looking like an optimal \mathcal{E}^* ?”. Importantly, the \mathcal{E}^* 's do not have to be physically realized events — they can be *any* radiation pattern measured on the detector wall, even continuous ones. For example, we can take \mathcal{E}^* to be events with a radiation pattern that look like the interior of a hexagon — then the event shape $\mathcal{O}_{\mathcal{M}}(\mathcal{E})$ is a measure of how far \mathcal{E} is, in “event space”, from an idealized hexagonal event. By taking our idealized events \mathcal{E}^* to have radiation patterns resembling literal geometric shapes, living in the parameterized manifold \mathcal{M} , the observable $\mathcal{O}_{\mathcal{M}}(\mathcal{E})$ can be used as a measure of how much \mathcal{E} “looks like” the shape of interest.

2.2 Measure-ing the energy flow

In order to make progress in determining the universal loss function in eq. (2.3), we must first understand the IRC-safe information available for us to use within the events \mathcal{E} and \mathcal{E}_θ . This information is represented by the *energy flow* of the event. We first briefly review energy flows, before proposing a new definition of the energy flow as a measure theoretic quantity, which enables a useful language for discussing “idealized” events such as those discussed in section 2.1.

The *energy flow* \mathcal{E} of an event is the distribution of energy within the event. At a very high level, in a collider experiment, one has a detector with geometry \mathcal{X} infinitely far away from the collision site — for instance, in pp collisions such as those at the LHC, one uses a cylindrical detector $\mathcal{X} = [y_{\min}, y_{\max}] \times S^1$, where $y \in [y_{\min}, y_{\max}]$ is the rapidity and $\phi \in S^1$ is the azimuthal angle. After a collision, particles hit the detector at a site with coordinate

$x_i \in \mathcal{X}$, where the energy E_i is recorded by a calorimeter. The energy flow \mathcal{E} for an event with M particles of energies E_i measured at locations x_i is given by:

$$\mathcal{E}(x) = \sum_{i=1}^M E_i \delta(x - x_i). \tag{2.4}$$

The energy flow quantifies the total amount of energy measured at position x , which can be thought of as an idealized calorimeter cell. Assuming that the particles are massless, this is the complete *accessible*³ kinematic information about the event, which therefore allows us to consider an event and its energy flow interchangeably. In the context of hadron colliders, the transverse momentum p_T is often used in place of the energy E . In this paper, however, we focus on energies to save on notational complexity, as the story is relatively unchanged when switching to p_T .

The energy flow operator is well-understood theoretically and in some cases, can even be computed analytically [71–83]. In terms of field-theoretic quantities, the energy flow is given as:

$$\mathcal{E}(x) = \lim_{r \rightarrow \infty} \int_{-\infty}^{\infty} dt n_i T^{0i}(t, rn_i), \tag{2.5}$$

where n_i is the unit 3-vector corresponding to the detector coordinate x . We assume for this work that the spectrum of energies is non-negative — that is, for all x , $\mathcal{E}(x) \geq 0$.

We now propose a natural generalization of the energy flow that captures its salient properties and is key to enabling our geometric analysis:

Definition 1 *The energy flow \mathcal{E} of an event in a detector geometry \mathcal{X} is a (positive) **measure** over subsets $X \subseteq \mathcal{X}$, such that $\mathcal{E}(\mathcal{X}) = E_{\text{tot}}$, the total energy of the event.*

In this new language, the energy flow $\mathcal{E}(X)$ is the total energy measured in *any* region $X \subseteq \mathcal{X}$ of the detector, rather than just probing a localized point $x \in \mathcal{X}$. The region X can be an extended set and does not need to be connected. Figure 1 illustrates an example of this on a cylindrical collider. This generalized notion of energy flow $\mathcal{E}(X)$ reduces to the usual energy flow $\mathcal{E}(x)$, which we now refer to as the *energy flow density*, and can be written as:

$$\mathcal{E}(X) = \int_X dx \mathcal{E}(x). \tag{2.6}$$

A particle measured at x_i will contribute energy to $\mathcal{E}(X)$ only if $x_i \in X$, which can be seen by carrying out the integration over the δ -functions in eq. (2.4).⁴ Unlike eq. (2.4), however, we do not restrict energy flows to just a finite sum of localized δ -functions —

³Here, we take *accessible* to mean the calorimeter information after infinite time has passed, preserving no timing information. We implicitly assume that the calorimeter’s detector response is linear, so that two photons entering the same calorimeter cell cannot be distinguished from a single photon with their summed energy, though even if the response is nonlinear, one cannot distinguish how many photons entered the calorimeter cell from the total energy alone.

⁴We will assume that energy flows can always be written as the integral of an associated energy flow density. Note that the energy flow density depends on the choice of coordinates x used on \mathcal{X} .




Figure 1. An illustration of the energy flow $\mathcal{E}(X)$, which is the total amount of radiation captured inside a subset X (in purple) of the total detector geometry \mathcal{X} . The red dots represent particles that hit the detector wall, with their size proportional to their energy.

they can be continuous, extended deposits of energy! In section 3, we will see energy flows with continuous energy distributions are key to defining generalized shape observables. We will refer to energy flows whose densities can be represented by a finite sum of weighted δ -functions (as in, for example, eq. (2.4)) as *atomic measures* or *atomic flows*. We will often write atomic measures as $\mathcal{E} \sim \sum_i E_i \delta_{x_i}$ for notational simplicity.

Under Definition 1, energy flows inherit a very rich and natural mathematical structure. The most important operation for our purposes is the integral of a function $\phi : \mathcal{X} \rightarrow \mathbb{R}$ against an energy flow \mathcal{E} , which we denote $\langle \mathcal{E}, \phi \rangle$, defined as:

$$\langle \mathcal{E}, \phi \rangle \equiv \int_{\mathcal{X}} dx \mathcal{E}(x) \phi(x) \tag{2.7}$$

$$= \sum_i E_i \phi(x_i) \text{ for atomic flows.} \tag{2.8}$$

This operation can be thought of as the energy-weighted expectation value of the random variable ϕ under the distribution \mathcal{E} . A brief review of this, and other salient measure-theoretic concepts and definitions we call upon in this paper, is presented in appendix A.

2.3 Geometrizing IRC safety

Infrared and colinear (IRC) safety is an incredibly powerful constraint on the form of observables — it ensures not only that an observable is well-defined in perturbation theory, but also that the observable is robust to detector effects. Using the language developed in section 2.2, IRC safety becomes a topological statement on the space of energy flows, which we may use to place constraints on the potential form of the universal loss function \mathcal{L} of eq. (2.3).

An observable \mathcal{O} is IRC safe if it satisfies:⁵

- **Infrared safety:** for any event atomic \mathcal{E} , adding or removing an ϵ -soft emission to \mathcal{E} leaves \mathcal{O} unchanged as $\epsilon \rightarrow 0$.
- **Collinear safety:** for any atomic event \mathcal{E} , splitting any particle into two particles at the same location with the same total energy leaves \mathcal{O} unchanged. Moreover, translating either particle by an ϵ -small displacement leaves \mathcal{O} unchanged as $\epsilon \rightarrow 0$.

Essentially, IRC safety means that observables should not change significantly if we change \mathcal{E} by slightly adjusting particle energies and positions. As with energy flows, we propose a generalization of IRC safety that captures all its salient features:

Definition 2 *An observable \mathcal{O} is IRC safe if it is continuous with respect to the weak* topology on energy flows.*

A function f on energy flows is continuous to the weak* topology if, for any sequence of energy flows \mathcal{E}_n that converges to \mathcal{E} , the function $F(\mathcal{E}_n)$ converges to $F(\mathcal{E})$ (see appendix A for more details). Note that this is actually a slightly *weaker* constraint than the one considered in ref. [25], which defines IRC safety through the metric topology induced by the EMD — the definition here does not require a metric on the space of events, or even a metric on the detector space, only a notion of continuity. In fact, there is a large class of metrics one can place on the space of events to metrize the weak* topology, not just the EMD.

An interesting consequence of this definition is that if an observable \mathcal{O} is IRC safe, then $\mathcal{O}(\mathcal{E})$ for *any* energy flow \mathcal{E} can be arbitrarily well-approximated by atomic energy flows. This implies, for example, that a continuous circle can be arbitrarily well approximated by a finite number of points arranged in a ring — this makes possible to not only encode continuous distributions numerically, but also to make broad statements about the behavior of IRC safe observables by considering their action only on simple atomic energy flows.

In order to be IRC safe, our universal loss function $\mathcal{L}(\mathcal{E}, \mathcal{E}')$ must be continuous in both of its arguments. This is very restrictive, and immediately implies that \mathcal{L} cannot have any terms that are discontinuous in either energy or distance, e.g. terms like E^{-1} or $d(x, y)^{-1}$, or any term of the form $\langle \mathcal{E}, \phi \rangle$ for noncontinuous ϕ . Recalling the discussion in section 2.1 that \mathcal{L} quantifies how close in the space of energy flows \mathcal{E} and \mathcal{E}' are, it is convenient (though not strictly necessary) to use \mathcal{L} to *metrize* the weak* topology — that is, if an observable \mathcal{O} is continuous with respect to the metric topology induced by \mathcal{L} , then it is also continuous with respect to the weak* topology, and therefore is IRC safe. This allows the same universal loss function \mathcal{L} to be used both to define shape observables and to define IRC safety.⁶ This is convenient, since it captures the very intuitive notion that if two events geometrically look similar (that is, \mathcal{L} is small), then IRC-safe observables evaluated on them should also be the same.

⁵There are several different statements of IRC-safety with different limit structures, each with different pathologies. A brief discussion of this can be found in section 2.1 of ref. [25].

⁶This is not required however — there are many different choices of \mathcal{L} that may be used to give the same definition of IRC safety, e.g. maximum mean discrepancies, which does not necessarily have to be the same function \mathcal{L} whose minimum defines shapes as in eq. (2.3).

2.4 The importance of being faithful

In order to encode geometric information about energy distributions, the universal loss function \mathcal{L} of eq. (2.3) must explicitly depend on the detector ground metric, $d(x, y)$. While there are many metrics on the space of measures that encode geometric information while also being IRC-safe (as defined in section 2.3), a natural choice is the family of Wasserstein metrics, which we denote $\mathcal{L}(\mathcal{E}, \mathcal{E}') = \text{EMD}^{(\beta, R)}(\mathcal{E}, \mathcal{E}')$ (for *Earth-Mover's* or *Energy-Movers Distance*, which we will use synonymously). We show in this section that unlike other potential candidates for \mathcal{L} , the Wasserstein metric will never warp distances between shapes — that is, the Wasserstein metric lifts the ground metric of the detector *faithfully*. A constructive proof of this can be found in appendix B.

The EMD between two measures \mathcal{E} and \mathcal{E}' is given by:

$$\begin{aligned} \text{EMD}^{(\beta, R)}(\mathcal{E}, \mathcal{E}') &= \min_{\pi \in \mathcal{M}(\mathcal{X} \times \mathcal{X})} \left[\frac{1}{\beta R^\beta} \langle \pi, d(x, y)^\beta \rangle \right] + |\Delta E_{\text{tot}}|, \\ \pi(\mathcal{X}, Y) &\leq \mathcal{E}'(Y), \quad \pi(X, \mathcal{X}) \leq \mathcal{E}(X), \quad \pi(\mathcal{X}, \mathcal{X}) = \min(E_{\text{tot}}, E'_{\text{tot}}), \end{aligned} \tag{2.9}$$

where $d(x, y)$ is the ground metric between points x and y on \mathcal{X} , and $\mathcal{M}(\mathcal{X} \times \mathcal{X})$ is the space of all positive measures on $\mathcal{X} \times \mathcal{X}$. The parameter $R > 0$ sets a distance scale for the EMD, and sets the relative scale for the two terms in eq. (2.9). The parameter $\beta \geq 1$ sets the distance norm.⁷ Note that our definition of the EMD differs from refs. [25, 26] by a factor of β , which we do to match the conventions of the GEOMLOSS [46] package. The additional energy difference term, $|\Delta E_{\text{tot}}|$, contributes whenever the two energy flows do not have the same total energy.

The Wasserstein metric is special in that it *faithfully* lifts the ground metric, d . To lift the ground metric means that the Wasserstein metric reduces to $d(x, y)^\beta$ when evaluated on two point measures $\mathcal{E} \sim \delta_x$ and $\mathcal{E}' \sim \delta_y$ — that is, the Wasserstein metric preserves distances between points. Moreover, to do so *faithfully* means that the Wasserstein metric preserves distances for entire extended shapes: if \mathcal{E} is *any* measure, and \mathcal{E}' is the same as \mathcal{E} whose density is translated by a vector t (that is, \mathcal{E}' has corresponding density $\mathcal{E}(x - t)$), then the metric between them is simply $d(0, t)^\beta$.

To see this explicitly, we can compare to two other potential candidates for our universal loss function \mathcal{L} : the class of Maximum Mean Discrepancies (MMDs) [84] and Chamfer

⁷In order to satisfy the triangle inequality and be a true metric, the first term in the EMD must be raised to the $1/\beta$ power, and either $2R$ should exceed the largest value of $d(x, y)$ [25, 26] or ΔE_{tot} must be guaranteed to be zero. In this paper, we will not need the triangle inequality, so we will not do this.

distances [85], which can respectively be written as:⁸

$$\text{MMD}^{(\beta)}(\mathcal{E}, \mathcal{E}') = -\frac{1}{2\beta} \int_{\mathcal{X} \times \mathcal{X}} dx dy d(x, y)^\beta (\mathcal{E}(x) - \mathcal{E}'(x)) (\mathcal{E}(y) - \mathcal{E}'(y)), \quad (2.10)$$

$$\begin{aligned} \text{CD}^{(\beta)}(\mathcal{E}, \mathcal{E}') &= \frac{1}{2\beta} \int_{\mathcal{X}} dx \min_{y \in \text{Supp}(\mathcal{E}')} [d(x, y)^\beta] \mathcal{E}(x) \\ &+ \frac{1}{2\beta} \int_{\mathcal{X}} dy \min_{x \in \text{Supp}(\mathcal{E})} [d(x, y)^\beta] \mathcal{E}'(y). \end{aligned} \quad (2.11)$$

These candidate loss functions are IRC safe, translationally invariant, and even lift the ground metric, but importantly, they do not do so faithfully. To see this, we choose the following example energy flows:

$$\mathcal{E}(x) \sim \frac{1}{2} \delta_0 + \frac{1}{2} \delta_a, \quad \mathcal{E}'(x) \sim \frac{1}{2} \delta_t + \frac{1}{2} \delta_{a+t}, \quad (2.12)$$

where a and t are arbitrary vectors (where \mathcal{E} consists of 2 points separated by a vector a , and \mathcal{E}' is \mathcal{E} translated by a vector t). A direct computation using eqs. (2.9), (2.10), and (2.11) (with $d(x, y) = |x - y|$) yields:

$$\text{EMD}^{(\beta, R=1)}(\mathcal{E}, \mathcal{E}') = \frac{1}{\beta} |t|^\beta, \quad (2.13)$$

$$\text{MMD}^{(\beta)}(\mathcal{E}, \mathcal{E}') = \frac{1}{2\beta} \left(|t|^\beta + \frac{1}{2} |t - a|^\beta + \frac{1}{2} |t + a|^\beta - |a|^\beta \right), \quad (2.14)$$

$$\text{CD}^{(\beta)}(\mathcal{E}, \mathcal{E}') = \frac{1}{2\beta} \left(\min \left[|t|^\beta, |t + a|^\beta \right] + \min \left[|t|^\beta, |t - a|^\beta \right] \right). \quad (2.15)$$

While eqs. (2.14) and (2.15) do indeed reduce to $\sim |t|^\beta$ when $a \rightarrow 0$ (that is, when \mathcal{E} reduces to a single point), in general the MMD and Chamfer distance effectively “distort” the shape.⁹ When $\beta = 1$, for instance, \mathcal{E}' appears slightly closer to \mathcal{E} , as measured using either MMD or the Chamfer distance, than its total displacement $|t|$ — energy in the interior of an extended distribution gets effectively “screened” by the energy in the rest of the distribution! Not only does this distortion ruin our ability to think of our observables as measuring the geometric distribution of energy in the detector, the screening effect also induces a practical issue, as it causes vanishing gradients when trying to optimize over \mathcal{L} [46] (in this case, by minimizing t , for example). The Wasserstein metric does not suffer these problems, making it the natural choice for our universal loss function.

3 Hearing shapes

Having constructed the Wasserstein metric and EMD in section 2 for event and jet shapes, we next generalize eqs. (1.1) and (1.2), which were originally introduced in ref. [25] as a

⁸We choose the MMD and Chamfer distance for comparison because, as shown in appendix B, the most general loss that is symmetric, IRC safe, and lifts the ground metric (though not necessarily faithfully) has terms individually resembling the Wasserstein, MMD, and Chamfer distances.

⁹Note that this is avoided for the MMD in the $\beta = 2$ case. This is in contrast to the Wasserstein metric, which is faithful for *any* $\beta \geq 1$. While some observables are defined with $\beta = 2$, this is not sufficient for all observables, and in particular we will focus on observables with $\beta = 1$, which is the only “true” metric satisfying the triangle inequality.

common form for many well-known observables. We treat eqs. (1.1) and (1.2) as definitions for *shape observables* $\mathcal{O}_{\mathcal{M}}$ and *shape parameters* $\theta_{\mathcal{M}}$, which together are the natural generalization of event and jet shapes. Moreover, we show how the manifold of energy flows \mathcal{M} can be chosen to construct new observables that probe specific geometric structures. We provide a prescription for building \mathcal{M} , which defines the shape observable, as well as prescriptions for *composing* shape observables together, allowing new shape observables to be defined from simpler ones in a geometrically intuitive way.¹⁰

This section proceeds as follows. First, we define generalized shape observables and shape parameters, and discuss their properties. Next, we discuss our prescription for shape composition, which can be used to define shape observables and parameters that probe complex geometric structure. Finally, we use our prescription to construct a large (but importantly, inexhaustive) suite of novel shape observables for jet substructure analysis to serve as an example of what can be done with this framework. Several examples of empirical studies using these new shape observables, evaluated using the SHAPER framework defined in section 4, can be found in section 5.

3.1 Shape observables and shape parameters

We define a *shape observable* $\mathcal{O}_{\mathcal{M}}$ as follows:

Definition 3 *A shape observable $\mathcal{O}_{\mathcal{M}}$, with associated shape parameters $\theta_{\mathcal{M}}$, on an energy flow \mathcal{E} is any function of the form:*

$$\mathcal{O}(\mathcal{E}) \equiv \min_{\mathcal{E}_{\theta} \in \mathcal{M}} \text{EMD}^{(\beta, R)}(\mathcal{E}, \mathcal{E}_{\theta}), \tag{3.1}$$

$$\theta(\mathcal{E}) \equiv \text{argmin}_{\mathcal{E}_{\theta} \in \mathcal{M}} \text{EMD}^{(\beta, R)}(\mathcal{E}, \mathcal{E}_{\theta}), \tag{3.2}$$

where \mathcal{M} is a manifold of positive measures on the detector space \mathcal{X} , and $\text{EMD}^{(\beta, R)}$ is the β -Wasserstein distance with length scale R .

Importantly, we return *both* the minimum EMD value, $\mathcal{O}(\mathcal{E})$, and the parameters of the shape that produced the minimum EMD, $\theta(\mathcal{E})$. For a manifold \mathcal{M} of generic parameterized shapes, we refer the former as the “*shapiness*” of \mathcal{E} , and the latter as the associated “*shape parameters*” of \mathcal{E} . For example, if \mathcal{M} is the manifold of N -(sub)jet events, then $\mathcal{O}(\mathcal{E})$ is the N -(sub)jettiness of \mathcal{E} , and $\theta(\mathcal{E})$ are the (sub)jet parameters of \mathcal{E} , highlighting that eq. (3.1) is really a generalization of the N -jettiness. Intuitively, $\mathcal{O}(\mathcal{E})$ answers the question, “How much like a *shape* does my event look like?”, while $\theta(\mathcal{E})$ answers the question, “Which *shape* does my event look most like?”.

For a fixed choice of β and scale R defining the EMD, shape observables are completely specified by the choice of the manifold of energy flows \mathcal{M} . This choice specifies the class of shapes being considered. Practically speaking, this manifold can be defined by choosing a set of coordinates θ on the manifold, that parameterize a set of constrained energy flows

¹⁰Like music, shapes are composed using the SHAPER paradigm by splicing and overlaying together smaller shapes, after which it may be “heard” by evaluating on event data — a posthoc rationalization for the title of this paper.

\mathcal{E}_θ . Since, as established in section 2.2, energy flows are positive measures on the detector space, \mathcal{M} can be built as a (weighted) parameterized probability distribution p_θ , which is realized by a finite sampling procedure for weighted points on \mathcal{X} .¹¹

In ref. [25], the manifold corresponding to several event and jet shapes were listed. However, the power of our framework is that *any* manifold of parameterized energy flows defines a valid shape observable, and moreover, this observable directly probes intuitive geometric information. By defining, for example, \mathcal{M} to be the manifold of energy flows resembling uniform rings of energy, uniform disks of energy, or even uniform ellipses of energy, the corresponding observables directly quantify the diskiness, circliness, and ellipsiness of events. To our knowledge, the event isotropy [40] is the first observable of this form with no known alternative formulation, as it quantifies the “uniformness” of events. Examples of analyses using these custom observables can be found in section 5.

Below, we list some useful properties of all shape observables:

1. **Monotonicity:** if \mathcal{O}_1 and \mathcal{O}_2 are defined by manifolds \mathcal{M}_1 and \mathcal{M}_2 , for a fixed choice of β and R , then $\mathcal{M}_1 \supseteq \mathcal{M}_2$ implies that for all energy flows \mathcal{E} , we have $\mathcal{O}_1(\mathcal{E}) \leq \mathcal{O}_2(\mathcal{E})$. This captures the monotonic nature of observables such as N -jettiness, which satisfy $\tau_{N+1} \leq \tau_N$.
2. **Closure:** $\mathcal{O}(\mathcal{E}) = 0$ if and only if $\mathcal{E} \in \mathcal{M}$. It then follows that the shape parameters $\theta(\mathcal{E})$ are such that $\mathcal{E}_\theta = \mathcal{E}$. In particular, this implies that for $\mathcal{M}_\mathcal{E}$, the space of *all* energy flows, $\mathcal{O}(\mathcal{E}) = 0$ for all events.
3. **Approximation Bounds:** for any L -Lipschitz function ϕ on \mathcal{X} , the “optimal shape” \mathcal{E}_θ corresponding to \mathcal{E} satisfies $\frac{1}{RL} |\langle \mathcal{E}, \phi \rangle - \langle \mathcal{E}_\theta, \phi \rangle| \leq \mathcal{O}(\mathcal{E})$ [26]. That is, the optimal shape \mathcal{E}_θ can be used to approximate additive Lipschitz observables on \mathcal{E} , up to a known bounded error.
4. **Upper Bounds:** if \mathcal{X} is bounded by a maximum distance scale R_{\max} , and both \mathcal{E} and all energy flows on \mathcal{M} satisfy $E_{\text{tot}} = 1$, then $\mathcal{O}(\mathcal{E})$ is bounded above by $\left(\frac{R_{\max}}{R}\right)^\beta$.¹² This can be seen by considering the extreme case where \mathcal{E} and \mathcal{E}_θ are singleton points located R_{\max} away from each other. Any configuration other than this will yield a lower EMD. This makes it always possible to normalize $\mathcal{O}(\mathcal{E}) \in [0, 1]$.

We call a manifold of energy flows \mathcal{M} *balanced* if all of the energy flows are all normalized (i.e. $E_{\text{tot}} = 1$), and unbalanced otherwise. Similarly, we call shape observable $\mathcal{O}_\mathcal{M}$ balanced if \mathcal{M} is balanced *and* it is only evaluated on normalized energy flows. For a balanced observable, the choice of scale R constitutes only a change of units for the distance metric, and is unimportant.¹³ Moreover, in the limit $R \rightarrow \infty$, the quantity $R^\beta \mathcal{O}(\mathcal{E})$ is only

¹¹Because energy flow densities depend on the choice of coordinates, the choice of p_θ corresponding to the class of shapes prescribed by \mathcal{M} is not unique. For example, to sample a Gaussian distribution, one can sample uniformly weighted points with positions distributed as a Gaussian, or sample uniformly spaced points with Gaussian weights depending on their position, or some mixture of both.

¹²This is not necessarily the least upper bound, which depends specifically on the choice of \mathcal{M} .

¹³Though we will see in section 4 that R matters again for our numerical approximations.

finite if $E_{\text{tot}} = E_{\theta_{\text{tot}}}$ [25]. In this limit, the (arg)min can be written over the submanifold $\mathcal{M}_E \subseteq \mathcal{M}$, which is the submanifold events with the same total energy as \mathcal{E} . Thus, the $R \rightarrow \infty$ limit effectively forces unbalanced shape observables to be balanced. We will primarily consider balanced shape observables for the remainder of this paper, though there are many important unbalanced observables (e.g. N -jettiness) that one may consider.

3.2 Composing shapes

Given two shape observables \mathcal{O}_1 and \mathcal{O}_2 , with the same choice of β and R , we can define a new *composite observable* $\mathcal{O} = \mathcal{O}_1 \oplus \mathcal{O}_2$. As with all shape observables above, we define the composite shape observable \mathcal{O} by specifying the corresponding manifold \mathcal{M} . We consider two possible scenarios to define \mathcal{M} :

- **Case 1 (Balanced):** if both \mathcal{O}_1 and \mathcal{O}_2 are balanced shape observables, then we define \mathcal{M} to be the manifold of all energy flows \mathcal{E} of the form $\mathcal{E} = z_1 \mathcal{E}_1 + z_2 \mathcal{E}_2$, where $\mathcal{E}_1 \in \mathcal{M}_1$ and $\mathcal{E}_2 \in \mathcal{M}_2$, and (z_1, z_2) is a point in Δ_1 , the 1-simplex.¹⁴
- **Case 2 (Totally Unbalanced):** if the energy flows on \mathcal{M}_1 and \mathcal{M}_2 have unconstrained E_{tot} 's (that is, E_{tot} is an independent free parameter for each manifold), then we define \mathcal{M} to be the manifold of all energy flows \mathcal{E} of the form $\mathcal{E} = \mathcal{E}_1 + \mathcal{E}_2$, where $\mathcal{E}_1 \in \mathcal{M}_1$ and $\mathcal{E}_2 \in \mathcal{M}_2$.

There exist other possible cases, such as if the manifolds only admit a small range of E_{tot} values, but we will not consider them in this work.¹⁵ In Case 1, the 1-simplex Δ_1 introduces 2 additional (constrained) parameters to the shape, z_1 and z_2 , that ensure that all energy flows are still normalized to 1, so that \mathcal{O} is balanced. Here, \mathcal{M} can be realized by generating points according to the parameterized sampling procedures for \mathcal{M}_1 and \mathcal{M}_2 , and multiplying the energy weights of the generated particles by z_1 and z_2 , respectively. The values of z_1 and z_2 control the relative contribution of each base shape to the composition. In Case 2, the total energy remains unconstrained, and the relative importance of each base shape to the contribution is controlled by the ratio of the $E_{1\text{tot}}$ and $E_{2\text{tot}}$ parameters. Note that monotonicity implies that, for both Case 1 and Case 2, $\mathcal{O}(\mathcal{E}) \leq \mathcal{O}_1(\mathcal{E})$ and $\mathcal{O}(\mathcal{E}) \leq \mathcal{O}_2(\mathcal{E})$.

It is easy to extend these definitions to a composition of N observables, $\mathcal{O} = \bigoplus_{i=1}^N \mathcal{O}_i$. For the balanced case, \mathcal{M} comprises of energy flows of the form $\mathcal{E} = \sum_{i=1}^N x_i \mathcal{E}_i$ for $\mathcal{E}_i \in \mathcal{M}_i$ and $z_i \in \Delta_{N-1}$. In the totally unbalanced case, \mathcal{M} comprises energy flows of the form $\mathcal{E} = \sum_{i=1}^N \mathcal{E}_i$ for $\mathcal{E}_i \in \mathcal{M}_i$. When the \mathcal{O}_i are all copies of the same shape observable \mathcal{O} , we define the notation $N \times \mathcal{O}$, and name the resulting composite shape observable the “ N -shapiness”.

Given this language, we can understand several existing shape observables as composites of basic “shape” building blocks. For instance, we can define an observable, the

¹⁴The N -simplex, Δ_N , is the set of all points $z_1, \dots, z_{N+1} \geq 0$ such that $\sum_{i=1}^{N+1} z_i = 1$.

¹⁵There is no natural way to define \mathcal{M} from \mathcal{M}_1 and \mathcal{M}_2 ; this is a choice that we make, and other choices could be considered. For example, one may consider defining \mathcal{M} to instead be the manifold of energy flows $\mathcal{E} = \mathcal{E}_1 \mathcal{E}_2$, which is a measure on $\mathcal{X} \times \mathcal{X}$, and could potentially be useful for measuring factorization properties. We leave the study of this class of shapes for potential future work.

(sub)pointiness $\tau_1(\mathcal{E})$, corresponding to the manifold of weighted (normalized) δ -function measures on \mathcal{X} . Geometrically, this observable to fitting a single weighted (normalized) point to an event \mathcal{E} , the former with a floating energy weight. We can then define the N -(sub)pointiness of \mathcal{E} to be $\tau_N(\mathcal{E}) = N \times \tau_1(\mathcal{E})$ — this is exactly the N -(sub)jettiness of \mathcal{E} , with N -subjettiness corresponding to balanced composition, and N -jettiness corresponding to totally unbalanced composition.

Shape composition provide a novel avenue for understanding event and jet substructure. For example, the observable $\tau_N \oplus \mathcal{I}$, where \mathcal{I} is the event isotropy, can be thought of as a “pileup-corrected” N -(sub)jettiness, where a uniform background is subtracted off. One can use this observable to probe the percentage of energy deposited in hard jets versus a uniform background due to pileup [86] and underlying event effects [87, 88]. In particular, the shape parameter $\theta(\mathcal{E}) = z_1(\mathcal{E})$ is an estimate of the percentage of energy in E due to hard jets. In section 5.7, we consider perform empirical studies of these types of shape observables.

3.3 Examples of novel shapes

In this subsection, we list (some) potentially phenomenologically interesting shape observables, all of which are defined for the first time in this work, that can be constructed using the prescription outlined above. For all these observables, we consider the detector geometry to be a rectangular patch of a cylinder, $\mathcal{X} = [y_{\min}, y_{\max}] \times [\phi_{\min}, \phi_{\max}]$, though one could consider the full cylinder as well. These observables are summarized in table 3, and we show empirical examples of these observables in action in section 5. We consider balanced observables with $\beta = 1$, leaving unbalanced observables for future work.

3.3.1 N -ringiness

We first consider a simple shape observable: *ringiness*, which probes how ring-like an event is. We begin by defining the manifold of all ring-like energy flows \mathcal{M}_R , which consist of energy flows corresponding to energy flow densities of the form:

$$\mathcal{E}_{x_0, R_0}(x) = \begin{cases} \frac{1}{2\pi R_0} & |x - x_0| = R \\ 0 & \text{otherwise} \end{cases}, \quad (3.3)$$

where the parameters x_0 and R_0 correspond to the center and radius of a ring, respectively. To build a sampler, we use the so-called *reparameterization trick* [89]. Drawing N samples from the unit uniform distribution $\phi \sim U(0, 1)$, the distribution of points:

$$x = R_0 (\cos(2\pi\phi), \sin(2\pi\phi)) + x_0, \quad (3.4)$$

where each particle has weight $\frac{1}{N}$, is a realization of $\mathcal{E}_{x_0, R_0}(x)$.

The corresponding observable, $\mathcal{O}_R(\mathcal{E})$, is the ringiness of the event \mathcal{E} . While most QCD jets are not expected to be ring-like, this observable can identify clumps of radiation scattered around a central point, as may be the case in a 3-pronged top quark decay. Additionally, observables that probe the boundary of a jet with an empty interior may prove useful in studies of the dead-cone effect [90–92] where collinear radiation is relatively suppressed.

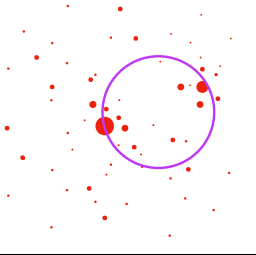
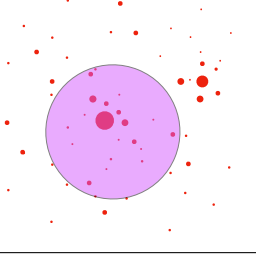
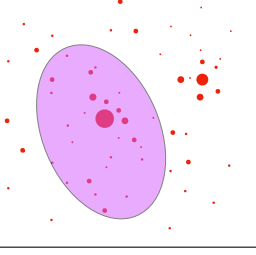
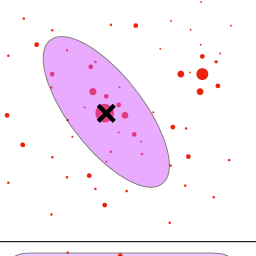
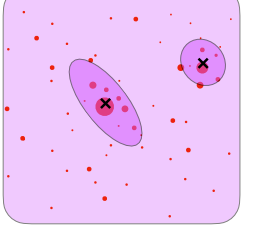
Section	Shape	Specification	Illustration
3.3.1	Ringiness \mathcal{O}_R	Manifold of Rings $\mathcal{E}_{x_0, R_0}(x) = \frac{1}{2\pi R_0}$ for $ x - x_0 = R_0$ $x_0 = \text{Center}, R_0 = \text{Radius}$	
3.3.2	Diskiness \mathcal{O}_D	Manifold of Disks $\mathcal{E}_{x_0, R_0}(x) = \frac{1}{\pi R_0^2}$ for $ x - x_0 \leq R_0$ $x_0 = \text{Center}, R_0 = \text{Radius}$	
3.3.3	Ellipsiness \mathcal{O}_E	Manifold of Ellipses $\mathcal{E}_{x_0, a, b, \varphi}(x) = \frac{1}{\pi ab}$ for $x \in \text{Ellipse}_{x_0, a, b, \varphi}$ $x_0 = \text{Center}, a, b = \text{Semi-axes}, \varphi = \text{Tilt}$	
3.3.4	(Ellipse +Point)iness	Composite Shape $\mathcal{O}_E \oplus \tau_1$ Fixed to same center x_0	
3.3.5	N-(Ellipse +Point)iness +Pileup	Composite Shape $N \times (\mathcal{O}_E \oplus \tau_1) \oplus \mathcal{I}$	

Table 3. Custom observables, defined using the SHAPER prescription, designed to probe jet substructure at increasing levels of complexity. For each observable, the manifold parameterization is given, either explicitly, or as a composition of previously defined objects. Here, τ_1 is the one-pointiness (1-subjettiness), and \mathcal{I} is the event isotropy. More details on these types of observables, plus explicit construction of sampling functions, can be found in section 3.3.

Having defined ringiness, we can next define N -ringiness, which probes how much an event looks like N rings, each of arbitrary center, radius, and weight. This observable is defined as $N \times \mathcal{O}_R$. Following the prescription outlined in section 3.2, we can build N weighted rings, by separately sampling eq. (3.4) for each ring's center and radius, and multiplying the weights by $z_i \in \Delta_{N-1}$.

For numerical methods, one must set initial values for these parameters in an IRC-safe way. While in principle, the choice of initialization should make no difference, in practice, the presence of numerical effects and local minima make the choice of initialization important. In our initialization scheme for N -rings, we perform k_T clustering to find N subjects. The location of the subjects is then taken to be the ring center, and the subject energy is taken to be the ring energy. We choose to initialize the radius of each ring to zero, so that the N -ringiness is guaranteed to only deviate from N -subjettiness only if it will make the event more ringlike, though it is also possible to initialize the radius to e.g. the distance of the $(N + 1)$ -th jet in the clustering history.

3.3.2 N -diskiness

Next, we define *diskiness* \mathcal{O}_D , which measures how much like a disk an event is. Similar to ringiness, we parameterize the manifold of energy flow densities:

$$\mathcal{E}_{x_0, R_0}(x) = \begin{cases} \frac{1}{\pi R_0^2} & |x - x_0| \leq R \\ 0 & \text{otherwise} \end{cases}, \quad (3.5)$$

where x_0 and R_0 are the center and radius of the disk.

To build a sampler, as with ringiness, we draw N samples from the unit uniform distribution $\phi \sim U(0, 1)$, and also N points $r \sim U(0, 1)$. Then, the distribution of points:

$$x = \sqrt{r} R_0 (\cos(2\pi\phi), \sin(2\pi\phi)) + x_0, \quad (3.6)$$

where each particle has weight $\frac{1}{N}$, is a realization of a uniform disk. The random variable r controls the radius of the point being sampled, and the square root is from a Jacobian factor to make the disk uniform.

Given the diskiness, we can easily compose the N -diskiness, $N \times \mathcal{O}_D$. The N -diskiness is analogous to the N -jettiness/XCONE jet algorithm, in that it returns the locations of conical clusters of particles. However, unlike XCONE, the radius R_0 is a learned, rather than fixed parameter, allowing for dynamic jet radii.¹⁶ To initialize the N -diskiness, the exact same procedure is used as described in section 3.3.1 for N -ringiness. Note that there are many ways to modify the N -diskiness to produce similar observables — for example, one can replace the uniform disks with Gaussians to probe different radiation patterns.

3.3.3 N -ellipsiness

Jets need not necessarily be circular! Indeed, many jet algorithms, such as the widely-used k_T [64] and Cambridge-Aachen [62, 63] algorithms, do not return circular jets. Motivated

¹⁶In place of fixing the radius R , as in the XCONE algorithm, one instead assumes that the jet energies are uniform across the disk, so the number of assumptions is conserved.

by this, we define a generalization of diskiness, the *ellipsiness* \mathcal{O}_E of a jet. The manifold of ellipses is given by energy flow densities of the form:

$$\mathcal{E}_{x_0,a,b,\varphi}(\vec{x}) = \begin{cases} \frac{1}{\pi ab} & \left(\frac{(x-x_0)\cos\varphi+(y-y_0)\sin\varphi}{a} \right)^2 + \left(\frac{(y-y_0)\cos\varphi-(x-x_0)\sin\varphi}{b} \right)^2 \leq 1, \\ 0 & \text{otherwise} \end{cases}, \quad (3.7)$$

x_0 is the center of the ellipse, a and b are the semi-major and semi-minor axes,¹⁷ and φ is the tilt of the x -axis. Here, we have restored vector notation \vec{x} to indicate that the x - and y -axes are treated differently. There are many equivalent alternate parameterizations of the ellipse, including in terms of its focal length $c = \sqrt{\max(a,b)^2 - \min(a,b)^2}$ and eccentricity $e = \sqrt{1 - \frac{\min(a,b)}{\max(a,b)}}$. Note that for $a = b$, the ellipse reduces to a disk, and the φ parameter becomes redundant.

The sampling procedure for disks can be recycled for ellipses, with some small modifications. Given N sampled points $\phi, r \sim U(0, 1)$, the distribution:

$$x = U_\varphi \cdot (a\sqrt{r}\cos(2\pi\phi), b\sqrt{r}\sin(2\pi\phi))^T + x_0, \quad (3.8)$$

where U_φ is the 2×2 rotation matrix corresponding to the angle φ and each particle has weight $\frac{1}{N}$, is a realization of a uniform ellipse. We can then easily compose the N -ellipsiness, which can serve as a jet algorithm that finds non-circular jets. In particular, this shape observable allows for the eccentricity e of the clustered jets to be extracted, allowing one to quantify how far from circular each jet is. As with the N -ringiness and N -diskiness, the centers of the ellipses are chosen using the k_T clustering algorithm. Both a and b are initialized to be zero, so that deviations from either N -subjettiness or N -diskiness occur if it makes the event more elliptical.

3.3.4 ... plus pointiness

Energy is not uniformly distributed within a jet! Indeed, to leading order in perturbative QCD, much of a jet's radiation will be soft and/or collinear with respect to the emitting parton. We can probe this by composing together shapes that explicitly target soft and collinear radiation separately. To this end, we construct a set of new observables, the (*shape+point*)iness, for $shape \in \{\mathcal{O}_R, \mathcal{O}_D, \mathcal{O}_E\}$. This is defined using the shape composition prescription described in section 3.2, as:

$$\mathcal{O}_i^\tau = \mathcal{O}_i \oplus \tau_1, \quad (3.9)$$

where τ_1 is the 1-pointiness (equivalently, the 1-subjettiness), and \mathcal{O}_i is any of $\{\mathcal{O}_R, \mathcal{O}_D, \mathcal{O}_E\}$ previously defined. Importantly, we fix the location of the δ -function in τ_1 to be x_0 , though one may consider letting the location of the δ -function float to define a recoil-free variant.¹⁸ We can then extend this definition to compose the N -(*shape+point*)iness.

¹⁷Non-respectively; a corresponds to the x -axis and b to the y -axis, and we make no distinction here which is the major versus minor axis.

¹⁸In the elliptical case, one may consider attaching δ -functions to one or both of the foci instead of the center. We leave the study of variants of these observables to future work.

When used as a jet algorithm, the N -(shape+point)iness provides a more physical picture of perturbative QCD than do the previously defined shapes. The base shapes, particularly disks and ellipses, capture wide-angle soft radiation, while the δ -functions capture both hard and soft collinear radiation at the center of the (sub)jet. Moreover, within each shape-point pair, the floating parameters z_1 and z_2 tell us the fraction of radiation in the wide-angle and collinear sectors, which in principle can be calculated in and compared to perturbative QCD.

When initializing observables of this type, the initialization occurs as described in previous sections using the k_T algorithm with all radii set to zero. However, we choose to split the k_T energy equally between the shape and the δ -function. Note that this is an IRC-safe choice, since at zero radius, the shape is indistinguishable from the δ -function.

3.3.5 ... plus pileup

In hadron-hadron collisions, there are many sources of contamination in jets, including underlying event contributions from proton remnants [87, 88], and pileup due to simultaneous hadron collisions [86]. We will collectively refer to these sources of contamination as *pileup* for simplicity. Pileup contamination biases and smears the “true” value of observables reconstructed from final state particles, driving the need for mitigation techniques.

Pileup is approximately uniformly distributed in the rapidity-azimuth plane. This is exactly the shape probed by the event isotropy, \mathcal{I} . Thus, in order to protect shape observables against pileup contamination, we can compose them with the event isotropy, which will soak up radiation uniform in the plane. This defines the *shapiness+pileup* observable:

$$\mathcal{O}_i^{\mathcal{I}} = \mathcal{O}_i \oplus \mathcal{I}, \tag{3.10}$$

where \mathcal{O}_i is any shape observable, including those previously defined. As a departure from ref. [40], we realize the uniform event by randomly sampling in the plane, rather than defining a grid. We also primarily focus on the $\beta = 1$ event isotropy. Unlike mitigation techniques such as area subtraction [93–95] or jet grooming [96, 97], where an implicit assumption is made about the pileup energy density (either explicitly as an input ρ , or implicitly through a soft scale z_{cut}), the shape observable $\mathcal{O}_i^{\mathcal{I}}$ makes no explicit energy scale assumptions.¹⁹ The uniform energy weight, z_2 , is optimized over, and so the observable “learns” its own pileup scale, which can then be extracted. We choose to initialize the pileup scale z_2 to zero, though one could choose any value of z_2 if they had a prior on the amount of pileup in events.

3.3.6 ... and more!

This has *not* been an exhaustive list — one can use any manifold \mathcal{M} of energy flows one can think of, with the only two limits being imagination and the ability to write down a sampling procedure. Other examples of shapes include polygons, hardcoded jet topologies (for example, two-pronged jets restricted to between $\Delta R = R_1$ and R_2 apart), Gaussian

¹⁹Of course, via the choice of \mathcal{I} , we are still making an explicit assumption about the shape of the pileup distribution, even if the overall energy scale is learned.

clusters, graph-based shapes, and so on. These observables can also be combined into more complex ones using shape composition. All of these can be constructed within the SHAPER framework (more details in section 4), and we encourage the community to use this prescription to develop their own observables.

4 The Shaper framework

Calculating the Wasserstein metric in eq. (2.9) is notoriously difficult; if both events have n particles, then the runtime needed by a brute force, generic Wasserstein solver can be as high as $\mathcal{O}(n^3 \log n)$ [98]. Generic solvers also make it difficult to extract the gradients of the metric with respect to one of the events, $\nabla_{\mathcal{E}} \text{EMD}(\mathcal{E}, \mathcal{E}')$, which are necessary for performing gradient descent over the space of events in \mathcal{M} . Fortunately, by using the (de-biased) Sinkhorn divergence, which uses an ϵ -regularization to approximate the Wasserstein metric, the total costs can be lowered all the way down to $\mathcal{O}(n^2 \log n)$ [46, 99–103].

In this section, we introduce the **Shape Hunting Algorithm for Parameterized Energy Reconstruction** — or SHAPER — to define and calculate shape observables. SHAPER is a PYTORCH-enabled [104] and parallelized computational framework for defining and composing shape observables and their corresponding energy flow manifolds, built using the GEOMLOSS [46] package. We start by outlining the SHAPER algorithm. Then, we provide details on the Sinkhorn divergence, before ending this section with implementation details. For the rest of this paper, we restrict ourselves to balanced observables, i.e. $E_{\theta, \text{tot}} = E_{\text{tot}} = 1$, leaving the unbalanced case for future work.

4.1 The Shaper algorithm

We now describe how to perform the minimization $(\arg)\min_{\mathcal{E}' \in \mathcal{M}} [\text{EMD}(\mathcal{E}, \mathcal{E}')]$ using SHAPER.²⁰ The SHAPER algorithm for estimating shape observables on an event is as follows:

1. **Define:** following the prescription of section 3.1, define a manifold \mathcal{M} and coordinates θ parametrizing the manifold. Define the ground metric $d(x, y)$, the exponent β , and the radius R . This fully defines the observable \mathcal{O} . Build a sampling function p_{θ} that uses the parameters θ to transform some base distribution into a realization of the energy flows $\mathcal{E}'_{\theta} \in \mathcal{M}$. Finally, choose an approximation parameter $\epsilon \ll 1$ and an annealing parameter $\Delta \in (0, 1)$.
2. **Initialize:** for each event \mathcal{E} , choose initial parameters θ . This initialization should be done in an IRC-safe way.
3. **Compute the EMD:** compute the de-biased Sinkhorn divergence, $S_{\epsilon}(\mathcal{E}, \mathcal{E}')$, as defined in section 4.3 below, as an estimate of the EMD. Save the corresponding de-biased Kantorovich potentials, F and G .

²⁰NEEMO [47] is another differentiable EMD estimator that works by parameterizing the space of Lipschitz-Kantorovich potentials.

4. **Gradient Update:** perform the gradient update:

$$\theta \leftarrow \theta - \alpha \left(\sum_{j=1}^M G(y_j) \frac{\partial E'_j}{\partial \theta} + \sum_{j=1}^M E'_j \frac{\partial G(y_j)}{\partial y_j} \frac{\partial y_j}{\partial \theta} \right), \quad (4.1)$$

where α is a learning rate hyper-parameter. The first term is the dependence of the EMD on particle energies due to θ , and the second is the dependence due on particle positions due to θ , both of which are implicit through the sampling function p_θ . This step can be replaced with any other gradient descent optimizer.

5. **Constrain:** if the manifold \mathcal{M} is nontrivial, impose any necessary constraints on the coordinates θ , such as wrapping angles between $-\pi$ and π , enforcing positivity, or a simplex projection.
6. **Converge:** repeat Steps 3–5 until convergence. Return the final value of the EMD and the final θ parameters.

The SHAPER framework contains modules to aid or automate each of these steps, which we describe further in section 4.4.

4.2 The dual formulation of Wasserstein

Observe that the EMD in eq. (2.9) falls into a generic class of problems called *linear programs*. A linear program involves minimizing a function $\mathcal{L}(x) = \langle c, x \rangle$ over vectors x , where c is some cost function linear in x . Furthermore, x satisfies some linear constraint of the form $b = Ax$, and we additionally require $x \geq 0$. In our case, x is the (flattened) transfer matrix π , c is the (flattened) distance matrix d^β , $b = (\mathcal{E}, \mathcal{E}')$ are the energy flows, and A is a matrix enforcing the simplex constraints on π .

The theory of linear programs is well-studied [105]. In particular, for every *primal* linear program, there exists a *dual* linear program, where the constraints and variables to be optimized switch roles, similar to the method of Lagrange multipliers. In the dual problem, one instead maximizes the function $\mathcal{L}(y) = \langle b, y \rangle$, subject to $A^T y \leq c$.²¹ For the Wasserstein metric, the dual formulation looks like:

$$\text{EMD}(\mathcal{E}, \mathcal{E}')^{(\beta, R)} = \max_{f, g: \mathcal{X} \rightarrow \mathbb{R}} [\langle \mathcal{E}, f \rangle + \langle \mathcal{E}', g \rangle], \text{ such that } f(x) + g(y) \leq \frac{1}{\beta R^\beta} d(x, y)^\beta, \quad (4.2)$$

where f and g are known as the *dual potentials* or *Kantorovich potentials*. This formulation of the EMD is known as the *Kantorovich-Rubinstein* metric [106].

In this form, the EMD has several nice properties. First, the arguments \mathcal{E} and \mathcal{E}' are explicit, rather than implicit in the form of constraints. This makes taking the gradient of the EMD with respect to either energy flow much easier. This property is incredibly useful for performing optimizations over energy flows, since it enables easy differentiation. Second, the optimization over an MN -dimensional object, π_{ij} , is replaced by an optimization over

²¹For problems showcasing strong duality, the existence of an optimal solution for the primal problem implies the existence of an optimal solution for the dual problem. The problems we consider in this work admit strong duality. See ref. [106] for a mathematically rigorous discussion.

the $(M + N)$ -dimensional object, f_i and g_j , making the simplex constraint structure more apparent. It can be shown that the optimal choice of f and g actually saturates the bound in eq. (4.2) [107]. Recalling that the ground metric satisfies $d(x, x) = 0$ for all x , we can see that the optimal f, g pair satisfies $f(x) = -g(x)$. This allows us to rewrite the constraint as:

$$|f(x) - f(y)| \leq \frac{1}{\beta R^\beta} d(x, y)^\beta. \tag{4.3}$$

That is, f is β -Hölder continuous. Note that for $\beta = 1$, this reduces to Lipschitz continuity on f .

4.3 Reviewing the Sinkhorn divergence

The source of the difficulty in evaluating eq. (4.2) is the highly nonconvex optimization. To alleviate this, we introduce a regulator [44, 45] to the dual Wasserstein metric:

$$\text{OT}_\epsilon^{(\beta, R)}(\mathcal{E}, \mathcal{E}') = \max_{f, g: \mathcal{X} \rightarrow \mathbb{R}} \left[\langle \mathcal{E}, f \rangle + \langle \mathcal{E}', g \rangle - \epsilon^\beta \log \left\langle \mathcal{E}(x) \otimes \mathcal{E}'(y), e^{\left(\frac{1}{\epsilon^\beta}(f(x)+g(y)) - \frac{1}{\beta R^\beta} d(x, y)^\beta\right)} \right\rangle \right], \tag{4.4}$$

where ϵ is a regulation parameter. The quantity $\text{OT}_\epsilon^{(\beta, R)}(\mathcal{E}, \mathcal{E}')$ is known as the *Sinkhorn divergence* [43] between measures \mathcal{E} and \mathcal{E}' . It reduces to the EMD as $\epsilon \rightarrow 0$.²² Notably, for any $\epsilon > 0$, the optimization over f and g is fully convex [107], making the minimum significantly easier to evaluate.

Note that there are *no* constraints on the functions f and g anymore. Instead, the maximum will only be achieved when $f(x) + g(y)$ is within order ϵ^β of $\frac{1}{\beta R^\beta} d(x, y)^\beta$, a softer version of the original simplex constraint. We can view the parameter ϵ as “blurring” the distance metric $d(x, y)$, where ϵ is a distance scale measured in units of R .²³

As an unconstrained, convex minimization problem, we can estimate the Sinkhorn divergence using simple gradient descent by taking derivatives of eq. (4.4) with respect to f and g . Given two atomic measures, \mathcal{E} and \mathcal{E}' , with M and N particles respectively, and an approximation parameter ϵ , we estimate the Kantorovich potentials $f(x_i)$ and $g(y_j)$ that give us the Sinkhorn divergence using the following algorithm:

1. **Initialize:** initialize $f(x_i) = 0$ and $g(y_j) = 0$.
2. **Gradient Update:** update f and g simultaneously as follows:

$$f(x_i) \leftarrow -\epsilon^\beta \log \left(\sum_{j=1}^N E'_j e^{\left(\frac{1}{\epsilon^\beta}(g(y_j)) - \frac{1}{\beta R^\beta} d(x_i, y_j)^\beta\right)} \right), \tag{4.5}$$

$$g(y_j) \leftarrow -\epsilon^\beta \log \left(\sum_{i=1}^M E_i e^{\left(\frac{1}{\epsilon^\beta}(f(x_i)) - \frac{1}{\beta R^\beta} d(x_i, y_j)^\beta\right)} \right). \tag{4.6}$$

²²In the $\epsilon \rightarrow \infty$ limit, we recover instead the maximum mean discrepancy (MMD) [84], another potential metric on collider events, which we have shown in section 2.4 is not faithful.

²³Even though the R parameter is unimportant for calculating the exact Wasserstein metric for balanced observables, beyond defining a unit scale, its importance re-emerges when defining the blurring scale ϵ .

3. **Converge:** repeat Step 2 until convergence. Return the Kantorovich potentials f and g , and the Sinkhorn Divergence eq. (4.4) evaluated on these potentials.

This algorithm is known to converge in finite time [43, 108, 109]. The runtime of each iteration scales as approximately $\mathcal{O}((M+N)^2)$, and the algorithm converges in approximately $\frac{1}{\epsilon^\beta}$ iterations. This can be further improved to only $\log\left(\frac{1}{\epsilon^\beta}\right)/\log\left(\frac{1}{\Delta}\right)$ iterations through the use of simulated annealing [110, 111], with a parameter $\Delta \in (0, 1)$. Beginning with a larger effective blurring radius $\epsilon' = 2R$, after every iteration of the Sinkhorn algorithm, we decrease $\epsilon' \leftarrow \Delta\epsilon'$, until finally reaching $\epsilon' = \epsilon$. Intuitively, we start with a large blurring scale R , and slowly “zoom in” to a distance scale of ϵ to refine the estimate of the Sinkhorn divergence.

However, the Sinkhorn divergence is biased, meaning that it is not generically the case that $\text{OT}_\epsilon(\mathcal{E}', \mathcal{E}') = 0$, which is an important property of the Wasserstein metric. We use therefore the de-biased Sinkhorn divergence, defined in ref. [46]:

$$S_\epsilon(\mathcal{E}, \mathcal{E}') = \text{OT}_\epsilon(\mathcal{E}, \mathcal{E}') - \frac{1}{2}\text{OT}_\epsilon(\mathcal{E}, \mathcal{E}) - \frac{1}{2}\text{OT}_\epsilon(\mathcal{E}', \mathcal{E}'). \quad (4.7)$$

The de-biased Sinkhorn divergence satisfies $S_\epsilon(\mathcal{E}, \mathcal{E}) = 0$ by construction. This can be easily realized algorithmically by simply substituting new de-biased Kantorovich potentials $f \rightarrow F$ and $g \rightarrow G$, where

$$F(x) = f(x) - \tilde{f}(x), \quad (4.8)$$

$$G(y) = g(y) - \tilde{g}(y). \quad (4.9)$$

Here, the notation $\tilde{f}(x)$ refers to the first Kantorovich potential corresponding to $\text{OT}_\epsilon(\mathcal{E}, \mathcal{E})$, and $\tilde{g}(y)$ refers to the second Kantorovich potential corresponding to $\text{OT}_\epsilon(\mathcal{E}', \mathcal{E}')$. For the rest of this paper, we refer to the de-biased Sinkhorn divergence as simply the Sinkhorn divergence wherever there is no chance for confusion.

In addition to returning the Sinkhorn divergence, this algorithm also returns the Kantorovich potentials, which allow us access to approximate gradients of the EMD, which can be used for shape parameter optimization. The gradients of the EMD with respect to the input measures can be read off of eq. (4.4):

$$\begin{aligned} \nabla_{\mathcal{E}}\text{EMD}(\mathcal{E}, \mathcal{E}') = F &\Rightarrow \begin{cases} \nabla_{E_i}\text{EMD}(\mathcal{E}, \mathcal{E}') = F(x_i), \\ \nabla_{x_i}\text{EMD}(\mathcal{E}, \mathcal{E}') = E_i\nabla F(x_i), \end{cases} \\ \nabla_{\mathcal{E}'}\text{EMD}(\mathcal{E}, \mathcal{E}') = G &\Rightarrow \begin{cases} \nabla_{E'_j}\text{EMD}(\mathcal{E}, \mathcal{E}') = G(y_j), \\ \nabla_{y_j}\text{EMD}(\mathcal{E}, \mathcal{E}') = E'_j\nabla G(y_j). \end{cases} \end{aligned} \quad (4.10)$$

4.4 Implementation details

Before turning to our case study, we discuss the specific details of the SHAPER implementation. SHAPER uses the GEOMLOSS package as a backend for computing Sinkhorn divergences, as described in section 4.3. By default, we use a relatively conservative annealing value of $\Delta = 0.9$, and choose $\epsilon^\beta = 10^{-3}$ for our estimates.²⁴ Currently, the SHAPER algorithm is implemented only for balanced observables, with either $\beta = 1$ or 2.

²⁴See sections 5.2 and 5.3 for studies involving the choice of ϵ .

Once the EMD and Kantorovich potentials are estimated using the GEOMLOSS backend, the gradient updates in Step 4 are then handled using automatic differentiation and backpropagation in PYTORCH. One may select from a suite of common machine learning optimizers to perform the gradient update — by default, we use the ADAM optimizer [112], with a learning rate of 0.01. Steps 3–5 can be easily parallelized over batches of hundreds of events at once. This is accomplished by treating the parameters θ for each \mathcal{E}_n to be completely independent. The Sinkhorn divergence can be computed on many events at once, and since the parameters are independent, we take advantage of highly parallelized PYTORCH operations to perform independent derivatives ∇_{θ_n} of the combined batch loss, $\sum_n \text{EMD}(\mathcal{E}_n, \mathcal{E}'_{\theta_n})$, all at once.

When using SHAPER, one must specify a maximum number of epochs, so that the program eventually halts even if convergence is not achieved. We set this number by default to be 500 epochs, though we observe that convergence happens far earlier than this. We define convergence through an early stopping procedure: if an event’s EMD has not decreased in at least N_{\max} epochs, stop early, and return the minimum EMD ever achieved during the training, and the parameters that achieved that minimum. When training on a large batch of events, we stop early when a certain fixed percentage (we choose 95% by default) have hit this condition — this is because there tends to be a handful of outlier events that take exceptionally long to converge. We choose $N_{\max} = 25$ epochs by default. Both this, and the batch stopping percentage, are adjustable user parameters.

To facilitate Steps 1 and 2 of the SHAPER algorithm, where the user input occurs, many common manifolds, such as sets of N points, hypercubes, simplices, and so on, are already pre-built. These pre-built “building-block” manifolds are listed in table 4.²⁵ From these, more complex manifolds can be easily constructed, such as the rings, disks, and ellipses described in section 3.3. The \oplus and $N \times$ operators make it very easy to define new composite observables from old ones and quickly build sophisticated shapes. Parameters can also be easily frozen for more customization. Furthermore, it is straightforward to define more building-block manifold as needed within the framework. Each of the observables described in section 3.3 is also pre-built into SHAPER. When defining custom manifolds that use the N -points as a building block, the custom shape will automatically use the same IRC-safe k_T -clustering initialization scheme as the observables in section 3.3, though it is possible to modify the initialization scheme as needed.

Each manifold contains instructions on how to enforce parameter constraints (as needed for Step 5). For most manifolds, this usually involves clipping the values to be within a desired range, though for the simplex Δ_N , which occurs in almost every shape for which energy weights must be normalized, the enforcement is nontrivial. We use a simplex projection algorithm, inspired by the K -Deep-Simplex framework [113] and its nonlinear ex-

²⁵Note that these manifolds are the *parameters* from which shapes \mathcal{E}_θ can be defined, not the shapes themselves. For example, a circular shape is defined by a point parameter (the center) and a positive real parameter (the radius), whereas the circle manifold S^1 whenever an angle parameter is needed to define a shape.

Manifold	Description	Constraints
Trivial	The set $\{0\}$	None
N -Points, \mathcal{X}^N	Set of N points, x_i , in \mathbb{R}^2	None
Positive Reals, $\mathbb{R}_{\geq 0}^N$	Set of N positive numbers, R_i	Clipped to $R_i \geq 0$
Hypercube, Λ^N	Set of N numbers, λ_i , between 0 and 1	Clipped to $0 \leq \lambda_i \leq 1$
Circle / Torus, S^N	Set of N angles, ϕ_i	Wrapped to $-\pi \leq \phi_i < \pi$
N -Simplex, Δ_N	Set of $N + 1$ numbers, $z_i \geq 0$, summed to 1	Simplex projection

Table 4. The building-block manifolds implemented in SHAPER, from which parameters can be defined and more complex manifolds can be built. For each manifold, a description is given, and constraint instructions are provided to project points onto the manifold. Note that we do not distinguish here between “positive” reals versus “non-negative” reals, since for continuous parameterizations there is no difference between including zero and getting arbitrarily close to zero.

tension [114], which solves the linear program [115]:

$$\min_{z_i} \left[\sum_i^{N+1} |z_i - y_i|^2 \right] \text{ such that } \sum_i^{N+1} z_i = 1, z_i \geq 0, \tag{4.11}$$

which finds the simplex $z_i \in \Delta_N$ closest to a set of unnormalized points y_i .

5 Empirical studies with jets

We now use the SHAPER framework and custom shape observable in example collider physics analyses. We begin by benchmarking the SHAPER algorithm, testing the performance of the Sinkhorn divergence and the optimization procedure by comparing the jet isotropy and N -subjettiness calculated using SHAPER to other methods. Next, we calculate the ring, disk, and ellipse-based observables defined in section 3.3 for a dataset of top and QCD jets, showing visualizations of each shape and analyzing the learned EMD’s and parameters. Finally, we explore the potential to use shape observables for automatic pileup removal.

5.1 Dataset

For our empirical studies, we use the top tagging benchmark of refs. [48, 49], which is a dataset consisting of a top quark jet signal and a mixed light-quark/gluon jet background. These samples are generated in PYTHIA 8.2.15 [116] at 14 TeV, and then passed through DELPHES 3.3.2 [117] to simulate the ATLAS detector. Jets are defined using the anti- k_T algorithm [64] in FASTJET 3.1.3 [118] with $R = 0.8$. Only the leading jet in any event is considered, and we select jets satisfying $p_{T,J} \in [475, 525]$ GeV and $|\eta_J| < 2$. The signal and background samples are generated using $t\bar{t}$ and QCD dijet events respectively. For signal top jets, a top parton, plus its decay products, are required to be within $\Delta R = 0.8$ of the jet axis. All events are translated such that the jet axis is at $(0, 0)$ on the rapidity-azimuth plane.




Figure 2. Example of a top jet (signal) event in our dataset, (a) before and (b) after the pileup mock-up procedure described in section 5.1. Each red point is a single particle, and the size of the point is proportional to the relative energy of the particle.

In this dataset, multiple parton interactions and pileup have not been included. To mock up the effects of pileup contamination in data, we add in pileup “by hand”. To each event, we add in N particles randomly distributed in an $R \times R = 0.8 \times 0.8$ square centered at the origin on the rapidity-azimuth plane, where N is Poisson-distributed with a mean of 75. Each particle is given an energy weight randomly sampled from a normal distribution with mean $\frac{E_{PU}}{N}$ and standard deviation $\frac{E_{fluct}}{N}$, where we take E_{PU} , which represents the total amount of pileup radiation, to be uniformly distributed between 50 and 250 GeV, and E_{fluct} , which represents per-particle fluctuations, to be 25 GeV.²⁶ Many refinements of this simplistic mockup could be considered, but this suffices to show qualitative features of SHAPER and the shape observables defined in section 3.3. For the purposes of calculating shape observables, all jets are normalized such that $E_{tot} = 1$, though we save the original total energy of each jet for the purpose of restoring units.

An example top jet is shown in figure 2, before and after pileup is added, to illustrate this procedure. This contamination procedure is performed for the benchmarking studies in sections 5.2 and 5.3 and the pileup studies in section 5.7. For the jet substructure studies in sections 5.5 and 5.6, we do not add any pileup, and instead we require that the jets have an invariant mass $m_J \in [145, 205]$ GeV to more closely match the analysis conditions of ref. [68].

5.2 Benchmarking Sinkhorn: jet isotropy

We first use SHAPER to compute the jet isotropy for the purposes of benchmarking the Sinkhorn divergence for runtime and accuracy. Jet isotropy is an ideal benchmark since

²⁶A floor of 0 GeV is set to avoid negative energies.

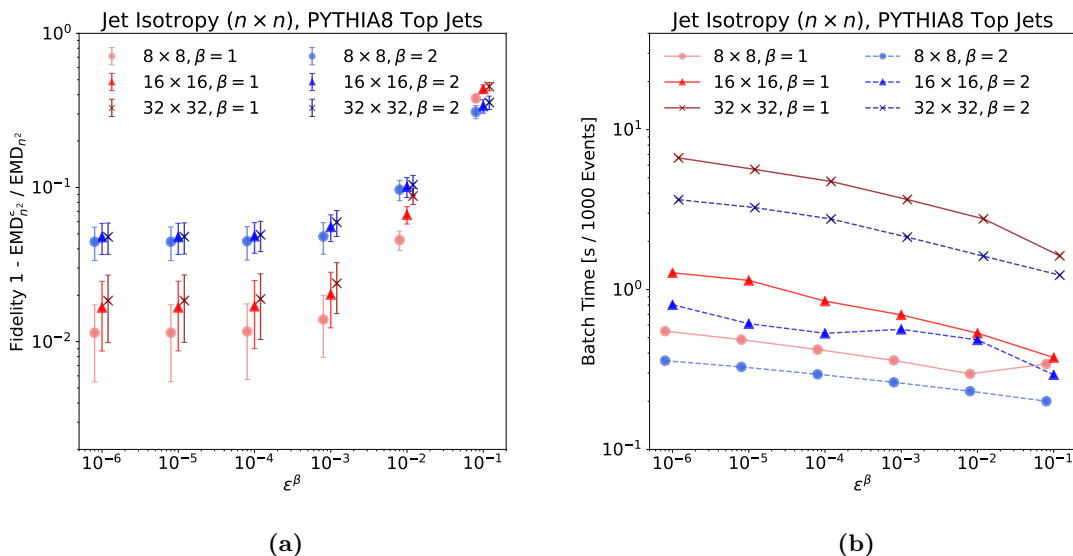


Figure 3. The (a) fidelity and (b) runtime of SHAPER when computing the $n \times n$ jet isotropy, for different values of ϵ , n , and β . The fidelity is defined as the ratio of the Sinkhorn divergence to the “true” Wasserstein metric, as computed using the POT library, across a batch of 1000 events. The runtime is the total time to evaluate the Sinkhorn divergences for the entire 1000 event batch, as computed using a NVIDIA A100. An annealing parameter of $\Delta = 0.9$ is used globally.

the minimization is trivial; for balanced isotropy, the parameterized manifold consists only of a single event. Therefore, no gradient descent is necessary, and this is purely a test of the Sinkhorn approximation. This can be viewed as a proxy for the per-epoch runtime and accuracy of the SHAPER algorithm.

We compute the $n \times n$ jet isotropy, which is the isotropy given by computing the EMD to the uniform event:

$$\mathcal{U}_{n \times n} \sim \text{particles in an } R \times R \text{ square, arranged in a uniform } n \times n \text{ grid}, \quad (5.1)$$

as defined in ref. [40]. We do this using SHAPER with many different values of ϵ , and compare to the same calculation done using the Python Optimal Transport (POT) [119] implementation of the EMD, which was used in refs. [40, 41].

In figure 3, we show the results of a runtime vs. accuracy study. We compute the jet isotropy of 1000 top jets for several different values of n , ϵ , and for $\beta = 1$ and 2. In figure 4, we show the learned 16×16 jet isotropies for $\epsilon^\beta = 10^{-3}$. For this experiment, we use a fixed (conservative) annealing parameter of $\Delta = 0.9$. SHAPER allows for events to be computed in parallelized batches; we run the entire computation in a single batch on a NVIDIA A100, and report the total runtime of the entire batch.²⁷

²⁷In principle, the only limiting factor to how many events SHAPER can process at once is the ability to fit everything on a single GPU. We find that we can run up to 10000 events in parallel on 32 GB of memory of a NVIDIA A100.

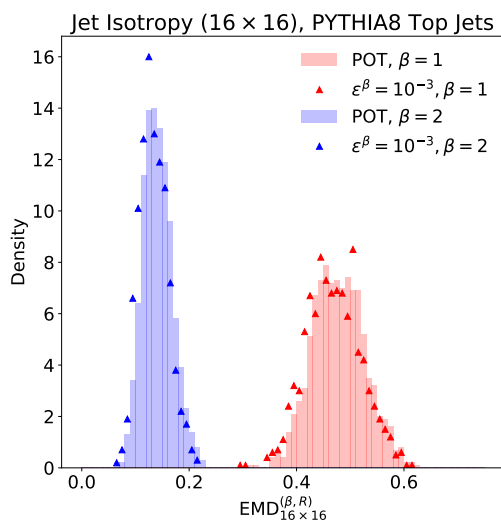


Figure 4. Distributions of the learned 16×16 jet isotropy, for $\beta = 1$ (red) and $\beta = 2$ (blue), as calculated using the POT library (filled) and SHAPER (points) with $\epsilon^\beta = 10^{-3}$.

We see from figure 3 that the accuracy of the Sinkhorn divergence, as an estimator for the Wasserstein metric, begins to saturate at $\epsilon^\beta = 10^{-3}$, and that there is no substantial gain from choosing a smaller ϵ . Picking $\epsilon^\beta = 10^{-3}$ ensures percent level accuracy for $\beta = 1$, and few-percent level accuracy for $\beta = 2$, which we can see visually in figure 4. Furthermore, for larger values of $n \times n$, the accuracy is mostly independent of n . We also observe from figures 3 and 4 that Sinkhorn tends to slightly underestimate the Wasserstein metric, which can be understood from the strictly negative ϵ^β -regulator in eq. (4.4). Note that for $n = 16$, it takes under 1 second to process 1000 events — this implies it is possible to process millions of events on the order of an hour, with further speedups possible by choosing a more aggressive value for the annealing parameter Δ .

5.3 Benchmarking optimization: N -subjettiness

We perform a second benchmark investigation using N -subjettiness. Unlike the jet isotropy, N -subjettiness requires a nontrivial minimization. This allows us to use it to estimate the fidelity of the SHAPER algorithm’s optimization step.

It is well known that the ratio $\tau_{32} \equiv \tau_3/\tau_2$ is a good discriminant between top and QCD jets, as top jets tend to have 3 prongs more often than QCD jets, and thus have lower expected values of τ_{32} [68]. We compute this ratio for several different values of ϵ to see if any discrimination power is lost (or gained) in the ϵ -approximation. Within the SHAPER framework, the N -subjettiness is given by the following manifold of parameterized events:

$$\mathcal{E}_{x_i, z_i}(x) = \sum_{i=1}^N z_i \delta(x - x_i), \quad x_i \in \mathcal{X}, z_i \in \Delta_{N-1}. \tag{5.2}$$

As a baseline, we compute N -subjettiness using FASTJET 3.4.0 with FJCONTRIB 1.050. The results of this study are shown in figure 5a as ROC curves, computed using an NVIDIA

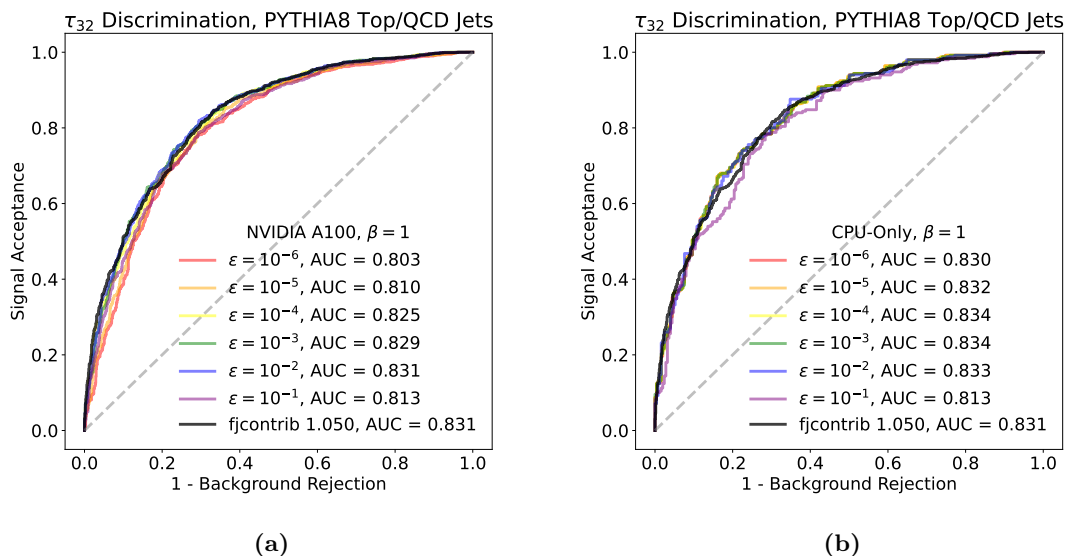


Figure 5. A ROC curve showing the performance of $\tau_{32} = \tau_3/\tau_2$ as a discriminator between top (signal) and QCD (background) jets, for several different values of ϵ , as calculated using (a) an NVIDIA A100 and (b) only on CPU. A baseline curve, calculated using the N -subjettiness routines in FASTJET 3.4.0, is shown in black.

A100 GPU. We see that the Sinkhorn approximations have roughly the same discriminatory power as the baseline for $\epsilon \sim 10^{-3}$. In figure 6, we show the distributions of τ_{32} for both datasets computed with FASTJET and SHAPER with $\epsilon = 10^{-3}$, and see good agreement between the two methods. As with the isotropy study in section 5.2, we observe that SHAPER tends to slightly underestimate the observable. In order to gauge the impact of float precision on our estimates, we repeat ROC curve calculation using only a CPU, which is shown in figure 5b. For values of $\epsilon \ll 10^{-3}$, we see that on the GPU architecture, the performance actually begins to degrade due to the lower machine precision due to the accumulation of floating-point errors during the optimization, while it saturates on the CPU. Therefore, it is recommended to use $\epsilon^\beta \sim 10^{-3}$, as this is the most stable compromise between fidelity and machine precision.

5.4 Hearing gradients

SHAPER can be used to not only estimate the shapiness of events, but also take derivatives of the shapiness with respect to the event. As discussed in section 4.3, this is completely automatic, since the gradients with respect to the energy flow are given manifestly by the Kantorovich potentials, allowing us to see precisely how our EMD calculations depend on the energies and positions of particles in an event.

Reading off of eq. (4.10), we obtain an expression for the gradient of the EMD with respect to the energy E_i of particle i :

$$\nabla_{E_i} \text{EMD}(\mathcal{E}, \mathcal{E}'_\theta) = F(x_i). \tag{5.3}$$

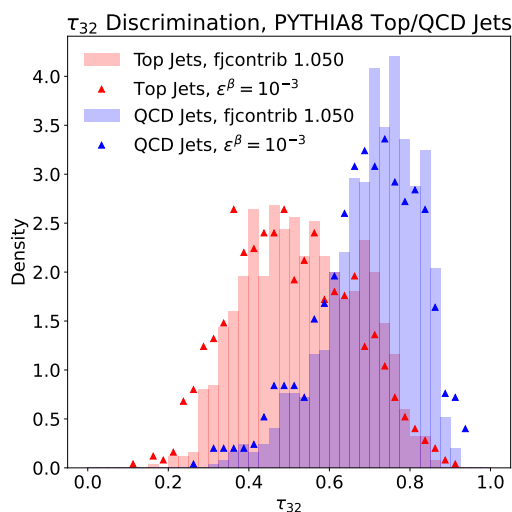


Figure 6. Distributions of the learned τ_{32} , for top jets (red) and QCD jets (blue), as calculated using FASTJET 3.4.0 (filled) and SHAPER (points) with $\epsilon^\beta = 10^{-3}$.

If the gradient at particle i is negative, then increasing the energy of that particle will decrease the EMD, making the event more \mathcal{M} -like. By adding “ghost” particles to \mathcal{E} , one can probe the energy dependence of the EMD from any point in \mathcal{E} . Similarly, we can take the gradient the EMD with respect to the position x_i of particle i :

$$\nabla_{x_i} \text{EMD}(\mathcal{E}, \mathcal{E}'_\theta) = E_i \nabla_{x_i} F(x_i). \tag{5.4}$$

This gradient (times -1) tells us where to move the particle i to decrease the EMD. Both the energy and position gradients answer the question, “If I want to decrease the EMD (make my energy look *more* like my shape), what should I do to a particle at site x_i ?” Moreover, because of the reparameterization invariance of the energy flow density, the energy and position gradients are both are valid ways to change the EMD: one can either change the energy of particles at the location x_i , move the particles at x_i somewhere else, or some combination of both.

In figures 7 and 8, the gradients from eqs. (5.3) and (5.4) are plotted for an example top jet, for the 3-subjettiness and $16 \times 16, \beta = 1$ jet isotropy, respectively. Using figure 7, we can see what parts of the event contribute to the 3-subjettiness — the three large clusters contribute negatively to the EMD (make the event look more like 3 subjets), while the rest of the event contribute positively to the EMD (makes the event deviate from 3 subjets). Similarly, we see from figure 8 that the overdensity of energy at the center of the event makes it less isotropic. In both figures, the vector quiver plot tells us which way particles should “flow” (against) to change the shape.

There are many potential applications of calculating gradients of a shape with respect to an energy flow. In experimental contexts, for example, one can use the fact that the gradients are practically instantaneous to compute to do easy Gaussian error propagation due

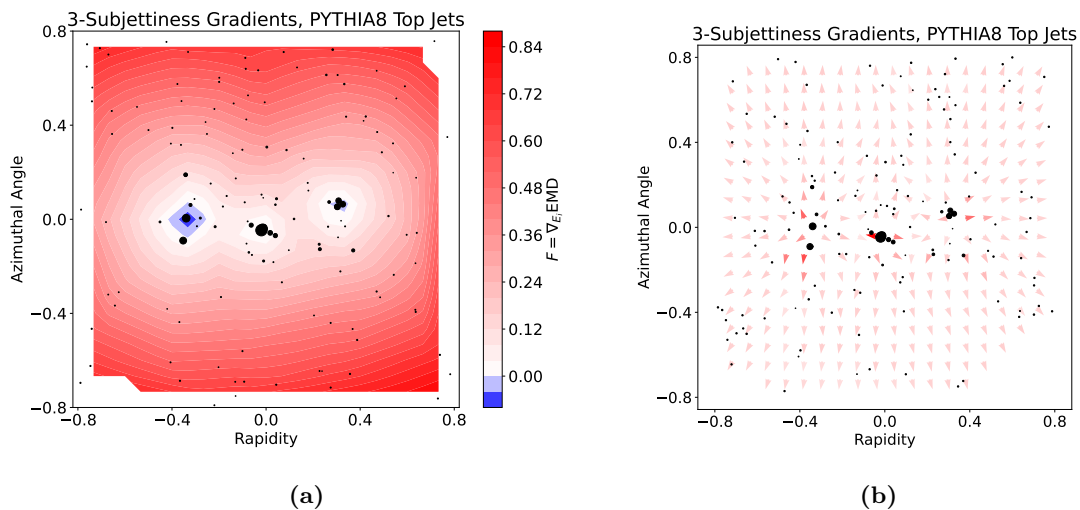


Figure 7. For an example top jet, (a) gradients of the 3-subjettiness with respect to the particle weights, $\nabla_{E_i} \text{EMD}$, and (b) gradients off the 3-subjettiness with respect to the particle positions, $\nabla_{x_i} \text{EMD}$, are shown. In (a), increasing the particle energy in red regions will increase the 3-subjettiness (look less like 3 subjets), and increasing particle energy in blue regions will decrease the 3-subjettiness (look more like 3 subjets). In (b), moving particles along or against the arrows will increase or decrease the 3-subjettiness, respectively, with the arrow's shading indicating the relative magnitude of the change.

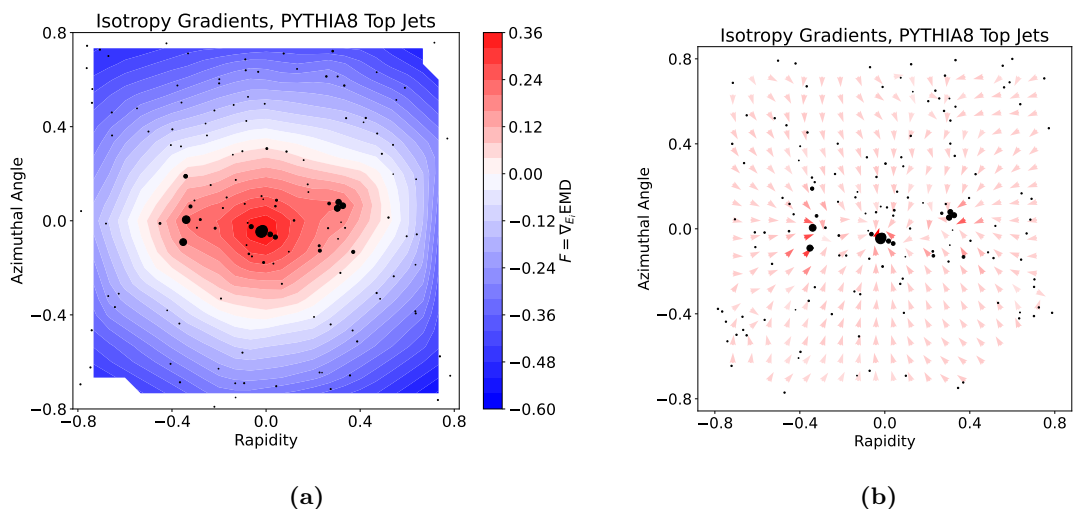


Figure 8. The same as figure 7, but for the 16×16 jet isotropy with $\beta = 1$.

to detector-induced uncertainties in particle energies and positions [120, 121]. This avoids having to do expensive re-sampling and recalculation of the event shape. Phenomenologically, these gradients can be used to probe the sensitivity of observables to certain radiation patterns. For example, the sensitivity of an observable to pileup can be measured by taking derivatives with respect to the pileup scale [95], which can be numerically realized in SHAPER, a potential avenue for future work.

5.5 Hearing jets ring (and disk, and ellipse)

We next use SHAPER to realize the custom shape observables defined in section 3.3, starting with some visualizations of these shape observables on an example event.

The three base observables we consider are the ringiness, diskiness, and ellipsiness, as defined in section 3.3. We calculate the shape observables $N \times \mathcal{O}$ (the N -ringiness, the N -diskiness, and N -ellipsiness) and $N \times \mathcal{O}^\tau$ (the N -(ring+point)iness, the N -(disk+point)iness, and N -(ellipse+point)iness), as defined in section 3.3 with $\beta = 1$. We also consider, for comparison, the N -subjettiness, τ_N , as defined in eq. (5.2). We use SHAPER to evaluate all of these shape observables on a single of top jet event from our dataset, restricted to $m_J \in [145, 205]$ GeV, though without any pileup contamination. Each extended shape is sampled with 100 points, with $\epsilon = 10^{-3}$ and $\Delta = 0.9$.

Geometric visualizations of each of the 21 event shapes, as evaluated on an example top jet, can be found in figures 9, 10, 11, and 12. From these visualizations, we note some interesting qualitative features of these shape observables. First, the point variants of each shape correspond more closely to clusters of energy. For example, while the $N = 3$ uniform rings (figure 10), disks (figure 11), or ellipses (figure 12) do not necessarily capture the regions of highest energy, the point variants of each shape align very well with the N -subjettinesses of figure 9. Correspondingly, the EMD's of the shapes are significantly reduced for the point variants — this suggests that this event in particular is not well modeled by uniform radiation profiles, but rather looks more like localized spikes with radiation clouds around them. Moreover, we can see that circular radiation clouds do not model the event as well as elliptical ones — this is reflected in the N -ellipsiness in figure 12, which learns extremely eccentric line-like structures in an attempt to best model the event, which results in lower EMD's than the corresponding N -diskinesses in figure 11. We also note that the N -subjettiness and the point shape variants all qualitatively find the same jet centers, suggesting that these shapes can be treated as perturbations to N -subjettiness.

5.6 Shapiness and shape parameters

Continuing the discussion in section 5.5, we now use SHAPER to compute distributions of these shape observables on a large sample of top and QCD jets, restricted to $m_J \in [145, 205]$ GeV, though without any pileup contamination. In particular, we show the utility of both the *shapiness* $\mathcal{O}(\mathcal{E})$ and the *shape parameters* $\theta(\mathcal{E})$ in describing the geometry of jets.

For each histogram in this section, we calculate an AUC score, showing the efficacy of a cut on that observable as a top/QCD discriminant. Note that these AUCs are for cuts on a single feature — the discrimination power can in principle be improved by transforming these features and combining many features per jet.




Figure 9. The (a) 1-, (b) 2-, and (c) 3-subjettiness of an example top jet event. Subjets are represented by a purple “ \times ”, with size proportional to the subjet’s energy weight.

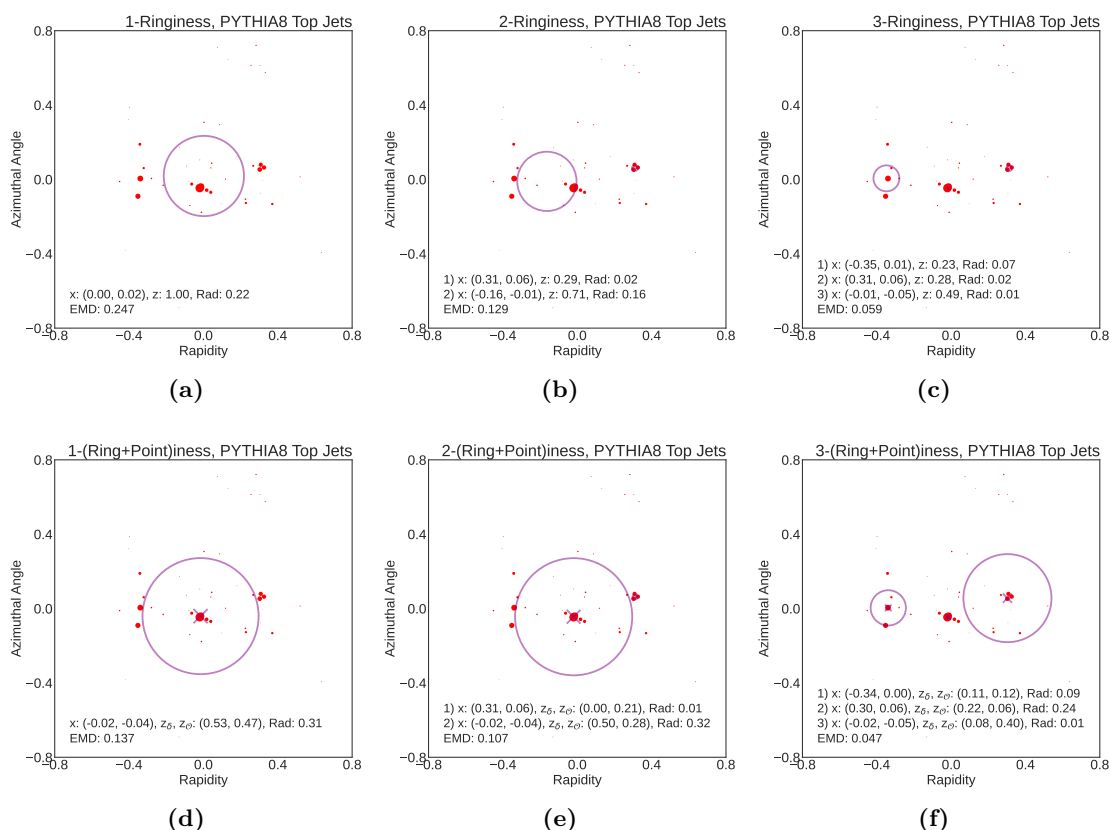


Figure 10. Top row: the (a) 1-, (b) 2-, and (c) 3-ringiness of an example top jet event. Bottom row: the (d) 1-, (e) 2-, and (f) 3-(ring+point)iness of the same top jet event. The point is represented by a “ \times ”, with size proportional to its energy weight.




Figure 11. Top row: the (a) 1-, (b) 2-, and (c) 3-diskiness of an example top jet event. Bottom row: the (d) 1-, (e) 2-, and (f) 3-(disk+point)iness of the same top jet event. The point is represented by a “x”, with size proportional to its energy weight.

As a representative sample of the “shapiness” observable, we plot the N -ellipsiness and N -(ellipse+point)iness of our top and QCD jet samples in figure 13. The EMD distributions for the ring and disk variants are qualitatively similar to the ellipse, and thus our discussion of these distributions carry over to them. We notice that the N -(ellipse+point)iness is not much lower than the corresponding N -ellipsiness, indicating that (at least in the absence of pileup), subjects can indeed be approximated as roughly uniform. In the test event visualization in figure 12, we can see qualitatively that the found ellipses have roughly the same center, comparing the N -ellipsiness and its corresponding point variant. We also note that the EMD decreases with N , as expected.

In figure 14, we show the *ratio* of the $N = 3$ shapiness to the $N = 2$ shapiness for each class of shape observable. For N -subjettiness, which we show for comparison, this is the classic observable τ_{32} , which is known to be a good top vs. QCD jet discriminant [68]. First, we observe that the uniform ring, disk, and ellipse observable, along with the point variants, each have an AUC of approximately 0.75, which is still considerably less than τ_{32} ’s AUC of 0.825. One should expect that, using just the EMD alone, a more complexly parameterized shape should have a lower AUC than a simpler shape. In the extreme case,

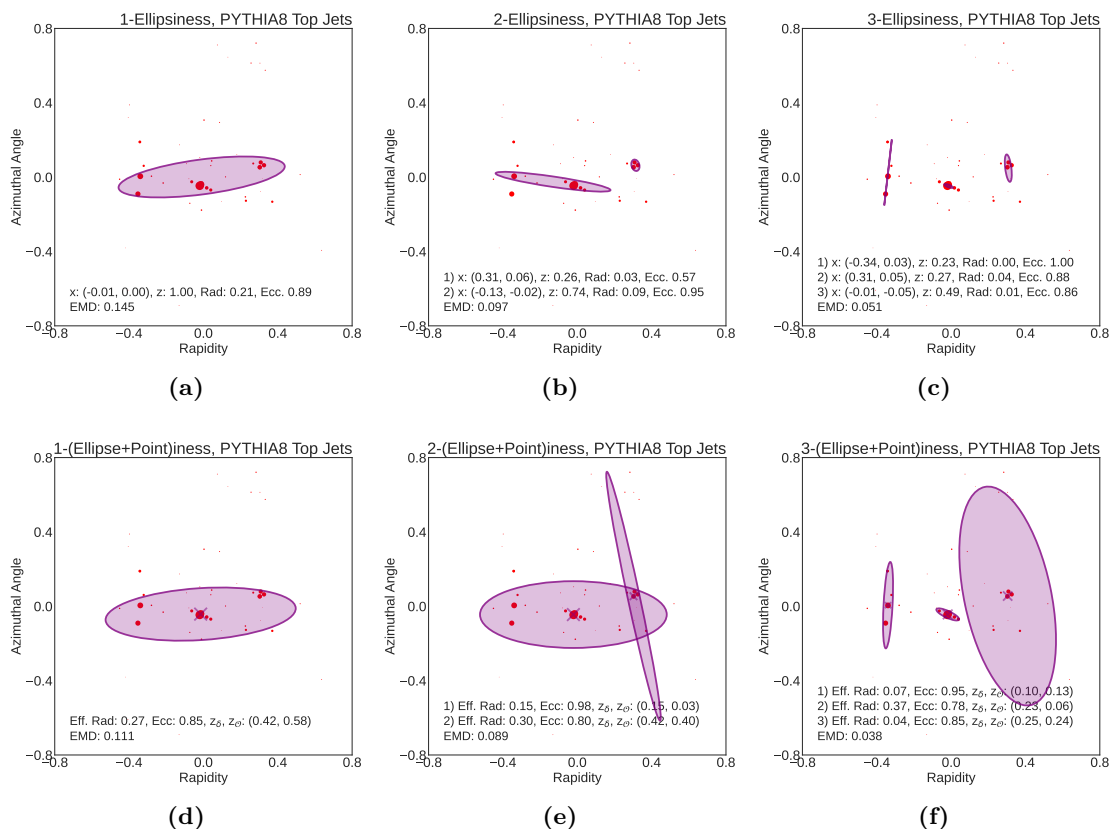


Figure 12. Top row: the (a) 1-, (b) 2-, and (c) 3-ellipsiness of an example top jet event. Bottom row: the (d) 1-, (e) 2-, and (f) 3-(ellipse+point)iness of the same top jet event. The point is represented by a “ \times ”, with size proportional to its energy weight. The effective radius of the ellipses are given by the geometric mean of the two axes, \sqrt{ab} .

where the parameterization is flexible enough to reproduce any event in the dataset, the EMD will always be zero and have no discriminatory power. However, this is not the end of the story — shape observables also include their learned parameters, and this information also contains multivariate discriminatory power. For the hypothetical infinitely flexible shape, the parameters contain the full event information even though the EMD is zero, and thus the combination of the shapiness and shape parameters together contain more information (and thus more discriminatory power) than just a simpler shape. We leave a full multivariate analysis of shape parameters for jet classification for potential future work.

We now move on to analyze the learned shape parameters themselves. In figure 15, the learned radius parameters of the $N = 1$ shapes are shown.²⁸ Considering the 1-(disk+point)iness and 1-(ellipse+point)iness, which as discussed earlier, we can think of as a “jet algorithm” looking for soft clusters of radius R with collinear radiation inside, we see that both learn an average radius of approximately $R = 0.4$, which is approximately half

²⁸For ellipses, we define an *effective radius*, taken to be \sqrt{ab} . This is the radius a circle with the same area as the ellipse.

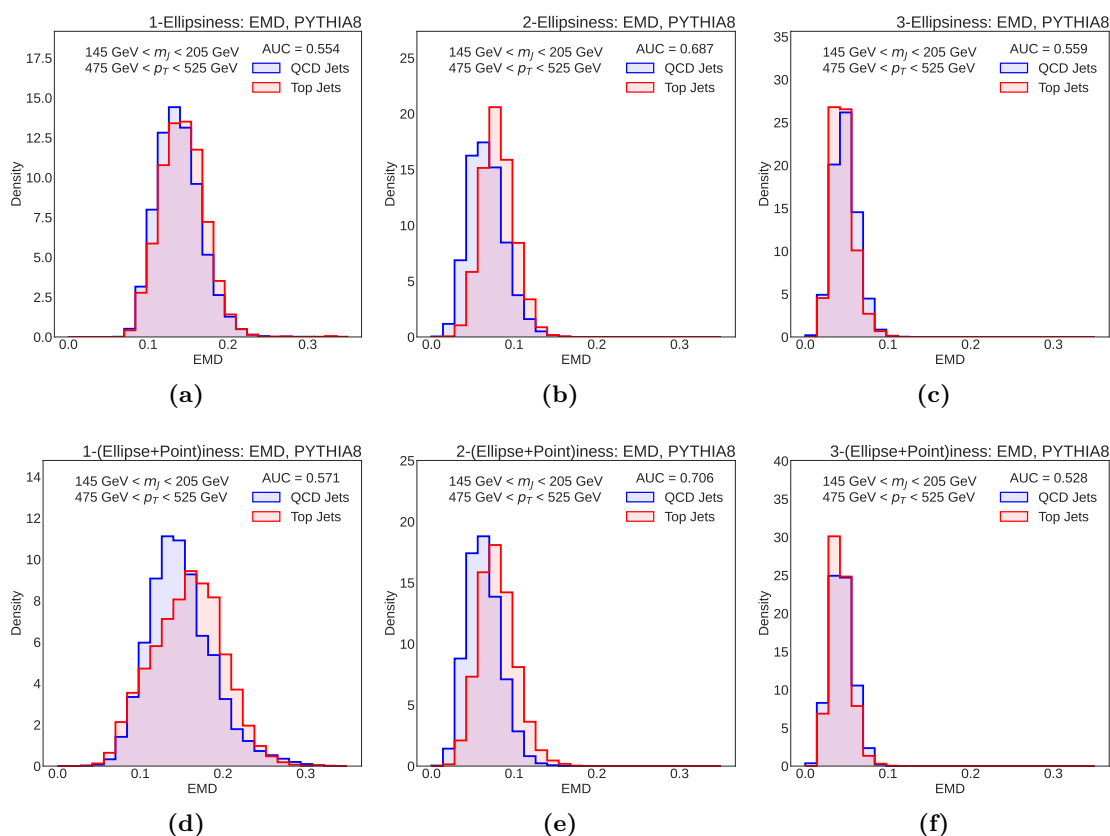


Figure 13. Distributions of the learned (a,d) 1-, (b,e) 2-, and (c,f) 3-ellipsiness (top row) and (ellipse+point)iness (bottom row) of the top (red) and QCD (blue) jet sample.

the radius of the original AK8 jet. One can think of the uniformity condition as imposing an radius effective cutoff. If one assumes that the energy density falls to zero with distance from the jet center, then there is some critical radius for which any larger uniform shape is an increasingly worse approximation, and in principle this critical radius can be computed in perturbative QCD.

For the point shape variants of the disk and ellipse, we note that there is only a small difference between the top and QCD jets — that is to say, the soft component of a top jet is about as “wide” as the soft component of a QCD jet, as seen by disks and ellipses. However, as top jets tend to have their energy distributed across 3 prongs, rather than 1 prong as in QCD jets, we should expect clusters of radiation away from the central one for top jets. Since rings are thin, we should expect rings to be able to better capture these localized prongs than area-filling disks and ellipses, which can be visualized in figure 10. This is especially apparent in figure 15d for the 1-(ring+point)iness shape, where top jets are significantly larger than QCD jets. Interestingly, we note a spike at $R = 0$ in all of these distributions. This means that SHAPER has determined that the best shape for these events is in fact point-like, and reduces to the N -subjettiness. This spike is reduced for the ellipse shapes, which implies that there are events better modeled by either an eccentric ellipse or a point rather than a uniform disk.




Figure 14. Top row: distributions of the $N = 3$ to $N = 2$ ratio of the (a) N -subjettiness of the top (red) and QCD (blue) jet sample. Middle row: the same ratio, but for the N - (b) ringiness, (c) diskiness, and (d) ellipsiness. Bottom row: the same, but for the N - (e) (ring+point)iness, (f) (disk+point)iness, and (g) (ellipse+point)iness.

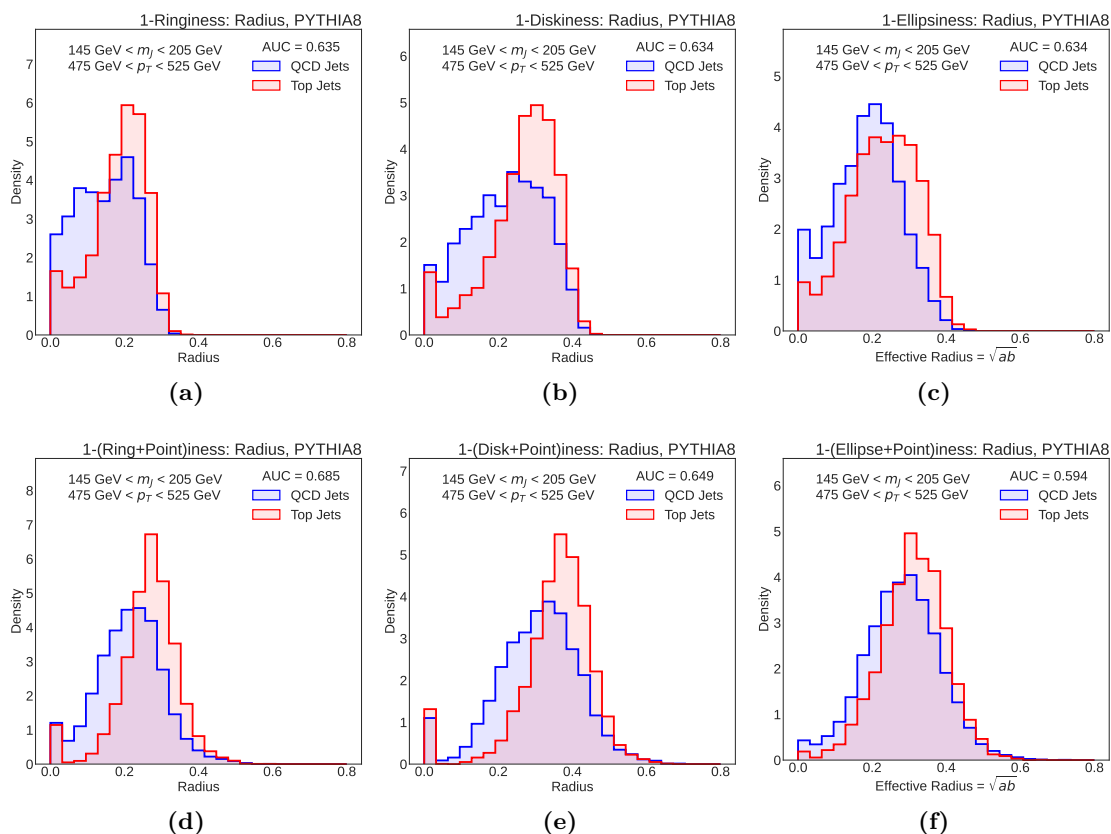


Figure 15. Top row: distributions of the learned radius parameter of the 1- (a) ringiness, (b) diskiness, and (c) ellipsiness of the top (red) and QCD (blue) jet sample. Bottom row: the radius parameter of the 1- (d) (ring+point)iness, (e) (disk+point)iness, and (f) (ellipse+point)iness of the same samples. For the ellipse, the effective radius is given by the geometric mean of the two axes.

As a last example of the information that shape observables and shape parameters can extract, we look at the learned eccentricity of ellipses. In figure 16, we show the *minimum eccentricity* across the N ellipses for the 1-, 2-, and 3- ellipsiness and (ellipse+point)iness for our top and QCD jet datasets, where the eccentricity is given by $\sqrt{1 - \min(a, b) / \max(a, b)}$. We can immediately see that a value of $\min(e) = 0$ is rare — our jets are much better described by ellipses than they are described by disks. With increasing N , the distributions of $\min(e)$ tend to shift leftwards for both the top and QCD jet samples. This indicates that the N -’th ellipse, which probes more substructure than the $(N - 1)$ -’th ellipse, is often less eccentric.

The eccentricity distributions for top jets are slightly different from QCD jets. This is a nontrivial result; top jets and QCD jets have very different prong structures (as demonstrated in figure 14), so the angular distribution of radiation in the rapidity-azimuth plane is less circular for top jets than QCD jets. Thus, jet eccentricity could be used as a discriminant.

The above plots show just a few examples of the information that can be extracted using generalized shape observables. It is important to emphasize there is much more information

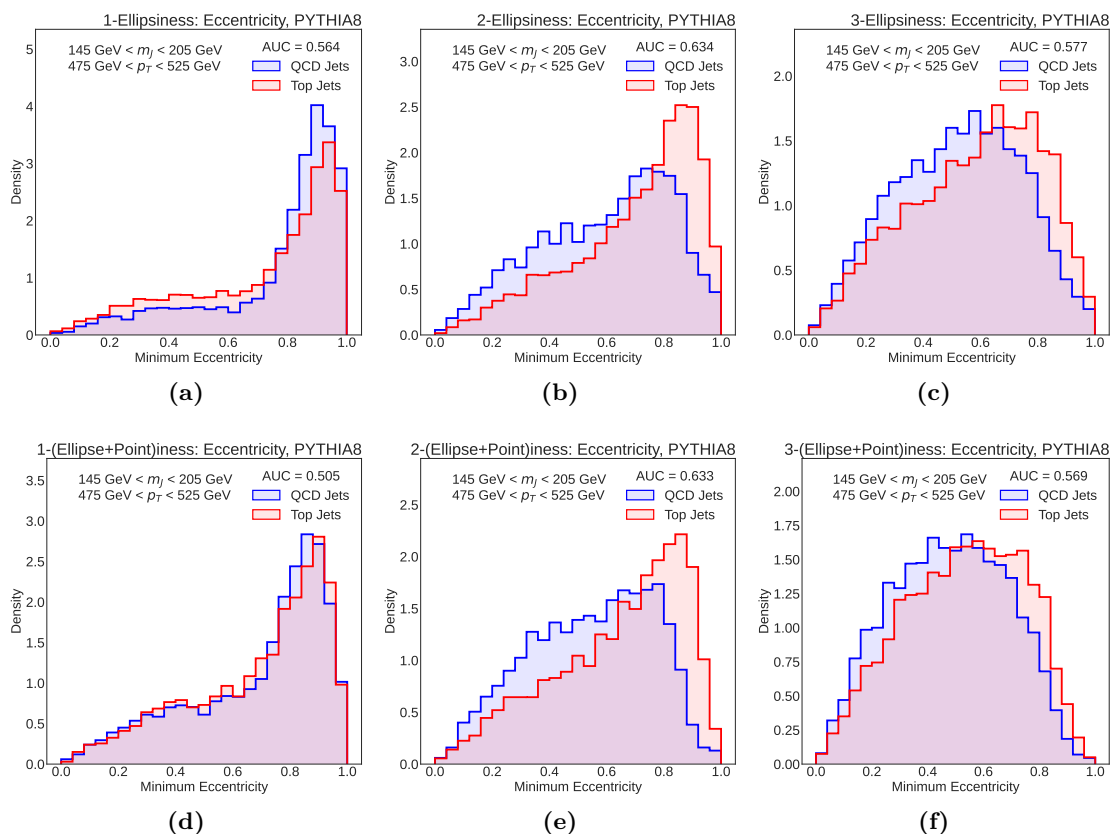


Figure 16. Distributions of the learned minimum eccentricities for (a,d) 1-, (b,e) 2-, and (c,f) 3-ellipsiness (top row) and (ellipse+point)iness (bottom row) of the top (red) and QCD (blue) jet sample.

beyond what we have shown here available for analysis, such as the learned weights of the points versus uniform shapes, the radii of shapes beyond the leading jet, and so on. This is to say nothing of the multivariate information contained in the correlations between the EMDs and shape parameters, both within the same shape and between different shapes. We leave a full quantitative study of the information contained within shapes for potential future work.

5.7 Pileup-mitigating shapes

Finally, we show a use case for shape observable composition by applying our custom shape observables to the task of pileup mitigation. We consider the top jet mass spectrum as an example, which is sharply peaked near $m_{\text{top}}^2 \approx (175 \text{ GeV})^2$. In our modified dataset, jets are contaminated with uniform pileup radiation, which significantly biases and smears the top mass peak. There exist many techniques to “groom” away extra contamination, such as area subtraction and soft drop [93–97], but these techniques often have external hyperparameters characterizing the contamination density (such as z_{cut} in soft drop).

In order to mitigate this bias even when the contamination density is unknown, we propose to build shapes that *factorize* the event into a uniform component, \mathcal{U} , and a

structure component, \mathcal{E} . As described in section 3.3.5, we accomplish this by building observables of the form:

$$\mathcal{O}_i^{\mathcal{I}} = \mathcal{O}_i \oplus \mathcal{I}, \tag{5.5}$$

where \mathcal{O}_i is any shape observable and \mathcal{I} is the jet isotropy.²⁹ We can then use the energy flow associated to \mathcal{O}_i to calculate a “pileup-corrected” mass, discarding the energy flow associated to \mathcal{I} . The weight z_2 associated with \mathcal{I} is then an estimate of the fraction of event energy due to pileup — the contamination density is an extracted observable, *not* an assumed hyperparameter!

For this study, we consider $\mathcal{O}_i^{\mathcal{I}}$, where \mathcal{O}_i is either the 3-subjettiness or the 3-(disk+point)iness, motivated by the 3-prong nature of top jets. We also consider the uncorrected observables, \mathcal{O}_i , without the uniform background for comparison. For each observable, the “shape-jet mass” m_J is calculated as the sum as the (massless) four vectors of particles comprising the shape. This can be numerically approximated by sampling the disk with 200 particles. To calibrate the shape-jet masses, we calculate the shape-jet mass on top jets without the addition of pileup for each observable. In figure 17a, we plot the shape-jet mass corresponding to each observable as evaluated on top jets. Each observable slightly undershoots the top mass peak — the exact discrepancies are given in table 5. These values are used to shift the means of each shape-jet mass curve to the right for the remainder of this study, such that the means give the truth mean top mass when evaluated on uncontaminated jets.

In figure 17b, we show the result of an empirical study using our pileup-mitigating shapes on our pileup-contaminated top jet sample. We also plot, for comparison, the mass corresponding to the *uncorrected* observables \mathcal{O}_i , with no uniform background. These distributions are all calibrated using the values in table 5. We see clearly that the two pileup-mitigating shapes indeed remove a significant amount of bias due to pileup, and bring m_{jet} much closer to m_{top} , suggesting that this is a viable strategy for pileup removal. Note, interestingly, that the distribution for the uncorrected 3-(disk+point)iness jet mass is similar to the uncorrected jet mass — this is because to best approximate the entire event, disks will grow large in an attempt to capture all of the pileup, since the disks themselves are also uniform radiation patterns.

In figure 18a, we show the distribution of extracted pileup energy fraction values, corresponding to the z_2 shape parameter. We see that there is qualitative agreement between the extracted pileup energy fractions and the true one, though the shapes tend to overestimate the pileup density slightly. This is consistent with figure 17b, where the shapes slightly underestimate the top mass without the calibration. In figure 18b, we use z_2 to compute the value of the extracted mass bias $\Delta\hat{m}_{\text{PU}}^2 = m_{\text{jet}}^2 - \hat{m}_{\text{top}}^2(z_2)$, compared to the “true” mass bias obtained by comparing the jet mass before and after contamination. To approximate the resolution of our estimator, we take $\hat{\sigma}^2 = \text{Var} \left[\sqrt{\hat{m}_{\text{PU}}^2} - \sqrt{m_{\text{PU}}^2} \right]$ as the average Gaussian uncertainty. We can see that both pileup-mitigating observables estimate

²⁹Unlike refs. [40, 41], we sample the isotropy using random points on the plane, with resampling each epoch, rather than use a uniform grid, for better statistical coverage of the plane.




Figure 17. The shape-jet mass for several different choices of shape observable, evaluated on (a) top jets without pileup, used to derive the calibration, and (b) top jets with pileup, calibrated using table 5. Two different shape observables are used — the 3-jettiness in blue and 3-(disk+point)iness in red. For each shape observable, the ordinary version is plotted with dark dashed lines while the pileup-mitigating variant is plotted with bright solid lines.

Observable	Bias [GeV]
3-(Disk+Point)iness+Pileup	-18.7
3-Subjettiness+Pileup	-23.1
3-(Disk+Point)iness	-5.5
3-Subjettiness	-12.5

Table 5. The difference between the mean shape-jet mass and the mean of the truth top jet distribution in figure 17a, for each of the four observables under consideration. These are used to calibrate the calculated mass distributions for each shape.

the pileup mass bias correctly on average, and that the 3-(disk+point)iness+pileup has a slightly better resolution than the 3-subjettiness+pileup, which is to be expected as the former has more freedom to better capture features of jets.

6 Conclusions and outlook

In this work, we generalized the notion of event and jet shapes into *shape observables*, which are a wide class of observables that can probe the geometric structure of collider events. We introduced a natural measure-theoretic language for describing events as energy flows, which encodes properties such as IRC safety as inherent topological information. Using this




Figure 18. (a) The learned pileup fractions z_2 for the 3-jettiness in blue and 3-(disk+point)iness in red. The actual fraction of energy due to pileup is in black. (b) The mass corrections to the top mass due to pileup, for the 3-jettiness in blue and 3-(disk+point)iness in red. On the x -axis, the change in jet mass $(\Delta m_{\text{PU}}^2)^{1/2}$ before and after contamination is plotted. On the y -axis, the value of $(\Delta \hat{m}_{\text{PU}}^2)^{1/2}$ as extracted from the corrected top mass is plotted for a small test set. As a proxy for the resolution of the estimate, we take $\hat{\sigma}$ to be the standard deviation of the residuals.

construction, we showed that the Wasserstein metric arises naturally from the requirement of geometric faithfulness. This post-hoc justifies its past use to define observables, IRC safety, and geometry on the space of events [25, 26, 39, 40].

We showcased how to define *arbitrary* shape observables, which can be specified solely by parameterizing the manifold of shapes one wants to fit to. Importantly, we can extract both the “shapiness” value (how much the event looks like the shape) and the “shape parameters” (which shape is the best fit). This is a very intuitive picture — if one wants to ask how “ellipsy” a jet is, for example, all one needs to do is parameterize the space of ellipses to define the ellipsiness observable. The SHAPER framework makes it easy to build new shape observables and evaluate them on events. It leverages the Sinkhorn approximation of the Wasserstein metric to enable fast, parallelizable, and differentiable ϵ -approximations of shape observables.

With the SHAPER framework in place, the natural question arises: *Which shapes should we study?* We introduced a few examples of shapes motivated by jet substructure, including rings, disks, and ellipses. Using the SHAPER prescription, it is easy to modify these shapes to probe additional jet structures, by including δ -functions to capture collinear radiation or uniform backgrounds to capture pileup. Our empirical studies showed how these custom observables might serve as an effective and simple pileup mitigation strategy in the study of top jets, where the amount of pileup does not need to be assumed beforehand. Many further analyses are possible: for example, in heavy quark decays, collinear bremsstrahlung is sup-

pressed (the so-called “dead-cone effect” [90–92]), leading to ring-like or annulus-like, rather than point-like, jets. Moreover, the ability to take derivatives with respect to entire events could allow one to compute observable sensitivities, such as pileup sensitivity [95]. The list of observables we have provided is certainly non-exhaustive, and we hope that the SHAPER framework can be used by the community to explore a wide range of brand new observables.

While SHAPER can be used to analyze *any* shape observable, and thus has broad generality, it is not necessarily the fastest way to compute or approximate any *specific* observable. For example, much faster, exact algorithms for computing N -jettiness and related observables are available in the FASTJET [118] package, which make use of the fact that the double optimization in eq. (3.1) can be simplified. Most of the new observables we have defined appear to be irreducible, though, making their analysis and theoretical treatment complicated. In principle, this class of observables is perturbatively calculable in QCD, at least numerically. We anticipate that these generalized shapes might provide a useful groundwork for further theoretical studies of QCD distributions.

Code and data. The code for the general-use SHAPER framework can be found at <https://github.com/rikab/SHAPER>, along with installation instructions. This directory also contains the analysis code for the top/QCD plots in section 5. In this directory is also a tutorialized example notebook, showing how to use the SHAPER framework.

Our dataset is a modified version of the top/QCD jet benchmark dataset described in refs. [48, 49]. The original version of this dataset can be found at <https://desycloud.desy.de/index.php/s/llbX3zpLhazgPJ6>.

Acknowledgments

Special thanks goes to Samuel Alipour-fard for helping to come up with the acronym SHAPER, and for useful discussions about pileup. We thank Cari Cesarotti and Matthew LeBlanc for useful discussions about event isotropy, and Ouail Kitouni, Niklas Nolte, and Mike Williams for useful discussions on EMD estimation with Kantorovich potentials. Finally, we would like to thank Eugene Wigner of ref. [122] for inspiring the title of section 2, and Mark Kac of ref. [123] for inspiring the title of this paper.

DB, ASD, RG, and JT are supported by the National Science Foundation under Cooperative Agreement PHY-2019786 ([The NSF AI Institute for Artificial Intelligence and Fundamental Interactions](#)). ASD’s research was also funded by the President’s PhD Scholarship at Imperial College London and supported by the EPSRC Centre for Doctoral Training in Mathematics of Random Systems: Analysis, Modelling and Simulation (EP/S023925/1). RG and JT are additionally supported by the U.S. DOE Office of High Energy Physics under grant number DE-SC0012567. AT’s research is supported by NSF DMS 2208392.

A Energy flows as measures

In this appendix, we review aspects of measure theory and topology related to energy flows, as introduced in section 2.2. To begin, we define a *measure*:

Definition 4 Given a set \mathcal{X} and a σ -algebra on \mathcal{X} denoted $\sigma(\mathcal{X})$, a (positive) **measure** \mathcal{E} over $\sigma(\mathcal{X})$ is a function $\mathcal{E} : \sigma(\mathcal{X}) \rightarrow \mathbb{R}$ satisfying:

- **Non-Negativity:** for all $X \in \sigma(\mathcal{X})$, we have $\mathcal{E}(X) \geq 0$.
- **Null Set:** for the empty set, we have $\mathcal{E}(\emptyset) = 0$.
- **Additivity:** for a (countable) collection of disjoint sets $X_i \in \sigma(\mathcal{X})$, we have:

$$\mathcal{E}\left(\bigcup_i X_i\right) = \sum_i \mathcal{E}(X_i). \tag{A.1}$$

If, additionally, we have $\mathcal{E}(\mathcal{X}) = 1$, we say \mathcal{E} is a **probability measure**.

Associating energy flows with measures is very natural: an energy flow answers the question ‘‘How much energy did I detect in the calorimeters located within the subregion X of my detector \mathcal{X} ?’’, and this answer satisfies all of the above properties. We refer to $\mathcal{E}(X)$ as the energy flow, whereas we call the object

$$\mathcal{E}(x) = \sum_i E_i \delta(x - x_i) \tag{A.2}$$

as the *energy flow density* or *measure density*. This language (while different from ref. [25]) is natural, since integrating over eq. (A.2) yields the energy flow:

$$\mathcal{E}(X) = \int_X dx \mathcal{E}(x). \tag{A.3}$$

Under this definition, the energy flow is independent of the choice of coordinates x used on \mathcal{X} , unlike the energy flow density. In particular, the form of energy flow is completely invariant under exactly 0-energy emissions or exactly collinear splittings, and additionally is also invariant under particle relabellings. Throughout this paper, we restrict ourselves to measures that can be written as the integral of a well-behaved associated density, which is the case for all physically realized energy flows. We refer to measures whose density is a finite sum of weighted δ -functions as *atomic measures*.

Measures can be used to formalize what we mean by integration. We begin by defining $\langle \mathcal{E}, \phi \rangle$ as the integral of an integrable function $\phi : \mathcal{X} \rightarrow \mathbb{R}$ against an energy flow \mathcal{E} .³⁰ Assuming coordinates x on \mathcal{X} , we define this quantity as the ordinary integral over the associated energy flow density:³¹

$$\langle \mathcal{E}, \phi \rangle \equiv \int_{\mathcal{X}} dx \mathcal{E}(x) \phi(x). \tag{A.4}$$

In the special case where \mathcal{E} is a probability measure, this quantity is the expectation value $\mathbb{E}_{X \sim \mathcal{E}}[\phi(X)]$ of the random variable $\phi(X)$ sampled over \mathcal{E} . The notation $\langle \mathcal{E}, \phi \rangle$ highlights

³⁰The integral of ϕ over \mathcal{E} is often denoted $\int d\mathcal{E} \phi$. When \mathcal{E} is the Lebesgue measure, uniform over \mathcal{X} , this becomes the ordinary integral.

³¹We will be satisfied here defining integration in terms of the ordinary Lebesgue integration on real numbers, which is always possible if an associated density exists.

the inner product structure of a measure \mathcal{E} “acting” on the space of functions on \mathcal{X} . Notably, this inner product is bilinear — though when considering positive measures, one must be careful to not produce energy flows where the energy is anywhere negative, as these are nonphysical.

Next, we discuss the important topological features of measures. First, we define *weak* convergence* (also sometimes referred to as *convergence in law*) of a sequence of measures:

Definition 5 *A sequence of measures \mathcal{E}_n converges with respect to the weak* topology to a measure \mathcal{E} , which we denote $\mathcal{E}_n \rightarrow \mathcal{E}$, if for any continuous test function ϕ on \mathcal{X} , the sequence of real numbers $\langle \mathcal{E}_n, \phi \rangle$ converges to the real number $\langle \mathcal{E}, \phi \rangle$.*

That is, we say a sequence of measures converges if every single expectation value converges. Associated to this definition of convergence is the *weak* topology*, which is the (weakest) topology for which the map $\langle \cdot, \phi \rangle$ is considered continuous for all continuous ϕ . Note that this definition does *not* make use of any metric on measures; it simply inherits the metric and topological structure of the real numbers. Building off of this, we can now define what it means for any function on energy flows to be continuous:

Definition 6 *A function F on the space of positive measures is **continuous with respect to the weak* topology** if, for every convergent sequence of measures $\mathcal{E}_n \rightarrow \mathcal{E}$, the sequence $F(\mathcal{E}_n)$ converges to $F(\mathcal{E})$.*

Continuity is a very powerful tool for dealing with energy flows. First, it immediately implies that continuous measures can be arbitrarily well approximated by atomic measures with an increasing number of samples. Moreover, if a function F on measures is continuous with respect to the weak* topology, this further implies it is only necessary to specify how F acts on atomic measures, since the action on all other measures is fixed by weak* continuity.

Intuitively, weak continuity captures the idea that a continuous distribution is well approximated by a discrete one with “enough samples” — as the number of sample increases, we approach the continuous one. Note, however, we have not yet defined what it means for two energy flows to be “close” to each other, so we cannot yet describe what happens to functions under small energy flow perturbations. This requires a metric on the space of energy flows that respects the weak* topology, which as argued in section 2 and further justified in appendix B, can be taken to be the Wasserstein metric.

We finally introduce one last piece of notation. For any two energy flows \mathcal{E}_1 and \mathcal{E}_2 , the *joint energy flow* $\mathcal{E} = \mathcal{E}_1 \otimes \mathcal{E}_2$ is the energy flow satisfying for any sets $X, Y \subseteq \mathcal{X}$:

$$\mathcal{E}(X, Y) = \int_{X \times Y} dx dy \mathcal{E}_1(x) \mathcal{E}_2(y). \tag{A.5}$$

The energy flows \mathcal{E}_1 and \mathcal{E}_2 are referred to as the *marginals* of \mathcal{E} . They can be written as $\mathcal{E}_1(X) = \mathcal{E}(X, \mathcal{X})$ and $\mathcal{E}_2(Y) = \mathcal{E}(\mathcal{X}, Y)$, with densities $\mathcal{E}_1(x) = \int_{\mathcal{X}} dy \mathcal{E}(x, y)$ and $\mathcal{E}_2(y) = \int_{\mathcal{X}} dx \mathcal{E}(x, y)$ respectively.

B Constructing the Wasserstein metric

In section 2.1, we proposed to write event and jet shapes $\mathcal{O}_{\mathcal{M}}$ using some universal loss function \mathcal{L} on energy flows, and in section 2.3, we showed that IRC safety implies that must

\mathcal{L} must be weakly continuous. Then, in section 2.4, we claimed that the condition of faithfully lifting the ground metric means that \mathcal{L} must be the Wasserstein metric, and showed examples of how similar metrics do not satisfy this property. In this appendix, we show how the Wasserstein metric arrives constructively. First, in appendix B.1, we briefly review the properties we use to construct \mathcal{L} , before outlining the construction in appendix B.2.

B.1 Shaping up the loss

To start, we list the properties we would like the universal loss function \mathcal{L} to satisfy. First, we would like \mathcal{L} to be a proper metric on the space of energy flows, which implies the following usual properties of metrics:

1. **Finiteness:** we require that \mathcal{L} is finite, even when evaluated on atomic measures. This immediately rules out the KL divergence [124] and its variants, including log-likelihoods, used in ref. [51] to define “fuzzy jet” observables. These functions are only finite on continuous measures with support almost everywhere.
2. **Positivity and Closure:** we require that $\mathcal{L}(\mathcal{E}, \mathcal{E}') \geq 0$. Moreover, we require that $\mathcal{L}(\mathcal{E}, \mathcal{E}') = 0$ if and only if $\mathcal{E} = \mathcal{E}'$. This captures the notion of the “optimal shape” as discussed in section 2.1.
3. **Symmetry:** we require that $\mathcal{L}(\mathcal{E}, \mathcal{E}') = \mathcal{L}(\mathcal{E}', \mathcal{E})$. This is to say, if the energy flow \mathcal{E} looks like \mathcal{E}' , then the energy flow \mathcal{E}' looks like \mathcal{E} as well.

We will not specifically require the triangle inequality, as we never use it throughout this work. Of course, the Wasserstein metric does indeed satisfy the triangle inequality when the appropriate powers of $1/\beta$ are included, making it a proper metric.

Not only is \mathcal{L} a metric, but it must be continuous to the weak* topology to be IRC safe, as established in section 2.3:

4. **Weak Continuity/IRC Safety:** we require that \mathcal{L} is weakly continuous in both of its arguments. This not only allows for continuous energy flows to be arbitrarily well approximated by atomic energy flows, but heavily constrains the energy- and position- dependence of \mathcal{L} .

Finally, as discussed in section 2.4, we demand that \mathcal{L} *faithfully lifts the ground metric*. This allows the space of events to inherit the geometry of the ground metric space with no distortions or warping:

5. **Faithfully Lifts the Ground Metric:** we require that if \mathcal{E} and \mathcal{E}' are atomic measures with a single normalized particle at positions x and y respectively, then $\mathcal{L}(\mathcal{E}, \mathcal{E}')$ is proportional to $d(x, y)^\beta$, where β is some fixed positive power.

Moreover, we require that \mathcal{L} *faithfully lifts the ground metric*. For any measure \mathcal{E} , we can “translate” the measure by t to \mathcal{E}_t , defined by the density $\mathcal{E}(x - t)$. We then require that $\mathcal{L}(\mathcal{E}, \mathcal{E}_t)$ is proportional to $d(0, t)^\beta$.

If the ground metric is translationally invariant, which is the case for Euclidean metrics, property 5 implies that \mathcal{L} must also be translationally invariant — if both \mathcal{E} and \mathcal{E}' are translated by a vector t , the metric distance between them is unchanged.

B.2 Why Wasserstein?

Having defined the properties we would like our universal loss function to satisfy, we can now finally construct $\mathcal{L}(\mathcal{E}, \mathcal{E}')$.

To begin, weak continuity allows us to drastically simplify the problem. Continuous energy flows may be approximated arbitrarily well by atomic energy flows, and thus it suffices to build our function only on atomic energy flows. This allows us to argue that \mathcal{L} can only depend on single powers of the distance $d(x, y)^\beta$. In particular, lifting the ground metric (though not necessarily faithfully, yet) forbids terms of the form $d(x, y)^\beta d(x', y')^\beta$ or other higher-order distance correlations, since in the single-particle case this would result in losses of the form $d^{\beta_1} + d^{\beta_2} + \dots$ with several differing exponents. This implies the following form for \mathcal{L} when evaluated on uniform atomic flows $\mathcal{E} \sim \sum_i \delta_{x_i}$ and $\mathcal{E}' \sim \sum_j \delta_{y_j}$:

$$\begin{aligned} \mathcal{L}(\mathcal{E}, \mathcal{E}') = & \Pi \left(\sum_{i,j} \pi_{ij} d(x_i, y_j)^\beta \right) \\ & + K \left(\sum_{i_1, i_2} k_{i_1 i_2} d(x_{i_1}, x_{i_2})^\beta \right) + K \left(\sum_{j_1, j_2} k_{j_1 j_2} d(y_{j_1}, y_{j_2})^\beta \right) \\ & + F \left(\sum_i f_i d(x_i, g(x_i))^\beta \right) + F \left(\sum_j f_j d(\tilde{g}(y_j), y_j)^\beta \right), \end{aligned} \quad (\text{B.1})$$

where Π , K , and F are universal functions, π , k , and f are coefficients that implicitly depend on the two measures, and g is a function from the domain of \mathcal{E} to the domain of \mathcal{E}' (and vice-versa for \tilde{g}). Note that this form is symmetric under swapping \mathcal{E} and \mathcal{E}' . Without loss of generality, $k_{i'i'}$ can be chosen to be a symmetric traceless matrix.

Still considering the single particle case, lifting the ground metric also implies that the functions Π and F must be linear functions (though the single particle case provides no information about K), as to reproduce the desired $d(x, y)^\beta$ behavior. Moreover, to still maintain only terms proportional to d^β even in the multiparticle case, K must also be linear. This linearity means that we can combine the arguments of each of the Π , K , and F functions.

Recall that energy flows are additive objects, and the energy flows defined by the densities $\mathcal{E} = (E_1 + E_2)\delta_x$ and $\mathcal{E} = E_1\delta_x + E_2\delta_x$ must be identical. This implies that each term of the terms in eq. (B.1) must be individually linear in the energy weights of \mathcal{E} and \mathcal{E}' ,

respectively. This implies the following constraints, up to an overall proportionality factor:

$$\sum_i \pi_{ij} = E'_j, \tag{B.2}$$

$$\sum_j \pi_{ij} = E_i, \tag{B.3}$$

$$\sum_{i_1} k_{i_1 i_2} = E_{i_2}, \tag{B.4}$$

$$\sum_{i_2} k_{i_1 i_2} = E_{i_1}, \tag{B.4}$$

$$\sum_{j_1} k_{j_1 j_2} = E'_{j_2}, \tag{B.5}$$

$$\sum_{j_2} k_{j_1 j_2} = E'_{j_1}, \tag{B.5}$$

$$f_i = E_i, \tag{B.6}$$

$$f_j = E'_j \tag{B.6}$$

Now, consider the two particle case, as in section 2.4, and define the energy flows for vectors a and t :

$$\mathcal{E}(x) \sim \frac{1}{2}\delta_0 + \frac{1}{2}\delta_a, \tag{B.7}$$

$$\mathcal{E}'(x) \sim \frac{1}{2}\delta_t + \frac{1}{2}\delta_{a+t}. \tag{B.8}$$

When $t \rightarrow 0$, the two energy flows are the same, and when $a \rightarrow 0$, the energy flows reduce to the single particle case. Evaluating eq. (B.1), we obtain:

$$\begin{aligned} \mathcal{L}(\mathcal{E}, \mathcal{E}') &= \Pi \left(\pi_{1_x 1_y} d(0, t)^\beta + \pi_{1_x 2_y} d(0, a+t)^\beta + \pi_{1_x 1_y} d(a, t)^\beta + \pi_{2_x 2_y} d(a, a+t)^\beta \right) \\ &+ K \left(2k_{1_x 2_x} d(0, a)^\beta + 2k_{1_y 2_y} d(t, a+t)^\beta \right) \\ &+ \frac{1}{2}F \left(d(0, g(0))^\beta + d(a, g(a))^\beta + d(\tilde{g}(t), t)^\beta + d(\tilde{g}(a+t), a+t)^\beta \right). \end{aligned} \tag{B.9}$$

In the limit $t \rightarrow 0$, this expression must evaluate to zero by closure, for any choice of a . This is possible if the off-diagonal components of π and k either cancel or are individually zero. Note that the function g must select a particle y_j given x_i (and vice-versa for \tilde{g}). Importantly, however, it must do this in a label-independent way, since the labels i and j are completely arbitrary. The only way to do so is to select y_j based on its distance from x_i — in principle, we can choose g to select the furthest particle, or the closest particle, or even the particle closest to the average distance amongst all particles. However, the only choice consistent with closure, as can be seen in eq. (B.9), is to choose the *closest* particle y_j to the given x_j . This allows us to write the third line of eq. (B.1) as:

$$\mathcal{L}(\mathcal{E}, \mathcal{E}') = [\dots] + F \left(\sum_i E_i \min_{y_j \in \text{Supp}(\mathcal{E}')} d(x_i, y_j)^\beta + \sum_j E'_j \min_{x_i \in \text{Supp}(\mathcal{E})} d(x_i, y_j)^\beta \right). \tag{B.10}$$

Eq. (B.10) is exactly the form of a Chamfer distance from eq. (2.11)! The Chamfer distance is the average closest distance from points in \mathcal{E} to \mathcal{E}' , plus the average closest distance from points in \mathcal{E}' to \mathcal{E} . Following the explicit counter-example given in eq. (2.15), the Chamfer distance does *not* lift the ground metric *faithfully*, and therefore, F must be zero.

Returning to eq. (B.9), we now want to show that $K = 0$. Since k is traceless by assumption, the constraints eqs. (B.4) and (B.5) imply that $k_{1_x 2_x} = E_1 E_2$ and $k_{1_y 2_y} =$

$E'_1 E'_2$. Assuming that $K \neq 0$, closure plus eqs. (B.2) and (B.3) allows us to solve for π_{ij} . From this, we deduce that π_{ij} reduces to $E_i E'_j$ as $t \rightarrow 0$ (with proportionality constant $\Pi = -2K$). Furthermore, π_{ij} has to be constant in t , since otherwise \mathcal{L} would depend on higher powers of the metric and therefore fail to lift the ground metric. Given this, eq. (B.1) reduces to:

$$\mathcal{L}(\mathcal{E}, \mathcal{E}') = K \left(\sum_{i_1, i_2} E_{i_1} E_{i_2} d(x_{i_1}, x_{i_2})^\beta + \sum_{j_1, j_2} E'_{j_1} E'_{j_2} d(y_{j_1}, y_{j_2})^\beta - 2 \sum_{ij} E_i E'_j d(x_i, y_j)^\beta \right) \tag{B.11}$$

However, this implies that \mathcal{L} is exactly the maximum mean discrepancy (MMD) of eq. (2.10)! The MMD can be thought of as the average potential energy of a system of springs with potential $V(r) \sim r^\beta$ connecting particles between \mathcal{E} and \mathcal{E}' , minus the self-energy of springs connecting particles within \mathcal{E} and within \mathcal{E}' . The explicit counter-example given in eq. (2.14) shows that this cannot be faithful, *except in the special case of $\beta = 2$* , which is not sufficient for our purposes.³² It follows that K must also be zero, and we have significantly constrained the form of our metric.

It now remains only to find the form of π . To do this, we now consider an energy flow \mathcal{E} , consisting of N particles each with energy $\frac{1}{N}$. By closure, we again must have $\mathcal{L}(\mathcal{E}, \mathcal{E}) = 0$, which occurs when π_{ij} is $\frac{1}{N^2}$ times the $N \times N$ identity matrix.

Next, consider \mathcal{E}' , an exact copy of \mathcal{E} , except the particles have been re-indexed, sending the original particle x_i to $x_{j=\sigma(i)}$. As measures are invariant under reordering, \mathcal{L} must still be zero, though π will be (proportional to) some $N \times N$ permutation matrix. We can write this as:

$$\begin{aligned} \mathcal{L}(\mathcal{E}, \mathcal{E}') &= \frac{1}{N} \sum_i d(x_i, y_{\sigma(i)})^\beta, \\ &= \sum_{ij} \pi_{ij} d(x_i, y_j)^\beta, \end{aligned} \tag{B.12}$$

This will only be zero if π_{ij} is the *correct* permutation matrix that undoes the index shuffling. If one instead guesses that π is a different permutation matrix, the point $y_{\sigma(i)}$ will not lie on top of the point x_i , leading to $\mathcal{L} > 0$.

Thus, even if we did not know how exactly the particle labels on \mathcal{E}' were shuffled, all we would have to do to find the correct π is to search through all possible permutation matrices, and take the one that gives the minimum answer:

$$\begin{aligned} \mathcal{L}(\mathcal{E}, \mathcal{E}') &= \min_{\sigma: [1, N] \rightarrow [1, N]} \frac{1}{N} \sum_i d(x_i, y_{\sigma(i)})^\beta \\ &= \min_{\pi_{ij} \in S_N} \frac{1}{N^2} \sum_{ij} \pi_{ij} d(x_i, y_j)^\beta, \end{aligned} \tag{B.13}$$

³²Generically, physical systems with potentials $V(r) \sim r^\beta$ experience screening, most notably electrostatic screening when $\beta = -1$. However, screening does not occur for ideal springs, i.e. $\beta = 2$. Screening and (un)faithfulness are related — as \mathcal{E} and \mathcal{E}' move relative to each other, the potential energy between them will not necessarily scale with r^β in extended systems due to screening.

where S_N is the set of all $N \times N$ permutation matrices. Moreover, the index-shuffling trick allows us to see that eq. (B.13) faithfully lifts the ground metric. If \mathcal{E}' is translated by a vector t , then π should still be the identity matrix before shuffling so that our result is proportional to $d(0, t)^\beta$. If we guess the wrong permutation matrix, while it may be possible for some distances to close (that is, $d(x_i, y_{\sigma(i)})$ to be less than $d(0, t)$), this will be made up for by other particles being further apart than $d(0, t)$, and so \mathcal{L} will be larger. For values of $\beta \geq 1$, this last fact follows from the triangle inequality. This can be seen explicitly in the two-particle case: if $|t| \gg |a|$, it follows that:

$$d(0, t)^\beta \leq \frac{1}{2}d(0, t - a)^\beta + \frac{1}{2}d(0, t + a)^\beta, \tag{B.14}$$

for $\beta \geq 1$. Therefore eq. (B.13) is a faithful metric for $\beta \geq 1$.

The problem in eq. (B.13) is referred to as the *combinatorial Monge problem*, which is the precursor to the Wasserstein metric. We can think of eq. (B.13) as finding the optimal map π to “transport” the points x to y , minimizing the total distance ^{β} needed to travel. These arguments are enough to fully fix the form of \mathcal{L} — even if \mathcal{E} and \mathcal{E}' are different energy flows, as long as they both have N particles with uniform weights, the minimization over permutation matrices is valid.

To finish our construction, we use the additivity of measures and the weak* topology to generalize eq. (B.13) to any energy flow. Since measures are additive, any atomic energy flow with *rational* energy weights can be thought of particles of a base denomination weight stacked on top of each other — for example, the two-particle energy flow $\mathcal{E} \sim 0.381\delta_{x_1} + 0.619\delta_{x_2}$ is exactly equal to the energy flow given by 381 particles at site x_1 and 619 particles at site x_2 , each with energy 0.001. One can then solve the combinatorial Monge problem using this uniform energy flow, and “collapse” the corresponding redundant subspace in π to produce the optimal transport map. In the case that energy weights are irrational, weak continuity ensures that these may be constructed as limits of rationally-weighted energy flows, so this is not an issue. Similarly, any continuous energy flow can be written this way.

Thus, the loss eq. (B.13) is our universal loss function! Accounting for arbitrary weights and particle numbers, as discussed above, this becomes the well-known Wasserstein metric. The β -Wasserstein metric, or “Energy/Earth Mover’s Distance” (EMD), is the metric on the space of measures that is positive, closed, metrizes the weak convergence, and faithfully lifts the ground metric. Repeating eq. (2.9) for convenience, the EMD between two measures \mathcal{E} and \mathcal{E}' is given by:

$$\begin{aligned} \text{EMD}^{(\beta, R)}(\mathcal{E}, \mathcal{E}') &= \min_{\pi \in \mathcal{M}(\mathcal{X} \times \mathcal{X})} \left[\frac{1}{\beta R^\beta} \langle \pi, d(x, y)^\beta \rangle \right] + |\Delta E_{\text{tot}}|, \\ \pi(\mathcal{X}, Y) &\leq \mathcal{E}'(Y), \quad \pi(X, \mathcal{X}) \leq \mathcal{E}(X), \quad \pi(\mathcal{X}, \mathcal{X}) = \min(E_{\text{tot}}, E'_{\text{tot}}). \end{aligned} \tag{B.15}$$

The parameter $R > 0$ sets a distance scale for the EMD. The additional energy difference term, $|\Delta E_{\text{tot}}| = |E_{\text{tot}} - E'_{\text{tot}}|$, contributes whenever the two energy flows do not have the same total energy. Our argument in this section does not fix this term uniquely, and in fact there are many different approaches to unbalanced metrics [39, 125, 126]. In the common case that both \mathcal{E} and \mathcal{E}' are atomic measures with M and N particles with energies E_i and

E'_j respectively, the EMD takes the form:

$$\text{EMD}^{(\beta,R)}(\mathcal{E}, \mathcal{E}') = \min_{\pi_{ij} \geq 0} \left[\frac{1}{\beta R^\beta} \sum_{i=1}^M \sum_{j=1}^N \pi_{ij} d_{ij}^\beta \right] + |\Delta E_{\text{tot}}|,$$

$$\sum_{i=1}^M \pi_{ij} \leq E'_j, \quad \sum_{j=1}^N \pi_{ij} \leq E_i, \quad \sum_{j=1}^N \pi_{ij} = \min(E_{\text{tot}}, E'_{\text{tot}}). \quad (\text{B.16})$$

Open Access. This article is distributed under the terms of the Creative Commons Attribution License ([CC-BY 4.0](https://creativecommons.org/licenses/by/4.0/)), which permits any use, distribution and reproduction in any medium, provided the original author(s) and source are credited. SCOAP³ supports the goals of the International Year of Basic Sciences for Sustainable Development.

References

- [1] E. Farhi, *A QCD test for jets*, *Phys. Rev. Lett.* **39** (1977) 1587 [[INSPIRE](#)].
- [2] D.P. Barber et al., *Tests of quantum chromodynamics and a direct measurement of the strong coupling constant α_S at $\sqrt{s} = 30$ GeV*, *Phys. Lett. B* **89** (1979) 139 [[INSPIRE](#)].
- [3] M. Dasgupta and G.P. Salam, *Event shapes in e^+e^- annihilation and deep inelastic scattering*, *J. Phys. G* **30** (2004) R143 [[hep-ph/0312283](#)] [[INSPIRE](#)].
- [4] G. Dissertori et al., *$e^+e^- \rightarrow 3$ jets and event shapes at NNLO*, *Nucl. Phys. B Proc. Suppl.* **183** (2008) 2 [[arXiv:0806.4601](#)] [[INSPIRE](#)].
- [5] L.G. Almeida et al., *Substructure of high- p_T jets at the LHC*, *Phys. Rev. D* **79** (2009) 074017 [[arXiv:0807.0234](#)] [[INSPIRE](#)].
- [6] G. Gur-Ari, M. Papucci and G. Perez, *Classification of energy flow observables in narrow jets*, [arXiv:1101.2905](#) [[INSPIRE](#)].
- [7] S. Brandt, C. Peyrou, R. Sosnowski and A. Wroblewski, *The principal axis of jets. An attempt to analyze high-energy collisions as two-body processes*, *Phys. Lett.* **12** (1964) 57 [[INSPIRE](#)].
- [8] A. De Rujula, J.R. Ellis, E.G. Floratos and M.K. Gaillard, *QCD predictions for hadronic final states in e^+e^- annihilation*, *Nucl. Phys. B* **138** (1978) 387 [[INSPIRE](#)].
- [9] C.F. Berger, T. Kucs and G.F. Sterman, *Event shape/energy flow correlations*, *Phys. Rev. D* **68** (2003) 014012 [[hep-ph/0303051](#)] [[INSPIRE](#)].
- [10] C.F. Berger and L. Magnea, *Scaling of power corrections for angularities from dressed gluon exponentiation*, *Phys. Rev. D* **70** (2004) 094010 [[hep-ph/0407024](#)] [[INSPIRE](#)].
- [11] T. Becher and M.D. Schwartz, *A precise determination of α_s from LEP thrust data using effective field theory*, *JHEP* **07** (2008) 034 [[arXiv:0803.0342](#)] [[INSPIRE](#)].
- [12] A. Banfi, H. McAslan, P.F. Monni and G. Zanderighi, *A general method for the resummation of event-shape distributions in e^+e^- annihilation*, *JHEP* **05** (2015) 102 [[arXiv:1412.2126](#)] [[INSPIRE](#)].
- [13] TASSO collaboration, *Jet production and fragmentation in e^+e^- annihilation at 12 GeV to 43 GeV*, *Z. Phys. C* **22** (1984) 307 [[INSPIRE](#)].

- [14] MARK-II collaboration, *First measurements of hadronic decays of the Z boson*, *Phys. Rev. Lett.* **63** (1989) 1558 [INSPIRE].
- [15] AMY collaboration, *Multi-hadron event properties in e^+e^- annihilation at $\sqrt{s} = 52$ GeV to 57 GeV*, *Phys. Rev. D* **41** (1990) 2675 [INSPIRE].
- [16] ALEPH collaboration, *Measurement of α_s from scaling violations in fragmentation functions in e^+e^- annihilation*, *Phys. Lett. B* **357** (1995) 487 [INSPIRE].
- [17] L3 collaboration, *Determination of α_s from hadronic event shapes measured on the Z^0 resonance*, *Phys. Lett. B* **284** (1992) 471 [INSPIRE].
- [18] TASSO collaboration, *Global jet properties at 14 GeV to 44 GeV center-of-mass energy in e^+e^- annihilation*, *Z. Phys. C* **47** (1990) 187 [INSPIRE].
- [19] SLD collaboration, *Measurement of $\alpha_s(M_Z^2)$ from hadronic event observables at the Z^0 resonance*, *Phys. Rev. D* **51** (1995) 962 [hep-ex/9501003] [INSPIRE].
- [20] ALEPH collaboration, *Studies of QCD at e^+e^- centre-of-mass energies between 91 GeV and 209 GeV*, *Eur. Phys. J. C* **35** (2004) 457 [INSPIRE].
- [21] DELPHI collaboration, *A study of the energy evolution of event shape distributions and their means with the DELPHI detector at LEP*, *Eur. Phys. J. C* **29** (2003) 285 [hep-ex/0307048] [INSPIRE].
- [22] L3 collaboration, *Studies of hadronic event structure in e^+e^- annihilation from 30 GeV to 209 GeV with the L3 detector*, *Phys. Rept.* **399** (2004) 71 [hep-ex/0406049] [INSPIRE].
- [23] OPAL collaboration, *Measurement of event shape distributions and moments in $e^+e^- \rightarrow$ hadrons at 91–209 GeV and a determination of α_s* , *Eur. Phys. J. C* **40** (2005) 287 [hep-ex/0503051] [INSPIRE].
- [24] A. Abdesselam et al., *Boosted objects: a probe of beyond the standard model physics*, *Eur. Phys. J. C* **71** (2011) 1661 [arXiv:1012.5412] [INSPIRE].
- [25] P.T. Komiske, E.M. Metodiev and J. Thaler, *The hidden geometry of particle collisions*, *JHEP* **07** (2020) 006 [arXiv:2004.04159] [INSPIRE].
- [26] P.T. Komiske, E.M. Metodiev and J. Thaler, *Metric space of collider events*, *Phys. Rev. Lett.* **123** (2019) 041801 [arXiv:1902.02346] [INSPIRE].
- [27] S. Peleg, M. Werman and H. Rom, *A unified approach to the change of resolution: space and gray-level*, *IEEE Trans. Pattern Anal. Machine Intell.* **11** (1989) 739.
- [28] Y. Rubner, C. Tomasi and L.J. Guibas, *A metric for distributions with applications to image databases*, in *Proceedings of the sixth international conference on computer vision, ICCV '98*, IEEE Computer Society, U.S.A. (1998), p. 59.
- [29] Y. Rubner, C. Tomasi and L.J. Guibas, *The earth mover's distance as a metric for image retrieval*, *Int. J. Comput. Vision* **40** (2004) 99.
- [30] O. Pele and M. Werman, *A linear time histogram metric for improved SIFT matching*, in the proceedings of the *ECCV*, (2008) [DOI:10.1007/978-3-540-88690-7_37].
- [31] O. Pele and B. Taskar, *The tangent earth mover's distance*, in *Lecture notes in computer science*, Springer, Berlin, Germany (2013), p. 397 [DOI:10.1007/978-3-642-40020-9_43].
- [32] L.N. Wasserstein, *Markov processes over denumerable products of spaces describing large systems of automata*, *Probl. Info. Transm.* **5** (1969) 47.

- [33] R.L. Dobrushin, *Prescribing a system of random variables by conditional distributions*, *Theor. Probabil. Appl.* **15** (1970) 458.
- [34] P.T. Komiske et al., *Exploring the space of jets with CMS open data*, *Phys. Rev. D* **101** (2020) 034009 [[arXiv:1908.08542](#)] [[INSPIRE](#)].
- [35] J.H. Collins, *An exploration of learnt representations of W jets*, [arXiv:2109.10919](#) [[INSPIRE](#)].
- [36] S.E. Park, P. Harris and B. Ostdiek, *Neural embedding: learning the embedding of the manifold of physics data*, [arXiv:2208.05484](#) [[INSPIRE](#)].
- [37] M. Crispim Romão et al., *Use of a generalized energy Mover's distance in the search for rare phenomena at colliders*, *Eur. Phys. J. C* **81** (2021) 192 [[arXiv:2004.09360](#)] [[INSPIRE](#)].
- [38] T. Cai, J. Cheng, N. Craig and K. Craig, *Linearized optimal transport for collider events*, *Phys. Rev. D* **102** (2020) 116019 [[arXiv:2008.08604](#)] [[INSPIRE](#)].
- [39] T. Cai, J. Cheng, K. Craig and N. Craig, *Which metric on the space of collider events?*, *Phys. Rev. D* **105** (2022) 076003 [[arXiv:2111.03670](#)] [[INSPIRE](#)].
- [40] C. Cesarotti and J. Thaler, *A robust measure of event isotropy at colliders*, *JHEP* **08** (2020) 084 [[arXiv:2004.06125](#)] [[INSPIRE](#)].
- [41] C. Cesarotti, M. Reece and M.J. Strassler, *The efficacy of event isotropy as an event shape observable*, *JHEP* **07** (2021) 215 [[arXiv:2011.06599](#)] [[INSPIRE](#)].
- [42] ATLAS collaboration, *Measurements of multijet event isotropies using optimal transport with the ATLAS detector*, [arXiv:2305.16930](#) [[INSPIRE](#)].
- [43] R. Sinkhorn, *A relationship between arbitrary positive matrices and stochastic matrices*, *Can. J. Math.* **18** (1966) 303.
- [44] M. Cuturi, *Sinkhorn distances: lightspeed computation of optimal transport*, *Adv. Neural Inf. Process. Syst.* **26** (2013) 2292.
- [45] C. Clason, D.A. Lorenz, H. Mahler and B. Wirth, *Entropic regularization of continuous optimal transport problems*, *J. Math. Anal. Appl.* **494** (2021) 124432.
- [46] J. Feydy et al., *Interpolating between optimal transport and MMD using Sinkhorn divergences*, in the proceedings of the 22nd international conference on artificial intelligence and statistics, (2019), p. 2681.
- [47] O. Kitouni, N. Nolte and M. Williams, *Finding NEEMo: geometric fitting using Neural Estimation of the Energy Mover's distance*, [arXiv:2209.15624](#) [[INSPIRE](#)].
- [48] A. Butter, G. Kasieczka, T. Plehn and M. Russell, *Deep-learned top tagging with a Lorentz layer*, *SciPost Phys.* **5** (2018) 028 [[arXiv:1707.08966](#)] [[INSPIRE](#)].
- [49] A. Butter et al., *The machine learning landscape of top taggers*, *SciPost Phys.* **7** (2019) 014 [[arXiv:1902.09914](#)] [[INSPIRE](#)].
- [50] D. Krohn, J. Thaler and L.-T. Wang, *Jets with variable R* , *JHEP* **06** (2009) 059 [[arXiv:0903.0392](#)] [[INSPIRE](#)].
- [51] L. Mackey, B. Nachman, A. Schwartzman and C. Stansbury, *Fuzzy jets*, *JHEP* **06** (2016) 010 [[arXiv:1509.02216](#)] [[INSPIRE](#)].
- [52] B. Mukhopadhyaya, T. Samui and R.K. Singh, *Dynamic radius jet clustering algorithm*, *JHEP* **04** (2023) 019 [[arXiv:2301.13074](#)] [[INSPIRE](#)].

- [53] A.J. Larkoski, D. Rathjens, J. Veatch and J.W. Walker, *Jet SIFT-ing: a new scale-invariant jet clustering algorithm for the substructure era*, [arXiv:2302.08609](#) [INSPIRE].
- [54] I.W. Stewart et al., *XCone: N -jettiness as an exclusive cone jet algorithm*, *JHEP* **11** (2015) 072 [[arXiv:1508.01516](#)] [INSPIRE].
- [55] H. Georgi and M. Machacek, *A simple QCD prediction of jet structure in e^+e^- annihilation*, *Phys. Rev. Lett.* **39** (1977) 1237 [INSPIRE].
- [56] A.J. Larkoski, D. Neill and J. Thaler, *Jet shapes with the broadening axis*, *JHEP* **04** (2014) 017 [[arXiv:1401.2158](#)] [INSPIRE].
- [57] I.W. Stewart, F.J. Tackmann and W.J. Waalewijn, *N -jettiness: an inclusive event shape to veto jets*, *Phys. Rev. Lett.* **105** (2010) 092002 [[arXiv:1004.2489](#)] [INSPIRE].
- [58] A. Banfi, G.P. Salam and G. Zanderighi, *Resummed event shapes at hadron-hadron colliders*, *JHEP* **08** (2004) 062 [[hep-ph/0407287](#)] [INSPIRE].
- [59] A. Banfi, G.P. Salam and G. Zanderighi, *Phenomenology of event shapes at hadron colliders*, *JHEP* **06** (2010) 038 [[arXiv:1001.4082](#)] [INSPIRE].
- [60] S. Catani, Y.L. Dokshitzer, M.H. Seymour and B.R. Webber, *Longitudinally invariant K_t clustering algorithms for hadron hadron collisions*, *Nucl. Phys. B* **406** (1993) 187 [INSPIRE].
- [61] S.D. Ellis and D.E. Soper, *Successive combination jet algorithm for hadron collisions*, *Phys. Rev. D* **48** (1993) 3160 [[hep-ph/9305266](#)] [INSPIRE].
- [62] Y.L. Dokshitzer, G.D. Leder, S. Moretti and B.R. Webber, *Better jet clustering algorithms*, *JHEP* **08** (1997) 001 [[hep-ph/9707323](#)] [INSPIRE].
- [63] M. Wobisch and T. Wengler, *Hadronization corrections to jet cross-sections in deep inelastic scattering*, in the proceedings of the *Workshop on Monte Carlo generators for HERA physics (plenary starting meeting)*, (1998), p. 270 [[hep-ph/9907280](#)] [INSPIRE].
- [64] M. Cacciari, G.P. Salam and G. Soyez, *The anti- k_t jet clustering algorithm*, *JHEP* **04** (2008) 063 [[arXiv:0802.1189](#)] [INSPIRE].
- [65] S.D. Ellis, Z. Kunszt and D.E. Soper, *Jets at hadron colliders at order α_s^3 : a look inside*, *Phys. Rev. Lett.* **69** (1992) 3615 [[hep-ph/9208249](#)] [INSPIRE].
- [66] CDF collaboration, *A measurement of jet shapes in $p\bar{p}$ collisions at $\sqrt{s} = 1.8$ TeV*, *Phys. Rev. Lett.* **70** (1993) 713 [INSPIRE].
- [67] S.D. Ellis et al., *Jet shapes and jet algorithms in SCET*, *JHEP* **11** (2010) 101 [[arXiv:1001.0014](#)] [INSPIRE].
- [68] J. Thaler and K. Van Tilburg, *Identifying boosted objects with N -subjettiness*, *JHEP* **03** (2011) 015 [[arXiv:1011.2268](#)] [INSPIRE].
- [69] M. Dasgupta and G.P. Salam, *Resummation of nonglobal QCD observables*, *Phys. Lett. B* **512** (2001) 323 [[hep-ph/0104277](#)] [INSPIRE].
- [70] A. Banfi, M. Dasgupta, K. Khelifa-Kerfa and S. Marzani, *Non-global logarithms and jet algorithms in high- p_T jet shapes*, *JHEP* **08** (2010) 064 [[arXiv:1004.3483](#)] [INSPIRE].
- [71] F.V. Tkachov, *Measuring multi-jet structure of hadronic energy flow or what is a jet?*, *Int. J. Mod. Phys. A* **12** (1997) 5411 [[hep-ph/9601308](#)] [INSPIRE].
- [72] N.A. Sveshnikov and F.V. Tkachov, *Jets and quantum field theory*, *Phys. Lett. B* **382** (1996) 403 [[hep-ph/9512370](#)] [INSPIRE].

- [73] G.P. Korchemsky, G. Oderda and G.F. Sterman, *Power corrections and nonlocal operators*, *AIP Conf. Proc.* **407** (1997) 988 [[hep-ph/9708346](#)] [[INSPIRE](#)].
- [74] C.L. Basham, L.S. Brown, S.D. Ellis and S.T. Love, *Energy correlations in electron-positron annihilation in quantum chromodynamics: asymptotically free perturbation theory*, *Phys. Rev. D* **19** (1979) 2018 [[INSPIRE](#)].
- [75] P.S. Chervor and N.A. Sveshnikov, *Jet observables and energy momentum tensor*, in the proceedings of the 12th international workshop on high-energy physics and quantum field theory (QFTHEP '97), (1997), p. 402 [[hep-ph/9710349](#)] [[INSPIRE](#)].
- [76] F.V. Tkachov, *A theory of jet definition*, *Int. J. Mod. Phys. A* **17** (2002) 2783 [[hep-ph/9901444](#)] [[INSPIRE](#)].
- [77] G.P. Korchemsky and G.F. Sterman, *Power corrections to event shapes and factorization*, *Nucl. Phys. B* **555** (1999) 335 [[hep-ph/9902341](#)] [[INSPIRE](#)].
- [78] A.V. Belitsky, G.P. Korchemsky and G.F. Sterman, *Energy flow in QCD and event shape functions*, *Phys. Lett. B* **515** (2001) 297 [[hep-ph/0106308](#)] [[INSPIRE](#)].
- [79] C.F. Berger et al., *Snowmass 2001: jet energy flow project*, *eConf C010630* (2001) P512 [[hep-ph/0202207](#)] [[INSPIRE](#)].
- [80] C.W. Bauer, S.P. Fleming, C. Lee and G.F. Sterman, *Factorization of e^+e^- event shape distributions with hadronic final states in soft collinear effective theory*, *Phys. Rev. D* **78** (2008) 034027 [[arXiv:0801.4569](#)] [[INSPIRE](#)].
- [81] D.M. Hofman and J. Maldacena, *Conformal collider physics: energy and charge correlations*, *JHEP* **05** (2008) 012 [[arXiv:0803.1467](#)] [[INSPIRE](#)].
- [82] V. Mateu, I.W. Stewart and J. Thaler, *Power corrections to event shapes with mass-dependent operators*, *Phys. Rev. D* **87** (2013) 014025 [[arXiv:1209.3781](#)] [[INSPIRE](#)].
- [83] A.V. Belitsky et al., *From correlation functions to event shapes*, *Nucl. Phys. B* **884** (2014) 305 [[arXiv:1309.0769](#)] [[INSPIRE](#)].
- [84] A. Ramdas, N. García Trillos and M. Cuturi, *On Wasserstein two-sample testing and related families of nonparametric tests*, *Entropy* **19** (2017) 47.
- [85] H.G. Barrow, J.M. Tenenbaum, R.C. Bolles and H.C. Wolf, *Parametric correspondence and chamfer matching: two new techniques for image matching*, in *Proceedings of the 5th international joint conference on artificial intelligence — volume 2, IJCAI '77*, Morgan Kaufmann Publishers Inc., San Francisco, CA, U.S.A. (1977), p. 659.
- [86] G. Soyez, *Pileup mitigation at the LHC: a theorist's view*, *Phys. Rept.* **803** (2019) 1 [[arXiv:1801.09721](#)] [[INSPIRE](#)].
- [87] CDF collaboration, *Charged jet evolution and the underlying event in $p\bar{p}$ collisions at 1.8 TeV*, *Phys. Rev. D* **65** (2002) 092002 [[INSPIRE](#)].
- [88] A.G. Agocs, G.G. Barnafoldi and P. Levai, *Jets and underlying events at LHC energies*, *J. Phys. Conf. Ser.* **270** (2011) 012017 [[arXiv:1011.5363](#)] [[INSPIRE](#)].
- [89] D.P. Kingma, T. Salimans and M. Welling, *Variational dropout and the local reparameterization trick*, [[arXiv:1506.02557](#)].
- [90] Y.L. Dokshitzer, V.A. Khoze and S.I. Troian, *Particle spectra in light and heavy quark jets*, *J. Phys. G* **17** (1991) 1481 [[INSPIRE](#)].

- [91] Y.L. Dokshitzer, V.A. Khoze and S.I. Troian, *On specific QCD properties of heavy quark fragmentation ('dead cone')*, *J. Phys. G* **17** (1991) 1602 [INSPIRE].
- [92] B.A. Schumm, Y.L. Dokshitzer, V.A. Khoze and D.S. Koetke, *MLLA and the average charged multiplicity of events containing heavy quarks in e^+e^- annihilation*, *Phys. Rev. Lett.* **69** (1992) 3025 [INSPIRE].
- [93] M. Cacciari and G.P. Salam, *Pileup subtraction using jet areas*, *Phys. Lett. B* **659** (2008) 119 [arXiv:0707.1378] [INSPIRE].
- [94] M. Cacciari, G.P. Salam and G. Soyez, *The catchment area of jets*, *JHEP* **04** (2008) 005 [arXiv:0802.1188] [INSPIRE].
- [95] G. Soyez et al., *Pileup subtraction for jet shapes*, *Phys. Rev. Lett.* **110** (2013) 162001 [arXiv:1211.2811] [INSPIRE].
- [96] A.J. Larkoski, S. Marzani, G. Soyez and J. Thaler, *Soft drop*, *JHEP* **05** (2014) 146 [arXiv:1402.2657] [INSPIRE].
- [97] M. Dasgupta, A. Fregoso, S. Marzani and G.P. Salam, *Towards an understanding of jet substructure*, *JHEP* **09** (2013) 029 [arXiv:1307.0007] [INSPIRE].
- [98] J. Altschuler, J. Niles-Weed and P. Rigollet, *Near-linear time approximation algorithms for optimal transport via Sinkhorn iteration*, in the proceedings of the *Advances in neural information processing systems*, (2017).
- [99] P. Dvurechensky, A. Gasnikov and A. Kroshnin, *Computational optimal transport: complexity by accelerated gradient descent is better than by Sinkhorn's algorithm*, in *Proceedings of the 35th international conference on machine learning*, J. Dy and A. Krause eds., *Proc. Machine Learn. Res.* **80** (2018) 1367.
- [100] T. Lin, N. Ho and M.I. Jordan, *On the efficiency of the Sinkhorn and Greenhorn algorithms and their acceleration for optimal transport*, in *International conference on machine learning* (2019).
- [101] J. Altschuler, F. Bach, A. Rudi and J. Niles-Weed, *Massively scalable Sinkhorn distances via the Nyström method*, in the proceedings of the *Advances in neural information processing systems*, (2019).
- [102] M. Scetbon and M. Cuturi, *Linear time Sinkhorn divergences using positive features*, in the proceedings of the *Advances in neural information processing systems*, (2020), p. 13468.
- [103] K. Atasu and T. Mittelholzer, *Linear-complexity data-parallel earth mover's distance approximations*, in the proceedings of the *36th international conference on machine learning*, (2019), p. 364.
- [104] A. Paszke et al., *PyTorch: an imperative style, high-performance deep learning library*, in *Advances in neural information processing systems* 32, H. Wallach et al. eds., Curran Associates Inc. (2019), p. 8024.
- [105] J. Matousek and B. Gartner, *Understanding and using linear programming*, Springer, Berlin, Heidelberg, Germany (2006).
- [106] C. Vilani, *Topics in optimal transportation*, American Mathematical Society, U.S.A. (2003).
- [107] J. Feydy, *Geometric data analysis, beyond convolutions*, ENS Paris-Saclay, France (2020).
- [108] R. Sinkhorn, *Diagonal equivalence to matrices with prescribed row and column sums*, *Am. Math. Mon.* **74** (1967) 402.

- [109] R. Sinkhorn and P. Knopp, *Concerning nonnegative matrices and doubly stochastic matrices*, *Pacific J. Math.* **21** (1967) 343.
- [110] J.J. Kosowsky and A.L. Yuille, *The invisible hand algorithm: solving the assignment problem with statistical physics*, *Neural Networks* **7** (1994) 477.
- [111] D.P. Bertsekas, *The auction algorithm: a distributed relaxation method for the assignment problem*, *Ann. Operat. Res.* **14** (1988) 105.
- [112] D.P. Kingma and J. Ba, *Adam: a method for stochastic optimization*, [arXiv:1412.6980](#) [[INSPIRE](#)].
- [113] P. Tankala, A. Tasissa, J.M. Murphy and D. Ba, *K-deep simplex: deep manifold learning via local dictionaries*, [arXiv:2012.02134](#).
- [114] M. Mueller, S. Aeron, J.M. Murphy and A. Tasissa, *Geometric sparse coding in Wasserstein space*, [arXiv:2210.12135](#).
- [115] W. Wang and M. Carreira-Perpiñán, *Projection onto the probability simplex: an efficient algorithm with a simple proof, and an application*, [arXiv:1309.1541](#).
- [116] T. Sjöstrand et al., *An introduction to PYTHIA 8.2*, *Comput. Phys. Commun.* **191** (2015) 159 [[arXiv:1410.3012](#)] [[INSPIRE](#)].
- [117] DELPHES 3 collaboration, *DELPHES 3, a modular framework for fast simulation of a generic collider experiment*, *JHEP* **02** (2014) 057 [[arXiv:1307.6346](#)] [[INSPIRE](#)].
- [118] M. Cacciari, G.P. Salam and G. Soyez, *FastJet user manual*, *Eur. Phys. J. C* **72** (2012) 1896 [[arXiv:1111.6097](#)] [[INSPIRE](#)].
- [119] R. Flamary et al., *POT: Python Optimal Transport*, *J. Mach. Learn. Res.* **22** (2021) 1.
- [120] CMS collaboration, *Jet energy scale and resolution in the CMS experiment in pp collisions at 8 TeV*, *2017 JINST* **12** P02014 [[arXiv:1607.03663](#)] [[INSPIRE](#)].
- [121] ATLAS collaboration, *Jet energy scale and resolution measured in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Eur. Phys. J. C* **81** (2021) 689 [[arXiv:2007.02645](#)] [[INSPIRE](#)].
- [122] E. Wigner, *The unreasonable effectiveness of mathematics in the natural sciences*, *Commun. Pure Appl. Math.* **13** (1960) 1 [[INSPIRE](#)].
- [123] M. Kac, *Can one hear the shape of a drum?*, *Am. Math. Mon.* **73** (1966) 1 [[INSPIRE](#)].
- [124] S. Kullback and R.A. Leibler, *On information and sufficiency*, *Annals Math. Statist.* **22** (1951) 79.
- [125] B. Piccoli and F. Rossi, *Generalized Wasserstein distance and its application to transport equations with source*, *Archive for Rational Mechanics and Analysis* **211** (2013) 335.
- [126] M. Liero, A. Mielke and G. Savaré, *Optimal transport in competition with reaction: the Hellinger-Kantorovich distance and geodesic curves*, *SIAM J. Math. Anal.* **48** (2016) 2869.