



HAL
open science

Wasserstein Auto-Encoders of Merge Trees (and Persistence Diagrams)

Mathieu Pont, Julien Tierny

► **To cite this version:**

Mathieu Pont, Julien Tierny. Wasserstein Auto-Encoders of Merge Trees (and Persistence Diagrams). IEEE Transactions on Visualization and Computer Graphics, In press, pp.1-16. 10.1109/TVCG.2023.3334755 . hal-04552330

HAL Id: hal-04552330

<https://hal.science/hal-04552330v1>

Submitted on 19 Apr 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Wasserstein Auto-Encoders of Merge Trees (and Persistence Diagrams)

Mathieu Pont and Julien Tierny

Abstract—This paper presents a computational framework for the Wasserstein auto-encoding of merge trees (MT-WAE), a novel extension of the classical auto-encoder neural network architecture to the Wasserstein metric space of merge trees. In contrast to traditional auto-encoders which operate on vectorized data, our formulation explicitly manipulates merge trees on their associated metric space at each layer of the network, resulting in superior accuracy and interpretability. Our novel neural network approach can be interpreted as a non-linear generalization of previous linear attempts [79] at merge tree encoding. It also trivially extends to persistence diagrams. Extensive experiments on public ensembles demonstrate the efficiency of our algorithms, with MT-WAE computations in the orders of minutes on average. We show the utility of our contributions in two applications adapted from previous work on merge tree encoding [79]. First, we apply MT-WAE to *merge tree compression*, by concisely representing them with their coordinates in the final layer of our auto-encoder. Second, we document an application to *dimensionality reduction*, by exploiting the latent space of our auto-encoder, for the visual analysis of ensemble data. We illustrate the versatility of our framework by introducing two penalty terms, to help preserve in the latent space both the Wasserstein distances between merge trees, as well as their clusters. In both applications, quantitative experiments assess the relevance of our framework. Finally, we provide a C++ implementation that can be used for reproducibility.

Index Terms—Topological data analysis, ensemble data, merge trees, persistence diagrams.



1 INTRODUCTION

WITH the recent advances in the development of computation hardware and acquisition devices, datasets are constantly increasing in size. This size increase induces an increase in the geometrical complexity of the features present in the datasets, which challenges interactive data analysis and interpretation. To address this issue, Topological Data Analysis (TDA) [30] has shown over the years its ability to reveal, in a generic, robust and efficient manner, the main structural patterns hidden in complex datasets, in particular for visual data analysis tasks [46]. Successful applications have been documented in multiple fields (turbulent combustion [19], [43], material sciences [45], [89], nuclear energy [64], fluid dynamics [53], [70], bioimaging [3], [14], quantum chemistry [71], [72] or astrophysics [87], [90]). Among the representations studied in TDA, the merge tree [23] (Fig. 2) has been prominent in data visualization [14], [19], [25].

In addition to the increase in geometrical complexity discussed above, a new challenge has recently emerged in many applications, with the notion of *ensemble datasets*. Such datasets encode a given phenomenon not only with a single dataset, but with a collection of datasets, called *ensemble members*. In that context, the topological analysis of an ensemble dataset consequently yields an ensemble of corresponding topological representations (e.g. one merge tree per ensemble member).

Then, developing statistical analysis tools to support the interactive analysis and interpretation of ensemble data becomes an important challenge. Recently, several works explored this direction, in particular with the notion of *average topological representation* [59], [78], [95], [97], [103]. These approaches can produce a topological representation which nicely summarizes the ensemble. Moreover, their application to clustering [78] reveal its

main trends. However, they do not provide any hints regarding the variability of the features in the ensemble. For this, Pont et al. [79] recently extended the notion of principal geodesic analysis to ensembles of merge trees. However, this approach implicitly assumes a linear relation between the merge trees of the ensemble. Specifically, it assumes that merge tree branches evolve linearly (in the birth/death space, Sec. 2) within the ensemble.

This paper addresses this issue with a novel formulation based on neural networks and introduces the first framework for the non-linear encoding of merge trees, hence resulting in superior accuracy. Specifically, we formulate merge tree non-linear encoding as an auto-encoding problem (Sec. 3). We contribute a novel neural network called *Wasserstein Auto-Encoder of Merge Trees*. This network is based on a novel layer model, capable of processing merge trees natively, without pre-vectorization. We believe this contribution to be of independent interest, as it enables an accurate and interpretable processing of merge trees by neural networks (without restrictions to auto-encoders). We contribute an algorithm for the optimization of such a network (Sec. 4). We illustrate the relevance of our contributions for visual analysis with two applications, data reduction (Sec. 5.1) and dimensionality reduction (Sec. 5.2). Similarly to previous linear attempts [79], since our approach is based on the Wasserstein distance between merge trees [78], which generalizes the Wasserstein distance between persistence diagrams [30], our framework trivially extends to persistence diagrams by simply adjusting a single parameter.

1.1 Related work

We classify the literature related to our approach into two categories: ensemble analysis and topological methods for ensembles. **(I) Ensemble analysis:** Typical approaches to ensemble visualization first characterize each member of the ensemble by extracting a geometrical object representing its features of interest (level

• Mathieu Pont and Julien Tierny are with the CNRS/Sorbonne Université.
E-mails: {firstname.lastname}@sorbonne-universite.fr

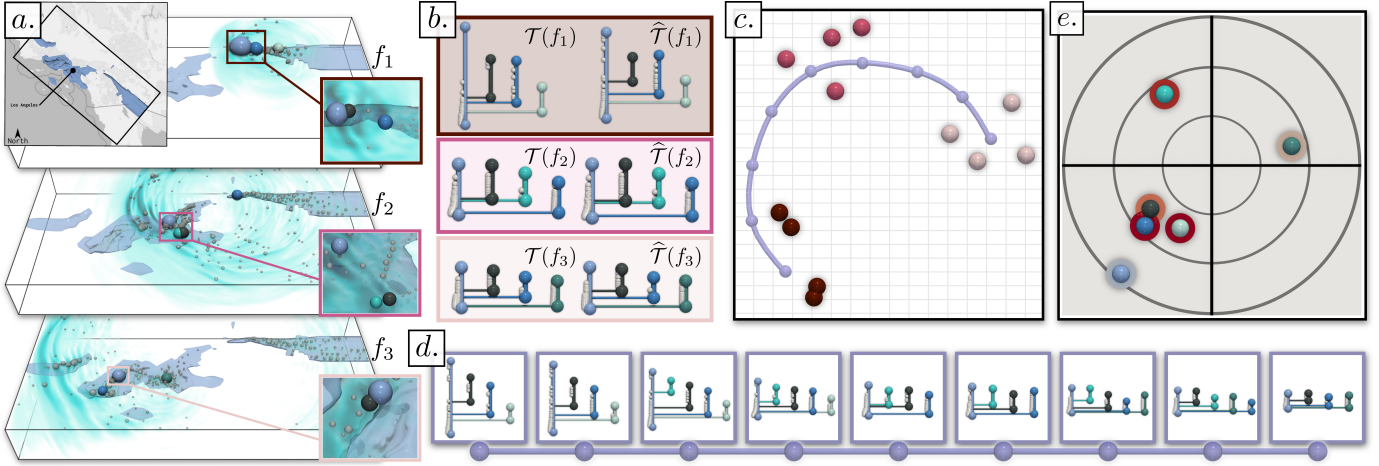


Fig. 1. Visual analysis of the *Earthquake* ensemble ((a) each ground-truth class is represented by one of its members), with our Wasserstein Auto-Encoder of Merge Trees (MT-WAE). We apply our contributions to merge tree compression ((b), right) by simply storing their coordinates in the last decoding layer of our network. We exploit the latent space of our network to generate 2D layouts of the ensemble (c). In contrast to classical auto-encoders, MT-WAE explicitly manipulates merge trees at each layer of the network, which results in improved accuracy and interpretability. Specifically, the reconstruction of user-defined locations ((c), purple) enables an interactive exploration of the latent space: the reconstructed curve (d) enables a continuous navigation between the clusters (from dark red to pink and light pink, (c)). MT-WAE also supports persistence correlation views (e) (adapted from [79]), which reveal the barycenter’s persistent features which exhibit the most variability in the ensemble (far from the center). Finally, by tracking the persistence evolution of individual features as they traverse the network down to its latent space, we introduce a *Feature Latent Importance* measure, which identifies the most *informative* features within the ensemble ((e), red circles).

sets, streamlines, etc). Next, a second step considers the ensemble of geometrical objects computed in the first step, and estimates a single *representative* object, representing an aggregate of the features of interest found in the ensemble. For example, level-set variability has been studied with spaghetti plots [29], with specific applications to weather data [80], [85]. More generally, the variability in curves and contours have been studied with the notion of box-plots [100] and its variants [65]. Hummel et al. [51] analyzed the variability in ensembles of flows with a Lagrangian approach. The main trends present in an ensemble have been studied for ensembles of streamlines [36] and level sets [37] via clustering techniques. Other approaches focused on visualizing the geometrical variability in the domain of the position of critical points [35], [42] or gradient separatrices [5]. For ensemble of contour trees, consistent planar layouts have also been studied [62], to support the direct visual comparison of the trees. While the latter approaches have a topological aspect, they focused on the direct visualization of the variability, and not on the statistical analysis of an ensemble of topological descriptors.

General purpose methods have been documented for non-linear encoding (e.g. *topological auto-encoders* [67], or *Wasserstein auto-encoders* [94]). Our work drastically differs from these methods, in terms of design and purpose. These methods [47], [67], [94] employ a classical auto-encoder (Sec. 3.1) to which they add specialized penalty terms. Then, their input is restricted to point sets (or vectorized data). In contrast, our work focuses on sets of merge trees (or persistence diagrams). This different kind of input requires a novel neural network model, capable of processing these topological objects natively (Sec. 3).

(2) Topological methods: Over the last two decades, the visualization community has investigated, adapted and extended [46], [102] several tools and notions from computational topology [30]. The persistence diagram [9], [30], [32], [41], the Reeb graph [12], [40], [74], the merge (Fig. 2) and contour trees [1], [23], [24], [39],

[63], and the Morse-Smale complex [18], [28], [44], [83], [86] are popular examples of topological representations in visualization.

In order to design a statistical framework for the analysis of ensembles of topological descriptors, one first needs to define a metric to measure distances between these objects. The Wasserstein distance [30] (Sec. 2.2) is now an established and well-documented metric for persistence diagrams. It is inspired by optimal transport [52], [66] and it is defined (Sec. 2.2) via a bipartite assignment problem (open-source software implementing exact computations [69] or fast approximations [11], [54] is available [93]). However, as discussed in previous work [10], [68], [78], [79], [91], the persistence diagram can lack specificity in its encoding of the features of interest, motivating more advanced descriptors, like the merge trees (Sec. A.2), which better distinguishes datasets. The comparison of Reeb graphs and their variants has been addressed with several similarity measures [48], [84]. Several works investigated the theoretical aspects of distances between topological descriptors, in particular with a focus towards stable distances [8], [15], [16], [68]. However, the computation of such distances rely on algorithms with exponential time complexity, which is not practicable for real-life datasets. In contrast, a distinct line of research focused on a balance between practical stability and computability, by focusing on polynomial time computation algorithms. Beketayev et al. [10] introduced a distance for the *branch decomposition tree* (BDT, Appendix A). Efficient algorithms for constrained edit distances [104] have been specialized to the specific case of merge trees, hence providing an edit distance for merge trees [91] which is both computable in practice and with acceptable practical stability. Pont et al. [78] extended this work to generalize the L_2 -Wasserstein distance between persistence diagrams [30] to merge trees, hence enabling the efficient computation of distances, geodesics and barycenters of merge trees. Wetzels et al. [98], [99] introduce metrics independent of a particular branch decomposition, but this comes

at the cost of a significantly larger computational effort (with quartic time complexity instead of quadratic), which prevents their practical computation on full-sized merge trees.

Once a metric is available, statistical notions can be developed for topological descriptors. Several methods [59], [95], [97] have been introduced for the estimation of barycenters of persistence diagrams (or vectorized variants [2], [22]). Similar approaches have been specifically derived for the merge trees [78], [103]. Another set of approaches [4], [61], [82] first considers *vectorizations* of topological descriptors (i.e. by converting them into high-dimensional Euclidean vectors) and then leverages traditional linear-algebra tools on these vectors (e.g. the classical PCA [77] or its variants from matrix sketching [101]). Several approaches in machine learning are constructed on top of vectorizations of topological descriptors [2], [22], [56] or kernel-based representations [26], [81]. However, such vectorizations have several limitations in practice. First, they are prone to approximation errors (resulting from quantization and linearization). Also, they can be difficult to revert (especially for barycenters), which makes them impractical for visualization tasks. Moreover, their stability is not always guaranteed. In contrast, Pont et al. [79] extended the generic notion of principal geodesic analysis to the Wasserstein metric space of merge trees, resulting in improved accuracy and interpretability with regard to the straightforward application of PCA on merge tree vectorizations. Similarly, Sisouk et al. [88] introduced a simpler approach for the linear encoding of persistence diagrams, with a less constrained framework based on dictionaries. However, these approaches implicitly assume a linear relation between the topological descriptors of the ensemble. For instance, it assumes that a given feature (i.e. a given branch of the merge tree) evolves linearly in the birth/death space (Sec. 2) within the ensemble. However, this hypothesis is easily challenged in practice (Figs. 5 and 7), potentially leading to inaccuracies. Our work overcomes this limitation with a drastically different formulation (based on auto-encoding neural networks) and introduces the first framework for the non-linear encoding of merge trees, resulting in superior accuracy. Several approaches [21], [55], [105] investigated the use of topological methods for the analysis of neural networks. In contrast, our work targets a different research problem, specifically the encoding of topological descriptors with neural networks.

1.2 Contributions

This paper makes the following new contributions:

- 1) *An approach to Merge tree non-linear encoding:* We formulate the non-linear parametrization of the Wasserstein metric space of merge trees (and persistence diagrams) as an auto-encoding problem. Our formulation (Sec. 3) generalizes and improves previous linear attempts [79].
- 2) *A vectorization-free neural network architecture for Merge Trees:* We contribute a novel neural network architecture called *Wasserstein Auto-Encoder of Merge Trees*, inspired by the classical auto-encoder, which can *natively* process merge trees (and persistence diagrams) *without* prior vectorization. For this, we contribute a novel layer model, which takes a set of merge trees on its input and produces another set of merge trees on its output, along with their coordinates in the layer’s parametrization. This results in superior accuracy (Sec. 6.2) and interpretability (Sec. 5.2). We contribute an algorithm for the optimization of this network (Sec. 4). We believe this contribution to be of independent interest.




Fig. 2. Illustration of the topological descriptors considered in this work, on a clean (a) and a noisy (b) variant of a 2D toy dataset. For all descriptors, the color code indicates the persistence of the corresponding saddle-maximum pair. Critical points are represented with spheres (larger ones for maxima). Persistence diagrams, merge trees and branch decomposition trees (BDTs) are respectively represented in the left, center and right insets. For both datasets, the four main features (the larger hills) are represented with salient pairs in the diagram and the merge tree. To avoid clutter in the visualization, the branches with low persistence (less than 10% of the function range) are rendered with small white arcs while larger, and colored arcs represent persistent branches (more than 10% of the function range). Figure adapted from [78], [79]

- 3) *An application to merge tree compression:* We describe how to adapt previous work [79] to our novel non-linear framework, in merge tree compression applications (Sec. 5.1). Specifically, the merge trees of the input ensemble are significantly compressed, by solely storing the final decoding layer of the network, as well as the coordinates of the input trees in this layer. We illustrate the interest of our approach with comparisons to linear encoding [79] in the context of feature tracking and ensemble clustering applications.
- 4) *An application to dimensionality reduction:* We describe how to adapt previous work [79] to our novel non-linear framework, in the context of dimensionality reduction applications (Sec. 5.2). Specifically, each tree of the ensemble is embedded as a point in a planar view, based on its coordinates in our auto-encoder’s latent space. To illustrate the versatility of our framework, we introduce two penalty terms, to improve the preservation of clusters and distances between merge trees.
- 5) *Implementation:* We provide a C++ implementation of our algorithms that can be used for reproducibility purposes.

2 PRELIMINARIES

This section presents the background of our work. First, we describe the input data and its topological representation. Second, we recap the Wasserstein metric space of merge trees [78], used by our approach. We refer to textbooks [30] for an introduction to computational topology.

2.1 Input data

The input data is an ensemble of N piecewise linear (PL) scalar fields $f_i: \mathcal{M} \rightarrow \mathbb{R}$, with $i \in \{1, \dots, N\}$, defined on a PL d -manifold \mathcal{M} , with $d \leq 3$ in our applications. Each ensemble member f_i is represented by a *topological descriptor*. In this work, we focus on the *Persistence Diagram* (PD), noted $\mathcal{D}(f_i)$, as well as a variant of the *Merge Tree* (MT), called the *Branch Decomposition Tree* (BDT), noted $\mathcal{B}(f_i)$. A formal description of these descriptors is given in Appendix A.




Fig. 3. Illustration of the computation of the Wasserstein distance W_2^T [78] between the BDTs $\mathcal{B}(f_i)$ (white, top) and $\mathcal{B}(f_j)$ (blue, bottom). W_2^T is computed by solving an optimal assignment problem in the birth/death plane (right), whose search space is constrained to partial rooted isomorphisms (the optimal solution is represented with a cyan halo on the BDTs and with cyan arrows in the birth/death plane). When $\varepsilon_1 = 1$, all the saddles are collapsed and the structure of the BDT is completely ignored. In that case, W_2^T is equal to the Wasserstein distance between persistence diagrams W_2^D (the optimal assignment is shown with red arrows, right). In this example, W_2^D reports a small distance whereas two hills have been swapped in the datasets. Figure adapted from [79].

In short, $\mathcal{D}(f_i)$ is a 2D point cloud (Fig. 2), where each off-diagonal point $p = (x, y)$ denotes a *topological feature* in the data filtration (e.g. a connected component, a cycle, a void, etc.). x and y denote the *birth* and the *death* of p (i.e. the scalar values for the creation and destruction of the corresponding topological feature in the data). The *persistence* of p is given by its height to the diagonal (Fig. 2, vertical cylinders, left insets). Important features are typically associated with a large persistence, while low amplitude noise is in the vicinity of the diagonal.

The Merge Tree (MT, Fig. 2, center insets) is a slightly more informative descriptor, as it additionally encodes the merge history between the topological features. To mitigate a phenomenon called *saddle swap* [78], [91], it is often pre-processed to merge adjacent forking nodes whose relative scalar value difference is smaller than a threshold $\varepsilon_1 \in [0, 1]$. The merge tree can be represented in a dual form called the *Branch Decomposition Tree* (BDT) $\mathcal{B}(f_i)$ (Fig. 2, right insets), where each persistent branch of the merge tree (vertical cylinder, center insets) is transformed into a node in the BDT (sphere, right insets) and where each horizontal segment of the merge tree is transformed into an arc. Given an arbitrary BDT \mathcal{B} , since each node $b \in \mathcal{B}$ embeds its own birth/death values (numbers in Fig. 2), it is possible to reconstruct the corresponding merge tree, as long as \mathcal{B} respects the *Elder rule* [30].

2.2 Wasserstein metric space

In this section, we first formalize the Wasserstein distance between persistence diagrams [30]. Next, we recall its generalization to merge trees [78]. This generalization enables our approach to support both topological descriptors (persistence diagrams and merge trees). This section includes elements adapted from [79], which have been considered to make this manuscript self-contained.

The evaluation of the distance between two diagrams $\mathcal{D}(f_i)$ and $\mathcal{D}(f_j)$ is typically preceded by a pre-processing step aiming at transforming the diagrams, such that they admit the same number of points, which will facilitate the evaluation of their distance. This procedure augments each diagram with the diagonal projection of

the off-diagonal points of the other diagram:

$$\begin{aligned} \mathcal{D}'(f_i) &= \mathcal{D}(f_i) \cup \{\Delta(p_j) \mid p_j \in \mathcal{D}(f_j)\} \\ \mathcal{D}'(f_j) &= \mathcal{D}(f_j) \cup \{\Delta(p_i) \mid p_i \in \mathcal{D}(f_i)\}, \end{aligned}$$

where $\Delta(p_i) = (\frac{x_i+y_i}{2}, \frac{x_i+y_i}{2})$ is the diagonal projection of the off-diagonal point $p_i = (x_i, y_i)$. Overall, this augmentation procedure inserts dummy features in the diagram (with zero persistence, on the diagonal), hence preserving the topological information of the diagrams, while guaranteeing that the two diagrams now have the same number of points ($|\mathcal{D}'(f_i)| = |\mathcal{D}'(f_j)|$).

In order to compare two points $p_i = (x_i, y_i) \in \mathcal{D}'(f_i)$ and $p_j = (x_j, y_j) \in \mathcal{D}'(f_j)$, a *ground distance* needs to be introduced in the birth/death plane. Specifically, we consider the distance $d_q (q > 0)$:

$$d_q(p_i, p_j) = (|x_j - x_i|^q + |y_j - y_i|^q)^{1/q} = \|p_i - p_j\|_q.$$

In the special case where both p_i and p_j are dummy features located on the diagonal (i.e. $x_i = y_i$ and $x_j = y_j$), $d_q(p_i, p_j)$ is set to zero (such that these dummy features do not intervene in the distance evaluation between the diagrams). Then, the L^q -Wasserstein distance W_q^D can be introduced as:

$$W_q^D(\mathcal{D}(f_i), \mathcal{D}(f_j)) = \min_{\phi \in \Phi} \left(\sum_{p_i \in \mathcal{D}'(f_i)} d_q(p_i, \phi(p_i))^q \right)^{1/q}, \quad (1)$$

where Φ is the set of all possible assignments ϕ mapping a point $p_i \in \mathcal{D}'(f_i)$ to a point $p_j \in \mathcal{D}'(f_j)$. Note that it is possible that ϕ maps a point $p_i \in \mathcal{D}'(f_i)$ to its diagonal projection (i.e. $\phi(p_i) = \Delta(p_i) = p_j \in \mathcal{D}'(f_j)$), which indicates the destruction of the corresponding feature (or symmetrically, its appearance).

Pont et al. recently generalized this metric to BDTs [78]. The expression of this distance, noted $W_2^T(\mathcal{B}(f_i), \mathcal{B}(f_j))$, is the same as Eq. 1 (for $q = 2$), with the important difference of the search space of possible assignments, noted $\Phi' \subseteq \Phi$. Specifically, Φ' is constrained to the set of (rooted) partial isomorphisms [78] between $\mathcal{B}(f_i)$ and $\mathcal{B}(f_j)$ (cyan halo on the BDTs of Fig. 3).

This novel metric comes with a clear interpretation. The control parameter ε_1 (Sec. A.2) balances the importance of the BDT structure in the distance. Specifically, when $\varepsilon_1 = 1$, all saddles are collapsed and we have $W_2^T(\mathcal{B}(f_i), \mathcal{B}(f_j)) = W_2^D(\mathcal{D}(f_i), \mathcal{D}(f_j))$. Generally speaking, as illustrated experimentally by Pont et al. [78], ε_1 acts as a control knob balancing the practical stability of the metric with its discriminative power. This generalized metric enables our framework to support both topological descriptors. For persistence diagrams, we set $\varepsilon_1 = 1$, while for merge trees, we set it to the default recommended value (i.e. $\varepsilon_1 = 0.05$ [78]). In the following, the metric space induced by this metric is noted \mathbb{B} .

For interpolation purposes, Pont et al. introduce a local normalization [78] in a pre-processing step (to guarantee the invertibility of any interpolated BDT into a valid MT). We will use the same procedure in this work to guarantee that the projected BDTs (Sec. 3.2) indeed describe valid MTs. Specifically, we normalize the persistence of each branch $b_i \in \mathcal{B}(f_i)$ with regard to that of its parent $b'_i \in \mathcal{B}(f_i)$, by moving b_i in the birth/death plane from (x_i, y_i) to $\mathcal{N}(b_i) = (\mathcal{N}_x(b_i), \mathcal{N}_y(b_i))$:

$$\begin{aligned} \mathcal{N}_x(b_i) &= (x_i - x'_i) / (y'_i - x'_i) \\ \mathcal{N}_y(b_i) &= (y_i - x'_i) / (y'_i - x'_i). \end{aligned} \quad (2)$$

This pre-normalization procedure guarantees that any interpolated BDT can indeed be reverted into a valid MT, by recursively applying Eq. 2, which explicitly enforces the *Elder rule* [30] on BDTs ($[x_i, y_i] \subseteq [x'_i, y'_i]$, Sec. A.2), hence the validity of the




Fig. 4. Geometric interpretation of Euclidean Auto-Encoders (EAE, Sec. 3.1). In its simplest form (one encoding and one decoding layer, $d' = 2$), an EAE can be viewed as the composition of a linear transformation ψ_1 ((a) to (b)) defined with respect to a first basis B_1 , followed by a non-linearity σ (here *ReLU*, (b) to (c)), followed by a second linear transformation ψ_2 ((c) to (d)) defined with respect to a second basis B_2 . Specifically, both ψ_1 and ψ_2 are optimized (via the optimization of B_1 and B_2) to minimize the reconstruction error between the input (a) and the output (d). In the case where σ is the identity, this reconstruction optimization is equivalent to Principal Component Analysis [17].

reconstructed MT. Two additional parameters were introduced by Pont et al. [78] in order to control the effect of this pre-normalization procedure (ε_2 balances the normalized persistence of small branches, selected via the threshold ε_3). These parameters are set to their default recommended values ($\varepsilon_2 = 0.95$, $\varepsilon_3 = 0.9$). In the following, we consider that all the input BDTs are pre-normalized with this procedure.

3 FORMULATION

This section describes our novel extension of the classical auto-encoder neural network architecture to the Wasserstein metric space of merge trees, with the novel notion of *Merge Tree Wasserstein Auto-Encoder* (MT-WAE). First, we describe a geometric interpretation of the classical auto-encoders (Sec. 3.1), which we call in the following *Euclidean Auto-Encoders* (EAE). Next, we describe how to generalize each geometrical tool used in EAE (Sec. 3.1) to the Wasserstein metric space of merge trees (Sec. 3.2). Finally, once these tools are available, we formalize our notion of MT-WAE with a novel neural network architecture (Sec. 3.3), for which we document an optimization algorithm in Sec. 4.

3.1 An interpretation of Euclidean Auto-Encoders

Let $P = \{p_1, p_2, \dots, p_N\}$ be a point set in a Euclidean space \mathbb{R}^d (Fig. 4a). The goal of Euclidean Auto-Encoders (EAE) is to define a d' -dimensional parameterization of P (with $d' \leq d$) which describes well the data (which enables its accurate *reconstruction*). Let $B_1 = \{b_1, b_2, \dots, b_{d'}\}$ be a basis of linearly independent vectors in \mathbb{R}^d (Fig. 4a). B_1 can be written in the form of a $d \times d'$ matrix, for which each of the d' columns is a vector of the basis. Then, one can express the coordinates $\psi_1(p_i) \in \mathbb{R}^{d'}$ of each point $p_i \in P$ with the basis B_1 :

$$\psi_1(p_i) = \arg \min_{\alpha^i} \|p_i - B_1 \alpha^i\|_2^2 + o_1, \quad (3)$$

where o_1 is an *offset* vector of $\mathbb{R}^{d'}$, and where $\alpha^i \in \mathbb{R}^{d'}$ can be seen as a set of d' coefficients, to apply on the d' vectors of the basis B_1 to best estimate p_i . Note that Eq. 3 can be re-written as a linear transformation:

$$\psi_1(p_i) = B_1^+ p_i + o_1, \quad (4)$$

where B_1^+ is the pseudoinverse of the matrix B_1 (Figs. 4a-b).

Given this new parameterization ψ_1 , one can estimate a *reconstruction* of the point p_i in \mathbb{R}^d , noted \hat{p}_i . For this, let us




Fig. 5. Comparison between PCA (a) and EAE (b) for the 1-dimensional encoding of a 2D point set sampling a 1-manifold (color: rotation angle). In its latent space ((a), bottom), PCA linearly projects the input points to a line, hence interleaving points from the upper and lower parts of the circle. This results in a poor reconstruction ((a), right), where points are interleaved along the axis B_1 . In contrast, EAE optimizes a composition of non-linear transformations, which consistently *unwraps* the circle onto a line in its latent space ((b), bottom), while nicely preserving the intrinsic parameterization of the circle (rotation angle). This results in an accurate reconstruction ((b), right): the embedding of the axis B_{n_e} in the data defines a faithful 1-dimensional parameterization of the circle.

consider another, similar, linear transformation ψ_2 (Figs. 4c-d), defined respectively to a second basis B_2 (given as a $d' \times d$ matrix) and a second offset vector o_2 (in \mathbb{R}^d). Then, the reconstruction \hat{p}_i of each point p_i is given by:

$$\hat{p}_i = \psi_2 \circ \psi_1(p_i) = B_2^+ (B_1^+ p_i + o_1) + o_2.$$

To get an accurate reconstruction \hat{p}_i (Fig. 4d) for all the points $p_i \in P$, one needs to optimize both ψ_1 and ψ_2 , to minimize the following data fitting energy:

$$E_{L_2} = \sum_{i=1}^N \|p_i - \hat{p}_i\|_2^2 = \sum_{i=1}^N \|p_i - \psi_2 \circ \psi_1(p_i)\|_2^2. \quad (5)$$

As discussed by Bourlard and Kamp [17], this formulation is a generalization of Principal Component Analysis (PCA) [77], a seminal statistical tool for variability analysis. However, PCA assumes that the input point cloud can be faithfully approximated via linear projections. As shown in Fig. 5, this hypothesis can be easily challenged in practice. This motivates a non-linear generalization of PCA, capable of faithfully approximating point clouds exhibiting non-linear structures (Fig. 5).

Specifically, to introduce non-linearity, the above linear transformations Ψ are typically composed with a non-linear function σ , called *activation function*, such that the transformation of each point p_i , noted $\Psi(p_i)$, is now given by: $\Psi(p_i) = \sigma(\psi(p_i))$. For example, the rectifier activation function (“ReLU”) will take the j^{th} coordinate of each data point (i.e. $(\psi(p_i))_j$) and *snap* it to zero if it is negative (Fig. 4b-c). We call the above non-linear transformation Ψ a *transformation layer*. It is characterized by its own vector basis B and its own offset vector o .

Next, to faithfully approximate complicated non-linear input distributions, the above transformation layer is typically composed with a number ($n_e + n_d$) of other transformation layers, defined similarly. Then, the initial data fitting energy (Eq. 5) can now be generalized into:

$$E_{L_2} = \sum_{i=1}^N \|p_i - \Psi_{n_e+n_d} \circ \dots \circ \Psi_{n_e+1} \circ \Psi_{n_e} \circ \dots \circ \Psi_2 \circ \Psi_1(p_i)\|_2^2, \quad (6)$$

where each transformation layer Ψ_k is associated with its own d_k -dimensional vector basis B_k and its offset vector o_k . Then,




Fig. 6. Low-level geometrical tools in the Wasserstein metric space of merge trees (Sec. 3.2). Given an origin MT and its BDT \mathcal{O} (counting $|\mathcal{O}|$ nodes (a)), a BDT vector $\mathcal{V}_1(\mathcal{O})$ is defined in the birth/death space as a concatenation of $|\mathcal{O}|$ 2D vectors (blue arrows (b)). Given a second BDT vector $\mathcal{V}_2(\mathcal{O})$ (cyan arrows (c)), a basis $B(\mathcal{O})$ can be defined ((d), here with $d' = 2$). For a given set of coefficients $\alpha \in \mathbb{R}^{d'}$, a new merge tree and its BDT (e) can be reconstructed by applying a sum of 2D displacements $\sum_{j=1}^{j=d'} \alpha_j (\mathcal{V}_j(\mathcal{O}))_i$ to each branch b_i of \mathcal{O} ((d), black dashes).

the notion of *Auto-Encoder* is a specific instance of the above formulation, with:

- 1) $d_1 > d_2 > \dots > d_{n_e}$, and
- 2) $d_{n_e} < d_{n_e+1} < \dots < d_{n_e+n_d} = d$, and
- 3) $\sigma_{n_e+n_d}$ is the identity.

Specifically, n_e and n_d respectively denote the number of *Encoding* and *Decoding* transformation layers, while d_{n_e} is the dimension of the so-called *latent space*. In practice, d_{n_e} is typically chosen to be much smaller than the input dimensionality (d), for non-linear dimensionality reduction purposes (each input point p_i is then represented in d_{n_e} dimensions, according to its coordinates in B_{n_e} , noted $\alpha_{n_e}^i \in \mathbb{R}^{d_{n_e}}$).

Eq. 6 defines an optimization problem whose variables (the $n_e + n_d$ bases and offset vectors) can be efficiently optimized (e.g. with gradient descent [57]) by composing the transformation layers within a neural network. Then, the gradient ∇E_{L_2} of E_{L_2} can be estimated by exploiting the automatic differentiation capabilities of modern neural network implementations [76], themselves based on the application of the chain-rule derivation on the above composition of transformation layers.

3.2 From EAE to MT-WAE

When the input data is not given as a point cloud in a Euclidean space (Sec. 3.1) but as an abstract set equipped with a metric, the above EAE formulation needs to be extended. For this, we redefine in this section the low-level geometrical tools used in EAE (Sec. 3.1), but within the context of the Wasserstein metric space \mathbb{B} [78]. In particular, we formalize the following notions:

- 1) BDT vector (Fig. 6b);
- 2) BDT basis (Fig. 6d);
- 3) BDT basis projection (Fig. 6d);
- 4) BDT transformation layer (Fig. 7b).

(1) BDT vector: Given a BDT \mathcal{B} with $|\mathcal{B}|$ branches, a *BDT vector* $\mathcal{V}(\mathcal{B}) \in \mathbb{B}$ is a vector in $\mathbb{R}^{2^{|\mathcal{B}|}}$, which maps each branch $b \in \mathcal{B}$ to a new location in the 2D birth/death space. \mathcal{B} is the origin of $\mathcal{V}(\mathcal{B})$. This is illustrated for example in Fig. 6b), where the branches of a given merge tree (Fig. 6a) are displaced in the birth-death plane (light blue vectors from the spheres of matching color).

(2) BDT basis: Given an origin BDT \mathcal{O} , a d' -dimensional basis of BDT vectors, noted $B(\mathcal{O})$, is a set $\{\mathcal{V}_1(\mathcal{O}), \mathcal{V}_2(\mathcal{O}), \dots, \mathcal{V}_{d'}(\mathcal{O})\}$ of d' linearly independent BDT vectors, having for common origin \mathcal{O} . This is shown in Fig. 6d), where two BDT vectors (from

Fig. 6b), and Fig. 6c), blue and green arrows) are combined into a basis. Sec. 4.3 clarifies the basis initialization by our approach.

(3) BDT basis projection: Given an arbitrary BDT \mathcal{B} , its *projection error* $e(\mathcal{B})$ in a d' -dimensional BDT basis $B(\mathcal{O})$ is:

$$e(\mathcal{B}) = \min_{\alpha} \left(W_2^T(\mathcal{B}, \mathcal{O} + B(\mathcal{O})\alpha) \right)^2, \quad (7)$$

where $\alpha \in \mathbb{R}^{d'}$ can be interpreted as a set of coefficients to apply on the d' BDT vectors of $B(\mathcal{O})$ to best estimate \mathcal{B} . Then, the projection of \mathcal{B} in $B(\mathcal{O})$, noted $\psi(\mathcal{B})$, is given by the optimal coefficients α associated to the projection error of \mathcal{B} (Eq. 7):

$$\psi(\mathcal{B}) = \arg \min_{\alpha} \left(W_2^T(\mathcal{B}, \mathcal{O} + B(\mathcal{O})\alpha) \right)^2. \quad (8)$$

The above equation is a re-interpretation of Eq. 3 (Sec. 3.1), where the L_2 norm (used for EAE) is replaced by the Wasserstein distance W_2^T . This projection procedure is illustrated in Fig. 6d). Given a BDT basis $B(\mathcal{O})$ (blue and green arrows), the projection of an arbitrary BDT \mathcal{B} is the linear combination $\psi(\mathcal{B})$ of the BDT vectors of the basis which minimizes its Wasserstein distance W_2^T to \mathcal{B} . In the birth-death space, the branches of the basis origin \mathcal{O} are represented by the colored spheres at the intersection between the blue and green arrows, while the branches of $\psi(\mathcal{B})$ are represented by the other spheres of matching color. In this example, the linear combination $\psi(\mathcal{B})$ which minimizes its distance to \mathcal{B} is obtained for the coefficients $\alpha = (0.5, 1)$. Then, to go from \mathcal{O} to $\psi(\mathcal{B})$, each branch b_i of \mathcal{O} is displaced by $0.5 \mathcal{V}_1(\mathcal{O})$ (intersection between the dashed lines and the blue arrows) and then by $1 \mathcal{V}_2(\mathcal{O})$ (intersection between the dashed lines and the green arrows). The resulting merge tree is shown in Fig. 6e). In contrast to the merge tree of the basis origin (Fig. 6a)), the persistence of the cyan branch has increased while that of the black branch has decreased.

(4) BDT transformation layer: Once the above tools have been formalized, we can introduce the novel notion of *BDT transformation layer*. Similarly to the Euclidean case (Sec. 3.1), a non-linear activation function σ (in our case, *Leaky ReLU*) can be composed with the above projection, yielding the new function $\Psi(\mathcal{B}) = \sigma(\psi(\mathcal{B}))$. Note that, at this stage, $\Psi(\mathcal{B})$ can be interpreted as a set of coefficients to apply on the BDT basis $B(\mathcal{O})$ to best estimate \mathcal{B} . In other words, $\Psi(\mathcal{B})$ is not a BDT yet, but simply a set of coefficients, which can be used later to reconstruct a BDT. Thus, a second transformation needs to be considered, to transform the set of coefficients $\Psi(\mathcal{B})$ back into a BDT. Then, we define the notion of *BDT transformation layer*, noted $\Pi(\mathcal{B})$, as the composition $\Pi(\mathcal{B}) = \Psi^{out} \circ \Psi^{in}(\mathcal{B})$ (Fig. 7b):

$$\Psi^{in}(\mathcal{B}) = \sigma \left(\arg \min_{\alpha} \left(W_2^T(\mathcal{B}, \mathcal{O}^{in} + B^{in}(\mathcal{O}^{in})\alpha) \right)^2 \right)$$

$$\Psi^{out}(\alpha) = \gamma(\mathcal{O}^{out} + B^{out}(\mathcal{O}^{out})\alpha),$$

where $\gamma(\mathcal{B})$ is a projection which transforms \mathcal{B} into a *valid* BDT, i.e. which respects the *Elder rule* (Sec. A.2). Given a branch $b \in \mathcal{B}$, γ enforces that: $\gamma(b)_x < \gamma(b)_y$ and $[\gamma(b)_x, \gamma(b)_y] \subseteq [0, 1]$ (Fig. 8).

BDT transformation layers can be seen as *local auto-encoders* (Fig. 7b): the first step Ψ^{in} converts an input BDT into a set of coefficients with a basis projection and a non-linearity, while the second step Ψ^{out} converts these coefficients into a BDT. Note that each BDT transformation layer is associated with its own input *and* output d' -dimensional bases $B^{in}(\mathcal{O}^{in})$ and $B^{out}(\mathcal{O}^{out})$.




Fig. 8. Projection γ ensuring the Elder rule.




Fig. 7. **Overview:** given an input ensemble, with its merge trees and BDTs (a), our Wasserstein Auto-Encoder of Merge Trees (MT-WAE) optimizes an auto-encoder (for this example, $n_e = n_d = 2$) where each layer Π_k natively processes BDTs (without pre-vectorization). Specifically, each layer Π_k can be interpreted as a local auto-encoder (b), where an *input* sub-layer Ψ_k^{in} transforms the input BDT $\mathcal{B}_{k-1}(f_i)$ into a set of coefficients $\alpha_k^i \in \mathbb{R}^{d_k}$ and where an *output* sub-layer Ψ_k^{out} transforms these coefficients back into a valid BDT $\mathcal{B}_k(f_i)$. The aggregated views ((c), (d), (e), (f), (g)), which overlap all the BDTs in the birth/death space (one color per BDT), show the ability of MT-WAE to progressively *unwrap* non-linear structures (circles) as the BDTs progress down the network, resulting in faithful local parameterizations in latent space ((e), the individual angular parameterizations of the circles are well preserved), as well as accurate reconstructions (g). This native support of BDTs results in a superior accuracy (Sec. 6.2) and an improved interpretability: individual features can now be tracked as they traverse the network, enabling new visual analysis capabilities (Sec. 5.2).

The processing of an ensemble of BDTs by a BDT transformation layer is illustrated in Fig. 7. Specifically, Fig. 7c) shows a zoom of the birth-death space, where all the BDTs have been aggregated (one color per BDT, each BDT has two branches, hence two patterns appear, one circle per branch). The left inset of Fig. 7b) shows the non-linearly transformed set of BDTs after the first BDT transformation layer, Π_1 . Next, the first step Ψ_2^{in} of the next BDT layer Π_2 converts each BDT $\Pi_1(\mathcal{B}(f_i))$ into a set of coefficients α_2^i . Finally, the second step Ψ_2^{out} of the layer Π_2 converts each of these set of coefficients into a new, non-linearly transformed BDT $\Pi_2 \circ \Pi_1(\mathcal{B}(f_i))$ (Fig. 7b), right inset).

3.3 MT-WAE formulation

Now that the above geometrical tools have been introduced for the Wasserstein metric space of merge trees, we can now formulate MT-WAE by direct analogy to the Euclidean case (Sec. 3.1). Given a set $\mathcal{S}_{\mathcal{B}} = \{\mathcal{B}(f_1), \dots, \mathcal{B}(f_N)\}$ of input BDTs, a MT-WAE is a composition of BDT transformation layers $\Pi_k(\mathcal{B})$ (Fig. 7), minimizing the following energy:

$$E_{W_2^T} = \sum_{i=1}^N \left(W_2^T \left(\mathcal{B}(f_i), \Pi_{n_e+n_d} \circ \dots \circ \Pi_{n_e} \circ \dots \circ \Pi_1(\mathcal{B}(f_i)) \right) \right)^2, \quad (9)$$

where each BDT transformation layer Π_k is associated with its own d_k -dimensional input *and* output vector bases $B_k^{in}(\mathcal{O}_k^{in})$ and $B_k^{out}(\mathcal{O}_k^{out})$. Moreover, the dimensions of the successive bases are chosen such that:

- 1) $d_1 > d_2 > \dots > d_{n_e}$, and
- 2) $d_{n_e} < d_{n_e+1} < \dots < d_{n_e+n_d}$,

where n_e and n_d denote the number of *Encoding* and *Decoding* layers and where d_{n_e} is the dimension of the MT-WAE latent space. Eq. 9 is a direct analog to the classical EAE (Eq. 6): the standard transformation layers Ψ_k have been replaced by BDT transformation layers Π_k and the L_2 norm by the distance W_2^T .

Algorithm 1 Wasserstein Auto-Encoder (algorithm overview).

Input: Set of BDTs $\mathcal{S}_{\mathcal{B}} = \{\mathcal{B}(f_1), \dots, \mathcal{B}(f_N)\}$.	
Output1: Set of $(n_e + n_d)$ input origins $\theta_{\mathcal{O}^{in}} = \{\mathcal{O}_1^{in}, \mathcal{O}_2^{in}, \dots, \mathcal{O}_{n_e+n_d}^{in}\}$;	
Output2: Set of $(n_e + n_d)$ input bases $\theta_{\mathcal{B}^{in}} = \{B_1^{in}(\mathcal{O}_1^{in}), B_2^{in}(\mathcal{O}_2^{in}), \dots, B_{n_e+n_d}^{in}(\mathcal{O}_{n_e+n_d}^{in})\}$;	
Output3: For each of the input BDTs ($i \in \{1, 2, \dots, N\}$), set of $(n_e + n_d)$ input coefficients: $\theta_{\alpha}^i = \{\alpha_1^i \in \mathbb{R}^{d_1}, \alpha_2^i \in \mathbb{R}^{d_2}, \dots, \alpha_{n_e+n_d}^i \in \mathbb{R}^{d_{n_e+n_d}}\}$;	
Output4: Set of $(n_e + n_d)$ output origins $\theta_{\mathcal{O}^{out}} = \{\mathcal{O}_1^{out}, \mathcal{O}_2^{out}, \dots, \mathcal{O}_{n_e+n_d}^{out}\}$;	
Output5: Set of $(n_e + n_d)$ output bases $\theta_{\mathcal{B}^{out}} = \{B_1^{out}(\mathcal{O}_1^{out}), B_2^{out}(\mathcal{O}_2^{out}), \dots, B_{n_e+n_d}^{out}(\mathcal{O}_{n_e+n_d}^{out})\}$;	
<hr/>	
1: $\theta \leftarrow \{\theta_{\mathcal{O}^{in}}, \theta_{\mathcal{B}^{in}}, \theta_{\mathcal{O}^{out}}, \theta_{\mathcal{B}^{out}}\}$;	// Overall set of optimization variables.
2: Initialize(θ);	// Initialization of the optimization variables (Sec. 4.3).
3: while $E_{W_2^T}(\theta)$ decreases do	
4: $\widehat{\mathcal{S}}_{\mathcal{B}} \leftarrow \text{Forward}(\mathcal{S}_{\mathcal{B}}, \theta)$;	// Forward propagation of the BDT ensemble $\mathcal{S}_{\mathcal{B}}$ (Sec. 4.4).
5: $\theta \leftarrow \text{Backward}(\mathcal{S}_{\mathcal{B}}, \widehat{\mathcal{S}}_{\mathcal{B}})$;	// Backward propagation (Sec. 4.5).
6: end while	

Fig. 7 illustrates a network of BDT transformation layers optimized on a synthetic ensemble (our optimization algorithm is described in Sec. 4). As mentioned in the previous section, each input BDT (spheres in the aggregated birth-death views, one color per BDT) is non-linearly transformed by the BDT transformation layers (see for instance Fig. 7d) and Fig. 7f). As a result, in this example, the BDT transformation layers progressively *unwrap* the non-linear structures in the birth-death plane (circles in Fig. 7c) as the BDTs traverse the network down to the latent space (Fig. 7e), where the resulting layout in the birth-death plane (line segments) manages to faithfully encode the purposely designed parameterization of the ensemble: the order of rotation angles (colors) is well preserved along the segments in latent space.

4 ALGORITHM

This section presents our algorithm for the minimization of Eq. 9.

4.1 Overview

Alg. 1 provides an overview of our main algorithm. The set of optimization variables, noted θ and declared line 1, includes the

$(n_e + n_d)$ input and output BDT bases, along with their origins. The optimization of these variables follows the standard overall procedure for the optimization of a neural network.

(1) **Initialization:** First, θ is initialized, as detailed in Sec. 4.3.

(2) **Forward propagation:** Then, the input ensemble of BDTs \mathcal{S}_B , illustrated in Fig. 7a, traverses the network to generate a reconstructed ensemble of BDTs, noted $\widehat{\mathcal{S}}_B$. This occurs line 4 of Alg. 1 and it is illustrated from Fig. 7c to Fig. 7g with the dark blue arrows. This is traditionally denoted as the *Forward* propagation. This part of our approach is documented in Sec. 4.4.

(3) **Backward propagation:** Given $\widehat{\mathcal{S}}_B$, the energy $E_{W_2^T}(\theta)$, detailed in Eq. 9, can be evaluated and its gradient $\nabla E_{W_2^T}(\theta)$ can be estimated by automatic differentiation, based on the application of the chain-rule on the composition of BDT transformation layers. Given the gradient $\nabla E_{W_2^T}(\theta)$, a step of gradient descent can be achieved to update the optimization variables θ . This occurs line 5 of Alg. 1. This is traditionally denoted as the *Backward* propagation. This part of our approach is detailed in Sec. 4.5.

(4) **Stopping condition:** The two steps of forward and backward propagations are then iterated until the energy stops decreasing, in practice until it decreases by less than 1% between two iterations.

4.2 Basis projection

We start by describing an efficient *Assignment/Update* algorithm for the projection of a BDT \mathcal{B} into a BDT basis $B(\mathcal{O})$ (Eq. 8), as it is a core geometrical component used throughout our approach.

The purpose of the projection (Eq. 8) is to find a set of coefficients $\alpha \in \mathbb{R}^{d'}$ to apply on the d' BDT vectors of $B(\mathcal{O})$ to best estimate \mathcal{B} . The geodesic analysis of merge trees [79] faces a similar issue, but its formulation (restricting α to $[0, 1]^{d'}$) allows for an iterative, brute-force optimization. Here, we introduce a more general and efficient strategy.

(1) **Assignment step:** Let us assume that we are given an initial set of coefficients α . Then, the *estimation* $\widehat{\mathcal{B}}$ of \mathcal{B} is given by $\widehat{\mathcal{B}} \leftarrow \mathcal{O} + B(\mathcal{O})\alpha$. The purpose of the assignment step is to refine the evaluation of $W_2^T(\mathcal{B}, \widehat{\mathcal{B}})$. For this, we first compute the optimal assignment¹ ϕ_* between \mathcal{B} and $\widehat{\mathcal{B}}$, w.r.t. Eq. 1. Then, $W_2^T(\mathcal{B}, \widehat{\mathcal{B}})$ can be re-written as:

$$W_2^T(\mathcal{B}, \widehat{\mathcal{B}}) = \sum_{i=1}^{|\mathcal{B}|} \|b_i - \phi_*(b_i)\|_2^2. \quad (10)$$

(2) **Update step:** Given the above estimation $\widehat{\mathcal{B}}$, the goal of the update step is to improve the coefficients α , in order to decrease $W_2^T(\mathcal{B}, \widehat{\mathcal{B}})$. Let $\widehat{\mathcal{B}}'$ be a vector representation of $\widehat{\mathcal{B}}$. Specifically, $\widehat{\mathcal{B}}'$ is a vector in $\mathbb{R}^{2|\widehat{\mathcal{B}}|}$ which concatenates the coordinates in the 2D birth/death plane of each branch b_i of $\widehat{\mathcal{B}}$. $\widehat{\mathcal{B}}'$ can be decomposed into $\mathcal{O}' + (B(\mathcal{O}'))'\alpha$, where $(B(\mathcal{O}'))'$ is a $(2|\widehat{\mathcal{B}}|) \times d'$ matrix. Additionally, let \mathcal{B}' be a similar vector representation of \mathcal{B} , but where the entries have been specifically re-ordered such that, for each of its 2D entries, we have:

$$(\mathcal{B}')_i = \phi_*^{-1}((\widehat{\mathcal{B}}')_i). \quad (11)$$

Intuitively, \mathcal{B}' is a re-ordered vector representation of \mathcal{B} , such that its i^{th} entry exactly matches though ϕ_* with the i^{th} entry of $\widehat{\mathcal{B}}'$. Given this vector representation, the Wasserstein distance

$W_2^T(\mathcal{B}, \widehat{\mathcal{B}})$ for a fixed optimal assignment ϕ_* (Eq. 10) can then be re-written as an L_2 norm:

$$W_2^T(\mathcal{B}, \widehat{\mathcal{B}}) = \|\mathcal{B}' - \widehat{\mathcal{B}}'\|_2^2. \quad (12)$$

Then, given the optimal assignment ϕ_* , the optimal $\alpha_* \in \mathbb{R}^{d'}$ are:

$$\begin{aligned} \alpha_* &= \arg \min_{\alpha} \|\mathcal{B}' - \widehat{\mathcal{B}}'\|_2^2 \\ \alpha_* &= \arg \min_{\alpha} \|\mathcal{B}' - (\mathcal{O}' + (B(\mathcal{O}'))'\alpha)\|_2^2 \\ \alpha_* &= \arg \min_{\alpha} \|\mathcal{B}' - \mathcal{O}' - (B(\mathcal{O}'))'\alpha\|_2^2. \end{aligned}$$

Similarly to the Euclidean case (Eq. 4), it follows then that α_* can be expressed as a function of the pseudoinverse of $(B(\mathcal{O}'))'$:

$$\alpha_* = \left((B(\mathcal{O}'))' \right)^+ (\mathcal{B}' - \mathcal{O}'). \quad (13)$$

At this stage, the estimation $\widehat{\mathcal{B}}$ can be *updated* with the above optimized coefficients α_* : $\widehat{\mathcal{B}} \leftarrow \mathcal{O} + B(\mathcal{O})\alpha_*$.

The above *Assignment/Update* sequence is then iterated. Each iteration decreases the *projection error* $e(\mathcal{B})$ constructively: while the *Update* phase (2) optimizes α (Eq. 7) to minimize the projection error under the current assignment ϕ_* , the next *Assignment* phase (1) further improves (by construction) the assignments (term W_2^T in Eq. 7), hence decreasing the projection error overall. In our implementation, this iterative algorithm stops after a predefined number of iterations n_{it} .

4.3 Initialization

Now that we have introduced the core low-level procedure of our approach (Sec. 4.2), we can detail the initialization step of our framework, which consists in identifying a relevant initial value for the overall optimization variable θ (line 2, Alg. 1). The BDT transformation layers Π_k are initialized one after the other, i.e. for increasing values of k .

(1) **Input initialization:** For each BDT transformation layer Π_k , its input origin $\mathcal{O}_k^{\text{in}}$ is initialized as the Wasserstein barycenter \mathcal{B}_* [78] of the BDTs on its input. Next, the first vector of B_k^{in} , is given by the optimal assignment (w.r.t. Eq. 1) between $\mathcal{O}_k^{\text{in}}$ and the layer's input BDT \mathcal{B} which maximizes $W_2^T(\mathcal{O}_k^{\text{in}}, \mathcal{B})$, i.e. which induces the worst projection error $e(\mathcal{B})$ (Eq. 7) given an empty basis. Next, the remaining $(d_k - 1)$ vectors of B_k^{in} are initialized one after the other, by including at each step the vector formed by the optimal assignment between $\mathcal{O}_k^{\text{in}}$ and the layer's input BDT \mathcal{B} which induces the maximum projection error $e(\mathcal{B})$ (Eq. 7), given the already initialized vectors. Note that this step makes an extensive usage of the projection procedure introduced in Sec. 4.2. Finally, if the dimension d_k of Π_k is greater than the number of input BDTs, the remaining vectors are initialized randomly, with a controlled norm (set to the mean of the already initialized vectors).

(2) **Output initialization:** For each BDT transformation layer Π_k , its output origin $\mathcal{O}_k^{\text{out}}$ and basis B_k^{out} are initialized as random linear transformations of its input origin and basis. Specifically, let W be a random matrix of size $(2|\mathcal{O}_k^{\text{out}}| \times 2|\mathcal{O}_k^{\text{in}}|)$. Given the vector representation $\mathcal{O}_k^{\text{in}'}$ of $\mathcal{O}_k^{\text{in}}$ (see Sec. 4.2), we initialize $\mathcal{O}_k^{\text{out}'}$ such that: $\mathcal{O}_k^{\text{out}'} \leftarrow W\mathcal{O}_k^{\text{in}'}$. Similarly, the output basis of Π_k is initialized such that: $B_k^{\text{out}'} \leftarrow WB_k^{\text{in}'}$.

4.4 Forward propagation

Alg. 2 presents the main steps of our forward propagation. This procedure follows directly from our formulation (Sec. 3.2). Each

1. This discussion describes the case where ϕ_* is a bijection between off-diagonal points of the 2D birth/death plane. Appendix B details the general case, where ϕ_* may send points of \mathcal{B} to the diagonal of $\widehat{\mathcal{B}}$, and reciprocally.

Algorithm 2 Forward propagation in our Wasserstein Auto-Encoder.

Input1: Set of input BDTs $\mathcal{S}_B = \{\mathcal{B}(f_1), \dots, \mathcal{B}(f_N)\}$.
Input2: Current value of the overall optimization variable θ .
Output: Set of reconstructed BDTs $\widehat{\mathcal{S}}_B = \{\widehat{\mathcal{B}}(f_1), \dots, \widehat{\mathcal{B}}(f_N)\}$.

```

1: for  $\mathcal{B} \in \mathcal{S}_B$  do
2:    $\mathcal{B}_0 \leftarrow \mathcal{B}$ .
3:   for  $k \in \{1, 2, \dots, n_e + n_d\}$  do
4:     // For each BDT transformation layer  $\Pi_k$ .
5:      $\Psi_k^{in}(\mathcal{B}_{k-1}) \leftarrow \text{basisProjection}(\mathcal{B}_{k-1}, \mathcal{O}_k^{in}, B_k^{in})$ . // Sec. 4.2.
6:      $\Psi_k^{in}(\mathcal{B}_{k-1}) \leftarrow \sigma(\Psi_k^{in}(\mathcal{B}_{k-1}))$ . // Sec. 3.2.
7:      $\mathcal{B}_k \leftarrow \Psi_k^{out}(\Psi_k^{in}(\mathcal{B}_{k-1})) = \gamma(\mathcal{O}_k^{out} + B_k^{out}(\mathcal{O}_k^{out})\Psi_k^{in}(\mathcal{B}_{k-1}))$ . // Sec. 3.2.
8:   end for
9:    $\widehat{\mathcal{S}}_B \leftarrow \widehat{\mathcal{S}}_B \cup \mathcal{B}_{n_e+n_d}$ 
10: end for

```

input BDT $\mathcal{B} \in \mathcal{S}_B$ is processed independently (line 1). Specifically, \mathcal{B} will traverse the network one layer Π_k at a time (line 3). Within each layer Π_k , the projection through the input sub-layer Ψ_k^{in} is computed (line 6) by composing a non-linearity σ with the basis projection (Sec. 3.2). This yields a set of coefficients representing the input BDT. Next, following Sec. 3.2, these coefficients are transformed back into a valid BDT with the output sub-layer (line 7). At the end of this process, a set of reconstructed BDTs $\widehat{\mathcal{S}}_B$ is available, for a fixed value of θ .

4.5 Backward propagation

Given the set of reconstructed BDTs $\widehat{\mathcal{S}}_B$ for the current value of θ (Sec. 4.4), the data fitting energy (Eq. 9) is evaluated. Specifically, for each input BDT $\mathcal{B} \in \mathcal{S}_B$, the optimal assignment ϕ_* w.r.t. Eq. 1 is computed between \mathcal{B} and its reconstruction, $\mathcal{B}_{n_e+n_d}$, provided on the output of the network. Next, similarly to Sec. 4.2 for basis projections, the vector representation B' of \mathcal{B} is constructed and re-ordered such that the i^{th} entry of this vector corresponds to the pre-image by ϕ_* of the i^{th} entry of $B'_{n_e+n_d}$ (c.f. Eq. 19). Then, given the optimal assignment ϕ_* , similarly to Sec. 4.2, $W_2^T(\mathcal{B}, \mathcal{B}_{n_e+n_d})$ can be expressed as an L_2 norm (Eq. 12). Given the set Φ_* of all the optimal assignments between the input BDTs and their output reconstructions, $E_{W_2^T}(\theta)$ is then evaluated:

$$E_{W_2^T}(\theta) = \sum_{i=1}^N \|\mathcal{B}(f_i)' - \mathcal{B}_{n_e+n_d}(f_i)'\|_2^2. \quad (14)$$

At this stage, for a given set Φ_* of optimal assignments, the evaluation of Eq. 14 only involves basic operations (as described in the previous sections: vector re-orderings, pseudoinverse computations, linear transformations, and compositions). All these operations are supported by the automatic differentiation capabilities of modern neural frameworks (in our case *PyTorch* [76]), enabling the automatic estimation of $\nabla E_{W_2^T}(\theta)$. Then, θ is updated by gradient descent [57].

Our overall optimization algorithm (Alg. 1) can be interpreted as a global instance of an *Assignment/Update* strategy. Each backward propagation updates the overall variable θ to improve the data fitting energy (Eq. 9), while the next forward propagation improves the network outputs and hence their assignments to the inputs. In the remainder, the terms PD-WAE and MT-WAE refer to the usage of our framework with persistence diagrams or merge trees respectively. We refer the reader to the Appendix C for a detailed discussion of the meta-parameters of our approach (e.g. layer number, layer dimensionality, etc).




Fig. 9. Application of our merge tree compression to feature tracking (experiment adapted from [78], [79] for comparison purposes). The 5 most persistent maxima (spheres) of three time steps (ion density during universe formation [73]) are tracked through time (left, f_1 , f_2 and f_3) by considering the optimal assignment (Eq. 1) between the corresponding merge trees (inset, left). The same tracking procedure is applied to the merge trees compressed by WAE (inset, right). Similarly to PGA [79], the resulting tracking is identical with the compressed trees. However, in comparison to PGA [79], for a target compression factor of 13.44, the relative reconstruction error (right) is clearly improved with WAE.

5 APPLICATIONS

This section illustrates the utility of our framework in concrete visualization tasks: merge tree compression and dimensionality reduction. These applications and use-cases are adapted from [79], to facilitate comparisons between previous work on the linear encoding of merge trees [79] and our novel non-linear framework.

5.1 Merge tree compression

As discussed by Pont et al. [79], like any data representation, merge trees can benefit from lossy compression. For example, for the in-situ analysis of high-performance simulations [6], each individual time-step of the simulation can be represented and stored to disk in the form of a topological descriptor [20]. In this context, this lossy compression eases the manipulation of the generated ensemble of topological descriptors (i.e. it facilitates its storage and transfer). Previous work has investigated the compression of an ensemble of merge trees via linear encoding [79]. In this section, we improve this application by extending it to non-linear encoding, thereby enabling more accurate compressions. Specifically, the input ensemble \mathcal{S}_B of BDTs is compressed, by only storing to disk:

- (1) the output sub-layer of the last decoding layer of the network, noted $\Psi_{n_e+n_d}^{out}$ (i.e. its origin, $\mathcal{O}_{n_e+n_d}^{out}$, as well as its basis, $B_{n_e+n_d}^{out}(\mathcal{O}_{n_e+n_d}^{out})$)
- (2) the corresponding N BDT coefficients $\alpha_{n_e+n_d}^i \in \mathbb{R}^{d_{n_e+n_d}}$.

Note that an alternative compression strategy would consist in storing the N BDT coefficients in latent space directly (i.e. $\alpha_{n_e}^i \in \mathbb{R}^{d_{n_e}}$), which would be typically more compact than the N BDT coefficients in the last output sub-layer ($\alpha_{n_e+n_d}^i \in \mathbb{R}^{d_{n_e+n_d}}$). However, in order to decompress this representation, one would need to store to disk the entire set of n_d decoding layers. This significant overhead would only be compensated for ensembles counting an extremely large number N of members. In our experiments (Sec. 6), $N = 48$ for the largest ensemble. Thus, we focus on the first strategy described above (storing potentially larger sets of coefficients, but a smaller number of decoding layers).

The compression factor can be controlled with two input parameters: (i) $d_{n_e+n_d}$ controls the dimensionality of the last decoding layer (hence its ability to capture small variabilities) and (ii) $|\mathcal{O}_{n_e+n_d}^{out}|$ controls the size of the origin of the last decoding sub-layer (hence its ability to capture small features). The resulting reconstruction error (Eq. 9) will be minimized for large values of

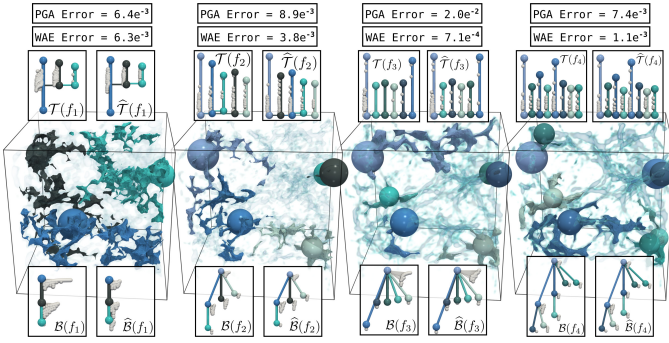


Fig. 10. Application of our merge tree compression to topological clustering (experiment adapted from [78], [79] for comparison purposes). The ensemble is clustered [78] based on the merge trees (top left insets) of the ensemble members (each column represents a member from one of the 4 clusters). The same clustering procedure [78] is applied to the merge trees compressed by WAE (top right insets). Similarly to PGA [79], the resulting clustering is identical with the compressed trees (it exactly matches the ground-truth classification [78]). However, in comparison to PGA [79], for a target compression factor of 15.09, the relative reconstruction error (top) is clearly improved with WAE. Visually, the compressed trees look very similar to the original ones: the prominent features (in colors, non-prominent features are shown with small white nodes) are well preserved in terms of number and persistence. The same holds for the compressed BDTs (bottom right insets) which are nearly isomorphic to the original BDTs (bottom left insets, the color shows the assignment between the original and compressed BDTs).

both parameters, while the compression factor will be minimized for low values. In the following experiments, we set both parameters to their default values (Appendix C). To decompress a BDT $\mathcal{B}(f_i)$, its stored coefficients $\alpha_{n_e+n_d}^i$ are simply propagated through the stored output sub-layer of the network, $\Psi_{n_e+n_d}^{out}$.

Figs. 9 and 10 show two examples of visualization tasks (feature tracking and ensemble clustering, use cases replicated from [78], [79] for comparison purposes). In these experiments, the BDTs have been compressed with the strategy described above. Next, the de-compressed BDTs have been used as an input to these two analysis pipelines. In both cases, the output obtained with the de-compressed BDTs is identical to the output obtained with the *original* BDTs. This shows the viability of the de-compressed BDTs and it demonstrates the utility of this compression scheme.

5.2 Dimensionality reduction

This section describes how to use MT-WAE to generate 2D layouts of the ensemble, for the global visual inspection of the ensemble. This is achieved by setting $d_{n_e} = 2$ and by embedding each BDT $\mathcal{B}(f_i)$ as a point in the plane, at its latent coordinates $(\alpha_{n_e}^i)_1$ and $(\alpha_{n_e}^i)_2$. This results in a summarization view of the ensemble, grouping similar BDTs together (Fig. 1c). The flexibility of our framework allows to further improve the quality of this 2D layout. Specifically, Appendix D introduces two penalty terms aiming at (1) improving the preservation of the Wasserstein metric W_2^T and (2) improving the preservation of the clusters of BDTs.

We augment our 2D layouts with *Persistence Correlation Views* (PCV) which were introduced in [79]. In short, the PCV embeds a branch b of the barycenter \mathcal{B}_* [78] as a point in 2D, in order to represent the variability of the corresponding feature in the ensemble, as a function of the coordinates in latent space. Specifically, the optimal assignments ϕ_{*i} between \mathcal{B}_* and each input BDT $\mathcal{B}(f_i)$ is first computed (Eq. 1). Next, for a given

branch $b \in \mathcal{B}_*$, the Pearson correlation $\rho(p_{b_i}, (\alpha_{n_e}^i)_1)$ between the persistence p_{b_i} of $\phi_{*i}(b) \in \mathcal{B}(f_i)$ and the first coordinate in latent space $(\alpha_{n_e}^i)_1$ is computed for the ensemble (i.e. for $i \in \{1, 2, \dots, N\}$). Next, the Pearson correlation $\rho(p_{b_i}, (\alpha_{n_e}^i)_2)$ is computed similarly with regard to the second coordinate in latent space $(\alpha_{n_e}^i)_2$. Finally, b is embedded in the PCV at the coordinates $(\rho(p_{b_i}, (\alpha_{n_e}^i)_1), \rho(p_{b_i}, (\alpha_{n_e}^i)_2))$. To avoid clutter in the visualization, we only report the most persistent branches of \mathcal{B}_* in the PCV. Intuitively, points in the PCV which are located far away from the center, along a given direction, indicate a strong correlation between that direction in latent space, and the persistence of the corresponding feature in the ensemble.

PCVs enable the identification of patterns of feature variability within the ensemble, as discussed in Fig. 11. This case study considers the *Isabel* ensemble, which consists of 12 scalar fields representing the wind velocity magnitude in a hurricane simulation. The ensemble comes with a ground-truth classification [78]: 4 members correspond to the formation of the hurricane (e.g. f_1 , Fig. 11), 4 other members to its drift (e.g. f_2 , Fig. 11) and 4 other members to its landfall (e.g. f_3 , Fig. 11). For this ensemble, our PD-WAE approach produces a 2D layout (Fig. 11b) which manages to recover the temporal coherency of the ensemble: the formation (dark red), drift (pink) and landfall (light pink) clusters are arranged in order along a line (direction $(1, -1)$). This shows the ability of PD-WAE to recover the intrinsic structure of the ensemble (here its temporal nature). The PCV (grey inset) further helps appreciate the variability of the features in the ensemble. There, each colored point indicates a persistent feature of the barycenter: the eye of the hurricane is represented by the blue sphere, while the cyan, black and white sphere represent peripheral gusts of wind (see the matching features in the data, Fig. 11a). The PCV clearly identifies two patterns of feature variability, along the direction $(1, -1)$, which coincides with the temporal alignment of the clusters in latent space. Specifically, it indicates that the persistence of the hurricane eye will be larger in the top left corner of the latent space, i.e. towards the beginning of the temporal sequence (dark red cluster). This is confirmed visually when inspecting the persistence diagrams of the individual members (the blue feature is less persistent in $\mathcal{D}(f_3)$ than in $\mathcal{D}(f_1)$ and $\mathcal{D}(f_2)$). In short, this visually encodes the fact that the strength of the hurricane eye decreases with time. In contrast, the features corresponding to peripheral wind gusts (cyan, black and white spheres) exhibit a common variability pattern, distinct from that of the hurricane eye: the persistence of the corresponding features increases as one moves along the direction $(1, -1)$ in latent space, i.e. as time increases (pink and light pink clusters). This is confirmed visually in the individual members, where the persistence of these features is larger in $\mathcal{D}(f_3)$ than in $\mathcal{D}(f_1)$ and $\mathcal{D}(f_2)$. In short, this visually encodes the fact that the strength of the peripheral wind gusts increases with time. Overall, while the 2D layout generated by PD-WAE enables the visualization of the intrinsic structure of the ensemble (here, its temporal nature), the PCV enables the visualization and interpretation of the variability in the ensemble at a feature level. The caption of Fig. 1 includes a similar discussion for MT-WAE.

Fig. 12 provides a qualitative comparison of the 2D layouts and Persistence Correlation Views (PCVs) between PD-PGA [79] and our novel non-linear framework, PD-WAE (see Sec. 6.2 for an extensive quantitative comparison). Specifically, it shows that, while PD-PGA (Fig. 12a) manages to isolate the clusters well,




Fig. 11. Visual analysis of the *Isabel* ensemble (one member per ground-truth class, top), with PD-WAE (in this example, $n_e = n_d = 2$). Our work enables diagram compression, while providing reconstructed diagrams ((a), top insets) which are highly similar to the input (below). The 2D layout generated by PD-WAE (b) recovers the temporal structure of the ensemble (the clusters are aligned in chronological order, from dark red to light pink). The PCV (grey inset) indicates that the peripheral gusts of wind in the data (white, black and cyan spheres) grow in importance when moving towards the bottom right corner of the latent space (their persistence increases over time). In contrast, the hurricane eye (light blue) exhibits less variability, but with stronger values towards the top left corner of the latent space (start of the sequence). The aggregated views (bottom, overlapping all diagrams, transparent: barycenter) for the input (c) and the latent space (d) show that PD-WAE nicely recovers a per-feature parameterization in the latent space (one ellipse per barycenter feature), which is locally consistent with the data temporal evolution (cluster color).




Fig. 12. Qualitative comparison of the 2D layouts (left, white inset) and PCVs (right, grey inset) between PD-PGA [79] (a) and PD-WAE (b) for the *Isabel* ensemble. PD-WAE manages to recover the intrinsic temporal structure of the ensemble and produces a linear alignment of the clusters in order of temporal appearance (dark red, pink, light pink).

its 2D layout does not recover the intrinsic, one-dimensional, temporal structure of the ensemble. In contrast, as discussed above, PD-WAE (Fig. 12b) manages to recover this intrinsic structure and produces a linear alignment of the clusters along the direction $(1, -1)$, in order of their temporal appearance. This alignment greatly facilitates the interpretation of the PCV, since time is now visually encoded there by the linear direction $(1, -1)$ (whereas it would be encoded by a curve in the case of PD-PGA).

In contrast to standard auto-encoders, our approach explicitly manipulates topological descriptors throughout the network. This results in improved interpretability and enables new capabilities:

(1) **Latent feature transformation:** As discussed in Fig. 11, it is now possible to visualize how topological descriptors are (non-linearly) transformed by the auto-encoder. Specifically, the

aggregated views of the birth/death space (bottom) illustrate how PD-WAE *unwraps* the diagrams in latent space, nicely recovering the data temporal evolution at a feature level (see the temporally consistent linear arrangements of points for each barycenter feature in Fig. 11d).

(2) **Latent space navigation:** Given a point in latent space, it is now possible to efficiently reconstruct its BDT/MT by propagating its latent coordinates through the decoding layers, enabling an interactive exploration of the merge tree latent space (Fig. 1d).

(3) **Feature traversal analysis:** For each consecutive layers Π_k and Π_{k+1} , we compute the optimal assignment (Eq. 1) between their input origins, $\mathcal{O}_k^{\text{in}}$ and $\mathcal{O}_{k+1}^{\text{in}}$. Next, we compute the optimal assignment between the barycenter \mathcal{B}_* [78] of the input ensemble and the first origin $\mathcal{O}_1^{\text{in}}$. This yields an explicit tracking of each branch b of \mathcal{B}_* down to the latent space. We introduce the notion of *Feature Latent Importance* (FLI), given by the persistence of b in latent space, divided by its original persistence. FLI indicates if a feature gains (or loses) importance in latent space. This enables the identification of the most *informative* features in the ensemble.

This is illustrated in Fig. 1e), where the cyan, white, dark blue and black features exhibit large FLI values (red circles). In the trees, these features are indeed present in most of the ensemble (Fig. 1d), with only moderate variations in persistence. Interestingly, the global maximum of seismic wave (light blue feature) is not a very informative feature (light blue circle): while it is also present throughout the sequence, its persistence decreases significantly (left to right). This visually encodes that, as the seismic wave travels from the epicenter, the strength of its global

TABLE 1

Running times (in seconds) of our algorithm for PD-WAE and MT-WAE computation (first sequential, then with 20 cores).

Dataset	N	B	PD-WAE			MT-WAE		
			1 c.	20 c.	Speedup	1 c.	20 c.	Speedup
Asteroid Impact (3D)	7	1,295	2,819.38	989.42	2.85	5,946.81	1,522.89	3.90
Cloud processes (2D)	12	1,209	11,043.20	1,318.63	8.37	16,150.90	2,566.98	6.29
Viscous fingering (3D)	15	118	1,345.31	268.78	5.01	3,727.64	417.31	8.93
Dark matter (3D)	40	316	125,724.00	10,141.30	12.40	135,962.00	8,051.02	16.89
Volcanic eruptions (2D)	12	811	4,925.15	638.58	7.71	4,151.04	449.14	9.24
Ionization front (2D)	16	135	627.04	95.31	6.58	1,140.44	144.57	7.89
Ionization front (3D)	16	763	17,285.10	1,757.71	9.83	101,788.00	5,350.46	19.02
Earthquake (3D)	12	1,203	14,272.10	2,074.52	6.88	9,888.68	1,024.45	9.65
Isabel (3D)	12	1,338	3,485.03	436.56	7.98	23,240.90	1,669.20	13.92
Starting Vortex (2D)	12	124	215.46	95.45	2.26	281.80	147.51	1.91
Sea Surface Height (2D)	48	1,787	41,854.70	2,901.39	14.43	222,594.00	13,540.20	16.44
Vortex Street (2D)	45	23	460.46	152.35	3.02	117.88	38.88	3.03

maximum is no longer significant in front of other local maxima (which illustrates the energy diffusion process).

6 RESULTS

This section presents experimental results obtained on a computer with two Xeon CPUs (3.2 GHz, 2x10 cores, 96GB of RAM). The input merge trees were computed with FTM [39] and pre-processed to discard noisy features (persistence simplification threshold: 0.25% of the data range). We implemented our approach in C++ (with OpenMP and PyTorch’s C++ API [76]), as modules for TTK [13], [93]. Experiments were performed on a set of 12 public ensembles described in [78], which includes a variety of simulated and acquired 2D and 3D ensembles extracted from previous work and past SciVis contests [73].

6.1 Time performance

The Wasserstein distance computation (Eq. 1) is the most expensive sub-procedure of our approach. It intervenes during energy evaluation (Sec. 4.5) but also at each iteration of basis projection (Sec. 4.2), itself occurring at each propagation iteration (Sec. 4.4), for each input BDT. To compute this distance, we use a fine-grain task-based parallel algorithm [78]. We leverage further parallelism to accelerate the process. Specifically, for each input BDT, its forward propagation (Sec. 4.4) can be run in a distinct parallel task. Similarly, when evaluating the overall energy (Sec. 4.5), the distance between an input BDT and its output reconstruction is computed in a distinct parallel task for each BDT. Tab. 1 evaluates the time performance of our framework for persistence diagrams (PD-WAE) and merge trees (MT-WAE). In sequential mode, the computation time is a function of the ensemble size (N), the tree sizes ($|\mathcal{B}|$), and the network size ($n_e + n_d$). Specifically, our approach computes, for each optimization iteration, $n_{it} \times (n_e + n_d) \times N$ Wasserstein distances, each of which requiring $\mathcal{O}(|\mathcal{B}|^2)$ steps in practice. In parallel, the iterative nature of our approach (Alg. 1) challenges parallel efficiency. However, timings are still improved after parallelization (orders of minutes on average), with a very good parallel efficiency for the largest ensembles (up to 95%).

6.2 Framework quality

Figs. 9 and 10 report compression factors for our application to merge tree compression (Sec. 5.1). These are ratios between the storage size of the input N BDTs and that of their compressed form. For a fixed target compression factor, WAE clearly improves the reconstruction error over linear encoding (PGA [79]). Appendix E extends this error comparison to all our test ensembles, and shows that, for identical compression factors, our framework improves the reconstruction error over PGA [79] by 37% for




Fig. 13. Comparison of planar layouts for typical dimensionality reduction techniques, on two merge tree ensembles. The color encodes the classification ground-truth [78]. For each quality score, the best value appears bold and the rank of the score is shown in parenthesis. This experimental protocol is adapted from [79] for comparison purposes.

TABLE 2

Comparison of *aggregated* layout quality scores (i.e. averaged over all merge tree ensembles, bold: best values). WAE-MC provides both superior metric (SIM) and cluster (NMI/ARI) preservation to pre-existing techniques (MDS [58], t-SNE [96], PGA [79], VEC-EAE).

Indicator	MDS	t-SNE	PGA	VEC-EAE	WAE	WAE-M	WAE-C	WAE-MC
NMI	0.78	0.83	0.82	0.71	0.84	0.76	0.96	0.87
ARI	0.68	0.75	0.74	0.55	0.77	0.63	0.95	0.82
SIM	0.86	0.75	0.79	0.74	0.78	0.87	0.78	0.86

persistence diagrams, and 52% for merge trees, hence confirming the accuracy superiority of WAE over PGA.

Fig. 13 provides a visual comparison for the planar layouts generated by a selection of typical dimensionality reduction techniques, applied on the input merge tree ensemble (i.e. each point is a merge tree). This experiment is adapted from [79] for comparison purposes. This figure reports quantitative scores. For a given technique, to quantify its ability to preserve the *structure* of the ensemble (i.e. its organization into ground-truth classes), we run k -means in the 2D layouts and evaluate the quality of the resulting clustering (given the ground-truth [78]) with the normalized mutual information (NMI) and adjusted rand index (ARI). To quantify its ability to preserve the *geometry* of the ensemble (i.e. to preserve its disposition within the Wasserstein metric space \mathbb{B}), we report the metric similarity indicator SIM [79], which evaluates the preservation of the Wasserstein metric W_2^T . All these scores vary between 0 and 1, with 1 being optimal.

MDS [58] and t-SNE [96] have been applied on the distance matrix of the input merge trees (Wasserstein distance, Eq. 1, default parameters [78]). By design, MDS preserves well the metric W_2^T (good SIM), at the expense of mixing ground-truth classes together (low NMI/ARI). t-SNE behaves symmetrically (higher NMI/ARI, lower SIM). We applied PGA [79] by setting the origin size parameter to a value compatible to our latent space $(0.1|\mathcal{S}_{\mathcal{B}}|$, Appendix C). As expected, PGA provides a trade-off between the extreme behaviors of MDS and t-SNE, with an improved cluster preservation over MDS (NMI/ARI), and an improved metric preservation over t-SNE (SIM). WAE also constitutes a trade-off between MDS and t-SNE, but with improved quality scores over PGA.

We compare our approach to a standard auto-encoder (EAE, Sec. 3.1) applied on the following vectorization of the input merge trees. Note that several vectorizations of persistence diagrams have been studied [2], [22], [56]. However we focus in this work on merge trees and only few vectorizations have been documented for these [61]. Hence we focus on the following strategy, inspired by [61]. Each input BDT $\mathcal{B}(f_i)$ is embedded in $\mathbb{R}^{2|\mathcal{B}^*|}$, such that the j^{th} entry of this vector corresponds to the birth/death location of

the branch of $\mathcal{B}(f_i)$ which maps to the j^{th} branch of the barycenter \mathcal{B}_* [78]. Next, we feed these vectorizations to an EAE, with the same meta-parameters as our approach (i.e. number of layers, dimensionality per layer). The corresponding results appear in the VEC-EAE column. Our approach (WAE) outperforms this straightforward application of EAE, with clearly higher clustering scores (NMI/ARI) and improved metric scores (SIM).

Fig. 13 also reports the layouts obtained with our approach after enabling the metric penalty term (WAE-M), the clustering penalty term (WAE-C) and both (WAE-MC), c.f. Appendix D. WAE-M (respectively WAE-C) significantly improves the metric (respectively cluster) preservation over MDS (respectively t-SNE). The combination of the two terms (WAE-MC) improves both quality scores *simultaneously*: it outperforms MDS (SIM) *and* it improves t-SNE (NMI/ARI). In other words, WAE-MC improves established methods by outperforming them on their dedicated criterion (SIM for MDS, NMI/ARI for t-SNE).

Appendix F extends our visual analysis to all our test ensembles. Tab. 2 also extends our quantitative analysis to all our test ensembles. It confirms the clear superiority of WAE over VEC-EAE. It also confirms that the combination of our penalty terms (WAE-MC) provides the best metric (SIM) and cluster (NMI/ARI) scores over existing techniques.

6.3 Limitations

As discussed in Sec. 2.2, the parameter ε_1 of the Wasserstein distance between merge trees (W_2^T) acts as a control knob, that balances the practical stability of the metric with its discriminative power. Specifically, for $\varepsilon_1 = 1$, we have $W_2^T = W_2^D$ and W_2^T is stable, but less discriminative. Pont et al. [78] showed experimentally that for relatively low values of ε_1 (0.05), W_2^T still behaved in a stable manner in practice for reasonable noise levels. Our overall MT-WAE framework behaves similarly. Appendix G provides a detailed empirical stability evaluation of our framework in the presence of additive noise. In particular, this experiment shows that for reasonable levels of additive noise ε (normalized with regard to the function range), typically $\varepsilon < 0.1$, the recommended default value of ε_1 (0.05) results in a stable MT-WAE computation. For larger noise levels ($\varepsilon > 0.1$), MT-WAE provides similar stability scores to PD-WAE, for values of ε_1 which are still reasonable in terms of discriminative power ($\varepsilon_1 = 0.1$).

A possible direction to improve the practical stability of the framework without having to deal with a control parameter such as ε_1 would be to consider branch decompositions driven by other criteria than persistence (such as hyper-volume [25] for instance). However, the persistence criterion plays a central role in the Wasserstein distance between merge trees, as discussed by Pont et al. [78] (Sec. 4), in particular to guarantee that interpolated BDTs computed during geodesic construction can indeed be inverted into a valid MT. Thus, other branch decomposition criteria than persistence would require to derive a completely new procedure for several key components of our framework, such as geodesic computation or barycenter estimation. This is an orthogonal research direction to this work, which we leave for future work.

Similarly to other optimization problems based on topological descriptors [78], [79], [95], [97], our energy is not convex. However, our experiments indicate that our initialization strategy (Sec. 4.3) leads to relevant solutions, which can be successfully applied for visualization (Sec. 5).

Since it is based on neural networks, our approach inherits from their intrinsic limitations. Specifically, the energy is not guar-




Fig. 14. Evolution of the normalized reconstruction error along the iterations, for PD-WAE (left) and MT-WAE (right).

anteed to monotonically decrease over the iterations. However, this theoretical limitation has never translated into a practical limitation in our experiments. Fig. 14 reports the evolution of the normalized reconstruction error for PD and MT-WAE computations, for all our test ensembles. In particular, it shows that typical, temporary energy increases can indeed be observed (as often reported when optimizing neural networks), but without preventing the network from converging overall (i.e. reaching a state where the energy decreases by less than 1% between consecutive iterations). Like other neural methods, our approach is conditioned by the meta-parameters defining the network (i.e. number of layers, dimensionality of each layer, etc.). However, we ran our experiments with fairly basic values for these meta-parameters (as detailed in Appendix C) and still obtained substantial improvements over linear encoding based on PGA [79]. This indicates that optimizing in the future these meta-parameters is likely to improve the quality of our framework, however possibly at the expense of longer computations. Finally, our application to merge tree compression does not guarantee any error bound in its current form, which we leave for future work.

7 CONCLUSION

In this paper, we presented a computational framework for the Wasserstein Auto-Encoding of merge trees (and persistence diagrams), with applications to merge tree compression and dimensionality reduction. Our approach improves previous linear attempts at merge tree encoding, by generalizing them to non-linear encoding, hence leading to lower reconstruction errors. In contrast to traditional auto-encoders, our novel layer model enables our neural networks to process topological descriptors natively, without pre-vectorization. As shown in our experiments, this contribution leads not only to superior accuracy (Sec. 6.2) but also to superior interpretability (Sec. 5.2): with our work, it is now possible to interactively explore the latent space and analyze how topological features are transformed by the network in its attempt to best encode the ensemble. Overall, the visualizations derived from our contribution (Figs. 1, 11) enable the interactive, visual inspection of the ensemble, both at a global level (with our 2D layouts) and at a feature level. Specifically, our novel notion of *feature latent importance* enables the identification of the most informative features in the ensemble.

In the future, we will continue our work towards the development of further statistical tools for the visual analysis of ensemble data, based on topological descriptors. In particular, there are several research avenues for improving our current approach. For example, the Wasserstein distance between merge trees is subject to several meta-parameters (Sec. 2.2), for which we

provide generic default values which have shown to be relevant in practice. A possible improvement could consist in letting the auto-encoding framework optimize these meta-parameters (on a per branch basis). However, this data-driven setup of the meta-parameters of our approach would come at the expense of extended running times. Another research avenue could consist in combining our framework with existing approaches on topological losses for image segmentation [49], [50], [92], in order to also auto-encode the scalar data. Moreover, another direction could consist in training a single neural network for auto-encoding multiple ensembles at once. However, this would require to derive new normalization strategies, since, as the scalar fields can take arbitrarily distinct value ranges from an ensemble to the next, the Wasserstein distances between their members can also take arbitrarily distinct values, which would challenge an efficient sampling of the metric space. Finally, we will investigate the usage of neural networks exploiting topological descriptors for further visual analysis tasks, such as trend analysis or anomaly detection or shape classification [60]. In that context, we believe that our new layer model (natively processing topological descriptors) sets the foundations for an accurate and interpretable usage of topological representations with neural networks.

ACKNOWLEDGMENTS

This work is partially supported by the European Commission grant ERC-2019-COG “TORI” (ref. 863464, <https://erc-tori.github.io/>).

REFERENCES

- [1] A. Acharya and V. Natarajan. A parallel and memory efficient algorithm for constructing the contour tree. In *IEEE PacificViz*, 2015.
- [2] H. Adams, S. Chepushtanova, T. Emerson, E. Hanson, M. Kirby, F. Motta, R. Neville, C. Peterson, P. Shipman, and L. Ziegelmeier. Persistence Images: A Stable Vector Representation of Persistent Homology. *Journal of Machine Learning Research*, 18(8):1–35, 2017.
- [3] K. Anderson, J. Anderson, S. Palande, and B. Wang. Topological data analysis of functional MRI connectivity in time and space domains. In *MICCAI Workshop on Connectomics in NeuroImaging*, 2018.
- [4] R. Anirudh, V. Venkataraman, K. N. Ramamurthy, and P. K. Turaga. A Riemannian Framework for Statistical Analysis of Topological Persistence Diagrams. In *IEEE CVPR Workshops*, 2016.
- [5] T. M. Athawale, D. Maljovec, C. R. Johnson, V. Pascucci, and B. Wang. Uncertainty Visualization of 2D Morse Complex Ensembles Using Statistical Summary Maps. *CoRR*, abs/1912.06341, 2019.
- [6] U. Ayachit, A. C. Bauer, B. Geveci, P. O’Leary, K. Moreland, N. Fabian, and J. Mauldin. ParaView Catalyst: Enabling In Situ Data Analysis and Visualization. In *ISAV*, 2015.
- [7] T. F. Banchoff. Critical points and curvature for embedded polyhedral surfaces. *The American Mathematical Monthly*, 45(1):245–256, 1967.
- [8] U. Bauer, X. Ge, and Y. Wang. Measuring distance between Reeb graphs. In *Symposium on Computational Geometry*, 2014.
- [9] U. Bauer, M. Kerber, and J. Reininghaus. Distributed computation of persistent homology. In *Algorithm Engin. and Exp.*, 2014.
- [10] K. Beketayev, D. Yeliussizov, D. Morozov, G. H. Weber, and B. Hamann. Measuring the distance between merge trees. In *TopolnVis*, 2014.
- [11] D. P. Bertsekas. A new algorithm for the assignment problem. *Mathematical Programming*, 21(1):152–171, 1981.
- [12] S. Biasotti, D. Giorgio, M. Spagnuolo, and B. Falcidieno. Reeb graphs for shape analysis and applications. *Theoretical Computer Science*, 392(1-3):5–22, 2008.
- [13] T. Bin Masood, J. Budin, M. Falk, G. Favelier, C. Garth, C. Gueunet, P. Guillou, L. Hofmann, P. Hristov, A. Kamakshidasan, C. Kappe, P. Klacansky, P. Laurin, J. Levine, J. Lukaszcyk, D. Sakurai, M. Soler, P. Steneteg, J. Tierny, W. Usher, J. Vidal, and M. Wozniak. An Overview of the Topology Toolkit. In *TopolnVis*, 2019.
- [14] A. Bock, H. Doraiswamy, A. Summers, and C. T. Silva. TopoAngler: Interactive Topology-Based Extraction of Fishes. *IEEE TVCG*, 24(1):812–821, 2018.
- [15] B. Bollen, E. Chambers, J. A. Levine, and E. Munch. Reeb Graph Metrics from the Ground Up. *CoRR*, 2021.
- [16] B. Bollen, P. Tennakoon, and J. A. Levine. Computing a Stable Distance on Merge Trees. *IEEE TVCG*, 2022.
- [17] H. Bourlard and Y. Kamp. Auto-Association by Multilayer Perceptrons and Singular Value Decomposition. *Biological Cybernetics*, 59, 1988.
- [18] P. Bremer, H. Edelsbrunner, B. Hamann, and V. Pascucci. A Multi-Resolution Data Structure for 2-Dimensional Morse Functions. In *Proc. of IEEE VIS*, 2003.
- [19] P. Bremer, G. Weber, J. Tierny, V. Pascucci, M. Day, and J. Bell. Interactive exploration and analysis of large scale simulations using topology-based data segmentation. *IEEE TVCG*, 17(9):1307–1324, 2011.
- [20] N. Brown, R. Nash, P. Poletti, G. Guzzetta, M. Manica, A. Zardini, M. Flatken, J. Vidal, C. Gueunet, E. Belikov, J. Tierny, A. Podobas, W. D. Chien, S. Markidis, and A. Gerndt. Utilising urgent computing to tackle the spread of mosquito-borne diseases. In *IEEE/ACM UrgentHPC@SC*, 2021.
- [21] R. Brüel Gabriëlsson and G. Carlsson. Exposition and interpretation of the topology of neural networks. In *ICMLA*, pp. 1069–1076, 2019. doi: 10.1109/ICMLA.2019.00180
- [22] P. Bubenik. Statistical topological data analysis using persistence landscapes. *J. Mach. Learn. Res.*, 16:77–102, 2015.
- [23] H. Carr, J. Snoeyink, and U. Axen. Computing contour trees in all dimensions. In *Symp. on Dis. Alg.*, 2000.
- [24] H. Carr, G. Weber, C. Sewell, and J. Ahrens. Parallel peak pruning for scalable SMP contour tree computation. In *IEEE LDAV*, 2016.
- [25] H. A. Carr, J. Snoeyink, and M. van de Panne. Simplifying Flexible Isosurfaces Using Local Geometric Measures. In *IEEE VIS*, 2004.
- [26] M. Carrière, M. Cuturi, and S. Oudot. Sliced wasserstein kernel for persistence diagrams. In *ICML*, 2017.
- [27] M. E. Celebi, H. A. Kingravi, and P. A. Vela. A comparative study of efficient initialization methods for the k-means clustering algorithm. *Expert Syst. Appl.*, 2013.
- [28] L. De Florian, U. Fugacci, F. Iuricich, and P. Magillo. Morse complexes for shape segmentation and homological analysis: discrete models and algorithms. *Computer Graphics Forum*, 34(2):761–785, 2015.
- [29] P. Diggle, P. Heagerty, K.-Y. Liang, and S. Zeger. *The Analysis of Longitudinal Data*. Oxford University Press, 2002.
- [30] H. Edelsbrunner and J. Harer. *Computational Topology: An Introduction*. American Mathematical Society, 2009.
- [31] H. Edelsbrunner, J. Harer, and A. Zomorodian. Hierarchical morse complexes for piecewise linear 2-manifolds. In *Symposium on Computational Geometry*, 2001.
- [32] H. Edelsbrunner, D. Letscher, and A. Zomorodian. Topological Persistence and Simplification. *Discrete Computational Geometry*, 28(4):511–533, 2002.
- [33] H. Edelsbrunner and E. P. Mücke. Simulation of simplicity: a technique to cope with degenerate cases in geometric algorithms. *ACM Transactions on Graphics*, 9(1):66–104, 1990.
- [34] C. Elkan. Using the triangle inequality to accelerate k-means. In *ICML*, 2003.
- [35] G. Favelier, N. Faraj, B. Summa, and J. Tierny. Persistence Atlas for Critical Point Variability in Ensembles. *IEEE TVCG*, 25(1):1152–1162, 2018.
- [36] F. Ferstl, K. Bürger, and R. Westermann. Streamline variability plots for characterizing the uncertainty in vector field ensembles. *IEEE TVCG*, 22(1):767–776, 2016.
- [37] F. Ferstl, M. Kanzler, M. Rautenhaus, and R. Westermann. Visual analysis of spatial variability and global correlations in ensembles of iso-contours. *Computer Graphics Forum*, 35(3):221–230, 2016.
- [38] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Prses, 2016.
- [39] C. Gueunet, P. Fortin, J. Jomier, and J. Tierny. Task-Based Augmented Contour Trees with Fibonacci Heaps. *IEEE TPDS*, 30(8):1889–1905, 2019.
- [40] C. Gueunet, P. Fortin, J. Jomier, and J. Tierny. Task-based Augmented Reeb Graphs with Dynamic ST-Trees. In *EGPGV*, 2019.
- [41] P. Guillou, J. Vidal, and J. Tierny. Discrete Morse Sandwich: Fast Computation of Persistence Diagrams for Scalar Data – An Algorithm and A Benchmark. *IEEE TVCG*, 2023.
- [42] D. Günther, J. Salmon, and J. Tierny. Mandatory critical points of 2D uncertain scalar fields. *Computer Graphics Forum*, 33(3):31–40, 2014.
- [43] A. Gyulassy, P. Bremer, R. Grout, H. Kolla, J. Chen, and V. Pascucci. Stability of dissipation elements: A case study in combustion. *Computer Graphics Forum*, 33(3):51–60, 2014.

- [44] A. Gyulassy, P. Bremer, and V. Pascucci. Shared-Memory Parallel Computation of Morse-Smale Complexes with Improved Accuracy. *IEEE TVCG*, 25(1):1183–1192, 2019.
- [45] A. Gyulassy, A. Knoll, K. Lau, B. Wang, P. Bremer, M. Papka, L. A. Curtiss, and V. Pascucci. Interstitial and Interlayer Ion Diffusion Geometry Extraction in Graphitic Nanosphere Battery Materials. *IEEE TVCG*, 22(1):916–925, 2016.
- [46] C. Heine, H. Leitte, M. Hlawitschka, F. Iuricich, L. De Floriani, G. Scheuermann, H. Hagen, and C. Garth. A survey of topology-based methods in visualization. *Computer Graphics Forum*, 35(3):643–667, 2016.
- [47] F. Hensel, M. Moor, and B. Rieck. A survey of topological machine learning methods. *Frontiers Artif. Intell.*, 2021.
- [48] M. Hilaga, Y. Shinagawa, T. Komura, and T. L. Kunii. Topology matching for fully automatic similarity estimation of 3D shapes. In *ACM SIGGRAPH*, 2001.
- [49] X. Hu, F. Li, D. Samaras, and C. Chen. Topology-Preserving Deep Image Segmentation. In *NeurIPS*, 2019.
- [50] X. Hu, Y. Wang, F. Li, D. Samaras, and C. Chen. Topology-Aware Segmentation Using Discrete Morse Theory. In *ICLR*, 2021.
- [51] M. Hummel, H. Obermaier, C. Garth, and K. I. Joy. Comparative visual analysis of lagrangian transport in CFD ensembles. *IEEE TVCG*, 19(12):2743–2752, 2013.
- [52] L. Kantorovich. On the translocation of masses. *AS USSR*, 1942.
- [53] J. Kasten, J. Reininghaus, I. Hotz, and H. Hege. Two-dimensional time-dependent vortex regions based on the acceleration magnitude. *IEEE TVCG*, 17(12):2080–2087, 2011.
- [54] M. Kerber, D. Morozov, and A. Nigmatov. Geometry helps to compare persistence diagrams. *ACM J. of Experimental Algorithmics*, 22, 2017.
- [55] V. Khrulkov and I. Oseledets. Geometry score: A method for comparing generative adversarial networks. In J. Dy and A. Krause, eds., *ICML*, vol. 80 of *Proc. of MLR*, pp. 2621–2629. PMLR, 10–15 Jul 2018.
- [56] K. Kim, J. Kim, M. Zaheer, J. S. Kim, F. Chazal, and L. Wasserman. PLLay: Efficient Topological Layer based on Persistence Landscapes. In *NeurIPS*, 2020.
- [57] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- [58] J. B. Kruskal and M. Wish. Multidimensional Scaling. In *SUPS*, 1978.
- [59] T. Lacombe, M. Cuturi, and S. Oudot. Large Scale computation of Means and Clusters for Persistence Diagrams using Optimal Transport. In *NIPS*, 2018.
- [60] B. Li, Y. Lu, C. Li, A. Godil, T. Schreck, M. Aono, Q. Chen, N. K. Chowdhury, B. Fang, T. Furuya, H. Johan, R. Kosaka, H. Koyanagi, R. Ohbuchi, and A. Tsumura. SHREC’ 14 Track: Large Scale Comprehensive 3D Shape Retrieval. In *3DOR*, 2014.
- [61] M. Li, S. Palande, and B. Wang. Sketching merge trees. In *IEEE TopoInVis*, 2023.
- [62] A. P. Lohfink, F. Wetzels, J. Lukaszcyk, G. H. Weber, and C. Garth. Fuzzy contour trees: Alignment and joint layout of multiple contour trees. *Computer Graphics Forum*, 39(3):343–355, 2020.
- [63] S. Maadasamy, H. Doraiswamy, and V. Natarajan. A hybrid parallel algorithm for computing and tracking level set topology. In *HiPC*, 2012.
- [64] D. Maljovec, B. Wang, P. Rosen, A. Alfonsi, G. Pastore, C. Rabiti, and V. Pascucci. Topology-inspired partition-based sensitivity analysis and visualization of nuclear simulations. In *IEEE PacificViz*, 2016.
- [65] M. Mirzargar, R. Whitaker, and R. Kirby. Curve boxplot: Generalization of boxplot for ensembles of curves. *IEEE TVCG*, 20(12):2654–2663, 2014.
- [66] G. Monge. Mémoire sur la théorie des déblais et des remblais. *Académie Royale des Sciences de Paris*, 1781.
- [67] M. Moor, M. Horn, B. Rieck, and K. M. Borgwardt. Topological autoencoders. In *ICML*, 2020.
- [68] D. Morozov, K. Beketayev, and G. H. Weber. Interleaving distance between merge trees. In *TopoInVis*. 2014.
- [69] J. Munkres. Algorithms for the assignment and transportation problems. *Journal of the Society for Industrial and Applied Mathematics*, 5(1):32–38, 1957.
- [70] F. Nauleau, F. Vivodtzev, T. Bridel-Bertomeu, H. Beaugendre, and J. Tierny. Topological Analysis of Ensembles of Hydrodynamic Turbulent Flows – An Experimental Study. In *IEEE Symposium on Large Data Analysis and Visualization*, 2022.
- [71] M. Olejniczak, A. S. P. Gomes, and J. Tierny. A Topological Data Analysis Perspective on Non-Covalent Interactions in Relativistic Calculations. *International Journal of Quantum Chemistry*, 120(8):e26133, 2019.
- [72] M. Olejniczak and J. Tierny. Topological Data Analysis of Vortices in the Magnetically-Induced Current Density in LiH Molecule. *Physical Chemistry Chemical Physics*, 2023.
- [73] Organizers. The IEEE SciVis Contest. <http://sciviscontest.ieeevis.org/>, 2004.
- [74] S. Parsa. A deterministic $o(m \log m)$ time algorithm for the reeb graph. In *Symposium on Computational Geometry*, 2012.
- [75] V. Pascucci, K. Cole-McLaughlin, and G. Scorzelli. Multi-resolution computation and presentation of contour trees. In *IASTED*, 2004.
- [76] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Köpf, E. Z. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *NeurIPS*, 2019. <https://pytorch.org/cppdocs/>.
- [77] K. Pearson. On Lines and Planes of Closest Fit to Systems of Points in Space. *Philosophical Magazine*, 2:559–572, 1901.
- [78] M. Pont, J. Vidal, J. Delon, and J. Tierny. Wasserstein Distances, Geodesics and Barycenters of Merge Trees. *IEEE TVCG*, 28(1):291–301, 2022. <https://github.com/MatPont/WassersteinMergeTreesData>.
- [79] M. Pont, J. Vidal, and J. Tierny. Principal Geodesic Analysis of Merge Trees (and Persistence Diagrams). *IEEE TVCG*, 2023.
- [80] K. Potter, A. Wilson, P. Bremer, D. Williams, C. Doutriaux, V. Pascucci, and C. R. Johnson. Ensemble-vis: A framework for the statistical visualization of ensemble data. In *2009 IEEE ICDM*, 2009.
- [81] J. Reininghaus, S. Huber, U. Bauer, and R. Kwitt. A stable multi-scale kernel for topological machine learning. In *IEEE CVPR*, 2015.
- [82] V. Robins and K. Turner. Principal Component Analysis of Persistent Homology Rank Functions with case studies of Spatial Point Patterns, Sphere Packing and Colloids. *Physica D: Nonlinear Phenomena*, 334:99–117, 2016.
- [83] V. Robins, P. J. Wood, and A. P. Sheppard. Theory and Algorithms for Constructing Discrete Morse Complexes from Grayscale Digital Images. *IEEE Trans. Pattern Anal. Mach. Intell.*, 33(8):1646–1658, 2011.
- [84] H. Saikia, H. Seidel, and T. Weinkauff. Extended Branch Decomposition Graphs: Structural Comparison of Scalar Data. *Computer Graphics Forum*, 33(3):41–50, 2014.
- [85] J. Sanyal, S. Zhang, J. Dyer, A. Mercer, P. Amburn, and R. Moorhead. Noodles: A tool for visualization of numerical weather model ensemble uncertainty. *IEEE TVCG*, 16(6):1421–1430, 2010.
- [86] N. Shivashankar and V. Natarajan. Parallel Computation of 3D Morse-Smale Complexes. *Computer Graphics Forum*, 31(3):965–974, 2012.
- [87] N. Shivashankar, P. Pranav, V. Natarajan, R. van de Weygaert, E. P. Bos, and S. Rieder. Felix: A topology based framework for visual exploration of cosmic filaments. *IEEE TVCG*, 22(6):1745–1759, 2016.
- [88] K. Sisouk, J. Delon, and J. Tierny. Wasserstein Dictionaries of Persistence Diagrams. *CoRR*, 2023. doi: 10.48550/arXiv.2304.14852
- [89] M. Soler, M. Petitfrere, G. Darche, M. Plainchault, B. Conche, and J. Tierny. Ranking Viscous Finger Simulations to an Acquired Ground Truth with Topology-Aware Matchings. In *IEEE Lдав*, 2019.
- [90] T. Sousbie. The Persistent Cosmic Web and its Filamentary Structure: Theory and Implementations. *Royal Astronomical Society*, 414:384–403, 2011.
- [91] R. Sridharamurthy, T. B. Masood, A. Kamakshidasan, and V. Natarajan. Edit Distance between Merge Trees. *IEEE TVCG*, 26(3):1518–1531, 2020.
- [92] N. Stucki, J. C. Paetzold, S. Shit, B. H. Menze, and U. Bauer. Topologically Faithful Image Segmentation via Induced Matching of Persistence Barcodes. In *ICML*, 2023.
- [93] J. Tierny, G. Favelier, J. A. Levine, C. Gueunet, and M. Michaux. The Topology Toolkit. *IEEE TVCG*, 24(1):832–842, 2017. <https://topology-tool-kit.github.io/>.
- [94] I. O. Tolstikhin, O. Bousquet, S. Gelly, and B. Schölkopf. Wasserstein auto-encoders. In *ICLR*, 2018.
- [95] K. Turner, Y. Mileyko, S. Mukherjee, and J. Harer. Fréchet Means for Distributions of Persistence Diagrams. *Discrete Computational Geometry*, 52(1):44–70, 2014.
- [96] L. P. van der Maaten and G. Hinton. Visualizing Data Using t-SNE. *JMLR*, 9(86):2579–2605, 2008.
- [97] J. Vidal, J. Budin, and J. Tierny. Progressive Wasserstein Barycenters of Persistence Diagrams. *IEEE TVCG*, 26(1):151–161, 2020.
- [98] F. Wetzels and C. Garth. A Deformation-based Edit Distance for Merge Trees. In *TopoInVis*, 2022.
- [99] F. Wetzels, H. Leitte, and C. Garth. Branch Decomposition-Independent Edit Distances for Merge Trees. *Computer Graphics Forum*, 2022.

- [100] R. T. Whitaker, M. Mirzargar, and R. M. Kirby. Contour boxplots: A method for characterizing uncertainty in feature sets from simulation ensembles. *IEEE TVCG*, 19(12):2713–2722, 2013.
- [101] D. P. Woodruff. *Sketching as a Tool for Numerical Linear Algebra*. Now Publishers, 2014.
- [102] L. Yan, T. B. Masood, R. Sridharamurthy, F. Rasheed, V. Natarajan, I. Hotz, and B. Wang. Scalar field comparison with topological descriptors: Properties and applications for scientific visualization. *Computer Graphics Forum*, 40(3):599–633, 2021.
- [103] L. Yan, Y. Wang, E. Munch, E. Gasparovic, and B. Wang. A structural average of labeled merge trees for uncertainty visualization. *IEEE TVCG*, 26(1):832–842, 2019.
- [104] K. Zhang. A Constrained Edit Distance Between Unordered Labeled Trees. *Algorithmica*, 15(3):205–222, 1996.
- [105] S. Zhou, E. Zelikman, F. Lu, A. Y. Ng, G. E. Carlsson, and S. Ermon. Evaluating the disentanglement of deep generative models through manifold topology. In *ICLR*, 2021.



Mathieu Pont is a Ph.D. student at Sorbonne Universite. He received a M.S. degree in Computer Science from Paris Descartes University in 2020. He is an active contributor to the Topology ToolKit (TTK), an open source library for topological data analysis. His notable contributions to TTK include distances, geodesics and barycenters of merge trees, for feature tracking and ensemble clustering.



Julien Tierny received the Ph.D. degree in Computer Science from the University of Lille in 2008. He is a CNRS research director at Sorbonne University. Prior to his CNRS tenure, he held a Fulbright fellowship and was a post-doctoral researcher at the University of Utah. His research expertise lies in topological methods for data analysis and visualization. He is the founder and lead developer of the Topology ToolKit (TTK), an open source library for topological data analysis.

APPENDIX

APPENDIX A TOPOLOGICAL DESCRIPTORS

This appendix provides a brief description of the topological descriptors considered in the main manuscript, namely the *Persistence Diagram* (PD, Sec. A.1) and the *Merge Tree* (MT), specifically, its variant called *Branch Decomposition Tree*, (BDT, Sec. A.2). We refer the reader to textbooks [30] for an introduction to computational topology.

A.1 Persistence diagrams

Given a piecewise linear (PL) scalar field $f_i : \mathcal{M} \rightarrow \mathbb{R}$, the *sub-level set* of f_i , noted $f_{i-\infty}^{-1}(w) = \{p \in \mathcal{M} \mid f_i(p) < w\}$, is defined as the pre-image of $(-\infty, w)$ by f_i . The *super-level set* of f_i is defined symmetrically: $f_{i+\infty}^{-1}(w) = \{p \in \mathcal{M} \mid f_i(p) > w\}$. As w continuously increases, the topology of $f_{i-\infty}^{-1}(w)$ changes at specific vertices of \mathcal{M} , called the *critical points* of f_i [7]. Critical points are classified by their *index* \mathcal{I}_i : 0 for minima, 1 for 1-saddles, $d-1$ for $(d-1)$ -saddles and d for maxima. In practice, f_i is enforced to contain only isolated, non-degenerate critical points [31], [33]. In 3D, connected components of $f_{i-\infty}^{-1}(w)$ are created at local minima and destroyed at 1-saddles. One-dimensional cycles are created at 1-saddles and destroyed at 2-saddles and voids are created at 2-saddles and destroyed at maxima.

The persistence diagram is a visual summary of the above topological features. As shown in Fig. 2 (main manuscript), it is closely related to the *merge tree*, which is the main topological representation studied in this paper. We first describe the persistence diagram though as the metric used in our work to measure distances between merge trees (Sec. 2.2, main manuscript) generalizes an established metric between persistence diagrams.

Specifically, in the domain, each topological feature of $f_{i-\infty}^{-1}(w)$ can be associated with a unique pair of critical points (c, c') , corresponding to its *birth* and *death*. The Elder rule [30] states that critical points can be arranged in pairs according to this observation, such that each critical point appears in only one pair (c, c') , with $f_i(c) < f_i(c')$ and $\mathcal{I}_i(c) = \mathcal{I}_i(c') - 1$. For instance, if two connected components of $f_{i-\infty}^{-1}(w)$ meet at a critical point c' , the *younger* component (created last, in c) *dies*, in favor of the *older* one (created first).

The persistence diagram $\mathcal{D}(f_i)$ embeds each pair to a single point in 2D at coordinates $(f_i(c), f_i(c'))$. The *persistence* of a pair is given by its height $f_i(c') - f_i(c)$. The persistence diagram provides a visual overview of the features of a dataset (Fig. 2, main manuscript), where salient features stand out from the diagonal while pairs corresponding to noise are located near the diagonal.

A.2 Merge trees

In the following, we introduce the main topological data representation studied in this paper: the *merge tree*. We also describe a specific representation of the merge tree called the *branch decomposition tree*, which can be interpreted as a generalization of the extremum persistence diagram, and which plays a central role in the computation of distances between merge trees (Sec. 2.2, main manuscript).

The *join tree*, noted $\mathcal{T}^-(f_i)$, is a visual summary of the connected components of $f_{i-\infty}^{-1}(w)$ [23]. It is a 1-dimensional

simplicial complex defined as the quotient space $\mathcal{T}^-(f_i) = \mathcal{M} / \sim$ by the equivalence relation \sim which states that p_1 and p_2 are equivalent if $f_i(p_1) = f_i(p_2)$ and if p_1 and p_2 belong to the same connected component of $f_{i-\infty}^{-1}(f_i(p_1))$.

The *split tree* (Fig. 2, main manuscript), noted $\mathcal{T}^+(f_i)$, is defined symmetrically and describes the connected components of the super-level set $f_{i+\infty}^{-1}(w)$. Each of these two *directed trees* is called a *merge tree* (MT), noted generically $\mathcal{T}(f_i)$ in the following. Intuitively, these trees track the creation of connected components of the sub (or super) level sets at their leaves, and merge events at their interior nodes. To mitigate a phenomenon called *saddle swap*, these trees are often post-processed [78], [91], by merging adjacent saddles in the tree if their relative difference in scalar value is smaller than a threshold $\epsilon_1 \in [0, 1]$.

Merge trees are often visualized via a persistence-driven *branch decomposition* [75], to make the persistence pairs captured by the tree stand out. In this context, a *persistent branch* is a monotone path on the tree connecting the nodes corresponding to the creation and destruction (according to the Elder rule, Sec. A.1) of a connected component of sub (or super) level set. Then, the branch decomposition provides a planar layout of the MT, where each persistent branch is represented as a vertical segment (center insets in Fig. 2, main manuscript).

The *branch decomposition tree* (BDT), noted $\mathcal{B}(f_i)$, is a directed tree whose nodes are the persistent branches captured by the branch decomposition and whose arcs denote adjacency relations between them in the MT. In Fig. 2 (main manuscript), the BDTs (right insets) can be interpreted as the dual of the branch decompositions (center insets, with matching colors): each vertical segment in the branch decomposition (center) corresponds to a node in the BDT (right) and each horizontal segment (center, denoting an adjacency relation between branches) corresponds to an arc in the BDT. The BDT can be interpreted as a generalization of the extremum persistence diagram: like $\mathcal{D}(f_i)$, $\mathcal{B}(f_i)$ describes the population of (extremum) persistence pairs present in the data. However, unlike the persistence diagram, it additionally captures adjacency relations between them (Fig. 2, main manuscript).

Note that, the birth and death of each persistent branch $b_i \in \mathcal{B}(f_i)$, noted (x_i, y_i) , span by construction an interval included in that of its parent $b'_i \in \mathcal{B}(f_i)$: $[x_i, y_i] \subseteq [x'_i, y'_i]$. This *nesting property* of BDTs [78] is a direct consequence of the Elder rule (Sec. A.1).

APPENDIX B GENERAL FORMULATION OF BASIS PROJECTION

The section 4.2 of the main manuscript presents an *Assignment/Update* algorithm to project an input BDT \mathcal{B} into a given BDT basis $\mathcal{B}(\mathcal{O})$.

In the *Assignment* phase, given an initial set of coefficients $\alpha \in \mathbb{R}^d$, the *estimation* $\widehat{\mathcal{B}}$ of \mathcal{B} is given by :

$$\widehat{\mathcal{B}} \leftarrow \mathcal{O} + \mathcal{B}(\mathcal{O})\alpha. \quad (15)$$

Given this estimation $\widehat{\mathcal{B}}$, the assignment step first evaluates the Wasserstein distance $W_2^{\mathcal{T}}(\mathcal{B}, \widehat{\mathcal{B}})$. For this, the optimal assignment ϕ_* between \mathcal{B} and $\widehat{\mathcal{B}}$ is computed with regard to the Wasserstein distance (Eq. 1 of the main manuscript).

Then, given the optimal assignment ϕ_* , the Wasserstein distance $W_2^{\mathcal{T}}(\mathcal{B}, \widehat{\mathcal{B}})$ can be re-written as:

$$W_2^{\mathcal{T}}(\mathcal{B}, \widehat{\mathcal{B}}) = \sum_{i=1}^{|\mathcal{B}|} \begin{cases} 0 & \text{if both } b_i \text{ and } \phi_*(b_i) \text{ are on the diagonal,} \\ ||b_i - \phi_*(b_i)||_2^2 & \text{otherwise.} \end{cases} \quad (16)$$




Fig. 15. When considering an assignment (purple arrows) in the 2D birth/death planes between *augmented* BDTs, four cases can occur (purple numbers). **Case (1)**: an *off-diagonal* branch (blue dot) can be mapped to an *off-diagonal* branch (green dot). **Case (2)**: an *off-diagonal* branch (blue dot) can be mapped to an augmented *diagonal* branch (green circle). **Case (3)**: an augmented *diagonal* branch (blue circle) can be mapped to an augmented *diagonal* branch (green circle). **Case (4)**: an augmented *diagonal* branch (blue circle) can be mapped to an *off-diagonal* branch (green dot). Sec. 4.2 (main manuscript) covers the cases (1) and (2). Sec. B generalizes this formulation to all cases.

The purpose of the subsequent *Update* step is precisely to optimize α in order to minimize the evaluation of $W_2^T(\mathcal{B}, \widehat{\mathcal{B}})$ by the above equation. For this, one needs to isolate from Eq. 16 all the terms involving α from those which do not.

In the simple case where ϕ_* describes a bijection between off-diagonal points (case covered in the main manuscript), no branch of \mathcal{B} depends on α . Then the isolation of the terms involving α is simple: only the branches of $\widehat{\mathcal{B}}$ depend on α (Eq. 15).

In the more general case, things are a bit more involved. As illustrated in Fig. 15, the computation of the optimal assignment $\phi_* : \mathcal{B} \rightarrow \widehat{\mathcal{B}}$ (purple arrows) implies a pre-processing phase of *augmentation* of the 2D birth/death plane. As described in the section 2.2 of the main manuscript, $\widehat{\mathcal{B}}$ is augmented with the diagonal projections of the branches of \mathcal{B} (Fig. 15, green circles) and \mathcal{B} is augmented with the diagonal projections of the branches of $\widehat{\mathcal{B}}$ (Fig. 15, blue circles). This augmentation phase enables the modeling of the destruction (or creation) of features during the assignment ϕ_* (between the blue and green items, Fig. 15). Then, the following four cases can occur (Fig. 15):

Case (1): An *off-diagonal* branch $b_1 \in \mathcal{B}$ is mapped to an *off-diagonal* branch $\phi_*(b_1) \in \widehat{\mathcal{B}}$. Then, the birth/death values of b_1 do not depend on α and only the birth/death values of $\phi_*(b_1)$ do. This corresponds to the case covered by the section 4.2 of the main manuscript.

Case (2): An *off-diagonal* branch $b_2 \in \mathcal{B}$ is mapped to a *diagonal* branch $\phi_*(b_2) \in \widehat{\mathcal{B}}$. Then, the birth/death values of b_2 do not depend on α . Then, this case is also covered by Sec. 4.2 (main manuscript).

Case (3): A *diagonal* branch $b_3 \in \mathcal{B}$ is mapped to a *diagonal*

branch $\phi_*(b_3) \in \widehat{\mathcal{B}}$. In that case, the ground distance $d_2(b_3, \phi_*(b_3))$ (section 2.2 of the main manuscript) is set to zero by convention (first line of Eq. 16). This models the fact that both b_3 and $\phi_*(b_3)$ are dummy features (with zero persistence) and that their ground distance, which is arbitrary, should not be taken into account in $W_2^T(\mathcal{B}, \widehat{\mathcal{B}})$. Therefore, we simply remove b_3 from \mathcal{B} and $\phi_*(b_3)$ from $\widehat{\mathcal{B}}$. This removal discards this first line of Eq. 16, which can then be re-written in the general form:

$$W_2^T(\mathcal{B}, \widehat{\mathcal{B}}) = \sum_{i=1}^{|\mathcal{B}|} \|b_i - \phi_*(b_i)\|_2^2. \quad (17)$$

Then, with this removal, the diagonal-diagonal assignments do not constitute a special case anymore.

Case (4): An *off-diagonal* branch $b_4 \in \mathcal{B}$ is mapped to a *diagonal* branch $\phi_*(b_4) \in \widehat{\mathcal{B}}$. In that case, b_4 turns out to be an *augmented* point of \mathcal{B} . In Fig. 15, these are reported with circles while original (i.e. non-augmented) points are reported with dots. Specifically, b_4 has been precisely inserted such that $b_4 = \Delta(\phi_*(b_4))$:

$$(b_4)_x = (b_4)_y = \frac{1}{2} \left((\phi_*(b_4))_x + (\phi_*(b_4))_y \right).$$

Given Eq. 15, $\phi_*(b_4)$ can be re-written as:

$$\phi_*(b_4) = o_4 + B(\mathcal{O})\alpha,$$

where o_4 is a branch of \mathcal{O} . Then, the birth/death values of b_4 are:

$$(b_4)_x = (b_4)_y = \frac{1}{2} \left((o_4 + B(\mathcal{O})\alpha)_x + (o_4 + B(\mathcal{O})\alpha)_y \right) \\ (b_4)_x = (b_4)_y = \frac{1}{2} \left((o_4)_x + (o_4)_y \right) + \frac{1}{2} \left((B(\mathcal{O})\alpha)_x + (B(\mathcal{O})\alpha)_y \right).$$

Then, it follows that b_4 can be re-written as:

$$b_4 = \Delta(o_4) + \Delta(B(\mathcal{O})\alpha). \quad (18)$$

From Eq. 18, it is clear that the coordinates of b_4 are dependent on α . In short, this is due to the fact that b_4 has been purposely inserted in \mathcal{B} as the diagonal projection of a branch of $\widehat{\mathcal{B}}$ which, itself, depends on α . Thus, to account for this special case, we need to further isolate the terms of Eq. 18 depending on α (i.e. $\Delta(B(\mathcal{O}))$), as described next.

Similarly to Sec. 4.2 (main manuscript), let $\widehat{\mathcal{B}}'$ be a vector representation of $\widehat{\mathcal{B}}$. Specifically $\widehat{\mathcal{B}}'$ is a vector in $\mathbb{R}^{2|\widehat{\mathcal{B}}|}$ which concatenates the coordinates in the birth/death plane of each branch b_i of $\widehat{\mathcal{B}}$. $\widehat{\mathcal{B}}'$ can be decomposed into $\mathcal{O}' + (B(\mathcal{O}'))'\alpha$, where $(B(\mathcal{O}'))'$ is a $(2|\widehat{\mathcal{B}}|) \times d'$ matrix. Also, let \mathcal{B}' be a similar vector representation of \mathcal{B} , but where the entries have been specifically re-ordered such that, for each of its 2D entries, we have:

$$(\mathcal{B}')_i = \phi_*^{-1}((\widehat{\mathcal{B}}')_i).$$

Then \mathcal{B}' can be decomposed as a sum of two vectors of $\mathbb{R}^{2|\widehat{\mathcal{B}}|}$:

$$\mathcal{B}' = \mathcal{B}'_1 + \mathcal{B}'_2,$$

such that \mathcal{B}'_1 has all its entries set to 0 except those covered by the above cases (1) and (2) (Fig. 15), and that \mathcal{B}'_2 has all its entries set to 0 except those covered by the above case (4) (purple, Fig. 15).

Given Eq. 18, each non-zero entry i of \mathcal{B}'_2 can be re-written as:

$$(\mathcal{B}'_2)_i = \Delta(o_i) + \Delta \left(\left((B(\mathcal{O}'))'\alpha \right)_i \right), \quad (19)$$

where o_i is the i^{th} entry of \mathcal{O} . Then, the vector \mathcal{B}'_2 can be further decomposed as follows:

$$\mathcal{B}'_2 = \mathcal{B}'_3 + \mathcal{B}'_4,$$

such that each non-zero entry i of \mathcal{B}'_3 is equal to $\Delta(o_i)$ and each non-zero entry i of \mathcal{B}'_4 is equal to $\Delta\left(\left((\mathcal{B}(\mathcal{O}))'\alpha\right)_i\right)$.

Let \mathcal{B}'_2 be a $(2|\widehat{\mathcal{B}}|) \times d'$ matrix such that:

$$\mathcal{B}'_4 = \mathcal{B}'_2 \alpha.$$

At this stage, we have:

$$\mathcal{B}' = \mathcal{B}'_1 + \mathcal{B}'_3 + \mathcal{B}'_4 = \mathcal{B}'_1 + \mathcal{B}'_3 + \mathcal{B}'_2 \alpha. \quad (20)$$

At this point, we managed to isolate the terms in \mathcal{B}' which are dependent on α (i.e. $\mathcal{B}'_2 \alpha$). Then, similarly to the section 4.2 of the main manuscript, for a fixed optimal assignment ϕ_* , the Wasserstein distance $W_2^T(\mathcal{B}, \widehat{\mathcal{B}})$ can be re-written as an L_2 norm:

$$W_2^T(\mathcal{B}, \widehat{\mathcal{B}}) = \|\mathcal{B}' - \widehat{\mathcal{B}}\|_2^2.$$

Then, given ϕ_* , by using Eq. 20, the optimal $\alpha_* \in \mathbb{R}^{d'}$ are:

$$\begin{aligned} \alpha_* &= \arg \min_{\alpha} \|\mathcal{B}' - \widehat{\mathcal{B}}\|_2^2 \\ \alpha_* &= \arg \min_{\alpha} \|\mathcal{B}'_1 + \mathcal{B}'_3 + \mathcal{B}'_2 \alpha - (\mathcal{O}' + (\mathcal{B}(\mathcal{O}'))'\alpha)\|_2^2 \\ \alpha_* &= \arg \min_{\alpha} \|\mathcal{B}'_1 + \mathcal{B}'_3 - \mathcal{O}' - \left(\left((\mathcal{B}(\mathcal{O}'))' - \mathcal{B}'_2\right)\alpha\right)\|_2^2. \end{aligned}$$

Then, similarly to the Euclidean case (Eq. 4 of the main manuscript), it follows then that α_* can be expressed as a function of the pseudoinverse of $\left(\left((\mathcal{B}(\mathcal{O}'))' - \mathcal{B}'_2\right)\right)$:

$$\alpha_* = \left(\left((\mathcal{B}(\mathcal{O}'))' - \mathcal{B}'_2\right)\right)^+ (\mathcal{B}'_1 + \mathcal{B}'_3 - \mathcal{O}'). \quad (21)$$

In short, the general expression of the optimal coefficients α_* (Eq. 21) is a generalization of the Eq. 13 of the main manuscript, such that the branches of \mathcal{B} dependent on α (case (4)) have been integrated within the pseudoinverse operation.

APPENDIX C COMPUTATIONAL PARAMETERS

The Wasserstein distance W_2^T is subject to three parameters (ε_1 , ε_2 and ε_3 , Sec. 2.2, main manuscript), for which we use the recommended default values ($\varepsilon_1 = 0.05$, $\varepsilon_2 = 0.95$, $\varepsilon_3 = 0.9$, [78]) when considering merge trees (MT-WAE). In contrast, when considering persistence diagrams, we switch ε_1 to 1 (ε_2 and ε_3 do not have any effect then) and W_2^T becomes equivalent to W_2^D (Sec. 2.2, main manuscript). Then our framework computes a Wasserstein Auto-Encoder of extremum persistence diagrams (PD-WAE for short).

Our main algorithm is subject to meta-parameters. n_{it} stands for the number of iterations in our basis projection procedure (Sec. 4.2, main manuscript). In practice, we set $n_{it} = 2$.

The number, size and dimensionality of the layers of our MT-WAE are also meta-parameters. Unless specified otherwise, we use only one encoding layer and one decoding layer, i.e. $n_e = n_d = 1$, with $d_{n_e} = 2$ (for dimensionality reduction purposes) and $d_{n_e+n_d} = 16$. For data reduction purposes and computational cost control, we also restrict the size of the origins and bases of the sub-layers of our network. Let $|\mathcal{S}_{\mathcal{B}}|$ be the total number of branches in the ensemble, i.e. $|\mathcal{S}_{\mathcal{B}}| = \sum_{i=1}^N |\mathcal{B}(f_i)|$. We restrict the maximum size of the following origins as follows: $|\mathcal{O}_1^m| \leq 0.2|\mathcal{S}_{\mathcal{B}}|$, $|\mathcal{O}_1^{out}| \leq 0.1|\mathcal{S}_{\mathcal{B}}|$, $|\mathcal{O}_2^m| \leq 0.1|\mathcal{S}_{\mathcal{B}}|$, $|\mathcal{O}_2^{out}| \leq 0.2|\mathcal{S}_{\mathcal{B}}|$. This origin size control also implicitly restricts the size of the corresponding bases. Overall, when integrating all these constraints, the number of variables in our networks is bounded by $((d_{n_e} + 1) \times 2 \times (0.2 + 0.1) + (d_{n_d} + 1) \times 2 \times (0.1 + 0.2)) \times |\mathcal{S}_{\mathcal{B}}| = 12|\mathcal{S}_{\mathcal{B}}|$. In practice, our networks optimized 68,902 variables on average (per ensemble).

APPENDIX D PENALTY TERMS

The flexibility of our framework allows to improve the quality of the dimensionality reduction computed by MT-WAE. Here, we introduce two penalty terms aiming at (1) improving the preservation of the Wasserstein metric W_2^T and (2) improving the preservation of the clusters of BDTs.

D.1 Metric penalty term

In order to improve the preservation of the Wasserstein metric W_2^T in the latent space, and hence in the 2D layout, we introduce the following penalty term $P_M(\theta)$:

$$P_M(\theta) = \sum_{\forall i \in \{1, \dots, N\}} \sum_{\forall j \neq i \in \{1, \dots, N\}} \left(W_2^T(\mathcal{B}(f_i), \mathcal{B}(f_j)) - \|\alpha_{n_e}^i - \alpha_{n_e}^j\|_2 \right)^2.$$

Concretely, given two BDTs $\mathcal{B}(f_i)$ and $\mathcal{B}(f_j)$, $P_M(\theta)$ penalizes the variations between their Wasserstein distances and the Euclidean distances between their coordinates $\alpha_{n_e}^i$ and $\alpha_{n_e}^j$ in the latent space.

The integration of the penalty term $P_M(\theta)$ in our optimization algorithm (Sec. 4, main manuscript) is straightforward. The Wasserstein distance matrix, which stores at its entry (i, j) the distance $W_2^T(\mathcal{B}(f_i), \mathcal{B}(f_j))$, is computed in a pre-processing stage. Since this matrix is a constant during the optimization, the expression of $P_M(\theta)$ only involves basic operations supported by automatic differentiation. Then, given a blending weight $\lambda_M \in [0, 1]$ (in practice, we set $\lambda_M = 1$), the penalty term $\lambda_M P_M(\theta)$ is simply added to the expression of the reconstruction energy $E_{W_2^T}(\theta)$ (Eq. 14, main manuscript). Next, the corresponding gradient is evaluated by automatic differentiation and the overall energy is optimized by gradient descent [57], as originally described in Sec. 4.5 (main manuscript).

D.2 Clustering penalty term

We introduce an additional penalty term to improve the preservation of the *natural* clusters of BDTs in the latent space, and hence in the 2D layout. Let $C \in \mathbb{R}^{kN}$ be the vector modeling the input k clusters: the entry $(ik) + j$ of this vector is equal to 1 if the BDT $\mathcal{B}(f_i)$ belongs to the cluster j , 0 otherwise. This *input clustering vector* can be provided either by a pre-defined ground-truth, by interactive user inputs or by any automatic clustering algorithm. In our experiments, to construct this vector C , we used the extension of the k -means clustering to the Wasserstein metric space of merge trees [78]. Next, in the latent space, we consider the classic k -means algorithm [27], [34], where each BDT $\mathcal{B}(f_i)$ is clustered according to its latent coordinates $\alpha_{n_e}^i \in \mathbb{R}^2$. This yields a set of k centroids in the 2D latent space $c_l \in \mathbb{R}^2$ with $l \in \{0, 1, \dots, k-1\}$. To evaluate the similarity between this clustering and the input clustering vector C , we use the celebrated *SoftMax* function [38]. Specifically, we consider the *latent clustering vector* $C' \in \mathbb{R}^{kN}$, such that the entry $(C')_{(ik)+j}$ denotes the probability that the 2D point $\alpha_{n_e}^i$ belongs to the cluster j (β is set to 5):

$$(C')_{(ik)+j} = \frac{e^{-\beta \|\alpha_{n_e}^i - c_j\|_2}}{\sum_{l=0}^{k-1} e^{-\beta \|\alpha_{n_e}^i - c_l\|_2}}.$$

Then, the clustering penalty term $P_C(\theta)$ is given by the Kullback-Leibler divergence (a standard indicator for probability similarity):

$$P_C(\theta) = KL(C, C') = \sum_{i=0}^{kN-1} C(i) \log \left(\frac{C(i)}{C'(i)} \right).$$

TABLE 3

Comparison of the Average Relative Reconstruction (ARR) Error, between PD-PGA [79] ($d_{max} = 3$ and $N_1 \leq 0.1|\mathcal{S}_B|$) and our approach PD-WAE ($d_{ne} = 3$ and $|\mathcal{O}_{n_e}^{out}| \leq 0.1|\mathcal{S}_B|$), for identical compression factors. Bold numbers in the *Ratio* column indicate instances where PD-WAE achieved a lower (hence better) reconstruction error.

Dataset	N	B	Compression Factor	ARR Error		Ratio
				PD-PGA [79]	PD-WAE	
Asteroid Impact (3D)	7	1,295	7.36	0.01	0.01	0.73
Cloud processes (2D)	12	1,209	7.99	0.12	0.10	0.81
Viscous fingering (3D)	15	118	7.87	0.02	0.01	0.67
Dark matter (3D)	40	316	8.68	0.01	4e-03	0.39
Volcanic eruptions (2D)	12	811	7.56	0.02	0.01	0.36
Ionization front (2D)	16	135	8.01	0.03	0.02	0.69
Ionization front (3D)	16	763	7.68	0.05	0.03	0.65
Earthquake (3D)	12	1,203	7.59	0.04	0.02	0.51
Isabel (3D)	12	1,338	7.58	0.08	0.08	0.91
Starting Vortex (2D)	12	124	7.39	0.01	3e-03	0.51
Sea Surface Height (2D)	48	1,787	24.66	0.16	0.16	1.02
Vortex Street (2D)	45	23	15.83	2e-03	9e-04	0.40

TABLE 4

Comparison of the Average Relative Reconstruction (ARR) Error, between MT-PGA [79] ($d_{max} = 3$ and $N_1 \leq 0.1|\mathcal{S}_B|$) and our approach MT-WAE ($d_{ne} = 3$ and $|\mathcal{O}_{n_e}^{out}| \leq 0.1|\mathcal{S}_B|$), for identical compression factors. Bold numbers in the *Ratio* column indicate instances where MT-WAE achieved a lower (hence better) reconstruction error.

Dataset	N	B	Compression Factor	ARR Error		Ratio
				MT-PGA [79]	MT-WAE	
Asteroid Impact (3D)	7	1,295	13.68	0.13	0.12	0.93
Cloud processes (2D)	12	1,209	13.84	2e-04	7e-07	3e-03
Viscous fingering (3D)	15	118	13.21	8e-04	7e-07	8e-04
Dark matter (3D)	40	316	15.09	2e-04	2e-05	0.08
Volcanic eruptions (2D)	12	811	13.83	0.01	2e-03	0.41
Ionization front (2D)	16	135	13.44	0.19	0.14	0.78
Ionization front (3D)	16	763	13.89	0.24	0.22	0.92
Earthquake (3D)	12	1,203	14.07	0.14	0.10	0.75
Isabel (3D)	12	1,338	14.03	3e-03	2e-03	0.72
Starting Vortex (2D)	12	124	11.92	2e-04	2e-06	0.01
Sea Surface Height (2D)	48	1,787	14.36	0.23	0.22	0.92
Vortex Street (2D)	45	23	20.27	3e-04	9e-05	0.26

Similarly to the metric penalty term, given a blending weight $\lambda_C \in [0, 1]$ (in practice, we set $\lambda_C = 1$), the penalty term $\lambda_C P_C(\theta)$ is added to the expression of the reconstruction energy $E_{W_2^T}(\theta)$ (Eq. 14, main manuscript). The corresponding gradient is estimated by automatic differentiation and the overall energy is optimized by gradient descent [57], as originally described in Sec. 4.5 (main manuscript).

APPENDIX E

DATA REDUCTION EXPERIMENTS

Tables 3 and 4 report a comparison between the reconstruction error generated by our Wasserstein Auto-Encoder (WAE) approach and the Principal Geodesic Analysis (PGA) approach by Pont et al. [79], for the application to data reduction (Sec. 5.1 of the main manuscript), in the case of persistence diagrams (Tab. 3) and merge trees (Tab. 4).

Specifically, we compute the reconstruction error of each input BDT $\mathcal{B}(f_i)$ via the distance W_2^T to its reconstruction (computed by the method under consideration, PGA or WAE). To be comparable across ensembles, this distance is then divided by the maximum W_2^T distance observed among two input BDTs in the ensemble. Finally, this *relative* reconstruction error is *averaged* over all the BDTs of the ensemble.

To enable a fair comparison, we set the number of axis of PGA, noted d_{max} , to 3 (as reported in the original data reduction description [79]) and we set the number of dimensions in the latent space of WAE to the same value (i.e. $d_{ne} = 3$). We also set the maximum size of the PGA origin, noted N_1 , to $0.1|\mathcal{S}_B|$, where $|\mathcal{S}_B|$ is the total number of branches in the ensemble, i.e. $|\mathcal{S}_B| = \sum_{i=1}^N |\mathcal{B}(f_i)|$. Similarly, for WAE, we set the maximum size of the latent output origin $|\mathcal{O}_{n_e}^{out}|$ to $0.1|\mathcal{S}_B|$.

For both methods (PGA and WAE), the compression factor is fixed to a common value on a per ensemble basis. As discussed in the section 5.1 of the main manuscript, the compression factor of WAE is controlled by adjusting, for the last decoding layer, its dimensionality noted $d_{n_e+n_d}$, and the maximum size of its output origin, noted $|\mathcal{O}_{n_e+n_d}^{out}|$.

Both tables show that WAE clearly outperforms PGA [79] in terms of average relative reconstruction error, with an average improvement of 37% for persistence diagrams, and 52% for merge trees.

Finally, note that for each ensemble, the merge tree based clustering [78] computed from the input BDTs is strictly identical to the clustering computed from the reconstructed BDTs. This confirms the viability of our reconstructed BDTs, and their usability for typical visualization and analysis tasks.

APPENDIX F

DIMENSIONALITY REDUCTION EXPERIMENTS

Fig. 16 extends the Figure 12 of the main manuscript to all our test ensembles. It confirms visually the conclusions of the table of aggregated scores (Table 2 of the main manuscript).

In particular, it confirms that WAE behaves as a *trade-off* between the respective advantages of standard techniques, such as MDS [58] and t-SNE [96]. Specifically, MDS is known to preserve the input metric well, while t-SNE tends to better preserve the global structure of the data (i.e. the ground-truth classification), at the expense of metric violation. Our approach (WAE) provides a trade-off between these two extreme behaviors: (i) it improves over MDS in terms of structure preservation (it provides equivalent or better NMI/ARI scores for 11 out of 12 ensembles) and (ii) it improves over t-SNE in terms of metric preservation (it provides an equivalent or better SIM score for 9 out of 12 ensembles). WAE also outperforms VEC-AE and improves PGA on most ensembles. Finally, the combination of our two penalty terms, WAE-MC, simultaneously outperforms MDS on metric preservation and t-SNE on cluster preservation (hence maximizing all criteria at once), for 8 of the 12 ensembles.

APPENDIX G

EMPIRICAL STABILITY EVALUATION

As documented in the original paper [78] introducing the Wasserstein distance between merge trees (W_2^T), saddle swap instabilities in the merge trees are commonly addressed with a *saddle-merging* pre-processing [78], [91]. This procedure consists in moving each branch b up the BDT $\mathcal{B}(f)$, if its saddle is too *close* to that of its parent branch (i.e. closer in normalized f values than a threshold ε_1 , see [78]). As documented by Pont et al. with practical stability evaluations (see Fig. 14 of [78]), this simple saddle-merging pre-processing drastically improves in practice the robustness of the metric W_2^T to additive noise. Thus, this saddle-merging pre-processing is of paramount importance for the practical usage of W_2^T on real-life datasets and Pont et al. recommend to use $\varepsilon_1 = 0.05$ as a default value. Note that this parameter ε_1 acts as a control knob, which balances the practical stability of the metric with its discriminative power (for $\varepsilon_1 = 1$, $W_2^T = W_2^D$).

In this appendix, we study the practical stability of our non-linear framework for merge tree encoding (WAE) to additive noise, in order to document the impact of the underlying metric's stability on the outcome of the analysis.

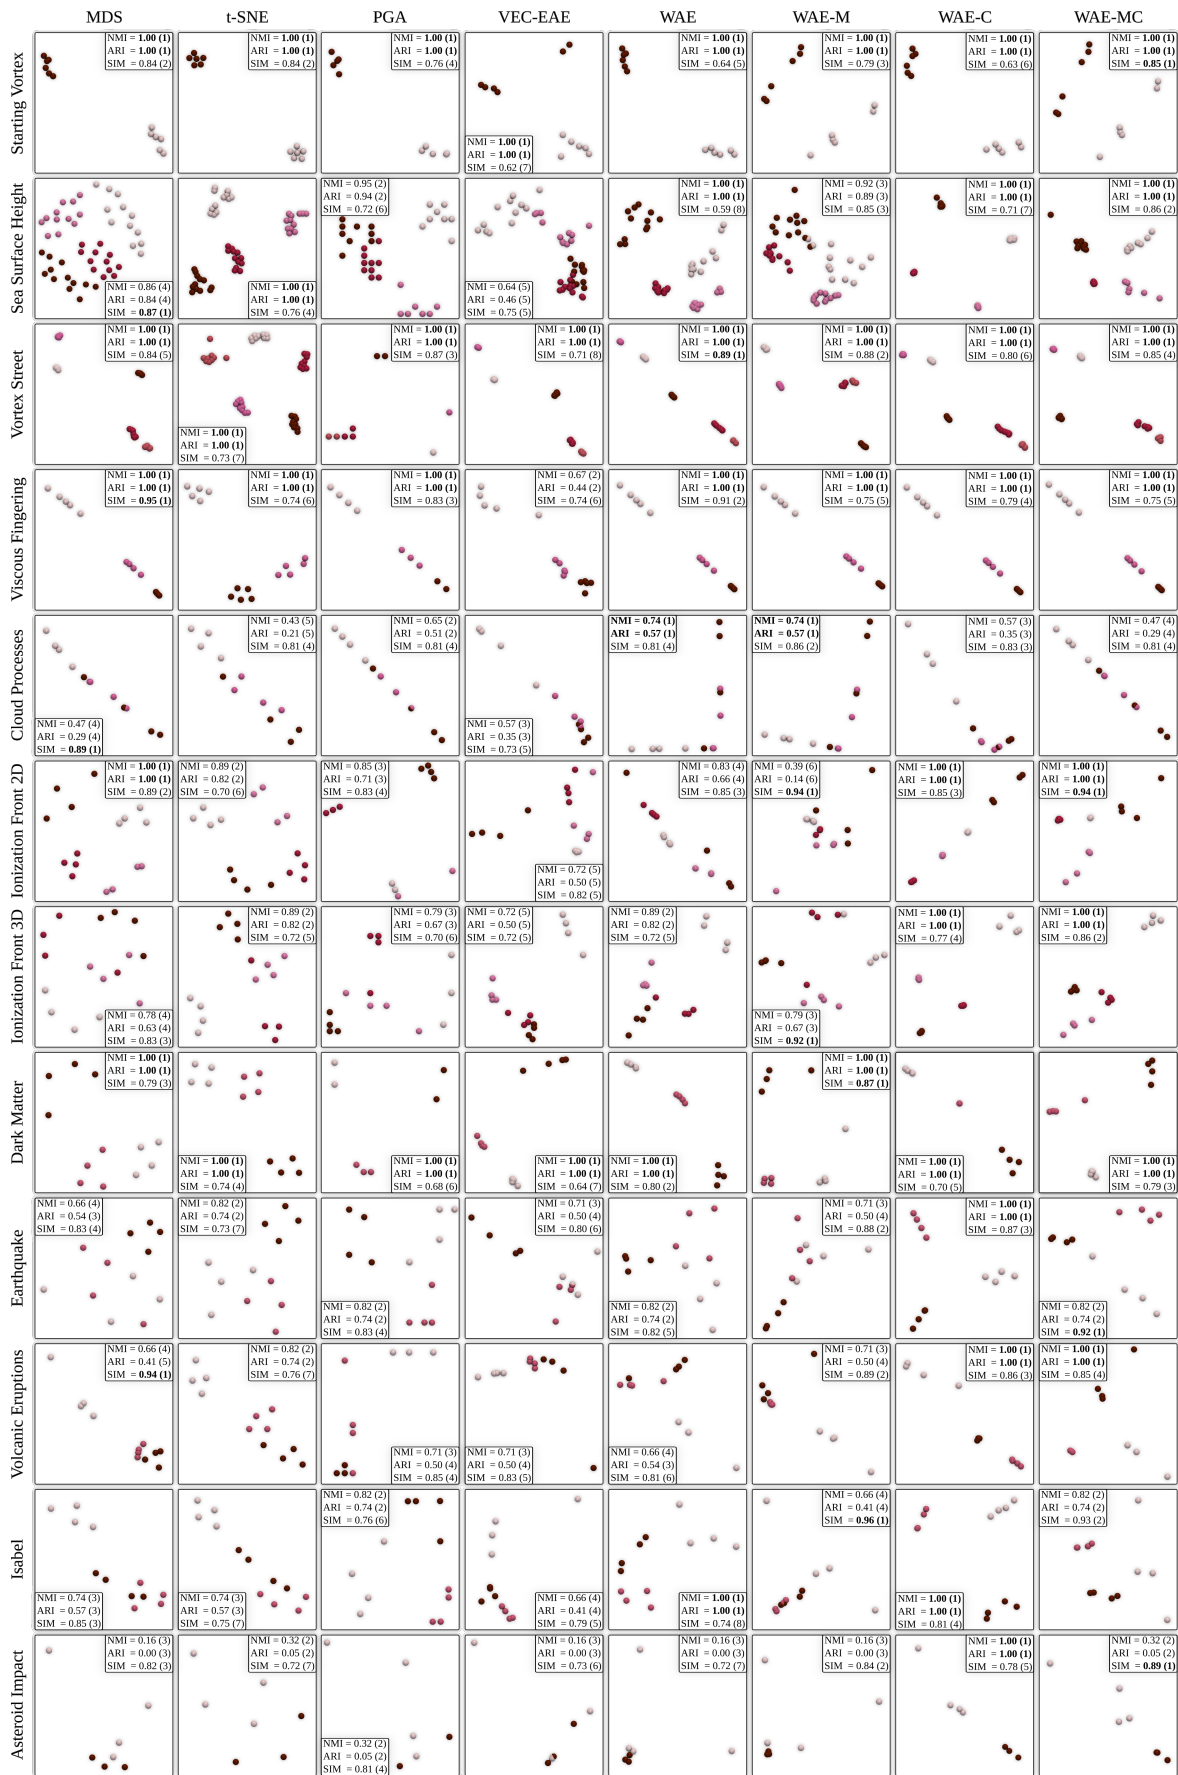


Fig. 16. Comparison of planar layouts for typical dimensionality reduction techniques on all our merge tree ensembles. The color encodes the classification ground-truth [78]. For each quality score, the best value appears bold and the rank of the score among all methods is in parenthesis.

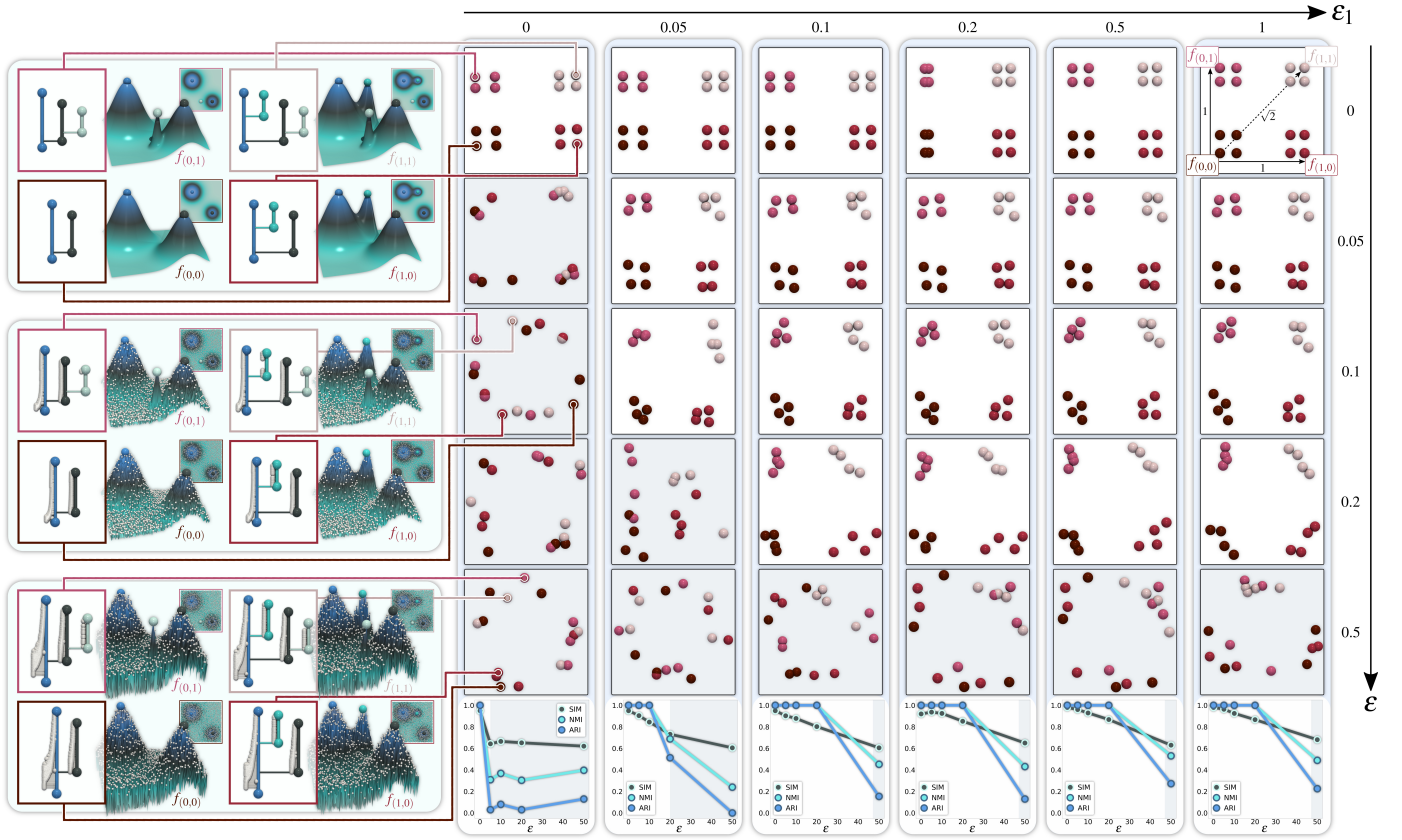


Fig. 17. **Empirical stability evaluation:** a synthetic ensemble of sixteen 2D scalar fields is specifically designed by sampling a 2D basis of Gaussian mixtures, with a controlled parameterization (see Sec. G.1 for a detailed specification). This yields a *ground-truth* parameterization and classification of the ensemble (four clusters: dark red, red, pink, light pink). Five versions of this ensemble are created, for increasing levels of additive noise (from $\epsilon = 0$ to $\epsilon = 0.5$, top to bottom). For each ensemble, a 2D layout is generated by our non-linear framework WAE (right insets), for increasing values of the parameter ϵ_1 from left to right (i.e. from the strict Wasserstein distance between merge trees, W_2^T for $\epsilon_1 = 0$, to progressive blends towards the Wasserstein distance between persistence diagrams, W_2^D for $\epsilon_1 = 1$). In the 2D layout and the quality scores (bottom curves), a grey background indicates an unstable computation (i.e. *NMI* and *ARI* are both below 1). For the default recommended value of the parameter ϵ_1 (0.05 [78]), WAE with W_2^T recovers well the ground-truth parameterization and classification (similarly to W_2^D), up to a level of additive noise of $\epsilon = 0.1$. For $\epsilon_1 \geq 0.1$, the 2D layouts generated by WAE provide a similar level of robustness for W_2^T and W_2^D (bottom curves).

G.1 Setup

For this experiment, we specifically generated a synthetic ensemble, in order to control both its intrinsic parameterization and its classification. For this, we proceeded as follows.

First, four 2D scalar fields (Fig. 17, top left inset) were generated by sampling a 2D basis of Gaussian mixtures with controlled parameterization. Specifically, the scalar field being the origin of the basis has two hills ($f_{(0,0)}$, dark red frame, top left inset). The extremity of the first (horizontal) axis ($f_{(1,0)}$, red frame, top left inset) has exactly the same hills, but with a first *additional* maximum (cyan sphere). The extremity of the second (vertical) axis ($f_{(0,1)}$, pink frame, top left inset) has a second additional maximum (white sphere). Finally, the fourth dataset ($f_{(1,1)}$, light pink frame, top left inset) has both extra maxima (cyan and white spheres). These Gaussian mixtures were generated by adjusting the height of the additional maxima (cyan and white spheres, Fig. 17) such that their diagrams describe a square on the Wasserstein metric space (see the top right 2D layout of Fig. 17):

$$\begin{aligned} W_2^D(\mathcal{D}(f_{(0,0)}), \mathcal{D}(f_{(1,0)})) &= W_2^D(\mathcal{D}(f_{(1,0)}), \mathcal{D}(f_{(1,1)})) \\ &= W_2^D(\mathcal{D}(f_{(1,1)}), \mathcal{D}(f_{(0,1)})) \\ &= W_2^D(\mathcal{D}(f_{(0,1)}), \mathcal{D}(f_{(0,0)})) \\ &= 1, \end{aligned}$$

and:

$$\begin{aligned} W_2^D(\mathcal{D}(f_{(0,0)}), \mathcal{D}(f_{(1,1)})) &= W_2^D(\mathcal{D}(f_{(1,0)}), \mathcal{D}(f_{(0,1)})) \\ &= \sqrt{2}. \end{aligned}$$

Next, we repeated this square generation process, around each corner of the above square, but this time with a smaller side length (equal to 0.15 in the Wasserstein metric space, instead of 1).

Overall, this results in a total of 16 scalar fields, specifically organized along a *ground-truth* 2-dimensional parameterization of the Wasserstein metric space, with a natural *ground-truth* classification (corresponding to the closest corner of the 2D grid, see the top right 2D layout of Fig. 17):

- Class 1 (bottom left corner, dark red spheres in Fig. 17):
– $f_{(0,0)}$, $f_{(0.15,0)}$, $f_{(0.15,0.15)}$, $f_{(0,0.15)}$;

- Class 2 (bottom right corner, red spheres in Fig. 17):
 - $f_{(0.85,0)}, f_{(1,0)}, f_{(1,0.15)}, f_{(0.85,0.15)}$;
- Class 3 (top right corner, light pink spheres in Fig. 17):
 - $f_{(0.85,0.85)}, f_{(1,0.85)}, f_{(1,1)}, f_{(0.85,1)}$;
- Class 4 (top left corner, bright pink spheres in Fig. 17):
 - $f_{(0,0.85)}, f_{(0.15,0.85)}, f_{(0.15,1)}, f_{(0,1)}$;

Given the above ground-truth parameterization, we call the *ground-truth distance matrix*, noted \mathbb{D} , the matrix defined such that each of its entries (i, j) is equal to $W_2^{\mathcal{D}}(\mathcal{D}(f_i), \mathcal{D}(f_j))$.

Next, we generated additional versions of the above ensemble, by introducing a random additive noise in the scalar fields, with a control on the maximum normalized amplitude $\varepsilon \in [0, 1]$ (i.e. the maximum amplitude of the noise is a fraction ε of the global function range of the input scalar field). Specifically, we considered the noise levels $\varepsilon \in \{0, 0.05, 0.1, 0.2, 0.5\}$. Overall, this results in 5 ensembles of 16 scalar fields each.

G.2 Protocol

Given the above ensembles, we first consider our non-linear framework for persistence diagrams, namely PD-WAE. Specifically, we generated, for each noise level, a 2D layout of the ensemble with PD-WAE (see Sec. 5.2, main manuscript). This is shown in the rightmost column of Fig. 17 ($\varepsilon_1 = 1$). We quantitatively evaluate the quality of this 2D layout along two criteria: metric preservation and cluster preservation.

First, given the 2D layout of the ensemble, we compute a distance matrix D in 2D, which we compare to the *ground-truth* distance matrix \mathbb{D} (see Sec. G.1) with the *SIM* indicator [79] (which varies between 0 and 1, 1 being optimal). Second, given the 2D layout of the ensemble, we compute a k -means clustering in 2D (with $k = 4$) and we compare the resulting classification to the *ground-truth* classification with the *NMI* and *ARI* indicators (which vary between 0 and 1, 1 being optimal).

To study the stability to additive noise of our framework when considering the Wasserstein distance between merge trees, we have replicated the above experiment for 5 more values of the control parameter ε_1 (in Fig. 17, from left to right: 0, 0.05, 0.1, 0.2 and 0.5). Overall this results in the 2D array represented in Fig. 17 where each column denotes a specific value of the control parameter ε_1 and where each line denotes a specific noise level ε .

G.3 Analysis

In the absence of noise ($\varepsilon = 0$, top row) and for arbitrary values of the parameter ε_1 , our non-linear WAE framework manages to produce a 2D layout of the ensemble which is faithful to the ground-truth parameterization (high *SIM* values, bottom curves in Fig. 17) and which preserves the ground-truth clusters (colors from dark red to light pink, Fig. 17, high *NMI/ARI* values).

As soon as noise is introduced ($\varepsilon \geq 0.05$), the strict distance $W_2^{\mathcal{T}}$ ($\varepsilon_1 = 0$, leftmost column) becomes unstable, as originally documented by Pont et al. [78]. As a consequence, both the ground-truth classification and parameterization are not recovered by MT-WAE in the 2D layout: spheres of different colors are mixed together (as assessed by the low *NMI/ARI* values, leftmost curves, Fig. 17) and the spheres are no longer organized along a 2D grid (as assessed by the lower *SIM* values, leftmost curve, Fig. 17). In contrast, with the original Wasserstein distance between persistence diagrams ($\varepsilon_1 = 1$, rightmost column), up to a significant level of noise ($\varepsilon = 0.2$), both the ground-truth

classification and parameterization are well preserved in the 2D layout generated by PD-WAE: the spheres with the same color remain clustered (high *NMI/ARI* values, rightmost curves) and the spheres are properly arranged along a 2D grid (high *SIM* values, rightmost curve).

For the recommended value of the control parameter ε_1 (0.05 [78]), MT-WAE still manages to recover well the ground-truth parameterization and classification, up to a noise level of $\varepsilon = 0.1$ (perfect clustering, with high *SIM* values). For a larger value of ε_1 ($\varepsilon_1 \geq 0.1$), the 2D layouts generated by MT-WAE are very similar to these generated with PD-WAE (rightmost column), with identical stability indicators (*SIM* and *NMI/ARI* curves, bottom).

In conclusion, this experiment shows that for mild levels of noise ($\varepsilon < 0.1$), the recommended value of ε_1 (0.05) results in a stable MT-WAE computation. For larger noise levels, MT-WAE provides similar stability scores to PD-WAE for values of ε_1 which are still reasonable in terms of discriminative power ($\varepsilon_1 = 0.1$).