## A. Appendix

## A.1. Persistent Homology Calculation Details

This section provides more details about the persistent homology calculation; it is more geared towards an expert reader and aims for a concise description of all required concepts.

Simplicial homology To understand persistent homology, we first have to understand simplicial homology. Given a simplicial complex $\mathfrak{K}$, i.e. a high-dimensional generalisation of a graph, let $C_{d}(\mathfrak{K})$ denote the vector space generated over $\mathbb{Z}_{2}$ whose elements are the $d$-simplices in $\mathfrak{K}^{7}$. For $\sigma=\left(v_{0}, \ldots, v_{d}\right) \in \mathfrak{K}$, let $\partial_{d}: C_{d}(\mathfrak{K}) \rightarrow C_{d-1}(\mathfrak{K})$ be the boundary homomorphism defined by

$$
\begin{equation*}
\partial_{d}(\sigma):=\sum_{i=0}^{d}\left(v_{0}, \ldots, v_{i-1}, v_{i+1}, \ldots, v_{d}\right) \tag{11}
\end{equation*}
$$

for a single simplex and linearly extended to $C_{d}(\mathfrak{K})$. The $d^{\text {th }}$ homology group $\mathrm{H}_{d}(\mathfrak{K})$ of $\mathfrak{K}$ is defined as the quotient group $\mathrm{H}_{d}(\mathfrak{K}):=\operatorname{ker} \partial_{d} / \operatorname{im} \partial_{d+1}$. The rank of the $d^{\text {th }}$ homology group is known as the $d^{\text {th }}$ Betti number $\beta_{d}$, i.e. $\beta_{d}(\mathfrak{K}):=\operatorname{rank} \mathrm{H}_{d}(\mathfrak{K})$. The sequence of Betti numbers $\beta_{0}, \ldots, \beta_{d}$ of a $d$-dimensional simplicial complex is commonly used to distinguish between different manifolds. For example, a 2 -sphere in $\mathbb{R}^{3}$ has Betti numbers $(1,0,1)$, while a 2 -torus in $\mathbb{R}^{3}$ has Betti numbers $(1,2,1)$. Betti numbers are of limited use for analysing real-world data sets, however, because their representation is too coarse and easily affected by small changes in the underlying simplicial complex. In an idealised, platonic setting, this does not pose a problem, because one assumes that the triangulation of a manifold is known a priori; for real-world data sets, however, we are typically dealing with point clouds and have no knowledge of the underlying manifold, making the calculation of the "proper" simplicial complex nigh impossible. These disadvantages prompted the development of persistent homology.

Persistent homology Let $\emptyset=\mathfrak{K}_{0} \subseteq \mathfrak{K}_{1} \subseteq \cdots \subseteq$ $\mathfrak{K}_{m-1} \subseteq \mathfrak{K}_{m}=\mathfrak{K}$ be a nested sequence of simplicial complexes, called filtration. Filtrations can be defined based on different functions; the Vietoris-Rips filtration that we discuss in the paper, for example, is defined by a distance function, such as the Euclidean distance between points of a point cloud. Notice that we may still calculate the simplicial homology of each $\mathfrak{K}_{i}$ in the filtration. The filtration provides more information, though: the family of boundary operators $\partial(\cdot)$, together with the inclusion homomorphism, induces a homomorphism between corresponding homology

[^0]groups of the filtration, i.e. $f_{d}^{i, j}: \mathrm{H}_{d}\left(\mathfrak{K}_{i}\right) \rightarrow \mathrm{H}_{d}\left(\mathfrak{K}_{j}\right)$. This homomorphism yields a sequence of homology groups
\[

$$
\begin{aligned}
& 0=\mathrm{H}_{d}\left(\mathfrak{K}_{0}\right) \xrightarrow{f_{d}^{0,1}} \mathrm{H}_{d}\left(\mathfrak{K}_{1}\right) \xrightarrow{f_{d}^{1,2}} \ldots \\
& \ldots \xrightarrow{f_{d}^{m-2, m-1}} \mathrm{H}_{d}\left(\mathfrak{K}_{m-1}\right) \xrightarrow{f_{d}^{m-1, m}} \mathrm{H}_{d}\left(\mathfrak{K}_{m}\right)=\mathrm{H}_{d}(\mathfrak{K})
\end{aligned}
$$
\]

for every dimension $d$. Given indices $i \leq j$, the $d^{\text {th }}$ persistent homology group is defined as

$$
\begin{equation*}
\mathrm{H}_{d}^{i, j}:=\operatorname{ker} \partial_{d}\left(\mathfrak{K}_{i}\right) /\left(\operatorname{im} \partial_{d+1}\left(\mathfrak{K}_{j}\right) \cap \operatorname{ker} \partial_{d}\left(\mathfrak{K}_{i}\right)\right) . \tag{12}
\end{equation*}
$$

It can be seen as the homology group that contains all homology classes created in $\mathfrak{K}_{i}$ that are still present ("active") in $\mathfrak{K}_{j}$. We define the $d^{\text {th }}$ persistent Betti number to be the rank of this group, i.e. $\beta_{d}^{i, j}:=\operatorname{rank} \mathrm{H}_{d}^{i, j}$, which generalises the previous definition for simplicial homology. Persistent homology results in a sequence of Betti numbers-instead of a single number-that permits a fine-grained description of topological activity. This activity is typically summarised in a persistence diagram, thus replacing the indices $i, j$ with real numbers based on the function that was used to calculate the filtration.

Persistence diagrams A filtration often has associated values (or weights) $w_{0} \leq w_{1} \leq \cdots \leq w_{m-1} \leq w_{m}$, such as the pairwise distances in a point cloud. These values permit the calculation of topological feature descriptors known as persistence diagrams: for each dimension $d$ and each pair $i \leq j$, one stores a pair $(a, b):=\left(w_{i}, w_{j}\right) \in \mathbb{R}^{2}$ with multiplicity

$$
\begin{equation*}
\mu_{i, j}^{(d)}:=\left(\beta_{d}^{i, j-1}-\beta_{d}^{i, j}\right)-\left(\beta_{d}^{i-1, j-1}-\beta_{d}^{i-1, j}\right) \tag{13}
\end{equation*}
$$

in a multiset (typically, $\mu_{i, j}^{(d)}=0$ for many pairs). The pair $(a, b)$ represents a topological feature that was created at a certain threshold $a$, and destroyed at another threshold $b$. In the case of the Vietoris-Rips filtration and connected components, we have $a=0$ because all connected components are present at the beginning of the filtration by definition. Similarly, $b$ will correspond to an edge used in the minimum spanning tree of the data set. In general, the resulting set of points is called the $d^{\text {th }}$ persistence diagram $\mathcal{D}_{d}$. Given a point $(a, b) \in \mathcal{D}_{d}$, the quantity $\operatorname{pers}(x, y):=|a-b|$ is referred to as its persistence.

## A.2. Proof of Theorem 1

Theorem 1. Let $X$ be a point cloud of cardinality $n$ and $X^{(m)}$ be one subsample of $X$ of cardinality $m$, i.e. $X^{(m)} \subseteq X$, sampled without replacement. We can bound the probability of $X^{(m)}$ exceeding a threshold in terms of the bottleneck distance as

$$
\begin{equation*}
\mathbb{P}\left(\mathrm{d}_{\mathrm{b}}\left(\mathcal{D}^{X}, \mathcal{D}^{X^{(m)}}\right)>\epsilon\right) \leq \mathbb{P}\left(\mathrm{d}_{\mathrm{H}}\left(X, X^{(m)}\right)>2 \epsilon\right) \tag{14}
\end{equation*}
$$



Figure A.1. Empirical convergence rate (mean) of the Hausdorff distance for a subsample of size $m$ of 100 points in a $d$-dimensional space, following a standard normal distribution.
where $\mathrm{d}_{\mathrm{H}}$ refers to the Hausdorff distance between the point cloud and its subsample, i.e.

$$
\begin{align*}
\mathrm{d}_{\mathrm{H}}(X, Y):=\max \{ & \sup _{x \in X} \inf _{y \in Y} \operatorname{dist}(x, y),  \tag{15}\\
& \left.\sup _{y \in Y} \inf _{x \in X} \operatorname{dist}(x, y)\right\}
\end{align*}
$$

for a baseline distance $\operatorname{dist}(x, y)$ such as the Euclidean distance.

Proof. The stability of persistent homology calculations was proved by Chazal et al. (2014a) for finite metric spaces. More precisely, given two metric spaces $X$ and $Y$, we have

$$
\begin{equation*}
\mathrm{d}_{\mathrm{b}}\left(\mathcal{D}^{X}, \mathcal{D}^{Y}\right) \leq 2 \mathrm{~d}_{\mathrm{GH}}(X, Y) \tag{16}
\end{equation*}
$$

where $\mathrm{d}_{\mathrm{GH}}(X, Y)$ refers to the Gromov-Hausdorff distance (Burago et al., 2001, p. 254) of the two spaces. It is defined as the infimum Hausdorff distance over all isometric embeddings of $X$ and $Y$. This distance can be employed for shape comparison (Chazal et al., 2009, Mémoli \& Sapiro, 2004), but is hard to compute. In our case, with $X=X$ and $Y=X^{(m)}$, we consider both spaces to have the same metric (for $Y$, we take the canonical restriction of the metric from $X$ to the subspace $Y$ ). By definition of the GromovHausdorff distance, we thus have $\mathrm{d}_{\mathrm{GH}}(X, Y) \leq \mathrm{d}_{\mathrm{H}}(X, Y)$, so Eq. 15 leads to

$$
\begin{equation*}
\mathrm{d}_{\mathrm{b}}\left(\mathcal{D}^{X}, \mathcal{D}^{Y}\right) \leq 2 \mathrm{~d}_{\mathrm{H}}(X, Y) \tag{17}
\end{equation*}
$$

from which the original claim from Eq. 14 follows by taking probabilities on both sides.

## A.3. Empirical Convergence Rates of $\mathrm{d}_{\mathrm{H}}\left(X, X^{(m)}\right)$

Figure A. 1 depicts the mean of the convergence rate (mean) of the Hausdorff distance for a subsample of size $m$ of 100 points in a $d$-dimensional space, following a standard normal distribution. We can see that the convergence rate is roughly similar, but shown on different absolute levels
that depend on the ambient dimension. While bounding the convergence rate of this expression is feasible (Chazal et al., 2015a,b), it requires more involved assumptions on the measures from which $X$ and $X^{(m)}$ are sampled. Additionally, we can give a simple bound using the diameter $\operatorname{diam}(X):=\sup \{\operatorname{dist}(x, y) \mid x, y \in X\}$. We have $\mathrm{d}_{\mathrm{H}}\left(X, X^{(m)}\right) \leq \operatorname{diam}(X)$ because the supremum is guaranteed to be an upper bound for the Hausdorff distance. This worst-case bound does not account for the sample size (or mini-batch size) $m$, though (see Theorem 2 for an expression that takes $m$ into account).

## A.4. Proof of Theorem 2

Prior to the proof we state two observations that arise from our special setting of dealing with finite point clouds.
Observation 1. Since $X^{(m)} \subseteq \quad X$, we have $\sup _{x^{\prime} \in X^{(m)}} \inf _{x \in X} \operatorname{dist}\left(x, x^{\prime}\right)=0$. Hence, the Hausdorff distance simplifies to:

$$
\begin{equation*}
\mathrm{d}_{\mathrm{H}}\left(X, X^{(m)}\right):=\sup _{x \in X} \inf _{x^{\prime} \in X^{m}} \operatorname{dist}\left(x, x^{\prime}\right) \tag{18}
\end{equation*}
$$

In other words, we only have to consider a "one-sided" expression of the distance because the distance from the subsample to the original point cloud is always zero.
Observation 2. Since our point clouds of interest are finite sets, the suprema and infima of the Hausdorff distance coincide with the maxima and minima, which we will subsequently use for easier readability.

Hence, the computation of $\mathrm{d}_{\mathrm{H}}\left(X, X^{(m)}\right.$ can be divided into three steps.

1. Using the baseline distance $\operatorname{dist}(\cdot, \cdot)$, we compute a distance matrix $\mathbf{A} \in \mathbb{R}^{n \times m}$ between all points in $X$ and $X^{(m)}$.
2. For each of the $n$ points in $X$, we compute the minimal distance to the $m$ samples of $X^{(m)}$ by extracting the minimal distance per row of $\mathbf{A}$ and gather all minimal distances in $\boldsymbol{\delta} \in \mathbb{R}^{n}$.
3. Finally, we return the maximal entry of $\boldsymbol{\delta}$ as $\mathrm{d}_{\mathrm{H}}\left(X, X^{(m)}\right)$.

In the subsequent proof, we require an independence assumption of the samples.

Proof. Using Observations 1 and 2 we obtain a simplified expression for the Hausdorff distance, i.e.

$$
\begin{equation*}
\mathrm{d}_{\mathrm{H}}\left(X, X^{(m)}\right):=\max _{i, 1 \leq i \leq n}\left(\min _{j, 1 \leq j \leq m}\left(a_{i j}\right)\right) \tag{19}
\end{equation*}
$$

The minimal distances of the first $m$ rows of $\mathbf{A}$ are trivially 0 . Hence, the outer maximum is determined by the
remaining $n-m$ row minima $\left\{\delta_{i} \mid m<i \leq n\right\}$ with $\delta_{i}=\min _{1 \leq j \leq m}\left(a_{i j}\right)$. Those minima follow the distribution $F_{\Delta}(y)$ with

$$
\begin{align*}
F_{\Delta}(y) & =\mathbb{P}\left(\delta_{i} \leq y\right)=1-\mathbb{P}\left(\delta_{i}>y\right)  \tag{20}\\
& =1-\mathbb{P}\left(\min _{1 \leq j \leq m} a_{i j}>y\right)  \tag{21}\\
& =1-\mathbb{P}\left(\bigcap_{j} a_{i j}>y\right)  \tag{22}\\
& =1-\left(1-F_{D}(y)^{m}\right)=F_{D}(y)^{m} . \tag{23}
\end{align*}
$$

Next, we consider $Z:=\max _{1 \leq i \leq n} \delta_{i}$. To evaluate the density of $Z$, we first need to derive its distribution $F_{Z}$ :

$$
\begin{align*}
F_{Z}(z) & =\mathbb{P}(Z \leq z)=\mathbb{P}\left(\max _{m<i \leq n} \delta_{i} \leq z\right)  \tag{24}\\
& =\mathbb{P}\left(\bigcap_{m<i \leq n} \delta_{i} \leq z\right) \tag{25}
\end{align*}
$$

Next, we approximate $Z$ by $Z^{\prime}$ by imposing i.i.d sampling of the minimal distances $\delta_{i}$ from $F_{\Delta}$. This is an approximation because in practice, the rows $m+1$ to $n$ are not stochastically independent because of the triangular inequality that holds for metrics. However, assuming i.i.d., we arrive at

$$
\begin{equation*}
F_{Z^{\prime}}(z)=F_{\Delta}(z)^{n-m} \tag{26}
\end{equation*}
$$

Since $Z^{\prime}$ has positive support its expectation can then be evaluated as:

$$
\begin{align*}
\mathbb{E}_{Z^{\prime} \sim F_{Z^{\prime}}}\left[Z^{\prime}\right] & =\int_{0}^{+\infty}\left(1-F_{Z^{\prime}}(z)\right) \mathrm{d} z  \tag{27}\\
& =\int_{0}^{+\infty}\left(1-F_{\Delta}(z)^{n-m}\right) \mathrm{d} z  \tag{28}\\
& =\int_{0}^{+\infty}\left(1-F_{D}(z)^{m(n-m)}\right) \mathrm{d} z  \tag{29}\\
& \geq \int_{0}^{+\infty}\left(1-F_{D}(z)^{(n-1)}\right) \mathrm{d} z \tag{30}
\end{align*}
$$

The independence assumption leading to $Z^{\prime}$ results in overestimating the variance of the drawn minima $\delta_{i}$. Thus, the expected maximum of those minima, $\mathbb{E}\left[Z^{\prime}\right]$, is overestimating the actual expectation of the maximum $\mathbb{E}[Z]$, which is why Eq. 27 to Eq. 29 constitute an upper bound of $\mathbb{E}[Z]$, and equivalently, an upper bound of $\mathbb{E}\left[\mathrm{d}_{\mathrm{H}}\left(X, X^{(m)}\right)\right]$. When increasing $m, \mathbb{E}\left[\mathrm{~d}_{\mathrm{H}}\left(X, X^{m}\right)\right]$ decreases monotonically since for a particular $m$, we draw $n-m$ samples from the minimal
distance distribution $F_{\Delta}$, and their maximum determines the Hausdorff distance. In contrast, our preliminary upper bound on the left-hand side of Eq. 29 forms a downwardsfacing parabola due to the quadratic form in the exponent. This indicates that a tighter bound is achieved for $m \neq n$ by using the minimal subsample size of $m=1$.

## A.5. Synthetic Data Set

SPHERES consists of eleven high-dimensional 100-spheres living in 101 -dimensional space. Ten spheres of radius $r=5$ are each shifted in a random direction (by adding the same Gaussian noise vector per sphere). To this end, we draw ten $d$-dimensional Gaussian vectors following $\mathcal{N}(\mathbf{0}, \mathbf{I}(10 / \sqrt{d}))$ for $d=101$. Crucially, to add interesting topological information to the data set, the ten spheres are enclosed by an additional larger sphere of radius $5 r$. The spheres were generated using the library scikit-tda.

## A.6. Architectures and Hyperparameter Tuning

Architectures for synthetic data set For the synthetically generated data set, we use a simple multilayer perceptron architecture consisting of two hidden layer with 32 neurons each both encoder and decoder and a bottleneck of two neurons such that the sequence of hidden-layer neurons is $32-32-2-32-32$. ReLU non-linearities and batch normalization were applied between the layers excluding the output layer and the bottleneck layer. The networks were fit using mean squared error loss.

Architectures for real world data sets For the MNIST, FASHION-MNIST, and CIFAR-10 data sets, we use an architecture inspired by DeepAE (Hinton \& Salakhutdinov, 2006). This architecture is composed of 3 layers of hidden neurons of decreasing size $(1000-500-250)$ for the encoder part, a bottleneck of two neurons, and a sequence of three layers of hidden neurons in decreasing size $(250-500-1000)$ for the decoder. In contrast to the originally proposed architecture, we applied ReLU non-linearities and batch normalization between the layers as we observed faster and more stable training. For the non-linearities of the final layer, we applied the $\tanh$ nonlinearity, such that the image of the activation matches the range of input images scaled between -1 and 1 . Also here, we applied mean squared error loss.

All neural network architectures were fit using Adam and weight-decay of $10^{-5}$.

Hyperparameter tuning For hyperparameter tuning we apply random sampling of hyperparameters using the scikit-optimize library (scikit-optimize contributers, 2018) with 20 calls per method on all data sets. We select the best model parameters in terms of $\mathrm{KL}_{0.1}$ on the vali-
dation split and evaluate and report it on the test split. To estimate performance means and standard deviations, we repeated the evaluation on an independent test split 5 times by using the best parameters (as identified in the hyperparameter search on the validation split) and refitting the models by resampling the train-validation split.

Neural networks For the neural networks, we sample the learning rate log-uniformly in the range $\left[10^{-4}, 10^{-2}\right]$, the batch size uniformly between $[16,128]$, and for our topological autoencoder method (TopoAE), we sample the regularisation strength $\lambda$ log-uniformly in the range $\left[10^{-1}, 3\right]$. Each model was allowed to train for at most 100 epochs and we applied early stopping with patience $=10$ based on the validation loss.

Competitor methods For t-SNE, we sample the perplexity uniformly in the range $5-50$ and the learning rate loguniformly in the range $10-1000$. For Isomap and UMAP, the number of neighbors included in the computation was varied between $15-500$. For UMAP, we additionally vary the min_dist parameter uniformly between 0 and 1 .

## A.7. Measuring the Quality of Latent Representations

Next to the reconstruction error (if available; please see the paper for a discussion on this), we use a variety of NLDR metrics to assess the quality of our method. Our primary interest concerns the quality of the latent space because, among others, it can be used to visualise the data set. We initially considered classical quality metrics from nonlinear dimensionality reduction (NLDR) algorithms (see Bibal \& Frénay (2019), Gracia et al. (2014), van der Maaten et al. (2009) for more in-depth descriptions), namely
(1) the root mean square error ( $\ell$-RMSE) between the distance matrix of the original space and the latent space (as mentioned in the main text, this is not related to the reconstruction error),
(2) the mean relative rank error ( $\ell$-MRRE), which measures the changes in ranks of distances in the original space and the latent space (Lee \& Verleysen, 2009),
(3) the trustworthiness ( $\ell$-Trust) measure (Venna \& Kaski, 2006), which checks to what extent the $k$ nearest neighbours of a point are preserved when going from the original space to the latent space, and
(4) the continuity ( $\ell$-Cont) measure (Venna \& Kaski, 2006), which is defined analogously to $\ell$-Trust, but checks to what extent neighbours are preserved when going from the latent space to the original space.

All of these measures are defined based on comparisons of the original space and the latent space; the reconstructed space is not used here. As an additional measure, we calculate the Kullback-Leibler divergence between density
distributions of the input space and the latent space. Specifically, for a point cloud $X$ with an associated distance dist, we first use the distance to a measure density estimator (Chazal et al., 2011, 2014b), defined as $\mathrm{f}_{\sigma}^{\mathcal{X}}(x):=$ $\sum_{y \in \mathcal{X}} \exp \left(-\sigma^{-1} \operatorname{dist}(x, y)^{2}\right)$, where $\sigma \in \mathbb{R}_{>0}$ represents a length scale parameter. For dist, we use the Euclidean distance and normalise it between 0 and 1 . Given $\sigma$, we evaluate $\mathrm{KL}_{\sigma}:=\operatorname{KL}\left(\mathrm{f}_{\sigma}{ }^{X} \| \mathrm{f}_{\sigma}{ }^{Z}\right)$, which measures the similarity between the two density distributions. Ideally, we want the two distributions to be similar because this implies that density estimates in a low-dimensional representation are similar to the ones in the original space.

## A.8. Assessing the Batch Size

As we used fixed architectures for the hyperparameter search, the batch size remains the main determinant for the runtime of TopoAE. In Figure A.2, we display trends (linear fits) on how loss measures vary with batch size. Addtionally, we draw runtime estimates. As we applied early stopping, for better comparability, we approximated the epoch-wise runtime by dividing the execution time of a run by its number of completed epochs. Interestingly, these plots suggest that the runtime grows with decreasing batch size (even though the topological computation is more costly for larger batch sizes!). In these experiments, sticking to 0 -dimensional topological features we conclude that the benefit of using mini-batches for neural network training still dominate the topological computations. The few steep peaks most likely represent outliers (the corresponding runs stopped after few epochs, which is why the effective runtime could be overestimated).

For the loss measures, we see that reconstruction loss tends to decrease with increasing batch size, while our topological loss tends to increase with increasing batch size (despite normalization). The second observation might be due to larger batch size enabling more complex data point arrangements and corresponding topologies.

## A.9. Extending to Variational Autoencoders

In Figure A. 3 we sketch a preliminary experiment, where we apply our topological constraint to variational autoencoders for the SPHERES data set. Also here, we observe that our constraint helps identifying the nesting structure of the enclosing sphere.

## A.10. Topological Distance Calculations

To assess the topological fidelity of the resulting latent spaces, we calculate several topological distances between the test data set (full dimensionality) and the latent spaces obtained from each method (two dimensions). More precisely, we calculate (i) the $1^{\text {st }}$ Wasserstein distance ( $W_{1}$ ),


Figure A.2. A scatterplot of batch sizes verses three measures of interest: Topological Loss, Reconstruction Loss, and $\mathrm{KL}_{0.1}$, our objective for the hyperparameter search. Additionally, we draw per-epoch runtime estimates.


Figure A.3. A depiction of latent spaces obtained for the Spheres data set with variational autoencoders (VAEs). Here, VAE represents a standard MLP-based VAE, whereas TopoVAE represents the same architecture plus our topological constraint.
(ii) the $2^{\text {nd }}$ Wasserstein distance $\left(W_{2}\right)$, and (iii) the bottleneck distance $\left(W_{\infty}\right)$ between the persistence diagrams obtained from the test data set of the Spheres data and their resulting 2D latent representations. Even though our loss function is not optimising this distance, we observe in Table A. 1 that the topological distance of our method ("TopoAE") is always the lowest among all the methods. In particular, it is always smaller than the topological distance of the latent space of the autoencoder architecture; this is true for all distance measures, even though $W_{\infty}$, for example, is known to be susceptible to outliers. Said experiment serves as a simple "sanity check" as it demonstrates that the changes induced by our method are beneficial in that they reduce the topological distance of the latent space to the original data set. For a proper comparison of topological features between the two sets of spaces, a more involved approach would be required, though.

## A.11. Alternative Loss Formulations

Our choice of loss function was motivated by the observation that only aligning the persistence diagrams between mini-batches of $\mathcal{X}$ and $\mathcal{Z}$ can lead to degenerate or "meaningless" latent spaces. As a simple example (see Figure A. 4 for a visualisation), imagine three non-collinear points in the input space and the triangle they are forming. Now assume

| Method | $W_{1}$ | $W_{2}$ | $W_{\infty}$ |
| :--- | ---: | ---: | ---: |
| Isomap | $4.32 \pm 0.037$ | $0.477 \pm 0.0045$ | $0.165 \pm 0.00096$ |
| PCA | $4.42 \pm 0.053$ | $0.476 \pm 0.0046$ | $0.158 \pm 0.00108$ |
| t-SNE | $4.38 \pm 0.038$ | $0.478 \pm 0.0045$ | $0.164 \pm 0.00094$ |
| UMAP | $4.47 \pm 0.042$ | $0.478 \pm 0.0045$ | $0.160 \pm 0.00092$ |
| AE | $3.99 \pm 0.037$ | $0.469 \pm 0.0053$ | $0.154 \pm 0.00128$ |
| TopoAE | $3.73 \pm 0.076$ | $0.459 \pm 0.0055$ | $0.152 \pm 0.00268$ |

Table A.1. Topological distances between the test data set and the corresponding latent space. We used subsamples of size $m=500$ and 10 repetitions (obtaining a mean and a standard deviation).


Figure A.4. An undesirable configuration of the latent space of three non-collinear points, resulting in equal persistence diagrams for $\mathcal{X}$ and $\mathcal{Z}$. Pairwise distances are shown as dotted lines. We prevent this by not explicitly minimising the distances between persistence diagrams but by including persistence pairings.
that the latent space consists of the same triangle (in terms of its side lengths) but with permuted labels. A loss term of the form

$$
\begin{equation*}
\mathcal{L}^{\prime}:=\left\|\mathbf{A}^{X}\left[\pi^{X}\right]-\mathbf{A}^{Z}\left[\pi^{Z}\right]\right\|^{2} \tag{31}
\end{equation*}
$$

only measures the distance between persistence diagrams (which would be zero in this situation) and would not be able to penalise such a configuration.


Figure A.5. A depiction of all latent spaces obtained for the Spheres data set. TopoAE used a batch size of 28 . This is an enlarged version of the figure shown in Section 5.2.


Figure A.6. Latent representations of the FASHION-MNIST data set. TopoAE used a batch size of 95. This is a larger extension of the figure shown in Section 5.2.


Figure A.7. Latent representations of the MNIST data set. TopoAE used a batch size of 126. This is a larger extension of the figure shown in Section 5.2.


Figure A.8. Latent representations of the CIFAR-10 data set. TopoAE used a batch size of 82 . This is a larger extension of the figure shown in Section 5.2.

| Data set | Method | $\mathrm{KL}_{0.001}$ | $\mathrm{KL}_{0}$. |
| :--- | :--- | :---: | :---: | :---: | :---: | :---: | :---: |

Table A.2. Extended version of the table from the main paper, showing more length scales and variance estimates.


[^0]:    ${ }^{7}$ It is also possible to describe this calculation with coefficients in other fields, but the case of $\mathbb{Z}_{2}$ is advantageous because it simplifies the implementation of all operations.

