

Visualization of AE’s Training on Credit Card Transactions with Persistent Homology

Jeremy Charlier^{1,2}, Francois Petit¹, Gaston Ormazabal², Radu State¹, and Jean Hilger³

¹ University of Luxembourg, L-1855 Luxembourg, Luxembourg
{name.surname}@uni.lu

² Columbia University, New York NY 10027, USA
{jjc2292,gso70}@columbia.edu

³ BCEE, L-1160 Luxembourg, Luxembourg
j.hilger@bcee.lu

Abstract. Auto-encoders are among the most popular neural network architecture for dimension reduction. They are composed of two parts: the encoder which maps the model distribution to a latent manifold and the decoder which maps the latent manifold to a reconstructed distribution. However, auto-encoders are known to provoke chaotically scattered data distribution in the latent manifold resulting in an incomplete reconstructed distribution. Current distance measures fail to detect this problem because they are not able to acknowledge the shape of the data manifolds, i.e. their topological features, and the scale at which the manifolds should be analyzed. We propose Persistent Homology for Wasserstein Auto-Encoders, called PHom-WAE, a new methodology to assess and measure the data distribution of a generative model. PHom-WAE minimizes the Wasserstein distance between the true distribution and the reconstructed distribution and uses persistent homology, the study of the topological features of a space at different spatial resolutions, to compare the nature of the latent manifold and the reconstructed distribution. Our experiments underline the potential of persistent homology for Wasserstein Auto-Encoders in comparison to Variational Auto-Encoders, another type of generative model. The experiments are conducted on a real-world data set particularly challenging for traditional distance measures and auto-encoders. PHom-WAE is the first methodology to propose a topological distance measure, the bottleneck distance, for Wasserstein Auto-Encoders used to compare decoded samples of high quality in the context of credit card transactions.

Keywords: Barcodes · Encoding-Decoding · Persistence Diagram.

1 Motivation

Dimension reduction techniques were initially driven by linear algebra with second order matrix decompositions such as the Singular Value Decomposition (SVD) [1] and, ultimately, with tensor decompositions [2, 3], a higher order analogue of the matrix decompositions. Although SVD is fast enough to be applied

on large data sets, it is limited to second order because it relies on matrices. While tensors are suitable for third-order experiments and above, the complexity of the decomposition and the time required for the computations become a limitation as soon as the size of the data set grows [4, 5]. Hence, in recent years, several architectures for dimension reductions based on neural networks have been proposed. Dense Auto-Encoders (AE) [6, 7] are one of the most well established approach. More recently, the Variational Auto-Encoders (VAE) presented by Kingma et al. in [8] constitute a well-known approach but they might generate poor target distribution because of the KL divergence.

We recall an AE is a neural network trained to copy its input manifold to its output manifold through a hidden layer. The encoder function sends the input space to the hidden space and the decoder function brings back the hidden space to the input space. As explained in [9], the points of the hidden space are chaotically scattered for most of the encoders, including the popular VAE. Even after proper training, groups of points of various sizes gather and cluster by themselves randomly in the hidden layer. Therefore, some features are missing in the reconstructed distribution $G(Z)$, $Z \in \mathcal{Z}$. The description of the scattered points is very complex using traditional distance measures, such as the Euclidean distance, because they are not able to acknowledge the shapes of the data manifolds. However, persistent homology is specifically designed to highlight the topological features of the data [10]. Therefore, building upon persistent homology, Wasserstein distance [11] and Wasserstein Auto-Encoders (WAE) [12], our main contribution is to propose qualitative and quantitative ways to evaluate the scattered hidden space and the overall performance of AE.

In this work we describe the persistent homology features of the encoded and decoded model distribution while minimizing the Optimal Transport (OT) function $W_c(P_X, P_G)$ for a squared cost $c(x, y) = \|x - y\|_2^2$ where P_X is the model distribution of the data, and P_G the latent variable model specified by the prior distribution P_Z of latent manifold $Z \in \mathcal{Z}$ and the generative model $P_G(X|Z)$ of the initial distribution $X \in \mathcal{X}$ given Z . The method is highlighted in figure 1. Our contributions are summarized below:

- A persistent homology procedure for WAE which we call PHom-WAE to highlight the topological properties of the encoded and decoded distribution of the data for different spatial resolutions. The objective is twofold: a persistent homology description of the encoded latent space $Q_z := \mathbb{E}_{P_X}[Q(Z|X)]$, and a persistent homology description of the decoded latent space following the generative model $P_G(X|Z)$.
- A distance measure for persistence diagrams, the bottleneck distance, applied to WAE to compare quantitatively the true and the target distributions on any data set. We measure the shortest distance for which there exists a perfect matching between the points of the two persistence diagrams. A persistence diagram is a stable summary representation of topological features of simplicial complex, a collection of vertices, associated to the data set.

- A persistent homology procedure to highlight the scattered latent distribution provoked by Variational Auto-Encoders (VAE) in comparison to WAE on a real-word data set. We use the concepts introduced for PHom-WAE to define PHom-VAE, Persistent Homology for VAE. We highlight the VAE’s hidden layer scattered distribution using persistent homology, confirming the original statement of VAE’s scattered distribution introduced in [13].
- Finally, we propose the first application of algebraic topology and WAE on a public data set containing credit card transactions, particularly challenging for traditional distance measures and auto-encoders.

The paper is structured as follows. We discuss related works in section 2. In section 3, we review the WAE formulation using OT derived by Tolstikhin et al. in [12]. By using persistence homology, we are able to compare the topological properties of the distributions P_X, P_Z and P_G , illustrated in figure 1. We highlight the experimental results in section 4 and we conclude in section 5 by addressing promising directions for future work.

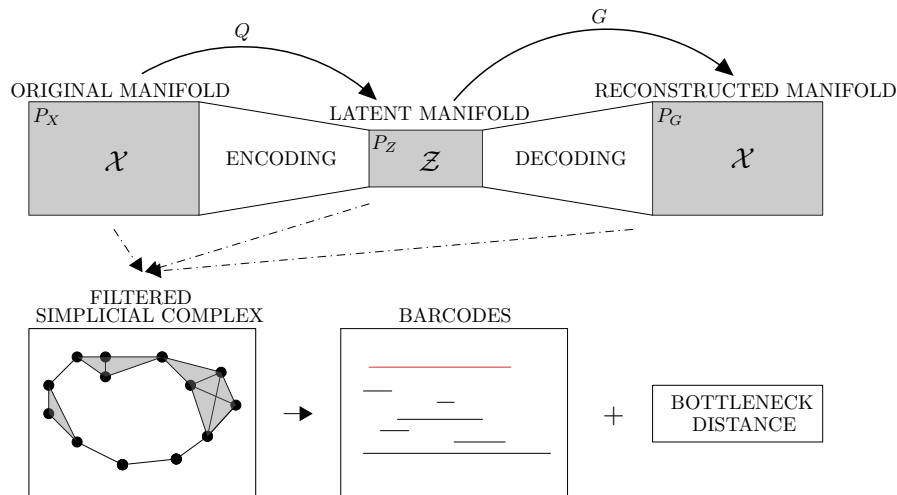


Fig. 1. In PHom-WAE, the function $W_c(P_X, P_G)$ between the true data distribution P_X and the latent variable model P_G of the manifold \mathcal{X} is minimized. P_G is specified by the prior distribution P_Z of the latent manifold \mathcal{Z} and the generative model $P_G(X|Z)$ with $X \in \mathcal{X}, Z \in \mathcal{Z}$. Then, the samples X, Z of each distribution P_X, P_Z and P_G are transformed independently in a metric data set to obtain filtered simplicial complex. It leads to the construction of the persistence diagram, summarized by the barcodes, to compute the bottleneck distance used to compare homologically the respective distribution P_X, P_Z and P_G .

2 Related Work

Literature on Persistence Homology and Topology A major trend in modern data analysis is to consider that the data has a shape, and more precisely, a topological structure. Topological Data Analysis (TDA) is a set of tools relying on computational algebraic topology allowing to obtain precise information about the data structure. Two of the most important techniques are persistence homology and the mapper algorithm [10].

Data sets are usually represented as points cloud of Euclidean spaces. Hence, the shape of a data set depends of the scale at which it is observed. Instead of trying to find an optimal scale, persistence homology, a method of TDA, studies the changes of the topological features (number of connected components, number of holes, ...) of a data set depending of the scale. The foundations of persistence homology have been established in the early 2000s in [14], [15], [16] and [17]. They provide a computable algebraic invariant of filtered topological spaces (nested sequences of topological spaces which encode how the scale changes) called persistence module. This module can be decomposed into a family of intervals called persistence diagram or barcodes. This family records how the topology of the space is changing when going through the filtration [18]. The space of barcodes is measurable through the bottleneck distance. Moreover, the space of persistence module is also endowed with a metric and under a mild assumptions these two spaces are isometric [19]. Additionally, the *Mapper* algorithm first appeared in [20]. It is a visualization algorithm which aims at produce a low dimensional representation of high-dimensional data sets in form of a graph, therefore, capturing the topological features of the points cloud.

Meanwhile, efficient and fast algorithms have emerged to compute persistence homology [15], [17] as well as to construct filtered topological spaces using, for instance, the Vietoris-Rips complex [21]. Therefore, it has already found numerous successful applications. For instance, Nicolau et al. in [22] detected subgroups of cancers with unique mutational profile. In [23], it has been shown that computational topology could be used in medicine, social sciences or sport analysis. More recently, Bendich et al. improved statistical analyses of brain arteries of a population [24] while Xia et al. were capable of extracting molecular topological fingerprints for proteins in biological science [25].

Literature on Auto-Encoders and Optimal Transport A large variety of AE have appeared in the last few years [13]. Although promising results were achieved, most of the solutions did not address the representation of the samples of the encoded and decoded manifolds. As outlined by Bengio et al. in [9], the points in the encoded manifold \mathcal{Z} for the majority of the encoders are chaotically scattered. Therefore, some features are missing in the reconstructed distribution $G(Z)$, $Z \in \mathcal{Z}$. Thus, sampling data points for the reconstruction with traditional AE is difficult. The added constraint of Variational Auto-Encoder (VAE) in [8] by the mean of a KL divergence, composed of a reconstruction cost and a regu-

larization term, provides a finer solution to generate adversarial data by reducing the impact of the chaotic scattered distribution P_Z of \mathcal{Z} .

Concurrently to the emergence of AE, Goodfellow et al. introduced the Generative Adversarial Network (GAN) model in [26] in 2014. Although the GAN did not have an encoder, it consists of two parts, a generator to generate adversarial samples and a discriminator to fit the generated data points to the true data distribution. However, the GAN suffers from a mode collapse between the generator and the discriminator [13]. As a solution, optimal transport [27] was applied to GAN in [28] with the Wasserstein distance, and therefore, introducing Wasserstein GAN (WGAN). By adding a gradient penalty to the Wasserstein distance, Gulrajani et al. in [29] introduced a new training for GANs while avoiding the mode collapse. As described in [30, 31] in the context of unbalanced optimal transport, Tolstikhin et al. applied these concepts to AE in [12]. They proposed to add one extra divergence to the objective minimization function in the context of generative modeling leading to Wasserstein Auto-Encoders (WAE).

In this paper, using persistent homology and the bottleneck distance, we propose qualitative and quantitative ways to evaluate the performance of the compression of AE. We build upon the work of WAE and unbalanced OT with persistent homology. We show that the barcodes, inherited from the persistence diagrams, are capable of representing the encoded manifold \mathcal{Z} generated by the WAE. Furthermore, we show that the bottleneck distance allows to compare quantitatively the topological features between the samples $Z \in \mathcal{Z}$ of the reconstructed distribution $G(Z)$ and the samples $X \in \mathcal{X}$ of the true distribution.

3 Proposed Method

Our method computes the persistent homology of both the latent manifold $Z \in \mathcal{Z}$ and the reconstructed manifold following the generative model $P_G(X|Z)$ based on the minimization of the optimal transport cost $W_c(P_X, P_G)$. In the resulting topological problem, the points of the manifolds are transformed to a metric space set for which a Vietoris-Rips simplicial complex filtration is applied (see definition 2). PHom-WAE achieves simultaneously two main goals: it computes the lifespan of the persistent homological features while measuring the bottleneck distance between the persistence diagrams of the WAE’s manifolds.

3.1 Preliminaries and Notations

We follow the notations used by Tolstikhin et al. in [12]. Sets are denoted by calligraphic letters such as \mathcal{X} , random variables by capital letters X , and their values by lower case letters x . Probability distributions are denoted by capital letters $P(X)$ and their corresponding densities by lower case letters $p(x)$.

3.2 Optimal Transport and Dual Formulation

Following the description of the optimal transport problem [27] and relying on the Kantorovich-Rubinstein duality, the Wasserstein distance is computed as

$$W_c(P_X, P_G) = \sup_{f \in \mathcal{F}_L} \mathbb{E}_{X \sim P_X}[f(X)] - \mathbb{E}_{Y \sim P_G}[f(Y)] \quad (1)$$

where (\mathcal{X}, d) is a metric space, $\mathcal{P}(X \sim P_X, Y \sim P_G)$ is a set of all joint distributions (X, Y) with marginals P_X and P_G respectively and \mathcal{F}_L is the class of all bounded 1-Lipschitz functions on (\mathcal{X}, d) .

3.3 Wasserstein Auto-Encoders

As described in [12], the WAE objective function is expressed such that

$$D_{\text{WAE}}(P_X, P_G) := \inf_{Q(Z|X) \in \mathcal{Q}} \mathbb{E}_{P_X} \mathbb{E}_{Q(Z|X)}[c(X, G(Z))] + \lambda \mathcal{D}_Z(Q_Z, P_Z) \quad (2)$$

where $c(X, G(Z)) : \mathcal{X} \times \mathcal{X} \rightarrow \mathcal{R}_+$ is any measurable cost function. In our experiments, we use a square cost function $c(x, y) = \|x - y\|_2^2$ for data points $x, y \in \mathcal{X}$. $G(Z)$ denotes the sending of Z to X for a given map $G : \mathcal{Z} \rightarrow \mathcal{X}$. Q , and G , are any nonparametric set of probabilistic encoders, and decoders respectively.

We use the Maximum Mean Discrepancy (MMD) for the penalty $\mathcal{D}_Z(Q_Z, P_Z)$ for a positive-definite reproducing kernel $k : \mathcal{Z} \times \mathcal{Z} \rightarrow \mathcal{R}$

$$\begin{aligned} \mathcal{D}_Z(P_Z, Q_Z) &:= \text{MMD}_k(P_Z, Q_Z) \\ &= \left\| \int_{\mathcal{Z}} k(z, \cdot) dP_Z(z) - \int_{\mathcal{Z}} k(z, \cdot) dQ_Z(z) \right\|_{\mathcal{H}_k} \end{aligned} \quad (3)$$

where \mathcal{H}_k is the reproducing kernel Hilbert space of real-valued functions mapping on \mathcal{Z} . For details on the MMD implementation, we refer to [12].

3.4 Vietoris-Rips Complex, Persistence Diagram and Barcodes

We explain the construction of the persistence module associated to a sample of a fixed distribution on a space. First, two manifold distributions are sampled from the WAE's training. Then, we construct the persistence modules associated to each sample of the points manifolds. We refer to the first subsection of section 2 for pointed reference on persistent homology.

We associate to our points manifold $\mathcal{C} \subset \mathbb{R}^n$, considered as a finite metric space, a sequence of simplicial complexes. For that aim, we use the Vietoris-Rips complex.

Definition 1 Let $V = \{1, \dots, |V|\}$ be a set of vertices. A simplex σ is a subset of vertices $\sigma \subseteq V$. A simplicial complex K on V is a collection of simplices $\{\sigma\}$, $\sigma \subseteq V$, such that $\tau \subseteq \sigma \in K \Rightarrow \tau \in K$. The dimension $n = |\sigma| - 1$ of σ is its number of elements minus 1. Simplicial complexes examples are represented in figure 2.

Definition 2 Let (X, d) be a metric space. The Vietoris-Rips complex at scale ϵ associated to X , denoted by $\text{VR}(X, \epsilon)$, is the abstract simplicial complex whose vertex set is X , and where $\{x_0, x_1, \dots, x_k\}$ is a k -simplex if and only if $d(x_i, x_j) \leq \epsilon$ for all $0 \leq i, j \leq k$.

We obtain an increasing sequence of Vietoris-Rips complex by considering the $\text{VR}(\mathcal{C}, \epsilon)$ for an increasing sequence $(\epsilon_i)_{1 \leq i \leq N}$ of value of the scale parameter ϵ

$$\mathcal{K}_1 \hookrightarrow \mathcal{K}_2 \hookrightarrow \mathcal{K}_3 \hookrightarrow \dots \hookrightarrow \mathcal{K}_{N-1} \hookrightarrow \mathcal{K}_N. \quad (4)$$

Applying the k -th singular homology functor $H_k(-, F)$ with coefficient in the field F [32] to (4), we obtain a sequence of F -vector spaces, called the k -th persistence module of $(\mathcal{K}_i)_{1 \leq i \leq N}$

$$H_k(\mathcal{K}_1, F) \xrightarrow{t_1} H_k(\mathcal{K}_2, F) \xrightarrow{t_2} \dots \xrightarrow{t_{N-2}} H_k(\mathcal{K}_{N-1}, F) \xrightarrow{t_{N-1}} H_k(\mathcal{K}_N, F). \quad (5)$$

Definition 3 $\forall i < j$, the (i, j) -persistent k -homology group with coefficient in F of $\mathcal{K} = (\mathcal{K}_i)_{1 \leq i \leq N}$ denoted $HP_k^{i \rightarrow j}(\mathcal{K}, F)$ is defined to be the image of the homomorphism $t_{j-1} \circ \dots \circ t_i : H_k(\mathcal{K}_i, F) \rightarrow H_k(\mathcal{K}_j, F)$.

Using the interval decomposition theorem [33], we extract a finite family of intervals of \mathbb{R}_+ called persistence diagram. Each interval can be considered as a point in the set $D = \{(x, y) \in \mathbb{R}_+^2 \mid x \leq y\}$. Hence, we obtain a finite subset of the set D . This space of finite subsets is endowed with a matching distance called the bottleneck distance and defined as follow

$$d_b(A, B) = \inf_{\phi: A' \rightarrow B'} \sup_{x \in A'} \|x - \phi(x)\|$$

where $A' = A \cup \Delta$, $B' = B \cup \Delta$, $\Delta = \{(x, y) \in \mathbb{R}_+^2 \mid x = y\}$ and the inf is over all the bijections from A' to B' .

Application We illustrate the construction of the barcodes and persistence diagram with the filtration parameter ϵ in figure 3, according to the previous definitions. For every data point, the size of the points is continuously and artificially increased using a filtration parameter ϵ . The points are, therefore, transformed to geometrical disks as the filtration parameter keeps growing. When two disks intersect, a line is drawn between the two corresponding original data points, creating a connected component defined as 1-simplex, while a barcode is drawn in a separate diagram. The barcodes highlight the birth-death cycles of each homological groups, H_0 for the connected components and H_1 for the loops. At the

end of the filtration procedure, a persistence diagram is drawn to recapitulate the birth-death events observed with the barcodes, as shown in figure 4. The persistence diagram is used to describe the topological properties of the original data points cloud using quantitative measures, such as the bottleneck distance.




Fig. 2. A simplicial complex is a collection of numerous “simplex” or simplices, where a 0-simplex is a point, a 1-simplex is a line segment, a 2-simplex is a triangle and a 3-simplex is a tetrahedron.

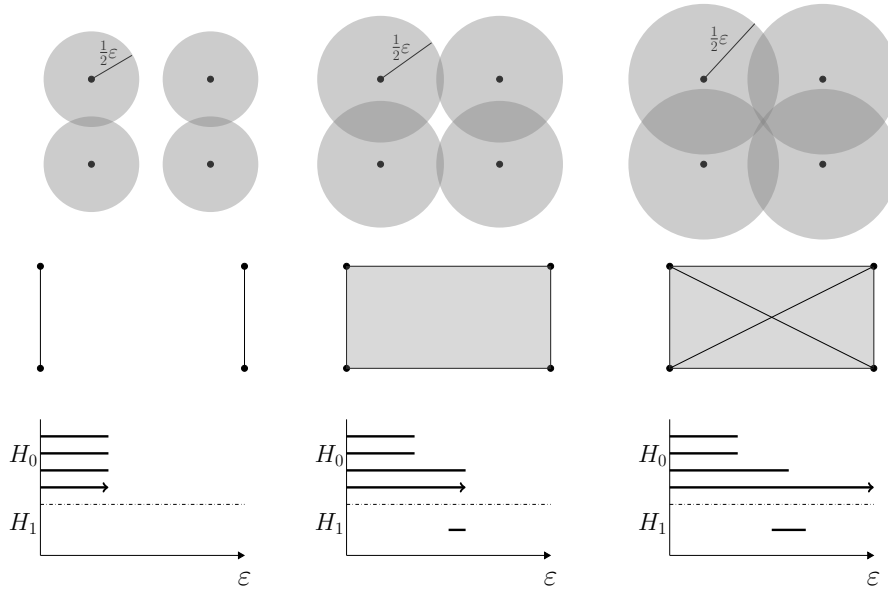


Fig. 3. To describe the persistent homological features of a data set, a filtration parameter ε grows around each data point, leading to geometrical disks. When two disks intersect, a line is drawn between the two original data points, creating a 1-simplex. The corresponding barcode is drawn to mark the birth-death cycle of the event according to the filtration parameter and the homological group observed. The 0-dimensional group H_0 denotes the connected components and the 1-dimensional group H_1 the loops.




Fig. 4. The filtration procedure leads to the construction of the barcodes of the homological groups H_0 and H_1 , a stable summary representation of the persistence diagrams.

3.5 PHom-WAE, Persistent Homology for Wasserstein Auto-Encoder

Bridging the gap between topology and neural networks, PHom-WAE uses a two-steps procedure. First, the minimization problem of WAE is solved for the encoder Q and the decoder G . We use Adam optimizer [34] for the optimization procedure. Then, the samples of the encoded and decoded distributions, P_Z and P_G , are mapped to persistence homology to describe their respective manifold.

We highlight the topological features of the WAE’s manifolds based on their respective distributions with persistent homology. First, the points contained in the manifold \mathcal{Z} inherited from P_Z , the manifold \mathcal{X} from P_X and $P_G(X|Z)$ are randomly selected into respective batches. Two samples, Y_1 from \mathcal{X} following P_X and Y_2 from \mathcal{X} following $P_G(X|Z)$, are selected to differentiate the topological features of the manifold \mathcal{X} before and after the encoding-decoding. Similarly, two other samples Y_1, Y_2 are randomly selected from \mathcal{Z} following P_Z to detect the scattered distribution of the manifold \mathcal{Z} after the encoding. The samples Y_1 and Y_2 are contained in the spaces \mathcal{Y}_1 and \mathcal{Y}_2 , respectively. Then, the spaces \mathcal{Y}_1 and \mathcal{Y}_2 are transformed to metric space sets $\hat{\mathcal{Y}}_1$ and $\hat{\mathcal{Y}}_2$ for computational purposes. Then, we filter the metric space sets $\hat{\mathcal{Y}}_1$ and $\hat{\mathcal{Y}}_2$ using the Vietoris-Rips simplicial complex filtration. Given a line segment of length ϵ , vertices between data points are created for data points respectively separated from a smaller distance than ϵ . It leads to the construction of a collection of simplices resulting in Vietoris-Rips simplicial complex $VR(\mathcal{C}, \epsilon)$ filtration. In our case, we decide to use the Vietoris-Rips simplicial complex as it offers the best compromise between the filtration accuracy and the memory requirement [10]. Subsequently, the persistence diagrams, dgm_{Y_1} and dgm_{Y_2} , are constructed. We recall a persistence diagram is a stable summary representation of topological features of simplicial complex. The persistence diagrams allow the computation of the bottleneck distance $d_b(dgm_{Y_1}, dgm_{Y_2})$. The barcodes represent the lifespan of the homological features detected by the persistence diagrams, for instance the holes. The barcodes are a collection of the interval modules. Furthermore, the lifespan of a homological feature is defined by the boundaries of its interval module, respec-

tively the birth time and the death time. Therefore, the barcodes illustrate in a simple way the birth-death of the pairing generators of the iterated inclusions. PHom-WAE is described in Algorithm 1.

4 Experiments

In this section, we empirically evaluate the proposed methodology PHom-WAE. We assess on a challenging data set for auto-encoders whether PHom-WAE can simultaneously achieve (i) accurate topological reconstruction of the data points and (ii) appropriate persistent homology mapping of the latent manifold.

Data Availability and Data Description We train PHom-WAE on one real-world open data set: the credit card transactions data set from the Kaggle database ⁴ containing 284807 transactions including 492 frauds. This data set is particularly interesting because it underlines the chaotically scattered points of the encoded manifold that are found during the AE training. Furthermore, this data set is challenging because of the strong imbalance between normal and fraudulent transactions while being of high interest for the banking industry. To preserve transactions confidentiality, each transaction is composed of 28 components obtained with PCA without any description and two additional features

⁴ The data set is available at <https://www.kaggle.com/mlg-ulb/creditcardfraud>.

Algorithm 1: Persistent Homology for Wasserstein Auto-Encoder

Data: training and validation sets, hyperparameter λ , encoder Q_ϕ , decoder G_θ
Result: persistent homology description of WAE’s manifolds

```

1 begin
2   /*Step 1: WAE Resolution*/
3   Select samples  $\{x_1, \dots, x_n\}$  from training set
4   Select samples  $\{z_1, \dots, z_n\}$  from validation set
5   Optimize until convergence  $Q_\phi$  and  $G_\theta$  using equation 2 and Adam gradient
      descent update ( $\text{lr} = 0.001, \beta_1 = 0.9, \beta_2 = 0.999$ )
6   /*Step 2: Persistence Diagram and Bottleneck Distance on WAE’s
      manifolds*/
7   Random selection of samples  $Y_1 \in \mathcal{Y}_1, Y_2 \in \mathcal{Y}_2$  from  $P_Z$  or  $P_X$  and  $P_G(X|Z)$ 
8   Transform  $\mathcal{Y}_1$  and  $\mathcal{Y}_2$  spaces into a metric space sets
9   Filter the metric space set with a Vietoris-Rips simplicial complex  $\text{VR}(\mathcal{C}, \epsilon)$ 
10  Compute the persistence diagrams  $\text{dgm}_{Y_1}$  and  $\text{dgm}_{Y_2}$ 
11  Evaluate the bottleneck distance  $d_b(\text{dgm}_{Y_1}, \text{dgm}_{Y_2})$ 
12  Build the barcodes with respect to  $Y_1$  and  $Y_2$ 
13 return
```

Time and *Amount* that remained unchanged. Each transaction is labeled as fraudulent or normal in a feature called *Class* which takes a value of 1 in the case of fraud or 0 otherwise.

Experimental Setup and Code Availability In our experiments, we use the Euclidean latent space $\mathcal{Z} = \mathcal{R}^2$ and the square cost function c previously defined as $c(x, y) = \|x - y\|_2^2$ for the data points $x, y \in \mathcal{X}$. The dimensions of the true data set is \mathcal{R}^{29} . We kept the 28 components obtained with PCA and the amount resulting in a space of dimension 29. For the error minimization process, we used Adam gradient descent [34] with the parameters $\text{lr} = 0.001, \beta_1 = 0.9, \beta_2 = 0.999$ and a batch size of 64. Different values of λ for Wasserstein penalty have been tested, we empirically obtained the lowest error reconstruction with $\lambda = 15$. The coefficients of persistence homology are evaluated within the field $\mathbb{Z}/2\mathbb{Z}$. We only consider homology groups H_0 and H_1 who represent the connected components and the loops, respectively. Higher dimensional homology groups did not noticeably improve the results quality while leading to longer computational time. The simulations were performed on a computer with 16GB of RAM, Intel i7 CPU and a Tesla K80 GPU accelerator. To ensure the reproducibility of the experiments, the code is available at the following address⁵.

Results and Discussions on PHom-WAE against PHom-VAE We test PHom-WAE against PHom-VAE, Persistent Homology for Variational Auto-Encoder. We recall a Variational Auto-Encoder [8] uses a KL divergence, denoted by $D_{KL}(P_X, P_G)$, composed of a reconstruction cost and a regularization term instead of an OT cost function. We use the same concepts introduced for PHom-WAE to define PHom-VAE. We compare the performance of PHom-WAE and PHom-VAE on two specificities: the latent manifold \mathcal{Z} and the reconstructed data distribution $G(Z)$ following the generative model distribution P_G for the samples $Z \in \mathcal{Z}$.

As pictured in figures 5 and 6, both the persistence diagram and the barcodes between the original and the reconstructed distributions, respectively P_X and $P_G(X|Z)$, of the manifold \mathcal{X} are more widely distributed for PHom-WAE than for PHom-VAE. Additionally, the persistence diagram and the barcodes of PHom-WAE are qualitatively closer to those associated with the original data manifold \mathcal{X} . It means the topological features are better preserved for PHom-WAE than for PHom-VAE. It highlights a better encoding-decoding process thanks to the use of an optimal transport cost function. Furthermore, in figure 7, a topological representation of the original and the reconstructed distributions is highlighted. We observe the iterated inclusion chains are more similar for PHom-WAE than for PHom-VAE. For PHom-VAE, the inclusions of the reconstructed distribution are randomly scattered through the space without connected vertices.

⁵ The code is available at <https://github.com/dagrate/PHom-WAE>.




Fig. 5. PHom-WAE and PHom-VAE’s rotated persistence diagrams in comparison to the persistence diagram of the original sample show the birth-death of the pairing generators of the iterated inclusions. Black points represent the 0-dimensional homology groups H_0 that is the connected components of the complex. Red triangles represent the 1-dimensional homology group H_1 that is the 1-dimensional features, the cycles.

In order to quantitatively assess the quality of the encoding-decoding process, we use the bottleneck distance between the persistent diagram of \mathcal{X} and the persistent diagram of $G(Z)$ of the reconstructed data points. We recall the strength of the bottleneck distance is to measure quantitatively the topological changes in the data, either the true or the reconstructed data, while being insensitive to the scale of the analysis. Traditional distance measures fail to acknowledge this as they do not rely on persistent homology and, therefore, can only reflect a measurement of the nearness relations of the data points without considering the overall shape of the data distribution. In table 1, we notice the smallest bottleneck distance, and therefore, the best result, is obtained with PHom-WAE. It means PHom-WAE is capable to better preserve the topological features of the original data distribution than PHom-VAE including the nearness measurements and the overall shape.

Last but not least, persistent homology and the bottleneck distance are used to highlight the scarcity of the distribution P_Z of the latent manifold \mathcal{Z} . Us-

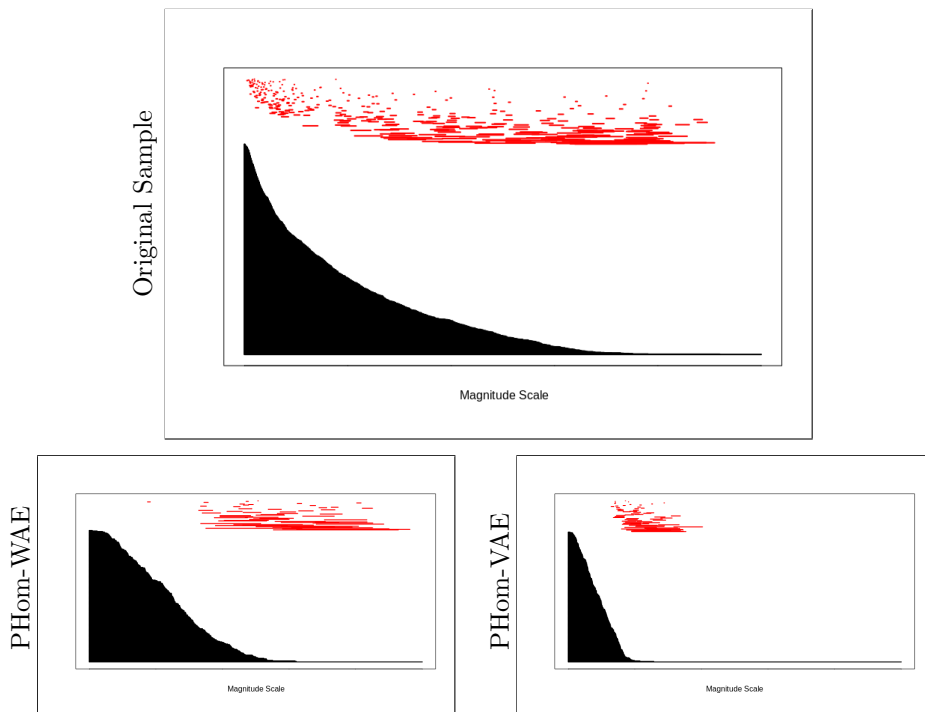


Fig. 6. PHom-WAE and PHom-VAE’s barcode diagrams in comparison to the barcode diagram of the original sample based on the persistence diagrams of figure 5. The barcodes diagrams are a simple way to represent the persistence diagrams. We refer to [10] for further details on their generation.

ing the bootstrapping technique of [35], we successively randomly select data samples Z_i contained in the manifold \mathcal{Z} . The total number of selected samples Z_i is at least 50% of the total number of points contained in the manifold \mathcal{Z} to ensure a reliable statistical representation. Assuming the data is not scattered, the bottleneck distance between the persistence diagrams of the samples Z_i is small. On the opposite, if the data is chaotically scattered in \mathcal{Z} , then the topological features between the samples Z_i are significantly different, and consequently, the bottleneck distance is large. In table 2, the bottleneck distance

Table 1. Bottleneck distance (smaller is better) for PHom-WAE and PHom-VAE between the samples X of the original manifold \mathcal{X} and the reconstructed manifold $G(Z|X)$ for $Z \in \mathcal{Z}$. Because of OT, PHom-WAE achieves better performance.

PHom-WAE	PHom-VAE	Difference (%)
0.0788	0.0878	10.25

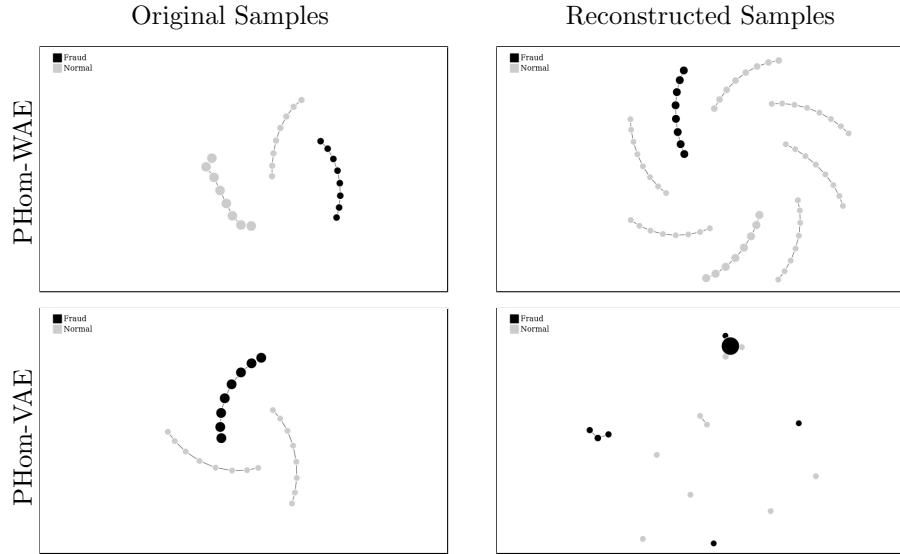


Fig. 7. Topological representation of the transaction samples, both fraudulent and normal, in black and gray respectively, of the reconstructed data distribution following $P_G(X|Z)$ for PHom-WAE and PHom-VAE. PHom-WAE does not alter the topological features of the manifold \mathcal{X} through the encoding-decoding process.

is significantly lower for PHom-WAE than for PHom-VAE. Therefore, the level of scattered chaos for PHom-WAE is lower than for PHom-VAE. It also means the distribution P_Z of the latent manifold \mathcal{Z} is better topologically preserved for PHom-WAE than for PHom-VAE. The reconstructed distribution $P_G(X|Z)$ of \mathcal{X} is, thus, less altered for PHom-WAE.

5 Conclusion

Building upon WAE and VAE, we introduce PHom-WAE and PHom-VAE, a new characterization of the manifolds of the WAE and VAE, respectively, that uses topology and persistence homology to highlight the manifold properties and the scattered points of the hidden space. We discussed their relations with other AE modeling techniques. Furthermore, relying on persistence homology, the bottleneck distance has been introduced to estimate quantitatively the alteration of the topological features occurring during the encoding-decoding process, a specificity that current traditional distance measures fail to acknowledge. We conducted experiments showing the performance of PHom-WAE in comparison to PHom-VAE using a challenging imbalanced real-world open data set containing credit card transactions, particularly suitable for the banking industry. We showed the superior performance of PHom-WAE in comparison to PHom-VAE. Future work will include further exploration of the topological features

Table 2. Bottleneck distance (smaller is better) for PHom-WAE and PHom-VAE between the samples Z_i of the latent manifold \mathcal{Z} following P_Z to detect scattered distribution. PHom-WAE better preserves the topological features during the encoding than PHom-VAE resulting in a manifold \mathcal{Z} less chaotically scattered, highlighted by the smaller bottleneck distance.

PHom-WAE	PHom-VAE	Difference (%)
0.0984	0.1372	28.28

such as the influence of the simplicial complex and the possibility to integrate a topological optimization function as a regularization term.

References

1. Golub, G.H.: Cf van loan. Matrix computation, Baltimore, MD, USA (1989)
2. Carroll, J.D., Chang, J.J.: Analysis of individual differences in multidimensional scaling via an n-way generalization of eckart-young decomposition. *Psychometrika* **35**(3) (1970)
3. Kolda, T.G., Bader, B.W.: Tensor decompositions and applications. *SIAM review* **51**(3) (2009)
4. Paatero, P.: A weighted non-negative least squares algorithm for three-way ‘PARAFAC’ factor analysis. *Chemometrics and Intelligent Laboratory Systems* **38**(2) (1997)
5. Bader, B.W., Harshman, R.A., Kolda, T.G.: Temporal analysis of semantic graphs using asalsan. In: *Data Mining, 2007. ICDM 2007. Seventh IEEE International Conference on. IEEE* (2007)
6. LeCun, Y., Boser, B., Denker, J.S., Henderson, D., Howard, R.E., Hubbard, W., Jackel, L.D.: Backpropagation applied to handwritten zip code recognition. *Neural computation* **1**(4), 541–551 (1989)
7. Hinton, G.E., Zemel, R.S.: Autoencoders, minimum description length and helmholtz free energy. In: *Advances in neural information processing systems*. pp. 3–10 (1994)
8. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* (2013)
9. Bengio, Y., Yao, L., Alain, G., Vincent, P.: Generalized denoising auto-encoders as generative models. In: *Advances in Neural Information Processing Systems*. pp. 899–907 (2013)
10. Chazal, F., Michel, B.: An introduction to Topological Data Analysis: fundamental and practical aspects for data scientists. *arXiv preprint arXiv:1710.04019* (2017)
11. Bousquet, O., Gelly, S., Tolstikhin, I., Simon-Gabriel, C.J., Schoelkopf, B.: From optimal transport to generative modeling: the vegan cookbook. *arXiv preprint arXiv:1705.07642* (2017)
12. Tolstikhin, I., Bousquet, O., Gelly, S., Schoelkopf, B.: Wasserstein auto-encoders. *arXiv preprint arXiv:1711.01558* (2017)
13. Goodfellow, I., Bengio, Y., Courville, A., Bengio, Y.: *Deep learning*, vol. 1. MIT press Cambridge (2016)

14. Frosini, P., Landi, C.: Size theory as a topological tool for computer vision. *Pattern Recognition and Image Analysis* **9**(4), 596–603 (1999)
15. Edelsbrunner, H., Letscher, D., Zomorodian, A.: Topological Persistence and Simplification. *Discrete & Computational Geometry* **28**(4), 511–533 (2002)
16. Robins, V., Abernethy, J., Rooney, N., Bradley, E.: Topology and intelligent data analysis. *Intelligent Data Analysis* **8**(5), 505–515 (2004)
17. Zomorodian, A., Carlsson, G.: Computing Persistent Homology. *Discrete & Computational Geometry* **33**(2), 249–274 (2005)
18. Ghrist, R.: Barcodes: the persistent topology of data. *Bulletin of the American Mathematical Society* **45**(1), 61–75 (2008)
19. Chazal, F., de Silva, V., Glisse, M., Oudot, S.: *The Structure and Stability of Persistence Modules*. SpringerBriefs in Mathematics, Springer International Publishing (2016)
20. Singh, G., Mémoli, F., Carlsson, G.E.: Topological methods for the analysis of high dimensional data sets and 3d object recognition. In: SPBG. pp. 91–100 (2007)
21. Zomorodian, A.: Fast construction of the Vietoris-Rips complex. *Computers & Graphics* **34**(3), 263–271 (2010)
22. Nicolau, M., Levine, A.J., Carlsson, G.: Topology based data analysis identifies a subgroup of breast cancers with a unique mutational profile and excellent survival. *Proceedings of the National Academy of Sciences* p. 201102826 (2011)
23. Lum, P.Y., Singh, G., Lehman, A., Ishkanov, T., Vejdemo-Johansson, M., Alagappan, M., Carlsson, J., Carlsson, G.: Extracting insights from the shape of complex data using topology. *Scientific reports* **3**, srep01236 (2013)
24. Bendich, P., Marron, J.S., Miller, E., Pieloch, A., Skwerer, S.: Persistent homology analysis of brain artery trees. *The annals of applied statistics* **10**(1), 198 (2016)
25. Xia, K., Wei, G.W.: Persistent homology analysis of protein structure, flexibility, and folding. *International journal for numerical methods in biomedical engineering* **30**(8), 814–844 (2014)
26. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: *Advances in neural information processing systems*. pp. 2672–2680 (2014)
27. Villani, C.: *Topics in optimal transportation*. No. 58, American Mathematical Soc. (2003)
28. Arjovsky, M., Chintala, S., Bottou, L.: Wasserstein generative adversarial networks. In: *International Conference on Machine Learning*. pp. 214–223 (2017)
29. Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., Courville, A.C.: Improved training of Wasserstein GANs. In: *Advances in Neural Information Processing Systems*. pp. 5767–5777 (2017)
30. Chizat, L., Peyré, G., Schmitzer, B., Vialard, F.X.: Unbalanced optimal transport: geometry and Kantorovich formulation. *arXiv preprint arXiv:1508.05216* (2015)
31. Liero, M., Mielke, A., Savaré, G.: Optimal entropy-transport problems and a new Hellinger–Kantorovich distance between positive measures. *Inventiones mathematicae* **211**(3), 969–1117 (2018)
32. Hatcher, A.: *Algebraic Topology*. Cambridge University Press (2002)
33. Oudot, S.Y.: *Persistence Theory: From Quiver Representations to Data Analysis*. No. 209 in *Mathematical Surveys and Monographs*, American Mathematical Society (2015)
34. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014)
35. Friedman, J., Hastie, T., Tibshirani, R.: *The elements of statistical learning*, vol. 1. Springer series in statistics New York, NY, USA: (2001)