
Geometry Score: A Method For Comparing Generative Adversarial Networks

Valentin Khruikov¹ Ivan Oseledets^{1 2}

Abstract

One of the biggest challenges in the research of generative adversarial networks (GANs) is assessing the quality of generated samples and detecting various levels of mode collapse. In this work, we construct a novel measure of performance of a GAN by comparing geometrical properties of the underlying data manifold and the generated one, which provides both qualitative and quantitative means for evaluation. Our algorithm can be applied to datasets of an arbitrary nature and is not limited to visual data. We test the obtained metric on various real-life models and datasets and demonstrate that our method provides new insights into properties of GANs.

1. Introduction

Generative adversarial networks (GANs) (Goodfellow et al., 2014) are a class of methods for training generative models, which have been recently shown to be very successful in producing image samples of excellent quality. They have been applied in numerous areas (Radford et al., 2015; Salimans et al., 2016; Ho & Ermon, 2016). Briefly, this framework can be described as follows. We attempt to mimic a given target distribution $p_{\text{data}}(\mathbf{x})$ by constructing two networks $G(\mathbf{z}; \theta^{(G)})$ and $D(\mathbf{x}; \theta^{(D)})$ called the generator and the discriminator. The generator learns to sample from the target distribution by transforming a random input vector \mathbf{z} to a vector $\mathbf{x} = G(\mathbf{z}; \theta^{(G)})$, and the discriminator learns to distinguish the model distribution $p_{\text{model}}(\mathbf{x})$ from $p_{\text{data}}(\mathbf{x})$. The training procedure for GANs is typically based on applying gradient descent in turn to the discriminator and the generator in order to minimize a loss function. Finding a good loss function is a topic of ongoing research, and several options were proposed in (Mao et al., 2016; Arjovsky et al., 2017).

One of the main challenges (Lucic et al., 2017; Barratt &

Sharma, 2018) in the GANs framework is estimating the quality of the generated samples. In traditional GAN models, the discriminator loss cannot be used as a metric and does not necessarily decrease during training. In more involved architectures such as WGAN (Arjovsky et al., 2017) the discriminator (critic) loss is argued to be in correlation with the image quality, however, using this loss as a measure of quality is nontrivial. Training GANs is known to be difficult in general and presents such issues as *mode collapse* when $p_{\text{model}}(\mathbf{x})$ fails to capture a multimodal nature of $p_{\text{data}}(\mathbf{x})$ and in extreme cases all the generated samples might be identical. Several techniques to improve the training procedure were proposed in (Salimans et al., 2016; Gulrajani et al., 2017).

In this work, we attack the problem of estimating the quality and diversity of the generated images by using the machinery of topology. The well-known Manifold Hypothesis (Goodfellow et al., 2016) states that in many cases such as the case of natural images the support of the distribution $p_{\text{data}}(\mathbf{x})$ is concentrated on a low dimensional manifold $\mathcal{M}_{\text{data}}$ in a Euclidean space. This manifold is assumed to have a very complex non-linear structure and is hard to define explicitly. It can be argued that interesting features and patterns of the images from $p_{\text{data}}(\mathbf{x})$ can be analyzed in terms of topological properties of $\mathcal{M}_{\text{data}}$, namely in terms of loops and higher dimensional *holes* in $\mathcal{M}_{\text{data}}$. Similarly, we can assume that $p_{\text{model}}(\mathbf{x})$ is supported on a manifold $\mathcal{M}_{\text{model}}$ (under mild conditions on the architecture of the generator this statement can be made precise (Shao et al., 2017)), and for sufficiently good GANs this manifold can be argued to be quite similar to $\mathcal{M}_{\text{data}}$ (see Fig. 1). This intuitive claim will be later supported by numerical experiments. Based on this hypothesis we develop an approach which allows for comparing the topology of the underlying manifolds for two point clouds in a stochastic manner providing us with a visual way to detect mode collapse and a score which allows for comparing the quality of various trained models. Informally, since the task of computing the precise topological properties of the underlying manifolds based only on samples is ill-posed by nature, we estimate them using a certain probability distribution (see Section 4).

We test our approach on several real-life datasets and popular GAN models (DCGAN, WGAN, WGAN-GP) and show that the obtained results agree well with the intuition and allow for comparison of various models (see Section 5).

¹Skolkovo Institute of Science and Technology ²Institute of Numerical Mathematics RAS. Correspondence to: Valentin Khruikov <khruikov.v@gmail.com>.




Figure 1. The Manifold Hypothesis suggests that in the case of natural images the data is supported on a low dimensional *data manifold* $\mathcal{M}_{\text{data}}$. Similarly, GANs sample images from an immersed manifold $\mathcal{M}_{\text{model}}$. By comparing topological properties of the manifolds $\mathcal{M}_{\text{data}}$ and $\mathcal{M}_{\text{model}}$ we can get insight in how strongly GAN captured intricacies in the data distribution $p_{\text{data}}(\mathbf{x})$, and quantitatively estimate the difference.

2. Main idea

Let us briefly discuss our approach before dwelling into technical details. As described in the introduction we would like to compare topological properties of $\mathcal{M}_{\text{data}}$ and $\mathcal{M}_{\text{model}}$ in some way. This task is complicated by the fact that we do not have access to the manifolds themselves but merely to samples from them. A natural approach in this case is to approximate these manifolds using some simpler spaces in such a way that topological properties of these spaces resemble those of $\mathcal{M}_{\text{data}}$ and $\mathcal{M}_{\text{model}}$. The main example of




Figure 2. Simplicial complex. Topological space X is constructed from several edges ($\sigma_1, \sigma_3, \sigma_4$) and a two dimensional face σ_2 .

such spaces are **simplicial complexes** (Fig. 2), which are build from intervals, triangles and other higher dimensional simplices. In order to reconstruct the underlying manifold using a simplicial complex several methods exist. In all such approaches proximity information of the data is used, such as pairwise distances between samples. Typically one chooses some threshold parameter ε and based on the value of this parameter one decides which simplices are added into the approximation (see Fig. 3). However a single value ε is not enough — for very small values the reconstructed space will be just a disjoint union of points and for very




Figure 3. A simplicial complex constructed on a sample X . First, we fix the proximity parameter ε . Then we take balls of the radius ε centered at each point, and if for some subset of X of size $k + 1$ all the pairwise intersections of the corresponding balls are non-empty, we add the k -dimensional simplex spanning this subset to the simplicial complex \mathcal{R}_ε .




Figure 4. Using different values of the proximity parameter ε we obtain different simplicial complexes (a). For $\varepsilon = \varepsilon_1$ the balls do not intersect and there are just 10 isolated components (b, [left]). For $\varepsilon = \varepsilon_2$ several components have merged and one *loop* appeared (b, [middle]). The filled triangle corresponding to the triple pairwise intersection is topologically trivial and does not affect the topology (and similarly darker tetrahedron on the right). For $\varepsilon = \varepsilon_3$ all the components merged into one and the same hole still exists (b, [right]). In the interval $[\varepsilon_2, \varepsilon_3]$ one smaller hole as on Fig. 3 appeared and quickly disappeared. This information can be conveniently summarized in the *persistence barcode* (c). The number of connected components (holes) in the simplicial complex for some value ε_0 is given by the number of intervals in H_0 (H_1) intersecting the vertical line $\varepsilon = \varepsilon_0$.

large ε it will be a single connected blob, while the correct approximation is somewhere in between. This issue is resolved by considering a *family* (Fig. 4, a) of simplicial complexes, parametrized by the (‘persistence’) parameter ε . It is also convenient to refer to the parameter ε as *time*, with the idea that we gradually throw more simplices into our simplicial complex as time goes by. For each value of ε we can compute topological properties of the corresponding simplicial complex, namely *homology* which encodes the number of holes of various dimensions in a space. Controlling the value of ε allows us to decide holes of which

size are meaningful and should not be discarded as a noise. For simplicial complex presented on Fig. 3 there are two one-dimensional holes, and for slightly bigger value of ε the lower hole disappeared (Fig. 4, b), while the top one remained intact, which suggests that the top hole is more important topological feature. Information about how homology is changing with respect to ε can be conveniently encoded in the so-called *persistence barcodes* (Ghrist, 2008; Zomorodian & Carlsson, 2005). An example of such barcode is given on (Fig. 4, c). In general, to find the rank of k -homology (delivering the number of k -dimensional holes) at some fixed value ε_0 one has to count intersections of the vertical line $\varepsilon = \varepsilon_0$ with the intervals at the desired block H_k .

These barcodes provide a way to compare topological properties of the underlying manifolds. In principle, we could obtain a metric of similarity of two datasets by comparing the barcodes of the simplicial complexes constructed based on each dataset (as described on Fig. 3), but there are disadvantages of this approach, such as a huge number of simplices for large datasets. Moreover, in order to extract interesting topological properties from such large simplicial complexes various tricks are required (Ghrist, 2008). To remedy these issues we can note that we are in fact interested in *topological* approximations rather than *geometrical*. The difference is that to obtain a correct estimate of the topological properties much smaller number of simplices is often sufficient, e.g., for any number of points sampled from a circle the correct answer could be obtained by taking just three points (thus obtaining a triangle which is topologically equivalent to the circle). Based on these ideas the so-called **witness complex** is introduced (De Silva & Carlsson, 2004), which provides a topological approximation with a small number of simplices. In order to achieve this a small subset of *landmark* points is chosen and a simplicial complex is constructed using these points as vertices (while also taking into account the proximity information about all the remaining points called *witnesses*).

To construct a numerical measure which could be compared across datasets we would like to estimate the correct values of homology. Comparing the computed barcodes is a challenging task since they are non-trivial mathematical objects (though some metrics exist they are hard to compute). We take the simpler route and to extract meaningful topological data from the barcode we propose computing **Relative Living Times** (RLT) of each number of holes that was observed. They are defined as the ratio of the total time when this number was present and of the value ε_{\max} when points connect into a single blob. These relative living times could be interpreted as a *confidence* in our approximation — if say for 50% of all period of topological activity we have observed that there is at least 1 one-dimensional hole (as on Fig. 4), then it is probably an accurate estimation of




Figure 5. Estimation of the topology of a dataset sampled from the 32-dimensional hyperplane in 784-dimensional space. With high confidence, we can say that there are no 1-dimensional holes. For details see Section 4.

topology of the underlying space.

Choosing the correct landmarks is a nontrivial task. We follow the discussion in (De Silva & Carlsson, 2004) which advises doing it randomly. To account for this randomness, we compute the RLT stochastically by repeating the experiment a large number of times. By averaging the obtained RLT we compute the **Mean Relative Living Times** (MRLT). By construction, they add up to 1 and employing Bayesian point of view we can interpret them as a probability distribution reflecting our confidence about the correct number of holes on *average*. An example of such distribution is given on Fig. 5, where we run our method for a simple planar dataset (in a high dimensional space). To quantitatively evaluate the topological difference between two datasets we propose computing the L_2 -error between these distributions. Note that in practice (when activation functions such as ReLU are used) the resulting space $\mathcal{M}_{\text{model}}$ may fail to be a manifold in precise mathematical sense, however, the analysis is still applicable since it deals with arbitrary topological spaces. Now let us introduce all the technical details.

3. Homology to the rescue

In this section we briefly discuss the important concepts of simplicial complexes and homology. For thorough introduction we refer the reader to the classical texts such as (Hatcher, 2002; May, 1999).

Simplicial complexes Simplicial complex is a classical concept widely used in topology. Formally it is defined as follows.

Definition 1. A simplicial complex \mathcal{S} (more precisely an *abstract simplicial complex*) is specified by the following data:

- The vertex set $Z = \{z_1, z_2, \dots, z_n\}$
- A collection of simplices Σ , where p -dimensional simplex σ_p is defined just as a $p + 1$ element subset of Z :

$$\sigma_p = \{z_{i_1}, z_{i_2}, \dots, z_{i_{p+1}}\}$$

- We require that the collection Σ is closed under taking

faces, that is for each p -dimensional simplex σ_p all the $(p-1)$ -dimensional simplices obtained by deleting one of the vertices z_{i_1}, \dots, z_{i_p} are also elements of Σ .

An example of a simplicial complex \mathcal{S} is presented on Fig. 2. It contains 5 vertices $\{z_1, z_2, \dots, z_5\}$ and several edges and faces: two-dimensional face σ_2 and one-dimensional edges $\sigma_1, \sigma_3, \sigma_4$. Note that these are *maximal* simplices, since by the third property all the edges of σ_2 are also elements of \mathcal{S} . Important topological properties of \mathcal{S} (such as connectedness, existence of one-dimensional loop) do not depend on in which Euclidean space \mathcal{S} is embedded or on precise positions of vertices, but merely on the combinatorial data — the number of points and which vertices together span a simplex.

As was described in Section 2 given a dataset X sampled from a manifold \mathcal{M} we would like to compute a family of simplicial complexes topologically approximating \mathcal{M} on various scales, namely witness complexes. This family is defined as follows. First we choose some subset $L \subset X$ of points called landmarks (whereas points in X are called witnesses) and some distance function $d(x, x')$, e.g., the ordinary Euclidean distance. There is not much theory about how to choose the best landmarks, but several strategies were proposed in (De Silva & Carlsson, 2004). The first one is to choose landmarks sequentially by solving a certain minimax problem, and the second one is to just pick landmarks at random (by uniformly selecting a fixed number of points from X). We follow the second approach since the minimax strategy is known to have some flaws such as the tendency to pick up outliers. The selected landmarks will serve as the vertices of the simplicial complex and witnesses will help to decide on which simplices are inserted via a predicate “is witnessed”:

$$\sigma \subset L \text{ is witnessed by } w \in X \text{ if } \forall l \in \sigma, \forall l' \in L \setminus \sigma \quad (1)$$

$$d(w, l)^2 \leq d(w, l')^2 + \alpha,$$

with α being a relaxation parameter which provides us with a sequence of simplicial complexes. The maximal value of α for the analysis is typically chosen to be proportional to the maximal pairwise distance between points in L . Witness complexes even for small values of α are good topological approximations to \mathcal{M} . The main advantage of a witness complex is that it allows constructing a reliable approximation using a relatively small number of simplices and makes the problem tractable even for large datasets. Even though it is known that in some cases it may fail to recover the correct topology (Boissonnat et al., 2009), it still can be used to compare topological properties of datasets, and if any better method is devised, we can easily replace the witness complex by this new more reliable simplicial complex.

Homology The precise definition of the homology is technical, and we have to omit it due to the limited space. We

refer the reader to [Chapter 2] (Hatcher, 2002) for a thorough discussion. The most important properties of homology can be summarized as follows. For any topological space X the so-called i^{th} homology groups H_i are introduced. The actual number of i -dimensional holes in X is given the *rank* of H_i , the concept which is quite similar to the dimension of a vector space. These ranks are called the *Betti numbers* β_i and serve as a coarse numerical measure of homology.

Homology is known to be one of the most easily computable topological invariants. In the case of X being a simplicial complex H_i can be computed by pretty much linear algebra, namely by analyzing kernels and images of certain linear maps. Dimensions of matrices appearing in this task are equal to the numbers d_k of simplices of specific dimension k in X , e.g. in the case of Fig. 2 we have $d_0 = 5, d_1 = 6, d_2 = 1$ and matrices will be of sizes 6×1 and 5×6 . Existing algorithms (Kaczynski et al., 2006) can handle extremely large simplicial complexes (with millions of simplices) and are available in numerous software packages. An important property of homology is that k^{th} homology depends only on simplices of dimension at most $k+1$, which significantly speeds up computations.

Persistent homology In Section 2 we discussed that to find a proxy of the correct topology of \mathcal{M} it is insufficient to use single simplicial complex but rather a family of simplicial complexes is required. As we transition from one simplicial complex to another, some holes may appear, and some disappear. To distinguish between which are essential and which should be considered noise the concept of persistence was introduced (Edelsbrunner et al., 2000; Zomorodian & Carlsson, 2005). The formal Structure Theorem (Zomorodian & Carlsson, 2005) states that for each generator of homology (“hole” in our notation) one could provide the time of its “birth” and “death”. This data is pictorially represented as (Fig. 4, [bottom]), with the horizontal axis representing the parameter and the vertical axis representing various homology generators. To perform the computation of these barcodes, an efficient algorithm was proposed in (Zomorodian & Carlsson, 2005). As an input to this algorithm one has to supply a sequence of $(\sigma_i, \varepsilon_i)$, with σ_i being a simplex and ε_i being its time of appearance in a family. This algorithm is implemented in several software packages such as Dionysus and GUDHI (Maria et al., 2014), but the witness complex is supported only in the latter.

4. Algorithm

Let us now explain how we apply these concepts to construct a metric to compare the topological properties of two datasets. First let us define the key part of the algorithm — the relative living times (RLT) of homology. Suppose that for a dataset X and some choice of landmarks L we have ob-

tained a persistence barcode with the persistence parameter α spanning the range $[0, \alpha_{\max}]$. Let us fix the dimension k in which we study the homology, and let $\mathcal{I}_k = \{[b_i, d_i]\}_{i=1}^n$ be the collection of persistence intervals in this dimension. Then in order to find the k^{th} Betti number for a fixed value α one has to count the number of persistence intervals containing α , and we obtain the integer valued function

$$\beta_k(\alpha) \triangleq |\{[b_i, d_i] \in \mathcal{I}_k : \alpha \in [b_i, d_i]\}|. \quad (2)$$

Then the RLT are defined as follows (for non-negative integers i):

$$\text{RLT}(i, k, X, L) \triangleq \frac{\mu(\{\alpha \in [0, \alpha_{\max}] : \beta_k(\alpha) = i\})}{\alpha_{\max}}, \quad (3)$$

that it is for each possible value of $\beta_k(\alpha)$ we find how long it existed *relatively* to the whole period of topological activity. Note that in our analysis we use witness complexes which depend on the choice of landmarks, which is random. Thus it is reasonable to consider the distribution of $\text{RLT}(i, k, X, L)$ on the set of landmarks (tuples of points), in other words, we repeatedly sample the landmarks and compute the RLT of the obtained persistence barcode. After sufficiently many experiments we can approximate the *Mean Relative Living Times* (MRLT):

$$\text{MRLT}(i, k, X) \triangleq \mathbb{E}_L[\text{RLT}(i, k, X, L)]. \quad (4)$$

We hypothesize that these quantities provide us with a good way to compare the topological properties of datasets, as they serve as measures of confidence in the estimation of the topology of the underlying manifolds. From Eq. (3) it follows that

$$\sum_i \text{MRLT}(i, k, X) = 1,$$

which suggest that for a fixed value of k we could interpret $\text{MRLT}(i, k, X)$ as a *probability distribution* (over integers). This distribution defines our certainty about the number of k -dimensional holes in the underlying manifold of X on *average*. In this work we consider the case $k = 1$, i.e. we study the first homology of datasets. We motivate this by drawing an analogy with the Taylor series: we can get a good understanding of behavior of a function by looking at the first term of the series (see also (Ghrist, 2008) for discussion). Based on the probabilistic understanding given two datasets X_1 and X_2 we define a measure of their topological similarity (*Geometry Score*) in the following way:

$$\text{GeomScore}(X_1, X_2) \triangleq \sum_{i=0}^{i_{\max}-1} (\text{MRLT}(i, 1, X_1) - \text{MRLT}(i, 1, X_2))^2, \quad (5)$$

with i_{\max} being an upper bound on $\beta_1(\alpha)$ for X_1 and X_2 (for typical datasets we found that $i_{\max} = 100$ suffices).

To construct the witness complex given the sets of landmarks L and witnesses X one has to provide the matrix of pairwise distances between L and X and the maximal value of persistence parameter α (see Eq. (1)). In our experiments, we have chosen α_{\max} to be proportional to the maximal pairwise distance between points in L with some coefficient γ . Since we only compute $\beta_1(\alpha)$ the simplices of dimension at most 2 are needed. In principle to compare two datasets any value of γ suffices, however in our experiments we found that to get a reasonable distribution for datasets of size ~ 5000 the value $\frac{1}{128}$ yields good results (for large γ a lot of time is spend in the regime of a single connected blob which shifts the distributions towards 0). We summarize our approach in Algorithm 1 and Algorithm 2. We also suggest that to obtain accurate results datasets of the same size should be used for comparison

Algorithm 1 The algorithm to compute RLT of a dataset. See Section 4 for details. Suggested default values of the parameters for a dataset $X \in \mathbb{R}^{N \times D}$ are $L_0 = 64$, $\gamma = \frac{1}{128} / \frac{N}{5000}$, $i_{\max} = 100$, $n = 10000$.

Require: X : $2D$ array representing the dataset

Require: L_0 : Number of landmarks to use

Require: γ : Coefficient determining α_{\max}

Require: i_{\max} : Upper bound on i in $\text{RLT}(i, 1, X, L)$

Require: n : Number of experiments

Require: $\text{dist}(A, B)$: Function computing the matrix of pairwise (Euclidean) distances between samples from A and B

Require: $\text{witness}(d, \alpha, k)$: Function computing the family of witness complexes using the matrix of pairwise distances d , maximal value of persistence parameter α and maximal dimension of simplices k

Require: $\text{persistence}(w, k)$: Function computing the persistence intervals of a family w in dimension k

Returns: An array of size $n \times i_{\max}$ of the obtained RLT for each experiment

Initialize: $\text{rlt} = \text{zeros}(n, i_{\max})$

for $i = 0$ **to** $n - 1$ **do**

$L^{(i)} \leftarrow \text{random_choice}(X, \text{size}=L_0)$

$d^{(i)} \leftarrow \text{dist}(L^{(i)}, X)$

$\alpha_{\max}^{(i)} \leftarrow \gamma \cdot \max(\text{dist}(L^{(i)}, L^{(i)}))$

$W^{(i)} \leftarrow \text{witness}(d^{(i)}, \alpha_{\max}^{(i)}, 2)$

$\mathcal{I}^{(i)} \leftarrow \text{persistence}(W^{(i)}, 1)$

for $j = 0$ **to** $i_{\max} - 1$ **do**

Compute $\text{RLT}(j, 1, X, L^{(i)})$ using Eqs. (2) and (3)

$\text{rlt}[i, j] \leftarrow \text{RLT}(j, 1, X, L^{(i)})$

end for

end for

Complexity Let us briefly discuss the complexity of each step in the main loop of Algorithm 1. Suppose that we have

Algorithm 2 *Geometry Score*, the proposed algorithm to compute topological similarity between datasets

Require: X_1, X_2 : arrays representing the datasets
Returns: s : a number representing the topological similarity of X_1 and X_2
Initialize: $s = 0$
 For X_1 and X_2 run Algorithm 1 with the same collection of parameters, obtaining arrays rlt_1 and rlt_2
 $\text{mrlt}_1 \leftarrow \text{mean}(\text{rlt}_1, \text{axis}=0)$
 $\text{mrlt}_2 \leftarrow \text{mean}(\text{rlt}_2, \text{axis}=0)$
 $s \leftarrow \text{sum}((\text{mrlt}_1 - \text{mrlt}_2)^2)$

a dataset $X \in \mathbb{R}^{N \times D}$. Computing the matrix of pairwise distances between all points in the dataset and the landmarks points requires $O(NDL_0)$ operations. The complexity of the next piece involving computing the persistence barcode is hard to estimate, however we can note that it does not depend on the dimensionality D of the data. In practice this computation is done faster than computing the matrix in the previous step (for datasets of significant dimensionality). All the remaining pieces of the algorithm take negligible amount of time. This linear scaling of the complexity w.r.t dimensionality of the data allows us to apply our method even for high-dimensional datasets. On a typical laptop (3.1 GHz Intel Core i5 processor) one iteration of the inner loop of Algorithm 1 for one class of the MNIST dataset takes approximately 900 ms.

5. Experiments

Experimental setup We have implemented Algorithms 1 and 2 in Python using GUDHI¹ for computing witness complexes and persistence barcodes. Our code is available on Github². Default values of parameters in Algorithm 1 were used for experiments unless otherwise specified. We test our method on several datasets and GAN models:

- **Synthetic data** — on synthetic datasets we demonstrate that our method allows for distinguishing the datasets based on their topological properties.
- **MNIST** — as the next experiment we test our approach on the MNIST dataset of handwritten digits. We compare two recently proposed models: WGAN (Arjovsky et al., 2017) and WGAN-GP (Gulrajani et al., 2017) in order to verify if the improved model WGAN-GP indeed produces better images.
- **CelebA** — to demonstrate that our method can be applied to datasets of large dimensionality we analyze the CelebA dataset (Liu et al., 2015) and check if we

can detect mode collapse in a GAN using MRLT.

- **CaloGAN** — as the final experiment we apply our algorithm to a dataset of a non-visual origin and evaluate the specific generative model CaloGAN (Paganini et al., 2017).

Synthetic data For this experiment we have generated a collection of simple 2D datasets $\{X_j\}_{j=1}^5$ (see Fig. 6) each containing 5000 points. As a test problem we would like to evaluate which of the datasets $\{X_j\}_{j=2}^5$ is the best approximation to the ground truth X_1 . For each of $\{X_j\}_{j=1}^5$ we ran Algorithm 1 using $L_0 = 32$, $n = 2000$, $i_{\max} = 3$, $\gamma = \frac{1}{8}$ and compute MRLT using Eq. (4). The resulting distributions are visualized on Fig. 6, [bottom]. We observe that we can correctly identify the number of 1-dimensional holes in each space X_j using the MAP estimate

$$\beta_1^*(X_j) = \arg \max_i \text{MRLT}(i, 1, X_j). \quad (6)$$

It is clear that X_4 is the most similar dataset to X_1 , which is supported by the fact that their MRLT are almost identical. Note that on such simple datasets we were able to recover the correct homology with almost 100% confidence and this will not be the case for more complicated manifolds in the next experiment.

MNIST In this experiment we compare topological properties of the MNIST dataset and samples generated by the WGAN and WGAN-GP models trained on MNIST. It was claimed that the WGAN-GP model produces better images and we would like to verify if we can detect it using topology. For the GAN implementations we used the code³ provided by the authors of (Gulrajani et al., 2017). We have trained each model for 25 epochs and generated 60000 samples. To compare topology of each class individually we trained a CNN classifier on MNIST (with 99.5% test accuracy) and split generated dataset into classes (containing roughly 6000 images each). For every class and each of the 3 corresponding datasets ('base', 'wgan', 'wgan-gp') we run Algorithm 1 and compute MRLT with $\gamma = \frac{1}{128}$. Similarly we evaluate MRLT for the entire datasets without splitting them into classes using $\gamma = \frac{1}{1000}$. The obtained MRLT are presented on Fig. 7 and the corresponding *Geometry Scores* for each model are given in Table 1. We observe that both models produce distributions which are very close to the ground truth, but for almost all classes WGAN-GP shows better scores. We can also note that for the entire datasets (Fig. 7, [right]) the predicted values of homology does not seem to be much bigger than for each individual digit. One possible explanation is that some samples (like say of class '7') fill the holes in the underlying manifolds of other classes (like class '1' in this case) since they look quite similar.

¹<http://gudhi.gforge.inria.fr/>

²<https://github.com/KhrulkovV/geometry-score>

³https://github.com/igul222/improved_wgan_training




Figure 6. Mean Relative Living Times (MRLT) for various 2D datasets. The number of one-dimensional holes is correctly identified in all the cases. By comparing MRLT we find that the second dataset from the left is the most similar to the ‘ground truth’ (noisy circle on the left).

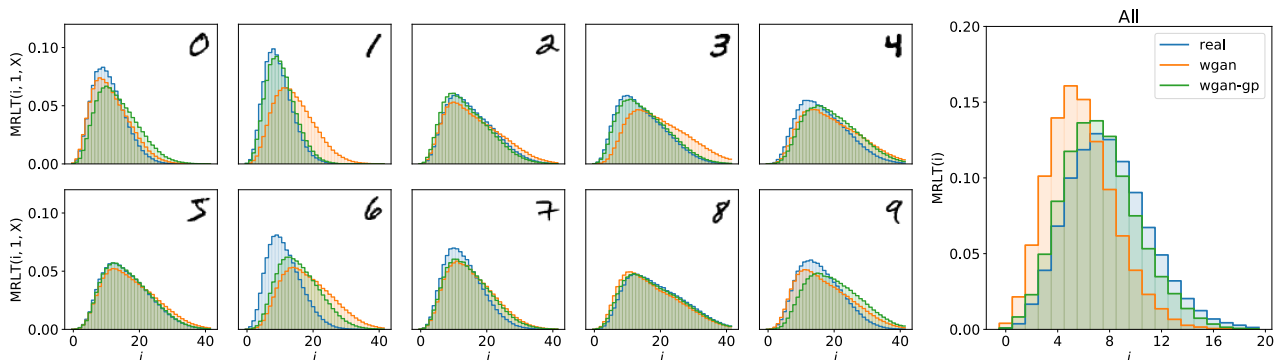


Figure 7. Comparison of MRLT of the MNIST dataset and of samples generated by WGAN and WGAN-GP trained on MNIST. MRLT match almost perfectly, however, WGAN-GP shows slightly better performance on most of the classes.

Table 1. *Geometry Scores* $\times 10^3$ of WGAN and WGAN-GP trained on the MNIST dataset (see also Fig. 7). Each class contained roughly 6000 images, except for ‘All’ which corresponds to the total datasets of 60000 images.

LABEL	0	1	2	3	4	5	6	7	8	9	ALL
WGAN	0.85	21.4	0.60	7.04	1.52	0.47	22.8	2.20	0.76	1.27	26.1
WGAN-GP	5.19	1.44	0.54	0.27	2.16	0.03	13.5	1.38	0.14	5.00	2.04

CelebA We now analyze the popular CelebA dataset consisting of photos of various celebrities. In this experiment we would like to study if we can detect mode collapse using our method. To achieve this we train two GAN models — a good model with the generator having high capacity and a second model with the generator much weaker than the discriminator. In this experiment we utilize the DCGAN model and use the implementation provided⁴ by the authors (Radford et al., 2015). For the first model (‘dcgan’) we use

⁴<https://github.com/carpedm20/DCGAN-tensorflow>

the default settings, and for the second (‘bad-dcgan’) we set the latent dimension to 8 and reduce the size of the fully connected layer in the generator to 128 and number of filters in convolutional layers to 4. Images in the dataset are of size 108×108 and to obtain faces we perform the central crop which reduces the size to 64×64 . We trained both models for 25 epochs and produced 10000 images for our analysis. Similarly, we randomly picked 10000 (cropped) images from the original dataset. We report the obtained results on Fig. 8. MRLT obtained using the good model matches the ground truth almost perfectly and *Geometry Score* of

the generated dataset is equal to 14×10^{-3} , confirming the good visual quality of the samples (Radford et al., 2015). MRLT obtained using the weak model are maximized for $i = 0$, which suggests that the samples are either identical or present very little topological diversity (compare with Fig. 5), which we confirmed visually. On Fig. 8, [right] we report the behavior of the *Geometry Score* and *Inception Score* (Salimans et al., 2016) w.r.t the iteration number. The *Inception Score* introduced uses the pretrained Inception network (Szegedy et al., 2015) and is defined as

$$I(\{x_n\}_{n=1}^N) \triangleq \exp \mathbb{E}_{\mathbf{x}}(D_{\text{KL}}(p(y|\mathbf{x})||p(y))),$$

where $p(y|\mathbf{x})$ is approximated by the Inception network and $p(y)$ is computed as $p(y) = \frac{1}{N} \sum_i p(y|\mathbf{x}_i)$. Note that the *Geometry Score* of the better model rapidly decreases and of the mode collapsed model stagnates at high values. Such behavior could not be observed in the *Inception Score*.




Figure 8. MRLT of the (cropped) CelebA dataset and samples generated using DCGAN and DCGAN with forced mode collapse. Plots on the right present the behavior of the *Geometry Score* and *Inception Score* for these two models during the training. Mode collapse in ‘bad-dcgan’ is easily observable using the *Geometry Score*.

CaloGAN In this experiment, we will apply our technique to the dataset appearing in the experimental particle physics. This dataset⁵ represents a collection of a calorimeter (an experimental apparatus measuring the energy of particles) responses, and it was used to create a generative model (Paganini et al., 2017) in order to help physicists working at the LHC. Evaluating the obtained model⁶ is a non-trivial task and was performed by comparing physical properties of the obtained and the real data. Since our method is not limited to visual datasets we can apply it in order to confirm the quality of this model. For the analysis we used ‘eplus’ dataset which is split into 3 parts (‘layer 0’, ‘layer 1’, ‘layer 2’) containing matrices of sizes 3×96 , 12×12 , 12×6 correspondingly. We train the CaloGAN model with defaults

settings for 50 epochs and generate 10000 samples (each sample combines data for all 3 layers). We then randomly pick 10000 samples from the original dataset and compare MRLT of the data and generated samples for each layer. Results are presented on Fig. 9. It appears that topological properties of this dataset are rather trivial, however, they are correctly identified by CaloGAN. Slight dissimilarities between the distributions could be connected to the fact that the physical properties of the generated samples do not exactly match those of the real ones, as was analyzed by the authors of (Paganini et al., 2017).




Figure 9. MRLT of the dataset used in experimental particle physics and of the samples generated using the corresponding CaloGAN model.

6. Related work and discussion

Several performance measures have been introduced to assess the performance of GANs used for natural images. *Inception Score* (Salimans et al., 2016) uses the outputs of the pretrained Inception network, and a modification called *Fréchet Inception Distance (FID)* (Heusel et al., 2017) also takes into account second order information of the final layer of this model. Contrary to these methods, our approach does not use auxiliary networks and is not limited to visual data. We note, however, that since we only take topological properties into account (which do not change if we say shift the entire dataset by 1) assessing the *visual* quality of samples may be difficult based only on our algorithm, thus in the case of natural images we propose to use our method in conjunction with other metrics such as *FID*. We also hypothesize that in the case of the large dimensionality of data *Geometry Score* of the features extracted using some network will adequately assess the performance of a GAN.

7. Conclusion

We have introduced a new algorithm for evaluating a generative model. We show that the topology of the underlying manifold of generated samples may be different from the topology of the original data manifold, which provides insight into properties of GANs and can be used for hyperparameter tuning. We do not claim however that the obtained metric correlates with the visual quality as estimated by humans and leave the analysis to future work. We hope that our research will be useful to further theoretical understanding of GANs.

⁵<https://data.mendeley.com/datasets/pvn3xc3wy5/1>

⁶<https://github.com/hep-lbdl/CaloGAN>

Acknowledgements

We would like to thank the anonymous reviewers for their valuable comments. We also thank Maxim Rakhuba for productive discussions and making our illustrations better. This study was supported by the Ministry of Education and Science of the Russian Federation (grant 14.756.31.0001).

References

- Arjovsky, M., Chintala, S., and Bottou, L. Wasserstein GAN. *arXiv preprint arXiv:1701.07875*, 2017.
- Barratt, S. and Sharma, R. A note on the Inception Score. *arXiv preprint arXiv:1801.01973*, 2018.
- Boissonnat, J.-D., Guibas, L. J., and Oudot, S. Y. Manifold reconstruction in arbitrary dimensions using witness complexes. *Discrete & Computational Geometry*, 42(1): 37–70, 2009.
- De Silva, V. and Carlsson, G. E. Topological estimation using witness complexes. *SPBG*, 4:157–166, 2004.
- Edelsbrunner, H., Letscher, D., and Zomorodian, A. Topological persistence and simplification. In *Foundations of Computer Science, 2000. Proceedings. 41st Annual Symposium on*, pp. 454–463. IEEE, 2000.
- Ghrist, R. Barcodes: the persistent topology of data. *Bulletin of the American Mathematical Society*, 45(1):61–75, 2008.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. In Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N. D., and Weinberger, K. Q. (eds.), *Advances in Neural Information Processing Systems 27*, pp. 2672–2680. Curran Associates, Inc., 2014.
- Goodfellow, I., Bengio, Y., Courville, A., and Bengio, Y. *Deep learning*, volume 1. MIT press Cambridge, 2016.
- Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., and Courville, A. Improved training of Wasserstein GANs. *arXiv preprint arXiv:1704.00028*, 2017.
- Hatcher, A. *Algebraic topology*. Cambridge University Press, 2002.
- Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Klambauer, G., and Hochreiter, S. GANs trained by a two time-scale update rule converge to a Nash equilibrium. *arXiv preprint arXiv:1706.08500*, 2017.
- Ho, J. and Ermon, S. Generative adversarial imitation learning. In *Advances in Neural Information Processing Systems 29*, pp. 4565–4573. Curran Associates, Inc., 2016.
- Kaczynski, T., Mischaikow, K., and Mrozek, M. *Computational homology*, volume 157. Springer Science & Business Media, 2006.
- Liu, Z., Luo, P., Wang, X., and Tang, X. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, 2015.
- Lucic, M., Kurach, K., Michalski, M., Gelly, S., and Bousquet, O. Are gans created equal? a large-scale study. *arXiv preprint arXiv:1711.10337*, 2017.
- Mao, X., Li, Q., Xie, H., Lau, R. Y., Wang, Z., and Smolley, S. P. Least squares generative adversarial networks. *arXiv preprint ArXiv:1611.04076*, 2016.
- Maria, C., Boissonnat, J.-D., Glisse, M., and Yvinec, M. The Gudhi library: Simplicial complexes and persistent homology. In *International Congress on Mathematical Software*, pp. 167–174. Springer, 2014.
- May, J. P. *A concise course in algebraic topology*. University of Chicago press, 1999.
- Paganini, M., de Oliveira, L., and Nachman, B. CaloGAN: Simulating 3D high energy particle showers in multi-layer electromagnetic calorimeters with generative adversarial networks. *arXiv preprint arXiv:1705.02355*, 2017.
- Radford, A., Metz, L., and Chintala, S. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.
- Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., Chen, X., and Chen, X. Improved techniques for training GANs. In Lee, D. D., Sugiyama, M., Luxburg, U. V., Guyon, I., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 29*, pp. 2234–2242. Curran Associates, Inc., 2016.
- Shao, H., Kumar, A., and Fletcher, P. T. The Riemannian geometry of deep generative models. *arXiv preprint arXiv:1711.08014*, 2017.
- Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., and Rabinovich, A. Going deeper with convolutions. In *Computer Vision and Pattern Recognition (CVPR)*, 2015. URL <http://arxiv.org/abs/1409.4842>.
- Zomorodian, A. and Carlsson, G. Computing persistent homology. *Discrete & Computational Geometry*, 33(2): 249–274, 2005.