

Generating survey databases with Wasserstein Generative Adversarial Networks.

Abstract

In a world increasingly surrounded by data, data privacy and anonymization are becoming more and more important. Under these circumstances, the need for synthetic databases that replicate the characteristics of the population while preserving privacy is arising. In this article, we investigate how we can use Wasserstein Generative Adversarial Networks (WGANs), developed by Arjovsky et al. (2017) in the context of image generation, to create synthetic survey databases. Survey data have both categorical and continuous variables and especially contain sampling weights that will be introduced in the proposed procedures. We evaluate the quality of the (generated) synthetic data through different indicators and especially a new practical and intuitive measure based on Support Vector Data Description (SVDD). All our analyses are achieved with the Labour Force Survey (LFS) data for Belgium.

Keywords:

Generative Adversarial Network, Wasserstein Generative Adversarial Network, Synthetic data set, Wasserstein distance, Sampling weights, Support Vector Data Description.

1. Introduction

In every domain, there exists a significant demand for data, yet this demand is juxtaposed with the necessity for privacy, thereby complicating access to data. There is an undeniable necessity for methods of anonymization, particularly in light of recent legislative developments such as the General Data Protection Regulation (GDPR) in Europe or the California Consumer Privacy Act (CCPA). The development of synthetic databases presents a viable solution to this dilemma, the demand for such data increasing alongside the growing volume of existing data.

Generative models can be constructed to generate synthetic data sets that share characteristics with the original data. Generative Adversarial Networks (GANs, Goodfellow et al., 2020) and Variational Autoencoders (VAEs, Kingma and Welling, 2022) are the two primary methods used in the context of synthetic data generation.

A comprehensive literature review providing a thorough exploration of the potential applications of GANs is presented by (Cheng et al., 2020). However, GANs have already been utilized for privacy purposes. Some efforts have been made to apply GANs, originally developed for generating fake images, to tabular data. For instance, Xu and Veeramachaneni (2018) employed GANs to forecast tabular data such as educational or medical records. Additionally, Camino et al. (2018) devised techniques to apply GAN and Wasserstein Generative Adversarial Network (WGAN) models to handle multiple categorical variables using softmax. Furthermore, Walia et al. (2020) employed the Wasserstein Conditional Generative Adversarial Network with Gradient Penalty (WCGAN-GP), utilizing the WGAN with the

Gradient Penalty framework (WGAN-GP) introduced by Gulrajani et al. (2017). WCGAN-GP is specifically tailored for tabular data, where critics (discriminators) incorporate label information.

In this article, we incorporate the methodologies proposed by Gulrajani et al. (2017) and Camino et al. (2018) for tabular data, specifically addressing scenarios where original data sets are associated with sampling weights. The inclusion of these weights is essential to uphold the representativeness of synthetic databases, particularly given that a significant portion of data sets incorporate sampling weights. This necessity is particularly evident in European surveys like Statistics on Income and Living Conditions (EU-SILC), Household Budget Survey (HBS), and Labor Force Survey (LFS), among others.

Assessing the quality of our synthetic data set is another pivotal aspect of this article. In Borji (2019) and Borji (2021), various methods are compared for assessing the quality of the results generated by GANs, particularly focusing on images. In this article, we have opted to utilize several relevant metrics from these studies, propose a new method for assessing the reliability of generated data sets, utilizing Support Vector Data Description (SVDD), and additionally incorporate the Wasserstein distance algorithm introduced by Schuhmacher et al. (2024). The advantage of our SVDD-based measure is that it provides a percentage indicating the extent to which synthetic data deviate from the original data.

The structure of the article is outlined as follows. Firstly, in Section 2 and 3, we offer background information about GANs and detail the methodologies that will be applied in this study. In Section 4, we introduce the different

measurement techniques, including our new approach. Finally, in Section 5, we present an application wherein we generate a synthetic version of the Labour Force Survey (LFS) data set for Belgium and compare the results using the various techniques outlined in Section 4.

2. GAN: literature review

Generative Adversarial Networks (GANs), initially proposed in Goodfellow et al. (2014), preprint of Goodfellow et al. (2020), were devised to generate synthetic images. In this article, GANs are introduced as a two-player minimax game between two multilayer perceptrons:

$$\min_G \max_D \mathbb{E}_{X \sim p_{\text{data}}(x)} [\log (D(X))] + \mathbb{E}_{Z \sim p_z(z)} [1 - \log (D(G(Z)))],$$

where G is the perceptron called generator, which generates synthetic data from noise Z coming from a distribution p_z and D is the perceptron called discriminator, which aims to correctly differentiate true data X , coming from the distribution p_{data} , from synthetic data $G(z)$, as represented in Figure 1.

Since their inception, GANs have been utilized in a growing number of applications, with image generation being the principal one. Additionally, Lu et al. (2019) demonstrated their application in discrete time series analysis.

In the context of anonymization, the study described in Fang et al. (2019) focuses on GANs for graphs, aiming to generate an anonymous graph with connection features similar to those of the original graph. For evaluation, they employed a combination of utility metrics and privacy metrics. In Feutry et al. (2018), GANs are employed to build a representation U from which




Figure 1: GAN

a specific type of label Y can be extracted. The construction of U also aims to prevent the extraction of private information contained in other labels Z , leading to a trade-off between these two objectives. This approach is applied to images to preserve crucial information (e.g., sentiment) while anonymizing the image. In the study conducted by Bae et al. (2019), GANs were employed to anonymize medical and DNA data by introducing noise through a generator from which the input is the true data x (and not only noise). The article outlines a procedure involving a generator (referred to as an encoder), a discriminator, and a target classifier. From an input x , an $\hat{x}$ is created. The learning objective of the encoder is to minimize the discriminator cross entropy (together with the distance between original and generated data) while maximizing the prediction score of the target classifier. The loss is calculated using the following formula:

$$L_E(\theta_E, x, r, \delta) = \lambda_e(d(x, \hat{x}) * \delta) + \lambda_d(L_{D1} + L_{D2})$$

where d is the Euclidean distance, δ controls the privacy level, L_{D1} is the loss of the discriminator (cross entropy) and L_{D2} is the target classifier score

$C(f(\hat{x}))$ (with a pretrained perceptron, for instance, one can train it to predict diseases for patients). In the context of tabular data, we can cite Park et al. (2018), an article that uses an adaptation of traditional Image-GAN (using 2-d convolutional networks) on table data sets. They transformed categorical data into matrices on which they use 2-d convolutional networks. This approach is criticized by Engelmann and Lessmann (2020). Indeed, convolutions are valuable for image processing but they appear unsuitable for tabular data where columns contain entirely heterogeneous content and their order is arbitrary. Engelmann and Lessmann (2020) use GAN and WGAN to create tabular data but in the context of oversampling methods (like SMOTE). They apply a GAN to estimate the distribution of the data and after, the generator provides additional samples of the minority classes using the Conditional Generative Adversarial Network structure (CGAN). CGAN was also examined by Xu et al. (2019) for its Conditional Tabular Generative Adversarial Network (CTGAN). In their approach, data are generated given each modality of certain categorical variables, which are randomly chosen. Moreover, the continuous variables are transformed through Gaussian mixture models according to their number of modes. Walia et al. (2020) also proposed a technique based on Wasserstein GAN with gradient penalty, similar to the approach outlined in Subsection 3.3, with the distinction that both the discriminator and generator are conditioned on additional information such as class labels (e.g., sick or not, payment default or not, etc.) and they do not use sampling weights. Vega-Márquez et al. (2020) have also explored a similar concept in the case of a simple CGAN.

In Camino et al. (2018), various methods adapted to multicategorical

variables are compared. Two methods employ GANs. First, Gumble-GAN is a traditional GAN utilizing a Gumble softmax layer for each categorical variable:

$$\sigma(\mathbf{x})_i = \frac{\exp\left(\frac{\log(x_i) + g_i}{\tau}\right)}{\sum_{j=1}^K \exp\left(\frac{\log(x_j) + g_j}{\tau}\right)}, \quad i = 1, \dots, K, \quad (1)$$

where, $\mathbf{x} = (x_1, \dots, x_K)^T$ represents the input, $g_1, \dots, g_K$ are *i.i.d.* random variables drawn from a Gumble $(0, 1)$ distribution, K is the number of levels of the categorical variable and τ is a positive hyperparameter. The second method involves a Wasserstein GAN utilizing the softmax layer:

$$\sigma(\mathbf{x})_j = \frac{\exp(x_j)}{\sum_{k=1}^K \exp(x_k)}, j = 1, \dots, K. \quad (2)$$

It is somewhat similar to what we present in Subsections 3.3 and 3.4, but with only categorical variables and without sampling weights. Additionally, they implemented the Gumbel softmax of Equation (1) on models utilizing autoencoders. Depending on the data sets and metrics employed, it is not possible to identify the best method. However, the Wasserstein GAN using the softmax method emerges as one of the top-performing approaches.

3. WGAN with sampling weights

3.1. Wasserstein GAN

A well-known problem in the GAN literature is mode collapsing (Metz et al., 2017; Zhang et al., 2018). It occurs when the generator produces data from a specific mode rather than from the entire original distribution.

The Wasserstein GAN is a distinctive method within the GAN framework, initially developed by Arjovsky et al. (2017). One of the aims of this method

is to counter this mode collapsing problem. This specific type of GAN utilizes the Wasserstein distance, also known as the Wasserstein-1 or Earth-Mover distance. Let $(\mathcal{X}, d)$ be a Polish space with metric d , and let $\mathbb{P}_r$ and $\mathbb{P}_\theta$ be two distributions in the space of probability measures defined on $\mathcal{X}$. Here, $\mathbb{P}_r$ (resp. $\mathbb{P}_\theta$) will denote the distribution of the original (resp. synthetic) data and we take $d(x, y) = \|x - y\|_2$. The Wasserstein distance between $\mathbb{P}_r$ and $\mathbb{P}_\theta$ is then given by

$$\begin{aligned} W(\mathbb{P}_r, \mathbb{P}_\theta) &= \inf_{\pi \in \Pi(\mathbb{P}_r, \mathbb{P}_\theta)} \mathbb{E}_{(X, Y) \sim \pi} [d(X, Y)] \\ &= \inf_{\pi \in \Pi(\mathbb{P}_r, \mathbb{P}_\theta)} \int_{\mathcal{X} \times \mathcal{X}} d(x, y) d\pi(x, y), \end{aligned} \quad (3)$$

where $\Pi(\mathbb{P}_r, \mathbb{P}_\theta)$ denotes the set of all joint distributions $\pi(x, y)$ whose marginals are respectively $\mathbb{P}_r$ and $\mathbb{P}_\theta$.

However, directly incorporating this distance into the GAN framework can pose challenges, particularly due to the computational and time-consuming nature of the Wasserstein distance. Yet, the following theorem can provide assistance.

Theorem 1. *Let $\mathbb{P}_r$ be any distribution. Let $\mathbb{P}_\theta$ be the distribution of $g_\theta(Z)$, where Z is a random variable with density p and $g_\theta(\cdot)$ is a locally Lipschitz function such that, for a certain probability distribution p over $\mathcal{Z}$, there are local Lipschitz constants $L(\theta, z)$ such that $\mathbb{E}_{Z \sim p} [L(\theta, Z)] < +\infty$. Then, there is a solution $f : \mathcal{X} \rightarrow \mathbb{R}$ to the problem:*

$$\max_{\|f\|_L \leq 1} \mathbb{E}_{X \sim \mathbb{P}_r} [f(X)] - \mathbb{E}_{X \sim \mathbb{P}_\theta} [f(X)] \quad (4)$$

and we have

$$\nabla_\theta W(\mathbb{P}_r, \mathbb{P}_\theta) = -\mathbb{E}_{Z \sim p} [\nabla_\theta f(g_\theta(Z))]$$

when both terms are well-defined and where $\max_{\|f\|_L \leq 1}$ denotes the maximum on the class of 1-Lipschitz functions $f : \mathcal{X} \rightarrow \mathbb{R}$.

The proof (Arjovsky et al., 2017) of the theorem is grounded on three key points:

1. Kantorovich-Rubinstein duality that tell us that:

$$W(\mathbb{P}_r, \mathbb{P}_\theta) = \sup_{\|f\|_L \leq 1} (\mathbb{E}_{x \sim \mathbb{P}_r} [f(x)] - \mathbb{E}_{x \sim \mathbb{P}_\theta} [f(x)]).$$

The proof of this duality can be found in Villani (2009);

2. the envelope theorem;
3. the dominate convergence theorem.

The aim is to substitute the Wasserstein distance that needs to be minimized by the difference between the two expectations in Equation (4). To accomplish this, it is imperative for our discriminator, commonly referred to as the critic in WGAN literature, to belong to the 1-Lipschitz class. It is worth noting that if f belongs to the K -Lipschitz class, the Kantorovich-Rubinstein duality becomes:

$$K \cdot W(\mathbb{P}_r, \mathbb{P}_\theta) = \sup_{\|f\|_L \leq K} \mathbb{E}_{x \sim \mathbb{P}_r} [f(x)] - \mathbb{E}_{x \sim \mathbb{P}_\theta} [f(x)],$$

and we approximate up to an irrelevant scaling factor.

To enforce this Lipschitz constraint in the literature, two techniques are available:

- the weight clipping, introduced for images by Arjovsky et al. (2017);
- the gradient penalty, introduced for images by Gulrajani et al. (2017).

3.2. WGAN with weight clipping

In this article, we adapt WGAN of Arjovsky et al. (2017) by incorporating sampling weights, see Algorithm 1. The concept behind weight clipping is to confine the weights of neural networks to a restricted space $\mathcal{W}$, thereby ensuring that the critic f remains K -Lipschitz. In the following algorithms, $f_\eta(\cdot)$ denotes the critic neural network, sampling is made with replacement, w_i , $i = 1, \dots, N$, is the sampling weight (such that their sum is equal to 1) and the chosen generalization of the stochastic gradient descent method is Adam algorithm (Kingma and Ba, 2017). ADAM is a first-order gradient-based optimization algorithm that uses exponentially smoothed gradient and gradient squared of the considered loss function. It combines the advantages of Root Mean Square Propagation (RMSProp), which stabilizes gradients across the iterations (by normalizing them) and Adaptive Gradient (AdaGrad), which deals with sparse gradients by keeping sufficiently large learning rates for rare parameters (which especially appear in NLP where inputs are often sparse vectors/matrices). Unlike AdaGrad, ADAM deals with sparse gradients through their exponential smoothing. In our experiments, ADAM provided the best results.

Remark: Another solution is to draw samples from the unweighted empirical distributions of the real data and then use the weights (see Appendix A).

3.3. WGAN with gradient penalty

We implement the approach introduced by Gulrajani et al. (2017). The gradient penalty capitalizes on the fact that a differentiable function on a convex set is 1-Lipschitz if and only if its gradient has a norm smaller or

Algorithm 1 WGAN with weight clipping and sampling weights

Require: α , the learning rate. c , the clipping parameter. m , the batch size.

n_{critic} , the number of iterations of the critic by generator iteration. β_1 and β_2 the parameter for the Adam algorithm.

Require: η_0 , the initial parameters of the critic. θ_0 , the initial parameters of the generator. Chosen randomly.

- 1: **while** θ has not converged **do**
 - 2: **for** $t = 0, \dots, n_{critic}$ **do**
 - 3: Sample $\{x^{(i)}\}_{i=1}^m \sim \mathbb{P}_r$ a batch from the weighted (with w_i) empirical distribution.
 - 4: Sample $\{z^{(i)}\}_{i=1}^m \sim p(z)$ a batch of prior data.
 - 5: $\gamma_\eta \leftarrow \nabla_\eta \left[\frac{1}{m} \sum_{i=1}^m f_\eta(g_\theta(z^{(i)})) - \frac{1}{m} \sum_{i=1}^m f_\eta(x^{(i)}) \right]$
 - 6: $\eta \leftarrow \text{Adam}(\eta, \gamma_\eta, \alpha, \beta_1, \beta_2)$
 - 7: $\eta \leftarrow \text{clip}(\eta, -c, c)$
 - 8: **end for**
 - 9: Sample $\{z^{(i)}\}_{i=1}^m \sim p(z)$ a batch of prior data.
 - 10: $\gamma_\theta \leftarrow -\nabla_\theta \frac{1}{m} \sum_{i=1}^m f_\eta(g_\theta(z^{(i)}))$
 - 11: $\theta \leftarrow \text{Adam}(\theta, \gamma_\theta, \alpha, \beta_1, \beta_2)$
 - 12: **end while**
-

equal to 1 everywhere. That explains the expression of $\Gamma^{(i)}$ in Algorithm 2 hereunder.

More details can be found in Gulrajani et al. (2017). One of the interests of this approach is that it allows for having fewer vanishing gradient problems¹ that occur when WGAN-CP has a small clipping parameter (Arjovsky et al., 2017; Gulrajani et al., 2017). It mitigates also the risk of exploding gradient² (Gulrajani et al., 2017).

3.4. Our GAN structure

First, all hidden layers in the generator utilize a LeakyReLU activation function with a specific ϵ hyperparameter. This activation function is defined as

$$h(x) = \begin{cases} x & \text{if } x > 0, \\ \epsilon x & \text{otherwise.} \end{cases}.$$

The use of LeakyReLU can help mitigate the vanishing gradient problems that can occur with sigmoidal or hyperbolic tangent activation functions (Maas et al., 2013). Additionally, elastic net regularization is incorporated to ensure the generalization of the model and prevent overfitting. The output layer is tailored to accommodate both continuous and categorical variables. For continuous variables, the output activation function is LeakyReLU; for categorical variables, the softmax activation function in Equation (2) is used,

¹Vanishing gradient occurs when the gradients of the loss function (here in Algorithm 2) become extremely small during backpropagation, causing the neural weights to update minimally and hindering effective learning.

²Exploding gradient occurs when the gradients of the loss function (here in Algorithm 2) become extremely large.

Algorithm 2 WGAN with gradient penalty and sampling weights

Require: α , the learning rate. m , the batch size. n_{critic} , the number of iterations of the critic by generator iteration. λ , the gradient penalty parameter. β_1 and β_2 the parameter for the Adam algorithm.

Require: η_0 , the initial parameters of the critic. θ_0 , the initial parameters of the generator. Chosen randomly.

```
1: while  $\theta$  has not converged do
2:   for  $t = 0, \dots, n_{critic}$  do
3:     Sample  $\{x^{(i)}\}_{i=1}^m \sim \mathbb{P}_r$  a batch from the weighted (with  $w_i$ ) empirical
       distribution.
4:     Sample  $\{z^{(i)}\}_{i=1}^m \sim p(z)$  a batch of prior data.
5:     for  $i = 1, \dots, m$  do
6:       Sample  $\kappa \sim U(0, 1)$ 
7:        $\tilde{x}^{(i)} = g_\theta(z^{(i)})$ 
8:        $\hat{x}^{(i)} = \kappa x^{(i)} + (1 - \kappa) \tilde{x}^{(i)}$ 
9:        $\Gamma^{(i)} = \left( f_\eta(\tilde{x}^{(i)}) - f_\eta(x^{(i)}) + \lambda (\|\nabla_{\hat{x}} f_\eta(\hat{x}^{(i)})\|_2 - 1)^2 \right)$ 
10:    end for
11:     $\gamma_\eta \leftarrow \nabla_\eta \left[ \frac{1}{m} \sum_{i=1}^m \Gamma^{(i)} \right]$ 
12:     $\eta \leftarrow \text{Adam}(\eta, \gamma_\eta, \alpha, \beta_1, \beta_2)$ 
13:  end for
14:  Sample  $\{z^{(i)}\}_{i=1}^m \sim p(z)$  a batch of prior data.
15:   $\gamma_\theta \leftarrow -\nabla_\theta \frac{1}{m} \sum_{i=1}^m f_\eta(g_\theta(z^{(i)}))$ 
16:   $\theta \leftarrow \text{Adam}(\theta, \gamma_\theta, \alpha, \beta_1, \beta_2)$ 
17: end while
```

as proposed by Camino et al. (2018). Second, our architecture of the critic includes several key features. Similarly to the generator, all of its hidden layers employ a LeakyReLU activation function with a specific ϵ parameter and elastic net regularization is applied to the model. Interestingly, unlike the conventional GAN where the output layer typically uses a sigmoid (standard logistic function) activation function, the output layer of our critic is a linear function.

Although dropout was experimented, it did not yield better results, so it was not included in the final architecture.

3.5. WGAN with quantile transformers

Others have employed quantile transformers within the framework of GANs (Xu and Veeramachaneni, 2018) or VAEs (Jamotton and Hainaut, 2023). In this article, we adapt the WGAN with gradient penalty to implement such a transformation. Our approach involves forcing all the continuous variables to adhere to a standard normal distribution by utilizing the quantile transformer formula $\Phi^{-1}(F(X_k))$. Here, $F(X_k)$, the cumulative distribution of the continuous variable X_k , follows a uniform distribution, resulting in $\Phi^{-1}(F(X_k))$ conforming to a normal distribution. In practice, for each point x_{kj} , we utilize the transformation:

$$\Phi^{-1} \left(\frac{\sum_{i=1}^N \left(w_i \mathbb{1}_{\text{rank}(x_{ki}) \leq \text{rank}(x_{kj})} - 0.5 * w_j \right)}{\sum_{i=1}^N w_i} \right),$$

where N is the size of the data set. Subsequently, we apply a min-max transformation to confine our variables within the range $[0, 1]$. We then employ a sigmoid function for the part of the output layer of the generator which

handles the continuous variables. To distinguish this technique from the one outlined in 3.3, we will refer to it as WGAN-QT in the remainder of the article. The purpose is to make the neural networks more effective. Variables with more complex distributions (non-normal, skewed, long-tailed, etc.) can disrupt the balance of our training data, impacting the generalization and performance of our models. Additionally, multimodal distributions can exacerbate vanishing gradient problems, as pointed out by Pascanu et al. (2013) and Jamotton and Hainaut (2023).

3.6. Other GANs to compare the results

We compare the methods explained in Sections 3.2, 3.3 and 3.5 with CopulaGANs and CTGANs. CopulaGAN first transforms the data before applying CTGAN. CTGANs (Xu et al., 2019) use Wasserstein GAN with gradient penalty, Gumbel softmax functions (as in Equation (1)), a conditional approach and Gaussian mixture models as explained in Section 2. The code of both (CTGAN and CopulaGAN) can be found in DataCebo (2024). CopulaGANs and CTGANs do not incorporate sampling weights. Additionally, we will assess the results by comparing them to those obtained from a vanilla GAN with a structure similar to the one described in Subsections 3.4 and 5.2, except for the output layer of the discriminator. We will provide two versions: with and without sampling weights.

4. Methods of evaluation

In Subsection 4.1, we present the measures found in the literature that we have decided to use, and in Subsection 4.2, we put forward a new measure.

4.1. Literature Review

In the studies by Borji (2019) and Borji (2021), various evaluation methods for GANs are outlined. From these papers, we have chosen the most suitable evaluation techniques for our specific task of generating survey data. We will also incorporate additional methods that appear relevant to our specific needs.

Among all these methods, we will provide: Fréchet Inception Distance, Density and Coverage and Kullback–Leibler divergence.

4.1.1. Fréchet Inception distance

Fréchet Inception Distance (Heusel et al., 2017) is one of the widely used measures in the context of GANs, that consists in assuming that the distribution of the data is close to multivariate normal distributions $\mathcal{N}(m_1, \mathbf{C}_1)$ and $\mathcal{N}(m_2, \mathbf{C}_2)$, for the original and the generated data respectively. This leads to an easy calculable 2-Wasserstein distance:

$$FID = \|m_1 - m_2\|_2^2 + \text{Tr} \left(\mathbf{C}_1 + \mathbf{C}_2 - 2(\mathbf{C}_1 \cdot \mathbf{C}_2)^{\frac{1}{2}} \right).$$

Although widely used, assuming this normality is unrealistic in practice. In our case, we only employ this measure for continuous variables.

4.1.2. Density and Coverage

Density and Coverage are proposed by Naeem et al. (2020):

$$\begin{aligned} \text{Density} &= \frac{1}{k} \sum_{j=1}^n \sum_{i=1}^n w_i \mathbb{1}_{y_j \in B(x_i, NDD_k(x_i))}; \\ \text{Coverage} &= \sum_{i=1}^n w_i \mathbb{1}_{\exists j \text{ s.t. } y_j \in B(x_i, NDD_k(x_i))}; \end{aligned}$$

where each x_i , $i = 1, \dots, n$, represents a data point from a test data set (with weights w_i such that their sum is equal to 1), while each y_j , $j = 1, \dots, n$, represents a synthetic data point. The notation $B(x, r)$ denotes a sphere in $\mathbb{R}^d$ (where d is the number of variables) centered at x with radius r , and $NDD_k(x_i)$ represents the distance from x_i to the k -th nearest neighbor among all x_t for $t \neq i$.

The fundamental concept involves constructing spheres around each real data point such that they contain the k real nearest neighbors. These spheres form a manifold around the original data set through their union. Then, for the density, we evaluate the average number of synthetic data points contained in each sphere divided by k . Meanwhile, for the coverage, we assess the proportion of spheres containing a point of the synthetic data set. The objective is to devise measures capable of assessing both the fidelity and diversity of the generated data set. These methods represent advancements over older techniques such as Precision and Recall, as proposed by Kynkäänniemi et al. (2019), offering increased robustness against outliers and mode dropping.

4.1.3. *Kullback–Leibler divergence*

We also provide Kullback–Leibler divergence, computed utilizing a knn-tree method as described in Boltz et al. (2007).

4.1.4. *Wasserstein-1 distance*

We also compute the Wasserstein-1 distance as defined in Equation (3) using the method of Schuhmacher et al. (2024). The closer it is to zero, the more similar the two multivariate distributions are.

4.1.5. Cramér-von Mises

We employ the empirical Cramér-von Mises criterion introduced in Annoye et al. (2024):

$$\widetilde{\text{CVM}} = n \sum_{i=1}^n w_i \left[\hat{F}_n(x_i^1, x_i^2) - \hat{G}_n(x_i^1, x_i^2) \right]^2, \quad (5)$$

where $\hat{F}_n$ (resp. $\hat{G}_n$) is the weighted empirical bivariate cumulative distribution function of any pair of variables constructed with the original (resp. synthetic) data set and the sum is over their values in the original data set. We then average the $\widetilde{\text{CVM}}$ s of all possible couples.

4.2. New measure using Support Vector Data Description

In this section, we propose to use Support Vector Data Description (SVDD, Tax and Duin, 2004) to obtain a new similarity measure. First, we describe the SVDD methodology in Subsection 4.2.1. For implementary reasons, we resort to One-Class Support Vector Machine (OC-SVM) in Subsection 4.2.2 and its link with SVDD in Subsection 4.2.3. Next, we introduce the sampling weights in Subsection 4.2.4 and the proposed measure in Subsection 4.2.5.

4.2.1. Support Vector Data Description

At its core, SVDD operates on the assumption that the densest region of our data set is represented by a single hypersphere with unknown radius R and center ζ .

Mathematically, SVDD aims to solve the following problem (Tax and Duin, 2004):

$$\min_{R \in \mathbb{R}, \zeta \in \mathcal{H}, \xi_1, \dots, \xi_n \in \mathbb{R}^+} R^2 + \frac{1}{\nu n} \sum_{i=1}^n \xi_i \quad (6)$$

subject to

$$\|\phi(x_i) - \zeta\|_{\mathcal{H}}^2 \leq R^2 + \xi_i, \quad \xi_i \geq 0, \text{ for } i = 1, \dots, n, \quad (7)$$

where ϕ is an implicit feature mapping (to apply the kernel trick) into an inner product space with norm $\|\cdot\|_{\mathcal{H}}$, ξ_i s are slack variables which allow for some outliers (i.e. points out of the hypersphere) and $\nu \in]0, 1]$ is a parameter controlling the influence of the slack variables. It is proved in Schölkopf et al. (2001) that ν is an upper bound on the fraction of outliers.

Constraints (7) can be incorporated into Equation (6) by using Lagrange multipliers; the resulting dual problem becomes

$$\min_{\mathbf{a} \in \mathbb{R}^n} \sum_{i,j=1}^n a_i a_j k(x_i, x_j) - \sum_{i=1}^n a_i k(x_i, x_i)$$

subject to

$$\begin{aligned} 0 \leq a_i &\leq \frac{1}{\nu n}, \quad i = 1, \dots, n, \\ \mathbf{e}^T \mathbf{a} &= 1, \end{aligned} \quad (8)$$

where $\mathbf{e} = (1, \dots, 1)^T$ and $\mathbf{a} = (a_1, \dots, a_n)^T$ is a vector of Lagrange multipliers.

This leads to the following decision function (Schölkopf et al., 2001):

$$\begin{aligned} f(x) &= R^2 - \|\phi(x) - \zeta\|_{\mathcal{H}}^2 \\ &= R^2 - k(x, x) + 2 \sum_{i=1}^n a_i k(x, x_i) - \sum_{i,j=1}^n a_i a_j k(x_i, x_j). \end{aligned} \quad (9)$$

In Figure 2, we represent the resulting hypersphere (as if it were two-dimensional).




Figure 2: SVDD

4.2.2. One-Class Support Vector Machine

OC-SVM is another algorithm for outlier detection. It consists in projecting the data in a hyperspace, where a hyperplane can be constructed to separate the normal points from the outliers. Given some vectors $x_i \in \mathbb{R}^n$, $i = 1, \dots, n$, the primal problem of OC-SVM is (Lampert et al., 2009)

$$\min_{\mathbf{v} \in \mathcal{H}, \xi_1, \dots, \xi_n \in \mathbb{R}^+, \rho \in \mathbb{R}} \frac{1}{2} \mathbf{v}^T \mathbf{v} - \rho + \frac{1}{\nu n} \sum_{i=1}^n \xi_i \quad (10)$$

subject to

$$\begin{aligned} \mathbf{v}^T \phi(x_i) &\geq \rho - \xi_i \\ \xi_i &\geq 0, \quad i = 1, \dots, n. \end{aligned} \quad (11)$$

In this context, $\mathbf{v}$ denotes the vector of coefficients of the hyperplane, ρ indicates the offset of the hyperplane from the origin, ϕ represents a mapping in the hyperspace, and the ξ_i s stand for the slack variables.

The dual problem of (10) and (11) is

$$\min_{\mathbf{a} \in \mathbb{R}^n} \frac{1}{2} \mathbf{a}^T Q \mathbf{a}$$

subject to

$$\begin{aligned} 0 &\leq a_i \leq \frac{1}{\nu n}, \quad i = 1 \dots, n \\ \mathbf{e}^T \mathbf{a} &= 1, \end{aligned}$$

where Q_{ij} , the element (i, j) of the matrix Q is equal to $k(x_i, x_j)$. The decision function is then (Lampert et al., 2009)

$$f(x) = \sum_{i=1}^n a_i k(x, x_i) - \rho. \quad (12)$$

4.2.3. Link between OC-SVM and SVDD

In practice, we opt for OC-SVM instead of SVDD, both algorithms having, with a Gaussian kernel, the same dual problem and thus the same a_i s.

For every $0 < a_j < \frac{1}{\nu n}$, ρ can be calculated from (12) with $x = x_j$. Next, from (9), we obtain

$$R^2 = 1 + \sum_{i,j=1}^n a_i a_j k(x_i, x_j) - 2\rho. \quad (13)$$

4.2.4. Sampling weights

If we introduce the sampling weights w_i s, (6) and (10) become

$$\min_{R \in \mathbb{R}, \zeta \in \mathcal{H}, \xi_1, \dots, \xi_n \in \mathbb{R}^+} R^2 + \frac{1}{\nu} \sum_{i=1}^n w_i \xi_i$$

and

$$\min_{\mathbf{v} \in \mathcal{H}, \xi_1, \dots, \xi_n \in \mathbb{R}^+, \rho \in \mathbb{R}} \frac{1}{2} \mathbf{v}^T \mathbf{v} - \rho + \frac{1}{\nu} \sum_{i=1}^n w_i \xi_i$$

respectively. This results in the same dual problems as in the unweighted case, except for the constraints in (8) that become

$$0 \leq a_i \leq \frac{w_i}{\nu}.$$

4.2.5. Proposed measure

Using SVDD, we calculate the radius in Equation (13) for both the original and the generated data sets and compare them through the relative difference

$$\frac{R_1 - R_2}{R_1}, \quad (14)$$

where R_1 (resp. R_2) is the radius for the original (resp. generated) data set.

The significance of this measure lies in its expression as a ratio, offering a more effective basis for interpretation. Unlike the Wasserstein distance, which has a clear interpretation at 0 but lacks clarity for values such as 1 or 10, this measure provides a more intuitive and practical comparison because it represents the percentage of deviation between the two data sets.

5. Application

5.1. Implementation

We utilize a combination of R and Python softwares, incorporating Keras and Tensorflow packages, to execute our experiments. Python serves as the primary tool for machine learning tasks, while R is employed specifically for cross-validation and various analyses such as Wasserstein distance and SVDD.

We have autonomously developed GANs using Keras and Tensorflow, except for two variants based on conditional tabular GANs (CTGAN and CopulaGAN), for which we utilize the implementation provided by the SDV package in Python (DataCebo, 2024).

5.2. Methodology

We have divided the Labour Force Survey for Belgium 2019 (LFS, 146 variables) into two parts: two-thirds of the data for training and validation ($N = 24266$) and one-third for the test set ($n = 12133$).

On the first part, we employ a validation procedure to select the hyperparameters. This involves further dividing this set into two parts, with two-thirds for training and one-third for validation. The list of hyperparameters for the three Wasserstein GANs and the vanilla GANs can be found in Table 1. Notice that, in the vanilla GANs, initial neural networks parameters are also chosen randomly, sampling is also made with replacement, $n_{critic} = 1$, a sigmoid function is used at the end of the discriminator, the loss function for both the generator and the discriminator is binary crossentropy, LeakyReLU and softmax functions are used as in the Wasserstein GANs and

λ and c are not employed. To determine the optimal hyperparameters, we select the combination that minimizes the Wasserstein distance (3), employing the methodology developed in Schuhmacher et al. (2024) for its calculation. Due to time limitations, we opt not to employ cross-validation in our process. Training neural networks of this nature is time-intensive, and incorporating cross-validation would extend the tuning period by 200%.

For the CTGAN and the CopulaGAN within the SDV package (DataCebo, 2024), we have kept the default parameters. However, these two methods take 13 times longer (than the above GANs) to be executed, making it challenging to fine-tune. These default values are used in Xu et al. (2019) and can be found in Table 2.

Next, for each GAN model, we generate a synthetic data set that matches the size of the test set. We then compare these synthetic and original data to assess their similarities and differences.

5.3. Results

The SVDD measure is provided for all the variables and also exclusively for continuous variables, both with ν equal to 0.05 and a Gaussian kernel with γ equal to one divided by the number of variables:

$$k(x_i, x_j) = \exp(\gamma \|x_i - x_j\|_2^2).$$

In Table 3, we compute all the quantities presented in Section 4. To prevent certain variables from dominating the measure, we have standardized all the continuous variables.

WGAN-GP appears to yield the most favorable results for Wasserstein distance and $SVDD/SVDD_{cont}$ while WGAN-QT has the lowest average

$\widetilde{\text{CVM}}$, KL and FID . Techniques employing sampling weights generally outperform CTGAN and CopulaGAN (except for the vanilla GAN with sampling weights). Among the three WGAN methods with sampling weights, those using gradient penalty perform better than the one using the clipping constraint. Unsurprisingly, the worst measures are obtained for the two vanilla GANs.

Hyperparameter	Type of search
Number of hidden layers	Fixed: 2
Number of neurons by hidden layer generator	Random search
Number of neurons by hidden layer discriminator	Random search
ϵ , LeakyReLU	Random search
α , learning rate	Random search
λ_{1g} , elastic net penalty of the generator loss (lasso)	Random search
λ_{2g} , elastic net penalty of the generator loss (ridge)	Random search
λ_{1d} , elastic net penalty of the discriminator loss (lasso)	Random search
λ_{2d} , elastic net penalty of the discriminator loss (ridge)	Random search
Epoch	Fixed: 500 (validation) 2000 (final)
m , batch size	Fixed: 256 (validation) 128 (final)
β_1 , Adam	Fixed: 0.5
β_2 , Adam	Fixed: 0.9
n_{critic}	Fixed: 5
z	Fixed: $\mathcal{N}(\mathbf{0}, \mathbf{I})$, with $\mathbf{I}$ the identity matrix of dimension 15.
λ , gradient penalty parameter	Random search
c , clipping parameter	Random search

Table 1: Hyperparameters for the Wasserstein and the vanilla GANs; final, final estimation (with N data) with optimal hyperparameters; one epoch corresponds to $\frac{N}{m}$ batches in the update part of the generator.

Hyperparameter	Value
Number of hidden layers	Fixed: 2
Number of neurons by hidden layer generator	Fixed: (256, 256)
Number of neurons by hidden layer discriminator	Fixed: (256, 256)
ϵ , LeakyReLU	Fixed: 0.2
α , learning rate	Fixed: $2 \cdot 10^{-4}$
λ_{2g} , ridge penalty of the generator loss	Fixed: $1 \cdot 10^{-6}$
λ_{2d} , ridge penalty of the discriminator loss	Fixed: $1 \cdot 10^{-6}$
Epoch	Fixed: 300
m , batch size	Fixed: 500
β_1 , Adam	Fixed: 0.5
β_2 , Adam	Fixed: 0.9
n_{critic}	Fixed: 1
z	Fixed: $\mathcal{N}(\mathbf{0}, \mathbf{I})$, with $\mathbf{I}$ the identity matrix of dimension 128.
λ , gradient penalty parameter	Fixed: 10
τ , Gumble softmax parameter	Fixed: 0.2

Table 2: Hyperparameters for the CTGAN and the CopulaGAN.

Measure	WGAN-GP	WGAN-CP	WGAN-QT	CTGAN	CopulaGAN	GAN weighted	GAN
Wasserstein	14.81	15.01	14.83	15.20	15.21	15.64	15.54
$SVDD$	0.001	-0.068	-0.035	0.020	0.017	0.195	0.250
$SVDD_{cont}$	0.001	-0.015	-0.006	0.001	-0.007	0.012	0.007
KL	6301	6243	6098	6317	6319	6501	6814
FID	8.17	9.50	2.28	14.55	15.45	37.06	34.50
Av. $\widetilde{CVM}$	0.005	0.007	0.002	0.015	0.012	0.048	0.061

Table 3: Results: Wasserstein, Wasserstein-1 distance; $SVDD$ (resp. $SVDD_{cont}$), relative difference of the radii (14) for all the variables (resp. the continuous variables); KL , Kullback-Leibler divergence; Av. $\widetilde{CVM}$, average of the $\widetilde{CVM}$; GAN weighted (resp. GAN), vanilla GAN with (resp. without) sampling weights.

Both $SVDD$ and $SVDD_{cont}$ measures appear quite interesting. Notably, some differences are substantial, whereas the variations in terms of Wasserstein distance are relatively small. Both vanilla GANs perform poorly for $SVDD$ while the CTGAN and CopulaGAN have better results than WGAN-QT (WGAN-QT, however, surpasses CTGAN and CopulaGAN for all other measures). Observing the performances of WGAN-QT and WGAN-CP for $SVDD$ and $SVDD_{cont}$, they appear to exhibit suboptimal results compared to WGAN-GP. Note that both methods (WGAN-CP and WGAN-QT) have negative $SVDD$ and $SVDD_{cont}$, implying that the synthetic data set has a certain number of points outside the original data set support, which is not the case with the other techniques.

Next, we provide some graphs about Coverage and Density for our different data sets. In Figure 3 (resp. 4), we illustrate Density (resp. Coverage) as a function of the number of neighbors within each ball, ranging up to 5000 (on a logarithmic scale).

We observe that WGAN-QT and WGAN-GP provide the best results while the vanilla GANs perform poorly. CTGAN and CopulaGAN have lower Coverage and Density than the Wasserstein GANs with sampling weights but outperform vanilla GANs (mainly for Coverage).

In Figure 4, WGAN-QT and WGAN-GP exhibit good Coverage for small balls with $k = 100$ while CTGAN and Copula GAN (resp. vanilla GANs) reach this Coverage level for k between 2000 and 3000 (resp. k close to 5000). As above, methods without sampling weights underperform compared to the weighted approaches (more precisely, Wasserstein GANs (resp. vanilla GAN) with weights outperform CTGAN and CopulaGAN (resp. vanilla GAN).




Figure 3: Density




Figure 4: Coverage

without weights).

Moreover, we also verified that none of the individuals within our synthetic data set were duplicates of individuals in our original data set and this was indeed confirmed. However, for privacy reasons, if such a duplicate appeared, a new synthetic individual would be generated and compared to the original ones.

To sum up, applied to our data, WGAN-QT produces results similar

(sometimes better) to WGAN-GP and both outperform the other methods. However, WGAN-QT is more time-consuming due to the quantile transform. In Appendix B, we provide frequency (resp. estimated density) plots to compare original and synthetic data for some categorical variables (resp. one continuous variable). Graphs for other variables lead to similar observations. Finally, in Appendix Appendix C, we provide an additional study on the standard deviation of the results in Table 3.

6. Conclusion

In this paper, we apply various data generation techniques using different innovations of the Wasserstein GAN. We have adjusted our algorithms to consider sampling weights and have also introduced new measures to evaluate quality of the generated data.

In our experiments, we have seen that the methods using gradient penalty (WGAN-GP and WGAN-QT) lead to the best results (using various metrics). Both algorithms use sampling weights and the Wasserstein distance, and outperform the other tested techniques that do not use them. WGAN-QT seems to have results that are comparable to WGAN-GP (slightly better or slightly worse, depending of the measures) but it takes more time. Furthermore, the *SVDD* measure we have computed appears promising. From the kernel trick, the SVDD algorithm provides a frontier that surrounds the data points more precisely than, for example, a hypersphere with a dimension equal to the number of variables. Next, as it reflects the relative difference between two radii (for the hyperspheres encompassing the original and synthetic data set, respectively), this measure can be relevant for the practitioner who has

at disposal a percentage of deviation of the generated data with respect to the original data. It is more intuitive than distance measures (like, for example, the Wasserstein distance) which can make it challenging to determine whether the obtained quantity is substantial or not. In our experiments, WGAN-GP achieves the best *SVDD* (and the lowest Wasserstein distance).

Other GAN techniques remain unexplored in our investigation. For instance, Baowaly et al. (2019) compare GAN, WGAN and BGAN (Boundary-Seeking GAN, where the generator is trained to produce data points lying on the decision boundary of the discriminator) using medical data. BGAN appears to outperform others in their study, although WGAN also performs well. Another example is provided by Bellemare et al. (2017) and Mottini et al. (2018), who utilize Cramér-GAN as an alternative to Wasserstein GAN. This technique employs the energy distance:

$\mathcal{E}(P, Q) := 2\mathbb{E}\|X - Y\|_2 - \mathbb{E}\|X - X'\|_2 - \mathbb{E}\|Y - Y'\|_2$ where X, X' (resp. Y, Y') are independent random variables distributed according to P (resp. Q). The interest of this distance lies in its unbiased sample gradient, unlike the Wasserstein distance. An interesting addition to our analysis could involve comparing our GAN algorithms with techniques based on Variational Autoencoders, which represent another prominent approach in the synthetic data generation literature. Another potential avenue for research involves improving the exploration of hyperparameters and enhancing the time efficiency of GANs, which currently represent a general bottleneck.

Appendix A. Alternative algorithm

Here we provide an alternative to Algorithm 1:

Algorithm 3 WGAN-CP with weight: Method 2

Require: α the learning rate. c , the clipping parameter. m , the batch size.

n_{critic} , the number of iterations of the critic by generator iteration. β_1 and β_2 the parameter for the Adam algorithm.

Require: η_0 , the initial parameters of the critic. θ_0 , the initial parameters of the generator. Chosen randomly.

```

1: while  $\theta$  has not converged do
2:   for  $t = 0, \dots, n_{critic}$  do
3:     Sample  $\{x^{(i)}\}_{i=1}^m \sim \mathbb{P}_r$  a batch from the unweighted empirical dis-
       tribution.
4:     Sample  $\{z^{(i)}\}_{i=1}^m \sim p(z)$  a batch of prior data.
5:      $\gamma_\eta \leftarrow \nabla_w \left[ \frac{1}{m} \sum_{i=1}^m f_\eta(g_\theta(z^{(i)})) - \frac{1}{\sum_{i=1}^m w_i} \sum_{i=1}^m w_i f_\eta(x^{(i)}) \right]$ 
6:      $\eta \leftarrow \text{Adam}(\eta, \gamma_\eta, \alpha, \beta_1, \beta_2)$ 
7:      $\eta \leftarrow \text{clip}(\eta, -c, c)$ 
8:   end for
9:   Sample  $\{z^{(i)}\}_{i=1}^m \sim p(z)$  a batch of prior data.
10:   $\gamma_\theta \leftarrow -\nabla_\theta \frac{1}{m} \sum_{i=1}^m f_\eta(g_\theta(z^{(i)}))$ 
11:   $\theta \leftarrow \text{Adam}(\theta, \gamma_\theta, \alpha, \beta_1, \beta_2)$ 
12: end while

```

Appendix B. Graphs

In this section, we provide some graphs comparing the original test set with a synthetic data set constructed using WGAN-GP. Plots with synthetic data generated with WGAN-QT are similar.




Figure B.5: Frequency of each gender in the synthetic (above) and original (below) data.
 x -axis: 1, male; 2, female.




Figure B.6: Frequency of people in each province in the synthetic (above) and original (below) data. x -axis, label of the province in LFS.




Figure B.7: Frequency of people in each educational level in the synthetic (above) and original (below) data. x -axis, label of the education level in LFS.




Figure B.8: Frequency of people with instruction outside the regular education system in the synthetic (above) and original (below) data. x -axis: 1, outside; 2, inside; 9, not applicable.




Figure B.9: Estimated density of age in the synthetic (purple) and original (green) data.

Appendix C. Standard deviation

In this section, we generate one hundred synthetic data sets using each method proposed in this chapter. Then, we study the standard deviation of the measures in Table 3 across the different synthetic data sets, except for the average $\widetilde{\text{CVM}}$ due to time constraints. We observe that all methods have similar standard deviations for the Wasserstein distance as well as for FID. For $SVVD$ and $SVDD_{cont}$, the standard deviations are quite small, confirming that the vanilla GANs underperform compared to the other methods and that WGAN-GP has an $SVDD$ value close to 0. Finally, the Kullback-Leibler divergence, although it may initially appear larger, is actually quite small relative to the value of the measure. These standard deviations show that the different predictions made with the same method all have similar distances to the original data.

Measure	WGAN-GP	WGAN-CP	WGAN-QT	CTGAN	CopulaGAN	GAN weighted	GAN
Wasserstein	0.0019	0.0016	0.0017	0.0019	0.0019	0.0010	0.0010
$SVDD$	0.0030	0.0045	0.0041	0.0033	0.0033	0.0040	0.0029
$SVDD_{cont}$	0.0006	0.0003	0.0003	0.0004	0.0004	0.0005	0.0004
KL	1.0216	1.4261	4.6379	0.02910	0.1161	0.0694	0.0180
FID	0.0061	0.0154	0.0336	0.0345	0.0340	0.0432	0.0384

Table C.4: Standard deviation of the results over 100 generations: Wasserstein, Wasserstein-1 distance; FID, Fréchet Inception Distance; $SVDD$ (resp. $SVDD_{cont}$), relative difference of the radii (14) for all the variables (resp. the continuous variables); KL , Kullback-Leibler divergence; GAN weighted (resp. GAN), vanilla GAN with (resp. without) sampling weights.

References

- Annoye, H., Beretta, A., Heuchenne, C., 2024. Statistical matching using kernel canonical correlation analysis and super-organizing map. *Expert Systems with Applications* , 123134doi:10.1016/j.eswa.2023.123134.
- Arjovsky, M., Chintala, S., Bottou, L., 2017. Wasserstein generative adversarial networks, in: *International Conference on Machine Learning*, PMLR. pp. 214–223.
- Bae, H., Jung, D., Choi, H.S., Yoon, S., 2019. Anomigan: Generative adversarial networks for anonymizing private medical data, in: *Pacific Symposium on Biocomputing 2020*, World Scientific. pp. 563–574.
- Baowaly, M.K., Lin, C.C., Liu, C.L., Chen, K.T., 2019. Synthesizing electronic health records using improved generative adversarial networks. *Journal of the American Medical Informatics Association* 26, 228–241.
- Bellemare, M.G., Danihelka, I., Dabney, W., Mohamed, S., Lakshminarayanan, B., Hoyer, S., Munos, R., 2017. The cramer distance as a solution to biased wasserstein gradients. *arXiv preprint arXiv:1705.10743* arXiv:1705.10743.
- Boltz, S., Debreuve, E., Barlaud, M., 2007. knn-based high-dimensional kullback-leibler distance for tracking, in: *Eighth International Workshop on Image Analysis for Multimedia Interactive Services (WIAMIS'07)*, IEEE. pp. 16–16.
- Borji, A., 2019. Pros and cons of gan evaluation measures. *Computer Vision and Image Understanding* 179, 41–65.

- Borji, A., 2021. Pros and cons of gan evaluation measures: New developments. arXiv preprint arXiv:2103.09396 **arXiv:2103.09396**.
- Camino, R., Hammerschmidt, C., State, R., 2018. Generating multi-categorical samples with generative adversarial networks. arXiv preprint arXiv:1807.01202 **arXiv:1807.01202**.
- Cheng, J., Yang, Y., Tang, X., Xiong, N., Zhang, Y., Lei, F., 2020. Generative adversarial networks: A literature review. KSII Transactions on Internet & Information Systems 14.
- DataCebo, 2024. Synthetic data vault (sdv). <https://github.com/sdv-dev/SDV>.
- Engelmann, J., Lessmann, S., 2020. Conditional wasserstein gan-based oversampling of tabular data for imbalanced learning. arXiv preprint arXiv:2008.09202 **arXiv:2008.09202**.
- Fang, J., Li, A., Jiang, Q., 2019. Gdagan: An anonymization method for graph data publishing using generative adversarial network, in: 2019 6th International Conference on Information Science and Control Engineering (ICISCE), IEEE. pp. 309–313.
- Feutry, C., Piantanida, P., Bengio, Y., Duhamel, P., 2018. Learning anonymized representations with adversarial neural networks. arXiv preprint arXiv:1802.09386 **arXiv:1802.09386**.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y., 2020. Generative adversarial networks. Communications of the ACM 63, 139–144.

- Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y., 2014. Generative adversarial networks. arXiv preprint arXiv:1406.2661 **arXiv:1406.2661**.
- Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., Courville, A., 2017. Improved training of wasserstein gans. Advances in neural information processing systems 30.
- Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S., 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. Advances in neural information processing systems 30.
- Jamotton, C., Hainaut, D., 2023. Variational autoencoder for synthetic insurance data. Technical Report. Université catholique de Louvain, ISBA.
- Kingma, D.P., Ba, J., 2017. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 **arXiv:1412.6980**.
- Kingma, D.P., Welling, M., 2022. Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114 **arXiv:1312.6114**.
- Kynkäänniemi, T., Karras, T., Laine, S., Lehtinen, J., Aila, T., 2019. Improved precision and recall metric for assessing generative models. Advances in neural information processing systems 32.
- Lampert, C.H., et al., 2009. Kernel methods in computer vision. Foundations and Trends® in Computer Graphics and Vision 4, 193–285.
- Lu, S., Yu, L., Feng, S., Zhu, Y., Zhang, W., 2019. Cot: Cooperative training

- for generative modeling of discrete data, in: International Conference on Machine Learning, PMLR. pp. 4164–4172.
- Maas, A.L., Hannun, A.Y., Ng, A.Y.T., et al., 2013. Rectifier nonlinearities improve neural network acoustic models, in: Proceedings of the International Conference on Machine Learning, Atlanta, GA. p. 3.
- Metz, L., Poole, B., Pfau, D., Sohl-Dickstein, J., 2017. Unrolled generative adversarial networks. arXiv preprint [arXiv:1611.02163](https://arxiv.org/abs/1611.02163) `arXiv:1611.02163`.
- Mottini, A., Lheritier, A., Acuna-Agost, R., 2018. Airline passenger name record generation using generative adversarial networks. arXiv preprint [arXiv:1807.06657](https://arxiv.org/abs/1807.06657) `arXiv:1807.06657`.
- Naeem, M.F., Oh, S.J., Uh, Y., Choi, Y., Yoo, J., 2020. Reliable fidelity and diversity metrics for generative models, in: International conference on machine learning, PMLR. pp. 7176–7185.
- Park, N., Mohammadi, M., Gorde, K., Jajodia, S., Park, H., Kim, Y., 2018. Data synthesis based on generative adversarial networks. Proceedings of the VLDB Endowment 11, 1071–1083. URL: <http://dx.doi.org/10.14778/3231751.3231757>, doi:10.14778/3231751.3231757.
- Pascanu, R., Mikolov, T., Bengio, Y., 2013. On the difficulty of training recurrent neural networks, in: International conference on machine learning, Pmlr. pp. 1310–1318.
- Schölkopf, B., Platt, J.C., Shawe-Taylor, J., Smola, A.J., Williamson, R.C., 2001. Estimating the support of a high-dimensional distribution. Neural computation 13, 1443–1471.

- Schuhmacher, D., Bähre, B., Gottschlich, C., Hartmann, V., Heinemann, F., Schmitzer, B., 2024. transport: Computation of Optimal Transport Plans and Wasserstein Distances. URL: <https://cran.r-project.org/package=transport>. r package version 0.15-0.
- Tax, D.M.J., Duin, R.P.W., 2004. Support vector data description. *Machine learning* 54, 45–66.
- Vega-Márquez, B., Rubio-Escudero, C., Riquelme, J.C., Nepomuceno-Chamorro, I., 2020. Creation of synthetic data with conditional generative adversarial networks, in: Martínez Álvarez, F., Troncoso Lora, A., Sáez Muñoz, J.A., Quintián, H., Corchado, E. (Eds.), 14th International Conference on Soft Computing Models in Industrial and Environmental Applications (SOCO 2019), Springer International Publishing, Cham. pp. 231–240.
- Villani, C., 2009. Optimal transport: old and new. volume 338 of *Grundlehren der mathematischen Wissenschaften A Series of Comprehensive Studies in Mathematics*. Springer, Heidelberg.
- Walia, M., Tierney, B., McKeever, S., 2020. Synthesising tabular data using wasserstein conditional gans with gradient penalty (wcgan-gp)., in: AICS, pp. 325–336.
- Xu, L., Skoularidou, M., Cuesta-Infante, A., Veeramachaneni, K., 2019. Modeling tabular data using conditional gan. *Advances in neural information processing systems* 32.

- Xu, L., Veeramachaneni, K., 2018. Synthesizing tabular data using generative adversarial networks. arXiv preprint arXiv:1811.11264 **arXiv:1811.11264**.
- Zhang, Z., Li, M., Yu, J., 2018. On the convergence and mode collapse of gan, in: SIGGRAPH Asia 2018 Technical Briefs, pp. 1–4.