
Using Generative Adversarial Networks to Synthesize Artificial Financial Datasets

Dmitry Efimov
Credit and Fraud Risk
American Express
200 Vesey St
New York, NY 10281
dmitry.efimov@aexp.com

Di Xu
Credit and Fraud Risk
American Express
200 Vesey St
New York, NY 10281
di.xu@aexp.com

Luyang Kong
Credit and Fraud Risk
American Express
200 Vesey St
New York, NY 10281
luyang.kong1@aexp.com

Alexey Nefedov
Credit and Fraud Risk
American Express
200 Vesey St
New York, NY 10281
alexey.nefedov@aexp.com

Archana Anandakrishnan
Credit and Fraud Risk
American Express
200 Vesey St
New York, NY 10281
archana.anandakrishnan@aexp.com

Abstract

Generative Adversarial Networks (GANs) became very popular for generation of realistically looking images. In this paper, we propose to use GANs to synthesize artificial financial data for research and benchmarking purposes. We test this approach on three American Express datasets, and show that properly trained GANs can replicate these datasets with high fidelity. For our experiments, we define a novel type of GAN, and suggest methods for data preprocessing that allow good training and testing performance of GANs. We also discuss methods for evaluating the quality of generated data, and their comparison with the original real data.

1 Introduction

Machine learning (ML) algorithms are ubiquitous in many practical domains, including the banking and financial industry, where some of their core applications are assessing creditworthiness of customers, offering customers optimal financial products, and identifying fraud. Measuring performance of these algorithms is critical to understanding their strengths and weaknesses. Comprehensive comparisons of algorithms require high quality datasets that can be used as standard benchmarks.

An integral part of the modern machine learning field is a corpus of publicly shared, high quality datasets from numerous domains, which are used to benchmark and validate ML algorithms developed by various researchers, academic groups and companies (see, for example, MNIST [1], ImageNet [2], Kaggle [3], University of Irvine Machine Learning Repository [4], etc. [5]). These datasets also play

an important role in advancing state of the field by greatly facilitating testing and development of novel ideas and methods.

In the long list of publicly available datasets, one can notice the shortage of datasets originating from the banking and financial industry, especially datasets associated with credit and fraud risk management operations. One of the key reasons for this is that making such data publicly available would be a violation of customers' privacy and trust.

For several available financial datasets (e.g. [6], [7]), this issue is resolved by some pre-treatment of data, for example by reducing data to principal components, or by taking only a small sample of data [8], [9]. However, these transformed or subsampled datasets may fail to capture some unique properties of financial data. For example, transactional data is usually large, highly structured, contains a wide range of both numerical (continuous and discrete) and categorical variables, variables with very skewed distributions, missing values, etc. Practical algorithms should be able to deal with such variables in a robust and efficient manner, and be scalable to the large size of this data. To develop state of the art ML methods, including methods for anomaly detection and model interpretation, ML researchers and practitioners need to have access to data that is as close to the real one as possible.

In this paper, we suggest using Generative Adversarial Networks (GANs) as a means to create synthetic, "artificial" financial data. We describe experiments with three American Express datasets used for risk modeling. We show that properly trained GANs can replicate these datasets with high fidelity. Specifically, we show that synthesized data follows the same distribution as the original data, and that ML models trained on synthesized data have the same performance as those trained on the original data. In our experiments, we use a novel type of GAN architecture combining conditional GAN and DRAGAN, which gives us better training convergence and testing performance.

Since GAN-generated data does not originate from real customers, it could be made public as a benchmark dataset for financial applications. Customers' personal data and privacy would be protected using this approach.

This paper is structured as follows. Section 2 gives a brief introduction to GANs and introduces a novel type of GAN that we used in our experiments. In Section 3, we discuss data preprocessing methods that can help to improve training and testing performance of GAN. In Section 4, we discuss methods for evaluation of the generated datasets and illustrate that any improvements made using the generated datasets scale and generalize to the real dataset. We conclude by discussing future steps in our research and other potential applications of GANs in the financial industry.

2 Data Generation with Generative Adversarial Networks

Generative Adversarial Network (GAN) is a type of neural network comprised of two connected networks, called generator and discriminator, which compete with each other during the training phase ([10]). The objective of the discriminator is to distinguish samples coming from a given training set ("real data") from samples created by generator ("synthesized data"). The error of discriminator is fed to the generator, which learns to produce samples that are increasingly difficult for discriminator to distinguish from real ones.

When training of GAN is complete, its generator can be used to synthesize data reproducing original, real data that was used for training. Since its introduction in 2014, GANs have mostly been used to generate realistically looking images in image classification and computer vision applications [11], [12], [13]. However, there has been an increasing number of GAN applications to non-image data, where the goal of GANs can be, for example, enhancing real training data with synthetic samples [14], [15], [16].

In our experiments, the primary goal of using GANs was to generate synthetic data that replicate distribution of original real data, and allow us to build ML models that perform on par with ML models built using original real data. We used three American Express datasets from three different use cases in risk modeling. Basic details of these datasets are given in Table 1. The datasets were chosen to represent some variety of data features, feature distributions, and type of classification problems. We omit further details about our use cases; however, they are not relevant in the scope of this paper.

Dataset	Number of Features (Numerical, Categorical)	Size of Data Sample	Target Variable
Dataset A	22 (21, 1)	120,990	Continuous
Dataset B	119 (113, 6)	2,197,762	Binary
Dataset C	471 (407, 64)	2,028,106	Binary

Table 1: Details of three datasets used in our experiments with GANs.

In the following paragraphs, we are going to give basic definitions related to GANs, and introduce a new type of GAN that we used in our experiments.

In a classic (vanilla) GAN, generator network G takes a random noise vector \mathbf{z} as input, and produces fake sample $G_\varphi(\mathbf{z})$, where φ is a set of generator’s parameters (Fig. 1 (left)). Discriminator network D takes in a sample \mathbf{x} , which can be either real or fake, and computes probability $D_\theta(\mathbf{x})$, where θ is a set of discriminator’s parameters. The goal of discriminator is to distinguish between real and fake (generated) samples:

$$D_\theta(\mathbf{x}) = \begin{cases} 1, & \text{if } \mathbf{x} \text{ is a real sample,} \\ 0, & \text{if } \mathbf{x} \text{ is a generated sample} \end{cases} \quad (1)$$

GAN training is performed by sequential minimization of discriminator’s loss function $J^{(D)}(\varphi, \theta)$ with respect to parameters θ , where $J^{(D)}(\varphi, \theta)$ is defined as

$$J^{(D)}(\varphi, \theta) = -\mathbb{E}_{\mathbf{x}} \log D_\theta(\mathbf{x}) - \mathbb{E}_{\mathbf{z}} \log(1 - D_\theta(G_\varphi(\mathbf{z}))), \quad (2)$$

and minimization of generator’s loss function $J^{(G)}(\varphi, \theta)$ with respect to parameters φ , where $J^{(G)}(\varphi, \theta)$ is defined as

$$J^{(G)}(\varphi, \theta) = -\mathbb{E}_{\mathbf{z}} \log D_\theta(G_\varphi(\mathbf{z})). \quad (3)$$

Training of vanilla GAN may be unstable due to gradient exploding or gradient vanishing effects [17]. These effects can be even stronger in applications with non-image, structured data.

To overcome this problem, Deep Regret Analytic Generative Adversarial Networks (DRAGANs) were introduced by Kodali et al. in [18]. In DRAGAN, discriminator’s loss is modified by adding a regularization term that prevents sharp gradients and improves convergence:

$$J_*^{(D)}(\varphi, \theta) = J^{(D)}(\varphi, \theta) + \lambda \cdot \mathbb{E}_{\mathbf{x}, \delta \sim N_d(0, cI)} [\|\nabla_{\mathbf{x}} D_\theta(\mathbf{x} + \delta)\| - k]^2, \quad (4)$$

where λ , c and k are hyperparameters. Fig. 2 shows the difference in convergence of regular GAN and DRAGAN on our data; the convergence of DRAGAN is much more stable.

Conditional GANs (CGANs) were introduced by Mirza et al. in [19] to better handle categorical features in training data. In CGAN, generator’s input contains two parts: a random noise vector \mathbf{z} and a dummy features vector \mathbf{y} generated from categorical features (Fig. 1 (right)). Generator’s output consists of vector \mathbf{x} of numerical features only. Before feeding to the discriminator, vectors \mathbf{y} and \mathbf{x} are concatenated and equations (2) and (3) are modified as

$$J^{(G)}(\varphi, \theta) = -\mathbb{E}_{\mathbf{z}} \log D_\theta(G_\varphi(\mathbf{z}, \mathbf{y}), \mathbf{y}) \quad (5)$$

and

$$J^{(D)}(\varphi, \theta) = -\mathbb{E}_{\mathbf{x}, \mathbf{y}} \log D_\theta(\mathbf{x}, \mathbf{y}) - \mathbb{E}_{\mathbf{z}} \log(1 - D_\theta(G_\varphi(\mathbf{z}, \mathbf{y}), \mathbf{y})) \quad (6)$$

accordingly.

Since our data contained categorical features, we decided to use a combination of DRAGAN with CGAN architecture, which we called conditional DRAGAN (CDRAGAN). In CDRAGAN, we add regularization term from (4) to the discriminator loss (6) in order to avoid exploding gradients. Just like in non-conditional case, convergence of CDRAGAN has better stability than convergence of CGAN (Fig. 3).

3 Preprocessing Data for GAN Training

Preprocessing of data used in GAN training is one of the crucial steps that determines success or failure of building a GAN that can accurately reproduce original data. In the domain of computer





Figure 1: Comparing architectures of GAN (left) and conditional GAN (right)



(a) Vanilla GAN

(b) DRAGAN

Figure 2: Comparing convergence of vanilla GAN and DRAGAN

vision, where GANs have been used the most since their introduction, all data features (image pixel values) are numerical, continuous, with similar range and distribution. In many other domains, however, data does not have these nice properties. In particular, financial data is usually comprised of features that have

- different types (numerical continuous, numerical discrete, categorical)
- different ranges
- different distributions (including very skewed ones, where the most frequent value occurs in more than 90% of samples)
- missing values (including variables with up to 90% of missing values)
- special values used to denote missing values, cap feature range on the left or on the right, etc.

These data properties pose unique challenges for GAN training, which should be addressed to ensure good training and testing performance of GAN.

For our datasets, introduced in Section 2, we found that the following preprocessing steps allowed good training and subsequently good testing performance of GANs.

1. One-hot encoding for categorical features.
2. Missing value indicator feature: a new feature that is equal to one in samples where the original feature has missing value, and zero – in all other samples.
3. Box-Cox transformation [20].
4. Standard scaling or min-max scaling.
5. Imputing missing values.

4 Evaluating data generated by GANs

When GANs are used to generate images, the quality of produced samples can be easily evaluated by visual observation: humans can easily determine if generated images look similar to the images from the training set, and are overall of good quality. For non-image data, there are no commonly accepted techniques for evaluating the quality of generated data.

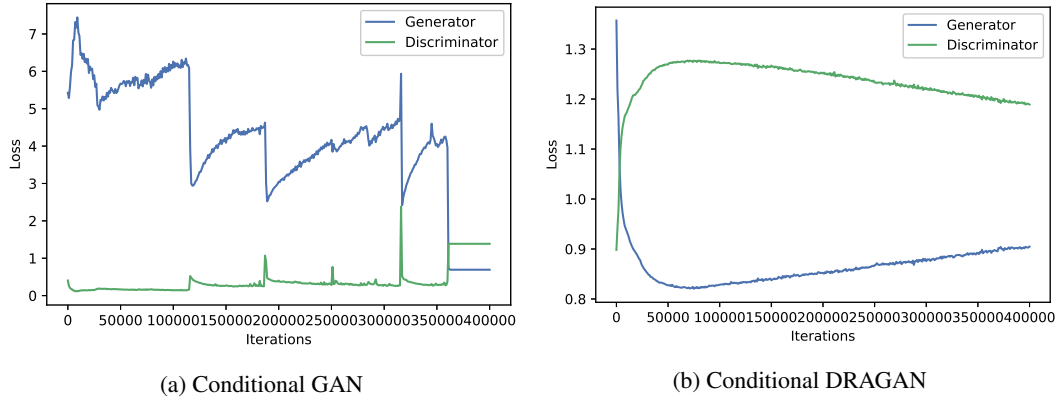


Figure 3: Comparing convergence of conditional GAN and conditional DRAGAN

We can start with checking whether histograms of generated features match those of real features. However, matching distributions of individual features and target variable do not guarantee that all existing interactions (relationships) *between features*, as well as *between features and target variable* were also replicated by the generator. Therefore, it is necessary to evaluate and compare overall distribution of generated and real data. Since relationship between features and target variable is the most important one for developing ML models, we can perform a separate test of this relationship by comparing supervised ML models trained on generated and real data. ML models with similar performance would be indicative of good replication of dependencies between features and target variable.

To summarize, we propose that a good generator should produce data that satisfy the following three criteria:

1. Distributions of individual features in generated data match those in real data
2. Overall distributions of generated and real data match each other
3. Relationship between features and the target variable in real data is replicated in generated data

In the following subsections, we will show how we performed above tests for the data generated by our GANs.

4.1 DataQC tool

DataQC [21] is an internal automated tool developed at American Express for data quality assessment. It allows users to evaluate similarities and differences between two provided datasets. The tool performs a comprehensive set of data quality tests to quickly highlight how one dataset is different from another. These tests include comparison of feature means, rates of missing values, uni- and multivariate distributions, and extreme values. The tool produces detailed findings and quantitative scores for all the tests that it runs.

We used DataQC to evaluate similarity between GAN-generated and real data. In particular, we used it to compare distributions of individual features in generated and real data. As an example, Fig. 4 shows histograms of four selected features in real and generated data for a GAN trained on Dataset C. For all other features we observed similar, very close match between real and generated distributions. This was the case for all three Datasets. We noted that our GANs were able to reproduce discrete distributions just as good as continuous ones, and that for some continuous variables GANs tended to produce slightly smoothed versions of their distributions. We found that smoothing of distributions can be reduced by increasing the number of hidden layers in the generator, or by increasing the number of training iterations.




Figure 4: Examples of feature histograms for real (blue) and generated (green) data. Experiment with Dataset C.

Dataset	Original data	Synthetic data
Dataset A	0.66	0.63
Dataset B	0.80	0.78
Dataset C	0.89	0.86

Table 2: The AUC scores of supervised model test for benchmark datasets.

4.2 Visualization with t-SNE algorithm

To compare overall distribution of real and generated data, we visualized the data using t-SNE algorithm [22]. t-SNE is a transductive algorithm: the model produced by this algorithm cannot be applied to out-of-sample data that was not used to build the model. Therefore, we combined the real and generated data together, obtained t-SNE representation for the combined data, and then split it into two parts corresponding to real and generated data. Fig. 5 shows t-SNE graphs for our experiment with Dataset C. It can be seen in the figure that t-SNE graph for generated data closely matches t-SNE graph for real data, reproducing most of its clusters and gaps with pretty good accuracy in shape and density.




Figure 5: tSNE visualization of real and generated data. Experiment with Dataset C.

4.3 Supervised models trained on real and generated data

To test relationship between features and target variable in real and generated data, we compared two supervised ML models: one trained on real data, and another one trained on data produced by GAN (containing the same number of samples as the real data). To train the first model, we used the actual target variable from the real data. To train the second model, we used target variable generated by GAN along with other features. Both ML models used identical hyperparameters during the training phase. Diagram 1 illustrates the idea of this approach.

Trained supervised models were validated on an out-of-sample real data with actual values of target variable. Area under the ROC curve (AUC, [23]) was used to compare performance of the two models. AUC scores were calculated for ground truth target values and predictions obtained from the models trained on real and generated datasets; the final AUC scores are presented in Table. 2. We found

these scores to be close enough to assume that our GAN model replicated the relationship between the target variables and the features with good accuracy.




Diagram 1. Supervised model test overview.

5 Results and conclusions

The final results are presented in Table 2. The models trained on synthetic data show slightly worse performances on out-of-sample validation data but the scores are very close. Generated data can be considered as good approximation of the original data though the performance of the models trained on the original data is still better than on synthetic data. Overall, we conclude that GANs can learn and accurately reproduce intricate features distribution and relationships between features of real modeling data.

We will continue our research in the following two directions. First, we would like to explore possible causes for lower performance of the models trained on synthetic data. We will also investigate different preprocessing approaches and GAN architectures in order to compare, and see if the baseline performance can be achieved or improved on pure GAN-generated data. We will also work on theoretical justification of the proposed approach, e.g. addressing the following question: given two datasets original and synthetic, what are the sufficient criteria for synthetic data to produce the model with the same performance on original data?

References

- [1] LeCun, Y. and Cortes, C. The MNIST database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>.
- [2] Deng, J et al. (2009) ImageNet: A Large-Scale Hierarchical Image Database. *CVPR09*
- [3] <http://kaggle.com/datasets>.
- [4] Dua, D. and Karra Taniskidou, E. (2017). UCI Machine Learning Repository *archive.ics.uci.edu/ml*, University of California, Irvine, School of Information and Computer Sciences.
- [5] Wikipedia contributors. List of datasets for machine learning research. *Wikipedia, The Free Encyclopedia*.
- [6] <https://www.kaggle.com/c/GiveMeSomeCredit/data>.
- [7] <https://www.kaggle.com/mlg-ulb/creditcardfraud>.
- [8] Dal Pozzolo, Andrea, et al. (2015) Calibrating probability with undersampling for unbalanced classification. *Computational Intelligence, IEEE Symposium Series*.
- [9] Yeh, I-Cheng, and Che-hui Lien. (2009) The comparisons of data mining techniques for the predictive accuracy of probability of default of credit card clients. *Expert Systems with Applications*. **36.2**: 2473-2480.
- [10] Goodfellow, I.J. et al. (2014) Generative adversarial nets. *NIPS'14 Proceedings of the 27th International Conference on Neural Information Processing Systems*, Vol.2:2672-2680

- [11] Perez, L. & Wang, J. (2017) The Effectiveness of Data Augmentation in Image Classification using Deep Learning. *arXiv.org*, arXiv:1712.04621
- [12] Shin, H.-C. et al. (2018) Medical Image Synthesis for Data Augmentation and Anonymization Using Generative Adversarial Networks. *arXiv.org*, arXiv:1807.10225
- [13] Zhu, X. et al. (2017) Data Augmentation in Emotion Classification Using Generative Adversarial Networks. *arXiv.org*, arXiv:1711.00648
- [14] Zheng, P. et al. (2018) One-Class Adversarial Nets for Fraud Detection. *CoRR*, abs/1803.01798
- [15] Fiore, U. et al. (2017) Using generative adversarial networks for improving classification effectiveness in credit card fraud detection. *Information Sciences*, doi:https://doi.org/10.1016/j.ins.2017.12.030
- [16] Kumar, A. et al. (2018) eCommerceGAN : A Generative Adversarial Network for E-commerce. *CoRR*, abs/1801.03244
- [17] Gulrajani, I. et al. (2017) Improved training of Wasserstein GANs. *arXiv.org*, arXiv:1704.00028
- [18] Kodali, N. et al. (2017) On Convergence and Stability of GANs. *arXiv.org*, arXiv:1705.07215
- [19] Mirza, M. & Osindero, S. (2014) Conditional generative adversarial nets. *arXiv.org*, arXiv:1411.1784
- [20] Box, G. and Cox D. (1964) An analysis of transformations. *Journal of the Royal Statistical Society, Series B*, 26 (2): 211–252. JSTOR 2984418. MR 0192611
- [21] Love, D. et al. (2018) An Automated System for Data Attribute Anomaly Detection. *Proceedings of the KDD 2017: Workshop on Anomaly Detection in Finance*, in PMLR 71:95-101
- [22] Maaten, L. & Hinton, G. (2008) Visualizing Data using t-SNE. *Journal of Machine Learning Research*, 9(Nov):2579-2605
- [23] Powers, D.M.W. (2011) Evaluation: From Precision, Recall and F-Measure to ROC, Informedness, Markedness and Correlation. *Journal of Machine Learning Technologies*, 2 (1): 37–63