

Tabular GANs for uneven distribution

Insaf Ashrapov¹[0000–0003–4938–0430]

Moscow Institute of Physics and Technology, Moscow, Russia intoffmipt.ru
<https://eng.mipt.ru/>

Abstract. GANs are well known for success in the realistic image generation. However, they can be applied in tabular data generation as well. We will review and examine some recent papers about tabular GANs in action. We will generate data to make train distribution bring closer to the test. Then compare model performance trained on the initial train dataset, with trained on the train with GAN generated data, also we train the model by sampling train by adversarial training. We show that using GAN might be an option in case of uneven data distribution between train and test data.

Keywords: Deep learning, Tabular data, Generative adversarial network

0.1 What is GAN [2]

“GAN composes of two deep networks: the **generator** and the **discriminator**”[4].

Both of them simultaneously trained. Generally, the model structure and training process presented like on Figure 1.

The task for the generator is to generate samples, which won't be distinguished from real samples by the discriminator. We won't give much detail here. You can read the medium post and the original paper [2]. Recent architectures such as StyleGAN 2 can produce outstanding photo-realistic images, examples in Figure 2.

Problems While face generation seems to be not a problem anymore, there are plenty of issues we need to resolve:

- **Training speed.** For training StyleGAN 2 you need 1 week and DGX-1 (8x NVIDIA Tesla V100).
- **Image quality** in specific domains. The state-of-the-art network still fails on other tasks. Such examples are shown in Figure 3.

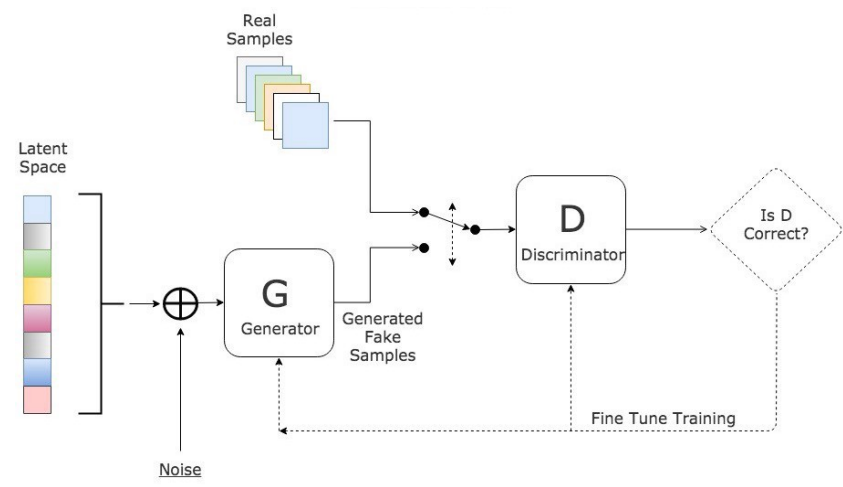


Fig. 1. GAN training pipeline [4]



Fig. 2. Hand-picked examples of human faces generated by StyleGAN 2 [6]




Fig. 3. Hand-picked bad examples of cars and cats generated by StyleGAN 2 [6]

0.2 Tabular GANs

Even cats and dogs generation seem heavy tasks for GANs because of not trivial data distribution and high object type variety. Besides such domains, the image background becomes important, which GANs usually fail to generate. Therefore, we've been wondering what GANs can achieve in tabular data. Unfortunately, there aren't many articles. The next two articles appear to be the most promising.

TGAN: Synthesizing Tabular Data using Generative Adversarial Networks [13] They raise several problems, why generating tabular data has its own challenges: the various data types (int, decimals, categories, time, text) different shapes of distribution (multi-modal, long tail, Non-Gaussian. . .) sparse one-hot-encoded vectors and highly imbalanced categorical columns.

Task formalizing Let us say table T contains n_c continuous variables and n_d discrete (categorical) variables and each row is C vector. These variables have an unknown joint distribution P . Each row is independently sampled from P . The object is to train a generative model M . M should generate new a synthetic table T_{synth} with the distribution similar to P . A machine learning model learned on T_{synth} should achieve a similar accuracy on a real test table T_{test} , as would a model trained on T .

Preprocessing numerical variables Neural networks can effectively generate values with a distribution centered over $(-1, 1)$ using \tanh . However, they show that nets fail to generate suitable data with multi-modal data. Thus they cluster a numerical variable by using and training a Gaussian Mixture Model (GMM) [9] with m ($m=5$) components for each of C .

Finally, GMM is used to normalize C to get V . Besides, they compute the probability of C coming from each of the m Gaussian distribution as a vector U .

Preprocessing categorical variables Due to usually low cardinality, they found the probability distribution can be generated directly using softmax. But it necessary to convert categorical variables to one-hot-encoding representation with noise to binary variables After prepossessing, they convert T with $n_c + n_d$ columns to V , U , D vectors. This vector is the output of the generator and the input for the discriminator in GAN. GAN does not have access to GMM parameters.

Generator They generate a numerical variable in 2 steps. First, generate the value scalar V , then generate the cluster vector U eventually applying \tanh . Categorical features are generated as a probability distribution over all possible labels with softmax. To generate the desired row LSTM [3] with attention mechanism is used. Input for LSTM in each step is random variable z , weighted context vector with previous hidden and embedding vector.

4 Insaf Ashrapov

Discriminator Multi-Layer Perceptron (MLP) with LeakyReLU [10] and Batch-Norm [5] is used. The first layer used concatenated vectors (V , U , D) along with mini-batch diversity with feature vector from LSTM. The loss function is the KL divergence term of input variables with the sum ordinal log loss function.

The detailed model structure is shown in Figure 4

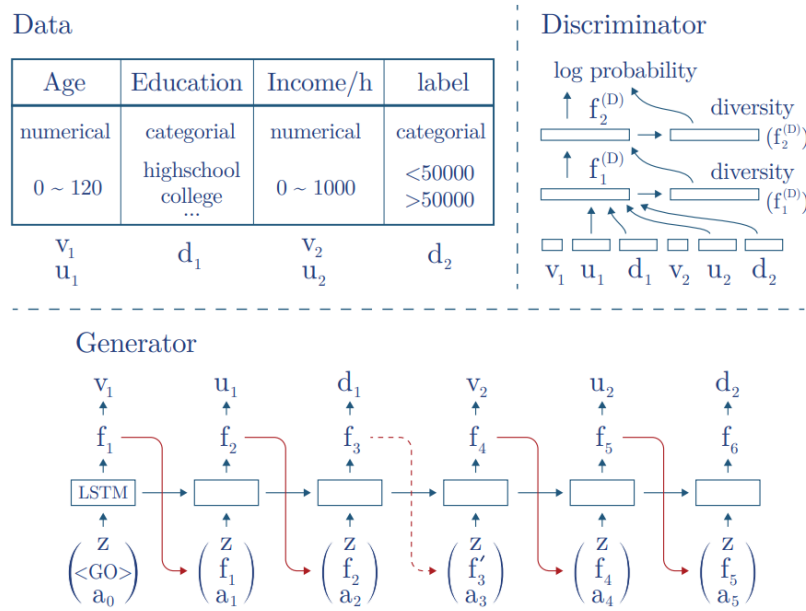


Fig. 4. Example of using TGAN to generate a simple census table. The generator generates T features one by one. The discriminator concatenates all features together. Then it uses Multi-Layer Perceptron (MLP) with LeakyReLU to distinguish real and fake data [13]

Results They evaluate the model on two datasets KDD99 and covtype. For some reason, they used weak models without boosting (xgboost, etc). Anyway, TGAN performs reasonably well and robust, outperforming bayesian networks. The average performance gap between real data and synthetic data is 5.7% [13].

Modeling Tabular Data using Conditional GAN (CTGAN) [12] The key improvements over previous TGAN are applying the mode-specific normalization to overcome the non-Gaussian and multimodal distribution. Then a conditional generator and training-by-sampling to deal with the imbalanced discrete columns.

Task formalizing The initial data remains the same as it was in TGAN. However, they solve different problems.

- **Likelihood of fitness.** Do columns in T_{syn} follow the same joint distribution as T_{train}
- **Machine learning efficacy.** When training the model to predict one column using other columns as features, can such model learned from T_{syn} achieve similar performance on T_{test} , as a model learned on T_{train}

Preprocessing Preprocessing for discrete columns keeps the same. For continuous variables, a variational Gaussian mixture model (VGM) is used. It first estimates the number of modes m and then fits a Gaussian mixture. After we normalize initial vector C almost the same as it was in TGAN, but the value is normalized within each mode. The mode is represented as one-hot vector β ($[0, 0, \dots, 1, 0]$). α is the normalized value of C . Example is shown in Figure 5.




Fig. 5. An example of mode-specific normalization[12].

As a result, we get our initial row represented as the concatenation of one-hot' ed discrete columns with representation discussed above of continues variables:

Training

The final solution consists of three key elements, namely: the conditional vector, the generator loss, and the training-by-sampling method as is shown in Figure 6 [12].

Conditional vector Represents concatenated one-hot vectors of all discrete columns but with the specification of only one category, which was selected. “For instance, for two discrete columns, $D1 = \{1, 2, 3\}$ and $D2 = \{1, 2\}$, the condition ($D2 = 1$) is expressed by the mask vectors $m1 = [0, 0, 0]$ and $m2 = [1, 0]$; so $cond = [0, 0, 0, 1, 0]$ ”.

6 Insaf Ashrapov

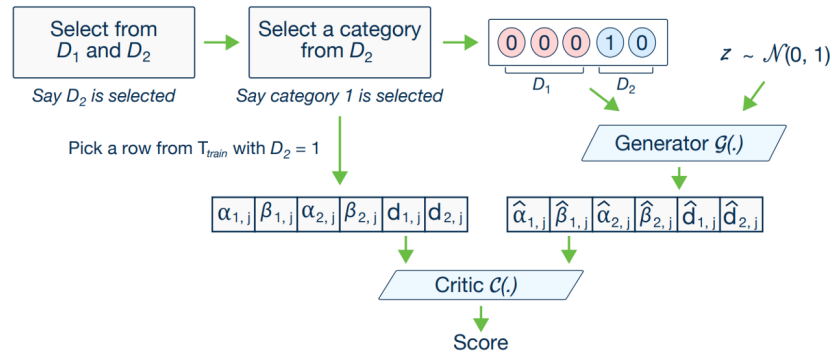


Fig. 6. CTGAN model. The conditional generator can generate synthetic rows conditioned on one of the discrete columns. With training-by-sampling, the cond and training data are sampled according to the log-frequency of each category, thus CTGAN can evenly explore all possible discrete values [12].

Generator loss

During training, the conditional generator is free to produce any set of one-hot discrete vectors.

But they enforce the conditional generator to produce d_i (generated discrete one-hot column) = m_i (mask vector) is to penalize its loss by adding the cross-entropy between them, averaged over all the instances of the batch.

Training-by-sampling

Specifically, the goal is to resample efficiently in a way that all the categories from discrete attributes are sampled evenly during the training process, as a result, to get real data distribution during the test.

In other words, the output produced by the conditional generator must be assessed by the critic, which estimates the distance between the learned conditional distribution $P_G(\text{row}|\text{cond})$ and the conditional distribution on real data $P(\text{row}|\text{cond})$.

The sampling of real training data and the construction of cond vector should comply to help critics estimate the distance.

Properly sample the cond vector and training data can help the model evenly explore all possible values in discrete columns. The model structure is given below, as opposite to TGAN, there is no LSTM layer. Trained with WGAN loss with gradient penalty.

Also, they propose a model based on Variational autoencoder (VAE), but it out of the scope of this article.

Results The proposed network CTGAN and TVAE outperform other methods. As they say, TVAE outperforms CTGAN in several cases, but GANs do have several favorable attributes. The generator in GANs does not have access to real data during the entire training process, unlike TVAE. Detailed results are shown in Figure 7.

Method	GM Sim.		BN Sim.		Real	
	\mathcal{L}_{syn}	\mathcal{L}_{test}	\mathcal{L}_{syn}	\mathcal{L}_{test}	clf	reg
Identity	-2.61	-2.61	-9.33	-9.36	0.743	0.14
CLBN	-3.06	-7.31	-10.66	-9.92	0.382	-6.28
PrivBN	-3.38	-12.42	-12.97	-10.90	0.225	-4.49
MedGAN	-7.27	-60.03	-11.14	-12.15	0.137	-8.80
VEEGAN	-10.06	-4.22	-15.40	-13.86	0.143	-6.5e6
TableGAN	-8.24	-4.12	-11.84	-10.47	0.162	-3.09
TVAE	-2.65	-5.42	-6.76	-9.59	0.519	-0.20
CTGAN	-5.72	-3.40	-11.67	-10.60	0.469	-0.43

Fig. 7. Benchmark results over three sets of experiments, namely Gaussian mixture simulated data (GM Sim.), Bayesian network simulated data (BN Sim.), and real data. They report the average of each metric. For real datasets (f1, etc). [12]

Besides, they published source code on GitHub [11], which with slight modification will be used further in the article.

1 Task formalization

Let say we have T_{train} and T_{test} (train and test set respectively). We need to train the model on T_{train} and make predictions on T_{test} . However, we will increase the train by generating new data by GAN, somehow similar to T_{test} , without using ground truth labels of it.

Experiment design The experiment design is shown in Figure 8. Let say we have T_{train} and T_{test} (train and test set respectively). The size of T_{train} is smaller and might have different data distribution. First of all, we train CTGAN on T_{train} with ground truth labels (step 1), then generate additional data T_{synth} (step 2). Secondly, we train boosting in an adversarial way on concatenated T_{train} and T_{synth} (target set to 0) with T_{test} (target set to 1) (steps 3 & 4). The goal is to apply newly trained adversarial boosting to obtain rows more like T_{test} . Note — original ground truth labels aren't used for adversarial training. As a result, we take top rows from T_{train} and T_{synth} sorted by

8 Insaf Ashrapov

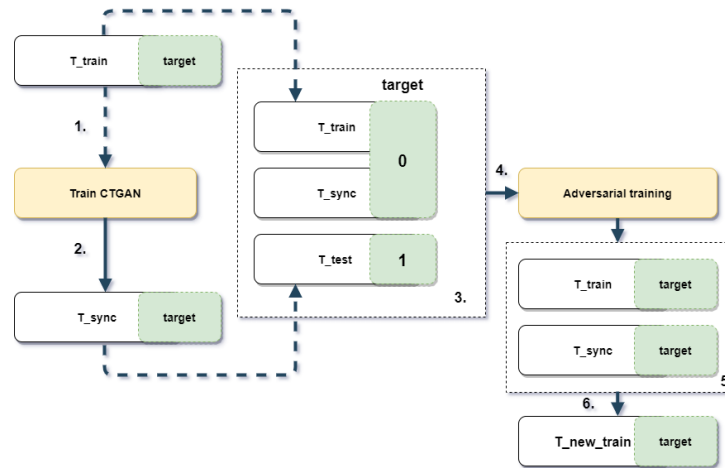


Fig. 8. Experiment design and workflow

correspondence to T_{test} (steps 5 & 6). Finally, train new boosting on them and check results on T_{test} .

Of course for the benchmark purposes we will test ordinal training without these tricks and another original pipeline but without CTGAN (in step 3 we won't use T_{sync}).

Code Experiment code and results are released as GitHub repo [1]. Pipeline and data preparation was based on Benchmarking Categorical Encoders' article and its repositories. We will follow almost the same pipeline, but for speed, only Single validation and Catboost encoder was chosen.

Datasets All datasets came from different domains. They have a different number of observations, several categorical and numerical features. The aim of all datasets is a binary classification. Preprocessing of datasets was simple: removed all time-based columns from datasets. The remaining columns were either categorical or numerical. In addition, while training results were sampled T_{train} — 5%, 10%, 25%, 50%, 75%. Some dataset characteristics such as the number of points, features, and categorical features are shown in Table 1.

From the first sight of view and in terms of metric and stability (std), GAN shows the worse results. However, sampling the initial train and then applying adversarial training we could obtain the best metric results and stability (sample original). To determine the best sampling strategy, ROC AUC scores of each dataset were scaled (min-max scale) and then averaged among datasets.

Name	Total points	Train points	Test points	Features	Categorical features
Telecom	7.0k	4.2k	2.8k	20	16
Adult income	48.8k	29.3k	19.5k	15	8
Employee	32.7k	19.6k	13.1k	10	9
Credit	307.5k	184.5k	123k	121	18
Mortgages	45.6k	27.4k	18.2k	20	9
Taxi	892.5k	535.5k	357k	8	5
Poverty - Predict household is poor	37.6k	22.5k	15.0k	41	38

Table 1. Datasets properties

1.1 Results

To determine the best validation strategy, we compared the top score of each dataset for each type of validation. As you can see from Table 2 - sampling with **gan** shows same results as **None** as baseline method, winning 2 out of 7 datasets. Same baseline.

dataset name	None	gan	sample original
credit	0.997	0.998	0.997
employee	0.986	0.966	0.972
mortgages	0.984	0.964	0.988
poverty_A	0.937	0.950	0.933
taxi	0.966	0.938	0.987
adult	0.995	0.967	0.998
telecom	0.995	0.868	0.992

Table 2. Different sampling results across the dataset, higher is better (100%-maximum per dataset ROC AUC)

sample type	mean	std
None	0.980	0.036
gan	0.969	0.06
sample original	0.981	0.032

Table 3. Different sampling results, higher is better for a mean (ROC AUC), lower is better for std (100%-maximum per dataset ROC AUC).

We can see that GAN outperformed other sampling types in 2 datasets. Whereas sampling from original outperformed other methods in 3 of 7 datasets. Of course, there isn't much difference. but these types of sampling might be an option. Detailed results in Table 3.

sample type	same target %	test score %
None	0	0.964
None	1	0.985
gan	0	0.966
gan	1	0.945
sample original	0	0.973
sample original	1	0.984

Table 4. Same target % is equal 1 then the target rate for train and test are different no more than 5%. Higher is better.

Let's define *same target %* is equal 1 then the target rate for train and test is different no more than 5%. So then we have almost the same target rate in train and test *None* and sample original are better. However, *gan* is starting performing noticeably better than target distribution changes as you can see in Table 4.

2 Acknowledgments

The author would like to thank Open Data Science community [7] for many valuable discussions and educational help in the growing field of machine and deep learning. Also, special big thanks to PAO Sberbank [8] for allowing solving such tasks and providing computational resources.

References

1. Ashrapov, I.: Github repo tabular gans for uneven distribution (2020), <https://github.com/Diyago/GAN-for-tabular-data>
2. Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial networks (2014)
3. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Comput.* **9**(8), 1735–1780 (Nov 1997). <https://doi.org/10.1162/neco.1997.9.8.1735>, <https://doi.org/10.1162/neco.1997.9.8.1735>
4. Hui, J.: Gan — what is generative adversarial networks gan (2018), https://medium.com/@jonathan_hui/gan-whats-generative-adversarial-networks-and-its-application-f39ed278ef09
5. Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift (2015)
6. Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., Aila, T.: Analyzing and improving the image quality of stylegan (2019)

7. ODS.ai: Open data science (2020), <https://ods.ai/>
8. Sberbank (2020), <https://www.sberbank.ru/>
9. Viroli, C., McLachlan, G.J.: Deep gaussian mixture models (2017)
10. Xu, B., Wang, N., Chen, T., Li, M.: Empirical evaluation of rectified activations in convolutional network (2015)
11. Xu, L., Skoularidou, M., Cuesta-Infante, A., Veeramachaneni, K.: Github repo ctgan (2019), <https://github.com/sdv-dev/CTGAN>
12. Xu, L., Skoularidou, M., Cuesta-Infante, A., Veeramachaneni, K.: Modeling tabular data using conditional gan (2019)
13. Xu, L., Veeramachaneni, K.: Synthesizing tabular data using generative adversarial networks (2018)