ESANN 2021 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Online event, 6-8 October 2021, i6doc.com publ., ISBN 978287587082-7.
Available from http://www.i6doc.com/en/.

Differentially Private Time Series Generation

Hiba Arnout[1,2] and Johanna Bronner[1] and Thomas Runkler[1,2]

1- Siemens AG - Corporate Technology
Otto-Hahn-Ring 6, 81739 Munich - Germany

2- Technical University of Munich - Department of Computer Science
Boltzmannstraße 3, 85748 Garching - Germany

Abstract. Privacy issues prevent data owner from improving Machine Learning (ML) performance as it makes external collaborations binding. To allow data sharing without confidentiality concerns, we propose in this work methods to generate time series in a privacy preserving manner. We combine the existing Generative Adversarial Networks (GAN) models for time series namely TimeGAN [1], ClaRe-GAN [2] and C-RNN-GAN [3] with differential privacy. This is achieved by changing their original discriminator with a private discriminator that relies on the differentially private stochastic gradient method (DPSGD) [4]. Our experiments show that the developed methods - in particular TimeGAN and ClaRe-GAN - outperform the existing and unique differentially private model for time series of RCGAN [5] in terms of privacy and accuracy.

1 Introduction

In many medical or industrial domains, the lack of data and privacy concerns prevent researchers from improving the efficiency of ML. In these cases, publishing synthetic privacy-preserving data that depict the behavior of the original dataset, could enlarge the scope of ML's applicability. There by, it will also preserve their privacy.

Let's consider a scenario, illustrated in Fig. 1, where data owners want to improve the performance of ML in a specific use case by collaborating with some external partners. For example, they want to find a better performing ML model for some medical data or a model that correctly predicts the state of a machine. In these cases, it will be enough to give the external partner some synthetic data with the same reactivity to any ML model. Thus, it will not reveal rare diseases that can be easily detected in the original dataset or some sensitive information about the machine parameters or its properties e.g., times when the machines where on/off... To this end, the data owner can use a privacy preserving version of GAN to generate new anonymous data and can share it with the external partner who will not have access to the original one.

GAN is a well-known technique to deal with the lack of data. It generates new synthetic data by sampling from a learnt distribution P_g that approximates the distribution of the real dataset P_r. One can assume that the generated samples differ from the original ones and don't contain their sensitive information. However, there is no guarantee that the generator by repeatedly sampling from P_g will not reproduce the training dataset or generate sensitive ones.

ESANN 2021 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Online event, 6-8 October 2021, i6doc.com publ., ISBN 978287587082-7.
Available from http://www.i6doc.com/en/.

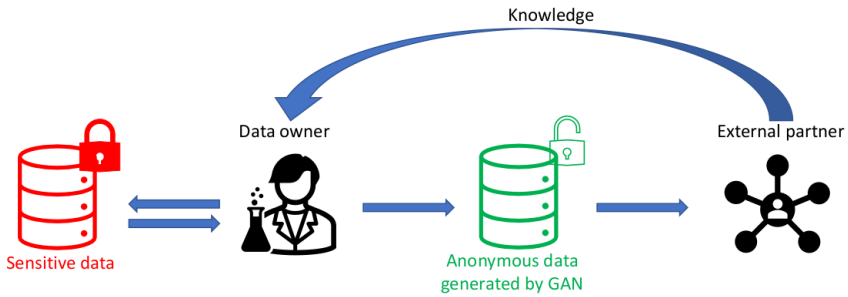

Fig. 1: Use Case Scenario: Data owners holding sensitive data can use a differentially private version of GAN to generate new anonymous time series. These data can be shared with external partners without confidentiality concerns and both parts can work together safely.

Recently, developing differentially private GAN models for images attracted a lot of researchers [6]. As a first attempt, DPGAN [7] proposed a differentially private version of WGAN [8]. Later, an improved version of this method was presented in Dp-GAN [9] and GANobfuscator [10]. On the other side, DP-CGAN [11] focused on another problem namely the generation of anonymous data and their corresponding labels. These models were designed to generate images in a privacy preserving manner. However, as stated in [1], when the task is to generate time series it is not sufficient to capture the distribution of the real dataset. The real challenge is to capture the temporal dynamic between the data point. Hence the previously described models cannot achieve the desired performance for time series. A unique differentially private generative model for time series was proposed by Esteban et al. while introducing RGAN [5]. Inspired by them and motivated by the increasingly need for privacy-preserving data, we will review in this work the existing generative models for time series and extend them to eradicate the privacy concerns. Concretely, we will extend the state-of-the-art GAN algorithms for times series namely -TimeGAN [1], ClaRe-GAN [2] and C-RNN-GAN [3]- with a differential privacy component to pull out the strengtheners and weaknesses of each method. Our approach relies on changing their discriminators with a private discriminator that uses DPSGD [4] and on tracking the spent privacy loss using the RDP accounting technique.

We conducted different experiments on a collection of publicly available datasets from the UCR repository [12] with different number of classes and time series length. We evaluate the results visually and computationally and assess the usefulness and the privacy of the data generated by the deferentially private models. Our experiments show that ClaRe-GAN and outperforms RCGAN in terms of privacy and accuracy while TimeGAN achieves the best accuracies for higher but still reasonable privacy values.

ESANN 2021 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Online event, 6-8 October 2021, i6doc.com publ., ISBN 978287587082-7.
Available from http://www.i6doc.com/en/.

2 Our Approach

In this work, we present a privacy preserving version for the sate-of-the-art GAN models for time series data. As privacy model we use differential privacy [13].

Differential privacy aims at minimizing the influence and the effect of each dataset's instance. Intuitively, the outcome of an algorithm should be insensitive to a small perturbation in the dataset. Given two identical datasets D and D' differing in one single instance, a randomized algorithm $\mathcal{M}$ is (ϵ, δ)-differentially private if for any subset of outputs S [13]:

$$\mathbb{P}\left[\mathcal{M}(D) \in S\right] \leq e^{\epsilon}\mathbb{P}\left[\mathcal{M}(D') \in S\right] + \delta \tag{1}$$

where ϵ and δ control the privacy, $\mathbb{P}$ is the randomness of the noise in the algorithm and $\mathcal{M}(D)$ and $\mathcal{M}(D')$ are the output of the algorithm $\mathcal{M}$ given the dataset D and D'. Lower ϵ and δ values lead to stricter privacy guarantees.

The post-processing theorem states that any randomized mapping on an (ϵ, δ)-differentially private algorithm is also differentially private:

Theorem post-processing [14] Given a randomized algorithm $M : D \to R$ that is (ϵ, δ)-differentially private and an arbitrary randomized mapping $f : R \to R'$, $f \circ M : D \to R'$ is (ϵ, δ)-differentially private.

Recently, a lot of researchers focused on finding a good generative model for time series data generating diverse samples of high-quality. In this context, different approaches have been proposed. While C-RNN-GAN [3] consists of a Long-Short Term Memory (LSTM) generator and discriminator, RCGAN [5] conditions both recurrent neural networks on an auxiliary information. A more complex architecture was proposed by TimeGAN [1] equipped with an embedding and a recovery function in addition to the classic GAN architecture. ClaReGAN [2] focuses on generating time series for multi-class datasets by learning the intra- and inter-class variation.

In this paper, we modify the architecture of the previously described generative models, in the following called Dp-TimeGAN, DP-CRNN-GAN and DP-ClareGAN, to generate time series in a privacy preserving manner. This is achieved by changing their original discriminators with a private discriminator equipped with a differentially private stochastic gradient descent [4]. Two techniques are used to achieve privacy namely clipping gradient and adding random noise. During the training procedure, the per-example gradients of the discriminator loss is computed for the real and generated data. Afterwards, both values are clipped to the minimum value between their L2-norm and a clipping value C. The clipped gradients are summed up and Gaussian noise $N(0, \sigma^2 C^2)$ is added where σ is noise multiplier. Based on the post-processing theorem [14], we guarantee that the use of the private discriminator in any generative model makes the generator and all other architecture's components (encoders etc..) private. The spent privacy loss is computed using the RDP accounting technique [15] as it enables a tighter privacy estimation than the moment accountant technique and an easy computation of the privacy budget curve for a composite mechanism.

ESANN 2021 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Online event, 6-8 October 2021, i6doc.com publ., ISBN 978287587082-7.
Available from http://www.i6doc.com/en/.

3 Experiments

We evaluate the performance of the designed differentially private frameworks visually and computationally. The computational evaluation is performed by computing the test accuracy of Train on Synthetic and Test on Real (TSTR) and Train on Real and Test on Synthetic (TRTS) [5]. TRTS denotes training a ML classifier on the real data and reporting the test accuracy when the model is tested on the synthetic ones. On the other side, TSTR denotes training a ML classifier on the synthetic data and reporting the test accuracy when the model is tested on the real ones. In both cases, the ML model is used to classify the time series of the datasets. We test the developed differentially private GANs on a collection of datasets from the UCR Repository[12] namely ItalyPowerDemand, TwoLeadECG, FreezerRegularTrain, Yoga and DistalPhalanxTW with time series length varying between 24 and 426 and number of classes equal to 2 except for DistalPhalanxTW where it is equal 5. Their performances are compared to the existing differentially private GAN model for time series data of RCGAN. In order to enable a fair comparison between all the frameworks, the same architecture is used for their generators and discriminators e 2-layers LSTM with 100 hidden units and the same number of iterations i.e., 100. For each framework, we take the iterations with the best performance (best TSTR and TRTS values). In contrast to the differentially private RCGAN and DP-ClareGAN, the labels are not generated with the data for DP-TimeGAN and Dp-CRNN-GAN. We label the data generated by these frameworks manually by finding the nearest real time series. In all the experiments we use $C = 0.3$, $\sigma = 0.3$ and $\delta = 10^{-3}$. For each dataset, we compute the spent privacy ϵ and assess the utility of the generated time series. This is achieved by computing the tstr and TRTS accuracy values for 100 time series generated by each framework. As ML model we use Random Forest [16]. In this set of generated time series all the classes of the original dataset are represented with the same number of time series. Our main goal is to find the best performing method i.e., the method that finds the right balance between privacy and utility of the generated time series.

Fig.2 illustrates the test accuracies values of TRTS and TSTR for the different datasets and different frameworks. The figures show that the best privacy values are achieved by C-RNN-GAN. At the same time, its TSTR and TRTS accuracies are really low. For all the datasets, ClaRe-GAN presents better privacy and TRTS TSTR values than the existing differential privacy algorithm of RGAN. It is to be noted that ClaRe-GAN and RCGAN generate labeled data. Especially, we noted a great improvement in terms of privacy for ItalyPowerDemand TwoLeadECG and DistalPhalanxTW datasets. TimeGAN is characterized by high TRTS and TSTR values for higher - but still reasonable - privacy values, by way of example 0.8 and 0.74 for TwoLeadECG dataset. Moreover, it achieves better privacy values for DistalPhalanxTW and ItalyPowerDemand. We have also noticed that the TSTR and TRTS values of RCGAN are around 0.5 which shows a limited utility of the generated data.

Fig. 3 illustrates the time series generated for the different dataset. TimeGAN

ESANN 2021 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Online event, 6-8 October 2021, i6doc.com publ., ISBN 978287587082-7.
Available from http://www.i6doc.com/en/.

and RCGAN generate noisy time series similar to the real dataset. The time series generated by ClaReGAN differ from the real ones and are more private. This corresponds to the privacy values presented in Fig. 2 i.e., for the TwoLead-ECG dataset Clare-GAN $\epsilon = 147.2$ compared to $\epsilon = 287.57$ and $\epsilon = 442.15$ for TimeGAN and RCGAN. C-RNN-GAN generates noise. This explain while C-RNN-GAN achieves the best ϵ values in Fig. 2.

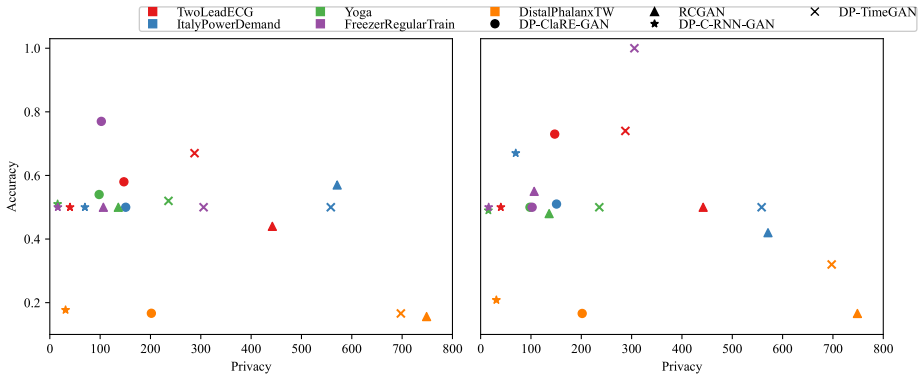

Fig. 2: Test accuracy values for TSTR and TRTS methods depicted in the left and right sub-figure respectively. While a higher accuracy value denotes better usefulness of the generated data, a lower ϵ value denotes a better privacy.

4 Conclusion

We presented in this work a method to generate time series in a private manner by combining the existing GAN frameworks for time series with differential privacy. We have shown that the developed frameworks achieve the desired behavior and that DP-TimeGAN and DP-Clare-GAN outperforms the existing differentially private RCGAN in terms of privacy and usefulness of the generated time series. Our experiments show also that DP-CRNN-GAN achieves the best privacy values. However, it also decreases drastically the quality of the generated data. In the future, we plan to investigate other differential privacy frameworks such as PATE.

References

[1] J. Yoon, D. Jarrett, and M. van der Schaar. Time-series generative adversarial networks. In *Advances in Neural Information Processing Systems*, pages 5508–5518, 2019.

[2] H. Arnout, J. Bronner, and T. Runkler. ClaRe-GAN: Generation of class-specific time series. `https://openreview.net/forum?id=whySRc6f5g_`. Accessed: 10-05-2021.

[3] O. Mogren. C-RNN-GAN: Continuous recurrent neural networks with adversarial training. *arXiv preprint arXiv:1611.09904*, 2016.

[4] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, pages 308–318, 2016.

[5] C. Esteban, S. L Hyland, and G. Rätsch. Real-valued (medical) time series generation with recurrent conditional gans. *arXiv preprint arXiv:1706.02633*, 2017.

[6] L. Fan. A survey of differentially private generative adversarial networks. In *The AAAI Workshop on Privacy-Preserving Artificial Intelligence*, 2020.

ESANN 2021 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Online event, 6-8 October 2021, i6doc.com publ., ISBN 978287587082-7.
Available from http://www.i6doc.com/en/.

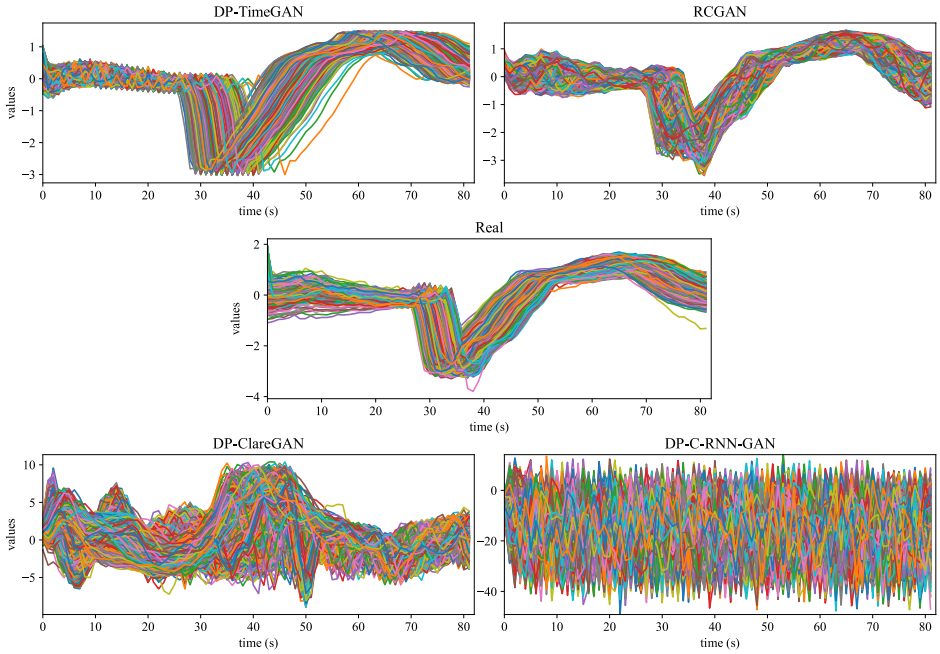

Fig. 3: Illustration of time series generated by the different models for the TwoLeadECG dataset.

[7] L. Xie, K. Lin, S. Wang, F. Wang, and J. Zhou. Differentially private generative adversarial network. *arXiv preprint arXiv:1802.06739*, 2018.

[8] M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein generative adversarial networks. In *International conference on machine learning*, pages 214–223. PMLR, 2017.

[9] X. Zhang, S. Ji, and T. Wang. Differentially private releasing via deep generative model (technical report). *arXiv preprint arXiv:1801.01594*, 2018.

[10] C. Xu, J. Ren, D. Zhang, Y. Zhang, Z. Qin, and K. Ren. Ganobfuscator: Mitigating information leakage under GAN via differential privacy. *IEEE Transactions on Information Forensics and Security*, 14(9):2358–2371, 2019.

[11] R. Torkzadehmahani, P. Kairouz, and B. Paten. DP-CGAN: Differentially private synthetic data and label generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019.

[12] A. Bagnall, J. Lines, W. Vickers, and E. Keogh. The UEA & UCR time series classification repository. *URL http://www. timeseriesclassification. com*, 2018.

[13] C. Dwork. Differential privacy: A survey of results. In *International conference on theory and applications of models of computation*, pages 1–19. Springer, 2008.

[14] C. Dwork, A. Roth, et al. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3-4):211–407, 2014.

[15] I. Mironov. Rényi differential privacy. In *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, pages 263–275. IEEE, 2017.

[16] L. Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.