Privacy-Preserving Synthetic Trajectory Data Generation

David Bergström, Md Fahim Sikder and Fredrik Heintz, Linköping University

1 Introduction

A major open research challenge is developing privacy-preserving machine learning methods that both achieve high performance and privacy guarantees even though the original training data contains sensitive personal information. The applications are abundant, from making cities safer, via on-demand public transportation systems to improved medical diagnosis.

The goal of our PhD projects is to develop new machine learning methods for creating synthetic spatio-temporal trajectory data sets preserving the privacy of the individuals in the original data. We will 1) extend generative adversarial network (GAN) methods to learn generative spatio-temporal trajectory models and 2) develop new Bayesian Optimization methods for creating tailored privacy-preserving synthetic data sets using these generative models.

2 Problem Statement

Accurately analyzing and predicting movements of objects through time and space is central to many applications including autonomous vehicles, transportation systems and city planning. Even though the movement is continuous in nature these trajectories are often associated with discrete properties and choices such as which exit to take in a crossing, the properties of the person or vehicle moving around, and the activity currently being performed.

The general problem is to take a set of noisy trajectories, where each trajectory is a timed sequence of noisy observations of a single object, learn a generative model from these trajectories and then use this model to generate synthetic trajectory data sets while preserving the privacy of the individuals in the original data. A generative model is a method for generating a target distribution with the desired statistics. Generative adversarial networks (GANs) are an expressive class of neural generative models with tremendous success in modeling high-dimensional continuous measures [1].

3 Approach

Generating trajectories poses a unique challenge as the model both have to capture the distributions of features *within* each time point **and** the potentially complex dynamics of those variables *across* time. One approach is to extend GANs with recurrent neural networks to handle sequences, such as C-RNN-GAN [2], Recurrent Conditional GAN (RC-GAN) [3], and Time-Conditional GAN (T-CGAN) [4]. However, these methods do not address the temporal correlations unique to trajectory data. Another line of work is to use autoregressive models to explicitly factor the distribution of sequences into a product of conditionals [5, 6].However, these approaches are deterministic and not really generative in the sense that they can generate several different sequences. A new and interesting approach is TimeGAN that learns an

embedding space jointly optimized with both supervised and adversarial objectives, that encourage the network to adhere to the dynamics of the training data during sampling [7]. This will be our starting-point.

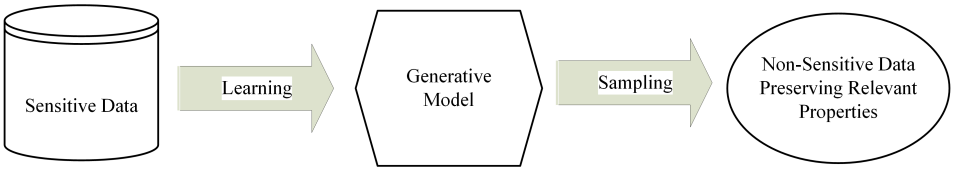

Figure 1: Proposed Model

The next step is to use these generative models to create privacy-preserving synthetic data sets. The power of synthetic data comes from its grounding in real data and real distributions, which make it almost indistinguishable from the original data. To further improve the privacy, synthetic data may be combined with differential privacy [8]. One approach to this is PATE-GAN [9], which combines the Private Aggregation of Teacher Ensembles (PATE) framework [10] with GANs. To construct the synthetic data set we propose to use Bayesian Optimization to select which samples to include [11].

Finally, we will also develop methods for verifying that the generated synthetic data is both accurate enough and privacy preserving.

References

[1] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *NeurIPS 27*, 2014.

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

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

[4] G. Ramponi, P. Protopapas, M. Brambilla, and R. Janssen. T-CGAN: Conditional generative adversarial network for data augmentation in noisy time series with irregular sampling. *arXiv 1811.08295*, 2018.

[5] S. Bengio, O. Vinyals, N. Jaitly, and N. Shazeer. Scheduled sampling for sequence prediction with recurrent neural networks. In *NeurIPS 28*, 2015.

[6] K. Xu A. Goyal R. Lowe J. Pineau A. Courville Y. Bengio D. Bahdanau, P. Brakel. An actor-critic algorithm for sequence prediction. In *ICLR*, 2018.

[7] J. Yoon, D. Jarrett, and M. van der Schaar. Time-series generative adversarial networks. In *NeurIPS 32*, 2019.

[8] S. Bellovin, P. Dutta, and N. Reitinger. Privacy and synthetic datasets. *Stanford Technology Law Review*, 22:1, 2019.

[9] Jinsung Yoon, James Jordon, and Mihaela van der Schaar. PATE-GAN: Generating synthetic data with differential privacy guarantees. In *Proc. ICLR*, 2019.

[10] N. Papernot, M. Abadi, U. Erlingsson, I. Goodfellow, and K. Talwar. Semi-supervised knowledge transfer for deep learning from private training data. *arXiv 1610.05755*, 2016.

[11] B. Shahriari, K. Swersky, Z. Wang, R. Adams, and N. De Freitas. Taking the human out of the loop: A review of bayesian optimization. *Proc. IEEE*, 104(1):148–175, 2015.