7a.013.UVA - Privacy Preserving Multi-Party Analytics

Peter Beling (PI) , Stephen Adams (Co-PI) , Alex Langevin (Researcher) , Steve Greenspan (IAB Project Mentor)
University of Virginia, CA Technologies

Project Start: 07/01/2018 | **End Date: 06/30/2019**

Project Summary:

Many domains and industries would benefit from sharing data for predictive modeling but for privacy, legal/regulatory, competitive, or other reasons, parties may be unwilling or unable to share information with each other. Differential privacy (DP), whereby carefully selected random noise is added to the development process, presents a way to share (approximate) information while guaranteeing mathematically a level of privacy for each party.

Year 6 of the project developed a DP method for jointly training a deep learning model applied to credit card fraud detection. Year 6 results successfully demonstrated multi-party deep learning under certain conditions and simplifying assumptions. In Year 7 we intend to leverage the results of Year 6 and examine alternate learning solutions, as well as relax some key assumptions, moving towards more realistic settings.

Year 7 will see the introduction of dataset heterogeneity – it will no longer be assumed that each party collects the same data on their customers and/or have the same customer (i.e. data) distribution. In this scenario it may no longer make sense to jointly develop a single general model that fails to account for local peculiarities in each party's data. To address this we plan to incorporate Generative Adversarial Networks (GANs), which are models that learn the underlying structure of a dataset, and can generate synthetic data that mimics the distribution of the real data.

The methods developed in Year 6 can be used to train DP GANs, which can then be shared with other parties and used to augment the training of local models. GANs can also be used to address issues with imbalanced data – our dataset has a fraud rate of 0.14%. We also plan to utilize the discriminator model, which is often overlooked in the literature outside of training, to pre-select data from other the GANs with a view of optimizing local model development.

The primary challenge foreseen will be the development of GANs for discrete data, for which there is currently no established method, and integrating this discrete model into a joint discrete/continuous GAN.

Details of Progress/Achievements:

Literature review and concept development
Currently in data pre-processing, feature extraction, and model building phase

Generative Adversarial Networks (GANs)

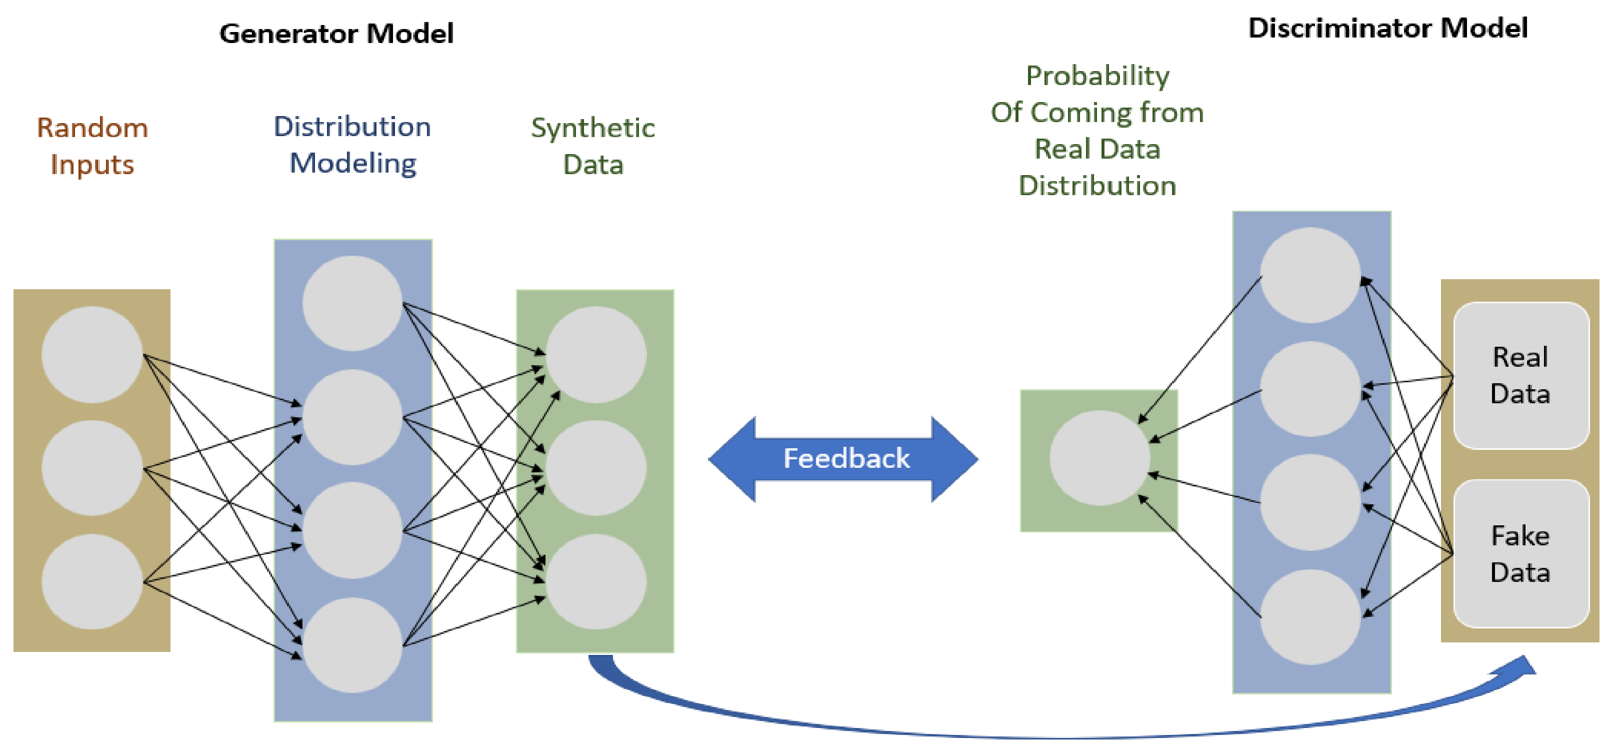

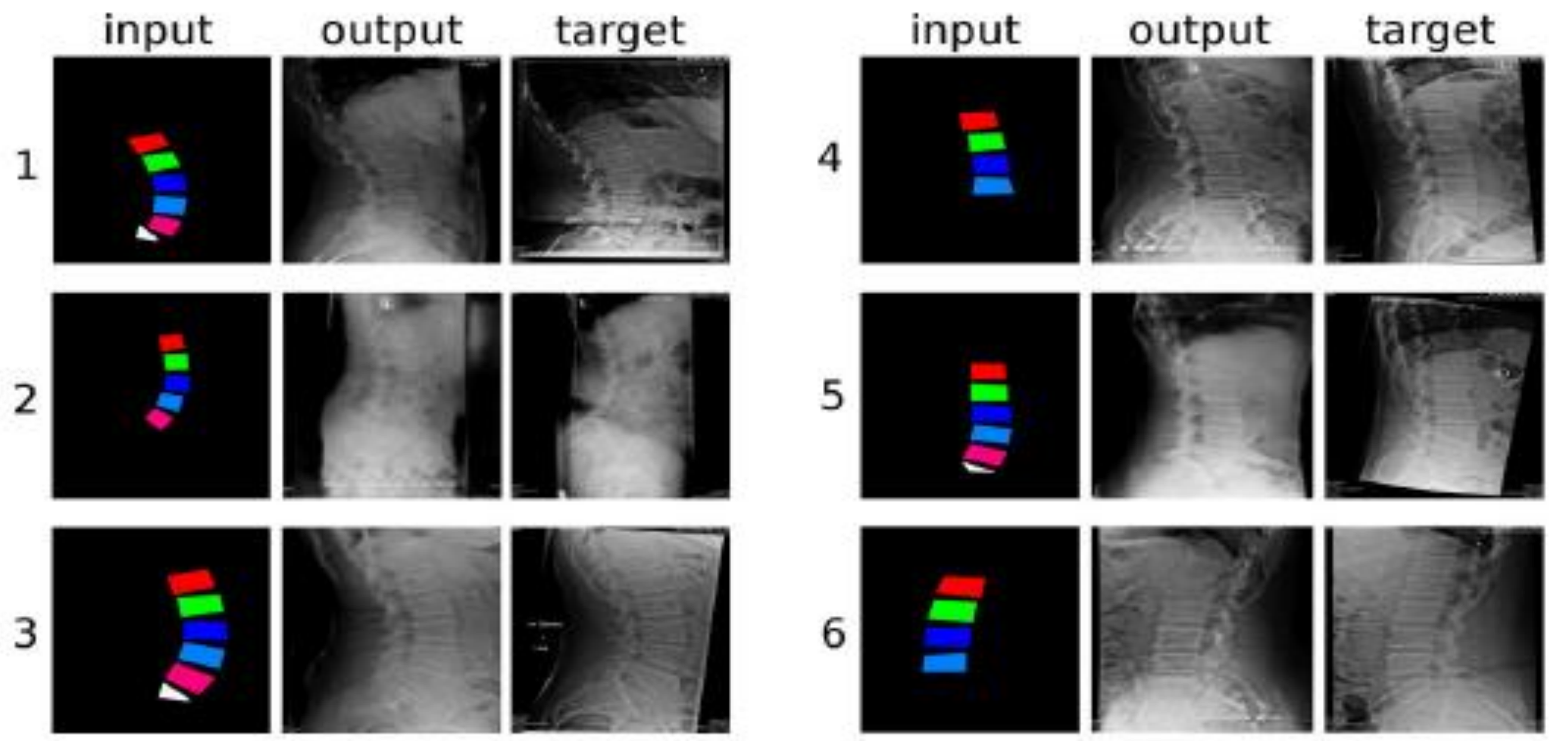

Source: Galbusera et al. (2018). Exploring the Potential of GANs for Synthesizing Radiological Images of the Spine to be Used in In Silico Trials. Frontiers in Bioengineering and Biotechnology. vol. 6 (53).

PROJECT DELIVERABLES

1. Literature Review	5. Test GAN performance in public (i.e. no differential privacy) setting
2. Develop and test GAN for continuous data	6. Build differential privacy into the discriminator model and re-evaluate
3. Develop and test alternative methods for discrete/categorical data GANs	7. Introduce multi-party setting with heterogeneous data
4. Develop architecture for joint continuous/categorical GANs	8. Performance benchmarking