DPD-fVAE: Synthetic Data Generation Using Federated Variational Autoencoders With Differentially-Private Decoder

Bjarne Pfitzner, Prof. Dr. Bert Arnrich

1 Motivation

- **Problem:**
  - Deep learning requires **lots of data**
  - **Local** datasets are often **small**
  - **Privacy regulations** restrict data sharing
- **Federated learning** [1] can solve this
- Training **data generators** enables future investigation of (previously unconsidered) research questions

3 Method: DPD-fVAE

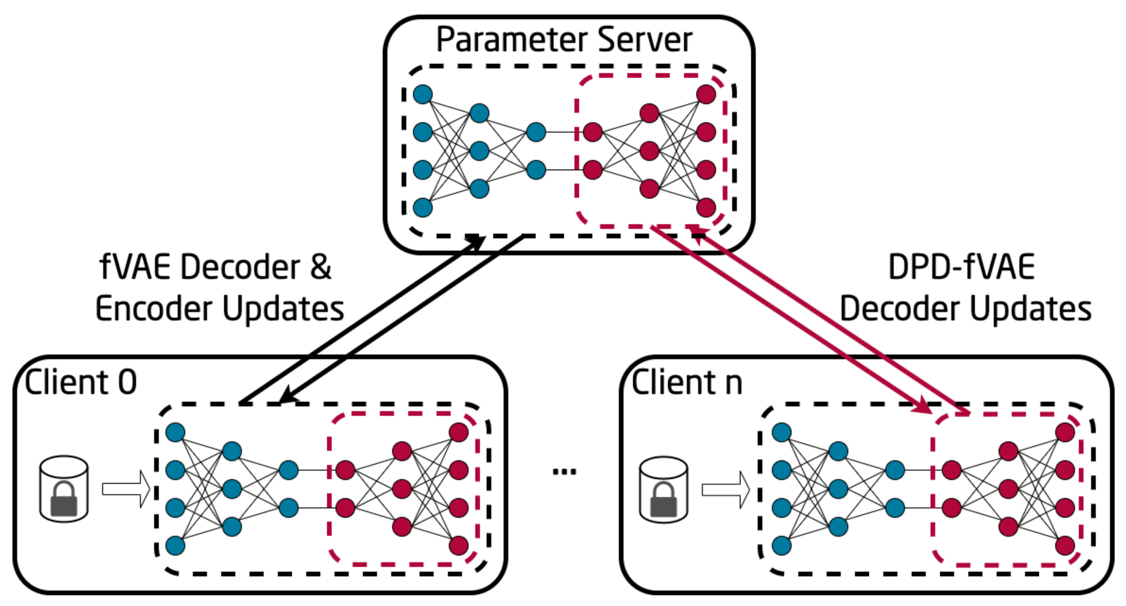

- Only synchronise **decoder**, not full model
- Reduces L_2-norms of updates

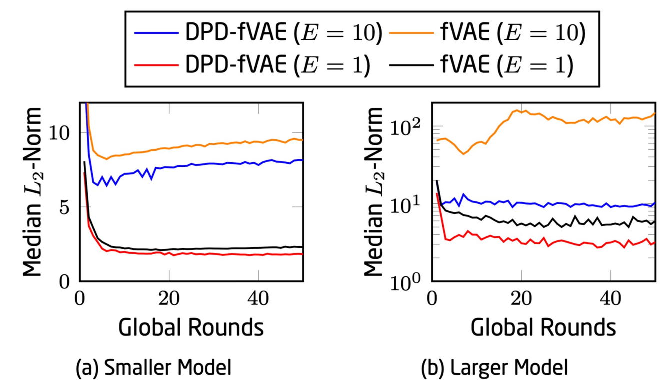

→**Reduces privacy spending** per round

2 Background

Federated Learning (FL) [1]

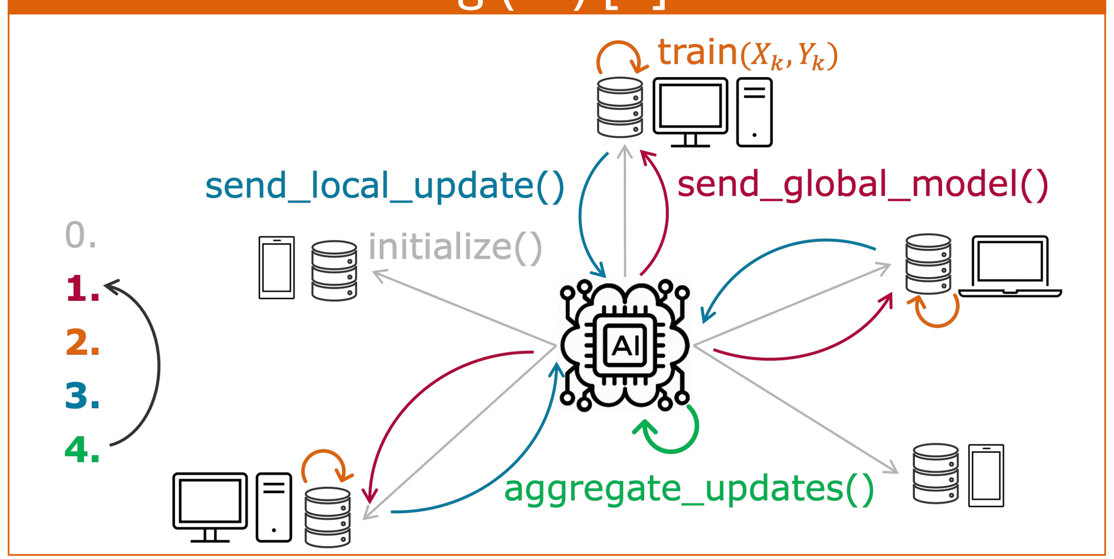

Differential Privacy (DP) [2]

- Formal guarantee of privacy
- Limits impact of single clients/data on model
- (ε, δ)-DP SGD [3]
  - ε: Budget
  - δ: Risk
  - Clips gradient updates' L_2-norms to S
  - Adds noise $\mathcal{N}(0, qS)$ to updates
- Two types of DP for FL:
  - **Central** DP (CDP): Server protects clients
  - **Local** DP (LDP): Clients protect data

Variational Autoencoders (VAEs) [4]

- Learns **latent space distribution** of data
- Capable of synthesising new data by sampling latent space

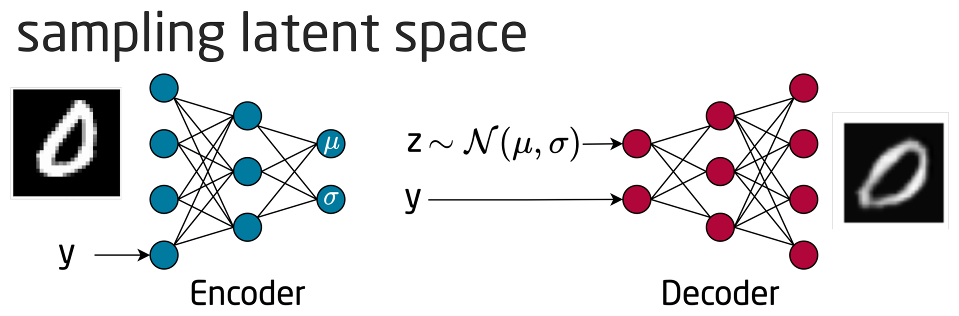

4 Results and Discussion

- **Synthetic images** with $(10, 10^{-5})$-CDP/-LDP

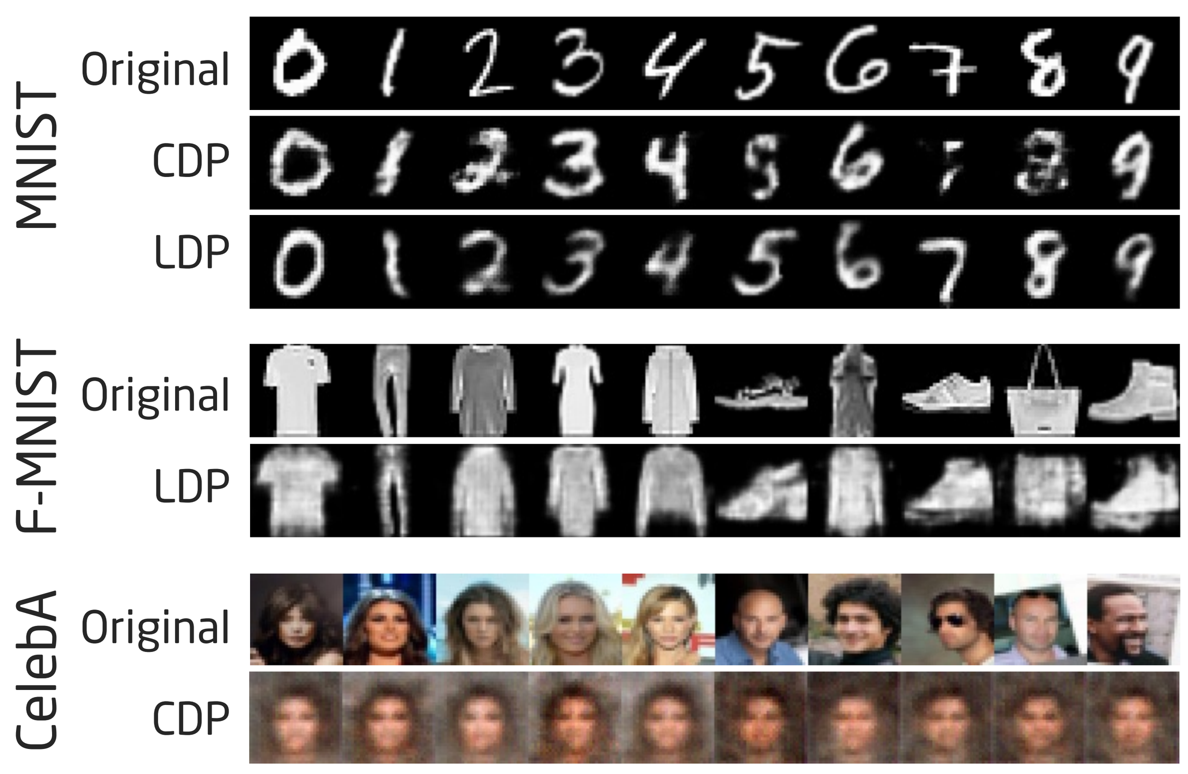

- DPD-fVAE converges where fVAE does not

- Comparison with SOTA
  - All other methods are **centralised**
  - DPD-fVAE is **federated**
  - Other FL methods are not comparable

- Evaluation of **different** privacy **budgets** ε
- **Quantitative evaluation** for MNIST data:
  - Fréchet Inception Distance (FID) [5]
  - Classifier (CNN) accuracy

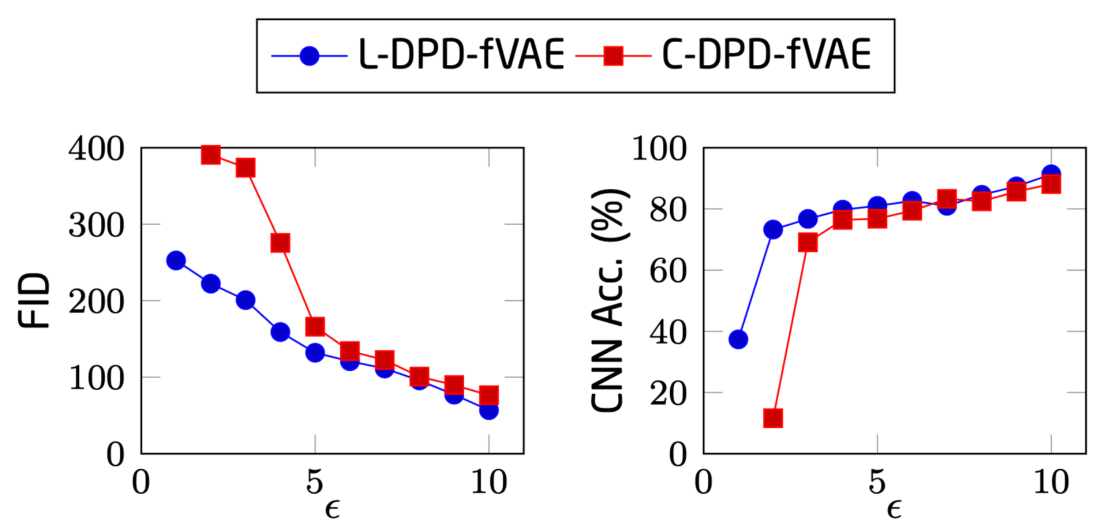

Key Takeaways

- DPD-fVAE performs in line with SOTA, even though FL is harder than centralised ML
- Base VAE struggles with sharpness and background information (CelebA)
- Generally, performance of DP-FL relies heavily on the scenario (# clients, size of local data, ...)

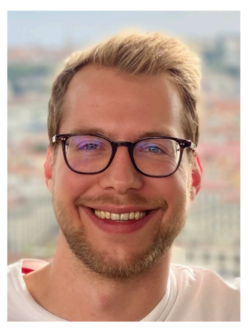

[1] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In Artificial intelligence and statistics, pages 1273–1282. PMLR, 2017.
[2] Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. Found. Trends Theor. Comput. Sci., 9(3-4):211–407, 2014.
[3] Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In Proceedings of the 2016 ACM SIGSAC conference on computer and communications security, pages 308–318, 2016.
[4] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114, 2013.
[5] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. Advances in neural information processing systems, 30, 2017.

Bjarne Pfitzner
Digital Health - Connected Healthcare
bjarne.pfitzner@hpi.de
+49 (0) 331 5509-1374

Prof. Dr. Bert Arnrich
Digital Health - Connected Healthcare
bert.arnrich@hpi.de
+49 (0) 331 5509-4850

Want to read more?

HPI Hasso Plattner Institut