Generative Federated Learning

Arthur A. Nijdam , Paul Stankovski Wagner,

***Abstract*— Federated Learning (FL) [1] was designed for scenarios where raw data collected by users should remain stored on edge devices for privacy reasons. During a FL iteration, each user receives a global model from an orchestrating server, trains the model on its local data, and sends the resulting weight update back to the server.**

FL suffers from convergence issues due to the inherent parallelization of the training process, particularly when the data is distributed non-IID or when annotator labels are unavailable.

In this work, we ask ourselves: What if it would be possible to share synthetically generated data among users? We employ Generative Federated Learning (GFL) for users to synthesize data like their own dataset, but without the privacy concerns associated with the original data. In contrast to other GFL approaches, our method is light-weight, model-agnostic, and does not rely on pre-trained generative models.

We investigate whether GFL improves model convergence, particularly for non-IID user populations. Furthermore, we show how generated data can be used to enhance contrastive Federated Learning.

***Index Terms*— Federated Learning, Generative Federated Learning, Non-IID, Unsupervised Learning.**

References

[1] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Artificial intelligence and statistics*, pp. 1273–1282, PMLR, 2017.

Arthur Nijdam and Paul Stankovski Wagner are with the Department of Electrical and Information Technology, Lund University, 221 00 Lund, Sweden (email: arthur_andreas.nijdam@eit.lth.se; paul.stankovski_wagner@eit.lth.se).