Decentralized Learning of GANs from Multi-Client Non-iid Data

Ryo Yonetani
OMRON SINIC X, Japan
ryo.yonetani@sinicx.com

Tomohiro Takahashi
OMRON Corporation, Japan
tomohiro.takahashi.2@omron.com

Atsushi Hashimoto
OMRON SINIC X, Japan
atsushi.hashimoto@sinicx.com

Yoshitaka Ushiku
OMRON SINIC X, Japan
yoshitaka.ushiku@sinicx.com

Read more The full version of this paper is available at https://arxiv.org/abs/1905.09684, which includes the complete formulation and theoretical analysis of the proposed algorithm and more experimental results.

1. Introduction

We are interested in learning a generative model from multiple image data collections that are each i) owned separately and privately by different clients and ii) drawn from non-identically distributions that comprise different classes. Examples of such *multi-client* and *non-iid* data include personal life-logging videos [3] and medical data [6], which are all owned independently and privately by multiple individuals or institutions and characterized differently due to geographical conditions or personal preferences. On the one hand, learning from such data resources distributed all over the world would improve the diversity of images the trained model can generate and ultimately benefit various applications such as anomaly detection and image-to-image translation. On the other hand, direct access to a collection of data captured exclusively by a single client creates the risk of private information leak, *e.g.*, a living area that could be inferred from a collection of life-logging videos taken by a single person. This dilemma between data utility and privacy will make it hard to aggregate all the client data in a central server, necessitating decentralized approaches.

Nonetheless, much work on decentralized learning has focused mainly on the supervised setting, such as Federated Learning that has been extensively studied [1]. It is hard to determine how such supervised approaches can be adopted for learning generative models from decentralized non-iid data. An exception proposed recently is decentralized learning of generative adversarial networks (GANs) [5], which lets each client train an individual discriminator with their own data while asking the central server to update a generator to fool those discriminators. While this approach allows clients to decentralize their data in each storage, it restricts all client data to be drawn independently from the identical distribution, and has no theoretical guarantee on what distribution will be learned otherwise.

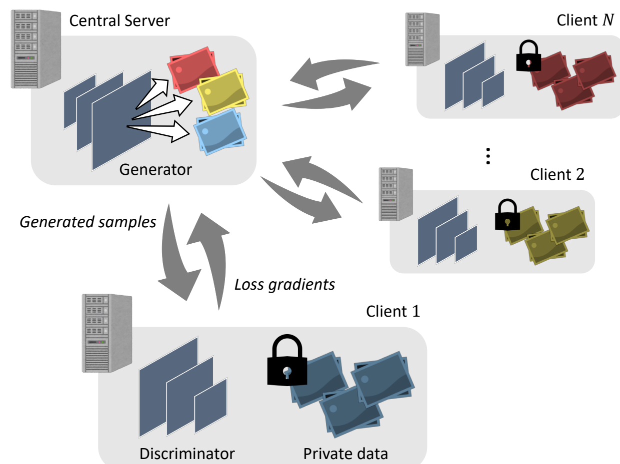

Figure 1. Decentralized learning of GANs.

2. Proposed Approach

Given this background, we develop a decentralized approach tailored for learning GANs from multi-client non-iid data. As shown in Figure 1, suppose that N clients independently train an individual discriminator with their own private data drawn from non-identical distributions with different classes, hereafter $p_1(x), \ldots, p_N(x)$. By training a generator stored in the central server to fool those discriminators, our approach aims at learning a distribution that comprises all the classes input data can belong to, more specifically, $p_{\max}(x) = \frac{1}{Z} \max_i p_i(x)$.

Forgiver First Aggregation The key technical contribution is in a strategy to coordinate multiple discriminators to inform a generator of multi-client non-iid data. Intuitively, when the discriminators are trained from non-identical data

Table 1. FID scores.

	MNIST		CIFAR10	
	Non-OVL	Mod-OVL	Non-OVL	Mod-OVL
MD-GAN [5]	38.42	34.33	56.64	50.30
GMAN* [4]	67.69	58.65	50.50	41.83
F2A (Ours)	**18.96**	**14.53**	**38.92**	**41.01**

distributions, they will judge each generated sample differently. The proposed approach, which we refer to as *Forgiver First Aggregation (F2A)*, i) aggregates such diverse judgments *while emphasizing more 'forgiving' ones who deemed the generated sample as more real and closer to what they have* and ii) updates the generator against the aggregated judgments. This will allow the generator to learn rare classes observed only by a small fraction of the discriminators as well as common ones shared by many clients. The aggregation of multi-discriminator judgments is implemented with a regularized weighted averaging function, where the averaging weights can also be updated with the generator end-to-end to better capture the non-iidness of given client data.

Learning Protocol As shown in Figure 1, the decentralized learning with F2A requires the central server and clients to exchange mini-batches of generated samples and their loss gradients. Another advantage of using F2A is that the computations of loss gradients, in which we need to aggregate judgments of multiple discriminators, can easily be combined with off-the-shelf secure-sum protocols such as [2]. This allows each discriminator to be provably secure under certain settings against malicious parties who try to intercept the training process and leak the private information from the judgments.

3. Experiments

Datasets We empirically evaluated our approach with image generation tasks on MNIST and CIFAR10. Each dataset was split into five subsets (*i.e.*, we considered five clients) with the following two conditions which split images into five such that $p_1(x), \dots, p_5(x)$ respectively comprised the images of classes $\{0,1\}$,$\{2,3\}$,$\{4,5\}$,$\{6,7\}$,$\{8,9\}$ (**non-overlapping; Non-OVL**), and $\{0,1,2,3\}$,$\{2,3,4,5\}$, $\{4,5,6,7\}$, $\{6,7,8,9\}$, $\{8,9,0,1\}$ (**moderately-overlapping; Mod-OVL**).

Implementation We implemented a variant of LS-GANs [7] based on a DCGAN-based architecture where client-wise discriminators had spectral normalization [8] instead of batch normalization.

Baselines, Evaluation Metrics, and Results Our decentralized learning with F2A consists of multiple discriminators to train a single generator. We therefore adopted MD-GAN [5] and GMAN [4] as a baseline method, which were both GANs with multiple discriminators but trained with completely different aggregation strategies. Fréchet Inception Distance (FID) was used as an evaluation metric. Importantly, we found that the choices of hyperparameters such as mini-batch size and the number of iterations affected FID greatly and differently for each method. Instead of picking out one specific hyperparameter, we tested each method with combinations of several mini-batch sizes and numbers of iterations and report the median FID scores. As shown in Table 1, we confirmed that F2A outperformed the other baselines.

References

[1] K. Bonawitz, H. Eichner, W. Grieskamp, D. Huba, A. Ingerman, V. Ivanov, C. Kiddon, J. Konecný, S. Mazzocchi, H. B. McMahan, T. V. Overveldt, D. Petrou, D. Ramage, and J. Roselander. Towards Federated Learning at Scale: System Design. *CoRR*, abs/1902.01046, 2019.

[2] K. Bonawitz, V. Ivanov, B. Kreuter, A. Marcedone, H. B. McMahan, S. Patel, D. Ramage, A. Segal, and K. Seth. Practical Secure Aggregation for Privacy-Preserving Machine Learning. In *CCS*, pages 1175–1191, 2017.

[3] S. Chowdhury, M. S. Ferdous, and J. M. Jose. Exploring Lifelog Sharing and Privacy. In *UbiComp*, pages 553–558, 2016.

[4] I. P. Durugkar, I. Gemp, and S. Mahadevan. Generative Multi-Adversarial Networks. In *ICLR*, 2016.

[5] C. Hardy, E. L. Merrer, and B. Sericola. MD-GAN: Multi-Discriminator Generative Adversarial Networks for Distributed Datasets. *CoRR*, abs/1811.03850, 2018.

[6] M. Li, R. Poovendran, and S. Narayanan. Protecting Patient Privacy against Unauthorized Release of Medical Images in a Group Communication Environment. *CMIG*, 29(5):367 – 383, 2005.

[7] X. Mao, Q. Li, H. Xie, R. Y. K. Lau, Z. Wang, and S. P. Smolley. Least Squares Generative Adversarial Networks. In *ICCV*, pages 2813–2821, 2017.

[8] T. Miyato, T. Kataoka, M. Koyama, and Y. Yoshida. Spectral Normalization for Generative Adversarial Networks. In *ICLR*, 2018.