

Security and Privacy in Federated learning: Challenges and Possible Solutions

Aikaterini Mitrokotsa

School of Computer Science
University of St. Gallen, Switzerland
{katerina.mitrokotsa}@unisg.ch

Mobile phones, wearables, autonomous vehicles and in general Internet of Things (IoT) devices are just some examples of distributed networks that create a wealth of data every day. This data is subsequently used as input in centralised machine learning models in order to achieve reliable user modeling and personalisation. The growing storage and computational power of mobile devices as well as increased privacy concerns have led to an increased interest in federated learning, which allows multiple clients to collaboratively train learning models under the orchestration of a central server, while the data remain located on the sources.

Distributed machine learning has many significant advantages compared to centralised machine learning, mainly regarding efficiency and privacy. However, some serious challenges remain:

- *Privacy*: Although only updates are sent to the server, research has shown that these updates may still leak sensitive information, thus, providing no formal guarantee of privacy. For instance, by having access to a gradient update and the previous model, it might be possible to infer a training example.
- *Security*: The central server represents a single point of failure or even a bottleneck. *How can a client be sure that the server has performed the aggregation correctly?* A “lazy” server might use a simpler model to reduce its computational load, or modify the aggregation result to bias the model.
- *Heterogeneity*: The federated learning process is massively parallel involving multiple clients (up to 10^{10}) with different resources/capabilities. Many of these devices (5% or more) will fail or drop (being controlled by different clients), creating thus, a highly stateless environment.

In this talk, we discuss the main security and privacy challenges in federated learning as well as how we may guarantee and secure and private dynamic aggregation of data which can be employed in the federated learning setting [2, 1]. More precisely, we discuss how by relying on verifiable homomorphic secret sharing, we can achieve secure and verifiable aggregation of multiple users’ secret data (*e.g.* parameters of the learning model), while employing multiple untrusted servers. The proposed solutions compute the sum of the users’ input and provides public verifiability, *i.e.*, anyone can be convinced about the correctness of the aggregated sum computed from a threshold amount of servers, while no communication between the users occurs.

References

1. Brunetta, C., Tsaloli, G., Liang, B., Banegas, G., Mitrokotsa, A.: Non-interactive, secure verifiable aggregation for decentralized, privacy-preserving learning. In: Baek, J., Ruj, S. (eds.) Information Security and Privacy - 26th Australasian Conference, ACISP 2021, Virtual Event, December 1-3, 2021, Proceedings. Lecture Notes in Computer Science, vol. 13083, pp. 510–528. Springer (2021). https://doi.org/10.1007/978-3-030-90567-5_26, https://doi.org/10.1007/978-3-030-90567-5_26
2. Tsaloli, G., Liang, B., Brunetta, C., Banegas, G., Mitrokotsa, A.: Deva: Decentralized, verifiable secure aggregation for privacy-preserving learning. In: Liu, J.K., Katsikas, S.K., Meng, W., Susilo, W., Intan, R. (eds.) Information Security - 24th International Conference, ISC 2021, Virtual Event, November 10-12, 2021, Proceedings. Lecture Notes in Computer Science, vol. 13118, pp. 296–319. Springer (2021). https://doi.org/10.1007/978-3-030-91356-4_16, https://doi.org/10.1007/978-3-030-91356-4_16