

Privacy in Distributed Learning

Advisor: **Phillip Rieger**

Motivation


Distributed learning approaches, such as federated learning (FL) and split learning (SL), aim to enable collaborative training without sharing private training data. In FL, existing approaches use cryptographic protocols and trusted hardware to prevent the server from accessing individual client updates. However, the resulting global model may still leak sensitive information, even after training is completed. Similarly, in SL, the server processes intermediate representations of private inputs, potentially revealing information during inference. Thus, protecting raw data or individual updates alone does not guarantee privacy.

Possible Topics

This research area offers several thesis topics on privacy in distributed learning. Depending on the topic, projects may investigate privacy attacks, develop methods to measure privacy leakage, or design novel privacy-preserving techniques. For example, one possible direction is to systematically evaluate what information can be inferred from aggregated models or intermediate representations. Specific research questions will be discussed with students based on ongoing research and current developments in the field.

Goals and Tasks

-  Study privacy attacks and defenses in distributed learning
-  Design methods to measure privacy leakage or improve privacy
-  Implement inference attacks, privacy metrics, or defense mechanisms
-  Experimentally evaluate privacy leakage and proposed approaches



Literature

- > [A. Salem et al.](#)
Updates-Leak: Data set inference and reconstruction attacks in online learning
[USENIX Security](#)
- > [F. Mo et al.](#)
PPFL: Privacy-preserving federated learning with trusted execution environments
[MobiSys](#)

Courses & Deliverables

- ☒ **Master Project**
Project code
Report
Presentation
 - OR –
 - ☒ **Master's Thesis**
Initial presentation
Project code
Thesis (60+ pages)
Final presentation

Recommended if you're studying

- ☒ CS ☒ ICE ☒ SEM

Prerequisites

- > Interest in AI security, privacy, and distributed learning
- > Practical experience with training deep neural networks
- > Familiarity with deep learning frameworks (e.g., PyTorch)

Advisor Contact

phillip.rieger@tugraz.at