

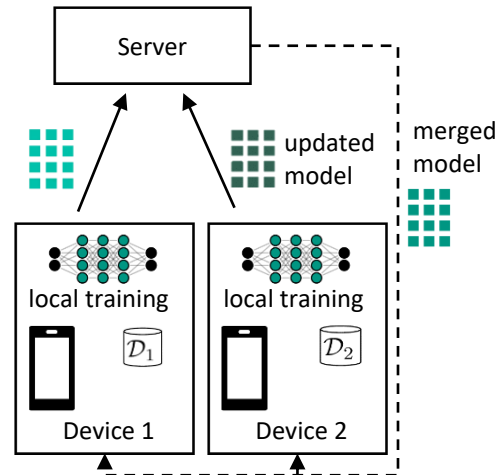
## Master Thesis

### Model Update Quality in Federated Learning

In federated learning, a set of **distributed devices** (e.g., smartphones or sensor nodes) cooperatively learn towards a specific goal. Thereby, each device trains a local **neural network (NN)** with its own locally collected training data. All devices periodically synchronize their knowledge via a server to be able to benefit from each other and jointly achieve a better performance. Different devices have different data (amount, quality, distribution), commonly called non-i.i.d. data. Therefore, the quality of the local model updates that a device sends to the server varies. However, measuring the quality of an update, which comprises millions of weight changes, is challenging.

The goal of this thesis is to develop methods on how the quality of a model update can be measured. This could for instance be based on separate test data, or more sophisticated analysis of statistical information about the weights or analysis of the loss surface. In addition, after being able to measure the quality of updates that different devices deliver, a prediction of the quality of future updates of each device could greatly improve the overall learning.

This thesis contributes to improve the distributed training of ML models, where a higher accuracy can be reached faster.



#### Skills required for the thesis

- Programming skills (Python)
- Background on machine learning
- Experience with ML frameworks such as PyTorch or tensorflow is beneficial but not required

#### Skills acquired within the thesis

- In depth knowledge of distributed ML with non-i.i.d. data
- Technical writing skills
- Work in a research environment

#### Language

The collaboration with the colleagues can be in English or German.

#### Contact

Dr. Martin Rapp – [martin.rapp@kit.edu](mailto:martin.rapp@kit.edu)

Kilian Pfeiffer – [kilian.pfeiffer@kit.edu](mailto:kilian.pfeiffer@kit.edu)