EPFL

Inria
inventeurs du monde numérique

MobNet: Online Learning at the Edge via Staleness Awareness and Performance Prediction

Georgios Damaskinos, Rachid Guerraoui, Anne-Marie Kermarrec, Vlad Nitu, Rhicheek Patra, Francois Taiani

The individual privacy

- Conflict of interests between service providers and users
  - Service providers need users' data to improve their services
  - Many users value their individual privacy

>$100 billion was knocked off

>$100 billion was knocked off

Private Machine Learning is the future

Private Machine Learning

- Secure cloud-based Machine Learning
  - Secure processors (e.g. SGX)
  - Homomorphic Encryption, Multi-party Computation
- Federated Learning

Cloud-based Machine Learning

- Instead of centralizing the users' data...

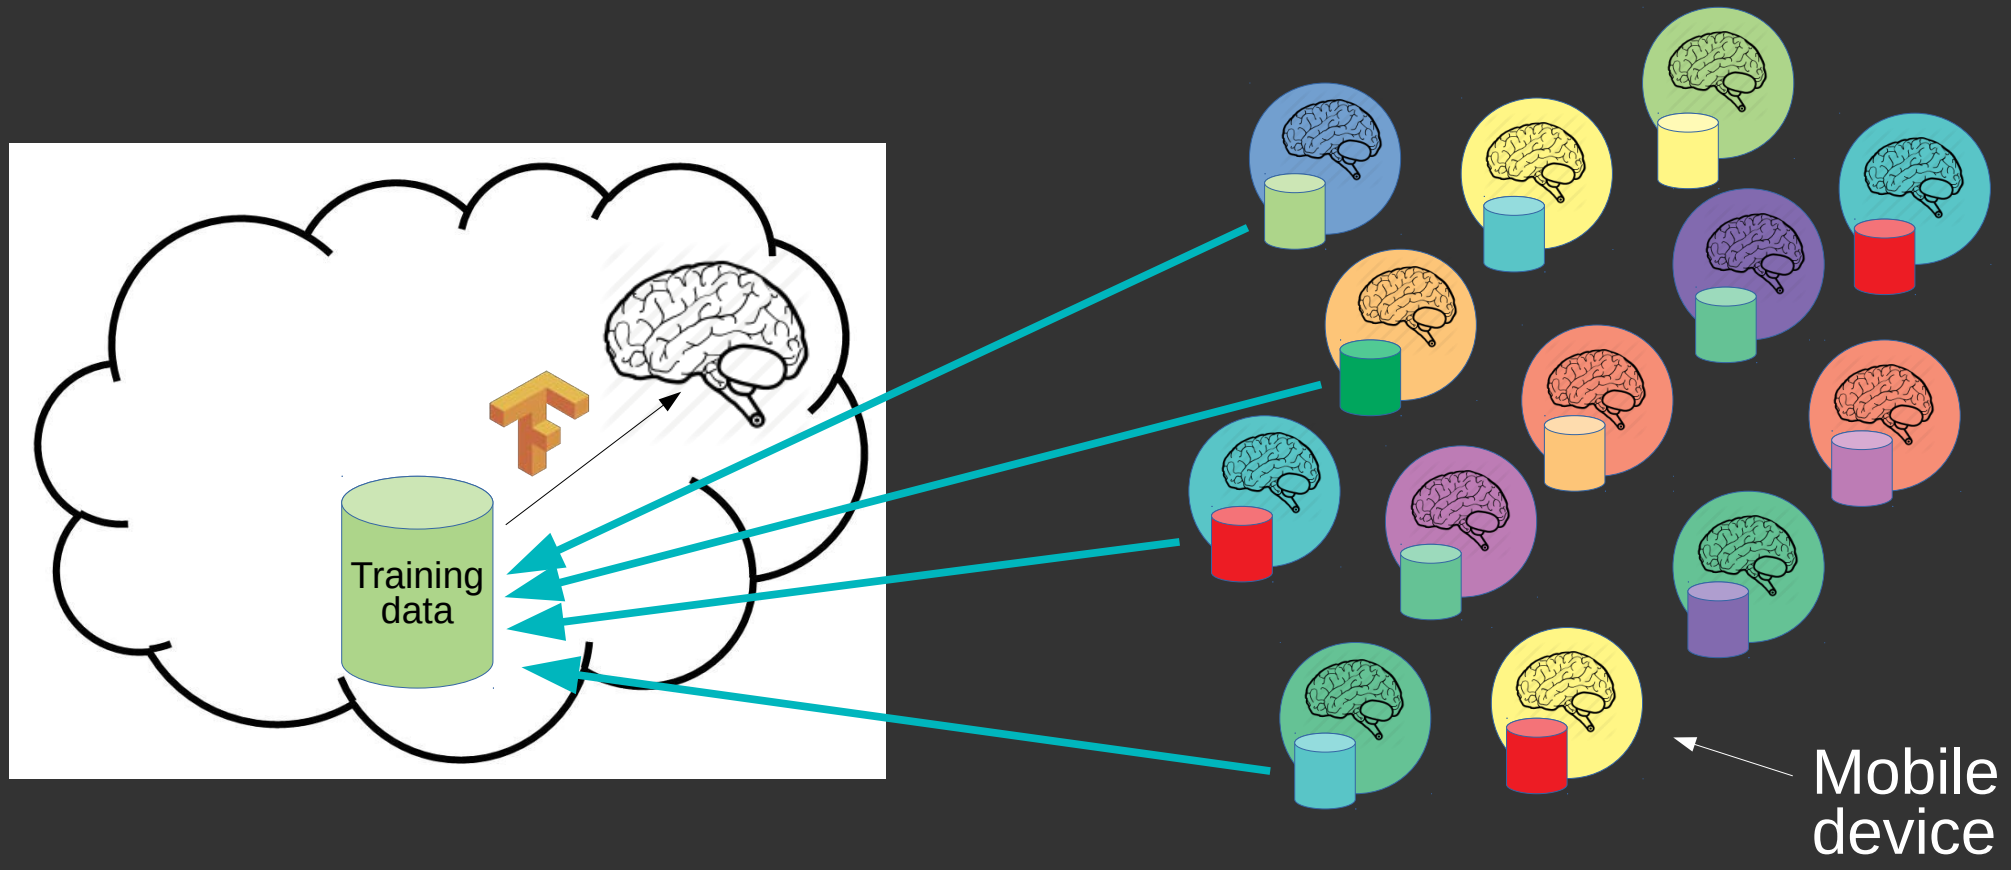

Federated Learning

- Train models right on the device.

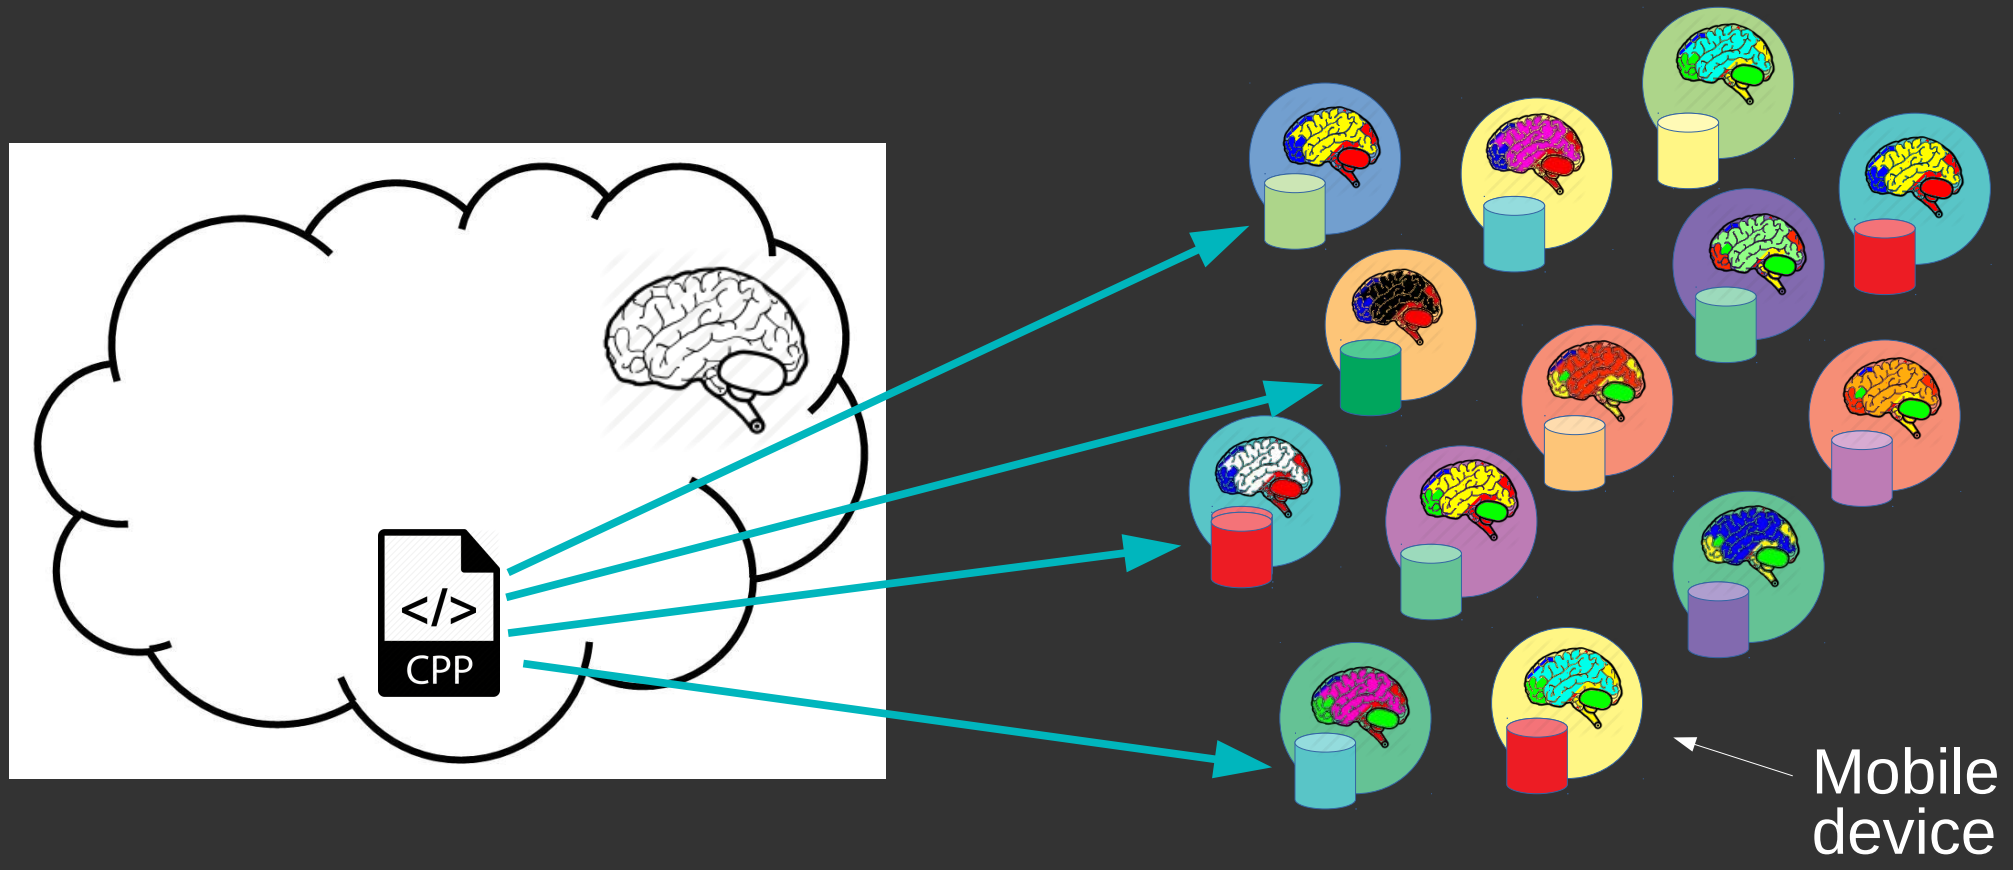

Federated Learning

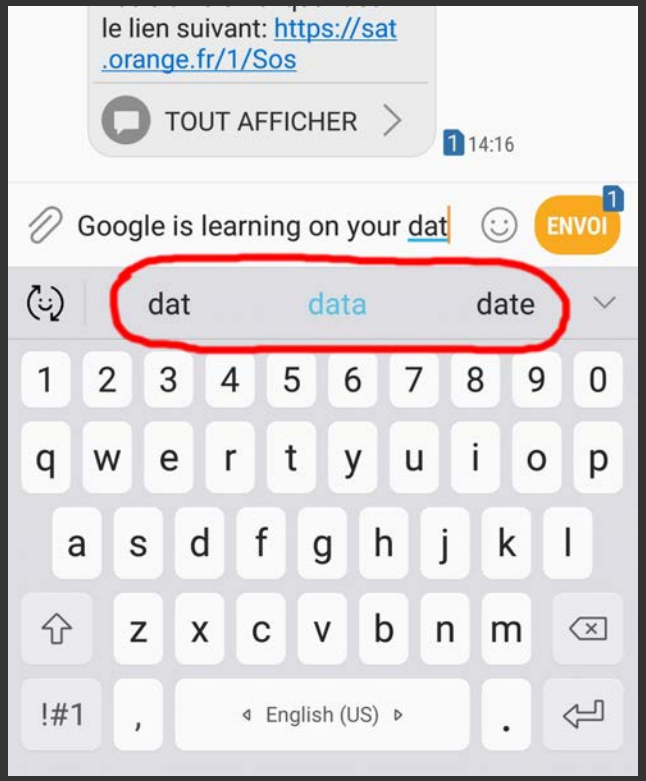

*Bonawitz, Keith, et al. "Towards Federated Learning at Scale: System Design." SysML 2019.

Federated Learning

- The network protocol:

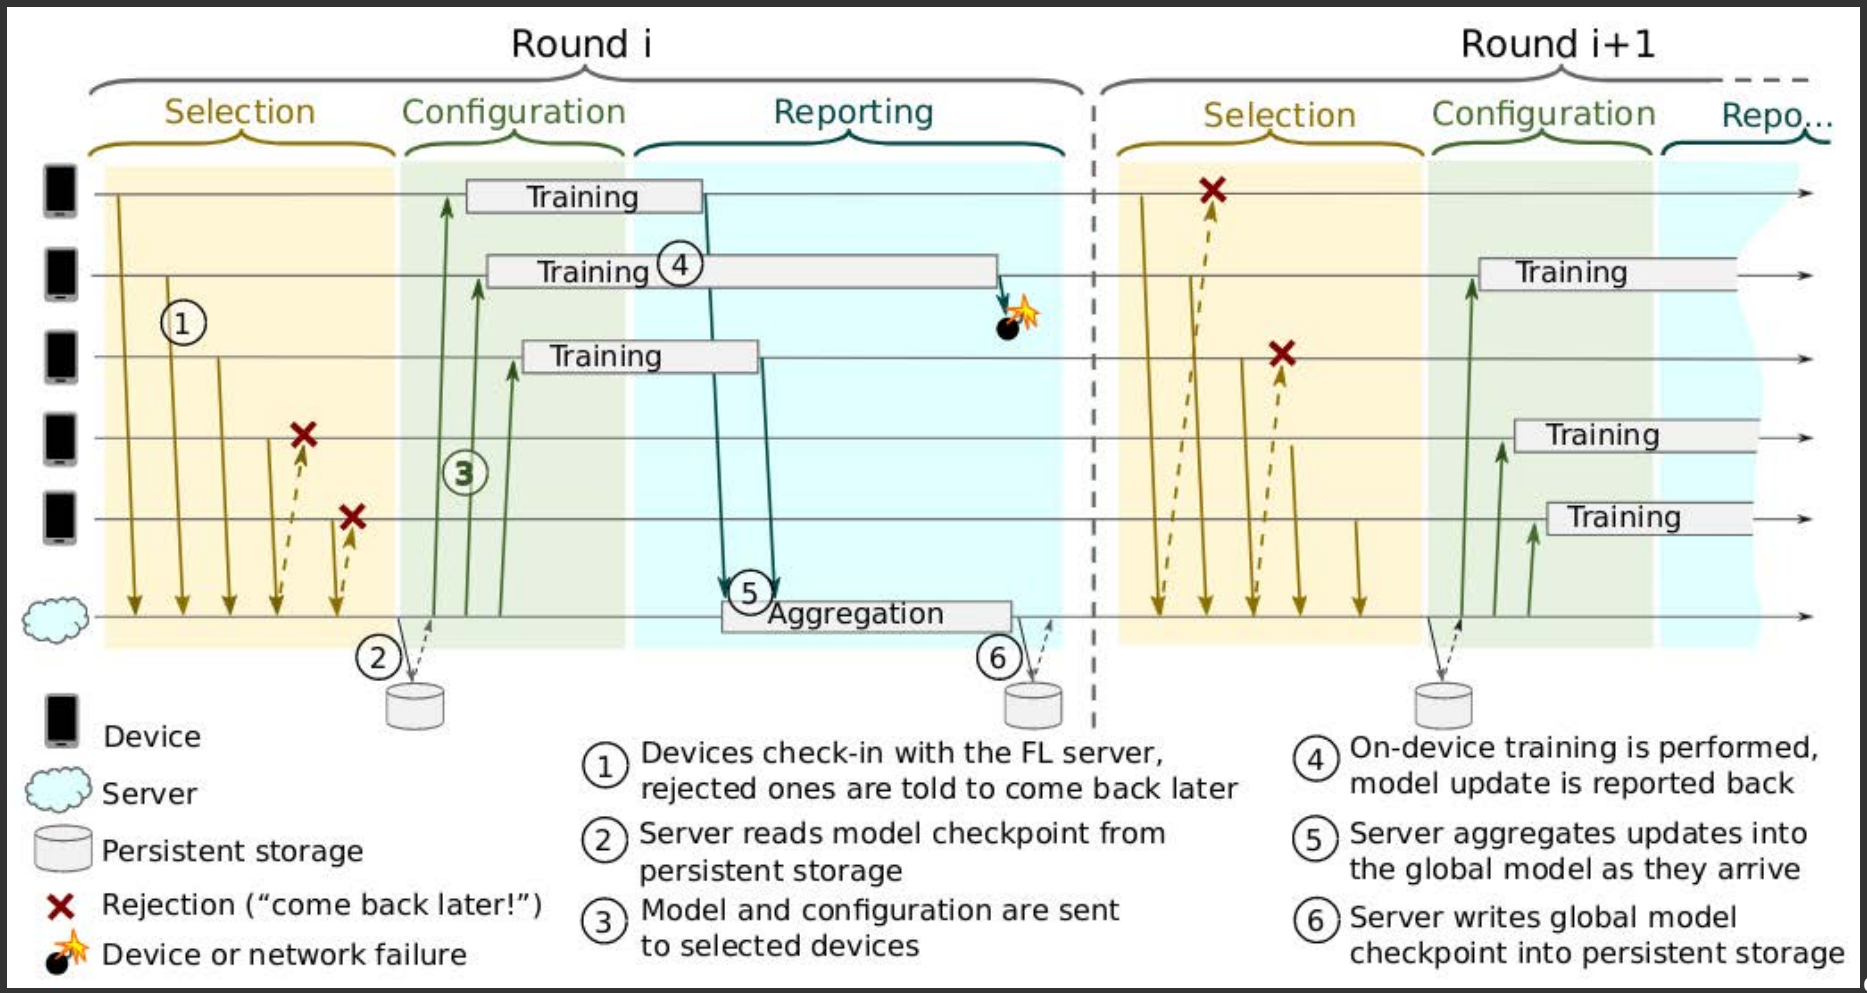

*Bonawitz, Keith, et al. "Towards Federated Learning at Scale: System Design." SysML 2019.

Federated Learning

- Designed for no impact on user experience
- Requirements for mobile devices
  - Idle, charging, connected to WiFi

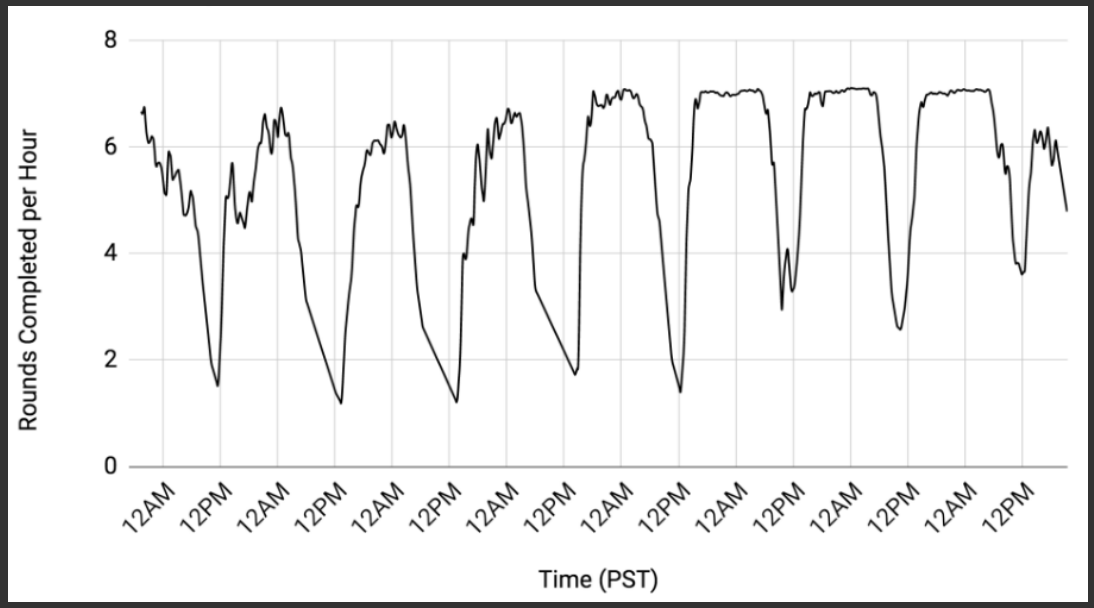

*Yang, Timothy, et al. "Applied federated learning: Improving google keyboard query suggestions." arXiv preprint arXiv:1812.02903 (2018).

Online vs. Offline FL

- Offline FL is not suitable for fast-data apps
  - e.g. social networks, recommenders
- Train 10 times more often → 23% better

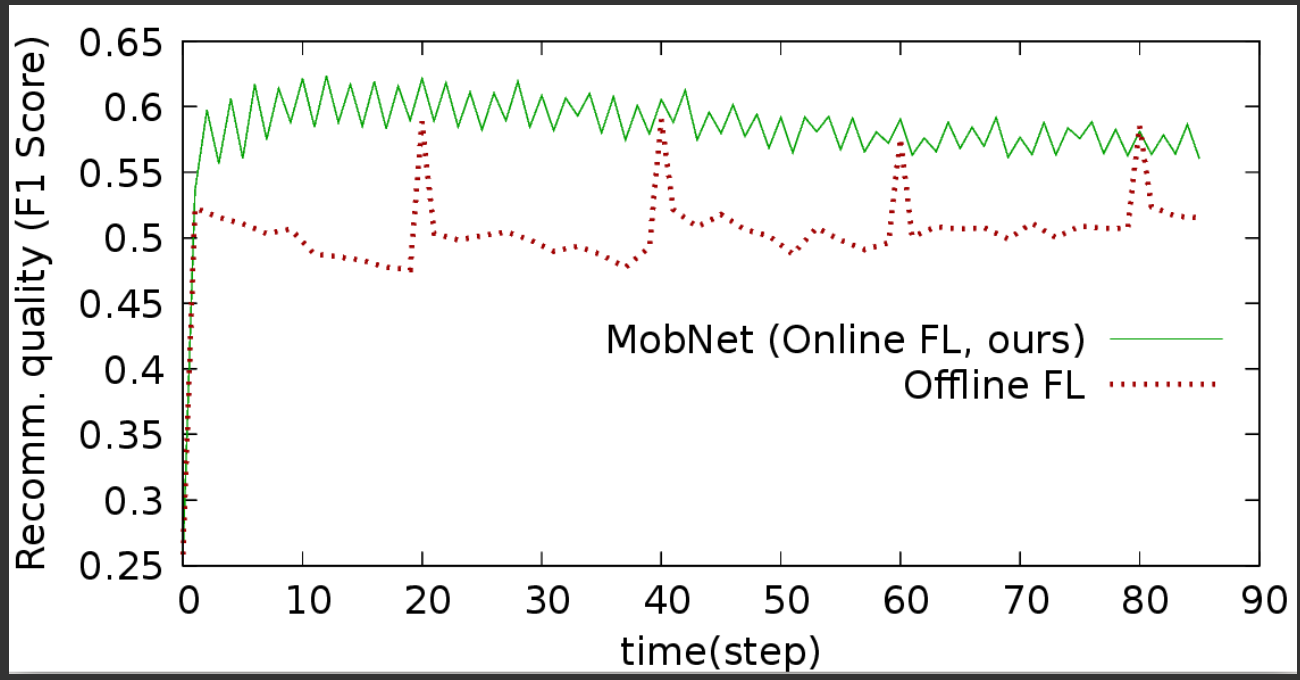

Drop the idle policy

- Gradient computation is compute-intensive
- Drop idle => gradient computation shares resources with the other applications
- AdaSGD: asynchronous rounds and shift decision-making to the mobile device

Drop the charging policy

- Drop charging => we need to predict and control the energy consumption
- Weak devices disrupt the training process

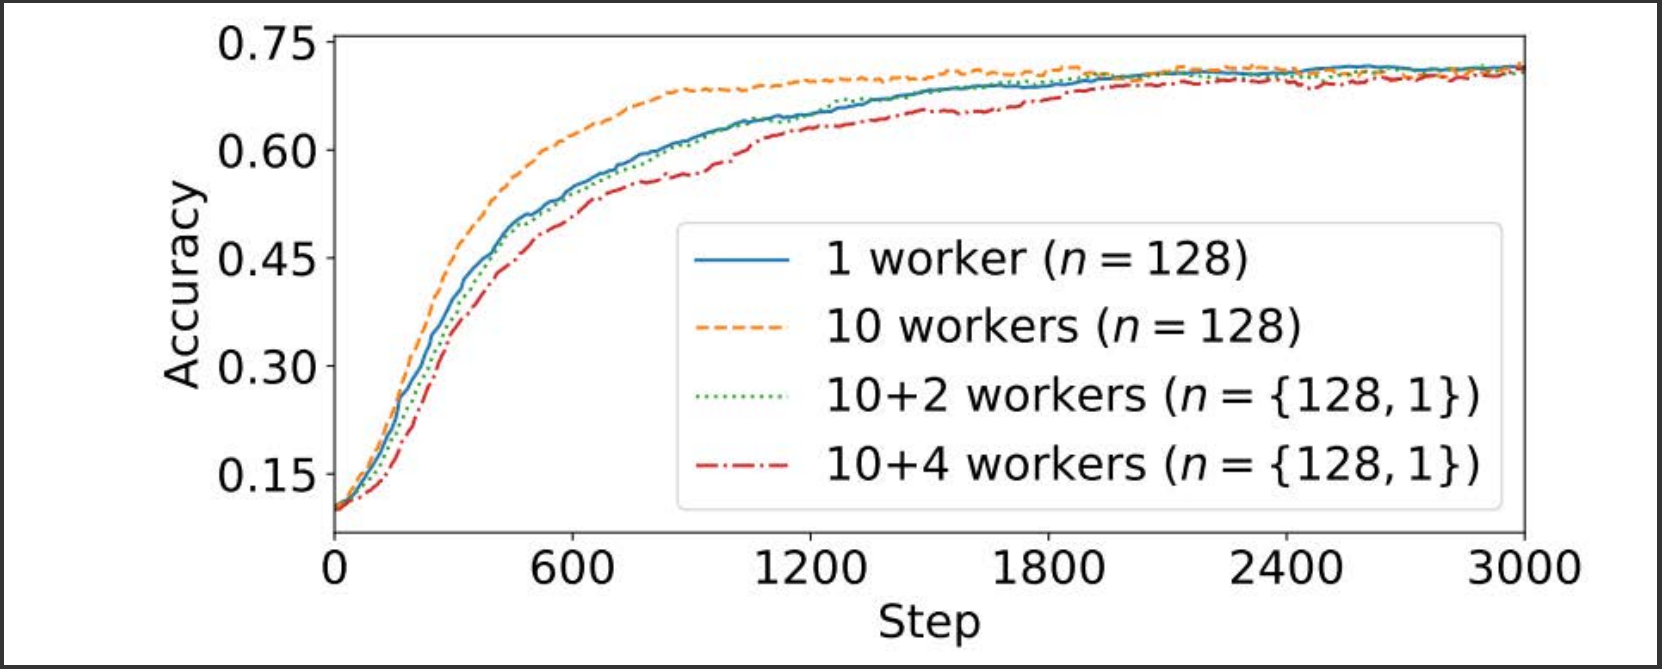

Conclusion

- The state-of-the-art federated learning is not suitable for fast-data applications
- We present MobNet:
  - AdaSGD – a staleness-aware learning algorithm
  - iProf – a regression-based profiler