The University of Texas at Austin

Opportunistic Collaborative Learning (OppCL): Theory, System, and Reality

Haoxiang (Steven) Yu (UT Austin), Sangsu Lee (UT Austin), Christine Julien (UT Austin)

Read & Share

Theory

In this work, we design ***Opportunistic Collaborative Learning* (OppCL)**, which allows resource constrained mobile devices to leverage locally available data, communication, and computation resources to train a machine learning model *on device* with no support of a centralized third party for training or even coordination.

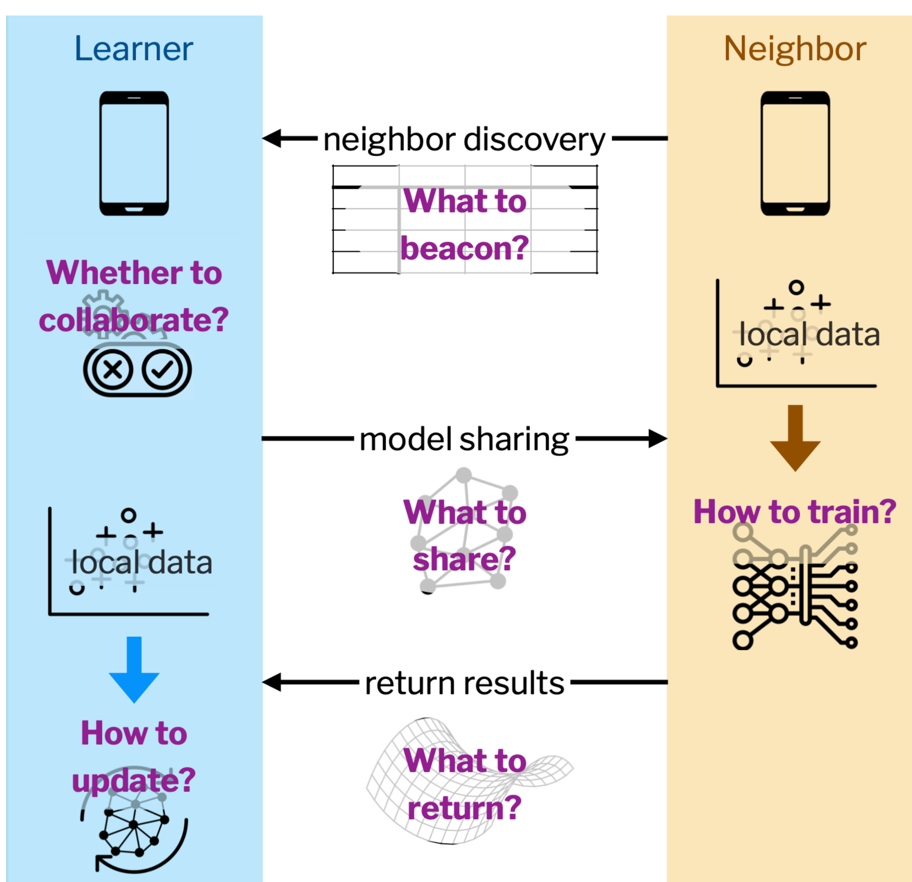

The **OppCL** framework allows devices to (1) *discover* neighbors opportunistically, including discovering the potential benefits of collaborating on a learning task; (2) *share* elements of the learning process across an encounter; (3) *update* a model in response to the collaboration results.

The ***Opportunistic Momentum*** algorithm is one example instantiation of the **OppCL** approach. Devices learn from their encountered neighbors' locally stored data while also remaining resilient to *overfitting* to common data.

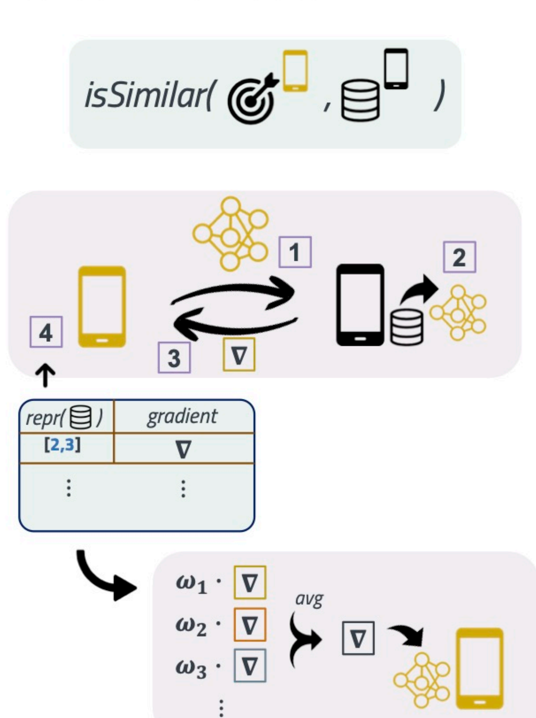

- **Opportunistically asks for gradients**
  - Based on the similarity of goal distribution and a neighbor's local distribution
- **Keeps gradients from prior encounters**
  - Stored with the representation of data distribution from which the gradient is obtained
- **Apply the weighted gradients**
  - Weighted with respect to the similarity of the corresponding data distribution
- **Adjusts learning rate to prevent overfitting**
  - Set a virtual boundary on the loss surface

Controlled evaluations with **OppCL** and **Opportunistic Momentum** demonstrate that they are responsive to learning opportunities presented by encounters and resilient to both non-iid data and to fluctuations in the available data sets. Devices can also completely personalize their needs; neighbor devices can learn completely different models based on their needs and encounters.

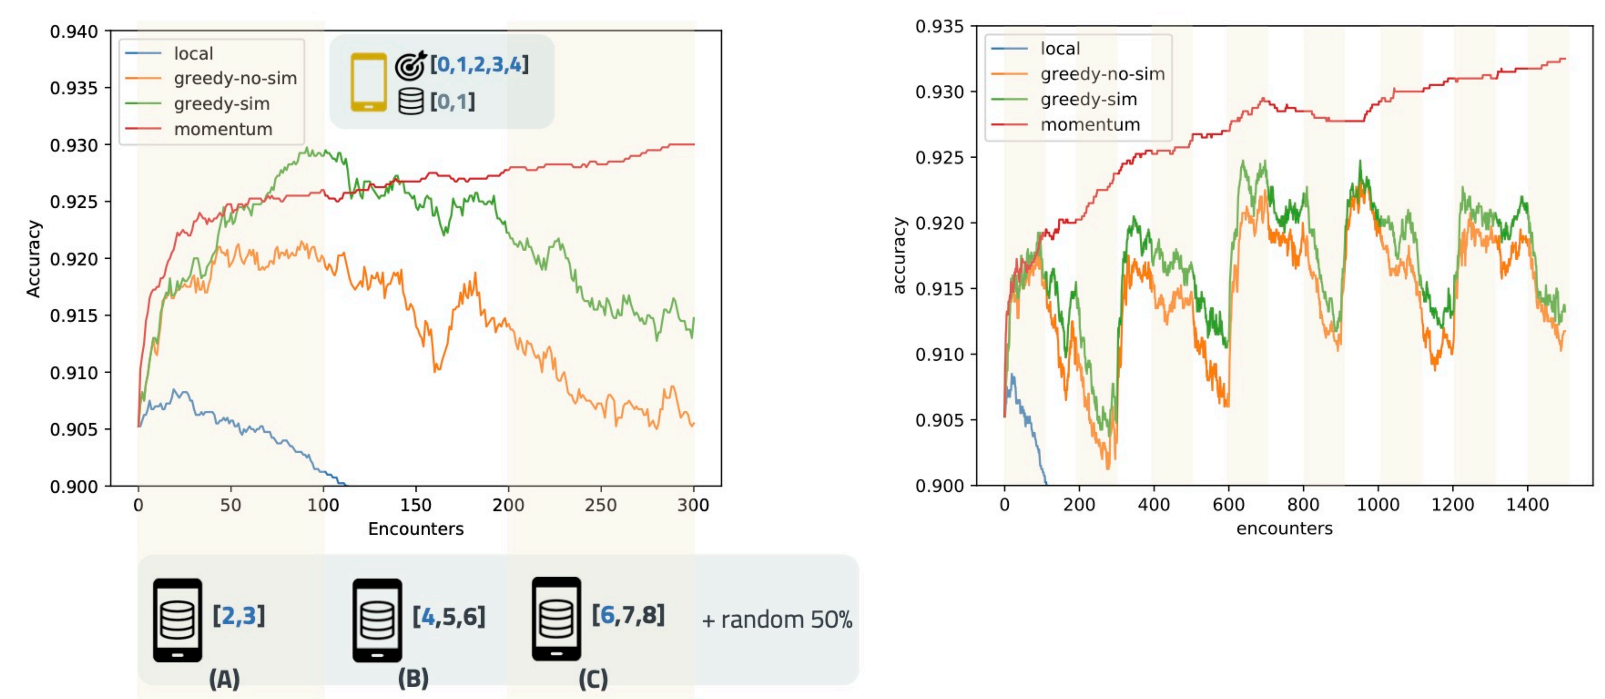

Next steps include considering heterogenous resource requirements (i.e., devices with different computational or communication capabilities) and dynamic application model demands (i.e., when devices' personalized goals change over time, the model should change with them).

System

Large-scale evaluation of decentralized learning, in general, is complex and difficult to support. We are building both **middleware** to support our approaches on-device and a large-scale **simulator** that can directly execute decentralized learning implementations on hundreds or thousands of emulated nodes while also simulating their real-world encounter patterns.

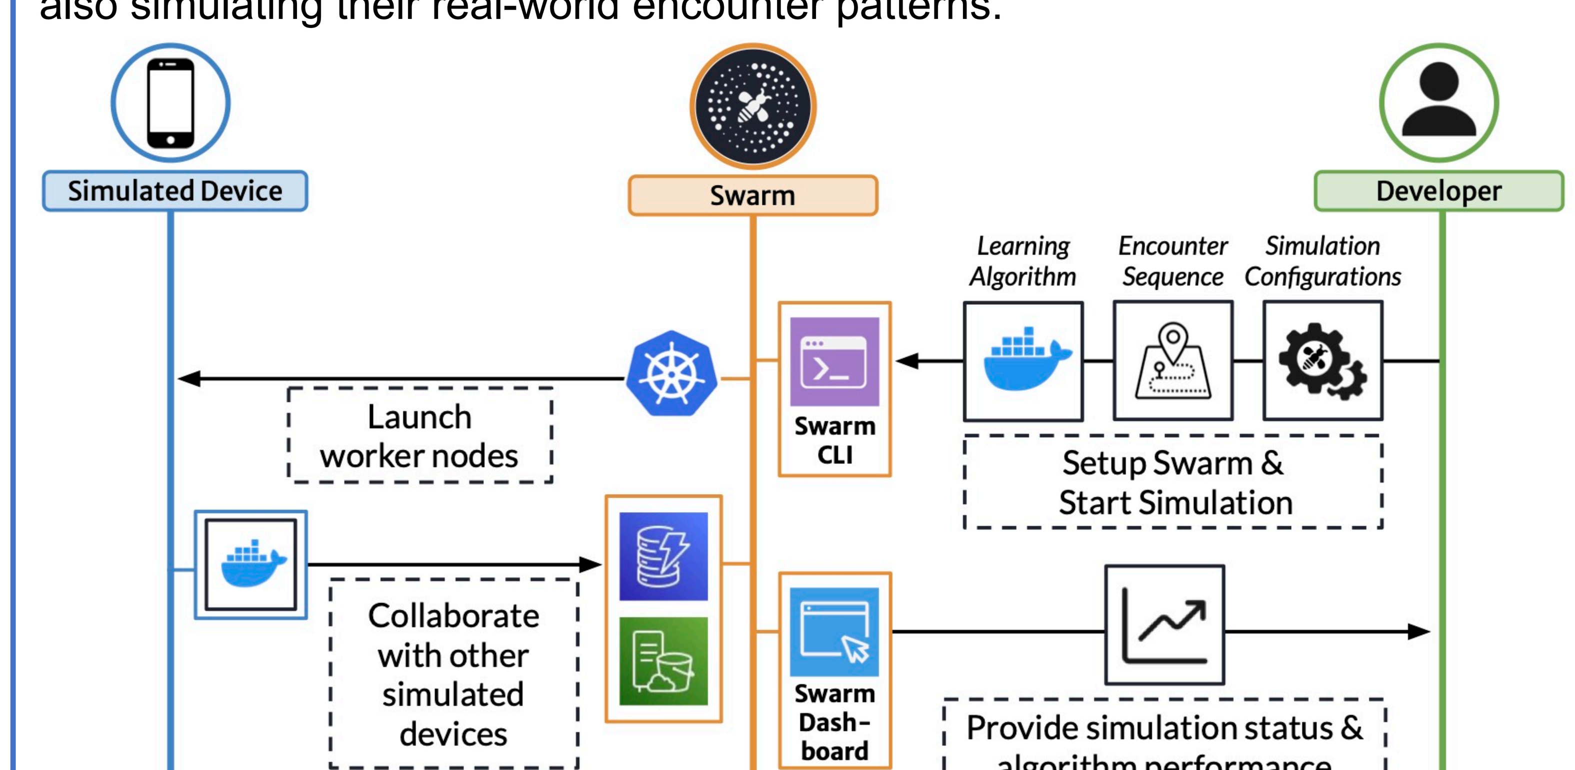

Reality

Ultimately, deploying **OppCL** and other decentralized learning models requires them to function on real hardware with real constraints. These concerns are commonly ignored by the theoretical work in this space. We consider memory constraints, communication constraints, energy costs, and operating system limitations in deploying **OppCL** approaches on smartphones and IoT devices.

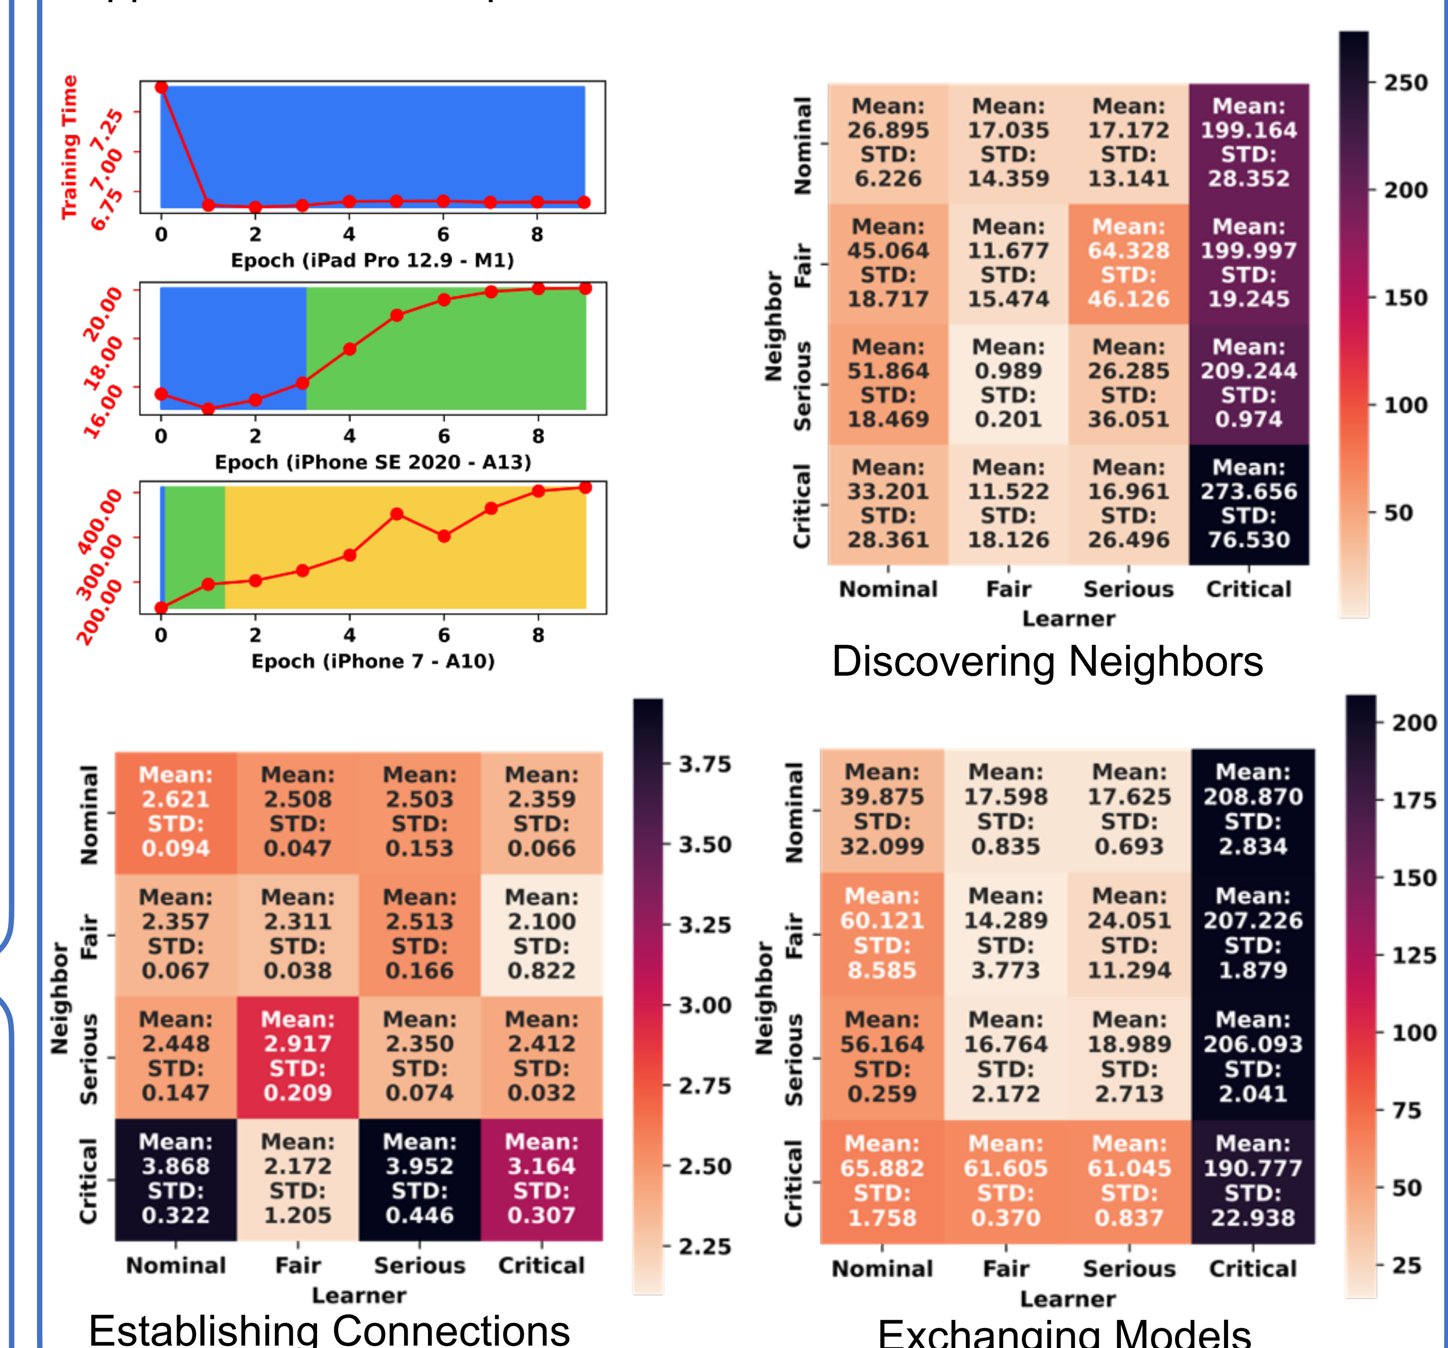

References

1. H. Yu, H. -Y. Chen, S. Lee, X. Zheng and C. Julien, "Prototyping Opportunistic Learning in Resource Constrained Mobile Devices," 2022 IEEE International Conference on Pervasive Computing and Communications Workshops and other Affiliated Events (PerCom Workshops), 2022, pp. 521-526,
2. S. Lee, H. Yu, X. Zheng and C. Julien, "Swarm: Playground for Large-scale Decentralized Learning Simulations," 2022 IEEE International Conference on Pervasive Computing and Communications Workshops and other Affiliated Events (PerCom Workshops), 2022, pp. 115-117.
3. S. Lee, X. Zheng, J. Hua, H. Vikalo and C. Julien, "Opportunistic Federated Learning: An Exploration of Egocentric Collaboration for Pervasive Computing Applications," *2021 IEEE International Conference on Pervasive Computing and Communications (PerCom)*, 2021, pp. 1-8,

Mobile & Pervasive Computing Group