Imperial College London

Privacy-Preserving Personal Model Training

Hamed Haddadi

Imperial College London

Joint work with: Sandra Servia-Rodriguez, Liang Wang, Jianxin R. Zhao, Richard Mortier

Imperial College London

Machine Learning today

- Get Users, collect loads of labelled data, train model, at a Data centre
- Privacy, cost, personalization challenges
- Think about sensitive images on FB, and voices on Echo, accountability, responsibility, GDPR, …

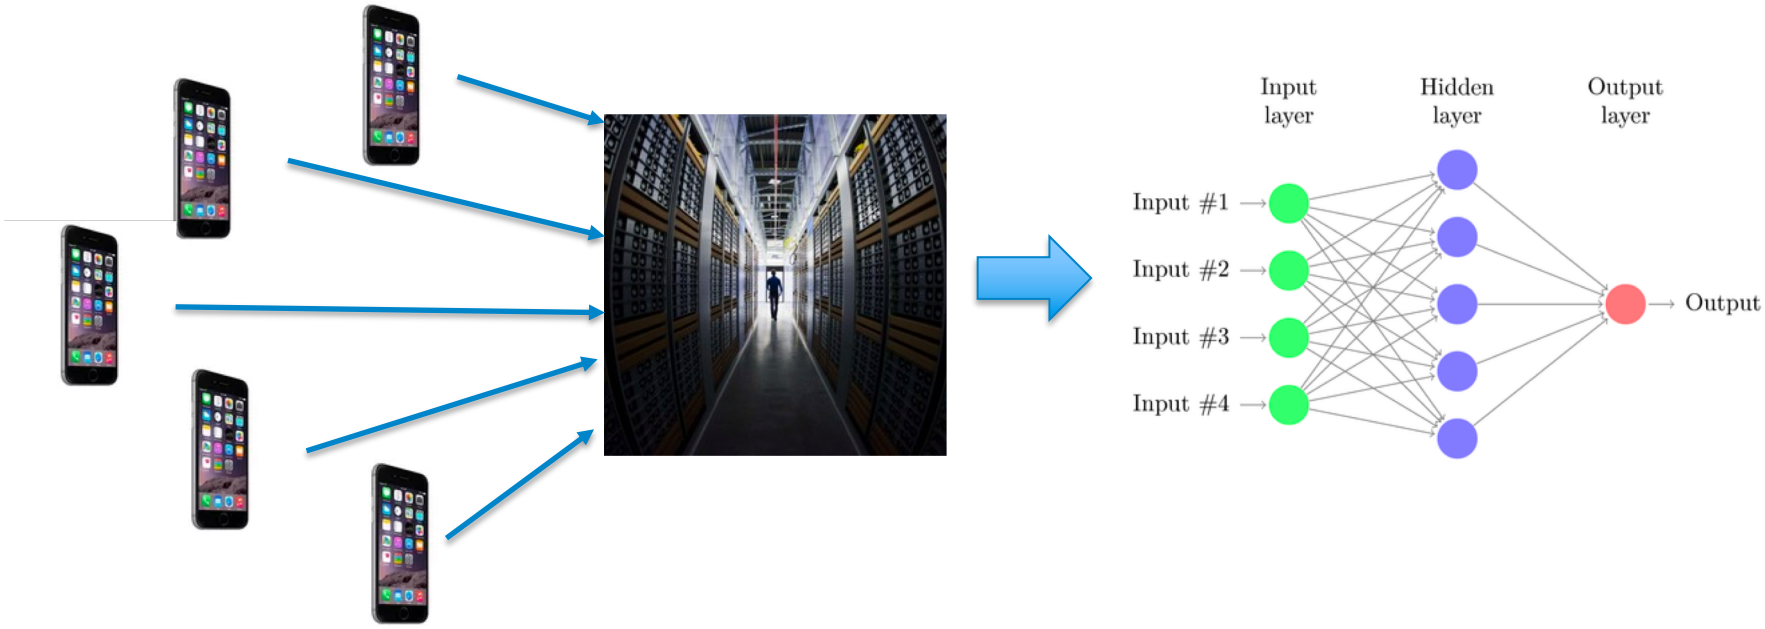

Imperial College London

Privacy during Usage of ML models

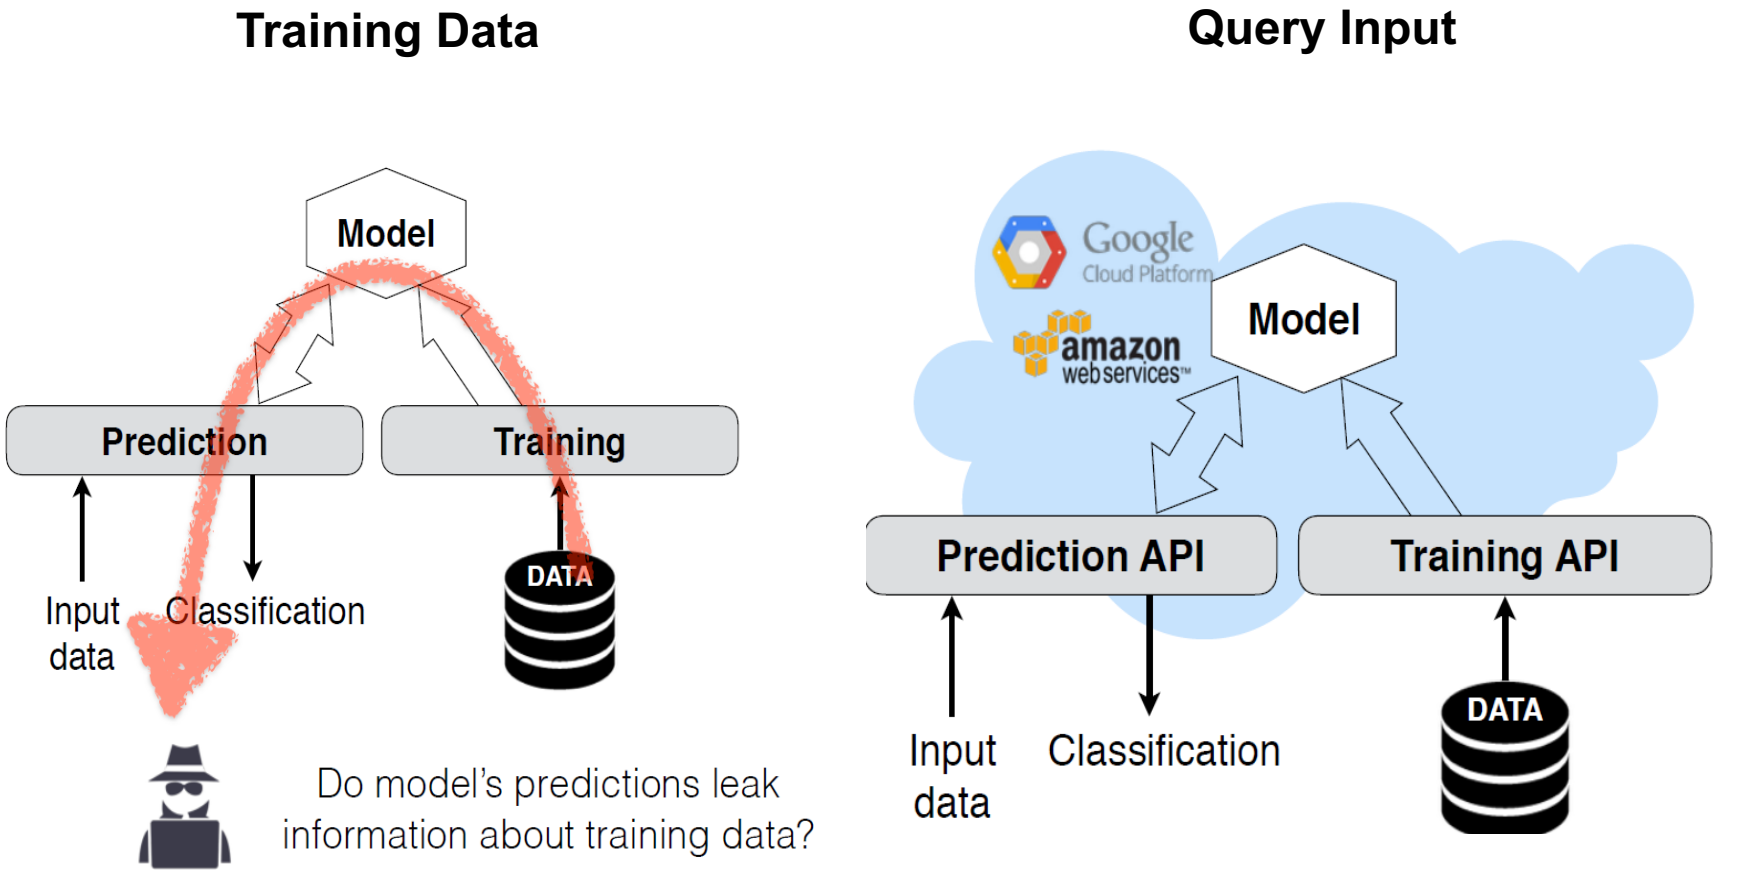

Imperial College London

Distributed Learning

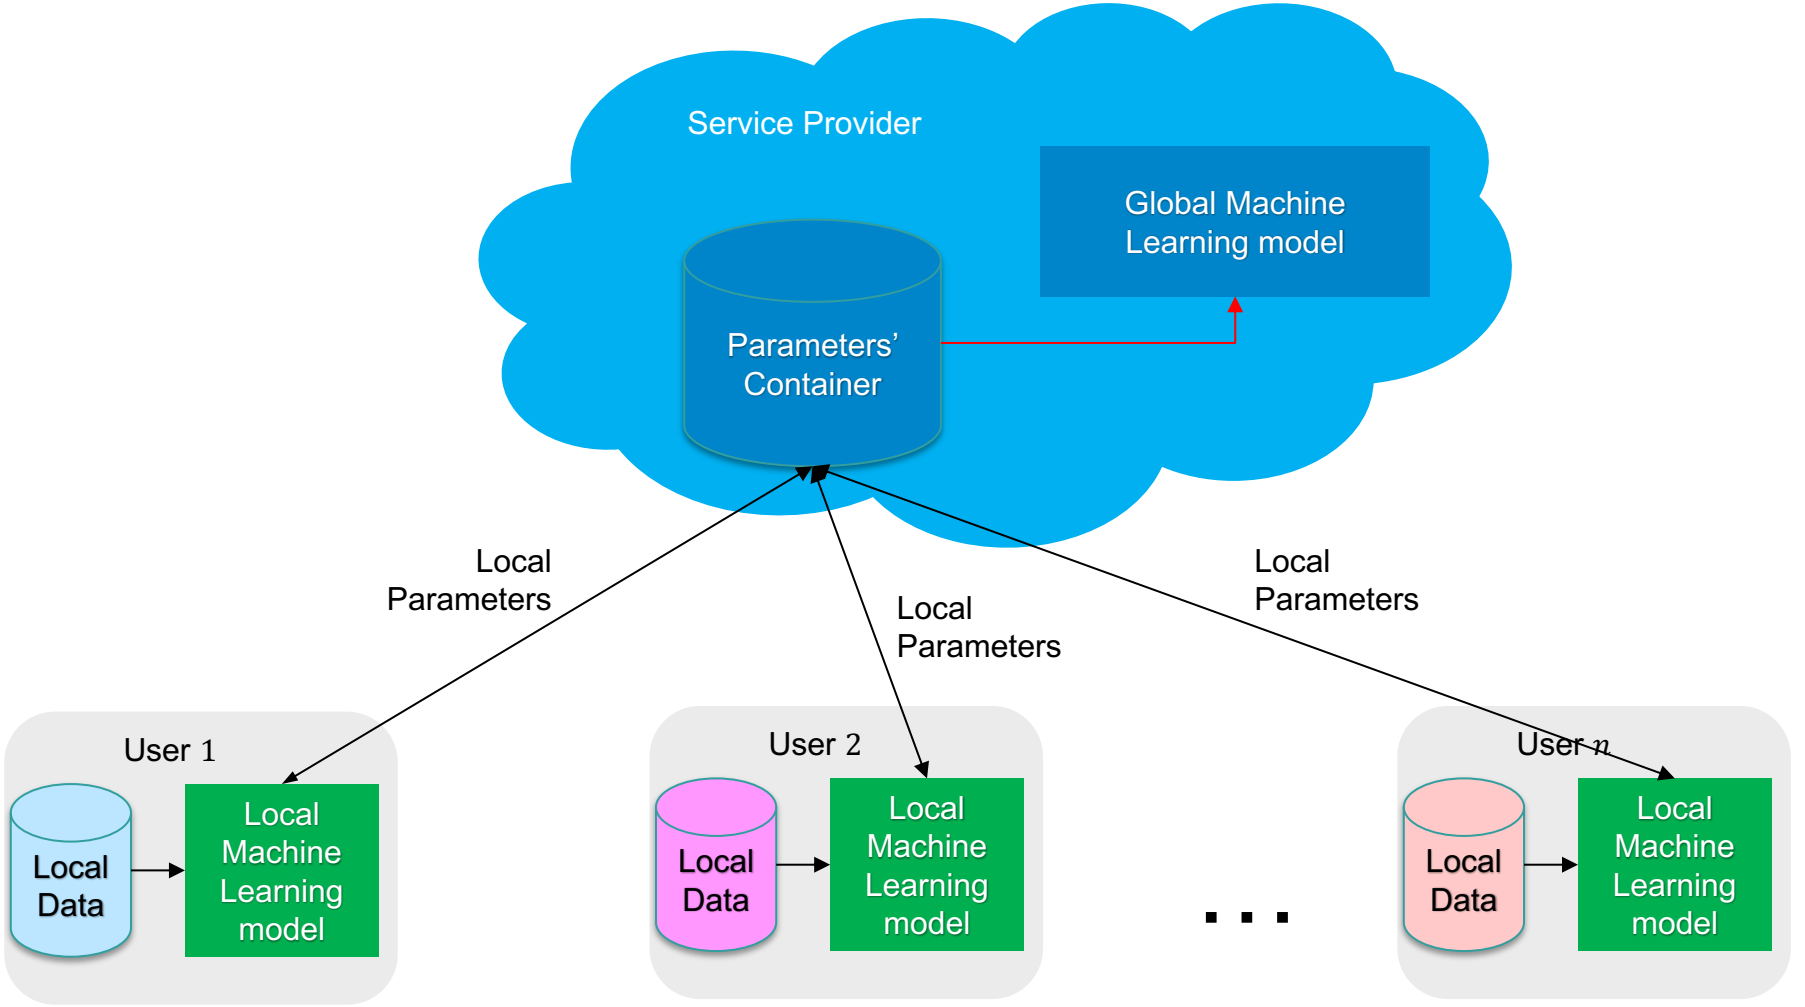

Imperial College London

Federated Learning

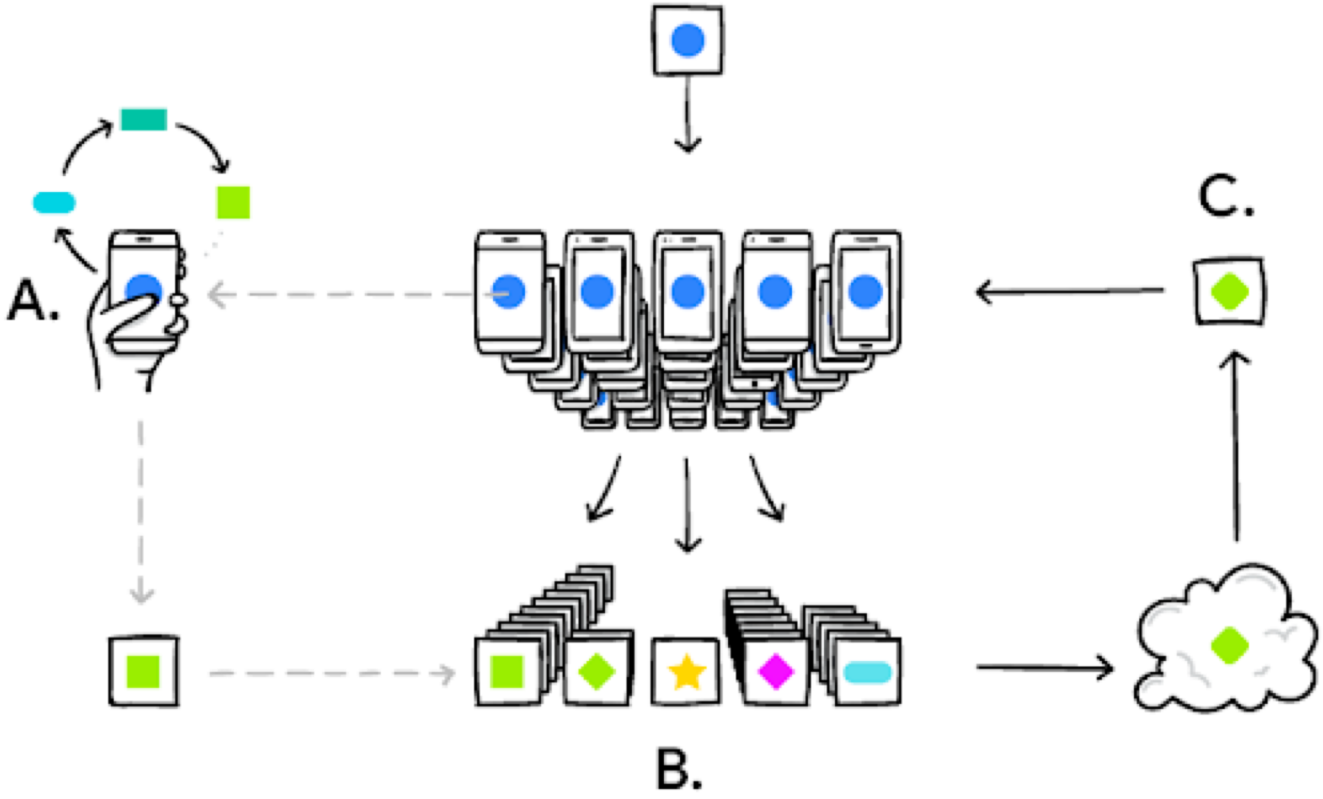

Federated Learning: Collaborative Machine Learning without Centralized Training Data (Google 2017)

Imperial College London

Databox

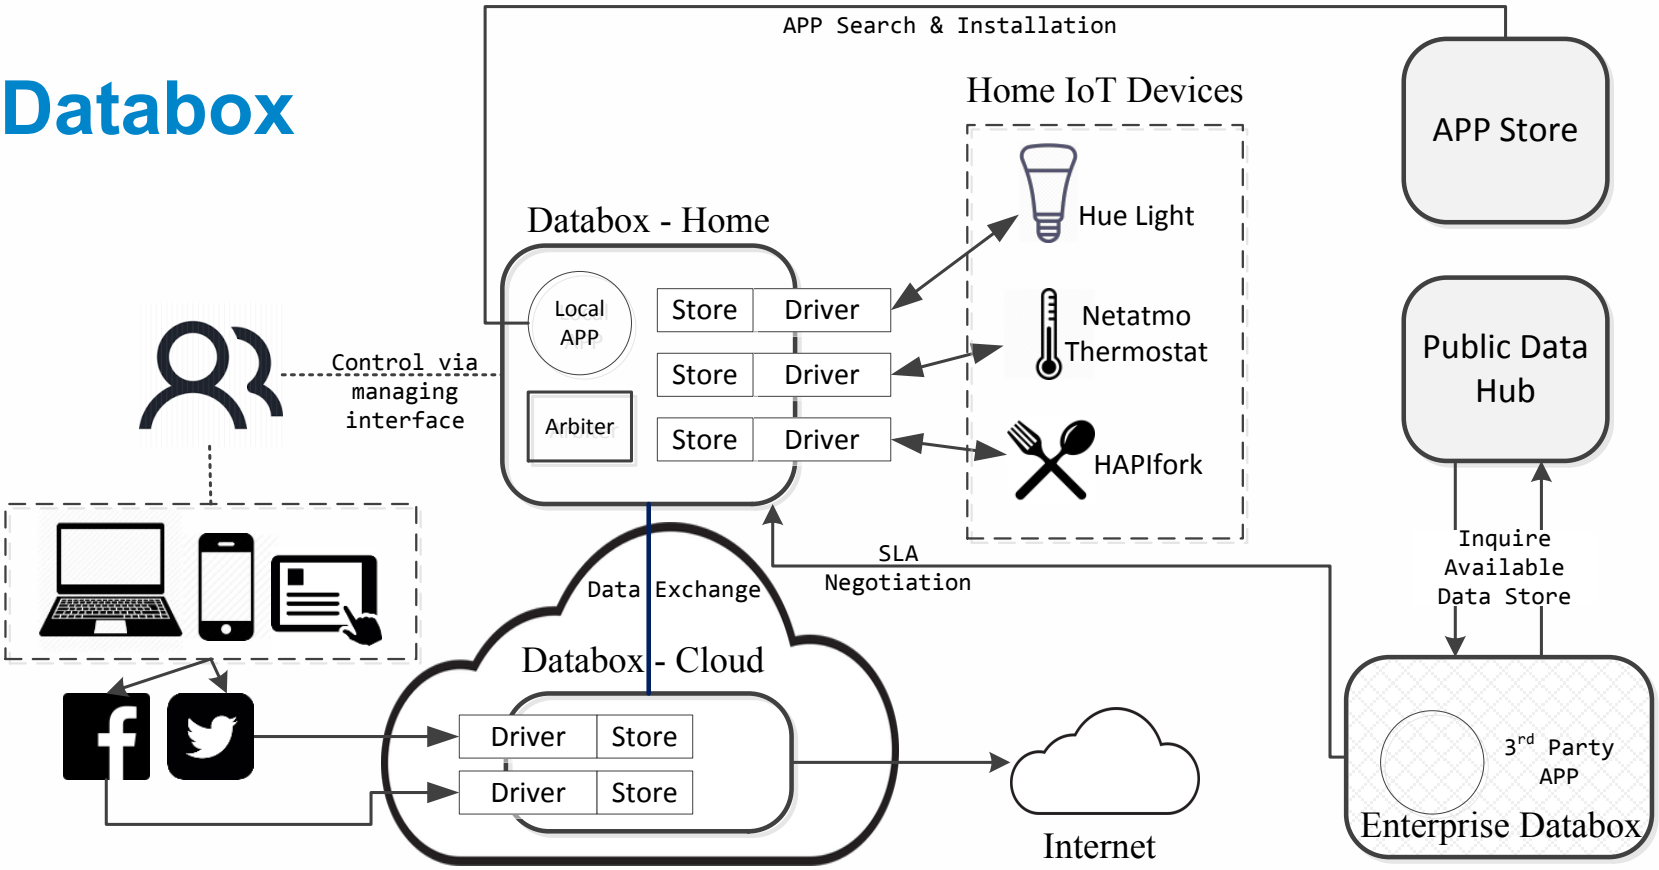

- See the details on www.databoxproject.uk & github.com/me-box

Imperial College London

Dataset 1

- Motion-based activity classifier on smartphone without revealing their data to others.
- *WISDM* Human Activity Recognition dataset, accelerometer data on an Android phone by 35 subjects performing 6 activities (*walking*, *jogging*, *walking upstairs*, *walking downstairs*, *sitting* and *standing*).
- Statistical measures obtained for every 10 seconds of accelerometer samples as the $d = 43$ dimensional features in our models.

Final sample: 5,418 accelerometer traces from 35 users, with on average 150.50 traces per user and standard deviation of 44.73.

Imperial College London

privacy-preserving activity recognition

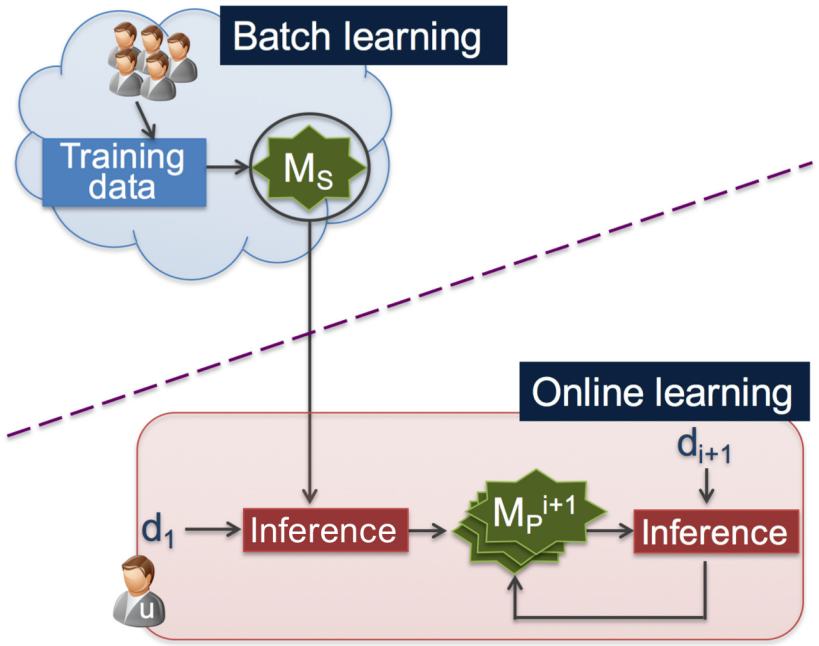

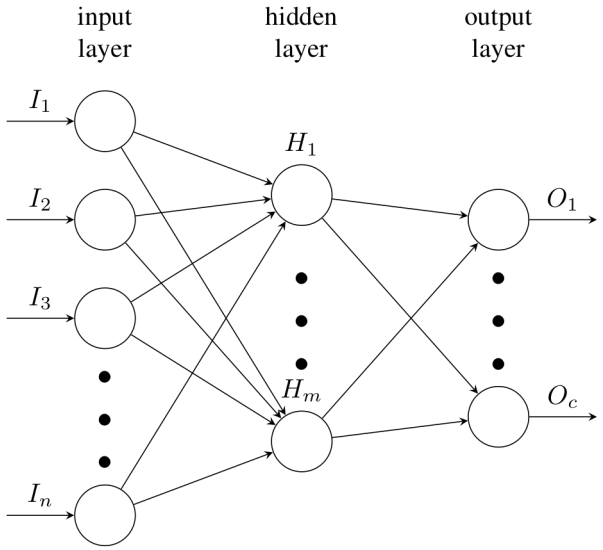

Two-layer feed-forward MLP

- Multilayer Perceptron with 2 layers for activity recognition, 1 hidden layer with 128 nodes, 1 logistic regression layer.
- 6, 406 parameters to be determined during training.
- All on Raspberry Pi 3 Model B

Imperial College London

More samples, good samples!

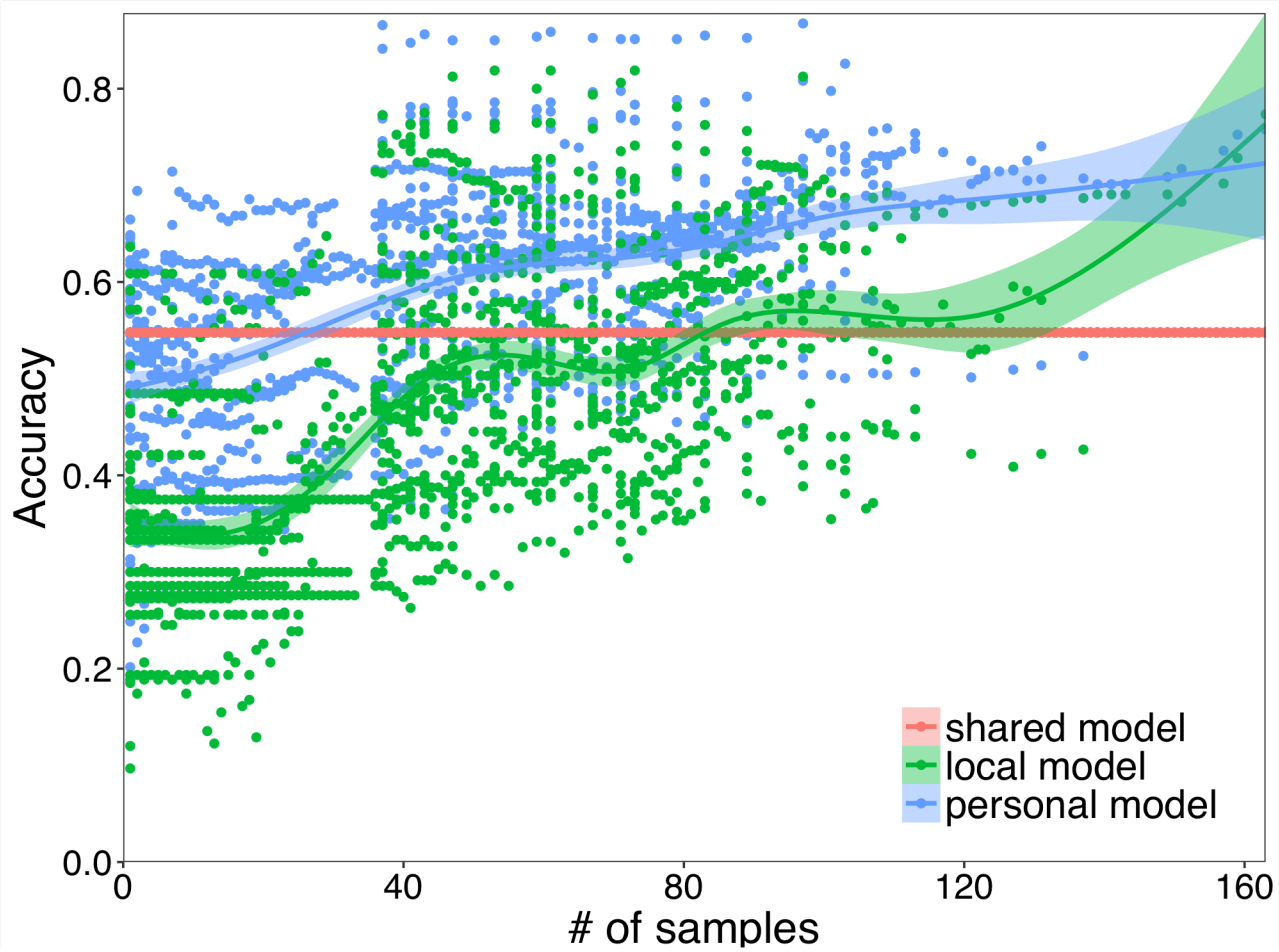

Imperial College London

Topic modelling from Wikipedia using LDA

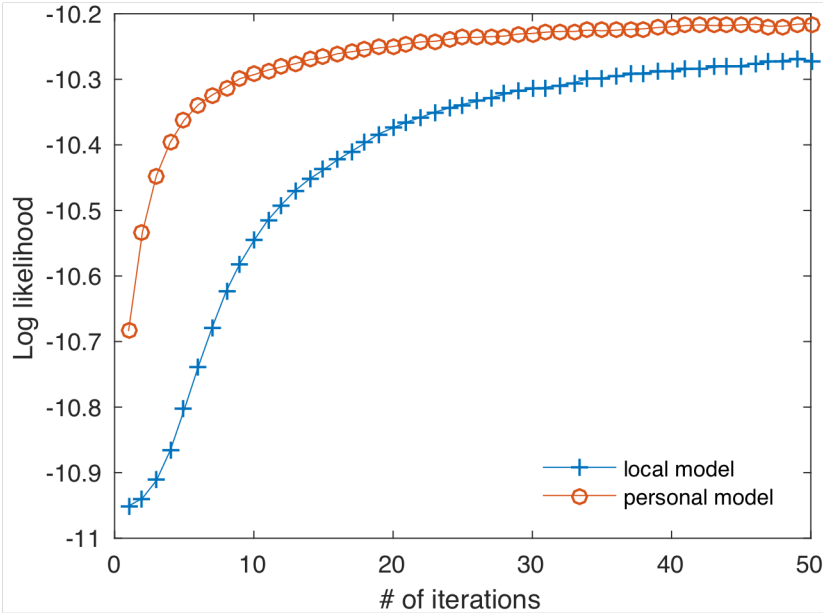

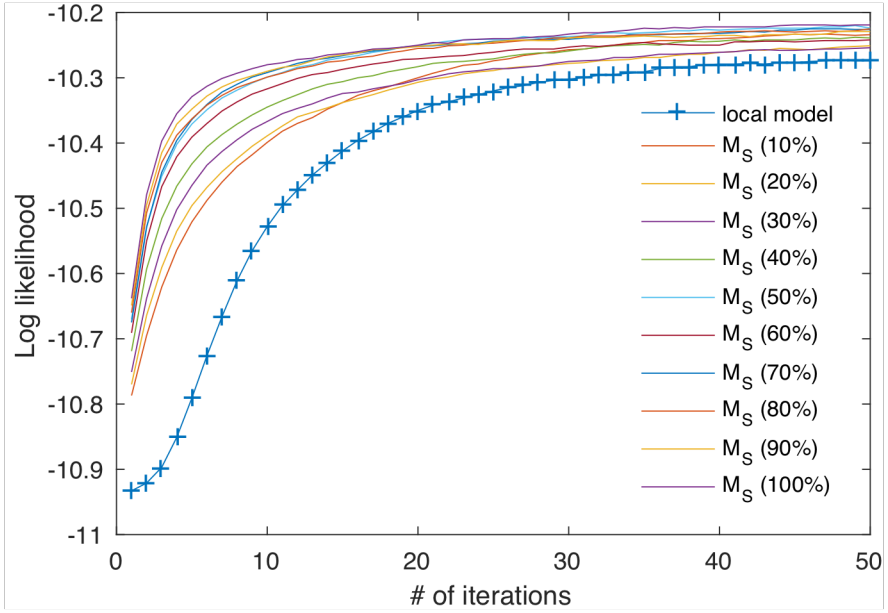

Using OWL Library: https://github.com/owlbarn

Imperial College London

Training time

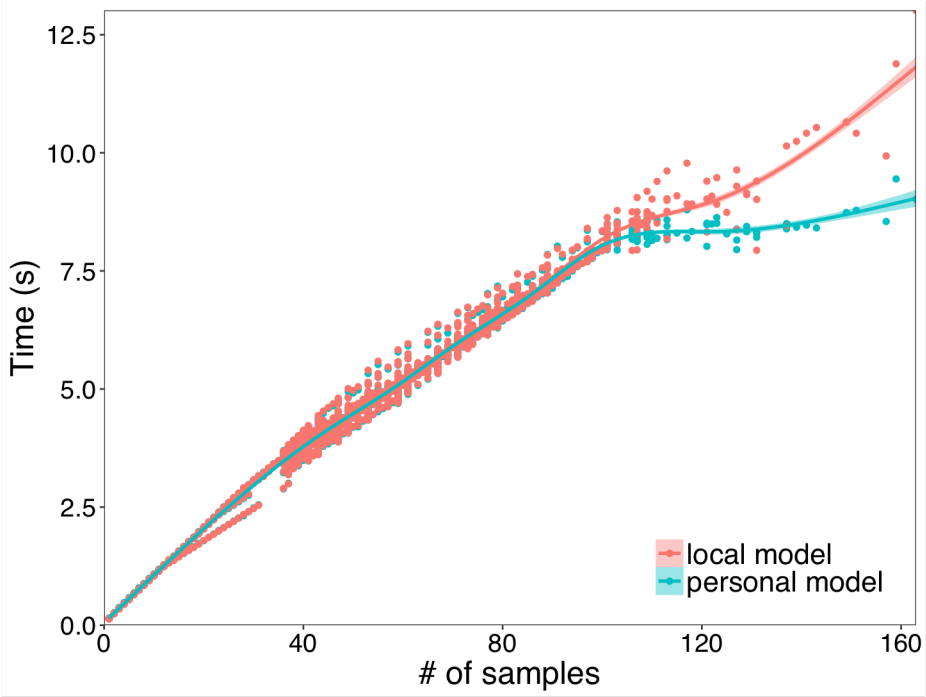

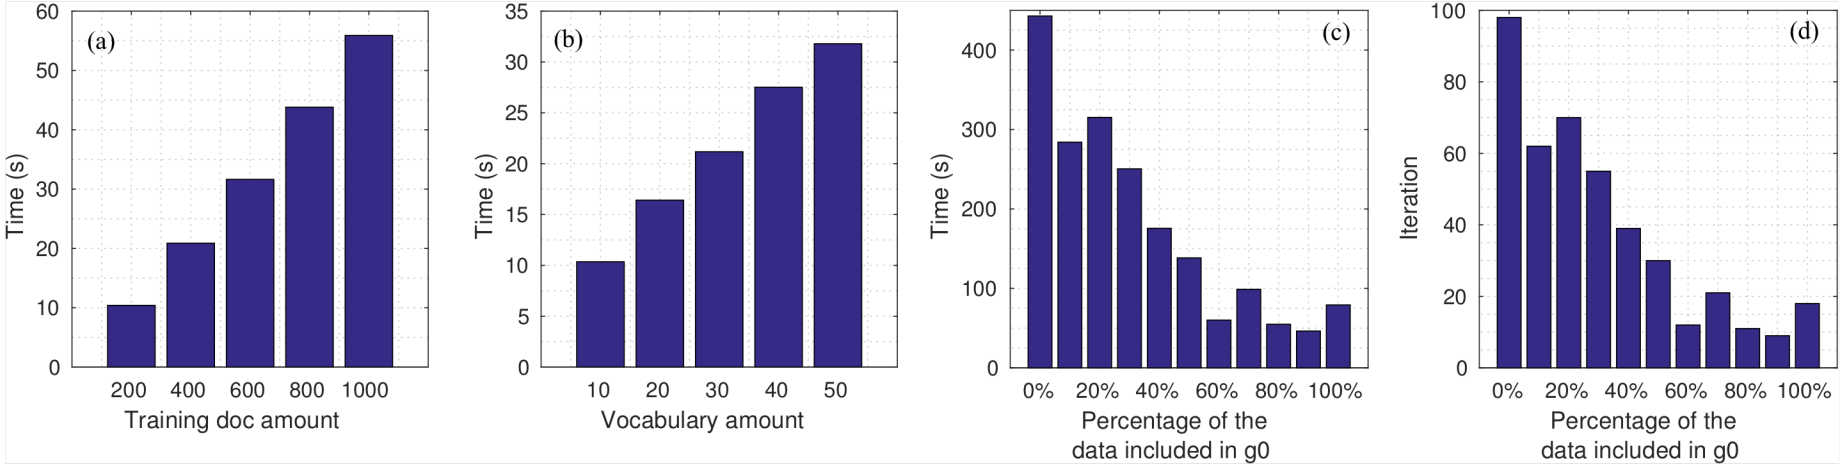

Imperial College London

Thank You!

Privacy-Preserving Personal Model Training

We are looking for postdocs and PhD students!

Hamed Haddadi

Imperial College London

 https://haddadi.github.io/

 h.haddadi@imperial.ac.uk

 @realhamed