Orchestra: Unsupervised Federated Learning via Globally Consistent Clustering

By: Ekdeep Singh Lubana[1], Chi Ian Tang[2], Fahim Kawsar[3], Robert P. Dick[1], and Akhil Mathur[3]

[1]University of Michigan, [2]University of Cambridge, [3]Nokia Bell Labs (UK)

UNIVERSITY OF MICHIGAN

UNIVERSITY OF CAMBRIDGE

NOKIA Bell Labs

Smart Devices: Perpetual Data Generators

Federated Learning: Collaborative Training

- Goal: Distributed, privacy-preserving, resource-efficient learning

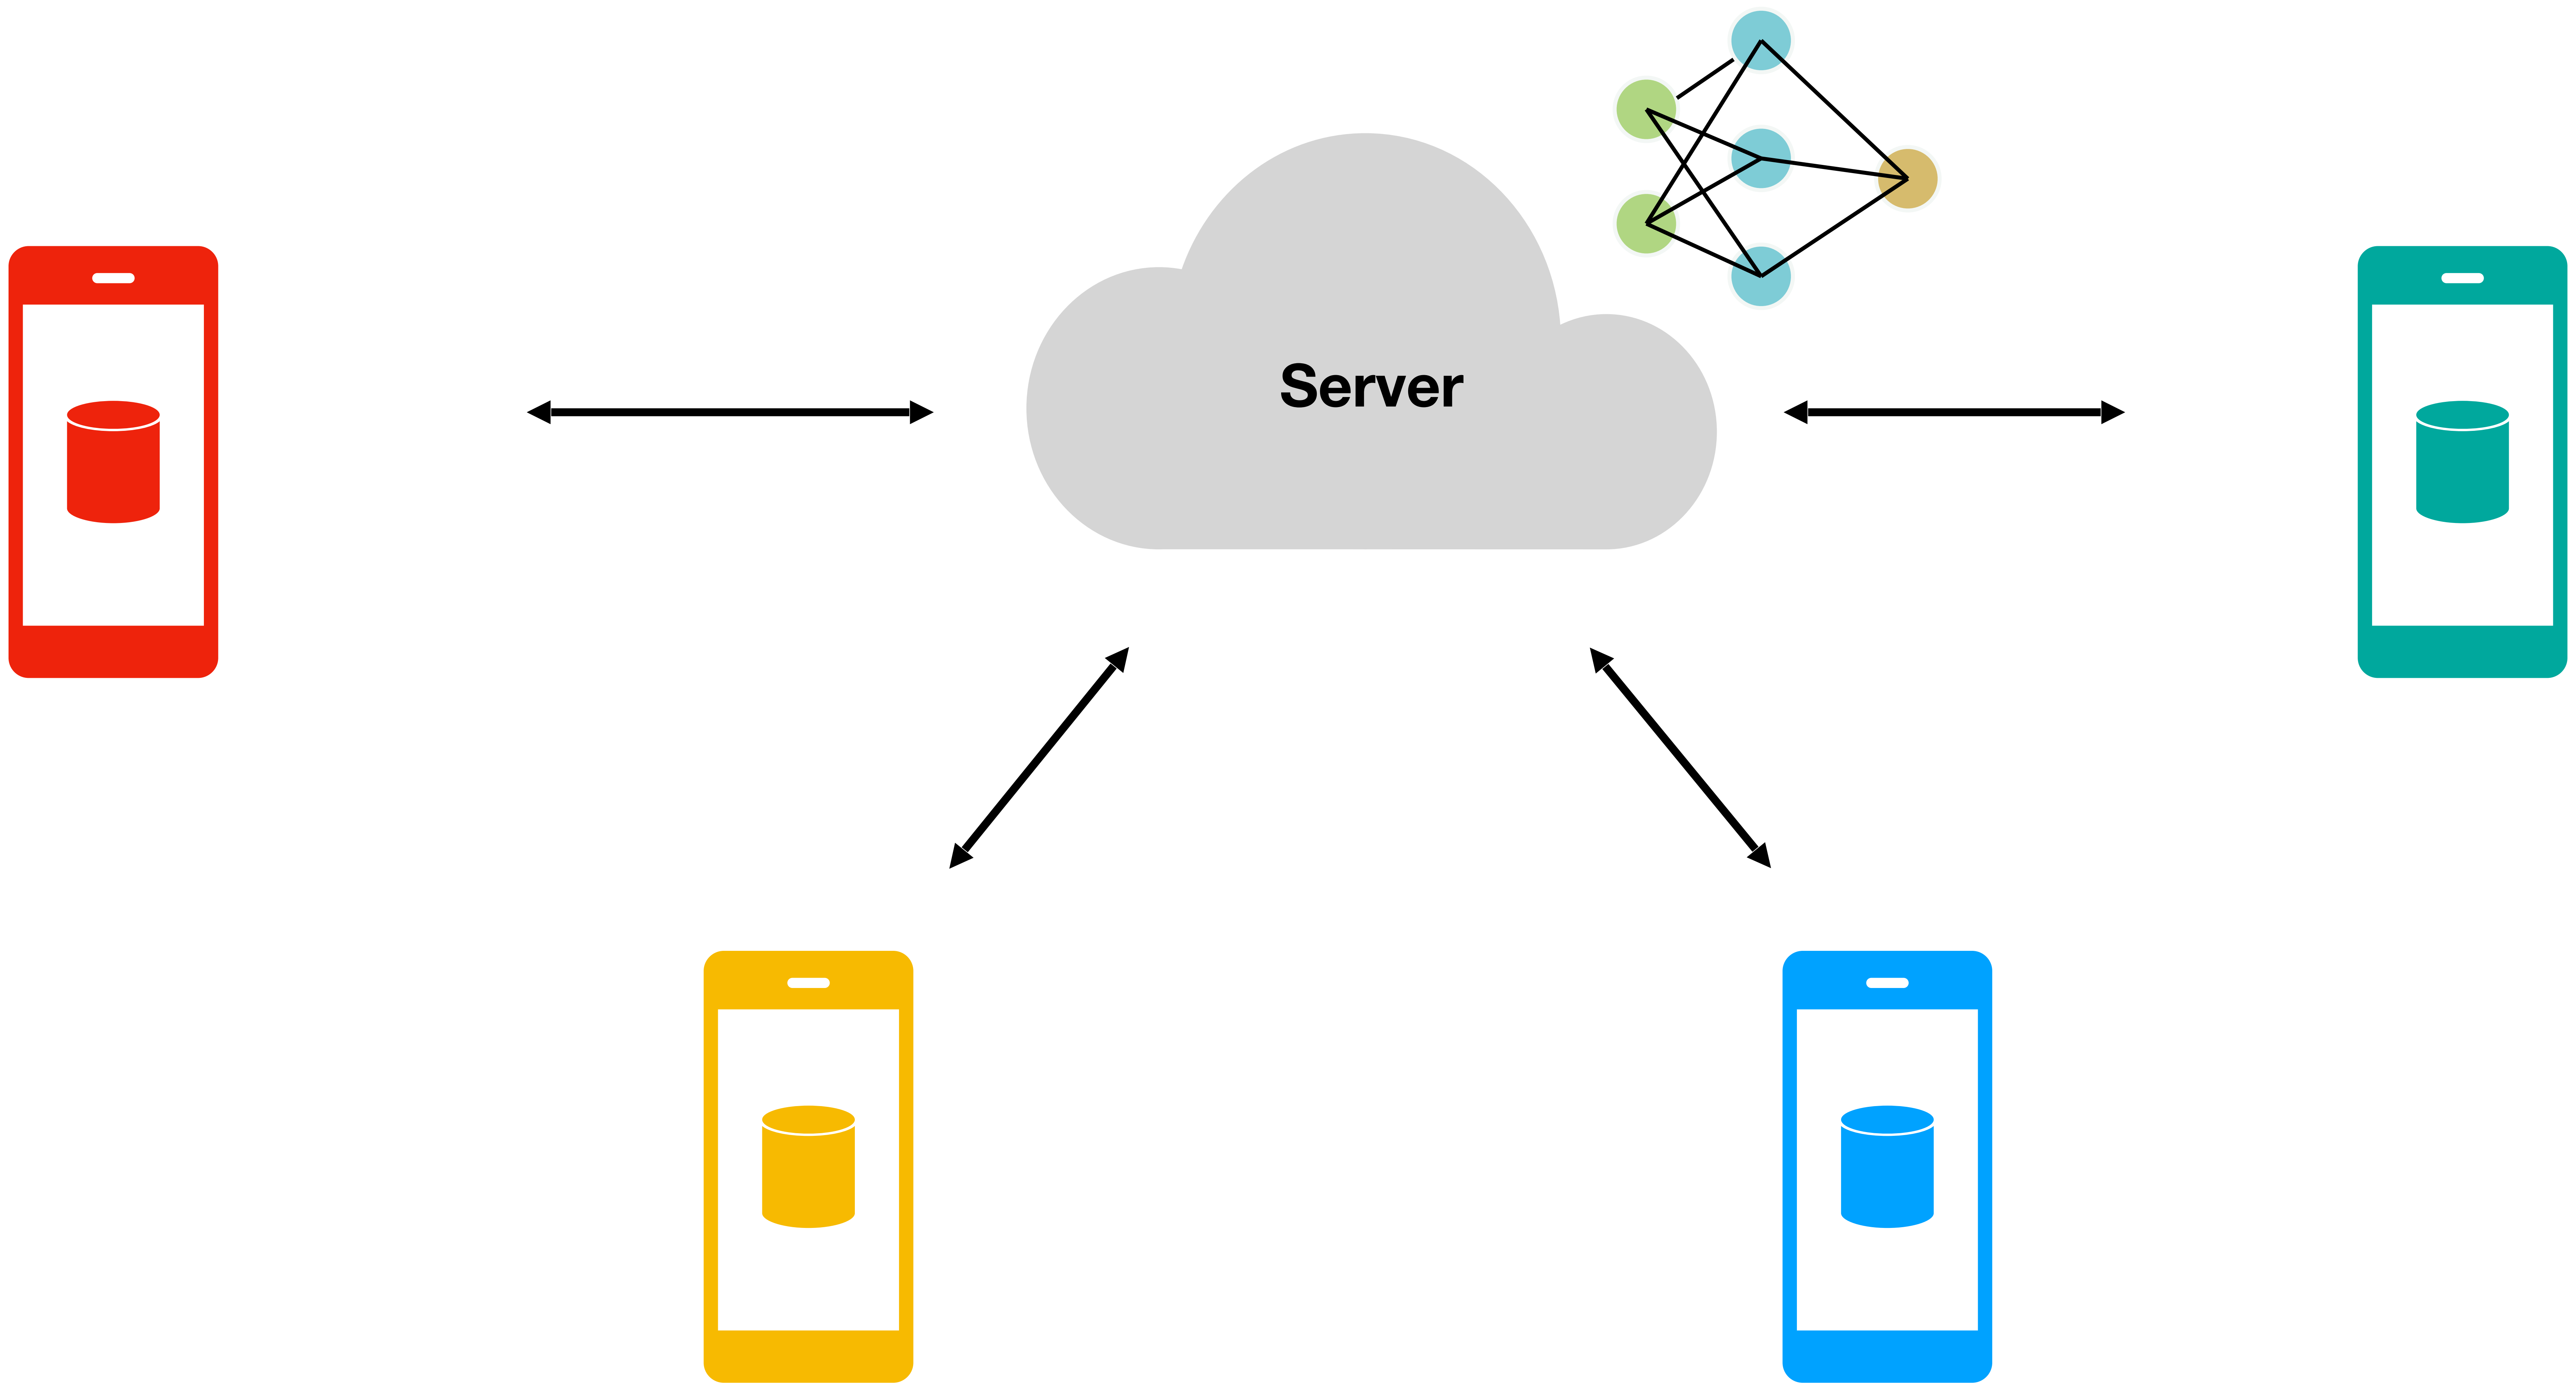

Step 1: Share Model with Clients

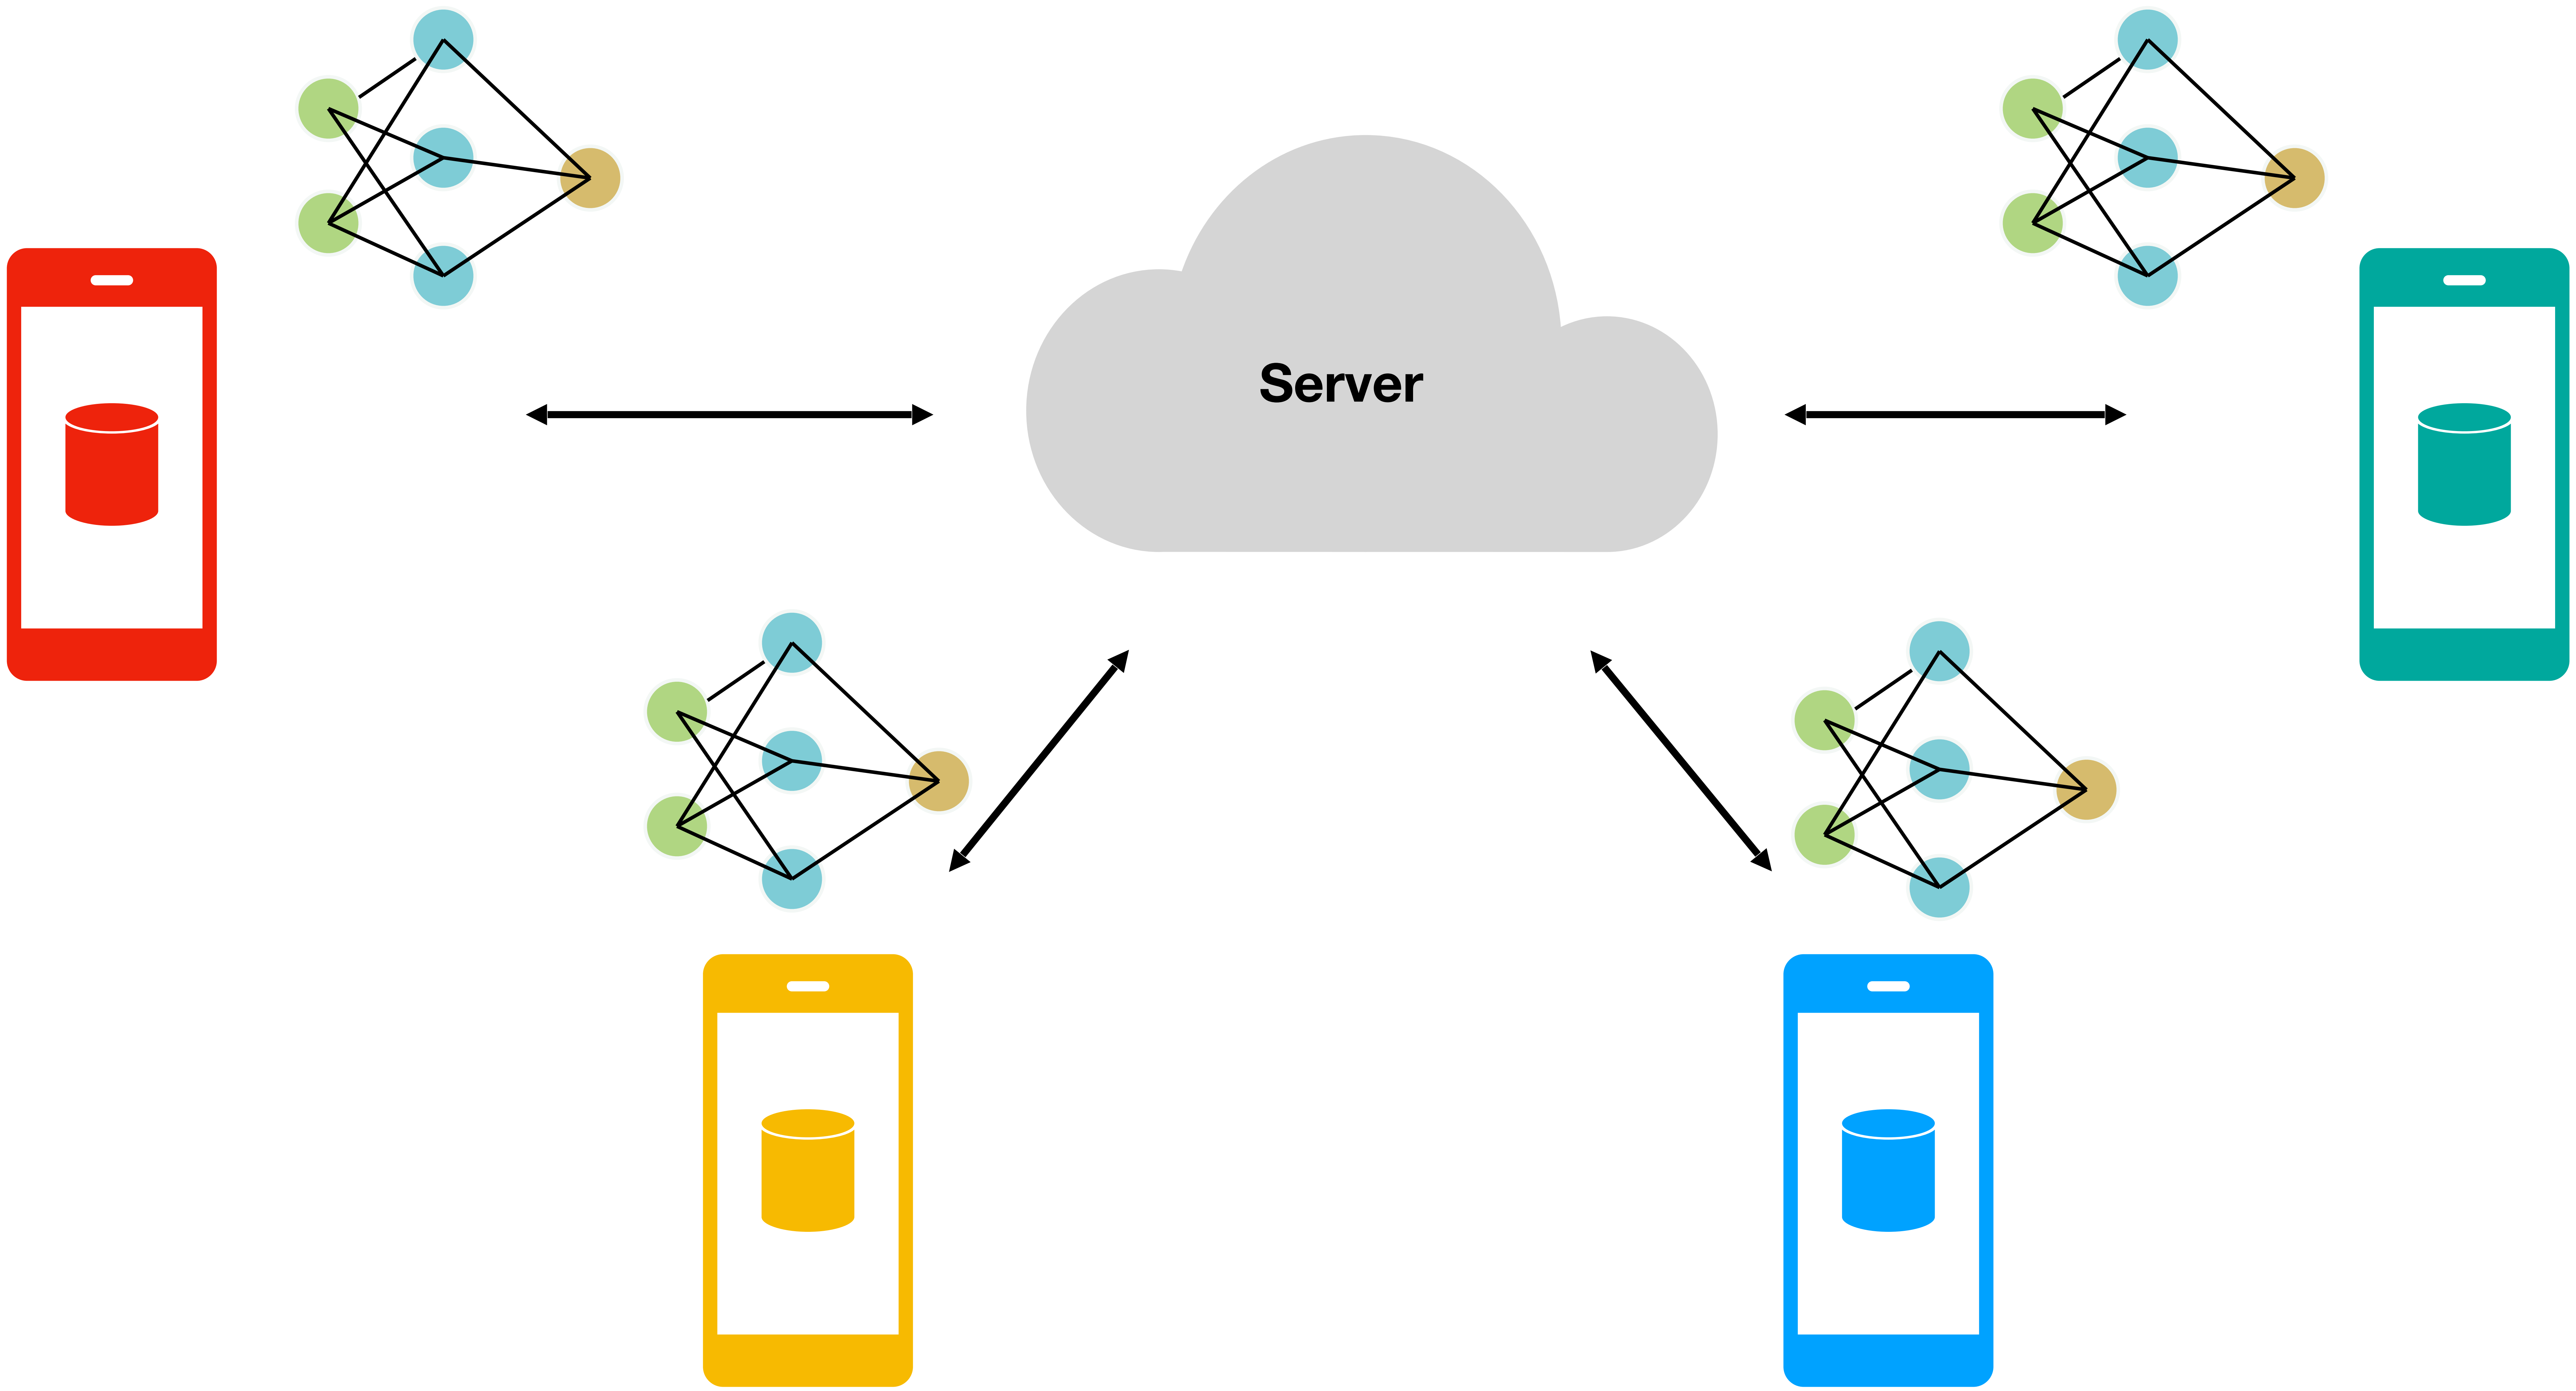

Step 2: Train Models on Local Data

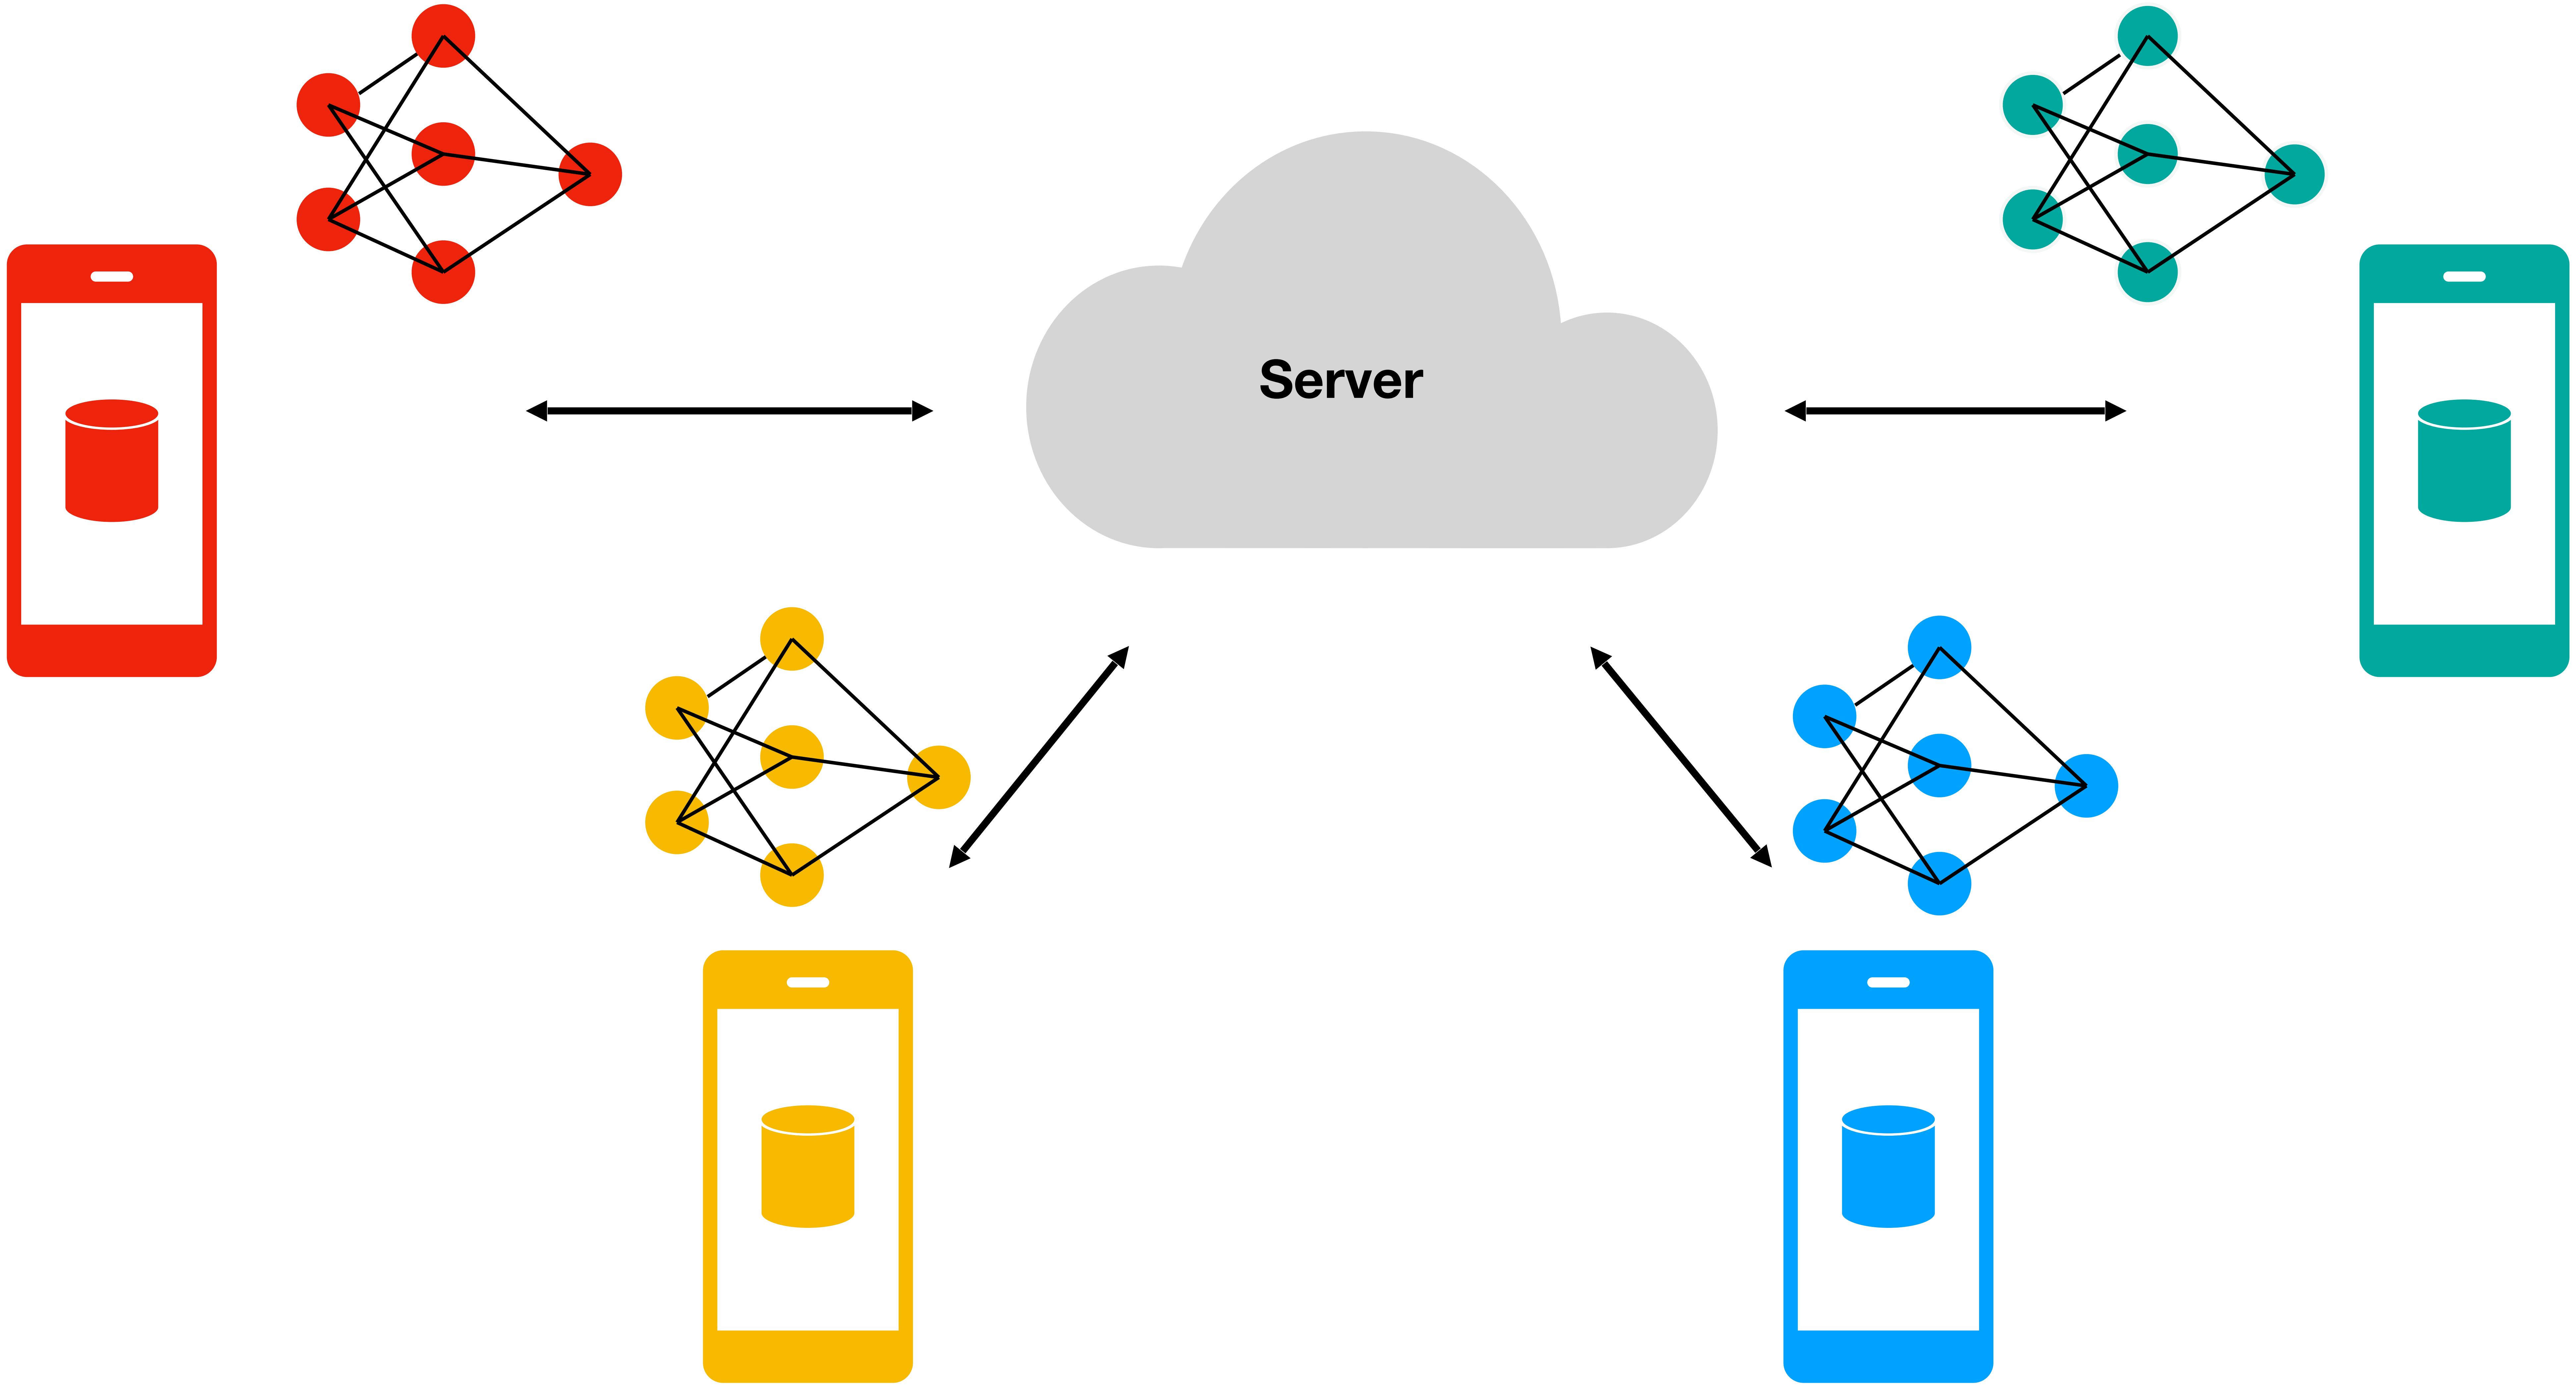

Step 3: Aggregate Models in a Secure Manner

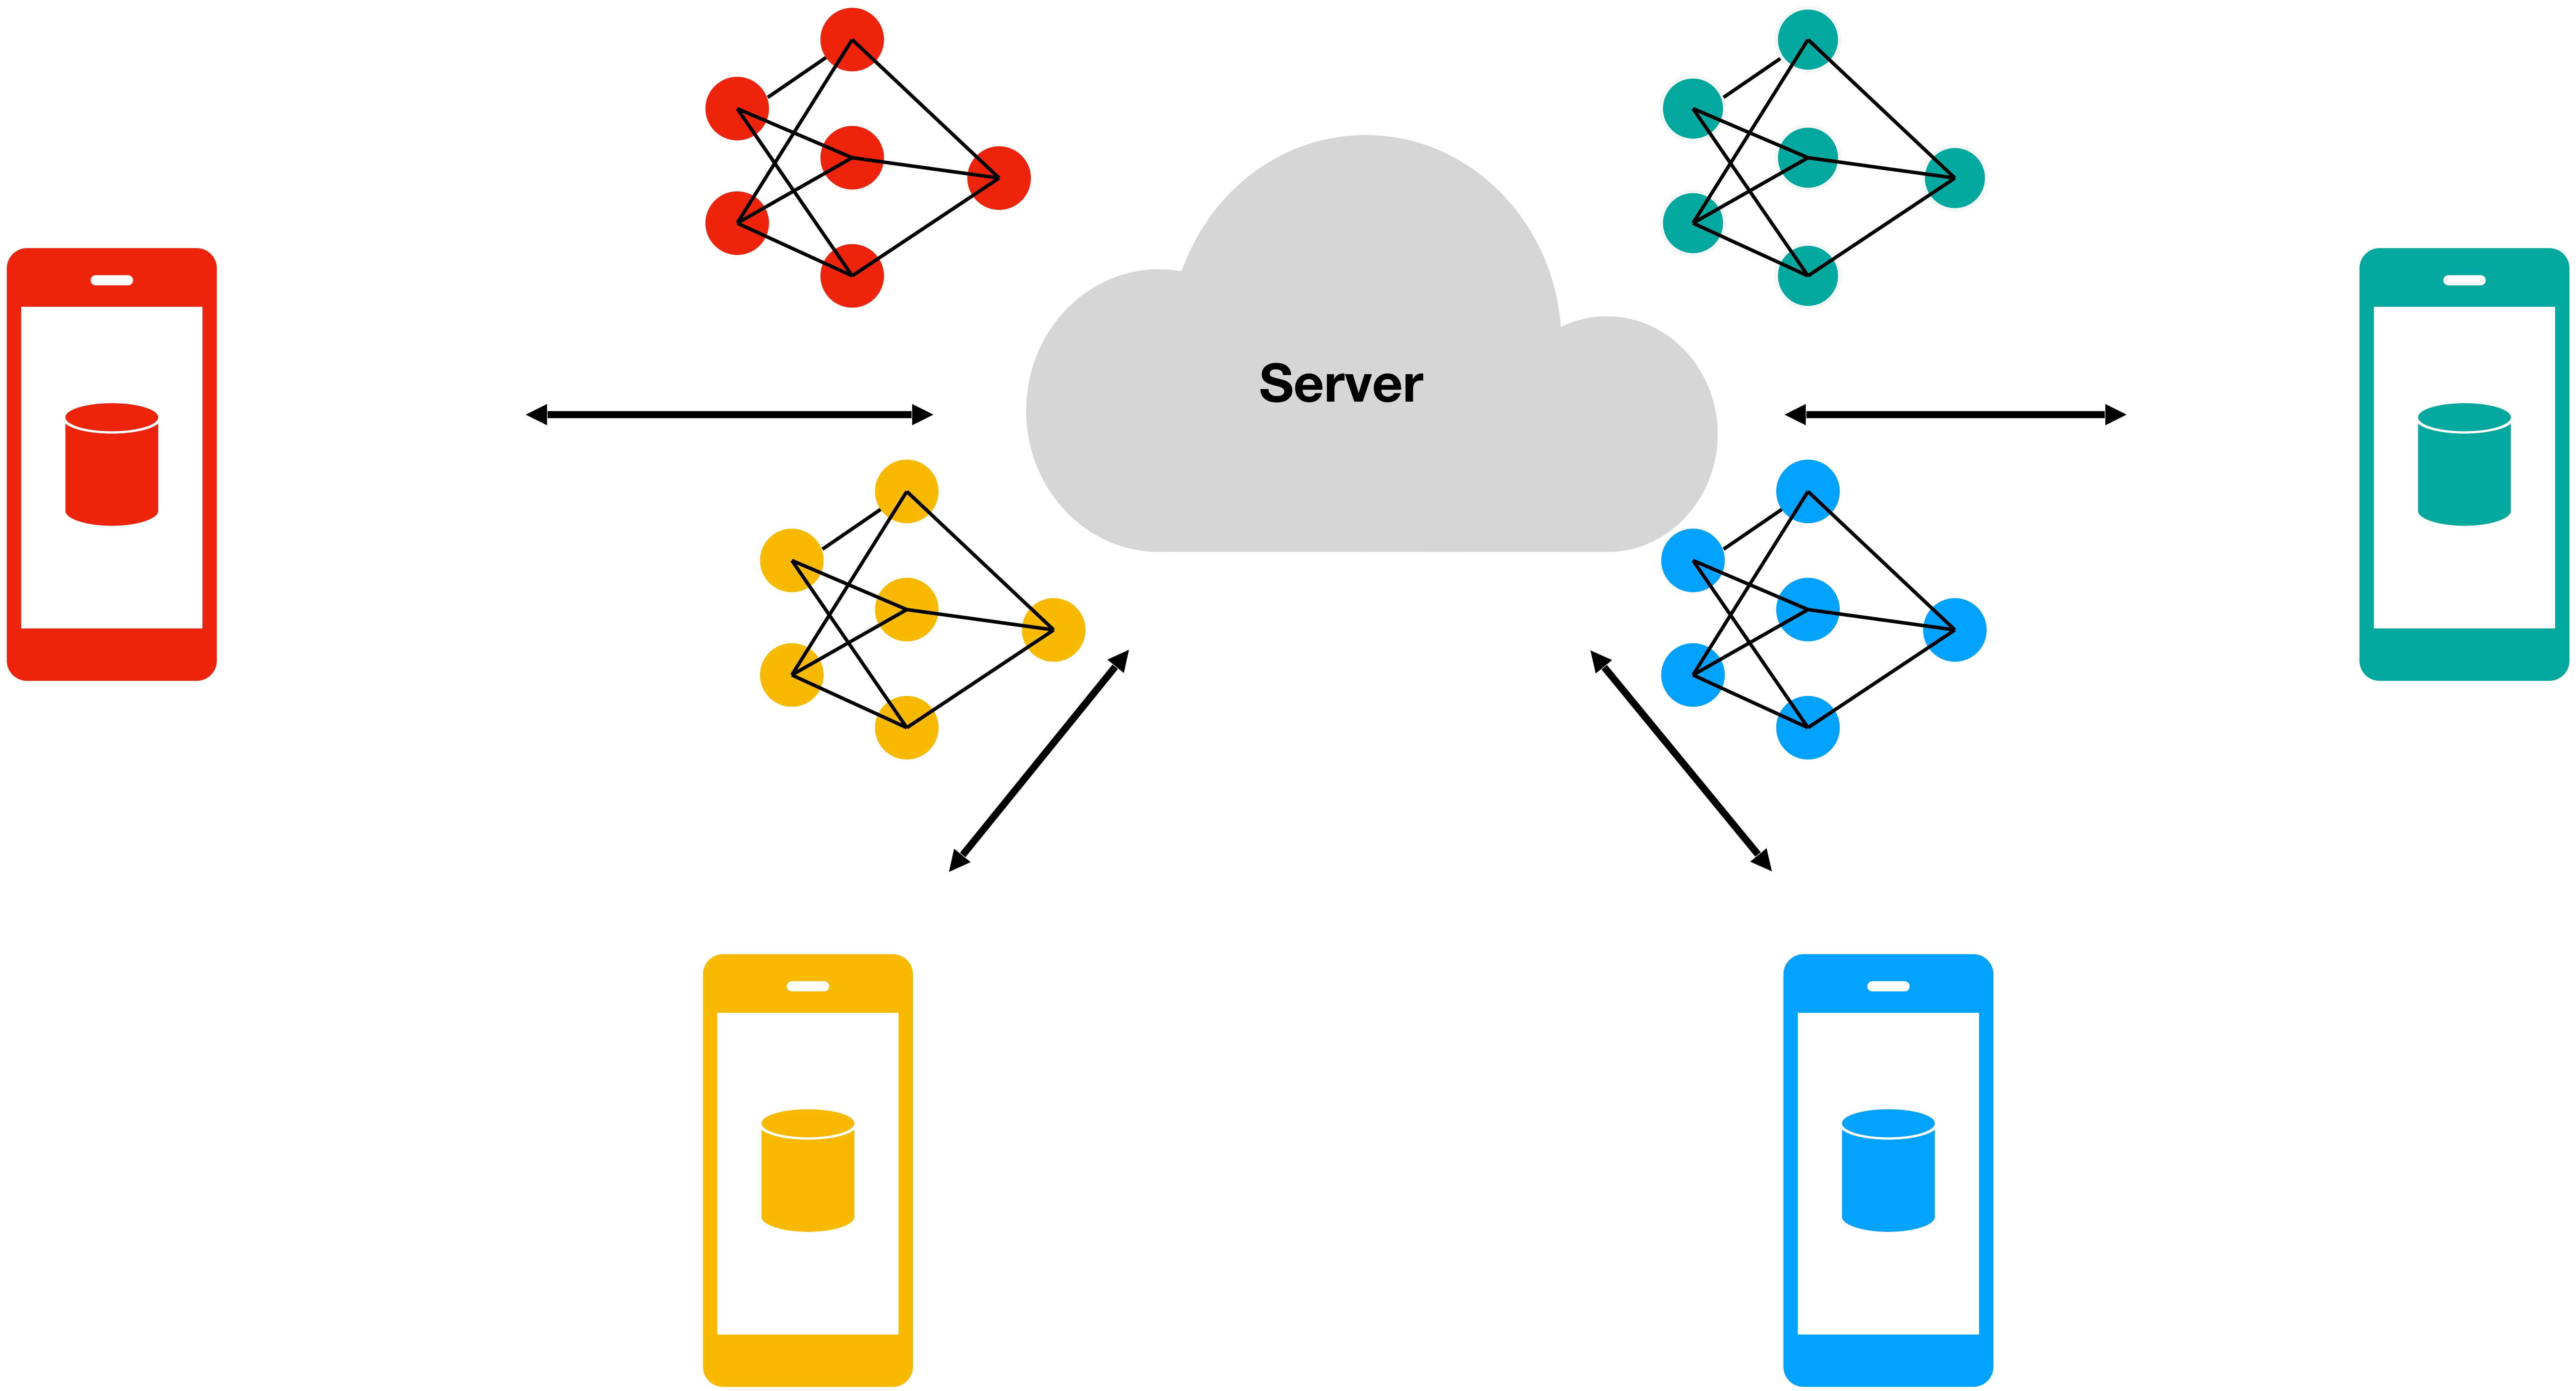

Step 4: Repeat Process with New Global Model

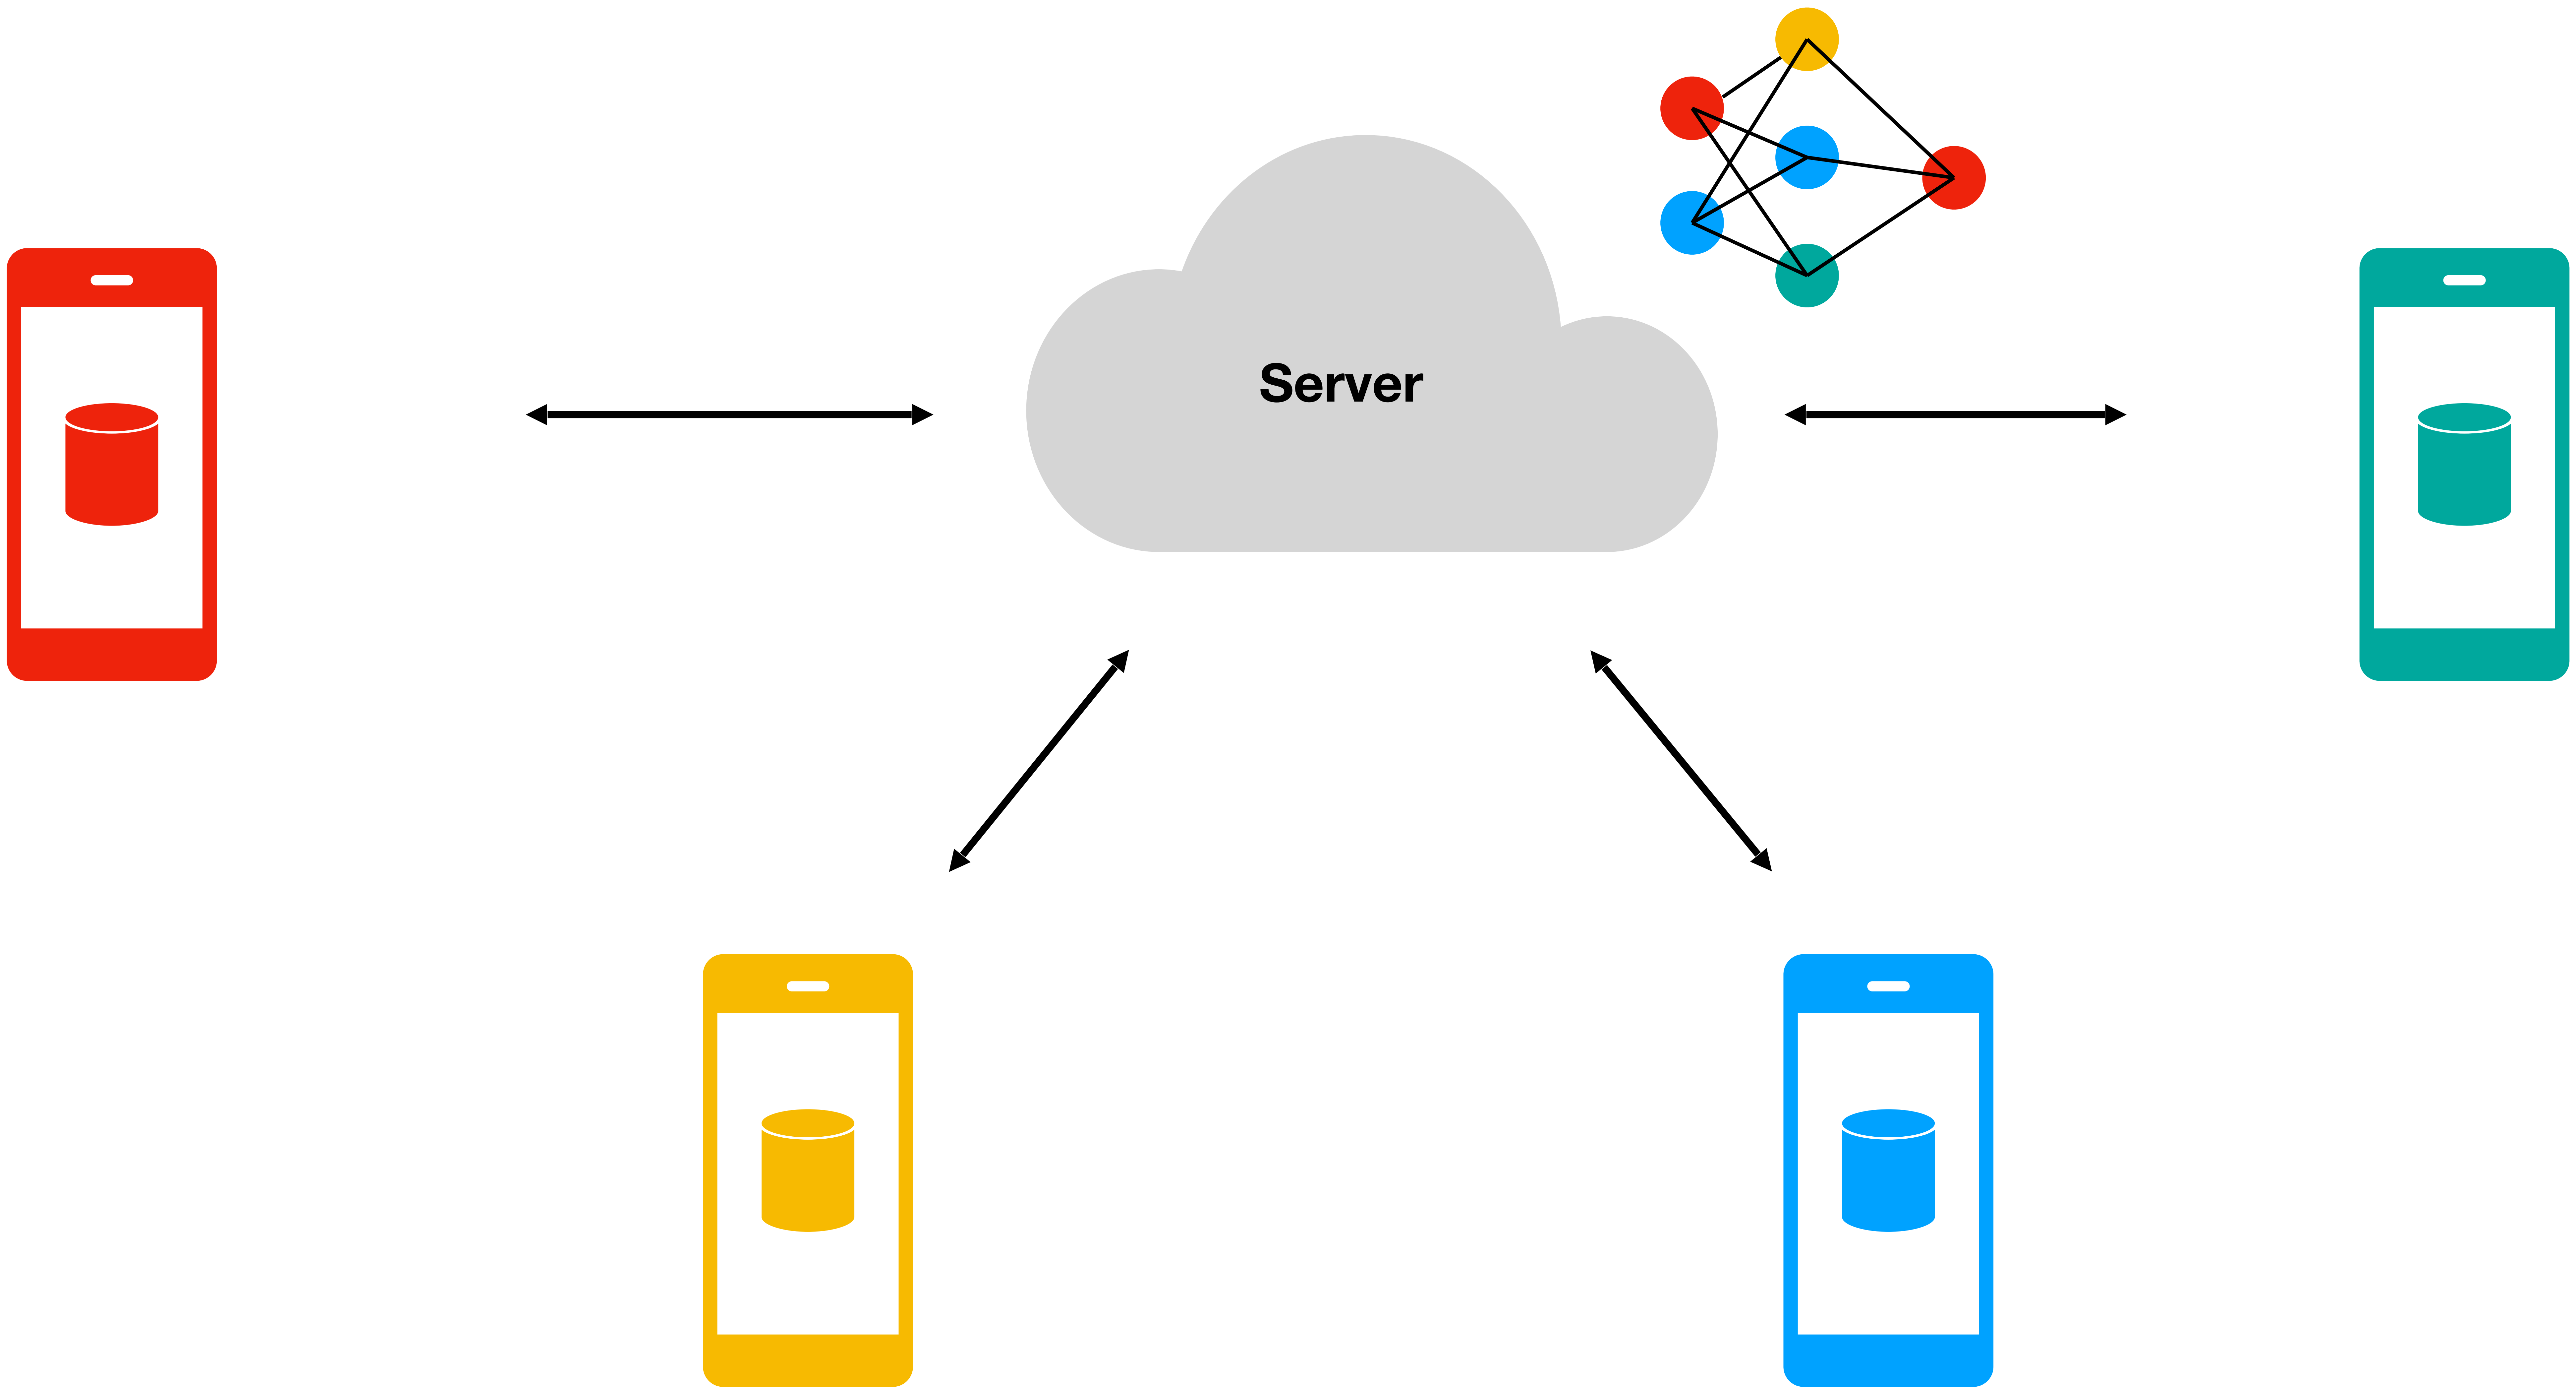

What's the catch?

- Needs **user interaction:** Interaction patterns form labels

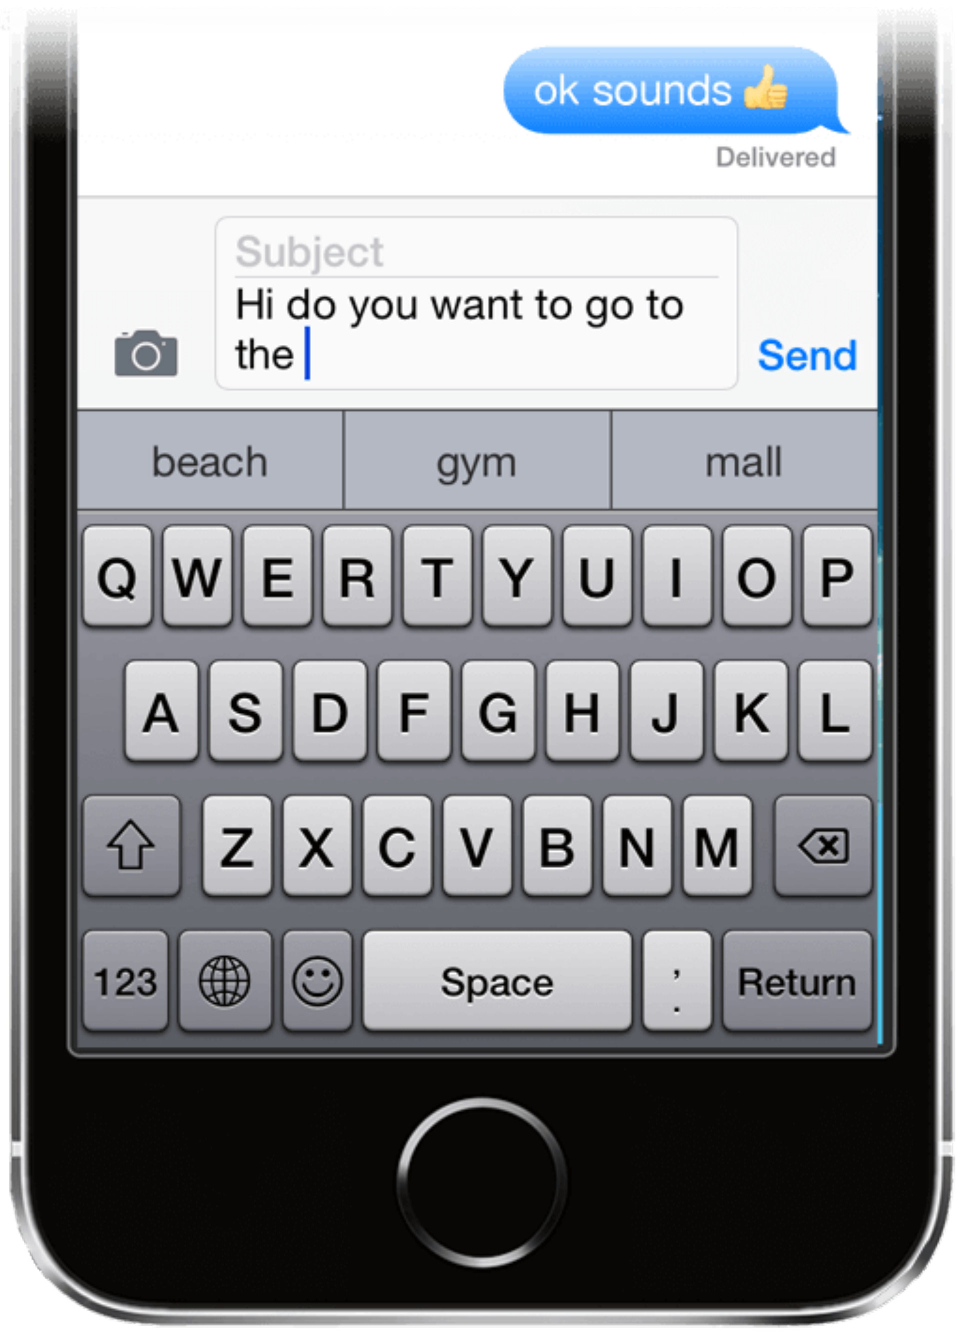

What’s the catch?

- Needs **user interaction:** Interaction patterns form labels
- Most sensors collect data in **non-interactive** manners

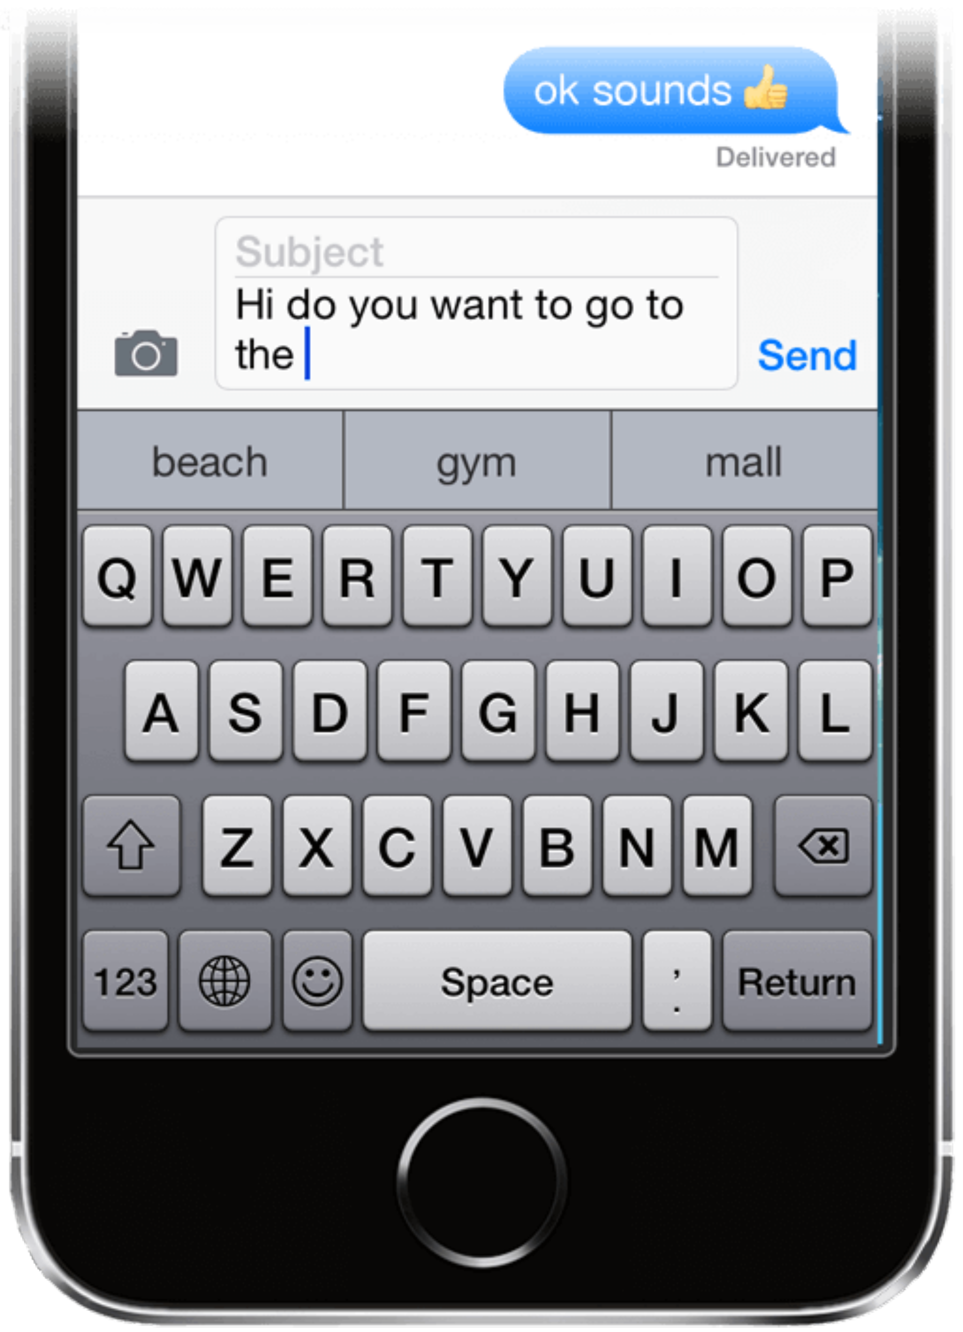

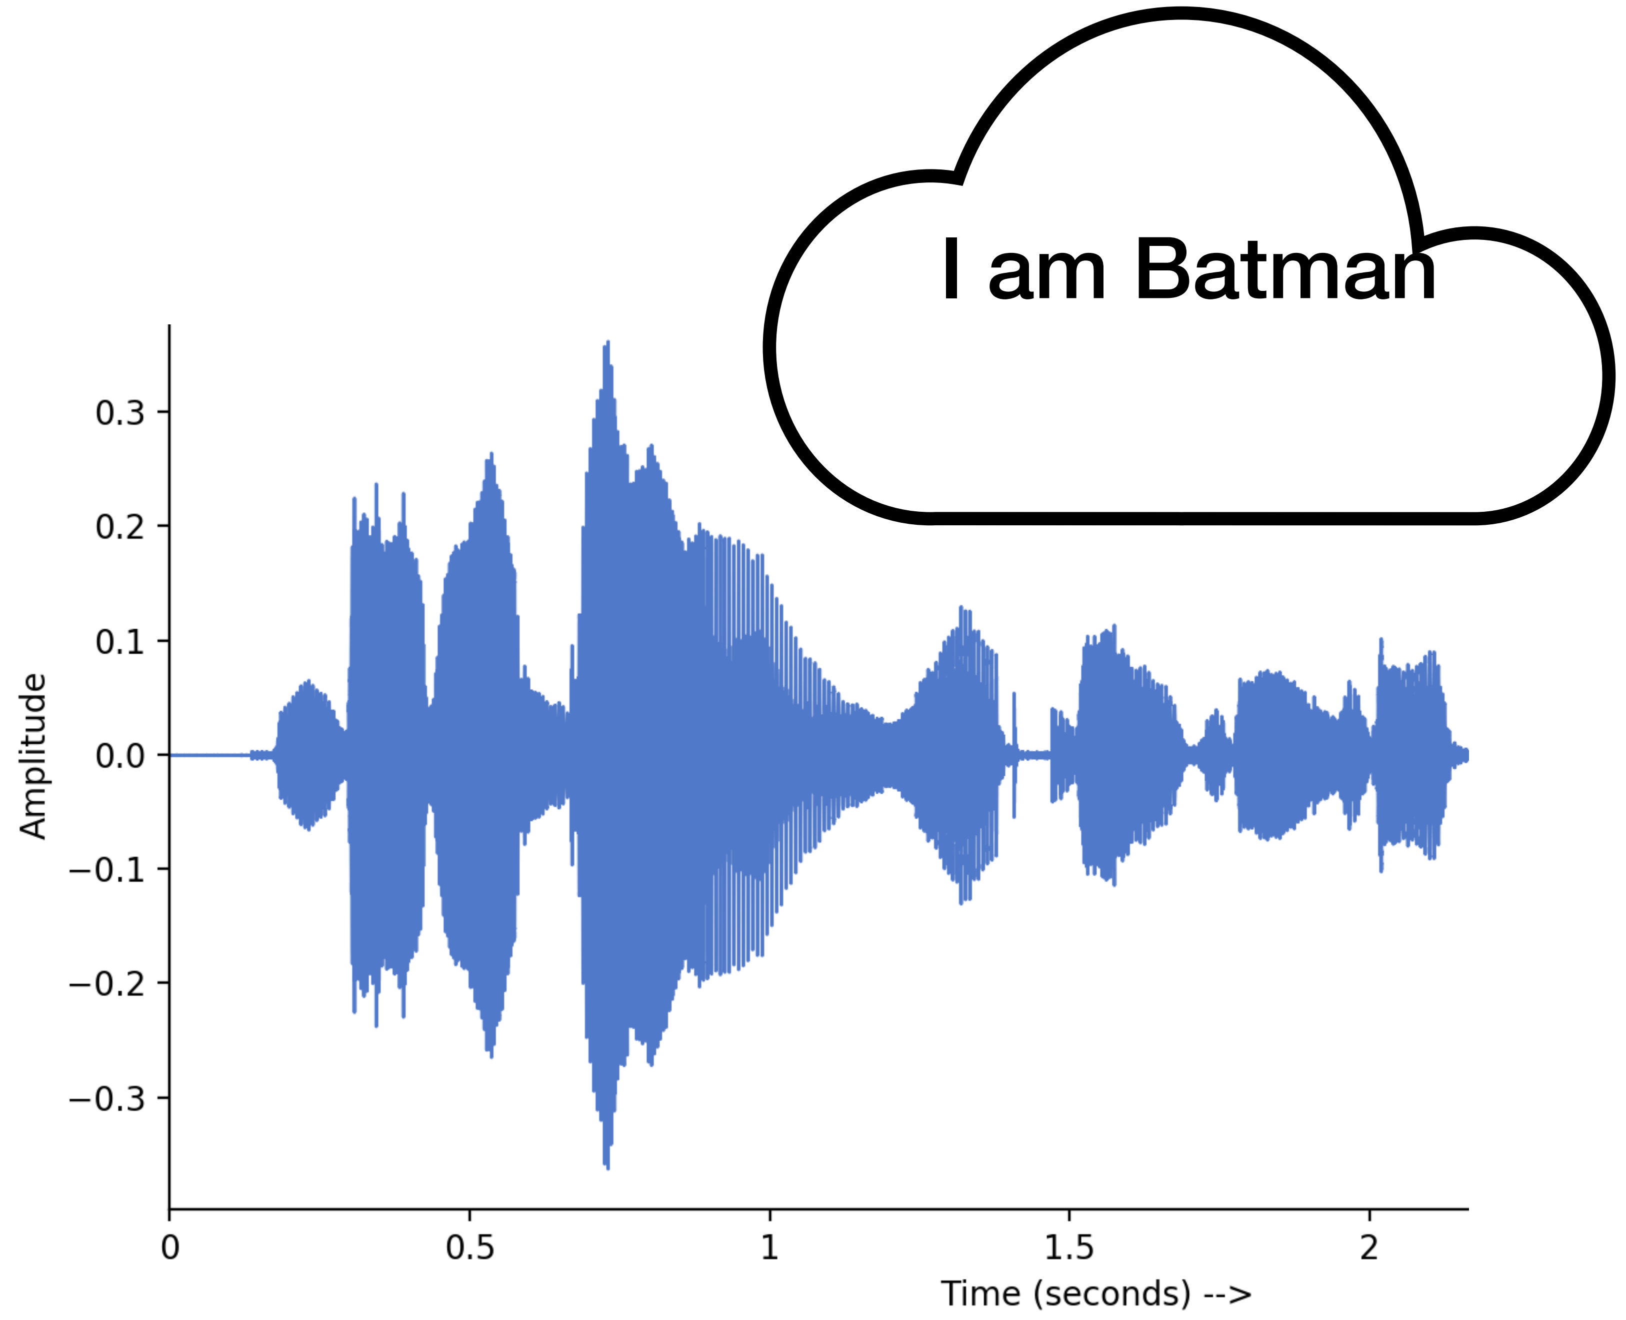

Seems too simple… what's the catch?

- Needs **user interaction:** Interaction patterns form labels
- Most sensors collect data in **non-interactive** manners

Need to go Unsupervised

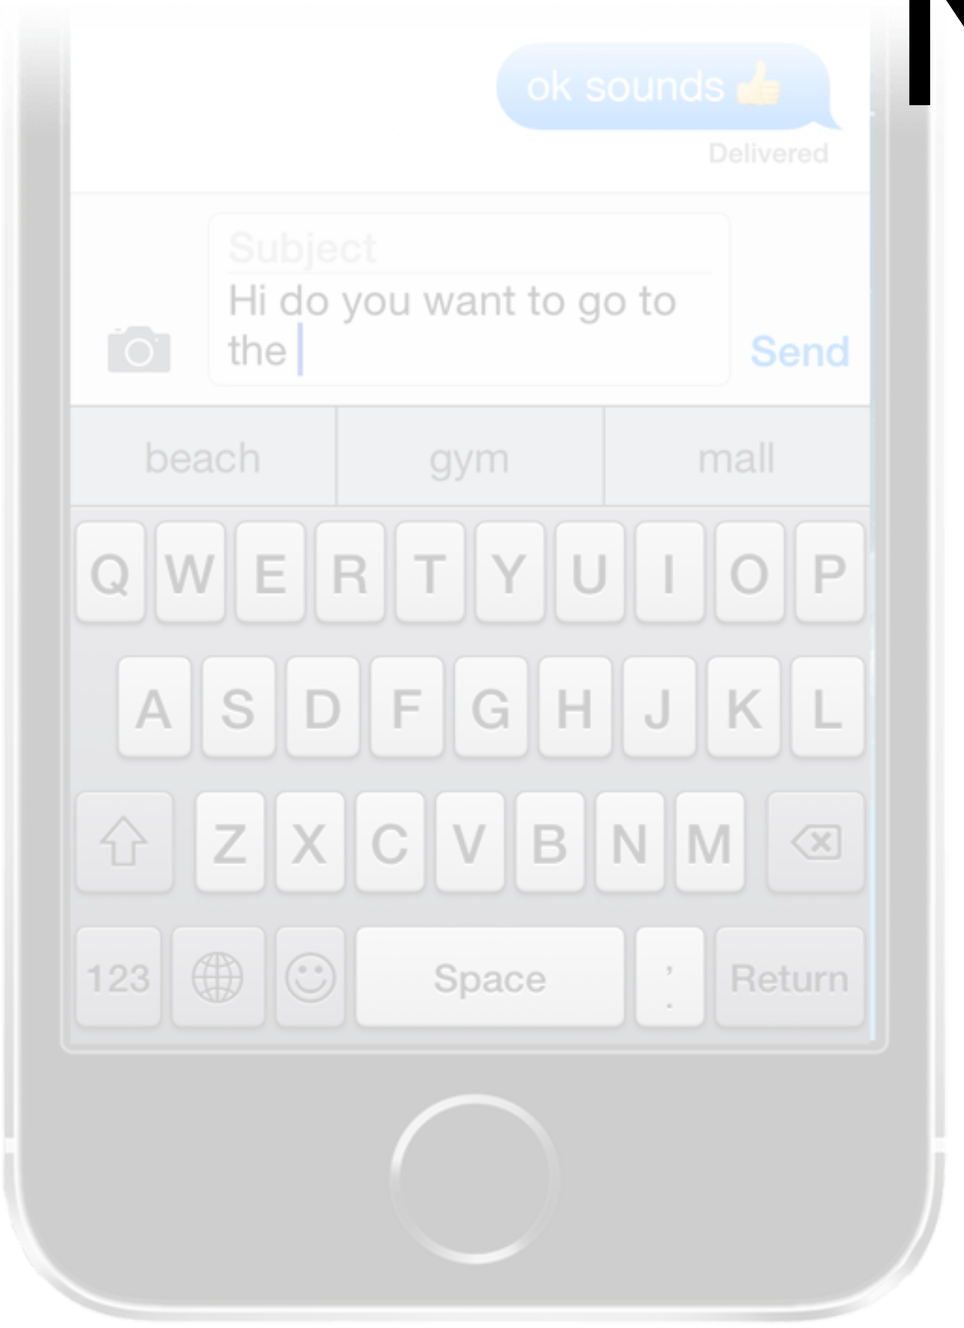

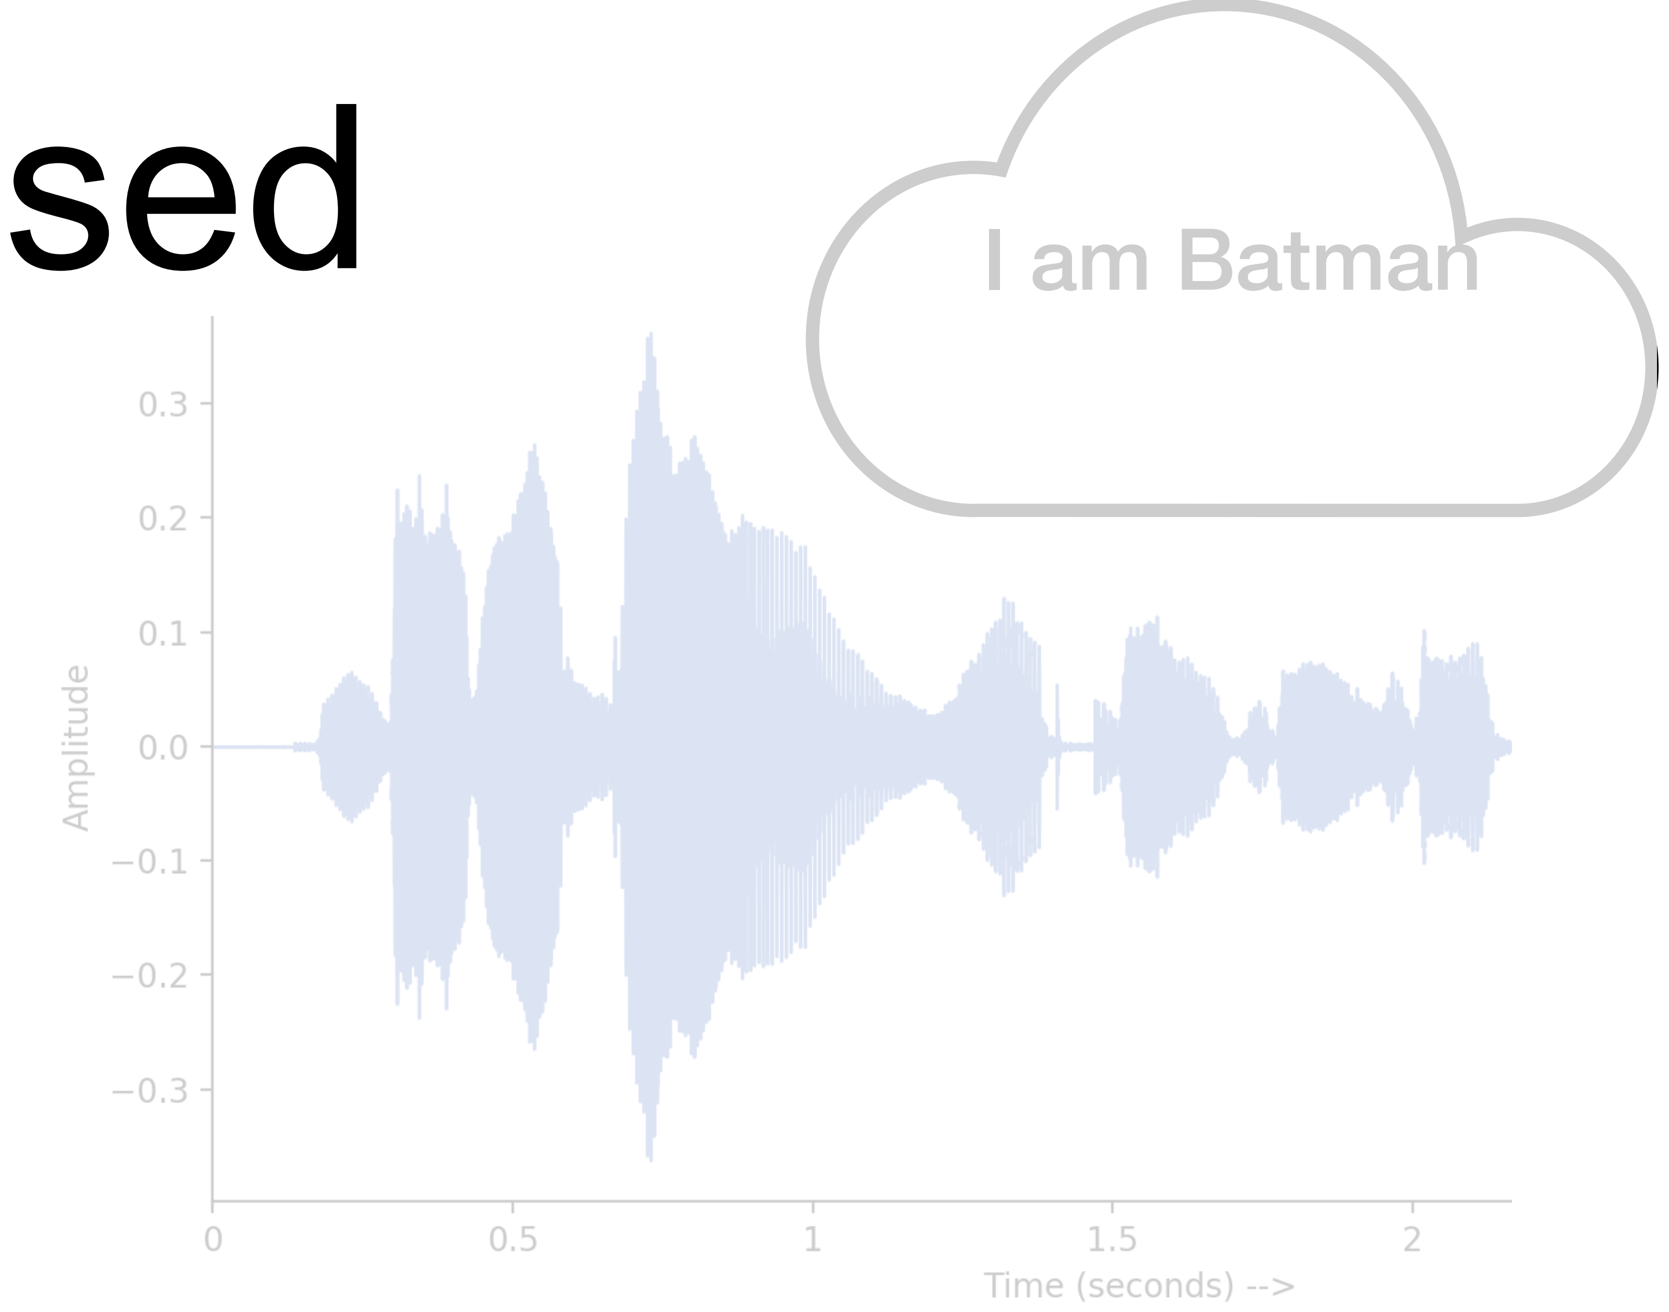

Limitations of Conventional Unsupervised Learning

- Useful unsupervised learning algorithms traditionally require **large amounts of compute and data**

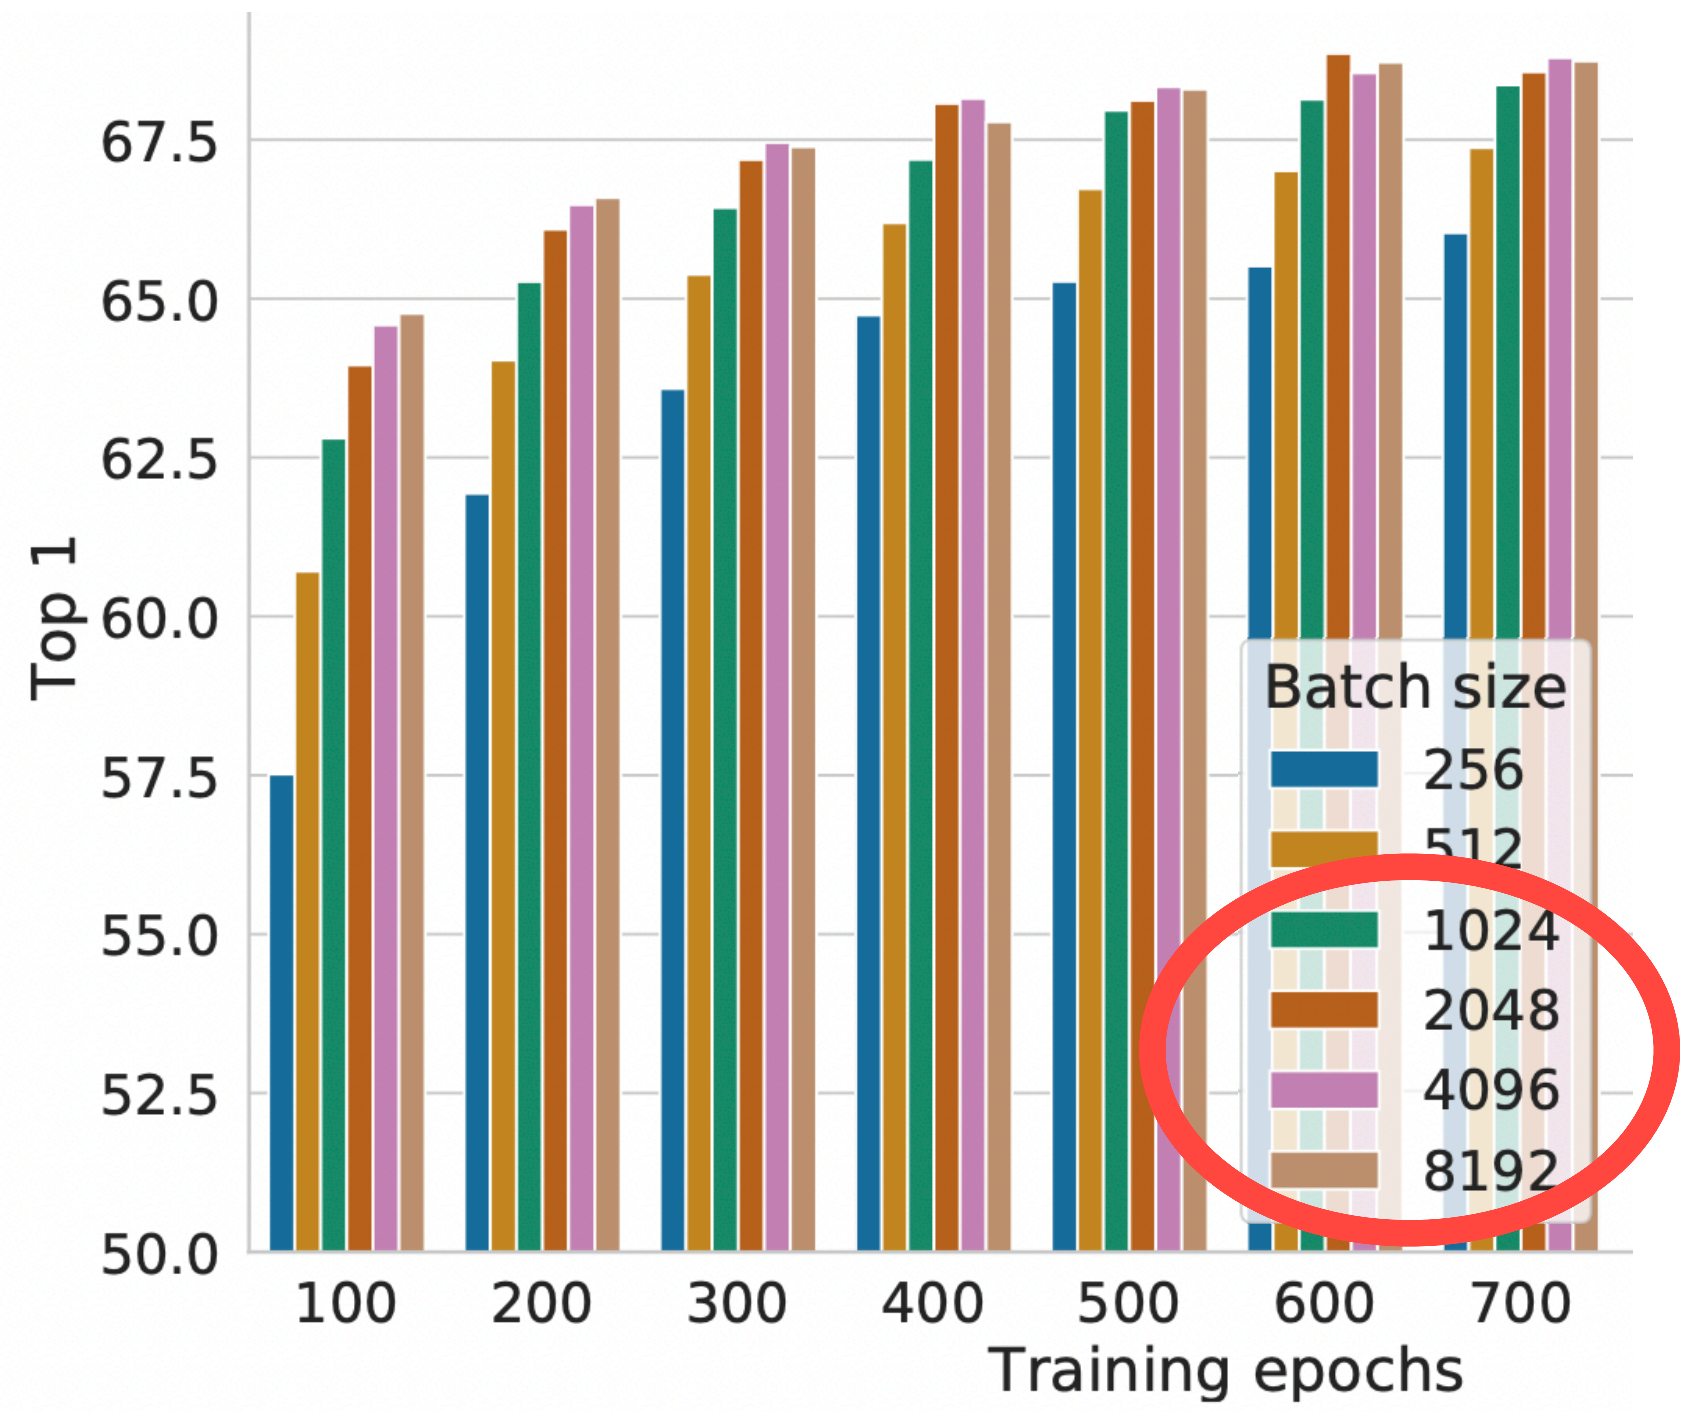

Source: Chen et al., A Simple Framework for Contrastive Learning of Visual Representations. ICML, 2020.

Limitations of Conventional Unsupervised Learning

- Useful unsupervised learning algorithms traditionally require **large amounts of compute and data**

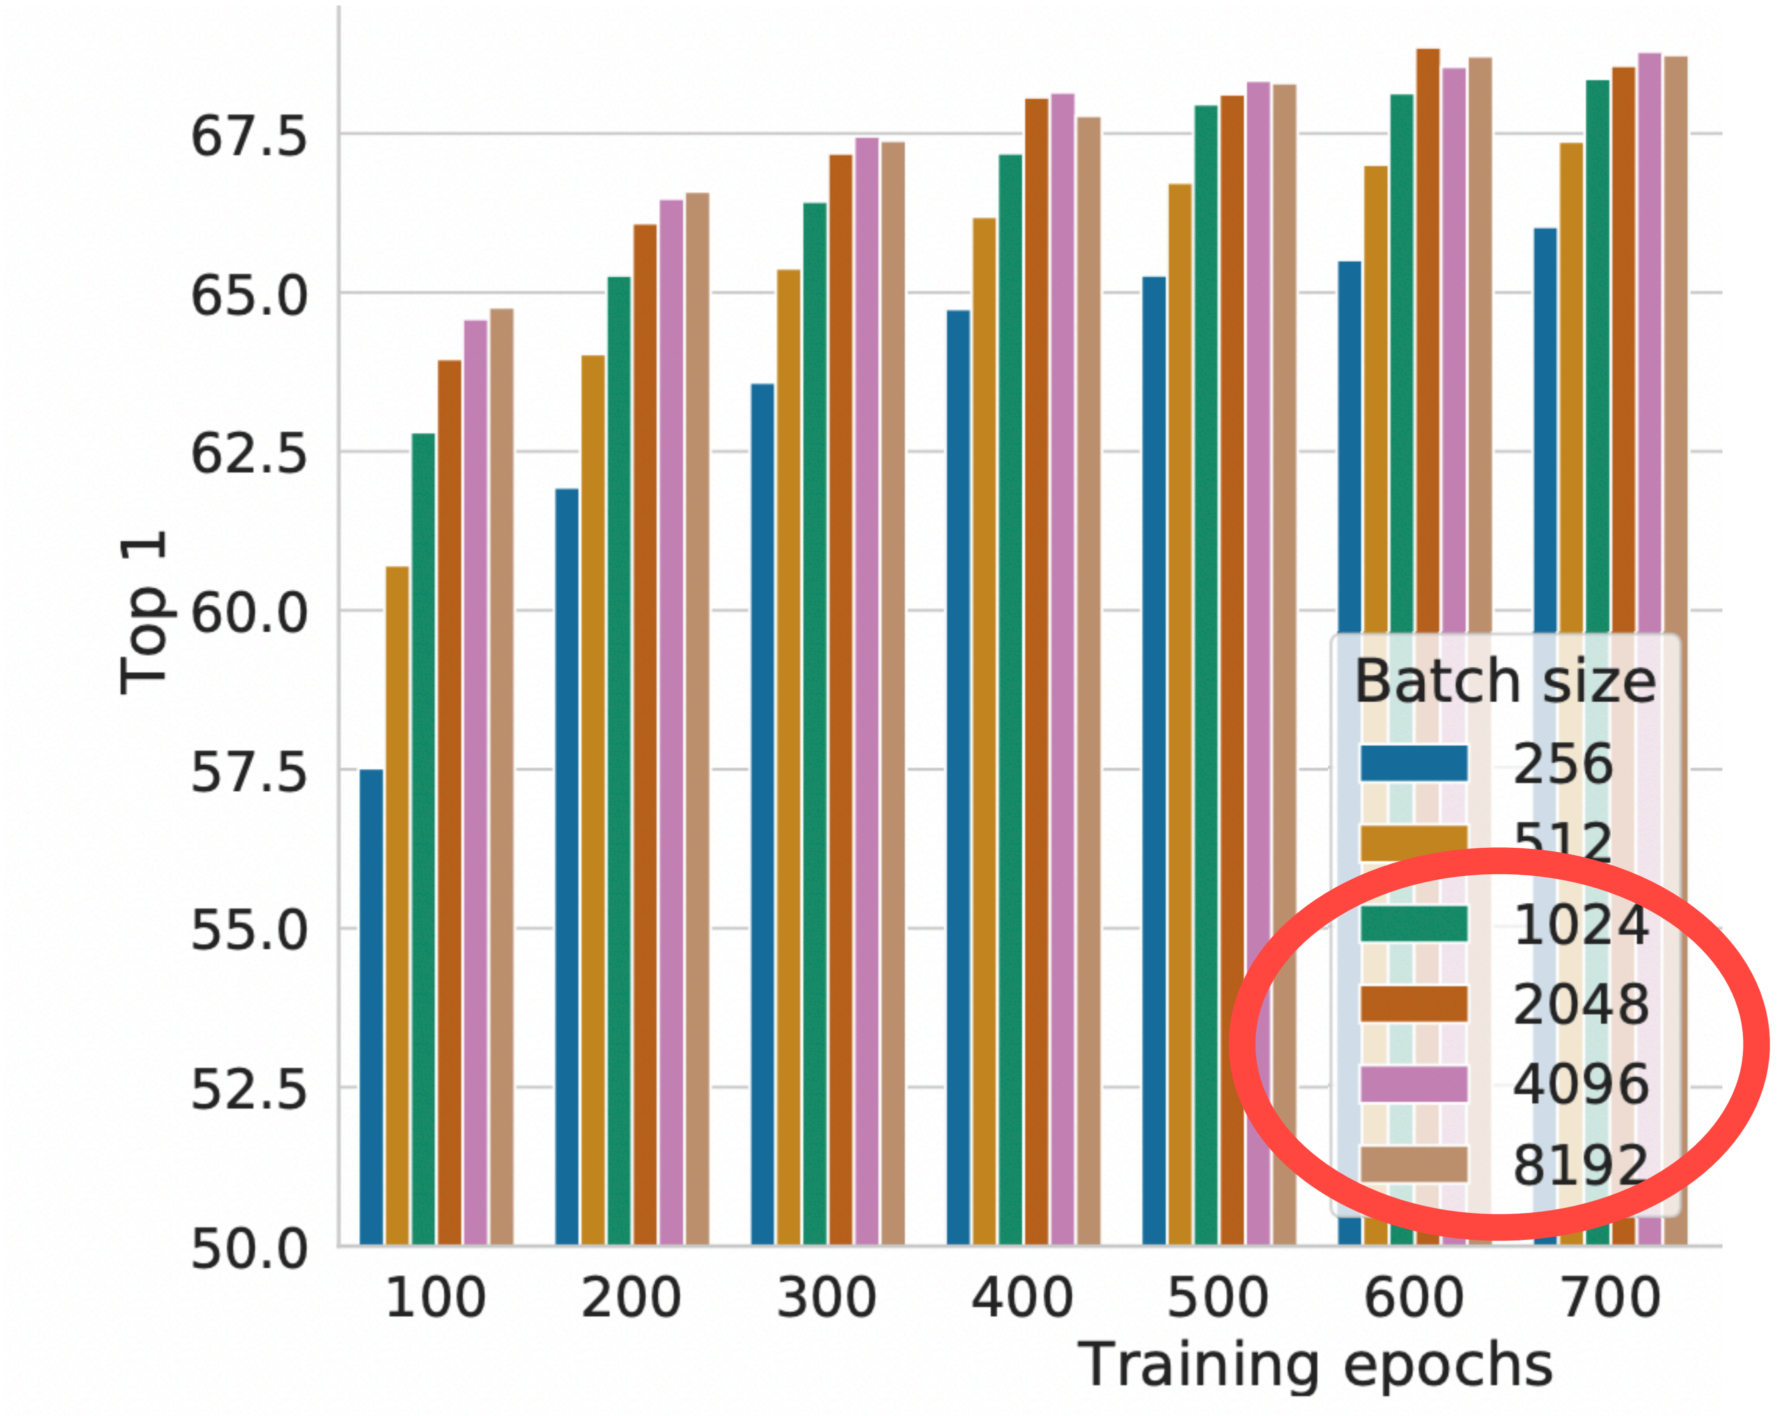

- Edge devices have small memories and few data points

- Minimal compute resources necessitate small batch-sizes

- Federated algorithms need to work with non-IID datasets

Source: Chen et al., A Simple Framework for Contrastive Learning of Visual Representations. ICML, 2020.

Our Proposition:

Orchestra: Unsupervised Federated Learning with Globally Consistent Clustering

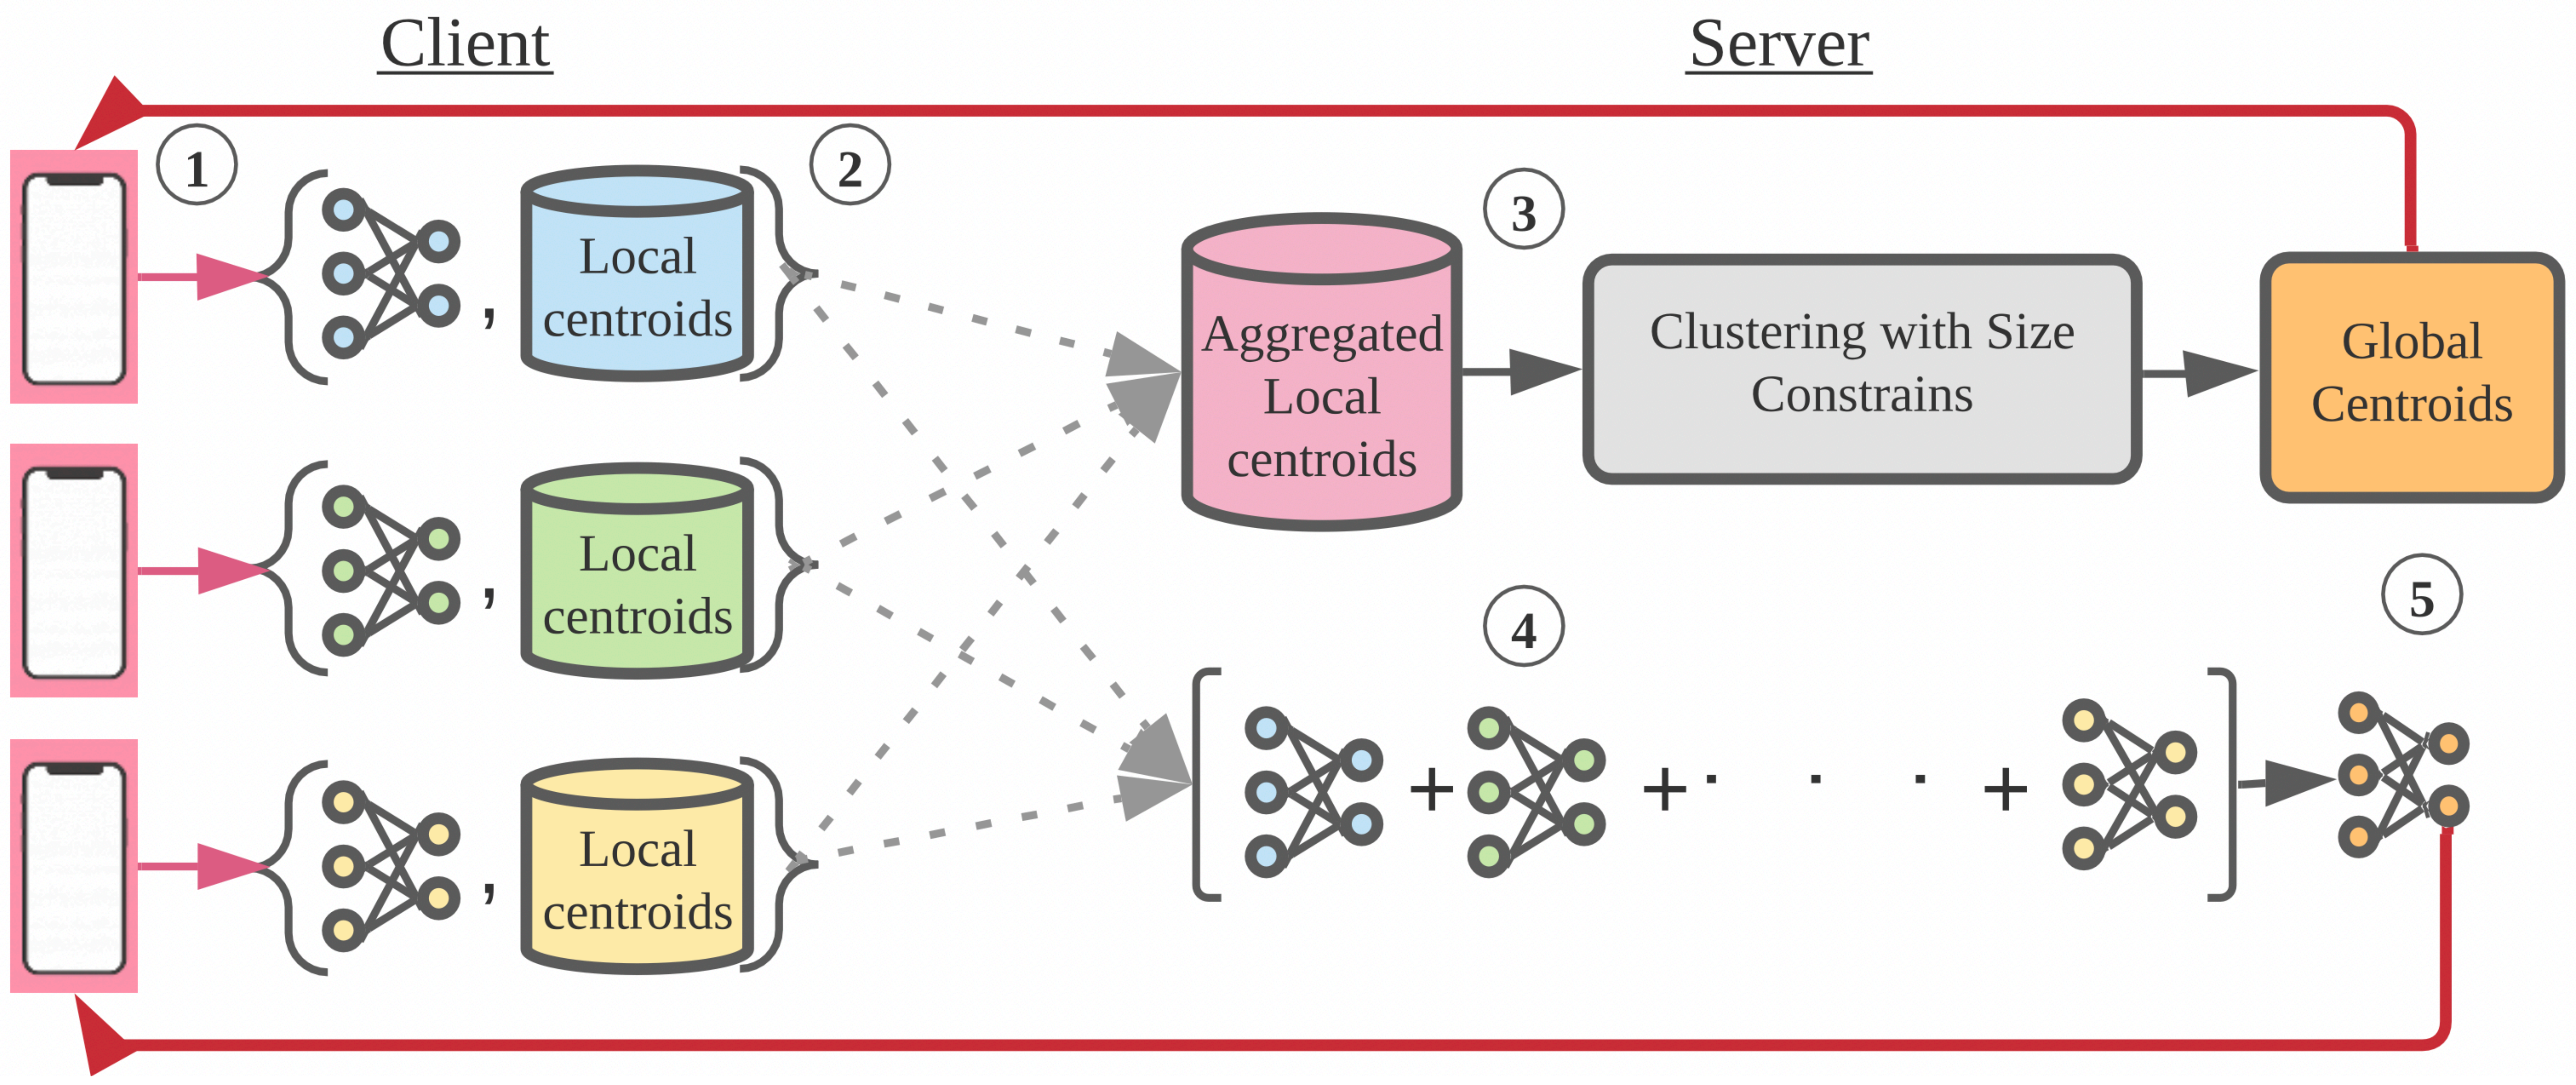

Step 1: Compute data representations

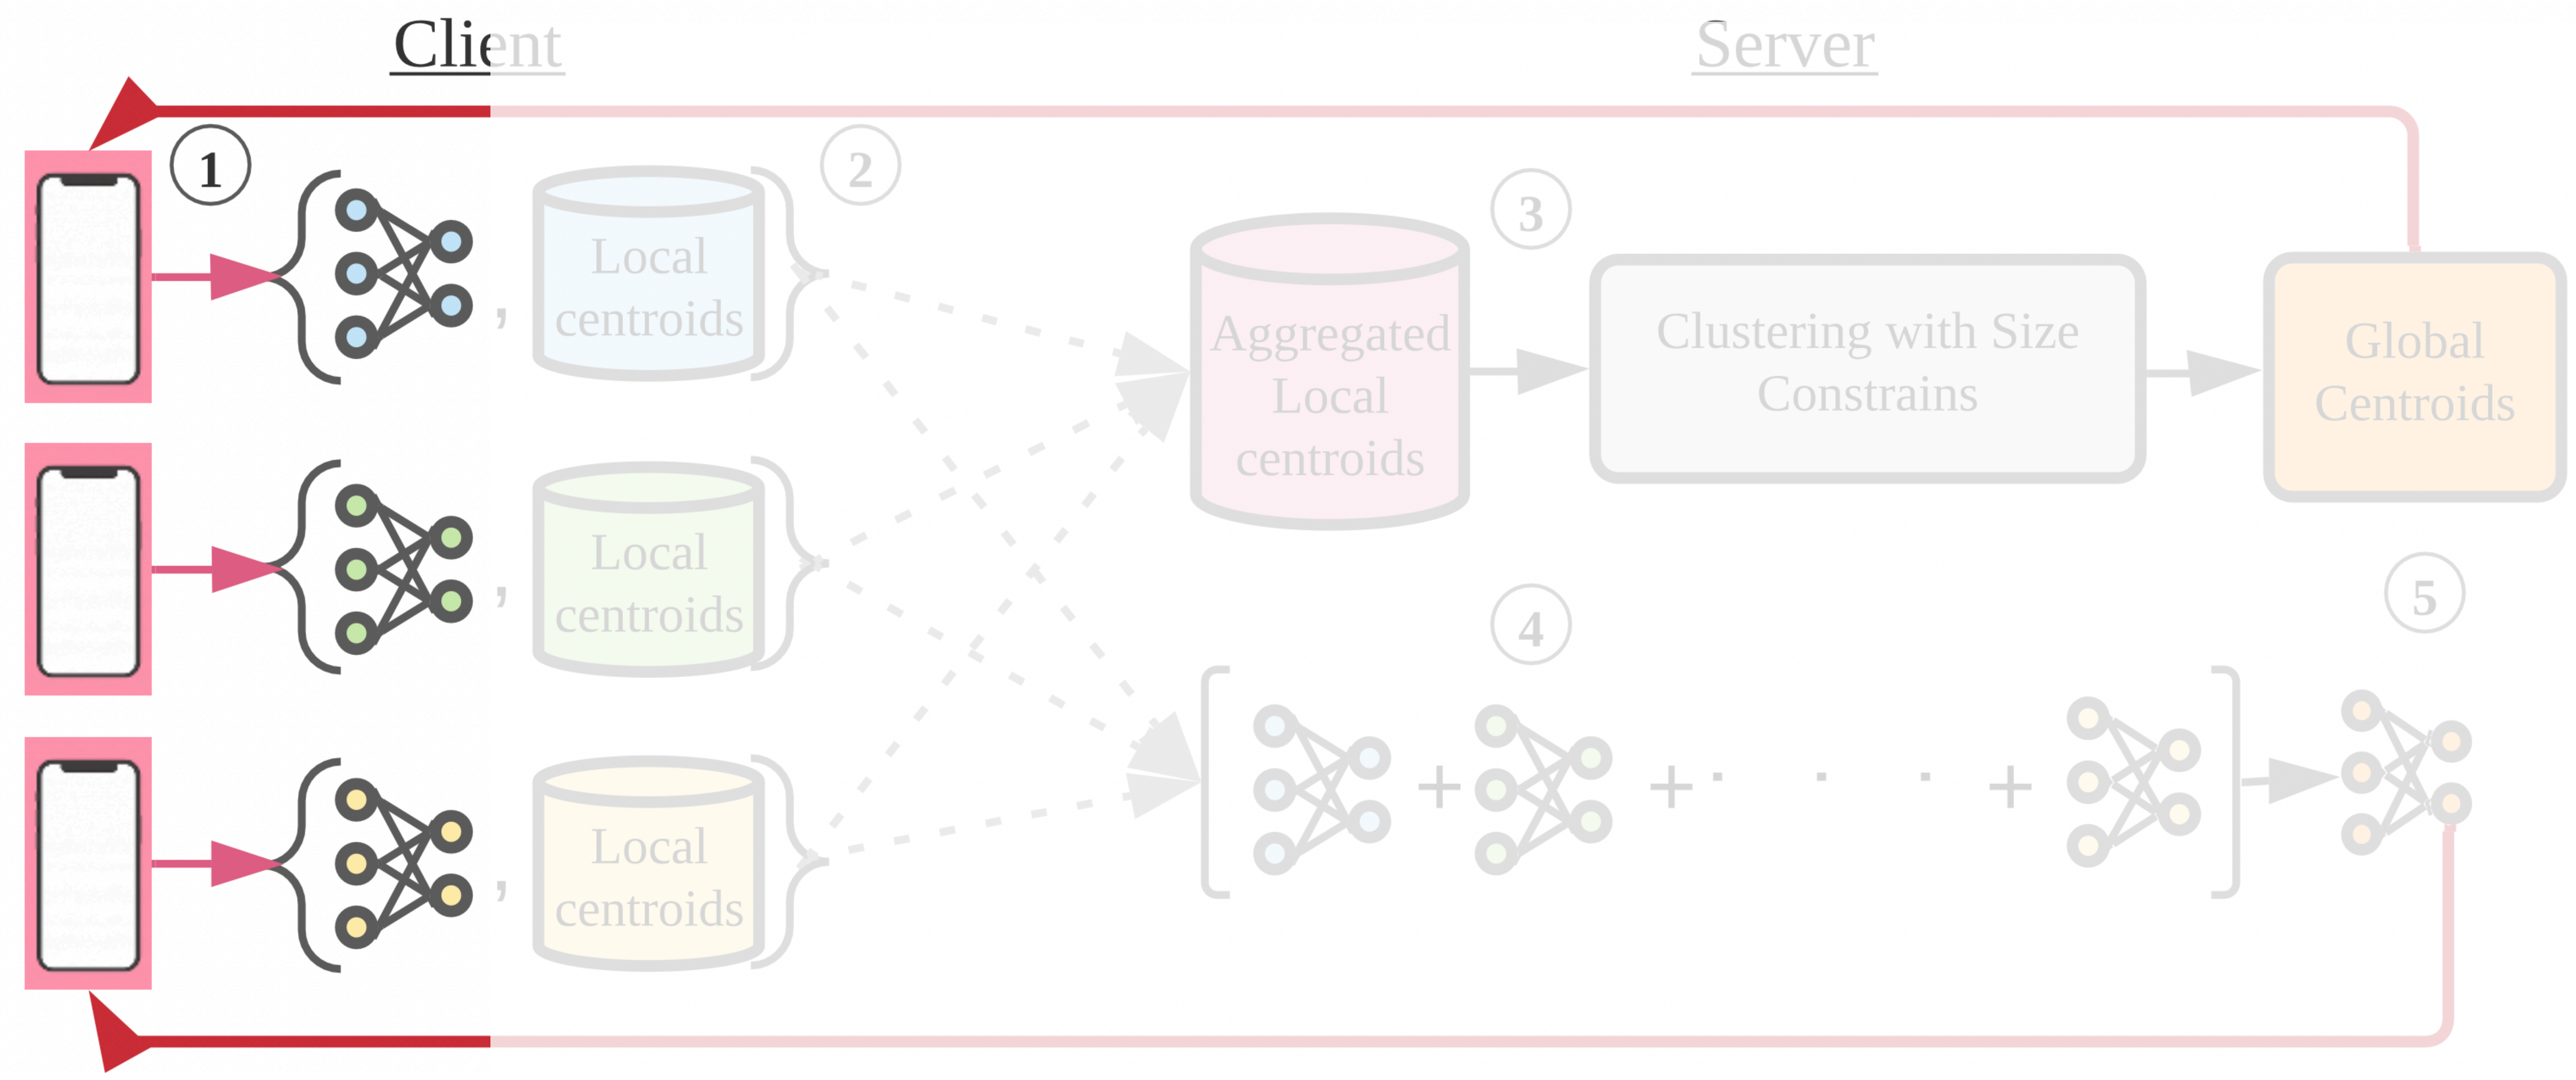

Step 2: Compute local centroids, share with server

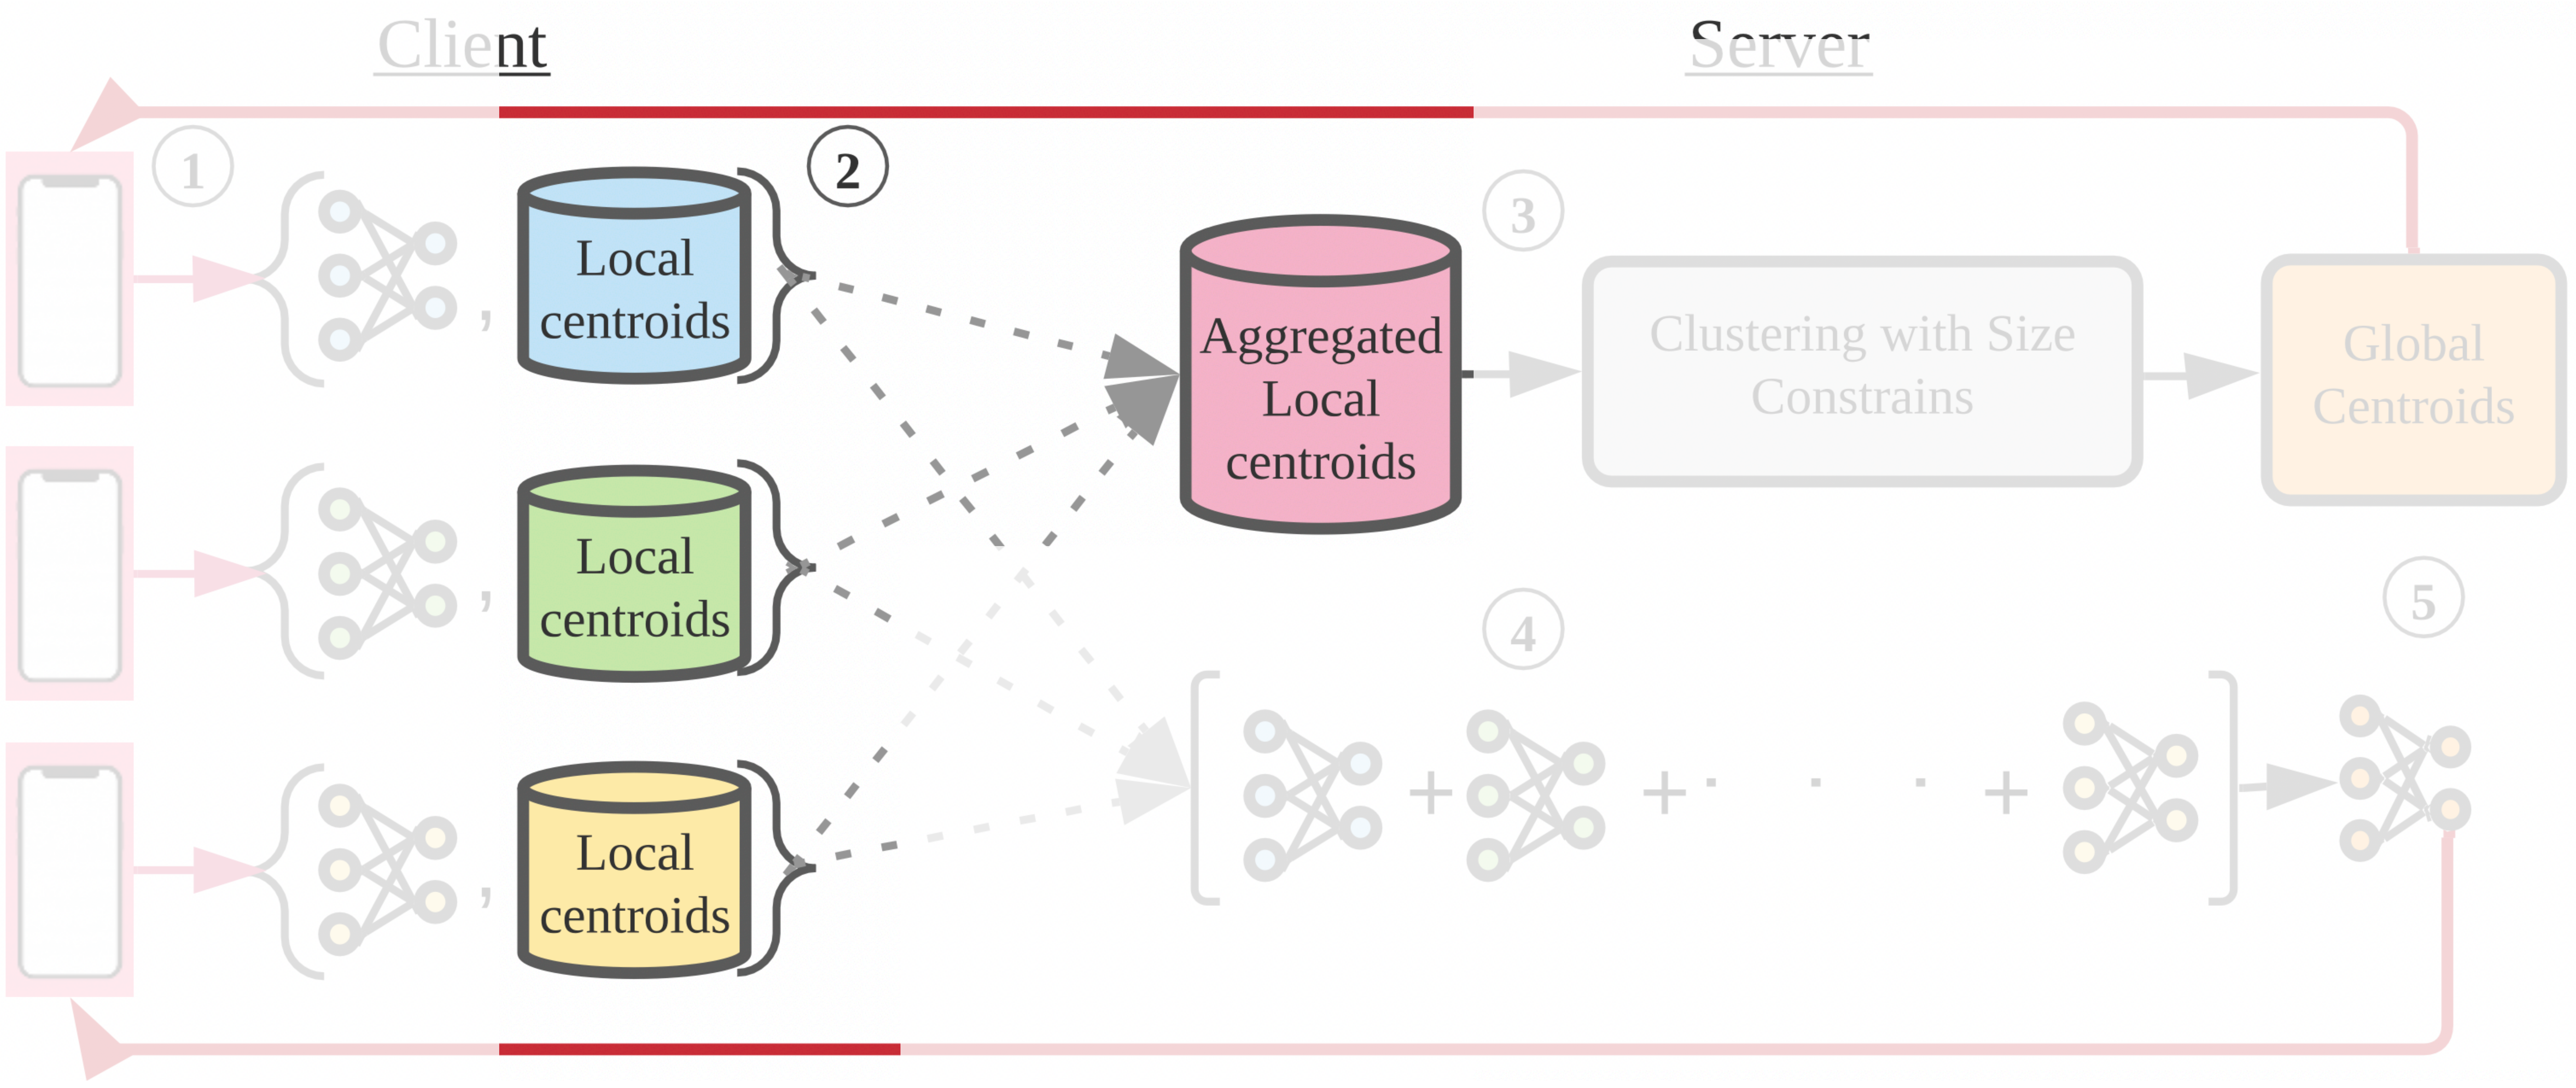

Step 3: Compute “global” centroids from “local” ones

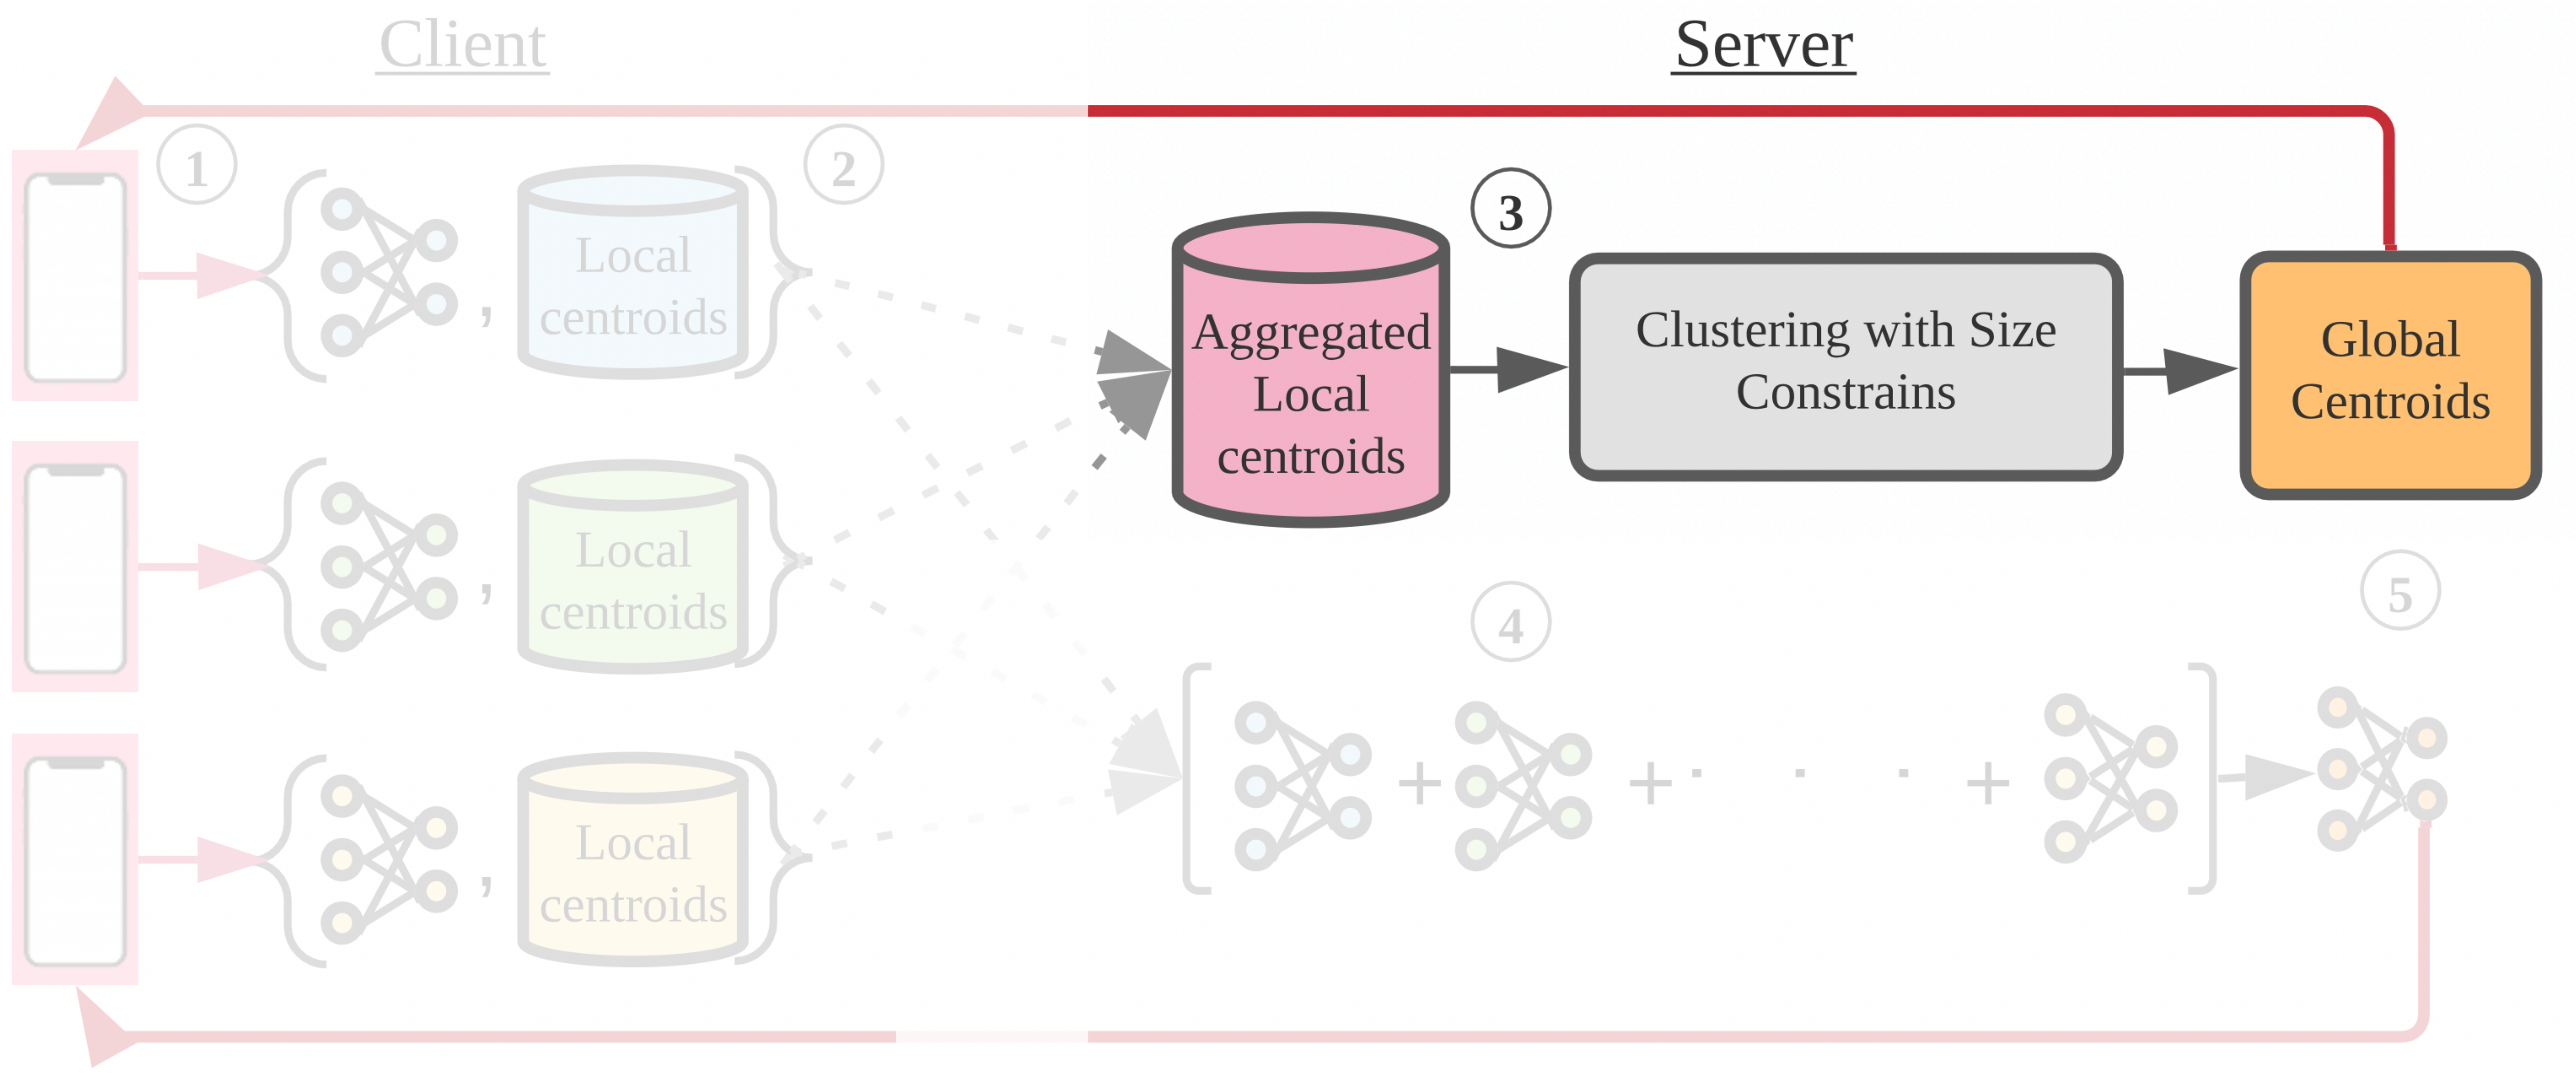

Step 4: Aggregate local models

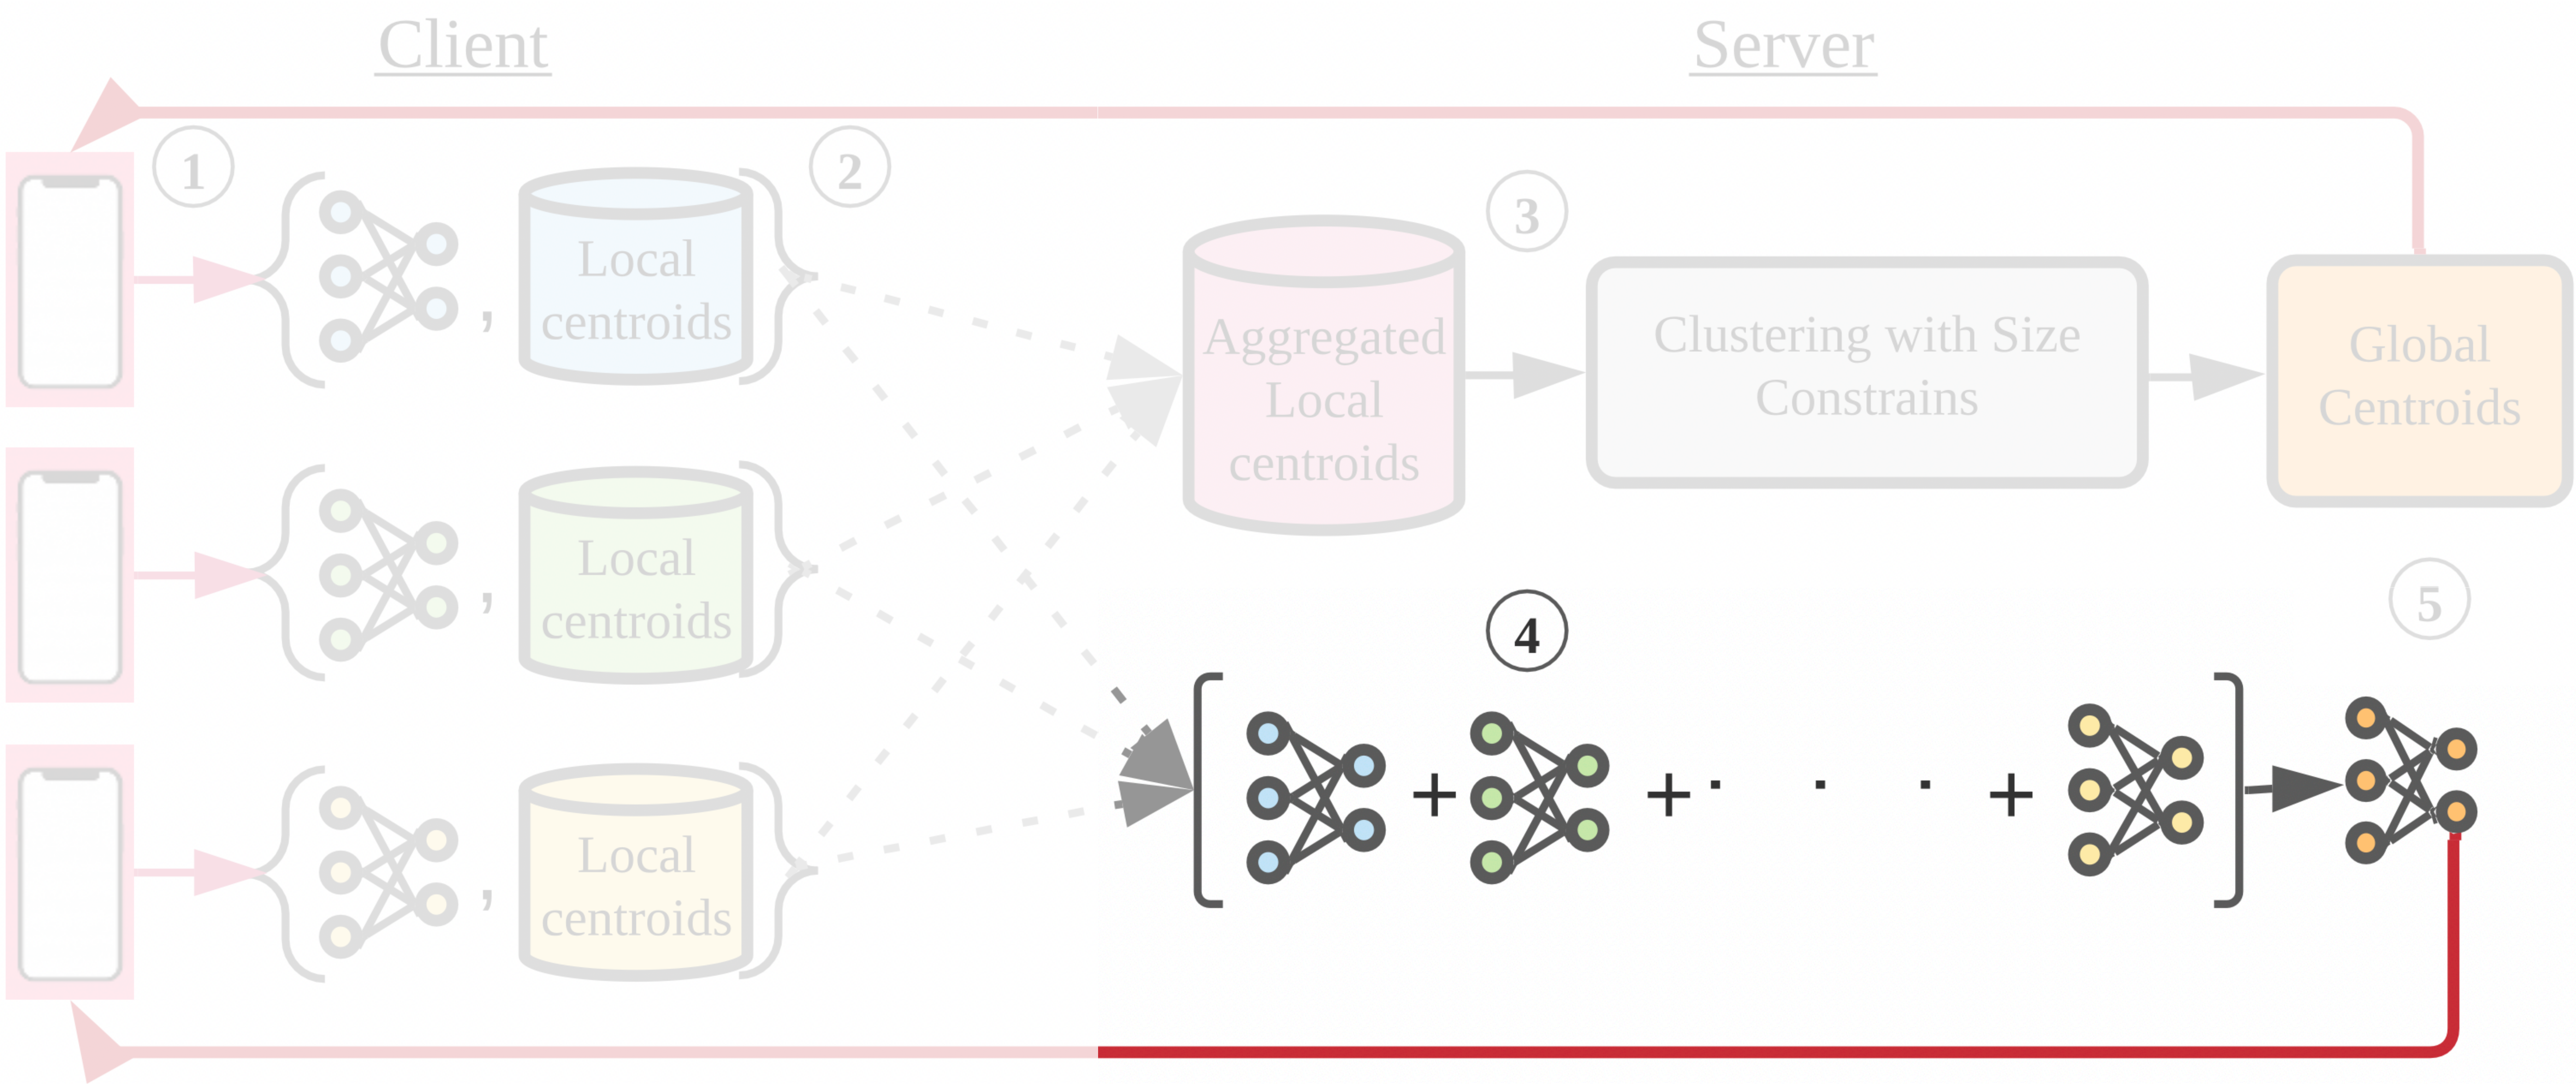

Step 5: Share global centroids and model with clients

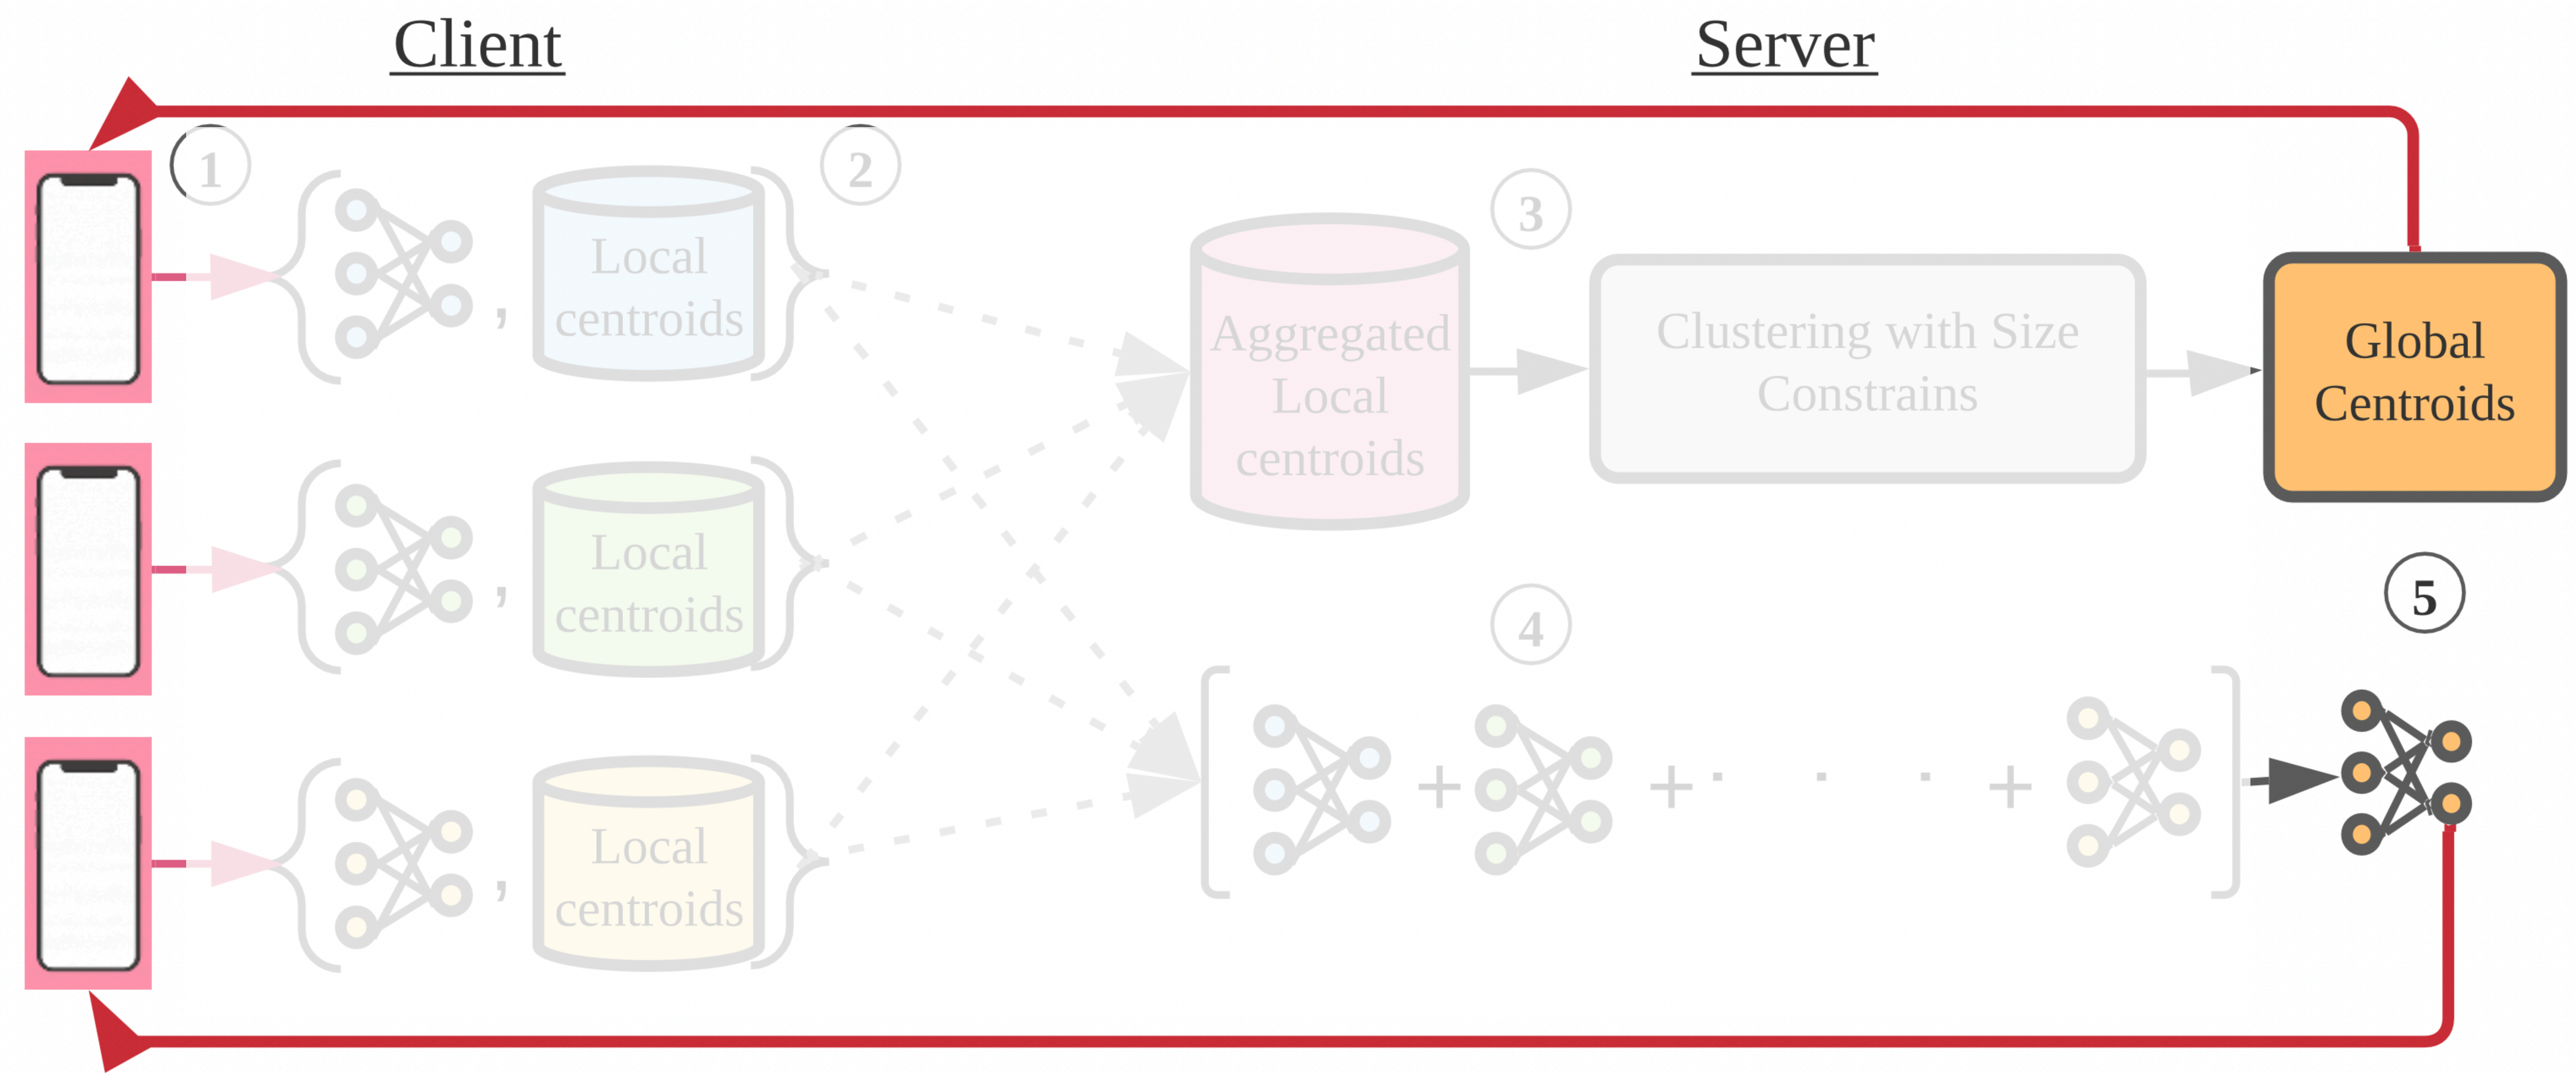

Step 6: Local training: Cluster Prediction Task

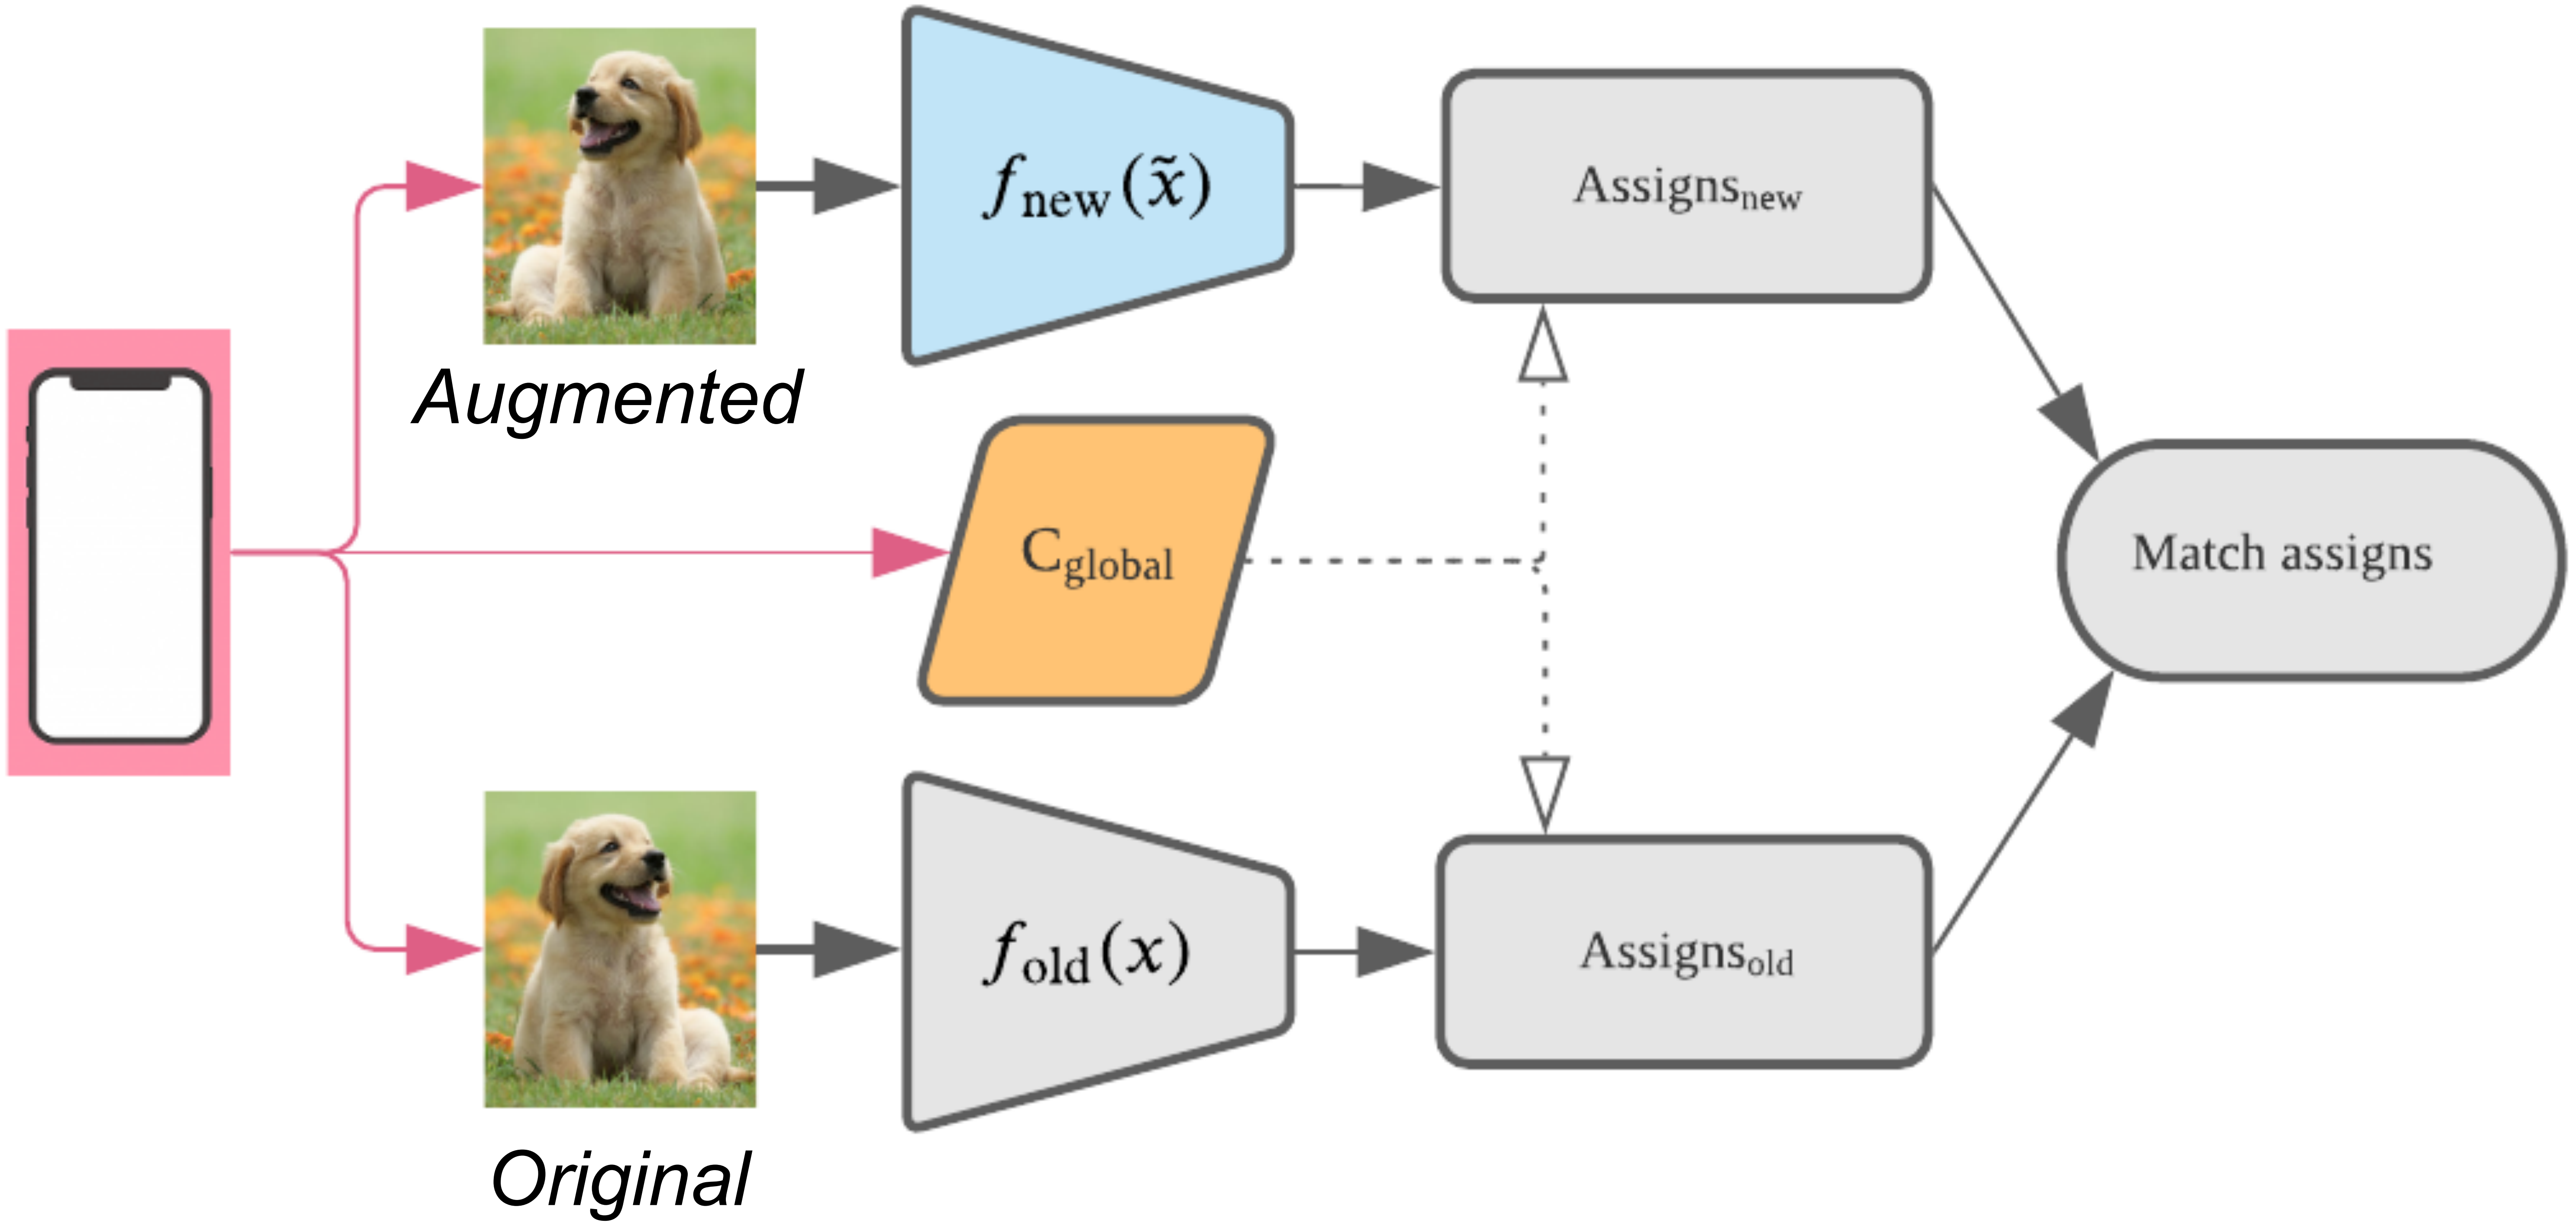

The Concert: Experimental Evaluation

- Datasets: CIFAR-10 (left) and CIFAR-100 (right)
- General Setting: Cross-device, 100 clients, 10 local epochs, 16 batch-size
- Baselines: Federated versions of SOTA unsupervised learning algorithms
- Implementation[1]: Flower (Beutel et al., 2020)

Code: https://github.com/akhilmathurs/orchestra

Evaluation 1: Sensitivity to Heterogeneity

- Better absolute accuracy
- Thrives under heterogeneity
- Robust under extreme settings too

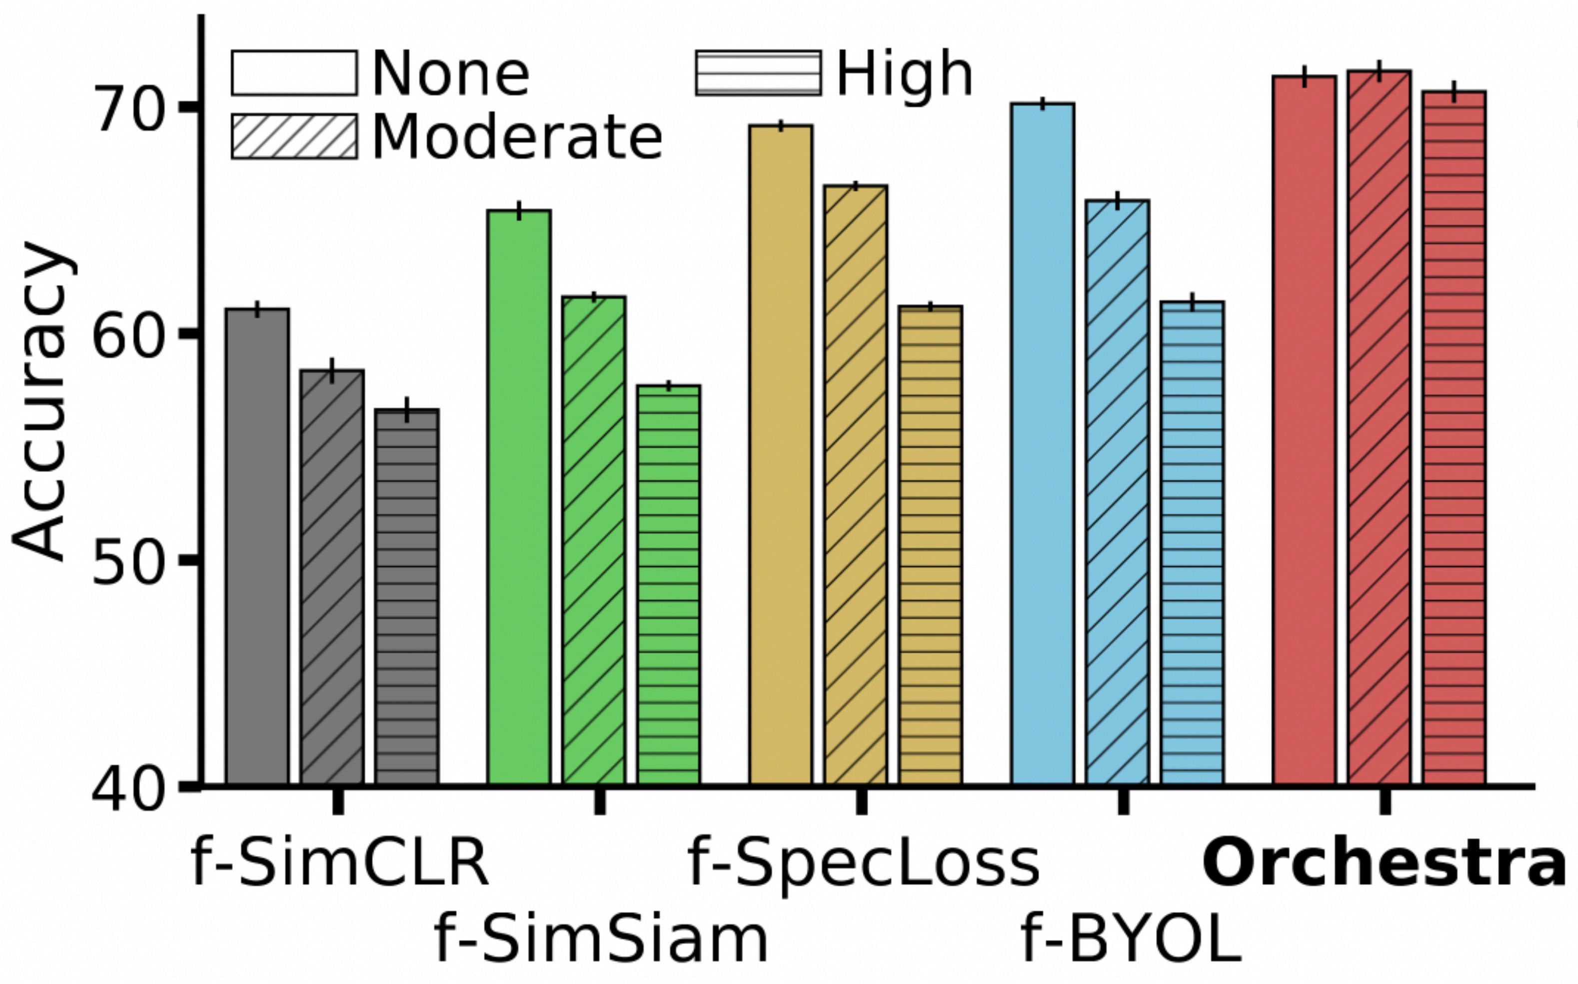

Left: CIFAR-10

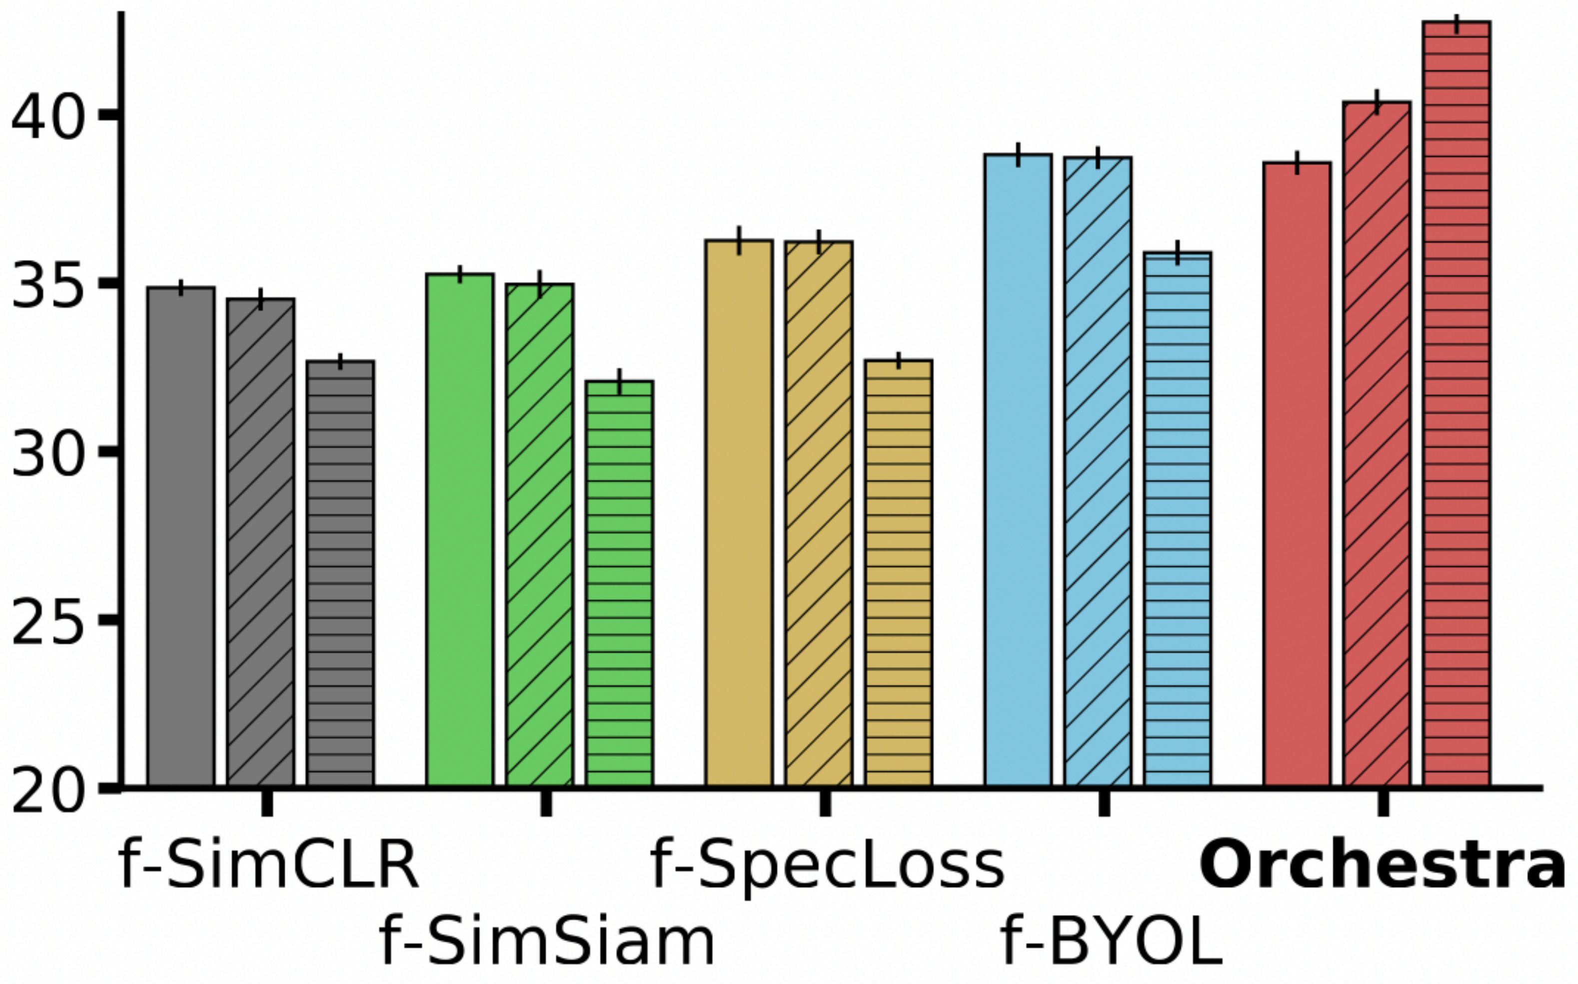

Right: CIFAR-100

Evaluation 1: Sensitivity to Heterogeneity

- Orchestra's robustness to heterogeneity arises from its use of local clustering, a task that becomes easier with more non-IID data!
  - See paper for detailed theoretical statements

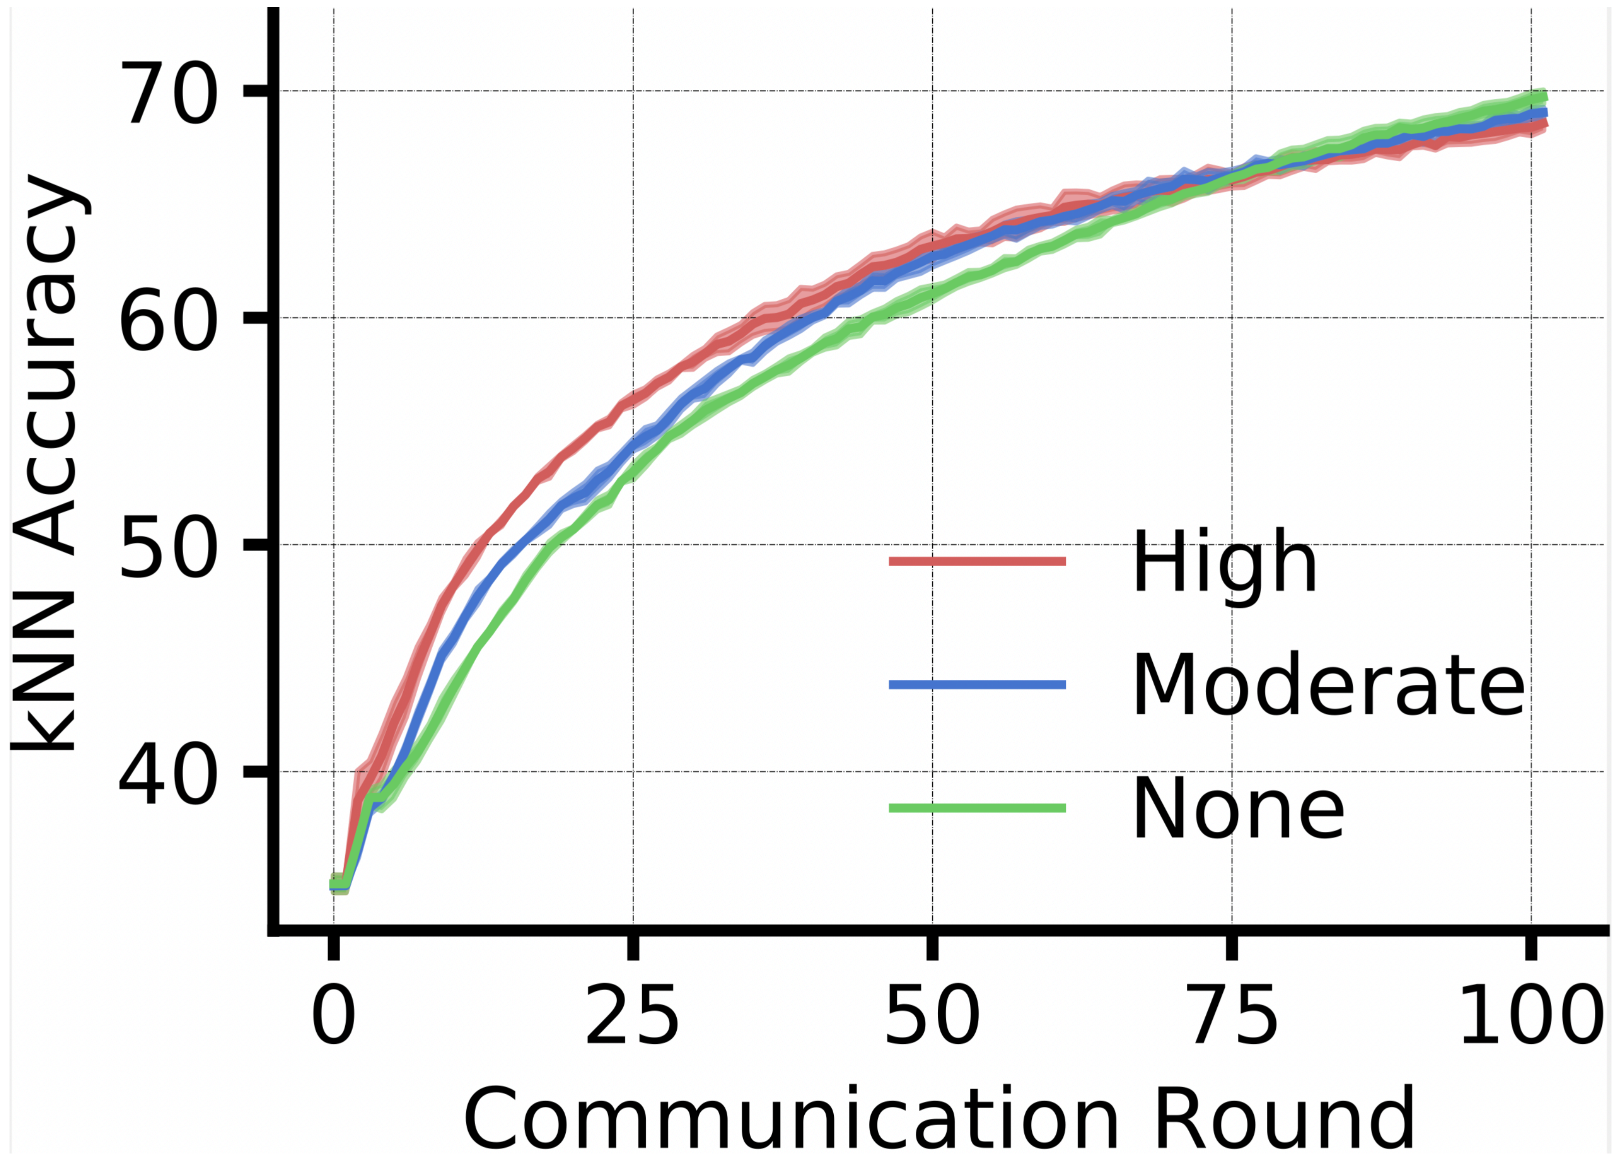

Left: CIFAR-10

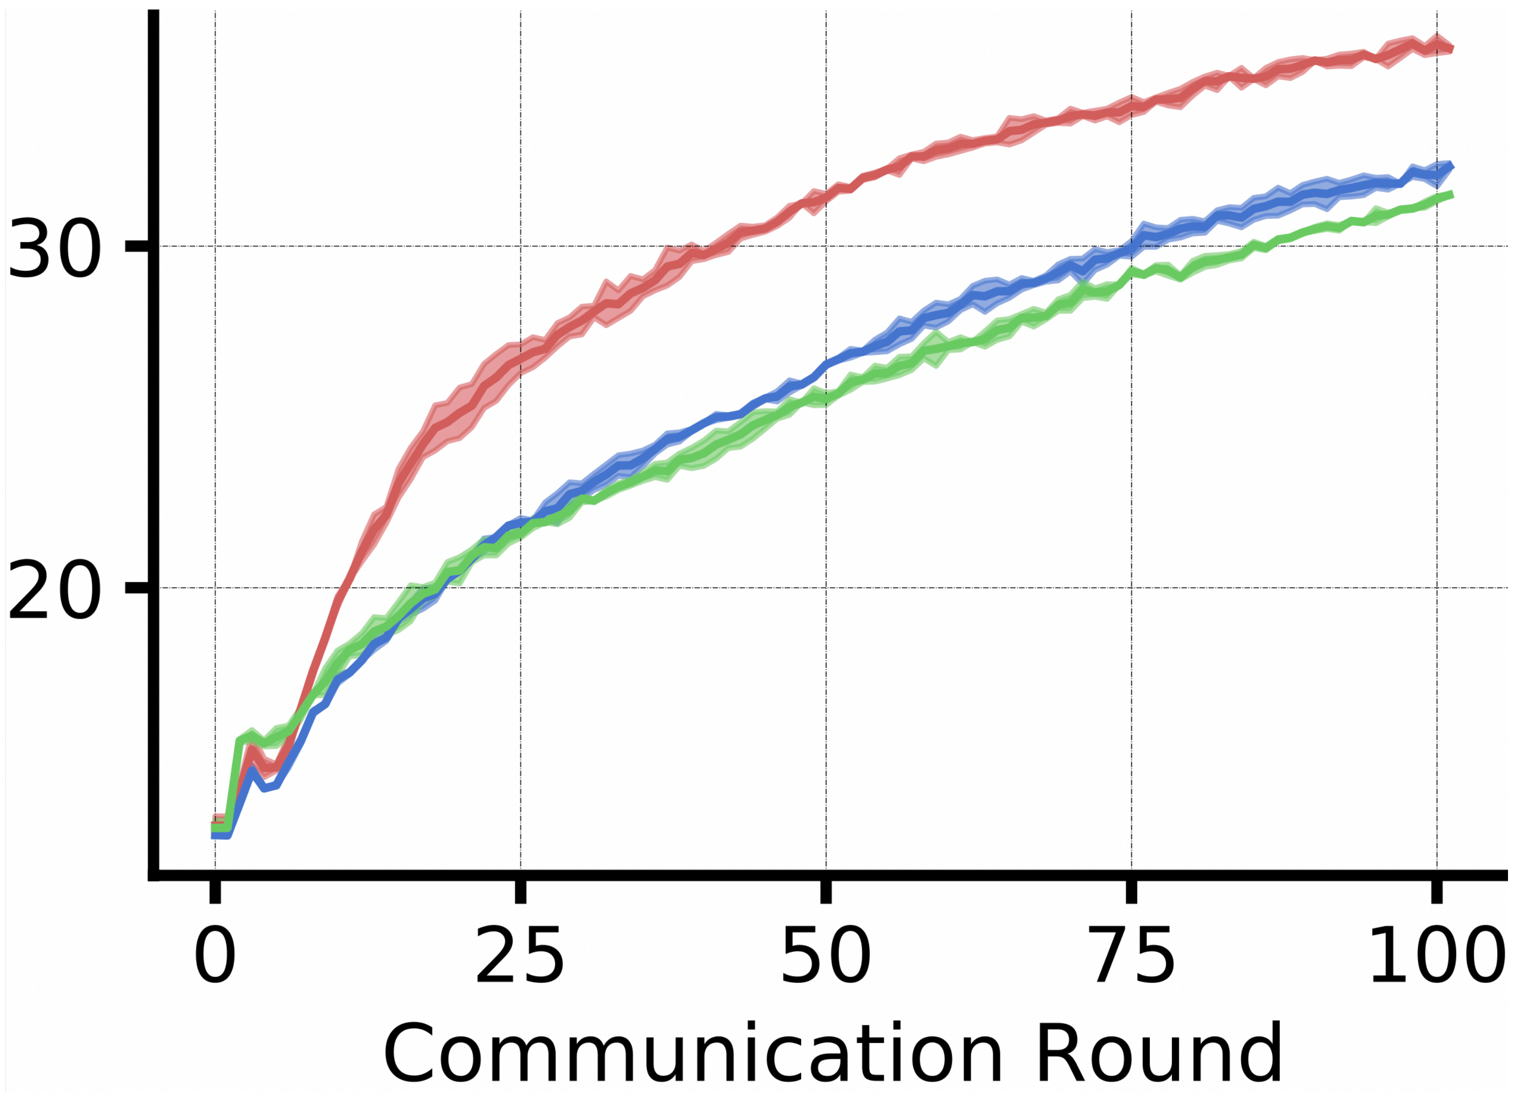

Right: CIFAR-100

Evaluation 2: Scalability with Number of Clients

- Much better absolute accuracy in all settings
- Scales well with large number of clients

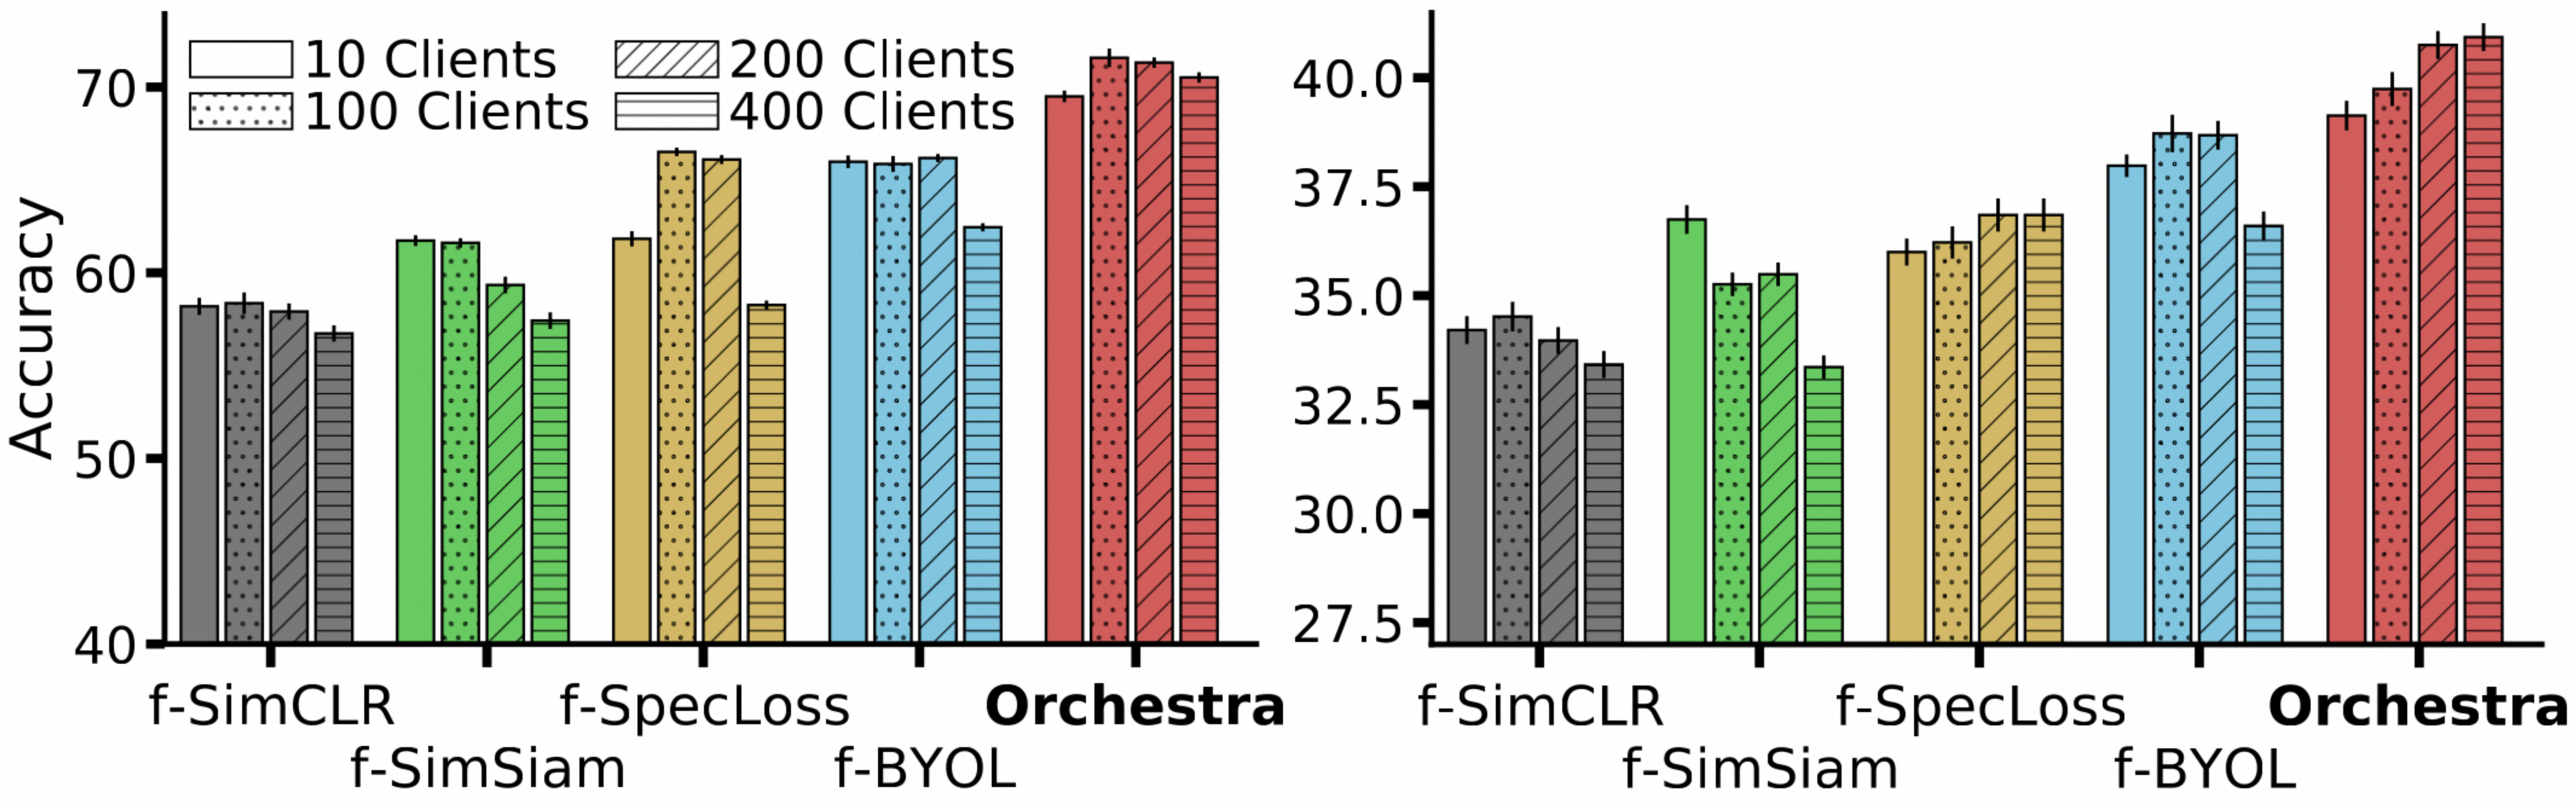

Left: CIFAR-10

Right: CIFAR-100

Evaluation 3: Participation Ratio

- Much higher robustness to participation ratio
- Especially more effective with smaller ratios, compared to other methods

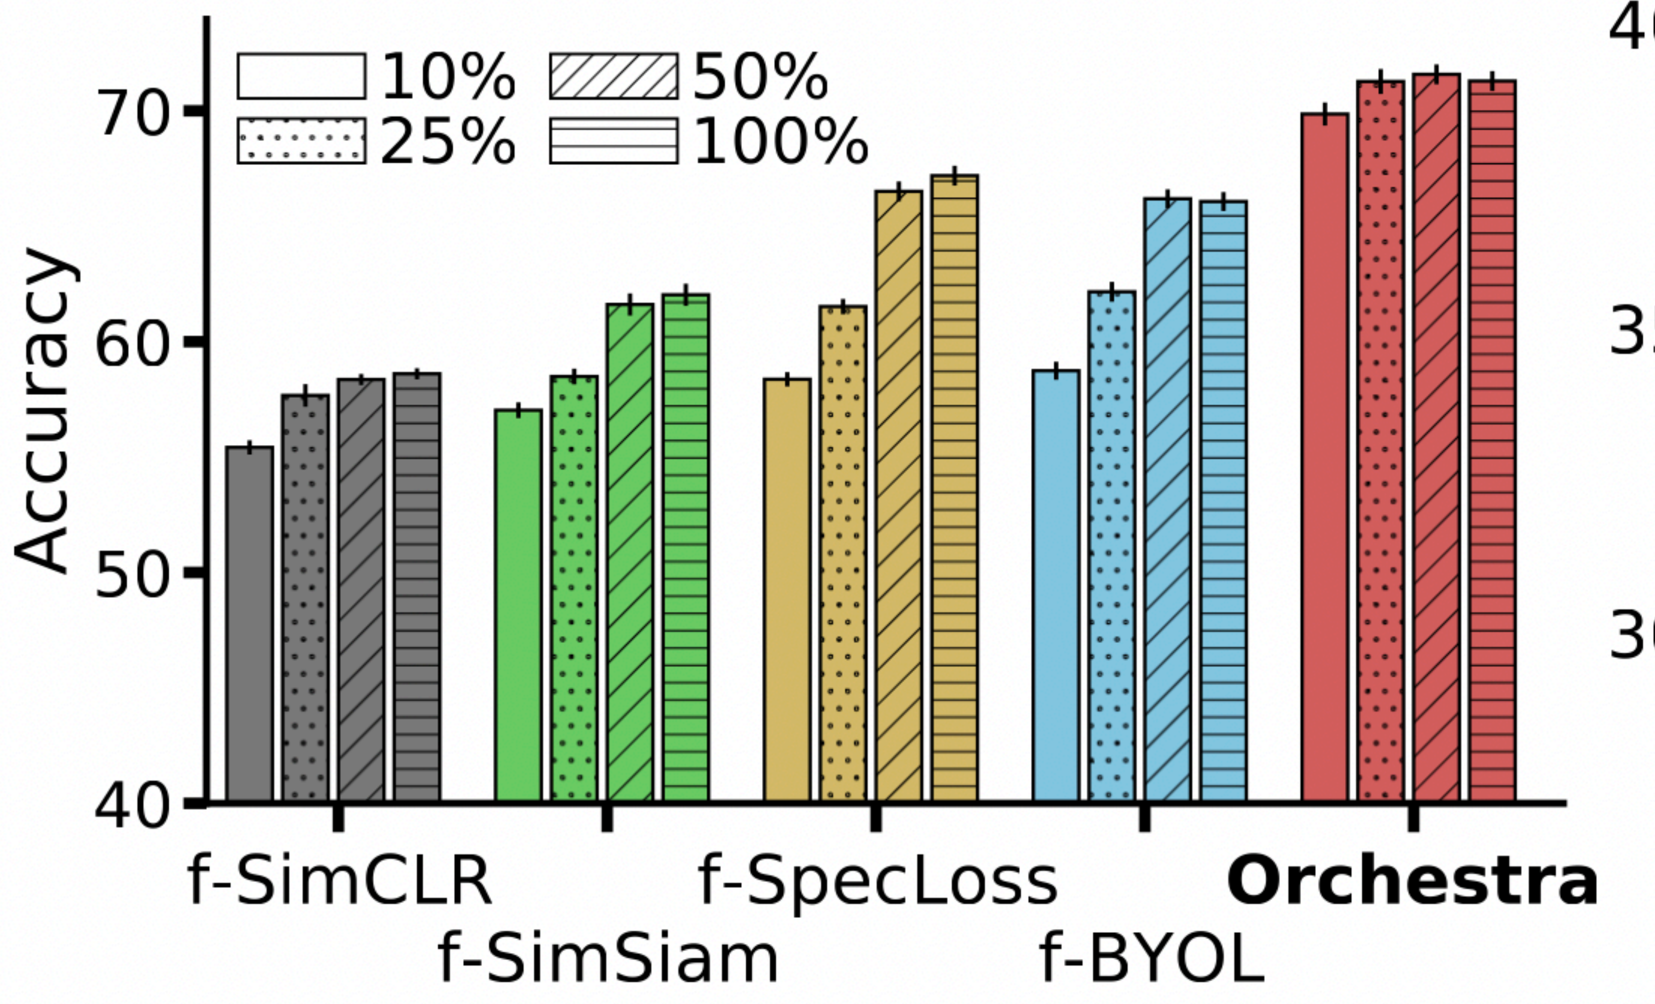

Left: CIFAR-10

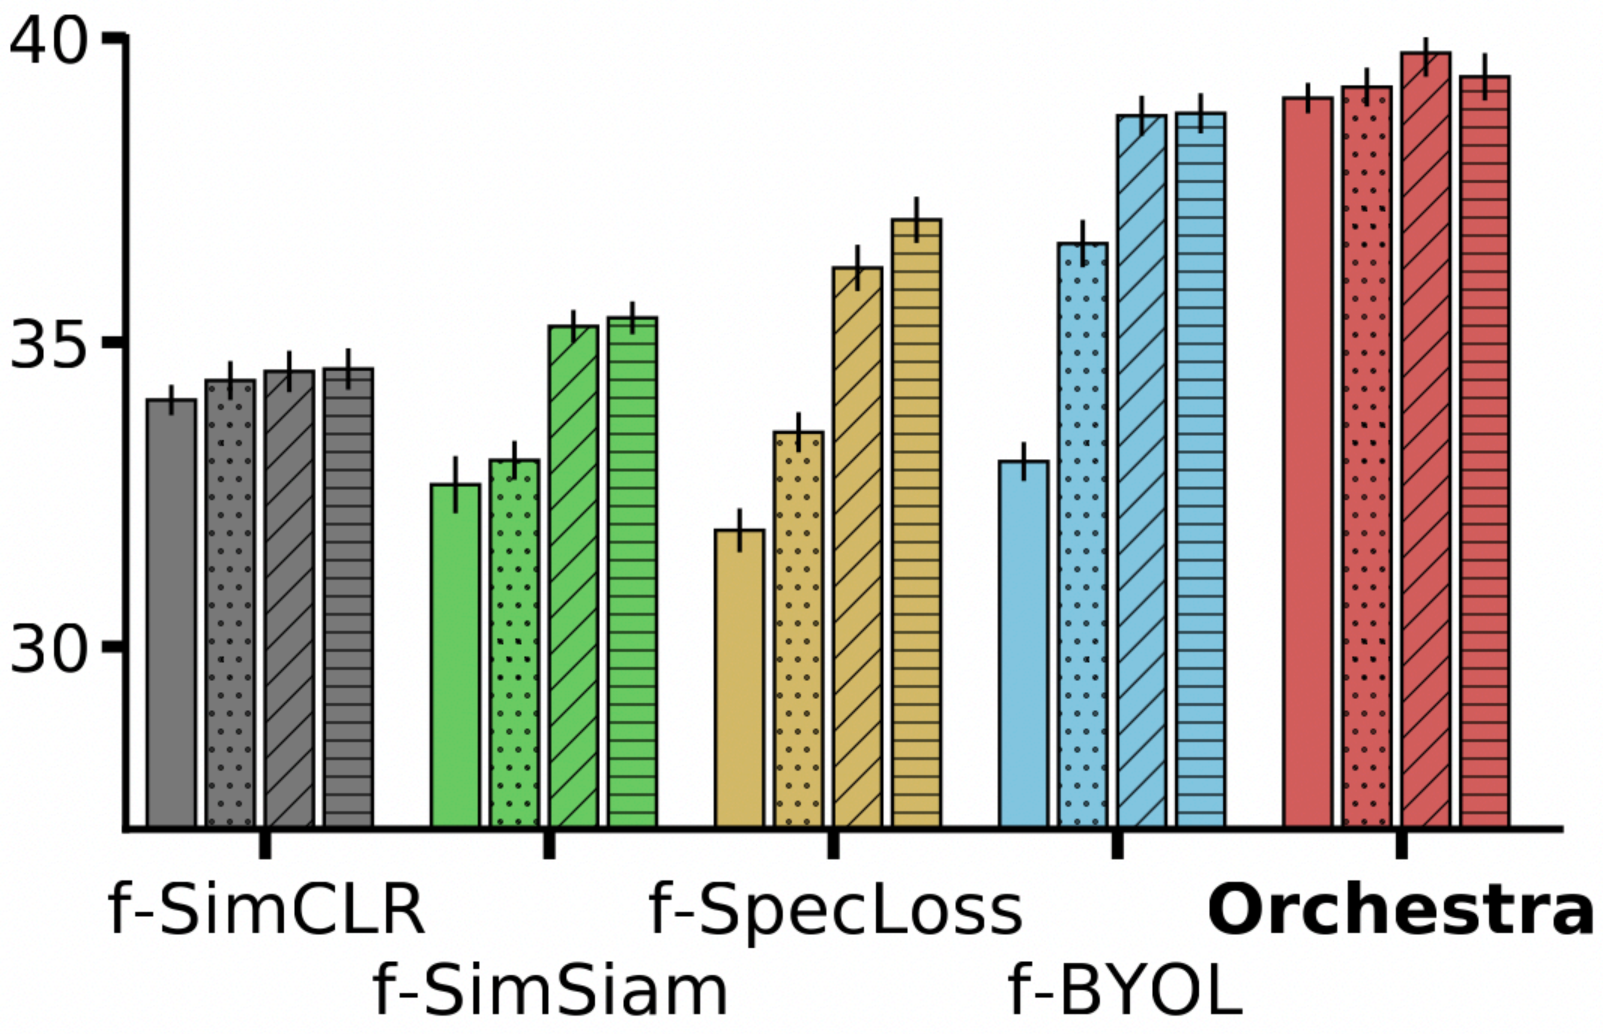

Right: CIFAR-100

Evaluation 4: Robustness to Local Epochs

- More compute efficient
- More communication efficient

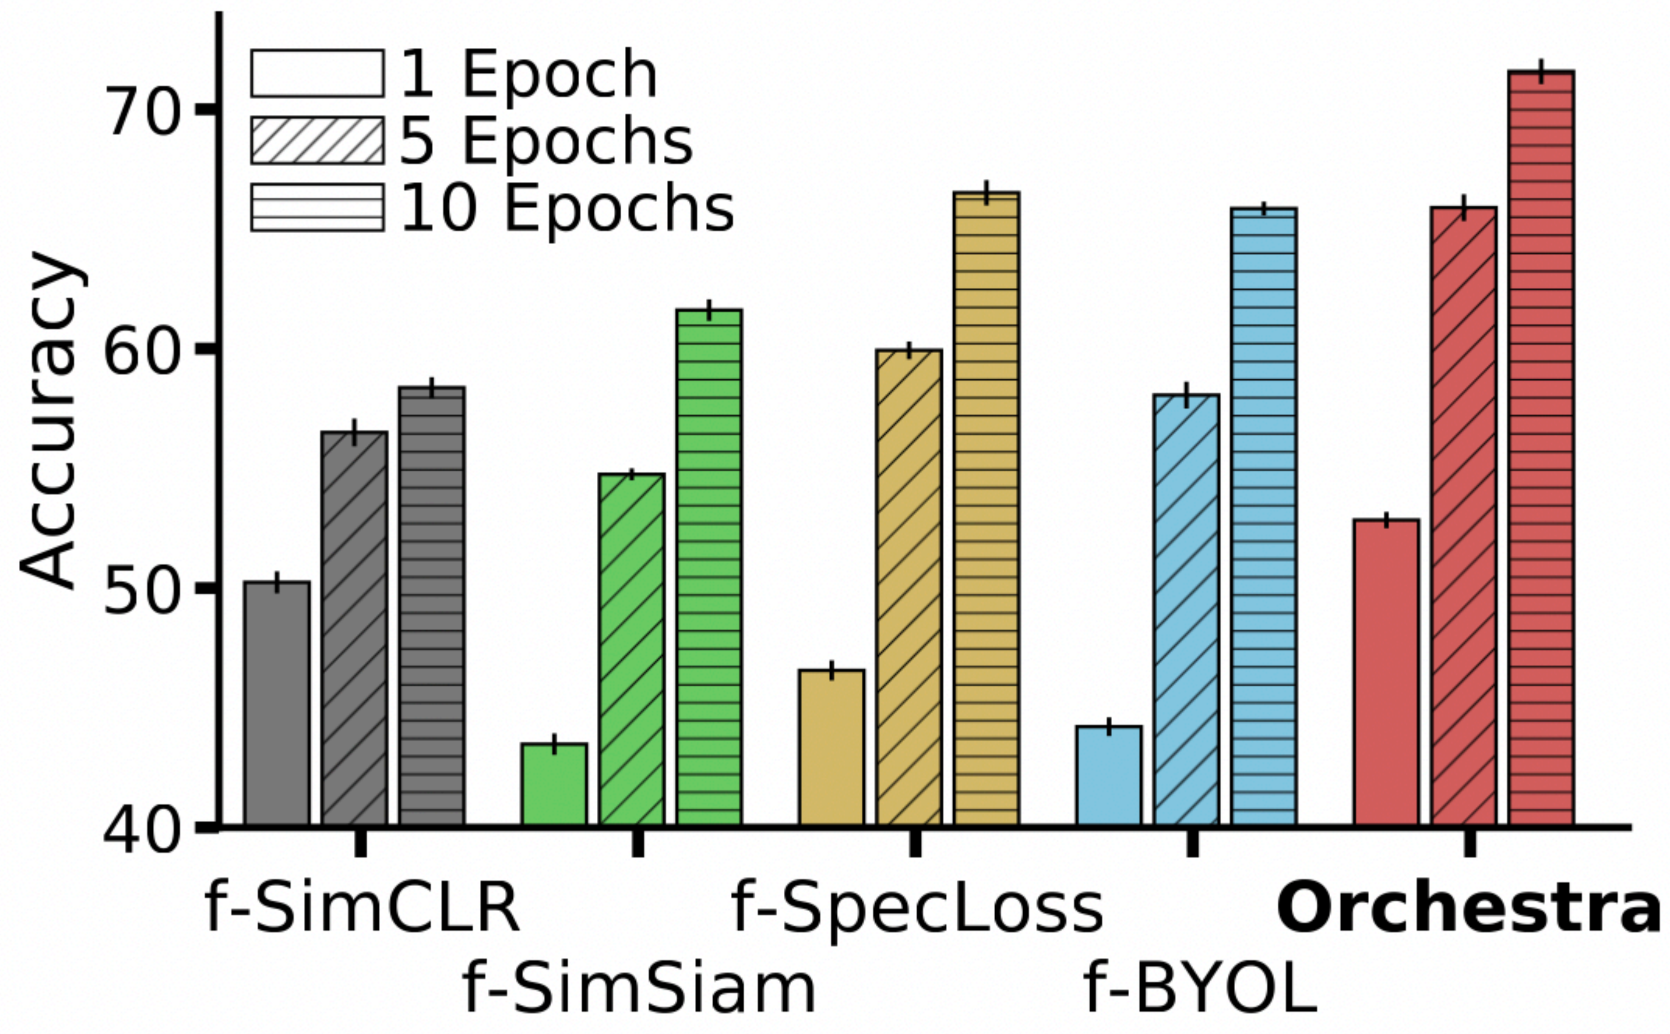

Left: CIFAR-10

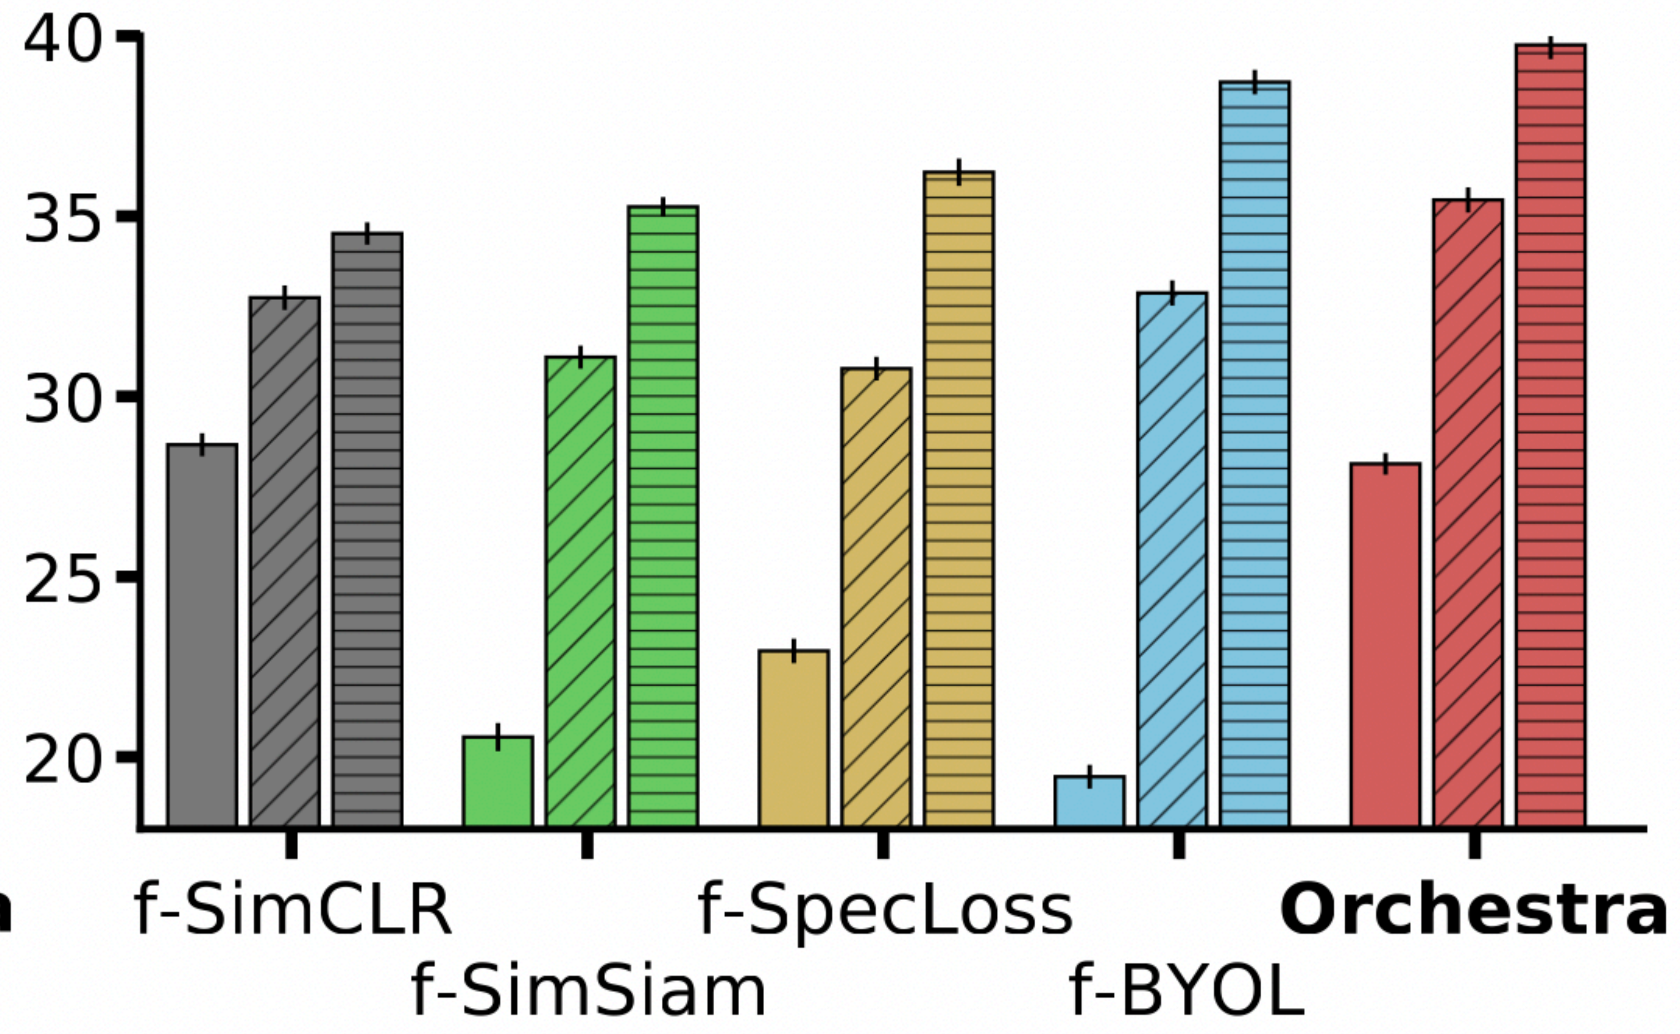

Right: CIFAR-100

Conclusion

- Orchestra provides a scalable, resource-efficient methodology for unsupervised federated learning
- The method is highly robust to heterogeneity and comes with guarantees
- For several more experiments, see the main paper

University of Michigan

University of Cambridge

Nokia Bell Labs