Murmura

An Evidential Trust-Aware Model Personalization Framework for Decentralized Federated Learning

Murtaza Rangwala, Richard O. Sinnott and Rajkumar Buyya

School of Computing and Information Systems
The University of Melbourne

1. Introduction to Federated Learning

Federated learning is a machine learning setting where **many clients** collaboratively **train a model** under the **orchestration of a central server**, while keeping the training **data decentralized**.

Iterative Protocol:

1. Eligible clients are selected by the server
2. Server sends the current model to selected clients
3. Clients train model on their private data
4. Server aggregates updates from each client
5. Global model is updated

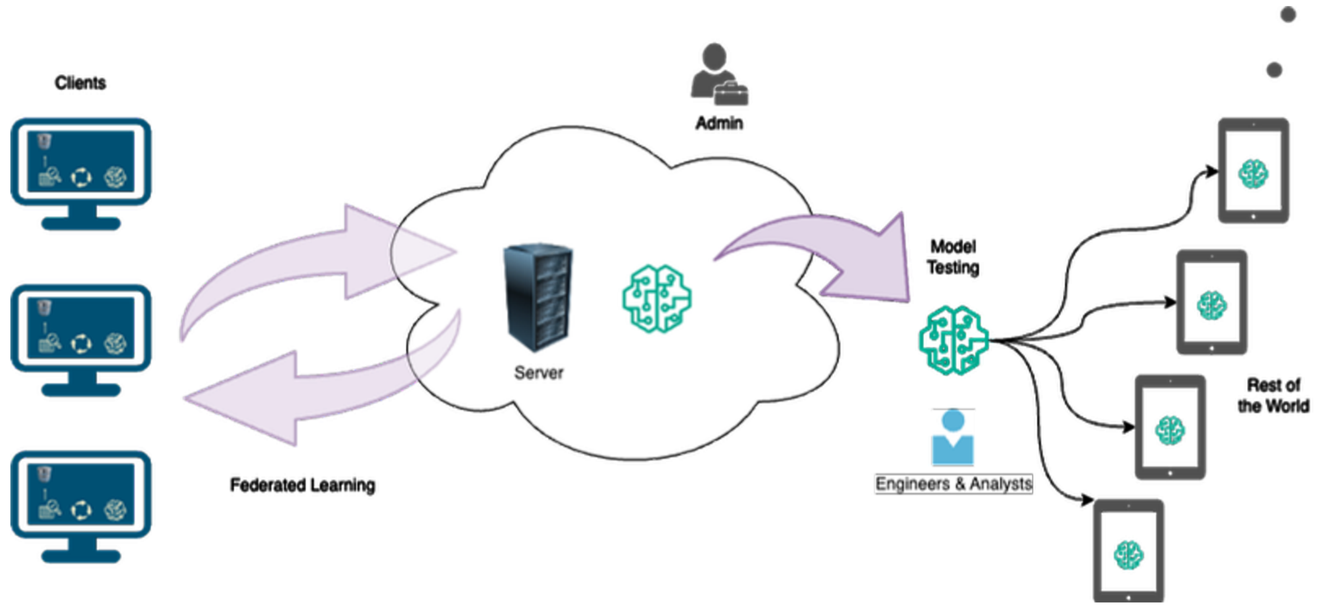

2. Problem with Standard Federated Learning

Reliance on a Centralized Entity

Single Point of Failure

- System-wide disruption if coordinator fails
- Critical bottleneck for scaling
- Performance bounded by central node

Trust Requirements

- Participants must trust the central entity
- Coordinator has privileged position
- Potential for manipulation of global model

Regulatory Challenges

- Cross-border data governance issues
- Compliance with regional regulations
- Questions of model ownership

3. Decentralized or P2P Learning

Decentralized learning is a machine learning paradigm where **multiple independent nodes** train and share model updates in a peer-to-peer network **without any central coordinator or server.**

- **Communication topology:** connected graph with clients as nodes, edges as channels
- **Rounds:** local updates + neighbor information exchange
- **Updates:** gradient steps; **communication:** parameter averaging
- No global model state, but local models converge to global solution
- Central authority may still define the learning task

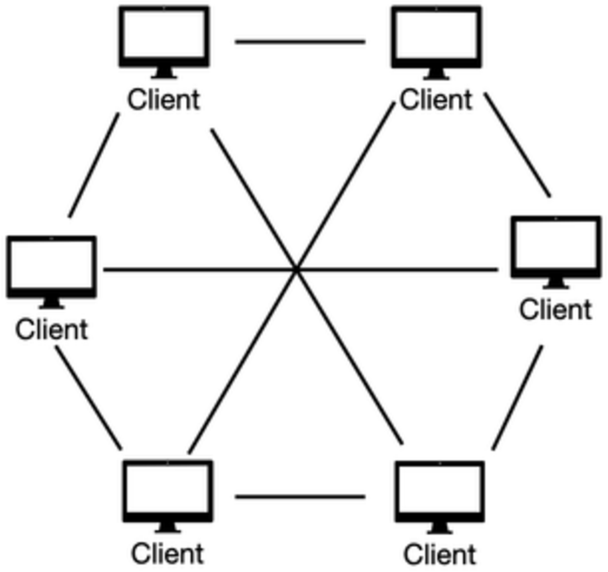

4. Core Challenges in Decentralization

Trust Establishment	• Establishing trust among distributed untrusting nodes
Resource Heterogeneity	• Dealing with nodes that have different computational power levels
Communication Efficiency	• Managing limited bandwidth for model updates
Model Convergence	• Achieving model consensus across distributed training nodes
Data Privacy	• Protecting node-specific data while enabling collaborative learning
Regulatory Compliance	• Navigating varied legal requirements across jurisdictions

5. The Trust Problem in DFL

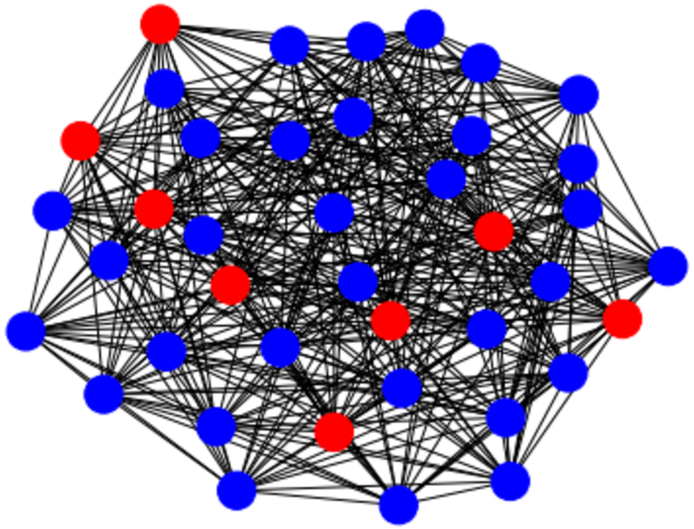

Untrustworthy nodes may be part of a DFL network

Each node needs to make a decision on which neighbors to trust

Wrong decisions can propogate bad updates across the network

In sparse network topologies, ignoring a malicious node can inadvertently isolate an honest subset of the network.

6. The Model Personalization Problem in DFL

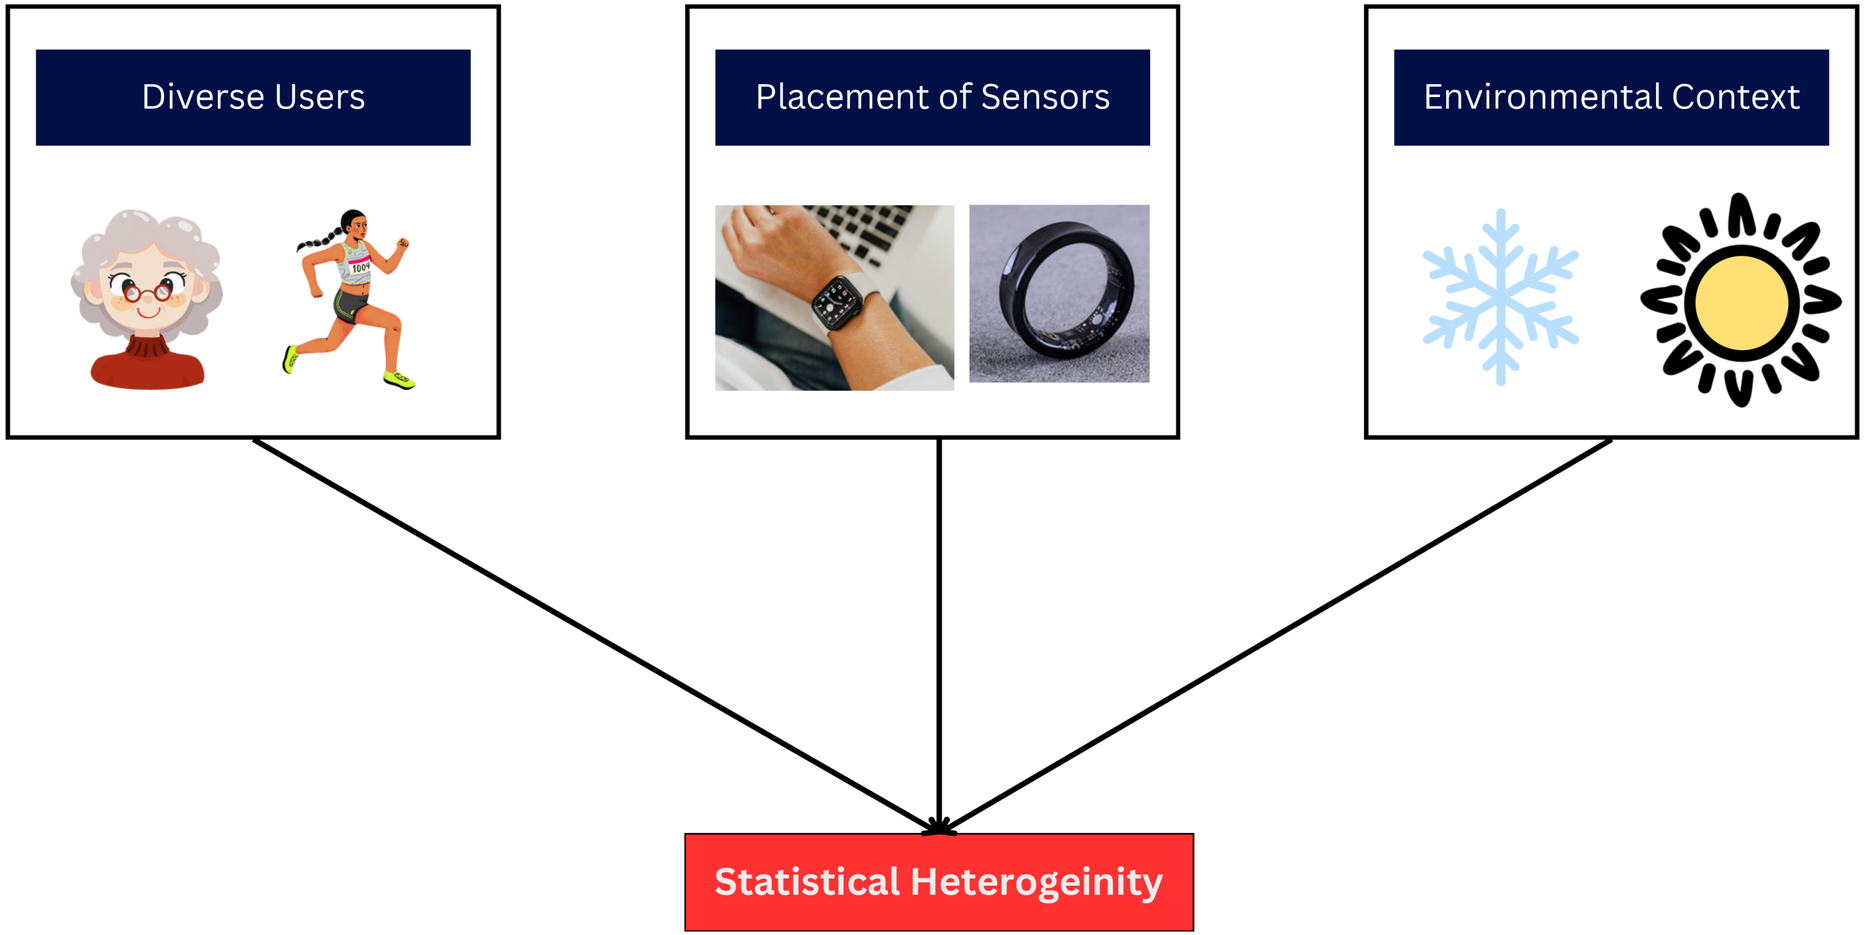

6. The Model Personalization Problem in DFL

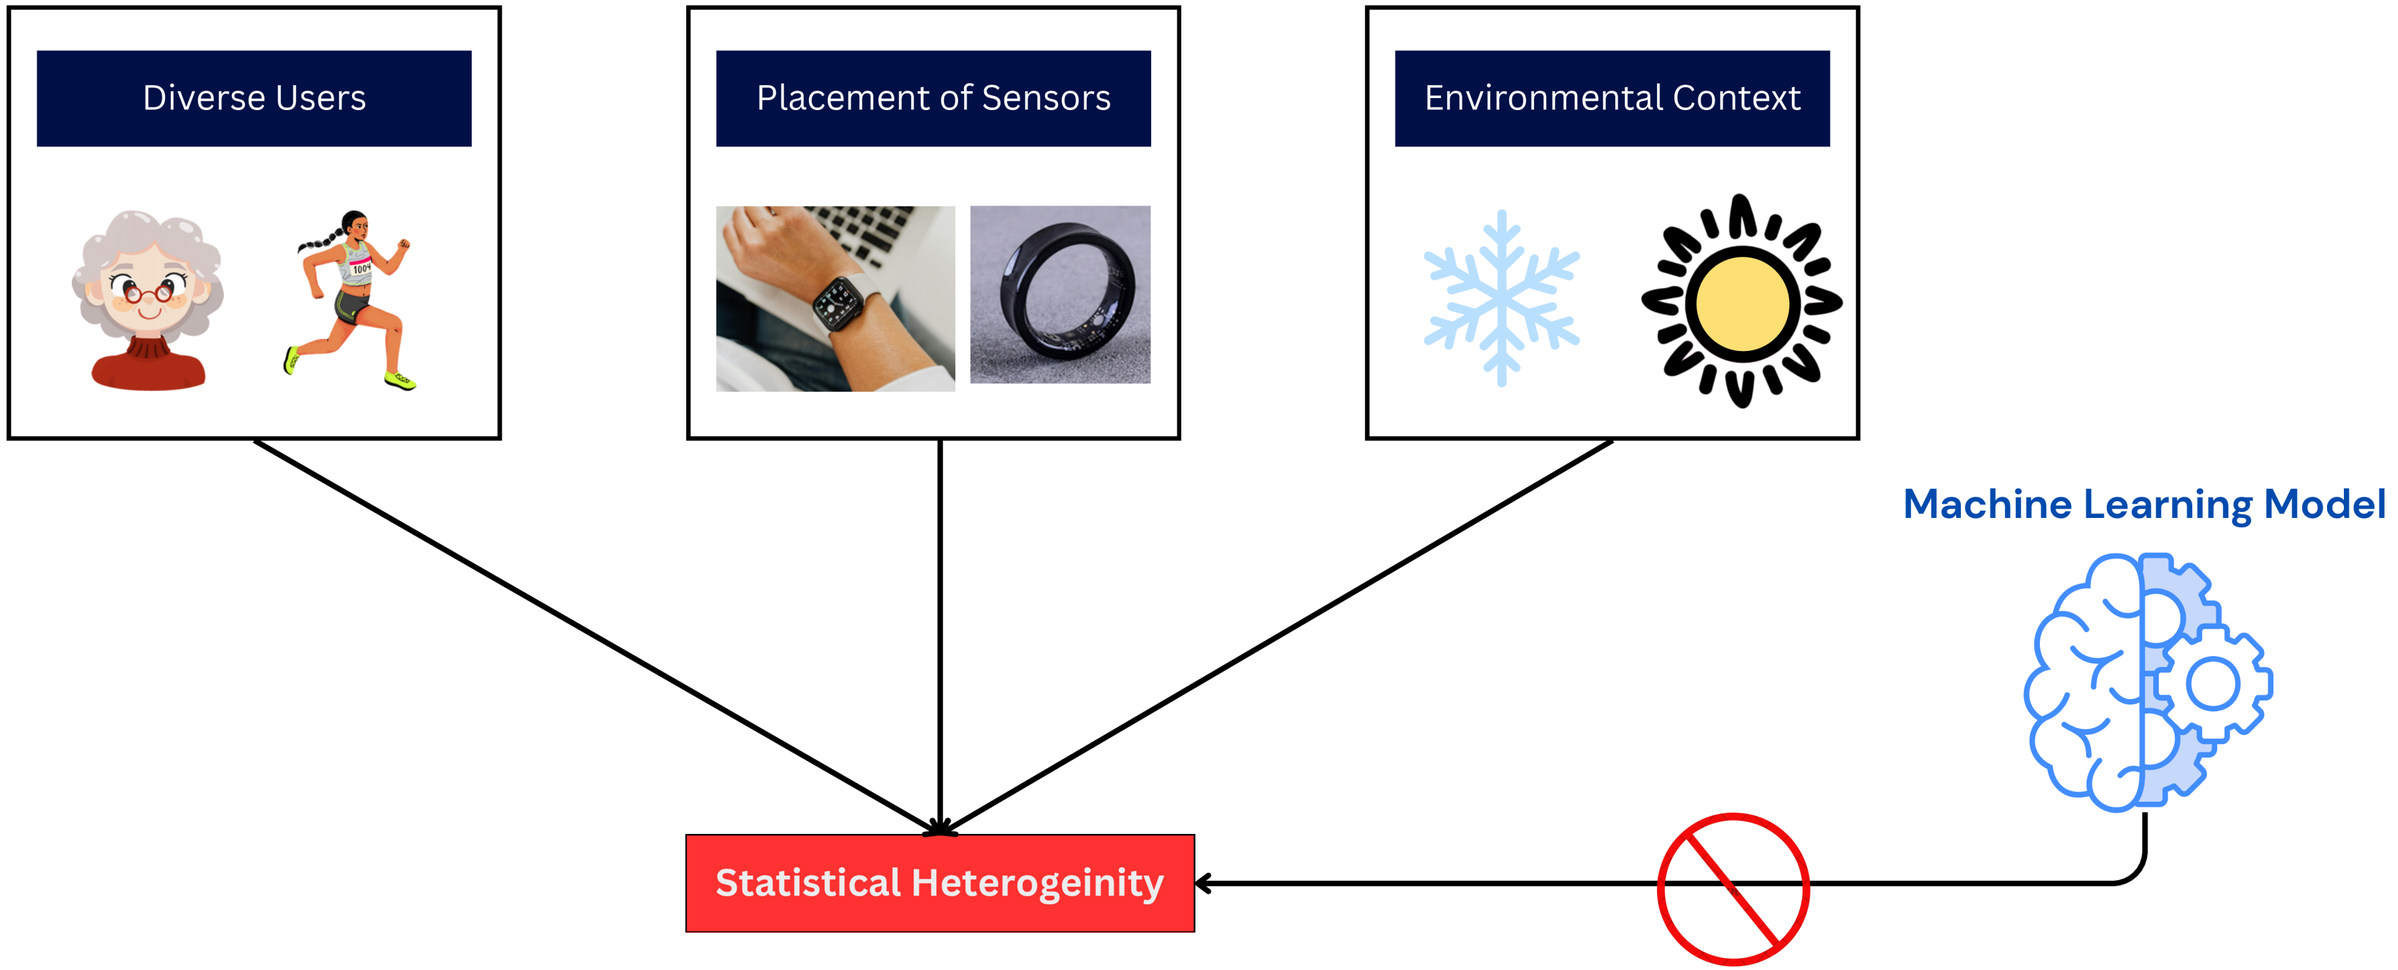

Each user needs a personalized model that works well on their data while still benefitting from collaboration with peers.

7. Existing Literature

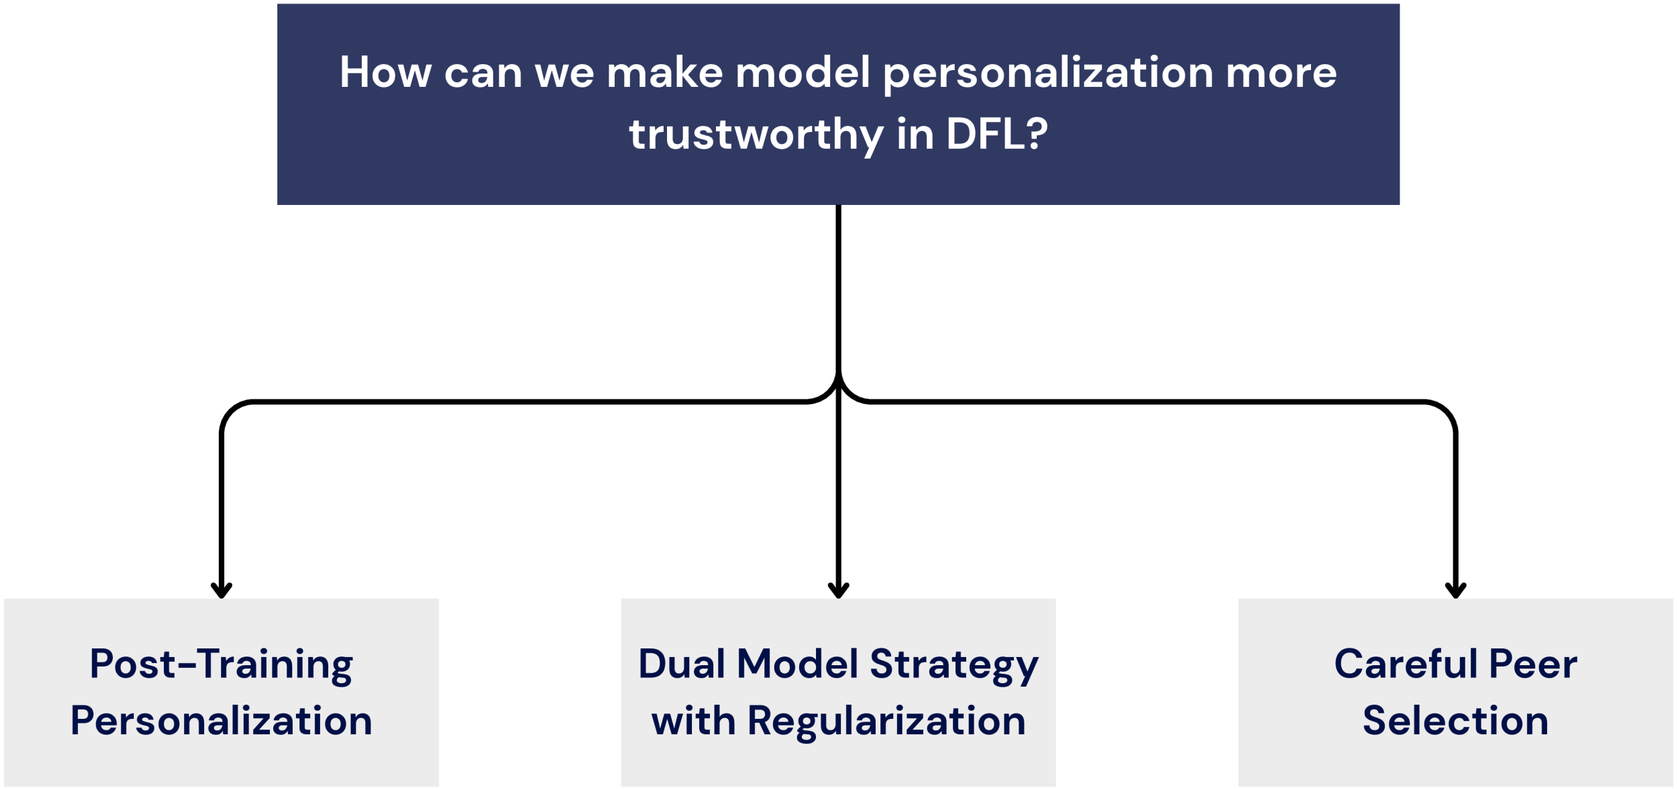

7. Existing Literature

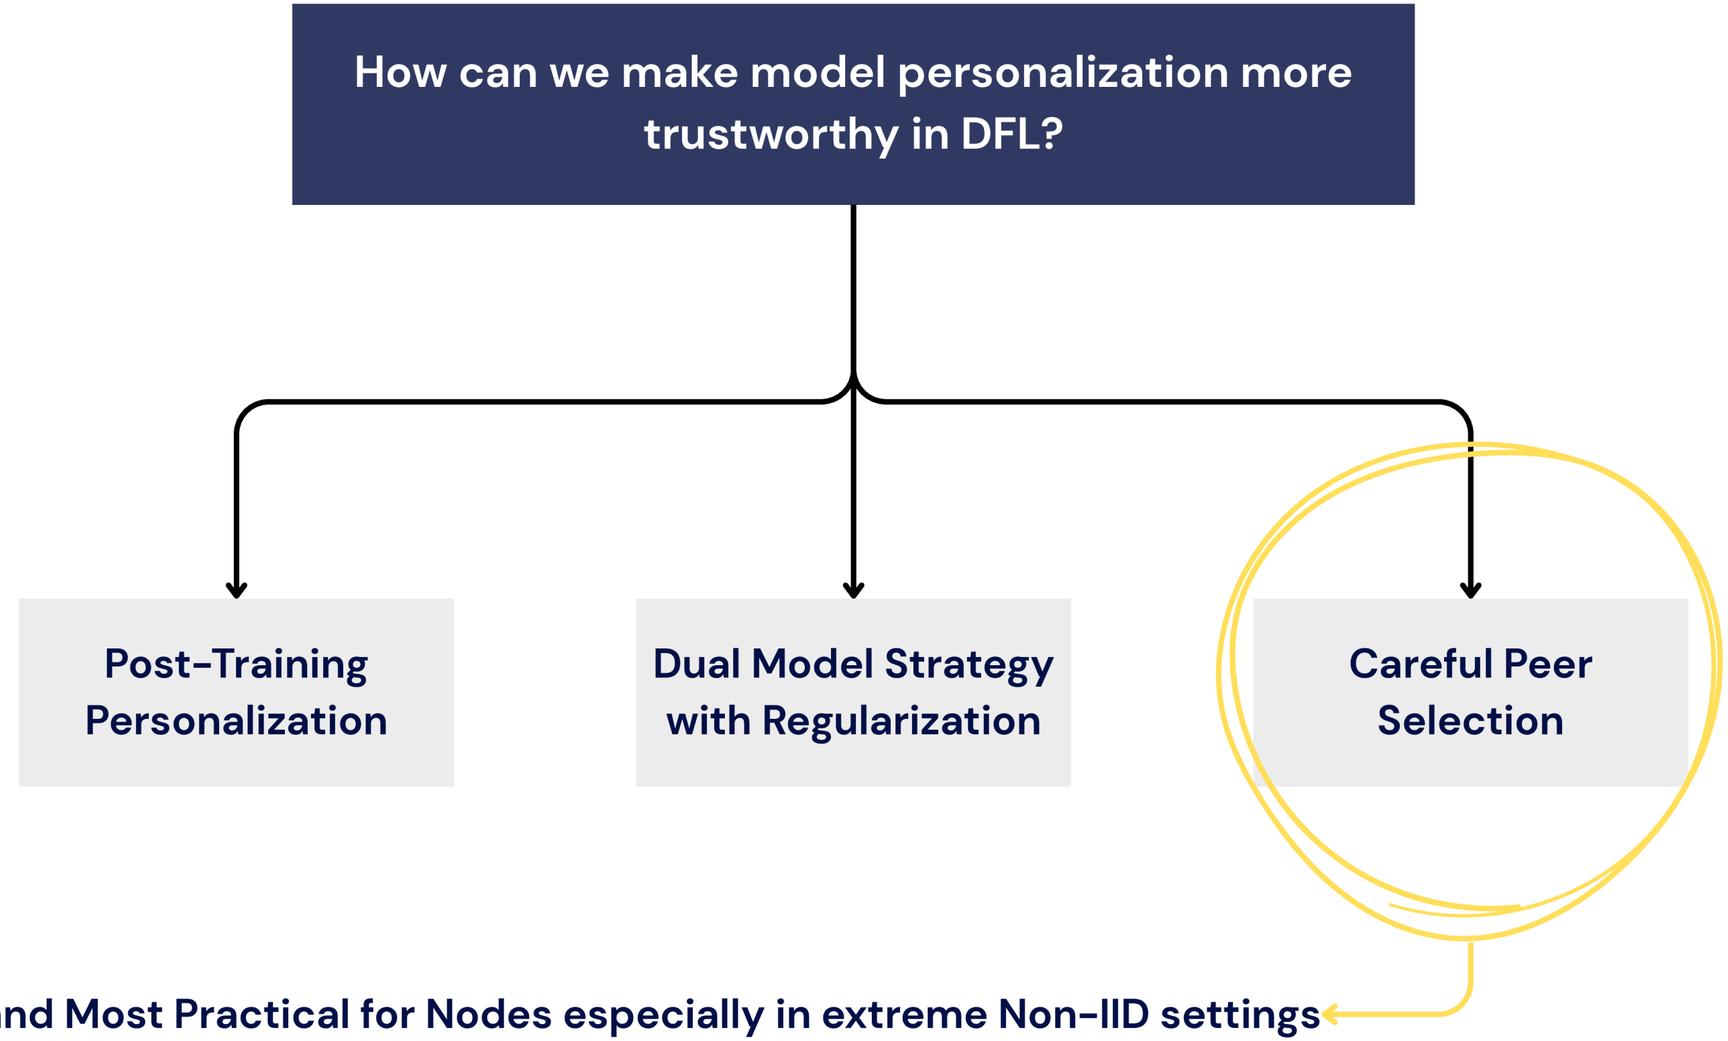

8. State of the Art – UBAR

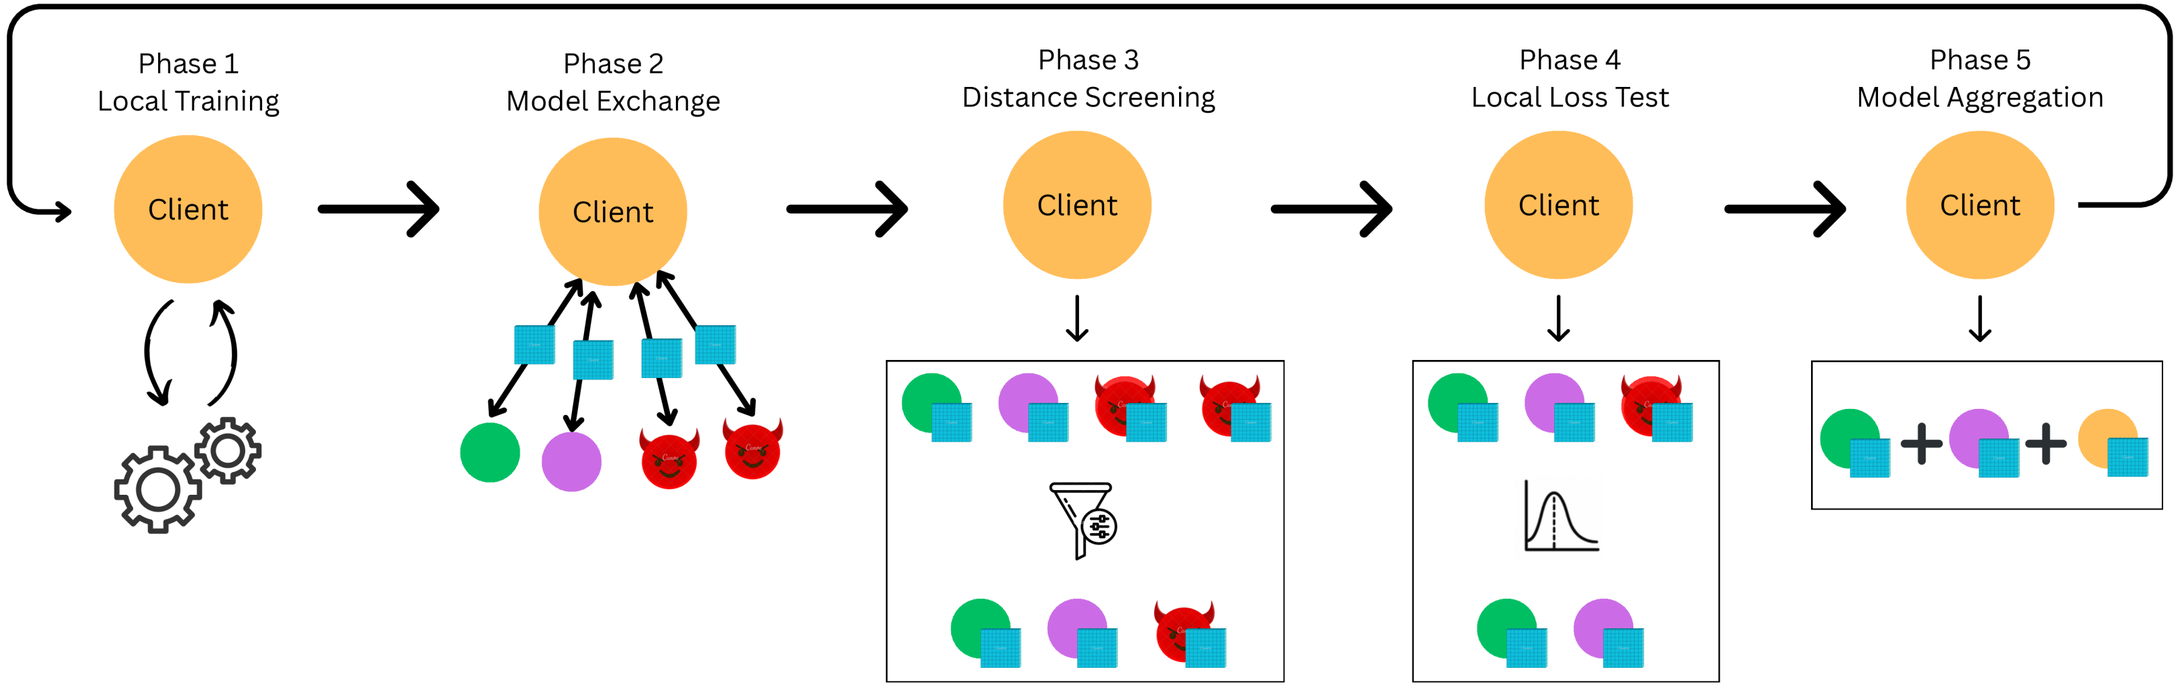

Distance Screening

Each client compares its model with neighbors and keeps only the closest subset. This fast, low-cost step removes distant or inconsistent models, cutting off obvious Byzantine outliers before deeper checks.

Local Loss Test

The shortlisted neighbors are re-evaluated using the client's own mini-batch. Only those whose models perform no worse than the client's local model are trusted for aggregation, ensuring robustness against subtle, data-aware attacks.

8. State of the Art – UBAR

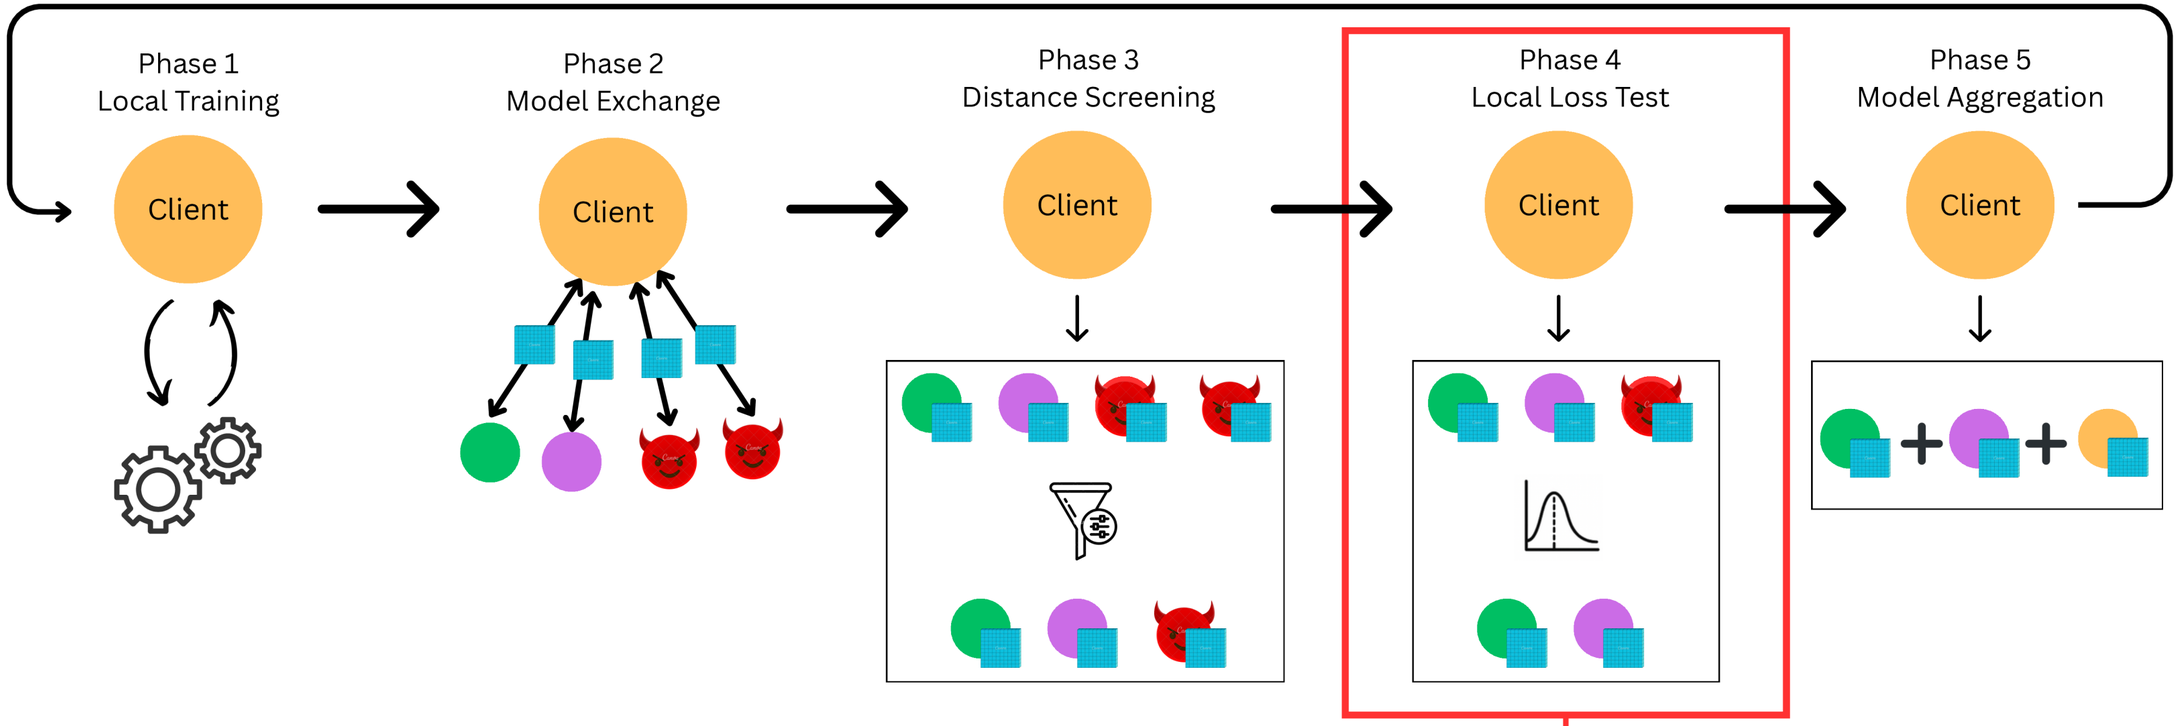

Computationally Heavy and Slow

9. State of the Art – BALANCE

Excludes the **Local Loss Test**

Introduces a **dynamic thresholding mechanism** with **a decay factor**

$$\|w_i^{t+\frac{1}{2}} - w_j^{t+\frac{1}{2}}\| \leq \gamma \cdot \exp(-\kappa \cdot \lambda(t)) \|w_i^{t+\frac{1}{2}}\|$$

Upper limit for accepting a model as benign

Determines the rate at which the exponential function decreases; a larger κ results in a faster decay, while a smaller κ leads to a slower decay

A monotonically increasing and non-negative function associated with the training round index

10. Problem with SOTA

Critical Challenge in Trust-Aware Model Personalization

- Distinguishing between harmful dissimilarity and beneficial diversity.
- Peers may have genuinely incompatible data distributions, insufficient training quality, or complementary diversity that could enhance generalization.

UBAR ⟶ Captures this distinction but computationally expensive

BALANCE ⟶ Unable to capture this distinction but computationally less expensive

11. An Introduction to Evidential Deep Learning

Regular Deep Learning Model

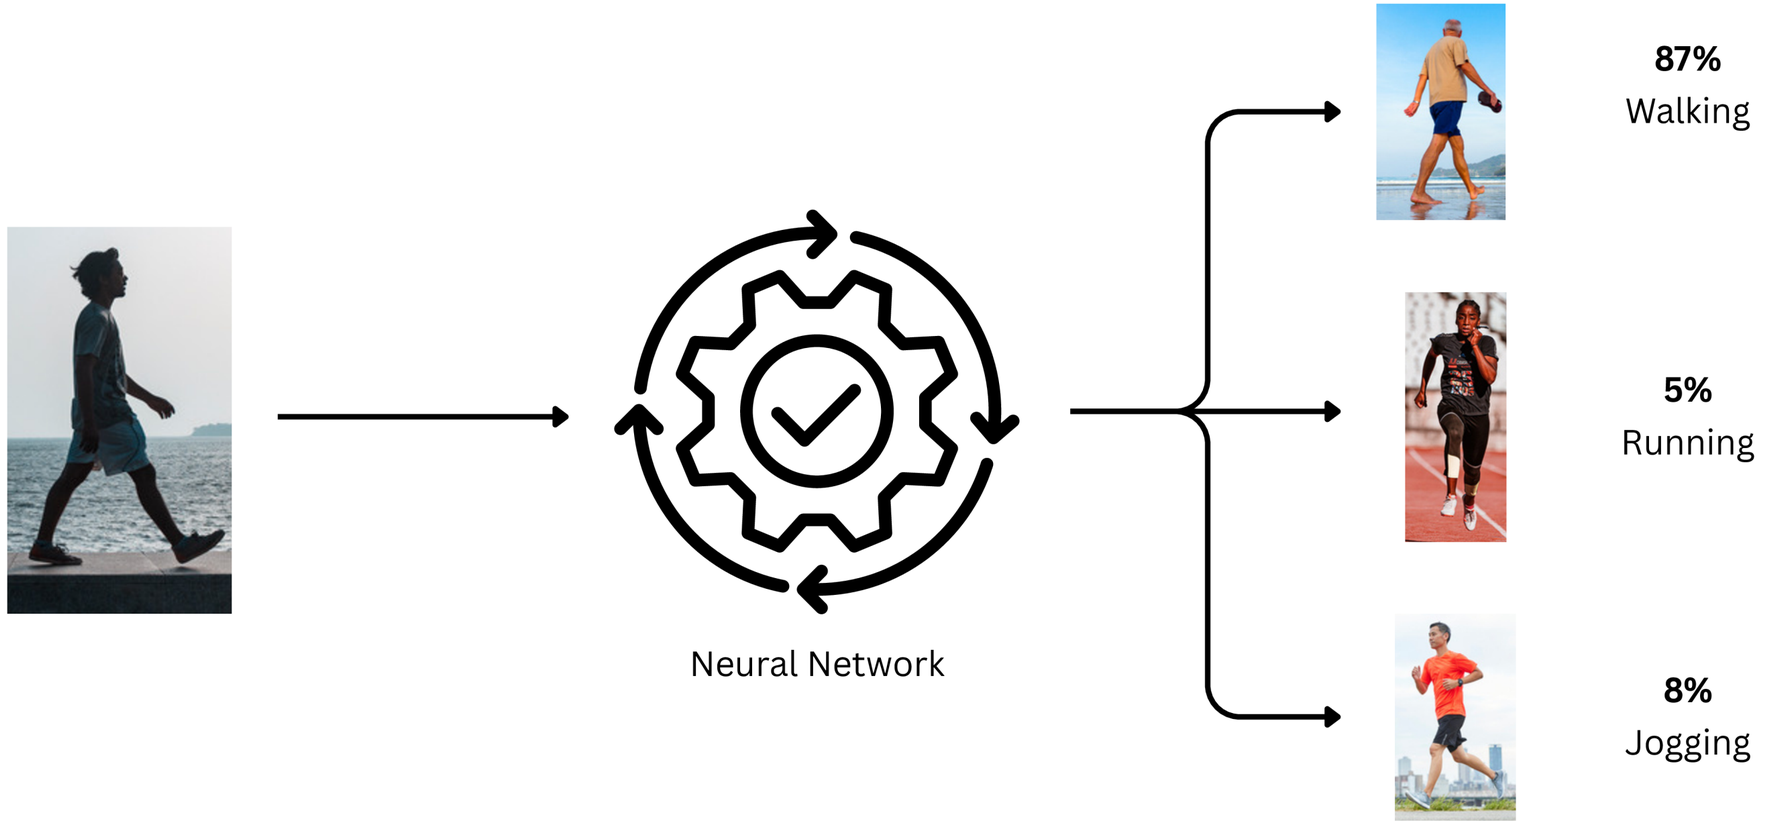

11. An Introduction to Evidential Deep Learning

Regular Deep Learning Model

Neural Network

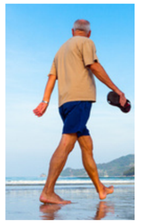

87%
Walking

5%
Running

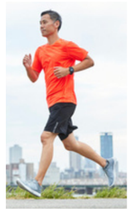

8%
Jogging

11. An Introduction to Evidential Deep Learning

Regular Deep Learning Model

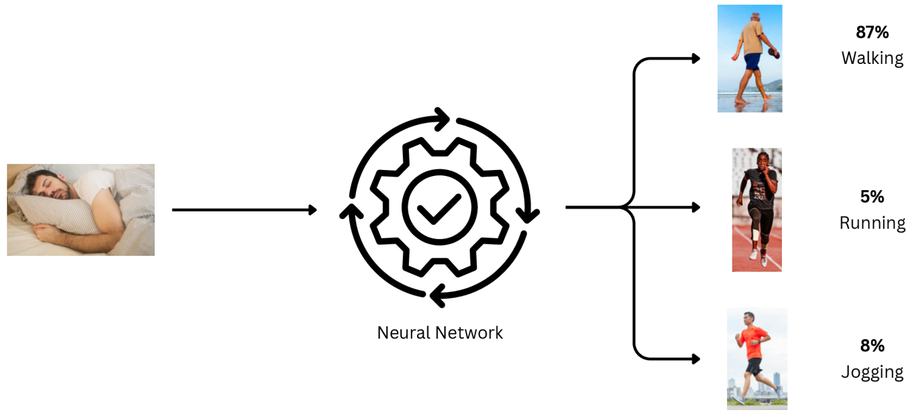

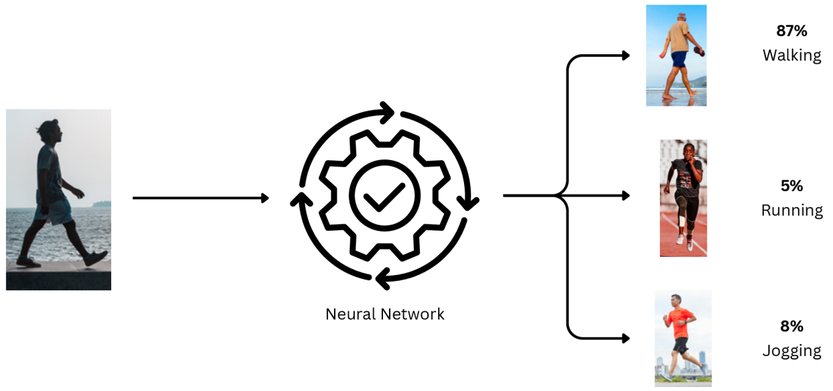

A regular deep learning model gives us NO information about confidence or uncertainty of the prediction

Is the model confident? Did it see tons of similar examples during training, or is it guessing?

Why is walking 87% for both images? **model is confident**, **data is ambiguous or the model has never seen the data before**

11. An Introduction to Evidential Deep Learning

THE UNIVERSITY OF MELBOURNE

What is Evidential Deep Learning?

EDL networks output evidence values that allow us to quantify why a model makes a prediction that it makes.

*It gives us an insight into **two** distinct types of uncertainty*

Epistemic Uncertainty

Lack of knowledge due to insufficient training

Has the model seen similar data?

Aleatoric Uncertainty

Inherent ambiguity in the data itself

Is the data genuinely challenging?

How can we leverage EDL for Trust-Aware Model Personalization?

12. EDL for Trust-Aware Model Personalization

If a neighbor's model shows ↑ aleatoric uncertainty on my data,

it means my data is genuinely challenging or ambiguous but it **does not indicate** that my neighbor is untrustworthy

If a neighbor's model shows ↑ epistemic uncertainty on my data,

it means my neighbor has not seen enough similar data before. This could be due to insufficient training, training on non-identical data or malicious intent. **My neighbor is untrustworthy** and should probably be excluded.

13. The Murmura EDL-Based Framework

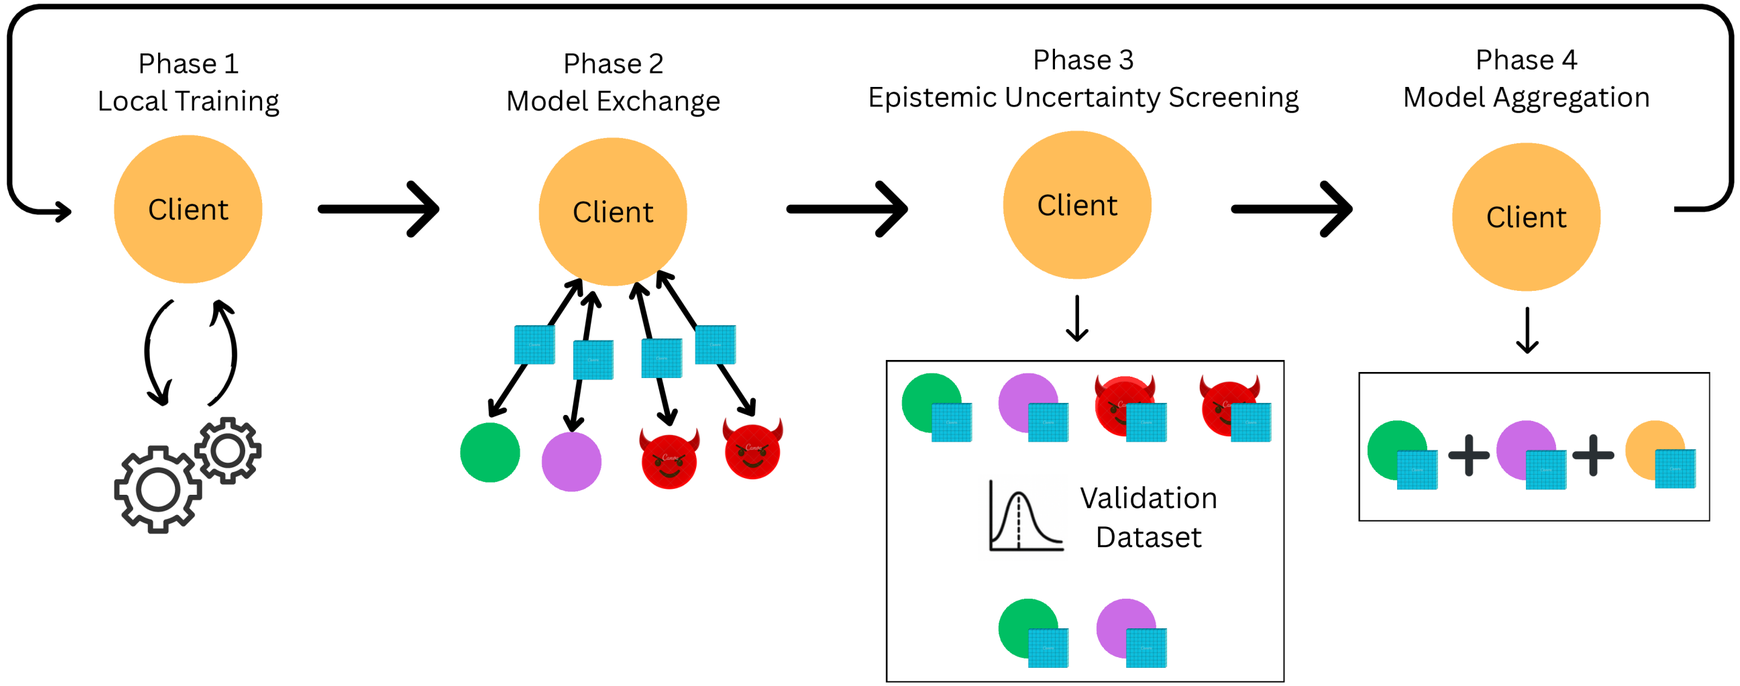

Murmura uses the same dynamic thresholding as BALANCE but for epistemic uncertainty instead of L2 (Euclidean) distance

This makes intuitive sense since the neighbors' models will start off with high epistemic uncertainty which might decrease as the training rounds progress. The threshold starts off relaxed and then tightens as the rounds progress.

14. Empirical Evaluation: Datasets

UCI HAR

Smartphone accelerometer and gyroscope data from 30 subjects performing 6 activities (walking, walking upstairs, walking downstairs, sitting, standing, lying). Contains 10,299 samples with 561 extracted features.

PAMAP2

IMU data from 9 subjects performing 12 activities including household and exercise tasks. We extract 40 features per sliding window from three body-worn sensors.

PPG-DaLiA

Photoplethysmography and accelerometer signals from 15 subjects under real-life conditions. We formulate activity classification from 8 activity types with time-frequency features.

15. Personalization Under Heterogeneity

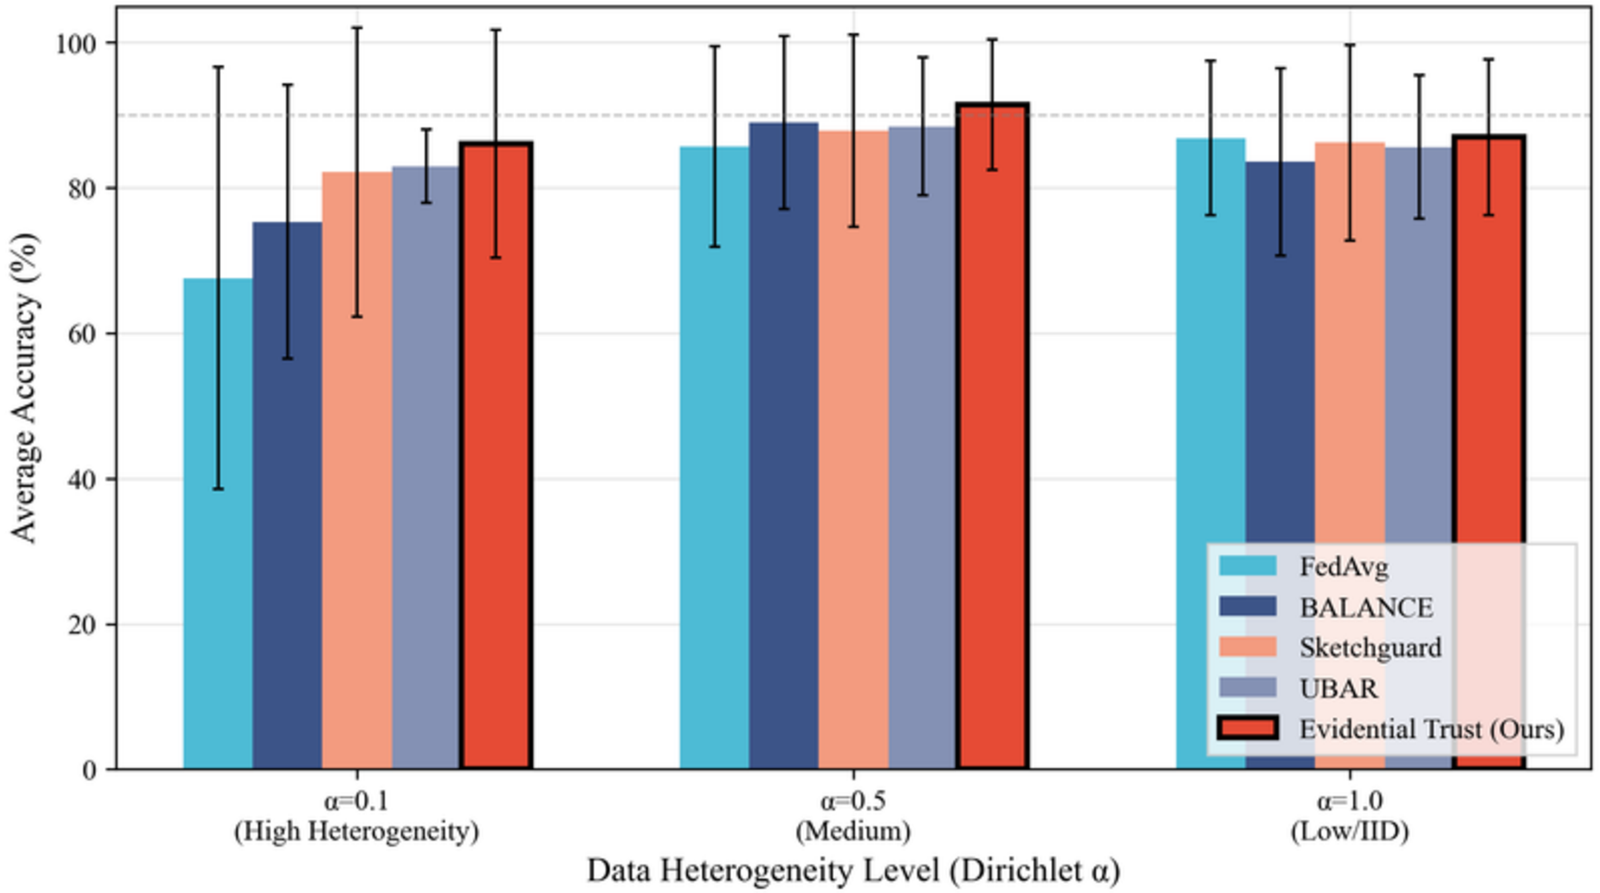

Model accuracy across data heterogeneity levels (Dirichlet α), averaged across all three datasets.

Lower α indicates higher heterogeneity.

Murmura (Evidential Trust) maintains consistent performance as heterogeneity increases.

15. Personalization Under Heterogeneity

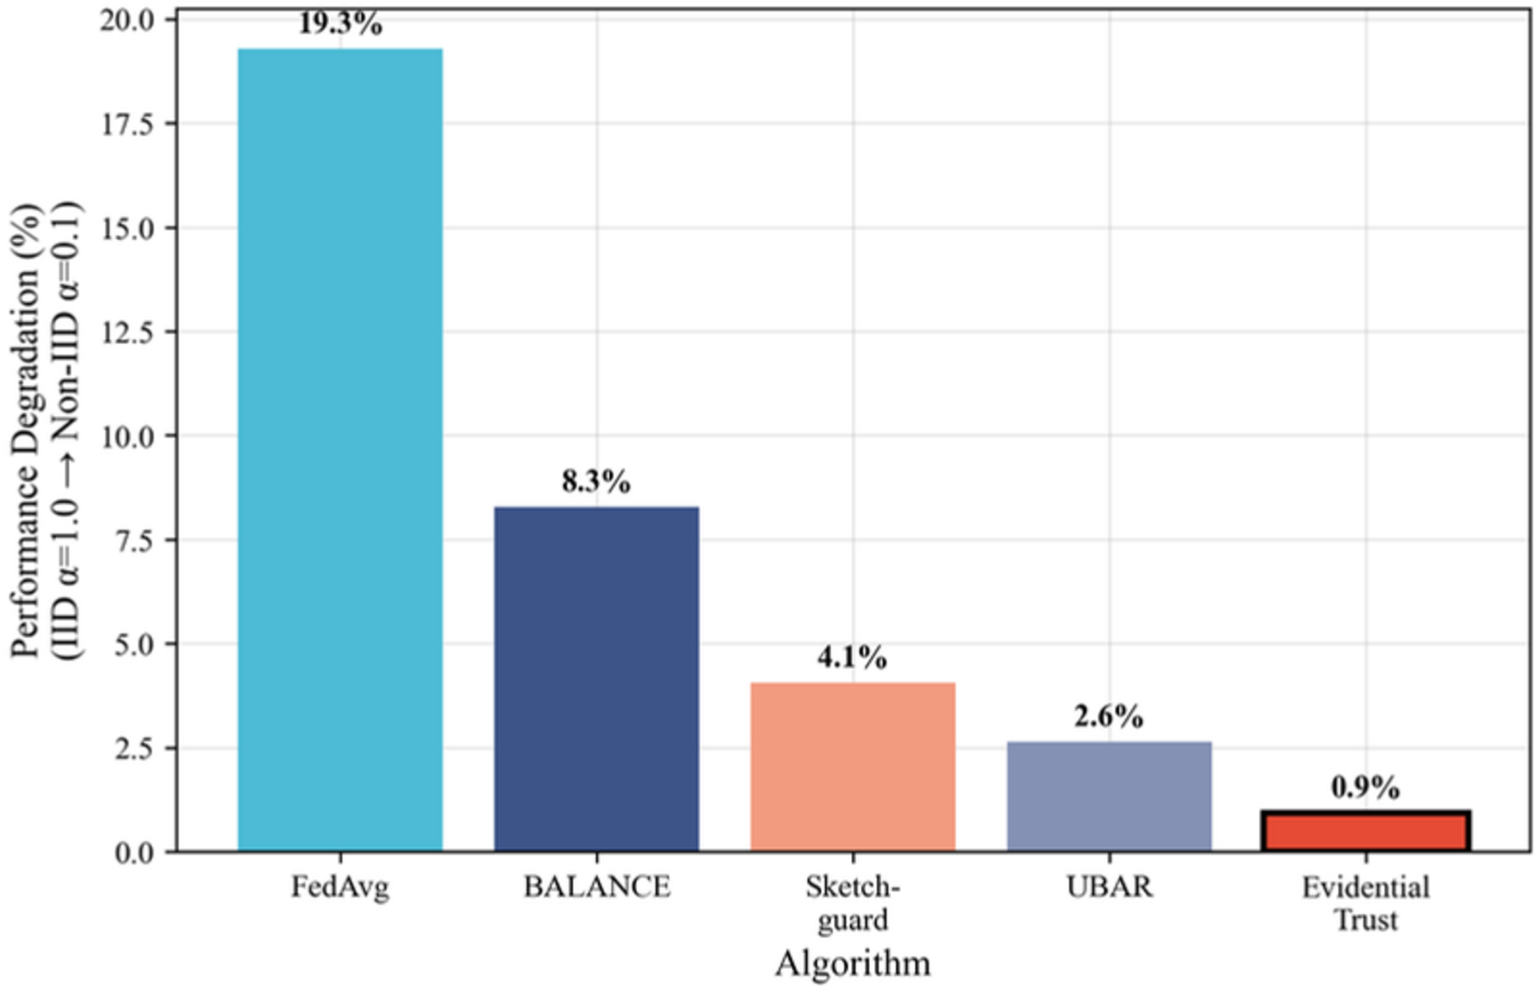

Performance degradation from IID (α = 1.0) to non-IID (α = 0.1) conditions.

Lower values indicate better robustness to heterogeneity.

MURMURA shows minimal degradation (0.9%) compared to baselines.

16. Convergence Speed

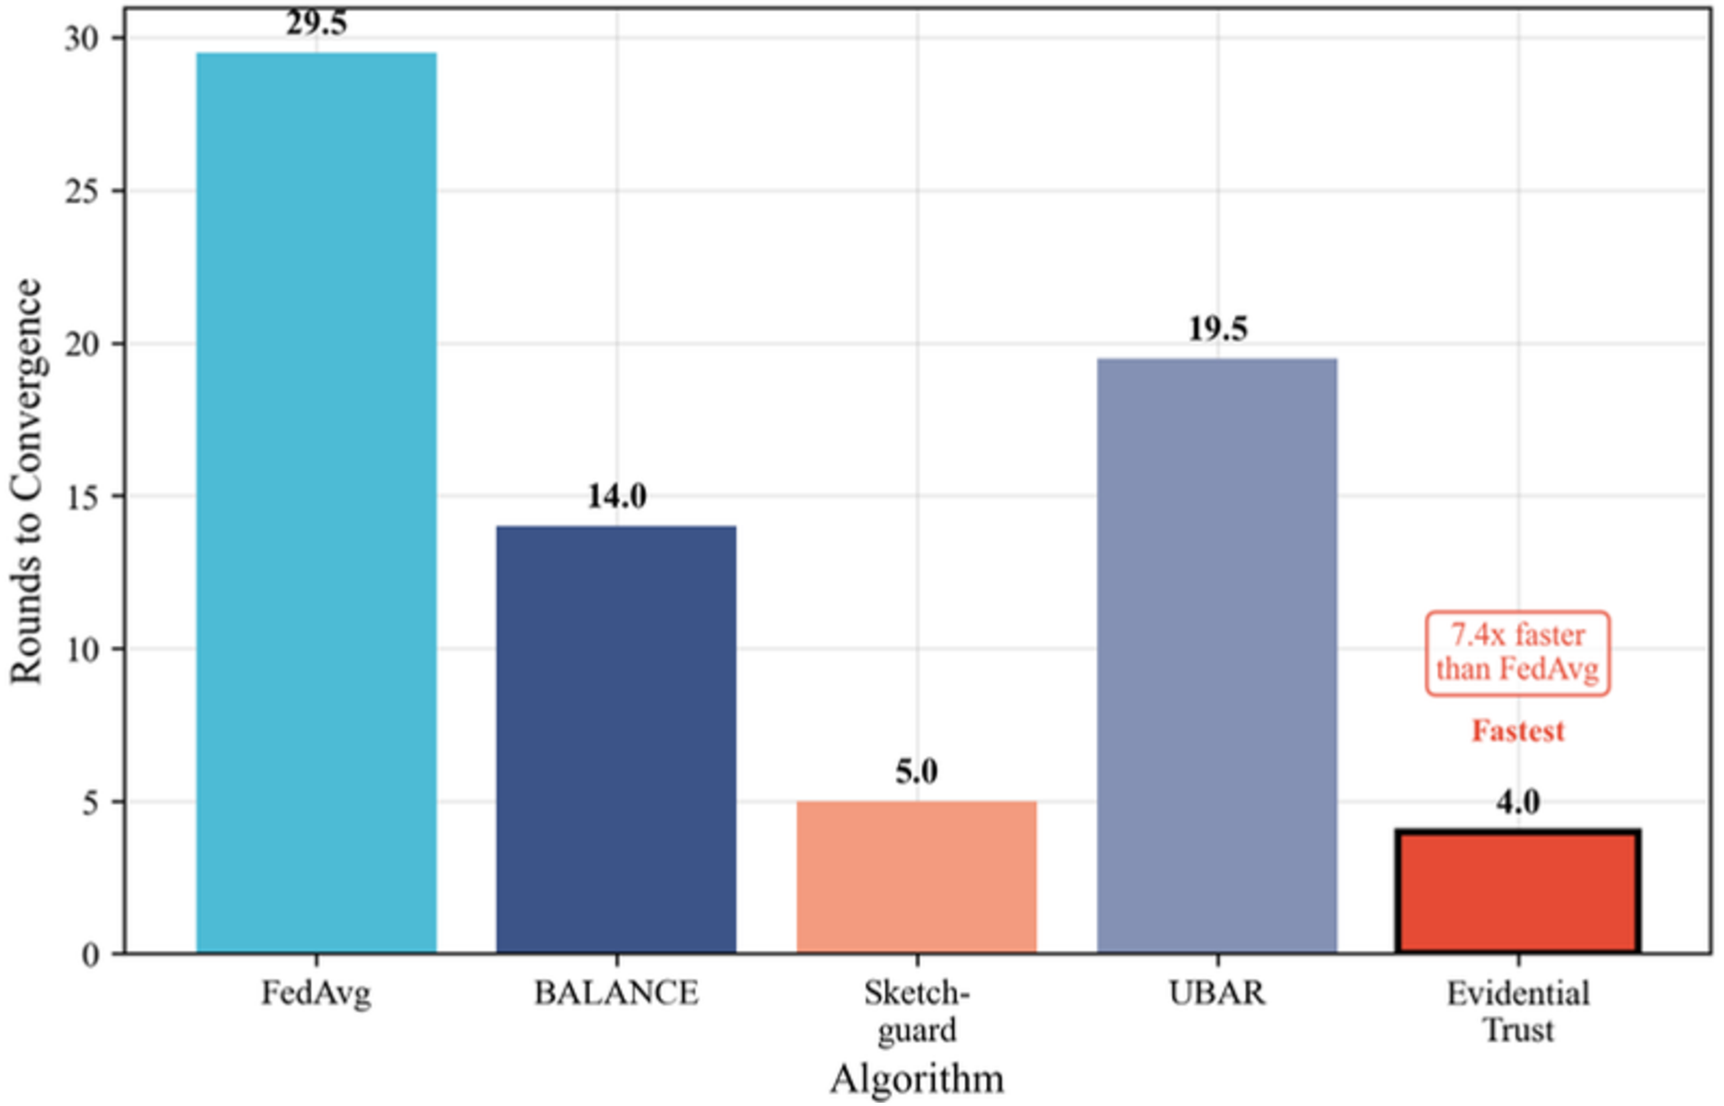

Convergence speed comparison showing rounds to reach peak accuracy.

MURMURA converges 7.4× faster than FedAvg by filtering incompatible peer updates.

Q&A

Paper Preprint

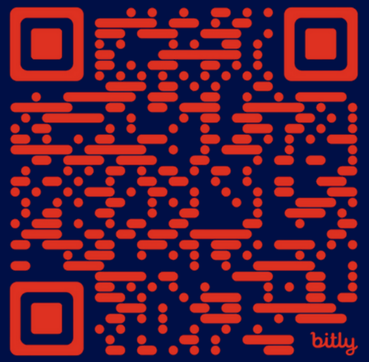

LinkedIn

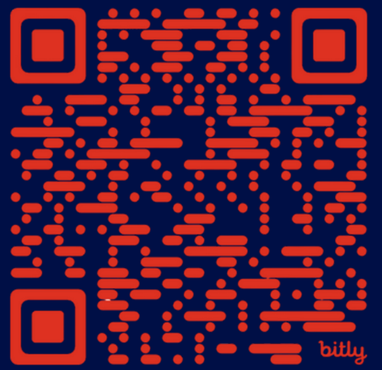

THE UNIVERSITY OF

MELBOURNE

Presentation Link