Knowledge & Data Engineering Systems

MOBILE: Mobility and Outage-Based Intelligent Federated Learning in Mobile Computing[1]

Qiyuan Wang, Qianyu Long and Christos Anagnostopoulos

Oct 30, 2025

[1] Supported by the EU Horizon Europe program, grant **101178099 (ELLIE)**.

The Promise and Reality of FL

Setting the Scene

Knowledge & Data Engineering Systems

The Promise:

- Federated Learning enables privacy-preserving distributed training
- Clients train locally, share only model updates
- Perfect for mobile devices with sensitive data

The Reality in Mobile Computing:

- Clients are constantly moving
- Network conditions fluctuate dramatically
- Connection drops are the norm, not the exception

Result: Not all selected clients successfully complete training!

The Energy Constraint Fallacy

Why Current Approaches Miss the Mark

What existing methods assume:

- Energy is the primary bottleneck
- OCEAN framework: 0.15J energy budget
- MADCA-FL: 0.006-0.0012J for vehicular networks

Reality check:

- Modern phones when transmitting: 1.5-2.5 Watts
- 0.15J depleted in less than 0.1 seconds!
- Phone batteries hold ∼66,600J
- GPU TDP can reach 450W

The real problem is mobility!

Introducing MOBILE

Mobility and Outage-Based Intelligent Federated Learning

Knowledge & Data Engineering Systems

Core Innovation:

- First framework to prioritize mobility patterns over energy
- Uses historical movement data to predict client reliability
- Formulated as Mixed-Integer Quadratic Programming

Key Properties:

- Orthogonal to existing FL algorithms
- No modifications to core FL operations needed
- Minimal computational overhead

Problem Formulation

From Intuition to Mathematics

The Challenge:

- At round t, select clients $\mathcal{N}_t' \subset \mathcal{N}$
- Allocate bandwidth fractions b_k^t to each client
- Maximize successful completions within deadline $\mathcal{T}$

Key Decisions:

- Binary selection: $a_k^t \in \{0, 1\}$ for each client k
- Bandwidth allocation: $b_k^t \in [b_{min}, 1]$ if selected
- Constraint: $\sum_{k \in \mathcal{N}} b_k^t = 1$

Return Vector: Capturing Proximity

Component 1 of MOBILE

Return Vector Definition:

$$r_k^t = 1 - \frac{d_{k,t}}{d_{\max,t}} \tag{1}$$

Intuition:

- $d_{k,t}$: Distance of client k from Base Station at time t
- $d_{\max,t}$: Maximum distance among all clients
- Higher return $\rightarrow$ Closer to BS $\rightarrow$ More reliable

Why this matters:

- Signal strength decreases with distance
- Clients near BS more likely to complete uploads
- Forms basis for optimization

Covariance Matrix: Learning from History

Knowledge & Data Engineering Systems

Component 2 of MOBILE

Covariance Matrix Construction:

$$Q_t = \frac{1}{W} \sum_{\tau=t-W+1}^{t} (\mathbf{r}_\tau - \bar{\mathbf{r}}_t)(\mathbf{r}_\tau - \bar{\mathbf{r}}_t)^\top \qquad (2)$$

What it captures:

- Spatio-temporal correlations between clients
- Movement patterns over sliding window W
- Clients that "move together" have high covariance

Strategic insight:

- Minimize portfolio risk through diversification
- Reduces collective failure probability

MIQP Optimization

Knowledge & Data Engineering Systems

Bringing It All Together

The MOBILE Objective:

$$\max_{\mathbf{a},\mathbf{b}} \sum_{t=0}^{T-1} \underbrace{\mathbf{r}_t^\top \mathbf{b}_t}_{\text{Expected return}} - \underbrace{\lambda \mathbf{b}_t^\top Q_t \mathbf{b}_t}_{\text{Risk penalty}} + \underbrace{\gamma \|\mathbf{a}_t\|_1}_{\text{Participation bonus}} \tag{3}$$

Interpretation:

- First term: Maximize expected successful completions
- Second term: Minimize risk from correlated failures
- Third term: Encourage sufficient participation

Constraints ensure:

- Binary client selection
- Valid bandwidth allocation

Cold-Start Strategy

Handling Limited Historical Data

The Bootstrap Problem:

- MOBILE needs historical data to build Q_t
- But what about initial rounds?

Adaptive Three-Phase Approach:

- **Phase 1** ($t < W^*$): Use distance-based baseline
- **Phase 2** ($W^* \leq t < W$): Partial MOBILE with available data
- **Phase 3** ($t \geq W$): Full MOBILE optimization

Experimental Setup

Datasets and Configuration

Mobility Dataset:

- YJMob100K from Japan
- 100,000 individuals
- 75 days of tracking
- 200×200 grid cells
- 30-minute intervals

Base Station:

- Position: (137, 78)
- Coverage: 10km radius
- Top-100 users selected

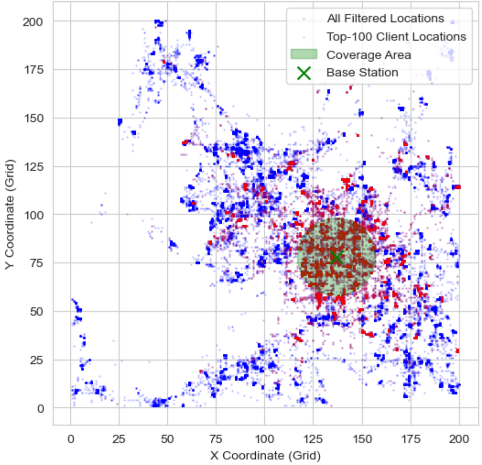

FL Datasets and Models

Diverse Evaluation Scenarios

HAM10000 - Medical Imaging:

- 7 classes of skin lesions
- Critical for healthcare applications
- Model: AlexNet

CIFAR10 - Natural Images:

- 10 object classes
- Standard FL benchmark
- Model: ResNet18

CIFAR100 - Fine-grained Classification:

- 100 object classes
- Tests scalability to many classes
- Model: VGG11

Main Results: Accuracy

Consistent Improvements Across Datasets

Method	CIFAR10	CIFAR100	HAM10000
FedAvg	56.65±0.69	26.35±0.31	59.13±1.22
Distance-based	57.76±0.30	26.75±0.13	61.08±0.40
MOBILE	**58.97±0.11**	**27.57±0.08**	**62.11±0.57**
MOBILE-CS	58.77±0.15	27.54±0.10	62.11±0.58

Key Observations:

- MOBILE achieves highest accuracy on all datasets
- Lower variance (more stable training)
- Cold-start version performs nearly as well

Learning Curve

Training Stability Over 100 Rounds

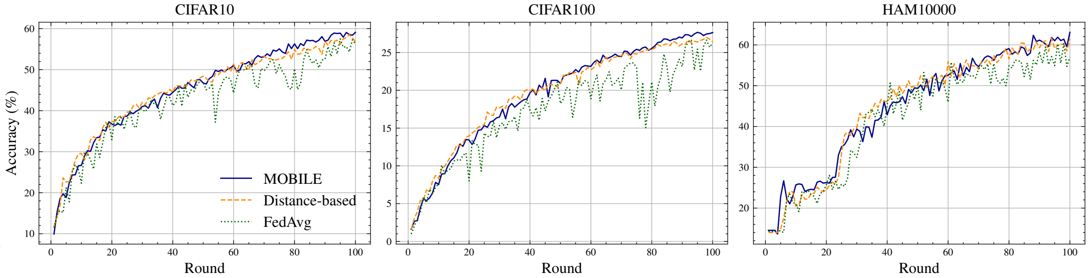

Observations:

- FedAvg shows high variance (unstable updates)
- MOBILE converges smoothly and consistently
- Distance-based shows intermediate behavior
- Stability comes from reliable client participation

Success Rates and Resource Efficiency

Knowledge & Data
Engineering Systems

The Core Achievement

Client Success Rates:

- MOBILE: 89.42%
- Distance-based: 84.50%
- FedAvg: 32.30%

Wasted Throughput:

- MOBILE: 30.65%
- Distance-based: 41.47%
- FedAvg: 73.52%

End-of-round distance:

- MOBILE: 8.53 grid units
- FedAvg: 26.02 grid units

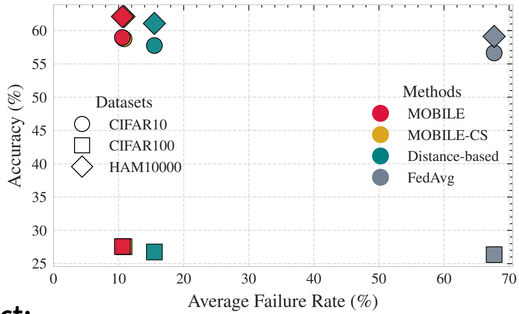

Impact:

- 177% improvement in success rate
- 58% reduction in wasted bandwidth
- Preserves client incentives
- Reduces failed contributions

Non-i.i.d Data Performance

Robustness to Data Heterogeneity

Knowledge & Data Engineering Systems

Setup: Dirichlet distribution with $\alpha = 0.5$

Method	CIFAR10	CIFAR100	HAM10000
FedAvg	44.31±0.34	22.25±0.24	58.60±0.58
Distance-based	47.19±0.32	22.92±0.16	60.54±0.37
MOBILE	**48.35±0.20**	**23.18±0.17**	**60.69±0.32**

Additional Strategy:

Top-ℓ selection by loss ranking

Future Directions

Where We Go From Here

Immediate Extensions:

- Incorporate computational heterogeneity
- Multi-tier edge computing architectures
- Adaptive parameter tuning

Longer-term Research:

- Predictive mobility modeling with deep learning
- Game-theoretic incentive mechanisms
- Privacy-preserving mobility pattern sharing

Knowledge & Data Engineering Systems

Thank you!

Questions?

Supported by the EU Horizon Europe program, grant 101178099 (ELLIE).