University of Glasgow

Knowledge & Data Engineering Systems

RADEL: Resilient and Adaptive Distributed Edge Learning in Dynamic Environments

Qiyuan Wang

November 2025

The Evolution of Edge Intelligence

Knowledge & Data Engineering Systems

Context and Challenges

The Promise of Edge Computing:

- Real-time analytics at the network periphery
- Decentralized machine learning where data originates
- Responsive services across distributed environments

Five Critical Challenges in Edge Learning:

- Node failures compromise service continuity
- Concept drift degrades model performance over time
- Irrelevant data reduces prediction accuracy
- Communication inefficiencies limit scalability
- Client mobility undermines learning assumptions

We need resilient and adaptive solutions!

The RADEL Framework

Five Interconnected Contributions

Core Philosophy:

- Resilience: Maintain service during failures
- Adaptability: Respond to changing conditions
- Efficiency: Minimize resource consumption

Contributions 1-2:

- Resilient predictive analytics
- Maintenance under drift

Contributions 3-5:

- Query-driven learning
- Dynamic personalized FL
- Mobility-aware FL

A comprehensive framework for practical edge deployment

Paper 1: The Resilience Challenge

Maintaining Service During Node Failures

Problem Statement:

- Edge nodes fail unpredictably (hardware, network, power)
- Replication-based methods are resource-intensive
- Need: Lightweight resilience mechanism

Core Innovation - Enhanced Local Models:

- Surrogate nodes serve requests of failing nodes
- Leverage statistical signatures of neighboring data
- Build enhanced models through information extraction

Key Strategies:

- **Information Adjacency**: Leverage spatial/statistical proximity

Paper 1: Results and Impact

Significant Performance Improvements

Experimental Setup:

- Multiple real-world datasets
- Various failure scenarios
- Comparison with centralized method

Key Results:

- Substantial accuracy improvements
- Minimal data transfer required
- Scalable to large networks

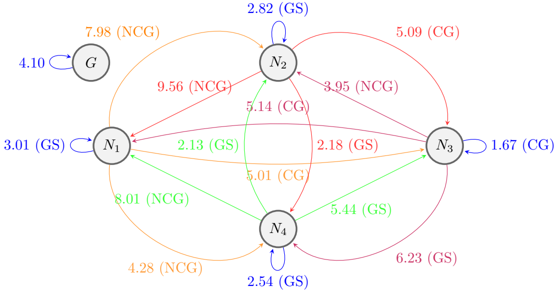

Contribution:

- Novel resilience without full replication
- Foundation for reliable edge services

Paper 2: The Concept Drift Challenge

Knowledge & Data Engineering Systems

VIA VERITAS VITA

Preserving Resilience Under Evolving Data

The Problem:

- Enhanced models degrade as data distributions evolve
- Concept drift is inevitable in dynamic edge environments
- Need: Efficient maintenance strategies

Analyzing Drift Impact:

- Studied effects of different drift types on enhanced models
- Quantified performance degradation patterns

Maintenance Strategies:

- **Mock Data (MD)**: Synthetic data generation approach
- **Enhanced Centroid Guided (ECG)**: Efficient targeted updates

Paper 2: Optimal Maintenance Trade-offs

Knowledge & Data Engineering Systems

Balancing Communication and Adaptation

Evaluation:

- Synthetic datasets with controlled drift
- Real-world datasets with natural evolution
- Heterogeneous data sources

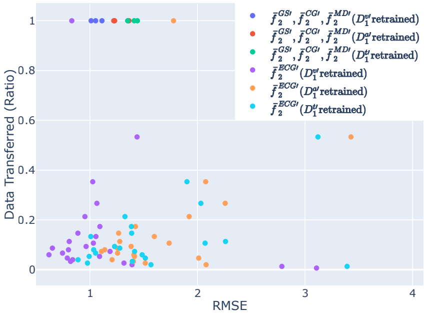

Key Achievements:

- Optimal trade-offs between transmission and effectiveness
- Maintained performance across heterogeneous sources
- Minimized overhead during adaptation

Contribution:

- Systematic approach to resilience maintenance
- Practical deployment guidelines

Paper 3: The Data Relevance Problem

Not All Data is Equally Useful

Key Insight:

- Traditional approaches: Use all available data
- Reality: Query patterns reveal relevant data regions
- Irrelevant data can actually degrade model performance

Challenge:

- Identify relevant data from queries
- Determine when to stop refinement
- Adapt to changing patterns

Solution:

- Query-driven identification
- Optimal stopping theory
- Adaptive update mechanisms

Paper 3: Framework Components

Data-Centric Learning Mechanism

Core Innovation: Query-Driven Identification

- Queries reveal which data regions are actually relevant
- Progressive discovery: Start with initial model f_0 on all data $\mathcal{D}$
- Each query identifies relevant samples → Build $\mathcal{D}_t$ incrementally
- Train refined models $\{f_t\}$ only on discovered relevant regions

Supporting Components:

1. Model Selection:

$$f_{t+1} = \begin{cases} f_t & \text{if } \epsilon_t < \gamma_t \epsilon_{t-1} \\ f_t^* & \text{otherwise} \end{cases}$$

- Explore vs. exploit: γ_t decays with t

2. Optimal Stopping:

- When to stop refinement?
- When to update for drift?
- Theory provides optimal T^* and k^*

3. Data Relevance Discriminator (SVM):

- Binary classifier $f_C(x) \in \{+1, -1\}$ separates relevant vs. irrelevant regions
- Trained on $\mathcal{D}_R$ (relevant) and $\mathcal{D}_I$ (irrelevant) from learning process

Paper 3: Dramatic Performance Gains

Learning from What Matters

Experimental Setup:

- Diverse datasets and query workloads
- Compared against full-data baselines

Impressive Results:

- Up to 63% improvement in accuracy
- Consistently outperforms traditional approaches
- Lower computational cost
- Robust under distribution drift

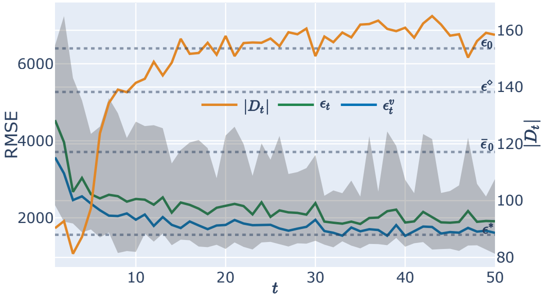

Contribution:

- Paradigm shift: Quality over quantity
- Practical query-aware learning

Paper 4: The Personalization Challenge

Knowledge & Data Engineering Systems

VIA VERITAS VITA

DA-DPFL Framework Overview

Challenges in Personalized Federated Learning:

- Data heterogeneity across clients (non-i.i.d.)
- High communication costs in model exchange
- Computational overhead at resource-constrained edge devices

DA-DPFL Innovation:

- Dynamic aggregation in decentralized topology
- Knowledge reuse within communication rounds
- Sparse-to-sparser pruning based on compressibility

Key Insight:

- Knowledge reuse can be achieved in each round if tolerating time cost
- Progressive pruning greatly reduces communication and computation

Paper 4: Technical Approach

Dynamic Aggregation Mechanism

System Model:

- Decentralized topology: No central server
- Peer-to-peer model exchange among neighbors
- Each client maintains personalized model

Sparse-to-Sparser Pruning:

- Pruning based on compressibility
- Progressive pruning: Start sparse, become sparser
- Preserve critical connections while reducing size

Benefits:

- Reduced training computation
- Lower communication overhead
- Personalized models for each client

Paper 4: Efficiency and Performance

Achieving Superior Trade-offs

Evaluation Setup:

- Non-i.i.d. data partitioning
- Multiple FL benchmarks
- Compared with SOTA DPFL methods

Outstanding Results:

- Superior test accuracy
- Up to 5× energy reduction
- Significant communication savings
- Effective personalization

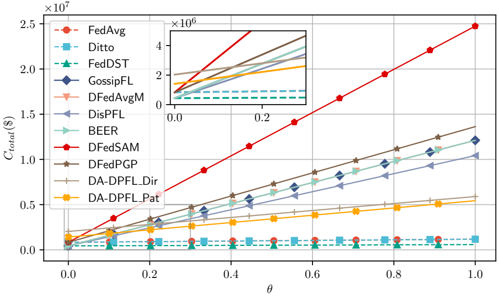

Contribution:

- Novel dynamic aggregation
- Practical for edge devices
- Scalable decentralized learning

Paper 5: The Mobility Challenge in FL

Knowledge & Data Engineering Systems

VIA VERITAS VITA

MOBILE - Mobility and Outage-Based Intelligent Learning

Reality of Mobile Federated Learning:

- Clients are constantly moving
- Network conditions fluctuate dramatically
- Connection drops are common, not exceptional

The Energy Constraint Fallacy:

- Existing work assumes: Energy is primary bottleneck
- Reality: Modern phones when transmitting: 1.5-2.5 Watts
- Phone batteries: $\sim$66,600J, GPU TDP: up to 450W
- The real problem is mobility, not energy!

MOBILE's Paradigm Shift:

- First framework prioritizing mobility patterns over energy
- Uses historical movement data for reliability prediction
- Joint optimization of client selection and bandwidth allocation

Paper 5: Mathematical Formulation

Portfolio Optimization Approach

Knowledge & Data Engineering Systems

Covariance Matrix - Learning from History:

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

- Captures spatio-temporal correlations between clients
- Clients that "move together" have high covariance

Return Vector - Capturing Proximity:

$$r_k^t = 1 - \frac{d_{k,t}}{d_{\max,t}} \qquad (1)$$

- Higher return $\rightarrow$ Closer to BS $\rightarrow$ More reliable

MIQP Optimization:

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

Paper 5: Transformative Results

Knowledge & Data Engineering Systems

VIA VERITAS VITA

Making FL Practical for Mobile Deployments

Experimental Setup:

- YJMob100K dataset
- 100K users, 75 days
- CIFAR10, CIFAR100, HAM10000

Game-Changing Performance:

- Success rate: 32% → 89%
- Wasted bandwidth: 73% → 30%
- Higher accuracy, stable convergence

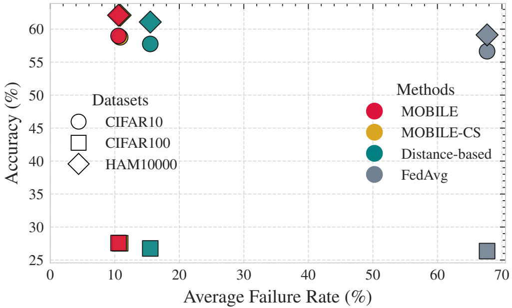

Key Features:

- Orthogonal to existing FL algorithms
- Cold-start strategy for bootstrap
- Minimal computational overhead

Contribution: Paradigm shift from energy to mobility; practical mobile FL framework

Conclusions and Future Directions

RADEL: Resilient and Adaptive Distributed Edge Learning

Knowledge & Data Engineering Systems

Key Contributions:

- **Resilience**: Service continuity during node failures
- **Adaptability**: Maintenance under concept drift
- **Query-driven**: 63% accuracy improvement
- **DA-DPFL**: 5× energy reduction
- **MOBILE**: 177% success rate improvement

Future Directions:

- *Short-term:*
  - Component integration
  - Multi-tier architectures
  - Adaptive parameter tuning
- *Long-term:*
  - Predictive mobility modeling
  - Privacy-preserving techniques
  - Federated continual learning
  - Game-theoretic incentives

RADEL: A robust foundation for reliable, adaptive machine learning at the edge!

Thank you!

Questions?