Summer School & Seminar 2026

Towards Trustworthy and Efficient One-Shot Federated Learning

Phd Student

Lingyu Qiu

P1: NeuroComputing | P2: IEEE BigData'25 | P3: ACM MobiCom'25 | P4: FLICS'26 | **P5: Computer Network**

Supervisor & Host Institution

Prof. Francesco Piccialli & UNINA

Aligning with the TUAI Programme

1 Federated AI Task

Focused on Task T4.1 – developing trustworthy, highly scalable decentralised machine-learning systems that overcome collaboration bottlenecks under constrained bandwidth and severe Non-IID heterogeneity across participants.

2 The One-Shot Federated Learning (OFL) Paradigm

Compresses hundreds of iterative synchronisation rounds into **a single communication round** – fundamentally bypassing frequent edge-gateway interaction and bandwidth limits, while providing an inherent barrier against gradient leakage.

3 Career Development Plan (CDP) Progress

Closely follows the technical roadmap of my CDP – progressing from theoretical foundations, through general computer-vision representation diagnostics, to a concrete smart-transportation spatiotemporal deployment, forming a closed research chain.

THE LOGICAL RESEARCH COHESION

Research Trajectory: Deconstructing the Pain Points of One–Shot FL

01 / FOUNDATIONS

What has prior work actually done?

A PRISMA-guided systematic **review** of 94 landmark OFL studies, proposing a three-axis taxonomy. Identifies pretrained models and decentralised generative synthesis as the two key paths to resolving manifold incompleteness.

Neurocomputing Journal

02 / ACTIVE SAMPLE GENERATION

No data — can we generate it?

Proposes **FedSDE** — a local conditional diffusion model generates boundary-challenging samples, combined with EMA self-distillation and soft-label upload, avoiding parameter transmission and resolving extreme Non-IID heterogeneity.

IEEE BigData'25

03 / MECHANISM DIAGNOSIS

What role do pretrained models play?

A deep **diagnostic study** finds frozen pretrained backbones dominate performance, but reveals that classical Gaussian aggregation collapses geometrically on Transformer feature spaces

FLEdge–AI MobiCom'25

04 / SOURCE-SIDE RECTIFICATION

How to fix the geometric conflict painlessly?

Proposes **FedOPAL** — without restructuring the server-side algorithm, 10 visual prompt tokens rectify the feature manifold at the source so closed-form (AFL) aggregation aligns perfectly, eliminating server-side iteration.

FLICS'26

05 / SPATIOTEMPORAL DEPLOYMENT

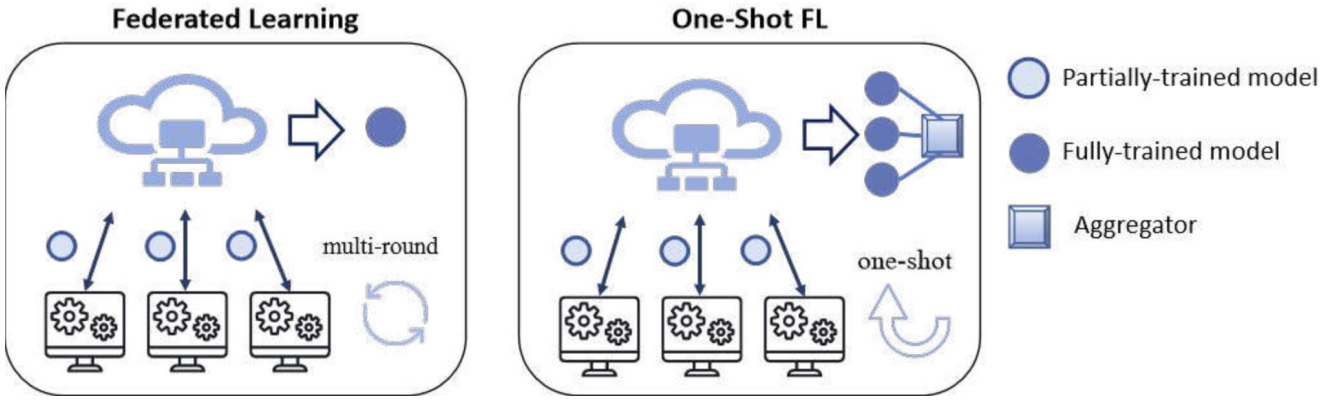
Does it hold up in spatiotemporal IoT?

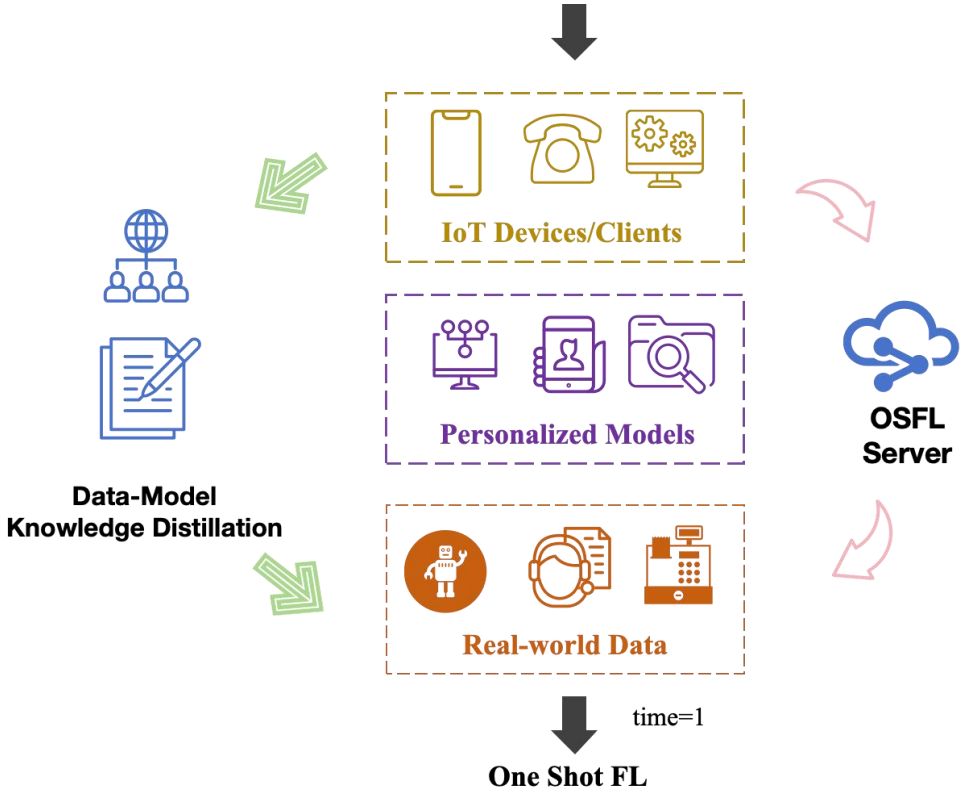
The latest contribution, **SWIFT.** Road-network topology is non-shareable and asymmetric across cities. Physical channel-wise decomposition isolates topological noise, synchronising only universal temporal dynamics for one-shot convergence.

Computer Network Journal

These five contributions build on one another： from theoretical foundations, through algorithmic redesign and mechanism diagnosis, to manifold rectification — closing the loop in a real physical system (spatiotemporal traffic IoT).

Research Question: What is One-Shot FL?





$$\min_{\theta_S} \mathcal{L}(\theta_S) \triangleq \frac{1}{|\mathcal{D}|} \sum_{k=1}^{K} \sum_{(x_i, y_i) \in \mathcal{D}_k} \ell(f_S(x_i; \theta_S), y_i)$$

Milestone 1 — Foundations: Survey of One–Shot FL

《Towards One–shot Federated Learning: Advances, Challenges, and Future Directions》

A Systematic Review: Deconstructing 94 Landmark OFL Studies

Published in Neurocomputing (2026). Following the **PRISMA** protocol, we searched and screened core One-Shot FL results from top AI venues (NeurIPS, ICLR, AAAI, ICML, etc.), organising them into a clear three-axis taxonomy.

Core Pain Point Identified

Conventional iterative FL requires **hundreds of communication** rounds. In one-shot FL, forcing server-side aggregation over highly Non-IID, fragmentary client data instead triggers severe **“representation collapse and forgetting.”**

Highest-Potential Directions

The survey concludes that **pretrained foundation** encoders and decentralised generative synthesis are the two most promising tracks for breaking through OFL's performance ceiling.

OFL Classification Dimensions

1. Client Initialization
Foundation Priors vs Random

2. Aggregation Method
Ensemble Distillation vs Closed–form

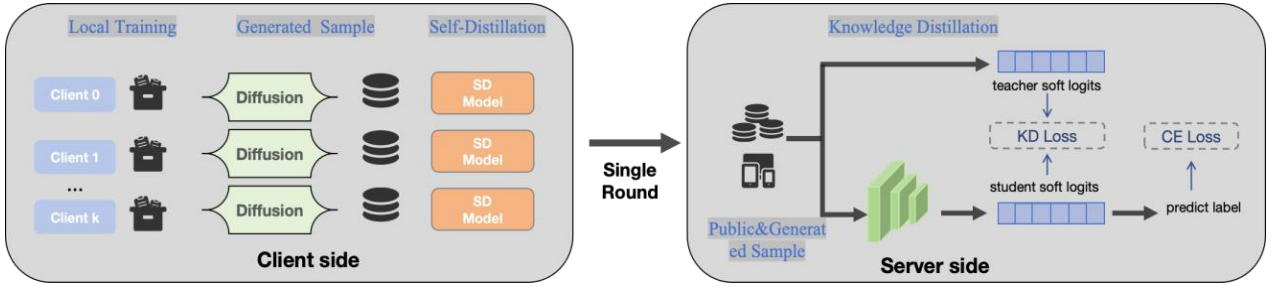
3. Heterogeneity Handle
Feature Align vs Generative AI.

The starting point of the research loop: this taxonomy **directly defines the methodological boundaries and focal points of the four exploration phases** that follow.

Milestone 2 — FedSDE: One-Shot FL via Local Conditional Diffusion Generation

《Deciphering One-Shot Federated Learning: The Pivotal Role of Pre-trained Models》

When a client holds only extremely skewed local data (α=0.01), uploading and merging local models directly triggers severe generalisation collapse on the server. Could the client instead generate data locally, completing its own manifold before transmitting knowledge?



Module 1: Classifier-Guided Hard Boundary Sample Generation

The client trains a conditional diffusion model. At inference, the local classifier forcing the diffusion model to generate decision-boundary-adjacent **"hard samples"**:

Module 2: Dual-Side EMA Self-Distillation

The local model performs self-distillation on the hybrid dataset; the server pools each client's synthetic **boundary data** for a single round of ensemble knowledge distillation.

FedSDE Solid Benchmark Gain

+16.57%

absolute accuracy gain **over ensemble baselines** (CIFAR-100, α=0.05)

52.34%
Full FedSDE (CIFAR-10, Dir 0.6)

+9.89
pts over standard diffusion alone

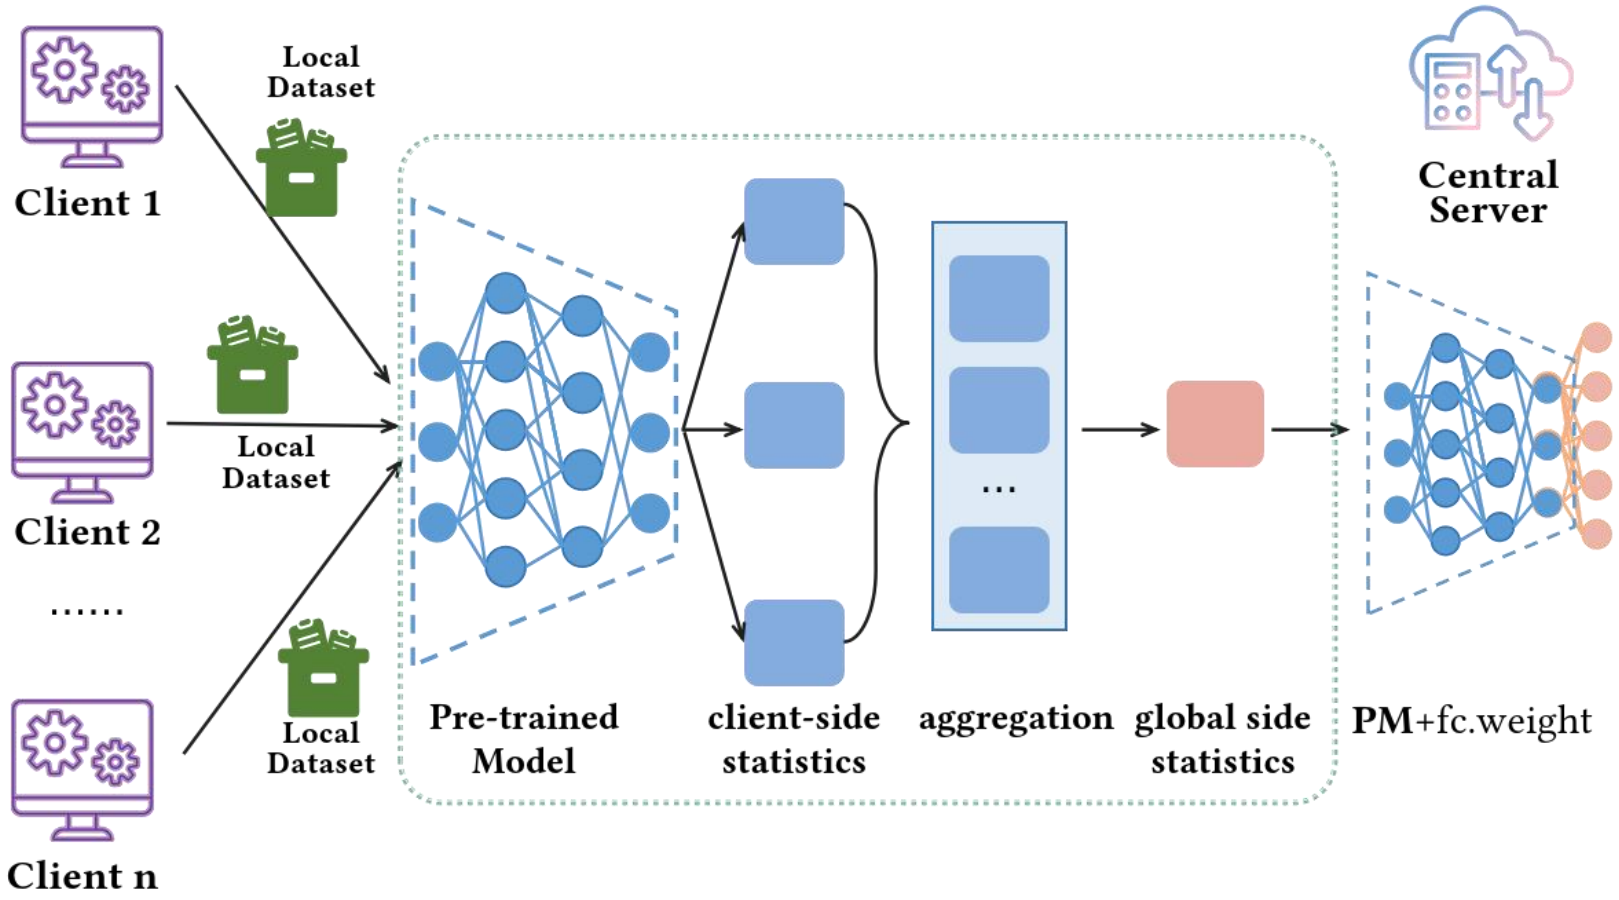
Client-Side: Self-Distillation

Server-Side: Knowledge Distillation

Published at IEEE BigData (2025); the full generation pipeline is open-sourced.

Milestone 3 — Pretrained Backbones and Geometric Collapse

《Deciphering One-Shot Federated Learning: The Pivotal Role of Pre-trained Models》

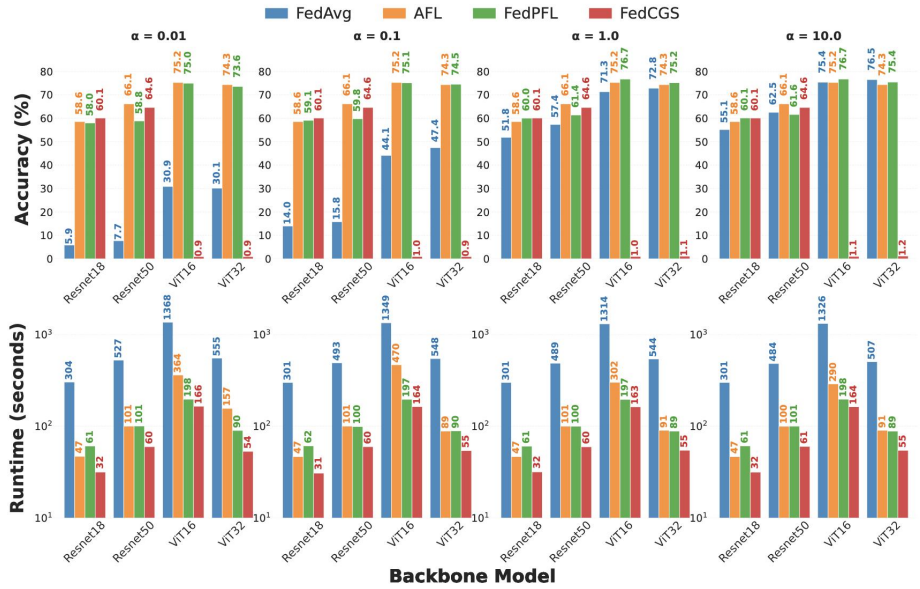


Milestone 3 — Pretrained Backbones and Geometric Collapse

《Deciphering One-Shot Federated Learning: The Pivotal Role of Pre-trained Models》

Although **local diffusion synthesis** produces high-quality samples, its local training cost reaches **275.13 seconds**

➡ Could we design an extremely lightweight OFL method requiring no iterative fine-tuning on either the server or client at all?



Building on the pretrained foundation models we designed a rigorous controlled study (ACM MobiCom 2025) to quantitatively answer: in one-shot FL, does the strength of the local backbone or the complexity of the server-side aggregation algorithm dominate performance?

Controlled Benchmarks on CIFAR-10 (α = 0.1)

The data show clearly: a frozen ViT-B/16 paired with the simple closed-form AFL aggregation sustains 93.95% robust accuracy — backbone capability dominates OFL performance entirely.

93.95%

AFL + ViT-B/16
(Frozen backbone)

10.61%

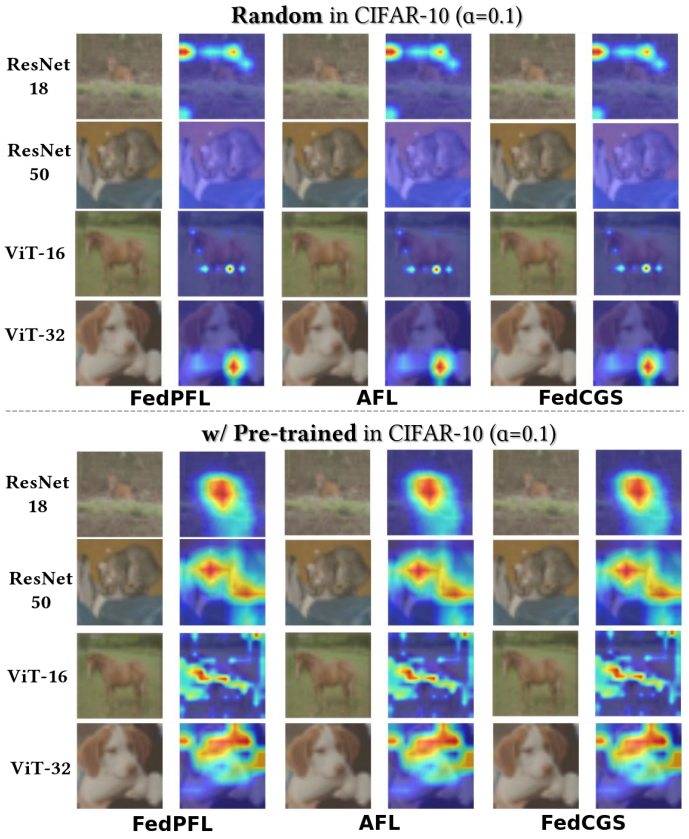
FedCGS + ViT-B/16
(Geometric collapse)

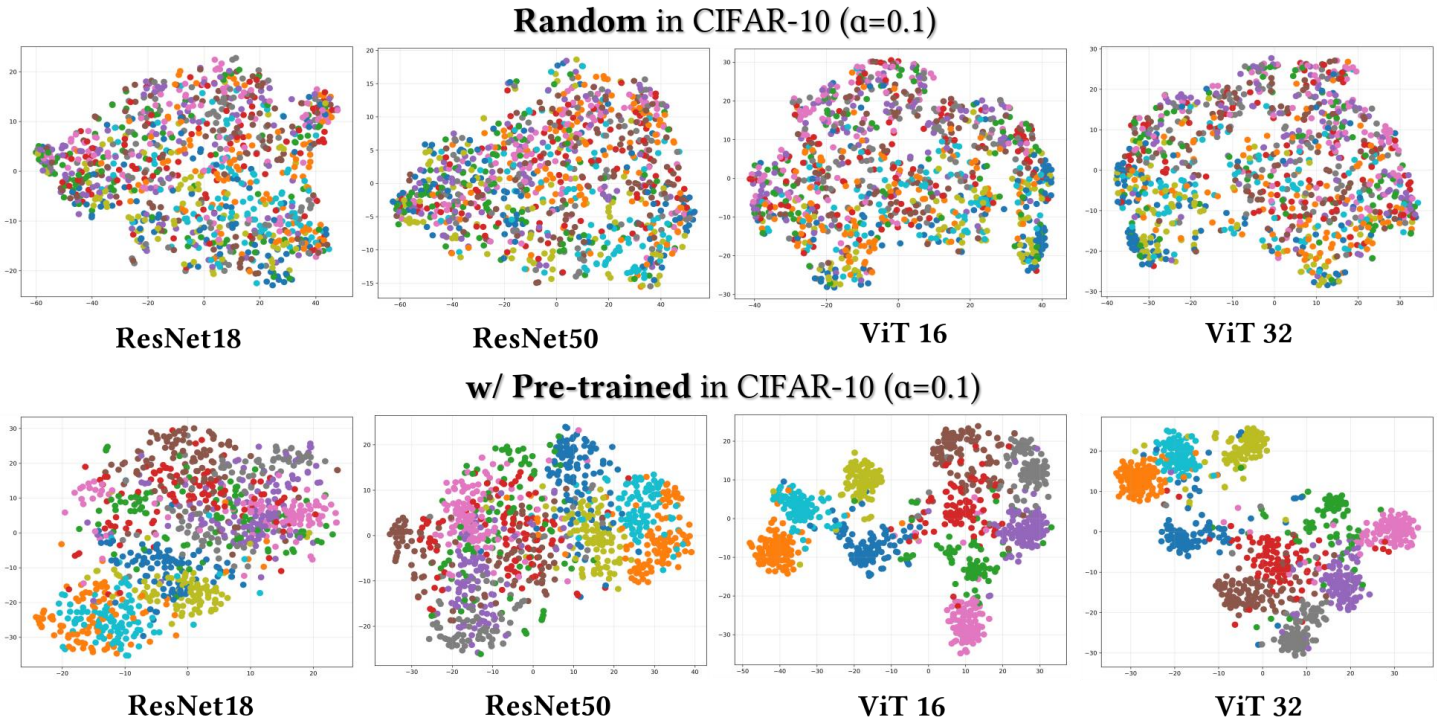
OFL Method	Backbone Used	Accuracy
FedCGS	ResNet-18	81.94%

** Published at FLEdge-AI, ACM MobiCom'25. This diagnosis directly motivated the next-stage manifold rectification work (FedOPAL).*

Milestone 3 — Pretrained Backbones and Geometric Collapse

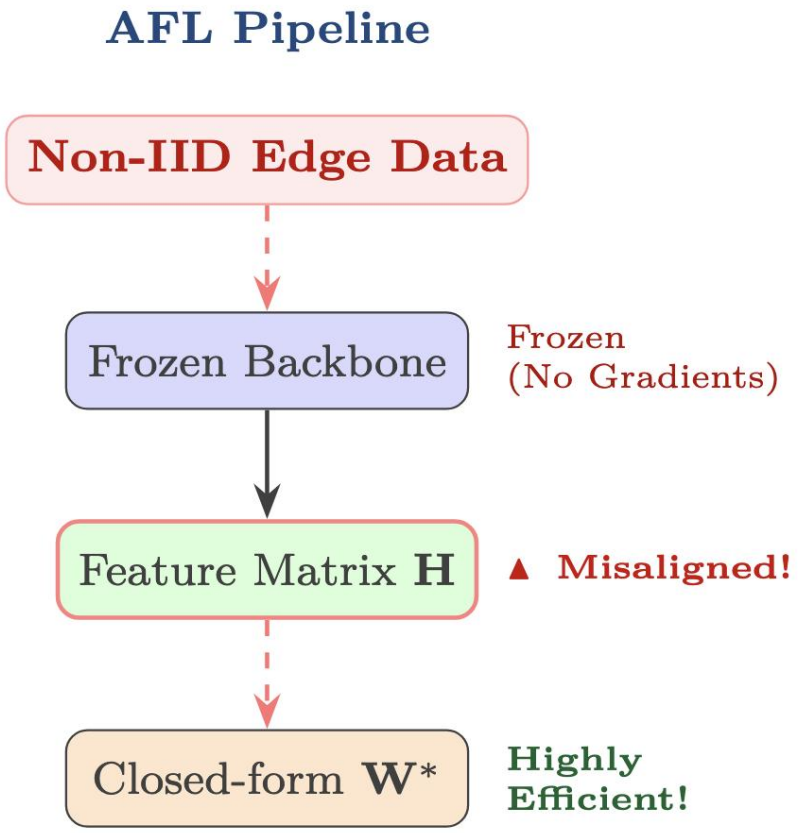
《Deciphering One-Shot Federated Learning: The Pivotal Role of Pre-trained Models》





Milestone 4 — Gap in Existing Analytical Methods

《FedOPAL: One–Shot Federated Learning via Analytic Visual Prompt Tuning》



The Critical Flaw

- Features from heterogeneous clients are **misaligned**.
- Frozen backbone cannot adapt to local shifts.
- The analytic solver operates on *distorted* statistics.

  ⇒ **Performance collapses** under extreme Non-IID.

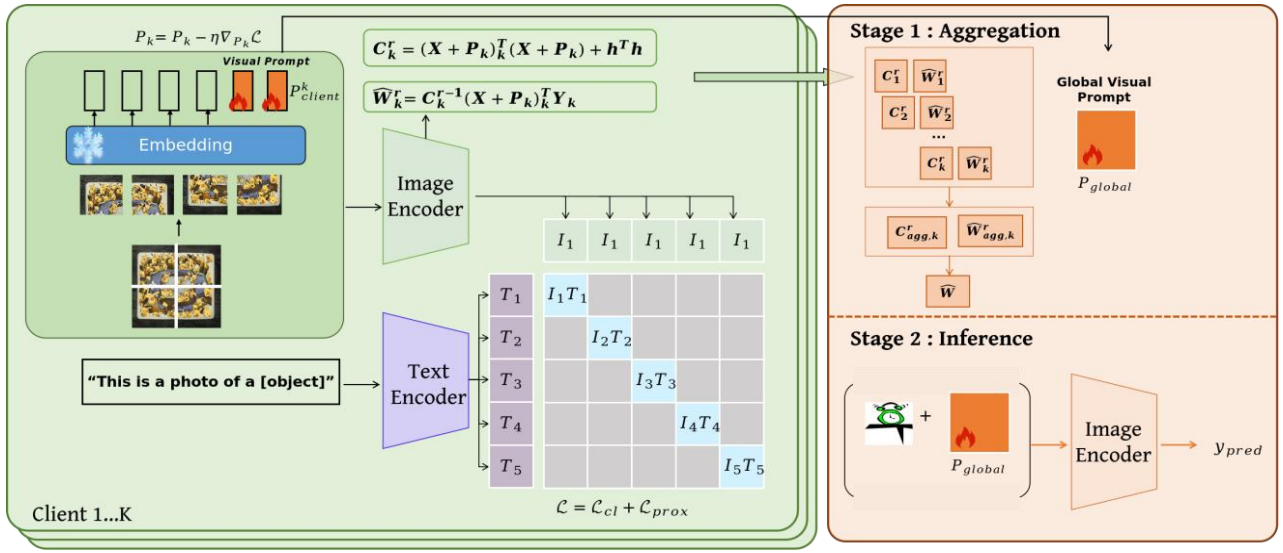
"A purely mathematical solver cannot adjust the feature extraction process itself."

Milestone 4 — FedOPAL: Manifold Rectification via Visual Prompt Tuning

《FedOPAL: One–Shot Federated Learning via Analytic Visual Prompt Tuning》

Since classical statistical FL aggregation is **incompatible** with the attention manifold of a frozen backbone, and redesigning the server-side aggregation law is costly and privacy-sensitive

---> could we **instead** rectify the distorted features directly at the source, using a **lightweight tunable prompt**?



Global Objective

$$\mathbf{W}^* = \arg\min_{\mathbf{W}} \sum_{k=1}^{K} \|\mathbf{H}_k\mathbf{W} - \mathbf{Y}_k\|_F^2 + \gamma\|\mathbf{W}\|_F^2 \quad (1)$$

Closed-Form Solution (Sufficient Statistics)

$$\mathbf{W}^* = \underbrace{\left(\sum_{k=1}^{K}\mathbf{R}_k + \gamma\mathbf{I}\right)^{-1}}_{\text{aggregated autocorr.}} \underbrace{\left(\sum_{k=1}^{K}\mathbf{C}_k\right)}_{\text{aggregated cross-corr.}} \quad (2)$$

where $\mathbf{R}_k = \mathbf{H}_k^\top\mathbf{H}_k$, $\mathbf{C}_k = \mathbf{H}_k^\top\mathbf{Y}_k$

No gradient descent on the server — pure linear algebra!

FedOPAL Performance Matrix (Zero Server–Training)

93.72%

CIFAR-10, α=0.1 — new SOTA among PM-based OFL

75.51%	+62.7pp	0
CIFAR-100 α=0.1	vs. FedCGS on DTD	server-side training rounds

Milestone 4 — Visual Prompt Tuning in ViT

《FedOPAL: One–Shot Federated Learning via Analytic Visual Prompt Tuning》

Standard VPT Sequence

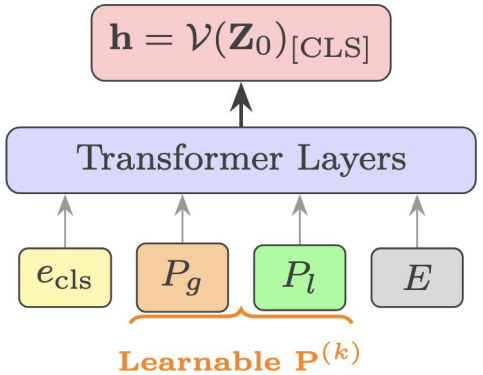
$$\mathbf{Z}_0 = [\mathbf{e}_{\text{cls}},\ \mathbf{P},\ \mathbf{E}]$$

**Baseline: Prompts are fully shared/centralised.*

FedOPAL Split-Prompt (Ours)

$$\mathbf{P}^{(k)} = \underbrace{\mathbf{P}_g^{(k)}}_{\text{Shared}} + \underbrace{\mathbf{P}_l^{(k)}}_{\text{Local}}$$

**Decouples global consensus from private context.*



Step 1 — Prompt Optimisation (Cross-Entropy + Proximal)

$$\min_{\mathbf{P}_g^{(k)}, \mathbf{P}_l^{(k)}} \mathcal{L}_{CE}\left(\text{Softmax}(\mathbf{W}_{head}^{\top}\mathbf{h}(\mathbf{x})),\ y\right) + \frac{\mu}{2}\left\|\mathbf{P}_g^{(k)} - \overline{\mathbf{P}}_g\right\|_F^2 \tag{3}$$

- μ: proximal coefficient — balances *local adaptation* vs *global alignment*.
- Backbone $\mathcal{V}$ is **frozen throughout** the optimization.

Step 2 — Sufficient Statistics Extraction

$$\mathbf{R}_k = \sum_{i=1}^{N_k} \mathbf{h}_i \mathbf{h}_i^{\top}, \qquad \mathbf{C}_k = \sum_{i=1}^{N_k} \mathbf{h}_i \mathbf{y}_{i,oh}^{\top} \tag{4}$$

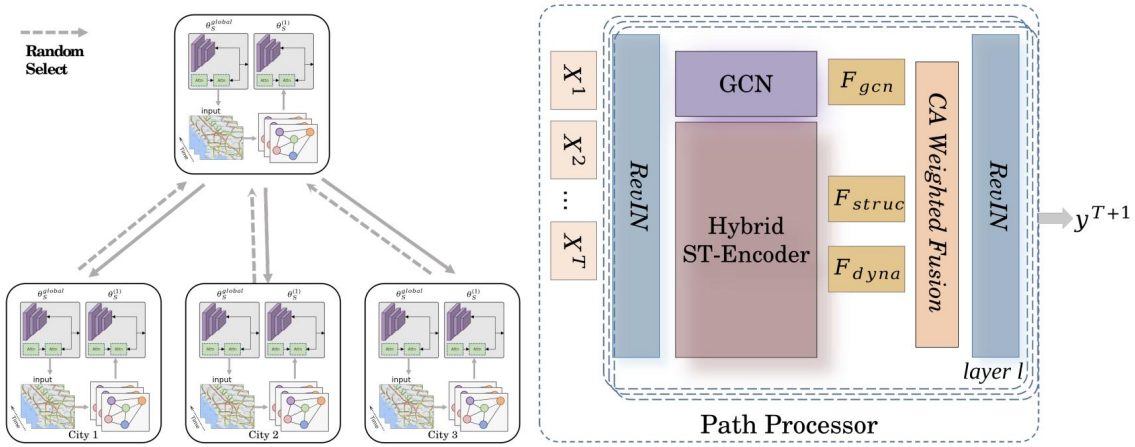
Upload to Server: $\left\{\mathbf{P}_g^{(k)},\ \mathbf{R}_k,\ \mathbf{C}_k\right\}$ | **Retain Locally:** $\mathbf{P}_l^{(k)}$

- ✓ Server: **0 training iterations**
- ✓ Client retains personalised context
- ✓ Single communication round

Milestone 5 — One–Shot FL for Spatiotemporal Cyber–Physical Traffic Networks

《Only One Fusion Matters: Spatiotemporal Weighted Integration for OneShot Federated Traffic Prediction》

The core bottleneck for smart-city CPS: each city's road topology graph G^(k) is physically non-shareable and asymmetric (node counts N_k ≠ N_j). Conventional FL forces synchronisation of the entire spatiotemporal GNN, entangling spatial topology with global temporal information — server-side averaging then injects severe "topological noise interference."



1. Hybrid Spatiotemporal Decomposition

Splits features channel-wise into Identity / Spatial / Temporal streams — only the universal temporal law is aggregated; city-specific topology (GCN) stays local.

2. Context-Adaptive Fusion

An adaptive gate re-weights the spatial vs. temporal streams in real time, based on current local traffic context.

Didi–Chengdu Real–world Benchmarks (3–Day Train)

300×

less communication bandwidth than iterative multi-round FL

2.885

SWIFT MAE (60m) one-shot

3.370

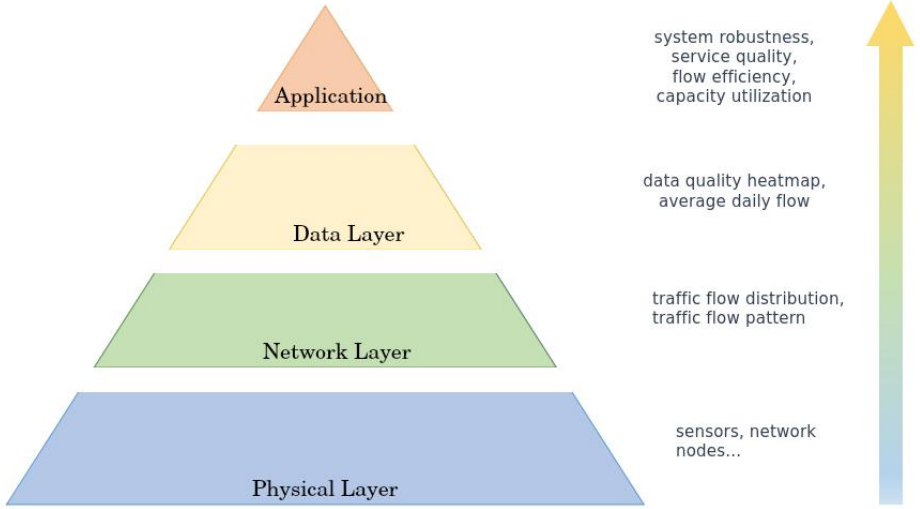
pFedCTP MAE (60m) one-shot

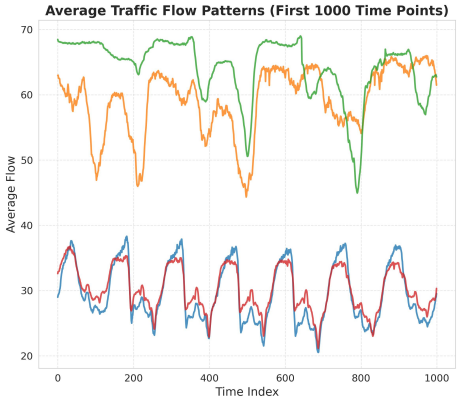
Table 2: Performance comparison of cross-city traffic prediction on Didi-Chengdu and Didi-Shenzhen datasets. (Note: The best results are highlighted in bold.)

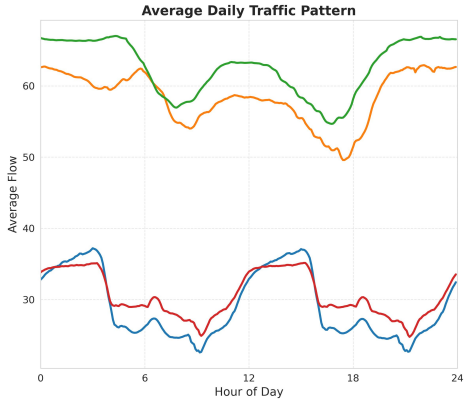
Methods	Didi-Chengdu (Target city)						Didi-Shenzhen (Target city)					
	MAE (↓)			RMSE (↓)			MAE (↓)			RMSE (↓)		
	10 min	30 min	60 min	10 min	30 min	60 min	10 min	30 min	60 min	10 min	30 min	60 min
Target-only baselines												
STGCN [22]	2.930	2.905	3.093	4.249	4.248	4.491	2.740	2.738	2.873	3.858	3.877	4.087
T-GCN [26]	2.340	2.850	3.350	3.369	4.136	4.845	2.109	2.589	3.010	3.177	3.920	4.572
ST-Net [25]	2.327	2.768	3.135	3.256	3.994	4.570	1.932	2.355	2.675	2.804	3.551	4.124
Federated cross-city baselines(multi-round)												
FedAvg [15]	3.181	3.56	4.005	4.323	4.852	5.428	2.621	3.131	3.571	3.623	4.380	5.009
FedProx [9]	3.086	3.645	4.088	3.940	4.777	5.400	2.473	2.904	3.290	3.290	4.003	4.613
pFedCTP [25]	**2.237**	2.790	3.246	**3.164**	3.993	4.622	**2.030**	2.491	2.882	**2.859**	**3.616**	4.206
Ours	2.34	**2.736**	**3.052**	3.281	**3.923**	**4.426**	2.098	**2.509**	**2.755**	2.935	3.62	**4.126**
One-shot federated learning												
FedAvg [15]	6.185	6.468	6.838	7.857	8.166	8.596	8.79	8.868	8.869	11.694	11.846	11.84
pFedCTP [25]	3.159	3.370	3.679	4.612	4.912	5.307	3.656	3.674	3.732	5.338	5.351	5.394
Ours	**2.545**	**2.867**	**3.138**	**3.516**	**4.097**	**4.542**	**2.161**	**2.516**	**2.885**	**3.048**	**3.695**	**4.323**

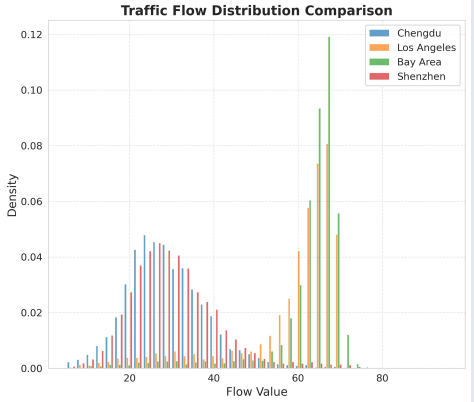
Milestone 5 — One–Shot FL for Spatiotemporal Cyber–Physical Traffic Networks

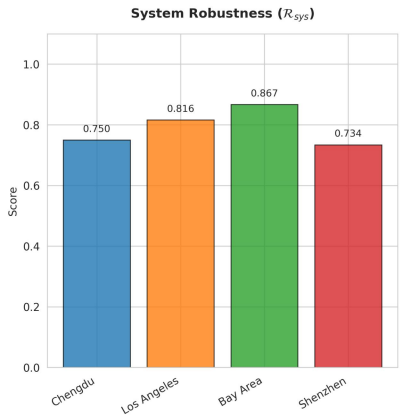
《Only One Fusion Matters: Spatiotemporal Weighted Integration for OneShot Federated Traffic Prediction》

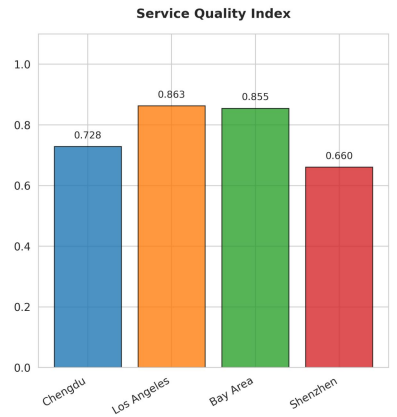


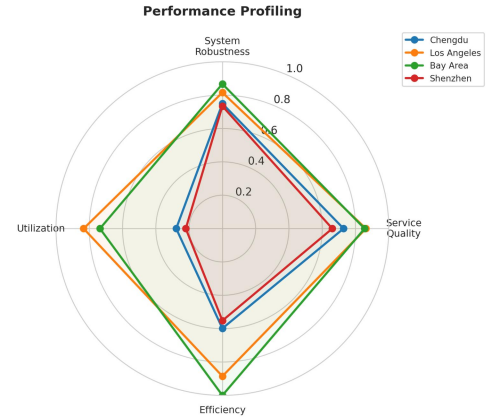










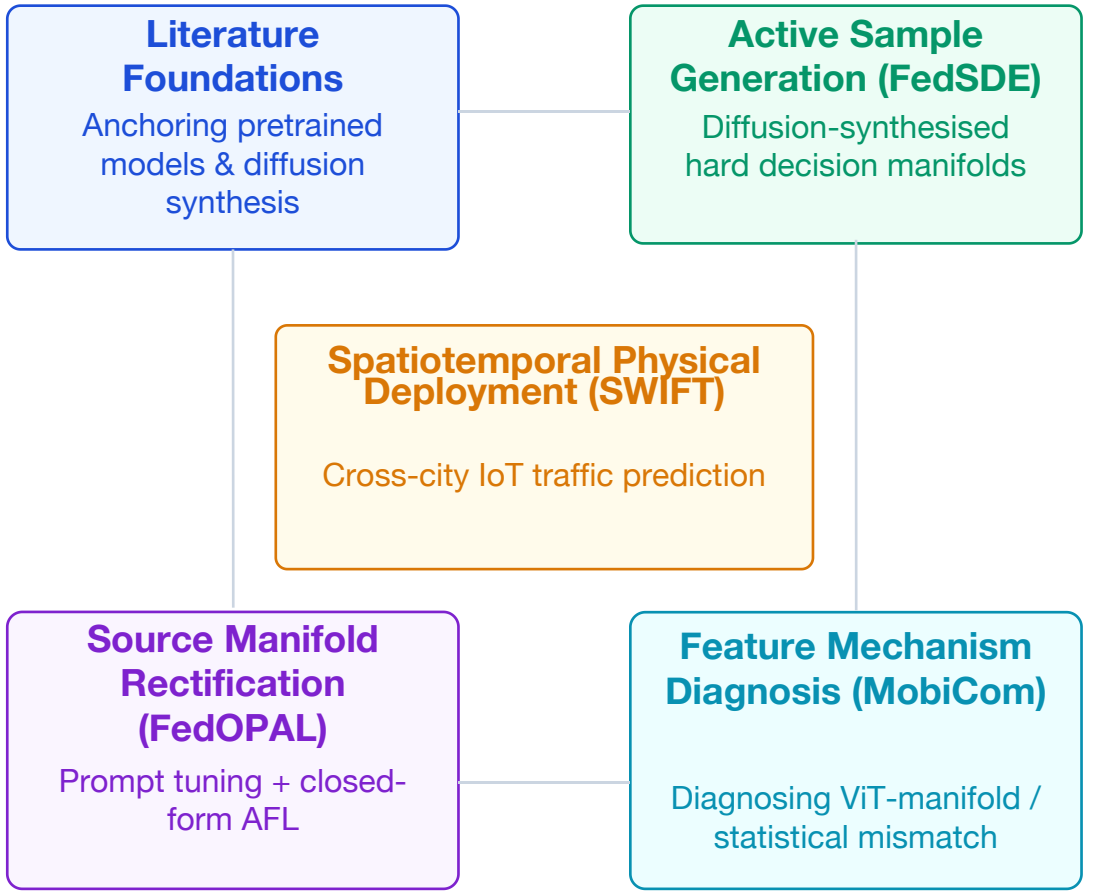


UNIFIED SCIENTIFIC SYNTHESIS

Five Dimensions, One Loop: The OFL Research Converges

A System-Level Interplay of Representation Isolation and Geometric Rectification

1. With no data to share, **FedSDE** (diffusion self-distillation) completes the decision boundary locally.
2. We then diagnosed the potential of **pretrained backbones in OFL** (MobiCom '25), catching the geometric mismatch between statistical averaging and ViT feature space.
3. The lightweight **FedOPAL** then irons heterogeneous manifolds into a linear state at the source, making closed-form AFL aggregation work perfectly.
4. Finally, “decoupled representation, shared core” is mapped onto complex traffic networks : **SWIFT** (spatiotemporal decoupling) blocks topological noise, completing an industrial IoT deployment.



Limitations and Open Technical Challenges

Unified OFL Framework — Dimensional Comparison

Method	Client Computation	Communication Payload	Privacy Level
FedSDE (Diffusion self-distill.)	High (local U-Net training)	Very low (soft-label logits only)	High (no gradient transfer)
FedOPAL (Prompt alignment)	Very low (frozen, prompt-only)	Very low (sufficient statistics only)	High (closed-form security barrier)
SWIFT (Spatiotemporal CPS)	Medium (dual-stream adaptive GCN)	Medium (temporal core only)	Very high (topology stays local)

Known Bottlenecks — Honest Self–Assessment

- **Local hardware burden:** FedSDE's 20-epoch conditional diffusion training on edge gateways is costly (275s on an RTX A6000).
- **Cross-domain generalisation gap:** FedOPAL is highly sensitive to the downstream domain — CLIP's natural-image pretraining degrades feature rectification on digit-only SVHN, trailing conventional methods by ~10%.
- **Static-topology assumption:** SWIFT's spatiotemporal decoupling assumes a static road network; under high-frequency topology changes, the adaptive GCN cannot rebuild the graph in real time.

** In line with the Marie Skłodowska-Curie review principle of academic honesty and transparent disclosure of limitations.*

Publications and Scientific Outputs

SURVEY	**Towards One-Shot Federated Learning: Advances, Challenges, and Future Directions** Amato, F., **Qiu, L.***, Tanveer, M., Cuomo, S., Annunziata, D., Giampaolo, F., & Piccialli, F. — Neurocomputing (2026), Vol. 664. DOI: 10.1016/j.neucom.2025.132088	[Published]
SYNTHESIS	**FedSDE: Self-Distillation with Diffusion Enhanced for One-Shot Federated Learning** **Qiu, L.**, Annunziata, D., Giampaolo, F., & Piccialli, F. — IEEE BigData Conference (2025). DOI: 10.1109/BigData63153.2025.11402460	[Published]
DIAGNOSIS	**Deciphering One-Shot Federated Learning: The Pivotal Role of Pre-trained Models** **Qiu, L.**, Annunziata, D., Giampaolo, F., & Piccialli, F. — FLEdge-AI MobiCom (2025), Hong Kong. DOI: 10.1145/3737899.3768517	[Published]
RECTIFICATION	**FedOPAL: One-Shot Federated Learning via Analytic Visual Prompt Tuning** **Qiu, L.**, Annunziata, D., Izzo, S., Giampaolo, F., & Piccialli, F. — FLICS 2026. Code: github.com/Lynn0925/FLICS	[Published]
CPS	**Only One Fusion Matters: Spatiotemporal Weighted Integration for One-Shot Federated Traffic Prediction** **Qiu, L.**, Annunziata, D., Izzo, S., Giampaolo, F., & Piccialli, F. — Cyber-Physical Systems / Computer Networks (2026). Code: github.com/Lynn0925/SWIFT	[Published]

All research code is fully open-sourced on GitHub, in strict compliance with the EU Horizon programme's Open Access requirements.

ROADMAP & FUTURE BEYOND

Next Research Steps

1. Extending OFL to Privacy-Sensitive Structured Tabular Data

Current validation focuses on images and road networks. Yet many of the **most sensitive IoT data types** (EHRs, health monitoring, financial transactions) are **tabular**. The next step is **transferring the visual-prompt rectification** principle to structured tabular foundation models.

2. Introducing Rigorous, Quantifiable Differential Privacy (ε,δ-DP)

Although only soft labels or first/second-order statistics are uploaded in a single round, no **formal noise guarantees currently exist**. The priority next step is **injecting calibrated Laplace or Gaussian noise** into FedSDE's and SWIFT's aggregated statistics, establishing an auditable privacy compliance moat.

Thank You
Questions & Discussion

Towards an Understanding of Artificial Intelligence (TUAI).

Email: lingyu.qiu@unina.it

EU

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Marie Skłodowska-Curie grant agreement No 101168344 (TUAI).