

A Temporal Stability Metric for Clustering-Aware Federated Learning in UAV-Assisted Vehicular Networks

Cherifa Boucetta^{*} and Oumaya Baala[†]

^{*}LIGM CNRS-UMR 8049, Université Gustave Eiffel, Champs-sur-Marne, France

[†]UTBM, CNRS, Institut FEMTO-ST, F-90010 Belfort Cedex, France

Abstract—Unmanned Aerial Vehicles (UAVs) are increasingly deployed as mobile edge aggregators in vehicular federated learning (FL) systems, yet many existing architectures primarily optimize communication, placement, or resource allocation without using temporal partition stability as an online reconfiguration signal. This paper presents an adaptive hierarchical FL framework in which UAV edge nodes dynamically reconfigure their clustering assignments based on traffic-derived behavioral features. Rather than grouping vehicles by geographic proximity, the proposed approach partitions nodes in an emission-aware latent space derived from Vehicle-Specific Power (VSP) proxies. A lightweight rule-based adaptation policy enables each UAV to select among three online clustering engines while issuing localized eco-advisory controls without exposing raw trajectory data, and five clustering paradigms are evaluated comparatively. Experiments on the NGSIM US-101 dataset show that the Silhouette Score exhibits a weaker association with FL convergence, whereas the proposed Cluster Drift Stability (CDS) metric exhibits a markedly stronger negative correlation with the number of rounds to convergence ($r = -0.960$, $p = 0.010$), reported as preliminary evidence over the five evaluated paradigms. Results further indicate that BIRCH-based partitioning enables targeted mitigation of high-emission subpopulations, achieving an estimated 29.1% localized CO_2 reduction under the emission model of Section III with a cluster of only 10 vehicles in free-flow conditions, while both GMM and BIRCH reach the lowest observed convergence count of 5 FL rounds.

Index Terms—Federated Learning, Unmanned Aerial Vehicles, Internet of Vehicles, Cluster Drift Stability, Emission Behavior Clustering, Green Corridor.

I. INTRODUCTION

The convergence of Intelligent Transportation Systems (ITS) and mobile edge computing offers new opportunities to mitigate the environmental footprint of arterial highways [1]. Stop-and-go waves arising during congestion dissolution create localized emission hotspots where exhaust output peaks sharply. Vehicular Ad-Hoc Networks (VANETs) and V2X infrastructures provide the communication fabric needed to smooth these micro-cycles, yet centralized traffic management remains penalized by prohibitive latency and by the exposure of identifiable mobility data.

Federated Learning (FL) has consequently been widely adopted in the Internet of Vehicles (IoV), letting edge nodes collaboratively train a shared model without uploading raw trajectories to a distant server [2]. Hierarchical FL (HFL) inserts Unmanned Aerial Vehicles (UAVs) as regional orchestrators between highly dynamic vehicle clusters and the cloud

[3]. Many UAV-assisted HFL designs, however, optimize communication, placement or resource-allocation variables while treating the learning partition as fixed or externally determined, which can limit adaptation to rapidly changing mobility and to strongly non-independent and identically distributed (Non-IID) data. In parallel, the paradigm of *agentic AI*, i.e. systems that perceive their environment, plan and act within a closed loop [4], has matured rapidly, but its application to the core mechanisms of federated optimization in vehicular environments remains largely unexplored.

This paper introduces Agentic-FL, a framework that turns UAV edge aggregators into adaptive orchestrators of emission-aware vehicular FL. Each UAV tracks vehicle kinematics, quantifies micro-emissions, and shapes its local FL space by clustering vehicles in a latent space derived from their emission behavior rather than from geographic proximity. We use *agentic* to denote the closed perceive–plan–act–learn loop and the reconfiguration authority delegated to the edge node: the decision policy itself is a deterministic rule table rather than a learned or planning-based controller, and the federated back-end is a standard two-stage FedAvg hierarchy. The contribution accordingly lies in the emission-aware partitioning, its temporal characterization, and the orchestration policy, not in a new aggregation algorithm. We further observe a methodological gap in the evaluation of clustered vehicular FL: static geometric metrics such as the Silhouette Score or the Davies-Bouldin Index do not capture the temporal fluidity of mobile nodes and may therefore provide limited information about FL convergence. We therefore formalize the Cluster Drift Stability (CDS) metric, which quantifies the temporal stationarity of clusters in the latent space.

Our core contributions are as follows:

- 1) We propose an adaptive hierarchical FL architecture for UAV-assisted vehicular edge coordination, in which lightweight edge policies dynamically reconfigure local clustering assignments.
- 2) We design an emission-aware latent feature representation and evaluate five clustering paradigms under dynamic vehicular conditions.
- 3) We introduce the CDS metric and provide preliminary empirical evidence that temporal cluster stability is more strongly associated with FL convergence than conven-

tional geometric clustering metrics in the evaluated configurations.

- 4) Using real NGSIM trajectories, we estimate that adaptive BIRCH-based partitioning enables targeted emission mitigation, yielding a model-based 29.1% CO₂ reduction in high-emission zones with only 10 vehicles.

Section II reviews related work, Section III introduces the system and emission models, Section IV presents Agentic-FL and the CDS metric, Section V discusses the empirical results, and Section VI concludes.

II. RELATED WORK

FL in VANETs enables distributed model training under privacy constraints, but single-tier designs cope poorly with high mobility and intermittent connectivity. HFL mitigates this through intermediate edge aggregators: HaghighiFard and Coleri [3] construct multi-hop vehicular clusters from a combination of relative mobility and model similarity to address mobility and non-IID data, while Pandya et al. [5] and Nguyen et al. [6] show that non-IIDness markedly destabilizes the global model. These works address mobility and hierarchical aggregation, but do not use temporal partition stability as an online signal for selecting among alternative emission-aware clustering strategies.

UAVs deployed as flexible edge aggregators offer considerable agility for transport systems [7], and Yuan et al. [8] couple hierarchical aggregation with cooperative caching for decentralized FL in UAV-assisted IoV. Closest to our setting, Zhang and Hanzo [9] formalize FL over multi-UAV networks and quantify how aerial link quality conditions model aggregation, while Pham et al. [10] jointly optimize UAV placement, power control and resource allocation to sustain FL under strict energy budgets. On the communication-efficiency side, Liu et al. [11] show that a client–edge–cloud hierarchy sharply reduces the number of costly cloud rounds, and Barbieri et al. [12] quantify the resulting energy footprint at the edge. Their primary objective is communication, aggregation or resource efficiency; they do not investigate online switching among alternative emission-aware vehicle partitions using a temporal-stability signal, which is where our contribution lies. The broader Agentic AI paradigm, which describes closed-loop multi-step autonomous reasoning, has meanwhile been applied to the autonomous management of network and application services [4], but not to the structure of federated optimization spaces in vehicular networks.

Clustering has long structured communication rather than learning in vehicular networks. Our earlier work showed that optimizing cluster formation improves data-collection efficiency [13], and that topology-based routing performance is tightly coupled to radio-connectivity dynamics [14]. The present study transposes this clustering perspective from the communication plane to the learning plane, where the quantity of interest is no longer delivery ratio but the temporal stability of the federated partition. Despite advances in eco-routing and vehicle-specific power estimation [15], few existing approaches jointly couple mobility, emission-aware partitioning,

and federated-learning dynamics within the same adaptive edge framework. We address this gap with a unified framework in which UAV edge nodes organize vehicles in an emission-aware latent space and adapt FL protocols to environmental and system objectives.

III. SYSTEM MODEL

A. Multi-Tier Edge Computing Network Architecture

We formalize a dynamic multi-tier mobile edge computing topology for privacy-preserving vehicular optimization, structured across the three layers of Fig. 1.

1) *Layer 1: Mobile Edge Vehicles ($\mathcal{V}$):* Let $\mathcal{V}(t) = \{v_1, \dots, v_{N(t)}\}$ denote the set of vehicles active at time t . Each v_i maintains a non-shareable local dataset $\mathcal{D}_i$ of kinematic observations sampled at 10 Hz and performs local optimization without transmitting its raw trajectories.

2) *Layer 2: Agentic Edge Aggregators ($\mathcal{U}$):* Let $\mathcal{U} = \{u_1, \dots, u_K\}$ be a swarm of $K = 3$ UAVs. Each u_k governs a behavior-driven cluster $\mathcal{C}_k(t) \subseteq \mathcal{V}(t)$, with $\bigcup_k \mathcal{C}_k(t) = \mathcal{V}(t)$ and $\mathcal{C}_k(t) \cap \mathcal{C}_m(t) = \emptyset$ for $k \neq m$, and hosts an on-board policy π that modulates the clustering and optimization space from computational and environmental indicators. The present evaluation abstracts the air-to-ground channel and assumes that active vehicles can communicate with their assigned aggregator; link-level reliability, mobility-induced connectivity and UAV trajectory optimization are outside the scope of this study.

3) *Layer 3: Central Cloud Coordination:* A cloud tier maintains the global parameters $w^{(t)}$, periodically aggregates UAV-level updates, and synchronizes highway segments without communicating with individual vehicles.




Fig. 1. Three-layer hierarchical architecture of the proposed Agentic-FL network.

B. Kinematic Micro-Emission Modeling Framework

For local emission-aware processing, vehicle kinematics are mapped to a behavioral emission representation using a Vehicle-Specific Power (VSP) proxy; privacy follows from the federated architecture, which never transmits raw trajectories, rather than from the transform itself. For vehicle v_i , the instantaneous VSP (kW/t) is

$$\text{VSP}_i(t) = v_i(t) \cdot (a_i(t) + g \cdot C_R + C_D \cdot v_i(t)^2) \quad (1)$$

with $v_i(t)$ the velocity (m/s), $a_i(t)$ the acceleration (m/s²), $g = 9.81 \text{ m/s}^2$, $C_R = 0.0135$ the rolling resistance and $C_D =$

0.000302 the aerodynamic drag coefficient. The mass emission rate proxy $\widehat{\text{CO2}}_i(t)$ (g/s) follows from the non-linear mapping

$$\widehat{\text{CO2}}_i(t) = \epsilon(c_i) \cdot \frac{\max(\text{VSP}_i(t), 0)}{1000} + \delta \quad (2)$$

where $c_i \in \{1, 2, 3\}$ types the vehicle as motorcycle ($\epsilon_1 = 0.08$), passenger vehicle ($\epsilon_2 = 0.21$) or Heavy Goods Vehicle (HGV, class 3, $\epsilon_3 = 0.72$), and $\delta = 0.02$ g/s is the idle emission baseline. Eq. (2) is a first-order behavioral proxy, not a certified emission inventory; it is applied consistently as a *relative* comparator across clustering paradigms, and every emission figure reported in Section V inherits its assumptions.

C. Hierarchical Federated Optimization (H-FedAvg)

We consider the two-stage H-FedAvg architecture targeted by the framework. In the full deployment, vehicles minimize a local loss $L_i(w)$ over their private history $\mathcal{D}_i$ and each global round t proceeds in two stages.

1) *Stage 1 (Vehicle-to-UAV)*: Each $v_i \in \mathcal{C}_k(t)$ downloads the cluster model $w_k^{(t)}$ and runs E local epochs of FedAvg-style SGD [2], yielding the delta $\Delta w_i^{(t)} = w_i^{(t,E)} - w_k^{(t)}$ sent to u_k , which applies uniform model averaging

$$w_k^{(t+1)} = w_k^{(t)} + \frac{1}{|\mathcal{C}_k(t)|} \sum_{v_i \in \mathcal{C}_k(t)} \Delta w_i^{(t)} \quad (3)$$

2) *Stage 2 (UAV-to-Cloud)*: After T cluster-level rounds the UAVs synchronize with the cloud, which aggregates cluster models weighted by active cluster sizes:

$$w^{(t+1)} = \sum_{k=1}^K \frac{|\mathcal{C}_k(t)|}{N} w_k^{(t+1)} \quad (4)$$

IV. AGENTIC-FL METHODOLOGY

A. The On-Board Perceive-Plan-Act-Learn (PPAL) Core

Each UAV hosts a goal-directed Perceive-Plan-Act-Learn (PPAL) loop whose state at epoch t is the tuple $s_k^{(t)} = (\widehat{\text{CO2}}_k^{(t)}, \text{CDS}_k^{(t)}, \phi_k^{(t)})$, built from three quantities the UAV observes locally: the normalized cluster emission level $\widehat{\text{CO2}}_k^{(t)} = \widehat{\text{CO2}}_k^{(t)} / \max_i \widehat{\text{CO2}}_i$, the Cluster Drift Stability $\text{CDS}_k^{(t)}$, and the boundary-ambiguity proxy $\phi_k^{(t)} = 1 - \bar{s}_k^{(t)}$, where $\bar{s}_k^{(t)}$ is the mean Silhouette coefficient, larger $\phi_k^{(t)}$ indicating weaker cluster separation. A rule-based policy maps this state to a discrete action, $\pi(s_k^{(t)}) \rightarrow a_k^{(t)} \in \{\mathcal{A}_{\text{clustering}}, \mathcal{A}_{\text{control}}\}$, where $\mathcal{A}_{\text{clustering}}$ reconfigures the clustering engine and $\mathcal{A}_{\text{control}}$ modulates the green speed advisories broadcast over V2I links.

Table I lists the four rules of π , evaluated by priority. The thresholds are heuristic operating points fixed a priori for interpretability, and were not fitted to maximize the reported convergence or emission results. Each is a fraction of an observable whose scale is set by the data, which makes the policy largely insensitive to multiplicative scaling of the emission proxy: urgency fires when the cluster mean exceeds 12% of the peak emission on the segment, instability when fewer than half the persistent vehicles keep their zone across snapshots, and negotiation when the boundary-ambiguity proxy exceeds 0.30. On the NGSIM regimes of Section V the observed values

stay far from these boundaries— ϕ spans 0.45 to 0.81 against a threshold of 0.30, and CDS spans 0.28 to 0.97 against 0.5—so displacing a threshold by a few percentage points changes neither the rule fired nor the algorithm selected; only the CDS rule is close enough for a ± 0.1 shift to matter.

TABLE I
RULE-BASED POLICY π — IF-THEN TRIGGERS, BY PRIORITY

Trigger condition	Action	Rationale
$\widehat{\text{CO2}}_k > 0.12$	GMM + urgent advisory	soft boundaries
$\text{CDS}_k < 0.5$	switch to BIRCH	$\mathcal{O}(N)$, incremental
$\phi_k > 0.30$	K-Means + negotiation	lowest $\mathcal{O}(NKD)$
otherwise	keep engine, periodic FL	no reconfiguration

B. Multi-Modal Latent Feature Extraction

Incoming high-frequency fields are consolidated into $\mathbf{f}_i(t) = [\widehat{\text{CO2}}_i(t), v_i(t), a_i(t), h_i(t), c_i(t)]^T$ (h_i : space headway), standardized, then projected by a local Principal Component Analysis (PCA) layer onto three orthogonal components, which limits multi-collinearity and on-board computing delay:

$$\mathbf{z}_i(t) = \mathbf{W}_{\text{PCA}}^T \left(\frac{\mathbf{f}_i(t) - \boldsymbol{\mu}_{\mathbf{f}}}{\boldsymbol{\sigma}_{\mathbf{f}}} \right) \quad (5)$$

where $\mathbf{W}_{\text{PCA}}$ captures 85% of the cumulative variance across the fleet.

C. The Five Clustering Paradigms

We evaluate five clustering paradigms to characterize the relationship between partition stability and FL convergence. The online policy of Table I selects among GMM, BIRCH and K-Means; Agglomerative and Spectral serve as comparative benchmarks and are not selected online. *K-Means* is a hard-partitioning baseline minimizing the within-cluster sum of squares $\sum_k \sum_{\mathbf{z}_i \in \mathcal{C}_k} \|\mathbf{z}_i - \boldsymbol{\mu}_k\|^2$ around centroids $\boldsymbol{\mu}_k$: complexity is only $\mathcal{O}(NKD)$, but hard boundaries are unstable under dense traffic fluctuations. *Gaussian Mixture Models* (GMM) model the latent emission distribution as $p(\mathbf{z}_i) = \sum_k \pi_k \mathcal{N}(\mathbf{z}_i | \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)$ fitted by Expectation-Maximization; the fractional memberships $\gamma_{ik} = p(k | \mathbf{z}_i)$ avoid abrupt reassignments near boundaries and yield stable gradient spaces. *BIRCH* incrementally builds a Clustering Feature tree $CF = (N, \mathbf{LS}, SS)$ over a continuous stream at $\mathcal{O}(N)$ cost; the policy selects it when CDS indicates partition instability, exploiting its incremental updates without requiring full recomputation. *Agglomerative* clustering merges clusters bottom-up under Ward's linkage $\Delta(A, B) = \frac{|A||B|}{|A|+|B|} \|\boldsymbol{\mu}_A - \boldsymbol{\mu}_B\|^2$, giving clean hierarchies at $\mathcal{O}(N^2 \log N)$, suited to low-density fluid states. Finally, *Spectral* clustering projects the data onto the eigenspace of the Laplacian $\mathbf{L} = \mathbf{D} - \mathbf{A}$ with affinity $\mathbf{A}_{ij} = \exp(-\|\mathbf{z}_i - \mathbf{z}_j\|^2 / 2\sigma^2)$, capturing non-convex behavioral distributions at high purity but $\mathcal{O}(N^3)$ cost.

D. Formalization of Cluster Drift Stability (CDS)

To address the shortcomings of static geometric metrics in mobile environments, we introduce the *Cluster Drift Stability* (CDS), which measures how persistently the partition assigns

the same vehicles to the same UAV zone as traffic evolves. Let t_0 be the reference snapshot and $\mathcal{P}_\ell = \mathcal{V}(t_0) \cap \mathcal{V}(t_0 + \ell)$ the vehicles observed at both t_0 and $t_0 + \ell$. For a lag ℓ , the assignment-stability score is

$$S_\ell = \frac{1}{|\mathcal{P}_\ell|} \sum_{v_i \in \mathcal{P}_\ell} \mathbb{I}[\hat{c}_i(t_0) = \hat{c}_i(t_0 + \ell)] \quad (6)$$

where $\hat{c}_i(\cdot)$ is the zone assigned to v_i and $\mathbb{I}[\cdot]$ the indicator function. The reported metric averages the two evaluated lags,

$$\text{CDS} = \frac{1}{2} (S_{20} + S_{40}) \quad (7)$$

with lags expressed in seconds, giving an effective observation window $W = 40$ s; each lag is also reported separately in Section V. For paradigms without a native out-of-sample predictor (Agglomerative, Spectral), labels at later snapshots are obtained from a nearest-centroid classifier fitted on the reference partition, which places all five paradigms behind a single interface. Higher $\text{CDS} \in [0, 1]$ indicates more persistent assignments, which motivates testing whether temporally stable partitions are associated with faster FL convergence in Section V.

E. Green Eco-Advisory Control Protocol

When an emission anomaly is detected, the advisory is issued per zone rather than per vehicle: the UAV ranks its K clusters by mean emission into LOW, MEDIUM and HIGH tiers, and broadcasts to each tier a speed cap and an acceleration-smoothing factor over low-latency V2I channels,

$$(v_i^{\text{adv}}, a_i^{\text{adv}}) = (\sigma_{\rho(k)} \cdot v_i(t), \alpha_{\rho(k)} \cdot a_i(t)) \quad (8)$$

where $\rho(k)$ is the emission rank of the cluster hosting v_i , and the tier coefficients are $(\sigma, \alpha) = (1.00, 0.95)$ for LOW, $(0.90, 0.70)$ for MEDIUM and $(0.80, 0.50)$ for HIGH. The advised kinematics are re-injected into Eqs. (1)–(2) to obtain the post-advisory emission. Damping the recorded traces directly amounts to assuming that targeted drivers follow the advisory immediately and fully; compliance is revisited in Section V-E.

V. SIMULATION SETUP AND RESULTS

A. Experimental Configuration and Dataset

The evaluation uses real trajectories from the public NGSIM US-101 dataset [16], recorded at 10 Hz over a 640-meter segment of the Hollywood Freeway, Los Angeles. Exploited fields include `v_Vel`, `v_Acc`, `Lane_ID`, `Space_Hdwy` and `v_Class`; all positional and kinematic values are converted from feet to metric units. To replicate mixed-fleet perturbations, the 2,169 original vehicle profiles are augmented with 45 synthetic class-3 HGV records, distributed across the three analyzed snapshots of the 953-s trace. The synthetic records are drawn per analyzed frame (15 per regime, 3 regimes) from a multivariate normal fitted by maximum likelihood on the empirical (speed, acceleration, headway) distribution of the real NGSIM class-3 vehicles, with the mean speed rescaled by regime (0.50 dense, 1.00 mixed, 1.80 free-flow) to respect the physics of each state, under fixed seed 42; no synthetic samples are used outside this class-3 augmentation. The 9,529 frames are ranked by mean speed into three velocity bands, and

the median frame of each is analyzed: a Dense State (frame 7679, 32.8 km/h, 162 vehicles, heavy congestion and severe engine cycling), a Mixed State (frame 249, 44.3 km/h, 82 vehicles, transitory patterns) and a Fluid State (frame 1529, 51.0 km/h, 122 vehicles, free flow, where speed optimization is most active). Feature vectors $\mathbf{f}_i$ are compressed by PCA to 3 components retaining 85% of the cumulative variance.

Federated learning setup. The local task regresses $\widehat{\text{CO}_2}(t+30)$, i.e. 3 seconds ahead, from the standardized vector $\mathbf{f}_i(t)$, the latent space of Eq. (5) serving only to derive the cluster assignments. The target is z -score normalized and each cluster fits a linear model through $E = 5$ full-batch gradient steps per round under a mean-squared-error objective, with $\eta = 0.008$ over $R = 50$ rounds. The reported loss is the global MSE normalized by its first-round value, and convergence is declared at the first round whose relative improvement drops below 0.5%. That loss is measured on the training population, with no held-out partition, so convergence rounds characterize optimization speed rather than generalization (Section V-E). They come from a single run at seed 42 shared by every stochastic component, whereas the Silhouette, DBI and CO_2 figures are averaged over five seeds $\{42, 0, 1, 7, 99\}$. Clustering and CDS use the augmented dense snapshot (162 vehicles) and the advisory the augmented free-flow one (122); each synthetic HGV exists at a single frame and thus carries no look-ahead target, so the FL dataset comprises 1,096,093 records from the 2,169 original NGSIM vehicles alone. A *random-partition* control assigns vehicles uniformly to $K = 3$ groups under the same model, optimizer, epochs and convergence criterion. To isolate the effect of the clustering structure, the evaluation abstracts the vehicle-to-UAV execution of Eq. (3): each cluster model is optimized on the pooled records of that cluster, while cloud aggregation follows Eq. (4). The results therefore characterize clustering-aware optimization behavior rather than a complete vehicle-level HFL deployment.

B. Latent Space Topology and Structural Validation

We select $K = 3$ as a compromise between geometric validity and semantic interpretability. For GMM, the Davies-Bouldin Index decreases from 1.200 at $K = 2$ to 1.024 at $K = 3$, whereas the Silhouette Score decreases only marginally from 0.556 to 0.549 (1.3%), as shown by the curves of Fig. 2. We therefore retain $K = 3$ to preserve distinct LOW (≈ 79 g/h), MEDIUM (≈ 109 g/h) and HIGH (≈ 124 g/h) emission tiers. The resulting latent topology, mapped over the top principal components in Fig. 3, shows that the system separates vehicles by environmental impact rather than geographic proximity.

On instantaneous geometric quality, GMM achieves the highest Silhouette Score (0.549), whereas Agglomerative yields the lowest DBI (0.864). K-Means and BIRCH reach Silhouette Scores of 0.388 and 0.387, while Spectral clustering drops to 0.186. The disagreement between the two geometric indices motivates evaluating temporal stability separately.




Fig. 2. Silhouette Score and Davies-Bouldin Index (DBI) across $K = 2$ to 6.




Fig. 3. Latent space visualization of behavior-driven vehicular clustering across the top principal components.

C. Association of CDS with FL Convergence

Fig. 4 maps both validation metrics onto the empirical FL convergence rounds. CDS yields the stronger sample correlation ($r_{\text{CDS}} = -0.960$, $p = 0.010$) compared with the Silhouette Score ($r_{\text{Sil}} = -0.670$, $p = 0.216$). A 2,000-replicate exploratory bootstrap gives a narrower percentile interval for CDS, $[-1.000, -0.905]$, than for the Silhouette Score, $[-1.000, -0.246]$; these intervals should be read cautiously, since only five clustering paradigms are available. The association is also consistent across the two tested lags: $r = -0.934$ at +20 s and $r = -0.969$ at +40 s, against -0.960 for their average, GMM and BIRCH remaining the two most stable paradigms at either lag.




Fig. 4. Association between FL convergence rounds and (a) the Silhouette Score and (b) CDS across the five evaluated clustering paradigms.

The Silhouette Score therefore provides a less informative association with convergence than CDS in the evaluated configurations. This observation is consistent with temporal cluster stability contributing to faster federated optimization, although the limited number of paradigms prevents establishing a general causal relationship. The two paradigms with

the highest CDS, GMM (0.968) and BIRCH (0.745), both converge in 5 rounds, as shown by the loss trajectories of Fig. 5. The trajectories also separate during the first few rounds, although the present experiment does not isolate the mechanism responsible for this difference. Agglomerative clustering is an informative control: it attains the best DBI (0.864) yet still needs 6 rounds, like K-Means and Spectral, and its CDS (0.361) is close to K-Means (0.340), indicating that geometric quality does not translate into temporal zone stability. The random-partition baseline of Fig. 5 converges more slowly than all five structured paradigms, suggesting that structured emission-aware partitioning contributes to the observed convergence improvement relative to that control.




Fig. 5. Global FL loss convergence per clustering paradigm, including the random-partition baseline control.

D. Edge Advisory Efficiency Across Traffic Regimes

The decision policy triggers localized velocity harmonization in a strongly regime-dependent way. In the Dense State, adaptation is physically constrained by headway: collective CO_2 reduction is compressed between 2.1% (GMM) and 3.8% (Spectral), while high-polluting zones saturate near 24%. In the Fluid State, by contrast, the active algorithm dictates the intervention strategy, as mapped by the bubble chart of Fig. 6.

A trade-off emerges between broad and targeted intervention. Spectral clustering yields the largest modeled segment-wide reduction (9.3%) but requires issuing advisories to 71 vehicles, making the intervention substantially less selective. BIRCH instead targets **10 vehicles** (the smallest footprint across all paradigms) and yields a modeled **29.1%** CO_2 reduction within that cluster, while leaving the remaining vehicle trajectories unmodified in the open-loop calculation. Both percentages derive from Eqs. (2) and (8) and therefore quantify the relative merit of the partitioning strategies rather than measured field savings.

E. Limitations and Scope of the Claims

Several limitations bound the conclusions. (i) *Statistical scope.* The CDS–convergence analysis covers only five clustering paradigms and one FL seed, and the convergence rounds take just two distinct values (5 or 6); the two-lag analysis therefore provides preliminary sensitivity evidence rather than broad statistical validation, and confirming generality requires a wider panel of algorithms, further segments and datasets, and




Fig. 6. Modeled advisory trade-off under dense and free-flow traffic regimes; bubble size scales with the number of high-emission vehicles receiving the advisory.

repeated seeds. (ii) *Generalization*. The FL loss is evaluated on the training population, so convergence reflects optimization stability rather than performance on unseen vehicles. (iii) *Agent autonomy*. The UAV policy is a deterministic rule table; autonomy refers to the closed PPAL loop and online reconfiguration, not to learned policy synthesis, and the rules are exercised on snapshots rather than over a continuously simulated mission. (iv) *Communication*. Air-to-ground links are abstracted and assumed available throughout the present evaluation. (v) *Emission estimates*. CO₂ reductions are model-based and depend on the VSP proxy of Eq. (2), the synthetic HGV augmentation, and full advisory compliance in Eq. (8); reaction delays and traffic feedback are not modeled.

VI. CONCLUSION

This paper introduced *Agentic-FL*, a decentralized framework that elevates UAV edge aggregators from passive communication relays to active orchestrators of federated learning in mobile vehicular networks. On real NGSIM US-101 trajectories, the Silhouette Score shows a weaker association with FL convergence, whereas CDS exhibits a stronger association across the five evaluated paradigms. CDS therefore emerges as a promising temporal indicator for adaptive cluster management, subject to broader validation. The results also expose a trade-off between broad and targeted intervention: Spectral clustering achieves the largest modeled segment-wide emission reduction, while BIRCH delivers localized gains by addressing a small subset of high-emission vehicles, both quantified under the emission model of Section III. These results suggest that adaptive UAV orchestration can complement communication-centric FL design with behavior-aware partitioning. Future work will replace the rule-based protocol with a Multi-Agent Deep Reinforcement Learning (MADRL) strategy for joint trajectory planning and algorithm selection, add a compliance-aware advisory model, and validate CDS across multiple seeds and datasets under a held-out evaluation protocol.

REFERENCES

- [1] S. Zeadally, R. Hunt, Y.-S. Chen, A. Irwin, and A. Hassan, "Vehicular ad hoc networks (VANETs): Status, results, and challenges," *Telecommunication Systems*, vol. 50, no. 4, pp. 217–241, 2012.
- [2] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2017, pp. 1273–1282.
- [3] M. S. HaghighiFard and S. Coleri, "Hierarchical federated learning in multi-hop cluster-based VANETs," *IEEE Transactions on Vehicular Technology*, vol. 74, no. 10, pp. 15 371–15 385, 2025.
- [4] A. Antonopoulos, E. Kartsakli, N. Bartzoudis, D. Brodimas, D. Vukobratovic, D. Tsolkas, F. Diego, G. Roumelas, and L. Tomaszewski, "AGILE-6G: Agentic AI for autonomous management of 6G network/application services," *IEEE Network*, 2025.
- [5] S. Pandya *et al.*, "Federated learning for smart cities: A comprehensive survey," *Sustainable Energy Technologies and Assessments*, vol. 55, p. 102987, 2023.
- [6] T. Nguyen, T. Do, G. Baudet, F. Alboaie, A. Menesatti, and P. Pace, "Federated learning for vehicular networks: a comprehensive survey," *IEEE Access*, vol. 12, pp. 4512–4538, 2024.
- [7] Y. Zeng, R. Zhang, and T. J. Lim, "Wireless communications with unmanned aerial vehicles: Opportunities and challenges," *IEEE Communications Magazine*, vol. 54, no. 5, pp. 36–42, 2016.
- [8] S. Yuan, C. Yang, J. Zhou, and S. Xie, "Hierarchical aggregation and cooperative caching for decentralized federated learning in UAV-assisted internet of vehicles," *IEEE Internet of Things Journal*, vol. 13, no. 1, pp. 1270–1282, 2026.
- [9] H. Zhang and L. Hanzo, "Federated learning assisted multi-UAV networks," *IEEE Transactions on Vehicular Technology*, vol. 69, no. 11, pp. 14 104–14 109, 2020.
- [10] Q.-V. Pham, M. Zeng, R. Ruby, T. Huynh-The, and W.-J. Hwang, "UAV communications for sustainable federated learning," *IEEE Transactions on Vehicular Technology*, vol. 70, no. 4, pp. 3944–3948, 2021.
- [11] L. Liu, J. Zhang, S. H. Song, and K. B. Letaief, "Client-edge-cloud hierarchical federated learning," in *Proceedings of the IEEE International Conference on Communications (ICC)*, 2020, pp. 1–6.
- [12] L. Barbieri, S. Kianoush, M. Nicoli, L. Serio, and S. Savazzi, "A close look at the communication efficiency and the energy footprints of robust federated learning in industrial IoT," *IEEE Internet of Things Journal*, vol. 12, no. 11, pp. 15 130–15 150, 2025.
- [13] C. Boucetta, F. Lassabe, and O. Baala, "Enhancing data collection in vehicular network through clustering optimization," in *Proceedings of the International Wireless Communications and Mobile Computing Conference (IWCMC)*, 2023, pp. 1070–1075.
- [14] C. Boucetta, O. Baala, K. Ait Ali, and A. Caminada, "Performance of topology-based data routing with regard to radio connectivity in VANET," in *Proceedings of the International Wireless Communications and Mobile Computing Conference (IWCMC)*, 2019, pp. 609–614.
- [15] H. Rakha, K. Ahn, and K. Moran, "INTEGRATION framework for modeling eco-routing strategies: Logic and preliminary results," *International Journal of Transportation Science and Technology*, vol. 1, no. 3, pp. 259–274, 2012.
- [16] U.S. Federal Highway Administration, "Next generation simulation (NGSIM) vehicle trajectories and supporting data," <https://data.transportation.gov/d/8ect-6jqj>, 2006, accessed: 2026.