A Three-Layer Game-Theoretic Framework for Soil-Moisture Clustered Federated Learning

Kinda Khawam
ROCS, LISN
Paris-Saclay University
Gif-sur-Yvette, France
kinda.khawam@universite-paris-saclay.fr

Samer Lahoud
Dalhousie University
Halifax, NS, Canada
sml@dal.ca

Ralph Touma
ROCS, LISN
Paris-Saclay University
Gif-sur-Yvette, France
ralph.touma@universite-paris-saclay.fr

Steven Martin
ROCS, LISN
Paris-Saclay University
Gif-sur-Yvette, France
steven.martin@lisn.fr

***Abstract*—Soil-moisture sensors individually collect little data, which makes collaborative learning attractive, but sensors operating under different soil textures and irrigation practices generate heterogeneous data distributions, so indiscriminate federation degrades accuracy. We propose a three-layer clustered federated-learning framework that separates cluster formation, per-cluster learning, and online admission of new or returning sensors. A closed-form prediction-error criterion estimated from local observations captures the trade-off between variance reduction through pooling and cross-sensor heterogeneity. This criterion drives both a cooperative merge-split mechanism and a symmetric hedonic clustering game, while online admission is formulated as an exact potential game with finite best-response convergence. Experiments on synthetic soil-moisture data show gains in the scarce-data regime and adaptive cluster granularity as noise and heterogeneity vary.**

***Index Terms*—Clustered Federated Learning; Game Theory.**

I. Introduction

Federated Learning (FL) enables collaborative training without sharing raw measurements [1], while Clustered FL (CFL) mitigates non-IID effects by grouping compatible clients [2]–[4]. In distributed sensing, collaboration must also accommodate short local histories, heterogeneous dynamics, and changes in active participation. Soil-moisture monitoring is a representative setting: measurements at different locations can follow distinct local dynamics, while a sensor may temporarily stop participating because of duty cycling, connectivity loss, maintenance, or failure/recovery.

We separate three functions across different time scales: (L1) formation of statistically coherent coalitions, (L2) learning within each fixed coalition, and (L3) online admission of new or returning sensors without full re-clustering. Building on the model-sharing analysis of [5], all decisions are driven by the same closed-form Mean Square Error (MSE) decomposition. Our contributions are: i) MSE-driven cooperative and decentralized coalition formation; ii) an exact-potential admission game with finite convergence; and iii) a controlled evaluation of the variance-heterogeneity trade-off in the scarce-data regime.

II. Related Work

CFL methods differ mainly in how client groups are identified. CFL [2] recursively separates clients from similarities between local updates; IFCA [3] alternates client-to-cluster assignment and model optimization; PACFL [4] uses subspace similarity to identify compatible data distributions. Our setting uses instead an explicit prediction-error decomposition to define coalition preferences. This makes cluster formation directly interpretable through the variance-heterogeneity trade-off and, importantly, allows the same statistics to be reused when sensors subsequently become active or inactive.

A recent survey [6] reviews the use of game theory in FL for incentive mechanisms, privacy management, and security. However, existing approaches primarily focus on participation incentives and resource management rather than the formation and maintenance of statistically coherent clusters. In contrast, our work leverages game theory to address dynamic admission control in clustered federated learning while preserving cluster coherence through prediction-aware decisions.

III. Prediction Criterion

Sensor j provides a soil-moisture series $\{M_{j,t}\}_{t=1}^{N_j}$. We fit a first-order local surrogate

$$\begin{aligned} &x_{j,t} = M_{j,t-1}, \qquad y_{j,t} = M_{j,t}, \\ &y_{j,t} = \theta_j x_{j,t} + \eta_{j,t}, \qquad \eta_{j,t} \sim \mathcal{N}(0, \varepsilon_j^2). \end{aligned} \tag{1}$$

where θ_j summarizes local persistence and ε_j^2 residual variability. Sensor j estimates θ_j by scalar OLS, $\widehat{\theta}_j = (\sum_t x_{j,t} y_{j,t})/(\sum_t x_{j,t}^2)$, and estimates ε_j^2 from its regression residuals.

For a coalition C with equal local history length n, uniform aggregation gives $\widehat{\theta}_C = |C|^{-1} \sum_{j \in C} \widehat{\theta}_j$. After local normalization, the expected prediction error is [5]

$$\mathbb{E}[\mathrm{MSE}_C] = \frac{\mu_e^C}{|C| n} + \frac{|C| - 1}{|C|} \sigma_C^2, \tag{2}$$

with $\mu_e^C = |C|^{-1} \sum_{j \in C} \varepsilon_j^2$ and $\sigma_C^2 = |C|^{-1} \sum_{j \in C} (\theta_j - \bar{\theta}_C)^2$. The first term captures the variance reduction due to pooling, whereas the second penalizes collaboration among sensors with dissimilar fitted dynamics. In practice, $\widehat{\theta}_j$ and $\widehat{\varepsilon}_j^2$ are extracted from each sensor trajectory. We use

$$\widehat{\varepsilon}_j^2 = \frac{1}{N_j - 1} \sum_{t=1}^{N_j} (y_{j,t} - x_{j,t} \widehat{\theta}_j)^2, \tag{3}$$

and maintain, for each coalition,

$$\widehat{\mu}_e^C = \frac{1}{|C|} \sum_{j \in C} \widehat{\varepsilon}_j^2, \qquad \widehat{\sigma}_C^2 = \frac{1}{|C|} \sum_{j \in C} (\widehat{\theta}_j - \widehat{\theta}_C)^2. \tag{4}$$

979-8-3195-0167-7/26/$31.00 ©2026 IEEE

Table I
ROLES OF THE THREE LAYERS.

Layer	Time scale	Main decision
L1	slow	initial coalition partition
L2	fast	model within each fixed coalition
L3	online	admission of new/returning sensors

Algorithm 1: L1 merge-split structure discovery

Initialize $\Pi \leftarrow \{\{1\}, \ldots, \{M\}\}$;
repeat
  update cluster statistics; propose candidate merges/splits; accept a move iff it decreases (6).
until *no accepted move*;

These data-derived statistics are used by both L1 and L3.

Each sensor j evaluates any coalition $C_i \ni j$ through the federated prediction error

$$c_j(C_i) = \mathbb{E}\big[\mathrm{MSE}_{C_i}\big], \tag{5}$$

given in (2). This quantity depends only on the coalition *size* $|C_i|$, its *aggregate noise level* $\mu_e^{C_i}$, and its *within-coalition dispersion* $\sigma^2_{C_i}$; it is invariant to the identity or ordering of the members. Thus, all members of a coalition incur exactly the same cost and the induced coalition-formation game is a *hedonic game* in the sense of [8]: preferences over coalitions depend only on the set of members, not on their labels.

IV. Three-Layer Game-Theoretic Framework

The architecture separates slow structural decisions from learning and online adaptation, as summarized in Table I.

L1: coalition formation. The initial partition Π is sought by minimizing $\min_\Pi \sum_{C_i \in \Pi} \sum_{j \in C_i} c_j(C_i)$ which boils down to

$$\min_{\Pi} \sum_{C \in \Pi} |C|\, \mathbb{E}[\mathrm{MSE}_C]. \tag{6}$$

Problem (6) is a partitioning problem over a finite set of sensors. Its exact search space grows with the Bell number and becomes rapidly impractical as M increases, which motivates local coalition updates rather than exhaustive enumeration. A cooperative merge-split procedure starts from singleton coalitions and accepts only transformations that reduce (6). Since the set of partitions is finite and every accepted move improves the objective, the procedure terminates [7].

For decentralized formation, define for $l \neq k$

$$e_{lk} = \frac{\varepsilon_l^2 + \varepsilon_k^2}{2n}, \qquad h_{lk} = \frac{(\theta_l - \theta_k)^2}{2}. \tag{7}$$

Using the pairwise variance identity, the coalition cost can be expressed in terms of symmetric pairwise noise and heterogeneity contributions. With the bounded approximation $(|C| - 1)/|C| \simeq 1$, the cost of sensor j becomes:

$$\widetilde{c}_j(C_i) = \frac{1}{|C_i|^2} \sum_{k \in C_i \setminus \{j\}} v_{jk}, \qquad v_{jk} = e_{jk} + h_{jk} = v_{kj}. \tag{8}$$

The induced hedonic game is symmetric and additively separable; hence a Nash-stable partition exists, and strict best-response dynamics terminate at a stable partition [8].

Algorithm 2: L1 hedonic best response

Initialize a partition Π;
repeat
  for each sensor i, evaluate all current coalitions and the singleton option using (8); move i only to a strict best response.
until *no sensor moves*;

L2: learning. Once Π is fixed, each coalition aggregates its local estimates into one model. In the scalar fixed-data setting, this is a single aggregation step. Thus L2 is separated from the slower structural decisions rather than introduced as a new FL optimizer.

L3: online admission. Let $\Pi^\star$ be the partition returned by L1. A new or returning sensor u evaluates candidate clusters through the marginal change of the same MSE criterion. Let $m_C = |C|$, $S_C = \sum_{j \in C} \varepsilon_j^2$, and $\bar{\theta}_C = m_C^{-1} \sum_{j \in C} \theta_j$. Adding u gives the rank-one variance update

$$\sigma^2_{C \cup \{u\}} = \frac{m_C}{m_C + 1} \sigma_C^2 + \frac{m_C}{(m_C + 1)^2} (\theta_u - \bar{\theta}_C)^2. \tag{9}$$

The corresponding marginal terms are

$$\Delta_u^{\mathrm{var}}(C) = \frac{1}{n} \left(\frac{S_C + \varepsilon_u^2}{(m_C + 1)^2} - \frac{S_C}{m_C^2} \right), \tag{10}$$

$$\Delta_u^{\mathrm{het}}(C) = \frac{m_C}{m_C + 1} \sigma^2_{C \cup \{u\}} - \frac{m_C - 1}{m_C} \sigma_C^2, \tag{11}$$

and the admission cost is $\Delta_u(C) = \Delta_u^{\mathrm{var}}(C) + \lambda \Delta_u^{\mathrm{het}}(C)$, with $\lambda \geq 0$. Equivalently, for active assignment profile a, define

$$\Phi(a) = \sum_{C \in \Pi^\star} \left[\frac{S_C(a)}{m_C(a)^2 n} + \lambda \frac{m_C(a) - 1}{m_C(a)} \sigma_C^2(a) \right]. \tag{12}$$

Charging sensor u its marginal contribution $c_u(a) = \Phi(a) - \Phi(a^{(-u)})$ yields $c_u(a') - c_u(a) = \Phi(a') - \Phi(a)$ for any unilateral deviation. Therefore L3 is an exact potential game and best-response admission converges in finitely many steps to a pure Nash equilibrium [9].

V. Numerical Results

We use a synthetic soil-moisture dataset with $M = 80$ heterogeneous sensor trajectories. The local surrogate is fitted to each trace, yielding the data-derived estimates $\widehat{\theta}_j$ and $\widehat{\varepsilon}_j^2$. These statistics are inserted into the closed-form MSE criterion used by L1 and L3, while predictive performance is evaluated independently on holdout observations through RMSE. Therefore, the study tests whether decisions guided by the analytical variance-heterogeneity criterion translate into better out-of-sample prediction. The revealed local history varies from 12 to 28 samples, and cold-start results are averaged over 20 runs with 95% confidence intervals.

Figure 1 shows that Clustered FL achieves the lowest RMSE among the tested references. The gain is largest with little local data: pooling reduces estimation variance while clustering avoids the persistent heterogeneity penalty of a single global model. Global FL changes only slightly with additional local samples because all sensors remain constrained to a common

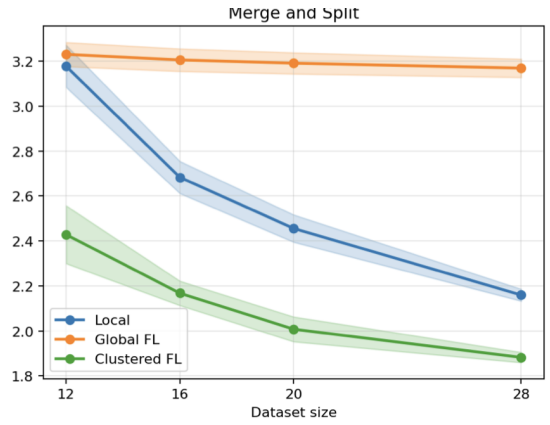


Figure 1. Cold-start holdout RMSE for merge-split clustering. Hedonic best response yields an almost identical RMSE trend.

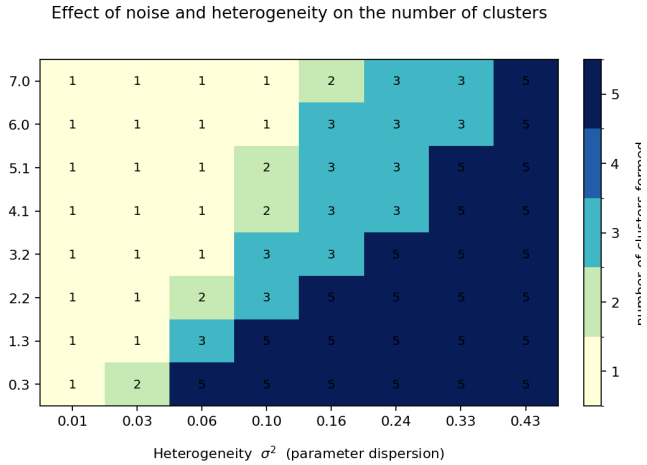


Figure 2. Number of clusters formed by L1 versus residual noise and fitted-parameter dispersion.

parameter, whereas Local learning improves steadily as estimation variance decreases. The decentralized hedonic partition gives nearly indistinguishable predictive performance, indicating that the two L1 solvers identify coalitions of comparable prediction quality despite their different update rules.

Figure 2 directly illustrates (2): higher residual noise makes pooling more valuable and leads to fewer, larger coalitions, whereas larger parameter dispersion strengthens the heterogeneity penalty and produces finer partitions. The staircase-like transition is consistent with the discrete nature of coalition formation: at low dispersion one large coalition is often sufficient, while increasing dispersion progressively favors additional splits. The number of clusters is therefore not fixed in advance but adapts to the estimated operating regime.

Figure 3 shows the different convergence behavior of the two L1 solvers: hedonic best response reaches a stable partition in only a few updates, while merge-split continues improving the global objective over a longer trajectory. Despite this difference, both yield essentially the same predictive RMSE in the cold-start experiments.

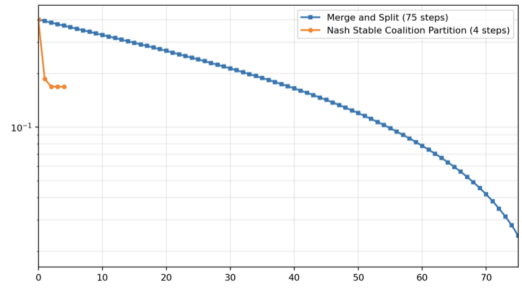


Figure 3. Clustering objective versus update steps for merge-split and hedonic best response.

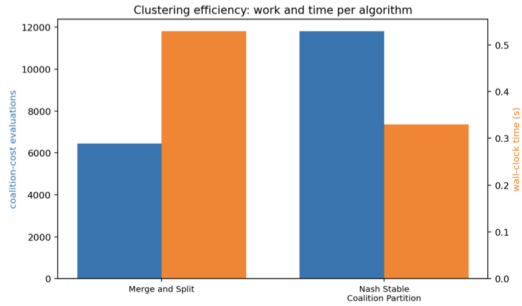


Figure 4. Existing L1 computational measurements: coalition-cost evaluations and wall-clock time.

Table II
SOLVER COMPARISON AT 28 REVEALED SAMPLES.

Solver	Steps	Final obj.	Assign.	RMSE
Hedonic BR	4	.02456	1.000	1.821
Merge-split	75	.02447	1.000	1.821

The existing computational measurements in Fig. 4 show a complementary trade-off: merge-split uses fewer coalition-cost evaluations, whereas hedonic best response completes faster in wall-clock time for the tested setting. At 28 revealed samples, the two solvers also reach the same clustered RMSE and assignment score despite very different update counts (Table II). This supports using merge-split as a cooperative reference and hedonic best response as the decentralized alternative.

Figure 5 provides an additional fixed-partition reference. At the shortest history the predefined labels remain competitive, but the data-driven partition improves as local parameter estimates become more reliable. The comparison is used only to show that prediction-oriented grouping need not coincide with a predefined category partition.

Figure 6 shows that prediction-aware clustered assignment remains beneficial for newly admitted sensors. In the same admission runs, the potential decreased monotonically and convergence required 3, 2, 2, and 1 sweeps at 12, 16, 20, and 28 revealed samples, respectively. This supports L3 as a lightweight online mechanism when local statistics become sufficiently reliable. At 28 samples, hedonic best response and merge-split achieve the same clustered RMSE (1.821) and assignment score (1.0), while requiring 4 and 75 update steps,

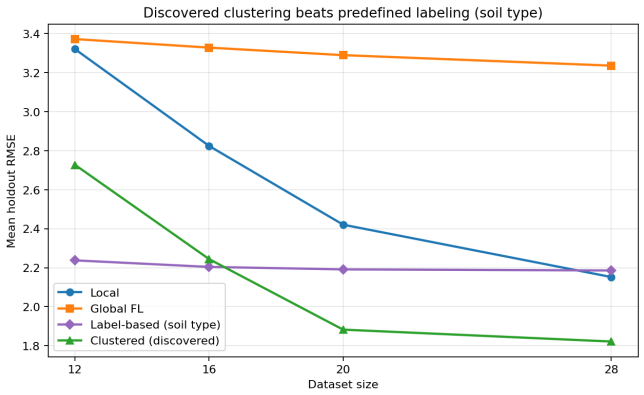


Figure 5. MSE-driven clustering versus a fixed partition based on predefined agronomic labels.

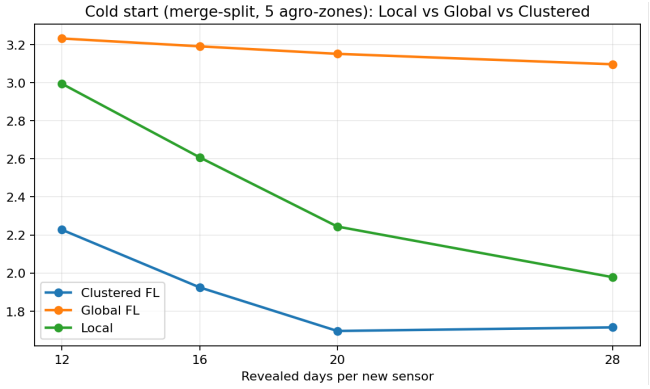


Figure 6. Holdout RMSE as progressively more local history is revealed for newly admitted sensors.

Table III
ADMISSION CONVERGENCE VERSUS REVEALED LOCAL HISTORY.

Samples	12	16	20	28
Sweeps	3	2	2	1
Φ_{start}	.341	.213	.134	.076
Φ_{final}	.322	.212	.134	.076

respectively.

The physical sensing deployment need not change whenever L3 is invoked. A soil probe can remain installed while its participation in FL changes because of duty cycling, temporary connectivity loss, maintenance, failure/recovery, or replacement. Conversely, a returning sensor may have only a short recent history from which to estimate its local statistics. In these cases, L3 decides whether the currently active sensor can be assigned to an existing learning coalition using the same MSE criterion as L1, avoiding a complete re-clustering after every participation event.

VI. Conclusion

We proposed a three-layer framework that uses one analytical prediction-error criterion to couple initial coalition formation and online admission in clustered FL. L1 combines cooperative merge-split with decentralized hedonic best response, while L3 is an exact potential game with finite convergence. The results show the benefit of selective collaboration in the scarce-data regime and the expected adaptation of cluster granularity to the noise-heterogeneity balance. Future work will consider richer predictors and field measurements.

Appendix: Pairwise Form of the L1 Cost

We detail the pairwise representation underlying (8). For a coalition C, the first term of (2), $\mu_e^C/(|C|n)$, captures the noise-variance reduction obtained by pooling samples. Since

$$\mu_e^C = \frac{1}{|C|}\sum_{j\in C}\varepsilon_j^2,$$

define, for any two sensors $l, k \in C$,

$$e_{lk} = \frac{1}{2n}(\varepsilon_l^2 + \varepsilon_k^2), \qquad e_{lk} = e_{kl}.$$

Averaging over all unordered pairs gives

$$\frac{1}{\binom{|C|}{2}}\sum_{l<k,\, l,k\in C} e_{lk} = \frac{1}{n|C|}\sum_{j\in C}\varepsilon_j^2 = \frac{\mu_e^C}{n}.$$

Hence

$$\frac{\mu_e^C}{|C|n} = \frac{1}{|C|}\,\frac{1}{\binom{|C|}{2}}\sum_{l<k,\, l,k\in C} e_{lk}. \tag{13}$$

This provides a symmetric pairwise representation of the noise contribution.

For the heterogeneity term, the classical variance identity yields

$$\sigma_C^2 = \frac{1}{2|C|^2}\sum_{l,k\in C}(\theta_l - \theta_k)^2. \tag{14}$$

With $h_{jk} = \frac{1}{2}(\theta_j - \theta_k)^2$, the double sum can equivalently be written through player-wise pairwise interactions. Substituting (13) and (14) into (2) gives

$$\widetilde{c}(C) = \frac{1}{|C|^2}\left[\frac{1}{n}\sum_{r\in C}\varepsilon_r^2 + \frac{|C|-1}{2|C|}\sum_{l,k\in C}(\theta_l - \theta_k)^2\right]. \tag{15}$$

Using the bounded approximation $(|C|-1)/|C| \simeq 1$ adopted for the equilibrium analysis, both contributions reduce to symmetric pairwise valuations $v_{jk} = e_{jk} + h_{jk}$. The resulting additively separable preference is therefore

$$\widetilde{c}_j(C) = \frac{1}{|C|^2}\sum_{k\in C\setminus\{j\}} v_{jk},$$

as used in (8). This pairwise structure is the basis for the Nash-stability result invoked by the decentralized L1 dynamics.

References

[1] P. Kairouz *et al.*, "Advances and open problems in federated learning," *Found. Trends Mach. Learn.*, 2021.
[2] F. Sattler, K.-R. Müller, and W. Samek, "Clustered federated learning," in *ICASSP*, 2020.
[3] A. Ghosh, J. Chung, D. Yin, and K. Ramchandran, "An efficient framework for clustered federated learning," in *NeurIPS*, 2020.
[4] S. Vahidian *et al.*, "Efficient distribution similarity identification in clustered federated learning via principal angles," in *AAAI*, 2023.
[5] K. Donahue and J. M. Kleinberg, "Model-sharing games: Analyzing federated learning under voluntary participation," in *AAAI*, 2021.
[6] R. Gupta and J. Gupta, "Federated learning using game strategies: State-of-the-art and future trends," *Computer Networks*, vol. 225, 2023.
[7] K. R. Apt and T. Radzik, "Stable partitions in coalitional games," *Artif. Intell.*, 2009.
[8] A. Bogomolnaia and M. O. Jackson, "The stability of hedonic coalition structures," *Games Econ. Behav.*, 2002.
[9] D. Monderer and L. S. Shapley, "Potential games," *Games Econ. Behav.*, 1996.