RELATIONALLY GUIDED HISTORY TRANSFER FOR STREAMING ENSEMBLE CLUSTERING

Anonymous authors
Paper under double-blind review

ABSTRACT

Streaming ensemble clustering seeks to maintain a consensus from base partitions as objects enter and leave successive windows. Existing static ensemble clustering methods, when applied independently to each window, discard historical structure carried by persistent objects. However, directly preserving the preceding consensus can retain outdated structure and provides no historical guidance for newly arrived objects. We propose AGILE, a streaming ensemble clustering framework that uses current base partitions to extend history beyond the window overlap and determine how strongly it should influence the new consensus. Pairwise intersections of base partitions define sparse relations through which historical information reaches arriving objects. To account for redundancy and disagreement among base partitions, AGILE draws complementary evidence from partitions corroborated by their peers and encodes relational corrections with relative-support-weighted signed factors. The corrected current structure supports a history-free consensus against which historical compatibility is assessed. An agreement–coverage gate then regulates the propagated history according to its consistency with the current consensus and cluster representation on the overlap. We establish monotone objective convergence, stationarity of fixed points, and linear time complexity in the window size. Experiments under different window overlap rates demonstrate the effectiveness of AGILE. The source code is available in the supplementary material.

1 INTRODUCTION

Ensemble clustering combines multiple base partitions into a consensus that captures their shared and complementary structure (Xu et al., 2025; Qi et al., 2026; Li et al., 2026). In streaming settings, this consensus must evolve as objects enter and leave the active window. Recent sliding-window clustering methods address efficient updates as data expire (Cohen-Addad et al., 2025; Feng et al., 2024a; Giannoulidis et al., 2025; Woodruff et al., 2023). Streaming ensemble clustering additionally requires integrating current base partitions with historical information from persistent objects shared by consecutive windows. This overlap offers an opportunity to retain useful structure, but its relevance must be reassessed as the population changes.

Existing static ensemble clustering methods improve consensus quality by refining relations among objects or integrating complementary base partitions (Zhou et al., 2024b; Li et al., 2025; Xu et al., 2025; 2026a). However, applying them independently to successive windows leaves the correspondence between persistent objects unused. Continual clustering methods use knowledge from previous data to guide clustering as new data arrive (Zhang et al., 2024; Wang et al., 2025; Cai et al., 2023; Arzeno & Vikalo, 2021). These methods carry historical knowledge through feature-based anchors or learned representations. In streaming ensemble clustering, history is available as consensus assignments on persistent objects. Extending this information to arriving objects therefore requires relations inferred from the current base partitions.

The window overlap makes the role of history asymmetric. Persistent objects have preceding consensus assignments, whereas newly arrived objects have none. Preserving these assignments can stabilize the clustering of persistent objects, but does not directly convey historical evidence to newcomers. Moreover, a grouping supported by the preceding window may no longer fit the current population. Discarding history loses potentially useful information; enforcing it without

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

checking current evidence can hinder adaptation. The central question is therefore *how to extend useful historical structure beyond the overlap while allowing the current consensus to change*.

Current base partitions connect persistent and newly arrived objects through their co-assignment relations. For example, an arriving object that repeatedly shares clusters with persistent objects can receive information about their preceding grouping. The quality of this connection matters, however. Several similar base partitions may repeat the same erroneous relation, so frequent agreement need not reflect independent support. Relation refinement has consequently become an important component of ensemble clustering (Zhou et al., 2024b; Xu et al., 2026b). In the streaming setting, unreliable relations can additionally carry historical errors into newly arrived objects. At the same time, a partition that disagrees with the majority may contain either complementary structure or noise. The challenge is therefore to combine repeated and potentially complementary co-assignment structure while controlling the influence of partial history.

We propose AGILE, which establishes current structure before determining the influence of history. It first calibrates the base partitions to a common scale and constructs sparse collision cells from their pairwise intersections. These cells express joint co-assignment support without forming a dense object–object matrix. To complement repeatedly observed relations, AGILE selects low-centrality partitions as candidate witnesses and weights their evidence by agreement with nearby partitions. Supported assignment corrections are encoded as signed relations, with relative support scores controlling their per-object weights. The resulting representation combines the original collision relations with the corrections and provides a current-only partition for evaluating historical compatibility.

Historical information is propagated through the original collision cells from persistent objects to arriving objects that share those cells. An agreement–coverage gate then determines its influence by comparing the preceding consensus with the current-only partition on the overlap and accounting for how many clusters are represented there. A robust consensus objective combines base-partition alignment, corrected current relations, and transported history. Sparse factor operations keep the computation linear in the window size for fixed ensemble dimensions and iteration budgets. The main contributions of this paper are:

- We develop a streaming ensemble clustering framework that uses current base partitions and persistent identities to extend partial history to arriving objects, with its contribution scaled by overlap agreement and coverage.
- We combine corroborated relation repair with collision-based history transport. Relative-support-weighted corrections contribute to a history-free current initialization, which provides the reference for an agreement–coverage gate.
- We establish monotone objective convergence, stationarity of fixed points, and linear time complexity in the window size for fixed ensemble dimensions and iteration budgets. Experiments under different window overlap rates demonstrate the effectiveness of AGILE.

2 Related Work

Temporal and continual clustering. Temporal clustering links clustering structures across successive time steps, typically balancing adaptation to current observations against consistency with preceding solutions (Chan et al., 2025; Epasto et al., 2022; Zhao et al., 2023; Yu et al., 2025). A related line of continual multi-view clustering studies how previously acquired knowledge can be retained as new views or observations become available (Wan et al., 2024; Zhang et al., 2024). Recent methods further develop adaptive knowledge preservation, cross-stage alignment, and incremental representation updating to mitigate forgetting during sequential learning (Wan et al., 2025; Wang et al., 2025; Feng et al., 2025).

Ensemble clustering. Ensemble clustering derives a consensus partition from multiple categorical base partitions, aiming to exploit their diversity while reducing unreliable decisions. Tensor and multi-level formulations model complementary dependencies among ensemble members (Jia et al., 2021; Zhou et al., 2021), whereas diffusion-based and partial-ensemble methods improve information propagation under heterogeneous or incomplete ensembles (Zhou et al., 2024b;a). Recent approaches further incorporate higher-order relations, graph-based structures, sparse consensus representations,

and local structural consistency to improve robustness and scalability (Xu et al., 2024b; Zhang et al., 2025c; Wei et al., 2025).

Incomplete multi-view clustering. Incomplete multi-view clustering seeks a consistent partition when some observations are unavailable (Yin et al., 2025; Dai et al., 2025; Jiang et al., 2025). Existing methods recover shared structure through prototype learning, similarity completion, and latent representation reconstruction (Yu et al., 2024; Feng et al., 2024b; Xu et al., 2024a). Other studies address correspondence uncertainty and unreliable view information through structural alignment, adaptive weighting, and relation correction (Liu et al., 2024; Yang et al., 2025; Zhou et al., 2026). Scalable formulations further employ robust consensus modeling and anchor-based representations to accommodate large incomplete data (Peng et al., 2026; Qin & Feng, 2026).

Our setting differs in both input and correspondence: each window provides base partitions, and adjacent windows are linked by persistent object identities. The preceding consensus is observed only on the overlap. AGILE uses current co-assignment relations to extend this information to arriving objects and calibrates its influence against a current-only partition.

3 AGILE: Current-First Relational Transfer

3.1 Problem formulation

Consider a stream processed through overlapping windows $\mathcal{W}_t$, with object identities $\mathcal{I}_t = \{i_1, \ldots, i_{n_t}\}$, where $n_t = |\mathcal{I}_t|$ is the number of objects in the current window. Stable identities determine the overlap $\mathcal{I}_t \cap \mathcal{I}_{t-1}$; each window requires a consensus before the next is processed. At window t, the method receives M base partitions $\Pi_t = \{\pi_t^{(m)}\}_{m=1}^{M}$ defined on all objects in $\mathcal{I}_t$. Partition m contains $c_{t,m}$ base clusters. The original features are not part of the consensus input. For $t > 1$, the immediately preceding consensus is available only through persistent identities. The desired output is a K_t-way consensus over all objects in the current window. For notational simplicity, we suppress the current window index t in the remainder of the method description, writing $n = n_t$, $c_m = c_{t,m}$, and $K = K_t$.

For a streaming window t, let $J = [j_{im}] \in \mathbb{N}^{n \times M}$ encode the current base partitions, where $j_{im} \in \{1, \ldots, c_m\}$ is the cluster assignment of object i in partition m. Let $S_i = 1$ if current object i also appears in window $t-1$, and $S_i = 0$ otherwise. When $S_i = 1$, $A_i \in \{1, \ldots, K_{\text{prev}}\}$ denotes its consensus assignment in the preceding window, where K_{prev} is the number of clusters in window $t-1$. Entries of A outside the overlap are ignored. The input at window t is therefore $(J, S, A, K, K_{\text{prev}})$, and no previous feature matrix or previous base-partition pool is required.

3.2 Robust consensus from base partitions

We formulate the consensus objective by first modeling alignment with the current base partitions, then incorporating current relations and historical information. Since base partitions may differ in both cluster size and cluster count, their indicator matrices require a common scale. Let $n_{mj} = |\{i : j_{im} = j\}|$ be the size of nonempty cluster j in partition m. We define $X_m \in \mathbb{R}^{n \times c_m}$ by $(X_m)_{ij} = \sqrt{K/c_m}/\sqrt{n_{mj}}$ when $j = j_{im}$ and zero otherwise. Thus $X_m^\top X_m = (K/c_m)I$ and $\|X_m\|_F^2 = K$, where I is the identity matrix and $\|\cdot\|_F$ denotes the Frobenius norm. Every partition has the same total energy.

Let $U \in \mathbb{R}^{n \times K}$ be the consensus embedding, with $U^\top U = I$. To compare partitions with different cluster labels and counts, we introduce an alignment matrix $R_m \in \mathbb{R}^{c_m \times K}$, constrained by $R_m^\top R_m = I$ if $c_m \geq K$ and $R_m R_m^\top = I$ otherwise. The disagreement between partition m and the consensus is measured by:

$$\delta_m = \frac{\|U - X_m R_m\|_F^2}{d_m}, \qquad d_m = K + \frac{K}{c_m}\min(c_m, K), \tag{1}$$

where d_m is the sum of the squared norms of U and $X_m R_m$, making the residual comparable across partitions. Minimizing the average residual treats all partitions equally, although some may agree poorly with the shared structure. We instead use the robust alignment score $\mathcal{A}(U, R) = -M^{-1}\sum_m \log(\varepsilon + \delta_m)$, where $R = \{R_m\}_{m=1}^{M}$ and $\varepsilon > 0$ is a numerical floor. Maximizing this

score gives smaller marginal influence to larger residuals. Equivalently,

$$\mathcal{A}(U, R) = \max_{w_m > 0} \frac{1}{M} \sum_{m=1}^{M} \big[\log w_m - w_m(\varepsilon + \delta_m) + 1 \big],$$

with $w_m = (\varepsilon + \delta_m)^{-1}$. These weights adapt direct partition alignment; the relations introduced next provide additional guidance about which objects should be grouped together.

3.3 Complementing consensus with current relations

Direct alignment compares each partition with the consensus separately. To capture co-assignments supported by pairs of partitions, we group objects by their joint labels. For an unordered pair $e = (m, r)$, let $g_e(i) = (j_{im}, j_{ir})$ identify the intersection cell containing object i. There are $L = M(M-1)/2$ partition pairs. For each pair, a sparse matrix C_e has one column per nonempty cell and one entry $\sqrt{K/(nL)}$ per row, at that object's cell. Concatenating these matrices gives:

$$B_{\text{col}} = [C_1, \ldots, C_L], \qquad (B_{\text{col}} B_{\text{col}}^\top)_{ij} = \frac{K}{nL} \sum_{m<r} \mathbf{1}[j_{im} = j_{jm}] \mathbf{1}[j_{ir} = j_{jr}], \tag{2}$$

where $\mathbf{1}[\cdot]$ is the indicator function. A relation contributes only when both partitions place the objects together. This collision factor has $\|B_{\text{col}}\|_F^2 = K$ and avoids materializing the dense object–object matrix.

Repeated agreement can nevertheless reflect similar partitions rather than complementary support. We therefore seek additional guidance from partitions with low overall agreement, while assessing whether nearby peers corroborate them. Let $a_{mr} = \text{ARI}(\pi^{(m)}, \pi^{(r)})$, where ARI is the adjusted Rand index. The q partitions with smallest raw centrality $b_m = (M-1)^{-1} \sum_{r \neq m} a_{mr}$ form the witness set $\mathcal{E}$. For each witness, let $\mathcal{N}_\nu(m)$ contain its ν most-agreeing other partitions. Its support is $\eta_m = \nu^{-1} \sum_{r \in \mathcal{N}_\nu(m)} [a_{mr}]_+$, where $[x]_+ = \max(x, 0)$, and its weight is $\omega_m = \eta_m / \sum_{\ell \in \mathcal{E}} \eta_\ell$. We use uniform weights if all selected supports vanish.

The witnesses guide reassignment through their weighted co-assignment affinity. For a candidate partition z, let $C_c(z) = \{j : z_j = c\}$ denote cluster c. We define:

$$W_{ij}^\omega = \sum_{m \in \mathcal{E}} \omega_m \mathbf{1}[j_{im} = j_{jm}], \qquad \ell_{ic}(z) = \frac{\sum_{j \in C_c(z),\, j \neq i} W_{ij}^\omega}{|C_c(z)| - \mathbf{1}[z_i = c]}, \tag{3}$$

where ℓ_{ic} measures the mean affinity of object i to cluster c and is defined as zero when the denominator is zero. The candidate quality is $\Phi(z) = n^{-1} \sum_i \ell_{i,z_i}(z)$. Each base partition provides a start: excess clusters are greedily merged by mean witness affinity, and starts with fewer than K clusters are skipped. Objects propose moves to their highest-scoring cluster only if their scores improve. A synchronous sweep preserves nonempty clusters and is accepted only if Φ does not decrease beyond numerical tolerance; otherwise the starting candidate is retained. We select z^* with the largest Φ across the retained candidates. Affinities are evaluated through contingency tables.

Assignments in z^* need not be equally supported. Let $u_i = \ell_{i,z_i^*}(z^*)$ and $v_i = \max_{c \neq z_i^*} \ell_{ic}(z^*)$ denote support for the assigned cluster and its strongest alternative. Their relative support determines the object's repair weight:

$$h_i = \frac{u_i + \epsilon_c}{u_i + v_i + 2\epsilon_c}, \qquad \bar{h} = \frac{1}{n} \sum_i h_i, \tag{4}$$

where $\bar{h}$ is the mean relative support and $\epsilon_c > 0$ is a numerical floor. To encourage agreement within repaired clusters and separation across them, we construct $B_{\text{rep}} \in \mathbb{R}^{n \times K(K-1)/2}$ with one column per unordered cluster pair (a, b). Its entries are $+s\sqrt{h_i/\bar{h}}$ for $z_i^* = a$, $-s\sqrt{h_i/\bar{h}}$ for $z_i^* = b$, and zero otherwise, where $s = \sqrt{K/[n(K-1)]}$. Consequently,

$$(B_{\text{rep}} B_{\text{rep}}^\top)_{ij} = \frac{\sqrt{h_i h_j}}{\bar{h}} \begin{cases} K/n, & z_i^* = z_j^*, \\ -K/[n(K-1)], & z_i^* \neq z_j^*, \end{cases} \qquad \|B_{\text{rep}}\|_F^2 = K. \tag{5}$$

The signs express these two preferences, while $h_i/\bar{h}$ redistributes the fixed repair strength across objects. We combine $B_{\text{str}} = [B_{\text{col}}, \sqrt{\beta} B_{\text{rep}}]$, where $\beta \geq 0$ controls repair relative to collision support. The resulting current-structure score is $\|B_{\text{str}}^\top U\|_F^2 / K$. Maximizing it favors embedding directions supported by these relations. The repair block is omitted if no valid candidate exists.

3.4 Extending and assessing historical information

Current relations also connect arriving objects to persistent objects whose preceding assignments are known. Let $Y \in \{0,1\}^{n \times K_{\text{prev}}}$ encode this partial history, with $Y_{ih} = 1$ exactly when $S_i = 1$ and $A_i = h$, and zero otherwise. Using Y alone leaves newcomers without historical guidance. We therefore propagate its assignments through the original collision cells. For $C_e(i) = \{j : g_e(j) = g_e(i)\}$, define:

$$(T_{\text{prop}})_{ih} = \frac{1}{L} \sum_{e=1}^{L} \frac{\sum_{j \in C_e(i)} Y_{jh}}{|C_e(i)|}, \qquad T = Y + T_{\text{prop}}, \tag{6}$$

where T_{prop} is propagated predecessor mass and T also retains the observed assignments. Dividing by the full cell size reduces contributions from cells with little historical support. Let $\bar{T}$ contain the nonzero columns of T, each normalized to unit Euclidean norm. With overlap fraction ρ_S and fraction ρ_C of objects receiving nonzero transported mass, set $\kappa = \rho_S \rho_C$ and $B_{\text{time}} = \sqrt{K\kappa}\,\bar{T}/\|\bar{T}\|_F$. Thus $\|B_{\text{time}}\|_F^2 = K\kappa$; the factor is zero if no nonzero column exists.

Transport makes history available, but does not establish its compatibility with the current window. We assess this against a current-only partition $z^{(0)}$, constructed from the current base partitions and corrected relations without historical information.

Agreement on an overlap representing only a few clusters gives limited evidence about the whole window. We therefore combine overlap agreement with cluster coverage:

$$g = \text{clip}\Big(\text{ARI}(A_S, z_S^{(0)}), 0, 1\Big) \sqrt{\frac{u_{\text{prev}}}{K_{\text{prev}}} \frac{u_{\text{cur}}}{K}}, \tag{7}$$

where g is the temporal compatibility gate, A_S and $z_S^{(0)}$ restrict the two partitions to persistent objects, u_{prev} and u_{cur} count their represented clusters, and $\text{clip}(x, 0, 1) = \min\{1, \max\{0, x\}\}$. ARI compares groupings without matching cluster labels; coverage reduces historical influence when few previous or current clusters appear in the overlap. We set $g = 0$ if fewer than two overlap objects exist or the temporal factor is inactive. Otherwise, g is fixed before optimization.

3.5 Integrating current structure and history

The alignment score preserves agreement with individual base partitions, while current relations provide joint grouping preferences. Transported history supplies a third source, weighted by its compatibility with the current-only partition. Combining these contributions yields:

$$\max_{\substack{U, \{R_m\}, \\ w_m > 0}} \mathcal{J}(U, R, w) = \frac{1}{M} \sum_{m=1}^{M} [\log w_m - w_m(\varepsilon + \delta_m) + 1] + \frac{\alpha}{K} \|B_{\text{str}}^\top U\|_F^2 + \frac{g}{K} \|B_{\text{time}}^\top U\|_F^2$$
$$\text{s.t.} \quad U^\top U = I, \qquad \begin{cases} R_m^\top R_m = I, & c_m \geq K, \\ R_m R_m^\top = I, & c_m < K, \end{cases} \quad m = 1, \ldots, M, \tag{8}$$

where $\alpha \geq 0$ controls the current-structure contribution, and $w = (w_m)_{m=1}^M$ collects the robust alignment weights. The structural contribution equals $\alpha(\|B_{\text{col}}^\top U\|_F^2 + \beta \|B_{\text{rep}}^\top U\|_F^2)/K$, separating collision support from repair. The temporal term favors embedding directions supported by shared historical assignments, without requiring $K = K_{\text{prev}}$. Its scale depends on both κ and g. At the first window or with inactive history, this term vanishes, leaving a current-only consensus. Overall, AGILE uses current ensemble structure to guide consensus learning, extend historical information to arriving objects, and assess its compatibility with the current window. Algorithm 1 summarizes the overall procedure.

3.6 Optimization

With the relation factors and gate fixed, we alternately update R, w, and U. At iteration r, rectangular Procrustes gives each alignment $R_m^{(r+1)}$, followed by $w_m^{(r+1)} = [\varepsilon + \delta_m(U^{(r)}, R_m^{(r+1)})]^{-1}$. The two relational terms have positive-semidefinite Gram matrices, allowing a tight affine lower bound for the

consensus update. Maximizing this bound reduces to a trace objective with coefficient matrix:

$$Y^{(r)} = \sum_m \frac{w_m^{(r+1)}}{Md_m} X_m R_m^{(r+1)} + \frac{\alpha}{K} B_{\mathrm{str}}(B_{\mathrm{str}}^\top U^{(r)}) + \frac{g}{K} B_{\mathrm{time}}(B_{\mathrm{time}}^\top U^{(r)}), \tag{9}$$

where $Y^{(r)} \in \mathbb{R}^{n\times K}$ is an update matrix, distinct from the historical indicator Y. For its thin singular value decomposition $Y^{(r)} = P\Sigma Q^\top$, P and Q contain singular vectors and Σ contains singular values; the update is $U^{(r+1)} = PQ^\top$. Appendix B.2 provides the derivations. After convergence, k-means on the row-normalized embedding produces the consensus partition.

4 Theoretical Properties and Complexity

The direct, collision, and active repair factors each have squared Frobenius norm K, while $\|B_{\mathrm{time}}\|_F^2 = K\kappa$. This common scale separates source strength from representation size. Each accepted repair also preserves its witness-agreement criterion up to numerical tolerance. Formal statements and proofs appear in Appendix D.

Theorem 1 (Monotone convergence in objective value)**.** *Fix the relational factors and gate, with $\alpha, \beta \geq 0$ and $\varepsilon > 0$. Each completed $R \to w \to U$ update cycle satisfies*

$$\mathcal{J}^{(r)} \leq \mathcal{J}^{(r+1)} \leq -\log\varepsilon + \alpha(1+\beta) + g\kappa.$$

Hence the objective-value sequence converges.

The alignment and weight updates solve their blocks exactly; the consensus update maximizes a tight lower bound. Every fixed point satisfies the first-order stationarity conditions (Corollary 1).

Time complexity. Let $L = M(M-1)/2$ and $q = O(\sqrt{M})$. Under bounded cluster counts, collision construction costs $O(nL)$, witness scoring over all starts costs $O(nMqK)$, and temporal transport costs $O(nLK_{\mathrm{prev}})$. Each optimization cycle costs $O(n(MK + LK + K^2 + KK_{\mathrm{prev}}))$ plus small alignment decompositions. Sparse initialization and final rounding are also sample-linear for fixed iteration budgets. The total time is therefore $O(n)$ for fixed ensemble dimensions, cluster counts, and iteration budgets. Appendix B.3 gives the full accounting. **Space complexity.** Storing the sparse relation factors, temporal information, consensus embedding, and partition-agreement matrix requires $O(n(M + L + K + K_{\mathrm{prev}}) + M^2)$ space. This is also linear in n for fixed ensemble dimensions and cluster counts.

5 Experiments

5.1 Experimental Setup

Datasets and Stream Construction. We evaluate AGILE on eight benchmark datasets. Their statistics are summarized in Table 1. Each dataset is organized into $T = 10$ active windows. We vary the adjacent-window overlap rate as $\rho \in \{0.2, 0.5, 0.8\}$, where $\rho = |\mathcal{I}_t \cap \mathcal{I}_{t-1}|/|\mathcal{I}_t|$ denotes the fraction of current objects that persist from the preceding window. Thus, the three settings expose the methods to sparse, moderate, and dense temporal correspondence. Persistent identities are used only to connect the current window to the immediately preceding consensus. Following the common ensemble-construction protocol (Xu et al., 2024b), we run k-means 100 times in each window to form a candidate pool of base partitions. In each trial, $M = 20$ base partitions are sampled from the pool using a fixed random seed. For every trial and window, all competing methods receive the same 20 base partitions, thereby excluding performance differences caused by ensemble sampling. **Comparison Methods.** We compare AGILE with fourteen representative clustering methods, including SEC (Liu et al., 2015), PTGP (Huang et al., 2016), LWGP

Table 1: Characteristics of the datasets.

ID	Dataset	#Samples	#Features	#Classes
D1	COIL100	7200	1024	100
D2	Nursery	12960	8	5
D3	TwentyNewsgroups	18846	5000	20
D4	NATICUSdroid	29333	86	2
D5	FashionMNIST	70000	784	10
D6	KannadaMNIST	70000	784	10
D7	MNIST	70000	784	10
D8	MHEALTH	343195	23	12

Table 2: ACC (%) across different overlap rates ρ, with the best results highlighted in bold.

Method	COIL100			Nursery			TwentyNewsgroups			NATICUSdroid		
	$\rho = 0.2$	$\rho = 0.5$	$\rho = 0.8$	$\rho = 0.2$	$\rho = 0.5$	$\rho = 0.8$	$\rho = 0.2$	$\rho = 0.5$	$\rho = 0.8$	$\rho = 0.2$	$\rho = 0.5$	$\rho = 0.8$
SEC [*KDD*, 2015]	45.73 ± 0.92	45.07 ± 0.24	44.68 ± 0.58	33.50 ± 0.98	34.38 ± 1.40	32.12 ± 1.06	11.65 ± 0.50	14.44 ± 0.44	19.08 ± 0.67	58.98 ± 4.88	61.34 ± 3.28	57.24 ± 2.47
PTGP [*TKDE*, 2016]	42.46 ± 0.38	40.29 ± 0.35	36.86 ± 0.45	33.69 ± 1.19	34.10 ± 1.53	32.04 ± 1.20	11.80 ± 0.42	14.29 ± 0.44	18.46 ± 0.59	53.34 ± 0.15	53.49 ± 0.07	53.58 ± 0.09
LWGP [*TCYB*, 2018]	46.63 ± 0.28	45.79 ± 0.51	43.76 ± 0.32	33.43 ± 0.90	33.61 ± 1.64	32.48 ± 1.06	6.55 ± 0.04	6.35 ± 0.05	6.19 ± 0.22	53.44 ± 0.11	53.54 ± 0.00	53.63 ± 0.05
USENC [*TKDE*, 2020]	53.36 ± 0.35	52.39 ± 0.17	50.81 ± 0.37	33.28 ± 1.08	33.80 ± 1.89	32.04 ± 1.30	12.75 ± 0.53	15.72 ± 0.55	20.95 ± 0.81	53.49 ± 0.01	53.54 ± 0.00	53.94 ± 0.88
ECPCS-MC [*TSMC*, 2021]	48.15 ± 0.29	47.48 ± 0.33	46.51 ± 0.25	32.52 ± 1.07	33.49 ± 2.08	31.14 ± 1.08	6.43 ± 0.03	6.16 ± 0.03	6.28 ± 0.31	53.28 ± 0.22	53.32 ± 0.09	53.41 ± 0.16
DenForest [*SIGMOD*, 2022]	12.07 ± 2.04	20.48 ± 3.67	20.70 ± 0.85	30.66 ± 3.77	28.84 ± 2.13	27.93 ± 3.49	6.41 ± 0.06	6.26 ± 0.10	6.17 ± 0.12	53.01 ± 0.82	53.46 ± 1.01	53.28 ± 1.00
FastMICE [*TKDE*, 2023]	53.20 ± 0.46	51.97 ± 0.44	50.03 ± 0.45	33.44 ± 1.42	34.15 ± 1.83	32.13 ± 1.20	13.06 ± 0.46	16.21 ± 0.40	21.14 ± 0.92	53.47 ± 0.06	53.54 ± 0.00	53.66 ± 0.02
PSIMVC-PG [*TNNLS*, 2024]	54.80 ± 0.41	53.54 ± 0.40	52.42 ± 0.47	32.49 ± 0.77	32.57 ± 0.96	31.38 ± 0.94	11.80 ± 0.25	13.66 ± 0.46	17.22 ± 0.93	53.29 ± 0.28	53.88 ± 1.09	53.51 ± 0.13
FCMVC-IV [*TIP*, 2024]	49.42 ± 0.64	48.54 ± 0.38	46.52 ± 0.73	33.64 ± 2.17	33.51 ± 2.81	31.88 ± 1.37	10.43 ± 0.44	12.12 ± 0.64	14.66 ± 0.95	56.15 ± 2.38	53.76 ± 1.81	53.99 ± 1.94
CEHM [*AAAI*, 2025]	47.13 ± 0.46	46.14 ± 0.44	44.25 ± 0.47	34.99 ± 0.27	$\mathbf{34.74 \pm 0.48}$	$\mathbf{34.41 \pm 0.37}$	6.03 ± 0.00	5.86 ± 0.00	5.64 ± 0.00	52.71 ± 0.24	52.77 ± 0.14	52.98 ± 0.22
FSEC [*Inf. Fusion*, 2025]	53.62 ± 0.34	52.59 ± 0.27	51.04 ± 0.44	33.43 ± 0.99	33.85 ± 2.12	31.81 ± 1.20	12.85 ± 0.52	16.10 ± 0.46	21.20 ± 0.70	53.49 ± 0.01	53.54 ± 0.00	53.65 ± 0.01
MVC-IIV [*TIP*, 2025]	53.47 ± 0.30	52.86 ± 0.17	51.68 ± 0.57	31.98 ± 1.23	31.67 ± 1.31	30.30 ± 0.85	12.43 ± 0.42	15.39 ± 0.43	20.20 ± 0.46	60.68 ± 4.83	60.21 ± 4.30	60.50 ± 3.73
RANGE [*CVPR*, 2026]	56.32 ± 0.38	56.16 ± 0.39	54.90 ± 0.19	32.82 ± 1.46	32.58 ± 1.19	30.54 ± 1.35	13.53 ± 0.38	16.57 ± 0.39	21.86 ± 0.77	50.68 ± 0.00	50.47 ± 0.00	50.33 ± 0.00
DREAM [*TKDE*, 2026]	52.28 ± 0.61	50.74 ± 0.35	49.16 ± 0.50	33.77 ± 0.09	33.74 ± 0.21	33.55 ± 0.14	6.03 ± 0.00	5.86 ± 0.00	5.93 ± 0.46	50.72 ± 0.09	50.49 ± 0.06	50.33 ± 0.01
AGILE (Ours)	$\mathbf{57.90 \pm 0.34}$	$\mathbf{56.65 \pm 0.41}$	$\mathbf{55.81 \pm 0.47}$	$\mathbf{35.05 \pm 2.25}$	33.89 ± 2.08	31.40 ± 0.98	$\mathbf{14.33 \pm 0.55}$	$\mathbf{17.82 \pm 0.55}$	$\mathbf{24.11 \pm 0.79}$	$\mathbf{63.19 \pm 2.54}$	$\mathbf{64.28 \pm 3.26}$	$\mathbf{61.31 \pm 1.83}$

Method	FashionMNIST			KannadaMNIST			MNIST			MHEALTH		
	$\rho = 0.2$	$\rho = 0.5$	$\rho = 0.8$	$\rho = 0.2$	$\rho = 0.5$	$\rho = 0.8$	$\rho = 0.2$	$\rho = 0.5$	$\rho = 0.8$	$\rho = 0.2$	$\rho = 0.5$	$\rho = 0.8$
SEC [*KDD*, 2015]	46.81 ± 1.57	46.85 ± 1.10	47.11 ± 1.46	51.25 ± 1.28	52.79 ± 2.31	53.96 ± 2.58	43.86 ± 1.48	44.62 ± 1.02	43.18 ± 1.58	37.19 ± 1.25	36.94 ± 1.12	36.60 ± 1.09
PTGP [*TKDE*, 2016]	47.04 ± 0.79	47.97 ± 0.67	47.17 ± 0.99	56.45 ± 1.43	58.09 ± 1.38	58.54 ± 1.34	47.70 ± 0.28	47.96 ± 0.15	48.04 ± 0.14	37.11 ± 0.66	36.87 ± 0.89	37.31 ± 0.74
LWGP [*TCYB*, 2018]	46.34 ± 0.47	46.67 ± 0.86	46.44 ± 0.87	50.79 ± 0.61	49.80 ± 1.12	49.49 ± 0.72	45.27 ± 0.51	45.52 ± 0.54	46.24 ± 0.31	38.79 ± 1.15	39.99 ± 0.60	40.65 ± 0.41
USENC [*TKDE*, 2020]	48.55 ± 0.82	49.00 ± 0.98	48.61 ± 0.88	57.22 ± 1.68	59.35 ± 1.48	59.76 ± 1.79	47.22 ± 0.36	47.72 ± 0.32	47.97 ± 0.30	37.72 ± 0.96	38.70 ± 1.06	38.09 ± 0.82
ECPCS-MC [*TSMC*, 2021]	46.19 ± 0.73	46.40 ± 0.72	46.18 ± 0.66	47.42 ± 1.87	46.19 ± 1.11	46.79 ± 1.71	38.95 ± 1.21	39.79 ± 2.18	39.39 ± 2.06	34.14 ± 0.52	34.56 ± 0.61	34.55 ± 0.70
DenForest [*SIGMOD*, 2022]	43.84 ± 0.47	43.44 ± 2.23	44.08 ± 0.75	37.62 ± 5.61	42.91 ± 1.24	36.28 ± 1.91	13.56 ± 0.96	12.32 ± 0.42	12.64 ± 1.32	22.25 ± 0.50	20.92 ± 0.72	20.10 ± 2.44
FastMICE [*TKDE*, 2023]	48.01 ± 0.64	48.51 ± 0.64	47.85 ± 0.95	57.78 ± 1.45	58.35 ± 1.54	60.07 ± 1.25	47.70 ± 0.28	47.04 ± 0.20	48.13 ± 0.17	38.25 ± 0.70	39.63 ± 1.01	39.10 ± 0.89
PSIMVC-PG [*TNNLS*, 2024]	47.55 ± 0.71	47.61 ± 0.48	47.72 ± 0.90	55.38 ± 1.53	56.77 ± 1.06	55.22 ± 1.32	46.99 ± 0.47	47.65 ± 0.26	47.64 ± 0.39	39.63 ± 0.72	39.58 ± 0.81	39.36 ± 0.66
FCMVC-IV [*TIP*, 2024]	48.81 ± 1.05	48.71 ± 1.06	48.72 ± 0.83	52.65 ± 1.07	52.02 ± 1.48	52.41 ± 1.24	46.18 ± 0.82	46.08 ± 0.52	46.35 ± 0.38	35.41 ± 0.81	35.81 ± 1.30	35.38 ± 1.34
CEHM [*AAAI*, 2025]	46.31 ± 0.41	47.07 ± 0.79	46.49 ± 0.81	45.62 ± 1.08	43.78 ± 0.85	44.92 ± 1.26	45.39 ± 0.60	45.93 ± 0.52	46.79 ± 0.63	34.25 ± 0.52	34.71 ± 0.55	34.72 ± 0.63
FSEC [*Inf. Fusion*, 2025]	47.97 ± 0.80	48.48 ± 0.66	48.07 ± 0.76	56.43 ± 1.21	58.88 ± 0.87	59.39 ± 0.84	47.79 ± 0.29	47.05 ± 0.21	48.14 ± 0.16	37.28 ± 0.54	38.01 ± 0.85	37.79 ± 0.63
MVC-IIV [*TIP*, 2025]	48.03 ± 0.75	49.13 ± 0.90	48.57 ± 0.78	58.91 ± 1.94	59.89 ± 1.53	60.05 ± 0.90	47.59 ± 0.29	47.98 ± 0.21	48.07 ± 0.27	39.45 ± 0.76	40.07 ± 0.50	39.89 ± 0.66
RANGE [*CVPR*, 2026]	50.00 ± 0.80	50.07 ± 0.77	50.50 ± 0.95	60.42 ± 1.18	$\mathbf{61.32 \pm 1.37}$	61.95 ± 1.21	47.67 ± 0.18	47.90 ± 0.24	47.93 ± 0.18	42.27 ± 0.62	41.61 ± 0.62	42.45 ± 0.50
DREAM [*TKDE*, 2026]	47.50 ± 1.05	48.02 ± 0.79	48.06 ± 1.17	54.24 ± 1.48	54.33 ± 1.51	54.34 ± 2.26	44.38 ± 0.41	45.20 ± 0.50	44.94 ± 0.49	34.30 ± 0.66	34.31 ± 0.96	34.49 ± 0.78
AGILE (Ours)	$\mathbf{52.07 \pm 0.68}$	$\mathbf{52.40 \pm 0.93}$	$\mathbf{52.80 \pm 0.91}$	$\mathbf{61.08 \pm 1.43}$	61.22 ± 0.84	$\mathbf{62.03 \pm 0.70}$	$\mathbf{48.10 \pm 0.23}$	$\mathbf{48.49 \pm 0.20}$	$\mathbf{48.77 \pm 0.32}$	$\mathbf{42.62 \pm 0.38}$	$\mathbf{42.67 \pm 0.51}$	$\mathbf{42.47 \pm 0.36}$

(Huang et al., 2018), USENC (Huang et al., 2020), ECPCS-MC (Huang et al., 2021), DenForest (Kim et al., 2022), FastMICE (Huang et al., 2023), PSIMVC-PG (Li et al., 2024), FCMVC-IV (Wan et al., 2024), CEHM (Li et al., 2025), FSEC (Zhang et al., 2025b), MVC-IIV (Zhang et al., 2025a), RANGE (Xu et al., 2026a), and DREAM (Xu et al., 2026c). The comparison covers classical and large-scale ensemble clustering, streaming or incremental clustering (DenForest, FCMVC-IV, and MVC-IIV), and incomplete multi-view clustering (PSIMVC-PG). Static ensemble methods are applied separately to each window. Methods requiring multi-view inputs are evaluated on representations derived from the same selected base partitions, without access to the original features. Their parameters follow the corresponding papers or released implementations. **Evaluation Metrics.** Clustering performance is evaluated using clustering accuracy (ACC), adjusted Rand index (ARI), and normalized mutual information (NMI), with higher values indicating better performance. For each trial, each metric is averaged over the ten window-level evaluations. We report the mean and standard deviation of these trial-level results over ten trials.

5.2 Comparison across overlap rates

Tables 2 and 3 present the ACC and ARI results, while the NMI results are reported in Appendix F.1. We observe: (i) AGILE achieves the best performance in 21/24 settings for ACC, 21/24 for ARI, and 19/24 for NMI, and obtains the best average rank under every metric and overlap rate. The advantage across metrics and overlap regimes indicates that the improvement is not tied to an evaluation criterion or degree of temporal correspondence. (ii) Compared with ensemble baselines, AGILE benefits from modeling relations beyond independent base-partition alignment. Pairwise collision structure captures jointly supported co-assignments, while corroborated repair supplies peer-supported relational corrections and robust weighting downweights poorly aligned partitions. These mechanisms can strengthen the current evidence used to form the consensus and assess historical compatibility. (iii) At $\rho = 0.2$, AGILE achieves the best ACC on all eight datasets and the best ARI on seven of them. Its advantage persists when a small fraction of objects carries predecessor assignments, showing that extensive temporal correspondence is not required for competitive performance in the evaluated settings. (iv) Increasing the overlap rate does not yield a uniform trend. As ρ increases from 0.2 to 0.8, AGILE's ACC rises from 14.33% to 24.11% on TwentyNewsgroups but decreases from 57.90% to 55.81% on COIL100. A larger overlap does not by itself ensure better clustering quality. Importantly, ρ measures only the amount of persistence, whereas the effective temporal contribution is further regulated by transported coverage κ and the agreement–coverage gate g; a larger overlap therefore does not necessarily imply stronger historical influence. This dataset-dependent behavior is consistent with the motivation for assessing historical information against current ensemble structure, while the leading average ranks across three overlap rates demonstrate the effectiveness of the framework across overlap conditions.

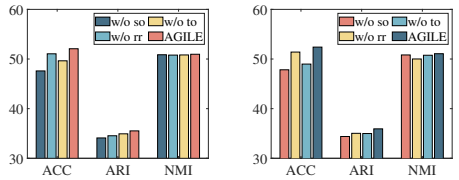
Table 3: ARI (%) across different overlap rates ρ, with the best results highlighted in bold.

Method	COIL100			Nursery			TwentyNewsgroups			NATICUSdroid		
	$\rho=0.2$	$\rho=0.5$	$\rho=0.8$	$\rho=0.2$	$\rho=0.5$	$\rho=0.8$	$\rho=0.2$	$\rho=0.5$	$\rho=0.8$	$\rho=0.2$	$\rho=0.5$	$\rho=0.8$
SEC [*KDD*, 2015]	27.08 ± 1.31	28.26 ± 0.79	32.16 ± 1.24	2.71 ± 1.05	4.60 ± 1.28	2.47 ± 0.95	1.27 ± 0.26	3.03 ± 0.30	6.93 ± 0.23	9.83 ± 6.31	12.81 ± 4.14	7.56 ± 2.75
PTGP [*TKDE*, 2016]	31.56 ± 0.50	31.37 ± 0.34	30.20 ± 0.45	2.88 ± 1.27	4.56 ± 1.98	1.94 ± 1.05	1.08 ± 0.19	2.46 ± 0.26	6.17 ± 0.45	0.45 ± 0.03	0.49 ± 0.02	0.51 ± 0.02
LWGP [*TCYB*, 2018]	33.84 ± 0.39	34.12 ± 0.62	34.70 ± 0.45	3.18 ± 1.09	4.77 ± 2.14	2.72 ± 1.26	0.00 ± 0.00	0.00 ± 0.00	0.10 ± 0.15	0.48 ± 0.02	0.50 ± 0.00	0.52 ± 0.01
USENC [*TKDE*, 2020]	43.45 ± 0.42	43.86 ± 0.24	44.06 ± 0.41	2.71 ± 0.89	4.42 ± 2.26	2.04 ± 1.25	1.32 ± 0.25	3.18 ± 0.32	7.24 ± 0.43	0.49 ± 0.00	0.50 ± 0.00	0.93 ± 1.24
ECPCS-MC [*TSMC*, 2021]	35.97 ± 0.28	36.93 ± 0.45	37.97 ± 0.28	3.16 ± 0.89	5.42 ± 2.52	2.86 ± 1.50	0.00 ± 0.00	0.00 ± 0.00	0.02 ± 0.04	0.44 ± 0.05	0.45 ± 0.02	0.47 ± 0.04
DenForest [*SIGMOD*, 2022]	0.77 ± 0.46	2.82 ± 1.45	2.15 ± 0.54	1.23 ± 0.66	3.43 ± 2.19	1.29 ± 0.72	0.00 ± 0.00	0.00 ± 0.00	0.01 ± 0.00	1.45 ± 1.26	1.75 ± 1.48	1.51 ± 1.31
FastMICE [*TKDE*, 2023]	43.22 ± 0.27	43.66 ± 0.44	43.71 ± 0.32	3.24 ± 1.54	5.13 ± 2.28	2.32 ± 1.19	1.47 ± 0.25	3.45 ± 0.32	7.55 ± 0.44	0.48 ± 0.01	0.50 ± 0.00	0.53 ± 0.01
PSIMVC-PG [*TNNLS*, 2024]	45.27 ± 0.53	45.59 ± 0.39	45.94 ± 0.32	3.38 ± 0.88	4.26 ± 1.20	3.03 ± 0.84	0.93 ± 0.17	2.38 ± 0.39	5.78 ± 0.58	0.45 ± 0.07	1.16 ± 1.45	0.50 ± 0.03
FCMVC-IV [*TIP*, 2024]	38.13 ± 0.62	38.82 ± 0.31	38.68 ± 0.48	5.28 ± 2.92	5.83 ± 3.90	3.53 ± 1.86	0.57 ± 0.20	1.73 ± 0.27	4.13 ± 0.47	4.44 ± 3.32	1.43 ± 2.28	1.78 ± 2.42
CEHM [*AAAI*, 2025]	32.96 ± 0.51	33.76 ± 0.59	34.08 ± 0.60	0.26 ± 0.05	0.36 ± 0.51	0.30 ± 0.18	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.33 ± 0.05	0.35 ± 0.03	0.38 ± 0.05
FSEC [*Inf. Fusion*, 2025]	43.76 ± 0.32	44.13 ± 0.30	44.10 ± 0.45	3.13 ± 1.28	5.01 ± 2.34	2.31 ± 1.35	1.35 ± 0.24	3.23 ± 0.32	7.24 ± 0.31	0.49 ± 0.00	0.50 ± 0.00	0.53 ± 0.00
MVC-IIV [*TIP*, 2025]	42.94 ± 0.36	43.89 ± 0.37	44.60 ± 0.42	2.61 ± 0.95	3.14 ± 1.29	1.90 ± 0.59	1.28 ± 0.21	3.02 ± 0.28	7.06 ± 0.20	9.36 ± 6.00	8.64 ± 5.16	8.95 ± 4.37
RANGE [*CVPR*, 2026]	47.22 ± 0.38	$\mathbf{48.59 \pm 0.32}$	48.80 ± 0.16	4.30 ± 2.34	4.32 ± 1.58	2.28 ± 1.50	1.32 ± 0.22	3.05 ± 0.27	7.15 ± 0.26	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00
DREAM [*TKDE*, 2026]	40.11 ± 0.50	40.09 ± 0.46	40.42 ± 0.68	0.00 ± 0.02	0.05 ± 0.10	0.02 ± 0.03	0.00 ± 0.00	0.00 ± 0.00	0.17 ± 0.28	0.00 ± 0.01	0.00 ± 0.01	0.00 ± 0.00
AGILE (Ours)	$\mathbf{48.16 \pm 0.44}$	48.38 ± 0.34	$\mathbf{48.95 \pm 0.30}$	$\mathbf{8.14 \pm 3.18}$	$\mathbf{6.47 \pm 2.63}$	$\mathbf{3.73 \pm 3.11}$	$\mathbf{1.56 \pm 0.19}$	$\mathbf{3.73 \pm 0.35}$	$\mathbf{9.00 \pm 0.62}$	$\mathbf{13.16 \pm 3.59}$	$\mathbf{14.28 \pm 4.32}$	$\mathbf{10.44 \pm 2.52}$

Method	FashionMNIST			KannadaMNIST			MNIST			MHEALTH		
	$\rho=0.2$	$\rho=0.5$	$\rho=0.8$	$\rho=0.2$	$\rho=0.5$	$\rho=0.8$	$\rho=0.2$	$\rho=0.5$	$\rho=0.8$	$\rho=0.2$	$\rho=0.5$	$\rho=0.8$
SEC [*KDD*, 2015]	30.73 ± 2.02	31.35 ± 1.07	31.54 ± 1.38	36.01 ± 1.71	37.29 ± 2.86	38.86 ± 2.87	26.57 ± 1.65	26.88 ± 0.88	25.93 ± 1.32	21.38 ± 1.19	21.17 ± 1.06	21.04 ± 1.00
PTGP [*TKDE*, 2016]	34.92 ± 0.62	35.54 ± 0.58	35.19 ± 0.64	42.89 ± 1.65	44.00 ± 1.51	45.19 ± 0.89	30.06 ± 0.29	30.13 ± 0.22	30.11 ± 0.21	22.66 ± 0.71	22.39 ± 0.85	22.66 ± 0.75
LWGP [*TCYB*, 2018]	33.70 ± 0.40	33.71 ± 0.47	33.82 ± 0.76	35.69 ± 1.08	34.66 ± 1.40	34.45 ± 0.60	27.86 ± 0.57	28.01 ± 0.44	28.41 ± 0.36	23.61 ± 1.39	25.02 ± 0.65	25.77 ± 0.45
USENC [*TKDE*, 2020]	34.91 ± 0.38	35.41 ± 0.86	35.11 ± 0.66	44.45 ± 1.66	45.53 ± 0.97	46.67 ± 1.53	29.79 ± 0.29	29.97 ± 0.35	30.08 ± 0.30	22.40 ± 1.03	23.37 ± 1.07	22.76 ± 0.82
ECPCS-MC [*TSMC*, 2021]	33.61 ± 0.41	33.66 ± 0.43	33.59 ± 0.37	31.15 ± 1.27	30.58 ± 1.00	31.26 ± 1.37	23.05 ± 1.16	24.47 ± 1.79	23.95 ± 2.32	19.11 ± 0.28	19.43 ± 0.28	19.40 ± 0.30
DenForest [*SIGMOD*, 2022]	31.85 ± 0.47	31.29 ± 3.31	32.44 ± 0.98	21.86 ± 7.09	28.10 ± 1.69	20.86 ± 1.88	0.48 ± 0.24	0.19 ± 0.11	0.28 ± 0.31	7.80 ± 0.36	5.55 ± 0.64	5.45 ± 2.64
FastMICE [*TKDE*, 2023]	34.61 ± 0.31	35.17 ± 0.48	34.89 ± 0.47	44.99 ± 1.17	44.64 ± 1.37	46.50 ± 1.16	30.08 ± 0.37	30.32 ± 0.28	30.28 ± 0.30	22.90 ± 0.75	24.28 ± 0.96	23.55 ± 0.81
PSIMVC-PG [*TNNLS*, 2024]	33.93 ± 0.42	34.34 ± 0.35	34.35 ± 0.55	41.64 ± 1.67	42.19 ± 1.20	40.93 ± 1.22	29.58 ± 0.32	29.98 ± 0.21	29.92 ± 0.41	24.63 ± 0.70	24.64 ± 0.82	24.38 ± 0.73
FCMVC-IV [*TIP*, 2024]	34.23 ± 0.51	34.25 ± 0.41	34.29 ± 0.52	38.94 ± 1.33	38.21 ± 1.34	38.35 ± 1.57	28.68 ± 0.53	28.86 ± 0.36	28.80 ± 0.39	20.75 ± 0.96	21.03 ± 1.10	20.59 ± 1.30
CEHM [*AAAI*, 2025]	33.11 ± 0.62	34.04 ± 0.49	33.73 ± 0.60	28.93 ± 1.63	26.99 ± 1.14	28.46 ± 1.24	28.09 ± 0.57	28.60 ± 0.62	29.22 ± 0.51	18.67 ± 0.32	19.13 ± 0.28	19.03 ± 0.28
FSEC [*Inf. Fusion*, 2025]	34.74 ± 0.33	35.17 ± 0.44	34.81 ± 0.48	43.76 ± 1.25	45.35 ± 1.00	46.27 ± 0.63	30.05 ± 0.36	30.29 ± 0.29	30.29 ± 0.19	21.91 ± 0.57	22.89 ± 0.95	22.47 ± 0.71
MVC-IIV [*TIP*, 2025]	34.32 ± 0.49	35.47 ± 0.72	35.16 ± 0.63	45.10 ± 1.76	45.88 ± 1.31	46.34 ± 0.71	29.35 ± 0.26	29.58 ± 0.24	29.50 ± 0.31	24.78 ± 0.58	25.35 ± 0.56	25.05 ± 0.51
RANGE [*CVPR*, 2026]	$\mathbf{35.64 \pm 0.69}$	35.03 ± 0.59	36.29 ± 0.73	45.77 ± 0.70	46.11 ± 1.20	$\mathbf{48.06 \pm 0.69}$	29.94 ± 0.24	29.96 ± 0.28	29.67 ± 0.36	27.34 ± 0.45	26.77 ± 0.36	27.57 ± 0.44
DREAM [*TKDE*, 2026]	33.86 ± 0.64	34.41 ± 0.51	34.62 ± 0.98	38.84 ± 1.92	38.59 ± 2.16	38.99 ± 2.72	27.15 ± 0.40	27.82 ± 0.63	27.54 ± 0.42	18.73 ± 0.55	18.73 ± 0.93	18.75 ± 0.76
AGILE (Ours)	35.52 ± 0.53	$\mathbf{35.93 \pm 0.59}$	$\mathbf{36.35 \pm 0.52}$	$\mathbf{46.62 \pm 0.78}$	$\mathbf{46.87 \pm 0.45}$	47.83 ± 0.38	$\mathbf{30.66 \pm 0.37}$	$\mathbf{31.07 \pm 0.17}$	$\mathbf{31.43 \pm 0.29}$	$\mathbf{27.77 \pm 0.30}$	$\mathbf{27.98 \pm 0.26}$	$\mathbf{27.74 \pm 0.29}$

5.3 Ablation Analysis

Figure 1 evaluates three variants of AGILE. *w/o so* denotes removing the current structural term $\frac{\alpha}{K}\|B_{\text{str}}^{\top}U\|_F^2$; *w/o rr* denotes removing the relation-repair term $\frac{\alpha\beta}{K}\|B_{\text{rep}}^{\top}U\|_F^2$; and *w/o to* denotes removing the temporal-transfer term $\frac{g}{K}\|B_{\text{time}}^{\top}U\|_F^2$. Removing the current structural term causes the largest performance degradation in the displayed settings, particularly in ACC and ARI, indicating its substantial contribution to clustering quality. Removing the relation-repair term also generally degrades performance, supporting its contribution alongside the original collision relations. Removing the temporal-transfer term produces smaller differences, consistent with a supplementary contribution from history in these settings. Overall, the displayed objective-term ablations support the contributions of current relational modeling, relation repair, and temporal transfer to clustering performance.



(a) FashionMNIST-0.2 (b) FashionMNIST-0.5

(c) KannadaMNIST-0.2 (d) KannadaMNIST-0.5

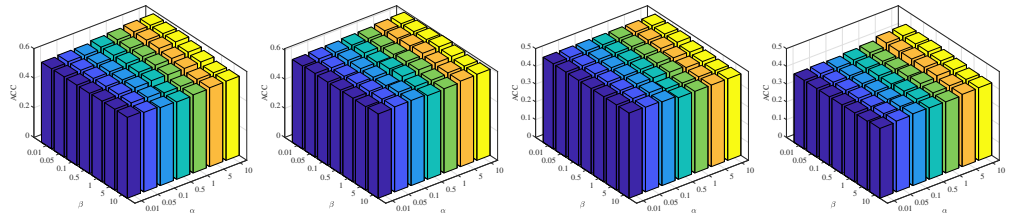
Figure 1: Ablation results under $\rho \in \{0.2, 0.5\}$.

5.4 Parameter Sensitivity

AGILE has only two searched hyperparameters, α and β, and we examine their influence on clustering performance. As shown in Figure 2, the ACC values remain relatively stable over a broad range of parameter combinations, indicating that the performance does not depend on a narrowly localized optimum. In the objective, α controls the contribution of the current relational structure, while β adjusts the relative contribution of relation repair. In practice, $\alpha, \beta \in \{0.5, 1, 5\}$ provide a stable and compact search range across the evaluated datasets. Results for the remaining datasets and overlap rates are reported in Appendix F.2.

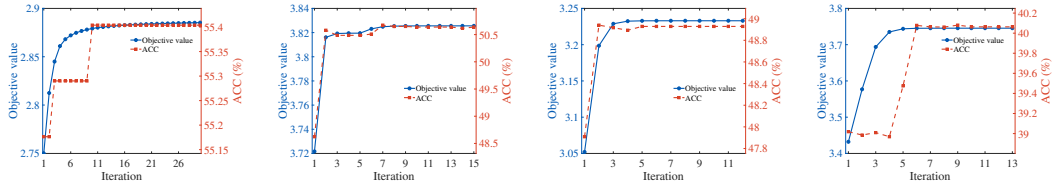
5.5 Convergence Analysis

Figure 3 illustrates the convergence behavior of AGILE by reporting the objective value and ACC over iterations. Across all four datasets, the objective value increases rapidly and converges within a few iterations, while the ACC shows an overall upward trend and gradually stabilizes as the



(a) COIL100 (b) KannadaMNIST (c) MNIST (d) MHEALTH

Figure 2: Sensitivity of ACC to α and β at $\rho = 0.2$. Results on the remaining datasets and overlap rates are reported in Appendix F.2.



(a) COIL100 (b) FashionMNIST (c) MNIST (d) MHEALTH

Figure 3: Convergence curves of AGILE over iterations on four representative datasets.

optimization proceeds. The objective curves are consistent with the monotone ascent property established in Section 4. Their stabilization within a limited number of iterations also supports the practicality of the iteration budget used in our experiments.

5.6 Computational Efficiency

Figure 4 compares the running time of AGILE with nine representative competing methods on the eight datasets. AGILE is not the fastest method because it additionally performs pair-collision construction, corroborated relation repair, and temporal transport. Nevertheless, these operations are implemented using sparse incidence factors and contingency statistics and therefore avoid constructing dense object–object matrices. With $M = 20$, each object contributes only $M(M-1)/2 = 190$ collision entries. Thus, the additional relational modeling increases the constant computational cost but does not change the sample-linear dependence on the window size established. The runtime results therefore reflect a trade-off between the additional relational processing and the consistent accuracy improvements reported above.

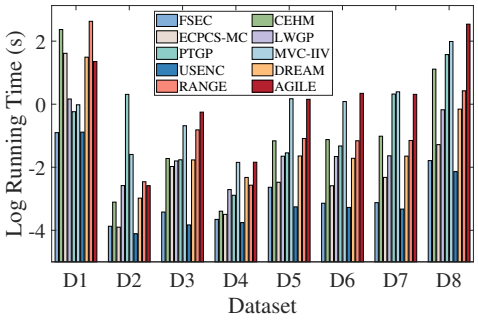


Figure 4: Running-time comparison of different methods.

6 Conclusion

We propose AGILE for streaming ensemble clustering over overlapping windows. Its central idea is to use current base partitions both to convey historical information to arriving objects and to construct a history-free current-only partition. Agreement with this reference and cluster coverage on the overlap determine the weight of the temporal contribution. We establish monotone objective convergence and stationarity of fixed points, with linear time and memory complexity in the window size for fixed ensemble dimensions, cluster counts, and iteration budgets. AGILE achieves the best average rank across all evaluated metrics and overlap rates. A limitation is that historical transfer depends on shared collision cells between persistent and arriving objects, which may restrict its benefit when the overlap is sparse.

AI USE STATEMENT

Generative AI tools were used only for language polishing and manuscript formatting checks. All AI-assisted edits were reviewed by the authors.

REFERENCES

Natalia M. Arzeno and Haris Vikalo. Evolutionary clustering via message passing. *IEEE Transactions on Knowledge and Data Engineering*, 33(6):2452–2466, 2021.

Hecheng Cai, Yuze Tan, Shudong Huang, and Jiancheng Lv. Lifelong multi-view spectral clustering. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pp. 3488–3496, 2023.

TH Chan, Shaofeng Jiang, Tianyi Wu, and Mengshi Zhao. Online clustering with nearly optimal consistency. In *International Conference on Learning Representations*, volume 2025, pp. 72115–72144, 2025.

Vincent Cohen-Addad, Shaofeng Jiang, Qiaoyuan Yang, Yubo Zhang, and Samson Zhou. Fair clustering in the sliding window model. In *International Conference on Learning Representations*, volume 2025, pp. 88088–88115, 2025.

Yuzhuo Dai, Jiaqi Jin, Zhibin Dong, Siwei Wang, Xinwang Liu, En Zhu, Xihong Yang, Xinbiao Gan, and Yu Feng. Imputation-free and alignment-free: Incomplete multi-view clustering driven by consensus semantic learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5071–5081. IEEE, 2025.

Alessandro Epasto, Mohammad Mahdian, Vahab Mirrokni, and Peilin Zhong. Improved sliding window algorithms for clustering and coverage via bucketing-based sketches. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*, pp. 3005–3042, 2022.

Chunhui Feng, Junhua Fang, Yue Xia, Pingfu Chao, Pengpeng Zhao, Jiajie Xu, and Xiaofang Zhou. Ocean: online clustering and evolution analysis for dynamic streaming data. In *Proceedings of the IEEE International Conference on Data Engineering*, pp. 4504–4517. IEEE, 2024a.

Wei Feng, Guoshuai Sheng, Qianqian Wang, Quanxue Gao, Zhiqiang Tao, and Bo Dong. Partial multi-view clustering via self-supervised network. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 2024b.

Yu Feng, Weixuan Liang, Xinhang Wan, Jiyuan Liu, Miaomiao Li, and Xinwang Liu. Incremental multi-view clustering: Exploring stream-view correlations to learn consistency and diversity. *IEEE Transactions on Knowledge and Data Engineering*, 37(12):7226–7239, 2025.

Apostolos Giannoulidis, Anastasios Gounaris, and John Paparrizos. Burst: Rendering clustering techniques suitable for evolving streams. *Proceedings of the VLDB Endowment*, 18(11):4054–4063, 2025.

Dong Huang, Jian-Huang Lai, and Chang-Dong Wang. Robust ensemble clustering using probability trajectories. *IEEE Transactions on Knowledge and Data Engineering*, 28(5):1312–1326, 2016.

Dong Huang, Chang-Dong Wang, and Jian-Huang Lai. Locally weighted ensemble clustering. *IEEE Transactions on Cybernetics*, 48(5):1460–1473, 2018.

Dong Huang, Chang-Dong Wang, Jian-Sheng Wu, Jian-Huang Lai, and Chee-Keong Kwoh. Ultra-scalable spectral clustering and ensemble clustering. *IEEE Transactions on Knowledge and Data Engineering*, 32(6):1212–1226, 2020.

Dong Huang, Chang-Dong Wang, Hongxing Peng, Jianhuang Lai, and Chee-Keong Kwoh. Enhanced ensemble clustering via fast propagation of cluster-wise similarities. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 51(1):508–520, 2021.

Dong Huang, Chang-Dong Wang, and Jian-Huang Lai. Fast multi-view clustering via ensembles: Towards scalability, superiority, and simplicity. *IEEE Transactions on Knowledge and Data Engineering*, 35(11):11388–11402, 2023.

Yuheng Jia, Hui Liu, Junhui Hou, and Qingfu Zhang. Clustering ensemble meets low-rank tensor approximation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 7970–7978, 2021.

Xiaorui Jiang, Buyun He, Peng Yuan Zhou, Xinyue Chen, Jingcai Guo, Jie Xu, and Yong Liao. A unified framework to BRIDGE complete and incomplete deep multi-view clustering under non-IID missing patterns. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 594–603, 2025.

Bogyeong Kim, Kyoseung Koo, Undraa Enkhbat, and Bongki Moon. Denforest: Enabling fast deletion in incremental density-based clustering over sliding windows. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*, pp. 296–309, 2022.

Feijiang Li, Jieting Wang, Liuya Zhang, Yuhua Qian, Shuai Jin, Tao Yan, and Liang Du. k-hyperedge medoids for clustering ensemble. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 18271–18278, 2025.

Feijiang Li, Jieting Wang, Yuhua Qian, Hongren Yan, and Tingjin Luo. How relation enrichment improves clustering ensemble performance: A second order induced relation view. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2026. doi: 10.1109/TPAMI.2026.3713520. Early Access.

Miaomiao Li, Siwei Wang, Xinwang Liu, and Suyuan Liu. Parameter-free and scalable incomplete multiview clustering with prototype graph. *IEEE Transactions on Neural Networks and Learning Systems*, 35(1):300–310, 2024.

Hongfu Liu, Tongliang Liu, Junjie Wu, Dacheng Tao, and Yun Fu. Spectral ensemble clustering. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 715–724, 2015.

Suyuan Liu, Siwei Wang, Ke Liang, Junpu Zhang, Zhibin Dong, Tianrui Liu, En Zhu, Kunlun He, and Xinwang Liu. Alleviate anchor-shift: Explore blind spots with cross-view reconstruction for incomplete multi-view clustering. In *Advances in Neural Information Processing Systems*, volume 37, 2024.

Shuangmei Peng, Zhe Chen, Tianyang Xu, and Xiao-Jun Wu. Relationship alignment for view-aware multi-view clustering. In *International Conference on Learning Representations*, 2026.

Te Qi, HaoYe Qiu, Yuheng Jia, Xiao Shi, and Ran Wang. Reliable co-occurrence guided partial clustering ensemble. *IEEE Transactions on Neural Networks and Learning Systems*, pp. 1–14, 2026.

Yalan Qin and Guorui Feng. Unified and efficient multi-view clustering from probabilistic perspective. In *International Conference on Learning Representations*, 2026.

Xinhang Wan, Bin Xiao, Xinwang Liu, Jiyuan Liu, Weixuan Liang, and En Zhu. Fast continual multi-view clustering with incomplete views. *IEEE Transactions on Image Processing*, 33:2995–3008, 2024.

Xinhang Wan, Jiyuan Liu, Hao Yu, Qian Qu, Ao Li, Xinwang Liu, Ke Liang, Zhibin Dong, and En Zhu. Contrastive continual multiview clustering with filtered structural fusion. *IEEE Transactions on Neural Networks and Learning Systems*, 36(8):14829–14842, 2025.

Jing Wang, Songhe Feng, Kristoffer Knutsen Wickstrøm, and Michael C. Kampffmeyer. AdaptCMVC: Robust adaption to incremental views in continual multi-view clustering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10285–10294, 2025.

Wei Wei, Jianguo Wu, Xinyao Guo, Jing Yan, and Jiye Liang. Self-constrained clustering ensemble. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(12):11946–11960, 2025.

David Woodruff, Peilin Zhong, and Samson Zhou. Near-optimal k-clustering in the sliding window model. In *Advances in Neural Information Processing Systems*, volume 36, pp. 22934–22960, 2023.

Gehui Xu, Jie Wen, Chengliang Liu, Bing Hu, Yicheng Liu, Lunke Fei, and Wei Wang. Deep variational incomplete multi-view clustering: Exploring shared clustering structures. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 16147–16155, 2024a.

Jiaxuan Xu, Taiyong Li, and Lei Duan. Enhancing ensemble clustering with adaptive high-order topological weights. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 16184–16192, 2024b.

Jiaxuan Xu, Lei Duan, Xinye Wang, and Liang Du. Dynamic anchor-based ensemble clustering via hypergraph reconstruction. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pp. 6740–6748, 2025.

Jiaxuan Xu, Lei Duan, Xinye Wang, and Liang Du. Large-scale robust enhanced ensemble clustering via outlier decoupling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 39691–39700, 2026a.

Jiaxuan Xu, Lei Duan, Xinye Wang, and Liang Du. Topological-aware multiple kernel learning for ensemble clustering. *Pattern Recognition*, pp. 113588, 2026b.

Jiaxuan Xu, Lei Duan, Xinye Wang, Liang Du, Yidan Zhang, and Zhen Guo. Dynamic anchor-based one-step hypergraph ensemble clustering. *IEEE Transactions on Knowledge and Data Engineering*, 38(9):5798–5810, 2026c.

Xihong Yang, Siwei Wang, Fangdi Wang, Jiaqi Jin, Suyuan Liu, Yue Liu, En Zhu, Xinwang Liu, and Yueming Jin. Automatically identify and rectify: Robust deep contrastive multi-view clustering in noisy scenarios. In *Proceedings of the International Conference on Machine Learning*, volume 267, pp. 71383–71404, 2025.

Jun Yin, Pei Wang, Shiliang Sun, and Zhonglong Zheng. Incomplete multi-view clustering via multi-level contrastive learning. *IEEE Transactions on Knowledge and Data Engineering*, 37(8): 4716–4727, 2025.

Shengju Yu, Zhibin Dong, Siwei Wang, Xinhang Wan, Yue Liu, Weixuan Liang, Pei Zhang, Wenxuan Tu, and Xinwang Liu. Towards resource-friendly, extensible and stable incomplete multi-view clustering. In *Proceedings of the International Conference on Machine Learning*, volume 235, pp. 57415–57440, 2024.

Shengju Yu, Pei Zhang, Siwei Wang, Suyuan Liu, Xinhang Wan, Zhibin Dong, Tiejun Li, and Xinwang Liu. Bit-swapping oriented twin-memory multi-view clustering in lifelong incomplete scenarios. In *Advances in Neural Information Processing Systems*, volume 38, 2025.

Chao Zhang, Deng Xu, Xiuyi Jia, Chunlin Chen, and Huaxiong Li. Continual multi-view clustering with consistent anchor guidance. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pp. 5434–5442, 2024.

Chao Zhang, Zhi Wang, Xiuyi Jia, Zechao Li, Chunlin Chen, and Huaxiong Li. Multi-view clustering with incremental instances and views. *IEEE Transactions on Image Processing*, 34:4203–4214, 2025a.

Runxin Zhang, Shuaijun Hang, Zhensheng Sun, Feiping Nie, Rong Wang, and Xuelong Li. Anchor-based fast spectral ensemble clustering. *Information Fusion*, 113:102587, 2025b.

Xu Zhang, Haoye Qiu, Weixuan Liang, Hui Liu, Junhui Hou, and Yuheng Jia. Generalization performance of ensemble clustering: From theory to algorithm. In *Proceedings of the International Conference on Machine Learning*, volume 267, pp. 75901–75927, 2025c.

Qi Zhao, Bai Yan, Jian Yang, and Yuhui Shi. Evolutionary robust clustering over time for temporal data. *IEEE Transactions on Cybernetics*, 53(7):4334–4346, 2023.

Haochen Zhou, Guofeng Ding, Mouxing Yang, Peng Hu, Yijie Lin, and Xi Peng. Uncover underlying correspondence for robust multi-view clustering. In *International Conference on Learning Representations*, 2026.

Peng Zhou, Liang Du, Yi-Dong Shen, and Xuejun Li. Tri-level robust clustering ensemble with multiple graph learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 11125–11133, 2021.

Peng Zhou, Liang Du, Xinwang Liu, Zhaolong Ling, Xia Ji, Xuejun Li, and Yi-Dong Shen. Partial clustering ensemble. *IEEE Transactions on Knowledge and Data Engineering*, 36(5):2096–2109, 2024a.

Peng Zhou, Boao Hu, Dengcheng Yan, and Liang Du. Clustering ensemble via diffusion on adaptive multiplex. *IEEE Transactions on Knowledge and Data Engineering*, 36(4):1463–1474, 2024b.

A Experimental Details

This appendix provides the experimental and implementation details that complement the main paper. Appendix B gives the complete optimization procedure and complexity accounting, Appendix C records additional properties of the streaming formulation, Appendix D provides the proofs, Appendix E summarizes the implementation choices, and Appendix F reports the remaining experimental results.

A.1 Stream and Ensemble Construction

The experiments use the eight datasets listed in Table 1. Each dataset is organized into ten active windows, with adjacent-window overlap rates $\rho \in \{0.2, 0.5, 0.8\}$. Object identities are retained only across adjacent windows. At the first window there is no preceding consensus; from the second window onward, only the immediately preceding consensus on persistent objects is available to the method. For every window, 100 k-means partitions form the candidate ensemble pool and $M = 20$ partitions are selected for each trial. The same selected partitions are supplied to all compared methods under a fixed random seed. Original features are not used by AGILE; methods that require vector or multi-view inputs are evaluated on representations constructed from the same selected base partitions, as described in Section 5.

A.2 Evaluation and Reproducibility

Each trial produces a clustering result for every window. ACC, ARI, and NMI are first averaged over the ten window-level evaluations, and the reported mean and standard deviation are then computed over ten trials. The selected base partitions and random seeds are shared across methods within the same trial. The two tuned hyperparameters are α and β. Both are evaluated on the fixed grid $\{0.01, 0.05, 0.1, 0.5, 1, 5, 10\}$. All experiments were performed on a workstation with an AMD EPYC 7763 64-Core Processor running at 2.45 GHz, two NVIDIA GeForce RTX 2080 Ti GPUs, and 128 GB of RAM under Windows 10 Professional. All methods were implemented in MATLAB R2024b, and the runtime results were obtained under the same hardware and software configuration. The experimental implementation is provided in the supplementary material.

B Optimization and Complexity

B.1 Complete procedure

Algorithm 1 One Streaming Window of AGILE

Input: Current base partitions J, overlap state (S, A), numbers of clusters (K, K_{prev}), and parameters α and β.
Output: Current-window consensus results $\widehat{\boldsymbol{z}}$.
Construct B_{col} and B_{rep}, and form $B_{\mathrm{str}} = [B_{\mathrm{col}}, \sqrt{\beta} B_{\mathrm{rep}}]$.
Compute the current initialization $U^{(0)}$.
Construct the temporal factor B_{time} and gate g.
Initialize $\{R_m\}_{m=1}^{M}$ and $\boldsymbol{w}$ from $U^{(0)}$.
while not converged **do**
    Update $\{R_m\}_{m=1}^{M}$ by Eq. (13).
    Update $\boldsymbol{w}$ by Eq. (15).
    Update U by Eq. (19).
end while
Row-normalize U and apply k-means to obtain the consensus results $\widehat{\mathbf{z}}$.

B.2 Closed-form Updates

We optimize R, w, and U alternately. At iteration r, the variables are updated in the order $R \to w \to U$.

Updating R. With $U^{(r)}$ and $w^{(r)}$ fixed, we update each R_m independently. The alignment matrices $\{R_m\}_{m=1}^M$ are independent of each other and can therefore be updated separately. Since $w_m^{(r)} > 0$, updating R_m is equivalent to minimizing δ_m. From $X_m^\top X_m = (K/c_m)I$ and the rectangular partial-isometry constraint on R_m, we have

$$\|X_m R_m\|_F^2 = \frac{K}{c_m} \min(c_m, K), \tag{10}$$

which is independent of R_m. Together with $U^\top U = I$, the normalized alignment distortion becomes

$$\delta_m(U, R_m) = 1 - \frac{2}{d_m} \operatorname{tr}(U^\top X_m R_m). \tag{11}$$

Thus, the R_m subproblem becomes

$$\max_{R_m} \operatorname{tr}\Big(R_m^\top X_m^\top U^{(r)}\Big), \quad \text{s.t. } \begin{cases} R_m^\top R_m = I, & c_m \geq K, \\ R_m R_m^\top = I, & c_m < K. \end{cases} \tag{12}$$

where c_m is the number of clusters in the m-th base partition and K is the number of consensus clusters. Let $X_m^\top U^{(r)} = P_m \Sigma_m Q_m^\top$. The closed-form solution is

$$R_m^{(r+1)} = P_m Q_m^\top. \tag{13}$$

Updating w. With $U^{(r)}$ and $R^{(r+1)}$ fixed, we update the partition weights w. The weights $\{w_m\}_{m=1}^M$ are also separable across base partitions. Using

$$-\log x = \max_{w>0} \{\log w - wx + 1\}, \tag{14}$$

whose unique maximizer is $w = x^{-1}$, and substituting $x = \varepsilon + \delta_m(U^{(r)}, R_m^{(r+1)})$, we obtain

$$w_m^{(r+1)} = \Big[\varepsilon + \delta_m\Big(U^{(r)}, R_m^{(r+1)}\Big)\Big]^{-1}. \tag{15}$$

Thus, a base partition receives a smaller weight when its aligned representation deviates more strongly from the current consensus.

Updating U. With $R^{(r+1)}$ and $w^{(r+1)}$ fixed, we optimize the consensus representation U. The terms involving U in Eq. (8) are

$$2 \operatorname{tr}\left(U^\top \sum_m \frac{w_m^{(r+1)}}{M d_m} X_m R_m^{(r+1)}\right) + \frac{\alpha}{K} \operatorname{tr}(U^\top A_{\text{str}} U) + \frac{g}{K} \operatorname{tr}(U^\top A_{\text{time}} U), \tag{16}$$

where $A_{\text{str}} = B_{\text{str}} B_{\text{str}}^\top \succeq 0$ and $A_{\text{time}} = B_{\text{time}} B_{\text{time}}^\top \succeq 0$. For any positive semidefinite matrix A,

$$\operatorname{tr}(U^\top A U) \geq 2 \operatorname{tr}(U^\top A U^{(r)}) - \operatorname{tr}((U^{(r)})^\top A U^{(r)}), \tag{17}$$

because the difference is $\operatorname{tr}((U - U^{(r)})^\top A (U - U^{(r)})) \geq 0$. Applying this lower bound to the two quadratic relation terms gives the surrogate subproblem

$$\max_{U^\top U = I} \operatorname{tr}(U^\top Y^{(r)}), \tag{18}$$

where $Y^{(r)}$ is defined in Equation (9). Let $Y^{(r)} = P \Sigma Q^\top$. The closed-form solution is

$$U^{(r+1)} = P Q^\top. \tag{19}$$

The above three steps complete one optimization sweep. We repeatedly update R, w, and U in this order until convergence.

B.3 Complexity Analysis

Let n denote the window size, M the number of base partitions, $L = M(M-1)/2$ the number of partition pairs, and $q = O(\sqrt{M})$ the number of witnesses.

Constructing the normalized base-partition matrices costs $O(nM)$. Computing pairwise partition agreements and constructing collision relations costs $O(nL)$, while selecting and weighting witnesses takes $O(M^2 \log M)$. Relation repair evaluates the assignment of each object to K clusters using q witnesses. This requires $O(nqK)$ per candidate and $O(nMqK)$ across all M initial partitions. Computing relative assignment support and constructing the signed relation factor takes an additional $O(nqK + nK)$. Historical propagation through the collision cells costs $O(nLK_{\mathrm{prev}})$, where K_{prev} is the preceding consensus cluster count.

The initial consensus is obtained by SVD, which costs $O(n(M+L+K)K)$ per iteration. During optimization, updating the alignment matrices and partition weights costs $O(nMK)$, apart from small SVDs whose dimensions are independent of n. Computing the current and historical contributions to the consensus update costs $O(nLK + nK^2 + nK_{\mathrm{prev}}K)$, and the SVD used to update the consensus embedding costs $O(nK^2)$. Each optimization iteration therefore requires $O(n(MK + LK + K^2 + K_{\mathrm{prev}}K))$ time. Applying K-means to the consensus embedding costs $O(nK^2)$ per iteration. The overall time complexity is thus linear in n for fixed ensemble dimensions, cluster counts, and iteration budgets.

The base-partition, collision, and repair matrices contain nM, nL, and $n(K-1)$ nonzero entries, respectively. Historical information requires $O(nK_{\mathrm{prev}})$ storage, while the consensus embedding and assignment scores require $O(nK)$. Together with the $O(M^2)$ partition-agreement matrix, the space complexity is $O(n(M + L + K + K_{\mathrm{prev}}) + M^2)$, which is also linear in n under the same assumptions.

B.4 Complexity and Streaming Capability Analysis

Table 4 compares AGILE with representative static and sequential clustering methods. Static ensemble methods such as GPEC, LEADEC, FCE, CEHM, YACHT, and CE-SoCo do not explicitly model information across successive windows. FCMVC-IV supports sequential learning and preserves information acquired from preceding stages, but operates in a continual multi-view setting.

Table 4: Complexity and streaming capability comparison of representative methods.

Method	Time Complexity	Space Complexity	Streaming Data	Historical Processing
GPEC [*ICML*, 25]	$\mathcal{O}(n^3)$	$\mathcal{O}(n^2)$	✗	✗
LEADEC [*TIP*, 25]	$\mathcal{O}(n^3)$	$\mathcal{O}(n^2)$	✗	✗
FCE [*TPAMI*, 25]	$\mathcal{O}(n^2)$	$\mathcal{O}(n)$	✗	✗
CEHM [*AAAI*, 25]	$\mathcal{O}(nc + c^2)$	$\mathcal{O}(nc)$	✗	✗
YACHT [*IJCAI*, 25]	$\mathcal{O}(nc^2 + c^3)$	$\mathcal{O}(nc + c^2)$	✗	✗
CE-SoCo [*TPAMI*, 26]	$\mathcal{O}(c^2 + n^2 \log n)$	$\mathcal{O}(\max\{c^2, n^2\})$	✗	✗
FCMVC-IV [*TIP*, 24]	$\mathcal{O}(T(mnc^2 + mc^3 + dc^2))$	$\mathcal{O}(nc + dc + c^2)$	✓	✓
AGILE (Ours)	$\mathcal{O}(nC^2 + C^3)$	$\mathcal{O}(nC + C^2)$	✓	✓

In contrast, AGILE performs sequential modeling directly from base partitions. It transfers the preceding consensus through persistent objects and regulates the historical contribution using current base-partition evidence. Let $C = \max\{K, K_{\mathrm{prev}}, \max_m c_m\}$ denote the maximum cluster-related dimension. With fixed ensemble size and iteration budgets, its $\mathcal{O}(nC^2 + C^3)$ time and $\mathcal{O}(nC + C^2)$ space complexities remain linear in the window size n.

C Additional Theoretical Analysis

C.1 Current-Only Reduction and Temporal Control

At the first window, or whenever no usable predecessor information is available, the temporal factor is inactive and $g = 0$. The objective then reduces to the current-window consensus formed from base-partition alignment and the corrected relational structure. This makes the first window well defined and allows the method to continue operating when temporal correspondence is weak.

The temporal contribution is also bounded. Since $0 \le g \le 1$, $0 \le \kappa \le 1$, and $\|B_{\text{time}}\|_F^2 = K\kappa$, its contribution to the objective is at most $g\kappa$. Thus, a larger overlap can provide more historical evidence without allowing the temporal influence to grow unchecked.

C.2 Supporting Properties

Proposition 1 (Common factor scale)**.** *For every base partition, $\|X_m\|_F^2 = K$. The initialization and collision factors satisfy $\|B_{\text{first}}\|_F^2 = \|B_{\text{col}}\|_F^2 = K$. Whenever a valid repair is constructed, $\|B_{\text{rep}}\|_F^2 = K$. If the temporal factor is active, $\|B_{\text{time}}\|_F^2 = K\kappa$.*

These normalizations keep the direct, collision, and repair representations on comparable scales. The number or size of collision cells therefore does not by itself increase their contribution. When repair is active, $\|B_{\text{str}}\|_F^2 = K(1+\beta)$.

Proposition 2 (Accepted repair and support weighting)**.** *Fix the corroboration weights $\{\omega_m\}_{m\in\mathcal{E}}$. Let z_{start} be a valid K-way candidate after deterministic merging, let z_{prop} be its synchronous reassignment proposal, and let $\tau \ge 0$ be the fixed numerical acceptance tolerance. The retained candidate $\widehat{z}$ satisfies $\Phi(\widehat{z}) \ge \Phi(z_{\text{start}}) - \tau$, and z^* maximizes the retained value of Φ over all valid base-partition starts. With the relative support scores in Equation* (4)*, its signed factor has the Gram form in Equation* (5)*, preserves $\|B_{\text{rep}}\|_F^2 = K$, and defines a positive-semidefinite operator. In exact arithmetic with $\tau = 0$, every accepted repair sweep is nondecreasing in Φ.*

The acceptance test prevents the repair sweep from materially reducing agreement with the selected witnesses. The relative support scores then concentrate more of the fixed repair strength on better-supported assignments.

Corollary 1 (Stationarity of fixed points)**.** *Every fixed point of the three update maps satisfies the first-order stationarity conditions of Equation* (8) *with respect to R, w, and U under their respective constraints.*

D Proofs

D.1 Proof of Proposition 1

For partition m, column j of X_m contains n_{mj} copies of $\sqrt{K/c_m}/\sqrt{n_{mj}}$, so its squared norm is K/c_m. Distinct columns have disjoint support, which gives $X_m^\top X_m = (K/c_m)I$ and $\|X_m\|_F^2 = K$. Because B_{first} concatenates $\sqrt{p_m}X_m$ and $\sum_m p_m = 1$, its squared Frobenius norm is $\sum_m p_m K = K$.

Every row of B_{col} contains exactly L nonzeros, each of squared value $K/(nL)$. Summing over n rows gives $\|B_{\text{col}}\|_F^2 = nL\,K/(nL) = K$. For the support-weighted repair, row i participates in exactly $K-1$ cluster-pair columns, each with squared magnitude $Kh_i/[n(K-1)\bar{h}]$. Hence its row energy is $Kh_i/(n\bar{h})$, and $\|B_{\text{rep}}\|_F^2 = (K/(n\bar{h}))\sum_i h_i = K$ by the definition of $\bar{h}$. Finally, B_{time} is defined by normalizing $\bar{T}$ to Frobenius norm $\sqrt{K\kappa}$, which proves the last identity. □

D.2 Proof of the support-weighted signed Gram identity

Let $r_i = \sqrt{h_i/\bar{h}}$. If $z_i^* = z_j^* = a$, rows i and j have identical signs in the $K-1$ columns associated with unordered pairs incident to a, so their inner product is $(K-1)s^2 r_i r_j = (K/n)\sqrt{h_i h_j}/\bar{h}$. If $z_i^* = a \ne b = z_j^*$, the rows are jointly nonzero only in the column corresponding to (a,b) and have opposite signs, giving $-s^2 r_i r_j = -K\sqrt{h_i h_j}/[n(K-1)\bar{h}]$. This proves Equation (5). Positive semidefiniteness follows directly because the operator is $B_{\text{rep}}B_{\text{rep}}^\top$. □

D.3 Proof of Proposition 2

The witness set and corroboration weights are computed once from the current base partitions and remain fixed during repair. For a valid start z_{start}, the algorithm constructs a synchronous reassignment proposal z_{prop}, keeping the member with the highest current-cluster affinity in place whenever all members of a cluster propose to leave. Any proposal with fewer than K nonempty

Table 5: NMI (%) across different overlap rates ρ, with the best results highlighted in bold.

Method	COIL100			Nursery			TwentyNewsgroups			NATICUSdroid		
	$\rho = 0.2$	$\rho = 0.5$	$\rho = 0.8$	$\rho = 0.2$	$\rho = 0.5$	$\rho = 0.8$	$\rho = 0.2$	$\rho = 0.5$	$\rho = 0.8$	$\rho = 0.2$	$\rho = 0.5$	$\rho = 0.8$
SEC [*KDD*, 2015]	75.92 ± 0.34	74.94 ± 0.16	74.31 ± 0.30	4.91 ± 1.77	7.77 ± 2.02	4.17 ± 1.46	11.13 ± 1.02	17.72 ± 1.08	27.01 ± 0.54	8.39 ± 4.70	10.22 ± 3.12	6.40 ± 2.16
PTGP [*TKDE*, 2016]	76.16 ± 0.19	74.39 ± 0.12	71.84 ± 0.18	5.35 ± 2.12	7.57 ± 2.99	3.54 ± 1.79	10.06 ± 0.71	15.03 ± 0.78	23.04 ± 1.01	5.21 ± 0.26	5.54 ± 0.11	5.60 ± 0.14
LWGP [*TCYB*, 2018]	78.68 ± 0.11	77.39 ± 0.16	75.66 ± 0.13	5.52 ± 1.90	7.44 ± 3.08	4.73 ± 2.01	1.80 ± 0.08	1.55 ± 0.09	1.74 ± 0.75	5.39 ± 0.18	5.61 ± 0.00	5.67 ± 0.07
USENC [*TKDE*, 2020]	80.85 ± 0.12	79.53 ± 0.09	77.78 ± 0.10	4.73 ± 1.52	7.05 ± 3.56	3.63 ± 1.98	12.97 ± 0.89	19.07 ± 0.81	28.01 ± 1.13	5.43 ± 0.23	5.61 ± 0.00	5.83 ± 0.87
ECPCS-MC [*TSMC*, 2021]	79.08 ± 0.09	77.89 ± 0.11	76.50 ± 0.11	4.66 ± 1.15	7.52 ± 3.23	4.26 ± 1.90	1.15 ± 0.07	0.86 ± 0.08	1.52 ± 0.65	5.19 ± 0.32	5.29 ± 0.14	5.32 ± 0.25
DenForest [*SIGMOD*, 2022]	26.47 ± 5.36	45.54 ± 8.12	45.45 ± 2.53	2.82 ± 1.47	6.99 ± 4.04	2.87 ± 1.40	0.81 ± 0.13	0.82 ± 0.20	1.13 ± 0.24	7.22 ± 0.90	6.93 ± 0.85	7.02 ± 0.81
FastMICE [*TKDE*, 2023]	80.42 ± 0.13	79.03 ± 0.14	77.22 ± 0.12	5.31 ± 2.51	7.80 ± 3.19	3.82 ± 1.72	12.72 ± 0.77	18.80 ± 0.91	27.83 ± 1.09	5.43 ± 0.16	5.61 ± 0.00	5.62 ± 0.17
PSIMVC-PG [*TNNLS*, 2024]	81.03 ± 0.13	79.58 ± 0.19	77.90 ± 0.13	4.37 ± 1.02	5.42 ± 1.40	3.81 ± 1.07	9.59 ± 0.44	13.67 ± 0.85	21.64 ± 1.45	5.04 ± 0.25	5.72 ± 1.03	5.48 ± 0.21
FCMVC-IV [*TIP*, 2024]	78.52 ± 0.28	77.29 ± 0.08	75.68 ± 0.16	7.37 ± 3.37	8.16 ± 4.63	$\mathbf{5.32 \pm 2.52}$	9.49 ± 0.75	13.36 ± 1.18	20.33 ± 0.96	7.16 ± 2.21	5.23 ± 1.60	5.21 ± 1.73
CEHM [*AAAI*, 2025]	77.82 ± 0.18	76.68 ± 0.21	75.03 ± 0.21	0.44 ± 0.10	0.66 ± 0.93	0.52 ± 0.32	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	4.26 ± 0.44	4.35 ± 0.20	4.62 ± 0.37
FSEC [*Inf. Fusion*, 2025]	80.99 ± 0.10	79.62 ± 0.08	77.86 ± 0.14	5.11 ± 1.98	7.68 ± 3.41	3.69 ± 1.95	13.23 ± 0.82	$\mathbf{19.37 \pm 0.77}$	$\mathbf{28.18 \pm 0.91}$	5.48 ± 0.10	5.61 ± 0.00	5.66 ± 0.13
MVC-IIV [*TIP*, 2025]	80.16 ± 0.09	78.80 ± 0.08	77.12 ± 0.21	3.35 ± 1.05	3.92 ± 1.55	2.44 ± 0.64	11.12 ± 0.88	16.95 ± 0.86	25.34 ± 0.50	10.46 ± 3.72	10.22 ± 3.04	9.72 ± 2.57
RANGE [*CVPR*, 2026]	81.78 ± 0.17	80.33 ± 0.10	78.58 ± 0.12	6.35 ± 3.30	6.53 ± 2.27	3.42 ± 2.13	11.79 ± 0.64	16.81 ± 0.71	24.72 ± 0.71	0.06 ± 0.00	0.04 ± 0.00	0.02 ± 0.00
DREAM [*TKDE*, 2026]	79.32 ± 0.21	77.70 ± 0.18	75.90 ± 0.17	0.04 ± 0.08	0.06 ± 0.10	0.04 ± 0.05	0.00 ± 0.00	0.00 ± 0.00	0.92 ± 1.41	0.08 ± 0.16	0.04 ± 0.13	0.01 ± 0.01
AGILE (Ours)	$\mathbf{81.82 \pm 0.15}$	$\mathbf{80.39 \pm 0.13}$	$\mathbf{78.76 \pm 0.12}$	$\mathbf{10.08 \pm 3.68}$	$\mathbf{8.75 \pm 3.13}$	5.07 ± 3.80	$\mathbf{13.27 \pm 0.67}$	18.29 ± 0.85	27.04 ± 0.73	$\mathbf{12.83 \pm 2.66}$	$\mathbf{13.38 \pm 2.89}$	$\mathbf{11.08 \pm 1.15}$

Method	FashionMNIST			KannadaMNIST			MNIST			MHEALTH		
	$\rho = 0.2$	$\rho = 0.5$	$\rho = 0.8$	$\rho = 0.2$	$\rho = 0.5$	$\rho = 0.8$	$\rho = 0.2$	$\rho = 0.5$	$\rho = 0.8$	$\rho = 0.2$	$\rho = 0.5$	$\rho = 0.8$
SEC [*KDD*, 2015]	47.56 ± 1.29	47.95 ± 0.74	48.08 ± 0.76	56.48 ± 1.02	57.33 ± 1.39	57.83 ± 1.37	38.91 ± 1.11	38.95 ± 0.56	38.15 ± 1.05	39.70 ± 1.24	39.53 ± 0.87	38.81 ± 0.90
PTGP [*TKDE*, 2016]	$\mathbf{51.31 \pm 0.31}$	$\mathbf{51.27 \pm 0.16}$	49.29 ± 0.19	60.92 ± 0.58	61.69 ± 0.51	61.18 ± 0.30	41.49 ± 0.29	41.56 ± 0.25	41.33 ± 0.21	40.79 ± 0.58	40.87 ± 0.73	40.86 ± 0.85
LWGP [*TCYB*, 2018]	51.04 ± 0.28	51.06 ± 0.34	51.05 ± 0.40	58.23 ± 0.37	57.91 ± 0.60	57.98 ± 0.29	39.73 ± 0.37	39.87 ± 0.30	40.23 ± 0.20	43.33 ± 1.06	44.34 ± 0.43	42.93 ± 0.50
USENC [*TKDE*, 2020]	50.93 ± 0.41	51.04 ± 0.35	51.07 ± 0.30	61.51 ± 0.58	62.05 ± 0.36	62.50 ± 0.53	41.04 ± 0.28	41.18 ± 0.28	41.17 ± 0.30	41.44 ± 0.71	42.25 ± 0.79	41.82 ± 0.78
ECPCS-MC [*TSMC*, 2021]	51.07 ± 0.23	51.12 ± 0.21	51.07 ± 0.23	55.60 ± 1.46	54.93 ± 1.05	55.76 ± 1.21	36.24 ± 1.05	37.52 ± 1.63	37.05 ± 2.26	38.97 ± 0.43	39.20 ± 0.44	39.20 ± 0.60
DenForest [*SIGMOD*, 2022]	49.05 ± 0.33	48.40 ± 2.40	48.99 ± 0.38	48.32 ± 5.97	53.92 ± 1.00	48.49 ± 1.57	3.77 ± 1.65	1.68 ± 0.76	2.36 ± 2.31	35.04 ± 0.73	29.58 ± 2.21	28.49 ± 6.75
FastMICE [*TKDE*, 2023]	51.08 ± 0.35	51.21 ± 0.14	51.25 ± 0.19	61.77 ± 0.35	61.78 ± 0.46	61.65 ± 0.41	41.31 ± 0.31	41.51 ± 0.30	41.36 ± 0.25	41.65 ± 0.57	42.70 ± 0.73	42.28 ± 0.69
PSIMVC-PG [*TNNLS*, 2024]	50.47 ± 0.51	50.67 ± 0.20	50.66 ± 0.32	59.96 ± 0.72	60.30 ± 0.50	60.07 ± 0.53	40.97 ± 0.29	41.16 ± 0.19	41.12 ± 0.33	42.36 ± 0.49	42.63 ± 0.57	42.23 ± 0.63
FCMVC-IV [*TIP*, 2024]	50.58 ± 0.52	50.64 ± 0.22	50.67 ± 0.35	58.80 ± 0.50	58.81 ± 0.73	58.86 ± 0.56	40.45 ± 0.44	40.57 ± 0.32	40.55 ± 0.34	40.39 ± 0.50	40.47 ± 0.87	40.45 ± 0.91
CEHM [*AAAI*, 2025]	50.71 ± 0.32	51.02 ± 0.21	51.06 ± 0.26	54.34 ± 1.00	53.47 ± 0.85	54.71 ± 0.69	40.50 ± 0.37	40.83 ± 0.39	41.13 ± 0.42	39.45 ± 0.37	39.76 ± 0.34	39.69 ± 0.38
FSEC [*Inf. Fusion*, 2025]	51.23 ± 0.28	51.20 ± 0.19	51.25 ± 0.26	61.32 ± 0.48	62.07 ± 0.41	61.42 ± 0.22	41.27 ± 0.33	41.45 ± 0.28	41.39 ± 0.19	41.26 ± 0.38	42.07 ± 0.74	41.78 ± 0.66
MVC-IIV [*TIP*, 2025]	50.43 ± 0.34	50.85 ± 0.35	50.78 ± 0.34	60.80 ± 0.64	61.49 ± 0.50	61.92 ± 0.25	40.33 ± 0.28	40.52 ± 0.21	40.40 ± 0.25	42.11 ± 0.38	42.58 ± 0.47	42.36 ± 0.36
RANGE [*CVPR*, 2026]	50.36 ± 0.38	50.78 ± 0.41	50.68 ± 0.34	62.03 ± 0.27	61.32 ± 0.44	62.50 ± 0.32	40.93 ± 0.21	40.92 ± 0.22	40.58 ± 0.34	44.57 ± 0.45	43.95 ± 0.39	43.59 ± 0.55
DREAM [*TKDE*, 2026]	49.51 ± 0.49	49.74 ± 0.43	49.98 ± 0.51	57.46 ± 0.86	57.67 ± 1.22	58.09 ± 1.31	38.63 ± 0.34	39.32 ± 0.51	39.04 ± 0.28	39.90 ± 0.51	39.65 ± 0.57	40.03 ± 0.69
AGILE (Ours)	50.96 ± 0.28	51.07 ± 0.33	$\mathbf{51.29 \pm 0.29}$	$\mathbf{62.12 \pm 0.48}$	$\mathbf{62.50 \pm 0.28}$	$\mathbf{62.98 \pm 0.29}$	$\mathbf{41.81 \pm 0.32}$	$\mathbf{42.17 \pm 0.16}$	$\mathbf{42.41 \pm 0.23}$	$\mathbf{44.92 \pm 0.41}$	$\mathbf{45.05 \pm 0.38}$	$\mathbf{44.79 \pm 0.45}$

clusters is discarded. Otherwise the same score Φ is recomputed, and the proposal is retained only when $\Phi(z_{\text{prop}}) + \tau \geq \Phi(z_{\text{start}})$, with the implementation using $\tau = 10^{-12}$. Hence $\Phi(\widehat{z}) \geq \Phi(z_{\text{start}}) - \tau$; for $\tau = 0$, the accepted sweep is exactly nondecreasing. The final repaired candidate is the retained candidate with the largest Φ over all valid starts. The remaining claims follow from Proposition 1 and the signed Gram identity above. □

D.4 Proof of Theorem 1

Appendix B.2 shows that the R and w updates globally maximize their respective subproblems and that the U update maximizes a tight lower bound coinciding with the original objective at $U^{(r)}$. Hence a complete $R \to w \to U$ cycle satisfies

$$\mathcal{J}^{(r+1)} \geq \mathcal{J}^{(r)}.$$

It remains to bound the objective. Equation (14) implies, for every $w_m > 0$, $\log w_m - w_m(\varepsilon + \delta_m) + 1 \leq -\log(\varepsilon + \delta_m) \leq -\log \varepsilon$. Moreover, $\|B^\top U\|_F^2 \leq \|B\|_F^2$ whenever $U^\top U = I$. Proposition 1 gives $\|B_{\text{str}}\|_F^2 \leq K(1 + \beta)$ and $\|B_{\text{time}}\|_F^2 = K\kappa$. Therefore

$$\mathcal{J}^{(r)} \leq -\log \varepsilon + \alpha(1 + \beta) + g\kappa,$$

which is finite. The objective sequence is monotone and bounded above, and hence converges. □

D.5 Proof of Corollary 1

At a fixed point (U^*, R^*, w^*), each R_m^* is a global Procrustes solution and each w_m^* is the unique maximizer of its scalar auxiliary term, so the first-order conditions of the R and w blocks hold. For the U block, the Euclidean gradient of the constraint-equivalent trace formulation at the fixed point is $2Y^*$, where Y^* is Equation (9) evaluated at U^*. Since $U^* = \text{polar}(Y^*)$ maximizes $\text{tr}(U^\top Y^*)$ on the Stiefel manifold, its projected gradient vanishes. The same condition therefore holds for the original objective, proving first-order stationarity of the fixed point. □

E Implementation Details

Current-only initialization. To construct a current-only reference for assessing historical compatibility, we assign greater initialization weight to base partitions with stronger ensemble support. Let $r_m = (M-1)^{-1} \sum_{r \neq m} [a_{mr}]_+$ denote the positive centrality of partition m, and let $\bar{r} = M^{-1} \sum_m r_m$ be its ensemble mean. We define $p_m = (r_m + \bar{r}/M) / \sum_\ell (r_\ell + \bar{r}/M)$, using

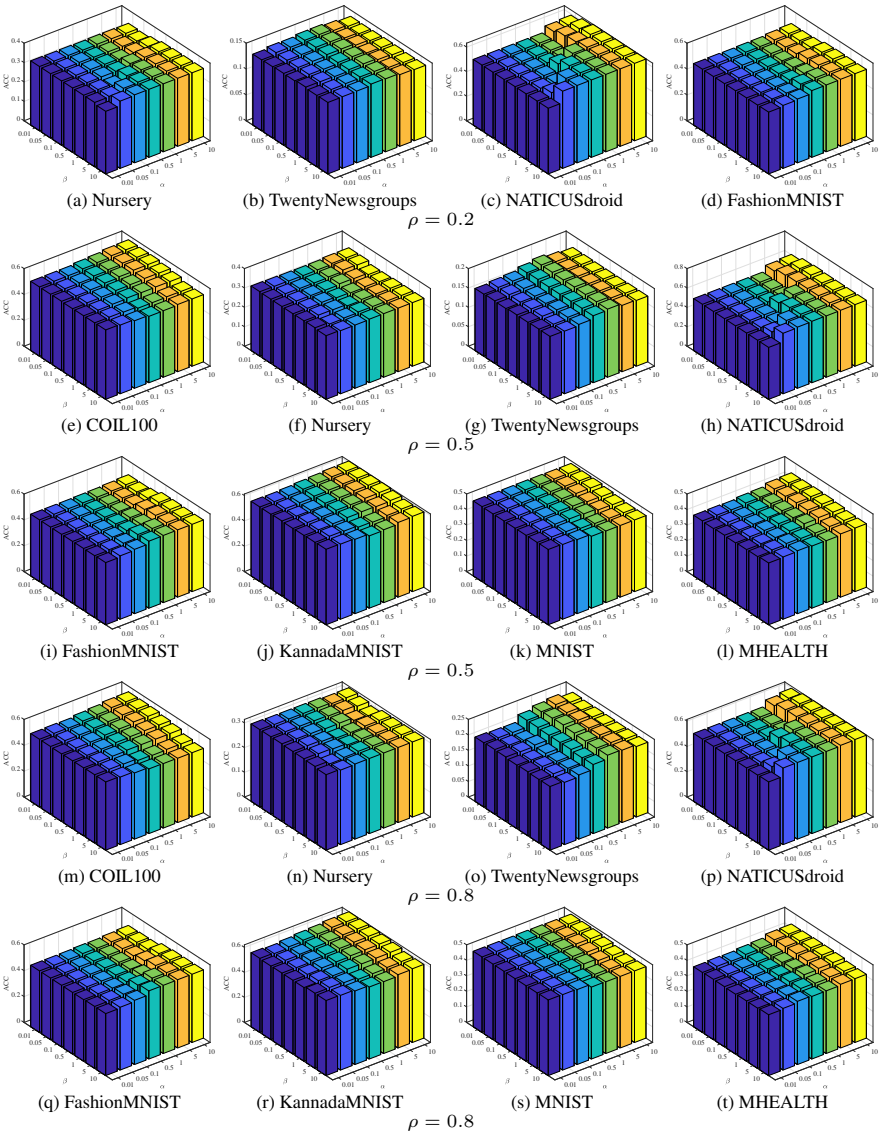


Figure 5: Additional sensitivity results of ACC with respect to α and β under different overlap rates.

$p_m = 1/M$ when all centralities vanish, and form $B_{\text{first}} = [\sqrt{p_1}X_1, \ldots, \sqrt{p_M}X_M]$. These weights emphasize agreement-supported partitions while preserving the common scale $\|B_{\text{first}}\|_F^2 = K$. The current-only partition $z^{(0)}$ is obtained through spectral initialization based on $[B_{\text{first}}, \sqrt{\alpha}B_{\text{str}}]$, which combines the weighted base partitions with the corrected current relations. Its cluster-indicator matrix, with each column normalized to unit Euclidean norm, provides the initial consensus embedding $U^{(0)}$. This construction uses no historical assignments, allowing $z^{(0)}$ to serve as the reference for the temporal compatibility gate. The weights p_m are used only for initialization and remain distinct from the alignment weights w_m learned during optimization.

We use $\varepsilon = 10^{-6}$ for robust weighting and perform at most 30 optimization iterations with a joint relative-change tolerance of 10^{-3}. The repair acceptance tolerance and support-score floor are both fixed at 10^{-12}. The witness count and corroboration neighborhood size are set to $q = \min\{M, \max(2, \lceil\sqrt{M}\rceil)\}$ and $\nu = \min\{M-1, \max(1, \lceil\sqrt{M}\rceil)\}$, respectively, without tuning. Both α and β are searched over $\{0.01, 0.05, 0.1, 0.5, 1, 5, 10\}$. In each trial, 20 base partitions are selected using a fixed random seed and shared by all methods. The same random-seed policy is applied across overlap settings and repeated trials.

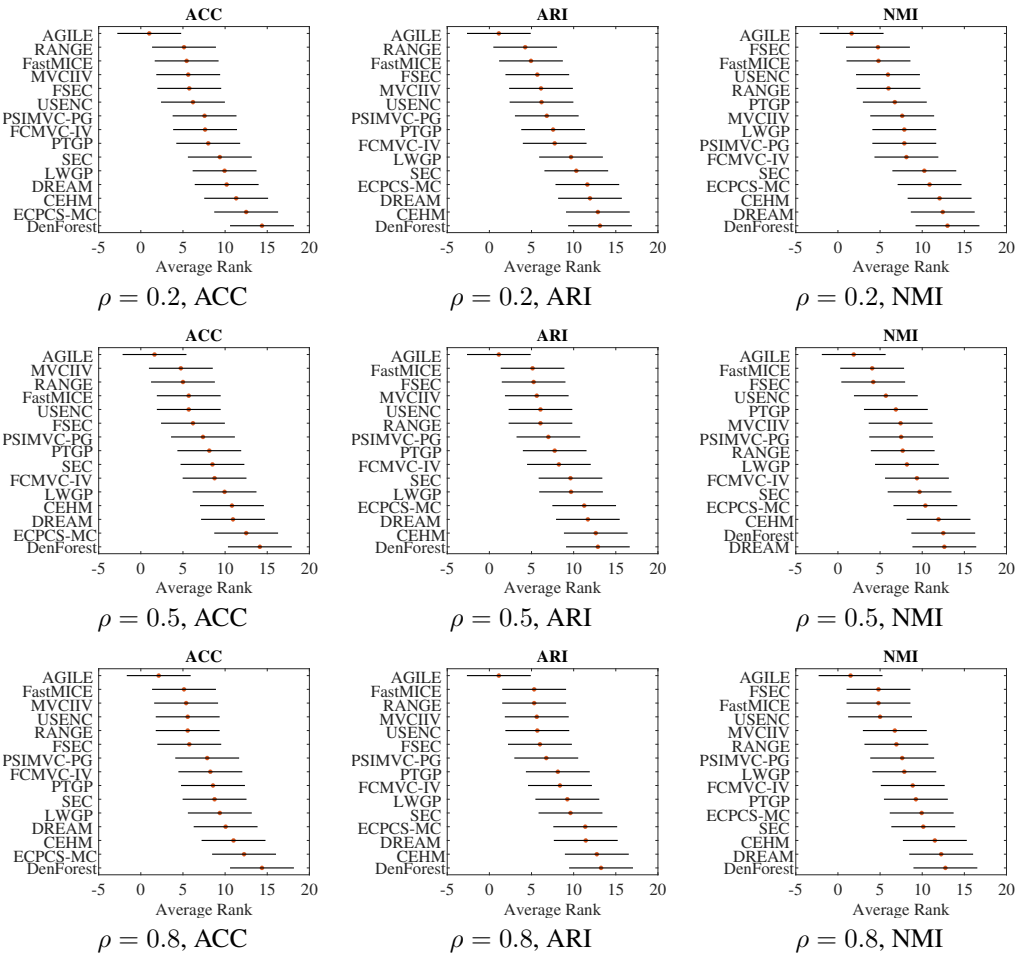


Figure 6: Average ranks across eight datasets. Rows correspond to $\rho = 0.2$, 0.5, and 0.8; columns correspond to ACC, ARI, and NMI. Lower ranks indicate better performance.

F Additional Experimental Results

F.1 Additional NMI Results

Table 5 reports the complete NMI comparison under all three overlap rates, complementing the ACC and ARI results in Tables 2 and 3. AGILE obtains the best NMI in 19 of the 24 settings and exceeds RANGE in all 24 dataset–overlap comparisons. The remaining best results are distributed across different competitors rather than concentrated in one method. The NMI trends also remain non-monotonic with respect to ρ on several datasets. Together with the ACC and ARI results, these observations show that clustering quality varies with the window configuration and is not determined by overlap rate alone.

F.2 Parameter Sensitivity on All Settings

Figure 5 complements Figure 2 with the remaining datasets and overlap rates. Most settings exhibit broad regions in which neighboring parameter choices give similar ACC, rather than a single isolated optimum. Nursery, TwentyNewsgroups, and NATICUSdroid show more visible changes in some overlap settings, indicating that the preferred balance between current structure and relation repair can vary across datasets. No single pair dominates all settings, but values in $\{0.5, 1, 5\}$ repeatedly lie in stable high-performing regions. This supports the compact practical search range recommended in Section 5.4 while retaining the full grid for the reported comparisons.

F.3 Statistical Rank Analysis

We conduct the Friedman test followed by the Nemenyi post-hoc test to compare all fifteen methods across the eight datasets. The analysis is performed separately for ACC, ARI, and NMI at each overlap rate. Methods are ranked by their mean performance on each dataset, with lower ranks indicating better results. Figure 6 summarizes the average ranks and the corresponding critical-distance intervals. Non-overlapping intervals indicate a statistically significant difference between the corresponding methods.

AGILE achieves the best average rank in all nine metric–overlap settings. Although the relative ordering of the competing methods changes across metrics and overlap rates, AGILE consistently maintains the leading position. Its advantage is observed at both low and high overlap rates, where the amount of historical information available through persistent objects differs substantially. The agreement across ACC, ARI, and NMI further shows that this favorable ranking is maintained under different measures of clustering quality. These results complement the dataset-level comparisons and support the overall effectiveness of AGILE across the evaluated window configurations.