PRELIMINARY PREPRINT VERSION – DO NOT CITE
This is not the version of record. The final version will appear in the Proceedings of IJCAI-ECAI 2026.

Semi-supervised Clustering via Adversarially Enhanced Intent Propagation

Wentao Zhong[1,2], **Ruina Bai**[1,2,*], **Jingjing Xue**[1,2], **Ying Nie**[1,2] and **Ruizhang Huang**[1,2,*]
[1] Text Computing & Cognitive Intelligence Engineering Research Center of National Education Ministry, Guizhou University
[2] College of Computer Science and Technology, Guizhou University, Guiyang, Guizhou 550025, China
gs.wtzhong24@gzu.edu.cn, {bairuina22453, jingjingxue.gz}@gmail.com, {gs.yzhe24, rzhuang}@gzu.edu.cn

Abstract

Semi-supervised clustering (SSC) enables personalized clustering under limited user supervision. Given the sparsity of initial user intents, constraint propagation has been proposed as a powerful approach to explore and generate new performance-enhancing constraints. However, existing methods struggle to handle a large number of "gray samples" that deviate from initial supervision signals and exhibit ambiguous semantics with indistinct boundaries. Compared with "distinct samples" that closely match the supervision, gray samples typically contain richer latent semantics, and accurately identifying their relational types can significantly improve clustering performance. To address this challenge, we propose Adversarially Enhanced Propagation-driven Intent-aware Clustering (AEPIC). Specifically, we design an Adversarially Enhanced Constraint Propagation (AECP) mechanism that leverages global adversarial learning over dual relational links to identify gray samples and expand them into meaningful pseudo-user intents. In addition, an intent-aware regularization strategy integrates these pseudo-user intents into representation learning and clustering optimization, further improving clustering performance. Experiments on 5 benchmark datasets demonstrate that, under sparse supervision, AEPIC consistently outperforms state-of-the-art semi-supervised clustering methods.

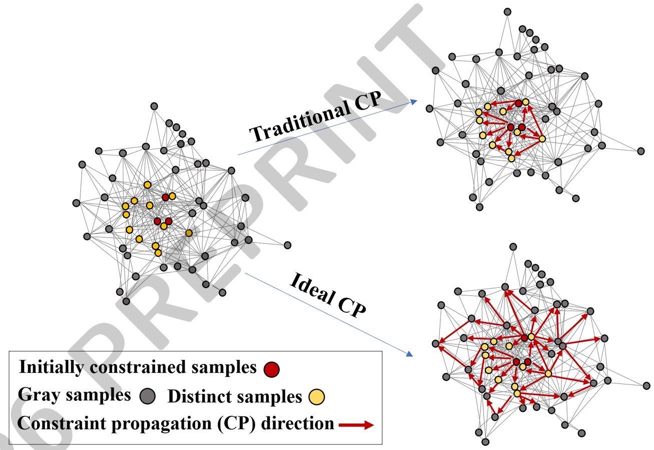

Figure 1: This figure illustrates the motivation of our study. In the feature space, ignoring gray samples during constraint propagation may cause samples that deviate from the initial constraints to lose key discriminative features. Our goal is to enhance model performance by effectively propagating gray samples to expand constraint information.

IJCAI-ECAI 2026 PREPRINT

1 Introduction

Semi-supervised clustering combines the strengths of semi-supervised learning and clustering analysis, leveraging limited prior knowledge (e.g., pairwise constraints or ground-truth labels) to guide the formation of cluster structures that better align with user preferences. In practice, determining whether two samples belong to the same class is often easier than obtaining ground-truth labels for individual samples. However, due to the scarcity and high cost of supervision, most existing methods rely on sparse and static user-intent constraints during training, which often leads to significant degradation in clustering performance under weak supervision.

In recent years, several semi-supervised models have attempted to alleviate the sparsity of initial user intents. Methods such as [Śmieja *et al.*, 2020; Wu *et al.*, 2024] propagate pseudo-labels across the dataset via neural networks, while others (e.g., [Sun, 2025]) focus on discovering new constraints from local neighborhood structures. Additionally, [Wang *et al.*, 2023] leverage both class labels and pairwise constraints to propagate user intents. Although these methods partially alleviate the sparsity of user intents, they fail to explicitly handle the propagation of gray samples that markedly deviate from the initial constraints. Ignoring or misappropriating these high-value samples can introduce a substantial number of noisy pseudo-user intents, severely compromising clustering performance.

Specifically, real-world data often includes gray samples that deviate from initial user intents, exhibiting ambiguous semantics and unclear boundaries. In contrast, distinct samples closely align with the initial intent and are easily recognized by clustering models. Although difficult to identify, gray samples often carry rich semantic information that can improve clustering structures. Figure 1 compares the directions of constraint propagation under different methods in the feature space. Traditional constraint propagation (CP) methods often overlook the propagation of supervision among gray samples, which can result in the loss of critical discriminative constraints. In contrast, an ideal propagation strategy should extend constraints to higher-order samples, rather than being limited to those highly similar to the initial constrained pairs. Under limited supervision, accurately identifying high-value gray samples and expanding them into reliable pseudo-user intents remains a core bottleneck in semi-supervised clustering.

To address these challenges, we propose the Adversarially Enhanced Constraint Propagation (AECP) module, which propagates supervision locally and employs an adversarial mechanism between two link types to globally enhance gray-sample confidence, producing higher-quality pseudo-user intents. Based on AECP, the semi-supervised AEPIC framework jointly optimizes these pseudo-user intents during representation learning and clustering to improve overall performance.

The main contributions of this paper are summarized as follows:

- We propose a novel Adversarially Enhanced Propagation-driven Intent-aware Clustering (AEPIC) framework, in which the Intent-Aware Regularization (IAR) module integrates the expanded pseudo-user intents into both representation learning and clustering. By jointly optimizing these two stages, AEPIC substantially improves clustering performance under sparse supervision.
- We design the Adversarially Enhanced Constraint Propagation (AECP) module to address limitations of existing propagation methods in identifying information-rich gray samples. AECP employs a global adversarial mechanism over supervisory connections to uncover constraint relationships among gray samples that are mispropagated or overlooked, enabling higher-quality expansion of pseudo-user intents.
- We conduct extensive experiments on 5 real-world datasets. Under identical sparse supervision, AEPIC consistently outperforms state-of-the-art semi-supervised clustering methods, demonstrating its effectiveness and generality.

2 Related Work

2.1 Semi-Supervised Clustering

Semi-supervised clustering leverages limited prior knowledge, such as must-link and cannot-link constraints, to guide clustering toward user-aligned structures. Early works extended classical algorithms with constraints: COP-Kmeans [Wagstaff *et al.*, 2001] incorporated pairwise constraints into K-means, exploiting symmetry and transitivity to improve clustering; MLC-Kmeans [Huang *et al.*, 2008] introduced the Must-Link Set for better centroid estimation; CLWC [Cheng *et al.*, 2008] jointly learned local weighted distance metrics and constraints.

With deep learning, constrained clustering has been embedded into neural networks. SDEC [Ren *et al.*, 2019] extends DEC [Xie *et al.*, 2016] by integrating pairwise constraints into deep representations; DCC [Zhang *et al.*, 2019] builds on IDEC [Guo *et al.*, 2017] to leverage pairwise constraints and richer side information, including continuous constraints and domain knowledge. Similarly, [Hsu and Kira, 2015] trained neural networks with prior constraints to learn cluster-friendly representations from raw data.

Despite incorporating user intent, these methods rely on sparse, static supervision. Acquiring high-quality prior knowledge is costly, limiting constraint effectiveness and motivating dynamic, scalable propagation of user intent under limited supervision.

2.2 Constraint Propagation

Constraint propagation extends limited pairwise user-supervised signals to unlabeled sample pairs, enhancing constraint utilization and model learning under sparse supervision. Early work focused on traditional frameworks: [Lu and Carreira-Perpinan, 2008] used Gaussian processes but was limited to binary classification; [Li *et al.*, 2008] formulated propagation as a semidefinite programming (SDP) problem, extendable to multi-class cases but computationally expensive; [Yu and Shi, 2004] supported only must-link constraints, limiting applicability.

To address these issues, [Lu and Peng, 2013a] proposed E^2CP, an exhaustive and efficient method jointly propagating must-link and cannot-link constraints while reducing computational cost. [Jia *et al.*, 2020] employed an adversarial mechanism to propagate the two links independently. However, both rely on traditional similarity learning and lack dynamic modeling of confidence matrices, limiting constraint exploitation. [Yin *et al.*, 2023] leverages limited supervision to construct richer similarity matrices. For multimodal data, [Lu and Peng, 2013b] proposed a unified intra- and inter-view propagation framework, whereas [Eaton *et al.*, 2010] propagates constraints across views using local similarity metrics, enabling partial cross-view transfer.

Despite these advances, challenges remain. For gray samples whose features deviate from initial supervision, existing methods often propagate constraints only to similar samples, overlooking or mispropagating informative gray samples, which introduces noise and reduces the effectiveness of user-intent propagation.

3 Methods

3.1 Network Architecture

We propose a novel Adversarially Enhanced Intent Propagation Clustering (AEPIC) model to tackle propagate in semi-supervised personalized clustering. As shown in Figure 2, AEPIC comprises two main components: the Adversarially Enhanced Constraint Propagation (AECP) module and the Intent-Aware Regularization (IAR) module. The AECP module includes Local Supervised Reconstruction (LSR) and

IJCAI-ECAI 2026 PREPRINT

PRELIMINARY PREPRINT VERSION – DO NOT CITE
This is not the version of record. The final version will appear in the Proceedings of IJCAI-ECAI 2026.

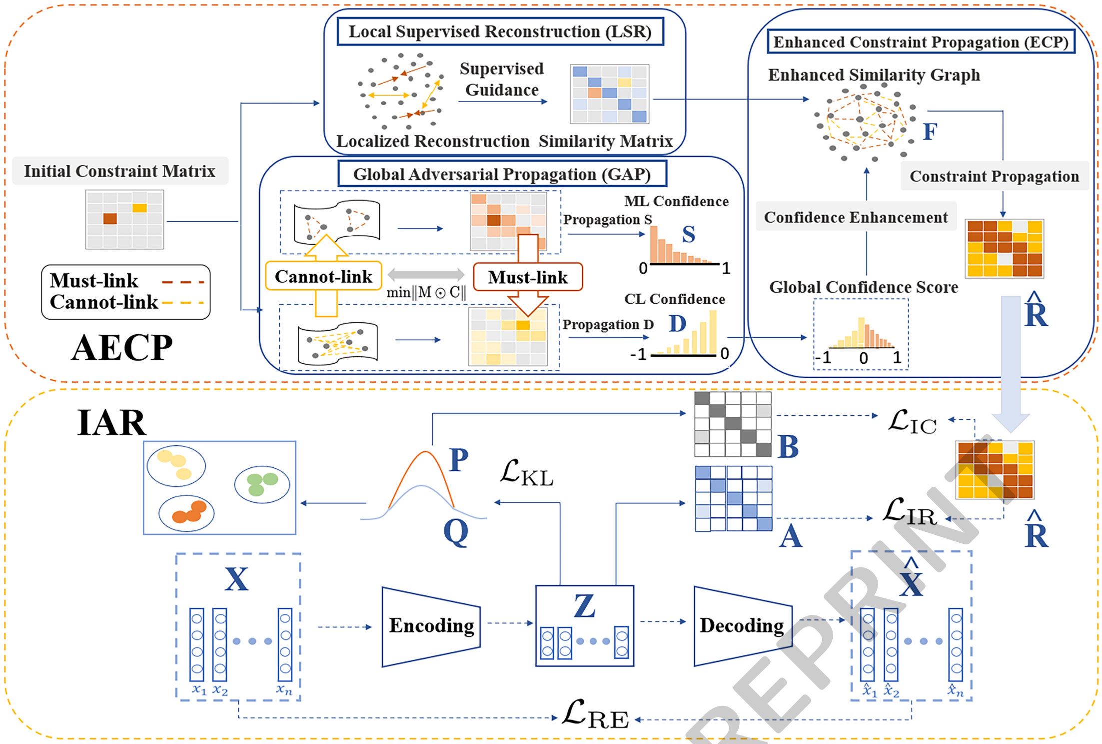

Figure 2: The AEPIC framework integrates the AECP module, which employs a dual-channel adversarial strategy to propagate two mutually exclusive link types and learn their confidence matrices, thereby enhancing gray-sample propagation and applying intent regularization in the IRC module to leverage the resulting pseudo-user constraints.

IJCAI-ECAI 2026 PREPRINT

Global Adversarial Propagation (GAP), with Enhanced Constraint Propagation (ECP) further improving propagation. The IAR module incorporates the propagated pseudo-user intents into model learning, enhancing clustering under sparse supervision.

3.2 AECP module

Local Supervised Reconstruction (LSR)

In intent propagation, high-similarity sample pairs do not always indicate must-link relations, and vice versa, since traditional similarity matrices are typically unsupervised and may not reflect user supervision accurately. To better leverage supervisory information, we reconstruct local similarities from the initial pairwise constraints, producing a graph that aligns more closely with user clustering preferences locally.

To encode the initial user-provided intentions, we define a pairwise constraint matrix $T = \{t_{ij}\}_{n\times n}$, where each entry indicates the type of relationship between instances x_i and x_j as follows:

$$t_{ij} = \begin{cases} 1, & \text{if } (x_i, x_j) \in \mathcal{M} \text{ (Must-link)} \\ -1, & \text{if } (x_i, x_j) \in \mathcal{C} \text{ (Cannot-link)} \\ 0, & \text{otherwise} \end{cases} \tag{1}$$

For each sample x_i, we learn a neighbor weight vector $\mathbf{w}_i = [w_{i1}, w_{i2}, \ldots, w_{ik}]^T$, where k denotes the number of neighbors. The objective is to learn neighborhood weights that are aligned with user intents. The corresponding optimization function is formulated as:

$$\begin{aligned} \min_{\mathbf{w}_i} \quad & \mathbf{w}_i^\top G_i \mathbf{w}_i + \lambda \sum_{j\in\mathcal{N}_i} H_{ij}(w_{ij} - t_{ij})^2 + \mu \|\mathbf{w}_i\|^2 \\ \text{s.t.} \quad & \sum_{j\in\mathcal{N}_i} w_{ij} = 1, \quad w_{ij} \geq 0 \end{aligned} \tag{2}$$

Each term in the objective function is defined as follows: $\mathcal{N}_i$ denotes the k-nearest neighbor set of sample x_i. $G_i \in R^{k\times k}$ is the local Gram difference matrix, defined as:

$$G_i(j,k) = \kappa_{ii} - \kappa_{ik} - \kappa_{ji} + \kappa_{jk}, \quad j,k \in \mathcal{N}_i \tag{3}$$

where $\kappa_{ij} = \langle x_i, x_j \rangle$ denotes the similarity between samples x_i and x_j computed using a kernel function.

H_{ij} is a constraint indicator defined as:

$$H_{ij} = \begin{cases} 1, & \text{if } (x_i, x_j) \in \mathcal{M} \cup \mathcal{C} \\ 0, & \text{otherwise} \end{cases} \tag{4}$$

λ controls the influence of the constraint alignment term, and μ is a regularization coefficient to prevent overfitting.

Once the neighbor weight vector $\mathbf{w}_i$ is learned for each sample x_i, we construct a new similarity matrix W, where each entry reflects the learned edge weight in the similarity graph:

$$W_{ij} = \begin{cases} w_{ij}, & x_j \in \mathcal{N}_i \\ 0, & \text{otherwise} \end{cases} \tag{5}$$

Although supervision is used to reconstruct the local graph, propagated signals remain mostly confined to samples similar to the initial supervised ones. Many gray samples, whose features differ substantially, receive lower confidence and are often ignored or mispropagated, degrading clustering performance. To address this, we propose an adversarial enhancement strategy to globally increase sample confidence.

Global Adversarial Propagation (GAP)

To address omissions or errors in propagating constraints among gray samples, we introduce a global dual-channel adversarial enhancement module. The main purpose of this module is to simultaneously propagate two semantically distinct types of constraints (Must-link and Cannot-link) while enforcing mutual exclusivity, so that each sample pair attains high confidence in at most one link. This effectively strengthens the propagation of constraints for gray samples.

To achieve this, we formulate the propagation as a dual-channel adversarial optimization problem:

$$\begin{gathered}\min_{\mathbf{D},\mathbf{S}} \| \mathbf{S} \odot \mathbf{D} \|_1 + \alpha\, \mathrm{Tr}(\mathbf{D}\mathbf{L}\mathbf{D}^\top) + \beta\, \mathrm{Tr}(\mathbf{S}\mathbf{L}\mathbf{S}^\top) \\ + \gamma \left(\| \mathbf{H} \odot (\mathbf{S} - \mathbf{M}) \|_F^2 + \| \mathbf{H} \odot (\mathbf{D} - \mathbf{C}) \|_F^2 \right) \\ \text{s.t. } \mathbf{D} > 0,\ \mathbf{S} > 0\end{gathered} \tag{6}$$

Here, $\mathbf{S}$ and $\mathbf{D} \in R^{n\times n}$ are the learned propagation graphs for must-link and cannot-link constraints, respectively. $\mathbf{M}$ and $\mathbf{C} \in R^{n\times n}$ denote the corresponding initial user intent matrices, defined similarly to Equation (1), and $\mathbf{H}_{ij}$ indicates whether a sample pair corresponds to the initial supervisory information.

The graph Laplacian $\mathbf{L}$ is constructed as $\mathbf{L} = \mathbf{A} - \mathbf{W}$, where $\mathbf{A}$ is the degree matrix and $\mathbf{W}$ is the k-nearest neighbor adjacency matrix. Hyperparameters α, β, and γ control the relative influence of the regularization, supervision, and adversarial terms.

Finally, the learned propagation confidence matrices $\mathbf{S}$ and $\mathbf{D}$ are optimized with an adversarial loss that ensures each sample pair exhibits high confidence in at most one matrix, while remaining consistent with the original graph structure. These matrices can be interpreted as global confidence estimators for the corresponding relational links, where higher values indicate a greater likelihood that the links should exist.

Enhanced Constraint Propagation (ECP)

To further refine the intent-aware similarity graph, we design an enhancement fusion mechanism. This module dynamically adjusts the edge weights of $\mathbf{W}$ based on the confidence comparison between $\mathbf{S}$ and $\mathbf{D}$, enabling selective enhancement or suppression of graph edges. Specifically, the edge update is computed as

$$\widetilde{\mathbf{W}}_{i,j} = \begin{cases} 1 - (1 - S_{i,j} + D_{i,j}) \cdot (1 - W_{i,j}) & \text{if } S_{i,j} > D_{i,j}, \\ (1 + S_{i,j} - D_{i,j}) \cdot W_{i,j} & \text{otherwise.} \end{cases} \tag{7}$$

We then adopt the Efficient Exhaustive Constraint Propagation (E^{2}CP) method, which models constraint propagation as a graph-based binary label propagation problem. This allows the enhanced propagation graph to extend user intents to unlabeled sample pairs, thereby globally propagating supervisory signals and improving constraint utilization. The propagation results are then discretized to generate pseudo-supervision information for semi-supervised clustering.

To ensure consistency in the graph structure, we first symmetrize the enhanced similarity matrix as $\widetilde{W} = \frac{1}{2}(\widetilde{W} + \widetilde{W}^\top)$.

The normalized graph Laplacian is then defined as

$$\widetilde{L} = I - \widetilde{D}^{-1/2}\widetilde{W}\widetilde{D}^{-1/2}, \tag{8}$$

where $\widetilde{D} = \mathrm{diag}(d_{ii})$ and $d_{ii} = \sum_{j=1}^{n} \widetilde{W}_{ij}$.

Let $t \in R^{n\times n}$ denote the initial user intent matrix. The vertical propagation result F_v is computed to propagate user intents along columns by minimizing

$$\min_{F_v} \tfrac{1}{2}\eta \|F_v - t\|_F^2 + \tfrac{1}{2}\mathrm{tr}(F_v^\top \widetilde{L} F_v), \tag{9}$$

which admits the closed-form solution

$$F_v = \eta(\eta I + \widetilde{L})^{-1} t. \tag{10}$$

Similarly, the horizontal propagation result F_h propagates intents along rows:

$$F_h = \eta t(\eta I + \widetilde{L})^{-1}. \tag{12}$$

By combining both vertical and horizontal directions, we obtain the final propagation result:

$$F = \eta^2(\eta I + \widetilde{L})^{-1} t (\eta I + \widetilde{L})^{-1}. \tag{13}$$

We treat F as a pseudo-constraint intent matrix, where F_{ij} indicates the propagated confidence for sample pair (x_i, x_j), normalized to $[-1, 1]$. The final pseudo-constraint matrix $\hat{\mathbf{R}}$ is obtained by thresholding F with τ and a dynamically adjusted secondary threshold $\hat{\tau}$ to control propagation quality:

$$\hat{\mathbf{R}}_{ij} = \begin{cases} 1, & F_{ij} > \tau + \hat{\tau}, \\ -1, & F_{ij} < -(\tau + \hat{\tau}), \\ 0, & \text{otherwise.} \end{cases} \tag{14}$$

3.3 IAR Module

To integrate the learned pseudo-user intents into clustering, we employ an intent regularization strategy that incorporates them into both representation learning (IRRL) and clustering (IRC), guiding the overall learning process.

Intent-Regularized Representation Learning (IRRL)

We build our model upon a deep auto-encoder (AE) architecture [Michelucci, 2022]. Through encoder–decoder pretraining, the AE learns a low-dimensional latent representation $\mathbf{Z}$ while reconstructing the input features $\hat{\mathbf{X}}$. The reconstruction loss is defined as

$$\mathcal{L}_{\mathrm{RE}} = \frac{1}{n}\sum_{i=1}^{n} \|\mathbf{x}_i - \hat{\mathbf{x}}_i\|^2, \tag{15}$$

where n denotes the number of samples, $\mathbf{x}_i$ is the input document representation, and $\hat{\mathbf{x}}_i$ is its reconstruction.

This module aims to incorporate propagated user intents into representation learning, encouraging the latent space to be consistent with user-specified clustering preferences. To

IJCAI-ECAI 2026 PREPRINT

this end, we construct an affinity matrix based on the learned latent representations $\mathbf{Z}$ as follows:

$$\mathbf{A} = \tilde{\mathbf{Z}} \cdot \tilde{\mathbf{Z}}^{\top}, \quad \tilde{\mathbf{Z}} = \mathrm{nor}(\mathbf{Z}) \tag{16}$$

where $\mathrm{nor}(\mathbf{Z})$ denotes the normalization operation applied to $\mathbf{Z}$ to reduce scale differences among features, and $a_{i,j}$ represents the (i,j)-th element of the affinity matrix $\mathbf{A}$.

Next, we process the learned pseudo-intent matrix $\hat{\mathbf{R}}$ to obtain a weighted affinity matrix $\mathbf{R} \in R^{n \times n}$, enabling better weighting of features across different dimensions:

$$\mathbf{R}_{i,j} = \begin{cases} 1, & \hat{\mathbf{R}}_{i,j} = 1 \\ a_{i,j}, & \hat{\mathbf{R}}_{i,j} = 0 \\ 0, & \hat{\mathbf{R}}_{i,j} = -1 \end{cases} \tag{17}$$

Subsequently, based on the pseudo-intent matrix learned through intent propagation, we dynamically expand the initial user intents to guide the representation learning process. The corresponding intent-regularized representation learning loss function is defined as:

$$\mathcal{L}_{\mathrm{IR}} = \|\mathbf{R} - \mathbf{A}\|_F = \sqrt{\sum_{i=1}^{n}\sum_{j=1}^{n}(R_{i,j} - a_{i,j})^2} \tag{18}$$

where $\|\cdot\|_F$ denotes the Frobenius norm, measuring the consistency between the learned representation features and the user-expanded intents.

In summary, the overall objective function of the IRRL module is defined as:

$$\mathcal{L}_{\mathrm{IRRL}} = \mathcal{L}_{\mathrm{RE}} + \lambda_1 \mathcal{L}_{\mathrm{IR}} \tag{19}$$

Where, λ_1 is a weighting parameter, set to 1.0 in our experiments, and $\mathcal{L}_{\mathrm{RE}}$ denotes the AE reconstruction loss. This module helps the model align learned feature representations with the expanded user intents.

Intent-Regularized Clustering (IRC)

This module builds upon representation learning by incorporating the propagated pseudo-user intents to guide cluster assignments, enabling the clustering structure to capture unique supervisory constraints among gray samples.

We first perform K-means clustering on the feature representations $\mathbf{Z}$ to obtain initial cluster centers $\{\boldsymbol{\mu}_k\}_{k=1}^{K}$, where K is the number of clusters. In each iteration, the predicted assignment matrix $\mathbf{Q}$ is computed as:

$$q_{ik} = \frac{\left(1 + \|\mathbf{z}_i - \boldsymbol{\mu}_k\|^2/\alpha\right)^{-\frac{\alpha+1}{2}}}{\sum_{k'}\left(1 + \|\mathbf{z}_i - \boldsymbol{\mu}_{k'}\|^2/\alpha\right)^{-\frac{\alpha+1}{2}}} \tag{20}$$

where q_{ik} is the probability of assigning $\mathbf{z}_i$ to cluster $\boldsymbol{\mu}_k$, with $\alpha = 1.0$ in our experiments.

Clustering is achieved by minimizing the KL divergence between $\mathbf{Q}$ and an auxiliary target distribution $\mathbf{P}$:

$$\mathcal{L}_{\mathrm{KL}} = \mathrm{KL}(\mathbf{P}\|\mathbf{Q}) = \sum_{i=1}^{n}\sum_{k=1}^{K} p_{ik}\log\frac{p_{ik}}{q_{ik}}, \tag{21}$$

where the target probabilities are

$$p_{ik} = \frac{q_{ik}^2/f_k}{\sum_{k'=1}^{K} q_{ik'}^2/f_{k'}}, \tag{22}$$

where $f_k = \sum_{i=1}^{n} q_{ik}$, and the final cluster assignment is $c_i = \arg\max_k q_{ik}$. To guide the generation of clustering structures using pseudo-user intents, we construct a cluster consistency matrix $\mathbf{B}$ from the assignment matrix $\mathbf{Q}$ as:

$$\mathbf{B} = \tilde{\mathbf{Q}} \cdot \tilde{\mathbf{Q}}^{\top}, \quad \tilde{\mathbf{Q}} = \mathrm{nor}(\mathbf{Q}). \tag{23}$$

To incorporate the propagated pseudo-user intents into the clustering process, we adopt the same intent regularization strategy used in the representation learning stage, and transform the pseudo-intent matrix $\mathbf{R}$ as follows:

$$R'_{ij} = \begin{cases} 1, & \hat{\mathbf{R}}_{i,j} = 1 \\ b_{ij}, & \hat{\mathbf{R}}_{i,j} = 0 \\ 0, & \hat{\mathbf{R}}_{i,j} = -1 \end{cases} \tag{24}$$

We then define the intent-regularized clustering loss $\mathcal{L}_{\mathrm{IC}}$ as:

$$\mathcal{L}_{\mathrm{IC}} = \|\mathbf{R}' - \mathbf{B}\|_F = \sqrt{\sum_{i=1}^{n}\sum_{j=1}^{n}\left(R'_{ij} - \bar{b}_{ij}\right)^2}. \tag{25}$$

This loss measures the consistency between the propagated pseudo-intents and the clustering assignments, enabling the model to accurately capture the supervisory constraints and achieve more precise clustering results. The overall loss of the IRC module is then defined as:

$$\mathcal{L}_{\mathrm{IRC}} = \mathcal{L}_{\mathrm{KL}} + \lambda_2 \mathcal{L}_{\mathrm{IC}}, \tag{26}$$

where λ_2 is set to 1.0 in our experiments. When $\lambda_2 = 0$, the model degenerates to standard clustering without pseudo-user intent guidance.

Dataset	Instances	Classes	Dimension
ACM	3,025	3	1,870
Abstract	4,306	3	10,000
Reuters-10k	10,000	4	2,000
Citeseer	3,327	6	3,703
USPS	9,298	10	256

Table 1: Dataset specifications with instance counts, class numbers, and feature dimensions.

4 Experiments

4.1 Experiments Settings

Datasets. We evaluate our model on five datasets. **Reuters-10k** contains 10,000 news articles across multiple topics. **ACM** consists of academic papers from various journals, and **Abstract** includes paper abstracts from the AMiner platform. **USPS** has 9,298 handwritten digit images, and **Citeseer** comprises scientific publications from six categories, represented by word vectors and annotated with citation links.

IJCAI-ECAI 2025 PREPRINT

Methods	ACM			Abstract			Reuters-10k			Citeseer			USPS		
	ACC	NMI	ARI	ACC	NMI	ARI	ACC	NMI	ARI	ACC	NMI	ARI	ACC	NMI	ARI
K-means [McQueen, 1967]	67.31	32.44	30.60	69.18	38.26	27.69	41.83	31.58	11.45	39.32	16.94	13.43	66.80	62.60	54.60
COP-Kmeans [Wagstaff *et al.*, 2001]	75.02	48.71	51.39	77.63	65.14	70.03	65.84	42.53	50.49	45.05	21.41	19.17	58.90	71.85	70.24
IDEC [Guo *et al.*, 2017]	85.13	56.16	62.16	83.83	60.98	62.03	75.03	51.98	58.26	60.49	27.17	25.70	76.20	75.60	67.90
GAE [Kipf and Welling, 2016]	84.52	55.38	59.46	87.37	59.00	65.32	54.40	25.92	19.61	61.35	34.63	33.55	63.10	60.70	50.30
SDCN [Bo *et al.*, 2020]	90.45	68.31	73.91	90.78	68.14	74.05	77.15	50.28	55.36	65.96	38.71	40.17	78.08	79.51	71.84
SDEC [Ren *et al.*, 2019]	85.53	55.37	61.43	90.78	68.14	74.05	77.15	50.28	55.36	59.73	40.17	32.10	76.39	77.68	69.95
DCC [Zhang *et al.*, 2019]	81.40	44.39	50.70	91.10	67.87	74.87	82.91	58.28	66.55	63.50	35.82	13.58	76.72	71.21	68.80
ADC [Sun *et al.*, 2022]	62.34	48.94	42.25	94.33	78.87	83.68	79.81	54.89	60.17	37.39	13.77	13.76	77.23	78.85	69.89
TBDSC [Duan *et al.*, 2024]	80.86	49.28	51.04	88.85	63.92	69.09	<u>90.22</u>	<u>69.76</u>	<u>80.94</u>	60.32	30.60	29.76	<u>85.61</u>	<u>86.53</u>	<u>82.88</u>
IPDC [Xu *et al.*, 2025]	<u>92.56</u>	<u>71.90</u>	<u>78.80</u>	<u>95.55</u>	<u>81.52</u>	<u>86.88</u>	87.65	65.41	71.16	<u>69.04</u>	<u>42.89</u>	<u>45.21</u>	83.73	85.43	80.15
AEPIC (Ours)	**94.48**	**78.31**	**84.15**	**96.70**	**85.11**	**90.25**	**92.70**	**77.68**	**84.37**	**72.32**	**45.89**	**49.33**	**85.75**	**87.09**	**83.03**
Variants of AEPIC															
AEPIC w/o LSR	93.72	76.17	82.06	96.42	84.04	89.46	90.05	74.16	79.24	67.81	43.51	44.65	84.81	86.47	81.72
AEPIC w/o GAP	93.88	76.39	82.56	96.26	83.63	89.00	91.74	76.21	82.82	70.69	42.97	46.53	85.11	86.66	82.46
AEPIC (Full)	**94.48**	**78.31**	**84.15**	**96.70**	**85.11**	**90.25**	**92.70**	**77.68**	**84.37**	**72.32**	**45.89**	**49.33**	**85.75**	**87.09**	**83.03**

Table 2: The table shows the clustering results of eleven methods on 5 datasets with an initial set of 500 pairwise constraints. The second-best results are underlined, and the best results are highlighted in bold. Results for variants of AEPIC are also included.

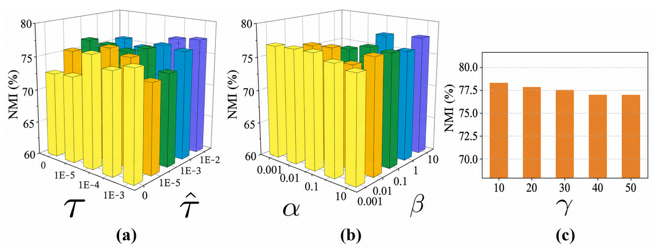

Figure 3: NMI scores under different hyperparameters with an initial set of 500 pairwise constraints on the ACM dataset.

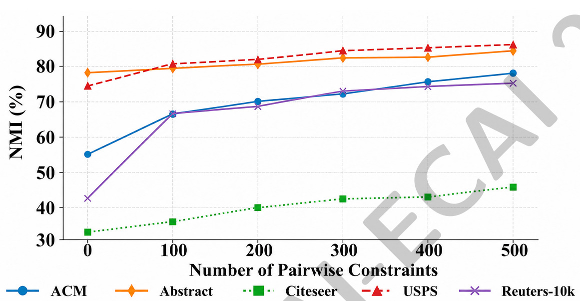

Figure 4: Clustering performance of AEPIC under different numbers of initial pairwise constraints across all experimental datasets.

Baselines.We compare our method with several state-of-the-art baselines, including classical clustering, deep unsupervised clustering, and deep semi-supervised clustering methods. The classical clustering methods are K-means [McQueen, 1967] and COP-Kmeans [Wagstaff *et al.*, 2001]. Deep unsupervised clustering methods include IDEC [Guo *et al.*, 2017], GAE [Kipf and Welling, 2016], and SDCN [Bo *et al.*, 2020]. Deep semi-supervised clustering methods consist of SDEC [Ren *et al.*, 2019], CPAC [Fogel *et al.*, 2019], DCC [Zhang *et al.*, 2019], ADC [Sun *et al.*, 2022], TBDSC [Duan *et al.*, 2024], and IPDC [Xu *et al.*, 2025].

evaluation metrics.We evaluate clustering performance using standard metrics: clustering accuracy (ACC) [Xu *et al.*, 2003], normalized mutual information (NMI) [Strehl *et al.*, 2000], and adjusted Rand index (ARI), where higher values indicate better results.

4.2 Comparison with State-of-the-Art

To evaluate AEPIC's clustering performance, we compare it with representative baselines using 500 initial pairwise constraints, evenly split between must-link and cannot-link pairs.

Table 2 presents results on 5 benchmark datasets. With the same constraints, AEPIC consistently outperforms classical deep clustering and recent semi-supervised methods, demonstrating the effectiveness and robustness of our constraint propagation strategy.

For example, on Reuters-10k, with only 500 constraints, AEPIC surpasses the best semi-supervised baseline by 2.75%, 11.35%, and 4.24% in ACC, NMI, and ARI, respectively, highlighting the efficacy of user intent expansion under limited supervision.

Table 2 (AEPIC variants) further compares the performance of different model variants. Here, AEPIC w/o ISR and AEPIC w/o GAP denote the removal of the reconstruction and adversarial modules, respectively. The results demonstrate that both modules are essential for achieving improved clustering performance.

4.3 Ablation Studies

We conducted ablation experiments on the AEPIC framework and its AECP module. In these experiments, the AECP constraint propagation was disabled in each of the two intent regularization stages, as shown in Table 3 ("CP" indicates the module is enabled). The results demonstrate that AECP plays a critical role in both intent regularization stages, and propagating user supervision significantly improves clustering performance.

4.4 Parameter Study Analysis

Sensitivity analyses on the ACM dataset evaluate AEPIC's stability across dual thresholds $(\tau, \hat{\tau})$, adversarial weights (α, β), and the consistency coefficient γ.As shown in

IJCAI-ECAI 2026 PREPRINT

PRELIMINARY PREPRINT VERSION – DO NOT CITE
This is not the version of record. The final version will appear in the Proceedings of IJCAI-ECAI 2026.

LOSS	ACM			Abstract			Reuters-10k			Citeseer			USPS		
	ACC	NMI	ARI	ACC	NMI	ARI	ACC	NMI	ARI	ACC	NMI	ARI	ACC	NMI	ARI
$L_{IRRL} + L_{IRC}$	92.79	72.55	79.61	96.49	84.40	89.67	81.87	68.65	65.40	67.84	42.68	43.91	83.72	85.41	80.13
L_{IRRL}(CP) $+ L_{IRC}$	94.38	77.96	83.90	96.66	85.06	90.12	84.33	70.07	68.35	70.00	44.53	46.25	84.35	86.16	81.04
$L_{IRRL} + L_{IRC}$(CP)	93.36	74.55	81.14	96.58	84.63	89.92	88.27	72.03	76.25	70.27	44.45	47.26	84.76	85.92	81.44
L_{IRRL}(CP) $+ L_{IRC}$(CP)	**94.48**	**78.23**	**84.15**	**96.70**	**85.11**	**90.26**	**92.70**	**77.68**	**84.37**	**72.32**	**45.89**	**49.33**	**85.75**	**87.09**	**83.03**

Table 3: Ablation study of AEPIC on Intent-Consistent Loss (L_{IRRL}) and Intent-Regularized Loss (L_{IRC}) across multiple datasets. "CP" denotes the incorporation of constraint propagation.

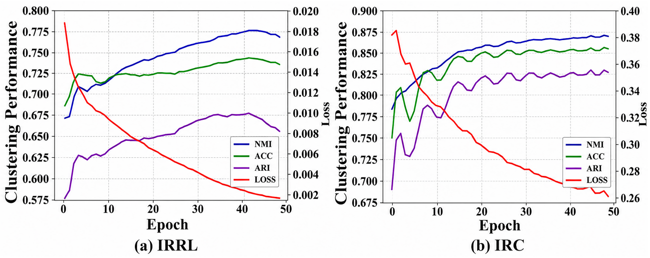

Figure 5: Convergence curves of the two intent regularization stages on the USPS dataset

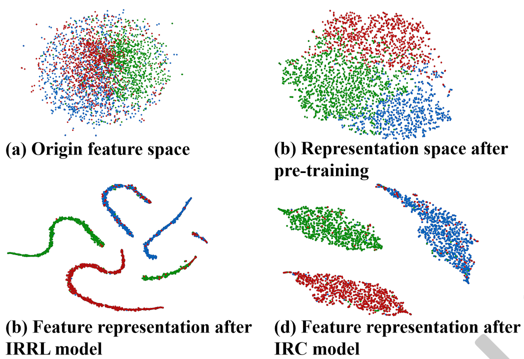

Figure 6: The ACM dataset is visualized using t-SNE.

Fig. 3(a), curriculum-based τ [Wang *et al.*, 2022] balances convergence speed and pseudo-constraint quality, while $\hat{\tau}$ further refines selection. Optimal NMI occurs at $\tau = 1 \times 10^{-4}$ (avoiding noise and sparsity), with smaller $\hat{\tau}$ consistently enhancing performance. Figs. 3(b) and 3(c) demonstrate the robustness of AEPIC's, showing minimal variation across wide ranges of $\alpha, \beta \in [10^{-3}, 10^{1}]$ and $\gamma \in [10, 50]$.

4.5 Analysis of Initial Constraints Experiments

To assess the impact of initial pairwise constraints on AEPIC, we conducted clustering experiments under varying settings. Notably, without initial constraints, the model receives no supervision.

As shown in Figure 4, even a small number of constraints (e.g., 100 pairs) substantially improves clustering, as propagation amplifies their value. Performance continues to rise with more constraints, but the improvement rate slows. This indicates that, as the number of constraints grows, the high-value information available for further optimization gradually decreases, and the model's reliance on additional constraints diminishes, leading to slower performance gains.

4.6 Convergence Analysis

Figure 5 illustrates the convergence curves of the two intent regularization modules in our model, namely the IRRL and IRC modules, visualized on the USPS dataset. In the IRRL module (Figure5 a), the clustering metrics exhibit noticeable fluctuations in the later training stages, likely due to the unstable quality of propagated pseudo-user constraints and their ineffective incorporation into the clustering structure. In contrast, the IRC module (Figure5 b) shows significantly improved convergence after approximately 30 training epochs, resulting in more stable clustering performance.

4.7 Visualization Results

Figure 6 illustrates the evolution of AEPIC's feature space on the ACM dataset: (a) the original features are sparse with unclear cluster boundaries; (b) pretraining enhances inter-cluster separability and feature structure; (c) intent-aware regularization during representation learning pulls similar samples together and separates dissimilar ones, producing clusters aligned with the underlying manifold; (d) further regularization during clustering yields more compact, well-separated clusters, guided by additional propagated constraints. These results show that the learned pseudo-intents effectively capture true semantics, improving the discrimination of ambiguous and boundary samples and enhancing overall clustering performance.

5 Conclusion

This paper proposes AEPIC, a semi-supervised clustering framework that incorporates user intents into both representation learning and clustering. To better capture ambiguous samples that deviate from supervision, we introduce the Adversarially Enhanced Constraint Propagation (AECP) module, which combines local structure reconstruction with global adversarial enhancement to improve constraint propagation and generate higher-quality pseudo-user intents. Experiments on multiple datasets show that AEPIC achieves superior clustering performance, with future work extending it to multimodal and large-scale scenarios.

Acknowledgments

This research was supported by the National Key R&D Program of China (No.2023YFC3304500), and the Guizhou Provincial Science and Technology Department Project (No. Qiankehe MS[2026] 077).

IJCAI-ECAI 2026 PREPRINT

References

[Bai *et al.*, 2020] Liang Bai, Junbin Wang, Jiye Liang, and Hangyuan Du. New label propagation algorithm with pairwise constraints. *Pattern Recognition*, 106:107411, 2020.

[Bai *et al.*, 2021] Ruina Bai, Ruizhang Huang, Yanping Chen, and Yongbin Qin. Deep multi-view document clustering with enhanced semantic embedding. *Information Sciences*, 564:273–287, 2021.

[Bair, 2013] Eric Bair. Semi-supervised clustering methods. *Wiley Interdisciplinary Reviews: Computational Statistics*, 5(5):349–361, 2013.

[Barták, 2001] Roman Barták. Theory and practice of constraint propagation. In *Proceedings of the 3rd Workshop on Constraint Programming in Decision and Control*, volume 50, 2001.

[Bo *et al.*, 2020] Deyu Bo, Xiao Wang, Chuan Shi, Meiqi Zhu, Emiao Lu, and Peng Cui. Structural deep clustering network. In *Proceedings of the web conference 2020*, pages 1400–1410, 2020.

[Cai *et al.*, 2023] Jianghui Cai, Jing Hao, Haifeng Yang, Xujun Zhao, and Yuqing Yang. A review on semi-supervised clustering. *Information Sciences*, 632:164–200, 2023.

[Chavoshinejad *et al.*, 2023] Jovan Chavoshinejad, Seyed Amjad Seyedi, Fardin Akhlaghian Tab, and Navid Salahian. Self-supervised semi-supervised nonnegative matrix factorization for data clustering. *Pattern Recognition*, 137:109282, 2023.

[Cheng *et al.*, 2008] Hao Cheng, Kien A Hua, and Khanh Vu. Constrained locally weighted clustering. *Proceedings of the VLDB Endowment*, 1(1):90–101, 2008.

[Duan *et al.*, 2024] Yu Duan, Zhoumin Lu, Rong Wang, Xuelong Li, and Feiping Nie. Toward balance deep semisupervised clustering. *IEEE Transactions on Neural Networks and Learning Systems*, 2024.

[Eaton *et al.*, 2010] Eric Eaton, Marie Desjardins, and Sara Jacob. Multi-view clustering with constraint propagation for learning with an incomplete mapping between views. In *Proceedings of the 19th ACM international conference on Information and knowledge management*, pages 389–398, 2010.

[Fogel *et al.*, 2019] Sharon Fogel, Hadar Averbuch-Elor, Daniel Cohen-Or, and Jacob Goldberger. Clustering-driven deep embedding with pairwise constraints. *IEEE computer graphics and applications*, 39(4):16–27, 2019.

[Fu and Lu, 2015] Zhenyong Fu and Zhiwu Lu. Pairwise constraint propagation: A survey. *arXiv preprint arXiv:1502.05752*, 2015.

[González-Almagro *et al.*, 2025] Germán González-Almagro, Daniel Peralta, Eli De Poorter, José-Ramón Cano, and Salvador García. Semi-supervised constrained clustering: An in-depth overview, ranked taxonomy and future research directions. *Artificial Intelligence Review*, 58(5):157, 2025.

[Guo *et al.*, 2017] Xifeng Guo, Long Gao, Xinwang Liu, and Jianping Yin. Improved deep embedded clustering with local structure preservation. In *Ijcai*, volume 17, pages 1753–1759, 2017.

[Hsu and Kira, 2015] Yen-Chang Hsu and Zsolt Kira. Neural network-based clustering using pairwise constraints. *arXiv preprint arXiv:1511.06321*, 2015.

[Huang *et al.*, 2008] Haichao Huang, Yong Cheng, and Ruilian Zhao. A semi-supervised clustering algorithm based on must-link set. In *International Conference on Advanced Data Mining and Applications*, pages 492–499. Springer, 2008.

[Jia *et al.*, 2020] Yuheng Jia, Hui Liu, Junhui Hou, and Sam Kwong. Pairwise constraint propagation with dual adversarial manifold regularization. *IEEE transactions on neural networks and learning systems*, 31(12):5575–5587, 2020.

[Kipf and Welling, 2016] Thomas N Kipf and Max Welling. Variational graph auto-encoders. *arXiv preprint arXiv:1611.07308*, 2016.

[Li *et al.*, 2008] Zhenguo Li, Jianzhuang Liu, and Xiaoou Tang. Pairwise constraint propagation by semidefinite programming for semi-supervised classification. In *Proceedings of the 25th international conference on Machine learning*, pages 576–583, 2008.

[Lin *et al.*, 2023] Mugang Lin, Kunhui Wen, Xuanying Zhu, Huihuang Zhao, and Xianfang Sun. Graph autoencoder with preserving node attribute similarity. *Entropy*, 25(4):567, 2023.

[Lu and Carreira-Perpinan, 2008] Zhengdong Lu and Miguel A Carreira-Perpinan. Constrained spectral clustering through affinity propagation. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2008.

[Lu and Peng, 2013a] Zhiwu Lu and Yuxin Peng. Exhaustive and efficient constraint propagation: A graph-based learning approach and its applications. *International journal of computer vision*, 103(3):306–325, 2013.

[Lu and Peng, 2013b] Zhiwu Lu and Yuxin Peng. Unified constraint propagation on multi-view data. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 27, pages 640–646, 2013.

[Maggini *et al.*, 2012] Marco Maggini, Stefano Melacci, and Lorenzo Sarti. Learning from pairwise constraints by similarity neural networks. *Neural Networks*, 26:141–158, 2012.

[Masud *et al.*, 2019] Md Abdul Masud, Joshua Zhexue Huang, Ming Zhong, and Xianghua Fu. Generate pairwise constraints from unlabeled data for semi-supervised clustering. *Data & Knowledge Engineering*, 123:101715, 2019.

[McQueen, 1967] James B McQueen. Some methods of classification and analysis of multivariate observations. In *Proc. of 5th Berkeley Symposium on Math. Stat. and Prob.*, pages 281–297, 1967.

IJCAI-ECAI 2026 PREPRINT

[Michelucci, 2022] Umberto Michelucci. An introduction to autoencoders. *arXiv preprint arXiv:2201.03898*, 2022.

[Qin *et al.*, 2024] Xiao Qin, Changan Yuan, Jianhui Jiang, and Long Chen. Deep semi-supervised clustering based on pairwise constraints and sample similarity. *Pattern Recognition Letters*, 178:1–6, 2024.

[Ren *et al.*, 2019] Yazhou Ren, Kangrong Hu, Xinyi Dai, Lili Pan, Steven CH Hoi, and Zenglin Xu. Semi-supervised deep embedded clustering. *Neurocomputing*, 325:121–130, 2019.

[Ren *et al.*, 2021] Pengzhen Ren, Yun Xiao, Xiaojun Chang, Po-Yao Huang, Zhihui Li, Brij B Gupta, Xiaojiang Chen, and Xin Wang. A survey of deep active learning. *ACM computing surveys (CSUR)*, 54(9):1–40, 2021.

[Ren *et al.*, 2024] Yazhou Ren, Jingyu Pu, Zhimeng Yang, Jie Xu, Guofeng Li, Xiaorong Pu, Philip S Yu, and Lifang He. Deep clustering: A comprehensive survey. *IEEE transactions on neural networks and learning systems*, 36(4):5858–5878, 2024.

[Śmieja *et al.*, 2020] Marek Śmieja, Łukasz Struski, and Mário AT Figueiredo. A classification-based approach to semi-supervised clustering with pairwise constraints. *Neural Networks*, 127:193–203, 2020.

[Strehl *et al.*, 2000] Alexander Strehl, Joydeep Ghosh, and Raymond Mooney. Impact of similarity measures on web-page clustering. In *Workshop on artificial intelligence for web search (AAAI 2000)*, volume 58, page 64, 2000.

[Sun *et al.*, 2022] Bicheng Sun, Peng Zhou, Liang Du, and Xuejun Li. Active deep image clustering. *Knowledge-Based Systems*, 252:109346, 2022.

[Sun, 2025] X. Sun. Semi-supervised clustering via constraints self-learning. *Mathematics*, 13(9):1535, 2025.

[Van Craenendonck *et al.*, 2018] Toon Van Craenendonck, Sebastijan Dumancic, and Hendrik Blockeel. Cobra: A fast and simple method for active clustering with pairwise constraints. *arXiv preprint arXiv:1801.09955*, 2018.

[Wagstaff *et al.*, 2001] Kiri Wagstaff, Claire Cardie, Seth Rogers, Stefan Schrödl, et al. Constrained k-means clustering with background knowledge. In *Icml*, volume 1, pages 577–584, 2001.

[Wang and Zhang, 2006] Fei Wang and Changshui Zhang. Label propagation through linear neighborhoods. In *Proceedings of the 23rd international conference on Machine learning*, pages 985–992, 2006.

[Wang *et al.*, 2015] Di Wang, Xinbo Gao, and Xiumei Wang. Semi-supervised nonnegative matrix factorization via constraint propagation. *IEEE transactions on cybernetics*, 46(1):233–244, 2015.

[Wang *et al.*, 2022] Yidong Wang, Hao Chen, Qiang Heng, Wenxin Hou, Yue Fan, Zhen Wu, Jindong Wang, Marios Savvides, Takahiro Shinozaki, Bhiksha Raj, et al. Freematch: Self-adaptive thresholding for semi-supervised learning. *arXiv preprint arXiv:2205.07246*, 2022.

[Wang *et al.*, 2023] Yalin Wang, Jiangfeng Zou, Kai Wang, Chenliang Liu, and Xiaofeng Yuan. Semi-supervised deep embedded clustering with pairwise constraints and subset allocation. *Neural Networks*, 164:310–322, 2023.

[Wu *et al.*, 2024] Song Wu, Yan Zheng, Yazhou Ren, Jing He, Xiaorong Pu, Shudong Huang, Zhifeng Hao, and Lifang He. Self-weighted contrastive fusion for deep multi-view clustering. *IEEE Transactions on Multimedia*, 26:9150–9162, 2024.

[Xie *et al.*, 2016] Junyuan Xie, Ross Girshick, and Ali Farhadi. Unsupervised deep embedding for clustering analysis. In *International conference on machine learning*, pages 478–487. PMLR, 2016.

[Xu *et al.*, 2003] Wei Xu, Xin Liu, and Yihong Gong. Document clustering based on non-negative matrix factorization. In *Proceedings of the 26th annual international ACM SIGIR conference on Research and development in informaion retrieval*, pages 267–273, 2003.

[Xu *et al.*, 2023] Xiao Xu, Haiwei Hou, and Shifei Ding. Semi-supervised deep density clustering. *Applied Soft Computing*, 148:110903, 2023.

[Xu *et al.*, 2025] Le Xu, Ruina Bai, Ruizhang Huang, Lina Ren, and Yongbin Qin. Discovering personalized document partition via intention-aware deep document clustering. *Knowledge-Based Systems*, 317:113407, 2025.

[Yang *et al.*, 2019] Xu Yang, Cheng Deng, Feng Zheng, Junchi Yan, and Wei Liu. Deep spectral clustering using dual autoencoder network. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4066–4075, 2019.

[Yang *et al.*, 2025] Xiaojun Yang, Tuoji Zhu, Siyuan Peng, Feiping Nie, and Zhiping Lin. Semi-supervised pivotal-aware nonnegative matrix factorization with label and pairwise constraint propagation for data clustering. *Pattern Recognition*, 157:110933, 2025.

[Yin *et al.*, 2023] Jingxing Yin, Siyuan Peng, Zhijing Yang, Badong Chen, and Zhiping Lin. Hypergraph based semi-supervised symmetric nonnegative matrix factorization for image clustering. *Pattern Recognition*, 137:109274, 2023.

[Yu and Shi, 2004] Stella X Yu and Jianbo Shi. Segmentation given partial grouping constraints. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(2):173–183, 2004.

[Zhang *et al.*, 2019] Hongjing Zhang, Sugato Basu, and Ian Davidson. A framework for deep constrained clustering-algorithms and advances. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 57–72. Springer, 2019.

[Zhou *et al.*, 2003] Dengyong Zhou, Olivier Bousquet, Thomas Lal, Jason Weston, and Bernhard Schölkopf. Learning with local and global consistency. *Advances in neural information processing systems*, 16, 2003.

[Zhu and Goldberg, 2009] Xiaojin Zhu and Andrew Goldberg. *Introduction to semi-supervised learning*. Morgan & Claypool Publishers, 2009.

IJCAI-ECAI 2026 PREPRINT