KNOWN-NOVEL RELATION LEARNING FOR MULTI-VIEW NOVEL CLASS DISCOVERY

Anonymous authors
Paper under double-blind review

ABSTRACT

Novel Class Discovery (NCD) aims to cluster unlabeled samples from novel classes by transferring knowledge from labeled known classes, where the known and novel label spaces are disjoint. Existing NCD methods have achieved promising progress in single-view scenarios, but they are less effective for real-world multi-view data, in which different views provide both consistent and complementary information. Although recent multi-view NCD methods exploit distributional correlations between known and novel classes and perform adaptive view fusion, the similarity relations between known and novel categories remain insufficiently explored. In fact, novel classes are not isolated from known classes in feature or semantic spaces; their relation patterns with respect to known categories can provide valuable structural references for novel-class partitioning. To address this issue, this paper proposes a Known-Novel Relation Learning framework for Multi-view Novel Class Discovery. Specifically, we first employ multi-view matrix factorization to learn view-specific latent structures coupled by a shared partition representation. By constraining its known-class part with ground-truth labels, discriminative supervision is injected into multi-view representation learning and adaptive view-weight estimation, enabling distributional knowledge transfer to novel classes. Furthermore, known-class prototypes are regarded as relation anchors to construct a known-novel relation matrix, which represents novel samples in the reference space defined by known categories. By associating sample-level anchor relations with novel-class partitioning, samples assigned to the same novel class are encouraged to preserve consistent relation patterns, thereby explicitly transferring relational knowledge to novel-class clustering. These components are integrated into a unified optimization framework with an efficient iterative solver. Experiments on multiple multi-view datasets demonstrate the effectiveness of the proposed method.

1 INTRODUCTION

Humans usually possess the ability to acquire new knowledge from existing knowledge (Yue et al., 2022; Wang et al., 2025; Yu et al., 2026). Specifically, when encountering unseen objects, humans can often understand and distinguish them by referring to the information of known categories. To endow learning models with a similar capability in open environments, Novel Class Discovery (NCD) has been proposed (Hsu et al., 2018; 2019; Li et al., 2025). This task aims to leverage labeled data from known classes to cluster unlabeled data from novel classes, where the label spaces of the two datasets are disjoint (Han et al., 2019; 2020).

With the increasing demand for open-environment learning, NCD has received extensive attention and has developed rapidly in recent years (Hou et al., 2024). Existing methods usually transfer knowledge from known classes to novel-class clustering through representation learning (Zhong et al., 2021), pseudo-label generation (Li et al., 2023), and relation-based knowledge transfer (Gu et al., 2023). These methods have achieved promising performance in single-view scenarios, but most of them overlook the multi-view nature that commonly exists in real-world data. In practical applications, the same object can often be described by multiple views. Different views contain not only cross-view consistent information but also view-specific discriminative information (Cui et al., 2024; Wu et al., 2024; Zhang et al., 2025a). Directly applying single-view NCD methods to multi-view data, or simply concatenating multi-view features, may fail to sufficiently exploit the con-

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.

sistency and complementarity among different views, thereby limiting the clustering performance on novel classes.

Although existing multi-view clustering methods can fuse information from different views (Li et al., 2022b; Wan et al., 2024b), they are mostly developed under an unsupervised assumption and thus cannot effectively utilize the label information from known classes (Wang et al., 2022; Ren et al., 2026; Kong et al., 2026). Wan et al. (2025) made a preliminary attempt to study multi-view novel class discovery, where matrix factorization was adopted to model the distributional correlation between known and novel classes, and view-weight learning was used to achieve multi-view information fusion. However, this method mainly exploits known-class information from the perspectives of multi-view representation learning and distributional knowledge transfer, while the similarity relations between known and novel classes have not been explicitly investigated.

The disjointness of label spaces does not imply that categories are completely independent in feature distributions or semantic structures. Novel classes may share potential similarities with certain known classes in terms of appearance, local structures, or high-level semantics (Li et al., 2023; Gu et al., 2023; Ma et al., 2025). Such similarities can provide important structural references for novel-class clustering. For example, in a visual recognition scenario where cats and dogs are known classes while wolves and tigers are novel classes, wolves usually exhibit stronger appearance-level similarity to dogs, whereas tigers show a closer visual relation to cats. In this case, the similarities of novel samples with respect to known classes can provide additional clues for distinguishing different novel categories. Therefore, how to further exploit known-novel similarities from the perspective of relation modeling and use them to guide novel-class partitioning remains an important problem in multi-view novel class discovery.

In light of this, this paper proposes a framework named Known-Novel Relation Learning for Multi-view Novel Class Discovery. Specifically, we first employ multi-view matrix factorization to learn view-specific latent structures coupled by a shared partition representation, so as to capture distributional consistency across known and novel classes. Meanwhile, by constraining the shared partition representation with known-class labels, the discriminative information from known classes can participate in multi-view representation learning and view-weight estimation, thus enabling distributional knowledge transfer from known classes to novel classes. Furthermore, we regard known-class prototypes as relation anchors and construct a relation matrix from novel samples to known-class anchors. This matrix characterizes the relational structure of novel samples in the reference space defined by known classes. We then associate the sample-level anchor relations with novel-class partitioning, encouraging samples belonging to the same novel class to preserve consistent relation patterns in the known-class anchor space. Through this design, the relation information between known and novel classes can be explicitly transferred to the clustering process of novel classes. Finally, the above components are integrated into a unified optimization framework.

The main contributions of this paper are summarized as follows:

1. To the best of our knowledge, this is the first attempt to explicitly exploit known–novel relational knowledge for multi-view novel class discovery. We reveal that the relation patterns of novel samples with respect to known classes provide complementary structural cues for novel-class partitioning.
2. We propose a prototype-based known–novel relation learning mechanism, where known-class prototypes serve as relation anchors to construct sample-level relation representations and cluster-level relation prototypes, thereby explicitly transferring relational knowledge to novel-class clustering.
3. We design an effective iterative optimization algorithm to solve the proposed unified objective function, and experimental results demonstrate that the proposed method achieves superior clustering performance compared with existing multi-view clustering and novel class discovery methods.

2 Related Work

Multi-view clustering. Multi-view clustering (MVC) integrates complementary information from multiple views to obtain a shared partition. Matrix-factorization methods learn view-specific latent factors and consensus representations (Fu et al., 2020; Liu et al., 2013; Wang et al., 2018; Gao et al.,

2019; Yang et al., 2021), while one-step and anchor-based methods further couple representation learning with clustering and improve scalability (Nie et al., 2020; Wan et al., 2024a; Li et al., 2022a; Yu et al., 2025). Recent deep and adaptive MVC methods also emphasize cross-view consistency and view reliability (Cui et al., 2024; Wu et al., 2024; Zhang et al., 2025a). However, these methods are predominantly unsupervised and cannot directly exploit labeled known classes for knowledge transfer to unseen categories.

Novel class discovery. Novel Class Discovery (NCD) clusters unlabeled samples from novel categories by transferring knowledge from labeled known categories with disjoint label spaces (Han et al., 2019; 2020; Zhong et al., 2021). Recent studies extend NCD to semantic segmentation, class-incremental discovery, and graph data (Joseph et al., 2022; Zhao et al., 2022; Roy et al., 2022; Xu et al., 2024). Relation-aware approaches model intra-class, inter-class, or class-level structure to improve known-to-novel transfer (Li et al., 2023; Gu et al., 2023), but are designed for a single feature space. IICMVNCD (Wan et al., 2025) studies the multi-view setting through distributional correlation and adaptive view fusion. In contrast, our method explicitly represents each novel sample by its relations to known-class prototypes in every view and couples these relations with the novel partition.

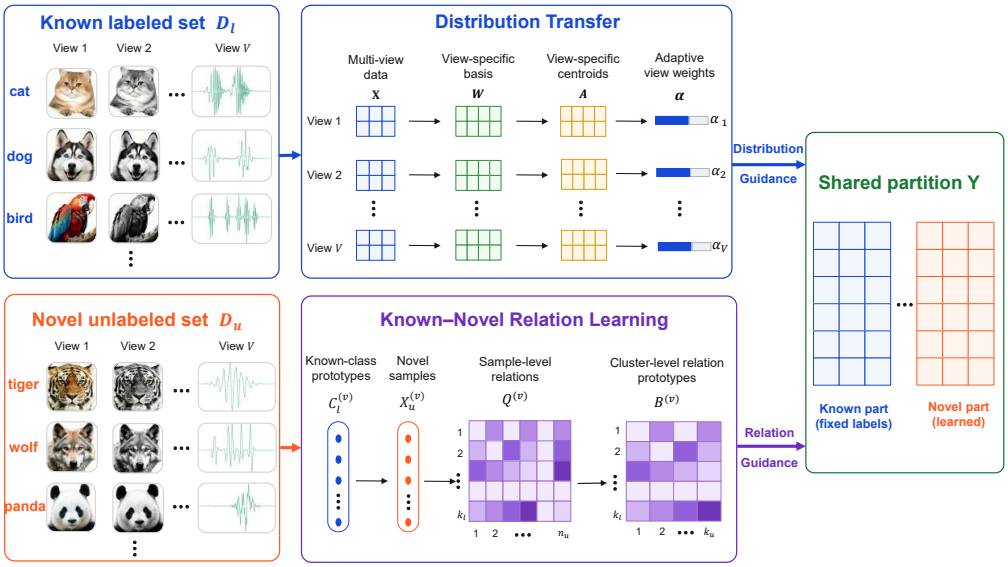


Figure 1: Framework of the proposed algorithm: We first exploit multi-view matrix factorization to transfer distributional knowledge from known classes to novel classes under known-class supervision and adaptive view weighting. Then, known-class prototypes are regarded as relation anchors to construct sample-level known–novel relations $\mathbf{Q}^{(v)}$, which are further summarized into cluster-level relation prototypes $\mathbf{B}^{(v)}$ to guide novel-class clustering. Finally, the distribution guidance and relation guidance are jointly incorporated into the shared partition $\mathbf{Y}$ for novel class discovery.

3 METHOD

3.1 OVERVIEW

Following the MVNCD setting introduced in (Wan et al., 2025), we consider a labeled known-class multi-view dataset and an unlabeled novel-class multi-view dataset with disjoint label spaces. Let $\mathbf{X}^v = [\mathbf{X}_l^v, \mathbf{X}_u^v] \in \mathbb{R}^{d_v \times n}$ denote the data matrix of the v-th view, where $\mathbf{X}_l^v \in \mathbb{R}^{d_v \times n_l}$ and $\mathbf{X}_u^v \in \mathbb{R}^{d_v \times n_u}$ correspond to the labeled and unlabeled samples, respectively, and $n = n_l + n_u$. The labeled set contains k_l known classes with the ground-truth label matrix $\mathbf{G}_l \in \{0,1\}^{k_l \times n_l}$, while the unlabeled set consists of samples from k_u novel classes. The goal is to infer the cluster partition of the unlabeled samples by transferring knowledge from the labeled known classes.

To this end, we propose a known-novel relation learning framework for multi-view novel class discovery. First, we employ multi-view matrix factorization to learn view-specific basis matrices and a shared partition representation, which captures the distributional structure across different views. The known-class part of the shared partition representation is constrained by $\mathbf{G}_l$, enabling labeled supervision to guide multi-view representation learning and adaptive view-weight estimation. Second, we use known-class anchors to reconstruct novel samples and learn their relation distributions with respect to known classes. In this way, each novel sample is represented by its relation pattern in the known-class anchor space. Third, we incorporate these relation patterns into novel-class partitioning, encouraging samples assigned to the same novel class to preserve consistent known-anchor relation structures. Finally, all components are integrated into a unified objective function. The principal notations used throughout this paper are summarized in Table 4, and the overall framework is illustrated in Fig. 1.

Notation. The principal symbols are defined when first used; a compact notation summary is provided in Appendix A.

3.2 Label-guided Distribution Transfer

MVNCD requires transferring knowledge from labeled known classes to unlabeled novel classes while exploiting multi-view complementary information. We adopt multi-view matrix factorization as the basic modeling backbone. For the v-th view, $\mathbf{X}^v$ is decomposed into a view-specific basis matrix $\mathbf{W}^v \in \mathbb{R}^{d_v \times k}$ and a latent factor matrix $\mathbf{Z}^v \in \mathbb{R}^{k \times n}$, where $k = k_l + k_u$. The preliminary objective is

$$\min_{\mathbf{W}^v, \mathbf{Z}^v} \quad \sum_{v=1}^{V} \|\mathbf{X}^v - \mathbf{W}^v \mathbf{Z}^v\|_F^2 \quad \text{s.t.} \quad \mathbf{W}^{v\top} \mathbf{W}^v = \mathbf{I}_k. \tag{1}$$

To connect representation learning with clustering, we further decompose the factor matrix as

$$\mathbf{Z}^v = \mathbf{A}^v \mathbf{Y}, \tag{2}$$

where $\mathbf{A}^v \in \mathbb{R}^{k \times k}$ is the view-specific centroid matrix, and $\mathbf{Y} \in \{0,1\}^{k \times n}$ is the shared partition matrix. Since known and novel classes are disjoint, $\mathbf{Y}$ is constructed by combining the fixed known-class labels and the learnable novel-class partition. Let $\widetilde{\mathbf{G}}_l \in \{0,1\}^{k \times n_l}$ be the embedded known-class label matrix:

$$\widetilde{\mathbf{G}}_l = \begin{bmatrix} \mathbf{G}_l \\ \mathbf{0}_{k_u \times n_l} \end{bmatrix}. \tag{3}$$

For novel samples, let $\mathbf{Y}_u \in \{0,1\}^{n_u \times k_u}$ denote their partition matrix, and let $\mathbf{E}_u \in \{0,1\}^{k \times k_u}$ map novel-class dimensions into the total class space. Then

$$\mathbf{Y} = \left[\widetilde{\mathbf{G}}_l, \mathbf{E}_u \mathbf{Y}_u^\top\right] \in \{0,1\}^{k \times n}. \tag{4}$$

Considering the varying reliability of different views, we introduce adaptive view weights $\boldsymbol{\alpha}$ and formulate the label-guided distribution transfer objective as

$$\begin{aligned} \min_{\mathbf{W}^v, \mathbf{A}^v, \mathbf{Y}_u, \boldsymbol{\alpha}} \quad & \sum_{v=1}^{V} \alpha_v^2 \|\mathbf{X}^v - \mathbf{W}^v \mathbf{A}^v \mathbf{Y}\|_F^2 \\ \text{s.t.} \mathbf{W}^{v\top} \mathbf{W}^v = \mathbf{I}_k, \boldsymbol{\alpha}^\top \mathbf{1}_V = 1, \boldsymbol{\alpha} \geq \mathbf{0}, \mathbf{Y}_u \in \Phi^{n_u \times k_u}. & \end{aligned} \tag{5}$$

Through the above formulation, the fixed known-class labels are embedded into the shared partition matrix and jointly participate in the learning of view-specific bases, centroids, and adaptive view weights. Therefore, the distributional knowledge contained in labeled known classes can be transferred to novel-class clustering.

3.3 Known-anchor Relation Reconstruction

The label-guided distribution transfer module injects known-class supervision into multi-view factorization. However, it mainly transfers knowledge through distribution modeling, while the relations between known and novel classes are still not explicitly explored. Although known and novel

classes have disjoint label spaces, they may still exhibit correlated structures in feature spaces. Such known-novel relations can provide additional references for novel-class discovery.

A straightforward way to exploit such relations is to construct a sample-level relation graph between unlabeled novel samples and all labeled known samples. However, this strategy is computationally expensive, since it requires measuring pairwise relations between the two sets in each view and leads to a relation graph with size proportional to $n_l n_u$. In addition, sample-level known anchors may contain redundant or noisy information. To obtain a compact and reliable reference space, we use known-class prototypes instead of all known samples as relation anchors. Specifically, let $\mathbf{E}_l \in \{0,1\}^{k\times k_l}$ be the selection matrix that extracts the known-class dimensions from the whole class space. For the v-th view, the known-class anchor matrix is defined as

$$\mathbf{C}_l^v = \mathbf{W}^v \mathbf{A}^v \mathbf{E}_l \in \mathbb{R}^{d_v \times k_l}. \tag{6}$$

Each column of $\mathbf{C}_l^v$ represents a known-class prototype learned under label guidance. Compared with sample-level anchors, these class-level anchors reduce the number of relation references from n_l to k_l, while preserving the semantic structure of known classes.

For unlabeled novel samples, we introduce a relation matrix $\mathbf{Q}^v \in \mathbb{R}^{k_l \times n_u}$, where each column records the relation distribution between one novel sample and all known-class prototypes. To make the relation representations comparable across samples, each column of $\mathbf{Q}^v$ is constrained on the probability simplex:

$$\mathbf{Q}^v \geq \mathbf{0}, \quad \mathbf{1}_{k_l}^\top \mathbf{Q}^v = \mathbf{1}_{n_u}^\top. \tag{7}$$

Under this constraint, $\mathbf{Q}^v$ can be regarded as the relational coordinate representation of novel samples in the known-prototype reference space.

The known-anchor relation reconstruction is formulated as

$$\begin{aligned} \min_{\mathbf{Q}^v} \quad & \sum_{v=1}^{V} \|\mathbf{X}_u^v - \mathbf{W}^v \mathbf{A}^v \mathbf{E}_l \mathbf{Q}^v\|_F^2 \\ \text{s.t.} \quad & \mathbf{Q}^v \geq \mathbf{0}, \quad \mathbf{1}_{k_l}^\top \mathbf{Q}^v = \mathbf{1}_{n_u}^\top. \end{aligned} \tag{8}$$

Through this formulation, known classes are used not only as supervised data for factorization, but also as compact relation references for novel-class discovery. The learned $\mathbf{Q}^v$ explicitly encodes how novel samples are associated with known-class prototypes in each view.

3.4 Relation-consistent Novel Partitioning

The relation matrix $\mathbf{Q}^v$ describes sample-level associations between novel samples and known-class prototypes. However, it does not directly determine the final novel-class partition. To connect known-novel relations with clustering, we introduce a relation dictionary $\mathbf{B}^v \in \mathbb{R}^{k_l \times k_u}$, where each column denotes the relation prototype of one novel cluster. The columns of $\mathbf{B}^v$ are constrained on the probability simplex:

$$\mathbf{B}^v \geq \mathbf{0}, \quad \mathbf{1}_{k_l}^\top \mathbf{B}^v = \mathbf{1}_{k_u}^\top. \tag{9}$$

Given the novel partition matrix $\mathbf{Y}_u$, $\mathbf{B}^v \mathbf{Y}_u^\top$ assigns the corresponding relation prototype to each novel sample according to its cluster indicator. Therefore, the relation-consistent novel partitioning objective is formulated as

$$\min_{\mathbf{B}^v, \mathbf{Y}_u} \sum_{v=1}^{V} \left\|\mathbf{Q}^v - \mathbf{B}^v \mathbf{Y}_u^\top\right\|_F^2 \quad \text{s.t.} \mathbf{B}^v \geq \mathbf{0}, \mathbf{1}_{k_l}^\top \mathbf{B}^v = \mathbf{1}_{k_u}^\top. \tag{10}$$

By minimizing Eq. (10), samples assigned to the same novel cluster are encouraged to share similar known-prototype relation patterns. Thus, the final partition is guided not only by feature reconstruction, but also by cluster-level known-novel relational consistency.

3.5 Loss Function

By integrating the above components, the final objective function is formulated as

$$\begin{aligned}\min_{\mathbf{W}^v,\mathbf{A}^v,\mathbf{Q}^v,\mathbf{B}^v,\mathbf{Y}_u,\boldsymbol{\alpha}} \quad & \sum_{v=1}^{V}\alpha_v^2\left\|\mathbf{X}^v-\mathbf{W}^v\mathbf{A}^v\mathbf{Y}\right\|_F^2 \\ & +\lambda_1\left\|\mathbf{X}_u^v-\mathbf{W}^v\mathbf{A}^v\mathbf{E}_l\mathbf{Q}^v\right\|_F^2+\lambda_2\left\|\mathbf{Q}^v-\mathbf{B}^v\mathbf{Y}_u^\top\right\|_F^2 \\ & \text{s.t.}\mathbf{W}^{v\top}\mathbf{W}^v=\mathbf{I}_k,\boldsymbol{\alpha}^\top\mathbf{1}_V=1,\boldsymbol{\alpha}\geq\mathbf{0},\mathbf{Y}_u\in\Phi^{n_u\times k_u}, \\ & \mathbf{Q}^v\geq\mathbf{0},\mathbf{1}_{k_l}^\top\mathbf{Q}^v=\mathbf{1}_{n_u}^\top,\mathbf{B}^v\geq\mathbf{0},\mathbf{1}_{k_l}^\top\mathbf{B}^v=\mathbf{1}_{k_u}^\top,\end{aligned} \tag{11}$$

where λ_1 and λ_2 are trade-off parameters.

The first term performs label-guided multi-view distribution transfer. The second term learns the relations between novel samples and known-class prototypes. The third term further summarizes these sample-level relations into novel-cluster relation prototypes and aligns them with the final partition. Through this unified formulation, known-class information is transferred to novel-class discovery from both distribution and relation perspectives.

4 Optimization

We develop an alternating optimization strategy for Eq. (11), updating one variable while fixing the others. Specifically, $\mathbf{W}^v$ is reduced to an orthogonal Procrustes problem and solved by SVD; $\mathbf{A}^v$ admits a closed-form least-squares update; $\mathbf{Q}^v$ is updated by projected gradient descent under the simplex constraint; $\mathbf{B}^v$ is obtained by averaging the relation vectors assigned to each novel cluster; $\mathbf{Y}_u$ is updated by discrete sample-wise assignment; and the adaptive view weights $\boldsymbol{\alpha}$ have a closed-form residual-based solution.

The relation update is a key step because it directly couples known-class anchors with novel-class partitioning. Let $\mathbf{C}_l^v=\mathbf{W}^v\mathbf{A}^v\mathbf{E}_l$. The subproblem for $\mathbf{Q}^v$ is convex, with gradient

$$\nabla_{\mathbf{Q}^v}=2\lambda_1\mathbf{C}_l^{v\top}(\mathbf{C}_l^v\mathbf{Q}^v-\mathbf{X}_u^v)+2\lambda_2(\mathbf{Q}^v-\mathbf{B}^v\mathbf{Y}_u^\top), \tag{12}$$

and is solved by a gradient step followed by column-wise projection onto the probability simplex. Given $\mathbf{Y}_u$, the c-th relation prototype is the mean of the relation vectors assigned to cluster c; conversely, updating $\mathbf{Y}_u$ assigns each sample to the cluster that minimizes the combined feature-reconstruction and relation-consistency cost across views. Hence, relation learning is explicitly involved in the final partition optimization rather than being used only as an auxiliary representation.

Complete derivations, all closed-form updates, the projected-gradient step, and the overall optimization procedure are provided in Appendix B. The per-iteration complexity is $\mathcal{O}(d(nk+k^2)+Vnk^2+Vk^3)$, which is linear in the sample size n when V and k are relatively small. Empirical convergence behavior is reported in Appendix F.2.

5 Experiment

In this section, we conduct comprehensive experiments on eight widely used benchmark datasets to evaluate the proposed method. We first introduce the experimental settings and report the comparison results with representative MVC and NCD methods. Then, ablation studies are conducted to verify the effectiveness of the key components. We further visualize the learned known–novel relation prototypes and investigate the relationship between the learned view weights and single-view clustering performance. Finally, parameter sensitivity and convergence analyses are provided to examine the robustness and optimization behavior of the proposed method.

5.1 Experimental Setting

5.1.1 Datasets

We conduct experiments on eight widely used multi-view benchmark datasets, including ProteinFold, KIPAN, Mfeat, Cora, Wiki, Citeseer, Reuters, and YTB20. These datasets cover protein structure data, multi-omics data, handwritten digits, citation networks, text-image data, multilingual text

Table 1: Dataset statistics and known/novel splits used in our experiments.

Dataset	Samples	Views	Classes	k_l	k_u	n_l	n_u
ProteinFold	694	12	27	13	14	312	382
KIPAN	897	3	3	1	2	256	641
Mfeat	2000	6	10	5	5	1000	1000
Cora	2708	4	7	3	4	1215	1493
Wiki	2866	2	10	5	5	1472	1394
Citeseer	3312	4	6	3	3	1965	1347
Reuters	18758	5	6	3	3	10667	8091
YTB20	63896	4	20	10	10	38654	25242

data, and video data, providing a comprehensive evaluation of the proposed method under diverse multi-view scenarios.

Following the standard novel class discovery setting, the categories of each dataset are divided into known and novel classes with an approximate ratio of 1:1. Specifically, for a dataset with C classes, the first $\lfloor C/2 \rfloor$ classes are treated as known classes, while the remaining classes are treated as novel classes. Samples from the known classes are used as labeled data, whereas samples from the novel classes are used as unlabeled data for discovering novel categories. The statistics and resulting data splits are summarized in Table 1.

5.1.2 Compared Algorithms

We compare the proposed method with representative MVC methods, including Best, Concat, OPMC (Liu et al., 2021), AEVC (Liu et al., 2024), and LASD (Zhang et al., 2025b), as well as NCD methods, including OpenNCD (Liu et al., 2023), SNCD (Wang et al., 2024), and IICMVNCD (Wan et al., 2025). MVC methods are evaluated on the unlabeled novel-class set. For the single-view NCD methods OpenNCD and SNCD, we concatenate multi-view features and feed them into an MLP projection network while keeping their original learning objectives and clustering strategies unchanged. Detailed descriptions of all compared methods are provided in Appendix D.

5.1.3 Implementation Details

We obtained the public codes of all compared methods from their official repositories and implemented all baselines using the parameter settings recommended in their original papers. We adopt ACC, NMI, and Purity as evaluation metrics, where higher values indicate better clustering performance.

For the proposed method, the trade-off parameters λ_1 and λ_2 are selected from $\{10^{-3}, 10^{-2}, \cdots, 10^{4}\}$. To examine the influence of parameter selection, we further report the sensitivity analysis of these two parameters.

All experiments were conducted on a machine equipped with an Intel Core i9-10850K CPU @ 3.60GHz, 128GB RAM, and an NVIDIA RTX 3090 GPU.

5.2 Experimental Results

We report the ACC performance of our method and the compared MVC and NCD methods on eight benchmark datasets in Table 2; the corresponding NMI and Purity results are provided in Appendix E. Our method achieves the best ACC on all eight datasets. Compared with the second-best result, the relative ACC improvements are 3.29%, 2.53%, 1.45%, 2.82%, 13.05%, 7.82%, 0.33%, and 1.95% on ProteinFold, KIPAN, Mfeat, Cora, Wiki, Citeseer, Reuters, and YTB20, respectively. The full results further show consistently strong NMI and Purity, indicating stable clustering quality across different types of multi-view data.

Compared with conventional MVC methods, our method can additionally exploit labeled known-class information for knowledge transfer. Compared with the single-view NCD methods OpenNCD and SNCD, it explicitly preserves view-specific information instead of relying only on concatenated

Table 2: ACC (%) on eight multi-view NCD benchmarks. The best results are highlighted in bold, while the second-best results are underlined. NMI and Purity are reported in Appendix E.

Dataset	MVC Methods					NCD Methods			
	Best	Concat	OPMC	AEVC	LASD	OpenNCD	SNCD	IICMVNCD	Ours
ProteinFold	30.37	33.51	28.53	32.18	<u>34.98</u>	26.44	30.37	33.77	**36.13**
KIPAN	92.36	89.55	89.86	91.67	<u>92.67</u>	74.73	92.20	92.51	**95.01**
Mfeat	92.10	94.80	93.60	92.70	91.75	96.40	<u>96.70</u>	95.60	**98.10**
Cora	50.77	39.99	42.33	63.20	<u>78.82</u>	22.44	30.74	76.36	**81.04**
Wiki	63.55	29.84	30.77	37.30	65.32	31.13	54.16	<u>65.42</u>	**73.96**
Citeseer	45.51	43.93	41.28	50.12	<u>60.25</u>	31.11	42.91	43.88	**64.96**
Reuters	59.85	59.49	62.77	–	<u>71.83</u>	53.58	46.79	71.56	**72.07**
YTB20	67.64	64.56	<u>81.12</u>	80.69	79.53	74.82	79.02	75.73	**82.70**

Table 3: Ablation study of the proposed method in terms of ACC. The best results are highlighted in bold, while the second-best results are underlined.

Methods	ProteinFold	KIPAN	Mfeat	Cora	Wiki	Citeseer	Reuters	YTB20
w/o λ_2	<u>33.25</u>	<u>94.51</u>	<u>96.70</u>	<u>70.60</u>	<u>64.92</u>	38.60	<u>67.19</u>	75.73
w/o λ_1, λ_2	31.41	91.11	95.60	55.26	62.80	35.86	66.35	72.53
w/o Sup.	28.27	92.20	96.40	48.56	54.16	38.38	62.33	**82.72**
Novel-only	32.20	92.20	96.00	49.97	53.87	<u>39.27</u>	58.98	76.04
Ours	**36.13**	**95.01**	**98.10**	**81.04**	**73.96**	**64.96**	**72.07**	<u>82.70</u>

features. IICMVNCD is the most closely related multi-view NCD baseline and generally performs better than many conventional MVC and single-view NCD methods. Nevertheless, our method consistently improves over IICMVNCD in ACC and Purity on all datasets, which supports the benefit of explicitly modeling known–novel relations in addition to distributional and inter-view correlation transfer.

5.3 Ablation Study

We evaluate four variants to examine the contribution of the main components: w/o λ_2 removes relation-consistent novel partitioning; w/o λ_1, λ_2 removes the entire known-anchor relation-learning branch; w/o Sup. removes known-class supervision; and Novel-only learns from novel-class samples without using known-class data. The ACC results are summarized in Table 3.

The complete model achieves the best or second-best performance on every dataset. Removing λ_2 degrades ACC on most datasets, showing that aligning sample-level relations with novel-cluster assignments is beneficial. Removing both λ_1 and λ_2 leads to a further drop, especially on structurally complex datasets such as Cora and Citeseer, which confirms the importance of the whole known-anchor relation-learning branch. The w/o Sup. and Novel-only variants are also generally inferior, indicating that known-class supervision provides transferable information for novel-class discovery. Although w/o Sup. is marginally higher on YTB20, the overall results consistently support the contribution of supervision and relation learning.

5.4 Relation Prototype Visualization

To further examine the learned relational knowledge, we visualize the relation prototypes on Mfeat and Wiki in Fig. 2. We average $\mathbf{B}^v$ over all views, where each column characterizes the relation pattern of a novel cluster with respect to known-class anchors. Different novel classes exhibit clearly different relation distributions over the known classes rather than sharing a uniform pattern. This observation is consistent with our motivation that novel categories can be characterized in a reference space defined by known categories, and it indicates that the learned relation prototypes provide complementary class-dependent structural cues for novel-class partitioning.

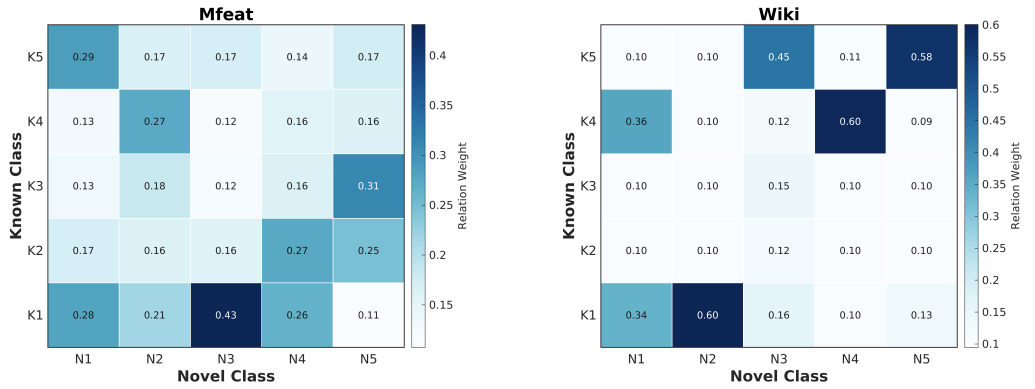


Figure 2: Visualization of the learned known–novel relation prototypes on two representative datasets.

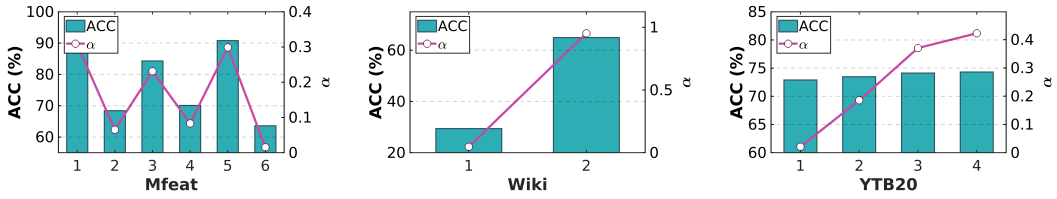


Figure 3: Relationship between learned view weights and single-view clustering performance.

5.5 View Weight Analysis

We further compare the learned weights α_v with single-view clustering performance on Mfeat, Wiki, and YTB20 in Fig. 3. Specifically, we perform k-means on the novel samples of each individual view and use ACC to characterize its clustering quality. The learned weights generally follow the same trend as the single-view performance: stronger views tend to receive larger weights, whereas weaker views receive smaller weights. This result suggests that the adaptive weighting mechanism can distinguish the relative reliability of different views during multi-view learning.

5.6 Parameter Sensitivity and Convergence

We further examine parameter sensitivity and practical convergence, with detailed results provided in Appendices F.1 and F.2. The method remains competitive over a reasonably broad range of λ_1 and λ_2, while the objective typically decreases rapidly and becomes stable within about 40 iterations.

6 Conclusion

In this paper, we investigate the Multi-view Novel Class Discovery problem. To sufficiently exploit both known-class supervision and multi-view complementary information, we propose a Known-Novel Relation Learning framework. Specifically, we first employ multi-view matrix factorization to capture the distributional structure across views, where known-class labels are incorporated to guide representation learning and adaptive view weighting. Moreover, to further utilize the available known-class information, we regard known-class prototypes as relation anchors and construct known-novel relation representations for novel samples. By aligning these sample-level relations with novel-cluster assignments, the proposed method transfers knowledge from known classes to novel classes from both distribution and relation perspectives. Experimental results on eight benchmark datasets demonstrate the effectiveness and robustness of the proposed method under diverse multi-view scenarios. This work provides a new perspective for relation-aware multi-view novel class discovery and may offer useful insights for future research in this direction.

AI USE STATEMENT

Generative AI tools were used to assist with code development and debugging, language polishing, LaTeX formatting, and feedback on the presentation and analysis of the work. All AI-assisted suggestions were critically reviewed and independently verified by the authors before incorporation. The authors take full responsibility for the methodology, implementation, experimental results, analysis, and final manuscript.

REPRODUCIBILITY STATEMENT

The main text specifies the objective, optimization strategy, datasets, baselines, and evaluation protocol. Complete optimization derivations, dataset splits, full metric tables, parameter analysis, and convergence results are provided in the appendices to facilitate independent reproduction.

REFERENCES

Jinrong Cui, Yuting Li, Han Huang, and Jie Wen. Dual contrast-driven deep multi-view clustering. *IEEE Transactions on Image Processing*, 33:4753–4764, 2024.

Lele Fu, Pengfei Lin, Athanasios V. Vasilakos, and Shiping Wang. An overview of recent multi-view clustering. *Neurocomputing*, 402:148–161, 2020. doi: 10.1016/j.neucom.2020.02.104.

Shengxiang Gao, Zhengtao Yu, Taisong Jin, and Ming Yin. Multi-view low-rank matrix factorization using multiple manifold regularization. *Neurocomputing*, 335:143–152, 2019.

Peiyan Gu, Chuyu Zhang, Ruijie Xu, and Xuming He. Class-relation knowledge distillation for novel class discovery. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 16428–16437, 2023. doi: 10.1109/ICCV51070.2023.01510.

Kai Han, Andrea Vedaldi, and Andrew Zisserman. Learning to discover novel visual categories via deep transfer clustering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.

Kai Han, Sylvestre-Alvise Rebuffi, Sebastien Ehrhardt, Andrea Vedaldi, and Andrew Zisserman. Automatically discovering and learning new visual categories with ranking statistics. In *International Conference on Learning Representations*, 2020.

Yue Hou, Xueyuan Chen, He Zhu, Ruomei Liu, Bowen Shi, Jiaheng Liu, Junran Wu, and Ke Xu. Nc2d: Novel class discovery for node classification. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, pp. 849–859, 2024.

Yen-Chang Hsu, Zhaoyang Lv, and Zsolt Kira. Learning to cluster in order to transfer across domains and tasks. In *International Conference on Learning Representations*, 2018.

Yen-Chang Hsu, Zhaoyang Lv, Joel Schlosser, Phillip Odom, and Zsolt Kira. Multi-class classification without multi-class labels. In *International Conference on Learning Representations*, 2019.

KJ Joseph, Sujoy Paul, Gaurav Aggarwal, Soma Biswas, Piyush Rai, Kai Han, and Vineeth N Balasubramanian. Novel class discovery without forgetting. In *European Conference on Computer Vision*, pp. 570–586. Springer, 2022.

Zisen Kong, Dongxia Chang, Yiming Wang, Pengyuan Li, and Yao Zhao. Tensorial multi-view clustering via alternative rank minimization and inter-view alignment. *IEEE Transactions on Knowledge and Data Engineering*, 2026.

Wenbin Li, Zhichen Fan, Jing Huo, and Yang Gao. Modeling inter-class and intra-class constraints in novel class discovery. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3449–3458, 2023. doi: 10.1109/CVPR52729.2023.00336.

Xiao Li, Min Fang, and Haixiang Li. Generalized category discovery with unknown sample generation. *IEEE Transactions on Image Processing*, 34:3366–3376, 2025.

Xuelong Li, Han Zhang, Rong Wang, and Feiping Nie. Multiview clustering: A scalable and parameter-free bipartite graph fusion method. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(1):330–344, 2022a. doi: 10.1109/TPAMI.2020.3011148.

Zhenglai Li, Chang Tang, Xiao Zheng, Xinwang Liu, Wei Zhang, and En Zhu. High-order correlation preserved incomplete multi-view subspace clustering. *IEEE Transactions on Image Processing*, 31:2067–2080, 2022b.

Jialu Liu, Chi Wang, Jing Gao, and Jiawei Han. Multi-view clustering via joint nonnegative matrix factorization. In *Proceedings of the 2013 SIAM international conference on data mining*, pp. 252–260. SIAM, 2013.

Jiaming Liu, Yangqiming Wang, Tongze Zhang, Yulu Fan, Qinli Yang, and Junming Shao. Open-world semi-supervised novel class discovery. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, pp. 4002–4010, 2023.

Jiyuan Liu, Xinwang Liu, Yuexiang Yang, Li Liu, Siqi Wang, Weixuan Liang, and Jiangyong Shi. One-pass multi-view clustering for large-scale data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 12344–12353, October 2021.

Suyuan Liu, Ke Liang, Zhibin Dong, Siwei Wang, Xihong Yang, Sihang Zhou, En Zhu, and Xinwang Liu. Learn from view correlation: An anchor enhancement strategy for multi-view clustering. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 26151–26161, 2024. doi: 10.1109/CVPR52733.2024.02471.

Shijie Ma, Fei Zhu, Xu-Yao Zhang, and Cheng-Lin Liu. Protogcd: Unified and unbiased prototype learning for generalized category discovery. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(7):6022–6037, 2025. doi: 10.1109/TPAMI.2025.3557502.

Feiping Nie, Shaojun Shi, and Xuelong Li. Auto-weighted multi-view co-clustering via fast matrix factorization. *Pattern Recognition*, 102:107207, 2020. doi: 10.1016/j.patcog.2020.107207.

Yazhou Ren, Zichen Wen, Junlong Ke, Chenhang Cui, Yonghao Huang, Xinyue Chen, Philip S. Yu, and Lifang He. Enhancing multi-view clustering: A sufficient information-theoretic approach for consistency acquisition and redundancy elimination. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 48(7):7420–7435, 2026. doi: 10.1109/TPAMI.2026.3665701.

Subhankar Roy, Mingxuan Liu, Zhun Zhong, Nicu Sebe, and Elisa Ricci. Class-incremental novel class discovery. In *European Conference on Computer Vision*, pp. 317–333. Springer, 2022.

Xinhang Wan, Jiyuan Liu, Xinbiao Gan, Xinwang Liu, Siwei Wang, Yi Wen, Tianjiao Wan, and En Zhu. One-step multi-view clustering with diverse representation. *IEEE Transactions on Neural Networks and Learning Systems*, pp. 1–13, 2024a. doi: 10.1109/TNNLS.2024.3378194.

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, 2024b. doi: 10.1109/TIP.2024.3388974.

Xinhang Wan, Jiyuan Liu, Qian Qu, Suyuan Liu, Chuyu Zhang, Fangdi Wang, Xinwang Liu, En Zhu, and Kunlun He. Intra-view and inter-view correlation guided multi-view novel class discovery. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 4114–4124, 2025.

Jing Wang, Feng Tian, Hongchuan Yu, Chang Hong Liu, Kun Zhan, and Xiao Wang. Diverse non-negative matrix factorization for multiview data representation. *IEEE Transactions on Cybernetics*, 48(9):2620–2632, 2018. doi: 10.1109/TCYB.2017.2747400.

Jinglong Wang, Xiawei Li, Jing Zhang, Qingyuan Xu, Qin Zhou, Qian Yu, Lu Sheng, and Dong Xu. Diffusion model is secretly a training-free open vocabulary semantic segmenter. *IEEE Transactions on Image Processing*, 34:1895–1907, 2025.

Rong Wang, Jitao Lu, Yihang Lu, Feiping Nie, and Xuelong Li. Discrete and parameter-free multiple kernel k-means. *IEEE Transactions on Image Processing*, 31:2796–2808, 2022.

Weishuai Wang, Ting Lei, Qingchao Chen, and Yang Liu. Semantic-guided novel category discovery. In *Proceedings of the Thirty-Eighth AAAI Conference on Artificial Intelligence and Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence and Fourteenth Symposium on Educational Advances in Artificial Intelligence*. AAAI Press, 2024. ISBN 978-1-57735-887-9.

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. doi: 10.1109/TMM.2024.3387298.

Ruijie Xu, Chuyu Zhang, Hui Ren, and Xuming He. Dual-level adaptive self-labeling for novel class discovery in point cloud segmentation. In *European Conference on Computer Vision*, pp. 288–305. Springer, 2024.

Zuyuan Yang, Naiyao Liang, Wei Yan, Zhenni Li, and Shengli Xie. Uniform distribution non-negative matrix factorization for multiview clustering. *IEEE Transactions on Cybernetics*, 51(6): 3249–3262, 2021. doi: 10.1109/TCYB.2020.2984552.

Shengju Yu, Suyuan Liu, Siwei Wang, Chang Tang, Zhigang Luo, Xinwang Liu, and En Zhu. Sparse low-rank multi-view subspace clustering with consensus anchors and unified bipartite graph. *IEEE Transactions on Neural Networks and Learning Systems*, 36(1):1438–1452, 2025.

Zhiqi Yu, Ke Lu, Kangkai Wu, Hongbo Chen, Fengling Li, and Jingjing Li. Training-free open-set domain adaptation with vision-language models. *IEEE Transactions on Image Processing*, 2026.

Jun Yue, Leyuan Fang, and Min He. Spectral-spatial latent reconstruction for open-set hyperspectral image classification. *IEEE Transactions on Image Processing*, 31:5227–5241, 2022.

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*, 2025a.

Junpu Zhang, Shengju Yu, Suyuan Liu, Siwei Wang, Miaomiao Li, Xinwang Liu, En Zhu, and Kunlun He. Learning the anchors with similar distributions to original data for multi-view clustering. In *Proceedings of the 33rd ACM International Conference on Multimedia*, pp. 857–866, 2025b.

Yuyang Zhao, Zhun Zhong, Nicu Sebe, and Gim Hee Lee. Novel class discovery in semantic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 4340–4349, 2022.

Zhun Zhong, Enrico Fini, Subhankar Roy, Zhiming Luo, Elisa Ricci, and Nicu Sebe. Neighborhood contrastive learning for novel class discovery. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10867–10875, 2021.

Table 4: Summary of principal notations.

Notation	Meaning
V	Number of views
n_l, n_u, n	Numbers of labeled, unlabeled, and total samples
k_l, k_u, k	Numbers of known, novel, and total classes
d_v	Feature dimension of the v-th view
$\mathbf{X}^v$	Data matrix of the v-th view
$\mathbf{X}_l^v, \mathbf{X}_u^v$	Known- and novel-class data matrices
$\mathbf{G}_l$	Known-class label matrix
$\mathbf{W}^v$	View-specific basis matrix
$\mathbf{A}^v$	View-specific centroid matrix
$\mathbf{Y}, \mathbf{Y}_u$	Shared and novel partition matrices
$\mathbf{E}_l, \mathbf{E}_u$	Known- and novel-class selection matrices
$\mathbf{C}_l^v$	Known-class anchor matrix
$\mathbf{Q}^v$	Sample-level relation matrix
$\mathbf{B}^v$	Cluster-level relation dictionary
$\boldsymbol{\alpha}$	Adaptive view-weight vector
λ_1, λ_2	Trade-off parameters

A Notation Summary

B Complete Optimization Details

We develop an alternating optimization algorithm to solve the problem in Eq. (11). In each iteration, one variable is updated while the others are fixed.

B.1 Update $\mathbf{W}^v$

Fixing other variables, the optimization problem w.r.t. $\mathbf{W}^v$ is formulated as

$$\begin{aligned} \min_{\mathbf{W}^v} \quad & \alpha_v^2 \left\| \mathbf{X}^v - \mathbf{W}^v \mathbf{A}^v \mathbf{Y} \right\|_F^2 + \lambda_1 \left\| \mathbf{X}_u^v - \mathbf{W}^v \mathbf{A}^v \mathbf{E}_l \mathbf{Q}^v \right\|_F^2 \\ \text{s.t.} \quad & \mathbf{W}^{v\top} \mathbf{W}^v = \mathbf{I}_k. \end{aligned} \tag{13}$$

It is equivalent to

$$\begin{aligned} \max_{\mathbf{W}^v} \quad & \operatorname{Tr}\left(\mathbf{W}^{v\top} \mathbf{M}^v \right) \\ \text{s.t.} \quad & \mathbf{W}^{v\top} \mathbf{W}^v = \mathbf{I}_k, \end{aligned} \tag{14}$$

where

$$\mathbf{M}^v = \alpha_v^2 \mathbf{X}^v \mathbf{Y}^\top \mathbf{A}^{v\top} + \lambda_1 \mathbf{X}_u^v \mathbf{Q}^{v\top} \mathbf{E}_l^\top \mathbf{A}^{v\top}. \tag{15}$$

Performing SVD on $\mathbf{M}^v$ as $\mathbf{M}^v = \mathbf{U}^v \mathbf{\Sigma}^v \mathbf{V}^{v\top}$, the optimal solution is given by

$$\mathbf{W}^v = \mathbf{U}^v \mathbf{V}^{v\top}. \tag{16}$$

B.2 Update $\mathbf{A}^v$

Fixing other variables, the subproblem w.r.t. $\mathbf{A}^v$ is

$$\min_{\mathbf{A}^v} \quad \alpha_v^2 \left\| \mathbf{X}^v - \mathbf{W}^v \mathbf{A}^v \mathbf{Y} \right\|_F^2 + \lambda_1 \left\| \mathbf{X}_u^v - \mathbf{W}^v \mathbf{A}^v \mathbf{E}_l \mathbf{Q}^v \right\|_F^2. \tag{17}$$

Taking the derivative w.r.t. $\mathbf{A}^v$ and setting it to zero, we obtain

$$\begin{aligned} \mathbf{A}^v = & \left(\alpha_v^2 \mathbf{W}^{v\top} \mathbf{X}^v \mathbf{Y}^\top + \lambda_1 \mathbf{W}^{v\top} \mathbf{X}_u^v \mathbf{Q}^{v\top} \mathbf{E}_l^\top \right) \\ & \left(\alpha_v^2 \mathbf{Y} \mathbf{Y}^\top + \lambda_1 \mathbf{E}_l \mathbf{Q}^v \mathbf{Q}^{v\top} \mathbf{E}_l^\top \right)^{-1}. \end{aligned} \tag{18}$$

B.3 UPDATE $\mathbf{Q}^v$

Fixing other variables, $\mathbf{Q}^v$ is updated by solving

$$\begin{aligned} \min_{\mathbf{Q}^v} \quad & \lambda_1 \left\|\mathbf{X}_u^v - \mathbf{C}_l^v \mathbf{Q}^v\right\|_F^2 + \lambda_2 \left\|\mathbf{Q}^v - \mathbf{B}^v \mathbf{Y}_u^\top\right\|_F^2 \\ \text{s.t.} \quad & \mathbf{Q}^v \geq \mathbf{0}, \quad \mathbf{1}_{k_l}^\top \mathbf{Q}^v = \mathbf{1}_{n_u}^\top, \end{aligned} \tag{19}$$

where $\mathbf{C}_l^v = \mathbf{W}^v \mathbf{A}^v \mathbf{E}_l$.

This is a convex quadratic problem with simplex constraints. Its gradient is

$$\nabla_{\mathbf{Q}^v} = 2\lambda_1 \mathbf{C}_l^{v\top} \left(\mathbf{C}_l^v \mathbf{Q}^v - \mathbf{X}_u^v\right) + 2\lambda_2 \left(\mathbf{Q}^v - \mathbf{B}^v \mathbf{Y}_u^\top\right). \tag{20}$$

Let

$$L_Q^v = 2\lambda_1 \left\|\mathbf{C}_l^v\right\|_2^2 + 2\lambda_2. \tag{21}$$

Then $\mathbf{Q}^v$ can be updated by projected gradient descent:

$$\mathbf{Q}^v \leftarrow \Pi_\Delta \left(\mathbf{Q}^v - \frac{1}{L_Q^v} \nabla_{\mathbf{Q}^v}\right), \tag{22}$$

where $\Pi_\Delta(\cdot)$ denotes column-wise projection onto the probability simplex.

B.4 UPDATE $\mathbf{B}^v$

Fixing other variables, the subproblem w.r.t. $\mathbf{B}^v$ is

$$\begin{aligned} \min_{\mathbf{B}^v} \quad & \left\|\mathbf{Q}^v - \mathbf{B}^v \mathbf{Y}_u^\top\right\|_F^2 \\ \text{s.t.} \quad & \mathbf{B}^v \geq \mathbf{0}, \quad \mathbf{1}_{k_l}^\top \mathbf{B}^v = \mathbf{1}_{k_u}^\top. \end{aligned} \tag{23}$$

Let $\mathcal{I}_c = \{j \mid [\mathbf{Y}_u]_{jc} = 1\}$ denote the sample index set assigned to the c-th novel cluster. The c-th column of $\mathbf{B}^v$ is updated by

$$\mathbf{b}_c^v = \frac{1}{|\mathcal{I}_c|} \sum_{j \in \mathcal{I}_c} \mathbf{q}_j^v, \tag{24}$$

where $\mathbf{q}_j^v$ is the j-th column of $\mathbf{Q}^v$.

Since each $\mathbf{q}_j^v$ lies on the simplex, their average also satisfies the simplex constraint. If $\mathcal{I}_c$ is empty, $\mathbf{b}_c^v$ is kept unchanged.

B.5 UPDATE $\mathbf{Y}_u$

Fixing other variables, the terms related to $\mathbf{Y}_u$ are reduced to

$$\begin{aligned} \min_{\mathbf{Y}_u} \quad & \sum_{v=1}^{V} \alpha_v^2 \left\|\mathbf{X}_u^v - \mathbf{W}^v \mathbf{A}^v \mathbf{E}_u \mathbf{Y}_u^\top\right\|_F^2 + \lambda_2 \left\|\mathbf{Q}^v - \mathbf{B}^v \mathbf{Y}_u^\top\right\|_F^2 \\ \text{s.t.} \quad & \mathbf{Y}_u \in \Phi^{n_u \times k_u}. \end{aligned} \tag{25}$$

Let $\mathbf{m}_c^v = \mathbf{W}^v \mathbf{A}^v \mathbf{E}_u \mathbf{e}_c$ and $\mathbf{b}_c^v = \mathbf{B}^v \mathbf{e}_c$. The cost of assigning the j-th novel sample to the c-th novel cluster is

$$\mathcal{C}_j(c) = \sum_{v=1}^{V} \left[\alpha_v^2 \left\|\mathbf{x}_{u,j}^v - \mathbf{m}_c^v\right\|_2^2 + \lambda_2 \left\|\mathbf{q}_j^v - \mathbf{b}_c^v\right\|_2^2\right]. \tag{26}$$

Then the assignment is updated by

$$c_j^* = \arg \min_{c \in \{1,\ldots,k_u\}} \mathcal{C}_j(c), \quad [\mathbf{Y}_u]_{jc_j^*} = 1. \tag{27}$$

All the other entries in the j-th row of $\mathbf{Y}_u$ are set to zero.

B.6 UPDATE $\boldsymbol{\alpha}$

Fixing other variables, $\boldsymbol{\alpha}$ is updated by solving

$$\begin{aligned} \min_{\boldsymbol{\alpha}} \quad & \sum_{v=1}^{V} \alpha_v^2 r_v \\ \text{s.t.} \quad & \boldsymbol{\alpha}^\top \mathbf{1}_V = 1, \quad \boldsymbol{\alpha} \geq \mathbf{0}, \end{aligned} \tag{28}$$

where

$$r_v = \|\mathbf{X}^v - \mathbf{W}^v \mathbf{A}^v \mathbf{Y}\|_F^2 . \tag{29}$$

Based on the Cauchy–Schwarz inequality, the closed-form solution is

$$\alpha_v = \frac{1/r_v}{\sum_{t=1}^{V} 1/r_t}. \tag{30}$$

The overall optimization procedure is summarized in Algorithm 1.

Algorithm 1 Optimization Algorithm

Input: Multi-view data $\{\mathbf{X}^v\}_{v=1}^V$, known-class label matrix $\mathbf{G}_l$, cluster numbers k_l and k_u, hyper-parameters λ_1 and λ_2.
Output: Novel partition matrix $\mathbf{Y}_u$.

```
Initialize α, W^v, A^v, Q^v, B^v, and Y_u.
while not converged do
    for v = 1 to V do
        Update W^v by solving Eq. (13);
        Update A^v by solving Eq. (17);
        Update Q^v by Eq. (22);
        Update B^v by Eq. (24);
    end for
    Update Y_u by Eq. (27);
    Update α by Eq. (30);
end while
```

B.7 COMPLEXITY ANALYSIS

Let $d = \sum_{v=1}^{V} d_v$ and $k = k_l + k_u$. In each iteration, updating $\mathbf{W}^v$ requires computing the matrix $\mathbf{M}^v$ and performing SVD on a $d_v \times k$ matrix, which costs $\mathcal{O}(d_v nk + d_v k^2)$ for the v-th view. Updating $\mathbf{A}^v$ is dominated by matrix multiplications and the inverse of a $k \times k$ matrix, with the cost of $\mathcal{O}(d_v nk + nk^2 + k^3)$. Updating $\mathbf{Q}^v$ involves the reconstruction over k_l known-class anchors and the relation dictionary term, whose cost is bounded by $\mathcal{O}(d_v n_u k_l + n_u k_l k_u)$. Updating $\mathbf{B}^v$ computes the relation prototype of each novel cluster, costing $\mathcal{O}(n_u k_l)$ per view. Updating $\mathbf{Y}_u$ evaluates k_u candidate assignments for each novel sample, involving both feature reconstruction and relation consistency terms, with the cost of $\mathcal{O}(d_v n_u k_u + n_u k_l k_u)$ for the v-th view. Finally, updating $\boldsymbol{\alpha}$ requires computing the reconstruction residuals of all views, costing $\mathcal{O}(dnk)$.

Combining the above terms over all views, the overall time complexity per iteration is

$$\mathcal{O}\left(d(nk + k^2) + Vnk^2 + Vk^3\right). \tag{31}$$

Since V and k are usually smaller than the sample size n, the proposed algorithm scales linearly with respect to n.

C DATASET STATISTICS

D COMPARED ALGORITHMS

We compare the proposed method with representative MVC and NCD methods. For MVC methods, we directly evaluate their clustering performance on the unlabeled novel-class set. For single-view

Table 5: Statistics of the benchmark datasets used in our experiments.

Dataset	Samples	Views	Classes
ProteinFold	694	12	27
KIPAN	897	3	3
Mfeat	2000	6	10
Cora	2708	4	7
Wiki	2866	2	10
Citeseer	3312	4	6
Reuters	18758	5	6
YTB20	63896	4	20

Table 6: Statistics of the dataset splits used in our experiments. We report the numbers of known and novel categories (k_l, k_u) and labeled and unlabeled samples (n_l, n_u).

Dataset	k_l	k_u	n_l	n_u
ProteinFold	13	14	312	382
KIPAN	1	2	256	641
Mfeat	5	5	1000	1000
Cora	3	4	1215	1493
Wiki	5	5	1472	1394
Citeseer	3	3	1965	1347
Reuters	3	3	10667	8091
YTB20	10	10	38654	25242

NCD methods, including OpenNCD and SNCD, we adapt them to the multi-view setting by concatenating features from all views and feeding the resulting representation into an MLP projection network. After obtaining the projected representation, their original learning objectives and clustering strategies are kept unchanged. The compared methods are summarized as follows:

1. **Best**: This baseline performs k-means clustering on each individual view of the unlabeled novel-class samples and reports the best result among all single views.
2. **Concat**: This baseline directly concatenates the features of all views for the unlabeled novel-class samples and then performs k-means clustering in the concatenated feature space.
3. **OPMC** (Liu et al., 2021): OPMC integrates multi-view matrix factorization and partition generation into a unified framework, where the learned clustering assignment guides the multi-view factorization process.
4. **AEVC** (Liu et al., 2024): AEVC enhances anchor quality by exploiting correlations among different views and constructs a more reliable consistent anchor graph for multi-view clustering.
5. **LASD** (Zhang et al., 2025b): LASD learns anchors with distributions similar to the original data from an optimal transport perspective and constructs a consensus transport matrix by exploiting complementary and consistent multi-view representations for clustering.
6. **OpenNCD** (Liu et al., 2023): OpenNCD adopts prototype-based contrastive learning and progressively groups prototypes to perform known-class recognition and novel-class discovery.
7. **SNCD** (Wang et al., 2024): SNCD introduces semantic prior information to guide representation learning and clustering, thereby improving novel-category discovery.
8. **IICMVNCD** (Wan et al., 2025): IICMVNCD exploits intra-view distribution consistency and inter-view correlations to transfer knowledge from known classes to novel classes in the multi-view setting.

Table 7: Comparison of ACC, NMI, and Purity for various MVC and NCD algorithms across eight benchmark datasets. The best results are highlighted in bold, while the second-best results are underlined.

Datasets	MVC Methods					NCD Methods			
	Best	Concat	OPMC	AEVC	LASD	OpenNCD	SNCD	IICMVNCD	Ours
ACC									
ProteinFold	30.37	33.51	28.53	32.18	<u>34.98</u>	26.44	30.37	33.77	**36.13**
KIPAN	92.36	89.55	89.86	91.67	<u>92.67</u>	74.73	92.20	92.51	**95.01**
Mfeat	92.10	94.80	93.60	92.70	91.75	96.40	<u>96.70</u>	95.60	**98.10**
Cora	50.77	39.99	42.33	63.20	<u>78.82</u>	22.44	30.74	76.36	**81.04**
Wiki	63.55	29.84	30.77	37.30	65.32	31.13	54.16	<u>65.42</u>	**73.96**
Citeseer	45.51	43.93	41.28	50.12	<u>60.25</u>	31.11	42.91	43.88	**64.96**
Reuters	59.85	59.49	62.77	–	<u>71.83</u>	53.58	46.79	71.56	**72.07**
YTB20	67.64	64.56	<u>81.12</u>	80.69	79.53	74.82	79.02	75.73	**82.70**
NMI									
ProteinFold	29.50	29.36	17.17	**33.18**	30.69	<u>31.45</u>	26.46	28.51	30.80
KIPAN	63.87	55.04	55.00	61.08	<u>64.45</u>	51.45	60.60	62.02	**71.47**
Mfeat	81.23	86.57	84.99	82.67	85.65	<u>91.58</u>	89.89	87.68	**93.76**
Cora	32.78	17.75	10.95	29.96	<u>53.38</u>	1.18	0.66	44.59	**55.16**
Wiki	64.98	6.04	6.06	13.95	61.26	13.31	36.20	<u>66.44</u>	**66.62**
Citeseer	2.90	0.70	1.89	10.39	<u>23.18</u>	0.21	2.00	2.87	**24.23**
Reuters	2.56	2.37	30.74	–	32.35	19.83	7.96	<u>38.52</u>	**38.69**
YTB20	70.23	70.40	<u>76.59</u>	74.52	74.80	75.10	73.89	73.45	**76.73**
Purity									
ProteinFold	41.88	38.48	34.29	39.97	<u>43.96</u>	39.27	34.03	40.05	**44.50**
KIPAN	92.36	89.55	89.86	91.67	<u>92.67</u>	88.14	92.20	92.51	**95.01**
Mfeat	92.10	94.80	93.60	92.70	91.79	<u>97.80</u>	96.70	95.60	**98.10**
Cora	51.77	39.99	42.40	63.20	<u>78.84</u>	32.15	32.02	76.36	**81.04**
Wiki	73.77	37.24	37.37	43.48	73.54	42.32	54.16	<u>75.47</u>	**75.54**
Citeseer	46.06	44.19	49.52	60.91	**70.73**	44.10	49.74	50.71	<u>70.08</u>
Reuters	60.08	59.86	71.83	–	71.83	69.50	61.29	<u>71.92</u>	**74.16**
YTB20	69.96	67.06	81.12	81.01	79.72	<u>81.72</u>	79.02	76.41	**82.70**

E Full Clustering Results

F Additional Analyses

F.1 Parameter Sensitivity

To investigate the influence of hyperparameters on the proposed method, we conduct parameter sensitivity experiments with respect to λ_1 and λ_2. The clustering performance under different parameter combinations is reported in Fig. 4. From the figure, it can be observed that the hyperparameters have a clear influence on the final clustering performance. Nevertheless, the proposed method can maintain relatively stable and competitive performance within a certain range of parameter values across different datasets and metrics. In particular, when λ_1 and λ_2 are selected from moderate ranges, the proposed method usually achieves better and more robust results, indicating that it is not overly sensitive to precise parameter tuning. These observations provide useful guidance for practical hyperparameter selection.

F.2 Convergence Analysis

As analyzed above, the proposed optimization algorithm is solved by an alternating optimization strategy. To further verify its practical convergence behavior, we plot the objective values over iterations on eight benchmark datasets, as shown in Fig. 5. It can be observed that the objective value decreases rapidly in the first several iterations and then gradually becomes stable. For most datasets, the objective function converges within 40 iterations.

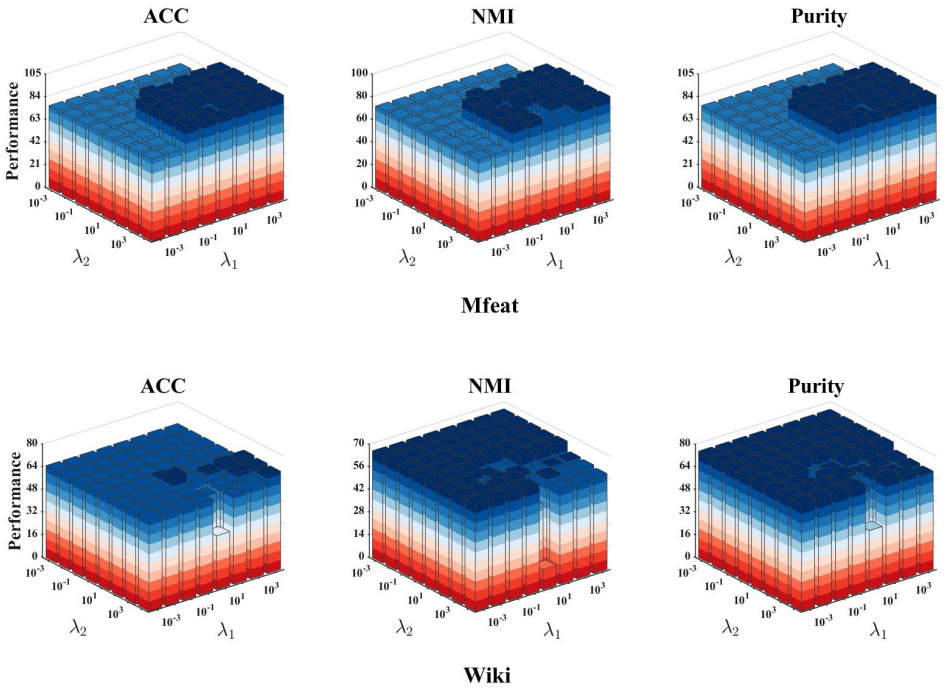


Figure 4: Parameter sensitivity analysis of the proposed method with respect to λ_1 and λ_2 on two representative datasets.

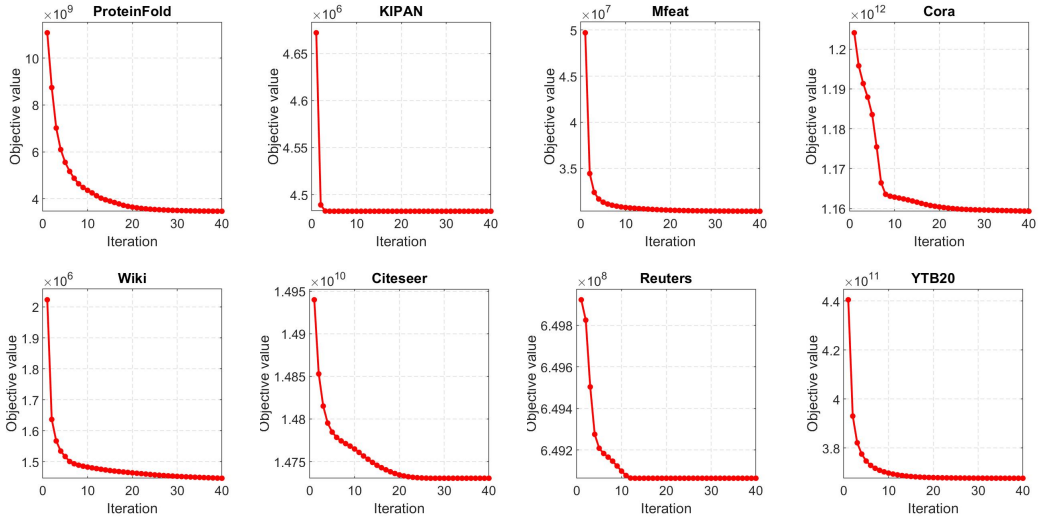


Figure 5: The objective values of the proposed method over iterations on eight benchmark datasets.