

Decoupled Contrastive Multi-View Clustering with High-Order Random Walks

Yiding Lu, Yijie Lin, Mouxing Yang, Dezhong Peng, Peng Hu*, Xi Peng*

College of Computer Science, Sichuan University
 {yidinglu.gm, linyijie.gm, yangmouxing, penghu.ml, pengx.gm}@gmail.com, pengdz@scu.edu.cn

Abstract

In recent, some robust contrastive multi-view clustering (MvC) methods have been proposed, which construct data pairs from neighborhoods to alleviate the false negative issue, *i.e.*, some intra-cluster samples are wrongly treated as negative pairs. Although promising performance has been achieved by these methods, the false negative issue is still far from addressed and the false positive issue emerges because all in- and out-of-neighborhood samples are simply treated as positive and negative, respectively. To address the issues, we propose a novel robust method, dubbed decoupled contrastive multi-view clustering with high-order random walks (DIVIDE). In brief, DIVIDE leverages random walks to progressively identify data pairs in a global instead of local manner. As a result, DIVIDE could identify in-neighborhood negatives and out-of-neighborhood positives. Moreover, DIVIDE embraces a novel MvC architecture to perform inter- and intra-view contrastive learning in different embedding spaces, thus boosting clustering performance and embracing the robustness against missing views. To verify the efficacy of DIVIDE, we carry out extensive experiments on four benchmark datasets comparing with nine state-of-the-art MvC methods in both complete and incomplete MvC settings. The code is released on <https://github.com/XLearning-SCU/2024-AAAI-DIVIDE>.

Introduction

Multi-view data, which can capture diverse and complementary aspects of the same object from different sources or modalities (Jiang et al. 2019; Lin et al. 2023), has become increasingly prevalent in many real-world applications such as human activity recognition (Yadav et al. 2021) and person re-identification (Yang et al. 2022a). As one of the most effective tools to analyze multi-view data, multi-view clustering (MvC) could group unlabeled instances into semantic clusters by exploring and exploiting the latent information hidden in different views. In other words, multi-view representation learning lies at the heart of MvC (Zhang et al. 2019b, 2020).

As a powerful unsupervised representation learning method, contrastive learning (He et al. 2020; Chen et al.




Figure 1: Our observation and basic idea. As illustrated, the existing robust contrastive MvC methods construct negative pairs resorting to ϵ - or k -nearest-neighborhood approaches, which will treat all in-neighborhood samples as positive and out-of-neighborhood samples as negative, resulting in the false positive and false negative issues. Different from these methods, our method identifies the data pairs in a global instead of local manner by performing multi-step random walks on an affinity graph. Thanks to the globality, on the one hand, our method will treat some in-neighborhood samples as negative if high-order neighbors are with a lower affinity score, thus avoiding the construction of FPs to some extent. On the other hand, our method will treat some out-of-neighborhood samples as positive, thus avoiding the construction of FNs in part.

2020), which aims to compact semantically similar pairs (*e.g.*, different views of the same instance) and scatter dissimilar pairs (*e.g.*, different views of different instances) in the embedding space, has achieved state of the arts in MvC (Huang et al. 2019; Zhang et al. 2019a; Trosten et al. 2021). Despite the success achieved by the contrastive MvC methods, they have suffered from the false negative (FN) issue (Yang et al. 2022b). To be specific, most existing contrastive MvC methods will randomly choose some samples to construct negative pairs, and thus some intra-cluster samples are wrongly treated as negative pairs with a high probability. In brief, suppose the data could be grouped into C clusters, there is a probability of $1/C$ to select the intra-cluster samples as negative. To alleviate the performance degradation caused by FNs, some studies (Zhong et al. 2021; Yang et al. 2022b; Huynh et al. 2022; Zheng et al. 2022; Sun et al. 2023) have proposed to rectify FNs through ϵ -

*Corresponding authors

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

neighborhood or k -nearest-neighbors. In brief, for each anchor, its ϵ -neighborhood or k -nearest-neighbors are treated as positive and the remainder is treated as negative. Although the methods have shown remarkable effectiveness in the FNs correction, they will suffer from under- and over-rectification problems. To be exact, the out-of-neighborhood intra-cluster samples will still be treated as negatives, and the in-neighborhood inter-cluster samples might be treated as positives. Notably, in the following, the out-of- and in-neighborhood refers to the first-order relation only.

To address the false negative and false positive (FP) challenges caused by the under- and over-rectification problems, we propose Decoupled contrastive multi-View clustering with high-order random walk (DIVIDE). As illustrated in Fig. 1, given an affinity graph calculated in the embedding space, DIVIDE first performs a random walk for each anchor by choosing a neighbor to move according to the affinity. By applying multi-step random walks, we progressively obtain the affinity between the anchor and the high-order neighbors, and thus the out-of-neighborhood intra-cluster samples could be identified and accordingly the FN issue is alleviated. As another merit of the high-order affinity, the in-neighborhood inter-cluster samples could also be identified and the FP issue is addressed accordingly because the first-order neighbors might be with a smaller high-order affinity. Besides the above contributions, we propose a novel MvC architecture that performs intra-view and inter-view contrastive learning in different spaces with the help of a cross-view decoder. Thanks to the decoupled architecture, the view-specific information could be preserved while learning the cross-view consistency, thus improving MvC performance. In brief, the main contributions of this work could be summarized as follows:

- To alleviate the under- and over-rectification problems faced by existing contrastive MvC methods, we propose using a random walk approach to obtain the high-order affinity so that the in-neighborhood FPs and out-of-neighborhood FNs could be discovered.
- We design a novel decoupled contrastive MvC framework. Different from existing methods, our framework performs intra- and inter-view contrastive learning in different instead of the same space, thus being in favor of the intra-view information and cross-view consistency.
- Extensive experiments verify the superiority of DIVIDE over nine methods on four benchmarks in complete and incomplete multi-view settings.

Related Work

In this section, we briefly review recent developments in deep MvC and contrastive MvC.

Deep Multi-view Clustering

From a self-supervised perspective, most existing deep multi-view clustering methods can be divided into two categories: reconstruction- and contrast-based methods. To be specific, the former typically focuses on preserving intra-view information (*i.e.*, complementarity) by using multiple

autoencoders (Zhang et al. 2019a, 2020; Zhang, Liu, and Fu 2019) to extract compact representations. As the reconstruction loss might retain clustering-agnostic redundancy-like backgrounds that are further used for reconstruction (Lin et al. 2021), the MvC performance would suffer from degradation. In a different way, the contrast-based methods (Lin et al. 2022b; Trosten et al. 2021, 2023) mainly attempt to explore consistent information across different views, which could preserve the discrimination while removing the redundancy from representations. Although contrast-based methods have achieved promising performance, the learned representations might collapse into nearly identical ones (Zolfaghari et al. 2021), hence losing view-specific information.

Different from these existing approaches, our decoupled contrastive framework could preserve view-specific information through cross-view reconstruction, thus enjoying the merits of both the reconstruction- and contrast-based MvC methods. Furthermore, our framework can naturally tackle the missing views thanks to the cross-view reconstruction paradigm. In other words, it does not require elaborate extensions for incomplete settings like most existing MvC methods (Xu et al. 2019; Fang 2023).

Contrastive Multi-view Clustering

Like the vanilla contrastive learning (He et al. 2020; Chen et al. 2020; Liu et al. 2022), contrastive MvC aims to align representations of the data from different views by maximizing the similarity between positive pairs and minimizing that of negative pairs, where the known cross-view samples are treated as positives and the others are treated as negative. Clearly, such a data pair construction method will result in the false negative issue because massive intra-cluster samples are treated as negatives. To address the FN issue, several works have been conducted. For example, Chuang et al. (2020) proposed modeling an unbiased negative probability distribution to reweight the negative terms on the denominator of InfoNCE so that the influence of FNs is alleviated. Zhong et al. (2021); Huynh et al. (2022) proposed using k -nearest-neighbors with the given sample to form the positive pairs, thus reducing the intra-cluster FNs. Moreover, Yang et al. (2022b) proposed dividing positive and negative samples through thresholding a provable similarity metric. In other words, it implicitly constructs positive pairs from a ϵ -neighborhood.

Although these methods have achieved robustness against FNs, most of them have suffered from under- and over-rectification problems as elaborated in Introduction. To solve the problems, unlike these robust methods that construct data pairs from the first-order neighborhood, our method constructs positive pairs from the high-order neighborhoods. More specifically, our DIVIDE employs a random walk approach to discover the out-of-neighborhood positives, thus alleviating the false negative issue. As another merit of the high-order neighborhood, our method will not treat the in-neighborhood negatives as positives, namely, alleviating the false positive issue.

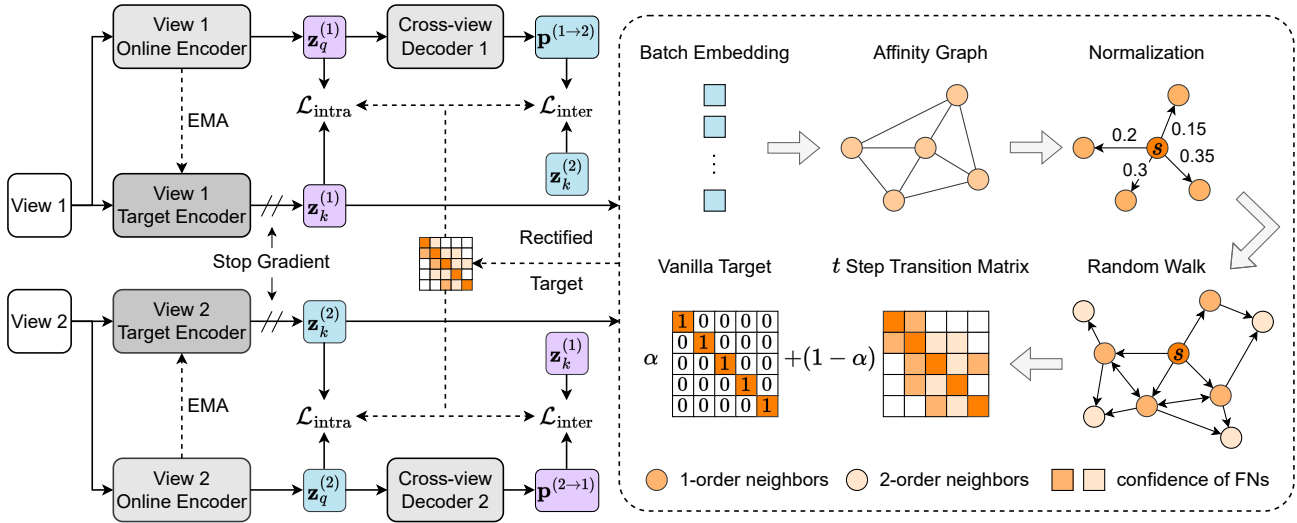


Figure 2: Overview of DIVIDE. For clarity, we use $V = 2$ as an example for illustration. DIVIDE consists of two modules: decoupled contrastive learning framework (the left part) and random-walk-based FNs correction (the right part). Left: data of each view is first passed into the view-specific Siamese encoders to obtain the dual-embedding $\mathbf{z}_q^{(v)}$ and $\mathbf{z}_k^{(v)}$. After that, the cross-view decoder projects $\mathbf{z}_q^{(v)}$ into a new subspace. Finally, the decoupled contrastive losses are applied in two subspaces to learn cross-view consistency and preserve view-specific information; Right: The affinity graph is constructed from each data batch and the corresponding adjacency matrix is normalized as the initial transition matrix. After that, DIVIDE performs a random walk for each anchor by choosing a neighbor to move according to the affinity. By adopting t -step random walks for all anchors, DIVIDE progressively identifies the corresponding high-order neighbors (*i.e.*, false negatives) with a probability matrix which is formulated as a t -step transition matrix \mathbf{M}^t . Finally, DIVIDE uses \mathbf{M}^t to rectify the target of decoupled contrastive losses so that the false negative issue can be addressed.

Method

In this section, we elaborate on the proposed method, *i.e.*, Decoupled contrastive multi-View clustering with high-order random walks (DIVIDE). As shown in Fig. 2, DIVIDE consists of two modules, namely, i) a decoupled contrastive learning framework that is designed to learn cross-view consistency while preserving view-specific information, and ii) a random-walk-based rectifier that is proposed to detect and correct FNs.

Decoupled Contrastive Learning Framework

Given a multi-view dataset $\{(\mathbf{x}_i^{(1)}, \dots, \mathbf{x}_i^{(V)})\}_{i=1}^N$ with N instances observed from V views, contrastive MvC methods (Li et al. 2023; Trosten et al. 2021; Lin et al. 2022b) aim at learning a common representation shared by different views through maximizing the consistency between the embedding extracted from different view-specific encoders. Due to overemphasizing cross-view consistency learning, the view-specific information would be discarded (Tsai et al. 2021; Wang et al. 2022), resulting in performance degradation (Yang et al. 2022b; Zhang et al. 2020). To overcome this challenge, we design a novel MvC framework that consists of view-specific Siamese encoders, cross-view decoders, and decoupled contrastive loss.

View-specific Siamese Encoders. For each view, the Siamese encoder consists of an online encoder $f_q^{(v)}$ and a

target encoder $f_k^{(v)}$ which is a momentum version (He et al. 2020) of $f_q^{(v)}$. More specifically, $f_k^{(v)}$ shares the same architecture with $f_q^{(v)}$ and is updated using the Exponential Moving Average (EMA) of $f_q^{(v)}$.

For a given mini-batch of instances, we first feed them into $f_q^{(v)}$ and $f_k^{(v)}$ to obtain the corresponding view-specific embedding, *i.e.*,

$$\begin{aligned} \mathbf{z}_{q,i}^{(v)} &= f_q^{(v)}(\mathbf{x}_i^{(v)}), \\ \mathbf{z}_{k,i}^{(v)} &= f_k^{(v)}(\mathbf{x}_i^{(v)}). \end{aligned}$$

Cross-view Decoders. As mentioned earlier, simply maximizing the consistency between cross-view embedding may result in the loss of view-specific information. As a remedy, we design a novel cross-view decoder that could preserve the view-specific information without overemphasizing the cross-view consistency. To be specific, we utilize the cross-view decoder $g^{(v \rightarrow u)}$ to project the embedding $\mathbf{z}_{q,i}^{(v)}$ into the embedding space of another view u , and output the embedding $\mathbf{p}_i^{(v \rightarrow u)}$, *i.e.*,

$$\mathbf{p}_i^{(v \rightarrow u)} = g^{(v \rightarrow u)}(\mathbf{z}_{q,i}^{(v)}).$$

Thanks to the cross-view decoders, our framework enjoys the following two merits. On the one hand, consistency learning is performed in two view-specific spaces instead of a single common space like existing methods (Lin

et al. 2022a). As a result, $\mathbf{z}_q^{(v)}$ would achieve the cross-view consistency while preserving the view-specific complementary information, which remarkably boosts the MvC performance as verified in Table 5. On the other hand, the samples could be recovered across different views via the cross-view decoders, thus endowing our method with incomplete multi-view clustering (Xu et al. 2019; Tang and Liu 2022) as shown in Table 1.

Decoupled Contrastive Loss. As elaborated above, the well-designed encoders and decoders aim to decouple the intra- and inter-view learning from a common space into two view-specific spaces. Accompanying the novel architecture, we propose the following decoupled contrastive loss:

$$\begin{aligned} \mathcal{L} &= \mathcal{L}_{\text{intra}} + \mathcal{L}_{\text{inter}} \\ &= \underbrace{\sum_{v=1}^V \mathcal{H}(\mathbf{T}^{(v)}, \rho(\mathbf{z}_q^{(v)}, \mathbf{z}_k^{(v)}))}_{\text{Intra-view Contrastive}} + \underbrace{\sum_{v \neq u}^V \mathcal{H}(\mathbf{T}^{(u)}, \rho(\mathbf{p}^{(v \rightarrow u)}, \mathbf{z}_k^{(u)}))}_{\text{Inter-view Contrastive}}, \end{aligned} \quad (1)$$

where $\mathcal{H}(p, q)$ denotes the cross entropy, $\mathbf{T} \in \mathbb{R}^{n \times n}$ is the pseudo target (Eq. 6) that indicates the given pair is positive or negative, and $\rho(\mathbf{a}, \mathbf{b})$ is the pairwise similarity $s(\cdot, \cdot)$ with the row-wise normalization operator, *i.e.*,

$$[\rho(\mathbf{a}, \mathbf{b})]_{ij} = \frac{\exp(s(\mathbf{a}_i, \mathbf{b}_j)/\tau)}{\sum_{l=1}^n \exp(s(\mathbf{a}_i, \mathbf{b}_l)/\tau)}, \quad (2)$$

where τ is the temperature fixed to 0.5 throughout our experiments.

In general, most existing contrastive MvC methods (Trosten et al. 2021; Li et al. 2023) usually use the off-the-shelf instances as positive pairs and construct negative pairs using random sampling, *i.e.*, \mathbf{T} is set as identity matrix \mathbf{I}_n . However, as mentioned in Introduction, it would inevitably introduce some FNs, thus misleading the model optimization. In the following, we will elaborate on how to endow our method with robustness against FNs.

False Negative Identification via Random Walks

To achieve robustness against FNs, one feasible solution is identifying the potential FNs from negative pairs and rectifying the pseudo target accordingly. To this end, most existing robust contrastive MvC methods (Zhong et al. 2021; Yang et al. 2022b) enlarge the domain of positives and accordingly reduce the amounts of FNs resorting to ϵ -neighborhood or k -nearest-neighbors.

However, such a vanilla neighborhood-based paradigm may cause under- and over-rectification. Namely, the out-of-neighborhood samples within the same cluster would still be treated as negatives while the neighboring inter-cluster samples would be mistakenly treated as positive. To reduce the amounts of FNs and avoid introducing FPs, we propose employing multi-step random walks to progressively identify the high-order neighbors for each anchor. With the high-order neighbors and the corresponding affinity, our method could model the probability of negative pairs being FNs. For ease of presentation, we briefly introduce the preliminary on the random walk. To be exact, a random walk on a graph is

a process that starts at some nodes and moves to other nodes according to the edge weights at each step (Lovász 1993). Formally, the random walk on a graph could be represented by a transition matrix defined as follows.

Definition 1 (Transition Matrix) For an undirected graph \mathcal{G} with n nodes, the random walk transition matrix on this graph is given by

$$\mathbf{M} = \mathbf{A}\mathbf{D}^{-1} \quad (3)$$

where \mathbf{A} is the adjacency matrix of this graph, \mathbf{A}_{ij} denotes the edge weight, $\mathbf{D} = \text{diag}(d_1, d_2, \dots, d_N)$, and $d_i = \sum_j \mathbf{A}_{ij}$.

Clearly, \mathbf{M}_{ij} herein denotes the probability of moving from the i -th node to the j -th node in one step. Similarly, the probability of the position after multi-step random walks could be denoted by:

$$p(t) = p(t-1)\mathbf{M} = \dots = p(0)\mathbf{M}^t \quad (4)$$

where \mathbf{M}^t is t -th power of the transition matrix \mathbf{M} and t is the transmission step.

To identify FNs from the negative pairs, we use each anchor as the starting node and perform t step random walks to find the potential FNs from the high-order neighbors. More specifically, we first construct a fully-connected affinity graph for the in-batch instances by regarding the embedding as nodes and defining the edge weights with the heat kernel similarity:

$$\mathbf{A}_{ij}^{(v)} = \exp(-\|\mathbf{z}_{k,i}^{(v)} - \mathbf{z}_{k,j}^{(v)}\|^2/\sigma), \quad (5)$$

where $\mathbf{z}_{k,i}^{(v)}$ is the i -th anchor embedding, $\mathbf{z}_{k,j}^{(v)}$ is the corresponding negative embedding, and σ is a hyper-parameter fixed as 0.1 throughout our experiments.

After that, one could obtain the random walk transition matrix $\mathbf{M}^{(v)}$ by normalizing the adjacency matrix $\mathbf{A}^{(v)}$ in a row-wise manner. Similarly, the multi-step transition matrix $\mathbf{M}^{(v)t}$ could be obtained, where $\mathbf{M}^{(v)t}_{ij}$ denotes the probability of the j -th negative being the t -th order neighbor (*i.e.*, FNs) for the i -th anchor. In other words, the entry of $\mathbf{M}^{(v)t}$ could serve as the soft target for the contrastive loss (Eq. 1), which denotes the correlation probability of given two samples.

Finally, we use \mathbf{M}^t as the pseudo target of Eq. 1 to achieve the robustness against FNs, *i.e.*,

$$\mathbf{T}^{(v)} = \alpha \mathbf{I}_n + (1 - \alpha)\mathbf{M}^{(v)t} \quad (6)$$

where α is a balanced parameter fixed to 0.5 in our experiments. In the implementation, we construct the affinity graphs for each view and accordingly obtain $\mathbf{T}^{(v)}$. For the intra-view contrastive learning of the v -th view, $\mathbf{T}^{(v)}$ is utilized as the target to improve intra-view discrimination. Such a self-rectified target could better preserve the view-specific information, thus benefitting MvC performance as evaluated in our experiments. In contrast, the affinity of the predicted view is used to construct cross-view target $\mathbf{T}^{(u)}$ for the inter-view contrastive learning. Such a swap-rectified target could improve the cross-view interactions, leading to better cross-view consistency. For a more comprehensive understanding, we conduct ablation study on different target construction strategies in Table 3.

Setting	Method	Scene15			Caltech101			Reuters			LandUse21		
		ACC	NMI	ARI	ACC	NMI	ARI	ACC	NMI	ARI	ACC	NMI	ARI
Incomplete	DCCAE(Wang et al. 2015)	29.0	29.1	12.9	29.1	58.8	23.4	47.0	28.0	14.5	14.9	20.9	3.7
	BMVC(Zhang et al. 2019b)	32.5	30.9	11.6	40.0	58.5	10.2	32.1	7.0	2.9	18.8	18.7	3.7
	PMVC(Li, Jiang, and Zhou 2014)	25.5	25.4	11.3	50.3	74.5	41.5	29.3	7.4	4.4	20.0	23.6	8.0
	DAIMC(Hu and Chen 2018)	27.0	23.5	10.6	<u>56.2</u>	<u>78.0</u>	41.8	40.9	18.7	15.0	19.3	19.5	5.8
	EERIMVC(Liu et al. 2020)	28.9	27.0	8.4	43.6	69.0	26.4	29.8	12.0	4.2	22.1	25.2	9.1
	SURE(Yang et al. 2022b)	39.6	41.6	23.5	34.6	57.8	19.9	47.2	30.9	23.3	<u>23.1</u>	<u>28.6</u>	<u>10.6</u>
	DSIMVC(Tang and Liu 2022)	30.6	35.5	17.2	16.4	24.8	9.2	39.9	19.6	17.1	18.6	18.8	5.7
	DGP(Lin et al. 2022b)	39.5	42.4	23.5	44.3	71.0	<u>45.3</u>	34.6	17.5	2.9	22.2	27.0	10.4
	ProImp(Li et al. 2023)	<u>41.6</u>	<u>42.9</u>	<u>25.3</u>	36.3	65.4	25.4	<u>51.9</u>	<u>35.5</u>	<u>28.5</u>	22.4	26.6	9.9
DIVIDE(Ours)	46.8	45.7	29.1	63.4	82.5	52.4	54.7	37.3	28.6	30.0	35.8	16.0	
Complete	DCCAE(Wang et al. 2015)	34.6	39.0	19.7	45.8	68.6	37.7	42.0	20.3	8.5	15.6	24.4	4.4
	BMVC(Zhang et al. 2019b)	40.5	41.2	24.1	50.1	72.4	33.9	42.4	21.9	15.1	25.3	28.6	11.4
	PMVC(Li, Jiang, and Zhou 2014)	30.8	31.1	15.0	40.5	63.8	28.3	32.5	11.1	7.5	25.0	31.1	12.2
	DAIMC(Hu and Chen 2018)	32.1	33.6	17.4	<u>57.5</u>	<u>78.7</u>	41.9	40.8	21.6	16.0	24.4	29.4	10.3
	EERIMVC(Liu et al. 2020)	39.6	39.0	22.1	49.0	74.2	34.2	33.2	14.3	3.9	24.9	29.6	12.2
	SURE(Yang et al. 2022b)	41.0	43.2	25.0	43.8	70.1	29.5	49.1	29.9	23.6	25.1	28.3	10.9
	DSIMVC(Tang and Liu 2022)	31.7	35.6	17.2	19.7	40.0	19.7	43.2	23.3	19.0	18.1	18.6	5.6
	DGP(Lin et al. 2022b)	41.1	<u>45.1</u>	24.8	51.3	74.8	<u>51.9</u>	36.2	18.9	4.8	<u>26.2</u>	<u>32.7</u>	<u>13.5</u>
	ProImp(Li et al. 2023)	<u>43.6</u>	45.0	<u>26.8</u>	37.6	67.0	25.0	<u>56.5</u>	<u>39.4</u>	<u>32.8</u>	23.7	27.9	10.8
DIVIDE(Ours)	49.1	48.7	31.6	62.2	83.0	50.5	59.3	39.5	29.0	32.3	39.7	18.1	

Table 1: The clustering performance on four multi-view benchmarks. The best and the second best results are denoted in bold and underline. ‘‘Complete’’ setting denotes multi-view data without view missing and ‘‘Incomplete’’ denotes 50% samples are with missing views.

Experiments

To evaluate the effectiveness of DIVIDE, we carry out extensive experiments to answer the following questions: (Q1) Does DIVIDE outperform state-of-the-art multi-view clustering methods? (Q2) Could DIVIDE effectively handle multi-view clustering with missing views? (Q3) Could the multi-step random walk effectively address the false negative issue? (Q4) Does each component of DIVIDE contribute to the overall performance?

Experimental Setup

We briefly introduce the experimental setup here, including the used datasets, settings, and implementation details.

Datasets. We conduct evaluations on four widely-used multi-view benchmarks. In brief,

- **Scene-15** (Li and Perona 2005) contains 4,485 images of 15 scene categories. Following Yang et al. (2022b), we use the PHOG and GIST features as two views.
- **Caltech-101** (Li et al. 2015) consists of 8,677 images of objects from 101 classes. Following Han et al. (2021), we utilize the deep features extracted by DECAF (Krizhevsky, Sutskever, and Hinton 2012) and VGG19 (Simonyan and Zisserman 2014) as two views.
- **Reuters** (Amini, Usunier, and Goutte 2009) is a multi-lingual news dataset with 18,758 samples from various languages. Following Huang et al. (2019), we project the English and French text into a 10-dim latent space using a standard autoencoder.

- **LandUse-21** (Yang and Newsam 2010) contains 2,100 satellite images from 21 classes. Following Lin et al. (2022b), we use PHOG and LBP features as two views.

Experimental Setting. Following most of the compared methods (Hu and Chen 2018; Liu et al. 2020; Lin et al. 2022b; Li et al. 2023), we obtain the clustering results by feeding the concatenated embedding of all views into k -means. For a comprehensive analysis, we use three widely used clustering metrics to evaluate the performance, including Normalized Mutual Information (NMI), Accuracy (ACC), and Adjusted Rand Index (ARI). A higher value of these metrics indicates a better clustering performance. Additionally, we verify the effectiveness of our method in incomplete MvC, *i.e.*, some views of a sample might be missing during data collection. Specifically, we randomly select $m = \eta \times n$ samples and remove one view from each to simulate incomplete data, where η is the missing rate and n is the total number of samples. To handle missing views, we use the cross-view decoders $g^{(v \rightarrow u)}$ to recover the representation of view u via $\mathbf{z}_i^{(u)} = g^{(v \rightarrow u)}(\mathbf{z}_{k,i}^{(v)})$, where v and u denote the observed view and missing view, respectively.

Implementation Details. We implement our method in PyTorch 1.13.0 and run all experiments on NVIDIA 3090 GPUs in Ubuntu 20.04 OS. We train our model for 200 epochs using the Adam optimizer without weight decay. The initial learning rate is set to 2×10^{-3} , and the batch size is fixed to 1024. To warm up, we set the target \mathbf{T} in Eq. (1) as an identity matrix \mathbf{I}_n for the first 100 epochs and then adopt the rectified target Eq. (6) in the remaining epochs. For the other hyper-parameters, we fix the contrastive temperature




Figure 3: Clustering performance under different missing rates on Scene-15. The colored regions denote the standard variances with five random experiments.

$\tau = 0.5$, the temperature of the kernel function $\sigma = 0.1$, the walking step $t = 5$, and the rectified weight $\alpha = 0.5$ throughout experiments. We employ a 4-layer fully connected network (FCN) with batch normalization and ReLU activation to build the view-specific encoder. The cross-view encoder is a 2-layer FCN and the details are placed in the supplementary material.

Comparisons with State of the Arts (Q1)

We compare our DIVIDE with nine state-of-the-art MvC methods, including DCCAE (Wang et al. 2015), BMVC (Zhang et al. 2019b), PMVC (Li, Jiang, and Zhou 2014), EERIMVC (Liu et al. 2020), DCP (Lin et al. 2022b), SURE (Yang et al. 2022b), DSIMVC (Tang and Liu 2022) and ProImp (Li et al. 2023). As shown in Table 1, DIVIDE achieves superior performance and greater robustness over the baselines in either complete or incomplete MvC settings. In particular, it significantly outperforms all baselines in the presence of missing views, indicating its capacity to handle incomplete MvC data. It is noteworthy that SURE (Yang et al. 2022b) is also an FN-tolerant contrastive MvC method that uses a robust contrastive loss to reverse the gradient of FNs. The superiority of DIVIDE over SURE further demonstrates its robustness against FNs.

Performance under Different Missing Rates (Q2)

To investigate the robustness of our proposed method, we conduct experiments by changing the missing rates η . To be specific, we vary the missing rate η from 0% to 90% with a gap of 10% on the Scene-15 dataset. As the missing rate increases, the total number of samples might be smaller than the batch size. To avoid such a case, we set batch size = $\min(1024, (1 - \eta)n)$. Fig. 3 shows the experimental results, which indicate that DIVIDE outperforms all baselines in all incomplete evaluations.

Analysis on Multi-step Random Walks (Q3)

In this section, we investigate the effectiveness of the proposed random walk strategy for FN identification and correction.

Random Walk vs. ϵ - and k -nearest-neighborhood. We compare our random walk strategy with the ϵ - and k -nearest-neighbor-based FN identification strategies in terms

Method	Scene15		Caltech101	
	ACC	KL	ACC	KL
Random Walk (Ours)	49.1	8.2	62.2	5.1
k -nearest-neighbors (Zhong et al. 2021)	48.5	22.3	57.0	17.2
ϵ -Neighborhood (Yang et al. 2022b)	48.4	19.1	60.7	8.5
None	48.0	23.0	54.3	22.1

Table 2: Clustering accuracy and KL divergence of different false negative identification strategies. “None” denotes the case without FN identification.




Figure 4: Clustering accuracy with different walking steps or different neighborhood sizes. “KNN” denotes k -nearest-neighbors and “RW” denotes our random walk paradigm. “-C” and “-I” refer to complete and incomplete multi-view clustering settings.

of clustering accuracy and KL divergence. The KL divergence is calculated between the rectified contrastive target \mathbf{T} (Eq. (6)) and the oracle \mathbf{T}^{GT} constructed by the ground-truth labels. Specifically, $\mathbf{T}_{ij}^{\text{GT}} = 1/|P_i|$ if the i -th and j -th samples belong to the same class, otherwise 0. Here, $|P_i|$ is the number of in-batch samples that share the same label with i -th sample. We use the KL divergence to measure the performance of the rectification, *i.e.*, $D_{\text{KL}}(\mathbf{T}^{\text{GT}} \parallel \mathbf{T})$. Note that the KL divergence not only measures the accuracy of FN identification but also reflects the amount of incorrect rectified FP. A lower KL divergence score indicates more accurate correction, closer to the upper bound \mathbf{T}^{GT} . As shown in Table 2, our method remarkably improves the robustness, enabling more accurate rectification of FNs while avoiding introduc-

Strategy	ACC	NMI	ACC
self & swap	49.1	48.7	31.6
self	48.7	48.5	31.2
swap	48.7	47.9	30.8
concat	48.1	48.1	30.5

Table 3: Influence of different targets for Eq. (1) on Scene-15. “self & swap” is the default strategy as discussed in Method. “self” denotes that the pseudo target $\mathbf{T}^{(v)}$ is obtained from the affinity graph of the given view v , “swap” denotes that the target is obtained from the graph of the predicted view u . “concat” means that the target is obtained from the fused affinity graph constructed by concatenating all views’ embedding.

$\mathcal{L}_{\text{inter}}$	$\mathcal{L}_{\text{intra}}$	decoder	rectify	Scene-15		
				ACC	NMI	ARI
	✓			39.6	39.7	22.5
✓				43.7	43.6	26.1
✓	✓			47.4	46.5	29.6
✓	✓		✓	48.4	47.3	30.4
✓	✓	✓		48.0	48.1	31.0
✓	✓	✓	✓	49.1	48.7	31.6

Table 4: Ablation study of our method on Scene-15, where ✓ denotes the components is adopted..

ing FPs. Besides, we compare our random walk strategy with the k -nearest-neighbor strategy by changing walking steps t for the former and neighborhood size k for the latter. Note that, $t = 0$ corresponds to the case without rectification, *i.e.* $\mathbf{T}^0 = \mathbf{I}_n$. Fig. 4 shows that the higher-order random walk achieves better performance than the lower ones, verifying the effectiveness of our method. In addition, one could observe that DIVIDE is insensitive to the walk step t , which achieves encouraging results with 3–7 steps. In contrast, the k -nearest-neighbor strategy will dramatically drop the performance if k is not elaborately selected.

Different Choices of the Target $\mathbf{T}^{(v)}$. We perform experiments by adopting different rectified targets $\mathbf{T}^{(v)}$ for the decoupled contrastive loss (Eq. (1)). As shown in Table 3, DIVIDE consistently performs well with various rectified targets, demonstrating the effectiveness of our random-walk-based false negative (FN) identification module. Specifically, the “self & swap” strategy achieves better results, which enjoys the merit of both the self target and the swap target, where the self target preserves the intra-view information better and the swap target facilitates the cross-view interactions.

Ablation Studies (Q4)

In this section, we evaluate the role of different components of DIVIDE on the Scene-15 dataset.

Ablation on Model Components. We first investigate the influence of different DIVIDE components and further explore some variations of the decoupled contrastive learning

decoder	target encoder	stop gradient	Scene-15		
			ACC	NMI	ARI
	share		44.0	44.3	26.7
	share	✓	46.7	46.3	29.7
	momentum	✓	48.4	47.3	30.4
✓	share		43.1	44.1	25.9
✓	share	✓	47.6	45.7	28.2
✓	momentum	✓	49.1	48.7	31.6

Table 5: Ablation study of momentum network and cross-view decoder.

paradigm. As shown in Table. 4, we isolate the effects of the intra-view contrastive loss ($\mathcal{L}_{\text{intra}}$), the inter-view contrastive loss ($\mathcal{L}_{\text{inter}}$), the cross-view decoder, and the false negative (FN) rectification strategy. When the cross-view decoder is removed, we directly apply $\mathcal{L}_{\text{inter}}$ to \mathbf{z}_q and \mathbf{z}_k . The experimental results show that $\mathcal{L}_{\text{intra}}$ could enhance the intra-view complementarity, leading to 3.7% and 2.9% improvement in terms of ACC and NMI, respectively. The cross-view decoder helps the model achieve a 0.6% improvement in ACC and a 1.6% improvement in NMI. Moreover, the rectified contrastive objective achieves a gain of 1% in ACC in all settings with and without the cross-view decoder. The results suggest that all components of DIVIDE contribute to the state-of-the-art performance in MvC tasks.

Ablation of Decoupled Contrastive Learning. We conduct further exploration of the decoupled contrastive learning framework by investigating the effectiveness of the momentum updating strategy and the stop-gradient operation of the target encoders. In the experiments, “share” refers to directly using the online encoder as the target encoder. The stop-gradient operation refers to stopping the gradient flow of the target encoder. As shown in Table 5, the performance of is significantly affected when either of these components is removed, indicating that the momentum updating strategy and stop-gradient operation play crucial roles in preserving view-specific information.

Conclusion

This paper proposes a robust contrastive multi-view clustering method that improves existing works in two aspects: i) achieving robustness against missing views through using a decouple contrastive learning framework. In addition, the framework could get an elegant balance between view-specific information preservation and cross-view consistency learning, thus facilitating the MvC performance; ii) achieving robustness against false negatives through using a multi-step random walk to identify the false negatives. As a result, both in-neighborhood negatives and out-of-neighborhood positives could be correctly identified, thus alleviating the false positive and false negative issues. Extensive experiments verify the effectiveness of the proposed method in both complete and incomplete MvC scenarios. In the future, we would like to explore how to extend our method to improve the robustness of contrastive representation learning by specifically incorporating its characteristics.

Acknowledgments

This work was supported in part by NSFC under Grant U21B2040, 62176171, and 62102274; in part by the Fundamental Research Funds for the Central Universities under Grant CJ202303.

References

- Amini, M. R.; Usunier, N.; and Goutte, C. 2009. Learning from multiple partially observed views—an application to multilingual text categorization. *Advances in neural information processing systems*, 22.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A Simple Framework for Contrastive Learning of Visual Representations. In *Proc. Int. Conf. Mach. Learn.*, 1597–1607.
- Chuang, C.-Y.; Robinson, J.; Lin, Y.-C.; Torralba, A.; and Jegelka, S. 2020. Debaised contrastive learning. *Advances in neural information processing systems*, 33: 8765–8775.
- Fang, S. 2023. Incomplete Multi-view Clustering via Diffusion Completion. *arXiv preprint arXiv:2305.11489*.
- Han, Z.; Zhang, C.; Fu, H.; and Zhou, J. T. 2021. Trusted Multi-View Classification. In *Proc. Int. Conf. Learn. Representations*.
- He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. B. 2020. Momentum Contrast for Unsupervised Visual Representation Learning. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 9726–9735.
- Hu, M.; and Chen, S. 2018. Doubly Aligned Incomplete Multi-view Clustering. In *Proc. Int. Joint Conf. Artif. Intell.*, 2262–2268.
- Huang, Z.; Zhou, J. T.; Peng, X.; Zhang, C.; Zhu, H.; and Lv, J. 2019. Multi-view Spectral Clustering Network. In *Proc. Int. Joint Conf. Artif. Intell.*, 2563–2569.
- Huynh, T.; Kornblith, S.; Walter, M. R.; Maire, M.; and Khademi, M. 2022. Boosting contrastive self-supervised learning with false negative cancellation. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, 2785–2795.
- Jiang, Y.; Xu, Q.; Yang, Z.; Cao, X.; and Huang, Q. 2019. DM2C: Deep Mixed-Modal Clustering. In *Proc. Int. Conf. Neural Inf. Process. Syst.*, 5880–5890.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. ImageNet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25.
- Li, F.; and Perona, P. 2005. A bayesian hierarchical model for learning natural scene categories. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 524–531.
- Li, H.; Li, Y.; Yang, M.; Hu, P.; Peng, D.; and Peng, X. 2023. Incomplete Multi-view Clustering via Prototype-based Imputation. *arXiv preprint arXiv:2301.11045*.
- Li, S.-Y.; Jiang, Y.; and Zhou, Z.-H. 2014. Partial multi-view clustering. In *Proceedings of the AAAI conference on artificial intelligence*, volume 28.
- Li, Y.; Nie, F.; Huang, H.; and Huang, J. 2015. Large-Scale Multi-View Spectral Clustering via Bipartite Graph. In *Proc. AAAI Conf. Artif. Intell.*, 2750–2756.
- Lin, F.; Bai, B.; Bai, K.; Ren, Y.; Zhao, P.; and Xu, Z. 2022a. Contrastive multi-view hyperbolic hierarchical clustering. *arXiv preprint arXiv:2205.02618*.
- Lin, Y.; Gou, Y.; Liu, X.; Bai, J.; Lv, J.; and Peng, X. 2022b. Dual contrastive prediction for incomplete multi-view representation learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(4): 4447–4461.
- Lin, Y.; Gou, Y.; Liu, Z.; Li, B.; Lv, J.; and Peng, X. 2021. COMPLETER: Incomplete Multi-view Clustering via Contrastive Prediction. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 11174–11183.
- Lin, Y.; Yang, M.; Yu, J.; Hu, P.; Zhang, C.; and Peng, X. 2023. Graph Matching with Bi-level Noisy Correspondence. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*.
- Liu, J.; Lin, Y.; Jiang, L.; Liu, J.; Wen, Z.; and Peng, X. 2022. Improve Interpretability of Neural Networks via Sparse Contrastive Coding. In *Findings of the Association for Computational Linguistics: EMNLP 2022*.
- Liu, X.; Li, M.; Tang, C.; Xia, J.; Xiong, J.; Liu, L.; Kloft, M.; and Zhu, E. 2020. Efficient and Effective Regularized Incomplete Multi-view Clustering. *IEEE Trans. Pattern Anal. Mach. Intell.*, 43(8): 2634–2646.
- Lovász, L. 1993. Random walks on graphs. *Combinatorics, Paul erdos is eighty*, 2(1-46): 4.
- Simonyan, K.; and Zisserman, A. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
- Sun, W.; Zhang, J.; Wang, J.; Liu, Z.; Zhong, Y.; Feng, T.; Guo, Y.; Zhang, Y.; and Barnes, N. 2023. Learning Audio-Visual Source Localization via False Negative Aware Contrastive Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6420–6429.
- Tang, H.; and Liu, Y. 2022. Deep safe incomplete multi-view clustering: Theorem and algorithm. In *International Conference on Machine Learning*, 21090–21110. PMLR.
- Trosten, D. J.; Lokse, S.; Jenssen, R.; and Kampffmeyer, M. 2021. Reconsidering representation alignment for multi-view clustering. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 1255–1265.
- Trosten, D. J.; Løkse, S.; Jenssen, R.; and Kampffmeyer, M. C. 2023. On the Effects of Self-supervision and Contrastive Alignment in Deep Multi-view Clustering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 23976–23985.
- Tsai, Y.-H. H.; Wu, Y.; Salakhutdinov, R.; and Morency, L.-P. 2021. Self-supervised learning from a multi-view perspective. In *ICLR*.
- Wang, H.; Guo, X.; Deng, Z.-H.; and Lu, Y. 2022. Rethinking minimal sufficient representation in contrastive learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16041–16050.
- Wang, W.; Arora, R.; Livescu, K.; and Bilmes, J. A. 2015. On Deep Multi-View Representation Learning. In *Proc. Int. Conf. Mach. Learn.*, 1083–1092.

- Xu, C.; Guan, Z.; Zhao, W.; Wu, H.; Niu, Y.; and Ling, B. 2019. Adversarial incomplete multi-view clustering. In *IJ-CAI*, volume 7, 3933–3939.
- Yadav, S. K.; Tiwari, K.; Pandey, H. M.; and Akbar, S. A. 2021. A review of multimodal human activity recognition with special emphasis on classification, applications, challenges and future directions. *Knowledge-Based Systems*, 223: 106970.
- Yang, M.; Huang, Z.; Hu, P.; Li, T.; Lv, J.; and Peng, X. 2022a. Learning with Twin Noisy Labels for Visible-Infrared Person Re-Identification. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*
- Yang, M.; Li, Y.; Hu, P.; Bai, J.; Lv, J. C.; and Peng, X. 2022b. Robust Multi-view Clustering with Incomplete Information. *IEEE Trans. Pattern Anal. Mach. Intell.*
- Yang, Y.; and Newsam, S. 2010. Bag-of-visual-words and spatial extensions for land-use classification. In *Proc. ACM SIGSPATIAL Int. Conf. Adv. Inf.*, 270–279.
- Zhang, C.; Cui, Y.; Han, Z.; Zhou, J. T.; Fu, H.; and Hu, Q. 2020. Deep Partial Multi-View Learning. *IEEE Trans. Pattern Anal. Mach. Intell.*
- Zhang, C.; Han, Z.; Cui, Y.; Fu, H.; Zhou, J. T.; and Hu, Q. 2019a. CPM-Nets: Cross Partial Multi-View Networks. In *Proc. Int. Conf. Neural Inf. Process. Syst.*, 557–567.
- Zhang, C.; Liu, Y.; and Fu, H. 2019. AE2-Nets: Autoencoder in Autoencoder Networks. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2577–2585.
- Zhang, Z.; Liu, L.; Shen, F.; Shen, H.; and Shao, L. 2019b. Binary Multi-View Clustering. *IEEE Trans. Pattern Anal. Mach. Intell.*, 41(7): 1774–1782.
- Zheng, L.; Xiong, J.; Zhu, Y.; and He, J. 2022. Contrastive learning with complex heterogeneity. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2594–2604.
- Zhong, H.; Wu, J.; Chen, C.; Huang, J.; Deng, M.; Nie, L.; Lin, Z.; and Hua, X.-S. 2021. Graph Contrastive Clustering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 9224–9233.
- Zolfaghari, M.; Zhu, Y.; Gehler, P.; and Brox, T. 2021. Crossclr: Cross-modal contrastive learning for multi-modal video representations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1450–1459.