

Deep Incomplete Multi-view Clustering with Distribution Dual-Consistency Recovery Guidance

Jiaqi Jin¹, Siwei Wang², Zhibin Dong¹, Xihong Yang¹, Xinwang Liu^{1,*}, En Zhu^{1,*}, Kunlun He³

¹National University of Defense Technology, Changsha, China

²Academy of Military Sciences, Beijing, China ³Chinese PLA General hospital, Beijing, China
{jinjiaqi, xinwangliu, enzhu}@nudt.edu.cn

Abstract

Multi-view clustering leverages complementary representations from diverse sources to enhance performance. However, real-world data often suffer incomplete cases due to factors like privacy concerns and device malfunctions. A key challenge is effectively utilizing available instances to recover missing views. Existing methods frequently overlook the heterogeneity among views during recovery, leading to significant distribution discrepancies between recovered and true data. Additionally, many approaches focus on cross-view correlations, neglecting insights from intra-view reliable structure and cross-view clustering structure. To address these issues, we propose *BURG*, a novel method for incomplete multi-view clustering with *distribUtion dUal-consistency Recovery Guidance*. We treat each sample as a distinct category and perform cross-view distribution transfer to predict the distribution space of missing views. To compensate for the lack of reliable category information, we design a dual-consistency guided recovery strategy that includes intra-view alignment guided by neighbor-aware consistency and cross-view alignment guided by prototypical consistency. Extensive experiments on benchmarks demonstrate the superiority of *BURG* in the incomplete multi-view scenario.

1. Introduction

In today’s data-driven research landscape, multi-view clustering[5, 6, 9, 22, 40, 47, 51] aims to integrate information from different perspectives in an unsupervised manner, enhancing the predictive capability of the model[23, 29]. However, in real-world scenarios, multi-view data is often incomplete due to high acquisition costs or privacy concerns[30, 37, 44, 53]. The missing information may hinder the discovery of underlying patterns and structures within the data, leading to inaccurate or unstable clustering.

*Corresponding author




Figure 1. Basic idea of cross-view distribution transfer recovery. The consistent distribution obtained from the complete views is transferred to the missing views using view-specific flow models. The recovered representation is then input into the view-specific decoder to generate the recovered data.

Therefore, the key to addressing this issue lies in how to utilize the available views to reasonably impute the missing ones[28, 32, 43], thereby restoring the integrity of the data and maintaining the natural alignment of multi-view data[21, 25, 30].

Currently, most mainstream deep incomplete multi-view clustering (DIMVC) methods first impute the missing instances in each view, and then perform clustering based on the complete and recovered views. These methods are commonly referred to as recovery-based DIMVC and can be broadly categorized into three types based on the recovery paradigm: GAN-based[39], graph-structured[4, 41], and information entropy-based methods[26, 27].

However, existing recovery-based methods often assume that all views are similar in the feature space or can be directly mapped to the same feature space. This perspective overlooks the fact that different views may exhibit diverse feature distributions and data characteristics due to their varied sources, which is commonly referred to as the heterogeneity between views. Furthermore, these methods treat the recovery of missing view instances as independent across samples. While this paradigm may be suitable for tasks that seek one-to-one correspondences, the learned representations are not optimal for clustering, as the goal is to map N representations into K clusters. In other words, dur-

ing the training process, these methods neglect the valuable guidance that various consistency structures—both within and across views—can provide for clustering.

Unlike previous methods for recovering missing data, we propose a deep incomplete multi-view clustering method with **d**istriBution **d**Ual-consistency **R**ecovery **G**uidance, abbreviated as BURG, to mitigate the large discrepancy between the data recovered from complete views and the true data. The approximate recovery process for missing instances is outlined in Fig. 1.

First, the complete views of each sample are processed through view-specific encoders and flow models to obtain representations that follow a prior distribution, which are then fused to form a consistent latent distribution. Finally, this representation undergoes the reverse process of the flow model corresponding to the view with missing instances, facilitating cross-view distribution transfer to recover the representation of the missing view. Specifically, due to the lack of manual annotations in an unsupervised setting, we initially treat each sample as independent and map its complete view features into a latent space that follows a Gaussian distribution through the forward process of the flow model. By leveraging the flow model’s ability to explicitly measure distributional discrepancies, we minimize the gap between views, ensuring that the generated data is realistic and reliable. Additionally, without guidance from clustering structure information during recovery, the learned features may lack distinguishability. To address this, we introduce a dual consistency strategy to guide the recovery process: neighbor-aware consistency and prototypical consistency. These ensure the alignment process’s reliability by enforcing consistency in cross-view neighbors and clustering assignments, respectively, enhancing the structural and clustering information within the features.

Further details of the proposed BURG are shown in Fig. 2. The overall model framework is broadly divided into three modules: multi-view feature extraction (MFE), distribution transfer learning (DTL), and dual-consistency guided recovery (DGR). We summarize the major contributions of our work as follows,

- We acknowledge the limitation of existing recovery-based DIMVC methods that neglect inter-view heterogeneity. In our proposed BURG, we employ a flow-based cross-view distribution transfer, which enables precise distribution transfer from observable views to missing views, thus minimizing the gap between the recovered and true data.
- We design a dual consistency strategy to guide both the data recovery process and representation learning. This strategy enriches the learned features with structural and clustering information, enhancing their discriminability and benefiting subsequent clustering tasks.
- Extensive experiments demonstrate the superiority of

the proposed view-specific flow model-based distribution transfer recovery method, as well as the effectiveness of the dual consistency strategy.

2. Related Work

2.1. Deep Incomplete Multi-view Clustering

Deep neural networks[2, 13, 20] enhance clustering performance by effectively modeling associations across views[3, 12, 24, 33, 36, 52], and various methods for deep incomplete multi-view clustering utilize this capability to tackle incomplete data. DIMVC methods can be classified into two main categories based on how they handle missing data: recovery-free and recovery-based approaches.

(1) Recovery-free methods learn representations solely from observable multi-view data, avoiding the complexities of imputing missing data.[10, 38, 49]. APADC[48] introduces an adaptive feature projection to bypass filling in missing data. DVIMC[45] employs the Product-of-Experts approach with consistency constraints to mitigate inter-view information imbalance, achieving a VAE-based recovery-free clustering method.

(2) Recovery-based methods focus on recovering missing data or its representations to mitigate the impact of incomplete data on clustering. These methods can be divided into three categories: *i*) GAN-based methods, which use adversarial strategies to ensure generated data in missing views approximates observed data, such as CPM-Nets[55] and GP-MVC[39]. *ii*) Graph-structured methods, like CRTCC[41], SPCC[8] and ICMVC[4], transfer graph structures from complete views to fill in missing data. *iii*) Information entropy-based methods, including COMPLETER[26] and its variant DCP[27], minimize conditional entropy through dual prediction to recover missing data. DIVIDE[31] follows a similar principle, using cross-view decoders to recover samples across views.

2.2. Normalizing Flow Model

Unlike generative models like GANs[11], Normalizing Flows[7, 16] explicitly model the probability distribution of data, allowing for precise likelihood estimation. Flow-based generative models map complex data distributions to simpler ones, such as Gaussian or logistic distributions, through a series of invertible and differentiable transformations. In practical applications, Normalizing Flow is widely used in areas such as image and text generation[1], anomaly detection[14] and image restoration[42].

Mathematically, let $\mathbf{z} \sim p_Z(\mathbf{z})$ be a latent variable that obeys a simple distribution, and \mathbf{x} is a complex data sample. A flow model defines a reversible transformation from \mathbf{z} to \mathbf{x} via a mapping function F . Using the variable transformation formula, the probability distribution $p(\mathbf{x})$ of \mathbf{x} can be

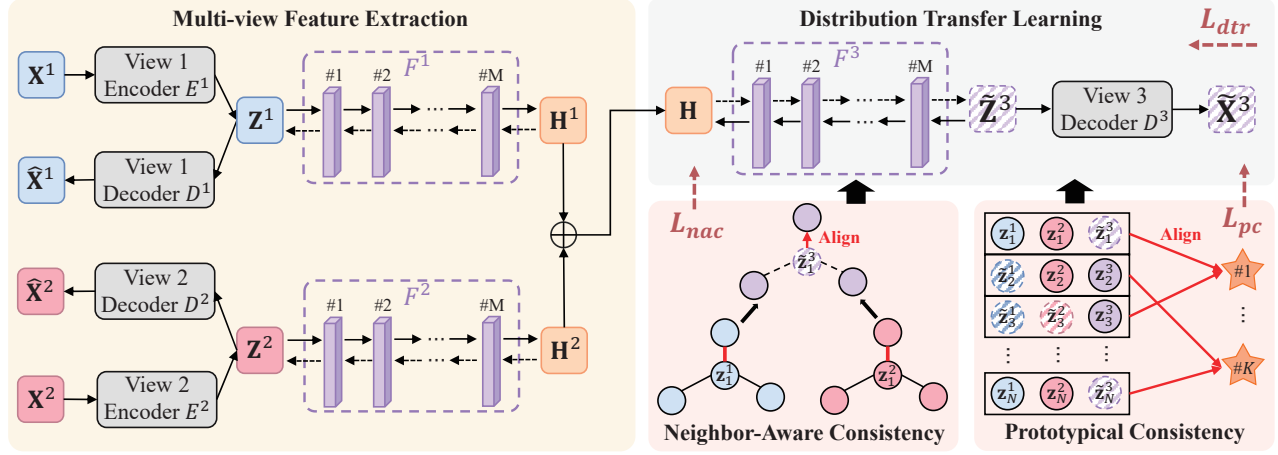


Figure 2. The framework of BURG. We use three views as an example, where the first and second views are complete, while the third is missing. As shown, BURG consists of a joint optimization of three modules: multi-view feature extraction(MFE), distribution transfer learning(DTL), and dual-consistency guided recovery(DGR). The two types of consistency are neighbor-aware and prototypical consistency. In DTL, the common embedding \mathbf{H} is formed by combining \mathbf{H}^1 and \mathbf{H}^2 , which are obtained through view-specific forward flows. Subsequently, \mathbf{H} is passed through reverse flow $(F^3)^{-1}$ to generate the latent representation \mathbf{Z}^3 of the missing view. Furthermore, DGR provides essential local intra-view structure and global inter-view clustering structure throughout the training process to ensure the discriminability of the recovered representation.

expressed as follows,

$$p_X(\mathbf{x}) = p_Z(F^{-1}(\mathbf{x})) \left| \det \left(\frac{\partial F^{-1}(\mathbf{x})}{\partial \mathbf{x}} \right) \right|, \quad (1)$$

where F^{-1} represents the inverse transformation, and $\det \left(\frac{\partial F^{-1}(\mathbf{x})}{\partial \mathbf{x}} \right)$ denotes the determinant of the Jacobian matrix of the inverse transformation.

Flow-based models are highly promising for IMVC because they can model complex distributions while providing precise likelihood estimates. By using invertible transformations, flow network can learn dependencies between different views and recover missing entries.

3. Method

In this section, we will first sequentially introduce the three modules in BURG, as shown in Fig. 2, and then present the objective function and optimization process of the model.

Notations. We denote the multi-view data as $\{\mathbf{X}^v\}_{v=1}^V$, where V is the number of views. Each view $\mathbf{X}^v = \{\mathbf{x}_1^v, \mathbf{x}_2^v, \dots, \mathbf{x}_N^v\} \in \mathbb{R}^{N \times d_v}$ consists of N instances with dimensionality d_v . The set $\mathbf{W} = \{\mathbf{w}^v\}_{v=1}^V$ indicate the presence of missing instances, where $\mathbf{w}^v = \{w_1^v, w_2^v, \dots, w_N^v\} \in \mathbb{R}^N$. Concretely,

$$w_i^v = \begin{cases} 1, & \text{if } i\text{-th sample is complete in } v\text{-th view} \\ 0, & \text{otherwise} \end{cases}. \quad (2)$$

The encoder, decoder and flow network for v -th view are denoted by E^v , D^v and F^v , respectively.

3.1. Multi-view Feature Extraction

3.1.1. Common and View-Specific Reconstruction

For the v -th view, the encoder-decoder pair E^v and D^v is utilized to project different views into a common subspace, with the resulting high-level representation denoted as \mathbf{Z}^v . The projection of encoder is expressed as $E^v(\mathbf{X}^v) : \mathbf{X}^v \in \mathbb{R}^{N \times d_v} \rightarrow \mathbf{Z}^v \in \mathbb{R}^{N \times d}$, and feature reconstruction of decoder is formulated as $D^v(\mathbf{Z}^v) : \mathbf{Z}^v \in \mathbb{R}^{N \times d} \rightarrow \hat{\mathbf{X}}^v \in \mathbb{R}^{N \times d_v}$, where d represents the dimension of the common subspace.

To enhance recovery and alignment in the next modules, we input both the view-specific and common representations into D^v to reconstruct complete instances. The reconstruction loss can be expressed as follows:

$$\ell_i^v = \|\mathbf{x}_i^v - D^v(\mathbf{z}_i^v)\|_2^2 + \|\mathbf{x}_i^v - D^v(\mathbf{z}_i)\|_2^2, \quad (3)$$

where the first and second terms correspond to decoding the view-specific/common representation $\mathbf{z}_i^v/\mathbf{z}_i$ of i -th sample, coefficient w_i^v ensures that reconstruction is performed only on complete instances. The calculation for the two latent representations in Eq. (3) are as follows,

$$\mathbf{z}_i^v = E^v(\mathbf{x}_i^v), \mathbf{z}_i = \frac{1}{\sum_{v=1}^V w_i^v} \sum_{v=1}^V w_i^v \mathbf{z}_i^v. \quad (4)$$

Based on Eqs. (3) and (4), the information reduction during the reconstruction process can be minimized, resulting in effective subspace representations.

3.1.2. Flow-based Distribution Learning

This section formalizes the forward generation process of the flow model, where the embedding \mathbf{z}_i^v obtained from the encoder E^v is input into F^v to produce a representation \mathbf{h}_i^v that follows a standard Gaussian distribution, expressed as $\mathbf{h}_i^v \leftarrow F^v(\mathbf{z}_i^v) \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. Our flow-based generative network F^v consists of M coupling layers and a scaling layer, represented as $F^v = f_1^v \circ f_2^v \circ \dots \circ f_M^v \circ s^v$. Unlike traditional Normalizing Flow, we incorporate affine transformations into each coupling layer to better adapt the model to complex data distributions.

First, we split $\mathbf{z}_i^v \in \mathbb{R}^d$ into two parts along its dimensions, i.e., $\mathbf{z}_i^v = [\mathbf{a}^\top; \mathbf{b}^\top]^\top$. The output of the m -th affine coupling layer is denoted as $\mathbf{h}_{i,m}^v = [\mathbf{a}_m^\top; \mathbf{b}_m^\top]^\top$. The non-linear transformation applied by the coupling layer varies with the parity of m . Concretely,

$$[\mathbf{a}_m; \mathbf{b}_m] = \begin{cases} \begin{bmatrix} \mathbf{a}_{m-1} \odot e^{s_c^m(\mathbf{b}_{m-1})} + t_c^m(\mathbf{b}_{m-1}) \\ \mathbf{b}_{m-1} \end{bmatrix}, & m \text{ is odd} \\ \begin{bmatrix} \mathbf{a}_{m-1} \\ \mathbf{a}_{m-1} \odot e^{s_c^m(\mathbf{b}_{m-1})} + t_c^m(\mathbf{b}_{m-1}) \end{bmatrix}, & m \text{ is even} \end{cases} \quad (5)$$

where $s_c^m(\cdot)$ and $t_c^m(\cdot)$ are the scaling and translation functions within the coupling layer, and \odot denotes element-wise multiplication.

In flow model, Jacobian determinant of the distribution transformation can be computed from the output of the coupling layer. For the m -th coupling layer, the Jacobian matrix J_m can be expressed in the following block form,

$$J_m = \begin{cases} \begin{bmatrix} \text{diag}(e^{s_c^m(\mathbf{b}_{m-1})}) & \frac{\partial t_c^m(\mathbf{b}_{m-1})}{\partial \mathbf{b}_{m-1}} \\ \mathbf{0} & \mathbf{I} \end{bmatrix}, & m \text{ is odd} \\ \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \frac{\partial t_c^m(\mathbf{a}_{m-1})}{\partial \mathbf{a}_{m-1}} & \text{diag}(e^{s_c^m(\mathbf{a}_{m-1})}) \end{bmatrix}, & m \text{ is even} \end{cases} \quad (6)$$

The flow model incorporates a scaling layer after coupling layers to control data scale and enhance expressive capability. The transformation process of the scaling layer is as follows,

$$\mathbf{h}_i^v = \exp(s_\theta^v) \odot [\mathbf{a}_M^\top; \mathbf{b}_M^\top]^\top, \quad (7)$$

where s_θ^v represents the learnable parameters of the scaling layer. Additionally, the Jacobian determinant of the scaling layer is simply the sum of the elements of s_θ^v .

According to the definition in Eq. (1), taking the logarithm of both sides yields the following log-likelihood of the difference between the distributions before and after the transformation,

$$\log p(\mathbf{z}_i^v) = \log p(\mathbf{h}_i^v) + \log |\det J|, \quad (8)$$

where $|\det J|$ is the sum of the products of the Jacobian determinants of all coupling layers, along with the influence of the scaling layer. In our model, we assume that the output \mathbf{h}_i^v of the forward process of the flow model follows a standard Gaussian distribution. Substituting the corresponding probability density function into Eq. (8) yields the following expression,

$$\begin{aligned} \log p(\mathbf{z}_i^v) = & -\frac{1}{2}(d \log 2\pi + \|\mathbf{h}_i^v\|_2^2) \\ & + \log |\det J_s| + \sum_{m=1}^M \log |\det J_m|. \end{aligned} \quad (9)$$

Eq. (9) can explicitly and accurately describe the data distribution, thereby facilitating an effective initialization for our view-specific flow model.

3.2. Distribution Transfer Learning

As shown in Fig. 2, the main goal of DTL module is to transfer the distribution of the common representation obtained from the observable views to the missing view based on the view-specific flow model. This enables the accurate estimation of missing data representations that are highly consistent with the real data.

Assuming that the v -th view instance of the i -th sample is missing while the others are observable, the rough process of recovering the missing representation $\tilde{\mathbf{z}}_i^v$ from the consistent embedding \mathbf{h}_i , which follows a standard Gaussian distribution, through $(F^v)^{-1}$ is as follows,

$$\mathbf{h}_i \xrightarrow{(s^v)^{-1}} \mathbf{h}_{i,M}^v \xrightarrow{(f_M^v)^{-1}} \dots \xrightarrow{(f_2^v)^{-1}} \mathbf{h}_{i,1}^v \xrightarrow{(f_1^v)^{-1}} \tilde{\mathbf{z}}_i^v, \quad (10)$$

where \mathbf{h}_i is derived from all available views of the i -th sample, denoted as $\{\mathbf{z}_i^{v'} | v' = 1, \dots, V, v' \neq v\}$. The process of obtaining the consistent distribution representation \mathbf{h}_i is as follows,

$$\mathbf{h}_i = \frac{1}{\sqrt{V-1}} \sum_{v'=1, v' \neq v}^V F^{v'}(\mathbf{z}_i^{v'}) \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad (11)$$

based on the scaling properties of the mean vector and covariance matrix of the Gaussian distribution, it can be straightforwardly derived that \mathbf{h}_i still follows a standard Gaussian distribution.

Therefore, the error in cross-view distribution transfer within the DTL module can be quantified by the reconstruction loss between the recovered representation and the true representation \mathbf{z}_i^v , which is formally expressed as follows,

$$\begin{aligned} \mathcal{L}_{dtl} = & \sum_{i=1}^N \sum_{v=1}^V \left(\prod_{v'=1}^V w_i^{v'} \right) \left\| (F^v)^{-1}(\mathbf{h}_i) - \mathbf{z}_i^v \right\|_2^2 \\ & + w_i^v (\ell_i^v - \log p(\mathbf{z}_i^v)), \end{aligned} \quad (12)$$

where the first loss term is used to minimize the error in cross-view distribution transfer learning, while the second term is aimed at preserving the multi-view feature extraction capability learned from the MFE module, thereby preventing catastrophic forgetting in the deep model.

3.3. Dual-Consistency Guided Recovery

3.3.1. Neighbor-Aware Consistency

Heterogeneity is reflected in the feature distributions and data characteristics between views, representing a global property. In contrast, similarity between views manifests in the local cross-view structure, representing a local neighbor property. In this work, this local similarity is embodied in the consistency of cross-view nearest neighbor relationships. When the v -th view of a sample is missing, we transfer the nearest neighbor relationships from other views to the v -th view, where the recovered representation $\tilde{\mathbf{z}}_i^v$ resides. The process of identifying cross-view nearest neighbor relationships is as follows,

$$n_a = \min_{\substack{v' \in \{1, \dots, V\} \\ v' \neq v}} \sigma(\Omega^{v'}(\tilde{\mathbf{z}}_i^v), \mathbf{z}_i^{v'}), \quad (13)$$

where $\sigma(\cdot, \cdot)$ denotes the similarity measure function, and $\Omega^{v'}(\tilde{\mathbf{z}}_i^v)$ represents the nearest neighbor representation of $\tilde{\mathbf{z}}_i^v$ in v' -th view.

Subsequently, the nearest neighbor relationship is transferred to the v -th view, and $\tilde{\mathbf{z}}_i^v$ is encouraged to approach its nearest neighbor representation within this view. The loss associated with this process can be formalized as follows,

$$\mathcal{L}_{nac} = \sum_{i=1}^N \sum_{v=1}^V (1 - w_i^v) \sigma(\tilde{\mathbf{z}}_i^v, \mathbf{z}_{n_a}^v). \quad (14)$$

3.3.2. Prototypical Consistency

Based on the assumption that the clusters to which samples belong across different views should be consistent, we propose prototypical consistency. This assumption is essential in multi-view clustering. First, we perform k -means on the consistent representations obtained from the encoder for the complete views, resulting in a set of consistent prototypes $\{\mathbf{c}_k\}_{k=1}^K$ for all samples. Subsequently, the normalized K -dimensional cluster assignment is obtained from the following formula,

$$p(\mathbf{z}_i^v, \mathbf{c}_k) = \frac{\exp(\sigma(\mathbf{z}_i^v, \mathbf{c}_k)/\tau)}{\sum_{k'=1}^K \exp(\sigma(\mathbf{z}_i^v, \mathbf{c}_{k'})/\tau)}, \quad (15)$$

where temperature coefficient τ is fixed at 1.0 in this work. For simplicity, we abbreviate $p(\mathbf{z}_i^v, \mathbf{c}_k)$ as $p_{i,k}^v$ in the following formulas.

In unsupervised scenario, the recovered views cannot be more reliable than complete ones. Therefore, we derive the

Algorithm 1 The Optimization of BURG

Input: Dataset $\{\mathbf{X}^v\}_{v=1}^V$ of size N ; number of clusters K ; networks $\{E^v, D^v, F^v\}_{v=1}^V$; max epoch of three stage E_1, E_2, E_3 ; trade-off parameters α and β .

Output: The predicted clustering results \hat{Y} .

```

1: for  $e = 1$  to  $E_1$  do
2:   Update  $\{E^v, D^v\}_{v=1}^V$  with Eq. (3)
3:   Update  $\{F^v\}_{v=1}^V$  with Eq. (9)
4: end for
5: for  $e = 1$  to  $E_2$  do
6:   Obtain  $\mathbf{h}_i$  with Eq. (11)
7:   Train the entire network with Eq. (12)
8: end for
9: for  $e = 1$  to  $E_3$  do
10:  Obtain the  $n_a$  and  $\tilde{y}_i$  with Eq. (13) and Eq. (16)
11:  Train the entire network with Eq. (18)
12: end for

```

prototype assignment for the entire sample using only the complete views,

$$\tilde{y}_i = \arg \max_{k \in \{1, \dots, K\}} \left(\frac{1}{\sum_{v=1}^V w_i^v} \sum_{v=1}^V p_{i,k}^v \right). \quad (16)$$

Finally, we align both the complete and recovered views of each sample towards its consistent prototype, resulting in the following prototypical consistency alignment loss,

$$\mathcal{L}_{pc} = - \sum_{i=1}^N \left[\sum_{v=1}^V \sum_{k=1}^K (\tilde{y}_{i,k} \log p_{i,k}^v + \gamma p_{i,k}^v \log p_{i,k}^v) \right], \quad (17)$$

where $\tilde{y}_{i,k} \in \tilde{y}_i$ and $\tilde{y}_i = \text{one_hot}(\tilde{y}_i)$. In Eq. (17), the latter of the two loss terms is Shannon entropy, which serves as a regularization term to prevent category collapse, ensuring that not all samples are assigned to the same cluster. The parameter γ is used to balance the strength of the regularization term.

3.4. Objective Function and Optimization

Our overall optimization objective consists of three components: the distribution transfer learning loss, the neighbor-aware consistency loss, and the prototypical consistency loss. In general, the objective loss function of BURG is formulated as follows,

$$\mathcal{L} = \mathcal{L}_{dttl} + \alpha \mathcal{L}_{nac} + \beta \mathcal{L}_{pc}, \quad (18)$$

where α and β are balanced hyperparameters used to adjust the influence of two types of consistency.

The whole optimization procedure is listed in Algorithm 1. After the entire optimization, we concatenate the missing view representations recovered by BURG with the complete views and then perform k -means to obtain the final clustering results.

4. Experiments

To validate the effectiveness of BURG, we conduct extensive experiments to answer the following questions: (Q1) Does BURG outperform current state-of-the-art methods in clustering performance on widely used datasets? (Q2) Do both types of consistency positively contribute to performance? (Q3) Does BURG effectively reduce the discrepancy between recovered data and true data? (Q4) How do the main hyperparameters affect BURG’s performance?

4.1. Experimental Settings

4.1.1. Datasets and Implementation Details

Six commonly-used multi-view datasets are adopted to validate the BURG. **CUB**¹ is a bird classification dataset using the first 10 categories, with two views: GoogLeNet[35] deep visual features and doc2vec[19] textual representations. **HandWritten**² includes digits ‘0’ to ‘9’, with 200 samples per class, using six views: profile correlations, Fourier coefficients, Karhunen-Loeve coefficients, morphological, pixel features, and Zernike moments. **CiteSeer**³ is a graph dataset with 3,312 samples in 6 categories. We use cites, content, inbound, and outbound as four view features instead of the adjacency graph. **Animal**[18] consists of 50 animal categories with 10,158 samples, using DECAF[17] and VGG19[34] deep features as two views. **Reuters**⁴ is a text dataset with 18,758 samples in 6 categories, using documents in German, English, French, Italian, and Spanish as views. **YouTubeFace10**[15] contains 10 person categories with 38,654 samples, using features from video frames: LBP, HOG, GIST, and Gabor.

We evaluate all models in our experiments using the three widely adopted clustering metrics of ACC, NMI and ARI. We adopt Adam optimizer with the learning rate of 0.0003 for all datasets. The epochs of the three stages are set to 200, 30, and 20 respectively. The batch size of the first two stages is 128, and it is increased to 512 in the last stage. Our model is implemented based on PyTorch 2.1.0 and trained on a desktop computer configured with NVIDIA GeForce RTX 3090 and 64G RAM.

4.1.2. Compared Methods

We compare BURG with nine state-of-the-art DIMVC methods. Concretely, **GP-MVC**[39] employs CycleGAN for stable cross-view generation through cyclic training. **COMPLETER**[26] enhances representation consistency via contrastive learning and mutual information maximization. **DCP**[27] combines mutual information maximization with cross-entropy minimization for consistent representations. **SURE**[50] introduces a noise-robust con-

trastive loss to mitigate false negatives. **DIMVC**[46] avoids imputation by using individual view embedding with EM-like optimization. **APADC**[48] presents an imputation-free approach aligning distributions in a common space. **SMILE**[54] discovers invariant semantic distributions without paired samples. **DIVIDE**[31] reduces false negatives using high-order random walks to identify in-cluster samples. **ICMVC**[4] extracts complementary information guided by high confidence. In these methods, only DIMVC and APADC are recovery-free approaches, while the others are recovery-based methods.

4.2. Performance Comparisons(Q1)

We employ three classic metrics to quantify the clustering performance of BURG compared to the other nine methods across all datasets at different missing rates (MR = 0.1, 0.3, 0.5, 0.7) in Table 1. From the results presented in these tables, we have the following observations:

- Overall, whether under low or high missing rates, our proposed BURG demonstrates strong superiority in clustering performance compared to both recovery-based and recovery-free methods. Particularly at higher missing rates, such as MR=0.5, from CUB to YouTubeFace10, BURG outperforms the second-best methods by 1.00%, 18.80%, 7.21%, 2.19%, 1.06%, and 1.45% respectively, showing more pronounced advantages.
- BURG exhibits remarkable performance across diverse datasets: On HandWritten with six views, it achieves over 90% in all metrics at MR=0.1; On CiteSeer, it outperforms GCN-based ICMVC without utilizing adjacency graphs; On deep feature datasets (Animal) and text datasets (Reuters), the advantages are relatively smaller due to their inherent strong representations. Notably, on the largest dataset YouTubeFace10, BURG maintains optimal performance, achieving over 80% in both ACC and NMI at MR=0.3 (ACC=80.69%, NMI=82.47%), demonstrating excellent scalability.

4.3. Model Analysis(Q2)

4.3.1. Ablation Study

We performed ablation studies on HandWritten and Animal datasets to evaluate our proposed neighbor-aware consistency (NAC) and prototypical consistency (PC) mechanisms. Table 2 compares four configurations at 0.5 missing rate: no consistency strategies (None), single strategies (NAC Only/PC Only), and the full model (NAC+PC).

NAC maintains local topology during missing view reconstruction, reducing noise from view heterogeneity. On the Animal dataset (50 clusters), NAC alone improves ACC from 28.44% to 48.61%. Meanwhile, PC ensures cross-view cluster compactness by aligning recovered embeddings with global prototypes. Notably, PC generally contributes more significantly than NAC - removing PC causes

¹<http://www.vision.caltech.edu/visipedia/CUB-200.html>

²<https://archive.ics.uci.edu/ml/datasets/Multiple+Features>

³<https://linqs-data.soe.ucsc.edu/public/lbc/citeseer.tgz>

⁴<https://archive.ics.uci.edu/ml/datasets.html>

	Missing Rates	0.1			0.3			0.5			0.7		
		ACC(%)	NMI(%)	ARI(%)	ACC(%)	NMI(%)	ARI(%)	ACC(%)	NMI(%)	ARI(%)	ACC(%)	NMI(%)	ARI(%)
CUB	GP-MVC	<u>67.33</u>	65.57	50.15	55.00	60.26	42.98	<u>56.33</u>	54.90	<u>39.03</u>	41.33	46.24	24.97
	COMPLETER	52.25	65.06	46.54	49.33	61.76	44.42	41.84	49.61	32.46	19.22	14.84	5.46
	DCP	61.25	66.28	47.80	56.50	59.08	<u>45.38</u>	32.08	35.71	16.92	28.25	25.14	13.65
	SURE	67.04	63.12	50.12	53.80	46.47	33.14	51.03	43.76	25.33	40.20	31.29	18.36
	DIMVC	60.98	60.77	44.86	58.26	50.66	36.82	42.76	37.20	30.37	37.58	27.64	21.63
	APADC	46.33	41.20	28.99	34.83	33.21	20.28	30.94	28.62	14.06	29.28	24.35	9.35
	SMILE	62.83	67.99	42.80	55.80	59.08	37.16	54.33	<u>56.16</u>	31.29	<u>44.17</u>	<u>48.62</u>	27.34
	DIVIDE	62.50	59.75	44.47	47.83	50.84	34.08	45.67	46.04	28.20	34.50	34.51	21.54
	ICMVC	62.83	66.54	<u>50.18</u>	57.00	<u>60.58</u>	43.13	45.67	50.79	33.48	42.19	45.80	<u>30.71</u>
BURG	71.50	71.60	60.48	<u>58.00</u>	61.85	45.86	57.33	58.89	45.12	53.83	56.58	45.12	
HandWritten	GP-MVC	91.20	83.52	<u>81.56</u>	<u>86.30</u>	76.67	72.51	75.40	66.34	57.86	69.55	61.80	51.41
	COMPLETER	70.50	75.46	62.42	63.45	63.80	49.40	50.57	50.80	38.64	35.65	32.12	21.23
	DCP	75.32	77.24	67.39	75.57	78.36	69.34	74.20	<u>72.34</u>	59.83	62.40	60.60	40.79
	SURE	72.40	73.62	62.19	70.21	71.15	60.64	65.59	63.28	53.23	62.95	54.59	44.46
	DIMVC	81.99	76.69	61.88	80.90	75.66	61.57	68.77	65.53	46.97	67.35	65.92	46.14
	APADC	89.88	84.94	76.45	74.63	72.83	61.89	70.35	68.41	54.77	68.92	<u>66.51</u>	50.99
	SMILE	87.95	82.66	70.25	85.40	<u>83.45</u>	68.92	71.05	65.30	55.07	70.20	<u>60.69</u>	<u>54.20</u>
	DIVIDE	91.85	86.77	78.61	78.75	65.24	58.01	75.45	61.56	52.02	<u>74.75</u>	57.11	50.42
	ICMVC	<u>92.55</u>	<u>87.56</u>	80.00	83.75	82.17	<u>73.70</u>	<u>75.50</u>	69.85	<u>60.04</u>	68.25	53.69	43.28
BURG	95.70	90.83	90.66	95.30	89.74	89.83	94.30	88.36	87.76	79.75	84.27	75.30	
CiteSeer	GP-MVC	23.28	13.96	1.85	22.16	13.07	1.54	20.89	7.81	0.94	20.56	7.10	0.55
	COMPLETER	28.08	11.16	5.47	27.05	11.99	4.91	24.43	8.93	2.21	22.31	5.03	1.64
	DCP	39.89	20.64	9.74	38.95	18.40	9.09	33.91	11.17	4.99	31.10	9.21	2.97
	SURE	32.61	15.18	8.76	30.59	11.22	6.04	28.77	8.89	2.65	24.55	4.72	0.42
	DIMVC	41.27	25.46	8.10	37.77	18.65	4.56	33.21	12.97	3.87	32.19	10.90	2.83
	APADC	37.44	23.96	11.97	36.93	19.07	8.93	30.10	17.81	6.03	27.81	12.10	<u>5.57</u>
	SMILE	41.73	25.34	7.76	37.77	18.53	4.59	36.59	14.56	4.50	32.52	10.87	2.91
	DIVIDE	43.36	<u>25.54</u>	<u>14.00</u>	38.01	18.57	8.42	35.75	14.44	4.74	30.22	10.11	4.14
	ICMVC	<u>44.19</u>	24.69	13.47	<u>40.40</u>	<u>20.02</u>	<u>11.43</u>	<u>40.22</u>	<u>20.53</u>	<u>10.09</u>	<u>33.85</u>	<u>12.35</u>	5.22
BURG	45.38	27.92	14.61	45.41	20.32	11.60	47.43	21.70	12.36	35.69	14.61	7.56	
Animal	GP-MVC	44.01	54.76	30.50	43.38	53.61	29.19	39.29	48.12	24.02	38.58	47.57	22.51
	COMPLETER	29.81	46.66	17.83	24.54	39.20	15.30	20.84	34.82	13.58	15.78	30.62	10.54
	DCP	28.86	47.14	20.00	28.14	47.04	18.84	26.34	44.00	16.25	24.79	43.62	14.98
	SURE	38.23	49.14	25.48	36.38	48.92	24.08	35.94	47.06	23.70	34.03	44.34	21.55
	DIMVC	48.29	60.20	35.79	<u>47.74</u>	<u>59.72</u>	<u>37.71</u>	<u>46.42</u>	<u>56.31</u>	33.46	45.27	53.06	28.31
	APADC	30.23	48.44	25.01	29.15	46.72	23.31	28.79	45.51	22.72	27.33	43.39	20.79
	SMILE	48.46	60.78	<u>39.48</u>	46.85	57.65	36.80	45.31	54.67	<u>34.10</u>	40.90	<u>53.83</u>	<u>30.70</u>
	DIVIDE	<u>49.18</u>	<u>61.64</u>	38.58	42.70	53.40	30.45	38.51	49.91	26.37	34.42	45.42	20.89
	ICMVC	45.10	60.73	33.62	41.17	57.19	31.55	38.34	54.30	28.76	34.67	50.71	25.14
BURG	49.24	62.84	41.60	49.11	61.89	40.28	48.61	59.08	39.71	<u>44.55</u>	55.62	34.80	
Reuters	GP-MVC	32.39	19.47	16.91	30.96	17.20	13.96	27.48	12.39	8.65	23.73	10.45	6.44
	COMPLETER	39.54	21.43	12.81	37.59	20.85	11.77	32.76	13.03	8.46	26.97	8.32	4.81
	DCP	<u>43.51</u>	23.95	16.96	41.89	22.95	15.97	40.19	21.64	14.42	39.03	20.91	13.23
	SURE	41.32	25.66	13.32	40.44	22.40	10.59	37.63	21.28	7.58	37.03	18.21	6.63
	DIMVC	43.47	24.90	18.33	<u>42.51</u>	<u>24.42</u>	17.94	<u>41.91</u>	<u>21.72</u>	<u>17.76</u>	<u>40.81</u>	19.99	13.81
	APADC	39.88	21.51	20.11	36.80	20.14	18.76	31.72	17.42	15.79	29.99	16.48	<u>15.08</u>
	SMILE	42.39	23.78	15.43	39.03	22.51	15.02	36.08	20.51	14.04	35.58	18.98	12.03
	DIVIDE	40.90	23.64	18.96	39.59	22.55	14.78	38.01	21.62	15.81	37.86	20.98	15.01
	ICMVC	42.95	<u>26.28</u>	<u>20.72</u>	40.89	23.89	<u>19.25</u>	39.48	21.01	16.65	37.64	<u>22.81</u>	14.17
BURG	44.01	26.39	21.15	42.86	26.38	20.22	42.97	24.77	19.61	41.84	23.63	20.60	
YouTubeFace10	GP-MVC	76.81	<u>82.44</u>	<u>72.73</u>	73.68	<u>80.70</u>	<u>70.81</u>	65.98	73.28	53.52	63.69	71.07	50.11
	COMPLETER	62.73	60.79	39.82	51.89	57.72	35.67	52.39	57.42	29.99	49.05	53.92	24.52
	DCP	64.46	72.97	56.48	59.39	59.65	40.50	57.62	60.50	27.54	51.39	56.29	27.29
	SURE	72.65	72.70	55.95	65.39	66.29	40.60	52.90	58.76	32.12	49.76	56.12	26.04
	DIMVC	74.94	79.34	68.96	73.87	76.42	63.97	<u>68.96</u>	<u>75.44</u>	60.29	61.77	61.32	54.59
	APADC	56.75	47.16	40.36	48.09	39.66	31.89	39.71	33.97	23.95	29.23	25.55	14.43
	SMILE	73.50	78.04	66.55	72.47	77.32	64.37	67.91	73.86	<u>60.68</u>	64.74	66.79	41.52
	DIVIDE	<u>77.67</u>	81.11	69.44	<u>75.07</u>	77.86	63.33	68.92	75.34	60.50	<u>67.12</u>	<u>73.10</u>	<u>57.22</u>
	ICMVC	68.90	74.17	60.06	67.44	74.59	60.02	62.76	65.44	53.03	58.00	58.67	47.14
BURG	78.92	83.68	72.97	80.69	82.47	72.79	70.41	76.08	61.17	67.86	75.00	59.23	

Table 1. The clustering results on six multi-view benchmarks under different missing rates. The optimal results are highlighted in **bold**, while the second-best are underlined.

substantial performance drops, as seen in HandWritten where ACC decreases from 94.30% to 75.90%. These findings highlight PC’s crucial role in bridging semantic gaps across heterogeneous views.

4.3.2. Comparative Verification of Dual Consistency

To ensure that the positive impact on clustering performance stems from our dual-consistency recovery guidance rather than simply from the additional 20 epochs of training,

we present in Fig. 3 the ACC curves across epochs with and without dual consistency (DC). These results clearly validate that the improvements in clustering performance are indeed attributable to our DC mechanism.

4.3.3. Analysis of Inter-View Correlation Enhancement

To investigate whether our proposed BURG effectively enhances the correlation consistency across different views, we evaluated its performance with varying numbers of

Dataset	Dual Consistency	ACC(%)	NMI(%)	ARI(%)
HandWritten	None	52.65	58.15	45.45
	NAC Only	75.90	67.36	46.57
	PC Only	75.80	81.19	70.60
	NAC + PC	94.30	88.36	87.76
Animal	None	28.44	38.69	19.34
	NAC Only	47.32	57.70	33.93
	PC Only	46.64	58.70	38.27
	NAC + PC	48.61	59.08	39.71

Table 2. Effect of the dual-consistency guided recovery strategy on HandWritten and Animal at missing rate of 0.5. The best results are highlighted in **bold**.




Figure 3. Comparison of clustering performance with and without the DC module from 230 to 250 epochs at missing rate of 0.5.

views on HandWritten and Reuters datasets, which contain the highest number of views among all datasets. The results are presented in Table 3. We observe that the performance of our model demonstrates an upward trend on both datasets as the number of views increases.

Dataset	View Number	ACC(%)	NMI(%)	ARI(%)
HandWritten	2V	72.00	71.23	60.67
	3V	78.15	77.04	69.14
	4V	81.10	82.27	75.09
	5V	92.65	87.56	85.07
	6V	94.30	88.36	87.76
Reuters	2V	40.78	18.52	16.29
	3V	39.21	19.91	17.77
	4V	42.48	24.07	22.14
	5V	42.97	24.77	19.61

Table 3. Experiment with different numbers of views on HandWritten and Reuters at missing rate of 0.5. The best results are highlighted in **bold**.

4.4. Visualization of the Discrepancy(Q3)

To validate whether our proposed BURG effectively reduces the discrepancy between recovered and ground-truth data, we conducted t-SNE visualization on the HandWritten dataset under a missing rate of 0.5, as illustrated in Fig. 4. In the visualization, blue points represent samples with missing views, while red points denote samples with complete views. The visualization results strongly substantiate




Figure 4. Visualization of the discrepancy between recovered and ground-truth data on the HandWritten dataset before and after training, with a missing rate of 0.5.

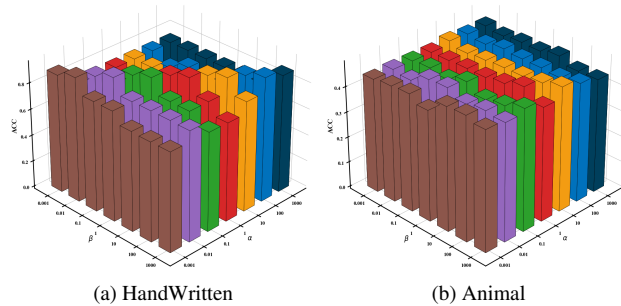


Figure 5. Sensitivity analysis of α and β for our method over HandWritten and Animal (The missing rate is 0.5).

our motivation.

4.5. Parameter Analysis(Q4)

We conducted parameter sensitivity analysis on two trade-off parameters, α and β , in Eq. (18) using the HandWritten and Animal datasets, as illustrated in Fig. 5. On the HandWritten dataset, we observed superior clustering performance when the ratio of α to β was relatively large, indicating that prototypical consistency plays a dominant role. In contrast, the Animal dataset exhibited minimal fluctuations in clustering results.

5. Conclusion

In this paper, we propose a distribution transfer recovery method with dual consistency to restore missing instances in multi-view data. Compared to existing methods, our approach introduces two key improvements: i) utilizing flow-based generative model, it integrates the distributions of available views and transfers them to missing views, effectively predicting the distribution space of the missing views and addressing inter-view heterogeneity; ii) during the recovery process, it incorporates neighbor-aware and consistent prototypical guidance, providing rich intra-view structural information and inter-view clustering information, ensuring the distinguishability of the recovered views.

Acknowledgments

This work is supported by National Science and Technology Innovation 2030 Major Project under Grant No. 2022ZD0209103, the Major Program Project of Xiangjiang Laboratory under Grant No. 24XJCYJ01002, the National Science Fund for Distinguished Young Scholars of China under Grant No. 62325604, the National Natural Science Foundation of China (project No. 62276271, 62406329, 62476281), and the National Natural Science Foundation of China Joint Found under Grant No. U24A20323.

References

- [1] Rameen Abdal, Peihao Zhu, Niloy J Mitra, and Peter Wonka. Styleflow: Attribute-conditioned exploration of stylegan-generated images using conditional continuous normalizing flows. *ACM Transactions on Graphics (ToG)*, 40(3):1–21, 2021. 2
- [2] Huayue Cai, Xiang Zhang, Long Lan, Guohua Dong, Chuanfu Xu, Xinwang Liu, and Zhigang Luo. Learning deep discriminative embeddings via joint rescaled features and log-probability centers. *Pattern Recognition*, 114:107852, 2021. 2
- [3] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In *Proceedings of the European conference on computer vision (ECCV)*, pages 132–149, 2018. 2
- [4] Guoqing Chao, Yi Jiang, and Dianhui Chu. Incomplete contrastive multi-view clustering with high-confidence guiding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 11221–11229, 2024. 1, 2, 6
- [5] Man-Sheng Chen, Jia-Qi Lin, Xiang-Long Li, Bao-Yu Liu, Chang-Dong Wang, Dong Huang, and Jian-Huang Lai. Representation learning in multi-view clustering: A literature review. *Data Science and Engineering*, 7(3):225–241, 2022. 1
- [6] Chenhang Cui, Yazhou Ren, Jingyu Pu, Jiawei Li, Xiaorong Pu, Tianyi Wu, Yutao Shi, and Lifang He. A novel approach for effective multi-view clustering with information-theoretic perspective. *Advances in Neural Information Processing Systems*, 36, 2024. 1
- [7] Laurent Dinh, David Krueger, and Yoshua Bengio. Nice: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516*, 2014. 2
- [8] Zhibin Dong, Jiaqi Jin, Yuyang Xiao, Bin Xiao, Siwei Wang, Xinwang Liu, and En Zhu. Subgraph propagation and contrastive calibration for incomplete multiview data clustering. *IEEE Transactions on Neural Networks and Learning Systems*, 2024. 2
- [9] Uno Fang, Man Li, Jianxin Li, Longxiang Gao, Tao Jia, and Yanchun Zhang. A comprehensive survey on multi-view clustering. *IEEE Transactions on Knowledge and Data Engineering*, 35(12):12350–12368, 2023. 1
- [10] Wei Feng, Guoshuai Sheng, Qianqian Wang, Quanxue Gao, Zhiqiang Tao, and Bo Dong. Partial multi-view clustering via self-supervised network. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 11988–11995, 2024. 2
- [11] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014. 2
- [12] Xifeng Guo, Long Gao, Xinwang Liu, and Jianping Yin. Improved deep embedded clustering with local structure preservation. In *Ijcai*, pages 1753–1759, 2017. 2
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 2
- [14] Or Hirschorn and Shai Avidan. Normalizing flows for human pose anomaly detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13545–13554, 2023. 2
- [15] Dong Huang, Chang-Dong Wang, and Jian-Huang Lai. Fast multi-view clustering via ensembles: Towards scalability, superiority, and simplicity. *IEEE Transactions on Knowledge and Data Engineering*, 35(11):11388–11402, 2023. 6
- [16] Durk P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. *Advances in neural information processing systems*, 31, 2018. 2
- [17] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25, 2012. 6
- [18] Christoph H Lampert, Hannes Nickisch, and Stefan Harmeling. Attribute-based classification for zero-shot visual object categorization. *IEEE transactions on pattern analysis and machine intelligence*, 36(3):453–465, 2013. 6
- [19] Quoc Le and Tomas Mikolov. Distributed representations of sentences and documents. In *International conference on machine learning*, pages 1188–1196. PMLR, 2014. 6
- [20] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436–444, 2015. 2
- [21] Lusi Li, Zhiqiang Wan, and Haibo He. Incomplete multi-view clustering with joint partition and graph learning. *IEEE Transactions on Knowledge and Data Engineering*, 35(1):589–602, 2021. 1
- [22] Liang Li, Yuangang Pan, Jie Liu, Yue Liu, Xinwang Liu, Kenli Li, Ivor W Tsang, and Keqin Li. Bgae: Auto-encoding multi-view bipartite graph clustering. *IEEE Transactions on Knowledge and Data Engineering*, 36(8):3682–3696, 2024. 1
- [23] Yingming Li, Ming Yang, and Zhongfei Zhang. A survey of multi-view representation learning. *IEEE transactions on knowledge and data engineering*, 31(10):1863–1883, 2018. 1
- [24] Yunfan Li, Mouxing Yang, Dezhong Peng, Taihao Li, Jiantao Huang, and Xi Peng. Twin contrastive learning for online clustering. *International Journal of Computer Vision*, 130(9):2205–2221, 2022. 2
- [25] Naiyao Liang, Zuyuan Yang, and Shengli Xie. Incomplete multi-view clustering with sample-level auto-weighted graph fusion. *IEEE Transactions on Knowledge and Data Engineering*, 35(6):6504–6511, 2022. 1

- [26] Yijie Lin, Yuanbiao Gou, Zitao Liu, Boyun Li, Jiancheng Lv, and Xi Peng. Completer: Incomplete multi-view clustering via contrastive prediction. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11174–11183, 2021. 1, 2, 6
- [27] Yijie Lin, Yuanbiao Gou, Xiaotian Liu, Jinfeng Bai, Jiancheng Lv, and Xi Peng. Dual contrastive prediction for incomplete multi-view representation learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(4): 4447–4461, 2022. 1, 2, 6
- [28] Cheng Liu, Rui Li, Hangjun Che, Man-Fai Leung, Si Wu, Zhiwen Yu, and Hau-San Wong. Latent structure-aware view recovery for incomplete multi-view clustering. *IEEE Transactions on Knowledge and Data Engineering*, 2024. 1
- [29] Xinwang Liu. Simplemkm: Simple multiple kernel k-means. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(4):5174–5186, 2022. 1
- [30] Xinwang Liu, Xinzong Zhu, Miaomiao Li, Lei Wang, Chang Tang, Jianping Yin, Dinggang Shen, Huaimin Wang, and Wen Gao. Late fusion incomplete multi-view clustering. *IEEE transactions on pattern analysis and machine intelligence*, 41(10):2410–2423, 2018. 1
- [31] Yiding Lu, Yijie Lin, Mouxing Yang, Dezhong Peng, Peng Hu, and Xi Peng. Decoupled contrastive multi-view clustering with high-order random walks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 14193–14201, 2024. 2, 6
- [32] Jingyu Pu, Chenhang Cui, Xinyue Chen, Yazhou Ren, Xiaorong Pu, Zhifeng Hao, S Yu Philip, and Lifang He. Adaptive feature imputation with latent graph for deep incomplete multi-view clustering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 14633–14641, 2024. 1
- [33] Yazhou Ren, Jingyu Pu, Zhimeng Yang, Jie Xu, Guofeng Li, Xiaorong Pu, S Yu Philip, and Lifang He. Deep clustering: A comprehensive survey. *IEEE Transactions on Neural Networks and Learning Systems*, 2024. 2
- [34] K Simonyan and A Zisserman. Very deep convolutional networks for large-scale image recognition. In *3rd International Conference on Learning Representations (ICLR 2015)*. Computational and Biological Learning Society, 2015. 6
- [35] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015. 6
- [36] Xinhang Wan, Jiyuan Liu, Hao Yu, Qian Qu, Ao Li, Xinwang Liu, Ke Liang, Zhibin Dong, and En Zhu. Contrastive continual multiview clustering with filtered structural fusion. *IEEE Transactions on Neural Networks and Learning Systems*, 2024. 2
- [37] Xinhang Wan, Bin Xiao, Xinwang Liu, Jiyuan Liu, Weixuan Liang, and En Zhu. Fast continual multi-view clustering with incomplete views. *IEEE Transactions on Image Processing*, 33:2995–3008, 2024. 1
- [38] Qianqian Wang, Huanhuan Lian, Gan Sun, Quanxue Gao, and Licheng Jiao. icmsc: Incomplete cross-modal subspace clustering. *IEEE Transactions on Image Processing*, 30:305–317, 2020. 2
- [39] Qianqian Wang, Zhengming Ding, Zhiqiang Tao, Quanxue Gao, and Yun Fu. Generative partial multi-view clustering with adaptive fusion and cycle consistency. *IEEE Transactions on Image Processing*, 30:1771–1783, 2021. 1, 2, 6
- [40] Siwei Wang, Xinwang Liu, Suyuan Liu, Wenxuan Tu, and En Zhu. Scalable and structural multi-view graph clustering with adaptive anchor fusion. *IEEE Transactions on Image Processing*, pages 1–1, 2024. 1
- [41] Yiming Wang, Dongxia Chang, Zhiqiang Fu, Jie Wen, and Yao Zhao. Incomplete multiview clustering via cross-view relation transfer. *IEEE Transactions on Circuits and Systems for Video Technology*, 33(1):367–378, 2022. 1, 2
- [42] Yufei Wang, Renjie Wan, Wenhan Yang, Haoliang Li, Lap-Pui Chau, and Alex Kot. Low-light image enhancement with normalizing flow. In *Proceedings of the AAAI conference on artificial intelligence*, pages 2604–2612, 2022. 2
- [43] Jie Wen, Zheng Zhang, Yong Xu, Bob Zhang, Lunke Fei, and Hong Liu. Unified embedding alignment with missing views inferring for incomplete multi-view clustering. In *Proceedings of the AAAI conference on artificial intelligence*, pages 5393–5400, 2019. 1
- [44] Chang Xu, Dacheng Tao, and Chao Xu. Multi-view learning with incomplete views. *IEEE Transactions on Image Processing*, 24(12):5812–5825, 2015. 1
- [45] Gehui Xu, Jie Wen, Chengliang Liu, Bing Hu, Yicheng Liu, Lunke Fei, and Wei Wang. Deep variational incomplete multi-view clustering: Exploring shared clustering structures. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 16147–16155, 2024. 2
- [46] Jie Xu, Chao Li, Yazhou Ren, Liang Peng, Yujie Mo, Xiaoshuang Shi, and Xiaofeng Zhu. Deep incomplete multi-view clustering via mining cluster complementarity. In *Proceedings of the AAAI conference on artificial intelligence*, pages 8761–8769, 2022. 6
- [47] Jie Xu, Yazhou Ren, Huayi Tang, Zhimeng Yang, Lili Pan, Yang Yang, Xiaorong Pu, S Yu Philip, and Lifang He. Self-supervised discriminative feature learning for deep multi-view clustering. *IEEE Transactions on Knowledge and Data Engineering*, 35(7):7470–7482, 2022. 1
- [48] Jie Xu, Chao Li, Liang Peng, Yazhou Ren, Xiaoshuang Shi, Heng Tao Shen, and Xiaofeng Zhu. Adaptive feature projection with distribution alignment for deep incomplete multi-view clustering. *IEEE Transactions on Image Processing*, 32:1354–1366, 2023. 2, 6
- [49] Zhe Xue, Yawen Li, Zhongchao Guan, Wenling Li, Meiyu Liang, and Hai Zhou. Robust multi-graph contrastive network for incomplete multi-view clustering. *IEEE Transactions on Multimedia*, 2024. 2
- [50] Mouxing Yang, Yunfan Li, Peng Hu, Jinfeng Bai, Jiancheng Lv, and Xi Peng. Robust multi-view clustering with incomplete information. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(1):1055–1069, 2022. 6
- [51] Xihong Yang, Siwei Wang, Fangdi Wang, Jiaqi Jin, Suyuan Liu, Yue Liu, En Zhu, Xinwang Liu, and Yueming Jin. Automatically identify and rectify: Robust deep contrastive multi-

- view clustering in noisy scenarios. In *Forty-second International Conference on Machine Learning*. 1
- [52] Xihong Yang, Jin Jiaqi, Siwei Wang, Ke Liang, Yue Liu, Yi Wen, Suyuan Liu, Sihang Zhou, Xinwang Liu, and En Zhu. Dealmvc: Dual contrastive calibration for multi-view clustering. In *Proceedings of the 31st ACM international conference on multimedia*, pages 337–346, 2023. 2
- [53] Shengju Yu, Zhibin Dong, Siwei Wang, Xinhang Wan, Yue Liu, Weixuan Liang, Pei Zhang, Wenxuan Tu, and Xinwang Liu. Towards resource-friendly, extensible and stable incomplete multi-view clustering. In *Forty-first International Conference on Machine Learning*, 2024. 1
- [54] Pengxin Zeng, Mouxing Yang, Yiding Lu, Changqing Zhang, Peng Hu, and Xi Peng. Semantic invariant multi-view clustering with fully incomplete information. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. 6
- [55] Changqing Zhang, Yajie Cui, Zongbo Han, Joey Tianyi Zhou, Huazhu Fu, and Qinghua Hu. Deep partial multi-view learning. *IEEE transactions on pattern analysis and machine intelligence*, 44(5):2402–2415, 2020. 2