Consistency Enhancement-Based Deep Multiview Clustering via Contrastive Learning

Hao Yang[1], **Hua Mao**[2], **Wai Lok Woo**[2], **Jie Chen**[1] and **Xi Peng**[1]
[1]College of Computer Science, Sichuan University
[2]Department of Computer and Information Sciences, Northumbria University
yhscu2018@gmail.com, {hua.mao, wailok.woo}@northumbria.ac.uk chenjie2010@scu.edu.cn, pengx.gm@gmail.com

arXiv:2401.12648v1 [cs.LG] 23 Jan 2024

Abstract

Multiview clustering (MVC) segregates data samples into meaningful clusters by synthesizing information across multiple views. Moreover, deep learning-based methods have demonstrated their strong feature learning capabilities in MVC scenarios. However, effectively generalizing feature representations while maintaining consistency is still an intractable problem. In addition, most existing deep clustering methods based on contrastive learning overlook the consistency of the clustering representations during the clustering process. In this paper, we show how the above problems can be overcome and propose a consistent enhancement-based deep MVC method via contrastive learning (CCEC). Specifically, semantic connection blocks are incorporated into a feature representation to preserve the consistent information among multiple views. Furthermore, the representation process for clustering is enhanced through spectral clustering, and the consistency across multiple views is improved. Experiments conducted on five datasets demonstrate the effectiveness and superiority of our method in comparison with the state-of-the-art (SOTA) methods. The code for this method can be accessed at https://anonymous.4open.science/r/CCEC-E84E/.

1 Introduction

With the diversity and increasing complexity of data sources, multiview clustering (MVC) has recently emerged as a key research area. Multiview data are characterized by distinct feature representations that capture different facets or modalities of a single entity [Chen *et al.*, 2022b]. The primary objective of MVC is to effectively segregate data samples into meaningful clusters. The potency of MVC is derived from its ability to leverage consistency information derived from various perspectives, which leads to enhanced clustering precision and resilience [Xu *et al.*, 2022]. MVC has attracted increasing attention for use in many machine learning tasks, including feature selection [Xu *et al.*, 2023], scene recognition [Tang *et al.*, 2023], and information retrieval [Han *et al.*, 2023; Chen *et al.*, 2022a].

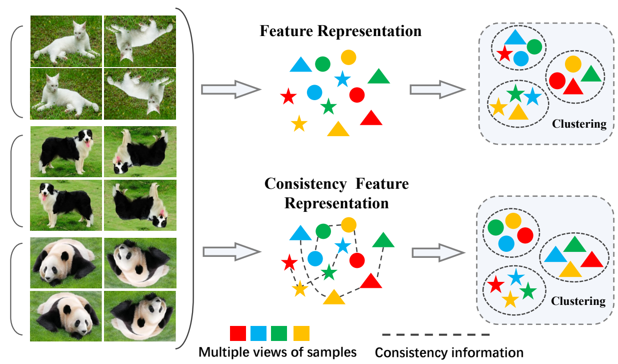

Figure 1: Feature representations lacking consistency reduce the accuracy of subsequent clustering representations.

In MVC, multiple feature representations are contained under multiple views for a sample. These feature representations are in different latent spaces, but they still have a certain level of consistency. As shown in Figure 1, we present the feature representations of three types of samples acquired from different views. If we directly cluster these feature representations, the clustering results will lose the consistency information contained in the original features, resulting in poor performance. However, if the consistency of the same sample across different views can be mined and this consistency is used to learn the sample representations for clustering, the ideal clustering effect will be attained.

In recent years, numerous MVC approaches have been proposed, including subspace-based methods [Tao *et al.*, 2023], matrix decomposition methods [Zhao *et al.*, 2021; Hu and Chen, 2019], graph-based methods [Li *et al.*, 2023], and multiple kernel-based methods [Liu *et al.*, 2021; Liu *et al.*, 2019]. However, these approaches possess poor representation capabilities and high computational complexity. Numerous deep learning-based methods have been proposed to alleviate the above problems. MFLVC [Xu *et al.*, 2022] learns different levels of features via a contrastive strategy. GCFagg [Yan *et al.*, 2023] integrates global and cross-view feature aggregation with structure-guided contrastive learning. DealMVC [Yang *et al.*, 2023] integrates a dual-contrastive calibration mechanism to align features acquired from global and local views. DMCE [Zhao *et al.*, 2023] integrates ensemble clustering to fuse similarity graphs derived from different views with a graph autoencoder. CVCL [Chen *et al.*, 2023b] aligns cluster centers across views to foster view-invariant representations.

These methods have shown that maintaining consistency among views is crucial in multiview learning scenarios. However, the existing methods [Yan *et al.*, 2023; Yang *et al.*, 2023; Zhao *et al.*, 2023] exhibit deficiencies in terms of extracting feature representations imbued with consistency information. Inconsistent feature representations reduce the accuracy of the subsequent clustering representations. Moreover, the prevailing contrastive learning approaches [Chen *et al.*, 2023b; Tang and Liu, 2022a; Lin *et al.*, 2022] for MVC overlook the consistency of the clustering representations during the clustering process.

To address these challenges, we propose a consistent enhancement-based deep MVC method via contrastive learning (CCEC), which aims at resolving the issues of consistency within MVC. We design a consistency preservation module to obtain the original multiview-consistent data representation and mine the consistency information of multiple views by introducing semantic connection blocks. In contrast with most existing contrastive learning methods, the proposed approach introduces spectral clustering to capture consistent semantic label information from multiple views. CCEC aligns the clustering representations between multiple views by cross-comparing the clustering labels generated by spectral clustering and a multilayer perceptron (MLP). Based on these view-invariant representations, the contrastive loss of the proposed CCEC method encourages the cluster assignments produced for positive pairs to be similar and pushes the cluster assignments provided for negative pairs apart.

Our major contributions are summarized as follows.

- We introduce an end-to-end deep MVC method termed CCEC, which enhances the consistency of MVC.
- CCEC provides a consistency-preserving discrepancy autoencoder structure that preserves the original data features of multiple views by introducing semantic connection blocks.
- CCEC designs a new cross-view contrastive clustering module to reinforce the interview consistency information within the framework.
- The proposed CCEC algorithm achieves a significant improvement in performance over that of state-of-the-art (SOTA) MVC methods on five datasets.

2 Related Work

In this section, we briefly review three topics related to this work, i.e., MVC, contrastive learning, and spectral clustering.

2.1 Multiview Clustering (MVC)

The existing MVC methods can be divided into five categories: subspace learning-based approaches [Chen *et al.*, 2023c], nonnegative matrix factorization-based approaches [Chen *et al.*, 2023d], graph-based approaches [Pan and Kang, 2021; Lin and Kang, 2021], multiple kernel learning-based approaches [Wang *et al.*, 2021], and deep learning-based approaches [Tang and Liu, 2022a; Lin *et al.*, 2022].

Although the traditional MVC methods are effective, they often capture shallow data representations, which limits their ability to discriminate among the derived data representations. To address this problem, recent developments have shifted toward deep MVC methods. These methods utilize deep neural networks to extract more detailed and hierarchical feature representations, effectively revealing the potential clustering patterns contained in multiview data. DSIMVC [Tang and Liu, 2022a] dynamically imputes missing views and selectively uses the imputed samples for training to ensure semantic consistency. DSMVC [Tang and Liu, 2022b] identifies and focuses on the most relevant features derived from each view to effectively balance the extraction of useful information from an increased number of views. DCP [Lin *et al.*, 2022] uses within-view reconstruction, dual cross-view contrastive learning, and cross-view dual prediction to address the challenges of consistency learning in multiview settings. Unlike these existing methods, the goal of this work is to present a new framework that can reduce the loss of consistency information incurred during clustering and ensure improved clustering performance.

2.2 Contrastive Learning

Contrastive learning, as a crucial paradigm in unsupervised learning, has substantially advanced the field of representation learning [Hadsell *et al.*, 2006; Tian *et al.*, 2020; Tsai *et al.*, 2020; Trosten *et al.*, 2021]. This approach is fundamentally based on the creation of a latent space in which the similarity between positive pairs is maximized, while that between negative pairs is minimized [Lin *et al.*, 2022; Chen *et al.*, 2023d]. A key element of this approach is the InfoNCE loss, a variant of noise-contrastive estimation (NCE), which serves as a lower mutual information bound [Misra and Maaten, 2020]. This concept has been effectively integrated into models such as MoCo [He *et al.*, 2020] and CPC [Oord *et al.*, 2018], focusing on maximizing the mutual information between different views of a sample.

Within the sphere of MVC, contrastive learning adeptly addresses the challenge of ensuring representational coherence across heterogeneous views [Chen *et al.*, 2023b]. While conventional methods predominantly rely on data augmentation to generate these variegated views, our approach diverges by adopting a novel pseudolabeling strategy.

2.3 Spectral Clustering

Spectral clustering is rooted in graph theory and uses graph representations of data to construct a clustering structure. This approach reveals the inherent clustering representation by refining the feature vectors of the input graph [Huang *et al.*, 2019; Chen *et al.*, 2023c]. Different views in MVC may produce unique graphical layouts, leading to consistency issues. Spectral clustering can unify these different graphs, providing a method for solving this consistency problem and improving the obtained clustering results. CSRF [Chen *et al.*, 2023a] improves upon the traditional spectral clustering approach by learning a fused affinity matrix at the spectral embedding feature level, expanding the applicability of spectral clustering in the MVC domain.

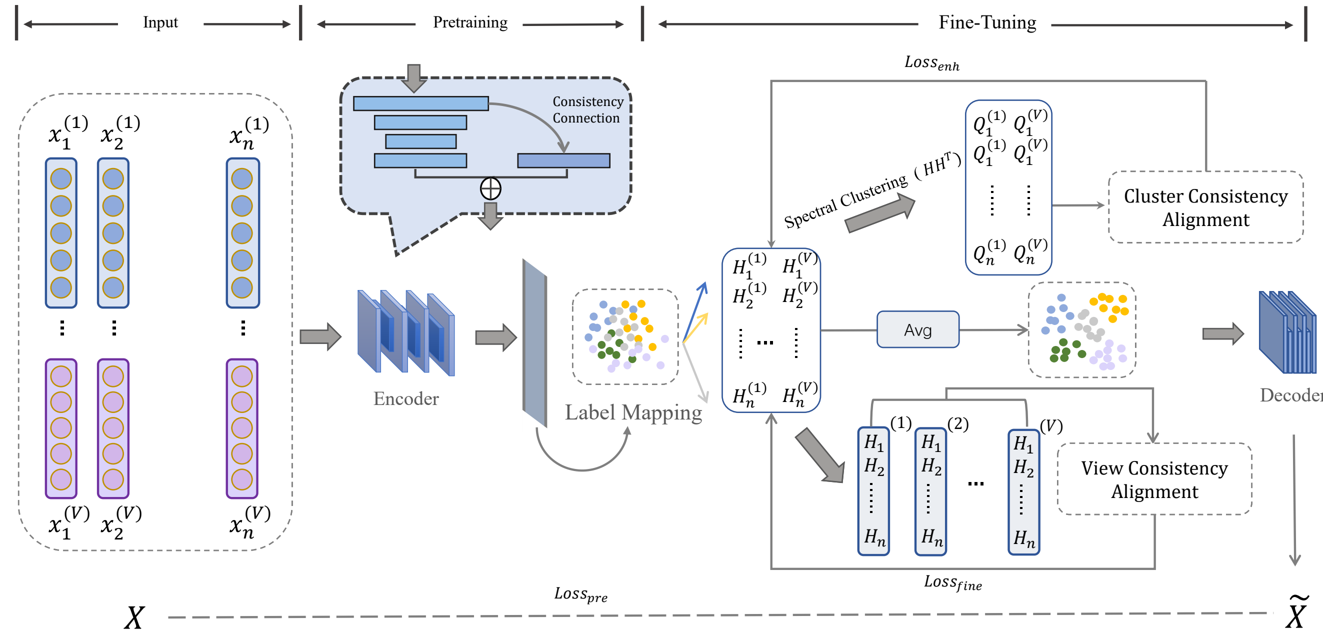

Figure 2: The framework of CCEC. Our module includes consistency-preserving autoencoder and cross-contrastive consistency learning modules. The former learns a consensus representation via semantic connection blocks, which fully explore the consistent information among multiple views. The latter integrates the pseudolabels generated by neural networks and spectral clustering into the contrastive learning process to capture potential consistency information.

3 METHODOLOGY

3.1 Motivation

In multiview learning, achieving feature consistency across disparate views is crucial. A traditional autoencoder predominantly focuses on representing a single view and might not effectively encapsulate the consistent features between views. Inspired by the architecture of ResNet [He *et al.*, 2016], we incorporate semantic connection blocks into the multiview feature extraction process to achieve feature consistency. The consistency information can be represented in a two-layer MLP-based feature extraction module as follows:

$$C = \sum_{i=1}^{n} f(f(x_i, (w_{1i}, w_{1i}^{'})), w_{2i}) - \sum_{i=1}^{n} f(f(x_i, w_{1i}), w_{2i}) \tag{1}$$

where: C is the consistency information, x_i is the input for the i-th view, f represents the transformation function, w_{1i}, w_{2i} and $w_{1i}^{'}$ are weight matrices.

We design a new contrastive learning method to enhance the consistency among multiple views. The cluster pseudolabels obtained through the spectral clustering method and the cluster pseudolabels generated by a neural network are used as positive pairs for contrastive learning. The mechanism for ensuring semantic alignment among the features in different views by minimizing the clustering distribution differences between different views of the same sample emphasizes the essence of our solution.

In pursuit of directly extracting semantic labels for end-to-end clustering from raw instances across multiple views, we present the CCEC framework. As depicted in Figure 2, the CCEC architecture is bifurcated into two primary modules: consistency-preserving autoencoder and cross-contrastive consistency learning modules. The task of the autoencoder, which preserves consistency, is to extract features with consistent information. The cross-contrastive consistency learning module refines the clustering outcome by juxtaposing cluster assignments, ensuring coherence and robustness.

3.2 Consistency Retention

A standard residual connection [He *et al.*, 2016] can be expressed as follows:

$$y = F(\mathbf{x}, \{W_1\}) + W_2\mathbf{x} \tag{2}$$

where: $\mathbf{x}$ is the input, y is the output, F represents the transformation function, W_1 and W_2 are weight matrices. This equation signifies that the transformed input $F(\mathbf{x}, W_1)$ and the original input $\mathbf{x}$ both contribute to the output. In a multiview scenario, ensuring that the main attributes of $\mathbf{x}$ remain in the output can help with retaining consistent features across views.

By using semantic connection blocks, we can mathematically represent the focus of the model on consistent features as follows:

$$C = \alpha \cdot F(\mathbf{x}, W_1) + (1 - \alpha) \cdot \mathbf{x}, \tag{3}$$

where C is the consistent feature vector and α is a weighting factor that determines the balance between the transformed features and the original input.

Leveraging the concept of residual connections, we design a novel encoder architecture. This design consists of multiple stacked semantic connection blocks, which are represented as follows:

$$E_i(\mathbf{x}) = F(E_{i-1}(\mathbf{x}), W_i^1) + W_i^2 E_{i-1}(\mathbf{x}), \quad (4)$$

where E_i is the output derived from the i-th semantic connection block and E_{i-1} is the output obtained from the previous semantic connection blocks (or the input for i=1). Each block retains the raw input data, ensuring that even as the data undergo transformations, the consistent features are accentuated and preserved.

3.3 Consistency Enhancement

Given a set of feature representations $Z = \{\mathbf{z}_1^m, \mathbf{z}_2^m, ..., \mathbf{z}_i^m\}$ (where $\mathbf{z}_i^m = E^m(\mathbf{x}_i^m)$) from multiple views, where m represents the m-th view, we stack an MLP and the softmax function on Z to obtain a clustering distribution matrix H. A similarity matrix S is then derived by taking the product of H with its transpose: $S = HH^T$. The resulting matrix S serves as a pivotal representation, capturing the intrinsic similarities among different cluster distributions. Then, we perform spectral clustering on S to obtain a clustering distribution matrix Q such that $Q = f(S)$, where f denotes the spectral clustering function.

To reinforce the consistency across various views, we harness the power of contrastive learning for the clustering distributions. For two given distributions C_1 and C_2, which are derived from different views, our objective function is defined as follows:

$$\mathcal{L} = \sum_{i=1}^{n} \sum_{j=1}^{n} d(C_1(i,j), C_2(i,j)), \quad (5)$$

where d represents a distance metric, which ensures that the differences between various clustering distributions are minimized. This mechanism ensures the semantic alignment of the features across distinct views, underpinning the essence of our approach.

3.4 Pretraining Phase with Consistency Feature Extraction

We first construct a pretraining network for optimizing the parameter initialization process. This network combines paired encoder-decoder modules, each of which is fine-tuned for distinct views from a given set. For every view v within V, a specific data sample $\mathbf{x}_i^v$ undergoes a transformation through the encoder to manifest as an embedded feature representation denoted by $\mathbf{z}_i^v$. Mathematically, this transformation is defined as follows:

$$\mathbf{z}_i^v = f_e^v(\mathbf{x}_i^v; W_e^v), \quad (6)$$

where f_e^v and W_e^v symbolize the encoder function and the associated weight parameters for view v, respectively.

This embedding representation $\mathbf{z}_i^v$ aims to reconstruct the original data samples through the decoder module. The outcome of this decoding process is represented as $\tilde{\mathbf{x}}_i^v$, which is described by the following equation:

$$\tilde{\mathbf{x}}_i^v = f_d^v(\mathbf{z}_i^v; W_d^v), \quad (7)$$

where f_d^v stands for the decoder function and W_d^v pertains to the decoder weights, both of which are tailored for view v.

As we continue through the pretraining phase, our overarching goal centers around minimizing the reconstruction loss spanning all views, which is captured succinctly by the loss function shown below:

$$\mathcal{L}_{pre} = \sum_{v=1}^{M} \sum_{i=1}^{N} ||\mathbf{x}_i^v - \tilde{\mathbf{x}}_i^v||^2. \quad (8)$$

3.5 Fine-Tuning the Contrastive Learning Process for Cross-View Consistency

Given a set of features $\{\mathbf{Z}_\mathrm{m}\}_{\mathrm{m}=1}^M$ obtained by Eq. (6), we notice that these features represent weakly consistent multiview feature representations. We denote them as weak consistency representations. To enhance these representations, our goal is to derive strong consistency clustering labels, which we term strong consistency representations. To achieve this goal, a three-layer linear MLP, represented as $F(\{\mathbf{Z}_m\}_{m=1}^M; W_H)$,is applied over $\{\mathbf{Z}_m\}_{m=1}^M$. This results in a set of cluster representations $\{\mathbf{H}_m\}_{m=1}^M$. Spectral clustering is then employed to reconstruct these labels, resulting in additional cluster representations denoted as $\{\mathbf{Q}_m\}_{m=1}^M$.

In the weak consistency space, the reconstruction objective given by Eq. (6) ensures that the representational capacity of $\{\mathbf{Z}_m\}_{m=1}^M$ is retained, thus mitigating the model collapse problem. In the strong consistency space, contrastive learning is employed to ensure that $\mathbf{H}_m$ and $\mathbf{Q}_m$ converge toward learning universal semantics across all views.

Both types of cluster representations $\mathbf{h}_m$ and $\mathbf{q}_m$ consist of $(2MN - 1)$ label pairs. Among these, $(M - 1)$ are positive feature pairs, while the remaining $(M(2N - 1))$ are negative feature pairs. Following the SwAV [Caron *et al.*, 2020] contrastive learning approach, we prioritize maximizing the similarity between the positive pairs and disregard the negative pairs. Taking inspiration from NT-Xent [Chen *et al.*, 2020], we use the cosine distance to measure the similarity between two features:

$$d(\mathbf{h}_i^m, \mathbf{q}_j^m) = \frac{\langle \mathbf{h}_i^m, \mathbf{q}_j^n \rangle}{\| \mathbf{h}_i^m \| \left\| \mathbf{q}_j^n \right\|}, \quad (9)$$

where $\langle \cdot, \cdot \rangle$ signifies the dot product operation.

Subsequently, the feature-contrastive loss between $\mathbf{H}_m$ and $\mathbf{Q}_n$ is expressed as:

$$\ell_{ec}^{(mn)} = -\frac{1}{N} \sum_{i=1}^{N} \frac{d(\mathbf{h}_i^m, \mathbf{q}_i^n)}{\tau_L}, \quad (10)$$

where τ_L represents the temperature parameter.

We further define an accumulated multiview feature-contrastive loss spanning all views as follows:

$$\mathcal{L}_{enh} = \frac{1}{2M} \sum_{m=1}^{M} \sum_{n \neq m} \ell_{ec}^{(mn)}. \quad (11)$$

However, in practical scenarios, some views may still have inconsistent clustering labels due to the influence of view-specific information. To ensure robustness, we aim to achieve clustering consistency, where identical clustering labels across all views should represent the same semantic clusters. In other words, $\{\mathbf{H}_j^m\}_{m=1}^M (\mathbf{H}_j^m \in \mathbb{R}^N)$ should remain consistent.

To achieve this consistency objective, we employ contrastive learning. For the m-th view, similar cluster labels $\mathbf{H}_j^m$ form $(MK-1)$ label pairs, i.e., $\{\mathbf{H}_j^m, \mathbf{H}_k^n\}_{k=1,\dots,K}^{n=1,\dots M}$, where the $\{\mathbf{H}_j^m, \mathbf{H}_j^n\}_{n\neq m}$ are constructed as $(M-1)$ positive label pairs and the remaining $M(K-1)$ label pairs are considered negative label pairs. We further define the label-contrastive loss between $\mathbf{H}^m$ and $\mathbf{H}^n$ as follows:

$$\ell_{fc}^{(mn)} = -\frac{1}{K}\sum_{j=1}^{K}\log\frac{e^{d(H_j^m,H_j^n)/\tau_S}}{\sum_{k=1}^{K}\sum_{v=m}^{n}e^{d(H_j^m,H_k^v)/\tau_S}-e^{1/\tau_L}}, \tag{12}$$

where τ_S represents the temperature parameter.

Thus, the clustering consistency objective is defined as follows:

$$\mathcal{L}_{fine} = \frac{1}{2}\sum_{m=1}^{M}\sum_{n\neq m}\ell_{fc}^{(mn)} + \sum_{m=1}^{M}\sum_{j=1}^{K}s_j^m\log s_j^m, \tag{13}$$

where $s_j^m = \frac{1}{N}\sum_{i=1}^{N}h_{ij}^m$ [Van Gansbeke *et al.*, 2020]. The first part of Eq. (13) aims to learn the clustering consistency across all views, while the second part serves as a regularization term, which is typically used to prevent all samples from being assigned to a single cluster.

The overall loss of the proposed method consists of three main components: the reconstruction loss of the pretrained network, the cross-contrastive consistency loss, and the consistency comparison fine-tuning loss:

$$\mathcal{L} = \mathcal{L}_{pre} + \mathcal{L}_{enh} + \mathcal{L}_{fine}. \tag{14}$$

4 Experiment

In this section, we conduct experiments on real-world datasets and report the results obtained by our model as well as some SOTA baselines to demonstrate the effectiveness of our model. This study aims to evaluate the performance and capabilities of our model with a focus on four key areas.

- Performance: The performance of our model is compared with that of the SOTA methods.
- Feature Extraction: The ability of our cross-contextual embedding consistency approach to extract consistent semantic features is assessed.
- Consistency Enhancement: We evaluating the effectiveness of the consistency enhancement module in our model.
- Hyperparameter Sensitivity: The effectiveness of our model under various hyperparameter settings is tested.

4.1 Experimental Setup

Datasets

To ascertain the efficacy of our proposed CCEC method, we carry out comprehensive experiments across five datasets: the MNIST-USPS [Peng *et al.*, 2019], MSRC-v1 [Winn and Jojic, 2005], Handwritten Digit [Asuncion and Newman, 2007], Outdoor Scene (O-Scene) [Oliva and Torralba, 2001], and Caltech Multiview Image datasets [Fei-Fei *et al.*, 2004]. The specifics and statistical details of each dataset are comprehensively summarized in Table 1.

Table 1: Descriptions of the employed multiview datasets.

Datasets	Samples	Views	Clusters
MNIST-USPS	5000	2	10
Handwritten	2000	6	10
MSRC-v1	210	5	7
Scene	2688	4	8
Caltech-2V	1400	2	7
Caltech-3V	1400	3	7
Caltech-4V	1400	4	7
Caltech-5V	1400	5	7

Baselines

We select representative methods to conduct an overall comparison, namely, DSIMVC [Tang and Liu, 2022a], DCP [Lin *et al.*, 2022], DSMVC [Tang and Liu, 2022b], MFLVC [Xu *et al.*, 2022], CVCL [Chen *et al.*, 2023b], DealMVC [Yang *et al.*, 2023], GCFaggMVC [Yan *et al.*, 2023] and DMCE [Zhao *et al.*, 2023]. For DCP, the best clustering result is reported from the combinations of different pair of individual views in each dataset.

Evaluation Metrics

The clustering effectiveness of each model is evaluated by four metrics, i.e., the clustering accuracy (ACC), normalized mutual information (NMI), automated readability index (ARI) and purity (PUR). The ACC measures the proportion of correctly labeled data points, while the NMI assesses the mutual agreement between the clustering results and the true labels (adjusted for the cluster size). The ARI offers a stringent clustering evaluation measure by considering the correct pairings of data points, and PUR measures the dominance of a single class within each cluster. For these metrics, larger values indicate better clustering performance.

4.2 Overall Performance Evaluation

The clustering outcomes generated by all the competing methodologies across the four multiview datasets are documented in Table 2, and the clustering results obtained on the Caltech datasets are presented in Table 3. The best and second-best values among the clustering results are emphasized in bold and underlined text, respectively. The methods that prioritize interview consistency, including CCEC, CVCL, and DSMVC, typically achieve substantial enhancements over the competing approaches on large-scale datasets, such as MNIST-USPS, Handwritten, and Scene. Furthermore, the CCEC method significantly outperforms the other contrastive learning-based approaches, including CVCL, DSMVC, DCP, DSIMVC, and MFLVC, across all the evaluated datasets. This result corroborates the importance of the consistency amplification strategy employed by CCEC.

Table 2: Results obtained on the MNIST-USPS, Scene, Handwritten, and MSRC-v1 datasets.

	MNIST-USPS				MSRC-v1				Scene				Handwritten			
Method	ACC	NMI	ARI	PUR	ACC	NMI	ARI	PUR	ACC	NMI	ARI	PUR	ACC	NMI	ARI	PUR
DealMVC	98.24	98.55	98.16	98.24	82.00	74.54	68.68	82.00	69.57	59.44	48.78	69.57	89.20	87.54	78.65	89.20
DSIMVC	99.17	98.13	98.42	99.17	79.05	69.00	66.17	79.05	70.15	62.51	60.06	70.15	87.20	80.39	76.51	87.20
DCP	99.02	97.29	99.25	99.02	78.57	74.84	80.04	79.43	76.15	63.19	52.49	76.15	85.75	85.05	88.64	85.75
GCFAgg	99.56	98.71	99.00	99.56	92.54	91.66	90.98	92.54	67.07	60.97	48.68	67.36	93.19	89.05	88.06	93.19
DMCE	99.04	89.94	86.71	99.04	90.47	81.99	79.19	90.47	73.33	65.13	**61.46**	73.33	90.00	81.63	78.60	90.00
DSMVC	99.34	99.07	99.00	99.34	90.42	86.72	85.63	90.42	75.45	61.26	57.19	75.45	96.85	94.07	94.29	96.80
MFLVC	99.66	99.01	99.25	99.66	94.29	89.17	87.43	94.29	64.57	54.58	43.16	64.57	83.55	82.62	74.42	84.65
CVCL	99.70	99.13	99.07	99.70	97.62	94.98	92.57	97.62	77.83	63.35	57.19	77.83	97.35	94.05	94.55	97.35
CCEC	**99.86**	**99.59**	**99.69**	**99.86**	**98.57**	**96.77**	**96.63**	**98.57**	**79.76**	**65.79**	59.99	**79.76**	**97.95**	**95.20**	**95.49**	**97.95**

Table 3: Results obtained on Caltech.

	Caltech-2V				Caltech-3V				Caltech-4V				Caltech-5V			
Method	ACC	NMI	ARI	PUR	ACC	NMI	ARI	PUR	ACC	NMI	ARI	PUR	ACC	NMI	ARI	PUR
DealMVC	60.00	52.96	40.42	61.14	68.57	60.28	51.00	68.57	76.64	76.28	67.82	76.64	88.71	80.95	82.34	88.71
DSIMVC	59.30	54.00	**59.77**	59.30	65.39	61.73	55.40	65.39	73.50	70.20	69.18	73.50	75.76	76.21	72.55	75.76
DCP	59.36	56.92	43.94	59.36	67.71	70.57	59.11	67.25	75.73	70.61	74.88	75.73	85.86	84.75	79.57	85.86
GCFAgg	66.43	50.08	55.60	66.43	64.00	53.45	46.11	65.29	73.43	66.10	60.50	73.43	83.36	73.31	69.75	83.36
DSMVC	61.00	50.01	42.36	61.00	74.93	64.83	59.05	74.93	83.79	81.81	77.25	84.43	92.57	83.37	84.16	92.57
DMCE	60.57	**64.60**	52.89	64.92	69.64	63.99	54.65	73.14	74.21	69.10	61.71	77.57	85.14	77.52	70.08	87.71
MFLVC	62.00	54.30	43.50	62.00	66.79	58.42	49.64	68.36	78.43	71.06	64.42	78.43	86.21	77.01	72.65	86.21
CVCL	63.79	51.78	44.06	64.57	77.93	70.92	65.19	78.14	82.50	71.85	66.10	82.50	90.67	86.82	76.04	90.67
CCEC	**76.00**	64.15	59.76	**76.00**	**79.29**	**72.52**	**65.76**	**79.29**	**90.21**	**83.74**	**78.71**	**90.21**	**94.36**	**89.90**	**88.29**	**94.36**

These findings validate the efficacy of our proposed CCEC method. The CCEC approach consistently attains the highest clustering results on all datasets in terms of the ACC metric. Notably, the CCEC method manifests performance gains of approximately 1.9%, 3.5%, 4.6%, and 1.9% over the second-best performing method on the Caltech-5V dataset with respect to the ACC, NMI, ARI, and purity metrics, respectively. Similarly, the performance of the CCEC method is markedly superior to that of the other competing methods on additional datasets, underscoring the preeminence of CCEC over the alternative techniques.

CCEC demonstrates notably better performance on datasets with greater numbers of views than on those with fewer views. The performance gaps between CCEC and the other methods are more pronounced on datasets such as Handwritten and Caltech-5V than for Scene and Caltech-2V. This is attributable to the capacity of CCEC to extract and amplify more consistent information in multiview scenarios, thereby enhancing its clustering performance.

4.3 Semantic Extraction Effectiveness

As shown in Figure 3, as the number of views increases, both the traditional feature extraction methods and the consistency information preservation methods can extract more consistency information to improve the accuracy of the subsequent clustering process. However, the traditional methods cannot fully preserve the complete consistency information in different views, which results in lower clustering accuracy than consistency clustering accuracy preservation methods. The consistency-preserving feature extraction process yields

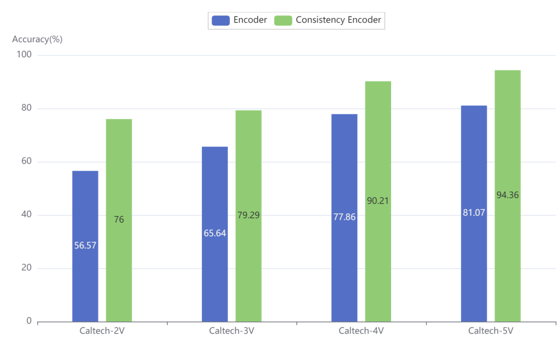

Figure 3: The clustering accuracies achieved by traditional feature extraction methods and consistency information preservation methods under different numbers of views.

a large amount of consistency information in the fourth view, significantly improving the resulting clustering accuracy and surpassing the accuracy attained by the traditional clustering methods with five views. This indicates that the consistency information preservation method is crucial for improving the clustering results.

4.4 The Role of Representation Enhancement

According to the overall reconstruction loss, three different loss components are included. To verify the importance of each component in CCEC, we perform ablation studies under the same experimental settings to isolate the necessity of each component. We conduct two experiments, one using traditional feature extraction methods and another us-

Table 4: Results of an ablation study concerning the main components of the proposed CCEC method conducted on all the datasets.

	LOSS			Caltech-2V				Caltech-3V				Caltech-4V				Caltech-5V			
Method	Lp	Le	Lf	ACC	NMI	ARI	PUR	ACC	NMI	ARI	PUR	ACC	NMI	ARI	PUR	ACC	NMI	ARI	PUR
CCEC(tra)		✓	✓	56.57	45.79	36.80	57.64	65.64	57.69	47.10	66.71	77.86	75.04	65.41	77.86	81.07	81.90	77.48	81.07
CCEC(con)	✓		✓	71.14	55.51	51.45	71.14	75.64	77.33	61.94	75.64	81.14	83.05	70.63	84.21	86.93	87.42	80.56	89.86
CCEC	✓	✓	✓	**76.00**	**64.15**	**59.76**	**76.00**	**79.29**	**72.52**	**65.76**	**79.29**	**90.21**	**83.74**	**78.71**	**90.21**	**94.36**	**89.90**	**88.29**	**94.36**

Table 5: Results of an ablation study concerning the main components of the proposed CCEC method conducted on all the datasets.

	LOSS			MNIST-USPS				MSRC-v1				Scene				Handwritten			
Method	Lp	Le	Lf	ACC	NMI	ARI	PUR	ACC	NMI	ARI	PUR	ACC	NMI	ARI	PUR	ACC	NMI	ARI	PUR
CCEC(tra)		✓	✓	99.22	97.85	98.28	99.22	88.57	89.41	89.79	84.76	70.80	60.06	49.75	70.80	84.36	82.35	80.36	84.36
CCEC(con)	✓		✓	99.60	99.03	99.14	99.60	89.05	85.26	85.40	90.00	73.58	71.63	51.68	73.58	90.55	88.63	88.52	90.55
CCEC	✓	✓	✓	**99.86**	**99.59**	**99.69**	**99.86**	**98.57**	**96.77**	**96.63**	**98.57**	**79.76**	**65.79**	**59.99**	**79.76**	**97.95**	**95.20**	**95.49**	**97.95**

ing consistency-preserving feature extraction modules but not cross view consistency enhancement modules. Table 4 and table 5 show the obtained clustering results in terms of the three metrics produced with different combinations of the loss components. The clustering results in the first two rows of Table 4 and table 5 are produced by the two special cases.

As expected, the best performance can be achieved when all loss terms are considered. Moreover, the clustering performance is significantly improved when the pretraining stage is employed in CCEC. For example, CCEC performs much better than CCEC(tra), with improvements of approximately 13.29%, 8%, 10.81% and 13.29% in terms of the ACC, NMI, ARI and purity metrics, respectively, achieved on the Caltech-5V dataset. In addition, we find that all the results produced by CCEC(con) are greater than those of CCEC(tra), indicating that the consistency enhancement module yields a lower overall performance improvement than the consistency preservation module. However, at the same time, we also find that the consistency preservation module does not significantly improve the process of extracting consistency information from each view. The consistency preservation module has a relatively flat effect on improving the results, while the consistency enhancement module can fully utilize consistency information. An increase by a single view substantially improves the resulting accuracy. Therefore, each component in the overall reconstruction loss plays a crucial role in learning view-invariant representations.

4.5 Hyperparameter Sensitivity Analysis

We conduct experiments on four representative datasets, i.e., the Caltech-5V, MNIST-USPS, Handwritten, and MSRC-v1 datasets, to investigate the sensitivity of the $\tau_L(\alpha)$ and $\tau_S(\beta)$ parameters in the proposed CCEC method. Figure 4 shows the clustering performance achieved by the CCEC method in terms of the ACC values obtained with different combinations of τ_L and τ_S. The clustering performance attained by the CCEC method on the Caltech dataset and MNIST-USPS dataset does not seriously fluctuate under different combinations of τ_L and τ_S. This finding indicates the robustness of the CCEC method. In addition, we find that on the Handwritten and MSRC-v1 datasets, the clustering results impact the

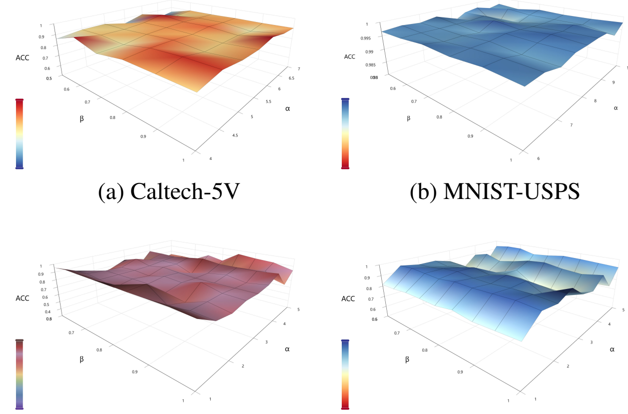

(a) Caltech-5V (b) MNIST-USPS

(c) Handwritten (d) MSRC-v1

Figure 4: The ACC values yielded by the CCEC method with different combinations of α and β on the four representative datasets.

$\tau_L(\alpha)$ parameter, making the model more sensitive to this parameter. The clustering results of CCEC are insensitive to the hyperparameters and exhibit stability.

5 Conclusion

In this paper, we have proposed CCEC to achieve enhanced consistency in MVC tasks. Specifically, we incorporate semantic connection blocks into the feature representation process to preserve the consistent information across multiple views. We have also implemented spectral clustering and contrastive learning on the obtained consistent feature representations to obtain a strong consistency clustering representation. The strong consistency clustering representation of CCEC contains rich consistency information, improves the accuracy of MVC, and demonstrates the importance of pseudolabels for improving the consistency of multiview learning. We have conducted extensive experiments and ablation studies on MVC datasets to validate the superiority of the proposed model and the effectiveness of its component in terms of the overall reconstruction loss metric.

References

[Asuncion and Newman, 2007] Arthur Asuncion and David Newman. Uci machine learning repository, 2007.

[Caron *et al.*, 2020] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. *Advances in neural information processing systems*, 33:9912–9924, 2020.

[Chen *et al.*, 2020] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020.

[Chen *et al.*, 2022a] Jie Chen, Shengxiang Yang, Hua Mao, and Conor Fahy. Multiview subspace clustering using low-rank representation. *IEEE Transactions on Cybernetics*, 52(11):12364–12378, 2022.

[Chen *et al.*, 2022b] Jie Chen, Shengxiang Yang, Xi Peng, Dezhong Peng, and Zhu Wang. Augmented sparse representation for incomplete multiview clustering. *IEEE Transactions on Neural Networks and Learning Systems*, pages 1–14, 2022.

[Chen *et al.*, 2023a] Jie Chen, Hua Mao, Dezhong Peng, Changqing Zhang, and Xi Peng. Multiview clustering by consensus spectral rotation fusion. *IEEE Transactions on Image Processing*, 32:5153–5166, 2023.

[Chen *et al.*, 2023b] Jie Chen, Hua Mao, Wai Lok Woo, and Xi Peng. Deep multiview clustering by contrasting cluster assignments. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 16752–16761, October 2023.

[Chen *et al.*, 2023c] Jie Chen, Zhu Wang, Hua Mao, and Xi Peng. Low-rank tensor learning for incomplete multiview clustering. *IEEE Transactions on Knowledge and Data Engineering*, 35(11):11556–11569, 2023.

[Chen *et al.*, 2023d] Jie Chen, Shengxiang Yang, Conor Fahy, Zhu Wang, Yinan Guo, and Yingke Chen. Online sparse representation clustering for evolving data streams. *IEEE Transactions on Neural Networks and Learning Systems*, pages 1–15, 2023.

[Fei-Fei *et al.*, 2004] Li Fei-Fei, Rob Fergus, and Pietro Perona. Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. In *2004 conference on computer vision and pattern recognition workshop*, pages 178–178. IEEE, 2004.

[Hadsell *et al.*, 2006] Raia Hadsell, Sumit Chopra, and Yann LeCun. Dimensionality reduction by learning an invariant mapping. In *2006 IEEE computer society conference on computer vision and pattern recognition (CVPR'06)*, volume 2, pages 1735–1742. IEEE, 2006.

[Han *et al.*, 2023] Zongbo Han, Changqing Zhang, Huazhu Fu, and Joey Tianyi Zhou. Trusted multi-view classification with dynamic evidential fusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(2):2551–2566, 2023.

[He *et al.*, 2016] 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.

[He *et al.*, 2020] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738, 2020.

[Hu and Chen, 2019] Menglei Hu and Songcan Chen. One-pass incomplete multi-view clustering. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01):3838–3845, Jul. 2019.

[Huang *et al.*, 2019] Zhenyu Huang, Joey Tianyi Zhou, Xi Peng, Changqing Zhang, Hongyuan Zhu, and Jiancheng Lv. Multi-view spectral clustering network. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 2563–2569. International Joint Conferences on Artificial Intelligence Organization, 7 2019.

[Li *et al.*, 2023] 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, 2023.

[Lin and Kang, 2021] Zhiping Lin and Zhao Kang. Graph filter-based multi-view attributed graph clustering. In *IJCAI*, pages 2723–2729, 2021.

[Lin *et al.*, 2022] 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.

[Liu *et al.*, 2019] Xinwang Liu, Xinzhong 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, 2019.

[Liu *et al.*, 2021] Xinwang Liu, Miaomiao Li, Chang Tang, Jingyuan Xia, Jian Xiong, Li Liu, Marius Kloft, and En Zhu. Efficient and effective regularized incomplete multi-view clustering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(8):2634–2646, 2021.

[Misra and Maaten, 2020] Ishan Misra and Laurens van der Maaten. Self-supervised learning of pretext-invariant representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.

[Oliva and Torralba, 2001] Aude Oliva and Antonio Torralba. Modeling the shape of the scene: A holistic representation of the spatial envelope. *International journal of computer vision*, 42:145–175, 2001.

[Oord *et al.*, 2018] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.

[Pan and Kang, 2021] Erlin Pan and Zhao Kang. Multi-view contrastive graph clustering. *Advances in neural information processing systems*, 34:2148–2159, 2021.

[Peng *et al.*, 2019] Xi Peng, Zhenyu Huang, Jiancheng Lv, Hongyuan Zhu, and Joey Tianyi Zhou. Comic: Multi-view clustering without parameter selection. In *International conference on machine learning*, pages 5092–5101. PMLR, 2019.

[Tang and Liu, 2022a] Huayi Tang and Yong Liu. Deep safe incomplete multi-view clustering: Theorem and algorithm. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 21090–21110. PMLR, 17–23 Jul 2022.

[Tang and Liu, 2022b] Huayi Tang and Yong Liu. Deep safe multi-view clustering: Reducing the risk of clustering performance degradation caused by view increase. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 202–211, June 2022.

[Tang *et al.*, 2023] Chang Tang, Zhenglai Li, Jun Wang, Xinwang Liu, Wei Zhang, and En Zhu. Unified one-step multi-view spectral clustering. *IEEE Transactions on Knowledge and Data Engineering*, 35(6):6449–6460, 2023.

[Tao *et al.*, 2023] Zhiqiang Tao, Jun Li, Huazhu Fu, Yu Kong, and Yun Fu. From ensemble clustering to subspace clustering: Cluster structure encoding. *IEEE Transactions on Neural Networks and Learning Systems*, 34(5):2670–2681, 2023.

[Tian *et al.*, 2020] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XI 16*, pages 776–794. Springer, 2020.

[Trosten *et al.*, 2021] Daniel J Trosten, Sigurd Lokse, Robert Jenssen, and Michael Kampffmeyer. Reconsidering representation alignment for multi-view clustering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1255–1265, 2021.

[Tsai *et al.*, 2020] Yao-Hung Hubert Tsai, Yue Wu, Ruslan Salakhutdinov, and Louis-Philippe Morency. Self-supervised learning from a multi-view perspective. *arXiv preprint arXiv:2006.05576*, 2020.

[Van Gansbeke *et al.*, 2020] Wouter Van Gansbeke, Simon Vandenhende, Stamatios Georgoulis, Marc Proesmans, and Luc Van Gool. Scan: Learning to classify images without labels. In *European conference on computer vision*, pages 268–285. Springer, 2020.

[Wang *et al.*, 2021] Rong Wang, Jitao Lu, Yihang Lu, Feiping Nie, and Xuelong Li. Discrete multiple kernel k-means. In *IJCAI*, pages 3111–3117, 2021.

[Winn and Jojic, 2005] John Winn and Nebojsa Jojic. Locus: Learning object classes with unsupervised segmentation. In *Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1*, volume 1, pages 756–763. IEEE, 2005.

[Xu *et al.*, 2022] Jie Xu, Huayi Tang, Yazhou Ren, Liang Peng, Xiaofeng Zhu, and Lifang He. Multi-level feature learning for contrastive multi-view clustering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16051–16060, June 2022.

[Xu *et al.*, 2023] Jie Xu, Yazhou Ren, Huayi Tang, Zhimeng Yang, Lili Pan, Yang Yang, Xiaorong Pu, Philip S. Yu, and Lifang He. Self-supervised discriminative feature learning for deep multi-view clustering. *IEEE Transactions on Knowledge and Data Engineering*, 35(7):7470–7482, 2023.

[Yan *et al.*, 2023] Weiqing Yan, Yuanyang Zhang, Chenlei Lv, Chang Tang, Guanghui Yue, Liang Liao, and Weisi Lin. Gcfagg: Global and cross-view feature aggregation for multi-view clustering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19863–19872, June 2023.

[Yang *et al.*, 2023] 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.

[Zhao *et al.*, 2021] Wei Zhao, Cai Xu, Ziyu Guan, and Ying Liu. Multiview concept learning via deep matrix factorization. *IEEE Transactions on Neural Networks and Learning Systems*, 32(2):814–825, 2021.

[Zhao *et al.*, 2023] Mingyu Zhao, Weidong Yang, and Feiping Nie. Deep multi-view spectral clustering via ensemble. *Pattern Recognition*, 144:109836, 2023.