

# Graph Cut-guided Maximal Coding Rate Reduction for Learning Image Embedding and Clustering

Wei He<sup>1</sup>[0000-0003-2222-8818], Zhiyuan Huang<sup>[0009-0004-4047-3328]</sup><sup>1</sup>, Xianghan Meng<sup>1</sup>[0009-0006-4042-917X], Xianbiao Qi<sup>2</sup>[0000-0002-8493-1966], Rong Xiao<sup>3</sup>[0000-0003-2207-5698], and Chun-Guang Li<sup>\*,1</sup>[0000-0002-5716-268X]

<sup>1</sup> Beijing University of Posts and Telecommunications, Beijing, P.R. China  
{wei.he, huangzhiyuan, mengxianghan, lichunguang}@bupt.edu.cn

<sup>2</sup> Visual Computing Group, International Digital Economy Academy  
(IDEA), Shenzhen, P.R. China

<sup>3</sup> Intellifusion, Shenzhen, P.R. China

**Abstract.** In the era of pre-trained models, image clustering task is usually addressed by two relevant stages: a) to produce features from pre-trained vision models; and b) to find clusters from the pre-trained features. However, these two stages are often considered separately or learned by different paradigms, leading to suboptimal clustering performance. In this paper, we propose a unified framework, termed graph Cut-guided Maximal Coding Rate Reduction (CgMCR<sup>2</sup>), for jointly learning the structured embeddings and the clustering. To be specific, we attempt to integrate an efficient clustering module into the principled framework for learning structured representation, in which the clustering module is used to provide partition information to guide the cluster-wise compression and the learned embeddings is aligned to desired geometric structures in turn to help for yielding more accurate partitions. We conduct extensive experiments on both standard and out-of-domain image datasets and experimental results validate the effectiveness of our approach.

## 1 Introduction

Image clustering, as a fundamental problem in computer vision and pattern recognition, aims to group images without annotated labels [18]. Image clustering is usually addressed by two successive stages: a) learning representation to generate features from images at first, and then b) finding clusters from the learned representation. State-of-the-art methods for image clustering, e.g., subspace clustering algorithms [7, 29, 47], demonstrate remarkable performance on simple datasets such as MNIST [25] and COIL [33] when proper features (e.g., Scatter Transform [3]) are provided. In general, different clustering algorithms implicitly employ different assumptions about the geometry of the clusters. For instance, in  $k$ -means algorithm [32] each cluster is modeled as a standard Gaussian distribution and thus characterized by the mean vector; in subspace clustering [45], each cluster is modeled as a low-dimensional subspace; in manifold

clustering [42] each cluster is modeled as a (low-dimensional) submanifold. Clustering performance will dramatically degenerate when the data distribution deviates from the assumption on the clusters.

When dealing with challenging and complex datasets, such as CIFAR [23] and ImageNet [9], the learned features play a more crucial role, as clustering based on conventional feature extraction typically fails to achieve satisfactory results. The key ingredient among the recent remarkable progress in clustering is learning features by pre-trained models (e.g., auto-encoders [22, 40] and contrastive learning [5]) that are suitable for the downstream clustering task. In the era of pre-trained models, deep clustering methods are designed by learning features via visual pretraining and then learning the cluster membership from refined features [26, 44]. More recently, when the large-scale pre-trained models, e.g., DINO [4, 37] and CLIP [39], are prevailing, deep clustering methods achieve impressive clustering performance by leveraging the rich representations produced by large-scale pre-trained models [1, 27]. While promising clustering accuracy has been reported, these deep clustering methods are usually designed to learn the pseudo labels without leveraging the potential effect from the structures of embeddings. There are a few attempts [8, 11, 28] to learn the embeddings with desired structures—*union of orthogonal subspaces*, by leveraging the framework of Maximal Coding Rate Reduction (MCR<sup>2</sup>) [48]. Nonetheless, none of these methods have developed a principled joint optimization framework to learn both the structured embeddings and the clustering.

In this paper, we present an effective joint optimizing framework, termed graph **C**ut-guided **M**aximal **C**oding **R**ate **R**eduction (**CgMCR<sup>2</sup>**), to learn both the structured embeddings and the clusters in principled way. Specifically, in CgMCR<sup>2</sup>, we integrate both the normalized cut based clustering learning and the MCR<sup>2</sup> based structured representations learning to form a unified optimization problem. Moreover, we design a two-stage training procedure, which consists of a one-shot initialization and a fine-tuning, to effectively train the proposed joint learning framework. We conduct extensive experiments on benchmark datasets to validate the superior performance of our proposed approach and also provide a set of ablation studies to evaluate the effect of each component.

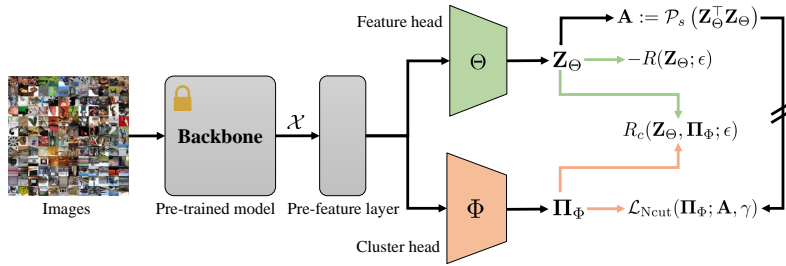
## 2 Relate Work

**Pre-trained Vision models.** Typically, pre-trained models leverage self-supervised pretext tasks to learn representations from unlabeled datasets. For example, Autoencoders [22] use an encoder-decoder architecture to learn latent low-dimensional representations by requiring the decoder to reconstruct the encoder inputs; contrastive learning (i.e., SimCLR [5]) exploits the self-supervision information via data augmentation and learns representations that maximize the agreement between positive pairs and the disagreement between negative pairs; and the following studies are proposed to improve the contrastive learning by enhancing the training stability (e.g., MoCo [6, 15]), reducing the requirement for negative samples (e.g., BYOL [13]), or avoiding the collapse solution of learned

features (e.g., VICReg [2]). More recently, large-scale pre-trained models based on large models such as BERT [10] and Vision Transformers (ViTs) [12] have showcased the capability to learn rich representations from diverse data sources. For instance, MAE [14] leverages the ViT as the backbone of auto-encoder and uses the input images with a high mask proportion for training. Inspired by self-supervised pretraining tasks in natural language processing, DINO [4, 37] implements a self-distillation framework with ViTs without using annotated labels. CLIP [39] pretrains a vision-language model on image-text data pairs to learn visual concepts from a text-guided contrastive learning task. Although the pre-trained models mentioned above successfully learn meaningful features, the structure of the clusters in the learned features remains unclear.

**Clustering via pre-trained models.** The success of pre-trained models has led to breakthroughs in image clustering. SCAN [44] suggests a three-stage clustering pipeline: learning embeddings from a pre-trained model, using a clustering module for label prediction, and fine-tuning the clustering module using pseudo-labels; RUC [38] and SPICE [35] enrich the pipeline of SCAN by utilizing robust training, refining network architecture or adjusting the fine-tuning strategy; CC [26] and GCC [49] propose unified frameworks for feature learning and clustering by optimizing the instance- and cluster-wise contrastive loss and graph contrastive loss, respectively. IMC-SwAW [36] integrates a discrete representation into the self-supervised learning via a classifier net to simultaneously learn the cluster membership. MiCE [43] introduces a probabilistic clustering framework that combines contrastive learning with a latent mixture of experts. ProPos [17] combines prototype scattering and positive sampling using EM-like steps to learn uniform, well-separated representations. More recently, TEMI [1] proposes a self-distillation clustering framework by leveraging large-scale pre-trained models; and TAC [27] brings in the pre-trained CLIP texture embedding as external guidance for image clustering. Unfortunately, the intrinsic structure of representations learned or refined by the deep clustering methods mentioned above is still unclear, and thus usually only the nearest neighbors' information can be leveraged for clustering.

**Clustering via MCR<sup>2</sup>.** The framework of Maximal Coding Rate Reduction (MCR<sup>2</sup>) [48] is designed for supervised learning, to learn compact and structured representations that enjoy both the diversity in each class and the discriminativity between classes. There are some attempts to use MCR<sup>2</sup> for deep clustering, e.g., NMCE [28] designs a specialized self-supervised learning module for initialization and subsequently optimizes the MCR<sup>2</sup> objective starting from the randomly initialized cluster membership; MLC [11] leverages the self-supervised learning module of NMCE for initialization and exploits a doubly stochastic affinity for partitioning the learned embeddings; then MLC is further evaluated on the CLIP pre-trained feature [8]. While remarkable performance has been obtained, none of them have developed a principled joint framework to learn both the structured embeddings and the clustering.



**Fig. 1:** Illustration for our CgMCR<sup>2</sup> architecture. We illustrate the forward pass (in black) and the gradient dependency of  $\mathbf{Z}_\Theta$  (in green) and  $\mathbf{\Pi}_\Phi$  (in orange) in different colored lines individually. For clarity, we exclude the parameters of the pre-feature layer from the visualization.

### 3 Our Approach: Graph Cut-guided Maximal Coding Rate Reduction (CgMCR<sup>2</sup>)

We begin with a brief review of the principle of MCR<sup>2</sup> in Section 3.1 and then present our framework for jointly learning structured embedding and clustering—CgMCR<sup>2</sup> in Section 3.2.

#### 3.1 Preliminaries of MCR<sup>2</sup>

**Supervised feature learning via MCR<sup>2</sup>.** Given a dataset  $\mathcal{X} = \{\mathbf{x}_i\}_{i=1}^N$  of  $N$  data points where  $\mathbf{x}_i \in \mathbb{R}^D$  and the ground-truth labels  $\mathbf{\Pi}^* \in \{0, 1\}^{N \times k}$  to assign these  $N$  points into  $k$  classes  $\mathcal{C} = \bigcup_{\ell=1}^k \mathcal{C}_\ell$ . The framework of Maximal Coding Rate Reduction (MCR<sup>2</sup>) [48] learns the embedding  $\mathbf{Z} \in \mathbb{R}^{d \times N}$  by maximizing the following objective:

$$\max_{\mathbf{Z}} R(\mathbf{Z}; \epsilon) - R_c(\mathbf{Z}, \mathbf{\Pi}^*; \epsilon) \quad (1)$$

where  $R(\mathbf{Z}; \epsilon) := \log \det \left( \mathbf{I} + \frac{d}{N\epsilon^2} \mathbf{Z}\mathbf{Z}^\top \right)$ ,  $R_c(\mathbf{Z}, \mathbf{\Pi}^*; \epsilon) := \frac{1}{N} \sum_{\ell=1}^k \log \det \left( \mathbf{I} + \frac{d}{N_\ell \epsilon^2} \mathbf{Z} \text{Diag}(\mathbf{\Pi}_\ell^*) \mathbf{Z}^\top \right)$ , and  $N_\ell$  is the number of data points in class  $\mathcal{C}_\ell$ . The first term  $R(\mathbf{Z}; \epsilon)$  measures the average coding length (a.k.a the coding rate) of the embeddings  $\mathbf{Z}$  subject to a prescribed rate distortion precision  $\epsilon > 0$ , and the second term  $R_c(\mathbf{Z}, \mathbf{\Pi}^*; \epsilon)$  measures the sum of the coding rate of each class indicated by  $\mathbf{\Pi}^*$ . Roughly, maximizing the second term  $-R_c(\mathbf{Z}, \mathbf{\Pi}^*; \epsilon)$  encourages the within-class embeddings to span a low-dimensional linear subspace; meanwhile, maximizing the first term  $R(\mathbf{Z}; \epsilon)$  encourages the embeddings as a whole to expand, and thus making the class-specific subspaces being orthogonal to each other. Such an arrangement of the class-specific subspaces is a desired property of the embeddings learned by MCR<sup>2</sup> and has been proved in [48].

**Deep Clustering based on MCR<sup>2</sup>.** The recent attempts [8, 11, 28] employing the MCR<sup>2</sup> framework for deep clustering try to optimize the rate reduction

objective to learn the representations and the partition jointly as follows:

$$\max_{\Theta, \Phi} R(\mathbf{Z}_\Theta; \epsilon) - R_c(\mathbf{Z}_\Theta, \mathbf{\Pi}_\Phi; \epsilon), \quad \text{s.t. } \mathbf{Z}_\Theta = f(\mathcal{X}; \Theta), \mathbf{\Pi}_\Phi = g(\mathcal{X}; \Phi), \quad (2)$$

where  $\mathbf{Z}_\Theta \in \mathbb{R}^{d \times N}$  and  $\mathbf{\Pi}_\Phi \in \mathbb{R}^{N \times k}$  are implemented by feature head  $f(\cdot; \Theta)$  and cluster head  $g(\cdot; \Phi)$  which are parameterized by  $\Theta$  and  $\Phi$ , respectively. Although the representations  $\mathbf{Z}_\Theta$  and the partition  $\mathbf{\Pi}_\Phi$  are jointly optimized, there lacks of an effective and principled mechanism to learn the partition (i.e., the cluster memberships) of the embeddings.

### 3.2 Graph Cut-guided Maximal Coding Rate Reduction

Following the prior attempts based on MCR<sup>2</sup> [8, 11, 28] to learn both the embeddings and the partition, we try to incorporate a principled way to learn the cluster membership  $\mathbf{\Pi}_\Phi$  into the framework (2).

**Spectral clustering for estimating  $\mathbf{\Pi}_\Phi$ .** Spectral Clustering [31], as a popular clustering algorithm with solid theoretical foundation, provides a principled approach to learn the partition  $\mathbf{\Pi}_\Phi$  based on an affinity matrix built from the structured representation  $\mathbf{Z}_\Theta$ . Recall that, given an affinity matrix  $\mathbf{A} \in \mathbb{R}^{N \times N}$  where entries  $a_{i,j}$  measures the similarities of paired data points  $(\mathbf{x}_i, \mathbf{x}_j)$ , spectral clustering [31] aims to find an ideal cluster membership  $\mathbf{\Pi} \in \{0, 1\}^{N \times k}$  as follows:

$$\min_{\mathbf{\Pi}} \text{trace}(\mathbf{\Pi}^\top \mathbf{L} \mathbf{\Pi}), \quad \text{s.t. } \mathbf{\Pi} \in \{0, 1\}^{N \times k} \quad (3)$$

where  $\mathbf{L} = \mathbf{D} - \mathbf{A}$  is known as the graph Laplacian corresponding to the affinity  $\mathbf{A}$  and  $\mathbf{D} = \text{Diag}(d_{1,1}, \dots, d_{N,N})$  is the degree matrix with diagonal entries  $d_{i,i} = \sum_{j=1}^N a_{i,j}$ . To avoid trivial partition, Normalized cut (Ncut) [41] introduces the volume of each partition to define the entries of  $\mathbf{\Pi}$  as

$$\tilde{\pi}_{i,\ell} := \begin{cases} \frac{1}{\sqrt{|\mathcal{V}_\ell|}} & \text{if } \mathbf{x}_i \in \mathcal{C}_\ell \\ 0 & \text{if } \mathbf{x}_i \notin \mathcal{C}_\ell, \end{cases} \quad (4)$$

where  $|\mathcal{V}_\ell| := \sum_{i:\mathbf{x}_i \in \mathcal{C}_\ell} d_{i,i}$  is the volume of  $\ell$ -th cluster. Rather than solving the combinatorial problem in (3), as a convention, spectral clustering methods [31] reformulate the problem with continuous relaxation and then solve the relaxed problem. For example, Normalized cut (Ncut) [41] solves the problem as follows:

$$\min_{\tilde{\mathbf{\Pi}}} \text{trace}(\tilde{\mathbf{\Pi}}^\top \mathbf{L} \tilde{\mathbf{\Pi}}), \quad \text{s.t. } \tilde{\mathbf{\Pi}}^\top \mathbf{D} \tilde{\mathbf{\Pi}} = \mathbf{I}. \quad (5)$$

Note that the solution is the ending  $k$  eigenvectors associated with the ending  $k$  minor eigenvalues of  $\tilde{\mathbf{L}} = \mathbf{D}^{-\frac{1}{2}} \mathbf{L} \mathbf{D}^{-\frac{1}{2}}$ . Unfortunately, the solution of the relaxed problem obtained by eigen-decomposition contains negative entries and can not be directly used as the cluster membership. Thus a  $k$ -means algorithm is adopted to generate the final clustering results.

**Unified formulation of CgMCR<sup>2</sup>.** Rather than directly using the conventional NCut, we follow [16] to relax the normalized cut problem as follows:

$$\begin{aligned} \min_{\mathbf{\Pi}} \quad & \text{trace}((\mathbf{\Pi}\mathbf{V})^\top \mathbf{L}(\mathbf{\Pi}\mathbf{V})) + \frac{\gamma}{2} \|(\mathbf{\Pi}\mathbf{V})^\top \mathbf{D}(\mathbf{\Pi}\mathbf{V}) - \mathbf{I}\|_F^2, \\ \text{s.t.} \quad & 0 \leq \mathbf{\Pi} \leq \mathbf{1}, \mathbf{\Pi} \cdot \mathbf{1} = \mathbf{1}, \end{aligned} \quad (6)$$

where  $\gamma$  is a trade-off parameter, and  $\mathbf{V}$  consists of estimated volumes for all the clusters, i.e.:

$$\mathbf{V} := \text{Diag} \left( \sum_{i=1}^N \pi_{i,1} d_{i,i}, \dots, \sum_{i=1}^N \pi_{i,k} d_{i,i} \right)^{-1/2}. \quad (7)$$

Note that the objective in Eq. (6) enforces strict numerical constraints to have an ideal membership  $\mathbf{\Pi}$  while relaxing the orthogonal constraints in Eq. (5) to a penalty term. This relaxation enable us to develop a differentiable approach for spectral clustering to learn the cluster membership directly (rather than the spectral embeddings). Interestingly, these strict numerical constraints can easily be satisfied by reparameterizing  $\mathbf{\Pi}$  with a neural network  $\mathbf{\Pi}_\Phi := g(\cdot; \Phi)$  where  $\Phi$  denotes all the parameters in  $g(\cdot)$  and a `softmax` function is used as the output.

By putting problems in (2) and (6) together, we have an effective and principled joint optimization framework as follows:

$$\min_{\Theta, \Phi} -R(\mathbf{Z}_\Theta; \epsilon) + R_c(\mathbf{Z}_\Theta, \mathbf{\Pi}_\Phi; \epsilon) + \mathcal{L}_{\text{Ncut}}(\mathbf{\Pi}_\Phi; \mathbf{A}, \gamma), \quad (8)$$

where  $\mathcal{L}_{\text{Ncut}}(\mathbf{\Pi}_\Phi; \mathbf{A}, \gamma) = \text{trace}(\tilde{\mathbf{\Pi}}_\Phi^\top \mathbf{L} \tilde{\mathbf{\Pi}}_\Phi) + \frac{\gamma}{2} \|\tilde{\mathbf{\Pi}}_\Phi^\top \mathbf{D} \tilde{\mathbf{\Pi}}_\Phi - \mathbf{I}\|_F^2$ ,  $\tilde{\mathbf{\Pi}}_\Phi = \mathbf{\Pi}_\Phi \mathbf{V}$ .

We use the cosine similarity of  $\mathbf{Z}_\Theta$  in current training iteration to define the affinity, i.e.,  $\mathbf{A} := \mathcal{P}_s(\mathbf{Z}_\Theta^\top \mathbf{Z}_\Theta)$ , where  $\mathcal{P}_s$  is a post-process operator. In practice, we simply reserve the  $s$  largest entries of each row in  $\mathbf{A}$ . We termed this framework in (8) as a graph **C**ut-**g**uided **M**aximal **C**oding **R**ate **R**eduction (**CgMCR<sup>2</sup>**).

**Remarks.** The advantages of the unified CgMCR<sup>2</sup> framework are *three-folds*.

- The cluster membership is obtained by a principled way via  $\mathcal{L}_{\text{Ncut}}(\mathbf{\Pi}_\Phi; \mathbf{A}, \gamma)$  and it can guide the optimization of  $R_c(\mathbf{Z}_\Theta, \mathbf{\Pi}_\Phi; \epsilon)$  for refining  $\mathbf{Z}_\Theta$ .
- The framework is differentiable and scalable, enabling the GPU acceleration for training in mini-batch mode with stochastic gradient descent.
- The framework can infer the representations and cluster memberships for unseen test data.

### 3.3 Implementations

**Network architecture.** For clarity, we illustrate our CgMCR<sup>2</sup> framework in Fig. 1. We utilize the frozen image encoder of CLIP [39] as a general feature extractor and add a single linear layer to generate the pre-trained CLIP feature. We then deploy a two-layers neural network as the feature head  $f(\cdot; \Theta)$  to

**Algorithm 1** Procedure for graph Cut-guided Maximal Coding Rate Reduction

- 
- 1: **Input:**  $N$  images,  $\epsilon, \gamma > 0$ , warm-up epochs  $T_1$ , fine-tuning epochs  $T_2$ , post-processing operator  $\mathcal{P}_s$  with sparsity  $s$ , batch size  $n$ , and learning rate  $\eta$
  - 2: **Initialization:**  $t = 0$ , pre-feature  $\mathcal{X}$ , and all trainable network parameters  $\Xi$
  - 3: **for** each  $t = 1, \dots, T_1 + T_2$  **do**
  - 4:   Pick a mini-batch sample  $\mathcal{X}^{(t)}$  with  $n$  data points from  $\mathcal{X}$
  - 5:   **# Forward pass**
  - 6:   Compute  $\mathbf{Z}_\Theta^{(t)}$  and  $\mathbf{\Pi}_\Phi^{(t)}$
  - 7:   Compute affinity  $\mathbf{A}^{(t)} := \mathcal{P}_s(\mathbf{Z}_\Theta^{(t)\top} \mathbf{Z}_\Theta^{(t)})$  and detach it from back propagation
  - 8:   **# Backward propagation**
  - 9:   **if**  $t \leq T_1$  **then**
  - 10:     Compute the gradients  $\nabla_\Xi$  with respect to objective in (9)
  - 11:   **else**
  - 12:     Compute the gradients  $\nabla_\Xi$  with respect to objective in (8)
  - 13:   **end if**
  - 14:    $\Xi^{(t+1)} \leftarrow \Xi^{(t)} - \eta \nabla_\Xi$
  - 15:    $t \leftarrow t + 1$
  - 16: **end for**
  - 17: **Output:**  $\Xi^{(t+1)}$
- 

learn the representation  $\mathbf{Z}_\Theta$ , and use a two-layers neural network with Gumbel-Softmax [19] as the cluster head  $g(\cdot; \Phi)$  to learn the cluster membership  $\mathbf{\Pi}_\Phi$ .

**Training procedure.** Taking into account the fact that optimizing  $R_c(\mathbf{Z}_\Theta, \mathbf{\Pi}_\Phi; \epsilon)$  for learning  $\mathbf{Z}_\Theta$  could be inefficient and challenging when  $\mathbf{\Pi}_\Phi$  is randomly initialized, we propose an efficient two-stage training strategy, which consists of a *one-shot initialization* stage and a *fine-tuning* stage.

In the *initialization* stage, we take  $T_1$  epochs for warm-up to learn the initial discriminative representation and the initial cluster membership by solving a simplified pretext task:

$$\min_{\Theta, \Phi} -R(\mathbf{Z}_\Theta; \epsilon) + \mathcal{L}_{\text{Ncut}}(\mathbf{\Pi}_\Phi; \mathbf{A}, \gamma). \quad (9)$$

Note that the second term  $R_c(\mathbf{Z}_\Theta, \mathbf{\Pi}_\Phi; \epsilon)$  in (8) is temporally ignored due to the lack of a good initialization for  $\mathbf{\Pi}_\Phi$ . Moreover, we detach the gradients with respect to  $\Theta$  (due to the affinity  $\mathbf{A}$ ) in the  $\mathcal{L}_{\text{Ncut}}(\mathbf{\Pi}_\Phi; \mathbf{A}, \gamma)$  term from the back propagation and thus the feature head and the cluster head can be separately trained via the loss  $-R(\mathbf{Z}_\Theta; \epsilon)$  and the loss  $\mathcal{L}_{\text{Ncut}}(\mathbf{\Pi}_\Phi; \mathbf{A}, \gamma)$ , respectively.

We then optimize the CgMCR<sup>2</sup> objective in Eq. (8) for fine-tuning. The well-initialized  $\mathbf{\Pi}_\Phi$  now provides a *self-supervised information* that guides the learning of within-cluster compact representation via the term  $R_c(\mathbf{Z}_\Theta, \mathbf{\Pi}_\Phi; \epsilon)$ . Consequently, the feature head and the cluster head are jointly learned to perform within-cluster compact and between-cluster discriminative representation learning as well as the optimal cluster membership in the fine-tuning. The whole procedure for training our CgMCR<sup>2</sup> is summarize in Algorithm 1.

**Comparison to NMCE, MLC and CPP.** In Table 1, we summarize the connections and differences between our CgMCR<sup>2</sup> and other deep clustering

**Table 1: Comparing our CgMCR<sup>2</sup> to deep clustering methods via MCR<sup>2</sup>.**

| Methods                  | Off-the-shelf pretrain | Initializing $\mathbf{Z}_\Theta$                                    | Initializing $\mathbf{\Pi}_\Phi$                        | Objective          |
|--------------------------|------------------------|---|---|--------------------|
| NMCE                     | ×                      | $\Theta \leftarrow \nabla_\Theta (R + \mathcal{L}_{\text{SimCLR}})$ | Random  | MCR <sup>2</sup>   |
| MLC                      | ×                      | $\Theta \leftarrow \nabla_\Theta (R + \mathcal{L}_{\text{SimCLR}})$ | $\Phi \leftarrow \Theta$                                | MCR <sup>2</sup>   |
| CPP                      | ✓                      | $\Theta \leftarrow \nabla_\Theta R$                                 | $\Phi \leftarrow \Theta$                                | MCR <sup>2</sup>   |
| <b>CgMCR<sup>2</sup></b> | ✓                      | $\Theta \leftarrow \nabla_\Theta R$                                 | $\Phi \leftarrow \nabla_\Phi \mathcal{L}_{\text{Ncut}}$ | CgMCR <sup>2</sup> |

methods based on the framework of MCR<sup>2</sup>. These works proceed by initializing representations through pre-trained models, and refining the representation and cluster membership after the initialization, in which NMCE [28] and MLC [11] leverage a specially designed objective that incorporates the contrastive learning objective ( $\mathcal{L}_{\text{SimCLR}}$ ) [5] for pretraining, whileas CPP [8] utilizes the pre-trained CLIP features. Among these works,  $\mathbf{\Pi}_\Phi$  is either randomly initialized (as in NMCE) or initialized by copying the parameters  $\Theta$  to  $\Phi$  (as in MLC and CPP). However, none of these methods employ a principled approach to learn the cluster membership in their optimization framework. As a result, NMCE is sensitive to the random initialization, whereas MLC and CPP rely on additional clustering algorithms to obtain cluster memberships.

## 4 Experimental Results

**Datasets.** We evaluate the performance of our CgMCR<sup>2</sup> on MNIST [25], Fashion MNIST (F-MNIST) [46], CIFAR-10, CIFAR-20, CIFAR-100 [23], Stanford Dogs-120 (Dogs-120) [20], Oxford Flowers-102 (Flowers-102) [34], TinyImageNet-200 (TinyImageNet) and ImageNet-1k. Among these datasets, CIFAR-20 contains the same data as CIFAR-100 with 20 super-classes as in [23], whileas Stanford Dogs-120, TinyImageNet and ImageNet-1k are subsets of ImageNet [9]. All data points are embedded to 768-dimensional vectors using a CLIP image encoder pre-trained on the ViT-L/14 backbone, or embedded to 512-dimensional vectors using the MoCov2 model pre-trained on the ResNet-18 backbone.

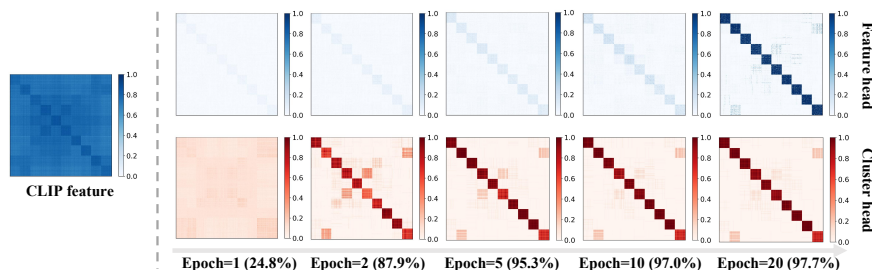
**Metrics.** To evaluate the clustering performance, we report the clustering accuracy (ACC) and the normalized mutual information (NMI). For ACC, we use the Hungarian matching algorithm [24] to find the best match between the pseudo-labels and ground-truth labels. By default, the pseudo-labels of CgMCR<sup>2</sup> are generated by the cluster head, i.e., the output of the `argmax` layer.

**Training settings.** We use Adam optimizer [21] with a fixed initial learning rate during the warm-up, and then using cosine annealing learning rate [30] during the fine-tuning. For all datasets, the output dimension  $k$  of cluster head is set to the number of true clusters.

### 4.1 Experiments on Standard Datasets

**Visualization of Affinity Matrices.** To demonstrate the ability of CgMCR<sup>2</sup> to produce both structured features and satisfactory clustering results during the





**Fig. 2: Similarity matrices** ordered by the ground-truth labels of the CLIP features computed by  $|\mathbf{X}^\top \mathbf{X}|$  (in the *left panel in blue*), representation computed by  $|\mathbf{Z}_\Theta^\top \mathbf{Z}_\Theta|$  (at the first row *in blue*) and cluster membership  $|\mathbf{\Pi}_\Phi \mathbf{\Pi}_\Phi^\top|$  (at the second row *in red*) of CgMCR<sup>2</sup> trained with  $\{1, 2, 5, 10, 20\}$  epochs on CIFAR-10, where the percentage number in bracket is ACC.

**Table 2: Clustering accuracy** (mean $\pm$ std) of feature head (*Top*) and the cluster head (*Bottom*) using CLIP features on five benchmark datasets over 3 trials.

| Methods                              | CIFAR-10       | CIFAR-20       | CIFAR-100      | TinyImageNet   | ImageNet-1k    |
|--------------------------------------|----------------|----------------|----------------|----------------|----------------|
| SC on $\mathbf{Z}_\Theta$            | 97.6 $\pm$ 0.1 | 68.8 $\pm$ 0.4 | 78.3 $\pm$ 0.3 | 72.7 $\pm$ 0.2 | 67.7 $\pm$ 0.2 |
| <b>argmax</b> on $\mathbf{\Pi}_\Phi$ | 97.7 $\pm$ 0.1 | 68.1 $\pm$ 0.4 | 77.8 $\pm$ 0.4 | 72.9 $\pm$ 0.2 | 67.5 $\pm$ 0.3 |

training process, we conduct a set of experiments on CIFAR-10 to visualize the affinity matrices of the original CLIP features  $\mathcal{X}$ , the feature head outputs  $\mathbf{Z}_\Theta$  and cluster head outputs  $\mathbf{\Pi}_\Phi$ . Visualization results are given as colored images in Fig. 2. As can be seen that, our framework learns structured representations (see, e.g., the block diagonal structure in *blue* at epoch 10) and produces acceptable initial cluster membership via the pretext task in (9). Then, fine-tuning the framework using the whole CgMCR<sup>2</sup> objective in (8) refines the structured representations (see, e.g., the block *diagonal* structure in *blue* at epoch 20). These structured representations, in turn, enable the cluster head to learn better cluster membership. Once trained, the cluster head can serve as a scalable predictor to produce satisfactory pseudo labels, achieving 97.7% accuracy on CIFAR-10.

**Performance Comparison between Feature Head and Cluster Head.** For a trained CgMCR<sup>2</sup>, rather than using the clustering head to yield the pseudo labels, we can also apply the conventional Spectral Clustering [41] to the affinity which is defined by  $\mathbf{A} := \mathcal{P}_s(\mathbf{Z}_\Theta^\top \mathbf{Z}_\Theta)$ . This method is denoted as “SC on  $\mathbf{Z}_\Theta$ ”. We report the comparison results in Table 2. As can be read, both the feature head and the cluster head perform almost equally good. In the following experiments, we also denote the performance of “SC on  $\mathbf{Z}_\Theta$ ” as “CgMCR<sup>2</sup>-SC” to distinguish it from the performance obtained by “argmax on  $\mathbf{\Pi}_\Phi$ ” (CgMCR<sup>2</sup>) for clarity.

**Comparison to Competing Clustering Methods using CLIP Features.** To evaluate the performance of our CgMCR<sup>2</sup>, we conduct experiments on five

**Table 3: Clustering Performance Comparison using CLIP Features on Five Benchmark Datasets.** ‘-’ denotes that the results are not available.

| Methods                          | CIFAR-10    |             | CIFAR-20    |             | CIFAR-100   |             | TinyImgNet  |             | ImageNet-1k |             |
|----------------------------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|
|                                  | ACC         | NMI         | ACC         | NMI         | ACC         | NMI         | ACC         | NMI         | ACC         | NMI         |
| <i>k</i> -means [32]             | 83.5        | 84.1        | 46.9        | 49.4        | 52.8        | 66.8        | 54.1        | 73.4        | 53.9        | 79.8        |
| Spectral [41]                    | 79.8        | 84.8        | 53.3        | 61.6        | 66.4        | 77.0        | 62.8        | 77.0        | 56.0        | 81.2        |
| EnSC [47]                        | 95.4        | 90.3        | 61.0        | 68.7        | 67.0        | 77.1        | 64.5        | 77.7        | 59.7        | 83.7        |
| SCAN [44]                        | 95.1        | 90.3        | 60.8        | 61.8        | 64.1        | 70.8        | 56.5        | 72.7        | 54.4        | 76.8        |
| TEMI [1]                         | 96.9        | 92.6        | 61.8        | 64.5        | 73.7        | 79.9        | -           | -           | 64.0        | -           |
| CPP [8]                          | 97.4        | 93.6        | 64.2        | 72.5        | 74.0        | 81.8        | 63.4        | 77.3        | 62.0        | 82.1        |
| <b>CgMCR<sup>2</sup></b>         | <b>97.7</b> | <b>94.3</b> | <b>68.1</b> | <b>73.8</b> | <b>77.8</b> | <b>81.9</b> | <b>72.9</b> | <b>81.4</b> | <b>67.5</b> | <b>87.0</b> |
| <b>CgMCR<sup>2</sup>-SC</b>      | <b>97.6</b> | <b>94.2</b> | <b>68.8</b> | <b>74.0</b> | <b>78.3</b> | <b>82.5</b> | <b>72.7</b> | <b>81.1</b> | <b>67.7</b> | <b>87.1</b> |
| <i>External texture guidance</i> |             |             |             |             |             |             |             |             |             |             |
| TAC [27]                         | 97.0        | 92.4        | 66.8        | 73.2        | 75.5        | 81.1        | 71.0        | 79.9        | 66.4        | 86.8        |

datasets and compare to a set of competing baseline methods. As the baseline, we choose the classical clustering algorithms, including *k*-means [32] and spectral clustering with normalized cut [41], subspace clustering method, Elastic Net Subspace Clustering (EnSC) [47], deep clustering methods, including SCAN [44], TEMI [1], CPP [8] and TAC [27]. All these methods are conducted on the CLIP features. We report the experimental results in Table 3. The results of CPP and TAC are reproduced with the released codes. The results of TEMI [1] are cited from the paper. We can read that, all methods yield promising performance owing to the CLIP feature. But, clearly, our CgMCR<sup>2</sup> achieves superior clustering performance. The performance improvements over CPP [8], which is also based on the framework of MCR<sup>2</sup>, are due to the principled way to produce the clustering membership. We note that TAC [27] also yields competitive results, but it leverages the external information brought by the CLIP text encoder.

**Comparison to State-of-the-art Deep Clustering Methods using MoCov2 Features.** We apply our CgMCR<sup>2</sup> framework on the pre-trained MoCov2 [6] features, and compare to state-of-the-art deep clustering methods, including CC [26], GCC [49], SCAN [44], SPICE [35], IMC-SwAV [36], NMCE [28] and MLC [11]. Also, we report the performance of *k*-means [32], Spectral Clustering [41] and Elastic Net Subspace Clustering (EnSC) [47]. All methods, except for NMCE, MLC and IMC-SwAV, are conducted on the pre-trained MoCov2 features; whereas NMCE, MLC and IMC-SwAV are performed with their specially designed pre-trained models, respectively. The results of TEMI [1] is reproduced by using the released codes on the MoCov2 features. For a relatively fair comparison, all method use ResNet18 as the backbone of the pre-trained models. Experimental results are listed in Table 4. As can be read that, our CgMCR<sup>2</sup> still achieves the leading clustering accuracy on majority cases.

**Table 4: Clustering Performance Comparison to State-of-the-art Deep Clustering Methods.**

| Methods                     | CIFAR-10    |             | CIFAR-20    |             | CIFAR-100   |             | TinyImageNet |             |
|-----------------------------|-------------|-------------|-------------|-------------|-------------|-------------|--------------|-------------|
|                             | ACC         | NMI         | ACC         | NMI         | ACC         | NMI         | ACC          | NMI         |
| <i>k</i> -means [32]        | 71.6        | 61.8        | 41.4        | 41.3        | 42.8        | 42.6        | 17.8         | 41.3        |
| Spectral [41]               | 80.7        | 69.2        | 45.4        | 42.6        | 40.9        | 56.9        | 20.3         | 36.8        |
| EnSC [47]                   | 83.2        | 74.1        | 50.6        | 45.5        | 41.2        | 62.7        | 25.3         | 40.4        |
| CC [26]                     | 79.0        | 70.5        | 42.9        | 43.1        | 36.9        | 58.1        | 24.0         | 44.0        |
| GCC [49]                    | 85.6        | 76.4        | 47.2        | 47.2        | 38.2        | 59.9        | 23.8         | 44.7        |
| SCAN [44]                   | 88.3        | 79.7        | 50.7        | 48.6        | 34.3        | 55.7        | -            | -           |
| SPICE [35]                  | 91.8        | 85.0        | 53.5        | <b>56.5</b> | -           | -           | 30.5         | 44.9        |
| IMC-SwAV [36]               | 89.7        | 81.8        | 51.9        | 52.7        | 45.1        | 67.5        | 27.9         | 48.5        |
| TEMI [1]                    | 88.7        | 84.1        | 47.3        | 48.7        | 45.4        | 66.7        | 27.8         | 42.3        |
| NMCE [28]                   | 83.0        | 76.1        | 43.7        | 48.8        | 40.0        | 53.9        | 21.6         | 40.0        |
| MLC [11]                    | 86.3        | 78.3        | 52.2        | <u>54.6</u> | 49.4        | <b>68.3</b> | 33.5         | <b>67.5</b> |
| <b>CgMCR<sup>2</sup></b>    | <b>92.8</b> | <b>88.6</b> | <u>55.6</u> | 54.3        | <u>49.8</u> | 67.9        | <u>35.9</u>  | 62.5        |
| <b>CgMCR<sup>2</sup>-SC</b> | <u>92.7</u> | <u>88.4</u> | <b>56.1</b> | 54.4        | <b>51.1</b> | <u>68.0</u> | <b>36.7</b>  | <u>62.9</u> |

## 4.2 Experiments on Out-of-Domain Datasets

To demonstrate the effectiveness of jointly learning both the structured representation and the clustering, we apply the pre-trained CLIP [39] to extract features for datasets MNIST, F-MNIST, Flowers-102 and Dogs-120, which are quite different from the training data for the pre-trained CLIP. We compare the performance of our CgMCR<sup>2</sup> to four representative clustering methods, including Spectral Clustering, EnSC, SCAN and CPP. Experimental results are listed in Table 5. We can observe that, our CgMCR<sup>2</sup> still yields satisfactory clustering accuracy on out-of-domain datasets, especially with a notable accuracy improvement of +10.3% on Dogs-120. Besides, we notice of that SCAN fails on Flowers-102, which is an *imbalanced* dataset. This is because that SCAN is designed on a class-balance assumption, which is unsatisfied on Flowers-102. On contrary, our CgMCR<sup>2</sup> still produces satisfactory clustering results.

**Table 5: Clustering Performance on Out-of-domain Datasets.**

| Methods                     | MNIST       |             | F-MNIST     |             | Flowers-102 |             | Dogs-120    |             |
|-----------------------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|
|                             | ACC         | NMI         | ACC         | NMI         | ACC         | NMI         | ACC         | NMI         |
| Spectral [41]               | 74.5        | 67.0        | 64.3        | 56.8        | 85.6        | 94.6        | 44.1        | 55.6        |
| EnSC [47]                   | 91.0        | 85.3        | 69.1        | 65.1        | 90.0        | 95.9        | 40.1        | 60.8        |
| SCAN [44]                   | 87.4        | 81.9        | 69.3        | 67.2        | 40.4        | 68.7        | 38.1        | <u>73.8</u> |
| CPP [8]                     | 95.7        | 90.4        | 70.9        | 68.8        | <u>91.3</u> | <u>96.4</u> | 51.0        | 69.5        |
| <b>CgMCR<sup>2</sup></b>    | <b>96.9</b> | <b>92.8</b> | <u>74.5</u> | <u>69.9</u> | 91.1        | 96.1        | <u>60.9</u> | 73.6        |
| <b>CgMCR<sup>2</sup>-SC</b> | <u>96.4</u> | <u>92.0</u> | <b>74.7</b> | <b>71.1</b> | <b>92.2</b> | <b>97.0</b> | <b>61.3</b> | <b>75.1</b> |




Fig. 3: Learning curves of each loss term, ACC, and NMI during training.

### 4.3 More Evaluation and Analysis

**Learning curves.** To show the effectiveness of the two-stage training procedure, we display the learning curves of the proposed CgMCR<sup>2</sup> on CIFAR-20, where we use 10 epochs (each epoch contains 33 mini-batch iterations) for warm-up and another 40 epochs for fine-tuning. In Fig. 3, we show the learning curve of each term in the objective of CgMCR<sup>2</sup>, i.e.,  $R(\mathbf{Z}_\Theta; \epsilon)$ ,  $R_c(\mathbf{Z}_\Theta, \mathbf{\Pi}_\Phi; \epsilon)$  and  $\mathcal{L}_{Ncut}(\mathbf{\Pi}_\Phi; \mathbf{A}, \gamma)$ , ACC and NMI during the training. As can be observed in Fig. 3 that, the initialization stage did perform a good warm-up, and during the fine-tuning stage, the value of  $R(\mathbf{Z}_\Theta; \epsilon)$  remains nearly constant, but  $R_c(\mathbf{Z}_\Theta, \mathbf{\Pi}_\Phi; \epsilon)$  is rapidly optimized toward to its minimum. Moreover, the curves of ACC and NMI improve rapidly and achieves the optimal clustering results with the self-supervision of well-initialized cluster membership in the fine-tuning stage.

**Evaluation on Different Training Strategy.** To validate the effectiveness of our two-stage training procedure, we conduct experiments on CIFAR-10 and CIFAR-100. Experimental results are shown in Table 6. The one-shot initialization learned by the pretext task (i.e.,  $-R + \mathcal{L}_{Ncut}$ ) evidently improves the performance of optimizing both MCR<sup>2</sup> (i.e.,  $-R + R_c$ ) and CgMCR<sup>2</sup> (i.e.,  $-R + R_c + \mathcal{L}_{Ncut}$ ) objectives in subsequent fine-tuning. We can see that combining the pretext-task with MCR<sup>2</sup> objective also serves as a competitive baseline, demonstrating the importance of a principled initialization for  $\mathbf{\Pi}_\Phi$ . The proposed CgMCR<sup>2</sup> objective also demonstrates better empirical performance compared to the MCR<sup>2</sup> objective, regardless of the training strategy employed.

**Evaluation on Sensitivity of Hyper-parameters.** To evaluate the sensitivity of the performance of our CgMCR<sup>2</sup> to the hyper-parameters, we conduct experiments on CIFAR-10 and CIFAR-100. We vary the value of hyper-parameters  $\epsilon$  and  $\gamma$  in (8) and report the clustering results. We set  $\epsilon$  in the range of  $\{0.2, 0.3, 0.4, 0.5\}$  while  $\gamma$  in a wide range of feasible region.<sup>4</sup> In Fig. 4 we show the experimental results. As can be observed that our method is less sensitive to

<sup>4</sup> We follow the practical method in [16] to find the feasible  $\gamma$  without using the ground-truth labels.

**Table 6:** Evaluation on Different Training Strategy on CIFAR-10 and CIFAR-100. We report ACC and NMI of the cluster head outputs.

| Warm-up                   | Fine-tune                       | CIFAR-10    |             | CIFAR-100   |             |
|---------------------------|---------------------------------|-------------|-------------|-------------|-------------|
|                           |                                 | ACC         | NMI         | ACC         | NMI         |
| N/A                       | $-R + R_c$                      | 86.5        | 86.8        | 60.4        | 67.4        |
| N/A                       | $-R + R_c + \mathcal{L}_{Ncut}$ | 91.2        | 89.3        | 66.0        | 70.8        |
| $-R + \mathcal{L}_{Ncut}$ | $-R + R_c$                      | <u>97.3</u> | <u>93.9</u> | <u>72.8</u> | <u>79.8</u> |
| $-R + \mathcal{L}_{Ncut}$ | $-R + R_c + \mathcal{L}_{Ncut}$ | <b>97.7</b> | <b>94.3</b> | <b>77.8</b> | <b>82.2</b> |

the hyper-parameters  $\epsilon$  and  $\gamma$ . Then, we fix the hyper-parameters  $\epsilon$  and  $\gamma$  and conduct experiments with varying model parameters of the feature head  $f(\cdot; \Theta)$  and the cluster head  $g(\cdot; \Phi)$ , where we set the number of hidden layers as  $\{1, 2, 3, 4\}$  and the number of neurons in each hidden layer as  $\{512, 1024, 2048, 4096\}$ . Experimental results are shown in Fig. 5. Again, we can see that CgMCR<sup>2</sup> is not sensitive to the model size whenever it contains at least 1 hidden layers with 1024 hidden neurons.

**Table 7: Evaluation on Other Choices for Clustering Module on CIFAR-10, -20, -100 and TinyImageNet.** ‘s/it’ denotes seconds per training iteration.

| Clustering module | Train time (s/it) | CIFAR-10    |             | CIFAR-20    |             | CIFAR-100   |             | TinyImageNet |             |
|-------------------|-------------------|-------------|-------------|-------------|-------------|-------------|-------------|--------------|-------------|
|                   |                   | ACC         | NMI         | ACC         | NMI         | ACC         | NMI         | ACC          | NMI         |
| $g(\cdot; \Phi)$  | 0.02              | <b>97.7</b> | <b>94.3</b> | <b>68.8</b> | <b>74.0</b> | <b>78.3</b> | <b>82.5</b> | <b>72.9</b>  | <b>81.4</b> |
| $k$ -means        | 6.8               | 96.8        | 92.5        | 53.6        | 66.4        | 74.2        | 80.1        | 68.5         | 79.7        |
| Spectral          | 4.9               | 97.1        | 93.0        | 60.2        | 70.9        | <b>78.3</b> | <u>82.4</u> | <u>71.9</u>  | <u>80.8</u> |
| EnSC              | 11.2              | <u>97.3</u> | <u>93.4</u> | <u>64.4</u> | <u>71.3</u> | 76.7        | 81.7        | 68.7         | 80.6        |

**Evaluation on Other Choices for Clustering Module.** Recall that the core idea of our CgMCR<sup>2</sup> is to introduce a clustering module to produce the partition that is able to dynamically guide the feature learning via MCR<sup>2</sup>. Here, we replace the differential spectral clustering module (i.e.,  $g(\cdot; \Phi)$ ) of CgMCR<sup>2</sup> with the conventional  $k$ -means [32], Spectral Clustering via Ncut [41] and Elastic Net Subspace Clustering (EnSC) [47] individually, to form corresponding learning frameworks. Specifically, we use the conventional clustering methods to produce the cluster membership  $\mathbf{\Pi}$  in each mini-batch training iteration to guide the fine-tuning of  $\mathbf{Z}_\Theta$ . We report the time cost and the clustering accuracy in Table 7. We can read that, using the cluster membership produced by conventional clustering algorithms to guide the feature learning still yield very competitive clustering accuracy, in most cases. Nevertheless, the time cost of using conventional clustering algorithms is  $100\times$  expensive than that of using the cluster head  $g(\cdot; \Phi)$  due to the inability to utilize GPU acceleration, and can not infer the un-




Fig. 4: Effect of hyper-parameters on CIFAR-10 (left) and CIFAR-100 (right).




Fig. 5: Effect of model parameters on CIFAR-10 (left) and CIFAR-100 (right).

seen data points in test sets directly. This confirms the advantage of integrating a differential spectral clustering module  $g(\cdot; \Phi)$  in the pipeline of CgMCR<sup>2</sup>.

## 5 Conclusion

We have proposed a joint framework—graph Cut-guided Maximal Coding Rate Reduction—for learning the structured embeddings and clustering both in principled way. To be specific, a differential spectral clustering module is employed to learn the clustering membership to guide the task of learning structured embeddings. We have conducted extensive experiments on five benchmark datasets and shown state-of-the-art clustering performance. In addition, we have also provided a set of ablation studies to validate the effectiveness of each component.

**Acknowledgments.** This study was funded by the National Natural Science Foundation of China under Grant No. 61876022.

**Disclosure of Interests.** The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## References

1. Adaloglou, N., Michels, F., Kalisch, H., Kollmann, M.: Exploring the limits of deep image clustering using pretrained models. In: British Machine Vision Conference. pp. 297–299. BMVA Press (2023)
2. Bardes, A., Ponce, J., LeCun, Y.: Vicreg: Variance-invariance-covariance regularization for self-supervised learning. In: International Conference on Learning Representations (2022)
3. Bruna, J., Mallat, S.: Invariant scattering convolution networks. *IEEE Trans. Pattern Anal. Mach. Intell.* **35**(8), 1872–1886 (2013)
4. Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: *IEEE/CVF International Conference on Computer Vision*. pp. 9630–9640. IEEE (2021)
5. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.E.: A simple framework for contrastive learning of visual representations. In: *International Conference on Machine Learning*. vol. 119, pp. 1597–1607. PMLR (2020)
6. Chen, X., Fan, H., Girshick, R., He, K.: Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297* (2020)
7. Chen, Y., Li, C.G., You, C.: Stochastic sparse subspace clustering. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 4155–4164 (2020)
8. Chu, T., Tong, S., Ding, T., Dai, X., Haeffele, B.D., Vidal, R., Ma, Y.: Image clustering via the principle of rate reduction in the age of pretrained models. In: *International Conference on Learning Representations* (2024)
9. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: *IEEE conference on computer vision and pattern recognition*. pp. 248–255 (2009)
10. Devlin, J., Chang, M., Lee, K., Toutanova, K.: BERT: pre-training of deep bidirectional transformers for language understanding. In: Burstein, J., Doran, C., Solorio, T. (eds.) *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. pp. 4171–4186. Association for Computational Linguistics (2019)
11. Ding, T., Tong, S., Chan, K.H.R., Dai, X., Ma, Y., Haeffele, B.D.: Unsupervised manifold linearizing and clustering. In: *IEEE/CVF International Conference on Computer Vision*. pp. 5427–5438 (2023)
12. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houshy, N.: An image is worth 16x16 words: Transformers for image recognition at scale. In: *Proceedings of the International Conference on Learning Representations* (2021)
13. Grill, J., Strub, F., Altché, F., Tallec, C., Richemond, P.H., Buchatskaya, E., Dohersch, C., Pires, B.Á., Guo, Z., Azar, M.G., Piot, B., Kavukcuoglu, K., Munos, R., Valko, M.: Bootstrap your own latent - A new approach to self-supervised learning. In: *Advances in Neural Information Processing Systems* (2020)
14. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.B.: Masked autoencoders are scalable vision learners. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 15979–15988 (2022)
15. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.B.: Momentum contrast for unsupervised visual representation learning. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 9726–9735 (2020)

16. He, W., Zhang, S., Li, C.G., Qi, X., Xiao, R., Guo, J.: Neural normalized cut: A differential and generalizable approach for spectral clustering. Submitted to Pattern Recognition (2024)
17. Huang, Z., Chen, J., Zhang, J., Shan, H.: Learning representation for clustering via prototype scattering and positive sampling. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **45**(6), 7509–7524 (2022)
18. Jain, A.K., Murty, M.N., Flynn, P.J.: Data clustering: a review. *ACM Computing Surveys* **31**(3), 264–323 (1999)
19. Jang, E., Gu, S., Poole, B.: Categorical reparameterization with gumbel-softmax. In: 5th International Conference on Learning Representations (2017)
20. Khosla, A., Jayadevaprakash, N., Yao, B., Fei-Fei, L.: Novel dataset for fine-grained image categorization. In: *IEEE Conference on Computer Vision and Pattern Recognition* (2011)
21. Kingma, D., Ba, J.: Adam: A method for stochastic optimization. In: *Int. Conf. Learn. Represent.* (2014)
22. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. In: *International Conference on Learning Representations* (2014)
23. Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images. Technical Report TR-2009, University of Toronto, Toronto (2009)
24. Kuhn, H.W.: The hungarian method for the assignment problem. *Naval research logistics quarterly* **2**(1-2), 83–97 (1955)
25. LeCun, Y., Bottou, L., Bengio, Y., Haffner, P., et al.: Gradient-based learning applied to document recognition. *Proceedings of the IEEE* **86**(11), 2278–2324 (1998)
26. Li, Y., Hu, P., Liu, Z., Peng, D., Zhou, J.T., Peng, X.: Contrastive clustering. In: *Proceedings of the AAAI conference on artificial intelligence* (2021)
27. Li, Y., Hu, P., Peng, D., Lv, J., Fan, J., Peng, X.: Image clustering with external guidance. In: *Proceedings of the International Conference on Machine Learning* (2024)
28. Li, Z., Chen, Y., LeCun, Y., Sommer, F.T.: Neural manifold clustering and embedding. *arXiv preprint arXiv:2201.10000* (2022)
29. Lim, D., Vidal, R., Haeffele, B.D.: Doubly stochastic subspace clustering. *arXiv preprint arXiv:2011.14859* (2020)
30. Loshchilov, I., Hutter, F.: SGDR: stochastic gradient descent with warm restarts. In: *Int. Conf. Learn. Represent.* (2017)
31. von Luxburg, U.: A tutorial on spectral clustering. *Statistics and Computing* **17**(4), 395–416 (2007)
32. MacQueen, J.: Some methods for classification and analysis of multivariate observations. In: *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*. pp. 281–297 (1967)
33. Nene, S.A., Nayar, S.K., Murase, H.: Columbia object image library (coil-100). Tech. Rep. CU-CS-006-96, Department of Computer Science, Columbia University (February 1996)
34. Nilsback, M., Zisserman, A.: Automated flower classification over a large number of classes. In: *Sixth Indian Conference on Computer Vision, Graphics & Image Processing*. pp. 722–729. *IEEE Computer Society* (2008)
35. Niu, C., Shan, H., Wang, G.: SPICE: semantic pseudo-labeling for image clustering. *IEEE Transactions on Image Processing* **31**, 7264–7278 (2022)
36. Ntelemis, F., Jin, Y., Thomas, S.A.: Information maximization clustering via multi-view self-labelling. *Knowledge-Based Systems* **250**, 109042 (2022)



37. Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., et al.: Dinov2: Learning robust visual features without supervision. arXiv preprint arXiv:2304.07193 (2023)
38. Park, S., Han, S., Kim, S., Kim, D., Park, S., Hong, S., Cha, M.: Improving unsupervised image clustering with robust learning. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12278–12287 (2021)
39. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., Sutskever, I.: Learning transferable visual models from natural language supervision. In: International Conference on Machine Learning. vol. 139, pp. 8748–8763. PMLR (2021)
40. Rumelhart, D.E., Hinton, G.E., Williams, R.J., et al.: Learning internal representations by error propagation. *Parallel Distributed Processing* pp. 318–362 (1986)
41. Shi, J., Malik, J.: Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **22**(8), 888–905 (2000)
42. Souvenir, R., Pless, R.: Manifold clustering. In: IEEE/CVF International Conference on Computer Vision. pp. 648–653 (2005)
43. Tsai, T.W., Li, C., Zhu, J.: Mice: Mixture of contrastive experts for unsupervised image clustering. In: Proceedings of the International Conference on Learning Representations (2021)
44. Van Gansbeke, W., Vandenhende, S., Georgoulis, S., Proesmans, M., Van Gool, L.: Scan: Learning to classify images without labels. In: European Conference on Computer Vision. pp. 268–285 (2020)
45. Vidal, R.: Subspace clustering. *IEEE Signal Processing Magazine* **28**(3), 52–68 (March 2011)
46. Xiao, H., Rasul, K., Vollgraf, R.: Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. arXiv preprint arXiv: 1708.07747 (2019)
47. You, C., Li, C.G., Robinson, D.P., Vidal, R.: Oracle based active set algorithm for scalable elastic net subspace clustering. In: IEEE Conference on Computer Vision and Pattern Recognition. pp. 3928–3937 (2016)
48. Yu, Y., Chan, K.H.R., You, C., Song, C., Ma, Y.: Learning diverse and discriminative representations via the principle of maximal coding rate reduction. In: Advances in Neural Information Processing Systems. pp. 9422–9434 (2020)
49. Zhong, H., Wu, J., Chen, C., Huang, J., Deng, M., Nie, L., Lin, Z., Hua, X.S.: Graph contrastive clustering. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 9224–9233 (2021)