

Stable Cluster Discrimination for Deep Clustering

Qi Qian

Alibaba Group, Bellevue, WA 98004, USA

qi.qian@alibaba-inc.com

Abstract

Deep clustering can optimize representations of instances (i.e., representation learning) and explore the inherent data distribution (i.e., clustering) simultaneously, which demonstrates a superior performance over conventional clustering methods with given features. However, the coupled objective implies a trivial solution that all instances collapse to the uniform features. To tackle the challenge, a two-stage training strategy is developed for decoupling, where it introduces an additional pre-training stage for representation learning and then fine-tunes the obtained model for clustering. Meanwhile, one-stage methods are developed mainly for representation learning rather than clustering, where various constraints for cluster assignments are designed to avoid collapsing explicitly. Despite the success of these methods, an appropriate learning objective tailored for deep clustering has not been investigated sufficiently. In this work, we first show that the prevalent discrimination task in supervised learning is unstable for one-stage clustering due to the lack of ground-truth labels and positive instances for certain clusters in each mini-batch. To mitigate the issue, a novel stable cluster discrimination (SeCu) task is proposed and a new hardness-aware clustering criterion can be obtained accordingly. Moreover, a global entropy constraint for cluster assignments is studied with efficient optimization. Extensive experiments are conducted on benchmark data sets and ImageNet. SeCu achieves state-of-the-art performance on all of them, which demonstrates the effectiveness of one-stage deep clustering.

1. Introduction

Clustering is a fundamental task in unsupervised learning. Given features of instances, the unlabeled data set will be partitioned into multiple clusters, where instances from the same cluster are similar according to the measurement defined by a distance function. With fixed features, most of research efforts focus on studying appropriate distance functions and ingenious algorithms have been proposed by different measurements, e.g., k-means clustering [24], spec-

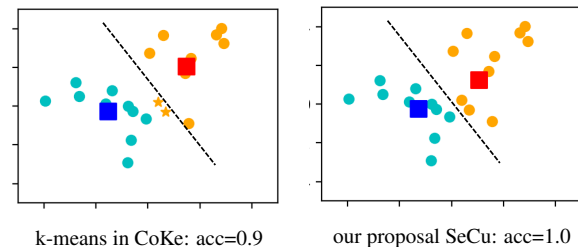


Figure 1: An illustration of the proposed method. 10 data points are randomly sampled from two different Gaussian distributions, respectively. Points with the same color are from the same distribution, while squares denote the corresponding cluster centers obtained by different methods. Star indicates the mis-classified data. Unlike k-means that assigns the uniform weight for different instances, our method considers the hardness of instances and is better for discrimination based clustering.

tral clustering [28], subspace clustering [13], etc.

With the development of deep neural networks, deep learning is capable of learning representations from raw materials and demonstrates the dominating performance on various supervised tasks [22]. Thereafter, it is also introduced to clustering recently [4, 11, 14, 19, 20, 23, 29, 30, 32]. Unlike the conventional clustering, representations of instances are learned with cluster assignments and centers simultaneously in deep clustering, which is more flexible to capture the data distribution. However, the coupled objective can result in a trivial solution that all instances collapse to the uniform features [4] and designing appropriate clustering criterion in the new scenario becomes challenging.

Many deep clustering algorithms [11, 14, 32] exploit a two-stage training strategy to decouple representation learning and clustering to avoid collapsing. The recent progress in unsupervised representation learning [2, 4, 7, 16, 17, 26] shows that informative features capturing the semantic similarity between instances without collapsing can be obtained by pre-training a large set of unlabeled data. Inspired by the observation, those methods leverage a pre-training stage and then focus on developing algorithms to optimize clus-

tering by fine-tuning a pre-trained model in the second stage. By refining the relations between nearest neighbors obtained from the pre-training stage, two-stage methods achieve a significantly better performance than one-stage ones without sufficient representation learning [19, 20].

However, the objective of pre-training can be inconsistent with that of clustering, which results in a sub-optimal performance for deep clustering. Note that different from supervised learning where the objective is explicitly defined by labels, that for unsupervised representation learning is arbitrary and various pretext tasks have been proposed, e.g., instance discrimination [7, 17], cluster discrimination [4, 26], masked modeling [2, 16], etc. Most of existing two-stage clustering methods adopt instance discrimination, i.e., SimCLR [7], for pre-training, whereas it aims to identify each instance as an individual class and the objective is different from clustering that aims to group instances from the same cluster. Moreover, SwAV [5] demonstrates that clustering itself is also an effective pretext task for representation learning. Therefore, we focus on facilitating one-stage deep clustering that optimizes representations and clustering simultaneously in this work.

Most of existing one-stage methods are proposed solely for representation learning. To tackle the collapsing issue, research efforts are mainly devoted to developing appropriate constraints for cluster assignments, especially for online deep clustering. For example, [1, 5] apply the balanced constraint that each cluster has the same number of instances for clustering. [26] further relaxes the balanced constraint to a lower-bound size constraint that limits the minimal size of clusters and demonstrates a more flexible assignment.

After obtaining cluster assignments as pseudo labels, representations and cluster centers can be optimized as in the supervised scenario. For example, [5, 6] learn the encoder network and cluster centers by solving a classification problem with the standard cross entropy loss. [26] has the same classification task for representation learning while adopting k-means for updating cluster centers. Although these methods achieve a satisfied performance on representation learning, a learning objective tailored for deep clustering has not attracted sufficient attentions.

In this work, we investigate the effective learning task for one-stage deep clustering. By analyzing the standard cross entropy loss for supervised learning, we find that it can be unstable for unsupervised learning. Concretely, the gradient for updating cluster centers consists of two ingredients: gradient from positive instances of the target cluster and that from irrelevant negative ones. However, with a limited mini-batch size in stochastic gradient descent (SGD), there can be no positive ones for a large proportion of clusters (e.g., 90% on ImageNet) at each iteration. Due to the lack of positive instances, the influence from negative ones is dominating. Unlike supervised learning where labels are fixed,

the cluster assignments can change during training in deep clustering. Therefore, the noise from the large variance of negative instances will be accumulated, which makes the optimization unstable.

To mitigate the problem, we propose a stable cluster discrimination (SeCu) task for deep clustering, which stops the gradient from negative instances for updating cluster centers in the cross entropy loss. Compared with k-means in [26], where positive instances have the uniform weight for updating centers, SeCu as a discrimination task considers the hardness of instances and a large weight will be assigned to hard instances for updating. Fig. 1 illustrates the hardness-aware clustering criterion implied by SeCu. Besides, we improve the cluster assignment by developing an entropy constraint that regularizes the entropy of assignments over the entire data set. Compared to the optimization with the size constraint [26], our method can reduce the number of variables and hyper-parameters, and thus make the learning more convenient. The main contributions of this work can be summarized as follows.

- A novel task is proposed for one-stage deep clustering. SeCu tailors the supervised cross entropy loss by eliminating the influence from the negative instances in learning cluster centers, which makes the training stable when ground-truth labels are unavailable.
- A global entropy constraint is exploited to balance the size of clusters and an efficient closed-form solution is developed for online assignment with the constraint.
- A simple framework with a single loss function and encoder is introduced for deep clustering in Eqn. 6. The proposed method is evaluated with the standard protocol on benchmark data sets and ImageNet [27]. The superior performance of SeCu confirms its effectiveness for deep clustering.

2. Related Work

In this section, we briefly review unsupervised representation learning and then focus on clustering.

2.1. Unsupervised Representation Learning

By leveraging massive unlabeled data, unsupervised representation learning can learn a pre-trained model that encodes the semantic information from data and helps downstream tasks with fine-tuning. Compared with supervised learning containing labels, one major challenge in unsupervised learning is defining appropriate positive pairs for representation learning. SimCLR [7] and MoCo [17] apply instance discrimination as the objective that generates positive pairs from diverse views of the same instance and considers other instances as the negative ones. Then, BYOL [15]

and SimSiam [8] demonstrate that with appropriate neural network architectures, the applicable models can be obtained with only positive pairs. Besides the work optimizing instance space, Barlow Twins [31] defines positive pairs in feature space and also learns effective pre-trained models. Recently, learning with masked modeling [2, 16] achieves success on the specific backbone of vision transformer [12] and shows the better fine-tuning performance than the generic objective on the ResNet [18].

2.2. Deep Clustering

Compared with instance discrimination, cluster discrimination aims to obtain positive pairs consisting of different instances for representation learning. Therefore, it has to learn representations and relations between instances simultaneously. DeepCluster [4] obtains membership of instances with an offline k-means and then optimizes representations alternately. Some work focuses on representation learning and introduces the balanced constraint that each cluster has the same number of instances for cluster assignments [1, 5] to mitigate the collapsing problem. Besides, DINO [6] applies an additional momentum encoder to further improve the representation. However, the constraint assumes a well balanced distribution, which is hard to capture the ground-truth data distribution. Thereafter, CoKe [26] relaxes the constraint to only lower-bound the minimal size of clusters to model the distribution more flexibly.

Besides representation learning, there are many works developed solely for better clustering [11, 14, 20, 23, 29, 30, 32]. To handle the collapsing issue, two-stage methods [11, 14, 32] leverage the representations from a pre-trained model to obtain nearest neighbors and fine-tune the model for clustering accordingly. In addition, the objective of pre-training can be included for clustering as multi-task learning [23]. The work closest to ours is CoKe [26] that optimizes representations and clustering simultaneously. Compared with the conventional k-means in CoKe, we propose a novel stable cluster discrimination objective tailored for one-stage deep clustering. Moreover, an entropy constraint is introduced to reduce the number of parameters while demonstrating the competitive performance.

3. Stable Cluster Discrimination

3.1. Cluster Discrimination

Given a data set with N images $\{x_i\}_{i=1}^N$, if the corresponding labels are accessible as $\{y_i\}_{i=1}^N$, the literature from distance metric learning [25] shows that appropriate representations of examples can be learned by optimizing a classification task as

$$\min_{\theta} \sum_{i=1}^N \ell(x_i, y_i; \theta)$$

where $\ell(\cdot)$ is the cross entropy loss with the normalized Softmax operator as suggested in [25]. θ denotes the parameters of the deep neural network.

For unsupervised learning that the label information is unavailable, the objective of cluster discrimination with K clusters can be written as

$$\min_{\theta_f, \{\mathbf{w}_j\}, y_i \in \Delta} \mathcal{L} = \sum_{i=1}^N \sum_{j=1}^K -y_{i,j} \log(p_{i,j}) \quad (1)$$

where y_i denotes the learnable label for x_i and $\Delta = \{y_i | \sum_{j=1}^K y_{i,j} = 1, \forall j, y_{i,j} \in \{0, 1\}\}$. With the normalized Softmax operator, the prediction $p_{i,j}$ can be computed

$$p_{i,j} = \frac{\exp(\mathbf{x}_i^\top \mathbf{w}_j / \lambda)}{\sum_{k=1}^K \exp(\mathbf{x}_i^\top \mathbf{w}_k / \lambda)} \quad (2)$$

where $\mathbf{x}_i = f(x_i)$ and $f(\cdot)$ denotes the encoder network. θ_f contains parameters of f . $\{\mathbf{w}_j\}_{j=1}^K$ consists of cluster centers. λ is the parameter of temperature and $\forall i, j, \|\mathbf{x}_i\|_2 = \|\mathbf{w}_j\|_2 = 1$ by normalization.

Compared with the supervised counterpart, the problem in Eqn. 1 has to optimize cluster assignments $\{y\}$, cluster centers $\{\mathbf{w}\}$, and representation encoder network f simultaneously. While most of existing work focus on optimizing $\{y\}$ with different constraints to avoid collapsing, this work investigates the learning objective for $\{\mathbf{w}\}$ and f , and a new clustering criterion is introduced accordingly.

3.2. Stable Loss for Optimization with Mini-batch

Unlike supervised learning, where labels are fixed for examples, cluster assignments $\{y\}$ in unsupervised learning are dynamic with the training of instance representations and cluster centers. Therefore, the original cross entropy loss becomes unstable for unsupervised discrimination. The problem can be elaborated by analyzing the updating criterion for cluster centers.

Letting y_i denote the label of \mathbf{x}_i , the gradient of \mathbf{w}_j from the standard cross entropy loss in Eqn. 1 can be computed as

$$\nabla_{\mathbf{w}_j} \mathcal{L} = \frac{1}{\lambda} \left(\sum_{i: y_i=j} (p_{i,j} - 1) \mathbf{x}_i + \sum_{k: y_k \neq j} p_{k,j} \mathbf{x}_k \right)$$

The former term is to pull the center \mathbf{w}_j to the assigned instances and the latter one is to push it away from instances of other clusters. However, the following Propositions show that this updating is unstable for deep clustering. The detailed proof can be found in the supplementary.

Proposition 1. *When sampling a mini-batch of b instances for K clusters, there are no positive instances for at least $K - b$ clusters.*

Proposition 2. Let Var_{pos} and Var_{neg} be the variance of sampled positive and negative instances, respectively. Assuming that each instance has unit norm and the norm of the cluster mean is a , we have $\text{Var}_{neg} = \mathcal{O}(\frac{1}{1-a^2})\text{Var}_{pos}$.

Remark Proposition 1 implies that with cross entropy loss, a large proportion of cluster centers will be updated only by negative instances. Proposition 2 indicates that the variance of sampled negative instances is much larger than that of positive ones when each cluster is compact as a approaching 1. Due to the small size of a mini-batch in training deep neural networks, the variance cannot be reduced sufficiently.

When $K = 10,000$ and $b = 1,024$ as in our multi-clustering settings for ImageNet, at least about 90% centers will only access a mini-batch of negative instances at each iteration. Without ground-truth labels, the bias in updating will be accumulated and mislead the learning process.

To mitigate the problem, we propose to eliminate the direction of negative instances from gradient for stable training as

$$\nabla_{\mathbf{w}_j} \mathcal{L} = \frac{1}{\lambda} \sum_{i:y_i=j} (p_{i,j} - 1) \mathbf{x}_i$$

The corresponding stable cluster discrimination loss becomes

$$\ell_{\text{SeCu}}(x_i, y_i) = -\log\left(\frac{\exp(\mathbf{x}_i^\top \mathbf{w}_{y_i} / \lambda)}{\exp(\mathbf{x}_i^\top \mathbf{w}_{y_i} / \lambda) + \sum_{k:k \neq y_i} \exp(\mathbf{x}_i^\top \tilde{\mathbf{w}}_k / \lambda)}\right) \quad (3)$$

where $\tilde{\mathbf{w}}_k$ denotes \mathbf{w}_k with the stop-gradient operator. Compared with the standard cross entropy loss, cluster centers in SeCu loss will be updated only by positive instances, which is more stable for deep clustering when optimizing with mini-batches.

In addition, the analysis in the following theorem demonstrates the novel hardness-aware clustering criterion for deep clustering implied by the proposed objective.

Theorem 1. When fixing $\{y_i\}$ and $\{\mathbf{x}_i\}$, let $\{\mathbf{w}^*\}$ be the optimal solution of the problem with loss function in Eqn. 3 and assume $\forall i, \|\mathbf{x}_i\|_2 = 1; \forall j, \|\mathbf{w}_j\|_2 = 1$, then we have

$$\mathbf{w}_j^* = \Pi_{\|\mathbf{w}\|_2=1} \left(\frac{\sum_{i:y_i=j} (1 - p_{i,j}) \mathbf{x}_i}{\sum_{i:y_i=j} 1 - p_{i,j}} \right) \quad (4)$$

where $\Pi_{\|\mathbf{w}\|_2=1}$ projects the vector to the unit norm.

Remark For k-means in CoKe [26], cluster centers will be averaged over all assigned instances with the uniform weight

$$\mathbf{w}_j = \Pi_{\|\mathbf{w}\|_2=1} \left(\frac{\sum_{i:y_i=j} \mathbf{x}_i}{\sum_i 1(y_i = j)} \right) \quad (5)$$

where $1(\cdot)$ is the indicator function. On the contrary, our proposal considers the weight of each instance by its hardness, i.e., p_{i,y_i} . By assigning a large weight to hard instances, the corresponding center can capture the distribution of hard instances better, which is illustrated in Fig. 1.

Our formulation also implies CoKe as a special case. Concretely, by increasing the temperature λ for updating \mathbf{w} to be infinite, we have $p_{i,j} = 1/K$ and Eqn. 4 will degenerate to Eqn. 5.

Finally, besides SGD, Eqn. 4 also suggests an alternative updating strategy for \mathbf{w} . Since $p_{i,j}$ is computed by \mathbf{w} , it may still require multiple iterations with Eqn. 4 for convergence. However, the representations of instances are also updated by training and we can keep a single update for centers at each iteration and improve centers along with the learning of representations. Compared with SGD for learning cluster centers, the strategy eliminates the learning rate for centers with a comparable performance as shown in the ablation study.

Given the learning objective SeCu, we will elaborate the proposed deep clustering method in next subsection.

3.3. SeCu for Deep Clustering

With the proposed loss function, the objective of stable cluster discrimination for deep clustering can be written as

$$\begin{aligned} \min_{\theta_f, \{\mathbf{w}_j\}, y_i \in \Delta} \sum_{i=1}^N \ell_{\text{SeCu}}(x_i, y_i) \\ \text{s.t. } h_m(Y) \geq b_m, \quad m = 1, \dots, M \end{aligned} \quad (6)$$

where $Y = [y_1, \dots, y_N]$, M denotes the number of constraints, and $h_m(\cdot)$ is the m -th constraint for cluster assignment. Considering that variables are entangled, the problem is solved alternately.

Update of θ_f First, when fixing $\{y^{t-1}\}$ and $\{\mathbf{w}^{t-1}\}$ from the last epoch, the sub-problem for representation learning in the t -th epoch becomes

$$\min_{\theta_f} \sum_{i=1}^N \ell_{\text{SeCu}}(x_i, y_i^{t-1})$$

where the SeCu loss degenerates to the standard cross entropy loss with fixed $\{\mathbf{w}\}$ and can be optimized with SGD. The one-hot label y^{t-1} is kept from the $(t-1)$ -th epoch, which makes the optimization consistent between adjacent epochs but the updating for representations may be delayed. To incorporate the information from the current epoch, we include two views of augmentations for the individual instance at each iteration, which is prevalent for representation learning [7, 17, 26].

Let x_i^1 and x_i^2 be two perturbed views of the original image with random augmentations, and the prediction is

$$p_{i,j}^s = \frac{\exp(\mathbf{x}_i^{s\top} \mathbf{w}_j^{t-1} / \lambda)}{\sum_j^K \exp(\mathbf{x}_i^{s\top} \mathbf{w}_j^{t-1} / \lambda)}; \quad s = \{1, 2\}$$

where $\mathbf{x}_i^s = f_t(x_i^s)$ is extracted by the current encoder. The soft label for each view can be obtained as [26]

$$y_i^1 = \tau y_i^{t-1} + (1 - \tau) p_i^2; \quad y_i^2 = \tau y_i^{t-1} + (1 - \tau) p_i^1 \quad (7)$$

The soft label contains the prediction from the other view of the same instance, which optimizes the consistency between different views at the same iteration. The problem with two-view optimization can be written as

$$\min_{\theta_f} \sum_{i=1}^N \ell_{\text{SeCu}}(x_i^1, y_i^1) + \ell_{\text{SeCu}}(x_i^2, y_i^2)$$

Update of y When fixing \mathbf{x}_i^t and $\{\mathbf{w}^t\}$, the cluster assignment can be updated by solving the problem

$$\begin{aligned} \min_{y_i \in \Delta} & - \sum_{i=1}^N \sum_{j=1}^K y_{i,j} \log(p_{i,j}) \\ \text{s.t.} & \quad h_m(Y) \geq b_m, \quad m = 1, \dots, M \end{aligned}$$

where $p_{i,j}$ is defined by \mathbf{x}_i^t and \mathbf{w}^t as in Eqn. 2. Without the constraints $\{h_m\}$, the objective implies a greedy solution that assigns each instance to the most related cluster. It can incur the trivial solution of collapsing, especially for online deep clustering, where each instance can only be accessed once in each epoch and cluster assignments cannot be refined with multiple iterations over the entire data. Therefore, many effective constraints are developed to mitigate the challenge.

First, we investigate the size constraint that is prevalent in existing methods [5, 26]. Following [26], the cluster size can be lower-bounded as

$$\begin{aligned} \min_{y_i \in \Delta} & - \sum_{i=1}^N \sum_{j=1}^K y_{i,j} \log(p_{i,j}) \\ \text{s.t.} & \quad \sum_i y_{i,j} \geq \gamma N / K, \quad j = 1, \dots, K \quad (8) \end{aligned}$$

where γ is the proportion to the average size and $\gamma = 1$ implies the balanced constraint. Let ρ_j denote dual variables for the j -th lower-bound constraint. The problem becomes

$$\max_{\rho: \rho \geq 0} \min_{y_i \in \Delta} - \sum_i \sum_j y_{i,j} \log(p_{i,j}) - \sum_j \rho_j \left(\sum_i y_{i,j} - \gamma N / K \right)$$

When a mini-batch of instances arrive, the cluster assignment for each instance can be obtained via a closed-form solution

$$y_{i,j}^t = \begin{cases} 1 & j = \arg \min_j - \log(p_{i,j}) - \rho_j \\ 0 & \text{o.w.} \end{cases} \quad (9)$$

Then, the dual variables will be updated by stochastic gradient ascent. More details can be found in [26].

The size constraint is capable of avoiding collapsing problem explicitly, but it introduces additional dual variables to help assignment and has to optimize them in training. To simplify the optimization, a global entropy constraint is investigated to balance the distribution over all clusters. Compared with the size constraint, the bound for the cluster size is implicit with the entropy constraint. Nevertheless, as we will demonstrate, it can help reduce the number of variables and hyper-parameters, which is more friendly for users.

Given the set of cluster assignments $\{y\}$, the entropy of the whole data set can be defined as

$$H(y) = - \sum_{j=1}^K \frac{\sum_i y_{i,j}}{N} \log\left(\frac{\sum_i y_{i,j}}{N}\right)$$

With the entropy as the regularization, the objective can be written as

$$\min_{y_i \in \Delta} - \sum_{i=1}^N \sum_{j=1}^K y_{i,j} \log(p_{i,j}) \quad \text{s.t.} \quad H(y) \geq \gamma$$

Compared with the problem consisting of K constraints in Eqn. 8, there is only a single constraint that controls the size of all clusters simultaneously. According to the dual theory [3], the problem is equivalent to

$$\min_{y_i \in \Delta} - \sum_{i=1}^N \sum_{j=1}^K y_{i,j} \log(p_{i,j}) - \alpha H(y) \quad (10)$$

Since one-hot labels of all instances are kept in memory, the optimal solution can be obtained by enumerating. When optimizing the label for x_i and fixing other labels, the closed-form solution for y_i is

$$y_{i,j}^t = \begin{cases} 1 & j = \arg \min_j - \log(p_{i,j}) - \alpha H(y^{t-1}, y_{i:j}) \\ 0 & \text{o.w.} \end{cases} \quad (11)$$

where $H(y^{t-1}, y_{i:j})$ replaces the previous label y_i of \mathbf{x}_i by j as $H(y^{t-1}, y_{i:j}) = H([y_1^{t-1}, \dots, y_{i-1}^{t-1}, y_{i+1}^{t-1}, \dots, y_n^{t-1}; y_{i,j} = 1])$. The convergence is guaranteed as in the following theorem.

Theorem 2. Let $\mathcal{L}(Y)$ denote the objective in Eqn. 10. If updating cluster assignments sequentially according to Eqn. 11, we have $\mathcal{L}(Y^t) \leq \mathcal{L}(Y^{t-1})$.

Proof. It is directly from the optimality of the closed-form solution in Eqn. 11. \square

Remark While some deep clustering methods apply the entropy regularization for learning, it is defined over a mini-batch of instances and optimized by SGD [14]. On the contrary, our method leverages the entropy of the entire data set that can capture the global distribution better. Moreover, the entropy in our method can be maximized with a closed-form solution, which is more efficient for learning.

With the two-view optimization, the updating becomes

$$j = \arg \min_j -(\log(p_{i,j}^1) + \log(p_{i,j}^2))/2 - \alpha H(y^{t-1}, y_{i:j})$$

Update of \mathbf{w} With fixed \mathbf{x}_i^t and the updated pseudo label $\{y^t\}$, cluster centers can be optimized by minimizing SeCu loss over two views of augmentations by SGD

$$\min_{\{\mathbf{w}_j\}} \sum_{i=1}^N \ell_{\text{SeCu}}(\mathbf{x}_i^1, y_i^t) + \ell_{\text{SeCu}}(\mathbf{x}_i^2, y_i^t)$$

Alternatively, \mathbf{w} can be updated by a closed-form solution

$$\mathbf{w}_j^{t:B} = \Pi_{\|\mathbf{w}\|_2=1} \left(\frac{\sum_{i:y_i=j}^B (1 - p_{i,j}) \mathbf{x}_i^t}{\sum_{i:y_i=j}^B 1 - p_{i,j}} \right)$$

where B denotes the total number of instances received in the current epoch.

The proposed method is easy to implement and the whole framework is illustrated in Alg. 1. Compared with the supervised discrimination, our main computational overhead is from cluster assignment, which is negligible due to the online strategy.

4. Experiments

The performance of the proposed method is evaluated on CIFAR-10 [21], CIFAR-100 [21], STL-10[10] and ImageNet [27]. To demonstrate the generalization performance of SeCu, we follow the setting in SCAN [14] that trains and evaluates models on the standard train/test split provided by the original data set. For STL-10 that contains an additional noisy unlabeled data set, we first learn the model on the training and noisy set, and then on the target training data as suggested by other methods [14, 32]. The standard metrics for clustering, including clustering accuracy (ACC), normalized mutual information (NMI) and adjusted rand index (ARI) are reported for evaluation. For a fair comparison, we follow the common practice to configure the architecture and training protocol for our method. Concretely, ResNet-50 [18] is applied on ImageNet, while ResNet-18 is adopted for other data sets. The key settings for ResNet-18 are elaborated as follows, while those of ResNet-50 follow [26]. More details can be found in the supplementary.

Algorithm 1 Pseudo-code of Stable Cluster Discrimination (SeCu) for One-stage Deep Clustering.

```
# f: encoder network
# w: cluster centers
# w_p: centers from the last epoch
# y: list of pseudo one-hot labels (Nx1)
# tau: weight for labels from last epoch
# lambda: temperature

# keep last cluster centers before each epoch
w_p = w.detach()
# train one epoch
for x in loader: # load a minibatch with b samples
    x_1, x_2 = f(aug(x)), f(aug(x)) # two random views
    y_x = y(x_id) # retrieve label from last epoch
    # compute prediction for each view
    p_1 = softmax(x_1 @ w_p / lambda)
    p_2 = softmax(x_2 @ w_p / lambda)
    # obtain soft label for discrimination
    y_1 = tau * y_x + (1-tau) * p_2
    y_2 = tau * y_x + (1-tau) * p_1
    # loss for representation learning
    loss_x = (SeCu(p_1, y_1) + SeCu(p_2, y_2)) / 2
    # update prediction for clustering
    p_1 = softmax(x_1.detach() @ w / lambda)
    p_2 = softmax(x_2.detach() @ w / lambda)
    # update cluster assignments with constraints
    y(x_id) = y_x = cluster_assign(p_1, p_2)
    # loss for cluster centers
    loss_w = (SeCu(p_1, y_x) + SeCu(p_2, y_x)) / 2
    # update encoder and cluster centers
    loss = loss_x + loss_w
    loss.backward()
```

Architecture Besides backbone network, a 2-layer MLP head is attached and the output dimension is 128. The similar architecture is adopted in [26, 32]. With the MLP head, the total number of parameters is almost the same as the original ResNet-18. The discussion about the MLP head can be found in Sec. 4.1.1. After the MLP head, the learned representation will be classified by a fully-connected (FC) layer that encodes cluster centers. The size of the FC layer is $128 \times K$. As a clustering method, we let K be the number of ground-truth classes as in [14] and have the direct prediction from the model as cluster assignments for evaluation. Following [1, 14], 10 different classification heads are applied as multi-clustering, which benefits the target clustering. To avoid the problem of selecting an appropriate head for evaluation when 10 heads have the same K , we have a varying cK as the number of clusters for each head, where $c \in \{1, \dots, 10\}$. The prediction from the head with $c = 1$ is reported for comparison.

Optimization To reduce efforts of parameter tuning and reuse the parameter of [26], we have $\mathbf{x}_i^\top \mathbf{w}_j$ in lie of $\log(p_{i,j})$ in Eqn. 9 and 11 for assignment. Searching the parameter for the latter one can achieve the similar performance. Before training the encoder network, one epoch is adopted to initialize cluster assignments and centers. The encoder network is optimized by SGD with a batch size of 128. The momentum and learning rate are set to be 0.9 and 0.2, respectively. Moreover, the first 10 epochs are applied for warm-up and the cosine decay for learning rate

is adopted subsequently. For CIFAR-10, the model is optimized by 400 epochs and cluster centers are learned by SGD with a constant learning rate of 1.2. For other challenging data sets, epochs and learning rate for centers are 800 and 0.8, respectively. τ for the soft label as in Eqn. 7 and the temperature λ are fixed as 0.2 and 0.05. For the size constraint, the lower-bound parameter γ and the learning rate of dual variables η_ρ are set to be 0.9 and 0.1, respectively. The only parameter α for the proposed global entropy constraint is set to be $6N/50$, where N is the total number of instances. The weight is scaled according to the ablation study on CIFAR-10 as illustrated in Sec. 4.1.4.

Augmentation Augmentation is essential for the success of unsupervised representation learning [7]. For a fair comparison, we apply the prevalent settings in existing work [7, 15]. Concretely, random crop, color jitter, random gray scale, Gaussian blur, solarize and random horizontal flip are included as augmentations. Considering that the resolution of images in CIFAR and STL-10 is much smaller than those in ImageNet, we let the minimal scale of random crop be 0.3, 0.2, and 0.2 for CIFAR-10, CIFAR-100, and STL-10, respectively. That for ImageNet is kept as 0.08. Other parameters are the same for all data sets and many recent work share the similar augmentation [9, 15, 26, 32].

All experiments on small data sets are implemented on a single V100, while 8 GPUs are sufficient for ImageNet.

4.1. Ablation Study

To illustrate the behavior of the proposed method, we study the effect of parameters in SeCu on CIFAR-10 with the size constraint.

4.1.1 Effect of MLP Head

The MLP head shows better generalization than the FC layer in representation learning [7], but its application in deep clustering has not been evaluated systematically. We compare different architectures of MLP head and summarize results in Table 1. Besides the standard MLP layer for projection, there can be another MLP layer attached to the projection MLP head, which is referred as the prediction head [15].

First, with the standard architecture denoted by 0 for projection and prediction, SeCu is already 3.7% better than SCAN on ACC and it confirms that the one-stage deep clustering strategy can learn more consistent features than two-stage methods. If including a single layer for SeCu, the performance can be improved by 2.1%, which shows that adding a layer with the reduced dimension in head is helpful for deep clustering. Moreover, due to the dimension reduction from the projection layer, the size of cluster centers decreases and the total number of parameters in the model

#Proj	#Pred	#Para	ACC	NMI	ARI
0	0	11.45M	85.5	76.2	73.1
1	0	11.30M	87.6	78.7	76.7
2	0	11.57M	88.1	79.4	77.6
3	0	11.83M	88.1	79.4	77.6
3	2	11.97M	88.2	79.6	77.8

Table 1: Comparison of MLP head on CIFAR-10. “#Proj” and “#Pred” indicate the number of layers in MLP for projection head and prediction head, respectively. ACC (%), NMI (%) and ARI (%) are reported. Baseline ACC of SCAN [14] is 81.8%.

is even smaller than that of the original network. When increasing the number of layers and having a 2-layer MLP head, ACC can be further improved by 0.5% and the model size is only slightly increased. After that, there is no significant gain by applying more complicated MLP head. It is because that the 2-layer MLP head is sufficient for small data sets and we will fix 2-layer MLP as the head for the rest experiments except ImageNet, where 3-layer projection and 2-layer prediction are applied as in [26].

4.1.2 Effect of Negative Instances

After the ablation experiments for the architecture, we study the effect of the proposed stable cluster discrimination loss in Eqn. 3. Table 2 compares it to the standard cross entropy loss (CE) that keeps the gradient from negative instances.

Loss	#Max	#Min	ACC	NMI	ARI
CE	6,486	4,501	51.6	50.2	35.0
SeCu	5,190	4,556	88.1	79.4	77.6

Table 2: Comparison of loss functions on CIFAR-10. “#Max” and “#Min” indicate the maximal and minimal size of obtained clusters.

First, we can observe that the proposed SeCu loss outperforms the standard cross entropy loss by a large margin of 36.5% on ACC, which confirms the benefit of stable training. Then, by investigating the distribution of learned clusters, we find that the minimal size of clusters obtained from different loss functions is almost the same due to the lower-bound constraint, while the maximal size varies. The largest cluster obtained by SeCu contains 5,190 instances, which is close to the ground-truth distribution of CIFAR-10. On the contrary, the dominating cluster found by CE loss has 6,486 instances that is significantly different from the uniform distribution in CIFAR-10. It is because that the negative instances in CE loss can perturb the learning of centers and result in sub-optimal cluster assignments.

4.1.3 Effect of Updating Criterion for Cluster Centers

The cluster centers in SeCu can be optimized by SGD or updated by a closed-form solution as in Eqn. 4. We compare these two strategies and summarize the performance in Table 3.

Methods	ACC	NMI	ARI
SeCu-CF	87.4±0.5	78.5±0.6	76.4±0.9
SeCu-SGD	88.1±0.3	79.3±0.4	77.5±0.4

Table 3: Comparison of optimization for cluster centers on CIFAR-10. Mean±std over 8 trails are reported.

Let “SeCu-CF” and “SeCu-SGD” be the closed-form update and the SGD optimization, respectively. We can observe that SeCu-CF has the similar performance to SeCu-SGD, which shows the effectiveness of the closed-form updating. Moreover, SeCu-SGD has a smaller standard deviation than SeCu-CF due to the smoothness from the momentum in SGD. Compared with SeCu-CF, SeCu-SGD has an additional parameter of the learning rate for centers. Hence, SeCu-CF can be applied to evaluate the preliminary performance of SeCu with less tuning efforts.

4.1.4 Effect of α in Entropy Constraint

Besides the size constraint, the proposed entropy constraint is also effective for balancing cluster assignments. Table 4 demonstrates the results with different α .

α	#Max	#Min	ACC	NMI	ARI
20,000	5,056	4,913	85.3	76.3	73.1
8,000	5,115	4,658	87.9	79.0	77.1
6,000	5,120	4,579	88.0	79.4	77.4
4,000	5,376	4,324	84.9	75.4	72.2
100	28,575	0	19.9	38.7	17.0
0	50,000	0	10.0	0.0	0.0

Table 4: Comparison of α in Eqn. 10 on CIFAR-10.

First, a large α will regularize the distribution to be uniform and result in a sub-optimal performance. By decreasing the weight, the assignment becomes more flexible and a desired performance can be observed when $\alpha = 6,000$. However, a small α will lead to an imbalanced distribution, which is inappropriate for the balanced data sets such as CIFAR and STL. $\alpha = 0$ discards the constraint and incurs collapsing. Evidently, the entropy constraint can balance the size of clusters effectively

Since the entropy is defined on the whole data set and CIFAR-10 contains 50,000 instances for training, α will be scaled as $6N/50$ for different data sets, where N denotes the total number of instances in the training set. More ablation experiments can be found in the supplementary.

4.2. Comparison with State-of-the-Art

After the ablation study, experiments are conducted on benchmark data sets to compare with state-of-the-art methods. The results averaged over 8 trials and the best performance among them are reported for our method. SeCu with size constraint and entropy constraint are denoted as “SeCu-Size” and “SeCu-Entropy”, respectively. Table 5 shows the performance of different methods, where two-stage methods with a pre-training phase are marked for “Two-stage”.

First, as a one-stage method, SeCu outperforms two-stage methods by a margin of about 3% over all metrics on CIFAR-10 and CIFAR-100-20. It demonstrates that learning representations and clustering simultaneously is essential for deep clustering. The two-stage method NNM[11] shows the comparable performance to SeCu on STL-10. It is because that STL-10 only constrains 5,000 target instances for training and both of NNM and SeCu already achieve the accuracy of supervised learning. Second, compared with the one-stage method without sufficient representation learning, i.e., PICA, SeCu demonstrates a better performance by about 20% on CIFAR-10 and CIFAR-100-20, and 10% on STL-10, which shows the importance of optimizing representations for instances. Third, with the same size constraint, SeCu surpasses CoKe by 4.3% for ARI on CIFAR-10, which confirms the effectiveness of the proposed learning objective. Finally, compared with size constraint in SeCu-Size, SeCu-Entropy demonstrates a competitive performance over all data sets. Since the entropy constraint is optimized by enumeration without introducing any auxiliary variable, it can be an substitute to the size constraint for cluster assignments.

After the clustering phase, self-labeling [14] is an effective strategy to further improve the performance. We try it for SeCu-Size with the setting in [14] and report the result in Table 5. Interestingly, self-labeling is also effective for SeCu and helps approach the supervised accuracy on CIFAR-10. As a one-stage method, SeCu has the strong augmentation for representation learning, which may introduce additional noise for clustering. Self-labeling has the weak augmentation for fine-tuning and refines clustering.

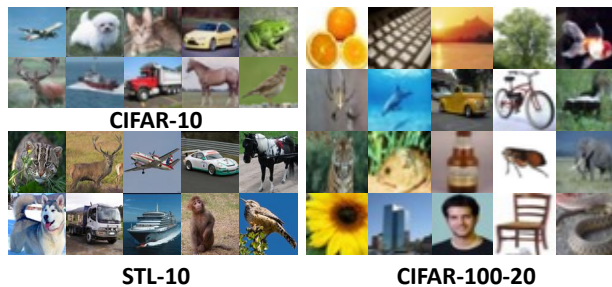


Figure 2: Illustration of exemplar images.

Methods	Two-stage	CIFAR10			CIFAR100-20			STL10		
		ACC	NMI	ARI	ACC	NMI	ARI	ACC	NMI	ARI
Supervised		93.8	86.2	87.0	80.0	68.0	63.2	80.6	65.9	63.1
DeepCluster [4]		37.4	-	-	18.9	-	-	33.4	-	-
IIC [20]		61.7	51.1	41.1	25.7	22.5	11.7	59.6	49.6	39.7
PICA [19] (best)		69.6	59.1	51.2	33.7	31.0	17.1	71.3	61.1	53.1
Pretext [7]+k-means	✓	65.9	59.8	50.9	39.5	40.2	23.9	65.8	60.4	50.6
SCAN [14] (mean)	✓	81.8	71.2	66.5	42.2	44.1	26.7	75.5	65.4	59.0
NNM [11]	✓	84.3	74.8	70.9	47.7	48.4	31.6	80.8	69.4	65.0
GCC [32]	✓	85.6	76.4	72.8	47.2	47.2	30.5	78.8	68.4	63.1
CoKe [26] (mean)		85.7	76.6	73.2	49.7	49.1	33.5	-	-	-
SeCu-Size (mean)		88.1	79.3	77.5	50.0	50.7	35.0	80.2	69.4	63.9
SeCu-Size (best)		88.5	79.9	78.2	51.6	51.6	36.0	81.4	70.7	65.7
SeCu-Entropy (mean)		88.0	79.2	77.2	49.9	50.6	34.1	79.5	68.7	63.0
SeCu-Entropy (best)		88.2	79.7	77.7	51.2	51.4	34.9	80.5	69.9	64.4
<i>with self-labeling:</i>										
SCAN [14]	✓	88.3	79.7	77.2	50.7	48.6	33.3	80.9	69.8	64.6
GCC [32]	✓	90.1	-	-	52.3	-	-	83.3	-	-
SeCu		93.0	86.1	85.7	55.2	55.1	39.7	83.6	73.3	69.3

Table 5: Comparison of clustering methods on benchmark data sets. “Two-stage” denotes an additional pre-training stage.

Exemplar Images Finally, we show exemplar images from clusters obtained by SeCu for different data sets. Fig. 2 illustrates the images that are close to cluster centers in CIFAR-10, CIFAR-100-20, and STL-10, respectively. Evidently, the proposed method can obtain the ground-truth class structure for CIFAR-10 and STL-10. Moreover, with a large intra-class variance, SeCu can still observe appropriate clusters for CIFAR-100-20.

Considering that labels of super-classes in CIFAR-100-20 contain vague concepts such as “vehicles 1” and “vehicles 2”, we include results on the target 100 classes in Table 6, which may benefit future research. The experiment settings for CIFAR-100 are identical to that for CIFAR-100-20, except the learning rate of centers, which is reduced to 0.01.

Methods	ACC	NMI	ARI
SeCu-Size (mean)	46.4	61.7	32.3
SeCu-Size (best)	47.1	61.9	32.6
SeCu-Entropy (mean)	45.3	60.7	31.5
SeCu-Entropy (best)	46.7	61.0	32.2
SeCu [†]	51.3	65.2	37.1

Table 6: Results of SeCu on CIFAR-100. [†] denotes a method with self-labeling.

4.3. Comparison on ImageNet

Now we evaluate SeCu on the challenging ImageNet data set in Table 7, where the cost of many two-stage methods becomes intractable and only available baselines are

included. First, SeCu achieves 51.4% ACC on ImageNet, which is better than SCAN with self-labeling by more than 11%. With the same 10 clustering heads and size constraint, SeCu is still 3.8% better than CoKe on ACC. Compared with the objective in CoKe, SeCu is tailored for deep clustering and the comparison confirms the efficacy of our method. With self-labeling, SeCu shows state-of-the-art performance of 53.5% ACC on ImageNet. It demonstrates that our method is applicable for large-scale data sets.

Methods	ACC	NMI	ARI
SCAN [†]	39.9	72.0	27.5
CoKe [‡]	47.6	76.2	35.6
SeCu	51.4	78.0	39.7
SeCu [†]	53.5	79.4	41.9

Table 7: Comparison on ImageNet. [†] denotes a method with self-labeling. [‡] is our reproduction with 10 clustering heads.

5. Conclusion

In this work, a novel stable cluster discrimination task that considers the hardness of instances is proposed for one-stage deep clustering. To avoid collapsing in representations, a global entropy constraint with theoretical guarantee is investigated. Finally, the comparison with state-of-the-art methods confirms the effectiveness of our proposal. While recent work shows that leveraging nearest neighbors improves deep clustering [11, 14, 32], incorporating such information in SeCu can be our future work.

References

- [1] Yuki Markus Asano, Christian Rupprecht, and Andrea Vedaldi. Self-labelling via simultaneous clustering and representation learning. In *ICLR*, 2020.
- [2] Hangbo Bao, Li Dong, and Furu Wei. Beit: BERT pre-training of image transformers. *CoRR*, abs/2106.08254, 2021.
- [3] Stephen Boyd, Stephen P Boyd, and Lieven Vandenbergh. *Convex optimization*. Cambridge university press, 2004.
- [4] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In *ECCV*, 2018.
- [5] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. In *NeurIPS*, 2020.
- [6] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, pages 9630–9640. IEEE, 2021.
- [7] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, volume 119, pages 1597–1607, 2020.
- [8] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *CVPR*, pages 15750–15758. Computer Vision Foundation / IEEE, 2021.
- [9] Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised vision transformers. *CoRR*, abs/2104.02057, 2021.
- [10] Adam Coates, Andrew Y. Ng, and Honglak Lee. An analysis of single-layer networks in unsupervised feature learning. In *AISTATS*, volume 15, pages 215–223. JMLR.org, 2011.
- [11] Zhiyuan Dang, Cheng Deng, Xu Yang, Kun Wei, and Heng Huang. Nearest neighbor matching for deep clustering. In *CVPR*, pages 13693–13702. Computer Vision Foundation / IEEE, 2021.
- [12] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*. OpenReview.net, 2021.
- [13] Ehsan Elhamifar and René Vidal. Sparse subspace clustering: Algorithm, theory, and applications. *IEEE Trans. Pattern Anal. Mach. Intell.*, 35(11):2765–2781, 2013.
- [14] Wouter Van Gansbeke, Simon Vandenhende, Stamatios Georgoulis, Marc Proesmans, and Luc Van Gool. SCAN: learning to classify images without labels. In *ECCV*, volume 12355, pages 268–285. Springer, 2020.
- [15] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre H. Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Ávila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, Bilal Piot, Koray Kavukcuoglu, Rémi Munos, and Michal Valko. Bootstrap your own latent - A new approach to self-supervised learning. In *NeurIPS*, 2020.
- [16] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross B. Girshick. Masked autoencoders are scalable vision learners. *CoRR*, abs/2111.06377, 2021.
- [17] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross B. Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, pages 9726–9735, 2020.
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.
- [19] Jiabo Huang, Shaogang Gong, and Xiatian Zhu. Deep semantic clustering by partition confidence maximisation. In *CVPR*, pages 8846–8855. Computer Vision Foundation / IEEE, 2020.
- [20] Xu Ji, Andrea Vedaldi, and João F. Henriques. Invariant information clustering for unsupervised image classification and segmentation. In *ICCV*, pages 9864–9873. IEEE, 2019.
- [21] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. 2009.
- [22] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. In *NeurIPS*, pages 1106–1114, 2012.
- [23] Yunfan Li, Peng Hu, Jerry Zitao Liu, Dezhong Peng, Joey Tianyi Zhou, and Xi Peng. Contrastive clustering. In *AAAI*, pages 8547–8555. AAAI Press, 2021.
- [24] Stuart P. Lloyd. Least squares quantization in PCM. *IEEE Trans. Inf. Theory*, 28(2):129–136, 1982.
- [25] Qi Qian, Lei Shang, Baigui Sun, Juhua Hu, Hao Li, and Rong Jin. Softtriple loss: Deep metric learning without triplet sampling. In *ICCV*, pages 6449–6457, 2019.
- [26] Qi Qian, Yuanhong Xu, Juhua Hu, Hao Li, and Rong Jin. Unsupervised visual representation learning by online constrained k-means. In *CVPR*, pages 16619–16628. IEEE, 2022.
- [27] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael S. Bernstein, Alexander C. Berg, and Fei-Fei Li. Imagenet large scale visual recognition challenge. *Int. J. Comput. Vis.*, 115(3):211–252, 2015.
- [28] Ulrike von Luxburg. A tutorial on spectral clustering. *Stat. Comput.*, 17(4):395–416, 2007.
- [29] Junyuan Xie, Ross B. Girshick, and Ali Farhadi. Unsupervised deep embedding for clustering analysis. In *ICML*, volume 48, pages 478–487. JMLR.org, 2016.
- [30] Bo Yang, Xiao Fu, Nicholas D. Sidiropoulos, and Mingyi Hong. Towards k-means-friendly spaces: Simultaneous deep learning and clustering. In *ICML*, 2017.
- [31] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction. In *ICML*, volume 139, pages 12310–12320. PMLR, 2021.
- [32] Huasong Zhong, Jianlong Wu, Chong Chen, Jianqiang Huang, Minghua Deng, Liqiang Nie, Zhouchen Lin, and Xian-Sheng Hua. Graph contrastive clustering. In *ICCV*, pages 9204–9213. IEEE, 2021.