



Twin Contrastive Learning for Online Clustering

Yunfan Li¹ · Mouxing Yang¹ · Dezhong Peng¹ · Taihao Li² · Jiantao Huang² · Xi Peng¹

Received: 31 August 2021 / Accepted: 15 June 2022

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2022

Abstract

This paper proposes to perform online clustering by conducting twin contrastive learning (TCL) at the instance and cluster level. Specifically, we find that when the data is projected into a feature space with a dimensionality of the target cluster number, the rows and columns of its feature matrix correspond to the instance and cluster representation, respectively. Based on the observation, for a given dataset, the proposed TCL first constructs positive and negative pairs through data augmentations. Thereafter, in the row and column space of the feature matrix, instance- and cluster-level contrastive learning are respectively conducted by pulling together positive pairs while pushing apart the negatives. To alleviate the influence of intrinsic false-negative pairs and rectify cluster assignments, we adopt a confidence-based criterion to select pseudo-labels for boosting both the instance- and cluster-level contrastive learning. As a result, the clustering performance is further improved. Besides the elegant idea of twin contrastive learning, another advantage of TCL is that it could independently predict the cluster assignment for each instance, thus effortlessly fitting online scenarios. Extensive experiments on six widely-used image and text benchmarks demonstrate the effectiveness of TCL. The code is released on <https://pengxi.me>.

Keywords Deep clustering · Online clustering · Unsupervised learning · Contrastive learning

1 Introduction

Clustering is one of the most fundamental tasks in machine learning and data mining. It aims to group data into different clusters without label information, such that the within-cluster data come from the same class or share similar semantics. Besides facilitating general representation learn-

ing (Caron et al., 2018, 2020), clustering is also helpful in a variety of real-world applications, such as face recognition (Shen et al., 2021), medical analysis (Thanh et al., 2017), and gene sequencing (Kiselev et al., 2019).

During past years, most clustering methods (Chen & Lerman, 2009; Nie et al., 2011; Zhang et al., 2012; Liu et al., 2016, 2017; Nie et al., 2019; Tang et al., 2019) mainly focus on developing different similarity metrics and clustering strategies. Though grounded in theory, their performance is limited by the adopted shallow models. Recently, deep clustering (Ghasedi Dizaji et al., 2017; Li et al., 2020) has shown promising results on various benchmarks by extracting representative features to facilitate downstream clustering. Early deep clustering methods (Peng et al., 2016; Yang et al., 2016; Caron et al., 2018; Xie et al., 2016) iteratively perform representation learning and clustering to bootstrap each other. However, this kind of method usually needs the entire dataset to perform offline clustering, which is less attractive for large-scale data and even impractical for streaming data. Luckily, the offline limitation could be solved by the idea of “label as representation” (Peng et al., 2015, 2019). By directly and independently predicting cluster assignment for each instance, large-scale and online clustering could be achieved (Hu et al., 2017; Ji et al., 2019; Huang et al., 2020).

Communicated by Frederic Jurie.

✉ Xi Peng
pengx.gm@gmail.com

Yunfan Li
yunfanli.gm@gmail.com

Mouxing Yang
yangmouxing@gmail.com

Dezhong Peng
pengdz@scu.edu.cn

Taihao Li
lith@zhejianglab.com

Jiantao Huang
jthuang@zhejianglab.edu.cn

¹ College of Computer Science, Sichuan University, Chengdu, China

² Zhejiang Lab, Hangzhou, China

Very recently, the rapid growth of contrastive learning (Chen et al., 2020a; Zbontar et al., 2021) significantly improves the performance of unsupervised representation learning. Motivated by their successes, some contrastive learning based clustering methods (Han et al., 2020; Li et al., 2021b; Niu & Wang, 2021; Van Gansbeke et al., 2020) are proposed, which achieve state-of-the-art results.

In this work, we propose an end-to-end online deep clustering method by conducting twin contrastive learning (TCL) based on the observation shown in Fig. 1. In brief, the rows and columns of the feature matrix correspond to the instance and cluster representations, respectively. Under this observation, TCL conducts contrastive learning in the row and column space of the feature matrix to jointly learn the instance and cluster representation. Specifically, TCL first constructs contrastive pairs through data augmentations. Different from most existing contrastive learning methods that use weak augmentations proposed in SimCLR (Chen et al., 2020a), we provide a new effective augmentation strategy by mixing weak and strong transformations. With the constructed pairs, TCL performs contrastive learning at both the instance and cluster level. The instance-level con-

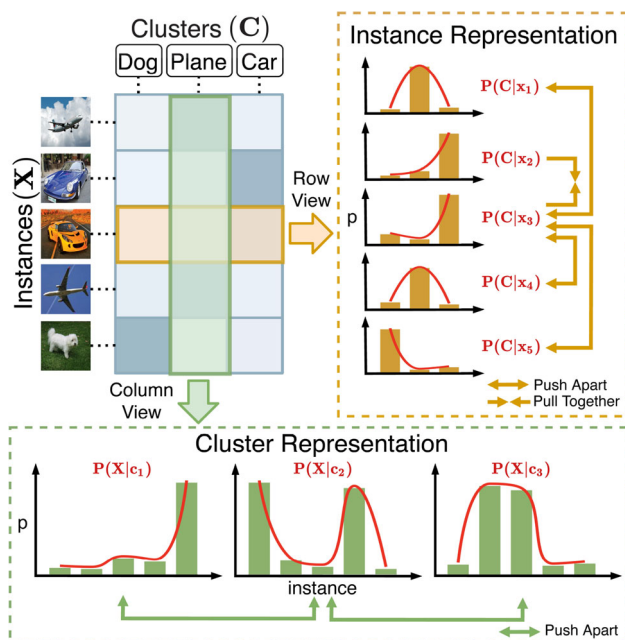


Fig. 1 Our key observation and basic idea. By projecting data into a feature space with the dimensionality of cluster number, the element in the i -th row and k -th column of the feature matrix represents the probability of instance i belonging to cluster k . Namely, rows correspond to the cluster assignment probabilities, which are special representations of instances. More interestingly, if we look at the feature matrix from the column view, each column actually corresponds to the cluster distribution over the data, which could be seen as a special representation of the cluster. As a result, the instance- and cluster-level representation learning (e.g., contrastive learning) could be conducted in the row and column space, respectively

trastive learning aims to pull within-class instances together while pushing between-class instances apart. And the cluster-level contrastive learning aims to distinguish distributions of different clusters while attracting distributions of the same cluster under different augmentations. To relieve the influence of intrinsic false-negative pairs and rectify cluster assignments, we progressively select the confident predictions (i.e., those with cluster assignment probability close to one-hot) to fine-tune the twin contrastive learning. Such a fine-tuning strategy is based on the observation that the predictions with high confidence are more likely to be correct and thus could be used as pseudo labels. Once the model converges, it could independently make cluster assignments for each instance in an end-to-end manner to achieve clustering. The major contributions of this work are summarized as follows:

- We reveal that the rows and columns of the feature matrix intrinsically correspond to the instance and cluster representations. On top of that, we propose TCL that achieves clustering by simultaneously conducting contrastive learning at the instance and cluster level;
- We provide a new data augmentation strategy by mixing weak and strong transformations, which naturally fits our TCL framework and is proved to be effective for both image and text data in our experiments;
- To alleviate the influence of intrinsic false negatives and rectify cluster assignments, we adopt a confidence-based criterion to generate pseudo-labels for fine-tuning both the instance- and cluster-level contrastive learning. Experiments show that such a fine-tuning strategy could further boost the clustering performance;
- The proposed TCL clusters data in an end-to-end and online manner, which only needs batch-wise optimization and thus could handle large-scale datasets. Moreover, TCL could handle streaming data since it could timely make cluster assignments for new coming data without accessing the whole dataset.

2 Related Work

In this section, we give a brief review on contrastive learning and deep clustering, followed by a discussion on the connection between these two topics.

2.1 Contrastive Learning

Recently, the contrastive learning paradigm shows its power in unsupervised representation learning (Caron et al., 2020; Chen et al., 2020a; Chen & He, 2020; Grill et al., 2020; He et al., 2020; Hu et al., 2020; Zbontar et al., 2021). It first constructs positive and negative pairs for each instance and then

projects them into a subspace to maximize the similarities of positive pairs and minimize those of the negatives (Hadsell et al., 2006). The most straightforward solution is to use labels to guide the pair construction (Khosla et al., 2020). However, in the unsupervised setting, other strategies are needed to construct and utilize contrastive pairs. For example, SimCLR (Chen et al., 2020a) constructs positive and negative pairs through augmentations within mini-batch. MoCo (He et al., 2020) recasts contrastive learning as a dictionary look-up task by building a dynamic dictionary with a queue and a moving-averaged encoder. To avoid the efforts in building negative pairs, BYOL (Grill et al., 2020) and SimSiam (Chen & He, 2020) replace negative pairs with an online predictor that prevents the network from collapsing into trivial solutions. As an alternative, AdCo (Hu et al., 2020) directly learns negative samples in an adversarial manner. Lately, Barlow Twins (Zbontar et al., 2021) performs contrastive learning from a redundancy-reduction perspective and achieves comparable results.

There are two major differences between our work and these contrastive learning methods. First, our method concurrently conducts row- and column-wise contrastive learning at both the instance and cluster level while most existing methods solely perform row-wise contrastive learning at the instance level. Such an elegant idea is based on our observation that rows and columns of the feature matrix correspond to the instance and cluster representations respectively. Second, the aforementioned methods adopt the weak augmentation strategy proposed in SimCLR (Chen et al., 2020a) because strong augmentations (RandAugment Cubuk et al. 2020 to be specific) experimentally show inferior performance (Wang & Qi, 2021). Though there are some works (Van Gansbeke et al., 2020) that use strong augmentations to fine-tune the network, it is still unclear how to directly facilitate contrastive learning with strong augmentations. From this perspective, this work could shed some light on how to effectively utilize weak and strong transformations by using the proposed TCL framework (more details in Table 6). The proposed augmentation strategy is suitable for various types of data, such as images and texts.

2.2 Deep Clustering

Both effective clustering strategies and discriminative features are essential in achieving good clustering. Thanks to the powerful representability of deep neural networks, deep clustering methods have attracted more and more attention in recent (Asano et al., 2019; Caron et al., 2018; Guo et al., 2017; Li et al., 2020, 2021a; Peng et al., 2016; Xie et al., 2016; Yang et al., 2016). To name a few, JULE (Yang et al., 2016) iteratively learns data representation and performs hierarchical clustering. DeepCluster (Caron et al., 2018) clusters data using the prior representation and uses the cluster

assignment of each sample as a classification target to learn the new representation. Though representation learning and clustering could bootstrap each other to some extent, this kind of two-stage method might suffer from errors accumulated during alternations. Another weakness of these methods is that they could not be applied in the online scenario, where data is presented in streams and only a batch of samples are accessible at one time. Specifically, JULE needs the global similarity to decide which sub-clusters should be merged, while DeepCluster and SL (Asano et al., 2019) need to perform offline k-means or solve a global optimal transport problem to acquire cluster assignments. To overcome the offline limitation, some online deep clustering methods are proposed (Dang et al., 2021; Huang et al., 2020; Ji et al., 2019; Li et al., 2021b; Zhong et al., 2020). For example, IIC (Ji et al., 2019) discovers clusters by maximizing mutual information between the cluster assignments of data pairs. PICA (Huang et al., 2020) learns the most semantically plausible data separation by maximizing the partition confidence of the clustering solution. Very recently, some studies (Niu & Wang, 2021; Park et al., 2020; Van Gansbeke et al., 2020) use pseudo-labels generated by preliminary clustering, namely, self-labeling, to further improve the clustering performance in a multi-stage manner.

Unlike most of the above works that perform representation learning and clustering in multiple stages, our method unifies these two tasks into the twin contrastive learning framework. Such a one-stage learning paradigm helps the model to learn more clustering-favorable representations compared with previous works that solely conduct instance-level contrastive learning (Niu & Wang, 2021; Van Gansbeke et al., 2020). In the boosting stage, despite rectifying the cluster assignments based on the features extracted in the early stage (Niu & Wang, 2021), we could also fine-tune the instance-level contrastive learning to alleviate the influence of false-negative pairs thanks to our one-stage learning paradigm.

2.3 Connection Between Contrastive Learning and Deep Clustering

Both representation learning and deep clustering share a common goal, namely, extracting discriminative features. Recently, a variety of works have shown that contrastive learning and deep clustering could bootstrap each other. On the one hand, the performance of representation learning could be enhanced by integrating the clustering property (Wang et al., 2021). For example, instead of constructing pairs between samples, SwAV (Caron et al., 2020) first performs clustering by solving an optimal transport problem and then contrasts the instances and the learned cluster centers. Similarly, PCL (Li et al., 2021a) pulls each sample to its corresponding cluster center with a prototypical contrastive loss.

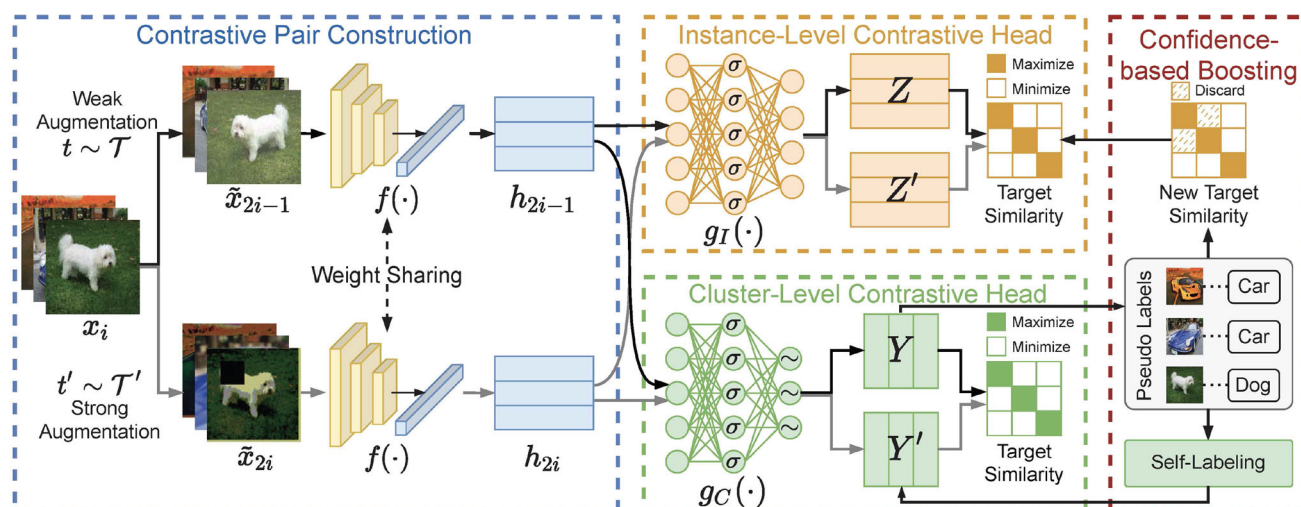


Fig. 2 The pipeline of Twin Contrastive Learning (TCL). First, it constructs data pairs through weak and strong augmentations. A shared backbone is used to extract features from augmented samples. Then, two independent MLPs (σ denotes the ReLU activation and \sim denotes the Softmax operation to produce soft labels) project the features into

the row and column space wherein the instance- and cluster-level contrastive learning are conducted, respectively. Finally, pseudo labels are selected based on the confidence of cluster predictions to alleviate the influence of false-negative pairs and rectify previous predictions, which further boosts the clustering performance

In addition, Van Gansbeke et al. (2021) and Dwivedi et al. (2021) show that leveraging the local neighborhood can be effective for contrastive learning. On the other hand, recent deep clustering methods such as SCAN (Van Gansbeke et al., 2020), CC (Li et al., 2021b), and SPICE (Niu & Wang, 2021) have achieved state-of-the-art performance thanks to the contrastive learning paradigm.

Though the combination of contrastive learning and deep clustering has brought promising results, most existing works still treat these two tasks separately. Different from existing works, this study elegantly unifies contrastive learning and deep clustering into the twin contrastive learning framework, which might bring some insights to both communities. Notably, this study is a significant extension of (Li et al., 2021b) with the following improvements:

- In this paper, we propose a confidence-based boosting strategy to fine-tune both the instance- and cluster-level contrastive learning. Specifically, most confident predictions are selected as pseudo labels based on the observation that they are more likely to be correct. Upon that, we use the generated pseudo labels to alleviate the influence of false-negative pairs (composed of within-class samples) in instance-level contrastive learning, and adopt cross-entropy loss to rectify cluster assignments in cluster-level contrastive learning. Notably, such a twin self-training paradigm is benefited from our TCL framework since the cluster assignments (from CCH) of instance features (from ICH) could be obtained in an online manner.

- In this paper, we propose a data augmentation strategy by mixing weak and strong transformations. Though such an augmentation strategy is seemingly simple, its effectiveness is closely correlated with the proposed TCL framework. Previous works have shown that directly introducing strong augmentation into the contrastive learning framework could lead to sub-optimal performance (Wang & Qi, 2021). Different from such a conclusion, we show that the mixed augmentation strategy naturally fits the proposed TCL framework (see Table 6 for more details).
- To investigate the generalization ability of the proposed method, we verify the effectiveness of our method in text clustering despite the difference in data augmentation. Experimental results demonstrate the superiority of the proposed TCL framework, mixed augmentation strategy, and confidence-based boosting strategy. A comparable performance gain is achieved by this journal extension compared with the previous conference paper (Li et al., 2021b).

3 Method

The pipeline of the proposed TCL is illustrated in Fig. 2. The model consists of three parts, namely, the contrastive pair construction (CPC), the instance-level contrastive head (ICH), and the cluster-level contrastive head (CCH), which are jointly optimized through twin contrastive learning and confidence-based boosting. Specifically, in the twin contrastive learning stage, CPC first constructs contrastive pairs

through data augmentations and then projects the pairs into a latent feature space. After that, ICH and CCH conduct instance- and cluster-level contrastive learning at the row and column space of the feature matrix respectively by minimizing the proposed twin contrastive loss. To alleviate the influence of intrinsic false-negative pairs in contrastive learning and to rectify cluster assignments, we propose a confidence-based boosting strategy (CB). In detail, some confident predictions are selected as pseudo labels to fine-tune the instance- and cluster-level contrastive learning with the self-supervised contrastive loss and self-labeling loss, which further improves the clustering performance.

Once the model converges, CCH could make cluster assignments for each instance to achieve online clustering. Notably, though the twin contrastive learning can be directly conducted on the same contrastive head as indicated in our basic idea, we experimentally find that decoupling it into two independent subspaces improves the clustering performance (see Sect. 4.6.4 for detailed discussion).

In this section, we first introduce the construction of contrastive pairs in CPC, then present the twin contrastive loss for training, and finally elaborate on our confidence-based boosting strategy.

3.1 Contrastive Pair Construction

Inspired by recent developments of contrastive learning (Caron et al., 2020; Chen et al., 2020a), the proposed TCL constructs contrastive pairs through data augmentation. Specifically, for each instance x_i , CPC stochastically samples and applies two groups of transformations t and t' from two augmentations families \mathcal{T} and \mathcal{T}' respectively, resulting in two correlated samples (i.e., a data pair) denoted as $\tilde{x}_{2i-1} = t(x_i)$ and $\tilde{x}_{2i} = t'(x_i)$.

Recent studies suggest that data augmentation is essential in contrastive learning methods (Caron et al., 2020; Chen et al., 2020a), and most of them adopt weak augmentations (Chen et al., 2020a; Grill et al., 2020; Zbontar et al., 2021) since directly using strong augmentations experimentally leads to inferior performance (Wang & Qi, 2021). In this work, we provide a novel data augmentation strategy by mixing weak and strong transformations, which achieves superior performance on both image and text data. To be specific, for image data, we adopt the transformations proposed by SimCLR (Chen et al., 2020a) and RandAugment (Cubuk et al., 2020) as the weak \mathcal{T} and strong \mathcal{T}' augmentation, respectively. For text data, we employ the synonym replacement strategy (Zhang et al., 2021a) as the weak augmentation \mathcal{T} and use the sentence operations (Wei & Zou, 2019) as the strong augmentation \mathcal{T}' .

Given the constructed pairs, a shared backbone $f(\cdot)$ is used to extract features h from the augmented samples through $h_{2i-1} = f(\tilde{x}_{2i-1})$ and $h_{2i} = f(\tilde{x}_{2i})$. Specific back-

bones are used to handle different types of data. In this work, we adopt ResNet (He et al., 2016) and Sentence Transformer (Reimers & Gurevych, 2019) as the backbone for image and text data, respectively.

3.2 Twin Contrastive Learning

In the training stage, the backbone, the instance-level contrastive head (ICH), and the cluster-level contrastive head (CCH) are jointly optimized according to the following twin contrastive loss, i.e.,

$$\mathcal{L}_{train} = \mathcal{L}_{ins} + \mathcal{L}_{clu}, \quad (1)$$

where \mathcal{L}_{ins} is the instance-level contrastive loss which is computed on ICH and \mathcal{L}_{clu} denotes the cluster-level contrastive loss computed on CCH.

In general, one may add a dynamic weight parameter to balance the two losses across the training process, but explicitly tuning the weight could violate the unsupervised constraint. In practice, we find a simple addition of the two contrastive losses already works well.

3.2.1 Instance-Level Contrastive Loss

The instance-level contrastive learning aims to maximize the similarities of positive pairs while minimizing those of negative ones. To achieve clustering, ideally, one could define pairs of within-class instances to be positive and those of between-class instances to be negative. However, since no prior label information is given, we construct instance pairs based on data augmentations as a compromise. To be specific, the positive pairs consist of samples augmented from the same instance, and the negative pairs otherwise.

Formally, for a mini-batch of size N , TCL performs two types of data augmentations on each instance x_i , resulting in $2N$ augmented samples $\{\tilde{x}_1, \tilde{x}_2, \dots, \tilde{x}_{2i-1}, \tilde{x}_{2i}, \dots, \tilde{x}_{2N}\}$. Each sample \tilde{x}_{2i-1} forms $2N - 1$ pairs with others, among which we choose the pair with its corresponding augmented sample $\{\tilde{x}_{2i-1}, \tilde{x}_{2i}\}$ to be positive and define other $2N - 2$ pairs to be negative.

As directly conducting contrastive learning on the feature matrix may cause information loss (Chen et al., 2020a), we stack a two-layer nonlinear MLP $g_I(\cdot)$ to map the features into a subspace via $z_i = g_I(h_i)$, where the instance-level contrastive learning is applied. The pair-wise similarity is measured using cosine distance, namely,

$$s(z_i, z_j) = \frac{z_i z_j^\top}{\|z_i\| \|z_j\|}, i, j \in [1, 2N]. \quad (2)$$

The InfoNCE loss (Oord et al., 2018) is adopted to optimize pair-wise similarities defined by Eq. (2). Without loss of

generality, the loss for a given augmented sample \tilde{x}_i (suppose it forms a positive pair with \tilde{x}_j) is defined as

$$\ell_i = -\log \frac{\exp(s(z_i, z_j)/\tau_I)}{\sum_{k=1}^{2N} \mathbb{1}_{[k \neq i]} \exp(s(z_i, z_k)/\tau_I)}, \quad (3)$$

where τ_I is the instance-level temperature parameter to control the softness, and $\mathbb{1}_{[k \neq i]}$ is an indicator function evaluating to 1 iff $k \neq i$. To identify the positive counterpart for each augmented sample, the instance-level contrastive loss is computed across all augmented samples, i.e.,

$$\mathcal{L}_{ins} = \frac{1}{2N} \sum_{k=1}^{2N} \ell_k. \quad (4)$$

3.2.2 Cluster-Level Contrastive Loss

When a sample is projected into a subspace whose dimensionality equals the cluster number, the i -th element of its feature represents its probability of belonging to the i -th cluster. In other words, the feature vector corresponds to its cluster assignment probability.

Suppose the target cluster number is M , similar to the instance-level contrastive head, we use another two-layer MLP $g_C(\cdot)$ to project the features into an M -dimensional space via $y_i = g_C(h_i)$. Here y_i corresponds to the cluster assignment probability of the augmented sample \tilde{x}_i . Formally, let $Y = [y_1, \dots, y_{2i-1}, \dots, y_{2N-1}] \in \mathcal{R}^{N \times M}$ be the cluster assignment probabilities of the mini-batch under the weak augmentation \mathcal{T} (and $Y' = [y_2, \dots, y_{2i}, \dots, y_{2N}]$ for those under the strong augmentation \mathcal{T}'). Based on the observation shown in Fig. 1, the columns of Y and Y' correspond to the cluster distributions over the mini-batch and could be interpreted as special cluster representations. We would like to point out that this observation still holds even when the dimension is larger than the ground-truth cluster number. In that case, a more fine-grained cluster structure is considered and its effectiveness is verified in Barlow Twins (Zbontar et al., 2021).

For clarity, we denote the i -th column of Y as \hat{y}_{2i-1} (and \hat{y}_{2i} for the i -th column of Y'), namely, the representation of cluster i under the weak (and strong) data augmentation. The representations of the same cluster under two augmentations form positive cluster pairs $\{\hat{y}_{2i-1}, \hat{y}_{2i}\}$, $i \in [1, M]$, while other pairs are defined to be negative. Again, we use cosine distance to measure the similarity between cluster \hat{y}_i and \hat{y}_j , that is

$$s(\hat{y}_i, \hat{y}_j) = \frac{\hat{y}_i^\top \hat{y}_j}{\|\hat{y}_i\| \|\hat{y}_j\|}, i, j \in [1, 2M] \quad (5)$$

Without loss of generality, the following loss function is adopted to identify cluster \hat{y}_i from all other $2M - 2$ clusters except its counter part \hat{y}_j , i.e.,

$$\hat{\ell}_i = -\log \frac{\exp(s(\hat{y}_i, \hat{y}_j)/\tau_C)}{\sum_{k=1}^{2M} \mathbb{1}_{[k \neq i]} \exp(s(\hat{y}_i, \hat{y}_k)/\tau_C)}, \quad (6)$$

where τ_C is the cluster-level temperature parameter to control the softness, and $\mathbb{1}_{[k \neq i]}$ is an indicator function evaluating to 1 iff $k \neq i$. By traversing all clusters, the cluster-level contrastive loss is computed through

$$\mathcal{L}'_{clu} = \frac{1}{2M} \sum_{k=1}^M \hat{\ell}_k. \quad (7)$$

As simply optimizing the above cluster-level contrastive loss might lead to trivial solution where most samples are assigned to a few clusters, we add a cluster entropy to prevent the model from collapsing and achieve more balanced clustering (Ghasedi Dizaji et al., 2017; Huang et al., 2020). Formally, let $P(\hat{y}_{2i-1}) = \frac{1}{N} \sum_{k=1}^N Y_{ki}$ be the assignment probability of cluster i within a mini-batch under the weak augmentation and $P(\hat{y}_{2i}) = \frac{1}{N} \sum_{k=1}^N Y'_{ki}$ be that under the strong augmentation, then the cluster entropy is computed by

$$H_{clu} = -\sum_{i=1}^{2M} [P(\hat{y}_i) \log P(\hat{y}_i)]. \quad (8)$$

To sum up, the cluster-level contrastive loss is finally defined as

$$\mathcal{L}_{clu} = \frac{1}{2M} \sum_{k=1}^{2M} \hat{\ell}_k - H_{clu}. \quad (9)$$

3.3 Confidence-Based Boosting

As the train progresses, we notice that the model tends to make more confident predictions (i.e., with cluster assignment probability close to one-hot). Those confident predictions are more likely to be correct (see Fig. 4). Based on this observation, in the boosting stage, we progressively select the most confident predictions as pseudo labels to fine-tune both the instance- and cluster-level contrastive learning. The pseudo labels are selected by the following criterion. Namely, for a mini-batch of size N , we use the raw data x as input to compute prediction pred with confidence conf for each

Algorithm 1 Twin Contrastive Learning

Input: dataset \mathcal{X} ; training iterations E_1 ; boosting iterations E_2 ; batch size N ; cluster number M ; temperature parameter τ_I and τ_C ; network f , g_I , and g_C ; augmentation strategies \mathcal{T} , \mathcal{T}' .

Output: cluster assignments.

```

// Training
for epoch = 1 to  $E_1$  do
    sample a mini-batch  $\{x_i\}_{i=1}^N$  from  $\mathcal{X}$ 
    sample two augmentations  $t \sim \mathcal{T}, t' \sim \mathcal{T}'$ 
    compute instance and cluster representations
    compute instance-level contrastive loss  $\mathcal{L}_{ins}$  through Eq. 2-4
    compute cluster-level contrastive loss  $\mathcal{L}_{clu}$  through Eq. 5-9
    compute training loss  $\mathcal{L}_{train}$  by Eq. 1
    update  $f, g_I, g_C$  through gradient descent to minimize  $\mathcal{L}_{train}$ 
end
// Boosting
for epoch = 1 to  $E_2$  do
    sample a mini-batch  $\{x_i\}_{i=1}^N$  from  $\mathcal{X}$ 
    sample two augmentations  $t \sim \mathcal{T}, t' \sim \mathcal{T}'$ 
    update pseudo labels with Eq. 10-11 and Eq. 16
    compute self-supervised contrastive loss  $\mathcal{L}_{scl}$  by Eq. 14
    compute self-labeling loss  $\mathcal{L}_{sl}$  by Eq. 15
    compute boosting loss  $\mathcal{L}_{boost}$  by Eq. 12
    update  $f, g_I, g_C$  through gradient descent to minimize  $\mathcal{L}_{boost}$ 
end
// Test
for  $x$  in  $\mathcal{X}$  do
    extract features by  $h = f(x)$ 
    compute cluster assignment by  $c = \arg \max g_C(h)$ 
end
    
```

instance by

$$\begin{aligned}
 y_i &= g_C(f(x_i)), \\
 \text{conf}_i &= \max(y_i), \\
 \text{pred}_i &= \arg \max(y_i).
 \end{aligned} \tag{10}$$

In every mini-batch, we select the top γ confident predictions from each cluster as pseudo labels, where γ is the confident ratio and we fix it to 0.5. To be specific, a prediction pred_i will be selected as pseudo label if it meets the following criteria, i.e.,

$$\begin{aligned}
 n &= \gamma \times N/M, \\
 \text{CONF}_k &= \text{sort}(\{\text{conf}_i \mid i \in [1, N], \text{pred}_i = k\})[n], \\
 \text{conf}_i &\geq \text{CONF}_{\text{pred}_i},
 \end{aligned} \tag{11}$$

where CONF_k is the n -th largest confidence of predictions on cluster $k \in [0, M - 1]$. Notably, selecting most confident predictions from each cluster leads to more class-balanced pseudo labels compared with threshold-based criterion (Niu & Wang, 2021; Park et al., 2020; Van Gansbeke et al., 2020). We store the pseudo labels for all instances, denoted as P , in the memory.

With the generated pseudo labels, we fine-tune the model with the following loss to further boost the clustering performance, namely,

$$\mathcal{L}_{boost} = \mathcal{L}_{scl} + \mathcal{L}_{sl}, \tag{12}$$

where \mathcal{L}_{scl} is the self-supervised contrastive loss used to alleviate the influence of false negative pairs in instance-level contrastive learning and \mathcal{L}_{sl} is the self-labeling loss for rectifying cluster assignments made by CCH.

Recall that in the instance-level contrastive learning, we treat pairs of samples augmented from different instances to be negative, since no label information is given. However, for downstream tasks such as classification and clustering, within-class samples should not be pushed apart. To this end, with the help of pseudo labels, we remove within-class samples from negatives pairs (i.e., the denominator in Eq. 4) and adopt a self-supervised contrastive loss to fine-tune the instance-level contrastive learning. Specifically, for each augmented sample x_i with pseudo label pred_i , the self-supervised contrastive loss is defined as

$$l'_i = -\log \frac{\exp(s(z_i, z_j)/\tau_I)}{\sum_{k=1}^{2N} \mathbb{1}_{[\text{pred}_k \neq \text{pred}_i]} \cdot \exp(s(z_i, z_k)/\tau_I)}, \tag{13}$$

where $\mathbb{1}$ is the indicator. Notably, inspired by the negative learning paradigm (Kim et al., 2019), here we only remove potential within-class pairs from negative ones without treating them as positive, considering that the latter strategy could be too strong when the pseudo labels are of inferior quality. By traversing all augmented samples, the self-supervised instance-level contrastive loss is computed through

$$\mathcal{L}_{scl} = \frac{1}{2N} \sum_{i=1}^{2N} l'_i. \tag{14}$$

For the cluster-level contrastive head, the self-labeling strategy is adopted to rectify previous predictions. Specifically, we define the self-labeling loss as the weighted cross-entropy on the strongly augmented samples $x' = t'(x)$, i.e.,

$$\begin{aligned}
 y'_i &= g_C(f(x'_i)), \\
 \mathcal{L}_{sl} &= -\frac{1}{N_p} \sum_{i=1, i \in P} w_{\text{pred}_i} \log \left(\frac{\exp(y'_i[\text{pred}_i])}{\sum_{k=0}^{M-1} \exp(y'_i[k])} \right),
 \end{aligned} \tag{15}$$

where N_p denotes the number of instances that have pseudo-labels in the mini-batch, and $w_c \propto 1/N_c$ is the weight parameter for cluster c of size N_c . The weighted loss could prevent large clusters from dominating the optimization.

Though the confident ratio is fixed to $\gamma = 0.5$, it does not mean that at most 50% of predictions will be selected as pseudo labels. As the boosting progresses and batch shuffles, more pseudo labels would be selected progressively. Furthermore, considering that the model might make some mistakes in selecting pseudo labels, we remove the pseudo labels from

P once their confidences decrease to below a certain threshold, namely,

$$\text{conf}_i < \alpha, \quad (16)$$

where α is the lower confidence bound of pseudo labels and is set to 0.99. This weeding out mechanism keeps the high quality of pseudo labels and gives the model a chance of rectifying previous predictions. The choice of α and γ is fixed across all the experiments, and we provide parameter analysis on them in Sect. 4.6.2.

The overall training, boosting, and test process of the proposed TCL is summarized in Algorithm 1.

4 Experiments

In this section, the clustering performance of the proposed TCL is evaluated on five image and two text datasets. A series of qualitative analyses and ablation studies are carried out to help comprehensively and intuitively understand the method.

4.1 Datasets

A brief description of the used datasets is summarized in Table 1. More specifically, the image datasets include CIFAR-10, CIFAR-100 (Krizhevsky & Hinton, 2009), STL-10 (Coates et al., 2011), ImageNet-10, and ImageNet-Dogs (Chang et al., 2017a). For CIFAR-100, its 20 super-classes rather than 100 fine-grained classes are taken as the ground truth. The text datasets include StackOverflow (Xu et al., 2017a) and Biomedical (Xu et al., 2017a). The former is a subset of the challenge data published by Kaggle, and the latter is a subset of the PubMed data distributed by BioASQ.

4.2 Experimental Settings

Different backbones could be used to handle different types of data, including but not limited to images and texts. For a fair comparison with previous works (Ji et al., 2019; Huang et al., 2020; Zhang et al., 2021b), we adopt ResNet34 (He et

al., 2016) as the backbone for images and the *distilbert-base-nli-stsb-mean-tokens* model from the Sentence Transformer library (Reimers & Gurevych, 2019) for texts. For ResNet34, its output dimension of the fully-connected classification layer (i.e., the dimension of h) is set to 512. For STL-10, its 100,000 unlabeled images are additionally used to compute the instance-level contrastive loss in the training stage. The ResNet34 is randomly initialized, while the Sentence Transformer is pre-trained to produce meaningful word embeddings to keep consistent with previous works (Xu et al., 2017a; Zhang et al., 2021b). For simplicity, instead of customizing the network to handle images of different sizes, we simply resize all images to 224×224 , and no additional modification is made on the standard ResNet34.

For images, we adopt data transformations proposed in SimCLR (Chen et al., 2020a) as weak augmentation, including ResizedCrop, ColorJitter, Grayscale, HorizontalFlip, and GaussianBlur. Notably, as small images already become blurred after up-scaling, we leave the GaussianBlur augmentation out for CIFAR-10/100 in the weak augmentation \mathcal{T} . The strong augmentation \mathcal{T}' is composed of four randomly selected transformations from RandAugment (Cubuk et al., 2020) with parameters listed in Table 2, followed by one Cutout (DeVries & Taylor, 2017) operation with a size of 75×75 .

For text data, we randomly substitute 20% words of each text with their top- n suitable words found by the pre-trained Roberta from the Contextual Augmenter Library (Ma, 2019) as weak augmentation, following the setting in SCCL (Zhang et al., 2021a). The four operations proposed by EDA (Wei & Zou, 2019) with a probability of 0.2 each are adopted as strong augmentations, including SynonymReplacement, RandomInsertion, RandomSwap, and RandomDeletion.

For the two contrastive heads, the dimensionality of ICH is set as 128 to keep more discriminative information (see ablation study in Sect. 4.6.4), and the dimensionality of CCH

Table 1 A summary of datasets used for evaluation

Dataset	Split	Samples	Classes
CIFAR-10	Train + Test	60,000	10
CIFAR-100	Train + Test	60,000	20
STL-10	Train + Test	13,000	10
ImageNet-10	Train	13,000	10
ImageNet-Dogs	Train	19,500	15
StackOverflow	–	20,000	20
Biomedical	–	20,000	20

Table 2 List of strong augmentations for images

Transformations	Parameter	Range
AutoContrast	–	–
Equalize	–	–
Identity	–	–
Brightness	B	[0.05, 0.95]
Color	C	[0.05, 0.95]
Contrast	C	[0.05, 0.95]
Posterize	B	[4, 8]
Rotate	θ	[– 30, 30]
Sharpness	S	[0.05, 0.95]
Shear X, Y	R	[– 0.3, 0.3]
Solarize	T	[0, 256]
Translate X, Y	λ	[– 0.3, 0.3]

is naturally set to the target cluster number. The temperature parameters are empirically set as $\tau_I = 0.5$, $\tau_C = 1.0$ for all datasets. The Adam optimizer with an initial learning rate of $1e-4$ and a weight decay of $1e-4$ is adopted to jointly optimize the two contrastive heads and the backbone network on image datasets. Since the backbone is pre-trained for text data, we set the learning rate of the optimizer as $5e-6$ for the backbone and $5e-4$ for two contrastive heads. The model is trained for 1000/500 epochs, followed by 200/100 boosting epochs for the image/text dataset with a batch size of 256. Experiments are carried out on Nvidia TITAN RTX 24 G and RTX 3090 on the Ubuntu 18.04 platform with CUDA 11.0 and PyTorch 1.8.0 (Paszke et al., 2019).

4.3 Compared Methods

The proposed TCL is evaluated on five image datasets and two text datasets. For image clustering, we take 21 representative state-of-the-art approaches for comparisons, including k-means (MacQueen, 1967), SC (Zelnik-Manor & Perona, 2005), AC (Gowda & Krishna, 1978), NMF (Cai et al., 2009), AE (Bengio et al., 2007), DAE (Vincent et al., 2010), DCGAN (Radford et al., 2015), DeCNN (Zeiler et al., 2010), VAE (Kingma & Welling, 2013), JULE (Yang et al., 2016), DEC (Xie et al., 2016), DAC (Chang et al., 2017b), ADC (Haeusser et al., 2018), DDC (Chang et al., 2019), DCCM (Wu et al., 2019), IIC (Ji et al., 2019), PICA (Huang et al., 2020), CC (Li et al., 2021b), SPICE (Niu & Wang, 2021), SCAN (Van Gansbeke et al., 2020), and PCL (Li et al., 2021a). For those representation-based methods, namely SC, NMF, AE, DAE, DCGAN, DeCNN, and VAE, clustering is achieved by applying the vanilla k-means on the learned features. To ensure the backbone is the same across all recent deep clustering methods, we reproduce SCAN with ResNet34 using its official released code. We would like to point out that SPICE further boosts the clustering performance under a semi-supervised framework, and it uses a deeper and wider ResNet backbone (e.g., WRN37-2) which enjoys a much better feature extraction ability (Chen et al., 2020). Thus, for a fair comparison, here we compare it with its self-trained results which are achieved on ResNet34. Besides, SPICE uses the model pre-trained on ImageNet for ImageNet-10/Dogs (denoted by “()” in Table 3), while all other methods including ours train the model from scratch.

For text clustering, we compare the proposed TCL with 11 benchmarks, including TF/TF-IDF (Jones, 1972), BagOf-Words (BOW) (Harris, 1954), SkipVec (Kiros et al., 2015), Para2Vec (Le & Mikolov, 2014), GSDPMM (Yin & Wang, 2016), RecNN (Socher et al., 2011), STCC (Xu et al., 2017b), HAC-SD (Rakib et al., 2020), ECIC (Rakib et al., 2020), and SCCL (Zhang et al., 2021a). Similarly, the vanilla k-means is conducted on the extracted features to cluster data for those

representation-based methods, including BOW, TF/TF-IDF, SkipVec, Para2Vec, and RecNN.

4.4 Evaluation Metrics

Three widely-used clustering metrics including Normalized Mutual Information (NMI) (Strehl & Ghosh, 2002), Clustering Accuracy (ACC) (Li & Ding, 2006), and Adjusted Rand Index (ARI) (Hubert & Arabie, 1985) are utilized to evaluate our method. Higher scores indicate better clustering performance.

4.5 Results

Both quantitative and qualitative studies are carried out to evaluate the proposed method. Specifically, we compare TCL with state-of-the-art baselines on image and text benchmarks, and visualize the clustering results across the training process.

4.5.1 Comparisons with State of the Arts

The clustering results on five image benchmarks shown in Table 3 demonstrate the promising performance of TCL. It is worth noting that our method even outperforms SPICE (Niu & Wang, 2021) on the ImageNet-Dogs dataset without ImageNet pre-training, which proves the effectiveness of TCL.

Table 4 shows the clustering results on two commonly-used text datasets. Because existing works seldom use the ARI to evaluate text clustering, here we just adopt NMI and ACC for comparisons. The results show that TCL achieves promising performance on both datasets. We would like to point out that SCCL (Zhang et al., 2021a) is also a contrastive learning based method, which achieves clustering by conducting DEC (Xie et al., 2016) on the representation learned by the instance-level contrastive learning. The dominance of TCL over SCCL (Zhang et al., 2021a) proves the effectiveness of the proposed twin contrastive learning and confidence-based boosting.

4.5.2 Evolution of Instance Features and Cluster Assignments

The instance- and cluster-level contrastive learning ought to help the model to learn discriminative representations and predict accurate cluster assignments, respectively. To experimentally study the convergence of TCL, we perform t-SNE (Van der Maaten & Hinton, 2008) on representations learned by ICH at four timestamps throughout the training and boosting stage, where the cluster assignments predicted by CCH are denoted in different colors. As shown in Fig. 3, features are all mixed and most instances are assigned to a few clusters at the beginning. As the training progresses, fea-

Table 3 The clustering performance on five object image benchmarks

Dataset Metrics	CIFAR-10			CIFAR-100			STL-10			ImageNet-10			ImageNet-Dogs		
	NMI	ACC	ARI	NMI	ACC	ARI	NMI	ACC	ARI	NMI	ACC	ARI	NMI	ACC	ARI
k-means	0.087	0.229	0.049	0.084	0.130	0.028	0.125	0.192	0.061	0.119	0.241	0.057	0.055	0.105	0.020
SC	0.103	0.247	0.085	0.090	0.136	0.022	0.098	0.159	0.048	0.151	0.274	0.076	0.038	0.111	0.013
AC	0.105	0.228	0.065	0.098	0.138	0.034	0.239	0.332	0.140	0.138	0.242	0.067	0.037	0.139	0.021
NMF	0.081	0.190	0.034	0.079	0.118	0.026	0.096	0.180	0.046	0.132	0.230	0.065	0.044	0.118	0.016
AE	0.239	0.314	0.169	0.100	0.165	0.048	0.250	0.303	0.161	0.210	0.317	0.152	0.104	0.185	0.073
DAE	0.251	0.297	0.163	0.111	0.151	0.046	0.224	0.302	0.152	0.206	0.304	0.138	0.104	0.190	0.078
DCGAN	0.265	0.315	0.176	0.120	0.151	0.045	0.210	0.298	0.139	0.225	0.346	0.157	0.121	0.174	0.078
DeCNN	0.240	0.282	0.174	0.092	0.133	0.038	0.227	0.299	0.162	0.186	0.313	0.142	0.098	0.175	0.073
VAE	0.245	0.291	0.167	0.108	0.152	0.040	0.200	0.282	0.146	0.193	0.334	0.168	0.107	0.179	0.079
JULE	0.192	0.272	0.138	0.103	0.137	0.033	0.182	0.277	0.164	0.175	0.300	0.138	0.054	0.138	0.028
DEC	0.257	0.301	0.161	0.136	0.185	0.050	0.276	0.359	0.186	0.282	0.381	0.203	0.122	0.195	0.079
DAC	0.396	0.522	0.306	0.185	0.238	0.088	0.366	0.470	0.257	0.394	0.527	0.302	0.219	0.275	0.111
ADC	–	0.325	–	–	0.160	–	–	–	–	–	–	–	–	–	–
DDC	0.424	0.524	0.329	–	–	–	0.371	0.489	0.267	0.433	0.577	0.345	–	–	–
DCCM	0.496	0.623	0.408	0.285	0.327	0.173	0.376	0.482	0.262	0.608	0.710	0.555	0.321	0.383	0.182
IIC	–	0.617	–	–	0.257	–	–	0.610	–	–	–	–	–	–	–
PICA	0.591	0.696	0.512	0.310	0.337	0.171	0.611	0.713	0.531	0.802	0.870	0.761	0.352	0.352	0.201
CC	0.705	0.790	0.637	0.431	0.429	0.266	0.746	0.850	0.726	0.859	0.893	0.822	0.445	0.429	0.274
SPICE	0.734	0.838	0.705	0.448	0.468	0.294	0.817	0.908	0.812	(0.927)	(0.969)	(0.933)	(0.498)	(0.546)	(0.362)
SCAN	0.796	0.861	0.750	0.485	0.483	0.314	0.703	0.818	0.661	–	–	–	–	–	–
PCL	<i>0.802</i>	<i>0.874</i>	<i>0.766</i>	<i>0.528</i>	<i>0.526</i>	0.363	0.410	0.718	0.670	0.841	0.907	0.822	0.440	0.412	0.299
TCL	0.819	0.887	0.780	0.529	0.531	<i>0.357</i>	<i>0.799</i>	<i>0.868</i>	<i>0.757</i>	0.875	<i>0.895</i>	0.837	<i>0.623</i>	0.644	0.516
TCL _{ICH}	0.792	0.867	0.737	0.522	0.517	0.337	0.732	0.792	0.564	<i>0.869</i>	0.891	<i>0.823</i>	0.624	<i>0.639</i>	<i>0.503</i>

The first and second best results are shown in bold and italic, respectively. “()” denotes that extra training data is used. “TCL_{ICH}” refers to clustering results achieved by conducting *k*-means on the ICH features

Table 4 The clustering performance on two text datasets.

Dataset Metrics	StackOverflow		Biomedical	
	NMI	ACC	NMI	ACC
TF	0.078	0.135	0.093	0.152
BOW	0.140	0.185	0.092	0.143
SkipVec	0.027	0.009	0.107	0.163
TF-IDF	0.156	0.203	0.254	0.280
Para2Vec	0.279	0.326	0.348	0.413
GSDPMM	0.306	0.294	0.320	0.281
RecNN	0.402	0.405	0.338	0.367
STCC	0.490	0.511	0.381	0.436
HAC-SD	0.595	0.648	0.335	0.401
ECIC	0.734	0.787	0.413	<i>0.478</i>
SCCL	0.745	0.755	0.415	0.462
TCL	<i>0.786</i>	0.882	0.429	0.498
TCL _{ICH}	0.788	<i>0.807</i>	<i>0.423</i>	0.470

The first and second best results are shown in bold and italic, respectively. “TCL_{ICH}” refers to clustering results achieved by conducting *k*-means on the ICH features

tures scatter more distinctly and cluster assignments become more balanced. Finally, more compact and well-separated clusters are achieved with the help of the confidence-based boosting.

4.6 Ablation Study

Five ablation studies are carried out to further show the importance of each component in the proposed method. To be specific, the effectiveness of the boosting strategy, the decoupling strategy, and the two contrastive heads are studied. In addition, we investigate the influence of the hyper-parameters in the boosting stage. Besides, different combinations of weak and strong augmentation are tried to verify the effectiveness of the proposed mixed augmentation.

4.6.1 Effectiveness of the Boosting Strategy

To verify the effectiveness of fine-tuning at both the instance and cluster level, we conduct ablation studies by removing one and both of the boosting losses and report the results




Fig. 3 The evolution of instance features and cluster assignments across the training and boosting stage on CIFAR-10. We perform t-SNE on the features learned by ICH and use different colors to indicate the cluster assignment predicted by CCH

in Table 5. The self-labeling loss on CCH is essential in performance boosting because it directly affects the cluster assignments and thus cannot be removed. The results show that both losses improve the clustering performance.

The success of fine-tuning relies on the quality of pseudo labels. As shown in Fig. 4, the confident predictions are more

Table 5 Effectiveness of the boosting strategy

Dataset	Boost	NMI	ACC	ARI
CIFAR-10	None	0.790	0.865	0.752
	SL	0.805	0.878	0.770
	SL + SCL	0.819	0.887	0.780
ImageNet-Dogs	None	0.518	0.549	0.381
	SL	0.562	0.600	0.441
	SL + SCL	0.623	0.644	0.516
StackOverflow	None	0.751	0.860	0.731
	SL	0.780	0.877	0.761
	SL + SCL	0.786	0.882	0.771

Best results are in bold
 “SL” refers to the self-labeling loss, and “SCL” refers to the self-supervised contrastive loss




Fig. 4 Evolution of clustering performance and confident predictions w.r.t. training epochs on CIFAR-10, where confident predictions are those with $CONF \geq 0.99$

likely to be correct after a period of training, and the model makes more confident predictions as the training progresses. The number of confident predictions significantly increases at the boosting stage due to the self-labeling loss. However, it does not mean that all the confident predictions are selected as pseudo labels since the selection is based on rating instead of an absolute threshold.

Note that as the following ablation studies only influence the training stage, we report their clustering performances without boosting for simplicity.

4.6.2 Hyper-Parameters for Pseudo Label Selection

Recall that in the boosting stage, we select most confident cluster assignments as pseudo labels based on the threshold α and the ratio γ . To investigate the influence of two hyper-parameters, we evaluate different choices of them in Fig. 5. From Fig. 5a and b, one could see that the boosting performance is stable across a reasonable range of α and γ , but degrades when one of the criteria is abandoned (i.e., $\alpha = 0.0$ or $\gamma = 1.0$). Such a result indicates the robustness and necessity of two pseudo label selection criteria. Besides, since pseudo labels are selected in a batch-wise manner, we also investigate the influence of batch size. To avoid the influence of batch size on contrastive learning, pseudo labels are selected at the beginning of each epoch within mini-batch of different sizes, and the batch size for optimization remains the default (i.e., 256). Figure 5c shows that the boosting performance is stable in general, but it slightly decreases for large batch sizes. A possible explanation is that for fixed α and γ , a larger N intrinsically tightens the pseudo label selection. In practice, a smaller γ is preferred when the batch size is small and vice versa.

4.6.3 Different Combinations of Data Augmentation

As discussed above, we provide a new effective augmentation strategy by mixing weak and strong transformations. To further show its superiority, we validate our model with different combinations of the weak and strong augmentation.

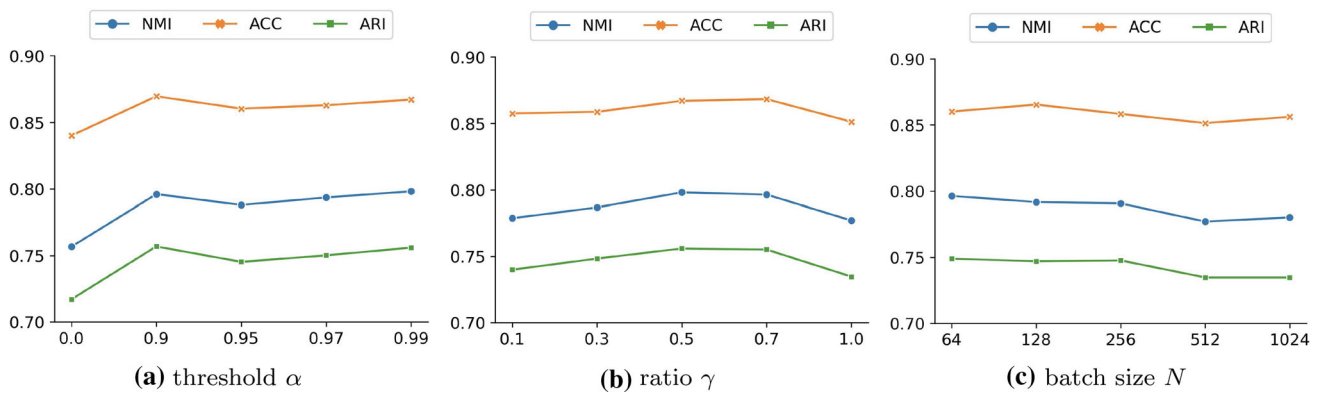


Fig. 5 Ablation study of different choices of hyper-parameters during the boosting stage on STL-10. One of the hyper-parameters is ablated from the default setting $\alpha = 0.99$, $\gamma = 0.5$, $N = 256$ each time

Table 6 Clustering performance gain on CIFAR-10 by introducing the mixed augmentation strategy

Method	Augmentation	NMI	ACC	ARI
SimCLR+k-means	$\mathcal{T} + \mathcal{T}$	0.699	0.782	0.616
	$\mathcal{T} + \mathcal{T}'$	0.734 (+ 0.035)	0.821 (+ 0.039)	0.675 (+ 0.059)
TCL	$\mathcal{T} + \mathcal{T}$	0.705	0.790	0.637
	$\mathcal{T} + \mathcal{T}'$	0.790 (+ 0.085)	0.865 (+ 0.075)	0.752 (+ 0.115)

Best results are in bold

\mathcal{T} denotes the weak augmentation, and \mathcal{T}' denotes the strong augmentation

Table 7 Importance of data augmentation

Dataset	Augmentation	NMI	ACC	ARI
CIFAR-10	$\mathcal{T} + \mathcal{T}$	0.712	0.805	0.648
	$\mathcal{T} + \mathcal{T}'$	0.790	0.865	0.752
	$\mathcal{T}' + \mathcal{T}'$	0.678	0.738	0.614
ImageNet-Dogs	$\mathcal{T} + \mathcal{T}$	0.441	0.413	0.262
	$\mathcal{T} + \mathcal{T}'$	0.518	0.549	0.381
	$\mathcal{T}' + \mathcal{T}'$	0.071	0.147	0.027
StackOverflow	$\mathcal{T} + \mathcal{T}$	0.546	0.580	0.420
	$\mathcal{T} + \mathcal{T}'$	0.751	0.860	0.731
	$\mathcal{T}' + \mathcal{T}'$	0.647	0.715	0.569

Best results are in bold

\mathcal{T} denotes the weak augmentation, and \mathcal{T}' denotes the strong augmentation

The results shown in Table 7 suggest that for both image and text data, such a mixed augmentation strategy results in the best clustering performance.

Notably, though such an augmentation strategy is seemingly simple, its effectiveness is closely correlated with the proposed TCL framework. Previous works have shown that directly introducing strong augmentation into the contrastive learning framework could lead to sub-optimal performance (Wang & Qi, 2021). Different from such a conclusion, we show that the mixed augmentation strategy naturally fits the proposed TCL framework. To support the claim, we compare the performance gain obtained by introducing mixed

augmentation to the “SimCLR+k-means” paradigm and our TCL framework. Table 6 shows the clustering performance on CIFAR-10, which demonstrates that our TCL benefits more from the mixed augmentation strategy.

4.6.4 Effectiveness of the Decoupling Strategy

Although the instance- and cluster-level contrastive learning could be directly conducted in the same subspace, in practice, we find that decoupling them into two separate subspaces leads to better clustering performance. Table 8 shows the ablation results, where “Yes” denotes that the twin contrastive learning is conducted in two separate subspaces and “No” denotes that both contrastive losses are computed in CCH.

Though decoupling the twin contrastive learning into two subspaces improves the performance, it does not contradict our motivation. No matter the two contrastive losses are computed jointly or separately, they act on the same representation h , and jointly optimize the network. In fact, it is consistent with the common contrastive learning framework (Chen et al., 2020a; He et al., 2020) by stacking projection heads on the representation to compute contrastive loss. The inferior performance of computing two contrastive losses in the same subspace could be attributed to two facts: (i) the dimension of rows that equals the cluster number is not high enough to contain much information for instance-level contrastive learning (see ablation studies in Table 9); and (ii) it would lead to a sub-optimal solution since the

Table 8 Effectiveness of the decoupling strategy

Dataset	Decoupling	NMI	ACC	ARI
CIFAR-10	Yes	0.770	0.854	0.730
	No	0.601	0.640	0.481
ImageNet-Dogs	Yes	0.518	0.535	0.385
	No	0.343	0.368	0.214
StackOverflow	Yes	0.747	0.857	0.725
	No	0.687	0.745	0.616

Best results are in bold

Table 9 Difference choices of the dimensionality of ICH

Dataset	Dimension	NMI	ACC	ARI
CIFAR-10	10	0.770	0.854	0.730
	32	0.789	0.867	0.754
	64	0.798	0.871	0.760
	128	0.790	0.866	0.752
	256	0.789	0.867	0.753
ImageNet-Dogs	15	0.518	0.535	0.385
	32	0.535	0.557	0.400
	64	0.529	0.551	0.389
	128	0.518	0.549	0.381
	256	0.541	0.562	0.411
StackOverflow	20	0.747	0.857	0.725
	32	0.749	0.858	0.726
	64	0.748	0.858	0.728
	128	0.751	0.860	0.731
	256	0.727	0.821	0.688

Best results are in bold

instance- and cluster-level contrast will influence each other in the same subspace. In brief, the instance-level contrastive learning aims at discriminating different instances instead of clusters, whose optimal solution is thus inferior to the cluster-level contrastive learning.

4.6.5 Importance of Twin Contrastive Heads

To investigate the effectiveness of the twin contrastive heads (i.e., ICH and CCH), we conduct ablation studies by removing one of them. Since the cluster assignments can no longer be obtained when CCH is removed, we perform k-means on the features learned by ICH instead.

The results in Table 10 prove the effectiveness of two contrastive heads and show that superior performance is obtained by jointly conducting instance- and cluster-level contrastive learning. Despite the performance improvement brought by CCH, we would like to emphasize that CCH is essential in achieving online clustering as it could directly and independently make cluster assignments for each instance.

Table 10 Effectiveness of two contrastive heads

Dataset	Contrastive Head	NMI	ACC	ARI
CIFAR-10	ICH + CCH	0.790	0.865	0.752
	w/o ICH	0.633	0.658	0.522
	w/o CCH	0.734	0.821	0.675
ImageNet-Dogs	ICH + CCH	0.518	0.549	0.381
	w/o ICH	0.376	0.366	0.221
	w/o CCH	0.504	0.535	0.336
StackOverflow	ICH + CCH	0.751	0.860	0.731
	w/o ICH	0.735	0.842	0.706
	w/o CCH	0.666	0.732	0.516

Best results are in bold

4.6.6 Over-Clustering Experiments

It is highly expected that the clustering methods could be robust to different choices of the target cluster number. Therefore, we conduct over-clustering experiments by doubling the ground-truth cluster number (and even 100 classes for CIFAR-100). In the experiments, all parameters except the CCH dimension remain the same as the default. As the Hungarian matching is not practical to find a many-to-one mapping for over-clustering evaluation, we adopt the majority voting mechanism as an alternative. Namely, we assign all samples of one cluster to the majority ground-truth class among the cluster, being consistent with the criterion used in SCAN (Van Gansbeke et al., 2020). The results in Table 11 demonstrate that TCL is robust against different target cluster numbers.

Essentially, the effect of over-clustering is to increase the intra-class variance. Such behavior is more favorable when the data is intrinsically fine-grained (i.e., composed of several subclasses). Forcing the model to produce more fine-grained partitions could break the cluster structure when the data is intrinsically coarse-grained, which explains the performance drop on CIFAR-10. On CIFAR-100 instead, the over-clustering significantly improves the clustering performance. However, one may note that over-clustering with 100 clusters is slightly worse than that with 40 clusters. Such a performance drop could attribute to the insufficiently large

Table 11 The robustness of TCL against different choices of the target cluster number

Dataset	Target cluster number	NMI	ACC	ARI
CIFAR-10	Standard (10)	0.790	0.865	0.752
	Over-cluster (20)	0.759	0.846	0.718
CIFAR-100	Standard (20)	0.477	0.481	0.303
	Over-cluster (40)	0.520	0.579	0.380
	Over-cluster (100)	0.493	0.566	0.359

Best results are in bold

batch size. To correctly represent clusters in CCH, it is necessary to include a reasonable number of samples in each cluster. With a batch size of 256, there are only two samples in each cluster on average. And some clusters might even be empty in some mini-batches, which would harm the cluster-level contrastive learning, leading to inferior clustering performance.

Note that in the proposed TCL framework, the target cluster number needs to be manually set. In practice, when the intrinsic cluster number is unknown, one could adopt some cluster number searching or community detection methods like X-means (Pelleg & Moore, 2000) and Louvain (Blondel et al., 2008) on the ICH output to estimate the cluster number. Moreover, as our TCL is robust to different choices of cluster numbers, the estimation does not necessarily need to be very precise. And one could adjust the dimensionality of CCH to cluster data under different resolutions depending on the practical needs.

4.6.7 Scaling-up to ImageNet

To see how TCL scales to large datasets with much more instances and clusters, we further conduct experiments on the ImageNet dataset. The training split with 1,281,167 images is used for both training and evaluation. As discussed above, cluster-level contrastive learning requires a large batch size to ensure a reasonable number of samples exist in each cluster. Thus, we set the batch size to 4096 on ImageNet with 1000 classes. However, we encounter the “out of GPU memory” problem even on eight RTX 3090 GPUs with such a large batch size. As a solution, we inherit the ResNet50 model learned by MoCo v2 Chen et al. (2020c) and freeze the first two blocks. In other words, we only optimize the last two blocks of ResNet50 and two contrastive heads. Due to the heavy computational burden, we train and boost the model for 100 and 20 epochs, respectively. We set $\tau_I = \tau_C = 0.2$ for training and use the default $\alpha = 0.99$ and $\gamma = 0.5$ for boosting. In the evaluation, we choose MoCo v2 (the model we inherited) as a baseline by conducting k-means on the extracted features. The clustering performance is shown in Table 12, which proves the effectiveness of the twin contrastive learning framework, the boosting strategy, and the scalability of our method.

Table 12 TCL scales up to ImageNet-1K with much more images and clusters

Metrics	NMI	ACC	ARI
MoCo (inherited model)	0.6186	0.3047	0.1428
TCL [†]	0.6332	0.3160	0.1901
TCL	0.6711	0.3789	0.2656

Best results are in bold

[†]Means without the boosting stage




Fig. 6 Ablation study of different choices of temperature τ_I and τ_C in TCL, trained for 50 epochs on ImageNet. One of the hyper-parameters is ablated from $\tau_I = \tau_C = 0.2$




Fig. 7 Ablation study of different choices of α and γ during the boosting stage on ImageNet. One of the hyper-parameters is ablated from the default setting $\alpha = 0.99$ and $\gamma = 0.5$

As suggested in SimCLR (Chen et al., 2020a), a smaller temperature τ_I in instance-level contrastive loss could speed up convergence with a large batch size. To investigate the influence of τ_I and τ_C in our TCL framework when scaling up to large batch size and cluster number, we ablate these two parameters in Fig. 6. As shown, a proper choice of τ_I could slightly improve the performance. More importantly, a smaller temperature τ_C is also preferred in cluster-level contrastive learning with a large cluster number, since it could sharpen the cluster assignments to obtain a more discriminative cluster representation. We further conduct ablation studies on two boosting hyper-parameters α and γ . The results in Fig. 7 show that the threshold α plays a more important role on ImageNet compared with the ratio γ . This is because all samples are likely to meet the ratio criterion eventually, as there are only four samples in each cluster on average.

5 Conclusion

Based on the observation that the rows and columns of the feature matrix intrinsically correspond to the instance and cluster representations, we propose an online deep clustering method termed Twin Contrastive Learning (TCL). By dually conducting contrastive learning at the instance and cluster level, TCL simultaneously learns representations and

performs clustering. In addition, to alleviate the influence of intrinsic false negative pairs in the instance-level contrastive learning and to rectify cluster assignments, we propose a confidence-based boosting strategy to further improve the performance by selecting some pseudo labels to fine-tune the twin contrastive learning. Extensive experiments demonstrate the effectiveness of TCL on five image benchmarks and two text datasets.

In the future, we plan to take a deeper look at the influence of different augmentations on contrastive learning. Furthermore, it is worthwhile to explore how to better utilize pseudo labels to identify false negative pairs for improving contrastive learning.

Acknowledgements The authors would like to thank the Associate editor and reviewers for the constructive comments and valuable suggestions that remarkably improve this study. This work was supported in part by the National Key R&D Program of China under Grant 2020YFB1406702; in part by NFSC under Grant 62176171, U21B2040, and U19A2078; and in part by Open Research Projects of Zhejiang Lab under Grant 2021KH0AB02.

References

- Asano, Y., Rupprecht, C., & Vedaldi, A. (2019). Self-labelling via simultaneous clustering and representation learning. In *International conference on learning representations*.
- Bengio, Y., Lamblin, P., Popovici, D., & Larochelle, H. (2007). Greedy layer-wise training of deep networks. In *Advances in neural information processing systems* (pp. 153–160).
- Blondel, V. D., Guillaume, J. L., Lambiotte, R., & Lefebvre, E. (2008). Fast unfolding of communities in large networks. *Journal of statistical mechanics: Theory and experiment*, 10, P10008.
- Cai, D., He, X., Wang, X., Bao, H., & Han, J. (2009). Locality preserving nonnegative matrix factorization. *IJCAI*, 9, 1010–1015.
- Caron, M., Bojanowski, P., Joulin, A., & Douze, M. (2018). Deep clustering for unsupervised learning of visual features. In *Proceedings of the European conference on computer vision (ECCV)* (pp. 132–149).
- Caron, M., Misra, I., Mairal, J., Goyal, P., Bojanowski, P., & Joulin, A. (2020). Unsupervised learning of visual features by contrasting cluster assignments. In *Thirty-fourth conference on neural information processing systems (NeurIPS)*.
- Chang, J., Guo, Y., Wang, L., Meng, G., Xiang, S., & Pan, C. (2019). Deep discriminative clustering analysis. arXiv preprint [arXiv:1905.01681](https://arxiv.org/abs/1905.01681)
- Chang, J., Wang, L., Meng, G., Xiang, S., & Pan, C. (2017a). Deep adaptive image clustering. In *Proceedings of the IEEE international conference on computer vision* (pp. 5879–5887).
- Chang, J., Wang, L., Meng, G., Xiang, S., & Pan, C. (2017b). Deep adaptive image clustering. In *Proceedings of the IEEE international conference on computer vision* (pp. 5879–5887).
- Chen, X., & He, K. (2020). Exploring simple siamese representation learning. arXiv preprint [arXiv:2011.10566](https://arxiv.org/abs/2011.10566)
- Chen, X., Fan, H., Girshick, R., & He, K. (2020c). Improved baselines with momentum contrastive learning. arXiv preprint [arXiv:2003.04297](https://arxiv.org/abs/2003.04297)
- Chen, T., Kornblith, S., Norouzi, M., Hinton, G. (2020a). A simple framework for contrastive learning of visual representations. arXiv preprint [arXiv:2002.05709](https://arxiv.org/abs/2002.05709)
- Chen, T., Kornblith, S., Swersky, K., Norouzi, M., & Hinton, G. E. (2020). Big self-supervised models are strong semi-supervised learners. *Advances in Neural Information Processing Systems*, 33, 22243–22255.
- Chen, G., & Lerman, G. (2009). Spectral curvature clustering (scc). *International Journal of Computer Vision*, 81(3), 317–330.
- Coates, A., Ng, A., & Lee, H. (2011). An analysis of single-layer networks in unsupervised feature learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics* (pp. 215–223).
- Cubuk, E. D., Zoph, B., Shlens, J., & Le, Q. V. (2020). Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops* (pp. 702–703).
- Dang, Z., Deng, C., Yang, X., Huang, H. (2021). Doubly contrastive deep clustering. arXiv preprint [arXiv:2103.05484](https://arxiv.org/abs/2103.05484)
- DeVries, T., & Taylor, G.W. (2017). Improved regularization of convolutional neural networks with cutout. arXiv preprint [arXiv:1708.04552](https://arxiv.org/abs/1708.04552)
- Dwivedi, D., Aytar, Y., Tompson, J., Sermanet, P., & Zisserman, A. (2021). With a little help from my friends: Nearest-neighbor contrastive learning of visual representations. arXiv preprint [arXiv:2104.14548](https://arxiv.org/abs/2104.14548)
- Ghasedi Dizaji, K., Herandi, A., Deng, C., Cai, W., & Huang, H. (2017). Deep clustering via joint convolutional autoencoder embedding and relative entropy minimization. In *Proceedings of the IEEE international conference on computer vision* (pp. 5736–5745).
- Gowda, K. C., & Krishna, G. (1978). Agglomerative clustering using the concept of mutual nearest neighbourhood. *Pattern Recognition*, 10(2), 105–112.
- Grill, J. B., Strub, F., Altché, F., Tallec, C., Richemond, P.H., Buchatskaya, E., Doersch, C., Pires, B. A., Guo, ZD., & Azar, M. G. et al. (2020). Bootstrap your own latent: A new approach to self-supervised learning. arXiv preprint [arXiv:2006.07733](https://arxiv.org/abs/2006.07733)
- Guo, X., Gao, L., Liu, X., & Yin, J. (2017). Improved deep embedded clustering with local structure preservation. In *IJCAI* (pp. 1753–1759).
- Hadsell, R., Chopra, S., & LeCun, Y. (2006). Dimensionality reduction by learning an invariant mapping. In *2006 IEEE computer society conference on computer vision and pattern recognition (CVPR'06)*, IEEE (Vol. 2, pp. 1735–1742).
- Haeusser, P., Plapp, J., Golkov, V., Aljalbout, E., & Cremers, D. (2018). Associative deep clustering: Training a classification network with no labels. In *German conference on pattern recognition* (pp. 18–32). Springer.
- Han, S., Park, S., Park, S., Kim, S., & Cha, M. (2020). Mitigating embedding and class assignment mismatch in unsupervised image classification. In *Computer Vision—ECCV 2020: 16th European conference, Glasgow, UK, August 23–28, 2020, proceedings, part XXIV 16* (pp. 768–784). Springer.
- Harris, Z. S. (1954). Distributional structure. *Word*, 10(2–3), 146–162.
- He, K., Fan, H., Wu, Y., Xie, S., & Girshick, R. (2020). Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 9729–9738).
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 770–778).
- Hu, W., Miyato, T., Tokui, S., Matsumoto, E., & Sugiyama, M. (2017). Learning discrete representations via information maximizing self-augmented training. In: *International conference on machine learning*, PMLR (pp. 1558–1567).
- Hu, Q., Wang, X., Hu, W., & Qi, G.J. (2020). Adco: Adversarial contrast for efficient learning of unsupervised representations from self-trained negative adversaries. arXiv preprint [arXiv:2011.08435](https://arxiv.org/abs/2011.08435)

- Huang, J., Gong, S., & Zhu, X. (2020). Deep semantic clustering by partition confidence maximisation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*.
- Hubert, L., & Arable, P. (1985). Comparing partitions. *Journal of Classification*, 2(1), 193–218.
- Ji, X., Henriques, J. F., & Vedaldi, A. (2019). Invariant information clustering for unsupervised image classification and segmentation. In *Proceedings of the IEEE international conference on computer vision* (pp. 9865–9874).
- Jones, K. S. (1972). A statistical interpretation of term specificity and its application in retrieval. *Journal of Documentation*.
- Khosla, P., Teterwak, P., Wang, C., Sarna, A., Tian, Y., Isola, P., Maschinot, A., Liu, C., & Krishnan, D. (2020). Supervised contrastive learning. *Advances in Neural Information Processing Systems*, 33.
- Kim, Y., Yim, J., Yun, J., & Kim, J. (2019). Nlnl: Negative learning for noisy labels. In *Proceedings of the IEEE/CVF international conference on computer vision* (pp. 101–110).
- Kingma, D. P., & Welling, M. (2013). Auto-encoding variational bayes. arXiv preprint [arXiv:1312.6114](https://arxiv.org/abs/1312.6114)
- Kiros, R., Zhu, Y., Salakhutdinov, R., Zemel, R. S., Torralba, A., Urtasun, R., & Fidler, S. (2015). Skip-thought vectors. arXiv preprint [arXiv:1506.06726](https://arxiv.org/abs/1506.06726)
- Kiselev, V. Y., Andrews, T. S., & Hemberg, M. (2019). Challenges in unsupervised clustering of single-cell rna-seq data. *Nature Reviews Genetics*, 20(5), 273–282.
- Krizhevsky, A., & Hinton, G. (2009). *Learning multiple layers of features from tiny images*. Master's thesis. Department of Computer Science, University of Toronto.
- Le, Q., & Mikolov, T. (2014). Distributed representations of sentences and documents. In *International conference on machine learning, PMLR* (pp. 1188–1196).
- Li, T., & Ding, C. (2006). The relationships among various nonnegative matrix factorization methods for clustering. In *Sixth International Conference on Data Mining (ICDM'06), IEEE* (pp. 362–371).
- Li, Y., Hu, P., Liu, Z., Peng, D., Zhou, J. T., & Peng, X. (2021b). *Contrastive clustering* (Vol. 35).
- Li, X., Zhang, R., Wang, Q., & Zhang, H. (2020). Autoencoder constrained clustering with adaptive neighbors. *IEEE Transactions on Neural Networks and Learning Systems*, 1–7.
- Li, J., Zhou, P., Xiong, C., & Hoi, S. C. (2021a). Prototypical contrastive learning of unsupervised representations. In *ICLR*.
- Liu, X., Dou, Y., Yin, J., Wang, L., & Zhu, E. (2016). Multiple kernel k-means clustering with matrix-induced regularization. In *Proceedings of the thirtieth AAAI conference on artificial intelligence* (pp. 1888–1894).
- Liu, W., Shen, X., & Tsang, I. (2017). Sparse embedded k-means clustering. In *Advances in neural information processing systems* (pp. 3319–3327).
- Ma, E. (2019). Nlp augmentation. <https://github.com/makcedward/nlpaug>
- MacQueen, J., et al. (1967). Some methods for classification and analysis of multivariate observations. In: *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability, Oakland, CA, USA* (Vol. 1, pp. 281–297).
- Nie, F., Wang, C. L., & Li, X. (2019). K-multiple-means: A multiple-means clustering method with specified k clusters. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining* (pp. 959–967).
- Nie, F., Zeng, Z., Tsang, I. W., Xu, D., & Zhang, C. (2011). Spectral embedded clustering: A framework for in-sample and out-of-sample spectral clustering. *IEEE Transactions on Neural Networks*, 22(11), 1796–1808.
- Niu, C., & Wang, G. (2021). Spice: Semantic pseudo-labeling for image clustering. arXiv preprint [arXiv:2103.09382](https://arxiv.org/abs/2103.09382)
- Oord, A., Li, Y., & Vinyals, O. (2018). Representation learning with contrastive predictive coding. arXiv preprint [arXiv:1807.03748](https://arxiv.org/abs/1807.03748)
- Park, S., Han, S., Kim, S., Kim, D., Park, S., Hong, S., & Cha, M. (2020). Improving unsupervised image clustering with robust learning. arXiv preprint [arXiv:2012.11150](https://arxiv.org/abs/2012.11150)
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., & Antiga, L., ... & Chintala, S. (2019). Pytorch: An imperative style, high-performance deep learning library. arXiv preprint [arXiv:1912.01703](https://arxiv.org/abs/1912.01703)
- Pelleg, D., & Moore, A. W. (2000). X-means: Extending k-means with efficient estimation of the number of clusters. *Icml*, 1, 727–734.
- Peng, X., Xiao, S., Feng, J., Yau, W. Y., & Yi, Z. (2016). Deep subspace clustering with sparsity prior. In *IJCAI* (pp. 1925–1931).
- Peng, X., Yi, Z., & Tang, H. (2015). Robust subspace clustering via thresholding ridge regression. *AAAI*, 25, 3827–3833.
- Peng, X., Zhu, H., Feng, J., Shen, C., Zhang, H., & Zhou, J. T. (2019). Deep clustering with sample-assignment invariance prior. *IEEE Transactions on Neural Networks and Learning Systems*, 31(11), 4857–4868.
- Radford, A., Metz, L., & Chintala, S. (2015). Unsupervised representation learning with deep convolutional generative adversarial networks. arXiv preprint [arXiv:1511.06434](https://arxiv.org/abs/1511.06434)
- Rakib, M. R. H., Zeh, N., Jankowska, M., & Milios, E. (2020). Enhancement of short text clustering by iterative classification. In *International conference on applications of natural language to information systems* (pp. 105–117). Springer.
- Reimers, N., & Gurevych, I. (2019). Sentence-bert: Sentence embeddings using siamese bert-networks. arXiv preprint [arXiv:1908.10084](https://arxiv.org/abs/1908.10084)
- Shen, S., Li, W., Zhu, Z., Huang, G., Du, D., Lu, J., & Zhou, J. (2021). Structure-aware face clustering on a large-scale graph with 10⁷ nodes. arXiv preprint [arXiv:2103.13225](https://arxiv.org/abs/2103.13225)
- Socher, R., Pennington, J., Huang, E. H., Ng, A. Y., & Manning, C. D. (2011). Semi-supervised recursive autoencoders for predicting sentiment distributions. In *Proceedings of the 2011 conference on empirical methods in natural language processing* (pp. 151–161).
- Strehl, A., & Ghosh, J. (2002). Cluster ensembles—a knowledge reuse framework for combining multiple partitions. *Journal of Machine Learning Research*, 3, 583–617.
- Tang, M., Marin, D., Ayed, I. B., & Boykov, Y. (2019). Kernel cuts: Kernel and spectral clustering meet regularization. *International Journal of Computer Vision*, 127(5), 477–511.
- Thanh, N. D., & Ali, M. (2017). Neutrosophic recommender system for medical diagnosis based on algebraic similarity measure and clustering. In *2017 IEEE international conference on fuzzy systems (FUZZ-IEEE), IEEE* (pp. 1–6).
- Van der Maaten, L., & Hinton, G. (2008). Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(11).
- Van Gansbeke, W., Vandenhende, S., Georgoulis, S., & Van Gool, L. (2021). Revisiting contrastive methods for unsupervised learning of visual representations. arXiv preprint [arXiv:2106.05967](https://arxiv.org/abs/2106.05967)
- Van Gansbeke, W., Vandenhende, S., Georgoulis, S., Proesmans, M., Van Gool, L. (2020). Scan: Learning to classify images without labels. In *European conference on computer vision* (pp. 268–285). Springer.
- Vincent, P., Larochelle, H., Lajoie, I., Bengio, Y., Manzagol, P. A., & Bottou, L. (2010). Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *Journal of Machine Learning Research*, 11(12).
- Wang, X., & Qi, G. J. (2021). Contrastive learning with stronger augmentations. arXiv preprint [arXiv:2104.07713](https://arxiv.org/abs/2104.07713)
- Wang, X., Liu, Z., & Yu, S. X. (2021). Unsupervised feature learning by cross-level instance-group discrimination. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 12586–12595).

- Wei, J., & Zou, K. (2019). Eda: Easy data augmentation techniques for boosting performance on text classification tasks. arXiv preprint [arXiv:1901.11196](https://arxiv.org/abs/1901.11196)
- Wu, J., Long, K., Wang, F., Qian, C., Li, C., Lin, Z., & Zha, H. (2019). Deep comprehensive correlation mining for image clustering. In *Proceedings of the IEEE international conference on computer vision* (pp. 8150–8159).
- Xie, J., Girshick, R., & Farhadi, A. (2016). Unsupervised deep embedding for clustering analysis. In *International conference on machine learning* (pp. 478–487).
- Xu, J., Xu, B., Wang, P., Zheng, S., Tian, G., & Zhao, J. (2017). Self-taught convolutional neural networks for short text clustering. *Neural Networks*, 88, 22–31.
- Xu, J., Xu, B., Wang, P., Zheng, S., Tian, G., & Zhao, J. (2017). Self-taught convolutional neural networks for short text clustering. *Neural Networks*, 88, 22–31.
- Yang, J., Parikh, D., & Batra, D. (2016). Joint unsupervised learning of deep representations and image clusters. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 5147–5156).
- Yin, J., & Wang, J. (2016). A model-based approach for text clustering with outlier detection. In *2016 IEEE 32nd international conference on data engineering (ICDE), IEEE* (pp. 625–636).
- Zbontar, J., Jing, L., Misra, I., LeCun, Y., & Deny, S. (2021). Barlow twins: Self-supervised learning via redundancy reduction. arXiv preprint [arXiv:2103.03230](https://arxiv.org/abs/2103.03230)
- Zeiler, M. D., Krishnan, D., Taylor, G. W., & Fergus, R. (2010). Deconvolutional networks. In *2010 IEEE computer society conference on computer vision and pattern recognition, IEEE* (pp. 2528–2535).
- Zelnik-Manor, L., & Perona, P. (2005). Self-tuning spectral clustering. *Advances in Neural Information Processing Systems*, 1601–1608.
- Zhang, D., Nan, F., Wei, X., Li, S., Zhu, H., McKeown, K., Nallapati, R., Arnold, A., & Xiang, B. (2021a). Supporting clustering with contrastive learning. arXiv preprint [arXiv:2103.12953](https://arxiv.org/abs/2103.12953)
- Zhang, D., Nan, F., Wei, X., Li, S., Zhu, H., McKeown, K., Nallapati, R., Arnold, A., & Xiang, B. (2021b). Supporting clustering with contrastive learning. arXiv preprint [arXiv:2103.12953](https://arxiv.org/abs/2103.12953)
- Zhang, W., Wang, X., Zhao, D., & Tang, X. (2012). Graph degree linkage: Agglomerative clustering on a directed graph. In *European conference on computer vision* (pp. 428–441). Springer.
- Zhong, H., Chen, C., Jin, Z., & Hua, X. S. (2020). Deep robust clustering by contrastive learning. arXiv preprint [arXiv:2008.03030](https://arxiv.org/abs/2008.03030)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.