

Modeling Inter-Class and Intra-Class Constraints in Novel Class Discovery

Wenbin Li¹, Zhichen Fan¹, Jing Huo^{1*}, Yang Gao¹

¹State Key Laboratory for Novel Software Technology, Nanjing University, China

Abstract

Novel class discovery (NCD) aims at learning a model that transfers the common knowledge from a class-disjoint labelled dataset to another unlabelled dataset and discovers new classes (clusters) within it. Many methods, as well as elaborate training pipelines and appropriate objectives, have been proposed and considerably boosted performance on NCD tasks. Despite all this, we find that the existing methods do not sufficiently take advantage of the essence of the NCD setting. To this end, in this paper, we propose to model both inter-class and intra-class constraints in NCD based on the symmetric Kullback-Leibler divergence (sKLD). Specifically, we propose an inter-class sKLD constraint to effectively exploit the disjoint relationship between labelled and unlabelled classes, enforcing the separability for different classes in the embedding space. In addition, we present an intra-class sKLD constraint to explicitly constrain the intra-relationship between a sample and its augmentations and ensure the stability of the training process at the same time. We conduct extensive experiments on the popular CIFAR10, CIFAR100 and ImageNet benchmarks and successfully demonstrate that our method can establish a new state of the art and can achieve significant performance improvements, e.g., 3.5%/3.7% clustering accuracy improvements on CIFAR100-50 dataset split under the task-aware-/agnostic evaluation protocol, over previous state-of-the-art methods. Code is available at <https://github.com/FanZhichen/NCD-IIC>.

1. Introduction

Deep learning has made great progress and achieved remarkable results in many computer vision fields, especially in image classification [13, 16, 22, 25, 27]. Unfortunately, these successes of deep learning heavily rely on a large amount of fully labelled data for training. On the other hand, in many realistic scenarios, it is difficult to collect or to annotate such a large-scale dataset. To address this problem, a new paradigm of *novel class discovery (NCD)* has

been proposed and attracted increasing attention in recent years [8, 10, 11, 33, 35, 37].

The goal of NCD is to train a classification model on a labelled dataset and simultaneously transfer the latent common knowledge to discover new classes (or clusters) in another unlabelled dataset. Different from *semi-supervised learning* [1, 2, 26, 36] that assumes the labelled and unlabelled datasets share the same label space, in the setting of NCD, the classes of the unlabelled dataset are disjoint with those of the labelled dataset, which is more challenging. In addition, NCD is also different from the generic *clustering* [3, 4, 18, 31] in that an additional labelled dataset is available in NCD. In general, for the standard clustering methods, the clustering results are not unique. That is to say, there may be multiple different and approximately correct results for a certain unlabelled dataset. In contrast, thanks to the available labelled dataset, NCD can eliminate the semantic ambiguity with the label guidance and finally makes the clustering be consistent with the real visual semantics [11]. Clearly, NCD is more realistic and more practical than unsupervised clustering.

In general, the existing NCD methods can be roughly divided into two categories, *i.e.*, two-stage based methods and single-stage based methods [20]. Most of the early methods are two-stage by using labelled and unlabelled data in different stages, such as KCL [14], MCL [15] and DTC [11]. Typically, the two-stage based methods first learn an embedding network on the labelled set through supervised learning and then use it on the unlabelled set to discover new clusters with little modifications. In contrast, the latest methods are almost single-stage, such as RS [10], NCL [37], UNO [8], DualRank [35] and ComEx [33], which use both labelled and unlabelled data in a single stage at the same time. The single-stage based methods can learn the feature representation and discover novel classes simultaneously, iteratively updating the learned feature embedding network and clustering results during the training process. Compared with the two-stage based methods, the single-stage methods can make more effective use of the similarity between labelled and unlabelled classes to achieve a better knowledge transfer between these two datasets. Therefore, in this paper, we will mainly focus on

*Corresponding author

the single-stage direction of NCD.

However, we find that the current single-stage based NCD methods do not sufficiently take advantage of the essence of the NCD setting, that is to say, overlooking the disjoint characteristic between the labelled and unlabelled classes. In this sense, on one hand, the labelled and unlabelled samples cannot be effectively separated, weakening the discriminability of the learned features. On the other hand, because the labelled data is learned under supervision while the unlabelled data has no supervision, *i.e.*, an imbalanced learning process (learning with different supervision strengths), it will make the learned feature representations biased toward the labelled data. In addition, we notice that although some methods [10, 35] have used data augmentation to generate additional samples and gained significant performance improvements, they generally employ the mean squared error (MSE) as the consistency regularization, which cannot constrain the consistency well with a good generalization ability.

To address the above two issues, we propose to model both *Inter-class and Intra-class Constraints* (IIC for short) built on the *symmetric Kullback-Leibler divergence (sKLD)* for discovering the novel classes. To be specific, an *inter-class sKLD constraint* is proposed to explicitly learn to separate different classes between labelled and unlabelled data, enhancing the discriminability of learned feature representations. Moreover, an *intra-class sKLD constraint* is presented to fully learn the intra-relationship between samples and their augmentations. According to our experiments, such an intra-class sKLD constraint can also stable the training process in the training phase. We have conducted extensive experiments on three benchmarks, including CIFAR10 [21], CIFAR100 [21] and ImageNet [7], and show that the proposed two constraints can significantly and consistently outperform the existing novel class discovery methods by a large margin.

To summarize, our contributions are as follows:

- We propose a new *inter-class Kullback-Leibler divergence constraint* to sufficiently model the relationship between the labelled and unlabelled datasets to learn more discriminative feature representations, which is somewhat overlooked in the literature.
- We propose a new *intra-class Kullback-Leibler divergence constraint* to effectively exploit the relationship between a sample and its different transformations to learn invariant feature representations.
- We evaluate the proposed constraints on three benchmark datasets for novel class discovery and obtain significant performance improvements over the state-of-the-art methods, which successfully demonstrates the effectiveness of the proposed method.

2. Related Work

Novel class discovery (NCD) is a new task attracted wide attention in recent years, which aims at discovering new classes in an unlabelled dataset given a class-disjoint labelled dataset as supervision. A variety of advanced NCD methods have been proposed and have tangibly improved the clustering performance on multiple benchmark datasets.

The early methods of NCD include KCL [14], MCL [15] and DTC [11]. In general, these methods first learn an embedding network of feature representations on the labelled data, and then use it directly for the unlabelled data. Specifically, KCL and MCL propose a framework for both cross-domain and cross-task transfer learning that leverages the pairwise similarity to represent categorical information, and learn the clustering network based on the pairwise similarity prediction through different objective functions, respectively. DTC extends the deep embedding clustering method [31] into a transfer learning setting and proposes a two-stage method. Importantly, Han *et al.* [11] formalize the task of *novel class discovery* for the first time.

Since then, the current NCD methods [8, 10, 19, 20, 33, 35, 37, 38] are almost single-stage and can take greater advantage of both labelled and unlabelled data. RS [10] introduces a three-step learning pipeline, which first trains the representation network with all labelled and unlabelled samples using self-supervised learning, and then uses ranking statistics to obtain pairwise similarity between unlabelled samples, and finally use the pairwise similarity to discover novel classes. DualRank [35] expands RS to a two-branch framework from both global and local levels. Similarly, DualRank uses dual ranking statistics and mutual knowledge distillation to generate pseudo labels and ensure the consistency between two branches. In order to generate pairwise pseudo labels, Joint [19] employs a Winner-Take-All (WTA) hashing algorithm [32] on the shared feature space for NCD.

NCL [37] and OpenMix [38] are largely motivated by contrastive learning [5, 12] and Mixup [34], respectively. NCL introduces the contrastive loss to learn more discriminative representations. On the other hand, OpenMix uses Mixup to mix labelled and unlabelled samples, building a learnable relationship between the two parts of data. Instead of using multiple objectives, UNO [8] introduces a unified objective function to transfer knowledge from the labelled set to unlabelled set. More recently, Joseph *et al.* [20] categorize the existing NCD methods into two classes (*i.e.*, two- and single-stage based methods), according to whether the labelled and unlabelled samples are available at the same time or not. They also propose a spacing loss to enforce separability between labelled and unlabelled points in the embedding space. ComEx [33] focuses on the generalized NCD (GNCD), aka generalized category discovery (GCD) [29], and proposes two groups of compositional ex-

perts to solve this problem.

To our best knowledge, the existing approaches do not make full use of the disjoint characteristic between labelled and unlabelled classes. In addition, we find that some methods utilizes the mean squared error (MSE) to constrain the learned representations of data augmentation, but it cannot achieve the desired effect. Instead, we propose inter-class and intra-class constraints based on the symmetric Kullback-Leibler divergence (sKLD) for NCD.

3. Method

3.1. Problem Formulation

In the NCD setting, given a labelled dataset $\mathcal{D}^l = \{(\mathbf{x}_1^l, y_1^l), \dots, (\mathbf{x}_N^l, y_N^l)\}$, the goal is to automatically discover C^u clusters (or classes) in an unlabelled dataset $\mathcal{D}^u = \{\mathbf{x}_1^u, \dots, \mathbf{x}_M^u\}$, where each \mathbf{x}_i^l in \mathcal{D}^l or \mathbf{x}_i^u in \mathcal{D}^u is an image and $y_i^l \in \mathcal{Y} = \{1, \dots, C^l\}$ is the corresponding class label of \mathbf{x}_i^l . In particular, we assume that the set of C^l labelled classes is disjoint with the set of C^u unlabelled classes. In this sense, the core of NCD is how to effectively learn transferable semantic knowledge from the disjoint labelled dataset \mathcal{D}^l to help performing clustering on the unlabelled dataset \mathcal{D}^u . Following the literature [8, 10, 37], we also assume the number of unlabelled classes C^u is known a priori in this paper.

To tackle this challenge problem, we propose two *symmetric Kullback-Leibler divergence (sKLD)* based constraints from both *inter-class* and *intra-class* perspectives to learn more discriminative feature representations for NCD models (see Fig. 1). In the following sections, we first introduce the *inter-class sKLD constraint* and *intra-class sKLD constraint* in Sec. 3.2, and then summarize the overall objective for training in Sec. 3.3.

3.2. Symmetric Kullback-Leibler Divergence for Novel Class Discovery

According to the above analyses, to effectively utilize the two parts of data in NCD, *i.e.*, a labelled set and an unlabelled set, we develop two *inter-class* and *intra-class symmetric Kullback-Leibler divergence (sKLD) constraints* to better accomplish the NCD task, especially in the single-stage paradigm.

Following UNO [8], as shown in Fig. 1, the architecture of our model consists of two parts: an encoder E and two classification heads h and g . The encoder E is implemented as a standard convolutional neural network (CNN), which converts an input images into a feature vector. The head h belonging to the labelled data is implemented as a linear classifier with C^l output neurons, and the head g belonging to the unlabelled data is composed of a multi-layer perceptron (MLP) and a linear classifier with C^u output neurons. In the training phase, each sample \mathbf{x}_i will

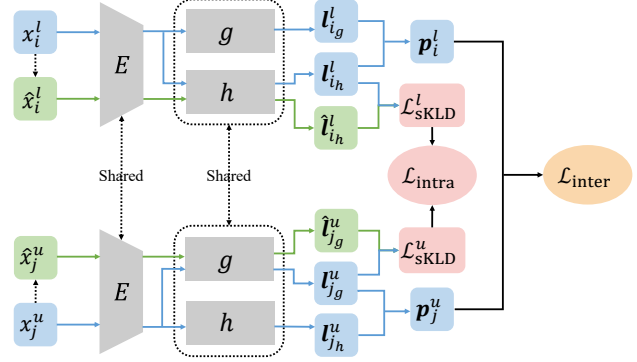


Figure 1. Architecture of the proposed method. We present the “raw samples” (labelled sample \mathbf{x}_i^l , unlabelled sample \mathbf{x}_j^u , corresponding logits and probability distributions) in blue and the “augmented counterparts” (labelled counterpart $\hat{\mathbf{x}}_i^l$, unlabelled counterpart $\hat{\mathbf{x}}_j^u$, corresponding logits and probability distributions) in green. The samples and their augmentations are inputted into the shared encoder E , and then fed into two classification heads h and g , obtaining predictions to calculate both inter-class sKLD $\mathcal{L}_{\text{inter}}$ and intra-class sKLD $\mathcal{L}_{\text{intra}}$. For brevity, we omit the calculation process of the standard cross-entropy loss \mathcal{L}_{CE} .

be first encoded as a feature vector by E , and then will be passed through both classification heads to obtain the corresponding logits $\mathbf{l}_{i_h} \in \mathbb{R}^{C^l}$ and $\mathbf{l}_{i_g} \in \mathbb{R}^{C^u}$, respectively. After that, the two logits are concatenated together as $\mathbf{l}_i = [\mathbf{l}_{i_h}, \mathbf{l}_{i_g}] \in \mathbb{R}^{C^l + C^u}$ and fed into a softmax layer σ with a temperature τ , obtaining the probability distribution $\mathbf{p}_i = \sigma(\mathbf{l}_i / \tau)$.

Inter-class Symmetric KLD Constraint. When solving the NCD problem, in the pipeline of the single-stage based methods, both labelled and unlabelled images will be accessed in each mini-batch during training. Although the distributions of these two parts of images are similar, in fact, the representations of them should be different from each other as much as possible, *i.e.*, the separability between the labelled and unlabelled classes. However, this point is somewhat overlooked in the existing single-stage based methods. Therefore, to address this issue, we propose an *inter-class sKLD constraint* to explicitly enlarge the distance between each labelled sample and each unlabelled sample in the current mini-batch using a symmetric Kullback-Leibler divergence distance. The formulation between a pair of samples is as follows:

$$\mathcal{L}_{\text{sKLD}} = \frac{1}{2} (D_{\text{KL}}(\mathbf{p}_i^l \| \mathbf{p}_j^u) + D_{\text{KL}}(\mathbf{p}_j^u \| \mathbf{p}_i^l)), \quad (1)$$

where \mathbf{p}_i^l and \mathbf{p}_j^u are the probability distributions generated for the labelled image $\mathbf{x}_i^l |_{i=1}^N$ and unlabelled image $\mathbf{x}_j^u |_{j=1}^M$ in a mini-batch, respectively. D_{KL} is the Kullback-Leibler

(KL) divergence defined as

$$D_{\text{KL}}(\mathbf{p}_i^l || \mathbf{p}_j^u) = \sum_{k=1}^{C^l+C^u} \mathbf{p}_i^l(k) \log \frac{\mathbf{p}_i^l(k)}{\mathbf{p}_j^u(k)}, \quad (2)$$

$$D_{\text{KL}}(\mathbf{p}_j^u || \mathbf{p}_i^l) = \sum_{k=1}^{C^l+C^u} \mathbf{p}_j^u(k) \log \frac{\mathbf{p}_j^u(k)}{\mathbf{p}_i^l(k)}. \quad (3)$$

Therefore, the inter-class sKLD over a mini-batch is

$$\mathcal{L}_{\text{inter-class}} = \frac{1}{NM} \sum_{i=1}^N \sum_{j=1}^M \mathcal{L}_{\text{sKLD}}, \quad (4)$$

where N and M are the numbers of labelled and unlabelled images in a mini-batch, respectively.

Intra-class Symmetric KLD Constraint. Recently, some NCD methods [8, 38] have introduced data-augmentation techniques, normally used in self-supervised learning [9], to help improve the clustering accuracy when discovering novel classes. In particular, UNO [8] leverages a multi-view strategy built on data augmentation and uses a swapped prediction mechanism through the Sinkhorn-Knopp algorithm [6] to generate pseudo labels for augmented unlabelled images. However, UNO does not consider the intra-class constraint between the augmented images, no matter labelled nor unlabelled. We argue that for different views (augmentations) of the same image, the outputs produced by the corresponding classification heads should be consistent. That is to say, the distance between any two probability distributions of different augmentations of the same image should be small.

Note that, because the logits of unlabelled images passing through the labelled head h and the logits of labelled images passing through the unlabelled head g will gradually approach to a zero vector, it is not necessary to calculate the consistency regularization loss for these two kinds of logits. In this paper, we only consider the logits of labelled images generated by the labelled head h and the logits of unlabelled images generated by the unlabelled head g :

$$\mathcal{L}_{\text{sKLD}}^l = \frac{1}{2} (D_{\text{KL}}(\mathbf{p}_{i_h}^l || \hat{\mathbf{p}}_{i_h}^l) + D_{\text{KL}}(\hat{\mathbf{p}}_{i_h}^l || \mathbf{p}_{i_h}^l)), \quad (5)$$

$$\mathcal{L}_{\text{sKLD}}^u = \frac{1}{2} (D_{\text{KL}}(\mathbf{p}_{j_g}^u || \hat{\mathbf{p}}_{j_g}^u) + D_{\text{KL}}(\hat{\mathbf{p}}_{j_g}^u || \mathbf{p}_{j_g}^u)), \quad (6)$$

$$\mathcal{L}_{\text{intra-class}} = \frac{1}{N} \sum_{i=1}^N \mathcal{L}_{\text{sKLD}}^l + \frac{1}{M} \sum_{j=1}^M \mathcal{L}_{\text{sKLD}}^u, \quad (7)$$

where $\mathbf{p}_{i_h}^l \in \mathbb{R}^{C^l}$ and $\hat{\mathbf{p}}_{i_h}^l \in \mathbb{R}^{C^l}$ indicate the probability distributions of the labelled image \mathbf{x}_i^l and its augmentation $\hat{\mathbf{x}}_i^l$ through the labelled head h and the softmax layer σ , respectively. Similarly, $\mathbf{p}_{j_g}^u \in \mathbb{R}^{C^u}$ and $\hat{\mathbf{p}}_{j_g}^u \in \mathbb{R}^{C^u}$ represent the probability distributions corresponding to the unlabelled image \mathbf{x}_j^u and its augmentation $\hat{\mathbf{x}}_j^u$, respectively.

Table 1. Details of dataset splits used in the experiments.

Dataset split	Labelled		Unlabelled	
	#Images	#Classes	#Images	#Classes
CIFAR10	25K	5	25K	5
CIFAR100-20	40K	80	10K	20
CIFAR100-50	25K	50	25K	50
ImageNet	1.25M	882	≈30K	30

3.3. Overall Objective

In addition to the two sKLD constraint loss terms mentioned in the above subsection, we also use the standard cross-entropy (CE) loss:

$$\mathcal{L}_{\text{CE}} = -\frac{1}{N+M} \sum_{i=1}^{N+M} \sum_{k=1}^{C^l+C^u} \mathbf{y}_i(k) \log \mathbf{p}_i(k), \quad (8)$$

where $\mathbf{y}_i(k)$ indicates the zero-padded ground truth label of image \mathbf{x}_i if \mathbf{x}_i is a labelled image, otherwise it is the zero-padded pseudo label of \mathbf{x}_i , predicted by the Sinkhorn-Knopp algorithm like UNO, and \mathbf{p}_i is the predicted probability distribution corresponding to \mathbf{x}_i . In summary, with the purpose of maximizing the inter-class sKLD and minimizing the intra-class sKLD at the same time, the overall objective function in our model is

$$\mathcal{L} = \mathcal{L}_{\text{CE}} - \alpha \mathcal{L}_{\text{inter-class}} + \beta \mathcal{L}_{\text{intra-class}}, \quad (9)$$

where $\alpha, \beta > 0$ are two hyperparameters of the two sKLD loss terms.

4. Experiments

4.1. Experimental Setup

Datasets. We perform our experiments on three benchmark datasets that are widely-used in the field of NCD, including CIFAR10 [21], CIFAR100 [21] and ImageNet [7]. Following the literature of NCD [10, 33, 37], each dataset is further divided into two parts for the NCD setting: a labelled subset and an unlabelled subset. Also, we assume the number of classes in the unlabelled subset is known a priori. In addition, according to the number of classes contained in the subsets, there are four different dataset splits: CIFAR10, CIFAR100-20, CIFAR100-50 and ImageNet, where CIFAR100-50 is first introduced in UNO [8] as a more challenging evaluation than CIFAR100-20. More details of the dataset splits are shown in Tab. 1.

Evaluation Metrics. Following the evaluation protocols used in the literature [8, 33], we also conduct our experiments using both task-aware and task-agnostic evaluation protocols. In the task-aware protocol, we know the task information in advance when evaluating a method, that is, we

Table 2. Ablation study of our method on three dataset splits. Results are reported in ACC (%), NMI and ARI that are averaged over 5 runs. “Inter-class” stands for the inter-class sKLD loss term and “Intra-class” is short for the intra-class sKLD. **Best** results are highlighted in each column.

#	Inter-class	Intra-class	CIFAR10			CIFAR100-20			CIFAR100-50		
			ACC	NMI	ARI	ACC	NMI	ARI	ACC	NMI	ARI
1	✗	✗	93.25	0.8717	0.8565	90.54	0.8475	0.8123	62.27	0.6781	0.4688
2	✓	✗	99.07	0.9644	0.9770	92.17	0.8695	0.8448	64.31	0.7046	0.5102
3	✗	✓	93.60	0.8753	0.8624	91.13	0.8506	0.8193	63.35	0.6800	0.4729
4	✓	✓	99.11	0.9657	0.9780	92.48	0.8727	0.8508	65.85	0.7106	0.5238

know whether an image is from the labelled subset or the unlabelled one. In contrast, in the task-agnostic protocol, we have no idea about such information. Following the literature, we use the average clustering accuracy (ACC) to evaluate the performance of our method, defined as

$$ACC = \max_{\text{perm} \in P} \frac{1}{N} \sum_{i=1}^N \mathbb{1}\{y_i = \text{perm}(\hat{y}_i)\}, \quad (10)$$

where y_i and \hat{y}_i are the ground-truth label and clustering assignment of a test sample $x_i^u \in D^u$, respectively. P is the set of all permutations and the optimal permutation can be calculated by the Hungarian algorithm [23].

Implementation Details. For a fair comparison with the existing methods, we use ResNet-18 [13] as the backbone and two steps to train our model. We first pre-train the model for 200 epochs on the labelled subset, and then fine-tune the pre-trained model for 500 epochs to discover novel classes. For all experiments, we fix the batch size to 512. Note that the proposed inter-class and intra-class symmetric Kullback-Leibler divergence (sKLD) terms are only used in the training phase and will be discarded in the test phase. Specifically, for the inter-class sKLD, we set the weight α to 0.02 for CIFAR100-20 and 0.05 for the other three dataset splits. As for the intra-class sKLD, we fix $\beta = 0.01$ on all benchmarks. Following the literature [3, 8, 18, 33], we use multi-head clustering and overclustering to boost the clustering performance. Besides, we apply multi-crop strategy and common data augmentations (*e.g.*, random crop, flip, color jittering, and grey-scale) to generate transformed samples. In addition, we conduct all the experiments on the Tesla V100 GPUs.

4.2. Ablation Study

Symmetric KLD vs. MSE for Intra-class Constraint.

As mentioned above, many previous works [8, 33, 37] have used data augmentation (*e.g.*, rotation, flipping and random crop) to generate randomly transformed samples for obtaining pseudo labels. To further ensure consistent predictions between different augmentations of the same image, *i.e.*, intra-class constraint, RS [10] and DualRank [35] intro-

Table 3. Ablation study on the intra-class constraint on CIFAR-10 and CIFAR-100. “sKLD” and “MSE” denote the intra-class sKLD and MSE consistency regularization term, respectively. Results are reported in ACC (%) averaged over 5 runs. **Best** results are highlighted in each column.

Method	CIFAR10	CIFAR100-20	CIFAR100-50
Baseline	93.25	90.54	62.27
+ MSE	93.11	90.96	62.47
+ sKLD	93.60	91.13	63.35

duce the mean squared error (MSE) as a regularization term, which is a common operation in semi-supervised learning. However, we find that the MSE-based intra-class constraint does not perform well in practice. Therefore, we employ an symmetric KLD measure instead of the common MSE to constrain such an intra-class consistency.

To verify the effectiveness of symmetric KLD, we compare the proposed intra-class sKLD with MSE and report the results averaged over 5 runs in Tab. 3. Note that the hyperparameters of both sKLD and MSE have been well selected through grid search. From the results, we can see that the MSE constraint can only slightly improve the performance of Baseline on CIFAR100-20 and CIFAR100-50, and even somewhat damages the performance on CIFAR10. On the contrary, the proposed sKLD constraint could consistently boost the performance of Baseline on all benchmarks. The intuitive reason is that the semantics of different augmentations (especially strong data augmentations) of the same image has been changed to some extent. Therefore, we should not strictly make their predictions be the same did as the MSE constraint. In contrast, the slacker sKLD constraint just making the predictions be similar will be more beneficial to the final clustering.

Symmetric KLD Constraints. In this subsection, we perform ablation study experiments on three dataset splits, including CIFAR10, CIFAR100-20 and CIFAR100-50, to evaluate the effectiveness of the proposed *inter-class sKLD* and *intra-class sKLD* components. Besides the ACC, we also report two extra evaluation criteria that are commonly

Table 4. Comparison with state-of-the-art methods on the training split of the unlabelled subset, using task-aware evaluation protocol. Results are reported in clustering accuracy (%) in the form of mean and standard deviation (averaged over 5 runs except for ImageNet). **Best** results are highlighted in each column. †Our reproduced result.

Method	Venue	Type	CIFAR10	CIFAR100-20	CIFAR100-50	ImageNet
<i>k</i> -means	Classic	-	72.5±0.0	56.3±1.7	28.3±0.7	71.9
KCL	ICLR'18	Two-stage	72.3±0.2	42.1±1.8	-	73.8
MCL	ICLR'19	Two-stage	70.9±0.1	21.5±2.3	-	74.4
DTC	ICCV'19	Two-stage	88.7±0.3	67.3±1.2	35.9±1.0	78.3
RS	ICLR'20	Single-stage	90.4±0.5	73.2±2.1	39.2±2.3	82.5
RS+	ICLR'20	Single-stage	91.7±0.9	75.2±4.2	44.1±3.7	82.5
OpenMix	CVPR'21	Single-stage	95.3	-	-	85.7
NCL	CVPR'21	Single-stage	93.4±0.5	86.6±0.4	-	90.7
Joint	ICCV'21	Single-stage	93.4±0.6	76.4±2.8	-	86.7
UNOv1	ICCV'21	Single-stage	96.1±0.5	85.0±0.6	52.9±1.4	90.6
UNOv2†	ICCV'21	Single-stage	93.3±0.4	90.5±0.7	62.3±1.4	90.7
DualRank	NeurIPS'21	Single-stage	91.6±0.6	75.3±2.3	-	88.9
ComEx	CVPR'22	Single-stage	93.6±0.3	85.7±0.7	53.4±1.3	90.9
IIC (Ours)	CVPR'23	Single-stage	99.1±0.0	92.4±0.2	65.8±0.9	91.9

Table 5. Comparison with state-of-the-art methods on CIFAR-10 and CIFAR-100, using task-agnostic evaluation protocol. Results are reported averaged over 5 runs in classification accuracy (%) and clustering accuracy (%) on test split of the labelled and unlabelled subsets, respectively. **Best** results are highlighted in each column.

Method	CIFAR10			CIFAR100-20			CIFAR100-50		
	Label	Unlabel	All	Label	Unlabel	All	Label	Unlabel	All
KCL	79.4	60.1	69.8	23.4	29.4	24.6	-	-	-
MCL	81.4	64.8	73.1	18.2	18.0	18.2	-	-	-
DTC	58.7	78.6	68.7	47.6	49.1	47.9	30.2	34.7	32.5
RS+	90.6	88.8	89.7	71.2	56.8	68.3	69.7	40.9	55.3
UNOv1	93.5	93.3	93.4	73.2	73.1	73.2	71.5	50.7	61.1
ComEx	95.0	92.6	93.8	75.2	77.3	75.6	75.3	53.5	64.4
IIC (Ours)	96.0	97.2	96.6	75.9	78.4	77.2	75.1	61.0	68.1

used in clustering: *i.e.*, normalized mutual information (NMI) [30] and adjusted rand index (ARI) [17]. MNI and ARI are generally used to measure the similarity between the clustering results and the ground-truth distributions. The closer the criteria are to 1, the better the clustering effect. The results on ablation study are shown in Tab. 2.

Considering that the trends of three evaluation criteria (*i.e.*, ACC, NMI and ARI) on all dataset splits are similar in Tab. 2, we will mainly take the ACC as an example for discussion. Specifically, as seen, the proposed inter-class sKLD alone can significantly improve the ACC of the baseline (*i.e.*, UNO) on all benchmarks, especially on CIFAR10 where the performance improvement is over 5.82%. Although the proposed intra-class sKLD alone do not distinctly boost the performance on CIFAR10, it tangibly gains 0.59% and 1.08% ACC improvements over the baseline on CIFAR100-20 and CIFAR100-50, respectively.

Notably, when combining the inter-class sKLD and

intra-class sKLD together, we observe that the results on all three dataset splits, especially on CIFAR100-20 and CIFAR100-50, are further improved, and the improvements are remarkable and stable. It is worth mentioning that because the CIFAR10 dataset is somewhat simple, the inter-class sKLD alone has made the ACC reach the saturation point and thus adding the additional intra-class sKLD does not bring distinct improvements. As expected, on the more challenging benchmarks, such as CIFAR100-20 and CIFAR100-50, integrating the inter-class sKLD with intra-class sKLD together can further obtains 0.31% and 1.54% ACC improvements over using the inter-class sKLD alone, respectively.

In summary, we are able to conclude that both inter-class and intra-class sKLD constraints are effective and each of them can consistently boost the performance of the baseline. In particular, it is clearly that the inter-class sKLD constraint alone can gain more improvements than the intra-




Figure 2. t-SNE visualization on CIFAR10. We compare our method with the reproduced UNOv2. Best viewed in color, each color indicates a different class and we also mark the index corresponding to each class.

class sKLD constraint alone. This is because the purpose of inter-class sKLD is to use prior information of NCD, *i.e.*, the unlabelled classes are naturally disjoint with the labelled classes, to learn a more discriminative embedding space. Differently, the intra-class sKLD is more like a consistency regularization term that is used to explicitly constrain samples and their augmentations. As discussed above, the intra-class sKLD, as a slacker regularization term than the MSE-based one, will bring a greater positive impact on the more difficult datasets or their splits.

4.3. Comparison with State of the Arts

We compare our proposed method with the current state-of-the-art methods in the field of NCD, including KCL [14], MCL [15], DTC [11], RS [10], RS+ (RS with incremental classifier) [10], OpenMix [38], DualRank [35], Joint [19], NCL [37], UNO [8] and ComEx [33] besides the classical k -means [24]. We report the results on four popular benchmarks in Tab. 4 and Tab. 5 using task-aware and task-agnostic evaluation protocols, respectively. It is worth noting that UNO has two versions, *i.e.*, UNOv1 and UNOv2. The results of UNOv1 are quoted from the original paper, while we report our reproduced results of UNOv2 using the officially released code.

In Tab. 4, we report the average clustering accuracy (*i.e.*, ACC) of all methods on the training split of the unlabelled set, using a task-aware evaluation protocol (commonly used in the literature). As we can see, our proposed method outperforms the current works on all benchmarks. In particular, our method surpasses the latest ComEx by a significant margin, such as, gaining 5.5% improvements on CIFAR10 and 12.4% improvements on CIFAR100-50, respectively. Compared with the reproduced UNOv2, our method still has a certain improvement, *i.e.*, about 5.8% improvements on CIFAR10, 1.9% improvements on CIFAR100-20, 3.5% improvements on CIFAR100-50 and 1.2% improvements on ImageNet. Note that the standard deviations of our method are also smaller than the previous works, indicating that the

proposed method does well in stability.

We also care about the performance on test split of the labelled and unlabelled sets with a task-agnostic evaluation protocol and report the results in terms of the classification accuracy and clustering accuracy in Tab. 5. As seen, our method outperforms the state-of-the-art methods on both labelled and unlabelled test set of all benchmarks except for CIFAR100-50, which successfully demonstrate its effectiveness and robustness. Specifically, for CIFAR10 and CIFAR100-20, our method not only is superior to the related works on the labelled set, but also obtains greater performance improvements on the unlabelled set, achieving the excellent comprehensive results. Moreover, for CIFAR100-50, our method is very close to the latest ComEx on the labelled subset (*i.e.*, with only 0.2% difference) and gains considerable improvements over ComEx on the unlabelled classes by 7.5%. Therefore, the results compared with the current state-of-the-art methods in Tab. 4 and Tab. 5 show that the proposed inter- and intra-class sKLDs indeed help to discriminate labelled points from unlabelled ones in the embedding space, both of which are simple yet effective.

4.4. Visualization

t-SNE Visualization. To better demonstrate the proposed method, we illustrate the t-SNE [28] visualization for all classes on CIFAR10, where the first five classes are labelled (*i.e.*, airplane, automobile, bird, cat and deer) and the remaining five classes are unlabelled (*i.e.*, dog, frog, horse, ship and truck). In Fig. 2, we compare our method with the reproduced UNOv2 [8] through its officially released code, to ensure a fair comparison, we follow the approach used in UNOv2 and use the concatenated logits l from both classification heads h and g to perform the t-SNE visualization.

As we can see, in general, two methods can well accomplish the NCD task, and there are only differences between them in the clustering ability on the certain classes. To be specific, for some labelled classes (*e.g.*, bird (2), cat (3) and

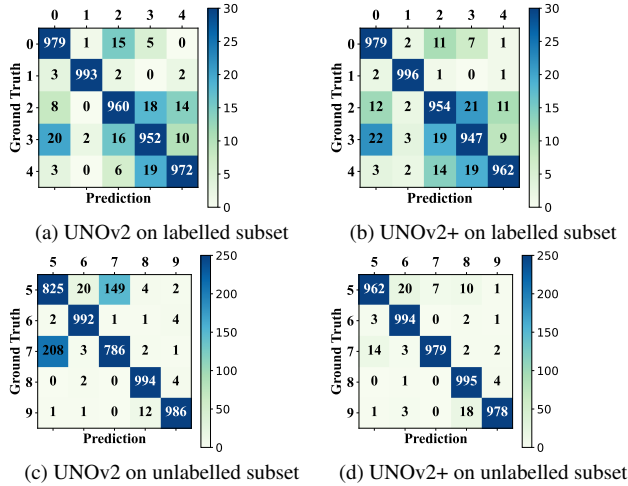


Figure 3. Confusion matrices of the reproduced UNO without (UNOV2) and with (UNOV2+) the inter-class sKLD constraint on CIFAR10, using the task-aware evaluation protocol. The two models are evaluated on test split of the labelled subset (Fig. 3a, Fig. 3b) and unlabelled subset (Fig. 3c, Fig. 3d), respectively.

deer (4)), UNOV2 shows more compact clustering results. Nevertheless, our method distinguishes better on the unlabelled classes (e.g., dog (5) and horse (7)), demonstrating that our method can learn more discriminative feature representations with better generalization. To further show this point, we draw attention to two unlabelled classes, i.e., dog (5) and horse (7). We can see that UNOV2 completely confuses these two classes, whose distributions are basically overlapped with each other in Fig. 2. We surmise the reason is that the poses of these two classes are visually similar, resulting in the poor clustering performance of UNOV2. In contrast, our method can clearly distinguish these two classes, because the proposed inter-class constraint can effectively separate different classes and the proposed intra-class constraint can compact the distribution of a class.

Confusion Matrix. Like RS [10] and UNO [8], we transform the original clustering task into a classification task for the NCD problem on the premise of knowing the number of categories of the unlabelled data. Therefore, we can validate the proposed method by the confusion matrix, which is commonly used in generic classification tasks. According to the ablation study in Sec. 4.2, it has been concluded that the proposed inter-class sKLD constraint plays a key role in our method. Therefore, it will be interesting to visually explain the effect of such a constraint. To be specific, we add the inter-class sKLD constraint into UNOV2 to obtain a new variant UNOV2+. Next, we establish the confusion matrices between the two models before and after using such an inter-class sKLD constraint.

The results are reported on the test split of CIFAR10 in Fig. 3 and Fig. 4 using task-aware and task-agnostic eval-

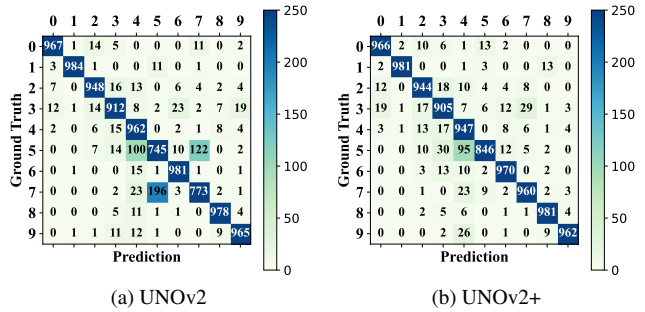


Figure 4. Confusion matrices of the reproduced UNO without (UNOV2) and with (UNOV2+) the inter-class sKLD constraint on the test split of CIFAR10, using the task-agnostic evaluation protocol, where each image is predicted into one of the 10 classes.

uation protocols, respectively. In Fig. 3, we compare the performance between UNOV2 and UNOV2+ on the labelled and unlabelled subsets. Specifically, we observe that the misclassification of UNOV2 is almost unchanged on the labelled subset after adding the inter-class sKLD constraint, but UNOV2+ obtains significant improvements for the unlabelled classes. With the help of the inter-class sKLD constraint, UNOV2+ will no longer confuse dog (5) and horse (7). In addition, Fig. 4 demonstrates that UNOV2+ can prevent misclassifying a labelled image into unlabelled classes or vice versa. These results show that the proposed inter-class sKLD constraint can effectively distinguish unlabelled data from labelled data to help discovering novel classes, and can be easily added into other single-stage based NCD methods as a plug-and-play loss term.

5. Conclusion

In this paper, we propose to model both inter-class and intra-class constraints built on the symmetric Kullback-Leibler divergence (sKLD) for novel class discovery (NCD). We conduct extensive experiments on four popular benchmarks and show that our method could outperform the existing state-of-the-art methods by a large margin. From the experimental results, we have the following findings: (1) making use of the disjoint characteristic between the labelled and unlabelled classes, i.e., constraining an inter-class constraint, is important and effective for NCD; (2) using a sKLD measure instead of the MSE for constraining the intra-class constraint is reasonable and beneficial to NCD.

Acknowledgements. This work is supported in part by the National Natural Science Foundation of China (62106100, 62192783, 62276128), Jiangsu Natural Science Foundation (BK20221441), CAAI-Huawei MindSpore Open Fund, the Collaborative Innovation Center of Novel Software Technology and Industrialization, and Jiangsu Provincial Double-Innovation Doctor Program (JSSCBS20210021).

References

- [1] David Berthelot, Nicholas Carlini, Ekin D. Cubuk, Alex Kurakin, Kihyuk Sohn, Han Zhang, and Colin Raffel. Remixmatch: Semi-supervised learning with distribution matching and augmentation anchoring. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2020.
- [2] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A Raffel. Mixmatch: A holistic approach to semi-supervised learning. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*, volume 32, 2019.
- [3] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 132–149, 2018.
- [4] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*, volume 33, pages 9912–9924, 2020.
- [5] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. A simple framework for contrastive learning of visual representations. In *Proceedings of the International Conference on Machine Learning (ICML)*, volume 119, pages 1597–1607, 2020.
- [6] Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*, pages 2292–2300, 2013.
- [7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 248–255, 2009.
- [8] Enrico Fini, Enver Sangineto, Stéphane Lathuilière, Zhun Zhong, Moin Nabi, and Elisa Ricci. A unified objective for novel class discovery. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9284–9292, 2021.
- [9] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2018.
- [10] Kai Han, Sylvestre-Alvise Rebuffi, Sébastien Ehrhardt, Andrea Vedaldi, and Andrew Zisserman. Automatically discovering and learning new visual categories with ranking statistics. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2020.
- [11] Kai Han, Andrea Vedaldi, and Andrew Zisserman. Learning to discover novel visual categories via deep transfer clustering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 8401–8409, 2019.
- [12] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9729–9738, 2020.
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016.
- [14] Yen-Chang Hsu, Zhaoyang Lv, and Zsolt Kira. Learning to cluster in order to transfer across domains and tasks. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2018.
- [15] Yen-Chang Hsu, Zhaoyang Lv, Joel Schlosser, Phillip Odom, and Zsolt Kira. Multi-class classification without multi-class labels. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019.
- [16] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4700–4708, 2017.
- [17] Lawrence Hubert and Phipps Arabie. Comparing partitions. *Journal of classification*, 2(1):193–218, 1985.
- [18] Xu Ji, Joao F Henriques, and Andrea Vedaldi. Invariant information clustering for unsupervised image classification and segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9865–9874, 2019.
- [19] Xuhui Jia, Kai Han, Yukun Zhu, and Bradley Green. Joint representation learning and novel category discovery on single-and multi-modal data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 610–619, 2021.
- [20] KJ Joseph, Sujoy Paul, Gaurav Aggarwal, Soma Biswas, Piyush Rai, Kai Han, and Vineeth N Balasubramanian. Spacing loss for discovering novel categories. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3761–3766, 2022.
- [21] Alex Krizhevsky, Geoffrey Hinton, et al. 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 *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*, volume 25, 2012.
- [23] Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955.
- [24] J MacQueen. Classification and analysis of multivariate observations. In *5th Berkeley Symp. Math. Statist. Probability*, pages 281–297, 1967.
- [25] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [26] Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin A Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*, volume 33, pages 596–608, 2020.

- [27] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–9, 2015.
- [28] Laurens Van Der Maaten. Accelerating t-sne using tree-based algorithms. *The Journal of Machine Learning Research*, 15(1):3221–3245, 2014.
- [29] Sagar Vaze, Kai Han, Andrea Vedaldi, and Andrew Zisserman. Generalized category discovery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7492–7501, 2022.
- [30] Nguyen Xuan Vinh, Julien Epps, and James Bailey. Information theoretic measures for clusterings comparison: Variants, properties, normalization and correction for chance. *The Journal of Machine Learning Research*, 11:2837–2854, 2010.
- [31] Junyuan Xie, Ross Girshick, and Ali Farhadi. Unsupervised deep embedding for clustering analysis. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 478–487, 2016.
- [32] Jay Yagnik, Dennis Strelow, David A Ross, and Ruei-sung Lin. The power of comparative reasoning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2431–2438, 2011.
- [33] Muli Yang, Yuehua Zhu, Jiaping Yu, Aming Wu, and Cheng Deng. Divide and conquer: Compositional experts for generalized novel class discovery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14268–14277, 2022.
- [34] Hongyi Zhang, Moustapha Cissé, Yann N. Dauphin, and David Lopez-Paz. Mixup: Beyond empirical risk minimization. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2018.
- [35] Bingchen Zhao and Kai Han. Novel visual category discovery with dual ranking statistics and mutual knowledge distillation. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*, volume 34, pages 22982–22994, 2021.
- [36] Mingkai Zheng, Shan You, Lang Huang, Fei Wang, Chen Qian, and Chang Xu. Simmatch: Semi-supervised learning with similarity matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14471–14481, 2022.
- [37] Zhun Zhong, Enrico Fini, Subhankar Roy, Zhiming Luo, Elisa Ricci, and Nicu Sebe. Neighborhood contrastive learning for novel class discovery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10867–10875, 2021.
- [38] Zhun Zhong, Linchao Zhu, Zhiming Luo, Shaozi Li, Yi Yang, and Nicu Sebe. Openmix: Reviving known knowledge for discovering novel visual categories in an open world. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9462–9470, 2021.