

Bridging the Gap: Learning Pace Synchronization for Open-World Semi-Supervised Learning

Bo Ye^{1,2}, Kai Gan^{1,2}, Tong Wei^{1,2†}, Min-Ling Zhang^{1,2}

¹School of Computer Science and Engineering, Southeast University, Nanjing 211189, China

²Key Lab. of Computer Network and Information Integration (Southeast University), MoE, China

{yeb, gank, weit, zhangml}@seu.edu.cn

Abstract

In open-world semi-supervised learning, a machine learning model is tasked with uncovering novel categories from unlabeled data while maintaining performance on seen categories from labeled data. The central challenge is the substantial learning gap between seen and novel categories, as the model learns the former faster due to accurate supervisory information. Moreover, capturing the semantics of unlabeled novel category samples is also challenging due to the missing label information. To address the above issues, we introduce 1) the adaptive synchronizing marginal loss which imposes class-specific negative margins to alleviate the model bias towards seen classes, and 2) the pseudo-label contrastive clustering which exploits pseudo-labels predicted by the model to group unlabeled data from the same category together in the output space. Extensive experiments on benchmark datasets demonstrate that previous approaches may significantly hinder novel class learning, whereas our method strikingly balances the learning pace between seen and novel classes, achieving a remarkable 3% average accuracy increase on the ImageNet dataset. Importantly, we find that fine-tuning the self-supervised pre-trained model significantly boosts the performance, which is overlooked in prior literature. Our code is available at <https://github.com/yebo0216best/LPS-main>.

1 Introduction

Over the past decade, Semi-Supervised Learning (SSL) algorithms [Zhu and Goldberg, 2009] have demonstrated remarkable performance across multiple tasks, even when presented with a meagre number of labeled training samples [Jia *et al.*, 2023]. These algorithms delve into the underlying data distribution by harnessing numerous unlabeled samples. Among the representative methods employed for this purpose are pseudo-labeling [Lee, 2013] and consistency regularization [Laine and Aila, 2016; Sajjadi *et al.*, 2016]. Pseudo-labeling involves utilizing model predictions as target labels,

[†]Corresponding author

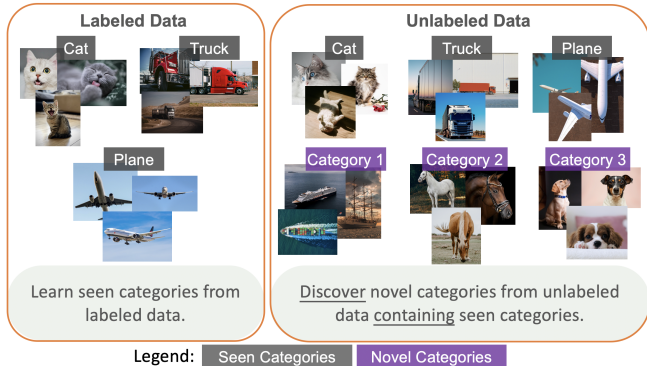


Figure 1: Summary of OpenSSL setting.

while consistency regularization encourages similar predictions for distinct views of an unlabeled sample. However, the majority of current approaches operate under the assumption that unlabeled data exclusively comprises samples belonging to *seen classes*, as observed within the labeled data [Bendale and Boult, 2015; Boult *et al.*, 2019]. In contrast, the presence of samples from *novel classes* in the unlabeled data is common, as it is challenging for human annotators to discern such instances amidst an extensive pool of unlabeled samples [Oliver *et al.*, 2018].

To aid this challenge, Open-World Semi-Supervised Learning, denoted as *OpenSSL*, has gained recent attention, leading to the proposition of several effective methodologies [Cao *et al.*, 2022; Guo *et al.*, 2022; Liu *et al.*, 2023]. Figure 1 demonstrates the problem setting of *OpenSSL* as an intuitive example. To tackle this issue, extant methods adopt a two-pronged strategy. On one front, they endeavour to identify unlabeled samples pertaining to seen classes and allocate pseudo-labels accordingly. On the other front, they automatically cluster unlabeled samples belonging to novel categories. Notably, *OpenSSL* shares an affinity with Novel Class Discovery (NCD) [Han *et al.*, 2019; Han *et al.*, 2020], particularly concerning the clustering of novel class samples. However, NCD methodologies presuppose that unlabeled samples originate exclusively from novel classes. *OpenSSL* relaxes this assumption to mirror real-world scenarios more accurately. Evidently, the central challenge of effectively clustering novel class samples hinges upon the acquisition of

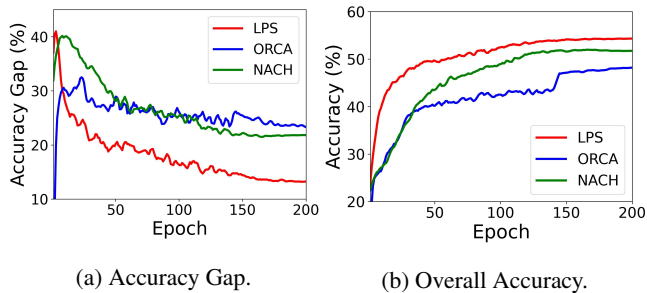


Figure 2: (a) Accuracy gap between seen and novel classes. (b) The overall accuracy of LPS, ORCA, and NACH. Experiments are conducted on the CIFAR-100 dataset with 50% seen classes (50% labeled) and 50% novel classes.

discriminative feature representations, given the absence of supervisory information. To mitigate this quandary, existing methods harness self-supervised learning paradigms (e.g., SimCLR [Chen *et al.*, 2020a]) which circumvent the need for labeled data during the training of feature extractors within deep neural networks. Subsequently, a linear classifier is cultivated by optimizing the cross-entropy loss for labeled data, in conjunction with specifically tailored unsupervised objectives for the unlabeled counterpart. Widely employed unsupervised objectives include entropy regularization and pairwise loss, both of which effectively enhance performance.

This paper introduces a novel OpenSSL algorithm. An initial observation reveals that the model exhibits faster learning of seen classes compared to novel classes. This discrepancy is intuitive because of accurate supervision within labeled data for seen classes, whereas novel classes are learned through unsupervised means. Figure 2 depicts a demonstration of the learning speed discrepancy between seen and novel classes. Motivated by this intrinsic problem, we propose an adaptive distribution-aware margin mechanism, designed to steer the model’s attention towards novel class learning. Notably, this margin diminishes as the model’s predicted class distribution approaches the underlying (class-balanced) distribution. To learn robust representations and facilitate the clustering of novel classes, we introduce pseudo-label contrastive clustering. This technique aggregates unlabeled samples sharing the same class, guided by model predictions. Importantly, we exploit multiple positive and negative pairs as supervisory signals, in contrast to the reliance on a single positive pair as seen in previous works. For unlabeled samples exhibiting low confidence, we integrate unsupervised contrastive learning to facilitate the acquisition of informative representations. This unsupervised contrastive objective operates as a complement to the pseudo-label contrastive clustering. Combining the aforementioned modules, we present, LPS, to address the OpenSSL challenge. Figure 2 showcases the efficacy of LPS compared with existing state-of-the-art approaches. Notably, we reveal that the conventional practice of freezing the feature extractor, previously trained via self-supervised learning in prior research, falls short of optimal.

In summary, our main contributions are:

- We propose a novel and simple method, LPS, to effectively synchronize the learning pace of seen and novel

classes for open-world semi-supervised learning.

- We conduct extensive experiments to verify the effectiveness of the proposed method against the previous state-of-the-art. Particularly, LPS achieves over 3% average increase of accuracy on the ImageNet dataset.
- We examine the effectiveness of the key components of the proposed method. Different from previous works, we discover that fine-tuning the pre-trained backbone allows the model to learn more useful features, which can significantly improve the performance.

2 Related Work

Semi-Supervised Learning. Within the realm of SSL, pseudo-labeling [Lee, 2013] and consistency regularization [Laine and Aila, 2016; Sajjadi *et al.*, 2016; Wei *et al.*, 2022a; Wei and Gan, 2023] are two widely used techniques. Pseudo-labeling converts model predictions on unlabeled samples into either soft labels or hard labels, subsequently employed as target labels. Consistency regularization strives to ensure model outputs exhibit a high degree of consistency when applied to perturbed samples. Recent advancements [Berthelot *et al.*, 2019; Sohn *et al.*, 2020; Xu *et al.*, 2021] combine pseudo-labeling with consistency regularization to yield further performance enhancements. In addition to the above techniques, the application of contrastive learning into SSL has also received substantial interest. For example, TCL [Singh *et al.*, 2021] introduces contrastive learning as a tool to enhance representation learning. TCL maximizes agreement between different views of the same sample while minimizing agreement for distinct samples. In consonance with this paradigm, we design a new complementary contrastive loss to explore the consistency of all samples effectively.

Novel Class Discovery. The setting of NCD aligns closely with the scenario investigated in this paper. NCD assumes a scenario where the labeled data consists of samples of seen classes, while the unlabeled data exclusively comprises samples of novel classes. [Han *et al.*, 2019] initially raised the NCD problem. Subsequent research such as [Han *et al.*, 2020; Zhong *et al.*, 2021a; Zhong *et al.*, 2021b] predominantly adopted multi-stage training strategies. The underlying principle is capturing comprehensive high-level semantic information from labeled data, subsequently propagated to unlabeled counterparts. The majority of NCD methods involve preliminary model pre-training, wherein several objective functions are invoked to minimize inter-sample distances for each class. However, in real-world scenarios, the assumption of unlabeled data solely comprising novel classes is unrealistic [Wei *et al.*, 2022b], as seen classes also significantly populate the unlabeled dataset. Our experimentation reveals that NCD algorithms struggle to match the performance of other leading methods in the context of OpenSSL.

Open-World Semi-Supervised Learning. While conventional SSL methods operate under the assumption of labeled and unlabeled data being associated with a predefined set of classes, recent advancements [Oliver *et al.*, 2018; Chen *et al.*, 2020b; Guo *et al.*, 2020; Saito *et al.*, 2021] challenge this notion and acknowledge the potential emergence of

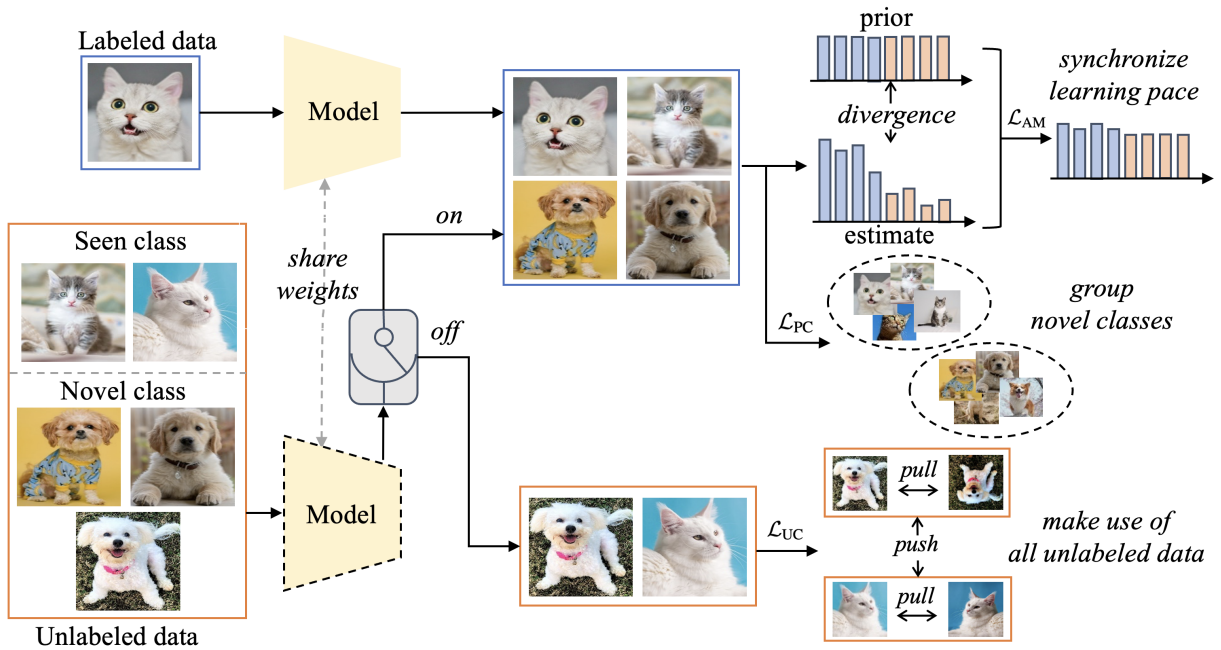


Figure 3: Overview of LPS framework. The LPS objective function is composed of an adaptive margin objective \mathcal{L}_{AM} , a pseudo-label contrastive clustering objective \mathcal{L}_{PC} , and an unsupervised contrastive learning objective \mathcal{L}_{UC} .

novel classes within the unlabeled data. Further, ORCA [Cao *et al.*, 2022] delves into the OpenSSL problem and attempts to cluster novel classes from unlabeled data. Both ORCA and NACH [Guo *et al.*, 2022] notice the discrepancy in learning pace between seen and novel classes. They propose to use an entropy regularizer and pair-wise loss to ameliorate this issue. More recently, OpenNCD [Liu *et al.*, 2023] introduces prototypes-based contrastive learning for recognizing seen classes and discovering novel classes. This paper introduces an adaptive margin loss and pseudo-label contrastive clustering mechanism to synchronize the learning pace and enhance novel class discovery.

3 The Proposed Method

Notations. The training dataset is composed of the labeled data $\mathcal{D}_l = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$ and the unlabeled data $\mathcal{D}_u = \{\mathbf{x}_i\}_{i=n+1}^{n+m}$. Within the context of OpenSSL, the classes in \mathcal{D}_l are designated as *seen classes*, constituting the set denoted as \mathcal{C}_s . The scenario of interest acknowledges a distribution mismatch, leading to \mathcal{D}_u comprising instances from both seen and *novel classes*. The collection of these novel classes is represented as \mathcal{C}_n . Additionally, we adopt the premise of a known number of novel classes akin to prior OpenSSL methodologies [Cao *et al.*, 2022; Guo *et al.*, 2022; Liu *et al.*, 2023]. The goal of OpenSSL is to classify samples originating from the \mathcal{C}_s and cluster samples emanating from the \mathcal{C}_n .

Overview. The fundamental challenge in OpenSSL arises from the pronounced discrepancy in learning paces between seen and novel classes, primarily due to the precise supervisory guidance for seen classes. This discrepancy results in a bias towards seen classes in the model’s predictions,

adversely impacting both the accurate classification of seen class samples and the effective clustering of novel class samples. To circumvent this challenge, we introduce *Learning Pace Synchronization* (LPS), a methodology with adaptive synchronizing loss and pseudo-label contrastive clustering as in Figure 3. The adaptive synchronizing loss aims to achieve a balance between the learning pace of seen and novel classes, and the pseudo-label contrastive clustering exploits pseudo-labels to group unlabeled data from the same class together in the output space.

3.1 Adaptive Synchronizing

To start with, we describe the proposed adaptive marginal loss which regularizes the learning pace of seen classes to synchronize the learning pace of the model. Conventionally, the margin is defined as the minimum distance of the data to the classification boundary. For a sample (\mathbf{x}, y) , we have:

$$\Delta(\mathbf{x}, y) = f(\mathbf{x})_y - \max_{j \neq y} f(\mathbf{x})_j \quad (1)$$

Instead of employing a fixed margin, LDAM [Cao *et al.*, 2020] introduces a class-specific margin, where the margin between rare classes and other classes is larger than the margin for frequent classes, for tackling class-imbalanced data. Specifically, it sets the margin of class j as:

$$\Delta_j = \frac{C}{n_j^{1/4}} \quad (2)$$

The constant C controls the intensity and n_j denotes the frequency of class j in the training data. Motivated by this, we propose a new variant of the margin loss to synchronize the learning pace of seen and novel classes. We apply adaptive

margin loss to demand a larger margin between the novel and other classes, so that scores for seen classes, towards which the model highly biased, do not overwhelm the novel classes. For each sample (\mathbf{x}, y) , the adaptive margin loss is defined as follows:

$$\ell_{\text{AM}}(\mathbf{x}, y) = -\log \frac{\exp(z_y - \Delta_y)}{\exp(z_y - \Delta_y) + \sum_{j \neq y} \exp(z_j)} \quad (3)$$

$$\text{where } \Delta_j = -KL(\hat{\pi} \parallel \pi) \frac{\hat{\pi}_j}{\max(\hat{\pi})} C, \quad j \in [K]$$

In this formulation, K represents the total number of classes, z_j signifies the model output for the j -th class, $\mathbf{z} = f(\mathbf{x}; \theta)$, and $\hat{\pi}$ denotes the estimated class distribution by the model. Additionally, we introduce an approximation of the true class distribution π , which is naturally inaccessible during training. In line with prior studies, we assume a uniform distribution for π , leaving the exploration of arbitrary distributions for future investigations. The hyper-parameter C is introduced to control the maximum margin, and we empirically set $C = 10$ across all experiments. We conduct a series of studies on the value of C in the supplementary material.

For the sake of simplicity, we assume that the mini-batch is comprised of labeled data B_l and unlabeled data B_u . Given that the computation of Eq. (3) relies on the class distribution, we proceed to estimate the complete class distribution through labeled data and unlabeled data exhibiting high predictive confidence. Specifically, we endeavour to achieve this estimation through:

$$\hat{\pi} = \text{Normalize} \left(\sum_{\mathbf{x}_i \in B_l} \hat{\mathbf{y}}_i + \sum_{\mathbf{x}_j \in B_u} \mathbb{I}(\max(\hat{\mathbf{y}}_j) \geq \lambda) \hat{\mathbf{y}}_j \right) \quad (4)$$

Here, $\hat{\mathbf{y}} = \text{softmax}(\mathbf{z})$. In view of the tendency for novel class samples to exhibit underconfidence, we empirically introduce a progressively evolving confidence threshold $\lambda_{\text{novel}} = 0.4 + 0.4 \times \frac{t}{T}$, where t and T signify the current training iteration and the total training iterations, respectively. For seen classes, a fixed confidence threshold $\lambda_{\text{seen}} = 0.95$ is employed.

Given that the estimated class distribution $\hat{\pi}$ mirrors the model’s confidence in class predictions, we harness this insight to regulate the learning pace of both seen and novel classes. Notably, in the early training phases, the model is inclined towards seen classes, with the logit adjustment term Δ_j assuming a larger negative margin for seen classes, thereby attenuating their learning pace. As training progresses, the model attains a more balanced capability across both seen and novel classes, as reflected by a diminishing Kullback-Leibler (KL) divergence between $\hat{\pi}$ and π .

In summary, the adaptive margin loss \mathcal{L}_{AM} for both labeled and pseudo-labeled data is defined as follows:

$$\mathcal{L}_{\text{AM}} = \frac{1}{|B_l|} \sum_{\mathbf{x}_i \in B_l} \ell_{\text{AM}}(\mathbf{z}_i^w, y) + \frac{1}{|B_u|} \sum_{\mathbf{x}_j \in B_u} \mathbb{I}(\max(\hat{\mathbf{y}}_j) \geq \lambda) \ell_{\text{AM}}(\mathbf{z}_j^s, \hat{p}_j) \quad (5)$$

In this context, \mathbf{z}^w and \mathbf{z}^s correspond to the output logits stemming from the weak and strong augmented versions of sample \mathbf{x} , respectively. The symbol $|\cdot|$ denotes the set cardinality operation. Additionally, we utilize $\hat{p} = \arg \max(\text{softmax}(\mathbf{z}^w))$ to represent the pseudo-label associated with the sample.

Distinctions and Connections with Alternatives. It is worth noting that the concept of adaptive margin has been used in prior literature [Li *et al.*, 2020; Ha and Blanz, 2021; Cao *et al.*, 2022]. Different from LPS, [Li *et al.*, 2020] leverages the semantic similarities between classes to generate adaptive margins with the motivation to separate similar classes in the embedding space, and [Ha and Blanz, 2021] utilizes the ground-truth distance between different samples to generate adaptive margins with the motivation to adapt to rating datasets. In OpenSSL, ORCA [Cao *et al.*, 2022] also integrates an adaptive margin mechanism based on the model’s predictive uncertainty, which can only equally suppress the learning pace of seen classes. However, there are still differences in the learning paces of different classes among seen classes. Our proposed adaptive margin is based on the current estimated distribution to reflect the learning pace of different classes, which offers increased flexibility for regulating the learning pace across classes by generating the class-specific negative margin. Furthermore, the inclusion of the KL divergence term effectively guards against the model converging to a trivial solution where all samples are arbitrarily assigned to a single class.

3.2 Pseudo-Label Contrastive Clustering

The basic idea of discovering novel classes is to explore the correlations between different samples and cluster them into several groups. Prior OpenSSL approaches often transform the clustering task into a pairwise similarity prediction task, wherein a modified form of binary cross-entropy loss is optimized. Different from existing works, we introduce a new clustering method to fully exploit reliable model predictions as supervisory signals.

Our approach involves the construction of a multi-viewed mini-batch by using weak and strong augmentations. Within each mini-batch, we group the labeled and confident unlabeled samples, which is denoted as $B_{l'}$. Concurrently, unlabeled samples exhibiting predicted confidence levels failing below the threshold λ are denoted as $B_{u'}$. Pseudo-label contrastive clustering only takes $B_{l'}$ as inputs. For each sample in $B_{l'}$, the set of the positive pairs contains samples with the same given label or pseudo-label. Conversely, the set of negative pairs contains samples of other classes. Formally, the objective of pseudo-label contrastive clustering is defined as follows:

$$\mathcal{L}_{\text{PC}} = -\frac{1}{|B_{l'}|} \sum_{\mathbf{x}_i \in B_{l'}} \log \frac{1}{|P(\mathbf{x}_i)|} \frac{\sum_{\mathbf{x}_p \in P(\mathbf{x}_i)} \exp(\mathbf{z}_i \cdot \mathbf{z}_p / \tau)}{\sum_{\mathbf{x}_a \in A(\mathbf{x}_i)} \exp(\mathbf{z}_i \cdot \mathbf{z}_a / \tau)}, \quad (6)$$

where \mathbf{z}_i denotes the output logits, $P(\mathbf{x}_i)$ denotes the set of positives of \mathbf{x}_i and $A(\mathbf{x}_i) \equiv B_{l'} \setminus \{\mathbf{x}_i\}$. τ is a tunable temperature parameter and we set $\tau = 0.4$ in all experiments.

In contrast to existing methods such as ORCA [Cao *et al.*, 2022] and NACH [Guo *et al.*, 2022], which establish a single positive pair for each sample by identifying its nearest neighbour, the objective in Eq. (6) adopts pseudo-labels to form multiple positive pairs. This approach offers dual advantages: firstly, the alignment of samples within the same class is more effectively harnessed through the utilization of multiple positive sample pairs; secondly, it leverages the consistency of distinct views of the same sample to mitigate the negative impact of erroneous positive pairs, while concurrently imparting a repulsion effect to samples from different classes through negative pairs.

Since Eq. (6) augments the labeled dataset by unlabeled samples of high predictive confidence, we ask whether unlabeled samples of low confidence can be used to enrich representation learning. In pursuit of this, we incorporate unsupervised contrastive learning [Wang and Isola, 2020] to encourage similar predictions rather than embeddings between a given sample and its augmented counterpart. This helps to signify the uniformity among unlabeled samples, ultimately leading to clearer separations. In detail, for each sample in the low-confidence set $B_{u'}$, the unsupervised contrastive learning couples it with its augmented view to constitute a positive pair. Simultaneously, a set of negative pairs is formulated, containing all the samples within the mini-batch except the sample itself. The unsupervised contrastive learning loss \mathcal{L}_{UC} is formulated as follows:

$$\mathcal{L}_{UC} = -\frac{1}{|B_{u'}|} \sum_{\mathbf{x}_j \in B_{u'}} \log \frac{\exp(\mathbf{z}_j \cdot \mathbf{z}_p / \tau)}{\sum_{\mathbf{x}_a \in A(\mathbf{x}_j)} \exp(\mathbf{z}_j \cdot \mathbf{z}_a / \tau)} \quad (7)$$

Here, \mathbf{x}_p is the positive sample of \mathbf{x}_j and $A(\mathbf{x}_j) \equiv B_{u'} \cup B_{u'} \setminus \{\mathbf{x}_j\}$ for the sample \mathbf{x}_j .

In essence, unsupervised contrastive learning complements the pseudo-label contrastive clustering by fully exploiting the unlabeled samples. In the experiments, the ablation studies underscore the pivotal role played by both types of contrastive losses in our approach.

Lastly, we incorporate a maximum entropy regularizer to address the challenge of converging during the initial training phases, when the predictions are mostly wrong (e.g., the model tends to assign all samples to the same class) [Arazo *et al.*, 2020]. Specifically, we leverage the KL divergence between the class distribution predicted by the model and a uniform prior distribution. It is worth noting that the integration of an entropy regularizer is a widespread practice in dealing with the OpenSSL problem, including approaches such as ORCA, NACH, and OpenNCD. The final objective function of LPS is articulated as follows:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{AM}} + \eta_1 \mathcal{L}_{\text{PC}} + \eta_2 \mathcal{L}_{\text{UC}} + \mathcal{R}_{\text{Entropy}} \quad (8)$$

where $\mathcal{R}_{\text{Entropy}}$ denotes the entropy regularizer, η_1 and η_2 are hyper-parameters set to 1 in all our experiments. We provide detailed analyses on the sensitivity of hyperparameters in the supplementary material.

4 Experiments

4.1 Experimental Setup

Datasets. We evaluate our method on three commonly used datasets, i.e., CIFAR-10, CIFAR-100 [Krizhevsky, 2009], and ImageNet [Russakovsky *et al.*, 2015]. Following prior works [Cao *et al.*, 2022; Guo *et al.*, 2022; Liu *et al.*, 2023], we assume that the number of novel classes is known. Specifically, we randomly select 50% of the classes as seen classes, and the remaining classes are regarded as novel classes, e.g., the number of novel classes is 50 for CIFAR-100. On each dataset, we consider two types of labeled ratios, i.e., only 10% or 50% of data in seen classes are labeled. For the ImageNet dataset, we subsample 100 classes to form the ImageNet-100 dataset for fair comparisons with existing works.

Following prior works [Cao *et al.*, 2022; Guo *et al.*, 2022; Liu *et al.*, 2023], we evaluate our method with respect to the accuracy of seen classes, novel classes, and all classes. For seen classes, the accuracy is calculated as the normal classification task. For novel classes, we first utilize the Hungarian algorithm [Kuhn, 1955] to solve the optimal prediction-target class assignment problem and then calculate the accuracy of novel classes. For overall accuracy, we also solve the optimal assignment in the entire unlabeled dataset to calculate the novel class accuracy, measuring the overall performance.

Implementation Details. Following [Cao *et al.*, 2022; Guo *et al.*, 2022; Liu *et al.*, 2023], we utilize the self-supervised learning method SimCLR [Chen *et al.*, 2020a] to pre-train the backbone and fix the first three blocks. In LPS, the weak augmentation contains random crop and horizontal flip, and the strong augmentation is RandAugment [Cubuk *et al.*, 2019]. For CIFAR-10 and CIFAR-100, we utilize ResNet-18 as our backbone which is trained by the standard SGD with a momentum of 0.9 and a weight decay of 0.0005. We train the model for 200 epochs with a batch size of 512. For the ImageNet dataset, we opt for ResNet-50 as our backbone. This choice also undergoes training via the standard SGD, featuring a momentum coefficient of 0.9 and a weight decay of 0.0001. The training process spans 90 epochs, with a batch size of 512. and The cosine annealing learning rate schedule is adopted on CIFAR and ImageNet datasets. These experiments are conducted on a single NVIDIA 3090 GPU.

4.2 Comparing with Existing Methods

Baselines. We compare LPS with SSL methods, open-set SSL methods, NCD methods, and OpenSSL methods. The NCD methods consider that the labeled data only has disjoint classes compared with the unlabeled data and aim at clustering novel classes without recognizing seen classes. For novel classes, clustering accuracy can be obtained directly. For seen classes, we first regard them as novel classes and leverage the Hungarian algorithm to match some of the discovered classes with seen classes, and then calculate the classification accuracy. We select two competitive NCD methods DTC [Han *et al.*, 2019] and RankStats [Han *et al.*, 2020] in the experiments. Moreover, we include GCD [Vaze *et al.*, 2022] for comparison, which is an extended NCD method.

Methods	CIFAR-10						CIFAR-100						ImageNet-100					
	10% labeled			50% labeled			10% labeled			50% labeled			10% labeled			50% labeled		
	Seen	Novel	All	Seen	Novel	All	Seen	Novel	All	Seen	Novel	All	Seen	Novel	All	Seen	Novel	All
FixMatch	64.3	49.4	47.3	71.5	50.4	49.5	30.9	18.5	15.3	39.6	23.5	20.3	60.9	33.7	30.2	65.8	36.7	34.9
DS ³ L	70.5	46.6	43.5	77.6	45.3	40.2	33.7	15.8	15.1	55.1	23.7	24.0	64.3	28.1	25.9	71.2	32.5	30.8
DTC	42.7	31.8	32.4	53.9	39.5	38.3	22.1	10.5	13.7	31.3	22.9	18.3	24.5	17.8	19.3	25.6	20.8	21.3
RankStats	71.4	63.9	66.7	86.6	81.0	82.9	20.4	16.7	17.8	36.4	28.4	23.1	41.2	26.8	37.4	47.3	28.7	40.3
SimCLR	44.9	48.0	47.7	58.3	63.4	51.7	26.0	28.8	26.5	28.6	21.1	22.3	42.9	41.6	41.5	39.5	35.7	36.9
ORCA	82.8	85.5	84.1	88.2	90.4	89.7	52.5	31.8	38.6	66.9	43.0	48.1	83.9	60.5	69.7	89.1	72.1	77.8
GCD	78.4	79.7	79.1	78.4	79.7	79.1	49.7	27.6	38.0	68.5	33.5	45.2	82.3	58.3	68.2	82.3	58.3	68.2
OpenNCD	83.5	86.7	85.3	88.4	90.6	90.1	53.6	33.0	41.2	69.7	43.4	49.3	84.0	65.8	73.2	90.0	77.5	81.6
NACH	86.4	89.4	88.1	89.5	92.2	91.3	57.4	37.5	43.5	68.7	47.0	52.1	86.3	66.5	71.0	91.0	75.5	79.6
LPS (ours)	86.3	90.6	88.6	90.2	93.4	92.4	55.2	41.0	47.5	64.5	49.9	54.3	87.0	73.6	78.0	91.3	81.3	84.5

Table 1: Accuracy comparison of seen, novel, and all classes on CIFAR-10, CIFAR-100, and ImageNet-100 datasets with 50% classes as seen and 50% classes as novel. We conducted experiments with 10% and 50% labeled data of seen classes.

For the SSL and open-set SSL methods, we leverage their capability in estimating out-of-distribution samples to extend to the OpenSSL setting. For comparison, we select FixMatch[Sohn *et al.*, 2020], which assigns pseudo-labels to unlabeled samples based on confidence. The classification accuracy of seen classes can be reported directly according to pseudo-labels. For novel classes, we first estimate samples without pseudo-labels as novel classes and then utilize k -means to cluster them. The open-set SSL methods maintain the classification performance of seen classes by rejecting novel classes. We compare with DS³L [Guo *et al.*, 2020] and calculate its accuracy in the same way as FixMatch.

For the OpenSSL methods, we compare with ORCA [Cao *et al.*, 2022], NACH [Guo *et al.*, 2022], and OpenNCD [Liu *et al.*, 2023]. We also compare the self-supervised pre-trained model SimCLR and conduct K-means on the primary features to calculate the accuracy.

Results. The results on three datasets are reported in Table 1. The mean accuracy is computed over three runs for each method. Although the non-OpenSSL methods perform well on their original tasks, their overall performance is unsatisfactory in the OpenSSL setting. The results of SimCLR are obtained by the pre-trained model without extra fine-tuning, and the OpenSSL methods are based on the pre-trained model. It is obvious that the OpenSSL methods achieve significant performance improvements compared to non-OpenSSL methods. Compared with the state-of-the-art OpenSSL methods, our method LPS achieves the best overall performance across all datasets. On the CIFAR-10 dataset, LPS outperforms NACH by 1.2% in novel class accuracy. Likewise, on the CIFAR-100 dataset, LPS demonstrates superiority, yielding a substantial 3.2% improvement. Particularly concerning the ImageNet-100 dataset, LPS has the capacity to surpass existing state-of-the-art methods, resulting in a 3.8% increase in overall accuracy. Experimental results demonstrate that LPS can effectively balance the learning of seen and novel classes.

Distribution Analysis. For further validation of our approach, we present a comprehensive analysis of the KL divergence trend between the estimated and prior class distributions, along with the estimated class distributions at the

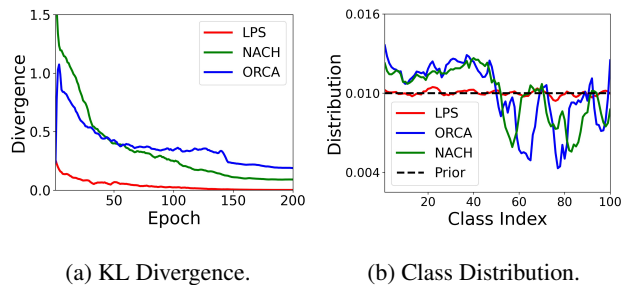


Figure 4: (a) The KL divergence between the estimated and prior class distributions. (b) The estimated class distribution of the final training iteration. Experiments are conducted on the CIFAR-100 dataset with 50% seen classes (50% labeled) and 50% novel classes.

concluding training iteration for LPS, NACH, and ORCA in Figure 4. It can be seen that LPS is able to give more attention to novel classes in the early training stage, in contrast to ORCA and NACH, which exhibit a preference for learning seen classes primarily. This becomes particularly evident in the estimated class distribution at the end of the training, where ORCA and NACH falter in effectively distinguishing novel classes, leading to the erroneous assignment of certain novel class samples. In contrast, the estimated distribution obtained by LPS closely aligns with the prior distribution, consequently fostering effective novel class differentiation.

Feature Visualization. We employ t-distributed stochastic neighbour embedding (t-SNE) [Van der Maaten and Hinton, 2008] to visualize the learned feature representations through LPS and prior methods. The results depicted in Figure 5 demonstrate that the representations obtained by LPS yield clearer classification boundaries and exhibit fewer instances of misclassifications. Conversely, ORCA and NACH manifest a higher frequency of misclassifications and unclear boundaries, resulting in comparatively inferior performance in contrast to LPS. In addition, we utilize the metric of normalized mutual information (NMI) to further assess the clustering performance and representations of novel classes. We report the results in Table 2 which clearly demonstrate the superiority of LPS. Given that all methods share the same

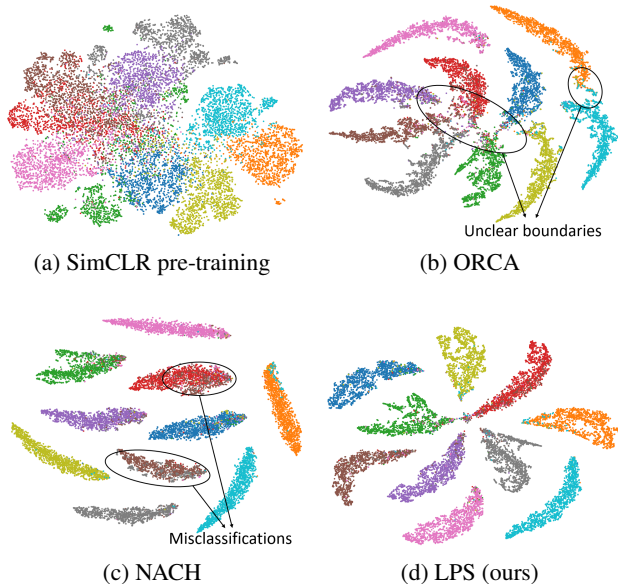


Figure 5: The t-SNE visualization of features for different methods on CIFAR-10 dataset with 50% seen classes (50% labeled) and 50% novel classes.

foundation of SimCLR pre-trained backbone, the visualization, and NMI results highlight the efficacy of our approach in enhancing representation learning.

Methods	CIFAR-10		CIFAR-100	
	10%	50%	10%	50%
ORCA	73.93	80.92	45.26	52.10
NACH	80.07	83.44	49.96	56.45
LPS (ours)	82.41	86.45	51.92	58.42

Table 2: The NMI of novel classes on CIFAR datasets.

Fine-tuning the Pre-trained Backbone. Furthermore, it is noteworthy that all previous OpenSSL methods adopt a practice of freezing the parameters within the first three blocks of the backbone, solely fine-tuning the last block, with the intention of mitigating overfitting. However, such an approach constrains the extent of performance enhancement, as the backbone’s parameters remain unmodifiable and unoptimized to better suit downstream tasks. To establish that our method is not susceptible to the overfitting dilemma, we conducted a series of experiments on the CIFAR dataset employing state-of-the-art OpenSSL methods while fine-tuning the backbone. The results are reported in Table 3. The experimental results reveal that existing OpenSSL methods manifest modest performance improvement, if any, in comparison to their initial frozen counterparts. In contrast, our proposed method, unaffected by overfitting concerns, consistently yields substantial performance gains across both seen and novel classes. Specifically, the overall accuracy for CIFAR-10 experiences a notable improvement of 2.9%, while an impressive 6.3% increase is observed for CIFAR-100. These results underscore

Methods	CIFAR-10			CIFAR-100		
	Seen	Novel	All	Seen	Novel	All
ORCA	89.8	90.8	90.5	69.4	42.5	48.2
OpenNCD	88.9	91.8	90.8	62.2	44.9	50.4
NACH	95.0	93.3	93.9	73.8	47.8	54.6
LPS (ours)	95.0	95.5	95.3	72.6	55.5	60.6

Table 3: Mean accuracy over three runs with removing the limitation of freezing the backbone on CIFAR datasets with 50% seen classes (50% labeled) and 50% novel classes.

Methods	CIFAR-10			CIFAR-100		
	Seen	Novel	All	Seen	Novel	All
w/o \mathcal{L}_{AM}	90.6	89.4	89.8	67.1	46.2	52.7
w/o \mathcal{L}_{PC}	89.7	89.1	89.3	60.4	48.1	51.8
w/o \mathcal{L}_{UC}	91.3	88.2	89.2	64.7	42.3	48.5
w/o $\mathcal{R}_{Entropy}$	90.8	55.4	66.2	73.4	28.8	30.8
LPS (ours)	90.2	93.4	92.4	64.5	49.9	54.3

Table 4: Accuracy when removing key components of our method. We report the average accuracy over three runs on CIFAR datasets with 50% seen classes (50% labeled) and 50% novel classes.

the effectiveness of LPS in harnessing the additional learnable parameters for further enhancing model performance.

Ablation Analysis. Moreover, we conduct a comprehensive analysis of the contributions of distinct components in our approach. The objective function of LPS comprises the adaptive margin loss (\mathcal{L}_{AM}), the pseudo-label contrastive clustering loss (\mathcal{L}_{PC}), the unsupervised contrastive learning loss (\mathcal{L}_{UC}), and the entropy regularizer ($\mathcal{R}_{Entropy}$). Concretely, the ablation study is mainly conducted by removing each term individually from the objective function except for the adaptive margin which is replaced by a standard cross-entropy. As observed in Table 4, the removal of any components leads to performance degradation. The substantial drop in novel performance after removing the entropy regularizer highlights its significant role in the process of novel class discovery. Moreover, the utilization of both pseudo-label contrastive loss and adaptive margin loss substantially improves the accuracy of novel classes.

5 Conclusion

In this study, we present Learning Pace Synchronization (LPS), a potent solution tailored to the OpenSSL problem. LPS introduces an adaptive margin loss to effectively narrow the learning pace gap that exists between seen and novel classes. Moreover, we formulate a pseudo-label contrastive clustering loss to augment the process of novel class discovery. Extensive evaluation is conducted across three benchmark datasets with distinct quantities of labeled data. We also discover that the conventional practice of freezing the self-supervised pre-trained backbone hinders the generalization performance. We hope our work can inspire more efforts towards this realistic setting.

References

- [Arazo *et al.*, 2020] Eric Arazo, Diego Ortego, Paul Albert, Noel E. O’Connor, and Kevin McGuinness. Pseudo-labeling and confirmation bias in deep semi-supervised learning. In *International Joint Conference on Neural Networks*, 2020.
- [Bendale and Boulton, 2015] Abhijit Bendale and Terrance E Boulton. Towards open world recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2015.
- [Berthelot *et al.*, 2019] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A Raffel. MixMatch: A holistic approach to semi-supervised learning. *Neural Information Processing Systems*, 2019.
- [Boulton *et al.*, 2019] Terrance E Boulton, Steve Cruz, Akshay Raj Dhamija, M Gunther, James Henrydoss, and Walter J Scheirer. Learning and the unknown: Surveying steps toward open world recognition. In *AAAI Conference on Artificial Intelligence*, 2019.
- [Cao *et al.*, 2020] Kaidi Cao, Colin Wei, Adrien Gaidon, Thang Luong, and Quoc Le. Learning imbalanced datasets with label-distribution-aware margin loss. In *Neural Information Processing Systems*, 2020.
- [Cao *et al.*, 2022] Kaidi Cao, Maria Brbic, and Jure Leskovec. Open-world semi-supervised learning. In *International Conference on Learning Representations*, 2022.
- [Chen *et al.*, 2020a] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International Conference on Machine Learning*, 2020.
- [Chen *et al.*, 2020b] Yanbei Chen, Xiatian Zhu, Wei Li, and Shaogang Gong. Semi-supervised learning under class distribution mismatch. In *AAAI Conference on Artificial Intelligence*, 2020.
- [Cubuk *et al.*, 2019] Ekin D. Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V. Le. Randaugment: Practical automated data augmentation with a reduced search space. *arXiv preprint arXiv:1909.13719*, 2019.
- [Guo *et al.*, 2020] Lan-Zhe Guo, Zhen-Yu Zhang, Yuan Jiang, Yu-Feng Li, and Zhi-Hua Zhou. Safe deep semi-supervised learning for unseen-class unlabeled data. In *International Conference on Machine Learning*, 2020.
- [Guo *et al.*, 2022] Lan-Zhe Guo, Yi-Ge Zhang, Zhi-Fan Wu, Jie-Jing Shao, and Yu-Feng Li. Robust semi-supervised learning when not all classes have labels. In *Neural Information Processing Systems*, 2022.
- [Ha and Blanz, 2021] Mai Lan Ha and Volker Blanz. Deep ranking with adaptive margin triplet loss. *arXiv preprint arXiv:2107.06187*, 2021.
- [Han *et al.*, 2019] Kai Han, Andrea Vedaldi, and Andrew Zisserman. Learning to discover novel visual categories via deep transfer clustering. In *IEEE International Conference on Computer Vision*, 2019.
- [Han *et al.*, 2020] Kai Han, Sylvestre-Alvise Rebuffi, Sebastian Ehrhardt, Andrea Vedaldi, and Andrew Zisserman. Automatically discovering and learning new visual categories with ranking statistics. In *International Conference on Learning Representations*, 2020.
- [Jia *et al.*, 2023] Bin-Bin Jia, Jun-Ying Liu, Jun-Yi Hang, and Min-Ling Zhang. Learning label-specific features for decomposition-based multi-class classification. *Frontiers of Computer Science*, 17(6):176348, 2023.
- [Krizhevsky, 2009] Alex Krizhevsky. Learning multiple layers of features from tiny images. Master’s thesis, 2009.
- [Kuhn, 1955] Harold W Kuhn. The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 1955.
- [Laine and Aila, 2016] Samuli Laine and Timo Aila. Temporal Ensembling for Semi-Supervised Learning. *arXiv: Neural and Evolutionary Computing*, 2016.
- [Lee, 2013] Dong-Hyun Lee. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, 2013.
- [Li *et al.*, 2020] Aoxue Li, Weiran Huang, Xu Lan, Jiashi Feng, Zhenguo Li, and Liwei Wang. Boosting few-shot learning with adaptive margin loss. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2020.
- [Liu *et al.*, 2023] Jiaming Liu, Yangqiming Wang, Tongze Zhang, Yulu Fan, Qinli Yang, and Junming Shao. Open-world semi-supervised novel class discovery. In *International Joint Conference on Artificial Intelligence*, 2023.
- [Oliver *et al.*, 2018] Avital Oliver, Augustus Odena, Colin A Raffel, Ekin Dogus Cubuk, and Ian Goodfellow. Realistic evaluation of deep semi-supervised learning algorithms. In *Neural Information Processing Systems*, 2018.
- [Russakovsky *et al.*, 2015] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. ImageNet large scale visual recognition challenge. *International Journal of Computer Vision*, 2015.
- [Saito *et al.*, 2021] Kuniaki Saito, Donghyun Kim, and Kate Saenko. Openmatch: Open-set semi-supervised learning with open-set consistency regularization. *Neural Information Processing Systems*, 2021.
- [Sajjadi *et al.*, 2016] Mehdi Sajjadi, Mehran Javanmardi, and Tolga Tasdizen. Regularization with stochastic transformations and perturbations for deep semi-supervised learning. In *Neural Information Processing Systems*, 2016.
- [Singh *et al.*, 2021] Ankit Singh, Omprakash Chakraborty, Ashutosh Varshney, Rameswar Panda, Rogerio Feris, Kate Saenko, and Abir Das. Semi-supervised action recognition with temporal contrastive learning. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2021.
- [Sohn *et al.*, 2020] Kihyuk Sohn, David Berthelot, Chun-Liang Li, Zizhao Zhang, Nicholas Carlini, Ekin D. Cubuk,

- Alex Kurakin, Han Zhang, and Colin Raffel. FixMatch: Simplifying semi-supervised learning with consistency and confidence. In *Neural Information Processing Systems*, 2020.
- [Van der Maaten and Hinton, 2008] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 2008.
- [Vaze *et al.*, 2022] Sagar Vaze, Kai Hant, Andrea Vedaldi, and Andrew Zisserman. Generalized category discovery. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2022.
- [Wang and Isola, 2020] Tongzhou Wang and Phillip Isola. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *International Conference on Machine Learning*, 2020.
- [Wei and Gan, 2023] Tong Wei and Kai Gan. Towards realistic long-tailed semi-supervised learning: Consistency is all you need. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023.
- [Wei *et al.*, 2022a] Tong Wei, Qian-Yu Liu, Jiang-Xin Shi, Wei-Wei Tu, and Lan-Zhe Guo. Transfer and share: Semi-supervised learning from long-tailed data. *Machine Learning*, 2022.
- [Wei *et al.*, 2022b] Tong Wei, Hai Wang, Weiwei Tu, and Yufeng Li. Robust model selection for positive and unlabeled learning with constraints. *Science China Information Sciences*, 65(11):212101, 2022.
- [Xu *et al.*, 2021] Yi Xu, Lei Shang, Jinxing Ye, Qi Qian, Yufeng Li, Baigui Sun, Hao Li, and Rong Jin. Dash: Semi-supervised learning with dynamic thresholding. In *International Conference on Machine Learning*, 2021.
- [Zhong *et al.*, 2021a] Zhun Zhong, Enrico Fini, Subhankar Roy, Zhiming Luo, Elisa Ricci, and Nicu Sebe. Neighborhood contrastive learning for novel class discovery. In *IEEE conference on Computer Vision and Pattern Recognition*, 2021.
- [Zhong *et al.*, 2021b] 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 *IEEE conference on Computer Vision and Pattern Recognition*, 2021.
- [Zhu and Goldberg, 2009] Xiaojin Zhu and Andrew B Goldberg. *Introduction to semi-supervised learning*. 2009.