

Class-Aware Contrastive Semi-Supervised Learning

Fan Yang^{2*} Kai Wu^{1*} Shuyi Zhang¹ Guannan Jiang⁴ Yong Liu¹
Feng Zheng^{3†} Wei Zhang⁴ Chengjie Wang^{1†} Long Zeng²

¹Tencent Youtu Lab ²Tsinghua University ³Southern University of Science and Technology ⁴CATL

Abstract

Pseudo-label-based semi-supervised learning (SSL) has achieved great success on raw data utilization. However, its training procedure suffers from confirmation bias due to the noise contained in self-generated artificial labels. Moreover, the model’s judgment becomes noisier in real-world applications with extensive out-of-distribution data. To address this issue, we propose a general method named Class-aware Contrastive Semi-Supervised Learning (CCSSL), which is a drop-in helper to improve the pseudo-label quality and enhance the model’s robustness in the real-world setting. Rather than treating real-world data as a union set, our method separately handles reliable in-distribution data with class-wise clustering for blending into downstream tasks and noisy out-of-distribution data with image-wise contrastive for better generalization. Furthermore, by applying target re-weighting, we successfully emphasize clean label learning and simultaneously reduce noisy label learning. Despite its simplicity, our proposed CCSSL has significant performance improvements over the state-of-the-art SSL methods on the standard datasets CIFAR100 [18] and STL10 [8]. On the real-world dataset Semi-iNat 2021 [27], we improve FixMatch [25] by 9.80% and CoMatch [19] by 3.18%. Code is available <https://github.com/TencentYoutuResearch/Classification-SemiCLS>.

1. Introduction

Raw data utilization is becoming a research focal point for its high accessibility and high affordability. The definition of raw data is the union set of in-distribution data (known classes and balanced distribution) and out-of-distribution data [32] (unknown classes or unbalanced distribution), as shown in Fig. 1. For in-distribution datasets, semi-supervised learning (SSL) has achieved excellent performance with the help of pseudo labeling [25] [3] [2] [13]. The primary process for the pseudo-label based SSL is it-

*Equal contribution.

†Corresponding author.

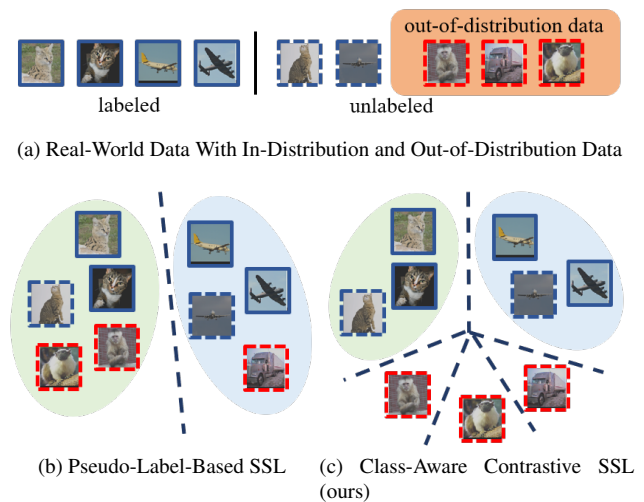


Figure 1. Intuition graph for class-aware contrastive semi-supervised learning. (a) represents real-world unlabeled data containing in-distribution and out-of-distribution data (unbalanced distribution or unknown classes). Unlike pseudo-label-based semi-supervised learning in (b), which wrongly clusters noisy out-of-distribution data, CCSSL in (c) reduces noise by image-level contrastive learning on out-of-distribution data while maintaining the class-aware clustering ability for in-distribution data.

erative training: 1) creating self-generated pseudo labels on raw data, 2) training on pseudo labels, 3) repeating 1 and 2. The underlying assumption for training on pseudo labels is that the distribution of the labeled data is close to the unlabeled, and the unlabeled dataset does not contain any novel categories. This assumption often does not hold in the real-world applications with extensive out-of-distribution data, which contains unbalanced distribution or unknown classes. SSL’s training and labeling loop collapses in the real-world because of the enormous noise introduced by pseudo labeling on out-of-distribution data. The confirmation bias [1] induced by noisy pseudo labels deteriorates SSL performance by a large margin [26].

SSL has used many techniques to alleviate the confirmation bias and achieved state-of-the-art results on standard in-

distribution benchmarks, like CIFAR [18] and STL10 [8]. Some methods achieve this goal by using the model’s self-correcting capability. In [25] [3] [2] [25], a high confidence threshold is used to filter incorrectly pseudo-labeled data. [34] [20] [29] use a self-ema teacher to generate pseudo labels with the assumption that the weighted averaged model produces more stable and less noisy predictions. Some methods [28] [31] try to narrow the distribution gap between raw data and labeled data by predictions’ uncertainty. However, justifying a model’s prediction by its output without introducing other information still suffers from confirmation bias [21]. Especially on the real-world data with unbalanced distribution or unknown classes, the effect of a model’s self-correcting is facing a huge challenge. By explicitly introducing contrastive information, our CCSSL alleviates the confirmation bias in the feature space and shows great de-noise ability on the real-world data.

Another trend is using prior or posterior information to co-rectify a model’s predictions. Both [11] [5] use a pre-trained model as a universal prior information. In [21], by knowledge embedding graph, semantic information is used to regularize feature learning. Although the above methods are proved helpful, fixed prior is hard to be blended into downstream tasks, resulting in inferior performance in practice. [19] utilizes posterior information by constructing a prediction graph and introducing contrastive learning from self-supervised learning (Self-SL). However, directly combining image-level feature repulsion interferes with SSL’s class-clustering ability. To solve this conflict, in our proposed CCSSL, a class-aware contrastive module is specifically explored for reliable in-distribution data clustering and noisy out-of-distribution data contrasting to better integrate with downstream classification tasks.

To this end, we propose a class-aware contrastive semi-supervised learning (CCSSL) method. CCSSL consists of a semi-supervised module and a class-aware contrastive module. Any end-to-end pseudo-label-based SSL can replace the semi-supervised module to benefit from the class-aware contrastive module’s confirmation bias alleviation ability. Unlike self-correcting methods that try to alleviate noise by model’s own output, CCSSL regularizes the training process by introducing information in the feature space. As shown in the Fig. 1, CCSSL constructs the feature space with high dimensional vectors from two strong augmented views. Rather than directly combining contrastive learning, CCSSL uses class-aware clustering on in-distribution data to maintain SSL’s clustering ability and image-level contrasting on out-of-distribution data for noise alleviation. Furthermore, to reduce the implicit noise introduced from the class-aware clustering, we incorporate a target re-weighting module that emphasizes learning on high probable in-distribution samples and reduces the effect of uncertain noisy samples. We found that with the

help of the re-weighting module, the prediction confidence of learned models is robust enough for roughly discriminating in-distribution and out-of-distribution data.

Our contributions can be described in three folds:

- We propose a novel SSL learning method CCSSL, which takes advantage of both SSL’s effective learning and Self-SL’s noise alleviation ability by class-wise clustering on in-distribution data and image-wise contrasting on out-of-distribution data.
- CCSSL is a drop-in helper that can improve upon any end-to-end pseudo-label-based SSL method for confirmation bias alleviation and faster convergence. CCSSL’s noise alleviation capability makes SSL methods more practical in the real-world setting.
- We thoroughly analyzed the effect of CCSSL with various state-of-the-art SSL methods on both in-distribution data and real-world data. By simply combining CCSSL, current state-of-the-art SSL methods can be further improved.

2. Related Work

In this section, we describe the recent trends of deep learning based SSL and contrastive based Self-SL. Then we analyze specific methods aiming at using on the real-world setting.

2.1. Semi-Supervised Learning

Deep learning based SSL mainly consists of pseudo-labeling and consistency regularization. Pseudo-label-based methods utilize unlabeled data by training on self-generated predictions. The self-training process validates the model itself and can also be named as self-confirming [9]. However, using pseudo labels suffers confirmation bias [1] because it is easy to overfit incorrect predictions during training. Both [25] [3] use high confidence predictions to filter noisy data. [15] makes a detour by label propagation to spread labeled data distribution on unlabeled data. Consistency based methods [16] [29] [25] aim to produce consistent predictions on different image views based on the manifold assumption that different views of the same image should lie on the same point in the high dimensional space. [25] [30] create different views of an image by strong augmentations to simulate image perturbations of an object. Lately, pseudo-label-based consistency training [25] [3] [2], which use a weak augmentation for creating pseudo labels and a strong augmentation for consistency training, has dominated and achieved excellent performance gain. However, the research on the real-world setting with extensive out-of-distribution data has been less studied. With evaluating in-distribution data and real-world data, we show CCSSL’s ability of alleviating confirmation bias in practical.

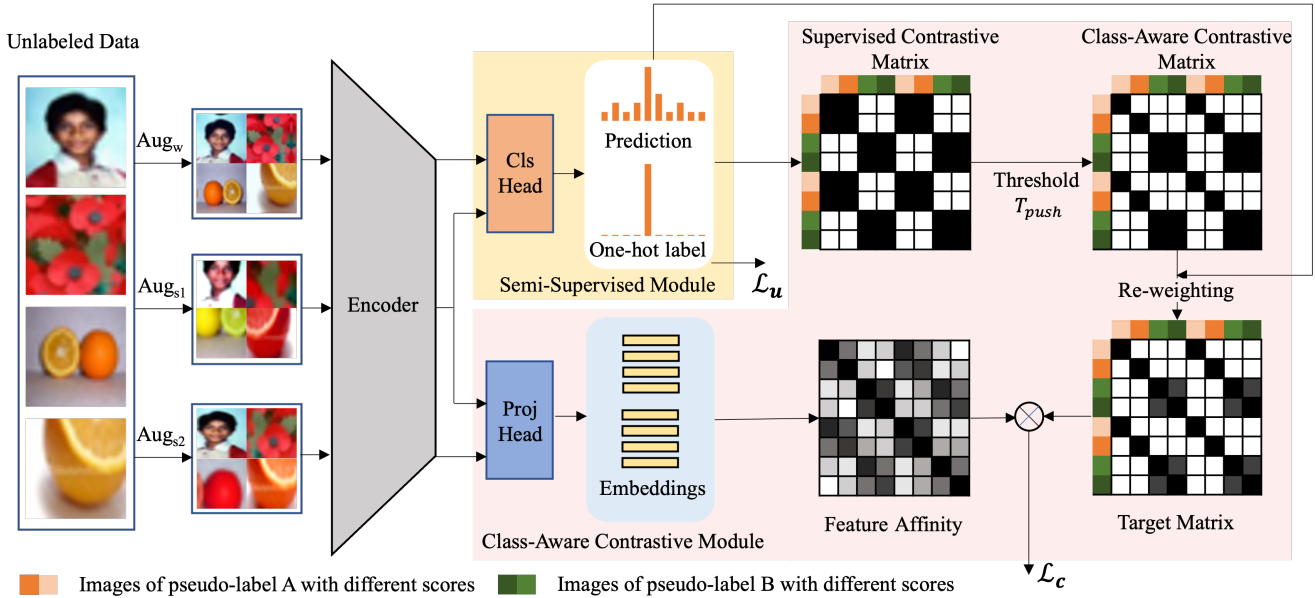


Figure 2. Framework of our proposed CCSL. Given a batch of unlabeled images, the weakly augmented views will go through a semi-supervised module that can be borrowed from any pseudo-label-based SSL to generate model predictions. With pseudo labels, we make a supervised contrastive matrix with only class-level information. Then, the class-aware contrastive matrix is formed by image-level contrasting on out-of-distribution data to reduce confirmation bias. By applying a re-weighting module, we focus learning on clean data and get the final target matrix. Besides, feature affinity matrix is made by two strong augmented views. Class-aware contrastive module for unlabeled data is formulated by minimizing cross-entropy between the affinity matrix and the target matrix.

2.2. Self-Supervised Learning

Self-SL has a long story for learning a universal representation without supervision. [12] [14] build the base for contrastive learning by proposing and refining the noise contrastive estimation (NCE). Based on NCE, [22] proposes infoNCE which inspires learning representations to discriminate different views of a sample among other samples. Moco [11] and simCLR [5] train a siamese network [6] to learn image-wise feature by pulling similar (positive) samples and pushing apart different (negative) samples [7]. Inspired by contrastive learning, we formulate positive pairs and negative pairs in CCSL to regularize feature training. However, directly combining contrastive learning ignores the internal conflict between image-level feature repulsion and SSL’s class-level feature clustering. By incorporating class-level information, CCSL can seamlessly blend into downstream tasks.

2.3. Real-World Unlabeled Data Learning

Real-world unlabeled data usually contains out-of-distribution data with unbalanced class distribution or unknown classes. Current pseudo-label-based SSL suffers from confirmation bias by incorrect pseudo-labeling on similar unknown classes or blurry images from out-of-distribution data, resulting in worse performance than supervised baseline [26]. Open set classification is a subfield

of real-world unlabeled data learning focusing on novel categories. [10] keeps tracking the effect of the supervised learning model to prevent performance hazards. By incorporating deep-construction-based representation learning, [33] learns to distinguish unknowns from knowns. Open set classification methods are proven to be effective on handwriting datasets or self-constructed datasets but seldom generalize on real-world data. [26] conducts a systematic evaluation of several recently proposed SSL techniques on a real-world setting by two challenging datasets and shows that recent SSL’s performance pales on real-world datasets in comparison to in-distribution standard datasets. Therefore, this paper further explores the real-world setting and narrows the gap between research and real-world applications by confirmation bias alleviation for out-of-distribution data.

3. Method

3.1. Problem Definition

In the setting of a semi-supervised classification task, let $\mathcal{X} = \{(x_i, y_i) : i \in (1, \dots, B)\}$ be a batch of labeled data, where x_i is the i th image and y_i represents its one-hot label, and $\mathcal{U} = \{u_i : i \in (1, \dots, \mu B)\}$ be a batch of unlabeled data, where μ is a hyperparameter that determines the relative ratio of \mathcal{X} and \mathcal{U} . In the real world, unlabeled data

\mathcal{U} contains in-distribution data \mathcal{U}_{in} and out-of-distribution data \mathcal{U}_{out} , which makes pseudo labels more susceptible to confirmation bias [21].

3.2. Framework

Our proposed CCSSL consists of a replaceable semi-supervised module and a class-aware contrastive module, as in Fig. 2. Our method focuses on introducing class-aware contrastive information to alleviate confirmation bias for semi-supervised module especially in real-world setting. Any pseudo-label-based SSL can be used as the semi-supervised module to enjoy the benefit. The notations are as follows:

Data augmentation provides different views of a single image. For a labeled image x_i , a weak augmentation is used for supervised learning. For an unlabeled image u_i , we apply a weak augmentation $Aug_w(\cdot)$ and two strong augmentations $Aug_s(\cdot)$. The Aug_w and one Aug_s are sent to semi-supervised module, while all Aug_s are used for class-aware contrastive learning.

Encoder $F(\cdot)$ is used to extract representation $r = F(Aug(x))$ for a given input x .

Semi-Supervised module can be replaced by any pseudo-label based semi-supervised learning method. The module generates pseudo labels during training and outputs the final prediction at inference time by a classification head $P_{cls}(\cdot)$. We use FixMatch [25] in the method for simplicity.

Class-Aware contrastive module constructs an embedding space by mapping the high-dimensional representation r to a low-dimensional embedding z through the projection head $P_{proj}(\cdot)$. We propose the class-aware contrastive learning in the embedding space to improve the robustness for out-of-distribution data and compatibility with downstream tasks.

3.3. Semi-Supervised Module

Following the general form of pseudo-label-based SSL, the module contains a supervised loss \mathcal{L}_x and an unsupervised loss \mathcal{L}_u . \mathcal{L}_x is the same as [25]. \mathcal{L}_u relies on pseudo labels $\hat{q}_i = \arg \max(p_i)$ from the model's prediction $p_i = P_{cls}(Aug_w(u_i))$ on the weak view of image u_i . Only pseudo labels with high confidence $q \geq t$, where $q = \max(p)$, will be retained for consistency training. H means cross entropy:

$$\mathcal{L}_u = \frac{1}{\mu B} \sum_{i=1}^{\mu B} \mathbb{1}(\max(p_i) \geq t) H(\hat{q}_i, P_{cls}(Aug_s(u_i))). \quad (1)$$

However, the manual set threshold t is difficult to guarantee the accuracy of pseudo labels \hat{q} , especially on real-world setting with out-of-distribution data. The phenomenon of overfitting on self-generated labels is called

confirmation bias [1], and it can hardly be alleviated by the model itself during training. Therefore, we propose a novel class-aware contrastive learning method, which introduces unbiased information in the feature space.

3.4. Class-Aware Contrastive Module

In Fig. 2, we seamlessly integrate clustering and contrasting in the feature space and apply re-weighting for attentive training. Given unlabeled images in a batch, we separate images into two parts based on the likelihood of in-distribution and out-of-distribution, in which T_{push} is used as the threshold for separation from the model's prediction confidence. For highly possible in-distribution data $\max(p) > T_{push}$, we apply class-aware clustering in the feature space to seamlessly blend into downstream tasks. For out-of-distribution data $\max(p) < T_{push}$, we follow contrastive learning mechanism [5] by taking a image's views as positive samples and others as negative. Then we apply foreground probability re-weighting for each sample to focus learning on high confidence clean data. Next, we will show how to formulate the method step by step.

Contrastive learning aims to learn a universal prior information for downstream tasks. As in [11] [5], we randomly sample a minibatch of N images and each image u_i has two strong augmentations Aug_s which are used to extract features r through encoder $F(\cdot)$. Then, r is mapped to obtain a square normalized low-dimensional embedding z by the projection head. Finally, we get an embedding's affinity matrix $S \in \mathbb{R}^{2N \times 2N}$ by the dot product of embeddings $\mathcal{Z} = \{z_i : i = 1, \dots, 2N\}$ with a temperature factor τ , as follow:

$$s_{ij} = \exp(z_i \cdot z_j / \tau), \quad (2)$$

Then, a contrastive matrix $\mathbf{W}_{con} \in \mathbb{R}^{2N \times 2N}$ for loss calculation with S is formulated as:

$$w_{ij}^{con} = \begin{cases} 1 & \text{if } j = i, \\ 1 & \text{if } z_i \text{ and } z_j \text{ are from the same image's} \\ & \text{views,} \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

The self-supervised contrastive loss InfoNCE [22] takes the following format:

$$\mathcal{L}^{\text{InfoNCE}} = - \sum_{i=1}^{2N} \log \frac{\exp(z_i \cdot z_i^* / \tau)}{\sum_{j=1}^{2N} \mathbb{1}_{(j \neq i)} \exp(z_i \cdot z_j / \tau)}, \quad (4)$$

where z_i^* is from the other Aug_s of the same image as z_i . The ideal solution for contrastive learning is to push each image far from others and pull same image views closer in the feature space. The property is effective to alleviate noise on the out-of-distribution data. However, it is not compatible with the classification task which requires to cluster images at the class level. To solve the conflict between contrastive learning and downstream tasks, we formulate the

class-aware contrastive module to take advantages of both Self-SL’s generalizability and SSL’s efficient training.

Class-Aware Contrastive learning considers class-aware information for in-distribution data while alleviating noise for out-of-distribution data. Inspired by [17], we consider the embeddings z from the same category as positive pairs instead of the same image with the help of the pseudo-labels \hat{q} from the semi-supervised module to form supervised contrastive matrix W_{scon} in Fig. 2. However, class-wise clustering by pseudo-labels is unsafe to use because a model’s prediction may contain a large amount of noises especially for out-of-distribution data, as explained in Sec. 1. Therefore, we introduce contrastive learning for noise regularization. We use T_{push} as the threshold for clustering and contrasting in the feature space, and then W_{scon} is transformed to class-aware contrastive matrix W_{clacon} as shown in Fig. 2. For model’s prediction $p > T_{push}$, we assume the images have a high probability to be in-distribution data U_{in} and should be pulled closer with the same class. For $p < T_{push}$, our formula degenerates to contrastive learning where only Aug_s of the same image are positive pairs. Each element w_{ij}^{clacon} in W_{clacon} has the following format:

$$w_{ij}^{clacon} = \begin{cases} 1 & \text{if } j = i, \\ 1 & \text{if } z_i \text{ and } z_j \text{ are from same category,} \\ & \text{and both } q_i \text{ and } q_j > T_{push}, \\ 0 & \text{otherwise,} \end{cases} \quad (5)$$

where i and j are indices of z for augmented sample. q_i denotes the pseudo label confidence score of the i th augmented views. Although class-aware contrastive is theoretically functional, directly using T_{push} to judge the cleanness of data is problematic since images just above the T_{push} might be out-of-distribution data too. To solve this issue, we propose a foreground re-weighting module to reduce the noise effect on out-of-distribution data.

Foreground Re-weighting is used to further emphasize training on in-distribution data with high confidence and weaken the potential bias brought from out-of-distribution data. We formulate the final target matrix W_{target} by re-weight W_{clacon} with the foreground probability from model’s prediction on Aug_w . Simply by multiplying confidence q of two samples, we get a re-weighting factor $w_{re} = q_i \cdot q_j$, where i and j represent the indices of row and column in W_{target} , respectively. Each element w_{ij}^{target} of W_{target} is defined as follows:

$$w_{ij}^{target} = \begin{cases} q_i \cdot q_j \cdot w_{ij}^{clacon} & \text{if } i \neq j, \\ w_{ij}^{clacon} & \text{otherwise.} \end{cases} \quad (6)$$

The loss of the class-aware contrastive module L_c is formulated by minimizing the cross-entropy between S in Eq. (2) and W_{target} . The loss of L_c can be defined as:

$$\mathcal{L}_c = H(S, W_{target}) = \sum_{i=1}^{2N} \frac{1}{1 + |P(i)|} \mathcal{L}_{c,i}, \quad (7)$$

where H denotes cross entropy and $P(i)$ represents indices of the views from other images of the same class with confidence $p > T_{push}$. $|P(i)|$ denotes its cardinality and $|P(i)| + 1$ represents all positive pairs. $\mathcal{L}_{c,i}$ as below:

$$\mathcal{L}_{c,i} = -\log \frac{\exp(z_i \cdot z_i^* / \tau)}{\sum_{j=1}^{2N} \mathbb{1}_{j \neq i} \exp(z_i \cdot z_j / \tau)} - \sum_{p \in P(i)} \log \frac{w_{ij} \cdot \exp(z_i \cdot z_p / \tau)}{\sum_{j=1}^{2N} \mathbb{1}_{j \neq i} \exp(z_i \cdot z_j / \tau)}, \quad (8)$$

where z_i^* is the embedding of Aug_s rather than z_i from the same image. We calculate the final total loss Eq. (9) using a weighted sum of supervised loss \mathcal{L}_x , semi-supervised loss \mathcal{L}_u and class-aware contrastive loss \mathcal{L}_c . \mathcal{L}_x and \mathcal{L}_u are from Sec. 3.3. λ_u and λ_c are the weights for semi-supervised loss and class-aware contrastive loss, respectively.

$$\mathcal{L} = \mathcal{L}_x + \lambda_u \mathcal{L}_u + \lambda_c \mathcal{L}_c \quad (9)$$

3.5. Applications

CCSSL method can be used with any pseudo-label-based SSL to alleviate confirmation bias in the training process. Applying CCSSL to SSL methods is simple. We applied the CCSSL to the mainstream SSL methods, like MixMatch [3], FixMatch [25] and CoMatch [19]. MixMatch and FixMatch are based on pseudo-labels and have a strong augmentation Aug_s for unlabeled data learning. We create another Aug_s for class-aware contrastive module, which trains parallelly with the original network. CoMatch also incorporates contrastive learning, but no class-wise clustering is involved. Since CoMatch already has two Aug_s , we directly use them for class-aware contrastive module. Our experiments in Sec. 4 show CCSSL not only improves the performance for SSL methods but also speeds up the training process.

4. Experiments

In this section, we first evaluate the effectiveness of CCSSL on in-distribution datasets, which may also contain blur out-of-distribution cases. Then, we show CCSSL’s ability of confirmations bias alleviation on real-world data. Last, a qualitative comparison is presented.

4.1. Datasets

We assume that the harder the dataset, the higher noise it has, so we conduct experiments on the following different noise level datasets.

Method	CIFAR100			CIFAR10			STL10
	400	2500	10000	40	250	4000	
Mixmatch [3]	32.39±1.32	60.06±0.37	71.69±0.33	52.46±11.5	88.95±0.86	93.58±0.10	38.02±8.29
ReMixMatch [2]	55.72±2.06	72.57±0.31	76.97±0.56	80.90±9.64	94.56±0.05	95.28±0.13	-
SSWPL [28]	-	73.48±0.45	79.12±0.85	-	-	-	-
LaplaceNet [23]	-	68.36±0.02	73.40±0.23	-	-	95.35±0.07	-
CoMatch [19]	58.11±2.34	71.63±0.35	79.14±0.36	93.09±1.39	95.09±0.33	95.44±0.20	79.80±0.38
FixMatch [19]	51.15±1.75	71.71±0.11	77.40±0.12	86.19±3.37	94.93±0.65	95.74±0.05	65.38±0.42
CCSSL(FixMatch)	61.19±1.65	75.7±0.63	80.68±0.16	90.83±2.78	94.86±0.55	95.54±0.20	80.01±1.39

Table 1. Top-1 Accuracy for in-distribution datasets including CIFAR100, CIFAR10, and STL10. On high noise-level datasets CIFAR100 and STL10, we achieve the best performance by simply adding CCSSL to Fixmatch. On the easier dataset CIFAR10 with less noise, CCSSL only provides marginal performance gain. '-' means not self-implemented.

Method	Semi-iNat 2021			
	From Scratch		From MoCo Pretrain	
	Top1	Top5	Top1	Top5
Supervised	19.09	35.85	34.96	57.11
MixMatch [3]	16.89	30.83	-	-
+CCSSL	19.65	35.09	-	-
FixMatch [25]	21.41	37.65	40.3	60.05
+CCSSL	31.21	52.25	41.28	64.3
CoMatch [19]	20.94	38.96	38.94	61.85
+CCSSL	24.12	43.23	39.85	63.68

Table 2. Accuracy for real-world dataset - Semi-iNat 2021. We evaluated two scenarios: random starting with serious noise and pre-trained starting with less noise. In the noisy random starting scenario, CCSSL achieves significant performance gain by noise alleviation. In the noiseless pretrained starting scenario, CCSSL achieves marginal improvement because noise has been reduced by the pretrained model in the beginning.

CIFAR10 [18] is a relatively easy-level in-distribution dataset consisting of 60K 32x32 colour images in 10 classes, with 6K images per class. There are 50K training images and 10K test images. We conduct three different experiments on 40, 250, 4000 random selected images in the class balanced way, following [25], as in Tab. 1.

STL10 [8] is a relatively medium-level in-distribution dataset comprises 10 classes. The training set has 5K images (10 predefined folds), 100k unlabeled images containing similar distribution with training images and 800 test images per class. All images are 96x96 pixels. We conduct experiments on all folds following [25], as shown in Tab. 1

CIFAR100 [18] is a relatively hard-level in-distribution dataset comprises 100 classes. The training set contains 50K images, while test contains 10K images. All images have a fixed resolution of 32x32. We conduct three different experiments on 400, 2500, 10000 random selected images in a class-balanced way, following [25], as shown in Tab. 1.

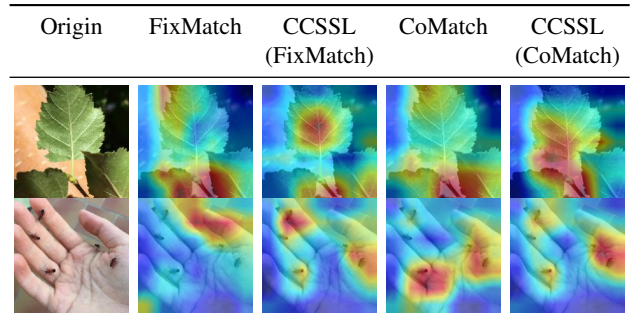


Table 3. Qualitative comparison on FixMatch and CoMatch with/without CCSSL by Grad-CAM [24]. With CCSSL’s de-noise effect, an SSL method focuses more on the foreground objects. CCSSL methods find all ants for the ant image and capture the correct position on the leaf image.

Semi-iNat 2021 [27] is a real-world dataset designed to expose some of the challenges encountered in a realistic setting, such as significant class imbalance, domain mismatch between the labeled and unlabeled data, and large unknown classes. The labeled training and val part each has 9721 and 4050 images. The unlabeled training part has 313248 images. Experiments in Tab. 2

4.2. In-distribution Data Evaluation

For CIFAR10, we set $\lambda_c = 0.2$ with $T_{push} = 0$ due to low noise level of the dataset and no out-of-distribution data. For high noise level STL10 and CIFAR100, we set $\lambda_c = 1.0$ with $T_{push} = 0$. We train 512 epochs for all experiments because of fast convergence speed. For other hyperparameters, optimizer and learning rate schedule, our setting is the same as FixMatch [25]. We compare CCSSL with MixMatch [3], ReMixMatch [2], FixMatch [25], SSWPL [28], LaplaceNet [23] and CoMatch [19]. Since the last three SSL methods can be viewed as an extension of FixMatch, we directly apply CCSSL on FixMatch for a fair comparison.

In Tab. 1, We find that the performance of CCSSL

is higher when the task is noisy and more confirmation bias involved. On CIFAR100 and STL10, the noisiest in-distribution datasets we use, we achieve the best performance only by adding CCSSL on FixMatch. CCSSL is also on par with state-of-the-art methods on CIFAR10. On CIFAR100, we achieve +10.04%/ +3.99%/ +3.26% over origin FixMatch and +3.08%/ +2.22%/ +1.54% over the best of all compared SSL methods for 400, 2500, and 10000 labels respectively. The performance gain decline can be explained by the confirmation bias level decrease with more training data. For the medium hard STL10, we achieve +14.62% over origin FixMatch and only +0.2% over the best of all compared SSL methods. CCSSL’s performance improvement decreases on the easier dataset with fewer categories. On CIFAR10, we achieve +4.64/ -0.07%/ -0.2% over FixMatch and -2.26%/ -0.23%/ -0.2% over the best of all compared SSL methods. With the decrease of noise level, our CCSSL’s performance improvement pales compared with other variations of SSL methods tuned for in-distribution datasets but is still helpful for FixMatch. Another thing that needs to be noticed is that when the noise level is the least (CIFAR10 with 4000 labels), FixMatch achieves the best performance by the simplest structure - which means that when the data is clear and enough, the structure of SSL becomes less critical.

4.3. Realistic Evaluation

On Semi-iNat 2021, we simulate two situations with different noise levels - (high: training from scratch, low: training from pre-trained model), as shown in Tab. 2. The results show that CCSSL can improve SSL in both situations but more noise more improvement.

For Semi-iNat 2021, we use an image size of 224x224, resnet50 backbone, projection head with a dimension of 64 like in [11], and training on 4 V100 with batchsize 64 for labeled data and batchsize 448 with $\mu = 7$ for unlabeled data. For FixMatch [25], pseudo label thresh $\tau = 0.8$ achieves the best results, but 0.6 when combining with our method. For MixMatch [3] and CoMatch [19], we directly follow their papers’ setting. We set $\lambda_c = 2.0$ with $T_{push} = 0.9$ due to high noise level on out-of-distribution data. When training from Self-SL pre-trained models, we freeze the first three blocks of resnet50 because of the forgetting problem. Other settings are the same as [26].

In Tab. 2 training from scratch, state-of-the-art SSL methods only have on-par performance with the supervised baseline. Compared to the performance on in-distribution datasets, the noise in pseudo-labels harms SSL methods by a large margin. With class-aware contrastive module, we are able to alleviate confirmation bias on out-of-distribution data while maintaining downstream tasks’ clustering ability. We improve state-of-the-art SSL by a large margin when using with CCSSL (+2.7%/ +9.8%/ +3.1% on MixMatch,




Figure 3. Convergence speed and performance difference on MixMatch and FixMatch with/without CCSSL. As shown in the figure, CCSSL methods achieves better performance with much fewer training epochs.

FixMatch and CoMatch separately). The experiments show that CCSSL can be a big helper for pseudo-label-based SSL to reduce the negative impact from noisy predictions when training from scratch on the real-world dataset.

In Tab. 2, training from MoCo Semi-iNat pre-trained model simulates the situation where SSL starts with less noise [4]. CCSSL only improves +1.05%/ +0.91% on FixMatch and CoMatch. The result shows that when a model starts with less noise, our proposed CCSSL can also reduce noise introduced in the training process.

4.4. Qualitative Evaluation

In Tab. 3, we show the qualitative results on FixMatch and CoMatch with/without CCSSL. With CCSSL’s de-noise effect, an SSL method focuses more on the foreground objects. CCSSL methods find all ants for the ant image and capture the correct position on the leaf image. We choose FixMatch and CoMatch for visualization because the two are the best on Semi-iNat among SSL methods we tested.

4.5. Convergence Speed

In Fig. 3, we study the convergence speed using CCSSL on MixMatch and FixMatch. We observe that CCSSL methods achieve the best performance with much fewer training iterations. This phenomenon is the same when training on clean pseudo-labels [19] because of the noise alleviation effect of CCSSL. More specifically, better pseudo labels in the early phase helps better guide the learning process and translate to faster convergence.

5. Ablation

5.1. Contrastive & Class-Aware & Re-weighting

In Tab. 4, we investigate three main techniques of CCSSL, including contrastive learning, class-aware, and re-weighting, based on FixMatch. We observe that all components are helpful. However, directly using contrastive

con- trastive	re- weighting	class- aware	Semi- iNat	CI- FAR100
			21.41	72.54
✓			27.73	72.37
✓	✓		29.48	72.67
✓		✓	30.49	75.89
✓	✓	✓	31.21	75.96

Table 4. Accuracy for different combinations of contrastive learning, class-aware and re-weighting.

Ratio (μ)	Semi-iNat	
	$T_{push} = 0$	$T_{push} = 0.9$
2	23.01	21.93
4	23.88	25.19
5	25.33	26.81
6	22.32	24.47
7	23.14	27.46

Table 5. Different ratios of unlabeled data on Semi-iNat with different degrees of class-aware (T_{push}). When $T_{push} = 0$ (no contrastive involved), smaller ratios achieves better results because of noise in unlabeled data. After setting $T_{push} = 0.9$ to regularize noise, larger ratio achieves better results.

learning deteriorates the performance on in-distribution dataset CIFAR100 while largely benefiting the real-world dataset Semi-iNat 2021 with extensive out-of-distribution data. CCSSL found the equilibrium point for contrastive and clustering by class-aware contrastive and re-weighting, and improved FixMatch by a large margin on both in-distribution and real-world datasets.

5.2. Ratio of Unlabeled Data

As shown in Tab. 5, we experiment on different values of μ , which defines the ratio of labeled and unlabeled data in a batch. We observe that when only using class-aware clustering without contrastive regularization ($T_{push} = 0$), CCSSL achieves the best result with a small ratio of unlabeled data ($\mu = 5$). However, after we increase the effect of contrastive learning by changing $T_{push} = 0.9$ (sample confidence < 0.9 will do contrastive learning), the larger ratio has better performance ($\mu = 7$). Less noisy dataset like CIFAR100 also benefits from large unlabeled data ratio [21]. This phenomenon can be explained by the higher noise amount introduced from self-generated labels with larger ratio of unlabeled data. With CCSSL’s noise alleviation, SSL methods achieve the best performance with a larger unlabeled ratio.

5.3. Threshold for Class-Aware and Contrastive

We found that our algorithm is robust with varying the threshold T_{push} from 0.6 to 0.9. However, the optimal

T_{push}	Semi-iNat	CIFAR100@I2500
0	23.88	75.59
0.4	23.38	75.96
0.6	24.49	75.76
0.7	24.64	75.9
0.8	23.63	75.7
0.9	25.21	75.79

Table 6. Experiments on different class-aware thresholds. Our algorithm is robust with varying T_{push} from 0.6 to 0.9. Semi-iNat has a larger noise level and needs larger T_{push} which proves CCSSL’s ability for noise alleviation.

threshold for class-aware and contrastive learning is inconsistent on different noise level datasets. For high noise level Semi-iNat with unbalanced distribution and unknown classes, $T_{push} = 0.9$ achieves the best performance. For low noise level CIFAR100, $T_{push} = 0.4$ achieves the best performance. The result proves CCSSL’s noise-reducing ability by needing to regularize more on noisy datasets.

6. Limitations

Potential harmful training fluctuation does exist in our experiments: performance on Semi-iNat fluctuates more than on CIFAR100 in the Fig. 3. The main reason lies in the lower quality of pseudo-labels for high noise-level data. Increasing the weight for the class-aware contrastive module or training with different seeds to average the effect of fluctuation can alleviate the problem. During our experiments, we fixed the weight and seed for a fair comparison.

7. Conclusion

We proposed CCSSL, a general confirmation bias alleviation method for pseudo-label-based SSL methods. CCSSL reduces noise by constructing another feature space instead of using a model’s output. By class-aware contrastive module, we apply image-level contrastive on out-of-distribution data for noise alleviation and class-level clustering on in-distribution data for blending into downstream tasks. We have conducted extensive experiments over different noise level datasets, including in-distribution datasets and a real-world dataset, to demonstrate the effectiveness of CCSSL. With simply replacing the semi-supervised module, we improve the state-of-the-art SSL [25] [3] [19] by a large margin. With the effort above, CCSSL is proved to be helpful in making SSL more practical in the real world.

Potential negative societal impact CCSSL is a fundamental technology that can generate robust semi-supervised classification models in the real world. Maliciously using a classification model to negatively impact society is possible.

References

- [1] Eric Arazo, Diego Ortego, Paul Albert, Noel E O'Connor, and Kevin McGuinness. Pseudo-labeling and confirmation bias in deep semi-supervised learning. In *2020 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2020. 1, 2, 4
- [2] David Berthelot, Nicholas Carlini, Ekin D Cubuk, Alex Kurakin, Kihyuk Sohn, Han Zhang, and Colin Raffel. Remixmatch: Semi-supervised learning with distribution alignment and augmentation anchoring. *arXiv preprint arXiv:1911.09785*, 2019. 1, 2, 6
- [3] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin Raffel. Mixmatch: A holistic approach to semi-supervised learning. *arXiv preprint arXiv:1905.02249*, 2019. 1, 2, 5, 6, 7, 8
- [4] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2021. 7
- [5] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020. 2, 3, 4
- [6] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15750–15758, 2021. 3
- [7] Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised vision transformers. *arXiv preprint arXiv:2104.02057*, 2021. 3
- [8] Adam Coates, Andrew Ng, and Honglak Lee. An analysis of single-layer networks in unsupervised feature learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 215–223. JMLR Workshop and Conference Proceedings, 2011. 1, 2, 6
- [9] Merriam-Webster Dictionary. Merriam-webster. *On-line at <http://www.mw.com/home.htm>*, 8, 2002. 2
- [10] 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*, pages 3897–3906. PMLR, 2020. 3
- [11] 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*, pages 9729–9738, 2020. 2, 3, 4, 7
- [12] Geoffrey E Hinton. Training products of experts by minimizing contrastive divergence. *Neural computation*, 14(8):1771–1800, 2002. 3
- [13] Zijian Hu, Zhengyu Yang, Xuefeng Hu, and Ram Nevatia. Simple: Similar pseudo label exploitation for semi-supervised classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15099–15108, 2021. 1
- [14] Aapo Hyvärinen and Peter Dayan. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(4), 2005. 3
- [15] Ahmet Iscen, Giorgos Tolias, Yannis Avrithis, and Ondrej Chum. Label propagation for deep semi-supervised learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5070–5079, 2019. 2
- [16] Jisoo Jeong, Seungeui Lee, Jeessoo Kim, and Nojun Kwak. Consistency-based semi-supervised learning for object detection. *Advances in neural information processing systems*, 32:10759–10768, 2019. 2
- [17] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschiot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *arXiv preprint arXiv:2004.11362*, 2020. 5
- [18] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 1, 2, 6
- [19] Junnan Li, Caiming Xiong, and Steven CH Hoi. Comatch: Semi-supervised learning with contrastive graph regularization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9475–9484, 2021. 1, 2, 5, 6, 7, 8
- [20] Yen-Cheng Liu, Chih-Yao Ma, Zijian He, Chia-Wen Kuo, Kan Chen, Peizhao Zhang, Bichen Wu, Zsolt Kira, and Peter Vajda. Unbiased teacher for semi-supervised object detection. *arXiv preprint arXiv:2102.09480*, 2021. 2
- [21] Islam Nassar, Samitha Herath, Ehsan Abbasnejad, Wray Buntine, and Gholamreza Haffari. All labels are not created equal: Enhancing semi-supervision via label grouping and co-training, 2021. 2, 4, 8
- [22] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 3, 4
- [23] Philip Sellars, Angelica I Aviles-Rivero, and Carola-Bibiane Schönlieb. Laplacenet: A hybrid energy-neural model for deep semi-supervised classification. *arXiv preprint arXiv:2106.04527*, 2021. 6
- [24] Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 618–626, 2017. 6
- [25] 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. *arXiv preprint arXiv:2001.07685*, 2020. 1, 2, 4, 5, 6, 7, 8
- [26] Jong-Chyi Su, Zezhou Cheng, and Subhansu Maji. A realistic evaluation of semi-supervised learning for fine-grained classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12966–12975, 2021. 1, 3, 7
- [27] Jong-Chyi Su and Subhansu Maji. The semi-supervised inaturalist challenge at the fgvc8 workshop, 2021. 1, 6
- [28] Fariborz Taherkhani, Ali Dabouei, Sobhan Soleymani, Jeremy Dawson, and Nasser M Nasrabadi. Self-supervised

- wasserstein pseudo-labeling for semi-supervised image classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12267–12277, 2021. 2, 6
- [29] Yihe Tang, Weifeng Chen, Yijun Luo, and Yuting Zhang. Humble teachers teach better students for semi-supervised object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3132–3141, 2021. 2
- [30] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *arXiv preprint arXiv:1703.01780*, 2017. 2
- [31] Zhenyu Wang, Yali Li, Ye Guo, Lu Fang, and Shengjin Wang. Data-uncertainty guided multi-phase learning for semi-supervised object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4568–4577, 2021. 2
- [32] Renzhe Xu, Peng Cui, Zheyang Shen, Xingxuan Zhang, and Tong Zhang. Why stable learning works? a theory of covariate shift generalization, 2021. 1
- [33] Ryota Yoshihashi, Wen Shao, Rei Kawakami, Shaodi You, Makoto Iida, and Takeshi Naemura. Classification-reconstruction learning for open-set recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4016–4025, 2019. 3
- [34] Qiang Zhou, Chaohui Yu, Zhibin Wang, Qi Qian, and Hao Li. Instant-teaching: An end-to-end semi-supervised object detection framework. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4081–4090, 2021. 2