

# Debiasing, calibrating, and improving Semi-supervised Learning performance via simple Ensemble Projector

Khanh-Binh Nguyen  
Sungkyunkwan University  
South Korea  
n.k.binh00@gmail.com

## Abstract

Recent studies on semi-supervised learning (SSL) have achieved great success. Despite their promising performance, current state-of-the-art methods tend toward increasingly complex designs at the cost of introducing more network components and additional training procedures. In this paper, we propose a simple method named Ensemble Projectors Aided for Semi-supervised Learning (EPASS), which focuses mainly on improving the learned embeddings to boost the performance of the existing contrastive joint-training semi-supervised learning frameworks. Unlike standard methods, where the learned embeddings from one projector are stored in memory banks to be used with contrastive learning, EPASS stores the ensemble embeddings from multiple projectors in memory banks. As a result, EPASS improves generalization, strengthens feature representation, and boosts performance. For instance, EPASS improves strong baselines for semi-supervised learning by 39.47%/31.39%/24.70% top-1 error rate, while using only 100k/1%/10% of labeled data for SimMatch, and achieves 40.24%/32.64%/25.90% top-1 error rate for CoMatch on the ImageNet dataset. These improvements are consistent across methods, network architectures, and datasets, proving the general effectiveness of the proposed methods. Code is available at <https://github.com/beandkay/EPASS>.

## 1. Introduction

Deep learning has shown remarkable success in a variety of visual tasks such as image classification [19], speech recognition [1], and natural language processing [34]. This success benefits from the availability of large-scale annotated datasets [20, 21, 27, 31, 32]. Large amounts of annotations are expensive or time-consuming in real-world domains such as medical imaging, banking, and finance. Learning without annotations or with a small number of annotations has become an essential problem in computer

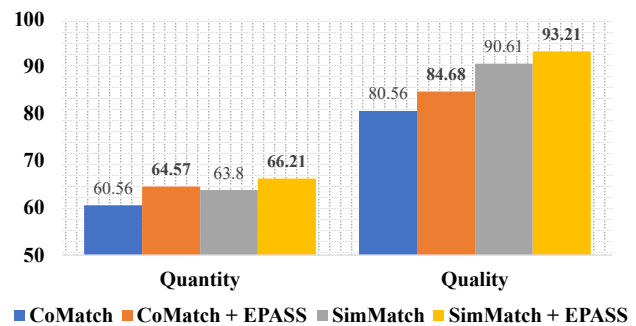


Figure 1. Quantity vs quality of pseudo-labels on ImageNet 10% with and without EPASS.

vision, as demonstrated by [5, 6, 10–12, 17, 18, 23, 24, 26, 35, 36, 42, 46, 51].

Contrastive self-supervised learning (CSL) is based on instance discrimination, which attracts positive samples while repelling negative ones to learn the representation [10, 18, 40]. Inspired by CSL, contrastive joint-training SSL methods such as CoMatch [26] and SimMatch [51] leverage the idea of a memory bank and momentum encoder from MoCo [18] to support representational learning. In the current mainstream contrastive joint-training SSL methods, a multi-layer perceptron (MLP) is added after the encoder to obtain a low-dimensional embedding. Training loss and accuracy evaluation are both performed on this embedding. The previously learned embeddings from a low-dimensional projector are stored in a memory bank. These embeddings are later used in the contrastive learning phase to aid the learning process and improve the exponential moving average (EMA) teacher [36]. Although previous approaches demonstrate their novelty with state-of-the-art benchmarks across many datasets, there are still concerns that need to be considered. For instance, conventional methods such as CoMatch [26] and SimMatch [51] are based

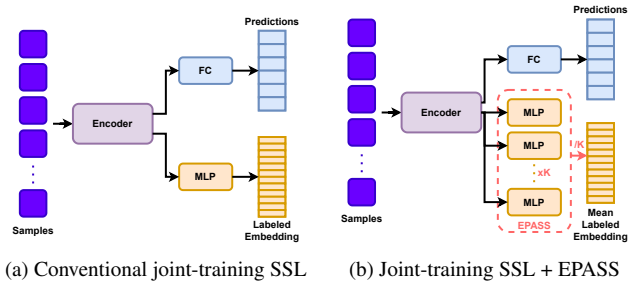


Figure 2. Training phase for contrastive joint-training SSL without/with the proposed EPASS. 2a represents the conventional training phase without EPASS [26, 51]. Unlike 2a, in 2b, instead of using **only one projector** to learn the embeddings, EPASS uses **multiple projectors** to ensemble the embeddings, which is less biased and more generalized.

on the assumption that **the learned embeddings are correct, regardless of confirmation bias**. This assumption is directly adopted from CSL; however, in a joint-training scheme, the easy-to-learn representation could easily dominate the hard-to-learn representation, leading to biased distributions and embeddings. This would become even worse when confirmation bias happens and the embeddings are driven away by the incorrect pseudo-labels. As a result, the embeddings stored in the memory bank are also affected, causing the confirmation bias issue and the erroneous EMA teacher.

The confirmation bias could be seen in Figure 1, where CoMatch only has **80.56%** correct pseudo-labels and SimMatch has **90.61%** correctness for pseudo-labels. When the embedding bias happens at the instance level and the confirmation bias happens at the semantic level, they degrade the performance of the EMA teacher. As a result, the well-learned embeddings at the instance level could be driven away by the confirmation bias at the semantic level during backward propagation, and vice versa.

Method	WRN-28-2	WRN-28-8
Original [6, 7, 35, 38, 43, 47]	1.4 M	23.4 M
Chen <i>et al.</i> [9]	3.7 M ( $\uparrow 2.3$ )	19.9 M (*, $\downarrow 3.5$ )
CoMatch [26]	1.5 M	23.71 M
SimMatch [51]	1.5 M	23.74 M
<b>CoMatch [26] + EPASS (3 projs)</b>	1.54 M ( $\uparrow 0.04$ )	24.30 M ( $\uparrow 0.59$ )
<b>SimMatch [51] + EPASS (3 projs)</b>	1.56 M ( $\uparrow 0.06$ )	24.39 M ( $\uparrow 0.65$ )

Table 1. Comparison with multi-head co-training. ‘\*’ indicates different architecture as Chen *et al.* [9] modified the number of channels of the final block from 512 to 256.

To address these limitations, we propose **Ensemble Projectors Aided for Semi-supervised Learning (EPASS)**, a plug-and-play module to strengthen the EMA teacher as well as to improve the generalization of the learned em-

beddings, as illustrated in Figure 2. Adding a projector helps mitigate the overfitting problem, and the generated features are more distinguishable for classification [26, 51]. Chen *et al.* [13] proves the strengths of ensemble projectors in teacher-student frameworks via knowledge distillation. Therefore, we leverage those strengths with SSL, especially contrastive joint-training frameworks. Although there has been study about ensemble for SSL [9], they only discover it in the classification head, thus resulting in a large number of parameter overheads as shown in Table 1. Unlike [9], we specifically enrich the learned embeddings from the model by employing multiple projectors rather than only one, as it is common in conventional methods. Using ensemble projectors in contrastive learning, where multiple projectors are used instead of a single one, may improve the performance and robustness of the learned representations. By using multiple projectors, the model can learn different feature representations from different perspectives, which can be combined to produce more informative representations of the data. Additionally, using ensemble projectors can help to improve the generalization performance of the model, by reducing the risk of overfitting to the specific characteristics of a single projector.

Using ensemble projectors can also increase the robustness of the model against variations in the data distribution, as the multiple projectors can learn different features that are less sensitive to changes in the data distribution. This can be especially useful in situations where the data distribution is not well-defined or changes over time. Therefore, the embeddings of the model would be the ensemble ones, which are less biased and more robust than conventional methods. Our comprehensive results show that such a simple ensemble design brings a sweet spot between model performance and efficiency.

By incorporating the ensemble projectors in a contrastive-based SSL fashion, the proposed EPASS makes better use of embeddings to aid contrastive learning as well as to improve the classification performance simultaneously. In addition, ensemble multiple projectors introduce a relatively smaller number of parameters compared with ensemble multiple classification heads. Extensive experiments justify the effectiveness of EPASS, which produces a less biased feature space. Specifically, EPASS achieves a state-of-the-art performance with **39.47%/31.39%/24.70%** top-1 error rate, while using only 100k/1%/10% of labeled data for SimMatch; and achieves **40.24%/32.64%/25.90%** top-1 error rate for CoMatch on ImageNet dataset.

The contributions of this paper are summarized as follows:

- We hypothesize that the conventional contrastive joint-training SSL frameworks are sub-optimal since the multi-objective learning could harm the learned em-

beddings when confirmation bias occurs.

- We propose EPASS, a simple plug-and-play module that improves a generalization of the model by imposing the ensemble of multiple projectors, which encourages the model to produce less biased embeddings.
- To the best of our knowledge, this is the first work to enhance the performance of contrastive joint-training SSL methods by considering the embedding bias.
- Extensive experiments on many benchmark datasets demonstrate that EPASS consistently improves the performance of contrastive joint-training methods.

## 2. Related Work

### 2.1. Semi-supervised Learning

Semi-supervised learning is an essential method to leverage a large amount of unlabeled data to enhance the training process. Pseudo-label [24] is the pioneer of nowadays popular methods, including self-training-based or consistency-based SSL approaches. In the pseudo-label-based methods, the model is first trained on a small amount of labeled data. Then, the model is used to make predictions for unlabeled data. The unlabeled data and their corresponding pseudo-labels are then used to train the model simultaneously with labeled data, forming the self-training-based methods [2, 3, 24, 28, 36, 42, 49]. Consistency-based methods [5, 6, 26, 29, 35, 47, 51] use a high threshold to determine the reliable predictions from weakly augmented samples. Then, they will be used as pseudo-labels for strongly augmented examples, and the low-confidence predictions will be discarded. However, those approaches suffer from confirmation bias [2] since they overfit the incorrect pseudo-labels during training. Moreover, methods using the high threshold to filter noisy data only use a small amount of unlabeled data during training, and when the model suffers from confirmation bias, it leads to the *Matthew effect*.

Sohn *et al.* [35] introduces a hybrid method named FixMatch, which combines pseudo-labeling with a consistency regularization method. By using a high threshold to filter out noisy pseudo-labels, FixMatch lets the model learn from only confident predictions, thus improving its performance. FlexMatch [47] introduces a Curriculum Pseudo Labeling (CPL) method based on the Curriculum Learning (CL) [4]. CPL configures a dynamic threshold for each class after each iteration, thus letting the model learn better for either hard-to-learn or easy-to-learn classes.

### 2.2. Contrastive joint-training SSL

Li *et al.* proposes CoMatch [26], which combines two contrastive representations on unlabeled data. However,

CoMatch is extremely sensitive to the hyperparameter setting. Especially during training, CoMatch requires a large memory bank to store the embedded features. Recently, Zheng *et al.* [51] published work that takes semantic similarity and instance similarity into account during training. It shows that forcing consistency at both the semantic level and the instance level can bring an improvement, thus achieving state-of-the-art benchmarks. Along this line of work, [44, 50] also leverage the benefit of Class-aware Contrastive loss to the training process of SSL.

Previous methods might fail to provide the correct embeddings due to confirmation bias. Conventionally, confirmation bias does not exist in CSL; however, it occurs in contrastive joint-training SSL by the use of a threshold. It leads to the degradation of the classifier and the projector, thus providing incorrect predictions and embeddings. Our EPASS exploits the ensemble strategy for multiple projectors, imposing consistency and improving generalization for the learned embeddings, thus enhancing the correctness of model predictions.

## 3. Method

### 3.1. Preliminaries

We first define notations used in the following sections. For semi-supervised image classification problem, let  $\mathcal{X} = \{(x_b, y_b) : b \in (1, \dots, B)\}$  be a batch of  $B$  labeled examples, where  $x_b$  is training examples and  $y_b$  is one-hot labels, and  $\mathcal{U} = \{u_b : b \in (1, \dots, \mu B)\}$  be a batch of  $\mu B$  unlabeled examples where  $\mu$  is a hyper-parameter determining the relative sizes of  $\mathcal{X}$  and  $\mathcal{U}$ . For labeled samples, we apply weak augmentation ( $\mathcal{A}_w$ ) to obtain the weakly augmented samples. Then, an encoder  $f(\cdot)$  and a fully-connected classifier  $h(\cdot)$  are applied to get the distribution over classes as  $p(y | x) = h(f(x))$ . The supervised cross-entropy loss for labeled samples is defined as:

$$\mathcal{L}_s = \frac{1}{B} \sum_{b=1}^B \mathcal{H}(y_b, p_b) \quad (1)$$

where  $\mathcal{H}$  is a standard cross-entropy loss function.

Conventionally, CoMatch and SimMatch apply a weak ( $\mathcal{A}_w$ ) and strong ( $\mathcal{A}_s$ ) augmentation on unlabeled samples, then use the trained encoder and fully-connected classifier to get the predictions as  $p_b^w = p(y | \mathcal{A}_w(u_b))$  and  $p_b^s = p(y | \mathcal{A}_s(u_b))$ . Following CoMatch [26] and SimMatch [51], the predictions that surpassing confidence threshold  $\tau$  would be directly used as pseudo-labels to compute the unsupervised classification loss as:

$$\mathcal{L}_u = \frac{1}{\mu B} \sum_{b=1}^{\mu B} \mathbb{1}(\max(\hat{p}_b^w) \geq \tau) \mathcal{H}(\hat{p}_b^w, p_b^s) \quad (2)$$

where  $\hat{p}_b^w = DA(p_b^w)$  is the pseudo-label for input  $\mathcal{A}_w(u_b)$  and  $DA$  is the distribution alignment strategy [26, 51] to balance the pseudo-labels distribution.

Besides, a non-linear projector head  $g(\cdot)$  is used to map the representation from encoder  $f(\cdot)$  into a low-dimensional embeddings space  $z = g \circ f$ . The embeddings then are used to compute contrastive loss, which we simplify as the cross-entropy between the two normalized graphs:

$$\mathcal{L}_c = \frac{1}{\mu B} \sum_{b=1}^{\mu B} \mathcal{H}(q_b^w, q_b^s) \quad (3)$$

where  $q = \phi(L_2 - norm(z))$  is the result after the transformation  $\phi(\cdot)$  of CoMatch or SimMatch on the  $L_2$  normalized vector. In CoMatch,  $\phi$  is building a pseudo-label graph to guide the representation learning as described in Section 3.2.1. For SimMatch,  $\phi$  is calculating the similarities between  $z^w$  and  $i$ -th instance by using a similarity function  $sim(\cdot)$ , which represents the dot product between  $L_2$  normalized vectors  $sim(u, v) = u^T v / \|u\| \|v\|$ . As a result, we have:

$$q_i^w = \frac{\exp(sim(z_b^w, z_i)/T)}{\sum_{k=1}^K \exp(sim(z_b^w, z^k)/T)}, \quad (4)$$

$$q_i^s = \frac{\exp(sim(z_b^s, z_i)/T)}{\sum_{k=1}^K \exp(sim(z_b^s, z^k)/T)}$$

The momentum embeddings stored in the memory bank and the EMA model are then defined as:

$$z_t \leftarrow m z_{t-1} + (1-m) z_t; \quad \theta_t \leftarrow m \theta_{t-1} + (1-m) \theta_t \quad (5)$$

where  $z$  is the embeddings,  $\theta$  is the model's parameters,  $t$  is the iteration, and  $m$  is the momentum parameter. The overall training objective is:

$$\mathcal{L} = \mathcal{L}_s + \lambda_u \mathcal{L}_u + \lambda_c \mathcal{L}_c \quad (6)$$

### 3.2. EPASS

We propose a simple yet effective method to boost the performance of the conventional contrastive-based SSL that maximizes the correctness of the embeddings from different projections by using the ensemble technique.

Unlike conventional methods such as CoMatch and SimMatch, which assume that the learned embeddings from one projector are absolutely correct, we propose using the ensemble embeddings from multiple projectors to mitigate the bias. While there could be diverse options to combine multiple embeddings (e.g., concatenation, summation), we empirically found that simply averaging the selected embeddings works reasonably well and is computationally efficient. As each projector is randomly initialized, it provides a different view of inputs, which benefits the generalization of the model. This intuition is similar to that of multi-view learning. However, since we generate views with multiple

projectors instead of creating multiple augmented samples, we introduce far less overhead to the pipeline. The ensemble of multiple projectors helps mitigate the bias in the early stages of training. In the joint-training scheme, the correct learned embeddings help improve the performance of the classification head and vice versa, thus reducing the confirmation bias effect. The embeddings stored in the memory bank by Equation 6 therefore are updated as:

$$z_t \leftarrow m z_{t-1} + (1-m) \bar{z}_t; \quad \bar{z}_t = \frac{\sum_{p=1}^P z_{t,p}}{P} \quad (7)$$

where  $P$  is the number of projectors.

#### 3.2.1 Application

**SimMatch:** Using our ensemble embeddings, we re-define instance similarity in SimMatch [51] and rewrite the Equation 4 as:

$$\bar{q}_i^w = \frac{\exp(sim(\bar{z}_b^w, \bar{z}_i)/T)}{\sum_{k=1}^K \exp(sim(\bar{z}_b^w, \bar{z}_k)/T)} \quad (8)$$

where  $T$  is the temperature parameter controlling the sharpness of the distribution,  $K$  is the number of weakly augmented embeddings, and  $i$  represents the  $i$ -th instance. Similarly, we can compute  $\bar{q}_i^s$  by calculating the similarities between the strongly augmented embeddings  $\bar{z}^s$  and  $\bar{z}_i$ .

$$\bar{q}_i^s = \frac{\exp(sim(\bar{z}_b^s, \bar{z}_i)/T)}{\sum_{k=1}^K \exp(sim(\bar{z}_b^s, \bar{z}_k)/T)} \quad (9)$$

The Equation 3 then is rewritten as:

$$\mathcal{L}_c = \frac{1}{\mu B} \sum_{b=1}^{\mu B} \mathcal{H}(\bar{q}_b^w, \bar{q}_b^s) \quad (10)$$

**CoMatch:** In CoMatch, the embeddings are used to construct a pseudo-label graph that defines the similarity of samples in the label space. Specifically, the instance similarity is also calculated as Equation 8 for weakly augmented samples. Then, a similarity matrix  $W^q$  is constructed as:

$$W_{bj}^q = \begin{cases} 1 & \text{if } b = j \\ \bar{q}_b \cdot \bar{q}_j & \text{if } b \neq j \text{ and } \bar{q}_b \cdot \bar{q}_j \geq \tau_c \\ 0 & \text{otherwise} \end{cases} \quad (11)$$

where  $\tau_c$  indicates the similarity threshold. Also, an embedding graph  $W^z$  is derived as:

$$W_{bj}^z = \begin{cases} \exp(\bar{z}_b \cdot \bar{z}_j'/t) & \text{if } b = j \\ \exp(\bar{z}_b \cdot \bar{z}_j/t) & \text{if } b \neq j \end{cases} \quad (12)$$



where  $z_b = g \circ f(\mathcal{A}_s(u_b))$  and  $z'_b = g \circ f(\mathcal{A}'_s(u_b))$ . The Equation 3 then is rewritten as:

$$\mathcal{L}_c = \frac{1}{\mu B} \sum_{b=1}^{\mu B} \mathcal{H}(\hat{W}_b^q, \hat{W}_b^z) \quad (13)$$

where  $\hat{W}_{bj} = \frac{W_{bj}}{\sum_j W_{bj}}$ ,  $\mathcal{H}(\hat{W}_b^q, \hat{W}_b^z)$  can be decomposed into:

$$\begin{aligned} \mathcal{H}(\hat{W}_b^q, \hat{W}_b^z) = & -\hat{W}_{bb}^q \log\left(\frac{\exp(\bar{z}_b \cdot \bar{z}'_b/T)}{\sum_{j=1}^{\mu B} \hat{W}_{bj}^z}\right) \\ & - \sum_{j=1, j \neq b}^{\mu B} \hat{W}_{bj}^q \log\left(\frac{\exp(\bar{z}_b \cdot \bar{z}'_j/T)}{\sum_{j=1}^{\mu B} \hat{W}_{bj}^z}\right) \end{aligned}$$

## 4. Experiments

### 4.1. Implementation Details

We evaluate EPASS on common benchmarks: CIFAR-10/100 [22], SVHN [30], STL-10 [14], and ImageNet [15]. We conduct experiments with varying amounts of labeled data, using previous work [5, 6, 26, 29, 35, 41, 43, 47, 51].

For a fair comparison, we train and evaluate all methods using the unified code base USB [37] with the same backbones and hyperparameters. We use Wide ResNet-28-2 [45] for CIFAR-10, Wide ResNet-28-8 for CIFAR-100, Wide ResNet-37-2 [52] for STL-10, and ResNet-50 [19] for ImageNet. We use SGD with a momentum of 0.9 as an optimizer. The initial learning rate is 0.03 with a cosine learning rate decay schedule of  $\eta = \eta_0 \cos\left(\frac{\tau \pi k}{16K}\right)$ , where  $\eta_0$  is the initial learning rate and  $k(K)$  is the total training step. We set  $K = 2^{20}$  for all datasets. During the testing phase, we employ an exponential moving average with a momentum of 0.999 on the training model to perform inference for all algorithms. The batch size for labeled data is 64, with the exception of ImageNet, which has a batch size of 128. The same weight decay value, pre-defined threshold  $\tau$ , unlabeled batch ratio  $\mu$  and loss weights are used for Pseudo-Label [24],  $\Pi$  model [33], Mean Teacher [36], VAT [29], MixMatch [6], ReMixMatch [5], UDA [41], FixMatch [35], FlexMatch [47], CoMatch [26], SimMatch [51], AdaMatch [7], and FreeMatch [38].

We use the same parameters as in [37, 43] for Dash method. For other methods, we follow the original settings reported in their studies. In Appendix ??, you can find a comprehensive description of the hyperparameters used. To ensure the robustness, we train each algorithm three times with different random seeds. Consistent with [47], we report the lowest error rates achieved among all checkpoints.

### 4.2. CIFAR-10/100, STL-10, SVHN

The best error rate of each method is evaluated by averaging the results obtained from three runs with different random seeds. The results are presented in Table 2,

where we report the classification error rates on the CIFAR-10/100, STL-10, and SVHN datasets. EPASS is shown to improve the performance of SimMatch and CoMatch significantly on all datasets. For instance, even though EPASS does not achieve state-of-the-art results in CIFAR-10/100, it still boosts the performance of conventional SimMatch and CoMatch. It should be noted that CIFAR-10/100 are small datasets where prior works have already achieved high performance, leaving little room for improvement. Moreover, ReMixMatch performs well on CIFAR-100 (2500) and CIFAR-100 (10000) due to the mixup technique and the self-supervised learning part. Additionally, on the SVHN and STL-10 datasets, SimMatch and CoMatch with EPASS surpass all prior state-of-the-art results by a significant margin, achieving a new state-of-the-art performance. These results demonstrate the effectiveness of EPASS in mitigating bias, particularly on imbalanced datasets such as SVHN and STL-10, where overfitting is a common issue.

### 4.3. ImageNet

EPASS is evaluated on the ImageNet ILSVRC-2012 dataset to demonstrate its effectiveness on large-scale datasets. In order to assess the performance of EPASS, we sample 100k/1%/10% of labeled images in a class-balanced manner, where the number of samples per class is 10, 13, or 128, respectively. The remaining images in each class are left unlabeled. Our experiments are conducted using a fixed random seed, and the results are found to be robust across different runs.

As presented in Table 3, EPASS outperforms the state-of-the-art methods, achieving a top-1 error rate of 39.47%/31.39%/24.70% for SimMatch and a top-1 error rate of 40.24%/32.64%/25.90% for CoMatch, respectively. The results clearly demonstrate the effectiveness of EPASS in improving the performance of SSL methods on large-scale datasets like ImageNet.

## 5. Ablation Study

### 5.1. ImageNet convergence speed

The convergence speed of the proposed EPASS is extremely noticeable through our extensive experiments. When training on ImageNet, we observe that EPASS achieves over 50% of accuracy in the first few iterations, indicating that the model is able to quickly learn meaningful representations from the unlabeled data. This is likely due to the fact that EPASS encourages the model to focus on the most informative and diverse instances during training, which helps the model learn more quickly and effectively. Additionally, we find that the accuracy of SimMatch and CoMatch with EPASS is consistently increasing with iterations, outperforming conventional SimMatch and CoMatch with the same training epochs. This suggests that the use

Dataset	CIFAR-10			CIFAR-100			SVHN			STL-10		
	Label Amount	40	250	4000	400	2500	10000	40	250	1000	40	250
UDA [41]	10.20±5.05	5.40±0.28	4.27±0.05	51.96±1.27	29.47±0.52	23.59±0.32	2.39±0.53	1.99±0.02	1.91±0.05	53.69±4.38	28.96±1.02	7.25±0.50
MixMatch [6]	38.84±8.36	20.96±2.45	10.25±0.01	80.58±3.38	47.88±0.21	33.22±0.06	26.61±13.10	4.48±0.35	5.01±0.12	52.32±0.91	36.34±0.84	25.01±0.43
ReMixMatch [5]	8.13±0.58	6.34±0.22	4.65±0.09	41.60±1.48	<u>25.72±0.07</u>	<b>20.04±0.13</b>	16.43±13.77	5.65±0.35	5.36±0.58	27.87±3.85	11.14±0.52	6.44±0.15
FixMatch [35]	12.66±4.49	4.95±0.10	4.26±0.01	45.38±2.07	27.71±0.42	22.06±0.10	3.37±1.01	1.97±0.01	2.02±0.03	38.19±4.76	8.64±0.84	5.82±0.06
FlexMatch [47]	5.29±0.29	4.97±0.07	4.24±0.06	40.73±1.44	26.17±0.18	21.75±0.15	5.42±2.83	8.74±3.32	7.90±0.30	29.12±5.04	9.85±1.35	6.08±0.34
Dash [43]	9.29±3.28	5.16±0.28	4.36±0.10	47.49±1.05	27.47±0.38	21.89±0.16	5.26±2.02	2.01±0.01	2.08±0.09	42.00±4.94	10.50±1.37	6.30±0.49
CoMatch [26]	6.51±1.18	5.35±0.14	4.27±0.12	53.41±2.36	29.78±0.11	22.11±0.22	8.20±5.32	2.16±0.04	2.01±0.04	13.74±4.20	<u>7.63±0.94</u>	5.71±0.08
SimMatch [51]	5.38±0.01	5.36±0.08	4.41±0.07	39.32±0.72	26.21±0.37	21.50±0.11	7.60±2.11	2.48±0.61	2.05±0.05	16.98±4.24	8.27±0.40	5.74±0.31
AdaMatch [7]	<u>5.09±0.21</u>	5.13±0.05	4.36±0.05	<u>38.08±1.35</u>	26.66±0.33	21.99±0.15	6.14±5.35	2.13±0.04	2.02±0.05	19.95±5.17	8.59±0.43	6.01±0.02
FreeMatch [38]	<b>4.90±0.12</b>	<b>4.88±0.09</b>	<b>4.16±0.06</b>	39.52±0.01	26.22±0.08	21.81±0.17	10.43±0.82	8.23±3.22	7.56±0.25	28.50±5.41	9.29±1.24	5.81±0.32
SoftMatch [8]	5.11±0.14	4.96±0.09	4.27±0.05	<b>37.60±0.24</b>	26.39±0.38	21.86±0.16	2.46±0.24	2.15±0.07	2.09±0.06	22.23±3.82	9.18±0.68	5.79±0.15
[26] + EPASS	5.55±0.21	5.31±0.13	4.23±0.05	50.73±0.33	29.51±0.16	22.16±0.12	2.98±0.02	<b>1.93±0.05</b>	<b>1.85±0.04</b>	<b>9.15±3.25</b>	<b>6.27±0.03</b>	<b>5.40±0.12</b>
[51] + EPASS	5.31±0.10	5.08±0.05	4.37±0.03	38.88±0.24	<b>25.68±0.33</b>	<u>21.32±0.14</u>	<b>2.31±0.04</b>	2.04±0.02	2.02±0.02	15.71±2.48	8.08±0.26	5.58±0.04
Fully-Supervised		4.62±0.05			19.30±0.09			2.13±0.02			None	

Table 2. Error rate on CIFAR-10/100, SVHN, and STL-10 datasets on 3 different folds. **Bold** indicates best result and Underline indicates the second best result.

Method	Top-1		Top-5		Top-1		Top-5	
	100k		1%		10%			
FixMatch [35]	43.66	21.80	-	-	28.50	10.90		
FlexMatch [47]	41.85	19.48	-	-	-	-		
CoMatch [26]	42.17	19.64	34.00	13.60	26.30	8.60		
SimMatch [51]	41.15	19.23	32.80	12.90	25.60	8.40		
FreeMatch [38]	40.57	18.77	-	-	-	-		
SoftMatch [8]	40.52	-	-	-	-	-		
[26] + EPASS	<u>40.24</u>	<u>18.40</u>	<u>32.64</u>	<u>12.71</u>	<u>25.90</u>	<u>8.48</u>		
[51] + EPASS	<b>39.47</b>	<b>18.24</b>	<b>31.39</b>	<b>12.41</b>	<b>24.70</b>	<b>7.44</b>		

Table 3. ImageNet error rate results. **Bold** indicates best result and Underline indicates the second best result.

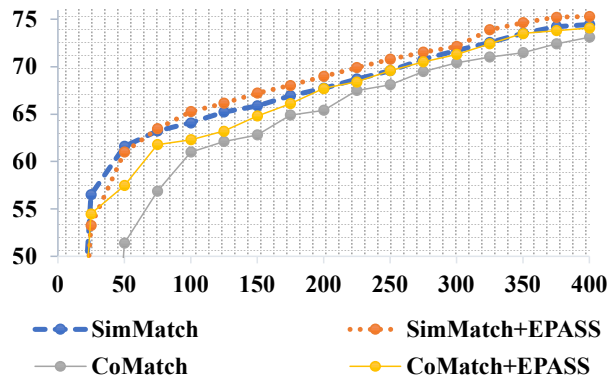


Figure 3. Convergence analysis of SimMatch with and without EPASS.

of EPASS enables the model to continue learning and improving over time, rather than plateauing or becoming overfitted. Overall, these results demonstrate the effectiveness of EPASS in improving the convergence speed and performance of SSL methods.

## 5.2. Calibration of SSL

Chen *et al.* [9] propose a method for addressing confirmation bias from the calibration perspective. To evaluate the effectiveness of EPASS in this regard, we measure the calibration of CoMatch and SimMatch on the ImageNet dataset with and without EPASS, using 10% labeled data<sup>1</sup>. Several common calibration indicators, including Expected Calibration Error (ECE), confidence histogram, and reliability diagram, are utilized in this study.

Figure 4 illustrates that when EPASS is used with 10% of labels, the ECE value of the model decreases. Moreover, under the 1% label scheme, CoMatch and SimMatch without EPASS are significantly overconfident and overfitted due to confirmation bias. However, when EPASS is employed, it helps to reduce the ECE by a large margin and also mitigate the overconfidence of the model. Notably, models with EPASS have average accuracy and average confidence that are approximately equal, whereas the average confidence of models without EPASS is usually higher than the accuracy.

It is worth mentioning that since CoMatch does not impose the interaction between semantic and instance similarity like SimMatch, the effect of introducing EPASS to CoMatch for calibration is not as significant as that for SimMatch. Additionally, the model with EPASS becomes underfit and may benefit from additional training.

## 5.3. Number of projectors

This section studies the effectiveness of the proposed projectors ensemble method and how different ensemble strategies affect performance. In this experiment, we study the effect of different numbers of projectors on performance. The top-1 classification accuracy of the proposed EPASS with different numbers of projectors is shown in Ta-

<sup>1</sup><https://github.com/hollance/reliability-diagrams>

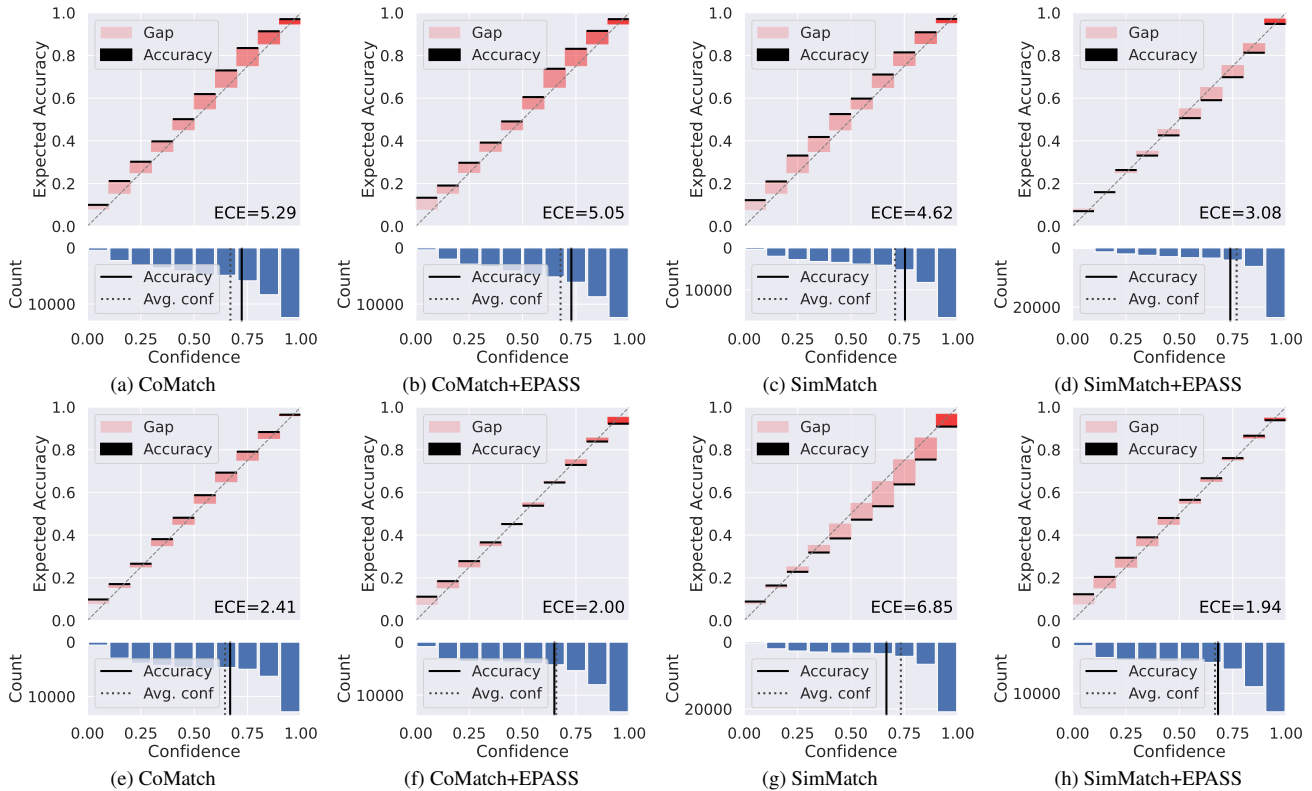


Figure 4. Reliability diagrams (top) and confidence histograms (bottom) for ImageNet dataset. The first row and second row are conducted with 10% and 1% of labels, respectively.

ble 4.

Method	# projectors			
	1	2	3	4
<b>CoMatch [26] + EPASS</b>	73.6	73.8	<b>74.1</b>	73.9
<b>SimMatch [51] + EPASS</b>	74.4	74.8	<b>75.3</b>	75.2

Table 4. Top-1 accuracy (%) on ImageNet 10% using different numbers of projectors.

In Table 5, we record the results of different ensemble strategies for EPASS. Overall, averaging the embeddings results in better performance than concatenation and summation.

Method	Ensemble strategy		
	Concatenate	Sum	Mean
<b>CoMatch [26] + EPASS</b>	74.0	73.9	<b>74.1</b>
<b>SimMatch [51] + EPASS</b>	75.1	74.8	<b>75.3</b>

Table 5. Top-1 accuracy (%) on ImageNet 10% using different ensemble strategies.

## 5.4. Imbalanced SSL

Dataset	CIFAR-10-LT		CIFAR-100-LT	
	$\lambda = 50$	$\lambda = 150$	$\lambda = 20$	$\lambda = 100$
FixMatch [35]	18.5±0.48	31.2±1.08	49.1±0.62	62.5±0.36
FlexMatch [47]	17.8±0.24	29.5±0.47	48.9±0.71	62.7±0.08
FreeMatch [38]	17.7±0.33	28.8±0.64	48.4±0.91	62.5±0.23
SoftMatch [8]	16.6±0.29	<b>27.4±0.46</b>	48.1±0.55	61.1±0.81
CoMatch [26]	<u>16.3±0.24</u>	30.1±0.31	46.2±0.41	60.0±0.21
SimMatch [51]	20.3±0.31	28.7±0.48	<u>45.4±0.55</u>	60.1±0.21
<b>[26] + EPASS</b>	<b>16.1±0.22</b>	29.6±0.41	45.9±0.45	<b>59.8±0.01</b>
<b>[51] + EPASS</b>	18.2±0.34	<u>28.4±0.43</u>	<b>45.2±0.51</b>	<b>59.6±0.11</b>
FixMatch + ABC [25]	14.0±0.22	22.3±1.08	46.6±0.69	<b>58.3±0.41</b>
FlexMatch + ABC [25]	14.2±0.34	23.1±0.70	46.2±0.47	58.9±0.51
FreeMatch + ABC [25]	<u>13.9±0.03</u>	22.3±0.26	45.6±0.76	58.9±0.55
[26] + ABC [25]	14.1±0.21	23.1±0.32	43.0±0.52	59.0±0.31
[51] + ABC [25]	14.5±0.25	<u>20.5±0.21</u>	43.3±0.44	58.9±0.50
<b>[26] + EPASS + ABC [25]</b>	14.0±0.19	22.4±0.41	<u>42.7±0.55</u>	<b>58.5±0.41</b>
<b>[51] + EPASS + ABC [25]</b>	<b>13.3±0.09</b>	<b>20.2±0.26</b>	<b>42.7±0.41</b>	58.8±0.37

Table 6. Error rates (%) of imbalanced SSL using 3 different random seeds. **Bold** indicates best result and Underline indicates the second best result.

To provide additional evidence of the effectiveness of EPASS, we assess its performance in the imbalanced semi-supervised learning scenario [16, 25, 39], where both the labeled and unlabeled data are imbalanced. Our experi-

Dataset	CIFAR-100		STL-10		Euro-SAT		TissueMNIST		Semi-Aves
Label Amount	200	400	20	40	20	40	100	500	3959
UDA [41]	30.75±1.03	19.94±0.32	39.22±2.87	23.59±2.97	11.15±1.20	5.99±0.75	55.88±3.26	51.42±2.05	32.55±0.26
MixMatch [6]	37.43±0.58	26.17±0.24	48.98±1.41	25.56±3.00	29.86±2.89	16.39±3.17	<b>55.73±2.29</b>	<b>49.08±1.06</b>	37.22±0.15
ReMixMatch [5]	<b>20.85±1.42</b>	<u>16.80±0.59</u>	30.61±3.47	18.33±1.98	4.53±1.60	4.10±0.37	59.29±5.16	52.92±3.93	<b>30.40±0.33</b>
FixMatch [35]	30.45±0.65	19.48±0.93	42.06±3.94	24.05±1.79	12.48±2.57	6.41±1.64	55.95±4.06	50.93±1.23	31.74±0.33
FlexMatch [47]	27.08±0.90	17.67±0.66	37.58±2.97	23.40±1.50	7.07±2.32	5.58±0.57	57.23±2.50	52.06±1.78	33.09±0.16
Dash [43]	30.19±1.34	18.90±0.420	43.34±1.46	25.90±0.35	9.44±0.75	7.00±1.39	57.00±2.81	50.93±1.54	32.56±0.39
CoMatch [26]	35.68±0.54	26.10±0.09	<u>29.70±1.17</u>	21.46±1.34	5.25±0.49	4.89±0.86	57.15±3.46	51.83±0.71	41.39±0.16
SimMatch [51]	23.26±1.25	16.82±0.40	34.12±1.63	22.97±2.04	6.88±1.77	5.86±1.07	57.91±4.60	51.14±1.83	34.14±0.30
AdaMatch [7]	<u>21.27±1.04</u>	17.01±0.55	36.25±1.89	23.30±0.73	5.70±0.37	4.92±0.87	57.87±4.47	52.28±0.79	<u>31.54±0.10</u>
<b>[26] + EPASS</b>	35.10±0.55	25.53±0.50	<b>29.56±2.50</b>	<b>21.14±0.31</b>	<b>3.41±0.24</b>	<b>2.91±0.41</b>	56.88±4.93	51.06±1.09	41.19±0.43
<b>[51] + EPASS</b>	22.52±0.83	<b>16.78±0.59</b>	30.03±0.71	22.65±1.94	5.35±0.81	<u>3.81±0.37</u>	57.22±5.97	<u>50.40±1.44</u>	33.83±0.04
Fully-Supervised	8.90±0.12		-		0.85±0.06		33.91±0.03		-

Table 7. Error rate on CIFAR-10/100, SVHN, and STL-10 datasets on 3 different folds. **Bold** indicates best result and Underline indicates second best result.

ments are conducted on CIFAR-10-LT and CIFAR-100-LT, using varying degrees of class imbalance ratios. For the CIFAR datasets, the imbalance ratio is defined as follows:  $\lambda = N_{max}/N_{min}$  where  $N_{max}$  is the number of samples on the head (frequent) class and  $N_{min}$  the tail (rare). Note that the number of samples for class  $k$  is computed as  $N_k = N_{max}\lambda^{-\frac{k-1}{C-1}}$ , where  $C$  is the number of classes. Following [16,25], we set  $N_{max} = 1500$  for CIFAR-10 and  $N_{max} = 150$  for CIFAR-100, and the number of unlabeled data is twice as many for each class. We use a WRN-28-2 [45] as the backbone. We use Adam as the optimizer. The initial learning rate is 0.002 with a cosine learning rate decay schedule as  $\eta = \eta_0 \cos\left(\frac{7\pi k}{16K}\right)$ , where  $\eta_0$  is the initial learning rate,  $k(K)$  is the current (total) training step and we set  $K = 2.5 \times 10^5$  for all datasets. The batch size of labeled and unlabeled data is 64 and 128, respectively. Weight decay is set as  $4e^{-5}$ . Each experiment is run on three different data splits, and we report the average of the best error rates.

The results are summarized in Table 6. Compared with other standard SSL methods, EPASS achieves the best performance across all settings. Especially on CIFAR-100 at an imbalance ratio 100, SimMatch with EPASS outperforms the second-best by 0.6%. Moreover, when plugged in the other imbalanced SSL method [25], EPASS still attains the best performance in most of the settings.

## 5.5. Result using USB

In this section, we evaluate the effectiveness of EPASS within the context of the USB [37] framework, adhering strictly to the USB settings for CV tasks that utilize pre-trained Vision Transformers (ViT). For a detailed overview of hyperparameters used in these experiments, please refer to Appendix ??.

As Table 7 indicates, EPASS improves the performance of SimMatch and CoMatch on all datasets, albeit marginally. These experiments utilize pre-trained ViT models, which provide a strong representation initialization on unlabeled data, leaving little room for improvement when applying SSL methods with this kind of model. Notably, ReMixMatch [5] achieves the highest performance among all SSL algorithms due to its usage of mixup [48], Distribution Alignment, and rotation self-supervised loss. However, on CIFAR-100, STL-10, Euro-SAT, and TissueMNIST datasets, EPASS outperforms ReMixMatch.

## 6. Future works

While our experiments have shown the effectiveness of EPASS in mitigating bias in computer vision tasks, it is unclear whether EPASS can be generalized to other domains, such as natural language processing or speech recognition. Therefore, future research could investigate the applicability of EPASS to these domains and explore how it can be adapted to different types of problems beyond classification tasks such as object detection or segmentation, which have different characteristics than classification problems and require more complex models.

## 7. Conclusion

Our proposed method, EPASS, enhances the performance and reliability of conventional contrastive joint-training SSL methods. EPASS achieves this by mitigating confirmation bias and embedding bias, which leads to simultaneous performance improvement and reduced overconfidence. EPASS outperforms strong competitors across a variety of SSL benchmarks, especially in the large-scale dataset setting. Additionally, EPASS introduces minimal overhead to the overall pipeline.



## References

- [1] Dario Amodei, Sundaram Ananthanarayanan, Rishita Anubhai, Jingliang Bai, Eric Battenberg, Carl Case, Jared Casper, Bryan Catanzaro, Qiang Cheng, Guoliang Chen, et al. Deep speech 2: End-to-end speech recognition in english and mandarin. In *International conference on machine learning*, pages 173–182. PMLR, 2016. [1](#)
- [2] 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, 2020. [3](#)
- [3] Philip Bachman, Ouais Alsharif, and Doina Precup. Learning with pseudo-ensembles. In *NIPS*, 2014. [3](#)
- [4] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *ICML ’09*, 2009. [3](#)
- [5] David Berthelot, Nicholas Carlini, Ekin Dogus Cubuk, Alexey Kurakin, Kihyuk Sohn, Han Zhang, and Colin Raffel. Remixmatch: Semi-supervised learning with distribution matching and augmentation anchoring. In *ICLR*, 2020. [1](#), [3](#), [5](#), [6](#), [8](#)
- [6] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A Raffel. Mixmatch: A holistic approach to semi-supervised learning. *Advances in neural information processing systems*, 32, 2019. [1](#), [2](#), [3](#), [5](#), [6](#), [8](#)
- [7] David Berthelot, Rebecca Roelofs, Kihyuk Sohn, Nicholas Carlini, and Alex Kurakin. Adamatch: A unified approach to semi-supervised learning and domain adaptation. *arXiv preprint arXiv:2106.04732*, 2021. [2](#), [5](#), [6](#), [8](#)
- [8] Hao Chen, Ran Tao, Yue Fan, Yidong Wang, Jindong Wang, Bernt Schiele, Xing Xie, Bhiksha Raj, and Marios Savvides. Softmatch: Addressing the quantity-quality trade-off in semi-supervised learning. 2023. [6](#), [7](#)
- [9] Mingcai Chen, Yuntao Du, Yi Zhang, Shuwei Qian, and Chongjun Wang. Semi-supervised learning with multi-head co-training. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 6278–6286, 2022. [2](#), [6](#)
- [10] 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. [1](#)
- [11] Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey E Hinton. Big self-supervised models are strong semi-supervised learners. *Advances in neural information processing systems*, 33:22243–22255, 2020. [1](#)
- [12] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020. [1](#)
- [13] Yudong Chen, Sen Wang, Jiajun Liu, Xuwei Xu, Frank de Hoog, and Zi Huang. Improved feature distillation via projector ensemble. In *Thirty-Sixth Conference on Neural Information Processing Systems*, 2022. [2](#)
- [14] 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. [5](#)
- [15] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, K. Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. [5](#)
- [16] Yue Fan, Dengxin Dai, Anna Kukleva, and Bernt Schiele. CossI: Co-learning of representation and classifier for imbalanced semi-supervised learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14574–14584, 2022. [7](#), [8](#)
- [17] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent-a new approach to self-supervised learning. *Advances in neural information processing systems*, 33:21271–21284, 2020. [1](#)
- [18] 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. [1](#)
- [19] 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*, pages 770–778, 2016. [1](#), [5](#)
- [20] Joel Hestness, Sharan Narang, Newsha Ardalani, Gregory Diamos, Heewoo Jun, Hassan Kianinejad, Md. Mostofa Ali Patwary, Yang Yang, and Yanqi Zhou. Deep learning scaling is predictable, empirically, 2017. [1](#)
- [21] Rafal Jozefowicz, Oriol Vinyals, Mike Schuster, Noam Shazeer, and Yonghui Wu. Exploring the limits of language modeling, 2016. [1](#)
- [22] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. [5](#)
- [23] Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*, 2017. [1](#)
- [24] Dong-Hyun Lee et al. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, volume 3, page 896, 2013. [1](#), [3](#), [5](#)
- [25] Hyuck Lee, Seungjae Shin, and Heeyoung Kim. Abc: Auxiliary balanced classifier for class-imbalanced semi-supervised learning. *Advances in Neural Information Processing Systems*, 34:7082–7094, 2021. [7](#), [8](#)
- [26] 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](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [27] Dhruv Mahajan, Ross Girshick, Vignesh Ramanathan, Kaiming He, Manohar Paluri, Yixuan Li, Ashwin Bharambe, and Laurens van der Maaten. Exploring the limits of weakly supervised pretraining, 2018. [1](#)
- [28] Geoffrey J McLachlan. Iterative reclassification procedure for constructing an asymptotically optimal rule of allocation

- in discriminant analysis. *Journal of the American Statistical Association*, 70(350):365–369, 1975. [3](#)
- [29] Takeru Miyato, Shin ichi Maeda, Masanori Koyama, and Shin Ishii. Virtual adversarial training: A regularization method for supervised and semi-supervised learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41:1979–1993, 2019. [3](#), [5](#)
- [30] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bis-sacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. 2011. [5](#)
- [31] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsuper-vised multitask learners. 2019. [1](#)
- [32] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer, 2020. [1](#)
- [33] Antti Rasmus, Harri Valpola, Mikko Honkala, Mathias Berglund, and Tapani Raiko. Semi-supervised learning with ladder networks. *arXiv preprint arXiv:1507.02672*, 2015. [5](#)
- [34] Richard Socher, Yoshua Bengio, and Christopher D Man-ning. Deep learning for nlp (without magic). In *Tutorial Abstracts of ACL 2012*, pages 5–5. 2012. [1](#)
- [35] 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. *Advances in neural information processing systems*, 33:596–608, 2020. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#), [8](#)
- [36] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *NIPS*, 2017. [1](#), [3](#), [5](#)
- [37] Yidong Wang, Hao Chen, Yue Fan, Wang Sun, Ran Tao, Wenxin Hou, Renjie Wang, Linyi Yang, Zhi Zhou, Lan-Zhe Guo, Heli Qi, Zhen Wu, Yu-Feng Li, Satoshi Nakamura, Wei Ye, Marios Savvides, Bhiksha Raj, Takahiro Shinozaki, Bernt Schiele, Jindong Wang, Xing Xie, and Yue Zhang. Usb: A unified semi-supervised learning benchmark for clas-sification. In *Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2022. [5](#), [8](#)
- [38] Yidong Wang, Hao Chen, Qiang Heng, Wenxin Hou, Yue Fan, , Zhen Wu, Jindong Wang, Marios Savvides, Takahiro Shinozaki, Bhiksha Raj, Bernt Schiele, and Xing Xie. Freematch: Self-adaptive thresholding for semi-supervised learning. 2023. [2](#), [5](#), [6](#), [7](#)
- [39] Chen Wei, Kihyuk Sohn, Clayton Mellina, Alan Yuille, and Fan Yang. Crest: A class-rebalancing self-training frame-work for imbalanced semi-supervised learning. In *Proceed-ings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10857–10866, 2021. [7](#)
- [40] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3733–3742, 2018. [1](#)
- [41] Qizhe Xie, Zihang Dai, Eduard Hovy, Thang Luong, and Quoc Le. Unsupervised data augmentation for consistency training. *Advances in Neural Information Processing Sys-tems*, 33:6256–6268, 2020. [5](#), [6](#), [8](#)
- [42] Qizhe Xie, Minh-Thang Luong, Eduard H. Hovy, and Quoc V. Le. Self-training with noisy student improves ima-genet classification. In *2020 IEEE/CVF Conference on Com-puter Vision and Pattern Recognition, CVPR 2020*, pages 10684–10695, 2020. [1](#), [3](#)
- [43] Yi Xu, Lei Shang, Jinxing Ye, Qi Qian, Yu-Feng Li, Baigui Sun, Hao Li, and Rong Jin. Dash: Semi-supervised learning with dynamic thresholding. In *International Conference on Machine Learning*, pages 11525–11536. PMLR, 2021. [2](#), [5](#), [6](#), [8](#)
- [44] Fan Yang, Kai Wu, Shuyi Zhang, Guannan Jiang, Yong Liu, Feng Zheng, Wei Zhang, Chengjie Wang, and Long Zeng. Class-aware contrastive semi-supervised learning. In *Pro-ceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14421–14430, June 2022. [3](#)
- [45] Sergey Zagoruyko and Nikos Komodakis. Wide residual net-works. *ArXiv*, abs/1605.07146, 2016. [5](#), [8](#)
- [46] Xiaohua Zhai, Avital Oliver, Alexander Kolesnikov, and Lu-cas Beyer. S4l: Self-supervised semi-supervised learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1476–1485, 2019. [1](#)
- [47] Bowen Zhang, Yidong Wang, Wenxin Hou, Hao Wu, Jin-dong Wang, Manabu Okumura, and Takahiro Shinozaki. Flexmatch: Boosting semi-supervised learning with curricu-lum pseudo labeling. *Advances in Neural Information Pro-cessing Systems*, 34:18408–18419, 2021. [2](#), [3](#), [5](#), [6](#), [7](#), [8](#)
- [48] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimiza-tion. *arXiv preprint arXiv:1710.09412*, 2017. [8](#)
- [49] Linfeng Zhang, Jiebo Song, Anni Gao, Jingwei Chen, Chen-glong Bao, and Kaisheng Ma. Be your own teacher: Im-prove the performance of convolutional neural networks via self distillation, 2019. [3](#)
- [50] Zhen Zhao, Luping Zhou, Lei Wang, Yinghuan Shi, and Yang Gao. Lssl: Label-guided self-training for semi-supervised learning. 2022. [3](#)
- [51] 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*, pages 14471–14481, 2022. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [52] Tianyi Zhou, Shengjie Wang, and Jeff Bilmes. Time-consistent self-supervision for semi-supervised learning. In *International Conference on Machine Learning*, pages 11523–11533. PMLR, 2020. [5](#)