

# Scaling Up Semi-supervised Learning with Unconstrained Unlabelled Data

Shuvendu Roy, Ali Etemad

Queen’s University, Canada  
{shuvendu.roy, ali.etemad}@queensu.ca

## Abstract

We propose UnMixMatch, a semi-supervised learning framework which can learn effective representations from unconstrained unlabelled data in order to scale up performance. Most existing semi-supervised methods rely on the assumption that labelled and unlabelled samples are drawn from the same distribution, which limits the potential for improvement through the use of free-living unlabelled data. Consequently, the generalizability and scalability of semi-supervised learning are often hindered by this assumption. Our method aims to overcome these constraints and effectively utilize unconstrained unlabelled data in semi-supervised learning. UnMixMatch consists of three main components: a supervised learner with hard augmentations that provides strong regularization, a contrastive consistency regularizer to learn underlying representations from the unlabelled data, and a self-supervised loss to enhance the representations that are learnt from the unlabelled data. We perform extensive experiments on 4 commonly used datasets and demonstrate superior performance over existing semi-supervised methods with a performance boost of 4.79%. Extensive ablation and sensitivity studies show the effectiveness and impact of each of the proposed components of our method. The code for our work is publicly available.

## 1 Introduction

Semi-supervised learning (SSL) uses large amounts of *unlabelled* data along with small amounts of *labelled* data to reduce the reliance on fully-labelled datasets. Most existing semi-supervised methods rely on the assumption that the labelled and unlabelled data belong to the same distributions, an assumption that is not necessarily true in real-world scenarios. Moreover, this assumption prohibits us from leveraging free-living unlabelled data with different distributions. In fact, it has been shown in previous studies that incorporating out-of-distribution data with the unlabelled set for SSL impairs performance (Oliver et al. 2018).

To adopt a less constrained approach regarding unlabelled data, *open set* SSL has been proposed (Saito, Kim, and Saenko 2021; Yu et al. 2020), which allows the unlabeled training set to contain samples from classes which are not necessarily present in the labelled set. Nevertheless, this setting still places certain restrictions on the unlabeled data,

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

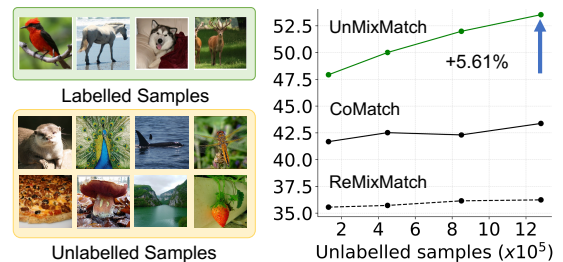


Figure 1: Learning from unconstrained unlabelled data (left). Scaling up unlabelled data provides a large improvement for UnMixMatch (right).

necessitating the inclusion of samples from **every** known class and ensuring that its data distribution is similar. These constraints create two important challenges. First, collecting an unlabelled dataset that necessarily includes samples from certain classes can be challenging in real-world settings. Second, this approach severely restricts the scalability of SSL to large, web-scale, unconstrained, unlabelled data since such data do not hold the aforementioned constraints. Most existing semi-supervised methods are not suitable for learning from unconstrained data as they rely on pseudo-label predictions, which require the ‘unlabeled’ set to have the same classes as the ‘labelled’ set.

In this paper, we propose a novel SSL approach called UnMixMatch, which can learn effective representations from unconstrained unlabelled data and effectively enable SSL to scale up using web-scale unlabelled data. UnMixMatch comprises three main components, which have some similarities to previous SSL methods but have been specifically modified and tailored toward the ‘scalability’ criteria: (1) A supervised learner with hard augmentations: We introduce a new hard augmentation module that combines RandAug with MixUp to prevent overfitting on the small labelled set. The convention of using soft augmentations of the existing literature does not perform well when learning from unconstrained unlabelled data. (2) A contrastive consistency regularizer: The primary unsupervised learning component of our method involves a contrastive regularizer, which learns the underlying data representations by enforcing the model to produce consistent predictions under strong perturbations.

In contrast to existing SSL methods, we do not regularize the class predictions, as the unlabeled set contains unknown classes. Instead, we enforce consistency in the predicted embedding space. (3) A self-supervised pre-text learning module: To further enhance the learned representations, we include a pre-text task called rotation prediction on the unlabeled data, where the model learns by predicting the degree of rotation applied to each sample.

We conduct extensive research on four common datasets: CIFAR-10, CIFAR-100, SVHN, and STL-10. First, we re-implement and benchmark 13 recent semi-supervised methods with unconstrained unlabelled data, using ImageNet-100 as the unlabelled set. We find that existing methods experience performance degradation in unconstrained settings. In comparison, UnMixMatch outperformed existing methods by an average of 4.79%. Additionally, UnMixMatch exhibits robust scaling capabilities regarding the size of the unlabelled datasets, as we observe an additional 5.61% improvement when we increase the unlabelled data size by a factor of 10 (see Fig. 1). Furthermore, we achieved state-of-the-art results in the open set SSL. Finally, we ablate each component of UnMixMatch and demonstrate the crucial role that each component plays in the performance.

In summary, we make the following contributions:

- We propose a novel semi-supervised method that can learn effective representations from unconstrained unlabelled data for the first time.
- We conduct an extensive study to benchmark the performance of existing semi-supervised methods when the unlabelled data are not constrained to match the distribution of the labelled data.
- We demonstrate that our method outperforms previous methods by a large margin in unconstrained learning and sets a new state-of-the-art for open set SSL.
- We also show that the performance of our method scales up by increasing the amount of unconstrained unlabelled data.
- To facilitate reproducibility, we release the code at [github.com/ShuvenduRoy/UnMixMatch](https://github.com/ShuvenduRoy/UnMixMatch).

## 2 Related Work

### 2.1 Constrained Semi-supervised Learning

Prior works on semi-supervised can be broadly divided into two main categories: pseudo-labelling and consistency regularization. Pseudo-labelling techniques (Lee et al. 2013; Xie et al. 2020b) mainly rely on the strategy of predicting pseudo-labels for the unlabeled data using the encoder being trained and learns with a combination of the labelled data and unlabeled data plus their pseudo-labels, and iterating on those predictions to gain gradual improvements. Consistency regularization techniques (Tarvainen and Valpola 2017; Sajjadi, Javanmardi, and Tasdizen 2016; Miyato et al. 2018) learn by forcing the embeddings of different augmented unlabeled samples to be similar. This takes place while the model also simultaneously learns via a supervised loss, which is optimized on the labelled samples. In the pseudo-labelling category, Lee et al. (2013) first introduced

the overall approach, and subsequent works improved the technique by adding various interesting elements. For example, in Noisy Student (Xie et al. 2020b), a pre-trained teacher was introduced to generate the pseudo-labels, and a student learned from the pseudo-labels along with the labelled data. In the consistency regularization category, Pi-model (Sajjadi, Javanmardi, and Tasdizen 2016) was one of the earliest works which used a consistency loss on the predictions of two augmentations of a sample. Later, Mean Teacher (Tarvainen and Valpola 2017) improved the performance by enforcing consistency between the predictions of an online encoder and an exponential moving average (EMA) encoder. VAT (Miyato et al. 2018) is another modification of Mean Teacher, which replaced the augmentations with adversarial perturbations. Later, Unsupervised Domain Adaptation (UDA) (Xie et al. 2020a) showed large improvements by using hard augmentations.

It should be noted that one of the key weaknesses of pseudo-label-based methods is the confidence-bias problem, which arises when the model generates confident wrong pseudo-labels. Yet, their ability to virtually increase the labelled set size by generating pseudo-labels for the unlabeled data has motivated researchers to combine them with consistency regularization methods within the same framework. For instance, MixMatch (Berthelot et al. 2019b) predicts pseudo-labels for unlabelled samples while enforcing consistency across augmented images. ReMixMatch (Berthelot et al. 2019a) improved upon MixMatch with several implementation tricks, such as augmentation anchoring and distribution alignment. FixMatch is another popular hybrid method known for its simplicity and performance. It predicts the pseudo-label of a sample from a weakly augmented image and treats it as a label for a heavily augmented sample when the confidence of the pseudo-label is above a certain threshold. Subsequently, several works have attempted to improve different aspects of FixMatch. FlexMatch (Zhang et al. 2021) employs an adaptive curriculum threshold for each class based on that class’s performance. CoMatch (Li, Xiong, and Hoi 2021) improves upon FixMatch by introducing a graph-based contrastive loss that learns both class representations and low-dimensional embeddings. ConMatch (Kim et al. 2022) also employs a similar concept, using pseudo-labels as supervision in a contrastive loss. Similarly, SimMatch (Zheng et al. 2022) improves FixMatch by introducing an instance similarity loss in addition to the semantic similarity loss imposed by pseudo-labels. ScMatch (Gui et al. 2022) utilizes the concept of clustering with SSL by dynamically forming super-classes.

### 2.2 Open Set Semi-supervised Learning

Open set SSL is a type of SSL where the unlabelled set includes samples from unknown classes. In prior work (Oliver et al. 2018), it was demonstrated that *the presence of unknown classes in the unlabelled dataset has a severe negative impact on the performance of semi-supervised methods*. Similar findings were also reported by Su et al. (2021) that analyzed the performance of more recent semi-supervised methods in open set settings. Nonetheless, some recent methods have attempted to effectively address the detri-

mental effect of unknown classes while learning from open set unlabelled data. For example, Yoshihashi et al. (2019) learned to distinguish known classes from unknown ones, effectively avoiding samples from unknown classes in the learning process. A similar approach was taken in Guo et al. (2020), which proposed a novel scoring function called energy discrepancy to detect and remove instances of unknown classes. OpenMatch (Saito, Kim, and Saenko 2021) used the concept of out-of-distribution to mitigate the negative impact of unknown classes. CCSSL (Yang et al. 2022) introduced a class-aware contrastive learning approach to improve performance in open set settings. In this study, we tackle a more challenging scenario where the unlabelled set contains instances of unknown classes and does not necessarily include all the known classes. Consequently, our goal is not to detect and remove the images of unknown classes, but rather to learn from them.

### 3 Method

#### 3.1 Preliminaries and Overview

Let,  $X_U = (x_i)_{i=1}^N$  be an unlabelled set, and  $X_L = (x_i, y_i)_{i=1}^n$  be a labelled set where  $n \ll N$ . In general, semi-supervised methods learn from  $X_L$  and  $X_U$  in supervised and unsupervised settings, respectively. Formally, SSL is formulated as:

$$\min_{\theta} \sum_{(x,y) \in X_L} \mathcal{L}_S(x, y, \theta) + \alpha \sum_{x \in X_U} \mathcal{L}_U(x, \theta), \quad (1)$$

where,  $\theta$  represents the learnable model,  $\mathcal{L}_S$  is the supervised loss, and  $\mathcal{L}_U$  is the unsupervised loss. Trivially, it is assumed that  $X_L$  and  $X_U$  come from the same **data** and **class** distributions. Let  $Y_l$  and  $Y_u$  be the set of classes for labelled and unlabelled data. In the *constrained* setting,  $Y_l = Y_u$ , and in the *open set* setting,  $Y_l$  is a proper subset of  $Y_u$ , i.e.,  $Y_l \subset Y_u$ . Regarding data distribution, both settings assume that the data comes from the same source and hence have similar distributions. These assumptions are hard to satisfy in a real-world task.

In this study, our objective is to learn from unconstrained data that does not have any particular constraints and can come from different data or class distributions, or both. The unlabelled set may consist of images of unknown classes, where  $Y_u \setminus Y_l \neq \emptyset$ . Additionally, the unlabelled set may not contain all the known classes, and in the extreme case,  $Y_u \cap Y_l = \emptyset$ . To address this challenge, we propose UnMixMatch, a method that can learn effective visual representations from **unconstrained unlabelled** data. UnMixMatch comprises three modules, which we discuss in detail in the following subsections.

#### 3.2 Supervised Module

As mentioned earlier, SSL requires  $X_L$  to be learned using a supervised component. The first contribution of our method is, therefore, to create a supervised learner suitable for our purpose of scalable SSL. Here, we hypothesize that given large amounts of unlabeled data in an unconstrained setting and relatively very small amounts of labelled data, the supervised module may overfit the small labelled set. Thus, unlike FixMatch (Sohn et al. 2020), MixMatch (Berthelot et al.




Figure 2: Overview of our proposed method.

2019b), and ReMixMatch (Berthelot et al. 2019a), which use weak augmentations for their supervised modules, we apply hard augmentations on the labelled samples in our supervised module. This acts as a regularizer for supervised learning and is better able to deal with overfitting compared to weaker augmentations. We utilize RandAug (Cubuk et al. 2020) as the hard augmentation followed by the MixUp operation (Berthelot et al. 2019b,a). We denote RandAug plus MixUp as the RandMixUp operation. Finally, a supervised loss is applied to a batch of samples.

**RandAug.** This is a hard augmentation technique for generating diverse samples by employing a sequence of transformations (Cubuk et al. 2020). More specifically, it applies  $R_n \in 1..13$  transformations randomly chosen from a list of 13 augmentations, including rotation, translation, and colour distortion. The magnitude of each transformation is sampled randomly from a pre-defined range. We denote the augmentation operation as  $\hat{x} = RandAug(x)$ .

**MixUp.** Let  $(x_1, y_1)$  and  $(x_2, y_2)$  be two pairs of samples and their class labels. MixUp operation interpolates between the data points to generate mixed samples and labels as:

$$\bar{x} = \lambda \cdot x_1 + (1 - \lambda) \cdot x_2 \quad (2)$$

$$\bar{y} = \lambda \cdot y_1 + (1 - \lambda) \cdot y_2, \quad (3)$$

where  $\lambda$  is the mixing coefficient. Following MixMatch (Berthelot et al. 2019b), we sample  $\lambda$  from a beta distribution ( $\lambda \sim Beta(\alpha, \alpha)$ ) with hyper-parameter  $\alpha$ .

**Supervised Loss.** For a batch of unlabeled samples  $X_u = ((\hat{x}_i); i \in (1, \dots, b))$ , with batch size  $b$ , we first generate the pseudo-label for each sample  $X_i$  as  $p_i = P_{\theta}(x_i)$ , where  $P_{\theta}$  is the encoder with a classification head.

Next, for a batch of labelled samples  $X_l = ((\hat{x}_i, y_i); i \in (1, \dots, b))$  and the unlabelled samples with pseudo-labels  $X_p = ((\hat{x}_i, p_i); i \in (1, \dots, b))$ , we augment all the samples using  $\bar{X} = RandAugMix(X_l, X_p)$ . Accordingly, we define the supervised loss of our method as:

$$\mathcal{L}_{sup} = \frac{1}{b} \sum_{\bar{x}, \bar{y} \in \bar{X}} \mathcal{H}(\bar{y}, P_{\theta}(y|\bar{x})), \quad (4)$$

where  $\mathcal{H}$  is the cross-entropy loss.

#### 3.3 Consistency Regularization Module

To deal with the unconstrained nature of existing unlabeled data and learn effective representations, we apply a consistency regularizer as our second contribution. A consis-

tency regularizer learns from the unlabelled data by enforcing consistency on its predictions under different augmentations. Prior works that have used consistency regularization for SSL (Sajjadi, Javanmardi, and Tasdizen 2016; Berthelot et al. 2019b,a) enforce consistency on the **class predictions** under different perturbations. However, regularization over class predictions is not useful for learning in *unconstrained* settings where unlabelled data do not necessarily come from the same classes as the labelled data. To address this, we enforce consistency in the low-dimensional embedding space using a contrastive loss. Using contrastive loss on the embedding space enables the model to learn class-agnostic representations from unconstrained unlabelled data. In UnMixMatch, we adopt the Noise Contrastive Estimation loss, a.k.a InfoNCE (Chen et al. 2020).

Contrastive learning learns from positive (perturbations of the same sample) and negative samples (all other samples) by bringing embeddings of the positive samples together and pushing them away for the negatives. For each sample,  $x_i \in X_U$ , two augmentations are applied to generate two augmented images  $\hat{x}_i = \text{RandAug}(x_i)$ . These are first passed through the encoder and a projection head (shallow linear layers with non-linearity and batch normalization) to obtain embeddings  $z_i = P_{\theta_p}(z|\hat{x}_i)$ . The contrastive loss is accordingly defined as:

$$\mathcal{L}_{con} = -\frac{1}{2b} \sum_{i=1}^{2b} \log \frac{\exp(z_i, z_{\kappa(i)}/\tau)}{\sum_{k=1}^{2b} \mathbf{1}_{[k \neq i]} \exp(z_i, z_k/\tau)}, \quad (5)$$

where,  $\kappa(i)$  is the index of the second augmented sample,  $\mathbf{1}_{[k \neq i]}$  is an indicator function which returns 1 when  $k$  is not equal to  $i$ , and 0 otherwise.  $\tau$  is a temperature parameter.

### 3.4 Self-supervised Module

Finally, we intend to enhance the quality of the representations extracted from the unconstrained unlabelled data using the consistency regularizer. It has been recently shown that self-supervised techniques can be employed to learn underlying domain-invariant representations for unlabelled data (Gidaris, Singh, and Komodakis 2018; Zhang, Isola, and Efros 2016). Moreover, this idea has already been demonstrated to be useful in conjunction with SSL (Gidaris, Singh, and Komodakis 2018; Zhai et al. 2019; Berthelot et al. 2019a). As a result, we integrate a straightforward yet effective self-supervised pre-text task called rotation prediction, which learns by predicting the degree of rotations applied to unlabelled images. In practice, a rotation module randomly samples one of the following rotations and applies it to an unlabelled image:  $0^\circ, 90^\circ, 180^\circ, 270^\circ$ . As a result, the rotation prediction task can be viewed as a four-way classification problem, represented as:

$$\mathcal{L}_{rot} = \frac{1}{b} \sum_{u \in U} \mathcal{H}(r, P_{\theta_r}(r|\text{Rotate}(u))) \quad (6)$$

Here,  $P_{\theta_r}$  is the encoder with a rotation head that predicts the rotation, and  $\mathcal{H}$  is the cross-entropy loss.

### 3.5 Total Loss

Finally, we incorporate the loss functions for the three modules above to create the total loss:

$$\mathcal{L}_{UnMixMatch} = \mathcal{L}_{sup} + \beta \mathcal{L}_{con} + \gamma \mathcal{L}_{rot}. \quad (7)$$

Here,  $\beta$  and  $\gamma$  are hyper-parameters that balance the significance of the contrastive and rotation losses. An overview of our proposed method is presented in Fig. 2.

## 4 Experiments and Results

### 4.1 Datasets and Implementation Details

For our main experiments, we follow the standard semi-supervised evaluation protocol from prior works (Sohn et al. 2020), and present the results for four datasets: CIFAR-10 (Krizhevsky, Hinton et al. 2009), CIFAR-100 (Krizhevsky, Hinton et al. 2009), SVHN (Netzer et al. 2011), and STL-10 (Coates, Ng, and Lee 2011). We present the results for different numbers of labelled samples, averaged over three runs. We use ImageNet-1K (Deng et al. 2009), and ImageNet-100 (a subset of ImageNet-1K) as the unconstrained unlabeled datasets since it has different class and data distribution in comparison to the four aforementioned datasets.

Our implementation and hyper-parameters closely follow FlexMatch (Zhang et al. 2021). For the encoder, following existing literature such as (Sohn et al. 2020; Zhang et al. 2021), we use Wide ResNet-28-2 (Zagoruyko and Komodakis 2016) for CIFAR-10 and SVHN, WRN-28-8 for CIFAR-100, and WRN-37-2 for STL-10. We train the method for  $2^{20}$  iterations with a batch size of 64, a learning rate of 0.03, and an SGD optimizer with a momentum of 0.9 and weight-decay of 0.0005. The code is implemented with Pytorch and built using TorchSSL (Zhang et al. 2021).

### 4.2 Results

Here, we present the main results of our method, including the performance on the four datasets (with ImageNet-100 for unlabeled data) in unconstrained settings, analysis of UnMixMatch’s performance with increasing the number of unlabelled data, its effectiveness in open set settings, and its performance in a barely supervised setting.

**Unconstrained Settings.** Table 1 presents the main results of our work on the four datasets. Here, we first reimplement 13 semi-supervised methods and report the results with *unconstrained* unlabelled data. We report the average accuracies and standard deviations across three individual runs for each setting. We also report the average accuracy across all settings for overall comparison. It should be noted that the performance of prior methods is considerably lower than what has been reported in the original papers, where the unlabeled and labelled samples came from the same datasets (unlabeled data were not unconstrained). Next, we observe that UnMixMatch demonstrates superior performance compared to other methods, with an average improvement of 4.79%. We obtain considerable improvement across all datasets and splits, except using CIFAR-100 with 2500 labels, where CCSSL achieves a better result.

When considering the number of labelled samples, we notice that the differences between UnMixMatch and other

Methods	CIFAR-10			CIFAR-100			SVHN			STL-10	Avg.
	40 labels	250 labels	4K labels	400 labels	2.5K labels	10K labels	40 labels	250 labels	1K labels	1K labels	
Supervised	24.24±1.1	43.33±2.2	83.76±0.4	10.39±0.3	39.57±0.4	63.40±0.1	24.67±2.1	75.42±2.4	87.74±0.6	60.96±1.3	51.35
Pi-Model	24.16±1.5	48.14±1.1	84.19±0.3	12.81±0.9	40.01±0.3	63.55±0.1	28.02±1.1	76.56±2.2	87.83±0.6	69.73±0.6	53.50
MeanTeacher	26.22±1.0	49.35±1.7	83.70±0.4	13.72±0.9	41.57±0.3	63.83±0.1	28.57±1.3	76.78±2.3	87.68±0.6	70.12±0.1	54.15
VAT	24.70±1.3	46.18±1.3	84.73±0.2	11.50±0.8	41.73±0.2	63.76±0.1	41.95±2.5	76.18±1.9	88.07±0.5	63.12±0.6	54.19
Pseudo-label	24.88±1.6	50.29±1.3	84.11±0.1	12.12±0.1	39.72±0.8	63.57±0.0	36.04±2.7	78.04±1.2	88.91±0.3	65.60±0.9	54.33
UDA	28.12±2.5	65.59±1.7	88.31±0.1	21.11±0.8	51.82±0.6	69.42±0.6	48.80±4.1	77.73±1.9	88.83±0.3	82.54±0.2	62.23
MixMatch	32.58±0.8	58.24±0.3	84.59±0.3	20.26±0.6	45.94±0.4	65.89±0.2	57.46±1.7	77.65±1.2	89.95±0.2	71.85±0.6	60.44
ReMixMatch	35.56±0.8	64.71±0.6	87.64±0.3	18.90±0.8	49.11±0.9	69.38±0.1	56.94±4.2	79.57±0.9	<u>90.57±0.3</u>	79.13±0.8	63.15
FixMatch	27.91±3.9	64.98±1.0	88.18±0.1	21.20±0.6	51.40±1.3	67.80±0.3	43.71±5.5	74.52±1.0	88.06±0.6	81.85±0.4	60.96
FlexMatch	32.80±0.8	63.22±1.0	87.82±0.1	20.87±0.7	51.28±0.8	69.52±0.5	<u>60.50±2.5</u>	79.72±0.6	88.85±0.4	<u>82.67±0.4</u>	63.72
CoMatch	<u>41.68±0.7</u>	62.31±0.7	84.52±0.3	22.60±0.6	44.00±1.0	58.55±0.3	45.87±2.8	73.19±0.3	86.45±0.2	82.00±0.0	60.12
CCSSL	30.89±5.9	<u>67.20±1.5</u>	<u>88.77±0.1</u>	<u>24.53±1.5</u>	<b>56.30±0.2</b>	<u>71.13±0.3</u>	50.02±6.6	<u>80.39±0.6</u>	88.60±0.3	82.00±0.0	63.98
SimMatch	23.77±1.8	57.72±1.3	84.12±0.7	18.65±1.2	47.33±1.0	66.54±0.8	51.23±1.6	74.48±1.1	88.57±1.0	77.23±1.2	58.96
ScMatch	27.81±1.1	56.78±0.6	83.09±0.2	18.14±1.1	46.21±0.7	64.24±0.2	56.59±0.9	75.08±0.8	89.23±0.3	79.44±0.7	59.66
UnMixMatch	<b>47.93±1.1</b>	<b>68.72±0.6</b>	<b>89.58±0.2</b>	<b>26.13±1.1</b>	<u>54.18±0.7</u>	<b>71.73±0.2</b>	<b>72.90±0.9</b>	<b>80.78±0.8</b>	<b>91.03±0.3</b>	<b>84.73±0.7</b>	<b>68.77</b>

Table 1: Comparison of our method against other SSL methods with *unconstrained unlabelled data* on 4 different datasets. The baseline methods are: Pi-Model (Sajjadi, Javanmardi, and Tasdizen 2016), MeanTeacher (Tarvainen and Valpola 2017), VAT (Miyato et al. 2018), Pseudo-label (Lee et al. 2013), UDA (Xie et al. 2020a), MixMatch (Berthelot et al. 2019b), ReMixMatch (Berthelot et al. 2019a), FixMatch (Sohn et al. 2020), FlexMatch (Zhang et al. 2021), CoMatch (Li, Xiong, and Hoi 2021), CCSSL (Yang et al. 2022), SimMatch (Zheng et al. 2022), and ScMatch (Gui et al. 2022).

Data	IN-100	Subset 1	Subset 2	IN-1K	Method	Accuracy (%)
No. of samples	130K	450K	850K	1.28M	FlexMatch (Zhang et al. 2021)	21.96±1.4
ReMixMatch	35.56	35.72	36.15	36.24	CCSSL (Yang et al. 2022)	15.63±1.7
CoMatch	41.68	42.52	42.31	43.38	UnMixMatch	<b>27.54±2.5</b>
UnMixMatch	<b>47.93</b>	<b>50.01</b>	<b>51.99</b>	<b>53.54</b>		

Table 2: The impact of unlabeled set size. Here, Subsets 1 & 2 are two random subsets of ImageNet-1K (IN-1K).

Methods	Labelled samples/class		
	50	100	400
Supervised	64.3±1.1	69.5±0.7	80.0±0.3
FixMatch (Sohn et al. 2020)	56.8±1.2	70.2±0.6	83.7±0.5
MTC (Yu et al. 2020)	79.7±0.9	86.3±0.9	91.0±0.5
OpenMatch (Saito et al. 2021)	89.6±0.9	92.9±0.5	94.1±0.5
UnMixMatch	<b>95.7±0.8</b>	<b>96.8±0.5</b>	<b>97.2±0.4</b>

Table 3: Performance comparison on open set SSL for CIFAR-10 with 6/4 known-unknown class split.

methods are more pronounced when the labelled set size is small. For example, with only 40 labelled samples from CIFAR-10, UnMixMatch achieves a 17.04% performance gain over CCSSL (which has the second highest overall average performance after ours), and 6.25% higher than the next best result for this specific setting, which was obtained by CoMatch. A similar pattern is observed for SVHN, where UnMixMatch outperforms CCSSL and FlexMatch by 22.88% and 12.4%, respectively.

**Scaling Up the Unlabelled Set.** Our main motivation for using unconstrained unlabelled data is to take advantage of the abundance of free-living unlabeled data. In this experiment, we evaluate the performance of UnMixMatch as the size of the unlabelled set is increased. The results of this study are summarized in Fig. 1 and presented in detail in Ta-

Table 4: Performance comparison on barely supervised learning. Only 1 sample per class is used for training.

ble 2. Specifically, we increase the size of the unlabelled set from 130K images of ImageNet-100 (a subset of ImageNet-1K) to 1.28M images of ImageNet-1K, with two more subsets of 450K and 850K images from ImageNet-1K. We perform this experiment on CIFAR-10 with 40 labelled samples with the two best methods (CoMatch and ReMixMatch) on this setting and observe an increasing trend in the accuracy of UnMixMatch as the number of images in the unlabelled set increases. With ImageNet-1K used as the unlabelled set, which is approximately 10 times larger than ImageNet-100, the accuracy of UnMixMatch improves from 47.93% to 53.54%, a significant improvement of 5.61% by simply increasing the size of the unlabelled set. CoMatch and ReMixMatch, on the other hand, show very small improvements with the increase in the unlabelled data, but the performance difference with our method further increases.

**Results on Open Set Setting.** Next, we investigate the performance of UnMixMatch on open set SSL. Open set SSL is a relatively less challenging setting than *unconstrained* settings, where the unlabelled set may contain images of unknown classes but must contain images of all known classes (Saito, Kim, and Saenko 2021; Yu et al. 2020). For learning in open set settings, we employ a variant of UnMixMatch, which takes advantage of the fact that the unlabelled set contains samples of all known classes from the labelled set and learns from the predicted pseudo-labels on the unlabelled set. In this variant, we replace the con-

trastive loss in our method with the class-aware contrastive loss of CCSSL (Yang et al. 2022). This method first predicts the pseudo-labels for the unlabelled samples and uses them with contrastive loss to learn clusters of known classes in the embedding space. For this experiment, we follow the experimental setup of OpenMatch (Saito, Kim, and Saenko 2021), which reports the results for CIFAR-10 with a 6/4 split. This split means that the labelled set contains images of 6 classes from CIFAR-10, while the unlabelled set includes images of 6 known and 4 unknown classes. Like OpenMatch, we take 6 animal classes as the known classes and 4 object classes as the unknown classes. We perform this experiment using three different splits with 50, 100, and 400 labelled samples per class in the known set. The results are presented in Table 3, where we find that our method outperforms the existing methods and sets a new state-of-the-art for *open set SSL*. Once again, UnMixMatch better demonstrates its effectiveness when the amount of labelled data is limited. With 50 labelled samples per class, our approach provides a 6.1% improvement over the second-best method, OpenMatch. For 100 and 400 labelled samples per class, UnMixMatch shows 3.9% and 3.1% improvements, respectively.

**Barely Supervised Setting.** In this section, we aim to test the performance of SSL in the extreme scenario where only one labelled sample per class is available. This *barely supervised* setting is considered to be very challenging, even with conventional SSL techniques that use *constrained* unlabelled data. Given the extremely low number of labelled samples, the results in such settings generally exhibit high variance, and therefore, we increase the number of folds to 5 to account for this variability. This is due to the fact that the quality of labelled data greatly influences the performance in such low data settings (Sohn et al. 2020; Roy and Etemad 2022). The results of this experiment are presented in Table 4. It shows the performance of CIFAR-10 for the three best methods identified previously in Table 1.

In this experiment, UnMixMatch achieves an accuracy of 27.54% and outperforms other methods by 5.58%. As before, CCSSL struggles to learn in low data settings, achieving a near chance-level accuracy of only 15.63%. However, FlexMatch shows relatively better performance with an accuracy of 21.96%. In general, we find that it is quite difficult to learn effective representations with just one labelled sample per class while using unconstrained unlabelled data. However, as previously shown in Table 1, UnMixMatch quickly gains significant improvements with a small increase in the number of labelled data and achieves an accuracy of 47.93% with four labelled samples per class.

### 4.3 Ablation Study

**Main Components.** Table 5a presents the main ablation results by removing each of the three main components of the proposed method: RandMixUp augmentation, consistency regularization, and rotation prediction. Note that we can not remove two components simultaneously, since semi-supervised learners require a minimum of one supervised and one unsupervised loss. Therefore, we can not perform ablation studies with only one component at a time. These

experiments are done on CIFAR-10 with 40 samples. The table demonstrates that all three components have a significant impact on the final performance of the model, with the removal of any one component resulting in a large drop in accuracy. In the first ablation experiment, we remove the RandMixUp augmentation module, effectively learning from the labelled samples with weak augmentations only (random resized crop and horizontal flip). This experiment results in the highest drop in accuracy across the ablation settings, with a 9.21% decrease. The second largest drop in accuracy is observed when removing consistency regularization, resulting in a 6.68% decrease. Similarly, removing the rotation prediction component results in a 6.1% drop in performance.

**RandMixUp.** As demonstrated in the main ablation study, the RandMixUp augmentation is a crucial component of our approach. In Table 5b, we present the accuracy of CIFAR-10 with 40 samples for different alternatives to RandMixUp. Recall that RandMixUp combines RandAug (Cubuk et al. 2020) and MixUp (Berthelot et al. 2019b). We first test the accuracy for RandAug and MixUp individually and observe a decrease in performance for both settings alone, with a 1.92% and 2.7% decrease, respectively. Here, using RandAug exhibits a lower drop than MixUp, suggesting the higher importance of RandAug in the RandMixUp. This analysis also shows that, unlike MixMatch and ReMixMatch, learning using MixUp does not have the most significant impact on the supervised component. Instead, the role of hard augmentation as a regularizer is the key to the impact of RandMixUp in our method.

We also examine the performance of other well-known hard augmentation techniques, namely CutMix and AugMix. Like MixUp, CutMix combines two samples, but in this case, a part of the second sample is cut and inserted into the first sample to create a new one. Despite the conceptual similarity, CutMix yields 2.11% worse accuracy than MixUp, which is 4.81% lower than the accuracy achieved by the proposed RandAugMix. On the other hand, AugMix shares a similar concept as RandAug. While RandAug applies multiple augmentations in sequence, AugMix applies multiple augmentations separately and generates the final sample by interpolating between them. AugMix achieves an accuracy of 44.75%, which is 1.26% lower than that of RandAug and 3.18% lower than that of RandAugMix. The study’s overall findings demonstrate the critical importance of RandAugMix in our method, with other conceptually similar augmentation techniques failing to achieve the same level of performance.

**Consistency Regularizer.** As shown in the main ablation study, the contrastive regularizer is the second most crucial component of our method. Recall that the proposed contrastive loss for UnMixMatch is a noise contrastive estimation loss that has gained popularity in the self-supervised learning literature (Oord, Li, and Vinyals 2018; Chen et al. 2020). However, other variants of contrastive loss have also been introduced in the literature (Kim et al. 2022; Lee et al. 2022; Li, Xiong, and Hoi 2021; Yang et al. 2022) and have shown improvements in different problem settings, including SSL. In this study, we investigate four variants of con-

Ablation	Accuracy	Augmentation	Accuracy	Contrastive loss	Accuracy
UnMixMatch	<b>47.93</b>	RandMixUp	47.93	UnMixMatch	47.93
w/o RandMixUp	38.72	RandAug	46.01	ConMatch	45.82
w/o Contrast Loss	41.25	MixUp	45.23	Contrastive Reg.	45.77
w/o Rotation Loss	41.83	CutMix	43.12	Graph Contrast	47.12
		AugMix	44.75	Class-aware Contrast	47.88

(a) Ablation of main components.      (b) Alternate hard augmentations.      (c) Variants of contrastive regularizers.

Regularization strategy	Accuracy	Self-sup. loss	Accuracy	Method	Accuracy
Cont. two strong (ours)	47.93	UnMixMatch	47.93	Rot. Pred.	47.93
Cont. on prediction	42.45	BYOL	46.89	Colorization	44.25
Cont. weak & strong	44.15	SimSiam	46.50	Jigsaw Solving	43.98
		VICReg	47.91	Image Inpaint	45.56

(d) Alternate regularization strategies.      (e) Alternate contrastive regularizers.      (f) Alternate pre-text tasks.

Table 5: Ablation studies on our method. All studies are on CIFAR-10 with 40 labelled samples.

trastive loss in the proposed framework. More specifically, we explore the contrastive variants of ConMatch (Kim et al. 2022) and Contrastive Regularization (Lee et al. 2022), as well as graph contrastive learning (Li, Xiong, and Hoi 2021), and class-aware contrastive learning (Yang et al. 2022). ConMatch (Kim et al. 2022) is based on a variant of contrastive loss that involves two hard augmentations and one weak augmentation. In Contrastive Regularization (Lee et al. 2022), high-confidence pseudo-labels were used for supervision while learning with a contrastive loss. Two other methods proposed slightly different versions of utilizing the pseudo-labels in contrastive learning settings: CoMatch (Li, Xiong, and Hoi 2021) and CCSSL (Yang et al. 2022), where CoMatch used a graph contrastive learning, and CCSSL utilized pseudo-labels with a contrastive loss.

We show the results for variants of contrastive loss in Table 5c. Using the contrastive variant of both ConMatch and Contrastive Regularizer result in a large drop in accuracy with 2.16% and 2.11%, respectively. Using the graph contrast concept of CoMatch with UnMixMatch results in a 47.12% accuracy, dropping by 0.81%. Finally, the class-aware contrastive loss of CCSSL gets the closest accuracy of the proposed UnMixMatch with 47.88% accuracy. In Table 6, we present expanded results across all datasets for the graph contrast from CoMatch (Li, Xiong, and Hoi 2021) and the class-aware contrast from CCSSL (Yang et al. 2022). Results in this table show that on average the proposed contrastive regularizer achieves 2.99% and 4.45% improvements over the graph contrast and class-aware contrast alternatives. The class-aware contrast obtains better results than the proposed method in one setting (CIFAR-100 with 400 labelled samples) only.

Next, we investigate different regularization strategies. First, we study the choice of applying the contrastive loss on the low-dimensional embedding space rather than the final prediction. In Table 5d, we present the results for this experiment, where we observe a 5.48% decrease in accuracy, indicating the importance of learning representations by regularizing the embedding space instead of the class predictions. Next, we examine another aspect of our regularization

approach by replacing two strong augmentations with one weak and one strong augmentation (similar to FixMatch). This variant results in a 3.78% drop in accuracy.

Finally, we show the performance of our method by completely replacing the contrastive loss with other important viable losses, namely, BYOL (Grill et al. 2020), SimSiam (Chen and He 2021), and VICReg (Bardes, Ponce, and LeCun 2021). BYOL (Grill et al. 2020) learns by predicting the representation of a target encoder (exponential moving average of the online encoder), using the online encoder for different augmentations of the same sample. SimSiam (Chen and He 2021) also learns by matching the representations of two augmented samples, but unlike previous methods, it doesn’t need negative samples or the target encoder. VICReg (Bardes, Ponce, and LeCun 2021) learns from unlabelled data by combining three terms: variance, invariance, and covariance regularizations. The results of this experiment are presented in Table 5e, where we see that the accuracy of BYOL and SimSiam is considerably lower than the proposed contrastive loss with 46.89% and 46.50%. However, we find VICReg to have a competitive performance. To fully understand the performance of VICReg, we present the results for all datasets in Table 7. This table shows that VICReg achieves better results for a few settings (CIFAR-10 with 250 labelled samples, SVHN with 250 and 1000 labelled samples, and STL-10 with 1000 labelled samples). However, the overall average accuracy with VICReg loss is 0.92% lower than the contrastive regularizer, and hence, we choose this as the default for the proposed UnMixMatch.

**Rotation Loss.** We also investigate a few other self-supervised approaches instead of the rotation prediction task, including colorization, jigsaw solving, and image inpainting. The results of this experiment are included in Table 5f, where rotation prediction shows a clear advantage over other pre-text tasks. Apart from the fact that rotation prediction performs well, it has two additional benefits over other methods. First, it can be integrated into our semi-supervised framework without adding considerable computational overhead. In contrast, tasks like colorization or in-

Methods	CIFAR-10			CIFAR-100			SVHN			STL-10	Avg.
	40 labels	250 labels	4K labels	400 labels	2.5K labels	10K labels	40 labels	250 labels	1K labels	1K labels	
UnMixMatch	<b>47.93±1.1</b>	<b>68.72±0.6</b>	<b>89.58±0.2</b>	26.13±1.1	<b>54.18±0.7</b>	<b>71.73±0.2</b>	<b>72.90±0.9</b>	<b>80.78±0.8</b>	<b>91.03±0.3</b>	<b>84.73±0.7</b>	<b>68.77</b>
Graph Con.	47.12±0.5	67.67±0.7	87.30±0.3	25.38±0.4	49.14±0.6	68.08±0.3	62.54±3.6	79.78±0.9	88.49±0.9	82.25±0.5	65.78
Class Con.	47.88±2.4	68.45±0.4	88.39±0.2	<b>26.59±0.3</b>	50.25±0.7	69.29±0.2	43.99±12.2	78.05±0.2	87.76±0.1	82.59±0.4	64.32

Table 6: Comparison of our method against different variants of contrastive regularizers on 4 different datasets.

Methods	CIFAR-10			CIFAR-100			SVHN			STL-10	Avg.
	40 labels	250 labels	4K labels	400 labels	2.5K labels	10K labels	40 labels	250 labels	1K labels	1K labels	
UnMixMatch	<b>47.93±1.1</b>	68.72±0.6	<b>89.58±0.2</b>	<b>26.13±1.1</b>	<b>54.18±0.7</b>	<b>71.73±0.2</b>	<b>72.9±0.9</b>	80.78±0.8	91.03±0.3	84.73±0.7	<b>68.77</b>
VICReg	47.91±1.3	<b>69.75±0.2</b>	89.08±0.1	24.74±0.5	50.25±0.8	68.51±0.6	68.94±2.9	<b>82.67±0.3</b>	<b>91.21±0.3</b>	<b>85.44±0.5</b>	67.85

Table 7: Comparison of our method for different alternate regularization strategies on 4 different datasets.

Setting	Accuracy (%)
UnMixMatch	47.93
MixUp only	45.23
w/ Match loss	44.25
w/ KL-div. loss	47.02

Table 8: Study on the relation of our method to prior works.

painting require an additional decoder to train, which increases the overall training cost. Second, rotation prediction has already demonstrated effectiveness with other semi-supervised methods such as ReMixMatch on many datasets.

**Relation to Existing Methods.** Our method has a few similarities with some of the previous semi-supervised methods. For example, MixMatch (Berthelot et al. 2019b) popularized the use of MixUp (Zhang et al. 2018) in the context of SSL. ReMixMatch (Berthelot et al. 2019a) built over MixMatch by using MixUp and a rotation prediction loss, making it the most similar to our proposed method. However, our method has certain distinctions from ReMixMatch. Firstly, ReMixMatch only uses MixUp with the supervised module, whereas ours uses RandMixUp. Moreover, our approach uses a contrastive loss as a regularizer, which is applied to the intermediate embeddings, whereas ReMixMatch uses a match loss on the predictions of mixed samples. Finally, ReMixMatch uses KL-divergence loss between the predictions of weakly and strongly augmented samples, which our method does not employ. Overall, recall from Table 1 that ReMixMatch, on average, shows 5.62% lower performance than our method. In Table 8, we further investigate the differences by incorporating the design choices of ReMixMatch into our method. First, by removing the RandAug from RandMixUp (using MixUp like ReMixMatch), we obtain an accuracy of 45.23% (2.7%↓). Next, using the match loss of ReMixMatch as a regularizer on the class predictions in our method results in an accuracy of 44.25% (3.68%↓). Finally, adding the KL-divergence loss to our method further results in a 0.91% drop.

#### 4.4 Sensitivity Study

Our method involves three important hyper-parameters, namely the  $\alpha$  value in MixUp, the contrastive loss weight  $\beta$ , and the rotation loss weight  $\gamma$ . Here, we present a sensitivity




Figure 3: Sensitivity study on important hyper-parameters.

analysis of different values for these hyper-parameters. As illustrated in Fig. 3a, the best accuracy is achieved with an alpha value of 0.1, while very large or small values of  $\alpha$  lead to a drop in accuracy. Fig. 3b indicates that the optimal performance is obtained with a  $\beta$  of 1.0. Furthermore, Fig. 3c reveals that a higher value of  $\gamma$  leads to better performance, highlighting the role of the self-supervised loss.

## 5 Conclusion

Existing semi-supervised methods struggle to learn when the assumption that the unlabeled data comes from the same distribution as the labelled data, is violated. This work proposes a new semi-supervised method called UnMixMatch for learning from unconstrained unlabelled data. Our method shows large improvements over existing methods and even larger improvements under low-labelled data settings. Our approach also outperforms existing methods on *open set* settings. Most importantly, UnMixMatch scales up in performance when the size of unlabelled data increases. We hope this research will draw attention to this more challenging and realistic SSL setting with unconstrained unlabelled data.

**Limitations.** While our comprehensive study on common benchmark datasets demonstrates the effectiveness of UnMixMatch across diverse image domains, some aspects require further investigation for universal applicability. Specifically, the use of hard augmentations and the rotation prediction task might necessitate additional tuning or minor modifications for optimal performance in specialized domains.

## Acknowledgements

This work was supported by Mitacs, BMO, and Ingenuity Labs Research Institute. We are also thankful to SciNet HPC Consortium for helping with the computation resources.

## References

- Bardes, A.; Ponce, J.; and LeCun, Y. 2021. Vi-creg: Variance-invariance-covariance regularization for self-supervised learning. In *International Conference on Learning Representations*.
- Berthelot, D.; Carlini, N.; Cubuk, E. D.; Kurakin, A.; Sohn, K.; Zhang, H.; and Raffel, C. 2019a. Remixmatch: Semi-Supervised learning with distribution matching and augmentation anchoring. In *International Conference on Learning Representations*.
- Berthelot, D.; Carlini, N.; Goodfellow, I.; Papernot, N.; Oliver, A.; and Raffel, C. A. 2019b. Mixmatch: A holistic approach to semi-supervised learning. In *Advances in Neural Information Processing Systems*, 5049–5059.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A simple framework for contrastive learning of visual representations. In *International Conference on Machine Learning*, 1597–1607.
- Chen, X.; and He, K. 2021. Exploring simple siamese representation learning. In *Computer Vision and Pattern Recognition*, 15750–15758.
- Coates, A.; Ng, A.; and Lee, H. 2011. An analysis of single-layer networks in unsupervised feature learning. In *International Conference on Artificial Intelligence and Statistics*, 215–223.
- Cubuk, E. D.; Zoph, B.; Shlens, J.; and Le, Q. V. 2020. Randaugment: Practical automated data augmentation with a reduced search space. In *Computer Vision and Pattern Recognition Workshops*, 702–703.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *International Conference on Computer Vision*, 248–255.
- Gidaris, S.; Singh, P.; and Komodakis, N. 2018. Unsupervised representation learning by predicting image rotations. In *International Conference on Learning Representations*.
- Grill, J.-B.; Strub, F.; Altché, F.; Tallec, C.; Richemond, P.; Buchatskaya, E.; Doersch, C.; Avila Pires, B.; Guo, Z.; Gheshlaghi Azar, M.; et al. 2020. Bootstrap your own latent—a new approach to self-supervised learning. In *Advances in Neural Information Processing Systems*, 21271–21284.
- Gui, G.; Zhao, Z.; Qi, L.; Zhou, L.; Wang, L.; and Shi, Y. 2022. Improving barely supervised learning by discriminating unlabeled samples with super-class. In *Advances in Neural Information Processing Systems*, 19849–19860.
- Guo, L.-Z.; Zhang, Z.-Y.; Jiang, Y.; Li, Y.-F.; and Zhou, Z.-H. 2020. Safe deep semi-supervised learning for unseen-class unlabeled data. In *International Conference on Machine Learning*, 3897–3906.
- Kim, J.; Min, Y.; Kim, D.; Lee, G.; Seo, J.; Ryoo, K.; and Kim, S. 2022. Conmatch: Semi-supervised learning with confidence-guided consistency regularization. In *European Conference on Computer Vision*, 674–690.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images. *Technical report, University of Toronto*.
- Lee, D.; Kim, S.; Kim, I.; Cheon, Y.; Cho, M.; and Han, W.-S. 2022. Contrastive regularization for semi-supervised learning. In *Computer Vision and Pattern Recognition*, 3911–3920.
- Lee, D.-H.; et al. 2013. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on Challenges in Representation Learning, ICML*.
- Li, J.; Xiong, C.; and Hoi, S. C. 2021. Comatch: Semi-supervised learning with contrastive graph regularization. In *International Conference on Computer Vision*, 9475–9484.
- Miyato, T.; Maeda, S.-i.; Koyama, M.; and Ishii, S. 2018. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, (8): 1979–1993.
- Netzer, Y.; Wang, T.; Coates, A.; Bissacco, A.; Wu, B.; and Ng, A. Y. 2011. Reading digits in natural images with unsupervised feature learning. In *NeurIPS Workshop on Deep Learning and Unsupervised Feature Learning*.
- Oliver, A.; Odena, A.; Raffel, C. A.; Cubuk, E. D.; and Goodfellow, I. 2018. Realistic evaluation of deep semi-supervised learning algorithms. In *Advances in Neural Information Processing Systems*, 3239–3250.
- Oord, A. v. d.; Li, Y.; and Vinyals, O. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Roy, S.; and Etemad, A. 2022. Impact of labelled set selection and supervision policies on semi-supervised learning. *arXiv preprint arXiv:2211.14912*.
- Saito, K.; Kim, D.; and Saenko, K. 2021. Openmatch: Open-set consistency regularization for semi-supervised learning with outliers. In *Advances in Neural Information Processing Systems*, 25956–25967.
- Sajjadi, M.; Javanmardi, M.; and Tasdizen, T. 2016. Regularization with stochastic transformations and perturbations for deep semi-supervised learning. In *Advances in Neural Information Processing Systems*, 1171–1179.
- Sohn, K.; Berthelot, D.; Carlini, N.; Zhang, Z.; Zhang, H.; Raffel, C. A.; Cubuk, E. D.; Kurakin, A.; and Li, C.-L. 2020. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. In *Advances in Neural Information Processing Systems*, 596–608.
- Su, J.-C.; Cheng, Z.; and Maji, S. 2021. A realistic evaluation of semi-supervised learning for fine-grained classification. In *Computer Vision and Pattern Recognition*, 12966–12975.

- Tarvainen, A.; and Valpola, H. 2017. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *Advances in Neural Information Processing Systems*, 1195–1204.
- Xie, Q.; Dai, Z.; Hovy, E.; Luong, T.; and Le, Q. 2020a. Unsupervised data augmentation for consistency training. In *Advances in Neural Information Processing Systems*, 6256–6268.
- Xie, Q.; Luong, M.-T.; Hovy, E.; and Le, Q. V. 2020b. Self-training with noisy student improves imagenet classification. In *Computer Vision and Pattern Recognition*, 10687–10698.
- Yang, F.; Wu, K.; Zhang, S.; Jiang, G.; Liu, Y.; Zheng, F.; Zhang, W.; Wang, C.; and Zeng, L. 2022. Class-aware contrastive semi-supervised learning. In *Computer Vision and Pattern Recognition*, 14421–14430.
- Yoshihashi, R.; Shao, W.; Kawakami, R.; You, S.; Iida, M.; and Naemura, T. 2019. Classification-reconstruction learning for open-set recognition. In *Computer Vision and Pattern Recognition*, 4016–4025.
- Yu, Q.; Ikami, D.; Irie, G.; and Aizawa, K. 2020. Multi-task curriculum framework for open-set semi-supervised learning. In *European Conference on Computer Vision*, 438–454.
- Zagoruyko, S.; and Komodakis, N. 2016. Wide residual networks. In *British Machine Vision Conference*.
- Zhai, X.; Oliver, A.; Kolesnikov, A.; and Beyer, L. 2019. S4I: Self-supervised semi-supervised learning. In *International Conference on Computer Vision*, 1476–1485.
- Zhang, B.; Wang, Y.; Hou, W.; Wu, H.; Wang, J.; Okumura, M.; and Shinozaki, T. 2021. Flexmatch: boosting semi-supervised learning with curriculum pseudo labeling. In *Advances in Neural Information Processing Systems*, 18408–18419.
- Zhang, H.; Cisse, M.; Dauphin, Y. N.; and Lopez-Paz, D. 2018. mixup: Beyond empirical risk minimization. In *International Conference on Learning Representations*.
- Zhang, R.; Isola, P.; and Efros, A. A. 2016. Colorful image colorization. In *European Conference on Computer Vision*, 649–666.
- Zheng, M.; You, S.; Huang, L.; Wang, F.; Qian, C.; and Xu, C. 2022. SimMatch: Semi-supervised learning with similarity matching. In *Computer Vision and Pattern Recognition*, 14471–14481.