

SimPLE: Similar Pseudo Label Exploitation for Semi-Supervised Classification

Zijian Hu* Zhengyu Yang* Xuefeng Hu Ram Nevatia
University of Southern California

{zijianhu, yang765, xuefengh, nevatia}@usc.edu

Abstract

A common classification task situation is where one has a large amount of data available for training, but only a small portion is annotated with class labels. The goal of semi-supervised training, in this context, is to improve classification accuracy by leverage information not only from labeled data but also from a large amount of unlabeled data. Recent works [2, 1, 26] have developed significant improvements by exploring the consistency constrain between differently augmented labeled and unlabeled data. Following this path, we propose a novel unsupervised objective that focuses on the less studied relationship between the high confidence unlabeled data that are similar to each other. The new proposed Pair Loss minimizes the statistical distance between high confidence pseudo labels with similarity above a certain threshold. Combining the Pair Loss with the techniques developed by the MixMatch family [2, 1, 26], our proposed SimPLE algorithm shows significant performance gains over previous algorithms on CIFAR-100 and Mini-ImageNet [31], and is on par with the state-of-the-art methods on CIFAR-10 and SVHN. Furthermore, SimPLE also outperforms the state-of-the-art methods in the transfer learning setting, where models are initialized by the weights pre-trained on ImageNet[15] or DomainNet-Real[23]. The code is available at github.com/zijian-hu/SimPLE.

1. Introduction

Deep learning has recently achieved state-of-the-art performance on many computer vision tasks. One major factor in the success of deep learning is the large labeled datasets. However, labeling large datasets is very expensive and often not feasible, especially in domains that require expertise to provide labels. Semi-Supervised Learning (SSL), on the other hand, can take advantage of partially labeled data, which is much more readily available, as shown in figure 1.

A critical problem in semi-supervised learning is how to generalize the information learned from limited label data

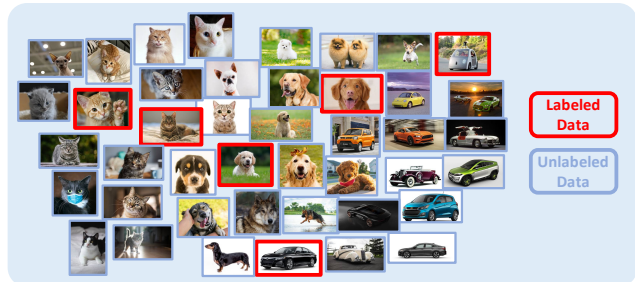


Figure 1: Illustration of an image set with a limited amount of labeled images among a large number of unlabeled images. Unlike unsupervised learning methods that only exploit the structure from unlabeled data, and supervised learning methods that only look at the limited amount of labeled data, semi-supervised learning utilizes information from both labeled and unlabeled data.

to unlabeled data. Following the *continuity assumption* that close data have a higher probability of sharing the same label [4], many approaches have been developed [27, 37, 8], including the recently proposed Label Propagation [14].

Another critical problem in semi-supervised learning is how to directly learn from the large amount of unlabeled data. Maintaining consistency between differently augmented unlabeled data has been recently studied and proved to be an effective way to learn from unlabeled data in both self-supervised learning [5, 11] and semi-supervised learning [16, 25, 2, 1, 26, 22, 28, 36, 32]. Other than consistency regularization, a few other techniques have also been developed for the semi-supervised learning to leverage the unlabeled data, such as entropy minimization [21, 17, 10] and generic regularization [13, 19, 35, 34, 30].

The recently proposed MixMatch[2] combined the above techniques and designed a unified loss function to let the model learn from differently augmented labeled and unlabeled data, together with the mix-up [35] technique, which encourages convex behavior between samples to increase models' generalization ability. ReMixMatch [1] further improves the MixMatch by introducing the Distribution Alignment and Augmentation Anchoring techniques, which al-

*Equal contributions; names ordered alphabetically.

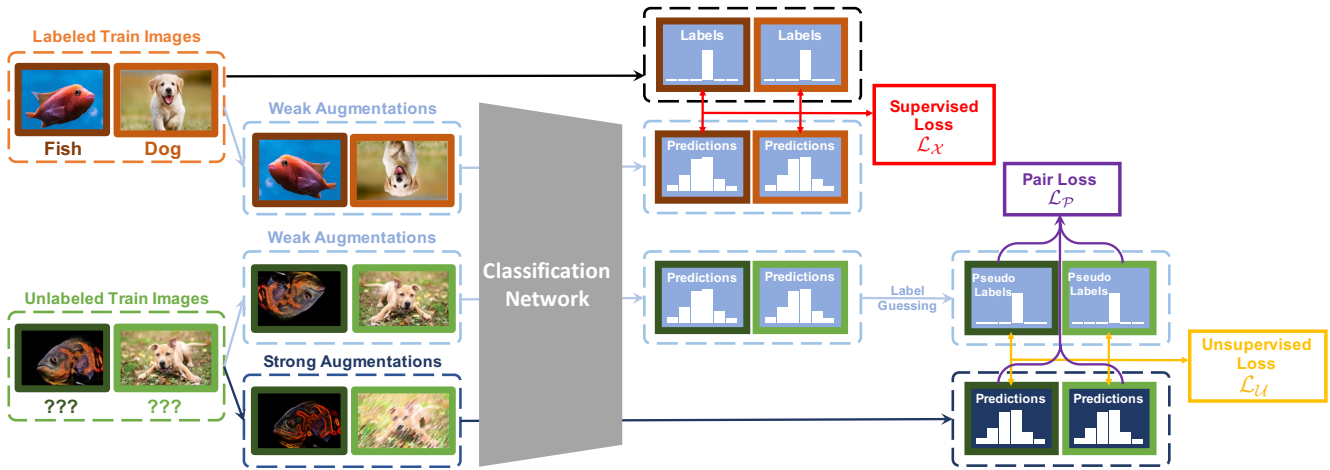


Figure 2: An overview of the proposed SimPLE algorithm. SimPLE optimizes the classification network with three training objectives: 1) supervised loss \mathcal{L}_X for augmented labeled data; 2) unsupervised loss \mathcal{L}_U that aligns the strongly augmented unlabeled data with pseudo labels generated from weakly augmented data; 3) Pair Loss \mathcal{L}_P that minimizes the statistical distance between predictions of strongly augmented data, based on the similarity and confidence of their pseudo labels.

lows the model to accommodate and leverage from the heavily augmented samples. FixMatch [26] simplifies its previous works by introducing a confidence threshold into its unsupervised objective function and achieves state-of-the-art performance over the standard benchmarks.

However, while Label Propagation [14] mainly focuses on the relationship between labeled data to unlabeled data, and the MixMatch family [2, 1, 26] primarily focuses on the relationship between differently augmented unlabeled samples, the relationship between different unlabeled samples is less studied.

In this paper, we propose to take advantage of the relationship between different unlabeled samples. We introduce a novel Pair Loss, which minimizes the distance between similar unlabeled samples of high confidence.

Combining the techniques developed by the MixMatch family [2, 1, 26], we propose the SimPLE algorithm. As shown in figure 2, the SimPLE algorithm generates pseudo labels of unlabeled samples by averaging and sharpening the predictions on multiple weakly augmented variations of the same sample. Then, we use both the labels and pseudo labels to compute the supervised cross-entropy loss and unsupervised L_2 distance loss. These two terms push the decision boundaries to go through low-density areas and encourage consistency among different variations of the same samples. Finally, with the newly proposed Pair Loss, we harness the relationships among the pseudo labels of different samples by encouraging consistency among different unlabeled samples which share a great similarity.

Our contribution can be described in four folds:

- We propose a novel unsupervised loss term that leverages the information from high confidence similar un-

labeled data pairs.

- Combining the techniques from MixMatch family [2, 1, 26] with the new Pair Loss, we developed the novel SimPLE algorithm for semi-supervised learning.
- We performed extensive experiments on the standard benchmarks and demonstrated the effectiveness of the proposed Pair Loss. SimPLE outperforms the state-of-the-art methods on CIFAR100 and Mini-ImageNet and on par with the state-of-the-art methods on CIFAR10, SVHN.
- We also evaluated our algorithm in a realistic setting where SSL methods are applied on pre-trained models, where the new proposed SimPLE algorithm also outperforms the current state-of-the-art methods.

2. Related Work

2.1. Consistency Regularization

Consistency regularization is widely used in the field of SSL. It refers to the idea that a model’s response to an input should remain consistent, when perturbations are used on the input or the model. The idea is first proposed in [16, 25]. In its simplest form, the regularization can be achieved via the loss term:

$$\|P_{\text{model}}(y|A(x); \theta) - P_{\text{model}'}(y|A(x); \theta)\|_2^2 \quad (1)$$

The stochastic transformation $A(x)$ can be either domain-specific data augmentation [2, 16, 25, 1], drop out [25], random max pooling [25], or adversarial transformation [22]. A further extension of this idea is to “perturb” the model,

p_{model} , instead of the input. The perturbation can be a time ensembling of model at different time step [16, 28], or an adversarial perturbation on model’s parameter θ [36]. Also, many works choose to minimize cross entropy instead of the L_2 norm [22, 32, 1, 26].

2.1.1 Augmentation Anchoring

Augmentation Anchoring is first proposed by ReMixMatch [1] and further developed in FixMatch [26]. It is a form of consistency regularization that involves applying different levels of perturbations to the input. A model’s response to a slightly perturbed input is regarded as the “anchor”, and we try to align model’s response to a severely perturbed input to the anchor. For example, we can slightly perturb the input by applying an “easy” augmentation such as horizontal flipping and severely perturb the input by applying a “hard” augmentation such as Gaussian blurring. As the model’s response to a slightly perturbed input is less unstable, including Augmentation Anchoring increases the stability of the regularization process.

2.2. Pseudo-labeling

Pseudo labels are artificial labels generated by the model itself and are used to further train the model. Lee [17] picks the class with the highest predicted probability by the model as the pseudo label. However, pseudo labels are only used during the fine-tuning phase of the model, which has been pre-trained. When we minimize the entropy on pseudo labels, we encouraged decision boundaries among clusters of unlabeled samples to be in the low-density region, which is requested by *low-density separation assumption* [4]. In this paper, for simplicity, we use a single lower case letter, $p \sim \Delta^N$ (the N -probability simplex), to represent either a hard label (a one-hot vector) or a soft label (a vector of probabilities). A simple yet powerful extension of pseudo-labeling is to filter pseudo labels based on a confidence threshold [9, 26]. We define the confidence of a pseudo label as the highest probability of it being any class (i.e., $\max_i(p_i)$). For simplicity, from now on, we will use $\max(p)$ as a shorthand notation for the confidence of any label p . With a predefined confidence threshold τ_c , we reject all pseudo labels whose confidence is below the threshold (i.e., $\max(p) < \tau_c$). The confidence threshold allows us to focus on labels with high confidence (low entropy) that are away from the decision boundaries.

2.3. Label Propagation

Label propagation is a graph-based idea that tries to build a graph whose nodes are the labeled and unlabeled samples, and edges are weighted by the similarity between those samples [4]. Although it is traditionally considered as a transductive method [27, 37], recently, it has been used in

an inductive setting as a way to give pseudo labels. In [14], the authors measure the similarity between feature representations of labeled and unlabeled samples embedded by a CNN. Then, each sample is connected with the K neighbors with the highest similarity to construct the affinity graph. After pre-training the model in a supervised fashion, they train the model and propagate the graph alternatively. The idea of using K nearest neighbors to build a graph efficiently is proposed in [8], as most edges in the graph should have a weight close to 0. Our similarity threshold, τ_c , takes a similar role.

3. Method

To take full advantage of the vast quantity of unlabeled samples in SSL problems, we propose the SimPLE algorithm that focuses on the relationship between unlabeled samples. In the following section, we first describe the semi-supervised image classification problem. Then, we develop the major components of our methods and incorporate everything into our proposed SimPLE algorithm.

3.1. Problem Description

We define the semi-supervised image classification problem as following. In a L -class classification setting, let $\mathcal{X} = ((x_b, y_b); b \in (1, \dots, B))$ be a batch of labeled data, and $\mathcal{U} = (u_b; b \in (1, \dots, B))$ be a batch of unlabeled data. Let $p_{\text{model}}(\tilde{y} | x; \theta)$ denote the model’s predicted softmax class probability of input x parameterized by weight θ .

3.2. Augmentation Strategy

Our algorithm uses Augmentation Anchoring [1, 26], in which pseudo labels come from weakly augmented samples act as “anchor”, and we align the strongly augmented samples to the “anchor”. Our weak augmentation, follows that of MixMatch[2], ReMixMatch [1], and FixMatch [26], contains a random cropping followed by a random horizontal flip. We use RandAugment [6] or a fixed augmentation strategy that contains difficult transformations such as random affine and color jitter as strong augmentation. For every batch, RandAugment randomly selects a fixed number of augmentations from a predefined pool; the intensity of each transformation is determined by a magnitude parameter. In our experiments, we find that method can adapt to high-intensity augmentation very quickly. Thus, we simply fix the magnitude to the highest value possible.

3.3. Pseudo-labeling

Our pseudo labeling is based on the label guessing technique used in [2]. We first take the average of the model’s predictions of several weakly augmented versions of the same unlabeled sample as its pseudo label. As the prediction is averaged from K slight perturbations of the same

Algorithm 1 SimPLE algorithm

```

1: Input: Batch of labeled examples and their one-hot labels  $\mathcal{X} = ((x_b, y_b); b \in (1, \dots, B))$ , batch of unlabeled examples  $\mathcal{U} = (u_b; b \in (1, \dots, B))$ , sharpening temperature  $T$ , number of weak augmentations  $K$ , number of strong augmentations  $K_{\text{strong}}$ , confidence threshold  $\tau_c$ , similarity threshold  $\tau_s$ .
2: for  $b = 1$  to  $B$  do
3:    $\tilde{x}_b = A_{\text{weak}}(x_b)$  ▷ Apply weak data augmentation to  $x_b$ 
4:   for  $k = 1$  to  $K$  do
5:      $\tilde{u}_{b,k} = A_{\text{weak}}(u_b)$  ▷ Apply  $k^{\text{th}}$  round of weak data augmentation to  $u_b$ 
6:   end for
7:   for  $k = 1$  to  $K_{\text{strong}}$  do
8:      $\hat{u}_{b,k} = A_{\text{strong}}(u_b)$  ▷ Apply  $k^{\text{th}}$  round of strong augmentation to  $u_b$ 
9:   end for
10:   $\bar{q}_b = \frac{1}{K} \sum_{k=1}^K \text{P}_{\text{model}'}(\tilde{y} | \tilde{u}_{b,k}; \theta)$  ▷ Compute average predictions across all weakly augmented  $u_b$  using EMA
11:   $q_b = \text{Sharpen}(\bar{q}_b, T)$  ▷ Apply temperature sharpening to the average prediction
12: end for
13:  $\hat{\mathcal{X}} = ((\tilde{x}_b, y_b); b \in (1, \dots, B))$  ▷ Weakly augmented labeled examples and their labels
14:  $\hat{\mathcal{U}} = ((\hat{u}_{b,k}, q_b); b \in (1, \dots, B), k \in (1, \dots, K_{\text{strong}}))$  ▷ Strongly augmented unlabeled examples, guessed labels
15:  $\mathcal{L}_{\mathcal{X}} = \frac{1}{|\mathcal{X}'|} \sum_{x, y \in \mathcal{X}'} H(y, \text{P}_{\text{model}}(\tilde{y} | x; \theta))$  ▷ Compute supervised loss
16:  $\mathcal{L}_{\mathcal{U}} = \frac{1}{L|\hat{\mathcal{U}}|} \sum_{u, q \in \hat{\mathcal{U}}} \mathbb{1}_{(\max(q) > \tau_c)} \|q - \text{P}_{\text{model}}(\tilde{y} | u; \theta)\|_2^2$  ▷ Compute thresholded unsupervised loss
17:  $\mathcal{L}_{\mathcal{P}} = \text{PairLoss}(\hat{\mathcal{U}}, \tau_c, \tau_s)$  ▷ Compute Pair Loss
18: return  $\mathcal{L}_{\mathcal{X}} + \lambda_{\mathcal{U}}\mathcal{L}_{\mathcal{U}} + \lambda_{\mathcal{P}}\mathcal{L}_{\mathcal{P}}$  ▷ Compute loss  $\mathcal{L}$  from  $\hat{\mathcal{X}}$  and  $\hat{\mathcal{U}}$ 

```

input instead of K severe perturbation [2] or a single perturbation [1, 26], the guessed pseudo label should be more stable. Then, we use the sharpening operation defined in [2] to increase the temperature of the label’s distribution:

$$\text{Sharpen}(p, T) := \frac{p^{\frac{1}{T}}}{\mathbf{1}^{\top} p^{\frac{1}{T}}} \quad (2)$$

As the peak of the pseudo label’s distribution is “sharpened”, the network will push this sample further away from the decision boundary. Additionally, following the practice of MixMatch [2], we use the exponential moving average of the model at each time step to guess the labels.

3.4. Loss

Our loss consists of three terms: the supervised loss $\mathcal{L}_{\mathcal{X}}$, the unsupervised loss $\mathcal{L}_{\mathcal{U}}$, and the Pair Loss $\mathcal{L}_{\mathcal{P}}$.

$$\mathcal{L} = \mathcal{L}_{\mathcal{X}} + \lambda_{\mathcal{U}}\mathcal{L}_{\mathcal{U}} + \lambda_{\mathcal{P}}\mathcal{L}_{\mathcal{P}} \quad (3)$$

$$\mathcal{L}_{\mathcal{X}} = \frac{1}{|\mathcal{X}'|} \sum_{x, y \in \mathcal{X}'} H(y, \text{P}_{\text{model}}(\tilde{y} | x; \theta)) \quad (4)$$

$$\mathcal{L}_{\mathcal{U}} = \frac{\sum_{u, q \in \hat{\mathcal{U}}} \mathbb{1}_{(\max(q) > \tau_c)} \|q - \text{P}_{\text{model}}(\tilde{y} | u; \theta)\|_2^2}{L|\hat{\mathcal{U}}|} \quad (5)$$

$\mathcal{L}_{\mathcal{X}}$ calculates the cross-entropy of weakly augmented labeled samples; $\mathcal{L}_{\mathcal{U}}$ represents the L_2 distance between strongly augmented samples and their pseudo labels, filtered by confidence threshold. Notice that $\mathcal{L}_{\mathcal{U}}$ only enforces

consistency among different perturbations of the same samples but not consistency among different samples.

3.4.1 Pair Loss

As we aim to exploit the relationship among unlabeled samples, we hereby introduce a novel loss term, Pair Loss, that allows information to propagate implicitly between different unlabeled samples. In Pair Loss, we use a high confidence pseudo label of an unlabeled point, p , as an “anchor.” All unlabeled samples whose pseudo labels are similar enough to p need to align their predictions under severe perturbation to the “anchor.” Figure 3 offers an overview of this selection process. During this process, the similarity threshold “extended” our confidence threshold in an adaptive manner, as a sample whose pseudo label confidence is below the threshold can still be selected by the loss and be pushed to a higher confidence level. Formally, we defined the Pair Loss as following:

$$\begin{aligned} \mathcal{L}_{\mathcal{P}} = & \frac{1}{\binom{K'B}{2}} \sum_{\substack{i, j \in [|\mathcal{U}'|], i \neq j \\ (v_i, q_i) = \mathcal{U}'_i \\ (v_r, q_r) = \mathcal{U}'_j}} \varphi_{\tau_c}(\max(q_i)) \\ & \cdot \varphi_{\tau_s}(f_{\text{sim}}(q_l, q_r)) \\ & \cdot f_{\text{dist}}(q_l, \text{P}_{\text{model}}(\tilde{y} | v_r; \theta)) \end{aligned} \quad (6)$$

Here, τ_c and τ_s denote the confidence threshold and similarity threshold respectively. $\varphi_t(x) = \mathbb{1}_{(x > t)}$ is a hard

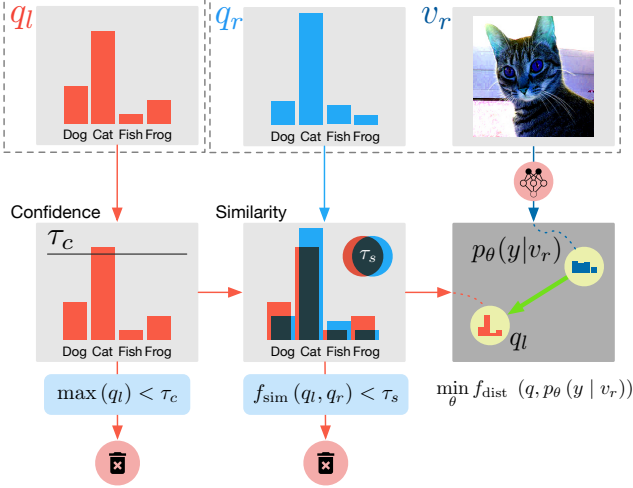


Figure 3: Pair Loss Overview. Given a pseudo label q_l (red) which is a probability vector representing the guessed class distribution, if the highest entry in q_l surpasses the confidence threshold τ_c , q_l will become an “anchor”. Then, for any pseudo label and image tuple q_r (light blue) and v_r (dark blue), if the overlapping proportion (i.e. similarity) between q_l and q_r is greater than the confidence threshold τ_s , this tuple (q_r, v_r) will contribute toward the Pair Loss by pushing model’s prediction of a strongly augmented version of v_r to the “anchor” q_l (green arrow). During this process, if either threshold can not be satisfied, q_l, q_r, v_r will be rejected.

threshold function controlled by threshold t . $f_{sim}(p, q)$ measures the similarity between two probability vectors p, q by Bhattacharyya coefficient [3]. The coefficient is bounded between $[0, 1]$, and represents the size of the overlapping portion of the two discrete distributions:

$$f_{sim}(p, q) = \sqrt{p^\top} \sqrt{q} \quad (7)$$

$f_{dist}(p, q)$ measures the distance between two probability vectors p, q . As $f_{sim}(p, q) \in [0, 1]$, we choose the distance function to be $f_{dist}(p, q) = 1 - f_{sim}(p, q)$.

Although based on analysis, we found that $\cos(\cos^{-1}(\sqrt{\tau_c}) + \cos^{-1}(\tau_s))^2$ is the infimal confidence a label need to have for it to be selected by both thresholds, such low confidence label are rarely selected in practice. Based on empirical evidence, we believe this is caused by the fact a label p that can pass through the high confidence threshold typically has a near one-hot distribution. Thus, for another label q to fall in the similarity threshold of q , it must also have relatively high confidence. Due to this property, the Pair Loss is not very sensitive to the choices of hyperparameters τ_s, τ_c , which we will show empirically in section 4.3.2.

3.4.2 Motivation for Different Loss Formulations

We follow MixMatch [2] in choosing supervised loss $\mathcal{L}_\mathcal{X}$ and unsupervised loss $\mathcal{L}_\mathcal{U}$ terms. We use the Bhattacharyya coefficient [3] in our Pair Loss because it measures the overlap between two distributions and allows a more intuitive selection of the similarity threshold τ_s . Although we believe that the Bhattacharyya coefficient [3] is more suitable than L_2 distance (or $2 - L_2$) to measure the similarity between two distributions, we keep the L_2 distance in unsupervised loss term to provide a better comparison with MixMatch [2]. Moreover, as cross-entropy measures the entropy and is asymmetric, it is not a good distance measurement between distributions. In our experiments, we observe that SimPLE with L_2 Pair Loss has 0.53% lower test accuracy than the original.

3.5. SimPLE Algorithm

By putting together all the components introduced in this section, we now present the SimPLE algorithm. During training, for a mini-batch of samples, SimPLE first augment labeled and unlabeled samples with both weak and strong augmentations. The pseudo labels of the unlabeled samples are obtained by averaging and then sharpening the models’ predictions on the weakly augmented unlabeled samples. Finally, we optimize the loss terms based on augmented samples and pseudo labels. During testing, SimPLE uses the exponential moving average of the weights of the model to make predictions, as the way done by MixMatch in [2]. Figure 2 gives an overview of SimPLE, and the complete algorithm is in algorithm 1.

CIFAR-100		
Method	10000 labels	Backbone
MixMatch*	64.01%	WRN 28-2
MixMatch Enhanced	67.12%	WRN 28-2
SimPLE	70.82%	WRN 28-2
MixMatch [†] [26]	71.69%	WRN 28-8
ReMixMatch [†] [26]	76.97%	WRN 28-8
FixMatch [26]	77.40%	WRN 28-8
SimPLE	78.11%	WRN 28-8

Table 1: CIFAR-100 Top-1 Test Accuracy. *: using our implementation. [†]: reported in FixMatch [26].

4. Experiments

Unless specified otherwise, we use Wide ResNet 28-2 [33] as our backbone and AdamW [20] with weight decay for optimization in all experiments. We also use the exponential moving average (EMA) of the network parameter of every training step for evaluation and label guessing.

Method	CIFAR-10		SVHN	
	1000 labels	4000 labels	1000 labels	4000 labels
VAT [†] [1]	81.36%	88.95%	94.02%	95.80%
MeanTeacher [†] [1]	82.68%	89.64%	96.25%	96.61%
MixMatch [2]	92.25%	93.76%	96.73%	97.11%
ReMixMatch [1]	94.27%	94.86%	97.17%	97.58%
FixMatch [26]	—	95.69%	97.64%	—
SimPLE	94.84%	94.95%	97.54%	97.31%
Fully Supervised ^{†‡}	95.75%		97.3%	

Table 2: CIFAR-10 and SVHN Top-1 Test Accuracy. All experiments use WRN 28-2. [†]: The accuracy is reported in ReMixMatch [1] and using its own implementation. [‡]: Fully supervised baseline using all the labels and simple augmentation (flip-and-crop).

Mini-ImageNet			
Method	4000 labels	Backbone	K
MixMatch*	55.47%	WRN 28-2	2
MixMatch Enhanced	60.50%	WRN 28-2	7
SimPLE	66.55%	WRN 28-2	7
MeanTeacher [†] [14]	27.49%	Resnet-18	—
Label Propagation [14]	29.71%	Resnet-18	—
SimPLE	49.39%	Resnet-18	7

Table 3: Mini-ImageNet Top-1 Test Accuracy. *: using our implementation. [†]: The score is reported in [14] and using its own implementation.

To have a fair comparison with MixMatch, we implemented an enhanced version of MixMatch by combining it with Augmentation Anchoring [1]. To report test accuracy, we take the checkpoint with the highest validation accuracy and report its test accuracy. By default, our experiments have fixed hyperparameters $\tau_c = 0.95$, $\tau_s = 0.9$ and EMA decay to 0.999.

4.1. Datasets and Evaluations

CIFAR-10: A dataset with 60K images of shape 32×32 evenly distributed across 10 classes. The training set has 50K images, and the test set contains 10K images. Our validation set size is 5000 for CIFAR-10. The results are available in table 2.

SVHN: SVHN consists of 10 classes. Its training set has 73257 images, and the test set contains 26032 images. Each image in SVHN is 32×32 . Our validation set size is 5000 for SVHN. The results are available in table 2.

CIFAR-100: Similar to CIFAR-10, CIFAR-100 also has 50K training images and 10K test images but with 100 classes. The image size is 32×32 , the same as CIFAR-10. Our validation set size is 5000 for CIFAR-100. The results are available in table 1.

Mini-ImageNet: Mini-ImageNet is first introduced in [31] for few-shot learning. The dataset contains 100 classes where each class has 600 images of size 84×84 . For SSL

evaluation, our protocol follows that of [14], in which 500 images are selected from each class to form the training set, and the leftover 100 images are used for testing. Since [14] do not specify its validation set split, we use a total of 7200 training images (72 per class) as validation set; this is of the same validation set size as [31].

DomainNet-Real [23]: DomainNet-Real has 345 categories with unbalanced numbers of images per class following a long tail distribution. We use this dataset for transfer learning experiments in section 4.3.1. For our evaluations, we resize the image to 84×84 and use 11-shot per class (a total of 3795) for the labeled training set.

4.2. Baseline Methods

We compare with the following baseline methods: FixMatch [26], MixMatch [2], ReMixMatch [1], VAT [21], MeanTeacher [29], and Label Propagation [14].

4.3. Results

For all datasets, our labeled and unlabeled set split is done by randomly sample the same number of images from all classes without replacement. In general, our hyperparameter choices follows that of MixMatch [2] and FixMatch [26].

CIFAR-100: We set the loss weight to $\lambda_{\mathcal{U}} = 150$, $\lambda_{\mathcal{P}} = 150$. As shown in table 1, we find that SimPLE has sig-

Transfer: DomainNet-Real to Mini-ImageNet		
Method	4000 labels	Convergence step
Supervised w/ EMA [§]	48.83%	4K
MixMatch* from scratch	50.31%	150K
MixMatch*	53.39%	69K
MixMatch Enhanced* from scratch	52.83%	734K
MixMatch Enhanced*	55.75%	7K
SimPLE from scratch	59.92%	338K
SimPLE	58.73%	53K

Table 4: DomainNet-Real pre-trained model transfer to Mini-ImageNet. All experiments use WRN 28-2. The model is converged when its validation accuracy reaches 95% of its highest validation accuracy. [§]: using labeled training set only. *: using our implementation.

nificant improvement on CIFAR-100. For better comparison with [26], we include experiments using the same optimizer (SGD), hyperparameters, and backbone network (WRN 28-8 with 23M parameters). With a larger backbone, our method still provides improvements over baseline methods. SimPLE is better than FixMatch by 0.7% and takes only 4.7 hours of training for convergence, while FixMatch takes about 8 hours to converge. We consider convergence is achieved when the validation accuracy reaches 95% of its highest value.

CIFAR-10, SVHN: For CIFAR-10, we set $\lambda_{\mathcal{U}} = 75$ and $\lambda_{\mathcal{P}} = 75$; we set $\lambda_{\mathcal{U}} = \lambda_{\mathcal{P}} = 250$ for SVHN. For both datasets, we use SGD with cosine learning rate decay [18] with decay rate set to $\frac{7\pi}{16}$ following that of FixMatch [26].

In table 2, we find that SimPLE is on par with ReMixMatch [1] and FixMatch [26]. ReMixMatch, FixMatch, and SimPLE are very close to the fully supervised baseline with less than 1% difference in test accuracy. SimPLE is less effective on these domains because the leftover samples are difficult ones whose pseudo labels are not similar to any of the high confidence pseudo labels. In this case, no pseudo labels can pass the two thresholds in Pair Loss and contribute to the loss. We observe that the percentage of pairs in a batch that passes both thresholds stabilizes early in the training progress (the percentage is 12% for SVHN and 10% for CIFAR-10). Thus, Pair Loss does not bring much performance gain as it does in the more complicated datasets.

Mini-ImageNet: To examine the scalability of our method, we conduct experiments on Mini-ImageNet. Mini-ImageNet is a more complex dataset because its categories and images are sampled directly from ImageNet. Although the image size is scaled down to 84×84 , it is still much more complicated than CIFAR-10, CIFAR-100, and SVHN. Therefore, Mini-ImageNet is an excellent candidate to illustrate the scalability of SimPLE.

In addition to WRN 28-2 experiments on Mini-ImageNet, we also apply the SimPLE algorithm on ResNet-

18 [12] for a fair comparison with prior works. The results are in table 3. In general, our method outperforms all other methods by a large margin on Mini-ImageNet regardless of backbones. Our method scales with the more challenging dataset.

Transfer: ImageNet-1K to DomainNet-Real		
Method	3795 labels	Convergence step
Supervised w/ EMA [§]	42.91%	4K
MixMatch*	35.34%	5K
MixMatch Enhanced*	35.16%	5K
SimPLE	50.90%	65K

Table 5: ImageNet-1K pre-trained model transfer to DomainNet-Real. All experiments use ResNet-50. The model is converged when its validation accuracy reaches 95% of its highest validation accuracy. [§]: using labeled training set only. *: using our implementation.

4.3.1 SSL for Transfer Learning Task

In real-world applications, a common scenario is where the target task is similar to existing datasets. Transfer learning is helpful in this situation if the target domain has sufficient labeled data. However, this is not guaranteed. Therefore, SSL methods need to perform well when starting from a pre-trained model on a different dataset. Another benefit of using a pre-trained model is having fast convergence, which is important for time-sensitive applications.

Since prior SSL methods often neglect this scenario, in this section, we evaluate our algorithm, MixMatch [2] and supervised baseline in the transfer setting. The supervised baseline only uses labeled training data and parameter EMA for evaluation. All transfer experiments use fixed augmentations.

Our first experiment is the adaptation from DomainNet-Real to Mini-ImageNet; the result is in table 4. We observe that the pre-trained models are on par with training from

Ablations: CIFAR-100						
Ablation	Augmentation Type	$\lambda_{\mathcal{P}}$	τ_c	τ_s	K	10000 labels
SimPLE	RandAugment	150	0.95	0.9	2	70.82%
SimPLE	RandAugment	150	0.95	0.9	<u>7</u>	73.04%
w/o Pair Loss	RandAugment	<u>0</u>	0.95	0.9	2	69.07%
w/o Pair Loss	RandAugment	<u>0</u>	0.95	0.9	<u>7</u>	69.94%
w/o RandAugment	<u>fixed</u>	150	0.95	0.9	2	67.91%
w/o RandAugment, w/o Pair Loss	<u>fixed</u>	<u>0</u>	0.95	0.9	2	67.41%
$\tau_c = 0.75$	RandAugment	150	<u>0.75</u>	0.9	2	71.96%
$\tau_s = 0.7$	RandAugment	150	0.95	<u>0.7</u>	2	70.85%
$\tau_c = 0.75, \tau_s = 0.7$	RandAugment	150	<u>0.75</u>	<u>0.7</u>	2	71.48%
$\lambda_{\mathcal{P}} = 50$	RandAugment	<u>50</u>	0.95	0.9	2	71.34%
$\lambda_{\mathcal{P}} = 250$	RandAugment	<u>250</u>	0.95	0.9	2	71.42%

Table 6: Ablation on CIFAR-100. All experiments use WRN 28-2

scratch but converge 5 \sim 100 times faster. Under transfer setting, SimPLE is 7.57% better than MixMatch and 9.9% better than the supervised baseline.

The experiment in table 5 is for transferring from ImageNet-1K [7] to DomainNet-Real. Since ImageNet-1K pre-trained ResNet-50 [12] is readily available in many machine learning libraries (e.g., PyTorch), we evaluate the performance and the convergence speed using ImageNet-1K pre-trained ResNet-50 to mimic real-world applications.

On DomainNet-Real, MixMatch is about 7% lower than the supervised baseline, while SimPLE has 8% higher accuracy than the baseline. MixMatch Enhanced, despite having Augmentation Anchoring, does not outperform MixMatch.

It is clear that SimPLE perform well in pre-trained setting and surpasses MixMatch and supervised baselines by a large margin. This behavior is consistent across datasets and network architectures. MixMatch, on the other hand, does not improve performance in the pre-trained setting.

Compared to training from scratch, the pre-trained models do not always provide performance improvements since the pre-trained models might have domain bias that is not easy to overcome. For example, in our DomainNet-Real to Mini-ImageNet experiment, the pre-trained test accuracy is slightly lower than training from scratch. However, the convergence speed is significantly faster (\sim 8 to 10 times) when starting from a pre-trained model.

4.3.2 Ablation Study over CIFAR-100

In this section, we conducted ablation studies on CIFAR-100 with WRN 28-2 to evaluate the effectiveness of different parts of our system. The results are available in table 6. We choose CIFAR-100 because it has a reasonable number of classes (reasonably complicated) and a small image size (fast enough for training).

We observe that Pair Loss significantly improves the performance. With a more diverse augmentation policy or in-

creasing the number of augmentations, the advantage of the Pair Loss is enhanced. Also, SimPLE is robust to threshold change. One possible explanation for the robustness is that since a pair must pass both thresholds to contribute to the loss, changing one of them may not significantly affect the overall number of pairs that pass both thresholds.

5. Conclusion

We proposed SimPLE, a semi-supervised learning algorithm. SimPLE improves on previous works [2, 1, 26] by considering a novel unsupervised objective, Pair Loss, which minimizes the statistical distance between high confidence pseudo labels with similarity above a certain threshold. We have conducted extensive experiments over the standard datasets and demonstrated the effectiveness of the SimPLE algorithm. Our method shows significant performance gains over previous state-of-the-art algorithms on CIFAR-100 and Mini-ImageNet [31], and is on par with the state-of-the-art methods on CIFAR-10 and SVHN. Furthermore, SimPLE also outperforms the state-of-the-art methods in the transfer learning setting, where models are initialized by the weights pre-trained on ImageNet [15], or DomainNet-Real [23].

6. Acknowledgement

This material is based on research sponsored by Air Force Research Laboratory (AFRL) under agreement number FA8750-19-1-1000. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation therein. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of Air Force Laboratory, DARPA or the U.S. Government.

References

- [1] David Berthelot, Nicholas Carlini, Ekin D. Cubuk, Alex Kurakin, Kihyuk Sohn, Han Zhang, and Colin Raffel. ReMix-Match: Semi-Supervised Learning with Distribution Matching and Augmentation Anchoring. In *International Conference on Learning Representations*, Apr. 2020. [1](#), [2](#), [3](#), [4](#), [6](#), [7](#), [8](#)
- [2] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A. Raffel. MixMatch: A Holistic Approach to Semi-Supervised Learning. In *Neural Information Processing Systems*, pages 5049–5059, Dec. 2019. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#), [11](#)
- [3] A. Bhattacharyya. On a measure of divergence between two multinomial populations. *Sankhyā: The Indian Journal of Statistics (1933-1960)*, 7(4):401–406, 1946. [5](#)
- [4] Olivier Chapelle, Bernhard Schlkopf, and Alexander Zien. Semi-Supervised Learning. *IEEE Transactions on Neural Networks*, 20(3), Mar. 2010. [1](#), [3](#)
- [5] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*, 2020. [1](#)
- [6] Ekin D. Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V. Le. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2020. [3](#), [11](#)
- [7] J. Deng, W. Dong, R. Socher, L. Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009. [8](#)
- [8] Matthijs Douze, Arthur Szlam, Bharath Hariharan, and Hervé Jégou. Low-shot learning with large-scale diffusion. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3349–3358, 2018. [1](#), [3](#)
- [9] Geoffrey French, Michal Mackiewicz, and Mark H. Fisher. Self-ensembling for visual domain adaptation. In *International Conference on Learning Representations*, Feb. 2018. [3](#)
- [10] Yves Grandvalet and Yoshua Bengio. Semi-supervised learning by entropy minimization. In *Advances in neural information processing systems*, pages 529–536, 2005. [1](#)
- [11] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum Contrast for Unsupervised Visual Representation Learning. In *Computer Vision and Pattern Recognition*, pages 9729–9738, June 2020. [1](#)
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. [7](#), [8](#)
- [13] Geoffrey E Hinton and Drew Van Camp. Keeping the neural networks simple by minimizing the description length of the weights. In *Proceedings of the sixth annual conference on Computational learning theory*, pages 5–13, 1993. [1](#)
- [14] Ahmet Iscen, Giorgos Tolias, Yannis Avrithis, and Ondrej Chum. Label propagation for deep semi-supervised learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. [1](#), [2](#), [3](#), [6](#)
- [15] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. ImageNet Classification with Deep Convolutional Neural Networks. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, pages 1097–1105. Curran Associates, Inc., 2012. [1](#), [8](#)
- [16] Samuli Laine and Timo Aila. Temporal Ensembling for Semi-Supervised Learning. *arXiv: Neural and Evolutionary Computing*, Oct. 2016. [1](#), [2](#), [3](#)
- [17] Dong-Hyun Lee. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, volume 3, 2013. [1](#), [3](#)
- [18] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016. [7](#), [11](#)
- [19] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. [1](#)
- [20] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. [5](#), [11](#)
- [21] T. Miyato, S. Maeda, M. Koyama, and S. Ishii. Virtual adversarial training: A regularization method for supervised and semi-supervised learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(8):1979–1993, 2019. [1](#), [6](#)
- [22] 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(8):1979–1993, Aug. 2019. [1](#), [2](#), [3](#)
- [23] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1406–1415, 2019. [1](#), [6](#), [8](#)
- [24] Edgar Riba, Dmytro Mishkin, Daniel Ponsa, Ethan Rublee, and Gary R. Bradski. Kornia: an open source differentiable computer vision library for pytorch. In *Winter Conference on Applications of Computer Vision*, 2020. [11](#), [12](#)
- [25] Mehdi Sajjadi, Mehran Javanmardi, and Tolga Tasdizen. Regularization with stochastic transformations and perturbations for deep semi-supervised learning. In *Neural Information Processing Systems*, pages 1171–1179, Dec. 2016. [1](#), [2](#)
- [26] Kihyuk Sohn, David Berthelot, Chun-Liang Li, Zizhao Zhang, Nicholas Carlini, Ekin D. Cubuk, Alex Kurakin, Han Zhang, and Colin Raffel. FixMatch: Simplifying Semi-Supervised Learning with Consistency and Confidence. *arXiv: Learning*, Jan. 2020. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#), [11](#)
- [27] Martin Szummer and Tommi Jaakkola. Partially labeled classification with markov random walks. In *Advances in neural information processing systems*, pages 945–952, 2002. [1](#), [3](#)
- [28] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve

- semi-supervised deep learning results. In *International Conference on Learning Representations*, Jan. 2017. [1](#), [3](#)
- [29] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30, pages 1195–1204. Curran Associates, Inc., 2017. [6](#)
- [30] Vikas Verma, Alex Lamb, Juho Kannala, Yoshua Bengio, and David Lopez-Paz. Interpolation Consistency Training for Semi-supervised Learning. In *International Joint Conference on Artificial Intelligence*, pages 3635–3641, Aug. 2019. [1](#)
- [31] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, koray kavukcuoglu, and Daan Wierstra. Matching networks for one shot learning. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29, pages 3630–3638. Curran Associates, Inc., 2016. [1](#), [6](#), [8](#)
- [32] Qizhe Xie, Zihang Dai, Eduard Hovy, Thang Luong, and Quoc Le. Unsupervised data augmentation for consistency training. *Advances in Neural Information Processing Systems*, 33, 2020. [1](#), [3](#)
- [33] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In Edwin R. Hancock Richard C. Wilson and William A. P. Smith, editors, *Proceedings of the British Machine Vision Conference (BMVC)*, pages 87.1–87.12. BMVA Press, September 2016. [5](#)
- [34] Guodong Zhang, Chaoqi Wang, Bowen Xu, and Roger Grosse. Three mechanisms of weight decay regularization. *arXiv preprint arXiv:1810.12281*, 2018. [1](#)
- [35] Hongyi Zhang, Moustapha Cisse, Yann N. Dauphin, and David Lopez-Paz. mixup: Beyond Empirical Risk Minimization. In *International Conference on Learning Representations*, Feb. 2018. [1](#), [11](#)
- [36] Liheng Zhang and Guo-Jun Qi. WCP: Worst-Case Perturbations for Semi-Supervised Deep Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. [1](#), [3](#)
- [37] Dengyong Zhou, Olivier Bousquet, Thomas N. Lal, Jason Weston, and Bernhard Schölkopf. Learning with Local and Global Consistency. In *Neural Information Processing Systems*, Dec. 2003. [1](#), [3](#)