
Dash: Semi-Supervised Learning with Dynamic Thresholding

Yi Xu ¹ Lei Shang¹ Jinxing Ye¹ Qi Qian¹ Yu-Feng Li ² Baigui Sun¹ Hao Li¹ Rong Jin¹

Abstract

While semi-supervised learning (SSL) has received tremendous attentions in many machine learning tasks due to its successful use of unlabeled data, existing SSL algorithms use either all unlabeled examples or the unlabeled examples with a fixed high-confidence prediction during the training progress. However, it is possible that too many correct/wrong pseudo labeled examples are eliminated/selected. In this work we develop a simple yet powerful framework, whose key idea is to select a subset of training examples from the unlabeled data when performing existing SSL methods so that only the unlabeled examples with pseudo labels related to the labeled data will be used to train models. The selection is performed at each updating iteration by only keeping the examples whose losses are smaller than a given threshold that is dynamically adjusted through the iteration. Our proposed approach, Dash, enjoys its adaptivity in terms of unlabeled data selection and its theoretical guarantee. Specifically, we theoretically establish the convergence rate of Dash from the view of non-convex optimization. Finally, we empirically demonstrate the effectiveness of the proposed method in comparison with state-of-the-art over benchmarks.

1. Introduction

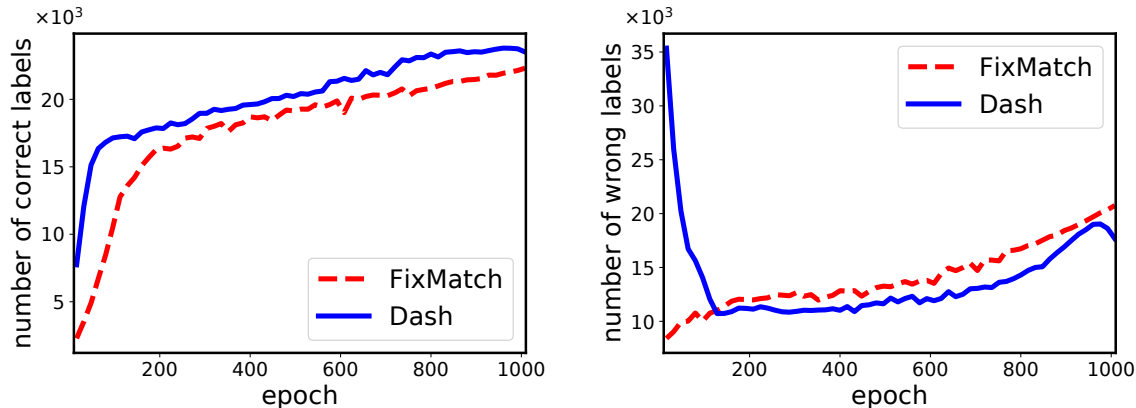
In spite of successful use in a variety of classification and regression tasks, supervised learning requires large amount of *labeled* training data. In many machine learning applications, labeled data can be significantly more costly, time-consuming and difficult to obtain than the *unlabeled* data (Zhu, 2005), since they usually require experienced human labors from experts (e.g., a doctor in detection of covid-19 by using X-ray images). Typically, only a small

amount of labeled data is available, but there is a huge amount of data without label. This is one of key hurdles in the development and deployment of machine learning models.

Semi-supervised learning (SSL) is designed to improve learning performance by leveraging an abundance of unlabeled data along with limited labeled data (Chapelle et al., 2006). In much recent work, SSL can be categorized into several main classes in terms of the use of unlabeled data: consistency regularization, pseudo labeling, generic regularization (e.g., large margin regularization, Laplacian regularization, etc (Chapelle et al., 2006)), and their combinations. With the image data augmentation technique, consistency regularization uses unlabeled data (Baird, 1992; Schmidhuber, 2015) based on the condition that the model predictions between different perturbed versions of the same image are similar. Another line of work is to produce artificial label for unlabeled data based on prediction model and add them to the training data set. With different approaches of artificial label production, various of SSL methods have been proposed in the literature including self-training (Yarowsky, 1995; Lee, 2013; Rosenberg et al., 2005; Sajjadi et al., 2016; Laine & Aila, 2017; Xie et al., 2020b) and co-training (Blum & Mitchell, 1998; Zhou & Li, 2005; Sindhwani & Rosenberg, 2008; Wang et al., 2008; Yu et al., 2008; Wang & Zhou, 2010; Chen et al., 2011). Due to its capability to handle both labeled data and unlabeled data, SSL has been widely studied in diverse machine learning tasks such as image classification (Sajjadi et al., 2016; Laine & Aila, 2017; Tarvainen & Valpola, 2017; Xie et al., 2020a; Berthelot et al., 2019b;a), natural language processing (Turian et al., 2010), speech recognition (Yu et al., 2010), and object detection (Misra et al., 2015).

Numerous empirical evidences show that unlabeled data in SSL can help to improve the learning performance (Berthelot et al., 2019b;a; Sohn et al., 2020), however, a series of theoretical studies (Ben-David et al., 2008; Singh et al., 2009; Li & Zhou, 2011b; Balcan & Blum, 2005) have demonstrated that this success is highly relying on a necessary condition that labeled data and unlabeled data with pseudo label come from the same distribution during the training process (Zhu, 2005; Van Engelen & Hoos, 2020). Unfortunately, it has been shown that this condition does not always hold in real applications and thus it could hurt

¹Machine Intelligence Technology, Alibaba Group. ²National Key Laboratory for Novel Software Technology, Nanjing University. Correspondence to: Yi Xu <yixu@alibaba-inc.com>, Yu-Feng Li <liyf@lamda.nju.edu.cn>.



(a) Number of selected unlabeled examples with correct pseudo labels (b) Number of selected unlabeled examples with wrong pseudo labels

Figure 1. An example of experimental results on Wide ResNet-28-8 for CIFAR-100 with 400 labeled images illustrates the reason of dynamically selecting unlabeled data to train learning models. Pseudo labels are generated based on the prediction models. FixMatch selects unlabeled example if its confidence prediction is greater than 0.95, while the proposed Dash algorithm selects unlabeled example based on a dynamic threshold through optimization iterations. (a) The proposed Dash selects more examples with correct pseudo labels than that of FixMatch. (b) The proposed Dash maintains much more examples with wrong pseudo labels at the beginning but it will drop off more examples with wrong pseudo labels after several epochs, comparing to FixMatch.

the performance (Li et al., 2017; Oliver et al., 2018). For example, the pseudo label of an unlabeled example generated by conventional SSL methods during the training progress is not correct (Hataya & Nakayama, 2019; Li et al., 2020). In this case, the degradation of model performance has been observed when using unlabeled data compared to the simple supervised learning model not using any unlabeled data at all (Chapelle et al., 2006; Oliver et al., 2018). Thus, not all unlabeled data are needed in SSL.

To improve SSL performance, multiple studies (Guo et al., 2020; Ren et al., 2020) examined the strategies of weighting different training unlabeled examples by solving a bi-level optimization problem. It is also a popular idea to select a subset of training examples from unlabeled examples for SSL. For example, FixMatch (Sohn et al., 2020) uses the unlabeled examples with a fixed high-confidence prediction (e.g., 0.95) in classification tasks using cross entropy loss. However, the fixed threshold may lead to eliminate too many unlabeled examples with correct pseudo labels (see Figure 1 (a)) and may lead to select too many unlabeled examples with wrong pseudo labels (see Figure 1 (b)). That is to say, the fixed threshold is possible not good enough during the training progress and thus it could degrade the overall performance.

Unlike the previous work, we aim to the proposed approach enjoys its adaptivity in terms of unlabeled data selection and its theoretical guarantee. This inspires us to consider answering the following question in this study.

Can we design a provable SSL algorithm that selects unlabeled data with dynamic thresholding?

To this end, we propose a generic SSL algorithm with **Dynamic Thresholding (Dash)** that can dynamically select unlabeled data during the training process. Specifically, Dash firstly runs over labeled data and obtains a threshold for unlabeled data selection. It then selects the unlabeled data whose loss values are smaller than the threshold to the training data-set. The value of threshold is gradually decreased over the optimization iterations. It can be integrated with existing SSL methods like FixMatch (Sohn et al., 2020). From the view of optimization, we show that eventually the proposed Dash can non-asymptotically converge with theoretical guarantee. Empirical evaluations on image benchmarks validate the effectiveness of Dash comparing with the state-of-the-art SSL algorithms.

2. Related Work

There has been growing interest in semi-supervised learning for training machine learning and deep learning (Flach, 2012; Goodfellow et al., 2016). A number of SSL methods have been studied by leveraging the structure of unlabeled data including consistency regularization (Bachman et al., 2014; Sajjadi et al., 2016; Laine & Aila, 2017; Tarvainen & Valpola, 2017; Miyato et al., 2018; Xie et al., 2020a), entropy minimization (Grandvalet & Bengio, 2005; Lee, 2013), and other interesting approaches (Berthelot et al., 2019b;a). In addition, several studies on SSL have been

proposed to keep SSL performing safe when using unlabeled data, which is known as safe SSL (Li & Zhou, 2015). An non-exhaustive list of those studies include (Cozman et al., 2003; Singh et al., 2009; Li & Zhou, 2011a; Balsubramani & Freund, 2015; Loog, 2015; Li et al., 2017; Krijthe & Loog, 2017; Li et al., 2021; Mey & Loog, 2019; Guo et al., 2020). For example, Ren et al. (2020) proposed a new SSL framework that uses an individual weight for each unlabeled example, and it updates the individual weights and models iteratively by solving a bi-level optimization problem approximately. In this paper, we mainly focus on improved deep SSL methods with the use of unlabeled data selection. Comprehensive surveys on SSL methods could be refer to (Zhu, 2005; Chapelle et al., 2006; Zhu & Goldberg, 2009; Hady & Schwenker, 2013; Van Engelen & Hoos, 2020).

The use of unlabeled data selection by a threshold is not new in the literature of SSL. As a simple yet widely used heuristic algorithm, pseudo-labeling (Lee, 2013) (a.k.a. self-training (McLachlan, 1975; Yarowsky, 1995; Rosenberg et al., 2005; Sajjadi et al., 2016; Laine & Aila, 2017; Xie et al., 2020b)) uses the prediction model itself to generate pseudo labels for unlabeled images. Then the unlabeled images whose corresponding pseudo label’s highest class probability is larger than a predefined threshold will be used for the training. Nowadays, pseudo-labeling has become an important component of many modern SSL methods (Xie et al., 2020a; Sohn et al., 2020).

With the use of weak and strong data augmentations, several recent works such as UDA (Xie et al., 2020a), ReMix-Match (Berthelot et al., 2019a) and FixMatch (Sohn et al., 2020) have been proposed in image classification problems. Generally, they use a weakly-augmented¹ unlabeled image to generate a pseudo label and enforce consistency against strongly-augmented² version of the same image. In particular, UDA and FixMatch use a fixed threshold to retain the unlabeled example whose highest probability in the predicted class distribution for the pseudo label is higher than the threshold. For example, UDA sets this threshold to be 0.8 for CIFAR-10 and SVHN, and FixMatch sets the threshold to be 0.95 for all data-sets. To encourage the model to generate high-confidence predictions, UDA and ReMix-Match sharpen the guessed label distribution by adjusting its temperature and then re-normalize the distribution. Sohn et al. (2020) have shown that the sharpening and thresholding pseudo-labeling have a similar effect.

By contrast, the proposed Dash method selects a subset

¹Both UDA and ReMixMatch use crop and flip as “weak” augmentation while FixMatch uses flip and shift.

²For “strong” augmentation, UDA uses RandAugment (Cubuk et al., 2020), ReMixMatch uses CTAugment (Cubuk et al., 2019), and FixMatch uses both.

of unlabeled data to be used in training models by a data-dependent dynamic threshold, and its theoretical convergence guarantee is established for stochastic gradient descent under the non-convex setting, which is applicable to deep learning.

3. Preliminary and Background

3.1. Problem Setting

We study the task of learning a model to map an input $\mathbf{x} \in \mathcal{X} \subseteq \mathbb{R}^d$ onto a label $\mathbf{y} \in \mathcal{Y}$. In many machine learning applications, \mathbf{x} refers to the feature and $\mathbf{y} \in \mathcal{Y}$ refers to the label for classification or regression. For simplicity, let ξ denote the input-label pair (\mathbf{x}, \mathbf{y}) , i.e. $\xi := (\mathbf{x}, \mathbf{y})$. We denote by \mathcal{P} the underlying distribution of data pair ξ , then $\xi \sim \mathcal{P}$. The goal is to learn a model $\mathbf{w} \in \mathbb{R}^d$ via minimizing an optimization problem whose objective function $F(\mathbf{w})$ is the expectation of random loss function $f(\mathbf{w}; \xi)$:

$$\min_{\mathbf{w} \in \mathbb{R}^d} F(\mathbf{w}) := \mathbb{E}_{\xi \sim \mathcal{P}} [f(\mathbf{w}; \xi)], \quad (1)$$

where $\mathbb{E}_{\xi}[\cdot]$ is an expectation taking over random variable $\xi \sim \mathcal{P}$. The optimization problem (1) covers most machine learning and deep learning applications. In this paper, we consider the classification problem with K -classes, whose loss function is the cross-entropy loss given by

$$\begin{aligned} f(\mathbf{w}; \xi_i) &= H(\mathbf{y}_i, \mathbf{p}(\mathbf{w}; \mathbf{x}_i)) \\ &:= \sum_{k=1}^K -y_{i,k} \log \left(\frac{\exp(p_k(\mathbf{w}; \mathbf{x}_i))}{\sum_{j=1}^K \exp(p_j(\mathbf{w}; \mathbf{x}_i))} \right), \end{aligned} \quad (2)$$

where $\mathbf{p}(\mathbf{w}; \mathbf{x})$ is the prediction function and $H(\mathbf{q}, \mathbf{p})$ is the cross-entropy between \mathbf{q} and \mathbf{p} . In this paper, we do not require the function $f(\mathbf{w}; \xi)$ to be convex in terms of \mathbf{w} , which is applicable to various deep learning tasks.

In SSL, it consists of labeled examples and unlabeled examples. Let

$$\mathbf{D}_l := \{(\mathbf{x}_i, \mathbf{y}_i), i = 1, \dots, N_l\} \quad (3)$$

be the labeled training data. Given unlabeled training examples $\{\mathbf{x}_i^u, i = 1, 2, \dots, N_u\}$, one can generate pseudo label $\hat{\mathbf{y}}_i^u$ based on the predictions of a supervised model on labeled data. Different SSL methods such as pseudo-labeling (Lee, 2013), Adversarial Training (Miyato et al., 2018), UDA (Xie et al., 2020a), and FixMatch (Sohn et al., 2020) have been proposed to generate pseudo labels. We denote by

$$\mathbf{D}_u := \{(\mathbf{x}_i^u, \hat{\mathbf{y}}_i^u), i = 1, \dots, N_u\} \quad (4)$$

the unlabeled data, where $\hat{\mathbf{y}}_i^u \in \mathcal{Y}$ is the pseudo label. Although it contains pseudo label, we still call \mathbf{D}_u unlabeled data for simplicity in our analysis. Usually, the number

of unlabeled examples is much larger than the number of labeled examples, i.e., $N_u \gg N_l$. Finally, the training data consists of labeled data \mathbf{D}_l and unlabeled data with pseudo label \mathbf{D}_u , and thus the training loss of an SSL algorithm usually contains supervised loss F_s and unsupervised loss F_u with a weight $\lambda_u > 0$: $F_s + \lambda_u F_u$, where F_s is constructed on \mathbf{D}_l and F_u is constructed on \mathbf{D}_u . In image classification problems, F_s is just the standard cross-entropy loss:

$$F_s(\mathbf{w}) := \frac{1}{N_l} \sum_{i=1}^{N_l} f(\mathbf{w}; \xi_i), \quad (5)$$

where $\xi_i \in \mathbf{D}_l$ and f is defined in (2). Thus, different constructions of the unsupervised loss F_u lead to different SSL methods. Typically, there are two ways of constructing F_u : one is to use pseudo labels to formulate a ‘‘supervised’’ loss such as cross-entropy loss (e.g., FixMatch), and another one is to optimize a regularization that does not depend on labels such as consistency regularization (e.g., Π -Model). Next, we will introduce a recent SSL work to interpret how to generate pseudo labels and construct unsupervised loss F_u .

3.2. FixMatch: An SSL Algorithm with Fixed Thresholding

Due to its simplicity yet empirical success, we select FixMatch (Sohn et al., 2020) as an SSL example in this subsection. Besides, we consider FixMatch as a warm-up of the proposed algorithm, since FixMatch uses a fixed threshold to retain unlabeled examples and it will be used as a pipeline in the proposed algorithm.

The key idea of FixMatch is to use a separate weak and strong augmentation when generating models predicted class distribution and one-hot label in unsupervised loss. Specifically, based on a supervised model \mathbf{w} and a weak augmentation α , FixMatch predict the class distribution

$$\mathbf{h}_i = \mathbf{p}(\mathbf{w}, \alpha(\mathbf{x}_i^u)) \quad (6)$$

for a weakly-augmented version of a unlabeled image \mathbf{x}_i^u , where $\mathbf{p}(\mathbf{w}, \mathbf{x})$ is the prediction function. Then it creates a pseudo label by

$$\hat{\mathbf{y}}_i^u = \arg \max(\mathbf{h}_i). \quad (7)$$

Following by (Sohn et al., 2020), the $\arg \max$ applied to a probability distribution produces a ‘‘one-hot’’ probability distribution. To construct the unsupervised loss, it computes the model prediction for a strong augmentation \mathcal{T} of the same unlabeled image \mathbf{x}_i^u :

$$\mathbf{p}(\mathbf{w}, \mathcal{T}(\mathbf{x}_i^u)). \quad (8)$$

The unsupervised loss is defined as the cross-entropy between $\hat{\mathbf{y}}_i^u$ and \mathbf{p}_i :

$$H(\hat{\mathbf{y}}_i^u, \mathbf{p}(\mathbf{w}, \mathcal{T}(\mathbf{x}_i^u))). \quad (9)$$

Eventually, FixMatch only uses the unlabeled examples with a high-confidence prediction by selecting based on a **fixed threshold** $\tau = 0.95$. Therefore, in FixMatch the cross-entropy loss with pseudo-label and confidence for unlabeled data is given by

$$F_u(\mathbf{w}) = \frac{1}{N_u} \sum_{i=1}^{N_u} I(\max(\mathbf{h}_i) \geq \tau) H(\hat{\mathbf{y}}_i^u, \mathbf{p}(\mathbf{w}, \mathcal{T}(\mathbf{x}_i^u))), \quad (10)$$

where $I(\cdot)$ is an indicator function.

As we discussed in introduction, this fixed threshold may lead to eliminate/select too many unlabeled examples with correct/wrong pseudo labels (see Figure 1), which eventually could drop off overall performance. It is a natural choice: the threshold is not fixed across the optimization iterations. Thus, in the next section, we are going to propose a new SSL scheme having a dynamic threshold.

4. Dash: An SSL Algorithm with Dynamic Thresholding

Before introducing the proposed method, we would like to point out the importance of unlabeled data selection in SSL from the theoretical view of optimization. Classical SSL methods (Zhu, 2005; Chapelle et al., 2006; Zhu & Goldberg, 2009; Hady & Schwenker, 2013; Van Engelen & Hoos, 2020) assume that labeled data and unlabeled data are from the same distribution. That is to say, $\xi \sim \mathcal{P}$ holds for $\xi \in \mathbf{D}_l \cup \mathbf{D}_u$. Then SSL methods aim to solve the optimization problem (1) by using a standard stochastic optimization algorithm like mini-batch stochastic gradient descent (SGD). Specifically, at iteration t , mini-batch SGD updates intermediate solutions by

$$\mathbf{w}_{t+1} = \mathbf{w}_t - \frac{\eta}{m} \sum_{i=1}^m \nabla f(\mathbf{w}_t; \xi_{t,i}), \quad (11)$$

where m is the mini-batch size, $\xi_{t,i}$ is sampled from training data $\mathbf{D}_l \cup \mathbf{D}_u$, $\nabla f(\mathbf{w}; \xi)$ is the gradient of $f(\mathbf{w}; \xi)$ in terms of \mathbf{w} . In this situation, the theoretical convergence guarantee of SSL algorithms can be simply established under mild assumptions on objective function $f(\mathbf{w}; \xi)$ such as smoothness and bounded variance (Ghadimi et al., 2016). If the labeled data and unlabeled data are not from the same distribution such as some of pseudo labels are not correct, classical SSL methods with standard stochastic optimization algorithm may lead to the performance drops (Chapelle et al., 2006; Oliver et al., 2018). Besides, the theoretical guarantee of the optimization algorithm for this case is not clear. This inspires us to design a new algorithm to overcome this issue. To this end, we proposed an SSL method that can dynamically select unlabeled examples during the training progress.

Algorithm 1 Dash: Semi-Supervised Learning with Dynamic Thresholding

Input: learning rate η_0 and mini-batch size m_0 for stage one, learning rate η and parameter m of mini-batch size for stage two, two parameters $C > 1$ and $\gamma > 1$ for computing threshold, and violation probability δ .

// Warm-up Stage: run SGD in T_0 iterations.

Initialization: $\mathbf{u}_0 = \mathbf{w}_0$

for $t = 0, 1, \dots, T_0 - 1$ **do**

 Sample m_0 examples $\xi_{t,i}$ ($i = 1, \dots, m_0$) from \mathbf{D}_l ,

$\mathbf{u}_{t+1} = \mathbf{u}_t - \eta_0 \tilde{\mathbf{g}}_t$ where $\tilde{\mathbf{g}}_t = \frac{1}{m_0} \sum_{i=1}^{m_0} \nabla f_s(\mathbf{u}_t; \xi_{t,i})$

end for

// Selection Stage: run SGD in T iterations.

Initialization: $\mathbf{w}_1 = \mathbf{u}_{T_0}$.

Compute the value of $\hat{\rho}$ as in (16). // In practice, $\hat{\rho}$ can be obtained as in (17).

for $t = 1, \dots, T$ **do**

 1) Sample $n_t = m\gamma^{t-1}$ examples from \mathbf{D}_u , where the pseudo labels in \mathbf{D}_u are generated by FixMatch

 2) Set the threshold $\rho_t = C\gamma^{-(t-1)}\hat{\rho}$.

 3) Compute truncated stochastic gradient \mathbf{g}_t as (18).

 4) Update solution by SGD using stochastic gradient \mathbf{g}_t and learning rate η : $\mathbf{w}_{t+1} = \mathbf{w}_t - \eta\mathbf{g}_t$.

end for

Output: \mathbf{w}_{T+1}

First, let us define the loss function for the proposed method. Same as FixMatch, the supervised loss $F_s(\mathbf{w})$ for the proposed method is the standard cross-entropy loss on labeled data \mathbf{D}_l :

$$F_s(\mathbf{w}) := \frac{1}{N_l} \sum_{i=1}^{N_l} f_s(\mathbf{w}; \xi_i), \quad (12)$$

where $\xi_i = (\mathbf{x}_i, \mathbf{y}_i)$ is sampled from \mathbf{D}_l , $f_s(\mathbf{w}; \xi_i) = H(\mathbf{y}_i, \mathbf{p}(\mathbf{w}; \alpha(\mathbf{x}_i)))$, and $\alpha(\mathbf{x})$ is the weakly-augmented version of x . Since the new dynamic threshold is not fixed, we let it rely on the optimization iteration t and it is denoted by ρ_t . Then the unsupervised loss is given by

$$F_u(\mathbf{w}) = \frac{1}{N_u} \sum_{i=1}^{N_u} I(f_u(\mathbf{w}; \xi_i^u) \leq \rho_t) f_u(\mathbf{w}; \xi_i^u), \quad (13)$$

where $\xi_i^u = (\mathbf{x}_i^u, \hat{\mathbf{y}}_i^u)$ is sampled from \mathbf{D}_u , $f_u(\mathbf{w}; \xi_i^u) = H(\hat{\mathbf{y}}_i^u, \mathbf{p}(\mathbf{w}; \mathcal{T}(\mathbf{x}_i^u)))$, $\mathcal{T}(x)$ is the strongly-augmented version of x , and the pseudo label $\hat{\mathbf{y}}_i^u$ is generated based on (6) and (7) using prediction model \mathbf{w} and an unlabeled image \mathbf{x}_i^u . The unsupervised loss (13) shows that the Dash will retain the unlabeled example whose loss is smaller than the threshold ρ_t . If we rewrite the indicator function $I(\max(\mathbf{h}_i) \geq \tau)$ in (10) to an equivalent expression

$$I(-\log(\max(\mathbf{h}_i)) \leq -\log(\tau)), \quad (14)$$

we can consider $-\log(\max(\mathbf{h}_i))$ as a cross-entropy loss for one-hot label. Roughly speaking, FixMatch retains the unlabeled images with the loss $-\log(\max(\mathbf{h}_i))$ smaller than $-\log(0.95) \approx 0.0513$. It is worth nothing that the loss $-\log(\max(\mathbf{h}_i))$ contains the information of weakly-augmented images, while the loss $f_u(\mathbf{w}; \xi_i^u)$ in (13) includes the information of both weakly-augmented and

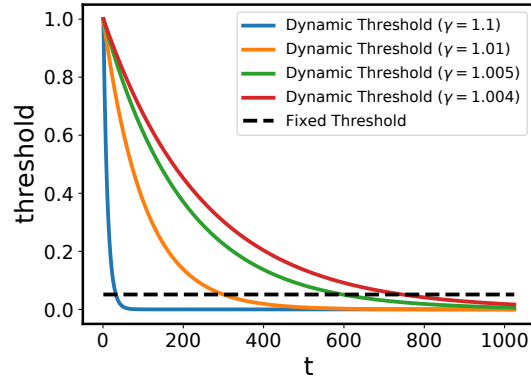


Figure 2. Comparison of fixed threshold and dynamic threshold. Fixed threshold is in the scale of negative log: $-\log(0.95)$, dynamic threshold $\rho_t = 1.0001\gamma^{-(t-1)}\hat{\rho}$.

strongly-augmented images, meaning that the proposed method considers the entire loss function.

We then need to construct ρ_t . Intuitively, with the increase of the optimization iteration t , the loss function would decrease in general, so that ρ_t is also required to decrease. Mathematically, we set the **dynamic threshold** ρ_t as a decreasing function of t , which is given by

$$\rho_t := C\gamma^{-(t-1)}\hat{\rho}, \quad (15)$$

where $C > 1, \gamma > 1$ are two constants. For example, we set $C = 1.0001$ in our experiments and thus at first iteration (i.e., $t = 1$) the unlabeled examples whose loss values are smaller than $\rho_t = 1.0001 \times \hat{\rho}$ will be used in the

training. Figure 2 shows a comparison of fixed threshold used in FixMatch and dynamic threshold ρ_t in (15) with $C = 1.0001$, $\hat{\rho} = 1$ and different γ , where the threshold in FixMatch is in the scale of negative log. It seems our thresholding strategy matches the curve of training loss in many real applications (e.g., see Figure 1 (a) of (Zhang et al., 2017)).

Next, it is important to estimate the value of $\hat{\rho}$. In theory, it can be estimated by

$$\hat{\rho} = \max \left\{ a, \frac{4G^2(1 + \delta b_0 m)}{\delta \mu a_0 m} \right\}, \quad (16)$$

where contains several parameters related to the property of considered problem (1) whose detailed definitions can be found in Theorem 1. Please note that the estimation of $\hat{\rho}$ in (16) is for the use of convergence analysis only. In practice, we can use the following averaged loss from the training set \mathbf{D}_l as the approximate $\hat{\rho}$:

$$\hat{\rho} \approx \frac{1}{|\mathbf{D}_l|} \sum_{\xi_i \in \mathbf{D}_l} f(\mathbf{w}_1; \xi_i), \quad (17)$$

where $|\mathbf{D}_l|$ is the number of examples in \mathbf{D}_l , and \mathbf{w}_1 can be learned on \mathbf{D}_l . We can see from (17) with (13) and (15) that the unlabeled examples whose losses are smaller than the averaged loss of labeled examples will be maintained during the training process.

Finally, it is ready to describe the proposed Dash algorithm in details that contains two stages: warm-up stage and selection stage. In the warm-up stage, it runs SGD to train a model over labeled data \mathbf{D}_l in certain steps. Not only for warm-up, this stage is also used for estimating $\hat{\rho}$ in (17). In the selection stage, we conduct SGD against \mathbf{D}_u using \mathbf{w}_1 as the initial solution. At each iteration t , we sample $n_t = m\gamma^{t-1}$ training examples from \mathbf{D}_u , where $m > 1$ is a parameter defined in (24). We compute the stochastic gradients according to (13):

$$\mathbf{g}_t = \frac{\sum_{i=1}^{n_t} I(f_u(\mathbf{w}_t; \xi_{t,i}^u) \leq \rho_t) \nabla f_u(\mathbf{w}_t; \xi_{t,i}^u)}{\sum_{i=1}^{n_t} I(f_u(\mathbf{w}_t; \xi_{t,i}^u) \leq \rho_t)}. \quad (18)$$

Since N_l is small, in practice we can also construct the stochastic gradient by using all labeled data as

$$\mathbf{g}_t = \frac{\sum_{i=1}^{n_t - N_l} I(f_u(\mathbf{w}_t; \xi_{t,i}^u) \leq \rho_t) \nabla f_u(\mathbf{w}_t; \xi_{t,i}^u)}{N_l + \sum_{i=1}^{n_t - N_l} I(f_u(\mathbf{w}_t; \xi_{t,i}^u) \leq \rho_t)} + \frac{\sum_{i=1}^{N_l} \nabla f_s(\mathbf{w}_t; \xi_{t,i})}{N_l + \sum_{i=1}^{n_t - N_l} I(f_u(\mathbf{w}_t; \xi_{t,i}^u) \leq \rho_t)}, \quad (19)$$

where $\xi_{t,i}^u = (x_{t,i}^u, y_{t,i}^u) \in \mathbf{D}_u$ and $\xi_{t,i} = (x_{t,i}, y_{t,i}) \in \mathbf{D}_l$. The solution is then updated by mini-batch SGD, whose update step is given by

$$\mathbf{w}_{t+1} = \mathbf{w}_t - \eta \mathbf{g}_t. \quad (20)$$

The detailed updating steps of the proposed algorithm are presented in Algorithm 1, which called SSL with **Dynamic Thresholding** (Dash).

5. Convergence Result

To establish the convergence result of the proposed Dash algorithm, we need to give some preliminaries. Recall that the training examples for the labeled data \mathbf{D}_l follow the distribution \mathcal{P} , and we aim to minimize the optimization problem (1). For the examples coming from the unlabeled data \mathbf{D}_u , suppose it is a mixture of two distributions, \mathcal{P} and \mathcal{Q} . More specifically, with a probability q , we will sample an example from \mathcal{P} and with a probability $1 - q$ sample from \mathcal{Q} :

$$\xi \sim q\mathcal{P} + (1 - q)\mathcal{Q}, \text{ where } \xi \in \mathbf{D}_u, q \in (0, 1). \quad (21)$$

We define the objective function $B(\mathbf{w})$ as the expected loss for distribution \mathcal{Q} , i.e.

$$B(\mathbf{w}) := \mathbb{E}_{\xi \sim \mathcal{Q}} [f(\mathbf{w}; \xi)]. \quad (22)$$

For the simplicity of convergence analysis, we do not consider the weak and strong augmentations, i.e., let $f_s = f_u = f$ in (12) and (13). Without loss of generality, we assume that our loss function is non-negative and is bounded by 1, i.e. $f(\mathbf{w}; \xi) \in [0, 1]$ for any \mathbf{w} and ξ . Then by (1) and (22), we have $F(\mathbf{w}) \in [0, 1]$ and $B(\mathbf{w}) \in [0, 1]$. In order to differentiate the two distribution, we follow the idea of Tsybakov noisy condition (Mammen et al., 1999; Tsybakov, 2004), and assume, for any solution \mathbf{w} , if $F(\mathbf{w}) \leq a$, then

$$\mathbb{E}_{\xi \sim \mathcal{Q}} [I_{\{\xi: f(\mathbf{w}; \xi) \leq F(\mathbf{w})\}}(\xi)] \leq bA^\theta(\mathbf{w}), \quad (23)$$

where $I_S(\xi)$ is an indicator function, $\theta \geq 1$, and b is constant.

Finally, we made a few more assumptions that are commonly used in the studies of non-convex optimization (e.g., deep learning) (Ghadimi & Lan, 2013; Yuan et al., 2019). Throughout this paper, we also make the assumptions on the problem (1) as follows.

Assumption 1. Assume the following conditions hold:

(i) The stochastic gradient $\nabla f(\mathbf{w}; \xi)$ is unbiased, i.e.,

$$\mathbb{E}_{\xi \sim \mathcal{P}} [\nabla f(\mathbf{w}; \xi)] = \nabla F(\mathbf{w}),$$

and there exists a constant $G > 0$, such that

$$\|\nabla f(\mathbf{w}; \xi)\| \leq G.$$

(ii) $F(\mathbf{w})$ is smooth with a L -Lipchitz continuous gradient, i.e., it is differentiable and there exists a constant $L > 0$ such that

$$\|\nabla F(\mathbf{w}) - \nabla F(\mathbf{u})\| \leq L\|\mathbf{w} - \mathbf{u}\|, \forall \mathbf{w}, \mathbf{u} \in \mathbb{R}^d.$$

Table 1. Comparison of top-1 testing error rates for different methods using Wide ResNet-28-2 for CIFAR-10, Wide ResNet-28-8 for CIFAR-100 (in %, mean \pm standard deviation).

Algorithm	CIFAR-10			CIFAR-100		
	40 labels	250 labels	4000 labels	400 labels	2500 labels	10000 labels
PI-Model	-	54.26 \pm 3.97	14.01 \pm 0.38	-	57.25 \pm 0.48	37.88 \pm 0.11
Pseudo-Labeling	-	49.78 \pm 0.43	16.09 \pm 0.28	-	57.38 \pm 0.46	36.21 \pm 0.19
Mean Teacher	-	32.32 \pm 2.30	9.19 \pm 0.19	-	53.91 \pm 0.57	35.83 \pm 0.24
MixMatch	47.54 \pm 11.50	11.05 \pm 0.86	6.42 \pm 0.10	67.61 \pm 1.32	39.94 \pm 0.37	28.31 \pm 0.33
UDA	29.05 \pm 5.93	8.82 \pm 1.08	4.88 \pm 0.18	59.28 \pm 0.88	33.13 \pm 0.22	24.50 \pm 0.25
ReMixMatch	19.10 \pm 9.64	5.44 \pm 0.05	4.72 \pm 0.13	44.28 \pm 2.06	27.43 \pm 0.31	23.03 \pm 0.56
RYS (UDA)	-	5.53 \pm 0.17	4.75 \pm 0.28	-	-	-
RYS (FixMatch)	-	5.05 \pm 0.12	4.35 \pm 0.06	-	-	-
FixMatch (CTA)	11.39 \pm 3.35	5.07 \pm 0.33	4.31 \pm 0.15	49.95 \pm 3.01	28.64 \pm 0.24	23.18 \pm 0.11
Dash (CTA, ours)	9.16 \pm 4.31	4.78 \pm 0.12	4.13 \pm 0.06	44.83 \pm 1.36	27.85 \pm 0.19	22.77 \pm 0.21
FixMatch (RA)	13.81 \pm 3.37	5.07 \pm 0.65	4.26 \pm 0.05	48.85 \pm 1.75	28.29 \pm 0.11	22.60 \pm 0.12
Dash (RA, ours)	13.22 \pm 3.75	4.56 \pm 0.13	4.08 \pm 0.06	44.76 \pm 0.96	27.18 \pm 0.21	21.97 \pm 0.14

Assumption 1 (i) assures that the stochastic gradient of the objective function is unbiased and the gradient of $f(\mathbf{w}; \xi)$ in terms of \mathbf{w} is upper bounded. Assumption 1 (ii) says the objective function is L -smooth, and it has an equivalent expression which is $\forall \mathbf{w}, \mathbf{u} \in \mathbb{R}^d$,

$$F(\mathbf{w}) - F(\mathbf{u}) \leq \langle \nabla F(\mathbf{u}), \mathbf{w} - \mathbf{u} \rangle + \frac{L}{2} \|\mathbf{w} - \mathbf{u}\|^2.$$

We now introduce an important property regarding $F(\mathbf{w})$, i.e. the Polyak-Łojasiewicz (PL) condition (Polyak, 1963) of $F(\mathbf{w})$.

Assumption 2. There exists $\mu > 0$ such that

$$2\mu(F(\mathbf{w}) - F(\mathbf{w}_*)) \leq \|\nabla F(\mathbf{w})\|^2, \forall \mathbf{w} \in \mathbb{R}^d.$$

This PL property has been theoretically and empirically observed in training deep neural networks (Allen-Zhu et al., 2019; Yuan et al., 2019). This condition is widely used to establish convergence in the literature of non-convex optimization, please see (Yuan et al., 2019; Wang et al., 2019; Karimi et al., 2016; Li & Li, 2018; Charles & Papailiopoulos, 2018) and references therein.

Now, we are ready to provide the theoretical result for Dash. Without loss of generality, let $F(\mathbf{w}_*) = 0$ in the analysis. Please note that this is a common property observed in training deep neural networks (Zhang et al., 2017; Allen-Zhu et al., 2019; Du et al., 2019; Arora et al., 2019; Chizat et al., 2019; Hastie et al., 2019; Yun et al., 2019). The following theorem states the convergence guarantee of the proposed Dash algorithm. We include its proof in the Appendix.

Theorem 1. Under Assumptions 1 and 2, suppose that $C > 1$ and $F(\mathbf{w}_*) = 0$, for any $\delta \in (0, 1)$, $\eta_0 L \leq 1$, $\eta L \leq 1$, let

$$T_0 = \frac{\log(2F(\mathbf{w}_0)/a)}{\log(1/(1-\eta_0\mu))}, m_0 = \frac{4G^2}{\delta\mu a},$$

$$m = \left[\max \left(\sqrt{\frac{\log(2/\delta)}{q^2}}, \sqrt{\frac{\log(2/\delta)}{(1-q)^2}}, \sqrt{\frac{\log(2/\delta)}{q(1-C^{-1})^2}} \right) \right], \quad (24)$$

$$\hat{\rho} = \max \left\{ a, \frac{4G^2(1 + \delta b_0 m)}{\delta\mu a_0 m} \right\} \quad (25)$$

in Algorithm 1, then with a probability $1 - (4T + 1)\delta$, we have $F(\mathbf{w}_{T+1}) \leq \hat{\rho}\gamma^{-T}$.

Remark. We can see from the above result that one can set the iteration number T to be large enough to ensure the convergence of Dash. Specifically, in order to have an ϵ optimization error, one can set $T = \log(\hat{\rho}/\epsilon)/\log(\gamma)$, then $F(\mathbf{w}_{T+1}) \leq \epsilon$. The total sample complexity of Dash is $T_0 m_0 + \sum_{t=1}^T m\gamma^{t-1} \leq T_0 m_0 + \frac{m\gamma^T}{\gamma-1} = T_0 m_0 + \frac{m\hat{\rho}}{\epsilon(\gamma-1)} = O(1/\epsilon)$. This rate matches the result of supervised learning in (Karimi et al., 2016) when analyzing the standard SGD under Assumptions 1 and 2.

6. Experiments

In this section, we present some experimental results for image classification tasks. To evaluate the efficacy of Dash, we compare it with several state-of-the-art (SOTA) baselines on several standard SSL image classification benchmarks including CIFAR-10, CIFAR-100 (Krizhevsky & Hinton, 2009), SVHN (Netzer et al., 2011), and STL-10 (Coates et al., 2011) data sets. Specifically, SOTA baselines are MixMatch (Berthelot et al., 2019b), UDA (Xie et al., 2020a), ReMixMatch (Berthelot et al., 2019a), FixMatch (Sohn

Table 2. Comparison of top-1 testing error rates for different methods using Wide ResNet-28-2 for SVHN and Wide ResNet-37-2 for STL-10 (in %, mean \pm standard deviation).

Algorithm	SVHN			STL-10
	40 labels	250 labels	1000 labels	1000 labels
II-Model	-	18.96 \pm 1.92	7.54 \pm 0.36	26.23 \pm 0.82
Pseudo-Labeling	-	20.21 \pm 1.09	9.94 \pm 0.61	27.99 \pm 0.83
Mean Teacher	-	3.57 \pm 0.11	3.42 \pm 0.07	21.43 \pm 2.39
MixMatch	42.55 \pm 14.53	3.98 \pm 0.23	3.50 \pm 0.28	10.41 \pm 0.61
UDA	52.63 \pm 20.51	5.69 \pm 2.76	2.46 \pm 0.24	7.66 \pm 0.56
ReMixMatch	3.34 \pm 0.20	2.92 \pm 0.48	2.65 \pm 0.08	5.23 \pm 0.45
RYS (UDA)	-	2.45 \pm 0.08	2.32 \pm 0.06	-
RYS (FixMatch)	-	2.63 \pm 0.23	2.34 \pm 0.15	-
FixMatch (CTA)	7.65 \pm 7.65	2.64 \pm 0.64	2.36 \pm 0.19	5.17 \pm 0.63
Dash (CTA, ours)	3.14 \pm 1.60	2.38 \pm 0.29	2.14 \pm 0.09	3.96\pm0.25
FixMatch (RA)	3.96 \pm 2.17	2.48 \pm 0.38	2.28 \pm 0.11	7.98 \pm 1.50
Dash (RA, ours)	3.03\pm1.59	2.17\pm0.10	2.03\pm0.06	7.26 \pm 0.40

et al., 2020) and the algorithm RYS from (Ren et al., 2020)³. Besides, other SSL algorithms such as II-Model (Rasmus et al., 2015), Pseudo-Labeling (Lee, 2013) and Mean Teacher (Tarvainen & Valpola, 2017) are included in the comparison.

6.1. Data-sets

The original CIFAR data-sets have 50,000 training images and 10,000 testing images of 32×32 resolutions, and CIFAR-10 has 10 classes containing 6,000 images each, while CIFAR-100 has 100 classes containing 600 images each. The original SVHN data-set has 73,257 digits for training and 26,032 digits for testing, and the total number of classes is 10. The original STL-10 data set has 5,000 labeled images from 10 classes and 100,000 unlabeled images, which contains out-of-distribution unlabeled images.

Following by (Sohn et al., 2020), we train ten benchmarks with different settings: CIFAR-10 with 4, 25, or 400 labels per class, CIFAR-100 with 4, 25, or 100 labels per class, SVHN with 4, 25, or 100 labels per class, and the STL-10 data set. For example, the benchmark CIFAR-10 with 4 labels per class means that there are 40 labeled images in CIFAR-10 and the remaining images are unlabeled, and then we denote this data set by CIFAR-10 with 40 labels. For fair comparison, same sets of labeled images from CIFAR, SVHN and STL-10 were used for the proposed Dash method and other baselines in all experiments.

³Since the authors did not name their algorithm, we use RYS to denote their algorithm for simplicity, where RYS is the combination of initials for last names of the authors.

6.2. Models and Hyper-parameters

We use the Wide ResNet-28-2 model (Zagoruyko & Komodakis, 2016) as the backbone for CIFAR-10 and SVHN, Wide ResNet-28-8 for CIFAR-100, and Wide ResNet-37-2 for STL-10. In the proposed Dash, we use FixMatch⁴ as our pipeline to generate pseudo labels and to construct supervised and unsupervised losses. We employ CTAugment (CTA) (Cubuk et al., 2019) and RandAugment (RA) (Cubuk et al., 2020) for the strong augmentation scheme following by (Sohn et al., 2020). Similar to (Sohn et al., 2020), we use the same training protocol such as optimizer, learning rate schedule, data preprocessing, random seeds, and so on.

The total number of training epochs is set to be 1024 and the mini-batch size is fixed as 64. For the value of weight decay, we use 5×10^{-4} for CIFAR-10, SVHN and STL-10, 1×10^{-3} for CIAR-100. The SGD with momentum parameter of 0.9 is employed as the optimizer. The cosine learning rate decay schedule (Loshchilov & Hutter, 2017) is used as (Sohn et al., 2020). The initial learning rate is set to be 0.06 for all data-sets. We use (18) to compute stochastic gradients.

At the first 10 epochs, we do not implement the selection scheme and thus the algorithm uses all selected training examples, meaning that the threshold is infinite, i.e., $\rho_t = \infty$. After that, we use the threshold to select unlabeled examples, and we choose $\gamma = 1.27$ in ρ_t to reduce the dynamic threshold until its value to be 0.05. That is to say, in practice we give a minimal value of dynamic threshold, which is 0.05⁵. We fix the constant C as 1.0001 and estimate the

⁴In our experiments, the FixMatch codebase is used: <https://github.com/google-research/fixmatch>

⁵In practice, we use $\rho_t = \max\{\rho_t, 0.05\}$.

Table 3. Comparison of top-1 testing error rates for different values of γ on CIFAR-10 (in %).

γ	1.01	1.1	1.2	1.3
250 labels	4.85	4.76	4.99	4.82
4000 labels	4.39	4.28	4.11	4.31

value of $\hat{\rho}$ by using (17). We decay the dynamic threshold every 9 epochs. We use the predicted label distribution as soft label during the training and it is sharpened by adjusting its temperature of 0.5, which is similar to MixMatch. Once the dynamic threshold is reduced to 0.05, we turn it to one-hot label in the training since the largest label probability is close to 1.

6.3. Results

We report the top-1 testing error rates of the proposed Dash along with other baselines for CIFAR in Table 1 and for SVHN and STL-10 in Table 2, where all the results of baselines are from (Sohn et al., 2020) except that the results of RYS are from (Ren et al., 2020). All top-1 testing error rates are averaged over 5 independent random trails with their standard deviations using the same random seeds as baselines used.

We can see from the results that the proposed Dash method has the best performance on CIFAR-10, SVHN and STL-10. For CIFAR-100, the proposed Dash is comparable to ReMixMatch, where ReMixMatch performs a bit better on 400 labels and Dash using RA is a bit better on 2500 labels and 10000 labels. This reason is that the proposed Dash uses FixMatch as its pipeline, and ReMixMatch uses distribution alignment (DA) to encourage the model to predict balanced class distribution (the class distribution of CIFAR-100 is balanced), while FixMatch and Dash do not use DA. We further conduct Dash with DA technique on CIFAR-100 with 400 labels, and the top-1 testing error rate is 43.31%, which is better than ReMixMatch (44.28%). We also find that Dash performs well on the data set with out-of-distribution unlabeled images, i.e., STL-10. The result in Table 2 shows that Dash with CTA has the SOTA performance of 3.96% on top-1 testing error rate.

Besides, the proposed Dash can always outperform FixMatch, showing that the use of dynamic threshold is important to the overall performance. We find the proposed Dash has large improvement when the labeled examples is small (the data-sets with 4 labels per classes), comparing to FixMatch. By using CTA, on CIFAR-100 with 400 labels, on CIFAR-10 with 40 labels, and on SVHN with 40 labels, the proposed Dash method outperforms FixMatch result more than 19%, 10%, and 58% in the terms of top-1 testing error rate, respectively. While by using RA, the corresponding improved rates are 4%, 8%, and 23% respectively.

Table 4. Comparison of top-1 testing error rates for PL and Dash with PL on CIFAR-10 (in %).

Algorithm	PL	Dash-PL
250 labels	49.78	46.90
4000 labels	16.09	15.59

These results reveal that the dynamic unlabeled example selection is an important term in SSL when the labeled data is small.

6.4. Ablation study

In this subsection, we provide two ablation studies using data sets CIFAR-10 with 250 labels and CIFAR-10 with 4000 labels. The first one is to use different γ in the dynamic threshold, and the second one is to change FixMatch to Pseudo-Labeling as the pseudo label generator in Dash.

Different values of γ . Since γ is a key component of the dynamic threshold, we conduct an ablation study on different values of γ in Dash. For simplicity, we only implement the CTA case. We try four different values of $\gamma \in \{1.01, 1.1, 1.2, 1.3\}$ and summarize the results in Table 3. Comparing these results with that in Table 1, we will find that the choice of $\gamma = 1.27$ in the previous subsection is not the best one. The results also show that Dash is not so sensitive to γ in a certain range.

Dash with Pseudo-Labeling. Since Dash can be integrated with many existing SSL methods, we use Pseudo-Labeling (PL) (Lee, 2013) as the pipeline to generate pseudo labels in Dash. The results are listed in Table 4, showing that Dash can improve PL, especially when the labeled images is small.

7. Conclusion

We propose a method Dash that dynamically selects unlabeled data examples to train learning models. Its selection strategy keeps the unlabeled data whose loss value does not exceed a dynamic threshold at each optimization step. The proposed Dash method is a generic scheme that can be easily integrated with existing SSL methods. We demonstrate the use of dynamically selecting unlabeled data can help to the performance of existing SSL method FixMatch in the semi-supervised image classification benchmarks, indicating the importance of dynamic threshold in SSL. The theoretical analysis shows the convergence guarantee of the proposed Dash under the non-convex optimization setting.

Acknowledgements

The authors would like to thank the anonymous reviewers for their helpful comments.

References

- Allen-Zhu, Z., Li, Y., and Song, Z. A convergence theory for deep learning via over-parameterization. In *International Conference on Machine Learning*, pp. 242–252, 2019.
- Arora, S., Du, S., Hu, W., Li, Z., and Wang, R. Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. In *International Conference on Machine Learning*, pp. 322–332, 2019.
- Bachman, P., Alsharif, O., and Precup, D. Learning with pseudo-ensembles. In *Advances in Neural Information Processing Systems*, pp. 3365–3373, 2014.
- Baird, H. S. Document image defect models. In *Structured Document Image Analysis*, pp. 546–556. Springer, 1992.
- Balcan, M.-F. and Blum, A. A pac-style model for learning from labeled and unlabeled data. In *International Conference on Computational Learning Theory*, pp. 111–126. Springer, 2005.
- Balsubramani, A. and Freund, Y. Optimally combining classifiers using unlabeled data. In *Conference on Learning Theory*, pp. 211–225, 2015.
- Ben-David, S., Lu, T., and Pál, D. Does unlabeled data provably help? worst-case analysis of the sample complexity of semi-supervised learning. In *Conference on Learning Theory*, pp. 33–44, 2008.
- Berthelot, D., Carlini, N., Cubuk, E. D., Kurakin, A., Sohn, K., Zhang, H., and Raffel, C. Remixmatch: Semi-supervised learning with distribution matching and augmentation anchoring. In *International Conference on Learning Representations*, 2019a.
- Berthelot, D., Carlini, N., Goodfellow, I., Papernot, N., Oliver, A., and Raffel, C. A. Mixmatch: A holistic approach to semi-supervised learning. In *Advances in Neural Information Processing Systems*, pp. 5050–5060, 2019b.
- Blum, A. and Mitchell, T. Combining labeled and unlabeled data with co-training. In *Proceedings of Annual Conference on Computational Learning Theory*, pp. 92–100, 1998.
- Chapelle, O., Scholkopf, B., and Zien, A. *Semi-Supervised Learning (1st edition)*. Cambridge: The MIT Press, 2006.
- Charles, Z. and Papailiopoulos, D. Stability and generalization of learning algorithms that converge to global optima. In *International Conference on Machine Learning*, pp. 745–754, 2018.
- Chen, M., Weinberger, K. Q., and Chen, Y. Automatic feature decomposition for single view co-training. In *International Conference on Machine Learning*, pp. 953–960, 2011.
- Chizat, L., Oyallon, E., and Bach, F. On lazy training in differentiable programming. In *Advances in Neural Information Processing Systems*, pp. 2937–2947, 2019.
- Coates, A., Ng, A., and Lee, H. An analysis of single-layer networks in unsupervised feature learning. In *International Conference on Artificial Intelligence and Statistics*, pp. 215–223, 2011.
- Cozman, F. G., Cohen, I., and Cirelo, M. C. Semi-supervised learning of mixture models. In *International Conference on Machine Learning*, pp. 99–106, 2003.
- Cubuk, E. D., Zoph, B., Mane, D., Vasudevan, V., and Le, Q. V. Autoaugment: Learning augmentation strategies from data. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 113–123, 2019.
- Cubuk, E. D., Zoph, B., Shlens, J., and Le, Q. V. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 702–703, 2020.
- Du, S., Lee, J., Li, H., Wang, L., and Zhai, X. Gradient descent finds global minima of deep neural networks. In *International Conference on Machine Learning*, pp. 1675–1685, 2019.
- Flach, P. *Machine learning: the art and science of algorithms that make sense of data*. Cambridge University Press, 2012.
- Ghadimi, S. and Lan, G. Stochastic first-and zeroth-order methods for nonconvex stochastic programming. *SIAM Journal on Optimization*, 23(4):2341–2368, 2013.
- Ghadimi, S., Lan, G., and Zhang, H. Mini-batch stochastic approximation methods for nonconvex stochastic composite optimization. *Mathematical Programming*, 155(1-2):267–305, 2016.
- Goodfellow, I., Bengio, Y., and Courville, A. *Deep learning*. MIT press, 2016.
- Grandvalet, Y. and Bengio, Y. Semi-supervised learning by entropy minimization. In *Advances in Neural Information Processing Systems*, pp. 529–536, 2005.
- Guo, L.-Z., Zhang, Z.-Y., Jiang, Y., Li, Y.-F., and Zhou, Z.-H. Safe deep semi-supervised learning for unseen-class unlabeled data. In *International Conference on Machine Learning*, pp. 3897–3906, 2020.

- Hady, M. F. A. and Schwenker, F. Semi-supervised learning. In *Handbook on Neural Information Processing*, pp. 215–239. Springer, 2013.
- Hastie, T., Montanari, A., Rosset, S., and Tibshirani, R. J. Surprises in high-dimensional ridgeless least squares interpolation. *arXiv preprint arXiv:1903.08560*, 2019.
- Hataya, R. and Nakayama, H. Unifying semi-supervised and robust learning by mixup. 2019.
- Karimi, H., Nutini, J., and Schmidt, M. Linear convergence of gradient and proximal-gradient methods under the polyak-lojasiewicz condition. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 795–811. Springer, 2016.
- Krijthe, J. H. and Loog, M. Projected estimators for robust semi-supervised classification. *Machine Learning*, 106(7):993–1008, 2017.
- Krizhevsky, A. and Hinton, G. Learning multiple layers of features from tiny images. *Master's thesis, Technical report, University of Tronto*, 2009.
- Laine, S. and Aila, T. Temporal ensembling for semi-supervised learning. In *International Conference on Learning Representations*, 2017.
- Lee, D.-H. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *In ICML Workshop on Challenges in Representation Learning*, 2013.
- Li, J., Socher, R., and Hoi, S. C. Dividemix: Learning with noisy labels as semi-supervised learning. *International Conference on Learning Representations*, 2020.
- Li, Y.-F. and Zhou, Z.-H. Improving semi-supervised support vector machines through unlabeled instances selection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 386–391, 2011a.
- Li, Y.-F. and Zhou, Z.-H. Towards making unlabeled data never hurt. In *International Conference on Machine Learning*, pp. 1081–1088, 2011b.
- Li, Y.-F. and Zhou, Z.-H. Towards making unlabeled data never hurt. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(1):175–188, 2015.
- Li, Y.-F., Zha, H.-W., and Zhou, Z.-H. Learning safe prediction for semi-supervised regression. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, pp. 2217–2223, 2017.
- Li, Y.-F., Guo, L.-Z., and Zhou, Z.-H. Towards safe weakly supervised learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(1):334–346, 2021.
- Li, Z. and Li, J. A simple proximal stochastic gradient method for nonsmooth nonconvex optimization. In *Advances in Neural Information Processing Systems*, pp. 5564–5574, 2018.
- Loog, M. Contrastive pessimistic likelihood estimation for semi-supervised classification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(3):462–475, 2015.
- Loshchilov, I. and Hutter, F. Sgdr: Stochastic gradient descent with warm restarts. *International Conference on Learning Representations*, 2017.
- Mammen, E., Tsybakov, A. B., et al. Smooth discrimination analysis. *The Annals of Statistics*, 27(6):1808–1829, 1999.
- McLachlan, G. J. Iterative reclassification procedure for constructing an asymptotically optimal rule of allocation in discriminant analysis. *Journal of the American Statistical Association*, 70(350):365–369, 1975.
- Mey, A. and Loog, M. Improvability through semi-supervised learning: a survey of theoretical results. *arXiv preprint arXiv:1908.09574*, 2019.
- Misra, I., Shrivastava, A., and Hebert, M. Watch and learn: Semi-supervised learning for object detectors from video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3593–3602, 2015.
- Miyato, T., Maeda, S.-i., Koyama, M., and Ishii, S. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(8):1979–1993, 2018.
- Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B., and Ng, A. Y. Reading digits in natural images with unsupervised feature learning. 2011.
- Oliver, A., Odena, A., Raffel, C. A., Cubuk, E. D., and Goodfellow, I. Realistic evaluation of deep semi-supervised learning algorithms. In *Advances in Neural Information Processing Systems*, pp. 3235–3246, 2018.
- Polyak, B. T. Gradient methods for minimizing functionals. *Zhurnal Vychislitel'noi Matematiki i Matematicheskoi Fiziki*, 3(4):643–653, 1963.
- Rasmus, A., Valpola, H., Honkala, M., Berglund, M., and Raiko, T. Semi-supervised learning with ladder networks. In *Neural Information Processing Systems*, pp. 3546–3554, 2015.
- Ren, Z., Yeh, R., and Schwing, A. Not all unlabeled data are equal: learning to weight data in semi-supervised

- learning. *Advances in Neural Information Processing Systems*, 33, 2020.
- Rosenberg, C., Hebert, M., and Schneiderman, H. Semi-supervised self-training of object detection models. In *Proceedings of the 7th IEEE Workshops on Application of Computer Vision*, pp. 29–36, 2005.
- Sajjadi, M., Javanmardi, M., and Tasdizen, T. Regularization with stochastic transformations and perturbations for deep semi-supervised learning. In *Advances in Neural Information Processing Systems*, pp. 1163–1171, 2016.
- Schmidhuber, J. Deep learning in neural networks: An overview. *Neural networks*, 61:85–117, 2015.
- Sindhwani, V. and Rosenberg, D. S. An rkhs for multi-view learning and manifold co-regularization. In *International Conference on Machine Learning*, pp. 976–983, 2008.
- Singh, A., Nowak, R., and Zhu, J. Unlabeled data: Now it helps, now it doesn't. In *Advances in Neural Information Processing Systems*, pp. 1513–1520, 2009.
- Sohn, K., Berthelot, D., Li, C.-L., Zhang, Z., Carlini, N., Cubuk, E. D., Kurakin, A., Zhang, H., and Raffel, C. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. In *Advances in Neural Information Processing Systems*, pp. 596–608, 2020.
- Tarvainen, A. and Valpola, H. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *Advances in Neural Information Processing Systems*, pp. 1195–1204, 2017.
- Tsybakov, A. B. Optimal aggregation of classifiers in statistical learning. *The Annals of Statistics*, 32(1):135–166, 2004.
- Turian, J., Ratinov, L., and Bengio, Y. Word representations: a simple and general method for semi-supervised learning. In *Proceedings of Annual Meeting of the Association for Computational Linguistics*, pp. 384–394, 2010.
- Van Engelen, J. E. and Hoos, H. H. A survey on semi-supervised learning. *Machine Learning*, 109(2):373–440, 2020.
- Wang, J., Luo, S.-w., and Zeng, X.-h. A random subspace method for co-training. In *Proceedings of IEEE International Joint Conference on Neural Networks*, pp. 195–200, 2008.
- Wang, W. and Zhou, Z.-H. A new analysis of co-training. In *International Conference on Machine Learning*, pp. 1135–1142, 2010.
- Wang, Z., Ji, K., Zhou, Y., Liang, Y., and Tarokh, V. Spiderboost and momentum: Faster variance reduction algorithms. In *Advances in Neural Information Processing Systems*, pp. 2403–2413, 2019.
- Xie, Q., Dai, Z., Hovy, E., Luong, T., and Le, Q. Unsupervised data augmentation for consistency training. In *Advances in Neural Information Processing Systems*, pp. 6256–6268, 2020a.
- Xie, Q., Luong, M.-T., Hovy, E., and Le, Q. V. Self-training with noisy student improves imagenet classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 10687–10698, 2020b.
- Yarowsky, D. Unsupervised word sense disambiguation rivaling supervised methods. In *Proceedings of Annual Meeting of the Association for Computational Linguistics*, pp. 189–196, 1995.
- Yu, D., Varadarajan, B., Deng, L., and Acero, A. Active learning and semi-supervised learning for speech recognition: A unified framework using the global entropy reduction maximization criterion. *Computer Speech & Language*, 24(3):433–444, 2010.
- Yu, S., Krishnapuram, B., Steck, H., Rao, R. B., and Rosales, R. Bayesian co-training. In *Advances in Neural Information Processing Systems*, pp. 1665–1672, 2008.
- Yuan, Z., Yan, Y., Jin, R., and Yang, T. Stagewise training accelerates convergence of testing error over sgd. In *Advances in Neural Information Processing Systems*, pp. 2604–2614, 2019.
- Yun, C., Sra, S., and Jadbabaie, A. Small relu networks are powerful memorizers: a tight analysis of memorization capacity. In *Advances in Neural Information Processing Systems*, pp. 15558–15569, 2019.
- Zagoruyko, S. and Komodakis, N. Wide residual networks. *British Machine Vision Conference*, 2016.
- Zhang, C., Bengio, S., Hardt, M., Recht, B., and Vinyals, O. Understanding deep learning requires rethinking generalization. *International Conference on Learning Representations*, 2017.
- Zhou, Z.-H. and Li, M. Tri-training: Exploiting unlabeled data using three classifiers. *IEEE Transactions on knowledge and Data Engineering*, 17(11):1529–1541, 2005.
- Zhu, X. and Goldberg, A. B. Introduction to semi-supervised learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 3(1):1–130, 2009.
- Zhu, X. J. Semi-supervised learning literature survey. Technical report, Technical Report, University of Wisconsin-Madison Department of Computer Sciences, 2005.