
Semi-Supervised Learning with Meta-Gradient

Xin-Yu Zhang*

TKLNDST, CS, Nankai University

Taihong Xiao*

University of California, Merced

Haolin Jia

Tongji University

Ming-Ming Cheng

TKLNDST, CS, Nankai University

Ming-Hsuan Yang

University of California, Merced

Abstract

In this work, we propose a simple yet effective meta-learning algorithm in semi-supervised learning. We notice that most existing consistency-based approaches suffer from overfitting and limited model generalization ability, especially when training with only a small number of labeled data. To alleviate this issue, we propose a learn-to-generalize regularization term by utilizing the label information and optimize the problem in a meta-learning fashion. Specifically, we seek the pseudo labels of the unlabeled data so that the model can generalize well on the labeled data, which is formulated as a nested optimization problem. We address this problem using the meta-gradient that bridges between the pseudo label and the regularization term. In addition, we introduce a simple first-order approximation to avoid computing higher-order derivatives and provide theoretic convergence analysis. Extensive evaluations on the SVHN, CIFAR, and ImageNet datasets demonstrate that the proposed algorithm performs favorably against state-of-the-art methods.

1 Introduction

The rapid advances of deep neural networks can be in part attributed to the availability of large-scale datasets with extensive annotations, which require considerable human labor. However, a typical real-world scenario is that only a small amount of data has the corresponding annotations while the majority of training examples are

unlabeled. Numerous semi-supervised learning (SSL) methods have since been developed in which the unlabeled data are exploited to facilitate generalizing the learned models.

Existing SSL algorithms include co-training (de Sa, 1994; Blum and Mitchell, 1998), label propagation (Szummer and Jaakkola, 2002), graph regularization (Blum and Chawla, 2001), and the consistency-enforcing approaches (Rasmus et al., 2015; Sajjadi et al., 2016; Laine and Aila, 2017; Miyato et al., 2018; Tarvainen and Valpola, 2017; Athiwaratkun et al., 2019; Yu et al., 2019). Notably, the consistency-based approaches treat SSL as a generalization problem and enforce consistent predictions against small perturbations of the input data or model parameters. The basic assumption is that similar training examples are more likely to belong to the same category, so that the predictions of the network in multiple passes should be consistent (Sajjadi et al., 2016). As such, the consistency-based approaches are essentially designing pseudo labels from the predictions of the same input signals, while the incorrect predictions may misguide the training process. To improve the quality of the pseudo labels, two orthogonal directions, *i.e.*, dedicating to carefully designed perturbations (Miyato et al., 2018; Yu et al., 2019) and delving into better role models (Laine and Aila, 2017; Tarvainen and Valpola, 2017), have been introduced. Aside from the aforementioned approaches, Athiwaratkun et al. (2019) analyze the training dynamics of the models trained with the consistency regularization and propose a variant of the stochastic weight averaging (SWA) (Izmailov et al., 2018), *i.e.*, fastSWA, to improve performance and accelerate convergence.

From the above, we can see that the label information in most SSL methods is commonly used for pseudo labeling or label propagation especially for the unlabeled data apart from providing the ground truth for the labeled data. Either pseudo labeling or label propagation is based on the assumption of inner structure of the

Proceedings of the 24th International Conference on Artificial Intelligence and Statistics (AISTATS) 2021, San Diego, California, USA. PMLR: Volume 130. Copyright 2021 by the author(s). * denotes equal contribution.

data manifold. For example, the regularization terms in II-Model (Laine and Aila, 2017) and VAT (Miyato et al., 2018) are assuming that the decision boundary should be flat near the data. However, these consistency regularization terms are quite empirical and generic, where the given label information is not explicitly exploited in the consistency regularization. As a result, the model may overfit to the limited labeled data at a bad local minimum, which results in limited model generalization ability.

To alleviate this issue, we propose a meta-learning algorithm in which the pseudo labels are designed explicitly for generalization on the task of interest. Specifically, We regard the labeled data as a validation set, and generate pseudo-labels of the unlabeled data by minimizing the validation loss. Thereby, the label information in the validation loss can influence the pseudo-labels via the meta-gradients to influence the end-to-end training of the model parameters. In this way, the validation loss term, as the proposed new regularization term in this paper, explicitly includes the label information and drives the model towards better generalization ability, as indicated by the decrease of validation loss (see Theorem 1). We further introduce a simple first-order approximation to alleviate the issue of computing higher-order derivatives, and an improved training protocol to address the sample bias problem. Under mild conditions, the proposed meta-learning algorithm enjoys a convergence rate of $O(1/\epsilon^2)$, which is identical to that of the stochastic gradient descent (SGD) algorithm. Extensive experimental results demonstrate that our method performs favorably against the state-of-the-art approaches on the SVHN (Netzer et al., 2011), CIFAR (Krizhevsky et al., 2009), and ImageNet (Rusakovsky et al., 2015) datasets, and the ablation studies validate the effectiveness of each component of our approach.

2 Related Work

Consistency-based Semi-Supervised Learning.

The consistency regularization term measures the discrepancy between the predictions and the pseudo labels, which are typically generated by the same data with small perturbations on the input signals (Sajjadi et al., 2016; Miyato et al., 2018) or model parameters (Tarvainen and Valpola, 2017). For the II-model (Sajjadi et al., 2016), the predictions and pseudo labels are generated by the same model with different data augmentations through different forward passes¹. Laine and Aila (2017) propose the temporal ensemble

¹Due to the existence of randomized operations, such as dropout (Hinton et al., 2012) and shake-shake regularization (Gastaldi, 2017), the outputs of the same input signal may be different in multiple forward passes.

approach to improve the quality of pseudo labels by keeping an exponential moving average (EMA) of the history predictions of each training example. However, the scalability of this method is limited since the memory footprint grows linearly with the number of training examples. Instead, Tarvainen and Valpola (2017) present the mean teacher method to track the model parameters and generate pseudo labels using the teacher model parameterized by the EMA of the history model parameters. On the other hand, Miyato et al. (2018) present the virtual adversarial training scheme to focus on disturbing the input data in an adversarial direction, and Yu et al. (2019) decouple the adversarial direction into the tangent and normal directions of the embedded training data manifold. With the dedicated perturbation directions, the robustness of the learned model can be significantly improved. Aside from the above-mentioned methods, Athiwaratkun et al. (2019) introduce the fastSWA method to average the model parameters along the training timeline.

Recently, there is a line of research focusing on mitigating the self-supervised algorithms to SSL, in which heavy data augmentation is used to promote the performance (Berthelot et al., 2019, 2020; Sohn et al., 2020). For example, MixMatch (Berthelot et al., 2019) enforces consistent predictions between the differently-augmented training data, and ReMixMatch (Berthelot et al., 2020) further introduces an automatic augmentation strategy to effectively generate the strongly-augmented data, for which the consistency term is enforced. Since our approach adopts a simple *mixup* augmentation, the direct comparison with these methods is unfair. The integration of our method with more sophisticated data augmentation is left for future work.

We notice that existing SSL methods do not exploit the label information when computing the consistency regularization, which leads to limited model generalization ability. To alleviate this, we relate the consistency loss with the label information by unfolding and differentiating through one optimization step. In this way, the update of the pseudo labels is guided by the meta-gradients of the labeled data loss, and the consistency loss is designed to improve the generalization ability specially for the underlying task. We also experimentally verify the important role of the label information in the effectiveness of consistency regularization.

Optimization-based Meta-Learning.

Numerous optimization-based meta-learning algorithms (Finn et al., 2017; Grant et al., 2018; Finn et al., 2018; Yoon et al., 2018; Rajeswaran et al., 2019) have been developed in recent years. Notably, Finn et al. (2017) formulate the meta-learning problem in a nested optimization format, where the inner loop imitates the process

of adaptation, while the outer loop focuses on optimizing the meta-objective. The inner-optimization is further replaced by a single SGD step so that the meta-objective can be optimized in an end-to-end manner. Thanks to its simplicity and effectiveness, optimization-based meta-learning algorithms have been applied to a wide range of vision and learning problems including example re-weighting (Ren et al., 2018), neural architecture search (Liu et al., 2018), and unrolled generative models (Metz et al., 2017). In this work, we develop a meta-learning algorithm in the semi-supervised settings and demonstrate the potential of meta-learning for these tasks. In addition, we present the theoretical convergence analysis of the proposed algorithm.

A concurrent work (Wang et al., 2020) proposes a similar meta-learning algorithm for SSL. Different from ours, however, their method learns a weight for each unlabeled data, and the pseudo-labels are inherited from previous consistency-based methods (*e.g.*, Tarvainen and Valpola (2017)). We instead infer the pseudo-labels within the meta-learning framework, making our method relatively independent of the previous consistency-based counterparts.

3 Proposed Algorithm

In a typical semi-supervised setting, we are given a few labeled data $\mathcal{D}^l = \{(\mathbf{x}_k^l, \mathbf{y}_k) : k = 1, \dots, N^l\}$ and a large amount of unlabeled data $\mathcal{D}^u = \{\mathbf{x}_i^u : i = 1, \dots, N^u\}$, where $N^l \ll N^u$. The goal is to train a classifier that generalizes well on the unseen test data drawn from the same distribution. In the following, we present the algorithmic details. For presentation clarity, a table of notations is provided in the supplementary material.

3.1 Learning to Generalize

Let $f(\mathbf{x}; \boldsymbol{\theta})$ be a generic classifier parameterized by $\boldsymbol{\theta}$ and $\Phi(\mathbf{p}, \mathbf{y})$ be a non-negative function that measures the discrepancy between distributions \mathbf{p} and \mathbf{y} . We further assume $\Phi(\mathbf{p}, \mathbf{y}) = 0$ if and only if $\mathbf{p} = \mathbf{y}$, and thus if \mathbf{y} is fixed, then $\mathbf{p} = \mathbf{y}$ is the (global) minima of the function $\Phi(\cdot, \mathbf{y})$. We formulate the loss of (\mathbf{x}, \mathbf{y}) as $\mathcal{L}(\mathbf{x}, \mathbf{y}; \boldsymbol{\theta}) = \Phi(f(\mathbf{x}; \boldsymbol{\theta}), \mathbf{y})$. The learning-to-generalize problem is then formulated as following:

$$\begin{aligned} \min_{\mathcal{Y}} \sum_{k=1}^{N^l} \mathcal{L}(\mathbf{x}_k^l, \mathbf{y}_k; \boldsymbol{\theta}^*(\mathcal{Y})) \\ \text{s.t. } \boldsymbol{\theta}^*(\mathcal{Y}) = \arg \min_{\boldsymbol{\theta}} \sum_{i=1}^{N^u} \mathcal{L}(\mathbf{x}_i^u, \hat{\mathbf{y}}_i; \boldsymbol{\theta}), \end{aligned} \quad (1)$$

where $\mathcal{Y} = \{\hat{\mathbf{y}}_i : i = 1, \dots, N^u\}$ denotes the pseudo labels of the unlabeled data. Here we consider the train-

ing labeled data $\{\mathbf{x}_k^l : k = 1, \dots, N^l\}$ as the validation set and seek the optimal pseudo labels \mathcal{Y} by minimizing the validation loss on the labeled data. However, the validation loss also depends on the optimal model parameters that are obtained by minimizing the unlabeled data loss using pseudo labels \mathcal{Y} . Thus the above formulation is a nest optimization problem.

Solving the nested minimization problem exactly is computationally prohibitive because calculating the gradients of the outer loop requires an entire inner optimization. Thus, we online approximate the outer loop gradients in a way similar to Ren et al. (2018). Specifically, we adapt the generated pseudo labels based on the current mini-batch and replace the inner optimization with a single SGD step. As such, the descent direction of the pseudo labels is guided by the back-propagated signals of the labeled data loss.

Consider the gradient-based deep learning framework in which gradients are calculated at the mini-batch level and the SGD-like optimizer is used to update model parameters. With a little bit abuse of notations, at the t^{th} training step, a mini-batch of labeled data $\{(\mathbf{x}_k^l, \mathbf{y}_k) : k = 1, \dots, B^l\}$ and a mini-batch of unlabeled data $\{\mathbf{x}_i^u : i = 1, \dots, B^u\}$ are sampled, where B^l and B^u denote the batch sizes of the labeled data and unlabeled data, respectively. The pseudo labels of the unlabeled data are initialized as the current predictions of the classifier:

$$\tilde{\mathbf{y}}_i = f(\mathbf{x}_i^u; \boldsymbol{\theta}_t). \quad (2)$$

We then compute the unlabeled data loss and the gradient *w.r.t.* the model parameters:

$$\begin{aligned} \mathcal{L}(\mathbf{x}_i^u, \tilde{\mathbf{y}}_i; \boldsymbol{\theta}_t) &= \Phi(f(\mathbf{x}_i^u; \boldsymbol{\theta}_t), \tilde{\mathbf{y}}_i), \\ \nabla \boldsymbol{\theta}_t &= \frac{1}{B^u} \sum_{i=1}^{B^u} \nabla_{\boldsymbol{\theta}} \mathcal{L}(\mathbf{x}_i^u, \tilde{\mathbf{y}}_i; \boldsymbol{\theta}_t). \end{aligned} \quad (3)$$

Note that since the initialized pseudo labels are precisely the predictions of the classifier, the unlabeled loss achieves minimum value - zero. Thus the gradient is zero, *i.e.*, $\boldsymbol{\theta}_t = 0$. However, the Jacobian matrix of $\nabla \boldsymbol{\theta}_t$ *w.r.t.* the pseudo labels is not necessarily a zero matrix, thus making optimization via differentiating $\nabla \boldsymbol{\theta}_t$ possible.

We apply one SGD step on the model parameters:

$$\tilde{\boldsymbol{\theta}}_{t+1} = \boldsymbol{\theta}_t - \alpha_t \nabla \boldsymbol{\theta}_t, \quad (4)$$

where α_t is the learning rate of the inner loop. The SGD step is then evaluated on the labeled data and the labeled data loss is treated as the meta-objective. We differentiate the meta-objective through the SGD step and compute the meta-gradient *w.r.t.* the pseudo

labels:

$$\begin{aligned}\mathcal{L}(\mathbf{x}_k^l, \mathbf{y}_k; \tilde{\boldsymbol{\theta}}_{t+1}) &= \Phi(f(\mathbf{x}_k^l; \tilde{\boldsymbol{\theta}}_{t+1}), \mathbf{y}_k), \\ \nabla_{\tilde{\mathbf{y}}_i} &= \frac{1}{B^l} \sum_{k=1}^{B^l} \nabla_{\tilde{\mathbf{y}}_i} \mathcal{L}(\mathbf{x}_k^l, \mathbf{y}_k; \tilde{\boldsymbol{\theta}}_{t+1}).\end{aligned}\quad (5)$$

Note that by unfolding one SGD step, the labeled data loss is related to the pseudo labels of the unlabeled data. Moreover, since the labeled data loss serves as the meta-objective to be differentiated, the update of the pseudo labels is guided by the label information, *i.e.*, the meta-gradients, and thus concerns the specific task on interest. Similar techniques are developed in the optimization-based meta-learning literature (Finn et al., 2017) and employed in a wide range of applications (Ren et al., 2018; Liu et al., 2018; Metz et al., 2017; Liu et al., 2019).

Finally, we perform one SGD step on the pseudo labels,

$$\hat{\mathbf{y}}_i = \tilde{\mathbf{y}}_i - \beta_t \nabla_{\tilde{\mathbf{y}}_i}, \quad (6)$$

where β_t is the meta learning rate, and compute the consistency loss from the unlabeled data and the updated pseudo labels. The meta-learning algorithm is summarized in Algorithm 1.

3.2 First-Order Approximation

In Section 3.1, the most computationally expensive operation is differentiation through the SGD step in Eq. (5), as the second-order derivative is involved. To avoid this, we apply the chain rule to the second-order derivative by substituting Eq. (3) and Eq. (4) into Eq. (5):

$$\begin{aligned}\frac{\partial \mathcal{L}}{\partial \tilde{y}_{i,j}}(\mathbf{x}_k^l, \mathbf{y}_k; \tilde{\boldsymbol{\theta}}_{t+1}) \\ = -\frac{\alpha_t}{B^u} \nabla_{\boldsymbol{\theta}}^\top \frac{\partial \mathcal{L}}{\partial \tilde{y}_{i,j}}(\mathbf{x}_i^u, \tilde{\mathbf{y}}_i; \boldsymbol{\theta}_t) \cdot \nabla_{\boldsymbol{\theta}} \mathcal{L}(\mathbf{x}_k^l, \mathbf{y}_k; \boldsymbol{\theta}_t).\end{aligned}\quad (7)$$

The gradient of the validation loss *w.r.t.* the pseudo labels can thus be formulated as

$$\begin{aligned}\nabla_{\tilde{y}_{i,j}} &= \frac{1}{B^l} \sum_{k=1}^{B^l} \frac{\partial \mathcal{L}}{\partial \tilde{y}_{i,j}}(\mathbf{x}_k^l, \mathbf{y}_k; \boldsymbol{\theta}_t) \\ &= -\frac{\alpha_t}{B^u} \nabla_{\boldsymbol{\theta}}^\top \frac{\partial \mathcal{L}}{\partial \tilde{y}_{i,j}}(\mathbf{x}_i^u, \tilde{\mathbf{y}}_i; \boldsymbol{\theta}_t) \cdot \left(\frac{1}{B^l} \sum_{k=1}^{B^l} \nabla_{\boldsymbol{\theta}} \mathcal{L}(\mathbf{x}_k^l, \mathbf{y}_k; \boldsymbol{\theta}_t) \right).\end{aligned}\quad (8)$$

Let

$$\nabla \boldsymbol{\theta}_t^l = \frac{1}{B^l} \sum_{i=1}^{B^l} \nabla_{\boldsymbol{\theta}} \mathcal{L}(\mathbf{x}_i^l, \mathbf{y}_i; \boldsymbol{\theta}_t), \quad (9)$$

and then it can be easily shown with Taylor expansion that as $\epsilon \rightarrow 0$,

$$\nabla_{\tilde{y}_{i,j}} = -\frac{\alpha_t}{2B^u\epsilon} \left(\frac{\partial \mathcal{L}}{\partial \tilde{y}_{i,j}}(\mathbf{x}_i^u, \tilde{\mathbf{y}}_i; \boldsymbol{\theta}_t + \epsilon \nabla \boldsymbol{\theta}_t^l) - \frac{\partial \mathcal{L}}{\partial \tilde{y}_{i,j}}(\mathbf{x}_i^u, \tilde{\mathbf{y}}_i; \boldsymbol{\theta}_t - \epsilon \nabla \boldsymbol{\theta}_t^l) \right). \quad (10)$$

Thus, we adopt the first-order approximation and use a sufficiently small ϵ to approximate $\nabla_{\tilde{y}_{i,j}}$. As suggested in Liu et al. (2018), we use $\epsilon = 0.01/\|\nabla \boldsymbol{\theta}_t^l\|_2$ in this work.

Furthermore, the gradients *w.r.t.* the pseudo labels can be calculated in the closed form. Here, following the common practice of consistency-based SSL (Tarvainen and Valpola, 2017), we adopt the Kullback–Leibler divergence loss $\Phi^{\text{KL}}(\mathbf{p}, \mathbf{y}) = \sum_n y_n \log(y_n/p_n)$ as the regular labeled data loss, and the mean squared error (MSE) loss $\Phi^{\text{MSE}}(\mathbf{p}, \mathbf{y}) = \|\mathbf{p} - \mathbf{y}\|_2^2$ for the consistency loss. For the MSE loss, the gradients *w.r.t.* the pseudo labels are approximated by:

$$\nabla_{\tilde{\mathbf{y}}_i} \approx \frac{\alpha_t}{B^u\epsilon} (f(\mathbf{x}_i^u; \boldsymbol{\theta}_t + \epsilon \nabla \boldsymbol{\theta}_t^l) - f(\mathbf{x}_i^u; \boldsymbol{\theta}_t - \epsilon \nabla \boldsymbol{\theta}_t^l)). \quad (11)$$

3.3 Improved Training Protocol

The above-discussed meta-learning algorithm utilizes the unlabeled examples to improve the generalization ability. However, there still remains the sampling bias issue in SSL. Motivated by the success of the *mixup* augmentation (Zhang et al., 2018) in SSL (Wang et al., 2019; Berthelot et al., 2019), we incorporate the cross-domain mixup augmentation in the proposed meta-learning algorithm. We first assume that the mini-batches of labeled and unlabeled data are of the same batch size, *i.e.*, $B^l = B^u = B$, and then interpolate between each pair of labeled and unlabeled examples to generate new training data. Note that when generating the corresponding labels, we interpolate between the actual labels \mathbf{y}_i of the labeled examples and the updated pseudo labels $\hat{\mathbf{y}}_i$ of the unlabeled examples,

$$\begin{aligned}\mathbf{x}_i^{\text{in}} &= \lambda_i \mathbf{x}_i^l + (1 - \lambda_i) \mathbf{x}_i^u, \\ \mathbf{y}_i^{\text{in}} &= \lambda_i \mathbf{y}_i + (1 - \lambda_i) \hat{\mathbf{y}}_i,\end{aligned}\quad i = 1, \dots, B, \quad (12)$$

where $\lambda_1, \dots, \lambda_B$ are *i.i.d.* samples drawn from the Beta(γ, γ) distribution. Finally, the total loss is formulated as

$$\mathcal{L} = \underbrace{\sum_{i=1}^B \mathcal{L}^{\text{KL}}(\mathbf{x}_i^{\text{in}}, \mathbf{y}_i^{\text{in}}; \boldsymbol{\theta}_t)}_{\text{classification loss}} + \underbrace{\sum_{i=1}^B \mathcal{L}^{\text{MSE}}(\mathbf{x}_i^u, \hat{\mathbf{y}}_i; \boldsymbol{\theta}_t)}_{\text{consistency loss}}, \quad (13)$$

and the algorithm with first-order approximation and mixup augmentation is illustrated in Algorithm 2.

3.4 Convergence Analysis

In this section, we present the convergence analysis of Algorithm 1. Due to of the scarcity of labeled examples, we assume all labeled data are sampled at each step,

Algorithm 1 Meta-Learning Algorithm.

Input: regular learning rates $\{\alpha_t\}$,
 meta learning rates $\{\beta_t\}$

for $t := 1$ to $\#iters$ do

$$\begin{aligned} & \{(\mathbf{x}_k^l, \mathbf{y}_k)\}_{k=1}^{B^l} \leftarrow \text{BatchSampler}(\mathcal{D}^l) \\ & \{\mathbf{x}_i^u\}_{i=1}^{B^u} \leftarrow \text{BatchSampler}(\mathcal{D}^u) \\ & \tilde{\mathbf{y}}_i = f(\mathbf{x}_i^u; \boldsymbol{\theta}_t) \\ & \mathcal{L}(\mathbf{x}_i^u, \tilde{\mathbf{y}}_i; \boldsymbol{\theta}_t) = \Phi(f(\mathbf{x}_i^u; \boldsymbol{\theta}_t), \tilde{\mathbf{y}}_i) \\ & \nabla \boldsymbol{\theta}_t = \frac{1}{B^u} \sum_{i=1}^{B^u} \nabla_{\boldsymbol{\theta}} \mathcal{L}(\mathbf{x}_i^u, \tilde{\mathbf{y}}_i; \boldsymbol{\theta}_t) \\ & \tilde{\boldsymbol{\theta}}_{t+1} = \boldsymbol{\theta}_t - \alpha_t \nabla \boldsymbol{\theta}_t \\ & \mathcal{L}(\mathbf{x}_k^l, \mathbf{y}_k; \tilde{\boldsymbol{\theta}}_{t+1}) = \Phi(f(\mathbf{x}_k^l; \tilde{\boldsymbol{\theta}}_{t+1}), \mathbf{y}_k) \\ & \nabla \tilde{\mathbf{y}}_i = \frac{1}{B^l} \sum_{k=1}^{B^l} \nabla_{\tilde{\mathbf{y}}_i} \mathcal{L}(\mathbf{x}_k^l, \mathbf{y}_k; \tilde{\boldsymbol{\theta}}_{t+1}) \\ & \hat{\mathbf{y}}_i = \tilde{\mathbf{y}}_i - \beta_t \nabla \tilde{\mathbf{y}}_i \\ & \mathcal{L}(\mathbf{x}_i^u, \hat{\mathbf{y}}_i; \boldsymbol{\theta}_t) = \Phi(f(\mathbf{x}_i^u; \boldsymbol{\theta}_t), \hat{\mathbf{y}}_i) \\ & \nabla \hat{\boldsymbol{\theta}}_t = \frac{1}{B^u} \sum_{i=1}^{B^u} \nabla_{\boldsymbol{\theta}} \mathcal{L}(\mathbf{x}_i^u, \hat{\mathbf{y}}_i; \boldsymbol{\theta}_t) \\ & \boldsymbol{\theta}_{t+1} = \text{Optimizer}(\boldsymbol{\theta}_t, \nabla \hat{\boldsymbol{\theta}}_t, \alpha_t) \end{aligned}$$

end

Algorithm 2 Algorithm with Mix-Up Augmentation.

Input: regular learning rates $\{\alpha_t\}$,
 meta learning rates $\{\beta_t\}$

for $t := 1$ to $\#iters$ do

$$\begin{aligned} & \{(\mathbf{x}_k^l, \mathbf{y}_k)\}_{k=1}^{B^l} \leftarrow \text{BatchSampler}(\mathcal{D}^l) \\ & \{\mathbf{x}_i^u\}_{i=1}^{B^u} \leftarrow \text{BatchSampler}(\mathcal{D}^u) \\ & \tilde{\mathbf{y}}_i = f(\mathbf{x}_i^u; \boldsymbol{\theta}_t) \\ & \mathcal{L}^{\text{KL}}(\mathbf{x}_k^l, \mathbf{y}_k; \boldsymbol{\theta}_t) = \Phi^{\text{KL}}(f(\mathbf{x}_k^l; \boldsymbol{\theta}_t), \mathbf{y}_k) \\ & \nabla \boldsymbol{\theta}_t^l = \frac{1}{B^l} \sum_{k=1}^{B^l} \nabla_{\boldsymbol{\theta}} \mathcal{L}^{\text{KL}}(\mathbf{x}_k^l, \mathbf{y}_k; \boldsymbol{\theta}_t) \\ & \epsilon = 0.01 \|\nabla \boldsymbol{\theta}_t^l\|^{-1} \\ & \nabla \tilde{\mathbf{y}}_i = \epsilon^{-1} (f(\mathbf{x}_i^u; \boldsymbol{\theta}_t + \epsilon \nabla \boldsymbol{\theta}_t^l) - f(\mathbf{x}_i^u; \boldsymbol{\theta}_t - \epsilon \nabla \boldsymbol{\theta}_t^l)) \\ & \hat{\mathbf{y}}_i = \tilde{\mathbf{y}}_i - \beta_t \nabla \tilde{\mathbf{y}}_i \\ & \lambda_i \leftarrow \text{Beta}(\gamma, \gamma) \\ & \mathbf{x}_i^{\text{in}} = \lambda_i \mathbf{x}_i^l + (1 - \lambda_i) \mathbf{x}_i^u \\ & \mathbf{y}_i^{\text{in}} = \lambda_i \mathbf{y}_i + (1 - \lambda_i) \hat{\mathbf{y}}_i \\ & \mathcal{L}_{\text{cls}}^{\text{KL}}(\mathbf{x}_i^{\text{in}}, \mathbf{y}_i^{\text{in}}; \boldsymbol{\theta}_t) = \Phi^{\text{KL}}(f(\mathbf{x}_i^{\text{in}}; \boldsymbol{\theta}_t), \mathbf{y}_i^{\text{in}}) \\ & \mathcal{L}_{\text{cons}}^{\text{MSE}}(\mathbf{x}_i^u, \hat{\mathbf{y}}_i; \boldsymbol{\theta}_t) = \Phi^{\text{MSE}}(f(\mathbf{x}_i^u; \boldsymbol{\theta}_t), \hat{\mathbf{y}}_i) \\ & \nabla \hat{\boldsymbol{\theta}}_t = \frac{1}{B^l} \sum_{i=1}^{B^l} (\nabla_{\boldsymbol{\theta}} \mathcal{L}_{\text{cls}}^{\text{KL}} + \nabla_{\boldsymbol{\theta}} \mathcal{L}_{\text{cons}}^{\text{MSE}}) \\ & \boldsymbol{\theta}_{t+1} = \text{Optimizer}(\boldsymbol{\theta}_t, \nabla \hat{\boldsymbol{\theta}}_t, \alpha_t) \end{aligned}$$

end

i.e., $B^l = N^l$, and that the MSE loss is used in the unlabeled consistency loss (Eq. (3)). Under mild conditions, we show that Algorithm 1 is guaranteed to converge to a critical point of the meta-objective (Theorem 1), and enjoys a convergence rate of $O(1/\epsilon^2)$ (Theorem 2), which is the same as the regular SGD. The proofs are presented in the supplementary material.

Theorem 1. *Let*

$$G(\boldsymbol{\theta}; \mathcal{D}^l) = \frac{1}{N^l} \sum_{k=1}^{N^l} \mathcal{L}(\mathbf{x}_k^l, \mathbf{y}_k; \boldsymbol{\theta}_t) \quad (14)$$

be the loss function of the labeled examples. Assume

- (i) the gradient function $\nabla_{\boldsymbol{\theta}} G$ is Lipschitz-continuous with a Lipschitz constant L_0 ; and
- (ii) the norm of the Jacobian matrix of f w.r.t. $\boldsymbol{\theta}$ is upper-bounded by a constant M , i.e.,

$$\|J_{\boldsymbol{\theta}} f(\mathbf{x}_i^u; \boldsymbol{\theta})\| \leq M, \quad \forall i \in \{1, \dots, N^u\}. \quad (15)$$

If the regular learning rate α_t and meta learning rate β_t satisfy $\alpha_t^2 \beta_t < (4M^2 L_0)^{-1}$, then each SGD step of Algorithm 1 will decrease the validation loss $G(\boldsymbol{\theta})$, regardless of the selected unlabeled examples, i.e.,

$$G(\boldsymbol{\theta}_{t+1}) \leq G(\boldsymbol{\theta}_t), \quad \text{for each } t. \quad (16)$$

Furthermore, the equality holds if and only if $\nabla \tilde{\mathbf{y}} = \mathbf{0}$ for the selected unlabeled batch at the t^{th} step.

Theorem 2. *Assume the same conditions as in Theorem 1, and*

$$\inf_t (\beta_t - 4\alpha_t^2 \beta_t^2 M^2 L_0) = D_1 > 0, \quad \inf_t \alpha_t = D_2 > 0. \quad (17)$$

We further assume that the unlabeled dataset contains the labeled dataset, i.e., $\mathcal{D}^l \subseteq \mathcal{D}^u$. Then, Algorithm 1 achieves $\mathbb{E} [\|\nabla_{\boldsymbol{\theta}} G(\boldsymbol{\theta}_t)\|^2] \leq \epsilon$ in $O(1/\epsilon^2)$ steps, i.e.,

$$\min_{1 \leq t \leq T} \mathbb{E} [\|\nabla_{\boldsymbol{\theta}} G(\boldsymbol{\theta}_t)\|^2] \leq \frac{C}{\sqrt{T}}, \quad (18)$$

where C is a constant independent of the training process.

Remarks. (i) The assumption in (15) is realistic. Here, we assume the neural network f is continuously differentiable w.r.t. $\boldsymbol{\theta}$. Due to the existence of norm-based regularization, i.e., weight decay, we can assume $\boldsymbol{\theta}$ is optimized within a compact set in the parameter space. The Jacobian function $J_{\boldsymbol{\theta}} f$ is thus bounded within the compact set due to its continuity. Furthermore, since there are finite training examples, the bound in (15) is plausible. (ii) The conditions in (17) specify that the learning rates α_t and β_t can neither grow too large nor decay to zero too rapidly. The step learning rate annealing strategy can satisfy this condition as long as the initial learning rate is sufficiently small. (iii) The condition $\mathcal{D}^l \subseteq \mathcal{D}^u$ can be satisfied by incorporating the labeled data into the unlabeled set.

Table 1: Semi-supervised classification error rates of the Conv-Large (Tarvainen and Valpola, 2017) architecture on the SVHN, CIFAR-10, and CIFAR-100 datasets. The numbers of labeled data are 1k, 4k, and 10k for these three datasets, respectively.

Method	SVHN	CIFAR-10	CIFAR-100
II-Model (Laine and Aila, 2017)	4.82%	12.36%	39.19%
TE (Laine and Aila, 2017)	4.42%	12.16%	38.65%
MT (Tarvainen and Valpola, 2017)	3.95%	12.31%	-
MT+SNTG (Luo et al., 2018)	3.86%	10.93%	-
VAT (Miyato et al., 2018)	5.42%	11.36%	-
VAT+Ent (Miyato et al., 2018)	3.86%	10.55%	-
VAT+Ent+SNTG (Luo et al., 2018)	3.83%	9.89%	-
VAT+VAAD (Park et al., 2018)	3.55%	9.22%	-
MA-DNN (Chen et al., 2018)	4.21%	11.91%	34.51%
Co-training (Qiao et al., 2018)	3.29%	8.35%	34.63%
MT+fastSWA (Athiwaratkun et al., 2019)	-	9.05%	33.62%
TNAR-VAE (Yu et al., 2019)	3.74%	8.85%	-
ADA-Net (Wang et al., 2019)	4.62%	10.30%	-
ADA-Net+fastSWA (Wang et al., 2019)	-	8.72%	-
DualStudent (Ke et al., 2019)	-	8.89%	32.77%
Ours	3.15%	7.78%	30.74%
Fully-Supervised	2.67%	4.88%	22.10%

Table 2: Semi-supervised classification error rates of the 26-layer ResNet (He et al., 2016) architecture with the shake-shake regularization (Gastaldi, 2017) on the CIFAR-10 and CIFAR-100 datasets.

Dataset	CIFAR-10			CIFAR-100		
	50k	50k	50k	50k	50k	50k
#Images	50k	50k	50k	50k	50k	50k
#Labels	1k	2k	4k	6k	8k	10k
fastSWA	6.6%	5.7%	5.0%	-	-	28.0%
Ours	6.3%	5.2%	4.1%	26.7%	25.1%	22.9%

4 Experiments

We evaluate the proposed algorithm on the SVHN (Netzer et al., 2011), CIFAR (Krizhevsky et al., 2009), and ImageNet (Russakovsky et al., 2015) datasets. The 13-layer Conv-Large (Tarvainen and Valpola, 2017) and 26-layer ResNet (He et al., 2016) with the shake-shake regularization (Gastaldi, 2017) are used as the backbone models. More implementation details can be found in the supplementary material. The training sources are available at <https://github.com/Sakura03/SemiMeta>.

4.1 Results on the SVHN and CIFAR

In Table 1, we report the semi-supervised classification error rates of the proposed algorithm and state-of-the-art methods on the SVHN, CIFAR-10, and CIFAR-100 datasets. The proposed meta-learning algorithm per-

Table 3: Semi-supervised classification error rates on the ImageNet (Russakovsky et al., 2015) dataset. 10% training images are used as the labeled data.

Method	Top-1	Top-5
Labeled-Only	53.65%	31.01%
MT	49.07%	23.59%
Co-training	46.50%	22.73%
ADA-Net	44.91%	21.18%
Ours	44.87%	18.88%
Fully-Supervised	29.15%	10.12%

forms favorably against the previous approaches on all three datasets. In addition, we explore the effectiveness of the proposed algorithm on different backbone architectures and evaluate on the 26-layer ResNet (He et al., 2016) with the shake-shake regularization (Gastaldi, 2017). Since only a few previous papers include experiments on this backbone, we just compare the performance with the “fastSWA” method (Athiwaratkun et al., 2019) which gives quite complete results and achieves the state-of-the-art accuracy. Table 2 shows that the proposed algorithm performs favorably under all different experimental configurations, even with fewer labeled examples, indicating the efficacy of the consistency loss guided by the meta-gradients.

4.2 Results on the ImageNet

The evaluation results with the ResNet-18 (He et al., 2016) backbone on the ImageNet dataset (Russakovsky

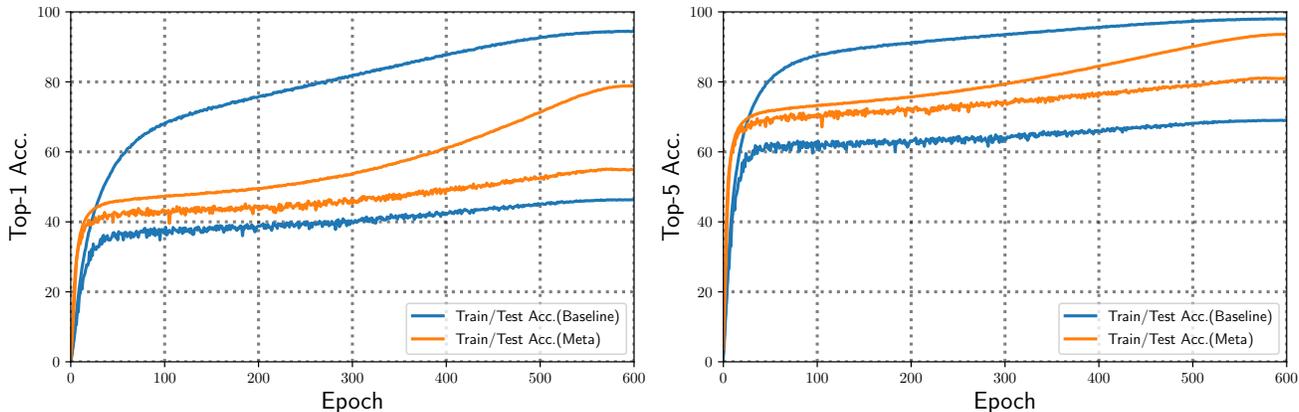


Figure 1: Accuracy curves of the baseline method and the proposed algorithm.

Table 4: Ablation study of the meta-learning component and the mixup augmentation. The same number of labeled data is used as in Table 1.

No.	Meta	Mix-Up	SVHN	CIFAR-10	CIFAR-100
1			9.76%	15.43%	38.74%
2	✓		3.68%	11.63%	35.40%
3		✓	5.60%	11.10%	32.67%
4	✓	✓	3.15%	7.78%	30.74%

et al., 2015) are summarized in Table 3. The proposed algorithm performs well against the ADA-Net (Wang et al., 2019) in terms of top-5 accuracy. In addition, we demonstrate the accuracy curves of the baseline setting and our approach in Figure 1, where the baseline setting means only 10% of the total training examples are used as the labeled data used during training. Figure 1 shows that merely involving 10% training examples will lead to severe overfitting, as the training accuracy is very high while testing accuracy is pretty low. The problem is alleviated in our approach thanks to the explicit learning-to-generalize training scheme. Though the training accuracy is not higher than the baseline model, the proposed algorithm achieves better testing accuracy. These results suggest that the consistency loss can effectively regularize the training and benefits to the generalization ability of the learned model.

4.3 Ablation Studies

Effectiveness of Components. We analyze the contributions of the meta-learning and mixup augmentation components of the proposed algorithm. The experimental settings are the same as those in Table 1. Table 4 shows both components can significantly improve the classification accuracy in the semi-supervised settings. The mixup augmentation is a simple trick to solve the sample bias issue in SSL. Also, we adapt the mixup formulation (*i.e.*, mixing the actual labels and updated pseudo labels in Eq. (12)) to make it compati-

Table 5: Comparison with ADA-Net (Wang et al., 2019) on both with and without mix-up setting. The number of provided labels are 4k and 1k in CIFAR-10 and SVHN, respectively.

Dataset	CIFAR-10		SVHN	
	w/o mixup	w mixup	w/o mixup	w/ mixup
ADA-Net	18.67%	8.87%	10.76%	5.90%
Ours	11.63%	7.78%	3.86%	3.15%

ble with the meta learning framework. Such adaptation is non-trivial as indicated by the performance improvement on top of the mixup-only setting (see the last two rows of Table 4). Moreover, the meta-learning component can further improve performance with the presence of the mix-up augmentation, indicating that meta-learning is orthogonal to the existing data augmentation techniques as a research direction. To further verify the effectiveness of the proposed meta-learning component, we compare our method with ADA-Net on both with and without mixup setting as shown in Table 5. We can conclude from the table that the proposed meta learning still outperforms than ADA-Net without using mixup augmentation.

Impact of #Labels. As shown in Figure 2, we evaluate the robustness of our method against the variation of the number of labels on the SVHN and CIFAR-10 datasets using the same experimental settings, except the number of labels as in Table 1. We can find that the accuracy of the labeled-only baseline degrades heavily when reducing the number of labels. In general, the scarcity of labeled data may lead to a worse generalization ability and more severe overfitting. However, the proposed SSL method can retain a relatively high performance under each setting thanks to our learning-to-generalize regularization. Therefore, the proposed method can effectively improve the generalization ability even with fewer labeled examples.

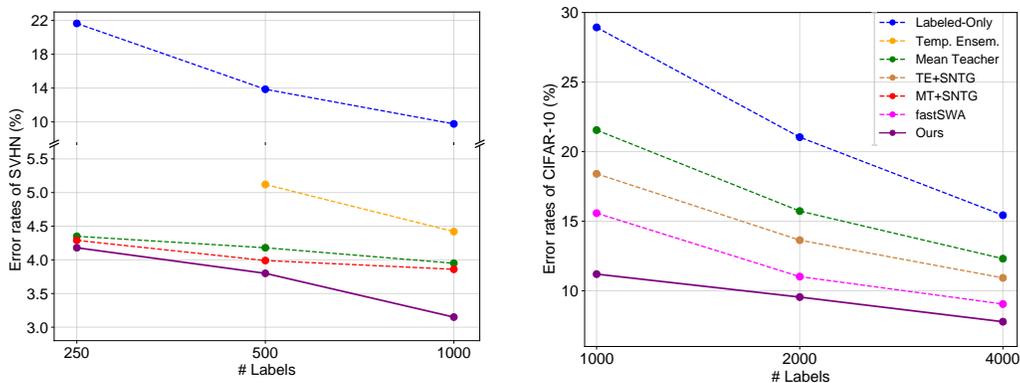
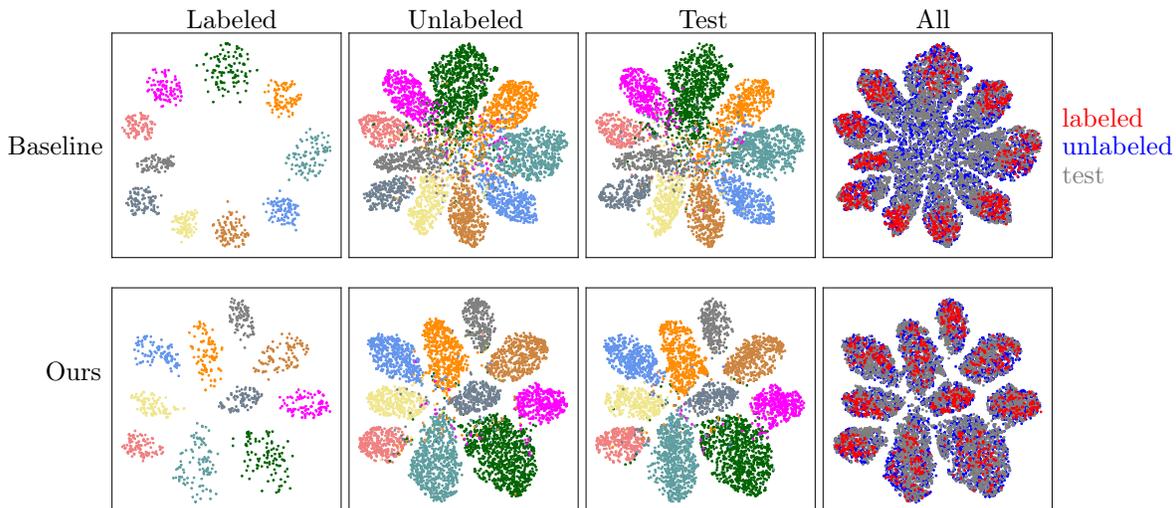
Figure 2: Error rate *v.s.* number of labeled examples.

Figure 3: Visualization of the SVHN features by the labeled-only baseline and our method. We extract features of the labeled (first column), unlabeled (second column), and test samples (third column) and project them to the two-dimensional space using t-SNE. In the first three columns, different categories are represented in different colors. In the fourth column, we plot the projected points all together to demonstrate the empirical distribution discrepancy. The labeled, unlabeled, and test examples are represented in red, blue, and gray, respectively.

Feature Visualization. To further analyse the efficacy of our method, we visualize the SVHN features by projecting 128-dimensional features onto a two-dimensional space using the t-SNE (van der Maaten and Hinton, 2008) technique. For comparison, we also present the feature visualization of the labeled-only baseline method. As displayed in Figure 3, there is a considerable empirical distribution discrepancy between labeled and unlabeled examples of the baseline method. In contrast, such discrepancy is reduced and the feature distributions of different domains are aligned to some extent by our approach, indicating the generalization ability is improved. Furthermore, considering the unlabeled and test samples, we observe the margin among features of different categories is obviously clearer under our approach, which results in more discriminative features and better classification accuracy.

5 Conclusion

In this work, we address the semi-supervised learning problem in a meta-learning fashion. Specifically, we propose a new regularization term that takes the model generalization ability into consideration in semi-supervised learning, which is different from most existing regularization terms that are based on the empirical assumption of the inner structure of the data manifold. Theoretically, we prove that the proposed algorithm is guaranteed to converge to a critical point at a convergence rate of $O(1/\epsilon^2)$. Extensive experimental results demonstrate that the proposed algorithm performs favorably against the state-of-the-art semi-supervised learning methods on the SVHN, CIFAR, and ImageNet datasets with better generalization ability.

References

- B. Athiwaratkun, M. Finzi, P. Izmailov, and A. G. Wilson. There Are Many Consistent Explanations of Unlabeled Data: Why You Should Average. In *International Conference on Learning Representations (ICLR)*, 2019. 1, 2, 6
- D. Berthelot, N. Carlini, I. Goodfellow, N. Papernot, A. Oliver, and C. A. Raffel. MixMatch: A Holistic Approach to Semi-Supervised Learning. In *Neural Information Processing Systems (NeurIPS)*, pages 5050–5060, 2019. 2, 4
- D. Berthelot, N. Carlini, E. D. Cubuk, A. Kurakin, K. Sohn, H. Zhang, and C. Raffel. Remixmatch: Semi-supervised learning with distribution alignment and augmentation anchoring. *International Conference on Learning Representations (ICLR)*, 2020. 2
- A. Blum and S. Chawla. Learning from Labeled and Unlabeled Data using Graph Mincuts. In *International Conference on Machine Learning (ICML)*, pages 19–26, 2001. 1
- A. Blum and T. Mitchell. Combining Labeled and Unlabeled Data with Co-Training. In *Conference on Learning Theory (COLT)*, pages 92–100, 1998. 1
- Y. Chen, X. Zhu, and S. Gong. Semi-Supervised Deep Learning with Memory. In *European Conference on Computer Vision (ECCV)*, pages 268–283, 2018. 6
- V. R. de Sa. Learning Classification with Unlabeled Data. In *Neural Information Processing Systems (NeurIPS)*, pages 112–119, 1994. 1
- C. Finn, P. Abbeel, and S. Levine. Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks. In *International Conference on Machine Learning (ICML)*, pages 1126–1135, 2017. 2, 4
- C. Finn, K. Xu, and S. Levine. Probabilistic Model-Agnostic Meta-Learning. In *Neural Information Processing Systems (NeurIPS)*, pages 9516–9527, 2018. 2
- X. Gastaldi. Shake-Shake Regularization. *arXiv preprint arXiv:1705.07485*, 2017. 2, 6
- E. Grant, C. Finn, S. Levine, T. Darrell, and T. Griffiths. Recasting Gradient-Based Meta-Learning as Hierarchical Bayes. In *International Conference on Learning Representations (ICLR)*, 2018. 2
- K. He, X. Zhang, S. Ren, and J. Sun. Deep Residual Learning for Image Recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. 6
- G. E. Hinton, N. Srivastava, A. Krizhevsky, I. Sutskever, and R. R. Salakhutdinov. Improving neural networks by preventing co-adaptation of feature detectors. *arXiv preprint arXiv:1207.0580*, 2012. 2
- P. Izmailov, D. Podoprikin, T. Garipov, D. Vetrov, and A. G. Wilson. Averaging Weights Leads to Wider Optima and Better Generalization. In *Uncertainty in Artificial Intelligence (UAI)*, pages 876–885, 2018. 1
- Z. Ke, D. Wang, Q. Yan, J. Ren, and R. W. Lau. Dual Student: Breaking the Limits of the Teacher in Semi-supervised Learning. In *IEEE International Conference on Computer Vision (ICCV)*, pages 6728–6736, 2019. 6
- A. Krizhevsky, G. Hinton, et al. Learning Multiple Layers of Features from Tiny Images. Technical report, Citeseer, 2009. 2, 6
- S. Laine and T. Aila. Temporal Ensembling for Semi-Supervised Learning. In *International Conference on Learning Representations (ICLR)*, 2017. 1, 2, 6
- H. Liu, K. Simonyan, and Y. Yang. DARTS: Differentiable Architecture Search. *International Conference on Learning Representations (ICLR)*, 2018. 3, 4
- L. Liu, M. Muelly, J. Deng, T. Pfister, and L.-J. Li. Generative Modeling for Small-Data Object Detection. In *IEEE International Conference on Computer Vision (ICCV)*, pages 6073–6081, 2019. 4
- Y. Luo, J. Zhu, M. Li, Y. Ren, and B. Zhang. Smooth Neighbors on Teacher Graphs for Semi-supervised Learning. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8896–8905, 2018. 6
- L. Metz, B. Poole, D. Pfau, and J. Sohl-Dickstein. Unrolled Generative Adversarial Networks. In *International Conference on Learning Representations (ICLR)*, 2017. 3, 4
- T. Miyato, S.-i. Maeda, M. Koyama, and S. Ishii. Virtual Adversarial Training: A Regularization Method for Supervised and Semi-Supervised Learning. *IEEE Transactions on Pattern Recognition and Machine Intelligence (PAMI)*, 41(8):1979–1993, 2018. 1, 2, 6
- Y. Netzer, T. Wang, A. Coates, A. Bissacco, B. Wu, and A. Y. Ng. Reading Digits in Natural Images with Unsupervised Feature Learning. In *Neural Information Processing Systems Workshops*, 2011. 2, 6
- S. Park, J. Park, S.-J. Shin, and I.-C. Moon. Adversarial Dropout for Supervised and Semi-supervised Learning. In *Association for the Advancement of Artificial Intelligence (AAAI)*, pages 3917–3924, 2018. 6
- S. Qiao, W. Shen, Z. Zhang, B. Wang, and A. Yuille. Deep Co-Training for Semi-Supervised Image Recognition. In *European Conference on Computer Vision (ECCV)*, pages 135–152, 2018. 6

- A. Rajeswaran, C. Finn, S. M. Kakade, and S. Levine. Meta-Learning with Implicit Gradients. In *Neural Information Processing Systems (NeurIPS)*, pages 113–124, 2019. 2
- A. Rasmus, M. Berglund, M. Honkala, H. Valpola, and T. Raiko. Semi-Supervised Learning with Ladder Networks. In *Neural Information Processing Systems (NeurIPS)*, pages 3546–3554, 2015. 1
- M. Ren, W. Zeng, B. Yang, and R. Urtasun. Learning to Reweight Examples for Robust Deep Learning. In *International Conference on Machine Learning (ICML)*, pages 4334–4343, 2018. 3, 4
- O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal on Computer Vision (IJCV)*, 115(3):211–252, 2015. 2, 6
- M. Sajjadi, M. Javanmardi, and T. Tasdizen. Regularization With Stochastic Transformations and Perturbations for Deep Semi-Supervised Learning. In *Neural Information Processing Systems (NeurIPS)*, pages 1163–1171, 2016. 1, 2
- K. Sohn, D. Berthelot, C.-L. Li, Z. Zhang, N. Carlini, E. D. Cubuk, A. Kurakin, H. Zhang, and C. Raffel. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *Neural Information Processing Systems (NeurIPS)*, 2020. 2
- M. Szummer and T. Jaakkola. Partially labeled classification with Markov random walks. In *Neural Information Processing Systems (NeurIPS)*, pages 945–952, 2002. 1
- A. Tarvainen and H. Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *Neural Information Processing Systems (NeurIPS)*, pages 1195–1204, 2017. 1, 2, 3, 4, 6
- L. van der Maaten and G. Hinton. Visualizing Data using t-SNE. *Journal of Machine Learning Research (JMLR)*, 9:2579–2605, 2008. 8
- Q. Wang, W. Li, and L. V. Gool. Semi-Supervised Learning by Augmented Distribution Alignment. In *IEEE International Conference on Computer Vision (ICCV)*, pages 1466–1475, 2019. 4, 6, 7
- Y. Wang, J. Guo, S. Song, and G. Huang. Meta-semi: A meta-learning approach for semi-supervised learning. *arXiv preprint arXiv:2007.02394*, 2020. 3
- J. Yoon, T. Kim, O. Dia, S. Kim, Y. Bengio, and S. Ahn. Bayesian Model-Agnostic Meta-Learning. In *Neural Information Processing Systems (NeurIPS)*, pages 7332–7342, 2018. 2
- B. Yu, J. Wu, J. Ma, and Z. Zhu. Tangent-Normal Adversarial Regularization for Semi-Supervised Learning. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10676–10684, 2019. 1, 2, 6
- H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz. mixup: Beyond Empirical Risk Minimization. In *International Conference on Learning Representations (ICLR)*, 2018. 4