

# Bi-Level Optimization for Semi-Supervised Learning with Pseudo-Labeling

Marzi Heidari<sup>1</sup>, Yuhong Guo<sup>1,2</sup>

<sup>1</sup>School of Computer Science, Carleton University, Ottawa, Canada

<sup>2</sup>CIFAR AI Chair, Amii, Canada

marziheidari@cmail.carleton.ca, yuhong.guo@carleton.ca

## Abstract

Semi-supervised learning (SSL) is a fundamental task in machine learning, empowering models to extract valuable insights from datasets with limited labeled samples and a large amount of unlabeled data. Although pseudo-labeling is a widely used approach for SSL that generates pseudo-labels for unlabeled data and leverages them as ground truth labels for training, traditional pseudo-labeling techniques often face challenges that significantly decrease the quality of pseudo-labels and hence the overall model performance. In this paper, we propose a novel Bi-level Optimization method for Pseudo-label Learning (BOPL) to boost semi-supervised training. It treats pseudo-labels as latent variables, and optimizes the model parameters and pseudo-labels jointly within a bi-level optimization framework. By enabling direct optimization over the pseudo-labels towards maximizing the prediction model performance, the method is expected to produce high-quality pseudo-labels. To evaluate the effectiveness of the proposed approach, we conduct extensive experiments on multiple SSL benchmarks. The experimental results show the proposed BOPL outperforms the state-of-the-art SSL techniques.

## 1 Introduction

Despite the remarkable advances achieved by deep learning models, their widespread application has been impeded by the cost of acquiring sufficient amount of labeled data (Kingma et al. 2014; LeCun, Bengio, and Hinton 2015). In light of this challenge, semi-supervised learning (SSL) has emerged as a highly promising research area by offering the capacity to leverage a small number of labeled samples and a sufficient number of unlabeled samples for effective learning (Van Engelen and Hoos 2020). The key for successful SSL lies in effectively exploiting the large number of unlabeled samples to remedy the shortage of the labeled data.

Many SSL methods have been developed to exploit the unlabeled data in various ways, including a variety of loss regularization based methods (Miyato et al. 2018; Zhang et al. 2020), and teacher-student model based methods (Athiwaratkun et al. 2019; Tarvainen and Valpola 2017). One popularly adopted SSL technique is pseudo-labeling that aims to effectively expand the labeled data by generating predicted pseudo-labels for the unlabeled samples and

then using them as ground-truth labels for model training. A number of recent works have explored diverse data augmentation based pseudo-labeling techniques such as MixMatch (Berthelot et al. 2019), ReMixMatch (Berthelot et al. 2020), FixMatch (Sohn et al. 2020), Dash (Xu et al. 2021) and FlexMatch (Zhang et al. 2021) to enhance the quality of the pseudo-labels and improve the model performance, while several other studies such as TSSDL (Shi et al. 2018) and LPD (Isken et al. 2019) have explored label propagation techniques to create pseudo-labels based on the density of the local neighborhood. Although these methods achieve enhanced SSL performance, they lack principled strategies to simultaneously maintain the quality of pseudo-labels and the prediction consistency between labeled and unlabeled samples. Such drawbacks can lead to poor suboptimal solutions, yielding prediction models that cannot generalize well.

In this paper, we propose a novel Bi-level Optimization method for Pseudo-label Learning (BOPL) to address the abovementioned limitations of existing pseudo-labeling techniques. Bi-level optimization provides a convenient framework for simultaneous optimization of two objectives that are at different levels and inter-dependent. We adopt bi-level optimization to address SSL through a novel and direct pseudo-labeling design. Specifically, we treat pseudo-labels as latent variables and formulate pseudo-labeling as a bi-level optimization problem to jointly learn the pseudo-labels and model parameters at different levels through a pair of bi-level objectives. By optimizing the pseudo-labels directly, the method is expected to produce more reliable labels for the unlabeled samples. By designing the bi-level objectives properly, the quality of pseudo-labels and the prediction consistency between labeled and unlabeled data can be simultaneously enhanced to promote the generalizability of the prediction model. To validate the proposed methodology, we conduct experiments on multiple SSL benchmark datasets and compare the proposed BOPL with multiple state-of-the-art SSL methods. The experimental results demonstrate BOPL outperforms the comparison methods and achieves state-of-the-art performance. The key contributions of this paper can be summarized as follows:

- We propose a novel bi-level optimization method, BOPL, for SSL, which treats pseudo-labels as latent variables and directly optimizes them.
- We design a proper pair of bi-level objectives to ensure

prediction consistency between labeled and unlabeled samples and promote generalizability of the model.

- Our empirical results validate the efficacy of the proposed BOPL by comparing it with a set of state-of-the-art SSL methods.

## 2 Related Works

**Semi-Supervised Learning** With the widespread application of deep learning models, SSL has been receiving growing attention for exploiting the unlabeled data to reduce demands for labeled samples. One popular line of this research is to develop regularization-based SSL methods that introduce additional loss terms based on the unlabeled data to enhance model training. Examples of such methods include the  $\Pi$ -model (Laine and Aila 2017) and the Temporal-Ensemble (Laine and Aila 2017), which add consistency regularizations to the loss function and leverage the exponential moving average of model predictions. Virtual Adversarial Training (VAT) (Miyato et al. 2018) is another regularization-based approach for SSL that trains a deep neural network with adversarial perturbation based regularization for all the training data samples. A more recent work, Consistency Regularization for Generative Adversarial Networks (CRGAN) (Zhang et al. 2020), combines a generative adversarial network (GAN) with a consistency regularization term to generate pseudo-labels for the unlabeled data. Another line of SSL research centers on the teacher-student based methods, which train a student network to match the predictions of a teacher network on unlabeled data. Mean Teacher (MT) model (Tarvainen and Valpola 2017) is a well-known method in this category. MT + Fast SWA (Athiwaratkun et al. 2019) combines Mean Teacher with Fast Stochastic Weight Averaging to further improve performance. Smooth Neighbors on Teacher Graphs (SNTG) (Luo et al. 2018) leverages a graph for the teacher model to regulate or control the distribution of features in the unlabeled samples. Interpolation Consistency Training (ICT) (Verma et al. 2022) is a regularization-based method built on teacher-student networks. It enforces consistency between the predictions on an interpolated set of unlabeled data points and the interpolation of the predictions on those points, pushing the decision boundary to low-density regions.

**Pseudo-Labeling** One early representative pseudo-labeling technique for SSL is Co-Training (Blum and Mitchell 1998), which trains two classifiers to generate pseudo-labels on unlabeled samples for each other. Recently, more SSL studies have focused on generating good pseudo-labels to support model training. Pseudo-Label (Lee 2013) generates labels for unlabeled samples based on model predictions while filtering out the low-confidence predictions. MixMatch (Berthelot et al. 2019) employs data augmentation to create multiple versions of each input data sample and obtains predictions for all of them. The predictions are then averaged to produce the pseudo-labels. Several other works, such as ReMixMatch (Berthelot et al. 2020), UDA (Xie et al. 2020), and FixMatch (Sohn et al. 2020), first generate pseudo-labels on weakly augmented samples based on confidence thresholds, and then use them

as annotations for strongly augmented samples. Another set of works known as label propagation methods, such as TSSDL (Shi et al. 2018) and LPD (Isken et al. 2019), assign pseudo-labels based on the density of the local neighborhood. DASO (Oh, Kim, and Kweon 2022), on the other hand, blends confidence-based pseudo-labels and density-based pseudo-labels in different ways for each class. Moreover, some approaches, such as Dash (Xu et al. 2021) and FlexMatch (Zhang et al. 2021), adjust the confidence thresholds dynamically in a curriculum learning manner to improve the quality of pseudo-labels for SSL. SimMatch (Zheng et al. 2022) uses similarity-based label propagation to improve pseudo-label quality. CoMatch (Li, Xiong, and Hoi 2021) combines consistency regularization and mutual learning for enhanced pseudo-labeling. Margin-Match (Sosea and Caragea 2023) integrates consistency regularization and pseudo-labeling. It tracks the model’s behavior over time to mask out low-quality predictions.

**Bi-Level Optimization Methods** Bi-level optimization has emerged as a powerful learning technique that allows the optimization of an inner (lower-level) objective while simultaneously optimizing an outer (higher-level) objective that depends on the solution to the lower-level problem. It has been deployed to solve various problems, including hyperparameter optimization (Pedregosa 2016), and neural architecture search (Liu, Simonyan, and Yang 2019). Recently, some researchers have adopted bi-level optimization for SSL. In particular, Meta Pseudo-Labels (Pham et al. 2021) formulates SSL as a bi-level optimization problem that optimizes the teacher network parameters at the outer level while optimizing the student network parameters at the inner level with the pseudo-labels determined by the teacher network. Meta-Semi (Wang et al. 2022) deploys a bi-level meta optimization framework, which optimizes the weights of the unlabeled instances at the outer level based on loss on the labeled samples while learning the model parameters at the inner level by minimizing the weighted loss on unlabeled samples with predicted pseudo-labels. By contrast, we propose to directly optimize pseudo-labels as latent variables at the outer level of an innovatively designed bi-level optimization framework for SSL.

## 3 Proposed Method

We consider the following SSL setting: The training dataset consists of a relatively small number of labeled samples,  $\mathcal{D}^l = \{(\mathbf{x}_i^l, \mathbf{y}_i)\}_{i=1}^{N^l}$ , and a large number of unlabeled samples,  $\mathcal{D}^u = \{\mathbf{x}_i^u\}_{i=1}^{N^u}$ , where  $\mathbf{x}_i^l \in \mathcal{X}$  and  $\mathbf{x}_i^u \in \mathcal{X}$  denote the  $i$ -th labeled and unlabeled instances respectively, and  $\mathbf{y}_i$  is the corresponding one-hot label vector of  $\mathbf{x}_i^l$  that indicates the class labels from the label set  $\mathcal{Y}$ . We assume the size of the unlabeled set greatly surpasses the labeled set:  $N^u \gg N^l$ . The goal is to train a classifier  $f : \mathcal{X} \rightarrow \mathcal{Y}$  that generalizes well to previously unseen test data drawn from the same distribution as the training data.

In this section, we present our proposed Bi-Level Optimization method for Pseudo-label Learning (BOPL), which directly optimizes the pseudo-labels of the unlabeled samples at the outer level by treating the model parameters

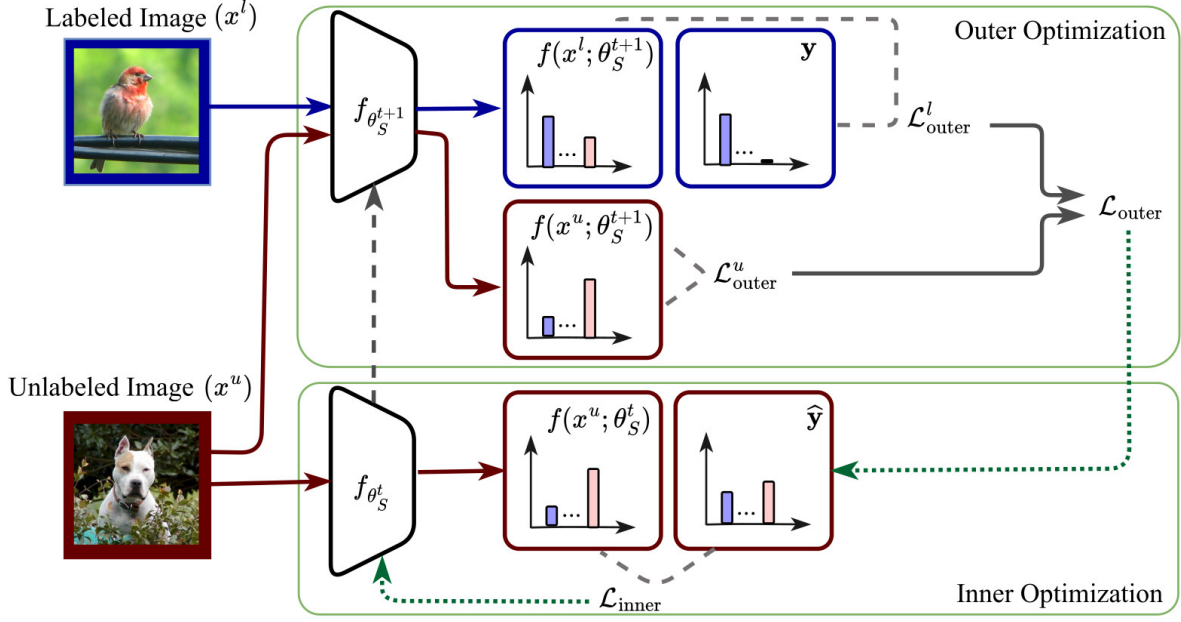


Figure 1: The proposed Bi-level Optimization framework for Pseudo-Labeling (BOPL). The inner loss  $\mathcal{L}_{\text{inner}}$  is used to optimize the model parameters by leveraging the unlabeled data with pseudo-labels. The outer loss  $\mathcal{L}_{\text{outer}}$  is used to optimize the pseudo-labels as latent variables by utilizing both the labeled and unlabeled samples; it treats the model parameters as a function of the pseudo-labels.

as a function of the pseudo-labels. The architecture of the proposed BOPL framework is illustrated in Figure 1. We present the bi-level optimization formulation in section 3.1, and provide the optimization algorithm in section 3.2. In the subsequent section 4, we present a model fine-tuning procedure given the learned pseudo-labels. Finally, we propose to further improve BOPL by incorporating interpolation consistency training into model fine-tuning in section 4.1.

### 3.1 Pseudo-Labeling with Bi-Level Optimization

We formulate the pseudo-labeling problem of SSL as a novel bi-level optimization problem by directly learning the pseudo-labels of unlabeled samples as latent variables through an outer level optimization, while learning the model parameters through an inner level optimization. To mitigate potential oscillations during training, we adopt a simple teacher-student concept to maintain two sets of model parameters: a student model  $\theta_S$  and a teacher model  $\theta_T$ . The student model is learned directly during training, while the teacher model is updated using an exponential moving average (EMA) of the student model; at the  $t$ -th iteration, the update is conducted as follows:

$$\theta_T^t = \beta \theta_T^{t-1} + (1 - \beta) \theta_S^t, \quad (1)$$

where  $\beta \in (0, 1)$  is a momentum coefficient hyperparameter. The student and teacher models are initialized by performing pre-training on the labeled and unlabeled data for some iterations using the Mean-Teacher approach (Tarvainen and Valpola 2017).

The bi-level loss functions of the proposed BOPL is designed as follows to induce high quality pseudo-labels while learning optimal model parameters as functions of the pseudo-labels. First, since we have a large amount of unlabeled samples, when the pseudo-labels of the unlabeled data are given, the optimal model parameters  $\theta_S^*$  can simply be learned by minimizing the standard cross-entropy loss on the pseudo-labeled data  $\mathcal{D}^u$  that are originally unlabeled:

$$\mathcal{L}_{\text{inner}}(\theta_S, \hat{Y}) = \frac{1}{N^u} \sum_{i=1}^{N^u} \ell_{\text{CE}}(\hat{y}_i, f(\mathbf{x}_i^u, \theta_S)) \quad (2)$$

where  $\ell_{\text{CE}}$  denotes the cross-entropy loss,  $f(\mathbf{x}_i^u, \theta_S)$  is the probabilistic prediction output of the student model for the unlabeled instance  $\mathbf{x}_i^u$ , and  $\hat{y}_i$  denotes the pseudo-label vector for  $\mathbf{x}_i^u$ . As different pseudo-labels will lead to different optimal model parameters  $\theta_S^*$ , we can treat  $\theta_S^*$  as a function of the pseudo-labels  $\hat{Y} = [\hat{y}_1, \dots, \hat{y}_i, \dots, \hat{y}_{N^u}]$ . The quality of the pseudo-labels  $\hat{Y}$  determines the quality of the model parameters  $\theta_S^*$ , while the latter reflects the former.

Due to the dependence of the optimal model parameters  $\theta_S^*$  on pseudo-labels, we propose to evaluate the quality of pseudo-labels by assessing the corresponding optimal model's performance. Specifically, given no separate validation set, we validate  $\theta_S^*$ 's performance using the whole training set— $\mathcal{D}^l$  and  $\mathcal{D}^u$ . On the labeled data  $\mathcal{D}^l$ , we use the standard cross-entropy loss as the validation loss:

$$\mathcal{L}_{\text{outer}}^l(\theta_S^*) = \frac{1}{N^l} \sum_{i=1}^{N^l} \ell_{\text{CE}}(\mathbf{y}_i, f(\mathbf{x}_i^l, \theta_S^*)). \quad (3)$$

This validation loss can naturally reflect the prediction consistency between the labeled (with true labels) and unlabeled (with pseudo-labels) samples. On the unlabeled data  $\mathcal{D}^u$ , we use the entropy loss as the validation loss:

$$\mathcal{L}_{\text{outer}}^u(\theta_S^*) = \frac{1}{N^u} \sum_{i=1}^{N^u} \ell_E(f(\mathbf{x}_i^u, \theta_S^*)), \quad (4)$$

where  $\ell_E$  denotes the entropy loss function. The entropy loss reflects the prediction uncertainty of the given model on the unlabeled samples. The overall validation loss on the training data integrates these two loss terms together:

$$\mathcal{L}_{\text{outer}}(\theta_S^*) = \mathcal{L}_{\text{outer}}^l(\theta_S^*) + \lambda \mathcal{L}_{\text{outer}}^u(\theta_S^*), \quad (5)$$

where  $\lambda \in [0, 1]$  is a trade-off hyperparameter. Since the labeled data have the true label information,  $\lambda$  is typically set to a value smaller than 1 to give more weight to the validation loss on  $\mathcal{D}^l$ .

To obtain high quality pseudo-labels, we then treat the pseudo-labels  $\hat{Y}$  as latent variables and learn the pseudo-labels by minimizing the corresponding optimal model's validation loss. This leads to the following bi-level optimization problem:

$$\begin{aligned} \min_{\hat{Y} \in \mathcal{C}} \quad & \mathcal{L}_{\text{outer}}(\theta_S^*(\hat{Y})) \\ \text{s.t.} \quad & \theta_S^* = \arg \min_{\theta_S} \mathcal{L}_{\text{inner}}(\theta_S, \hat{Y}) \end{aligned} \quad (6)$$

where the pseudo-label variables  $\hat{Y}$  are subject to the soft label distribution constraints, such as:

$$\mathcal{C} = \{\hat{Y} : \hat{y}_i \geq 0, \sum_j \hat{y}_{ij} = 1, \forall i\}. \quad (7)$$

By optimizing the pseudo-labels directly, the proposed bi-level optimization is expected to produce more reliable high quality pseudo labels for the unlabeled samples through the outer validation loss. Meanwhile, it also provides a principled mechanism to ensure prediction consistency between the labeled and pseudo-labeled samples. These desirable properties can foreseeably enhance the generalization capacity of the subsequent prediction model.

### 3.2 Optimization Procedure

We propose to solve the bi-level optimization problem in Eq.(6) using stochastic gradient descent. The key for designing the optimization algorithm lies in deriving the gradient of the outer validation loss function  $\mathcal{L}_{\text{outer}}$  regarding the pseudo-label variables  $\hat{Y} = [\hat{y}_1, \dots, \hat{y}_i, \dots, \hat{y}_{N^u}]$ .

As the pseudo-label variables only impact the validation loss via the model parameters  $\theta_S^*$ , we derive the gradient of  $\mathcal{L}_{\text{outer}}$  w.r.t. each  $\hat{y}_i$  using the following chain rule:

$$\nabla_{\hat{y}_i} \mathcal{L}_{\text{outer}} = \nabla_{\theta_S^*} \mathcal{L}_{\text{outer}} \cdot \nabla_{\hat{y}_i} \theta_S^*. \quad (8)$$

Moreover, we propose to approximate  $\theta_S^*$  by updating the model parameters  $\theta_S$  with a single gradient descent step over the inner loss function  $\mathcal{L}_{\text{inner}}$ . Specifically, at the  $t$ -th iteration, we set

$$\theta_S^* = \theta_S^{t+1} = \theta_S^t - \alpha \nabla_{\theta_S} \mathcal{L}_{\text{inner}}(\theta_S^t, \hat{Y}^t) \quad (9)$$

where  $\alpha$  is the learning rate. For convenience of notation, let's define:

$$\delta(\theta_S^{t+1}) = \nabla_{\theta_S} \mathcal{L}_{\text{outer}}^l(\theta_S^{t+1}) + \lambda \nabla_{\theta_S} \mathcal{L}_{\text{outer}}^u(\theta_S^{t+1}). \quad (10)$$

We then compute the target gradient as follows.

**Proposition 3.1.** *With the chain rule in Eq.(8) and the approximation in Eq.(9), the gradient of  $\mathcal{L}_{\text{outer}}$  w.r.t.  $\hat{y}_i$  can be expressed as:*

$$\nabla_{\hat{y}_i} \mathcal{L}_{\text{outer}} = -\alpha \cdot \delta(\theta_S^{t+1}) \cdot \nabla_{\hat{y}_i} \nabla_{\theta_S} \mathcal{L}_{\text{inner}}(\theta_S^t, \hat{Y}) \quad (11)$$

*Proof.* Given the approximation in Eq.(9), we have:

$$\begin{aligned} \nabla_{\hat{y}_i} \theta_S^* &= \nabla_{\hat{y}_i} \theta_S^{t+1} = \nabla_{\hat{y}_i} (\theta_S^t - \alpha \nabla_{\theta_S} \mathcal{L}_{\text{inner}}(\theta_S^t, \hat{Y})) \\ &= -\alpha \cdot \nabla_{\hat{y}_i} \nabla_{\theta_S} \mathcal{L}_{\text{inner}}(\theta_S^t, \hat{Y}) \end{aligned} \quad (12)$$

Then following the chain rule in in Eq.(8) and the definition of  $\mathcal{L}_{\text{outer}}$  in Eq.(5), we can immediately derive:

$$\begin{aligned} \nabla_{\hat{y}_i} \mathcal{L}_{\text{outer}} &= (\nabla_{\theta_S} \mathcal{L}_{\text{outer}}^l(\theta_S^{t+1}) + \lambda \nabla_{\theta_S} \mathcal{L}_{\text{outer}}^u(\theta_S^{t+1})) \cdot \nabla_{\hat{y}_i} \theta_S^* \\ &= -\alpha \cdot \delta(\theta_S^{t+1}) \cdot \nabla_{\hat{y}_i} \nabla_{\theta_S} \mathcal{L}_{\text{inner}}(\theta_S^t, \hat{Y}) \end{aligned} \quad (13)$$

□

Although it is convenient to compute the first order derivatives of the loss functions, the second order derivative  $\nabla_{\hat{y}_i} \nabla_{\theta_S} \mathcal{L}_{\text{inner}}(\theta_S^t, \hat{Y})$  is not easy to compute. We hence further leverage a finite difference approximation method (Bottou 2012) for the partial derivative  $\nabla_{\theta_S} \mathcal{L}_{\text{inner}}(\theta_S^t, \hat{Y})$  to provide a convenient solution to the second order derivative.

**Proposition 3.2.** *Let  $\epsilon$  be a very small constant. By using a finite difference approximation for  $\nabla_{\theta_S} \mathcal{L}_{\text{inner}}(\theta_S^t, \hat{Y})$ , we can approximate the target gradient as follows:*

$$\nabla_{\hat{y}_i} \mathcal{L}_{\text{outer}} \approx \frac{\alpha}{2\epsilon} \left( \log(f(\mathbf{x}_i^u; \theta_S^+)) - \log(f(\mathbf{x}_i^u; \theta_S^-)) \right) \quad (14)$$

where  $\theta_S^+ = \theta_S^t + \epsilon \cdot \delta(\theta_S^{t+1})$  and  $\theta_S^- = \theta_S^t - \epsilon \cdot \delta(\theta_S^{t+1})$ .

*Proof.* Let's use the following finite difference approximation for the partial derivative  $\nabla_{\theta_S} \mathcal{L}_{\text{inner}}(\theta_S^t, \hat{Y})$ :

$$\nabla_{\theta_S} \mathcal{L}_{\text{inner}}(\theta_S^t, \hat{Y}) \approx \frac{\mathcal{L}_{\text{inner}}(\theta_S^+, \hat{Y}) - \mathcal{L}_{\text{inner}}(\theta_S^-, \hat{Y})}{2\epsilon \cdot \delta(\theta_S^{t+1})} \quad (15)$$

By substituting this back to Eq.(11), we can express the target gradient as:

$$\begin{aligned} \nabla_{\hat{y}_i} \mathcal{L}_{\text{outer}} &\approx -\frac{\alpha}{2\epsilon} \left( \nabla_{\hat{y}_i} \mathcal{L}_{\text{inner}}(\theta_S^+, \hat{Y}) - \nabla_{\hat{y}_i} \mathcal{L}_{\text{inner}}(\theta_S^-, \hat{Y}) \right) \\ &= -\frac{\alpha}{2\epsilon} \left( \begin{array}{c} \nabla_{\hat{y}_i} \ell_{\text{CE}}(\hat{y}_i, f(\mathbf{x}_i^u, \theta_S^+)) - \\ \nabla_{\hat{y}_i} \ell_{\text{CE}}(\hat{y}_i, f(\mathbf{x}_i^u, \theta_S^-)) \end{array} \right) \\ &= \frac{\alpha}{2\epsilon} \left( \log(f(\mathbf{x}_i^u; \theta_S^+)) - \log(f(\mathbf{x}_i^u; \theta_S^-)) \right) \end{aligned} \quad (16)$$

□

---

**Algorithm 1: Training Algorithm for BOPL**


---

**Input:** training dataset:  $\mathcal{D}^l$  and  $\mathcal{D}^u$ ; hyperparameters; initialized model parameters:  $\theta_S, \theta_T$

**Output:** learned model parameters  $\theta_S$

**Set:**  $\theta_S^1 \leftarrow \theta_S; \theta_T^1 \leftarrow \theta_T; \hat{Y} = f(X^u; \theta_T^1)$

**for**  $t = 1$  to maxiters **do**

**for** minibatch  $B^u \in \mathcal{D}^u$  **do**

    Compute loss  $\mathcal{L}_{\text{inner}}(\theta_S^t, \hat{Y})$  on  $B^u$  via Eq.(2)

    Calculate  $\theta_S^{t+1}$  using Eq.(9)

    Compute  $\mathcal{L}_{\text{outer}}^u(\theta_S^{t+1})$  on  $B^u$  via Eq.(4)

    Set  $\mathcal{L}_{\text{outer}}^l = 0$

**for** minibatch  $B^l \in \mathcal{D}^l$  **do**

$\mathcal{L}_{\text{outer}}^l = \mathcal{L}_{\text{outer}}^l + \frac{1}{N^l} \ell_{\text{CE}}(B^l; \theta_S^{t+1})$

**end for**

    Calculate  $\delta(\theta_S^{t+1})$  using Eq.(10)

    Calculate gradient  $\nabla_{\hat{Y}} \mathcal{L}_{\text{outer}}$  on  $B^u$  with Eq.(14)

    Update  $\hat{Y}$  on  $B^u$  with Eq.(17) and Eq.(18)

**end for**

$\theta_S^{t+1} \leftarrow \theta_S^t - \alpha \nabla_{\theta_S} \mathcal{L}_{\text{inner}}(\theta_S^t, \hat{Y})$  on  $\mathcal{D}^u$

$\theta_T^{t+1} \leftarrow \beta \theta_T^t + (1 - \beta) \theta_S^{t+1}$

**end for**

Fine-tune the model parameters  $\theta_S$  with the input data and the learned pseudo-labels  $\hat{Y}$

---

**Update of Pseudo-Labels** The pseudo-labels  $\hat{Y}$  can be initialized by applying the prediction model with the initial teacher model parameters  $\theta_T^1$ , such as  $\hat{y}_i = f(\mathbf{x}_i^u; \theta_T^1)$ . Then in each iteration, we update each pseudo-label vector  $\hat{y}_i$  with a gradient descent step. Moreover, to ensure a valid label distribution vector, we have to renormalize each updated pseudo-label vector to satisfy the constraints  $\mathcal{C}$  in Eq.(7). In particular, we adopt the following gradient descent and renormalization process with a learning rate  $\alpha$  by deploying the ReLU operator:

$$\tilde{y}_i = \frac{\text{ReLU}(\hat{y}_i - \alpha \nabla_{\hat{y}_i} \mathcal{L}_{\text{outer}})}{\sum_j \text{ReLU}(\hat{y}_i - \alpha \nabla_{\hat{y}_i} \mathcal{L}_{\text{outer}})_j} \quad (17)$$

where  $\nabla_{\hat{y}_i} \mathcal{L}_{\text{outer}}$  is calculated using Eq.(14). Furthermore, to mitigate oscillations of gradient updates across iterations, we integrate the prediction outputs from the current teacher model to determine the final updated pseudo-labels. For the  $t$ -th iteration, we eventually update the pseudo-labels as follows:

$$\hat{y}_i = \gamma \tilde{y}_i + (1 - \gamma) f(\mathbf{x}_i^u; \theta_T^t), \quad (18)$$

where the hyperparameter  $\gamma \in [0, 1]$  determines the linear combination weight. The overall batch-wise training algorithm for the proposed BOPL is presented in Algorithm 1.

## 4 Model Fine-Tuning

After bi-level optimization, we can obtain high quality pseudo-labels  $\hat{Y}$  for the unlabeled samples in the training set. To ensure the model parameters are well trained given the pseudo-labels, we propose to further fine-tune the model parameters on both the labeled and pseudo-labeled data. In

particular, we use the standard cross-entropy loss as the supervised loss  $\mathcal{L}_{\text{sup}}$  on the labeled data  $\mathcal{D}^l$  and use a mean squared error as the pseudo-label supervised loss  $\mathcal{L}_{\text{pseudo}}$  on the pseudo-labeled data  $(\mathcal{D}^u, \hat{Y})$  for model fine-tuning:

$$\mathcal{L}_{\text{sup}}(\theta_S) = \mathcal{L}_{\text{outer}}^l(\theta_S) \quad (19)$$

$$\mathcal{L}_{\text{pseudo}}(\theta_S) = \frac{1}{N^u} \sum_{i=1}^{N^u} \|f(\mathbf{x}_i^u; \theta_S) - \hat{y}_i\|_2^2 \quad (20)$$

The mean squared error measures the stability of the model predictions on the pseudo-labeled data.

We perform model fine-tuning to maintain prediction consistency between labeled and unlabeled data by minimizing the following joint loss function:

$$\mathcal{L}_{\text{ft}} = \mathcal{L}_{\text{sup}} + \eta \mathcal{L}_{\text{pseudo}} \quad (21)$$

where  $\eta$  is a trade-off hyperparameter that balances the contribution of the two loss terms.

### 4.1 Fine-Tuning with ICT

Interpolation Consistency Training (ICT) (Verma et al. 2022) facilitates model training and improves model’s robustness and generalizability by enforcing prediction consistency across interpolated points. We extend the idea of ICT to improve our fine-tuning procedure. As part of this approach, we generate interpolated data points by employing mix-up augmentation on unlabeled samples with their learned pseudo-labels,  $(\mathcal{D}^u, \hat{Y})$ . This involves linearly combining a pair of randomly selected unlabeled samples  $\mathbf{x}_i^u$  and  $\mathbf{x}_j^u$  and their pseudo-labels using a mixing parameter,  $\mu$ , sampled from a Beta distribution:

$$\begin{aligned} \mathbf{x}^m &= \mu \mathbf{x}_i^u + (1 - \mu) \mathbf{x}_j^u, \\ \hat{\mathbf{y}}^m &= \mu \hat{\mathbf{y}}_i + (1 - \mu) \hat{\mathbf{y}}_j \end{aligned} \quad (22)$$

where  $\mathbf{x}_m^u$  and  $\hat{\mathbf{y}}_m$  are the mix-up (or interpolated) sample and pseudo-label, respectively. By pairing the samples in  $\mathcal{D}^u$  with a random shuffled version of it, the same number (i.e.,  $N^u$ ) of mix-up samples can be generated. We then compute the mean squared consistency loss on the mix-up samples and their corresponding pseudo-labels as follows

$$\mathcal{L}_{\text{cons}}^{\text{ICT}} = \frac{1}{N^u} \sum_{i=1}^{N^u} \|f(\mathbf{x}_i^m; \theta_S) - \hat{\mathbf{y}}_i^m\|_2^2. \quad (23)$$

We fine-tune the model parameters  $\theta_S$  by replacing the  $\mathcal{L}_{\text{pseudo}}$  loss in Eq.(21) with this ICT consistency loss, and refer to this extended fine-tuning approach based on BOPL as BOPL+ICT.

## 5 Experiments

### 5.1 Experimental Setup

**Datasets** We conducted comprehensive experiments on four commonly used image classification benchmarks: CIFAR-10, CIFAR-100 (Krizhevsky, Hinton et al. 2009), SVHN (Netzer et al. 2011) and STL-10 (Coates, Ng, and Lee 2011). Following previous works, on each dataset we randomly select a subset of samples with equal sizes from each

Dataset Number of Labeled Samples	CIFAR-10			CIFAR-100		SVHN
	250	1000	4000	2500	10000	1000
Mean Teacher (Tarvainen and Valpola 2017)	32.32 <sub>(2.30)</sub>	17.32 <sub>(4.00)</sub>	10.36 <sub>(0.25)</sub>	53.91 <sub>(0.57)</sub>	35.83 <sub>(0.24)</sub>	5.65 <sub>(0.45)</sub>
ICT (Verma et al. 2022)	-	-	7.66 <sub>(0.17)</sub>	-	-	3.53 <sub>(0.07)</sub>
MixMatch (Berthelot et al. 2019)	11.05 <sub>(0.15)</sub>	7.75 <sub>(0.32)</sub>	6.24 <sub>(0.06)</sub>	39.94 <sub>(0.37)</sub>	28.31 <sub>(0.33)</sub>	3.27 <sub>(0.31)</sub>
UDA (Xie et al. 2020)	8.82 <sub>(1.08)</sub>	5.87 <sub>(0.13)</sub>	4.29 <sub>(0.07)</sub>	33.13 <sub>(0.22)</sub>	24.50 <sub>(0.25)</sub>	1.89 <sub>(0.01)</sub>
ReMixMatch (Berthelot et al. 2020)	5.44 <sub>(0.05)</sub>	6.27 <sub>(0.34)</sub>	6.24 <sub>(0.06)</sub>	-	-	3.27 <sub>(0.31)</sub>
FixMatch (Sohn et al. 2020)	5.07 <sub>(0.35)</sub>	-	4.26 <sub>(0.05)</sub>	28.29 <sub>(0.11)</sub>	22.60 <sub>(0.12)</sub>	2.28 <sub>(0.11)</sub>
FlexMatch (Zhang et al. 2021)	4.98 <sub>(0.09)</sub>	-	4.19 <sub>(0.01)</sub>	26.49 <sub>(0.20)</sub>	21.90 <sub>(0.15)</sub>	6.72 <sub>(0.01)</sub>
CoMatch (Li, Xiong, and Hoi 2021)	4.91 <sub>(0.33)</sub>	-	4.56 <sub>(0.20)</sub>	28.37 <sub>(0.35)</sub>	20.86 <sub>(0.36)</sub>	-
SimMatch (Zheng et al. 2022)	4.84 <sub>(0.36)</sub>	-	3.96 <sub>(0.01)</sub>	25.07 <sub>(0.32)</sub>	20.58 <sub>(0.11)</sub>	-
Meta-Semi (Wang et al. 2022)	-	7.34 <sub>(0.22)</sub>	6.10 <sub>(0.10)</sub>	-	-	-
MarginMatch (Sosea and Caragea 2023)	4.73 <sub>(0.12)</sub>	-	3.98 <sub>(0.02)</sub>	-	-	1.93 <sub>(0.01)</sub>
Meta Pseudo-Labels (Pham et al. 2021)	-	-	3.89 <sub>(0.07)</sub>	-	-	1.99 <sub>(0.07)</sub>
BOPL (Ours)	<b>4.65</b> <sub>(0.27)</sub>	<b>4.42</b> <sub>(0.18)</sub>	<b>3.12</b> <sub>(0.08)</sub>	<b>24.84</b> <sub>(0.29)</sub>	<b>19.92</b> <sub>(0.27)</sub>	<b>1.81</b> <sub>(0.07)</sub>
BOPL+ICT (Ours)	<b>4.12</b> <sub>(0.26)</sub>	<b>4.02</b> <sub>(0.13)</sub>	<b>3.03</b> <sub>(0.08)</sub>	<b>23.16</b> <sub>(0.27)</sub>	<b>18.12</b> <sub>(0.24)</sub>	<b>1.78</b> <sub>(0.07)</sub>

Table 1: Comparison results in terms of mean test error and standard deviation using WRN-28-2 as the backbone on CIFAR-10 and SVHN and using WRN-28-8 as the backbone on CIFAR-100.

class as labeled data and keep the remaining samples unlabeled. We conducted experiments on CIFAR-10 with 250, 1,000, 2,000, and 4,000 labeled samples, on CIFAR-100 with 2,500, 4,000, and 10,000 labeled samples, on SVHN with 1000 and 500 labeled samples, and on STL-10 with 1,000 images as the labeled data.

**Implementation Details** Following previous works (Luo et al. 2018; Tarvainen and Valpola 2017), we adopted random  $2 \times 2$  translation and random horizontal flip to augment the training set. To maintain fair comparisons with a great number of previous studies on the multiple datasets, we conducted comprehensive experiments by using four types of backbone networks: a 13-layer CNN (CNN-13), a Wide-RestNet-28-2 (WRN-28-2) (Zagoruyko and Komodakis 2016), a Wide-RestNet-37-2 (WRN-37), and a Wide-RestNet-28-8 (WRN-28-8). The WRN models are chosen based on previous works for comparability (Berthelot et al. 2019). For training CNN-13, we employed the SGD optimizer with a Nesterov momentum (Nesterov 1983) of 0.9, an L2 regularization coefficient of  $1e-4$  for CIFAR-10 and CIFAR-100 datasets and  $5e-5$  for SVHN, and an initial learning rate  $\alpha$  of 0.1. To schedule the learning rate, we utilized the cosine learning rate annealing technique (Loshchilov and Hutter 2017; Verma et al. 2022), which reduces the learning rate in a cosine-like fashion to help the model converge at better minima. The WRN-28-2 model was trained using SGD as the optimizer as well. An L2 regularization coefficient of  $5e-4$  and an initial learning rate of 0.01 were employed. For the WRN-37-2 model, the training configuration includes the SGD optimizer, an L2 regularization coefficient of  $5e-4$ , and an initial learning rate of 0.01. Lastly, for the WRN-28-8 model, the training setup involves the SGD optimizer, an L2 regularization coefficient of  $1e-3$ , and an initial learning rate of 0.01. Specifically for BOPL, we set the batch size to 128,  $\lambda = 1e-2$ ,  $\epsilon = 1e-2$ ,  $\gamma = 0.5$ ,  $\beta = 0.999$ , and  $\eta = 1$ . We pre-train the model for 50 epochs using the Mean-Teacher algorithm and then proceed to train

Dataset Number of Labeled Samples	STL-10 1000
MeanTeacher (Tarvainen and Valpola 2017)	21.43 <sub>(2.39)</sub>
MixMatch (Berthelot et al. 2019)	10.41 <sub>(0.61)</sub>
UDA (Xie et al. 2020)	7.66 <sub>(0.56)</sub>
ReMixMatch (Berthelot et al. 2020)	5.23 <sub>(0.45)</sub>
FixMatch (Sohn et al. 2020)	5.17 <sub>(0.63)</sub>
MarginMatch (Sosea and Caragea 2023)	5.52 <sub>(0.15)</sub>
BOPL (Ours)	<b>4.93</b> <sub>(0.54)</sub>
BOPL+ICT (Ours)	<b>4.11</b> <sub>(0.52)</sub>

Table 2: Comparison of mean test error and standard deviation by using WRN-37-2 on STL-10.

BOPL for 400 epochs. Finally, we fine-tune the model for 100 epochs using both the labeled data and the unlabeled data with learned pseudo-labels. For each experiment, we repeat five independent runs and report the mean test error with standard deviation.

## 5.2 Comparison Results

We compare the proposed BOPL approach with a great set of state-of-the-art SSL methods, including Mean Teacher (Tarvainen and Valpola 2017), MixMatch (Berthelot et al. 2019), FixMatch (Sohn et al. 2020), ReMixMatch (Berthelot et al. 2020), FlexMatch (Zhang et al. 2021), UDA (Xie et al. 2020), CoMatch (Li, Xiong, and Hoi 2021), SimMatch (Zheng et al. 2022), Meta Pseudo-Labels (Pham et al. 2021), Meta-Semi (Wang et al. 2022), ICT (Verma et al. 2022) and MarginMatch (Sosea and Caragea 2023). We conducted comprehensive experiments on benchmark datasets using different backbone networks separately to provide fair comparisons with these previous SSL methods.

Table 1 reports the comparison results on three datasets by using WRN-28-2 as the backbone network on CIFAR-10 and SVHN and using WRN-28-8 as the backbone network on CIFAR-100. We can see that BOPL and BOPL+ICT

Dataset Number of Labeled Samples	CIFAR-100	
	4000	10000
BOPL	<b>36.78</b> <sub>(0.29)</sub>	<b>28.92</b> <sub>(0.18)</sub>
–w/o EMA	40.01 <sub>(0.54)</sub>	31.75 <sub>(0.43)</sub>
–w/o $\mathcal{L}_{\text{outer}}^u$	38.76 <sub>(0.49)</sub>	30.12 <sub>(0.42)</sub>
–w/o $\mathcal{L}_{\text{outer}}^l$	45.34 <sub>(0.63)</sub>	39.92 <sub>(0.52)</sub>
–w/o fine-tuning	42.12 <sub>(0.55)</sub>	35.80 <sub>(0.45)</sub>
–w/o pre-training	44.15 <sub>(0.57)</sub>	37.52 <sub>(0.43)</sub>

Table 3: Ablation study results in terms of mean test error and standard deviation by using CNN-13 on CIFAR-100.

outperform all the other state-of-the-art comparison methods such as Meta-Semi, SimMatch, CoMatch, Meta Pseudo-Labels, and MarginMatch in terms of mean test error across all cases. In particular, when using only 250 labeled samples on CIFAR-10, BOPL and BOPL+ICT achieve impressively low test errors of 4.65% and 4.12%, respectively. Furthermore, BOPL and BOPL+ICT achieve mean test errors of 24.84% and 23.16% on CIFAR-100 with 2,500 labeled samples, respectively. Lastly, on SVHN, BOPL and BOPL+ICT attain mean test errors of 1.81% and 1.78%, respectively, using 1,000 labeled samples.

Table 2 provides a comprehensive comparison of various SSL methods on the STL-10 dataset, utilizing WRN-37 as the backbone network. With a fixed number of 1,000 labeled samples, our proposed BOPL and BOPL+ICT achieve remarkable results with mean test errors of 4.93% and 4.11%, respectively. These results outperform the previous state-of-the-art methods, including FixMatch, UDA, and MarginMatch, showcasing again the effectiveness of our proposed approach.

Overall, the proposed BOPL and BOPL+ICT models outperform the state-of-the-art semi-supervised learning (SSL) methods across various experimental settings that have been widely adopted in prior studies. This highlights the exceptional effectiveness of our framework, particularly in handling SSL tasks with a limited number of labeled samples, demonstrating its potential for practical applications in resource-constrained scenarios.

### 5.3 Ablation Study

We conducted an ablation study to investigate the contribution of different components of BOPL on CIFAR-100 by using CNN-13 as the backbone network. In particular, we compared the full model BOPL with the following variants: (1) “–w/o EMA”, which drops the teacher model parameters by disabling the EMA update; (2) “–w/o  $\mathcal{L}_{\text{outer}}^u$ ”, which drops the unlabeled data from the outer loss; (3) “–w/o  $\mathcal{L}_{\text{outer}}^l$ ”, which drops the labeled data from the outer loss; (4) “–w/o fine-tuning”, which drops fine-tuning, and (5) “–w/o pre-training”, which drops pre-training. The comparison results are reported in Table 3. We can see the variant “–w/o  $\mathcal{L}_{\text{outer}}^l$ ” produces the largest test error increase, which indicates the labeled data is critical for assessing the quality of the pseudo-labels and ensuring prediction consistency between labeled and pseudo-labeled data under the proposed bi-level optimization framework. Meanwhile,

Dataset Number of Labeled Samples	CIFAR-100	
	4000	10000
ReLU+Teacher (BOPL)	<b>36.78</b> <sub>(0.29)</sub>	<b>28.92</b> <sub>(0.18)</sub>
ReLU	37.11 <sub>(0.38)</sub>	30.25 <sub>(0.27)</sub>
Softmax	38.12 <sub>(0.41)</sub>	31.95 <sub>(0.39)</sub>
Softmax+Teacher	37.57 <sub>(0.30)</sub>	29.43 <sub>(0.28)</sub>

Table 4: Results of different pseudo-label update methods in terms of mean test error and standard deviation by using CNN-13 on CIFAR-100.

all the variants have higher test errors than the full model BOPL, which suggests that all the components contribute to the effective performance of BOPL.

### 5.4 Impact of Pseudo-Label Update

Our proposed bi-level optimization method aims to learn high quality pseudo-labels to support SSL. It is important to properly update the pseudo-labels to ensure valid pseudo-label vectors along the stochastic gradient descent training process of the BOPL framework. We conducted experiments on CIFAR-100 with both 4,000 and 10,000 labeled samples by using CNN-13 as the backbone network to investigate the impact of pseudo-label update strategies. In particular, we considered the following pseudo-label update strategies: (1) “ReLU+Teacher”, which is the pseudo-label update strategy we adopted for BOPL. It first applies ReLU based normalization on the updated pseudo-label vector via Eq.(17) and then integrates the teacher model’s prediction via Eq.(18). (2) “ReLU”, which only uses the ReLU based update in Eq.(17) and drops the teacher model predictions. (3) “Softmax”, which uses the softmax normalization in the following form:  $\hat{y}_i = \text{softmax}(\tau(\hat{y}_i - \alpha \nabla_{\hat{y}_i} \mathcal{L}_{\text{outer}}))$ , where  $\tau$  is the sharpening hyper-parameter and is set to 7 in the experiments. (4) “Softmax+Teacher”, which further integrates the teacher model’s prediction after applying softmax normalization, similar to Eq.(18). The comparison results are reported in Table 4. We can see integrating teacher model’s prediction can help reduce test errors in both cases—“ReLU+Teacher” and “Softmax+Teacher”, which highlights the effectiveness of combining the learned pseudo-labels with the teacher model predictions. Meanwhile, “ReLU” normalization works better than “Softmax” normalization, and our proposed pseudo-label update strategy “ReLU+Teacher” produces the best performance.

## 6 Conclusion

In this paper, we proposed BOPL, a novel bi-level optimization approach for SSL, which produces high-quality pseudo-labels for unlabeled data by directly learning pseudo-labels at the outer level of the bi-level optimization, while jointly optimizing model parameters at the inner level. This approach provides a principled framework to simultaneously improve the quality of pseudo-labels and learn the prediction model. Our experimental evaluations on several SSL benchmarks have demonstrated that BOPL outperforms the state-of-the-art SSL methods, achieving remarkable results.

## References

- Athiwaratkun, B.; Finzi, M.; Izmailov, P.; and Wilson, A. G. 2019. There are many consistent explanations of unlabeled data: Why you should average. In *International Conference on Learning Representations (ICLR)*.
- Berthelot, D.; Carlini, N.; Cubuk, E. D.; Kurakin, A.; Sohn, K.; Zhang, H.; and Raffel, C. 2020. ReMixMatch: Semi-supervised learning with distribution matching and augmentation anchoring. In *International Conference on Learning Representations (ICLR)*.
- Berthelot, D.; Carlini, N.; Goodfellow, I.; Papernot, N.; Oliver, A.; and Raffel, C. A. 2019. Mixmatch: A holistic approach to semi-supervised learning. *Advances in Neural Information Processing Systems (NeurIPS)*.
- Blum, A.; and Mitchell, T. 1998. Combining labeled and unlabeled data with co-training. In *annual conference on Computational learning theory*.
- Bottou, L. 2012. *Stochastic gradient descent tricks*, volume 7700 of *Lecture Notes in Computer Science (LNCS)*, 430–445. Springer, neural networks, tricks of the trade, reloaded edition.
- Coates, A.; Ng, A.; and Lee, H. 2011. An analysis of single-layer networks in unsupervised feature learning. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*.
- Isen, A.; Tolias, G.; Avrithis, Y.; and Chum, O. 2019. Label propagation for deep semi-supervised learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Kingma, D. P.; Mohamed, S.; Jimenez Rezende, D.; and Welling, M. 2014. Semi-supervised learning with deep generative models. *Advances in Neural Information Processing Systems (NeurIPS)*.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images.
- Laine, S.; and Aila, T. 2017. Temporal Ensembling for Semi-Supervised Learning. In *International Conference on Learning Representations (ICLR)*.
- LeCun, Y.; Bengio, Y.; and Hinton, G. 2015. Deep learning. *Nature*, 521: 436–444.
- Lee, D.-H. 2013. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *ICML Workshop: Challenges in Representation Learning*.
- Li, J.; Xiong, C.; and Hoi, S. C. 2021. CoMatch: Semi-supervised learning with contrastive graph regularization. In *International Conference on Computer Vision (ICCV)*.
- Liu, H.; Simonyan, K.; and Yang, Y. 2019. DARTS: Differentiable architecture search. In *International Conference on Learning Representations (ICLR)*.
- Loshchilov, I.; and Hutter, F. 2017. SGDR: Stochastic gradient descent with warm restarts. In *International Conference on Learning Representations (ICLR)*.
- Luo, Y.; Zhu, J.; Li, M.; Ren, Y.; and Zhang, B. 2018. Smooth neighbors on teacher graphs for semi-supervised learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Miyato, T.; Maeda, S.-i.; Koyama, M.; and Ishii, S. 2018. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*.
- Nesterov, Y. 1983. A method for unconstrained convex minimization problem with the rate of convergence  $O(1/k^2)$ . In *the USSR Academy of Sciences*.
- Netzer, Y.; Wang, T.; Coates, A.; Bissacco, A.; Wu, B.; and Ng, A. Y. 2011. Reading digits in natural images with unsupervised feature learning. In *NIPS Workshop on Deep Learning and Unsupervised Feature Learning*.
- Oh, Y.; Kim, D.-J.; and Kweon, I. S. 2022. Distribution-aware semantics-oriented pseudo-label for imbalanced semi-supervised learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Pedregosa, F. 2016. Hyperparameter optimization with approximate gradient. In *International Conference on Machine Learning (ICML)*.
- Pham, H.; Dai, Z.; Xie, Q.; and Le, Q. V. 2021. Meta pseudo labels. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Shi, W.; Gong, Y.; Ding, C.; Tao, Z. M.; and Zheng, N. 2018. Transductive semi-supervised deep learning using min-max features. In *European Conference on Computer Vision (ECCV)*.
- Sohn, K.; Berthelot, D.; Carlini, N.; Zhang, Z.; Zhang, H.; Raffel, C. A.; Cubuk, E. D.; Kurakin, A.; and Li, C.-L. 2020. FixMatch: Simplifying semi-supervised learning with consistency and confidence. *Advances in Neural Information Processing Systems (NeurIPS)*.
- Sosea, T.; and Caragea, C. 2023. MarginMatch: Improving Semi-Supervised Learning with Pseudo-Margins. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Tarvainen, A.; and Valpola, H. 2017. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Advances in Neural Information Processing Systems (NeurIPS)*.
- Van Engelen, J. E.; and Hoos, H. H. 2020. A survey on semi-supervised learning. *Machine learning*, 109: 373–440.
- Verma, V.; Kawaguchi, K.; Lamb, A.; Kannala, J.; Solin, A.; Bengio, Y.; and Lopez-Paz, D. 2022. Interpolation consistency training for semi-supervised learning. *Neural Networks*, 145: 90–106.
- Wang, Y.; Guo, J.; Wang, J.; Wu, C.; Song, S.; and Huang, G. 2022. Meta-Semi: A Meta-Learning Approach for Semi-Supervised Learning. *CAAI Artificial Intelligence Research*.
- Xie, Q.; Dai, Z.; Hovy, E.; Luong, T.; and Le, Q. 2020. Unsupervised data augmentation for consistency training. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Xu, Y.; Shang, L.; Ye, J.; Qian, Q.; Li, Y.-F.; Sun, B.; Li, H.; and Jin, R. 2021. Dash: Semi-supervised learning with dynamic thresholding. In *International Conference on Machine Learning (ICML)*.

- Zagoruyko, S.; and Komodakis, N. 2016. Wide residual networks. In *British Machine Vision Conference (BMVC)*.
- Zhang, B.; Wang, Y.; Hou, W.; Wu, H.; Wang, J.; Okumura, M.; and Shinozaki, T. 2021. FlexMatch: Boosting semi-supervised learning with curriculum pseudo labeling. *Advances in Neural Information Processing Systems (NeurIPS)*.
- Zhang, H.; Zhang, Z.; Odena, A.; and Lee, H. 2020. Consistency regularization for generative adversarial networks. In *International Conference on Learning Representations (ICLR)*.
- Zheng, M.; You, S.; Huang, L.; Wang, F.; Qian, C.; and Xu, C. 2022. SimMatch: Semi-supervised learning with similarity matching. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.