BidMatch: Boosting Semi-Supervised Learning by Bi-Dimensional Sample Weight Guidance

Xianling Yang[1], Zhiwen Yu[1, 2], Sun Song[3], Kaixiang Yang[1*]
[1]School of Computer Science and Engineering, South China University of Technology
[2]Pengcheng Laboratory
[3]College of Computer and Information Science, Chongqing Normal University
xl_yang2023@163.com, zhwyu@scut.edu.cn, sunsong@cqnu.edu.cn, yangkx@scut.edu.cn

Abstract

Semi-supervised learning (SSL) based on pseudo-label and consistency has achieved significant success. The core idea behind these methods is to assign sample weights based on pseudo-label probabilities, thereby guiding the model toward biased learning. However, existing research still faces two major challenges in guiding learning: (1) how to evaluate learning states across different classes in the absence of labels, and (2) how to construct an effective sample weight space that provides precise guidance throughout training. To address these challenges, we propose the Bi-Dimensional Sample Weight Guidance algorithm, BidMatch. BidMatch introduces **C**lass **I**nformation **E**ntropy (CIE), which utilizes pseudo-label information entropy to capture inter-class learning relationships, thereby enriching the representation of learning states across different classes under unlabeled conditions. Additionally, **P**seudo-label **P**robability **R**edistribution (PPR) is proposed to maintain distribution invariance and sparsity during training, thereby emphasizing differences in instance importance. By leveraging CIE and PPR, BidMatch generates sample weights that account for both class and instance dimensions, effectively guiding the model toward balanced and efficient learning across classes. BidMatch has demonstrated state-of-the-art performance on various SSL datasets. Notably, it achieved a 6.45% error rate on CIFAR-10 with only one label per class, significantly outperforming baseline methods.

Introduction

Semi-supervised learning(SSL) leverages limited labeled data and abundant unlabeled data to optimize the learning process, thereby improving model performance on small labeled sample datasets. Deep SSL combines the SSL paradigm with the powerful feature extraction capabilities of neural networks, enhancing the model's learning ability by fully utilizing large amounts of unlabeled data, ultimately achieving better results than supervised learning(Yang et al. 2022b; Oliver et al. 2018). These learning methods demonstrate significant application value in scenarios where resources are limited or acquiring labeled data is challenging(Zhang et al. 2025; Ferman, Garrido, and Bharaj 2024; Li et al. 2024; Wang et al. 2024).

In recent years, combination methods based on pseudo-label and consistency regularization, along with their extensions (Sohn et al. 2020; You et al. 2024; Li, Xiong, and Hoi 2021; Zheng et al. 2022), have achieved great success. These methods primarily rely on setting sample weights based on pseudo-label probabilities to guide the model toward consistent alignment with a bias. Among these, FixMatch(Sohn et al. 2020), which first proposed this idea, uses the prediction values of weakly augmented views of unlabeled samples as class confidence and selects high-confidence unlabeled samples for consistency regularization. This approach significantly improves semi-supervised learning performance without increasing computational complexity. However, FixMatch simply maps the sample weights of unlabeled data to a discrete space based on a pre-set threshold, neglecting the varying importance of classes and samples across different training stages. To address this issue, SoftMatch (Chen et al. 2023) assigns a weight of 1 to samples with pseudo-label probabilities greater than the mean, while mapping the weights of other samples to a Gaussian space. FlexMatch (Zhang et al. 2021) assigns probabilities to different classes based on the pseudo-label probabilities of weakly augmented views of unlabeled samples, but still assigns sample weights of either 0 or 1. While these methods have advanced combination SSL, they still face two key limitations: (1) current methods rely solely on pseudo-label to determine the learning states of different classes, lacking inter-class relationships and struggling to balance the learning difficulty across classes, and (2) current methods map sample weights to a simple discrete space, failing to reflect the varying importance of different samples under different training conditions, thereby hindering effective guidance.

To address the limitations of combination SSL in sample weight guidance, we propose a bi-dimensional sample weight guidance algorithm, BidMatch, which simultaneously considers both the learning state of each class and the importance of individual instances. BidMatch introduces Class Information Entropy (CIE) to address the challenge of varying learning difficulty across classes by characterizing learning states among samples of different classes. Unlike directly using pseudo-label probabilities, CIE captures inter-class learning relationships, offering a more nuanced understanding of class-level learning difficulty. Addition-

*Corresponding author.

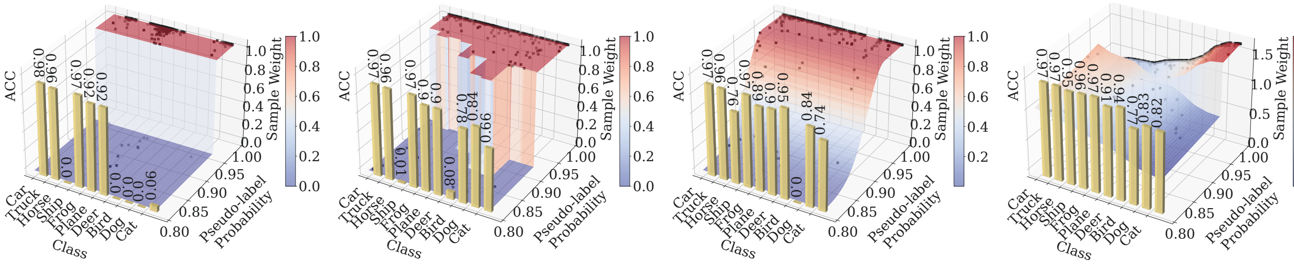

(a) FixMatch Avg=0.48. (b) FlexMatch Avg=0.71. (c) SoftMatch Avg=0.79. (d) BidMatch Avg=0.91.

Figure 1: Illustration of the training state, performance metrics, and sample weight distribution for different models after 100k iterations on the CIFAR-10 dataset with 10 labels. In this visualization, the three-dimensional surface represents the sample weights, while the scatter points depict the distribution of pseudo-labels. The bar chart illustrates the model's classification performance across each class on the test set, with Avg indicating the average accuracy. Compared to baseline methods, BidMatch achieves a well-balanced pseudo-label distribution and demonstrates consistently superior overall performance.

ally, BidMatch proposes Pseudo-label Probability Redistribution(PPR) for weakly augmented views, ensuring both distribution invariance and sparsity in training, which enhances sample importance assessment. Based CIE and PPR, BidMatch constructs a continuous sample weight space, allowing the model to engage in targeted learning, guided by refined sample weights, for improved learning performance. Figure 1 illustrates the sample weights, pseudo-label distribution, and performance across different classes during training for weight-guided semi-supervised methods. Compared to baseline methods, BidMatch distributes unlabeled samples more evenly across different classes during training, effectively mitigating semantic drift and achieving exceptional performance.

In summary, this paper presents the following contributions:

- We propose **C**lass **I**nformation **E**ntropy, which utilizes pseudo-label information entropy to capture inter-class learning relationships, thereby enriching the representation of learning states across different classes under unlabeled conditions.
- We introduce **P**seudo-label **P**robability **R**edistribution to evaluate the importance of individual instances under varying training status. By incorporating both class and instance dimensions, we construct a sample weight space that effectively guides the model's learning process.
- Extensive experiments on various SSL datasets demonstrate that BidMatch achieves state-of-the-art performance and effectively guides the model to maintain stable training while balancing learning across different classes.

Related Work

This section begins with a review of classical SSL, including pseudo-label and consistency regularization, and then discusses related research that combines these two approaches through sample weight guidance.

Pseudo-Label and Consistency Regularization.

The core idea of pseudo-label methods(Lee et al. 2013; Zhou and Li 2010)is to use the model's predictions on unlabeled samples as labels, treating these pseudo-labeled data as part of the training set, thus enhancing the information available to the model during training. For example, SimCLRv2(Chen et al. 2020) generates pseudo-labels using a pre-trained model, while MPL(Pham et al. 2021) uses a validation set to prevent distribution drift in unlabeled data. Consistency regularization methods(Ke et al. 2019; Pezeshki et al. 2016), based on the manifold hypothesis or smoothness assumption, argue that samples with similar distributions in feature space should produce identical outputs, describing methods in which small perturbations of data should not alter the model's output. Mean Teacher(Tarvainen and Valpola 2017) performs alignment via a teacher-student model, VAT(Miyato et al. 2018) trains by applying adversarial perturbations to the samples, and NICE(Ni and Koniusz 2023) uses generative networks to augment the learning data. Recent advancements in deep learning have spurred extensive research, building on these fundamental techniques to explore areas like data augmentation(Cubuk et al. 2020; Kurakin et al. 2020; Miyato et al. 2018; Sajjadi, Javanmardi, and Tasdizen 2016; Xie et al. 2020), model architecture(Donahue, Krähenbühl, and Darrell 2016; Dong-DongChen and WeiGao 2018; Gan et al. 2017; Laine and Aila 2016; Abuduweili et al. 2021), and loss function design(Iscen et al. 2019; Ren, Yeh, and Schwing 2020; Lee et al. 2022; Yang et al. 2023).

Sample Weight-Guided Combination Methods.

FixMatch(Sohn et al. 2020) initially integrated pseudo-label with consistency regularization by mapping sample weights to a discrete space based on pseudo-label probabilities, which led to substantial performance improvements in semi-supervised algorithms. Subsequent studies expanded on this approach by incorporating techniques like graph embeddings(Li, Xiong, and Hoi 2021; Zheng et al. 2023; Sun, Shi, and Li 2023), generative models(You et al. 2024; Hong et al. 2023) and contrastive learning(Yang et al. 2022a; Zhou et al.

2024), further enhancing the combination-based method. In sample weight guidance, influence functions(Koh and Liang 2017; Ren, Yeh, and Schwing 2020; Kong, Shen, and Huang 2022) can compute sample weights based on model performance, but the lack of labeled data and high computational costs limit their scalability. Methods such as FlexMatch(Zhang et al. 2021) and SoftMatch(Chen et al. 2023) extend FixMatch by improving sample weight. They set different thresholds for various classes based on the class pseudo-label probability, aiming to address the issue of varying learning difficulties across classes. The sample weights are mapped to a truncated Gaussian space to improve the utilization of unlabeled samples.

However, these methods still face challenges, such as overly simplistic sample weight assignments and difficulty in distinguishing the learning states between classes. Our work is precisely centered around these challenges, offering a more effective sample weight guidance method.

Methodology

In this section, we first review the preliminary work on sample weight-guided combination methods, followed by a detailed description of BidMatch. We begin by defining a series of notations as follows: let $\mathcal{D}_l = \{x_{li}, y_{li}\}_{i=1}^{N_L}$ denote the set of labeled samples and $\mathcal{D}_u = \{x_{ui}\}_{i=1}^{N_U}$ the set of unlabeled samples. The unlabeled samples undergo two types of augmentations, resulting in weakly augmented sample x_{ui}^w and strongly augmented sample x_{ui}^s. We represent the predicted probability distribution for a sample after t iterations as $p_t\left(y\left|x\right.\right)$ and denote the argmax of this distribution as a one-hot label $\hat{p}_t\left(y\left|x\right.\right)$. The cross-entropy between the two distributions is represented as $\mathcal{H}(\cdot,\cdot)$.

Preliminary

Based on pseudo-label methods, the model's predicted results are utilized as labels for unlabeled data during supervised training. The loss function for consistency regularization methods encompasses both labeled loss and consistency regularization loss: $\mathcal{L} = \mathcal{L}_s + \mathcal{L}_u$. Specifically, for a training batch with a batch size of B, the supervised loss is defined as:

$$\mathcal{L}_s = \frac{1}{B}\sum_{i=1}^{B}\mathcal{H}\left(\hat{y}_i, p_t\left(y\left|x_i\right.\right)\right), \tag{1}$$

where $\hat{y}_i$ represents the one-hot vector of the true label. The consistency regularization method posits that the predicted probability distributions for weakly and strongly augmented unlabeled samples should be similar, defining the consistency regularization loss as the cross-entropy between the two distributions(Sajjadi, Javanmardi, and Tasdizen 2016; Laine and Aila 2016):

$$\mathcal{L}_u = \frac{1}{\eta B}\sum_{i=1}^{\eta B}\mathcal{H}\left(p_t\left(y\left|x_{ui}^w\right.\right), p_t\left(y\left|x_{ui}^s\right.\right)\right), \tag{2}$$

where η is the ratio of unlabeled samples to labeled samples.However, during the training phase, the model may perform poorly on certain samples. Applying consistency regularization indiscriminately to all unlabeled samples can mislead the model, thus reducing convergence speed and overall performance. To address this, FixMatch introduces a confidence threshold τ, mapping sample weights to a discrete space where samples exceeding the threshold τ are assigned a value of 1, followed by consistency regularization.

$$\begin{aligned}\mathcal{L}_u = &\frac{1}{\eta B}\sum_{i=1}^{\eta B}\mathcal{H}\left(\hat{p}_t\left(y\left|x_{ui}^w\right.\right), p_t\left(y\left|x_{ui}^s\right.\right)\right)\\ &\times \mathbb{1}\left(\max\left(p_t\left(y\left|x_{ui}^w\right.\right)\right) > \tau\right).\end{aligned} \tag{3}$$

FixMatch ignores the varying importance of classes and samples across different training stages. FlexMatch designs adaptive class thresholds based on the differences in pseudo-label probabilities across iteration t and class c, assigning different threshold $\tau(c)$ to each class:

$$\begin{aligned}\sigma\left(c\right) = &\sum_{i=1}^{\eta B}\mathbb{1}\left(p_t\left(y\left|x_{ui}^w\right.\right) > \tau\right)\\ &\times \mathbb{1}\left(\arg\max p_t\left(y\left|x_{ui}^w\right.\right) = c\right),\end{aligned} \tag{4}$$

$$\tau_c = \frac{\sigma\left(c\right)}{\max\limits_{c}\sigma}\tau. \tag{5}$$

SoftMatch maps sample weights to a Gaussian space, employing a truncated Gaussian distribution for weight assignment:

$$\lambda_{ui,t} = \begin{cases}\exp\left(-\frac{\left(\max\left(p_t\left(y\,|\,x_{ui}^w\right)\right)-\mu_t\right)^2}{2\hat{\sigma}_t^2}\right) & \\ \qquad\qquad \text{if } \max\left(p_t\left(y\,|\,x_{ui}^w\right)\right) < \mu_t \\ 1 \qquad\qquad\qquad\qquad \text{otherwise,}\end{cases} \tag{6}$$

where μ_t is the mean and σ_t is the standard deviation of the distribution, and $\lambda_{ui,t}$ represents the weight for x_{ui}.

BidMatch builds upon prior research and addresses the limitations of existing sample weight guidance approaches. To handle varying learning difficulties across classes, BidMatch introduces CIE, which characterizes the learning states of different classes and captures inter-class learning relationships. Furthermore, BidMatch redistributes pseudo-label probabilities to assess the importance of individual instances. Based the measures of class-level and instance-level, BidMatch constructs a bi-dimensional sample weight space that provides targeted guidance to the model. The overall process of BidMatch is depicted in Figure 2, with each component detailed in the following sections.

Class Information Entropy

The differential learning difficulties among classes presents a fundamental challenge, particularly in SSL scenarios with extremely limited labeled data. In existing research, the learning state of a class is typically measured by the expectation $\mathbb{E}(\max(p_c^t(y|x_{ui}^w)))$ of the pseudo-label probability distribution across all pseudo-labels of class c within a batch,

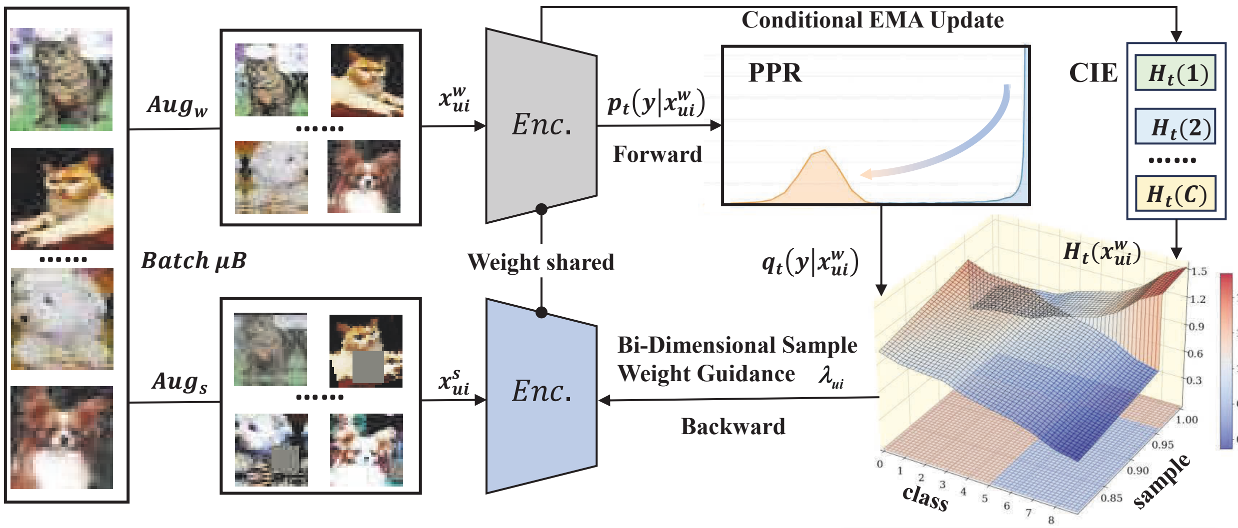

Figure 2: The BidMatch framework overview. For each training batch, BidMatch computes CIE and PPR based on the prediction probabilities of weakly-augmented views for each sample. Conditional updates of CIE are then performed using an EMA. With the calculated CIE and PPR values, BidMatch constructs a sample weight space to effectively guide the model's learning process.

which fails to capture the learning relationships among unlabeled samples across different classes. To address this limitation, we propose using CIE to assess the learning difficulty of different classes during training.

CIE relies on the predicted probabilities of weakly augmented views. Within each training batch, we first calculate the information entropy for each unlabeled sample:

$$h\left(x^w_{ui}\right) = -p_t\left(y\left|x^w_{ui}\right.\right) \cdot \log\left(p_t\left(y\left|x^w_{ui}\right.\right)\right). \tag{7}$$

The class with the maximum predicted probability from the weakly augmented view is assigned as the sample's pseudo-label, and the sample's information entropy is used to gauge the learning status of that class. Due to individual variations among samples, using information entropy for each instance can lead to frequent updates of the class learning status, which may be unstable and inaccurate. To mitigate this issue, we draw inspiration from the stochastic gradient descent optimization method. Within each training batch, we first count the number of pseudo-labels for each class:

$$n_t\left(c\right) = \sum_{i=1}^{\eta B} \mathbb{1}\left(\arg\max p_t\left(y\left|x^w_{ui}\right.\right) = c\right). \tag{8}$$

The CIE $h_c(t)$ for class c is defined as the average information entropy of all pseudo-labeled samples belonging to class c in the current batch, and is calculated based on the sample information entropy $h(x^w_{ui})$ and the number of pseudo-labels n_t:

$$h_t\left(c\right) = \frac{1}{n_t\left(c\right)} \sum_{i=1}^{\eta B} h\left(x^w_{ui}\right) \times \mathbb{1}\left(\arg\max p_t\left(y\left|x^w_{ui}\right.\right) = c\right). \tag{9}$$

Due to significant variations in information entropy values across different datasets and training stages, the generalization ability of class information entropy needs to be enhanced. To address this, the method normalizes the CIE for each batch:

$$H_t\left(c\right) = \frac{h_t\left(c\right)}{\max\limits_{c} h_t}, \tag{10}$$

where $H_t(c)$ represents the learning status of class c in the CIE H_t. Compared to the class pseudo-label probability $\mathbb{E}(\max(p^t_c(y|x^w_{ui})))$, $H_t(c)$ incorporates information from all classes and captures the learning relationships between classes, which offers more comprehensive information to guide model training.

Bi-Dimentional Sample Weight

Current combination methods typically map sample weights to a simple, discrete space based solely on pseudo-label confidence, which fails to capture the nuanced importance of samples across different classes and training stages. In contrast, BidMatch first redistributes pseudo-label probabilities to help the model better distinguish the significance of each sample. It then constructs sample weights by incorporating the CIE module, which effectively guides the model's learning dynamics.

Pseudo-Label Probability Redistribution. The purpose of redistributing pseudo-label probabilities is to help the model identify the relative importance of individual instances. Due to the absence of true instance labels, previous research typically uses the maximum prediction probability of a weakly augmented view as an indicator of pseudo-label quality. However, directly relying on maximum prediction probability does not suffice across different training stages.

For instance, during early training, when the model has not yet fully learned, maximum prediction probabilities tend to be low, limiting the effective use of unlabeled samples and making it difficult to select those that positively guide the model. In later stages, most maximum prediction values for unlabeled samples approach 1, leading to a densely packed distribution that fails to effectively emphasize sample importance.

To address this, BidMatch standardizes and adjusts the distribution of maximum prediction probabilities. Following prior studies, it first applies Distribution Alignment $DA(\cdot)$ on weakly augmented views to mitigate early-stage imbalances in pseudo-label distribution(Kurakin et al. 2020; Chen et al. 2023). Then, it normalizes the aligned pseudo-label probabilities to calculate a sample importance coefficient:

$$q_t\left(y\left|x_{ui}^w\right.\right) = \frac{\max\left(DA\left(p_t\left(y\left|x_{ui}^w\right.\right)\right)\right) - \mu_t}{\sigma_t}, \tag{11}$$

where μ_t and σ_t represent the mean and variance of the pseudo-label probabilities, updated per batch using EMA. This standardized redistribution offers two primary benefits for assessing sample importance: (1)it ensures that pseudo-label probability distribution remains consistent as training progresses, and (2) it sparsifies the previously dense distribution, enabling the model to differentiate between samples more effectively.

Sample Weight Construction. BidMatch assesses the learning difficulty across classes and the importance of individual instances based on the CIE and PPR modules. First, the unlabeled sample x_{ui} is matched to the CIE value of the corresponding class based on its pseudo-label. Since the original CIE values are all positive, they are insensitive to the actual learning status of different classes. Inspired by the advantage function in reinforcement learning, BidMatch enhances sensitivity by subtracting a baseline value from the original CIE. Specifically, it uses the expected entropy $\mathbb{E}(H_t)$ as the baseline.

$$H_t\left(x_{ui}\right) = H_t \cdot \hat{p}_t\left(y\left|x_{ui}^w\right.\right) - \mathbb{E}\left(H_t\right). \tag{12}$$

The instance value $q_t(y|x_{ui}^w)$ represents the confidence in the pseudo-label. For low-confidence samples, where $q_t(y|x_{ui}^w) < 0$ indicates low pseudo-label reliability, the sample weight should be adjusted accordingly in the opposite direction. Therefore, when constructing sample weights, the corresponding CIE value associated with the instance should also be reversed to reflect this uncertainty. By considering both individual instances and class learning states, BidMatch assigns weights as follows:

$$\lambda_{ui,t} = \begin{cases} \exp\left(q_t\ \left(y\left|x_{ui}^w\right.\right) T + H_t\left(x_{ui}\right)\right) \\ \qquad\qquad\qquad\qquad \text{if } q_t\left(y\left|x_{ui}^w\right.\right) > 0 \\ \exp\left(q_t\left(y\left|x_{ui}^w\right.\right) T^{-1} - H_t\left(x_{ui}\right)\right) \ \text{otherwise,} \end{cases} \tag{13}$$

where T is the temperature factor used to balance the guidance provided by PPR and CIE, and to mitigate interference caused by misclassified samples during training. Table 1 uses q and H to denote the instance and class dimensions, respectively. Note that $|H|$ is generally larger than $|q|$, and a larger value of H indicates that the corresponding class is less well learned. BidMatch adopts a progressive optimization strategy, the sample weight λ_{ui} approximately highlights two key regions: (1) When $H > 0$, the class is underlearned, and BidMatch promotes learning from high-confidence instances where $q > 0$ to improve model reliability. (2) When $H < 0$, the corresponding class is already well learned. In this case, BidMatch encourages the model to focus on harder instances with $q < 0$ within that class. This progressive learning strategy, moving from easy to difficult individual instances, aligns closely with the human cognitive process. Further analysis can be found in the Supplementary Material.

Dimentional		**Instance**	
		$q > 0$	$q < 0$
Class	$H > 0$	$\lambda_{ui} = e^{\lvert q\rvert + \lvert H\rvert}$ high weight	$\lambda_{ui} = e^{-\lvert H\rvert - \lvert q\rvert}$ low weight
	$H < 0$	$\lambda_{ui} = e^{\lvert q\rvert - \lvert H\rvert}$ low weight	$\lambda_{ui} = e^{\lvert H\rvert - \lvert q\rvert}$ high weight

Table 1: Approximate visualization of the sample weight distribution guided by bidMatch as defined in Equation 13

Implementation

To enhance model stability during gradient descent optimization, BidMatch applies Exponential Moving Average (EMA) updates to both the model parameters and the CIE within each training batch. However, in the early stages of training, some classes may lack pseudo-labels, resulting in a CIE value of zero and potentially leading to erroneous updates. To address this issue, BidMatch proposes a conditional EMA updating strategy, in which only the classes with pseudo-labels present in the current batch are updated.

$$H_t(c) = \begin{cases} \alpha H_{t-1}(c) + (1-\alpha) H_{t'}(c) & \text{if } H_{t'}(c) > 0 \\ H_t(c) & \text{otherwise,} \end{cases} \tag{14}$$

where $H_{t'}(c)$ denotes the CIE for class c calculated in the current batch, while $H_{t-1}(c)$ represents the EMA result from the previous iteration. In each training batch, we assign exclusive weights to the samples, with the unsupervised loss function and the complete loss function defined as follows:

$$\mathcal{L}_u = \frac{1}{\mu B}\sum_{i=1}^{\eta B} \lambda_{ui,t} \cdot \mathcal{H}\left(\hat{p}_t\left(y\left|x_{ui}^w\right.\right), p_t\left(y\left|x_{ui}^s\right.\right)\right), \tag{15}$$

$$\mathcal{L} = \mathcal{L}_s + \lambda \cdot \mathcal{L}_u. \tag{16}$$

The overall implementation flow of the method is outlined in the supplementary material.

Experiments

Experimental Setup

Experiments on the BidMatch method were conducted using four datasets: CIFAR-10, CIFAR-100(Krizhevsky, Hinton et al. 2009), STL-10(Coates, Ng, and Lee 2011),

Dataset	CIFAR-10			CIFAR-100			SVHN		STL-10	
Labels	**10**	**40**	**250**	**400**	**2500**	**10000**	**40**	**1000**	**40**	**1000**
Fully-Supervised	4.62±0.05			19.30±0.09			2.13±0.02		NONE	
MeanTeacher	79.18±1.11	70.09±1.60	37.46±3.30	81.11±1.44	45.17±1.06	31.75±0.23	36.09±3.98	3.27±0.05	71.72±1.45	33.90±1.37
PseudoLabel	80.21±0.55	74.61±0.26	46.49±2.20	87.45±0.85	57.74±0.28	36.55±0.24	64.61±5.60	9.40±0.32	74.68±0.99	34.14±0.71
VAT	79.81±1.17	74.66±2.12	41.03±1.79	85.20±1.40	46.84±0.79	32.14±0.19	74.74±0.38	4.11±0.20	74.74±0.38	37.95±1.12
MixMatch	65.76±7.06	36.19±6.48	13.63±0.59	67.59±0.66	39.76±0.48	27.78±0.29	30.60±8.39	3.69±0.37	54.93±0.96	24.70±0.68
UDA	34.53±10.69	10.62±1.75	5.16±0.06	46.39±1.59	27.73±0.21	22.49±0.23	5.12±4.27	**1.89±0.01**	37.42±8.44	6.64±0.17
FixMatch	24.79±7.65	7.47±0.28	8.68±0.05	46.42±1.76	28.03±0.16	22.20±0.12	3.81±1.18	1.96±0.03	35.97±4.14	6.25±0.33
MPL	23.55±6.01	6.62±0.91	5.76±0.24	46.26±1.84	27.71±0.19	21.74±0.09	9.33±8.02	2.28±0.04	35.76±4.83	6.66±0.00
FlexMatch	21.08±7.23	4.97±0.06	4.98±0.09	39.94±1.62	26.49±0.20	21.93±0.15	8.19±3.20	6.72±0.30	29.15±4.16	5.77±0.18
SimMatch	15.81±10.36	5.38±0.01	5.05±0.08	39.32±0.72	26.21±0.37	21.50±0.11	7.06±2.21	2.05±0.05	16.98±4.34	5.84±0.31
SoftMatch	15.77±5.15	4.91±0.12	4.82±0.09	37.10±0.77	26.66±0.25	22.03±0.03	2.33±0.25	2.01±0.01	21.42±3.48	5.73±0.24
ReFixMatch	32.62±14.89	4.95±0.05	4.85±0.06	46.73±1.37	27.25±0.25	21.78±0.04	2.63±1.46	2.01±0.05	28.66±4.40	5.76±0.42
EPASS	22.69±4.31	5.31±0.10	5.08±0.05	38.88±0.24	**25.68±0.33**	21.32±0.14	2.31±0.04	2.02±0.02	15.71±2.48	5.58±0.04
BidMatch(Ours)	**6.45±1.11**	**4.78±0.05**	**4.71±0.04**	**36.27±0.63**	25.91±0.33	**21.24±0.10**	**2.22±0.36**	2.01±0.06	**13.75±1.45**	**5.26±0.11**

Table 2: Test error rates(%) from comparative experiments with varying numbers of labels, evaluated on the CIFAR-10, CIFAR-100, SVHN, and STL-10 datasets. The best number is in bold.

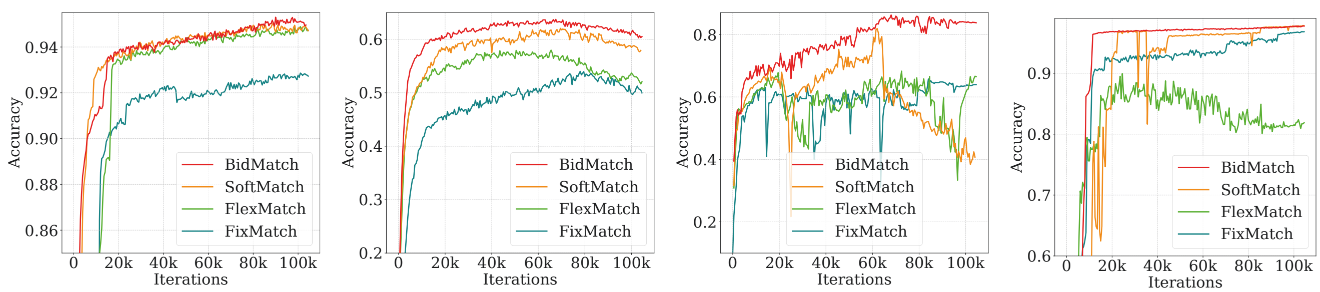

(a) CIFAR-10 (b) CIFAR-100 (c) STL-10 (d) SVHN

Figure 3: The performance of four sample weight guided semi-supervised methods was evaluated across different datasets. During the model training, testing was conducted on the test set every 5,000 iterations, for a total of 2^{20} iterations. Each dataset used in the experiments contained 40 labels.

and SVHN(Netzer et al. 2011), along with the ImageNet dataset(Deng et al. 2009). The network architecture utilized for the smaller datasets was a Wide Residual Network(Zagoruyko and Komodakis 2016) . Specifically, WRN-28-2 was applied to CIFAR-10 and SVHN, WRN-28-8 was used for CIFAR-100, and WRN-37-2 was employed for STL-10. For ImageNet, the ResNet-50 architecture(He et al. 2016) was adopted. Both weak and strong augmentations were applied to the unlabeled images(Cubuk et al. 2020). Detailed experimental settings are provided in the supplementary material.

CIFAR-10/100, SVHN, and STL-10.

The experiment compared BidMatch with classical deep SSL approaches, such as UDA(Xie et al. 2020) and MPL(Pham et al. 2021). Furthermore, we included semi-supervised methods that incorporate sample weighting, including SoftMatch(Chen et al. 2023), FlexMatch(Zhang et al. 2021), and FixMatch(Sohn et al. 2020), as well as current state-of-the-art methods in standard SSL, such as RefixMatch (Nguyen and Yang 2023) and EPASS (Nguyen 2024). All comparison methods were evaluated under the same experimental settings as BidMatch.

The results of the comparative experiments are presented in Table 2, our method achieved superior performance in 9 out of the 10 configurations. The only exceptions were on the CIFAR-100 dataset with 2500 labels and the SVHN dataset with 1000 labels, where our method performed comparably to the best methods. Furthermore, BidMatch demonstrated greater advantages as the label ratio decreased. For instance, on the CIFAR-10 dataset with only 10 labels, our method outperformed FixMatch by 18.34%. On the STL-10 dataset with 40 labels, BidMatch achieved a 21.60% improvement over FixMatch. These findings indicate that our method effectively mitigates semantic drift issues in SSL scenarios with extremely scarce labels.

Figure 3 presents a comparative analysis of the complete training trajectories of BidMatch and baseline methods. It is evident that BidMatch exhibits strong training stability. In contrast, the baseline methods suffer from performance collapse in the later stages on multiple datasets. However, BidMatch maintains consistently high performance, indicating that its weight-guided learning strategy is beneficial to the overall training process.

ImageNet

The experiment was further extended to the ImageNet dataset, which contains 100k labeled samples, to assess the

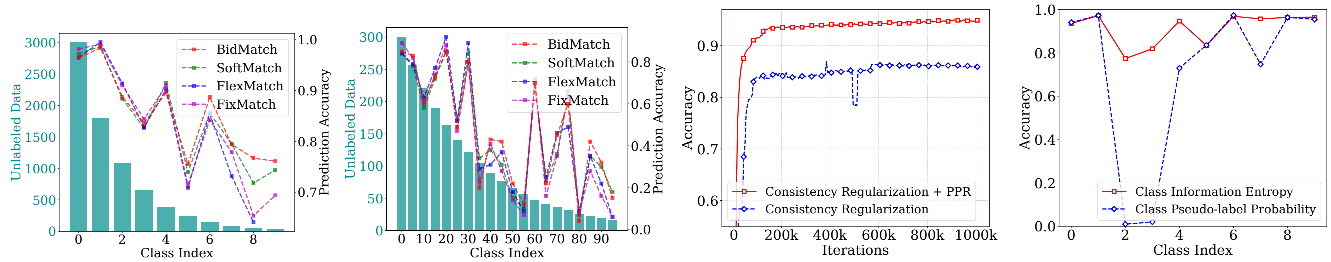

(a) (b) (c) (d)

Figure 4: (a) and (b) experimental results comparison on CIFAR-10 and CIFAR-100 under long-tail settings. (c) the performance of the CR and PPR methods on the CIFAR-10 dataset with 40 labels. (d) the classification accuracy of CIT and class pseudo-label probability across different classes on the CIFAR-10 dataset with 10 labels.

Method	Top-1	Top-5
FixMatch	43.66	20.80
FlexMatch	41.85	19.48
SoftMatch	40.52	-
BidMatch	**40.04**	**18.81**

Table 3: Top1 error rate (%) on ImageNet with 100k lables. The best number is in bold.

performance of BidMatch under high data complexity. Owing to computational resource constraints, this experiment was conducted only once. Table 3 presents the comparative results between BidMatch and other sample weight-guided semi-supervised methods, demonstrating that BidMatch achieved the best performance in both Top-1 and Top-5 accuracy, highlighting its robustness in handling complex datasets.

	CIE	PPR	CIFAR10	CIFAR100
CR			13.08 $\pm$2.13	51.56 $\pm$ 0.22
CR+CIE	✓		8.69 $\pm$0.26	47.97 $\pm$ 0.87
CR+PPR		✓	4.91 $\pm$0.08	36.71 $\pm$ 0.62
BidMatch	✓	✓	**4.78$\pm$0.05**	**36.27$\pm$ 0.63**

Table 4: Experimental results on CIFAR-10 with 40 labels and CIFAR-100 with 400 labels, highlighting the effectiveness of CIT and PPR in conjunction with consistency regularization(CR).

Ablation Study

We conducted an ablation study focusing on the contributions of CIE and PPR in the following methods: (1) the consistency regularization method in which all sample weights are set to 1, (2) the CIE module, which retains only $H_t\left(x_{ui}\right)$ in Eq.13, (3) the PPR module, which retains only $q_t\ \left(y\left|x_{ui}^w\right.\right)$ in Eq.13, and (4) the complete BidMatch method.

The results in Table 4 indicate that integrating both CIE and PPR is essential to achieve optimal performance. Although the contribution of CIE is relatively smaller than that of PPR, this is primarily because PPR is applied to all unlabeled samples, and its fine-grained weighting plays a critical role throughout the training process, serving as the main driver of performance improvements in BidMatch. Nevertheless, the inclusion of CIE enables more targeted training, which can further enhance model performance. As shown in Figure 4a and 4b, BidMatch can effectively balance learning across different classes even under long-tailed conditions, further validating the necessity of the CIE module.

Pseudo-Label Probability Redistribution: To further assess the effectiveness of PPR during training, we compare the performance of the consistency regularization method with that of a model trained with the PPR module. The results in Figure 4c show that when the model uses PPR to assess the importance of different samples, it significantly improves the performance of the traditional consistency regularization method. This suggests that the impact of different samples on the model varies, and assigning distinct weights to samples during training is crucial.

Class Information Entropy: To further demonstrate the superiority of CIE in leveraging inter-class learning information, we compare it with a classical class-state learning computation method, class pseudo-label probability. Specifically, we replace the $H_t(c)$ term in Equation 13 with $1 - \mathbb{E}(\max(p_c^t(y|x_{ui}^w)))$. The results in Figure 4d indicate that when labeled data is extremely limited, using class pseudo-label probability may even lead to semantic drift for all samples in certain classes. This finding further confirms that relying solely on pseudo-label probabilities makes it challenging to accurately guide the learning status of different classes. Further analysis and visualizations are provided in the supplementary material.

Conclusion

This paper addresses two major challenges in SSL and proposes a bi-dimensional sample weight guidance method. Built upon CIE and PPR, BidMatch constructs a comprehensive sample weight space. Extensive experiments demonstrate the superiority of BidMatch in guiding the model training process effectively. Future work could involve a deeper theoretical exploration of sample weighting, as well as its extension to other learning paradigms.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China (No. 62476101, 92467109, 62572199, U21A20478), National Key R&D Program of China 2023YFA1011601, the Guangdong Basic and Applied Basic Research Foundation under Grant 2024A1515140137, the Guangzhou Science and Technology Plan Project 2024A04J3749, Project Supported by Scientific and Technological Research Program of Chongqing Municipal Education Commission(Grant No.KJQN202300547) and also supported by Chongqing Municipal Construction Science and Technology Plan Project(Chengke Zi 2024 No. 8-7) and the Fundamental Research Funds for the Central Universities (2024ZYGXZR062).

References

Abuduweili, A.; Li, X.; Shi, H.; Xu, C.-Z.; and Dou, D. 2021. Adaptive Consistency Regularization for Semi-Supervised Transfer Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 6923–6932.

Chen, H.; Tao, R.; Fan, Y.; Wang, Y.; Wang, J.; Schiele, B.; Xie, X.; Raj, B.; and Savvides, M. 2023. SoftMatch: Addressing the Quantity-Quality Tradeoff in Semi-supervised Learning. In *The Eleventh International Conference on Learning Representations*.

Chen, T.; Kornblith, S.; Swersky, K.; Norouzi, M.; and Hinton, G. E. 2020. Big self-supervised models are strong semi-supervised learners. *Advances in neural information processing systems*, 33: 22243–22255.

Coates, A.; Ng, A.; and Lee, H. 2011. An analysis of single-layer networks in unsupervised feature learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, 215–223. JMLR Workshop and Conference Proceedings.

Cubuk, E. D.; Zoph, B.; Shlens, J.; and Le, Q. V. 2020. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, 702–703.

Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.

Donahue, J.; Krähenbühl, P.; and Darrell, T. 2016. Adversarial feature learning. *arXiv preprint arXiv:1605.09782*.

Dong-DongChen, W.; and WeiGao, Z. 2018. Tri-net for semi-supervised deep learning. In *Proceedings of twenty-seventh international joint conference on artificial intelligence*, 2014–2020.

Ferman, D.; Garrido, P.; and Bharaj, G. 2024. FaceLift: Semi-supervised 3D Facial Landmark Localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1781–1791.

Gan, Z.; Chen, L.; Wang, W.; Pu, Y.; Zhang, Y.; Liu, H.; Li, C.; and Carin, L. 2017. Triangle generative adversarial networks. *Advances in neural information processing systems*, 30.

He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.

Hong, C.; Cha, B.; Kim, B.; and Oh, T.-H. 2023. Enhancing Classification Accuracy on Limited Data via Unconditional GAN. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, 1057–1065.

Iscen, A.; Tolias, G.; Avrithis, Y.; and Chum, O. 2019. Label propagation for deep semi-supervised learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5070–5079.

Ke, Z.; Wang, D.; Yan, Q.; Ren, J.; and Lau, R. W. 2019. Dual student: Breaking the limits of the teacher in semi-supervised learning. In *Proceedings of the IEEE/CVF international conference on computer vision*, 6728–6736.

Koh, P. W.; and Liang, P. 2017. Understanding black-box predictions via influence functions. In *International conference on machine learning*, 1885–1894. PMLR.

Kong, S.; Shen, Y.; and Huang, L. 2022. Resolving Training Biases via Influence-based Data Relabeling. In *International Conference on Learning Representations*.

Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images.

Kurakin, A.; Raffel, C.; Berthelot, D.; Cubuk, E. D.; Zhang, H.; Sohn, K.; and Carlini, N. 2020. Remixmatch: Semi-supervised learning with distribution matching and augmentation anchoring. *In International Conference on Learning Representations*.

Laine, S.; and Aila, T. 2016. Temporal ensembling for semi-supervised learning. *arXiv preprint arXiv:1610.02242*.

Lee, D.; Kim, S.; Kim, I.; Cheon, Y.; Cho, M.; and Han, W.-S. 2022. Contrastive Regularization for Semi-Supervised Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 3911–3920.

Lee, D.-H.; et al. 2013. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, volume 3, 896. Atlanta.

Li, G.; Yu, Z.; Yang, K.; Chen, C. P.; and Li, X. 2024. Ensemble-Enhanced Semi-Supervised Learning With Optimized Graph Construction for High-Dimensional Data. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.

Li, J.; Xiong, C.; and Hoi, S. C. 2021. Comatch: Semi-supervised learning with contrastive graph regularization. In *Proceedings of the IEEE/CVF international conference on computer vision*, 9475–9484.

Miyato, T.; Maeda, S.-i.; Koyama, M.; and Ishii, S. 2018. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE transactions* *on pattern analysis and machine intelligence*, 41(8): 1979–1993.

Netzer, Y.; Wang, T.; Coates, A.; Bissacco, A.; Wu, B.; Ng, A. Y.; et al. 2011. Reading digits in natural images with unsupervised feature learning. In *NIPS workshop on deep learning and unsupervised feature learning*, volume 2011, 4. Granada.

Nguyen, K.-B. 2024. Debiasing, calibrating, and improving Semi-supervised Learning performance via simple Ensemble Projector. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2441–2451.

Nguyen, K.-B.; and Yang, J.-S. 2023. Boosting Semi-Supervised Learning by bridging high and low-confidence predictions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1028–1038.

Ni, Y.; and Koniusz, P. 2023. NICE: NoIse-modulated Consistency rEgularization for Data-Efficient GANs. In Oh, A.; Naumann, T.; Globerson, A.; Saenko, K.; Hardt, M.; and Levine, S., eds., *Advances in Neural Information Processing Systems*, volume 36, 13773–13801. Curran Associates, Inc.

Oliver, A.; Odena, A.; Raffel, C. A.; Cubuk, E. D.; and Goodfellow, I. 2018. Realistic evaluation of deep semi-supervised learning algorithms. *Advances in neural information processing systems*, 31.

Pezeshki, M.; Fan, L.; Brakel, P.; Courville, A.; and Bengio, Y. 2016. Deconstructing the ladder network architecture. In *International conference on machine learning*, 2368–2376. PMLR.

Pham, H.; Dai, Z.; Xie, Q.; and Le, Q. V. 2021. Meta pseudo labels. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11557–11568.

Ren, Z.; Yeh, R.; and Schwing, A. 2020. Not all unlabeled data are equal: Learning to weight data in semi-supervised learning. *Advances in Neural Information Processing Systems*, 33: 21786–21797.

Sajjadi, M.; Javanmardi, M.; and Tasdizen, T. 2016. Regularization with stochastic transformations and perturbations for deep semi-supervised learning. *Advances in neural information processing systems*, 29.

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*, 33: 596–608.

Sun, Y.; Shi, Z.; and Li, Y. 2023. A Graph-Theoretic Framework for Understanding Open-World Semi-Supervised Learning. In Oh, A.; Naumann, T.; Globerson, A.; Saenko, K.; Hardt, M.; and Levine, S., eds., *Advances in Neural Information Processing Systems*, volume 36, 23934–23967. Curran Associates, Inc.

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*, 30.

Wang, H.; Zhang, Z.; Gao, J.; and Hu, W. 2024. A-Teacher: Asymmetric Network for 3D Semi-Supervised Object Detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14978–14987.

Xie, Q.; Dai, Z.; Hovy, E.; Luong, T.; and Le, Q. 2020. Unsupervised data augmentation for consistency training. *Advances in neural information processing systems*, 33: 6256–6268.

Yang, F.; Wu, K.; Zhang, S.; Jiang, G.; Liu, Y.; Zheng, F.; Zhang, W.; Wang, C.; and Zeng, L. 2022a. Class-aware contrastive semi-supervised learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14421–14430.

Yang, L.; Zhao, Z.; Qi, L.; Qiao, Y.; Shi, Y.; and Zhao, H. 2023. Shrinking Class Space for Enhanced Certainty in Semi-Supervised Learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 16187–16196.

Yang, X.; Song, Z.; King, I.; and Xu, Z. 2022b. A survey on deep semi-supervised learning. *IEEE Transactions on Knowledge and Data Engineering*, 35(9): 8934–8954.

You, Z.; Zhong, Y.; Bao, F.; Sun, J.; Li, C.; and Zhu, J. 2024. Diffusion models and semi-supervised learners benefit mutually with few labels. *Advances in Neural Information Processing Systems*, 36.

Zagoruyko, S.; and Komodakis, N. 2016. Wide residual networks. *arXiv preprint arXiv:1605.07146*.

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*, 34: 18408–18419.

Zhang, Q.; Qi, Y.; Tang, X.; Yuan, R.; Lin, X.; Zhang, K.; and Yuan, C. 2025. Rethinking pseudo-label guided learning for weakly supervised temporal action localization from the perspective of noise correction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 10085–10093.

Zheng, M.; You, S.; Huang, L.; Luo, C.; Wang, F.; Qian, C.; and Xu, C. 2023. SimMatchV2: Semi-Supervised Learning with Graph Consistency. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 16432–16442.

Zheng, M.; You, S.; Huang, L.; Wang, F.; Qian, C.; and Xu, C. 2022. Simmatch: Semi-supervised learning with similarity matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14471–14481.

Zhou, J.; Sun, J.; Li, C.; Jiang, Q.; Zhou, M.; Lam, K.-M.; Zhang, W.; and Fu, X. 2024. HCLR-net: Hybrid contrastive learning regularization with locally randomized perturbation for underwater image enhancement. *International Journal of Computer Vision*, 1–25.

Zhou, Z.-H.; and Li, M. 2010. Semi-supervised learning by disagreement. *Knowledge and Information Systems*, 24: 415–439.