Image-Feature Weak-to-Strong Consistency: An Enhanced Paradigm for Semi-Supervised Learning

Zhiyu Wu[®] and Jinshi Cui[®]

National Key Laboratory of General Artificial Intelligence, School of Intelligence Science and Technology, Peking University wuzhiyu@pku.edu.cn, cjs@cis.pku.edu.cn

Abstract. Image-level weak-to-strong consistency serves as the predominant paradigm in semi-supervised learning (SSL) due to its simplicity and impressive performance. Nonetheless, this approach confines all perturbations to the image level and suffers from the excessive presence of naive samples, thus necessitating further improvement. In this paper, we introduce feature-level perturbation with varying intensities and forms to expand the augmentation space, establishing the image-feature weak-to-strong consistency paradigm. Furthermore, our paradigm develops a triple-branch structure, which facilitates interactions between both types of perturbations within one branch to boost their synergy. Additionally, we present a confidence-based identification strategy to distinguish between naive and challenging samples, thus introducing additional challenges exclusively for naive samples. Notably, our paradigm can seamlessly integrate with existing SSL methods. We apply the proposed paradigm to several representative algorithms and conduct experiments on multiple benchmarks, including both balanced and imbalanced distributions for labeled samples. The results demonstrate a significant enhancement in the performance of existing SSL algorithms.

Keywords: Feature-Level Perturbation · Image-Feature Weak-to-Strong Consistency · Naive Sample Identification

1 Introduction

Semi-supervised learning (SSL) [4,30,32,46], a research topic that aims to reduce manual labeling costs by capitalizing on unlabeled data, has received considerable attention in recent years. Among the proposed methods, the image-level weak-to-strong consistency paradigm introduced by FixMatch [32] has proven highly effective in leveraging the potential of unlabeled data. Specifically, Fix-Match applies varying intensities of image-level perturbations to raw samples and encourages consistent predictions between corresponding weakly and strongly augmented views. Subsequent studies have refined this paradigm primarily in two aspects: (1) improving the utilization ratio of unlabeled data by replacing



Fig. 1: Overview of the old (left) and proposed (right) paradigms. Our paradigm introduces feature-level perturbation to expand the augmentation space and facilitates direct interactions between both types of perturbations to boost their synergy.

the constant threshold in FixMatch with a dynamic threshold [35,42] and (2) facilitating balanced predictions across different classes via refining pseudo-labels based on historical predictions [3,5]. Extensive experiments have demonstrated the effectiveness of these methods in enhancing the image-level weak-to-strong consistency paradigm.

Despite the success of the established paradigm, two challenges persist. Firstly, the existing paradigm confines all perturbations to the image level, thus impeding the exploration of a broader augmentation space and consistency at non-image levels. Secondly, as revealed in [12], a notable proportion of samples, even undergoing strong image-level perturbation, continues to be accurately classified with high confidence, resulting in a loss close to zero. These instances (some examples are shown in Fig. 5), referred to as *naive samples*, are well-learned and thus fail to boost the model's performance. Therefore, relying solely on image-level perturbation proves insufficient to fully exploit the potential of unlabeled data.

To address the aforementioned challenges, we introduce a novel SSL paradigm named Image-Feature Weak-to-Strong Consistency (IFMatch). An overview of the conventional and proposed approaches is shown in Fig. 1. The primary innovation of our method lies in feature-level perturbation, which is motivated by earlier findings that regularizers effective in the input space can be extended to hidden representations [14, 15, 34]. Specifically, feature-level perturbation randomly alters intermediate features in the backbone, establishing the basis for featurelevel consistency regulation. Furthermore, to achieve effective feature-level perturbation, we present refined designs focusing on position and strategy. For the perturbation position, we identify two positions to induce weak and strong feature-level perturbations, respectively. Regarding the perturbation strategy, we develop a series of approaches from three perspectives ('movement', 'dropout', and 'value'), thereby comprehensively exploring the feature augmentation space. Note that the proposed feature-level perturbation is sample-agnostic and thus differs from previous class-based feature augmentation techniques [20, 36].

Building upon well-crafted augmentation techniques, we introduce the integration of two types of perturbations as another innovation. Our preliminary investigation (Tab. 3) reveals that a direct fusion of strong image-level and featurelevel augmentation engenders destructive perturbations. To mitigate this issue, we propose a triple-branch structure to combine these two perturbation forms in a milder manner, as depicted in Fig. 1. Specifically, the teacher branch, mirroring its counterpart in the existing paradigm, employs predictions from samples subjected to weak image-level perturbation as pseudo-labels. Moreover, the first student branch integrates weak image-level perturbation and strong feature-level perturbation to thoroughly explore the feature augmentation space. Given limited prior research on feature-level perturbation, we delve deeper into the optimal threshold for this branch, revealing a preference for a high constant threshold over a dynamic one. In the second student branch, we combine weak feature-level perturbation with strong image-level perturbation, posing additional challenges for recognizing naive samples. However, the fused perturbation proves excessively arduous for classifying hard samples. Accordingly, we devise a confidence-based identification strategy to dynamically distinguish between naive and challenging samples, thereby introducing weak feature-level perturbation exclusively for naive samples. As the second branch predominantly relies on image-level perturbation, it can seamlessly integrate with existing threshold mechanisms. Prior to our study, UniMatch [40] uses feature-level perturbation and a multi-branch structure in semi-supervised semantic segmentation. However, their method is limited to the channel-wise dropout strategy and fails to achieve interactions between the two types of perturbations within a single branch, setting it apart from our paradigm. In summary, our contributions can be outlined as follows.

(1) We introduce the image-feature weak-to-strong consistency paradigm for semi-supervised learning, where the feature-level perturbation and the integration of two types of perturbation within each branch serve as the fundamental enhancement compared to previous methods.

(2) We meticulously design the position and strategy of feature-level perturbation to diversify its intensity and form.

(3) We devise a confidence-based identification strategy to discriminate between naive and challenging samples, thereby introducing extra challenges solely for naive samples in the second student branch.

(4) We conduct experiments on multiple benchmarks, including both balanced and imbalanced distributions for labeled samples. The results demonstrate significant promotion of existing algorithms when following our paradigm.

2 Related Work

Consistency regulation [2] represents a fundamental approach to leveraging the potential of unlabeled data. Its core objective is to ensure consistent predictions across diverse perturbed views of the same sample. Various perturbation methods have been explored, such as stochastic augmentation [22, 31], adversarial attack [25], and mixup [43]. Additionally, FixMatch [32] applies strong

data augmentation techniques [7] to raw images, establishing the image-level weak-to-strong consistency paradigm. This paradigm significantly streamlines the framework and marks a pivotal milestone in semi-supervised learning.

Follow-up studies refine FixMatch along two dimensions: the utilization ratio of unlabeled data and fairness considerations. Regarding the former, numerous dynamic threshold strategies have been proposed. For example, FlexMatch [42] maps the predefined threshold to class-specific thresholds based on class-wise learning status. SoftMatch [5] models sample weights by a dynamic Gaussian function, maintaining a soft margin between unlabeled samples of different confidence levels. FreeMatch [35] utilizes the confidence on unlabeled data as the adaptive threshold. Moreover, the concept of fairness stems from entropy-based regulation [1,11,16,18,44]. The most representative strategy in this context, Distribution Alignment [3], encourages equal-frequency predictions across classes by refining pseudo-labels based on overall predictions within the unlabeled set.

Concurrently, several studies express similar concerns to the old paradigm. FeatMatch [20] and ISDA [36] prevent the exclusive reliance on image-level perturbation by devising class-based feature augmentation approaches. In contrast, our proposed feature-level perturbation is sample-agnostic, circumventing harmful perturbations derived from erroneous pseudo-labels. Additionally, UniMatch [40] uses channel-wise dropout to establish an auxiliary branch for feature-level perturbation. In comparison, we present feature-level perturbations with varying intensities and forms, further combining them with image-level perturbations within each branch to boost their synergy. Besides, [12] observes excessive naive samples and applies more aggressive image-level augmentation to these samples. To address this issue, we propose a confidence-based identification strategy to effectively distinguish between naive and challenging samples and introduce weak feature-level perturbation exclusively for naive samples.

3 Methodology

3.1 Image-Level Weak-to-Strong Consistency

We begin by reviewing the image-level weak-to-strong consistency paradigm, as depicted in the left part of Fig. 1. Let $\mathcal{D}_L = \{(x_i, y_i)\}_{i=1}^{N_L}$ and $\mathcal{D}_U = \{u_i\}_{i=1}^{N_U}$ represent the labeled and unlabeled datasets, respectively. Here, x_i and u_i denote the labeled and unlabeled training samples, and y_i refers to the one-hot label for the labeled sample x_i . We denote the prediction of sample x as p(y|x). Given a batch of labeled and unlabeled data, the model is optimized with the objective $\mathcal{L} = \mathcal{L}_s + \lambda_u \mathcal{L}_u$. Specifically, \mathcal{L}_s indicates the cross-entropy loss (\mathcal{H}) for the labeled batch of size B_L .

$$\mathcal{L}_s = \frac{1}{B_L} \sum_{i=1}^{B_L} \mathcal{H}(y_i, p(y|x_i)) \tag{1}$$

On the other hand, \mathcal{L}_u signifies the image-level consistency regulation between the prediction of the strongly augmented view $\mathcal{A}^{\mathcal{I}_s}(u)$ and the pseudo-label



Fig. 2: Candidate positions for introducing feature-level perturbation. Position A perturbs the output of the residual block, representing strong feature-level perturbation $\mathcal{A}^{\mathcal{F}_s}$. Position B perturbs the output of a random convolution within the residual component, corresponding to weak feature-level perturbation $\mathcal{A}^{\mathcal{F}_w}$. x^{in}/x^{out} denotes input/output feature maps of the residual block.

sourced from the corresponding weakly augmented view $\mathcal{A}^{\mathcal{I}_w}(u)$. To mitigate incorrect pseudo-labels, FixMatch [32] introduces a constant threshold, and subsequent studies propose various dynamic threshold strategies. Without loss of generality, we denote the threshold as τ_t and define \mathcal{L}_u as follows.

$$\mathcal{L}_u = \frac{1}{B_U} \sum_{i=1}^{B_U} \mathbb{1}(\max(p_i^{\mathcal{I}_w}) \ge \tau_t) \mathcal{H}(\hat{p}_i^{\mathcal{I}_w}, p_i^{\mathcal{I}_s})$$
(2)

Here, $p_i^{\mathcal{I}_w}$ is the abbreviation of $DA(p(y|\mathcal{A}^{\mathcal{I}_w}(u_i)))$, where DA indicates the distribution alignment strategy [3]. $\hat{p}_i^{\mathcal{I}_w}$ denotes the one-hot pseudo-label obtained from $argmax(p_i^{\mathcal{I}_w})$. Moreover, $p_i^{\mathcal{I}_s}$ stands for $p(y|\mathcal{A}^{\mathcal{I}_s}(u_i))$. Lastly, B_U corresponds to the batch size of unlabeled data.

3.2 Feature-Level Perturbation

To avoid the exclusive reliance on image-level perturbation, this paper introduces feature-level perturbation varied in intensity and form, thereby expanding the augmentation space. In the following parts, we will elaborate on the implementation of feature-level perturbation regarding its position and strategy.

Position refers to where the feature-level perturbation is introduced. Given that WideResNet [41] is the prevailing backbone in SSL, we employ its basic component, the residual block [13] with pre-activation, as an illustrative example. Building upon previous studies that exploit feature-level regularizers directly after the convolution [14, 15, 34], we identify two positions for introducing feature-level perturbation, as depicted in Fig. 2. In position A, the method perturbs the summation of the residual and identity connections, *i.e.*, the output of the residual block. As position A acts as the bottleneck for feedforward, perturbation at position A exhibits high intensity and is therefore regarded as strong feature-level perturbation $\mathcal{A}^{\mathcal{F}_s}$. Additionally, position B perturbs the output of a random convolution within the residual component, which is relatively milder than position A. Therefore, perturbation at position B is considered as weak feature-level perturbation $\mathcal{A}^{\mathcal{F}_w}$. It is worth noting that both

position choices are applicable to all residual blocks, and the algorithm randomly selects one position at each iteration to insert feature-level perturbation.

Strategy denotes the operation of feature-level perturbation. We first review strong image-level perturbation strategies $\mathcal{A}^{\mathcal{I}_s}$ [7], which can be roughly divided into three groups: 'movement', 'dropout', and 'value'. Specifically, 'movement' comprises spatial perturbations like translation and shearing, 'dropout' corresponds to the Cutout technique [9], and 'value' represents operations that modify pixel values while preserving pixel relationships, such as adjusting contrast, brightness, and histogram. Considering the effectiveness of $\mathcal{A}^{\mathcal{I}_s}$, we develop feature-level perturbation strategies from the same perspectives. For 'movement', we implement feature map translation and shearing along the X-axis and Y-axis. Regarding 'dropout', we apply channel-wise and spatial-wise dropout on hidden representations. Concerning 'value', we employ random-sized convolutions to achieve local smoothing, subsequently using the weighted sum of smoothed and input feature maps as the output. During training, the model randomly selects one candidate strategy at each iteration to perturb hidden representations at the selected position. For the detailed implementation of featurelevel perturbation strategies, please refer to Appendix Sec.1.

3.3 Image-Feature Weak-to-Strong Consistency

Building upon well-designed image-level and feature-level perturbations, we can integrate them to achieve consistency regulation in both dimensions. Intuitively, a trivial approach is combining strong image-level and feature-level perturbations in one branch to construct a dual-branch structure similar to the old paradigm. However, such a toy paradigm engenders destructive perturbations according to our pilot study (Tab. 3). Accordingly, we integrate these two forms of perturbations in a more moderate manner, developing the image-feature weak-to-strong consistency paradigm (IFMatch) depicted in the right part of Fig. 1.

Our paradigm retains the traditional teacher branch, utilizing predictions $\hat{p}^{\mathcal{I}_w}$ of samples subjected to weak image-level perturbation as pseudo-labels. Furthermore, we introduce two student branches to integrate both types of perturbations. Mathematically, given a batch of labeled and unlabeled data, the model is optimized using the objective $\mathcal{L} = \mathcal{L}_s + \lambda_u(\mathcal{L}_{u_1} + \mathcal{L}_{u_2})$, where \mathcal{L}_s denotes the cross-entropy loss for labeled data, and \mathcal{L}_{u_1} and \mathcal{L}_{u_2} represent the consistency regulation for unlabeled data in the two student branches respectively.

In the first student branch, we integrate weak image-level perturbation $\mathcal{A}^{\mathcal{I}_w}$ and strong feature-level perturbation $\mathcal{A}^{\mathcal{F}_s}$ to comprehensively explore the feature perturbation space. Furthermore, we employ the constant threshold τ adopted in FixMatch to filter out incorrect pseudo-labels. This choice arises from the observation that existing dynamic threshold mechanisms are less suitable for $\mathcal{A}^{\mathcal{F}_s}$, especially in challenging tasks. Thus, the unsupervised loss \mathcal{L}_{u_1} in the first student branch can be expressed as follows.

$$\mathcal{L}_{u_1} = \frac{1}{B_U} \sum_{i=1}^{B_U} \mathbb{1}(\max(p_i^{\mathcal{I}_w}) \ge \tau) \mathcal{H}(\hat{p}_i^{\mathcal{I}_w}, p_i^{\mathcal{I}_w, \mathcal{F}_s})$$
(3)

Algorithm	1:	Image-Feature	Weak-to-Strong	Consistency	(IFMatch)
-----------	----	---------------	----------------	-------------	-----------

Ι	nput: labeled data $\mathcal{D}_L = \{(x_i, y_i)\}_{i=1}^{N_L}$, unlabeled data $\mathcal{D}_U = \{u_i\}_{i=1}^{N_U}$, batch
	sizes B_L and B_U , loss weight λ_u , total iterations T, perturbations
	$\{\mathcal{A}^{\mathcal{I}_w}, \mathcal{A}^{\mathcal{I}_s}, \mathcal{A}^{\mathcal{F}_w}, \mathcal{A}^{\mathcal{F}_s}\}, \text{ constant threshold } \tau, \text{ threshold mechanism}$
	τ_t , target confidence $\{h_1, h_2, \cdots, h_{N_U}\}$, mask $\{\mathcal{M}_1, \mathcal{M}_2, \cdots, \mathcal{M}_{N_U}\}$
1 f	or $t \leftarrow 1$ to T do
	<pre>// the same parts as the old paradigm</pre>
2	Compute supervised loss \mathcal{L}_s , pseudo-labels $p^{\mathcal{I}_w}$ and $\hat{p}^{\mathcal{I}_w}$, and threshold τ_t ;
	// the first student branch
3	Compute predictions $p^{\mathcal{I}_w, \mathcal{F}_s}$ and loss \mathcal{L}_{u_1} using Eq. (3);
	// the second student branch
4	Decide whether to introduce $\mathcal{A}^{\mathcal{F}_w}$ for u_i in the second student branch
	according to the mask \mathcal{M}_i using Eq. (6) and Eq. (7);
5	Compute predictions $p^{\mathcal{I}_s,\mathcal{F}_w}$ and loss \mathcal{L}_{u_2} using Eq. (4);
	// update model
6	Back-propagate loss $\mathcal{L} = \mathcal{L}_s + \lambda_u (\mathcal{L}_{u_1} + \mathcal{L}_{u_2})$ and update parameters ;
	<pre>// Confidence-based identification strategy</pre>
7	Record target confidence $h_i = p_{i,j}^{\mathcal{I}_s, \mathcal{F}_w}(j = \hat{p}_i^{\mathcal{I}_w})$ and mask $\mathcal{M}_i = \mathbb{1}(h_i \ge \tau_t);$
8 e	nd

Here, $p_i^{\mathcal{I}_w, \mathcal{F}_s}$ denotes the prediction of u_i in the first branch.

In the second student branch, we combine strong image-level perturbation $\mathcal{A}^{\mathcal{I}_s}$ and weak feature-level perturbation $\mathcal{A}^{\mathcal{F}_w}$ to provide additional challenges in recognizing naive samples. However, a direct fusion of $\mathcal{A}^{\mathcal{I}_s}$ and $\mathcal{A}^{\mathcal{F}_w}$ poses obstacles in classifying hard samples. Accordingly, we present a confidence-based identification strategy to distinguish between naive and challenging samples, thus introducing extra challenges exclusively for naive samples, which will be detailed in Sec. 3.4. Moreover, given the central role of $\mathcal{A}^{\mathcal{I}_s}$, we leverage existing threshold mechanisms τ_t to exclude misleading pseudo-labels, facilitating seamless integration with previous algorithms. Consequently, the unsupervised loss \mathcal{L}_{u_2} in the second branch can be defined as follows.

$$\mathcal{L}_{u_2} = \frac{1}{B_U} \sum_{i=1}^{B_U} \mathbb{1}(\max(p_i^{\mathcal{I}_w}) \ge \tau_t) \mathcal{H}(\hat{p}_i^{\mathcal{I}_w}, p_i^{\mathcal{I}_s, \mathcal{F}_w})$$
(4)

where $p_i^{\mathcal{I}_s,\mathcal{F}_w}$ represents the prediction of sample u_i in the second branch. Particularly, the use of distinct thresholds in the two student branches is motivated not only by experiments (Tab. 5) but also by an inherent rationale. Specifically, the threshold serves to balance the utilization ratio of unlabeled data and pseudo-label accuracy [5, 35]. Furthermore, $\mathcal{A}^{\mathcal{F}_s}$ and $\mathcal{A}^{\mathcal{I}_s}$ act as the central driving forces behind the respective student branch. Due to the unique characteristics of $\mathcal{A}^{\mathcal{F}_s}$ and $\mathcal{A}^{\mathcal{I}_s}$, they exhibit varying preferences in the trade-off between the utilization and pseudo-label accuracy, resulting in the different desired threshold mechanisms in the two student branches. For a further understanding of our paradigm, please refer to Algorithm 1.



Fig. 3: Overview of the confidence-based identification strategy. The approach records sample-wise target confidence in the second student branch and identifies naive samples by comparing the confidence with the threshold in the second student branch. w/ and w/o abbreviates with and without, respectively.

3.4 Confidence-Based Identification

As analyzed in Sec. 3.3, the combination of strong image-level perturbation $\mathcal{A}^{\mathcal{I}_s}$ and weak feature-level perturbation $\mathcal{A}^{\mathcal{F}_w}$ poses supplementary challenges for naive samples yet incurs exceeding difficulty for hard ones. Likewise, previous research has encountered similar challenges in expanding the augmentation space. For example, UniMatch [40] observes a performance degradation caused by the concatenation of channel-wise dropout and $\mathcal{A}^{\mathcal{I}_s}$ and finally denies the feasibility of combining $\mathcal{A}^{\mathcal{I}}$ and $\mathcal{A}^{\mathcal{F}}$ within one branch. Besides, [12] imposes more aggressive image-level perturbations exclusively on naive samples. To distinguish between naive and challenging samples, their approach records sample-wise loss and utilizes the OTSU [28] method for distinction. However, OTSU is conventionally applied to the segmentation of bimodal gray-scale images, whereas the histogram of loss values typically exhibits a descending trend. As shown in Fig. 4, their approach tends to produce identification biased toward naive samples.

To effectively identify the naive samples, we propose a confidence-based identification strategy, which is depicted in Fig. 3. Similar to the threshold-based rule for pseudo-label filtering that only predictions exceeding the confidence threshold qualify as pseudo-labels, the proposed strategy can be succinctly concluded as: 'only strongly augmented samples $\mathcal{A}^{\mathcal{I}_s}(u)$ with target confidence exceeding the threshold are considered naive and necessitate additional weak feature-level perturbation $\mathcal{A}^{\mathcal{F}_w}$ '. Specifically, the approach comprises two steps. Firstly, we record the confidence linked to the pseudo-label within the predicted distribution of the second branch. Mathematically, the target confidence h_i for sample u_i in the second student branch can be obtained as follows.

$$h_i = p_{i,j}^{\mathcal{I}_s, \mathcal{F}_w} \ (j = \hat{p}_i^{\mathcal{I}_w}) \tag{5}$$

It is important to clarify that $p_i^{\mathcal{I}_s, \mathcal{F}_w}$ solely represents the prediction of sample u_i in the second branch and has nothing to do with the introduction of $\mathcal{A}^{\mathcal{F}_w}$ for u_i . Secondly, the algorithm generates a binary mask for each sample to determine whether to introduce additional $\mathcal{A}^{\mathcal{F}_w}$ by comparing the target confidence with the threshold τ_t in the second student branch. Consequently, the mask for sample u_i , denoted as \mathcal{M}_i , and the perturbations applied to u_i in the second student

Table 1: Accuracy (%) on four benchmarks with varying numbers of labels. **Bold** indicates the best performance, and <u>underline</u> denotes the second best. We omit the standard deviation of performance, as the proposed paradigm performs stable across different random seeds. IFMatch (*) represents the proposed paradigm using the threshold mechanism of * in the second student branch. In Particular, given that SoftMatch proposes a sample weight strategy rather than a threshold scheme, IFMatch (Soft) inherits its weight scheme when computing \mathcal{L}_{u_2} and replaces the τ_t in Eq. (6) with the mean value of the dynamic Gaussian in SoftMatch when identifying naive samples.

Datasets	C	IFAR-1	10	CI	FAR-1	00	SV	HN	STI	2-10	Imag	eNet
# Label	40	250	4000	400	2500	10000	40	1000	40	1000	100k	400k
VAT [25]	25.34	58.97	89.49	14.80	53.16	67.86	25.25	95.89	25.26	62.05	-	-
MeanTeacher [33]	29.91	62.54	91.90	18.89	54.83	68.25	63.91	96.73	28.28	66.10	-	-
MixMatch [4]	63.81	86.37	93.34	32.41	60.24	72.22	69.40	96.31	45.07	78.30	-	-
ReMixMatch [3]	90.12	93.70	95.16	57.25	73.97	79.98	75.96	94.84	67.88	93.26	-	-
UDA [38]	89.38	94.84	95.71	53.61	72.27	77.51	94.88	98.11	62.58	93.36	-	-
Dash [39]	91.07	94.84	95.64	55.18	72.85	78.12	97.81	98.03	65.48	93.61	-	-
MPL [29]	93.38	94.24	95.45	53.74	72.29	78.26	90.67	97.72	64.24	93.34	-	-
FixMatch [32]	92.53	95.14	95.79	53.58	71.97	77.80	96.19	98.04	64.03	93.75	56.34	67.72
IFMatch (Fix)	95.82	95.97	96.49	66.26	74.65	$\underline{80.01}$	97.75	98.15	78.54	95.28	61.26	$\underline{71.48}$
FlexMatch [42]	95.03	95.02	95.81	60.06	73.51	78.10	91.81	93.28	70.85	94.23	58.15	68.69
IFMatch (Flex)	95.90	$\underline{95.99}$	96.42	66.28	$\underline{75.03}$	79.87	97.76	98.04	79.83	95.24	61.75	71.14
SoftMatch [5]	95.09	95.18	95.96	62.90	73.34	77.97	97.67	97.99	78.58	94.27	59.43	70.51
IFMatch (Soft)	95.84	95.93	96.36	67.02	74.99	79.96	98.04	98.22	82.46	95.22	62.44	71.66
FreeMatch [35]	95.10	95.12	95.90	62.02	73.53	78.32	98.03	98.04	84.44	94.37	58.48	69.48
IFMatch (Free)	95.83	96.02	$\underline{96.44}$	66.65	75.16	80.08	98.08	$\underline{98.18}$	84.75	95.29	61.69	71.30

branch can be expressed as follows.

$$\mathcal{M}_i = \mathbb{1}(h_i \ge \tau_t) \tag{6}$$

Perturbations for
$$u_i = \begin{cases} \{\mathcal{A}^{\mathcal{I}_s}, \mathcal{A}^{\mathcal{F}_w}\} & \mathcal{M}_i = 1\\ \{\mathcal{A}^{\mathcal{I}_s}\} & \mathcal{M}_i = 0 \end{cases}$$
 (7)

The approach dynamically updates \mathcal{M}_i throughout the training process, thereby considering both the learning status of individual samples and model performance when introducing $\mathcal{A}^{\mathcal{F}_w}$.

4 Experiments

4.1 Balanced Semi-Supervised Learning

Settings. We conduct experiments on multiple benchmarks, including CIFAR-10/100 [19], SVHN [26], STL-10 [6], and ImageNet [8], with various numbers of labeled samples. In the context of balanced SSL, we maintain a balanced distribution for labeled samples. The choice of backbone architecture adheres to established practices, employing specific models for different datasets: WRN-28-2 [41] for CIFAR-10 and SVHN, WRN-28-8 for CIFAR-100, WRN-37-2 [45] for STL-10, and ResNet-50 [13] for ImageNet. As for training parameters, we follow the unified codebase TorchSSL [42] to ensure fair comparisons. Specifically, batch

sizes B_L and B_U are set to 128 and 128 on ImageNet and 64 and 448 on the remaining datasets. Moreover, the model is trained using the SGD optimizer with an initial learning rate of 0.03 and a momentum decay of 0.9. The learning rate is adjusted by a cosine decay scheduler over a total of 2^{20} steps. Besides, we train the EMA model with a momentum decay of 0.999 for inference. The fixed threshold τ in the first branch is set to 0.95. For SVHN, we follow previous studies [35, 39] and constrain the dynamic threshold within the range of [0.9, 1.0]. To suppress randomness, we repeat each experiment three times and report the average accuracy in Tab. 1. We omit the standard deviation given the stable performance of the proposed paradigm across different random seeds.

Performance. We apply the proposed paradigm (IFMatch) to several algorithms and compare their performance with existing methods on multiple benchmarks. The results presented in Tab. 1 underscore the significant contribution of our paradigm to promoting the performance of existing SSL algorithms. Specifically, the four algorithms involved achieve the average improvement of 4.05%, 3.22%, 1.62%, and 1.39% across all datasets. Moreover, previous methods struggle to enhance performance on less challenging tasks, such as CIFAR-10. In contrast, when integrated with the proposed paradigm, IFMatch (Fix) outperforms FixMatch by 2.29%/0.83%/0.70% on CIFAR-10 with 40/250/4000labeled samples, respectively. Furthermore, our approach effectively addresses the challenges posed by extremely limited labeled samples, with IFMatch (Fix) achieving enhancements of 14.51% and 12.68% on STL-10 with 40 labels and CIFAR-100 with 400 labels, respectively, compared to FixMatch. Moreover, the four algorithms achieve an average promotion of 3.69%/2.30% on ImageNet with 100k/400k labels when following the proposed paradigm, indicating its effectiveness on large-scale datasets. Additionally, when employing our paradigm, the performance gap between dynamic and constant thresholds considerably narrows on challenging tasks, while IFMatch (Fix) even emerges as the preferred method on some less demanding tasks. The shifted performance gap implies a reduced reliance on complex threshold strategies of our paradigm, which primarily benefits from the expanded augmentation space. In summary, the impressive performance observed in Tab. 1 confirms the effectiveness of the proposed paradigm.

4.2 Imbalanced Semi-Supervised Learning

Settings. We evaluate the proposed paradigm in the context of imbalanced SSL, where both labeled and unlabeled data exhibit a long-tailed distribution. Consistent with prior studies [5,10,21,23,27,37], we construct labeled and unlabeled sets using the configurations of $N_c = N_1 \cdot \gamma^{-\frac{c-1}{C-1}}$ and $M_c = M_1 \cdot \gamma^{-\frac{c-1}{C-1}}$. Specifically, for CIFAR-10-LT, we set N_1 to 1500, M_1 to 3000, and γ to range from 50 to 150. Moreover, for CIFAR-100-LT, we set N_1 to 150, M_1 to 300, and γ to span from 20 to 100. In all experiments, we employ WRN-28-2 as the backbone and utilize the Adam optimizer [17] with a weight decay of 4e-5. Besides, batch sizes B_L and B_U are fixed as 64 and 128 respectively. Additionally, the learning rate is initially set to 2e-3 and adjusted by a cosine decay scheduler during

Dataset	CIFAR-10-LT			CIFAR-100-LT			
γ	50	100	150	20	50	100	
FixMatch [32] IFMatch (Fix)	$\frac{81.54{\pm}0.30}{85.32{\pm}0.33}$	$74.89{\scriptstyle\pm1.20}\atop79.46{\scriptstyle\pm0.35}$	$\frac{70.38 \pm 0.88}{75.59 \pm 0.48}$	$\frac{49.58 \pm 0.78}{53.61 \pm 0.25}$	$\frac{42.11{\pm}0.33}{44.58{\pm}0.38}$	37.60±0.48 39.92±0.37	
FlexMatch [42] IFMatch (Flex)	$\substack{81.87 \pm 0.19 \\ 84.36 \pm 0.12}$	$74.49{\scriptstyle\pm 0.92}\\79.04{\scriptstyle\pm 0.42}$	$70.20{\scriptstyle\pm 0.36}\atop{\scriptstyle73.02{\scriptstyle\pm 0.47}}$	50.89 ± 0.60 51.83 ± 0.33	$\substack{42.80 \pm 0.39 \\ 43.49 \pm 0.35}$	$\begin{array}{c} 37.30 {\pm} 0.47 \\ 37.98 {\pm} 0.36 \end{array}$	
SoftMatch [5] IFMatch (Soft)	83.45 ± 0.29 85.56 ± 0.15	$77.07{\scriptstyle\pm 0.37}\\\textbf{80.28}{\scriptstyle\pm 0.24}$	72.60 ± 0.46 76.18 \pm 0.34	51.91 ± 0.55 53.97 ± 0.32	$\substack{43.76 \pm 0.51 \\ \textbf{44.98} \pm \textbf{0.41}}$	$\frac{38.92 {\pm} 0.81}{39.68 {\pm} 0.41}$	
FreeMatch [35] IFMatch (Free)	$\substack{83.30 \pm 0.33 \\ 84.88 \pm 0.20}$	$\frac{75.96 \pm 0.49}{80.12 \pm 0.29}$	$71.20{\scriptstyle \pm 0.64} \\ 74.85{\scriptstyle \pm 0.47}$	$ \begin{vmatrix} 51.60 \pm 0.91 \\ 52.77 \pm 0.46 \end{vmatrix} $	$\substack{42.95 \pm 0.48 \\ 43.68 \pm 0.45}$	$\begin{array}{c} 37.50 {\pm} 0.23 \\ 38.00 {\pm} 0.55 \end{array}$	

Table 2: Performance (%) on CIFAR-10/100-LT with varying imbalance ratio.

training. We set the fixed threshold τ in the first branch to 0.95. To account for randomness, we report the mean and standard deviation of the performance across three different random seeds.

Performance. We apply the proposed paradigm to numerous algorithms and compare the performance with their counterparts following the old paradigm. The results in Tab. 2 demonstrate that our paradigm substantially improves existing models, enabling them to achieve state-of-the-art performance. Specifically, when integrated with the enhanced paradigm, the four algorithms achieve the average promotion of 3.72%, 2.03%, 2.16%, and 1.92% across all benchmarks. Notably, IFMatch (Fix) outperforms FixMatch by 5.21% and 2.32% on CIFAR-10-LT (γ =150) and CIFAR-100-LT (γ =100), highlighting the robustness of the proposed paradigm when facing severe imbalance. Moreover, the performance gap between dynamic-threshold-based methods and FixMatch substantially narrows within our paradigm. Accordingly, the proposed approach promotes the utilization of unlabeled data, reducing reliance on complex threshold mechanisms. Overall, the impressive performance observed in imbalanced SSL suggests the potential of our approach to be deployed in real-world applications.

4.3 Component Analysis

This section presents the contribution of each component within the proposed paradigm. Unless otherwise stated, we conduct experiments on IFMatch (Fix). Additionally, we provide running speed analysis, feature visualization results [24], and performance on semi-supervised semantic segmentation in the Appendix. The subsequent analysis presents the revealed findings.

The proposed approach to integrating image-level and feature-level perturbations yields optimal performance. As the simplest idea, we can integrate $\mathcal{A}^{\mathcal{F}_s}$ and $\mathcal{A}^{\mathcal{I}_s}$ within one branch to establish a dual-branch structure similar to the old paradigm. Nonetheless, such a toy design engenders disruptive perturbations that exceed acceptable augmentation levels for consistency regulation, resulting in the performance degradation presented in lines 1-2 of Tab. 3. An alternative approach, proposed by UniMatch [40], combines $\mathcal{A}^{\mathcal{F}}$ and $\mathcal{A}^{\mathcal{I}}$ in separate branches, each employing a distinct type of perturbation (line 3).

Table 3: Experiments (%) on the strategies for integrating image-level and featurelevel perturbations. CBI abbreviates the confidence-based identification strategy. CIFAR-m-n denotes CIFAR-m with n labeled samples. Line 3 combines the structure setting of UniMatch [40] and our proposed feature-level perturbation techniques.

Branch I	Branch II	$\big \operatorname{CIFAR-10-40}$	CIFAR-100-400
$egin{array}{c} ext{FixMatch} \ \mathcal{A}^{\mathcal{F}_S} + \mathcal{A} \end{array}$	(baseline, no $\mathcal{A}^{\mathcal{F}}$) \mathcal{I}_s in one branch	92.53 88.26	$53.58 \\ 51.47$
$egin{array}{c} \mathcal{A}^{\mathcal{F}_{S}} & \ \mathcal{A}^{\mathcal{F}_{S}} + \mathcal{A}^{\mathcal{I}_{W}} \ \mathcal{A}^{\mathcal{F}_{S}} + \mathcal{A}^{\mathcal{I}_{W}} \end{array}$	$\mathcal{A}^{\mathcal{I}_{S}} \ \mathcal{A}^{\mathcal{F}_{W}} + \mathcal{A}^{\mathcal{I}_{S}} \ \mathcal{A}^{\mathcal{F}_{W}} + \mathcal{A}^{\mathcal{I}_{S}} + ext{CBI}$	95.53 95.47 95.82	$ \begin{array}{r} 64.54 \\ 63.56 \\ 66.26 \end{array} $

Table 4: Ablation study (%) for the strategies of feature-level perturbation.

$\mathcal{A}^{\mathcal{F}}$ Strategy	CIFAR-10-40	CIFAR-100-400
All strategies w/o channel-wise dropout w/o spatial-wise dropout w/o translation w/o shearing	$\begin{array}{c} \textbf{95.82} \\ 95.66 (\downarrow 0.16) \\ 95.57 (\downarrow 0.25) \\ 95.59 (\downarrow 0.23) \\ 95.58 (\downarrow 0.24) \end{array}$	$\begin{array}{c} \textbf{66.26} \\ 65.68 \ (\downarrow 0.58) \\ 65.77 \ (\downarrow 0.49) \\ 65.70 \ (\downarrow 0.56) \\ 65.81 \ (\downarrow 0.47) \end{array}$
w/o value modification	95.62 (10.20)	$65.75 (\downarrow 0.51)$

Despite its effectiveness, this method fails to boost interactions between the two types of perturbations. Accordingly, we attempt to achieve the co-occurrence of $\mathcal{A}^{\mathcal{F}}$ and $\mathcal{A}^{\mathcal{I}}$ within one branch, thereby constructing a more comprehensive augmentation space. The proposed paradigm adopts a milder combination strategy, specifically using $\mathcal{A}^{\mathcal{F}_s} + \mathcal{A}^{\mathcal{I}_w}$ and $\mathcal{A}^{\mathcal{F}_w} + \mathcal{A}^{\mathcal{I}_s}$ in the two student branches respectively. However, as shown in line 4 of Tab. 3, the introduction of $\mathcal{A}^{\mathcal{F}_w}$ leads to a performance drop compared to UniMatch, probably due to the derived difficulties in recognizing hard samples. To address this issue, we devise a confidence-based identification strategy to distinguish between naive and challenging samples, posing additional challenges exclusively for naive samples. Finally, the proposed paradigm (IFMatch), as presented in line 5, achieves the optimal performance among all candidate designs for integrating $\mathcal{A}^{\mathcal{F}}$ and $\mathcal{A}^{\mathcal{I}}$, demonstrating its effectiveness and advantages.

All feature-level perturbation strategies contribute to the overall performance. As shown in Tab. 4, we conduct an ablation study on feature-level perturbation strategies. The baseline model in line 1 integrates all proposed strategies and achieves state-of-the-art performance. In subsequent configurations (line 2-6), we sequentially remove each strategy. Compared to the baseline, the absence of any proposed perturbation technique leads to a decline in performance on both datasets. Therefore, the proposed perturbation strategies exhibit a synergistic effect, collaboratively exploring the feature perturbation space.

Feature-level perturbation $\mathcal{A}^{\mathcal{F}}$ favors a high constant threshold over a dynamic one. As detailed in Sec. 3.3, $\mathcal{A}^{\mathcal{F}}$ and $\mathcal{A}^{\mathcal{I}}$ may require different threshold mechanisms due to their distinct characteristics. Considering the central role of $\mathcal{A}^{\mathcal{F}}$ ($\mathcal{A}^{\mathcal{F}_s}$) in the first student branch, we further delve into the

Branch I	Branch II	CIFAR-10-40	CIFAR-100-400
$ \begin{array}{l} {\rm FlexMatch} \\ \tau = 0.95 \end{array} $	FlexMatch	95.28	63.38
	FlexMatch	95.90	66.28
$\begin{array}{l} \text{SoftMatch} \\ \tau = 0.95 \end{array}$	SoftMatch	95.39	63.80
	SoftMatch	95.84	67.02
$ \begin{array}{c} \text{FreeMatch} \\ \tau = 0.95 \end{array} $	FreeMatch	95.56	64.45
	FreeMatch	95.83	66.65

Table 5: Exploration (%) for the threshold strategy in the first branch.

Table 6: Comparison (%) between SAA [12] and confidence-based identification.

Identification Strategy	CIFAR-10-40	CIFAR-100-400
SAA [12] Confidence-Based Identification	95.54 95.82	65.32 66.26

threshold strategy best suited for it. In general, existing threshold mechanisms can be categorized into constant thresholds and dynamic thresholds. The former prioritizes high pseudo-label accuracy at the expense of diminished utilization of unlabeled data, while the latter performs conversely. Accordingly, we determine the optimal threshold for the first branch from two candidates: a dynamic threshold mirroring the one used in the second branch and the constant threshold used in FixMatch. Combining Tab. 5 and Tab. 1, the proposed paradigm consistently outperforms the old paradigm regardless of the threshold adopted in the first branch. Additionally, the constant threshold beats all dynamic threshold mechanisms, suggesting that feature-level perturbation $\mathcal{A}^{\mathcal{F}}$ is better suited to a constant threshold of 0.95, rather than existing dynamic threshold mechanisms.

IFMatch outperforms previous methods from the perspective of the naive sample ratio. Firstly, as shown in the left part of Fig. 4, our approach (orange) achieves a lower naive sample ratio than UniMatch (green). This suggests more effective utilization of unlabeled data in our paradigm, primarily benefiting from combining both types of perturbation within each branch. Secondly, we compare the confidence-based identification strategy and SAA [12] within our paradigm to discern a better naive sample identification method. Specifically, SAA utilizes loss values as the criteria, recording historical loss for each sample and performing naive-challenging division by OTSU [28]. As shown in Tab. 6, the proposed module consistently outperforms SAA. Additionally, the naive sample ratio presented in the left part of Fig. 4 indicates that SAA (blue) tends to produce identification biased toward naive samples (e.q. starting with)a ratio of 0.5 in the first iteration). Furthermore, we analyze the identification process of both strategies in the later training stages, as shown in the right part of Fig. 4. Notably, when employing the cross-entropy loss, the loss threshold of 1.79 in SAA corresponds to a confidence threshold of 0.17 (computed as $e^{-1.79}$), which is a detrimental configuration. The limitations of SAA can be ascribed to the ill-suited nature of the OTSU technique for segmenting a histogram with a monotonically decreasing trend. In contrast, our strategy remains effective in

14 Wu and Cui



Fig. 4: Visualization for naive sample ratio and the identification process of SAA [12] and confidence-based identification on CIFAR-100-400.



Fig. 5: Visualization of perturbations. We provide the feature map for samples that undergo feature-level augmentation. A naive sample u exhibits the characteristic of $\forall \mathcal{A}^{\mathcal{I}_s}(u), \mathcal{H}(\hat{p}_i^{\mathcal{I}_w}, p_i^{\mathcal{I}_s}) \approx 0$, thus contributing trivially to the model's performance.

distinguishing between naive and challenging samples even if the unlabeled set is predominantly composed of naive samples.

Visualization analysis. We present the visualization for image-level and feature-level perturbations in Fig. 5, where the four parts (from left to right) correspond to the raw images, the images subjected to $\mathcal{A}^{\mathcal{I}_s}$, and the feature maps that undergo $\mathcal{A}^{\mathcal{F}_w} \circ \mathcal{A}^{\mathcal{I}_s}$ and $\mathcal{A}^{\mathcal{F}_s} \circ \mathcal{A}^{\mathcal{I}_w}$. We can observe that the exclusive reliance on $\mathcal{A}^{\mathcal{I}_s}$ engenders easy-to-recognize (naive) samples. In contrast, the proposed paradigm effectively combines $\mathcal{A}^{\mathcal{I}}$ and $\mathcal{A}^{\mathcal{F}}$ to expand the augmentation space. For additional visualization results, please refer to Appendix Sec.2.

5 Conclusion

This paper introduces feature-level perturbation to expand the augmentation space of the conventional semi-supervised learning framework, establishing the image-feature weak-to-strong consistency paradigm (IFMatch). To achieve effective feature-level perturbation, we propose refined designs that consider both perturbation position and strategy, facilitating diversity in terms of intensity and form. Furthermore, our paradigm develops two parallel student branches to seamlessly integrate the two types of perturbations and comprehensively explore the augmentation space. Additionally, we devise a confidence-based identification strategy to distinguish between naive and challenging samples, posing additional challenges exclusively for naive samples. Extensive experiments are conducted on multiple benchmarks, including both balanced and imbalanced label distributions. The experimental results demonstrate the effectiveness of our approach.

15

References

- Arazo, E., Ortego, D., Albert, P., O'Connor, N.E., McGuinness, K.: Pseudolabeling and confirmation bias in deep semi-supervised learning. In: 2020 International Joint Conference on Neural Networks (IJCNN). pp. 1–8. IEEE (2020)
- 2. Bachman, P., Alsharif, O., Precup, D.: Learning with pseudo-ensembles. Advances in neural information processing systems **27** (2014)
- Berthelot, D., Carlini, N., Cubuk, E.D., Kurakin, A., Sohn, K., Zhang, H., Raffel, C.: Remixmatch: Semi-supervised learning with distribution alignment and augmentation anchoring. arXiv preprint arXiv:1911.09785 (2019)
- Berthelot, D., Carlini, N., Goodfellow, I., Papernot, N., Oliver, A., Raffel, C.A.: Mixmatch: A holistic approach to semi-supervised learning. Advances in neural information processing systems **32** (2019)
- Chen, H., Tao, R., Fan, Y., Wang, Y., Wang, J., Schiele, B., Xie, X., Raj, B., Savvides, M.: Softmatch: Addressing the quantity-quality trade-off in semi-supervised learning. arXiv preprint arXiv:2301.10921 (2023)
- Coates, A., Ng, A., Lee, H.: An analysis of single-layer networks in unsupervised feature learning. In: Proceedings of the fourteenth international conference on artificial intelligence and statistics. pp. 215–223. JMLR Workshop and Conference Proceedings (2011)
- Cubuk, E.D., Zoph, B., Shlens, J., Le, Q.V.: Randaugment: Practical automated data augmentation with a reduced search space. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops. pp. 702–703 (2020)
- Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A largescale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009)
- DeVries, T., Taylor, G.W.: Improved regularization of convolutional neural networks with cutout. arXiv preprint arXiv:1708.04552 (2017)
- Fan, Y., Dai, D., Kukleva, A., Schiele, B.: Cossl: Co-learning of representation and classifier for imbalanced semi-supervised learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 14574– 14584 (2022)
- 11. Grandvalet, Y., Bengio, Y.: Semi-supervised learning by entropy minimization. Advances in neural information processing systems **17** (2004)
- Gui, G., Zhao, Z., Qi, L., Zhou, L., Wang, L., Shi, Y.: Enhancing sample utilization through sample adaptive augmentation in semi-supervised learning. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 15880–15889 (2023)
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
- Hinton, G.E., Srivastava, N., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.R.: Improving neural networks by preventing co-adaptation of feature detectors. arXiv preprint arXiv:1207.0580 (2012)
- Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: International conference on machine learning. pp. 448–456. pmlr (2015)
- Kim, J., Hur, Y., Park, S., Yang, E., Hwang, S.J., Shin, J.: Distribution aligning refinery of pseudo-label for imbalanced semi-supervised learning. Advances in neural information processing systems 33, 14567–14579 (2020)

- 16 Wu and Cui
- 17. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
- Krause, A., Perona, P., Gomes, R.: Discriminative clustering by regularized information maximization. Advances in neural information processing systems 23 (2010)
- 19. Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images (2009)
- Kuo, C.W., Ma, C.Y., Huang, J.B., Kira, Z.: Featmatch: Feature-based augmentation for semi-supervised learning. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVIII 16. pp. 479–495. Springer (2020)
- Lai, Z., Wang, C., Gunawan, H., Cheung, S.C.S., Chuah, C.N.: Smoothed adaptive weighting for imbalanced semi-supervised learning: Improve reliability against unknown distribution data. In: International Conference on Machine Learning. pp. 11828–11843. PMLR (2022)
- Laine, S., Aila, T.: Temporal ensembling for semi-supervised learning. arXiv preprint arXiv:1610.02242 (2016)
- Lee, H., Shin, S., Kim, H.: Abc: Auxiliary balanced classifier for class-imbalanced semi-supervised learning. Advances in Neural Information Processing Systems 34, 7082–7094 (2021)
- Van der Maaten, L., Hinton, G.: Visualizing data using t-sne. Journal of machine learning research 9(11) (2008)
- Miyato, T., Maeda, S.i., Koyama, M., Ishii, S.: Virtual adversarial training: a regularization method for supervised and semi-supervised learning. IEEE transactions on pattern analysis and machine intelligence 41(8), 1979–1993 (2018)
- 26. Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B., Ng, A.Y.: Reading digits in natural images with unsupervised feature learning (2011)
- Oh, Y., Kim, D.J., Kweon, I.S.: Daso: Distribution-aware semantics-oriented pseudo-label for imbalanced semi-supervised learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9786– 9796 (2022)
- 28. Otsu, N.: A threshold selection method from gray-level histograms. IEEE transactions on systems, man, and cybernetics 9(1), 62-66 (1979)
- Pham, H., Dai, Z., Xie, Q., Le, Q.V.: Meta pseudo labels. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 11557– 11568 (2021)
- Rosenberg, C., Hebert, M., Schneiderman, H.: Semi-supervised self-training of object detection models (2005)
- Sajjadi, M., Javanmardi, M., Tasdizen, T.: Regularization with stochastic transformations and perturbations for deep semi-supervised learning. Advances in neural information processing systems 29 (2016)
- Sohn, K., Berthelot, D., Carlini, N., Zhang, Z., Zhang, H., Raffel, C.A., Cubuk, E.D., Kurakin, A., Li, C.L.: Fixmatch: Simplifying semi-supervised learning with consistency and confidence. Advances in neural information processing systems 33, 596–608 (2020)
- 33. Tarvainen, A., Valpola, H.: Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. Advances in neural information processing systems **30** (2017)
- Verma, V., Lamb, A., Beckham, C., Najafi, A., Mitliagkas, I., Lopez-Paz, D., Bengio, Y.: Manifold mixup: Better representations by interpolating hidden states. In: International conference on machine learning. pp. 6438–6447. PMLR (2019)

- Wang, Y., Chen, H., Heng, Q., Hou, W., Savvides, M., Shinozaki, T., Raj, B., Wu, Z., Wang, J.: Freematch: Self-adaptive thresholding for semi-supervised learning. arXiv preprint arXiv:2205.07246 (2022)
- Wang, Y., Huang, G., Song, S., Pan, X., Xia, Y., Wu, C.: Regularizing deep networks with semantic data augmentation. IEEE Transactions on Pattern Analysis and Machine Intelligence 44(7), 3733–3748 (2021)
- Wei, C., Sohn, K., Mellina, C., Yuille, A., Yang, F.: Crest: A class-rebalancing self-training framework for imbalanced semi-supervised learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10857–10866 (2021)
- Xie, Q., Dai, Z., Hovy, E., Luong, T., Le, Q.: Unsupervised data augmentation for consistency training. Advances in neural information processing systems 33, 6256–6268 (2020)
- Xu, Y., Shang, L., Ye, J., Qian, Q., Li, Y.F., Sun, B., Li, H., Jin, R.: Dash: Semisupervised learning with dynamic thresholding. In: International Conference on Machine Learning. pp. 11525–11536. PMLR (2021)
- 40. Yang, L., Qi, L., Feng, L., Zhang, W., Shi, Y.: Revisiting weak-to-strong consistency in semi-supervised semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7236–7246 (2023)
- 41. Zagoruyko, S., Komodakis, N.: Wide residual networks. arXiv preprint arXiv:1605.07146 (2016)
- Zhang, B., Wang, Y., Hou, W., Wu, H., Wang, J., Okumura, M., Shinozaki, T.: Flexmatch: Boosting semi-supervised learning with curriculum pseudo labeling. Advances in Neural Information Processing Systems 34, 18408–18419 (2021)
- Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: Beyond empirical risk minimization. arXiv preprint arXiv:1710.09412 (2017)
- Zhao, Z., Zhou, L., Duan, Y., Wang, L., Qi, L., Shi, Y.: Dc-ssl: Addressing mismatched class distribution in semi-supervised learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9757– 9765 (2022)
- Zhou, T., Wang, S., Bilmes, J.: Time-consistent self-supervision for semi-supervised learning. In: International Conference on Machine Learning. pp. 11523–11533. PMLR (2020)
- 46. Zhu, X.J.: Semi-supervised learning literature survey (2005)