

RDA: Reciprocal Distribution Alignment for Robust Semi-supervised Learning

Yue Duan¹, Lei Qi², Lei Wang³, Luping Zhou⁴, and Yinghuan Shi^{1*}

¹ Nanjing University, China ² Southeast University, China

³ University of Wollongong, Australia ⁴ University of Sydney, Australia

Abstract. In this work, we propose Reciprocal Distribution Alignment (RDA) to address semi-supervised learning (SSL), which is a hyperparameter-free framework that is independent of confidence threshold and works with both the matched (conventionally) and the mismatched class distributions. Distribution mismatch is an often overlooked but more general SSL scenario where the labeled and the unlabeled data do not fall into the identical class distribution. This may lead to the model not exploiting the labeled data reliably and drastically degrade the performance of SSL methods, which could not be rescued by the traditional distribution alignment. In RDA, we enforce a reciprocal alignment on the distributions of the predictions from two classifiers predicting pseudo-labels and complementary labels on the unlabeled data. These two distributions, carrying complementary information, could be utilized to regularize each other without any prior of class distribution. Moreover, we theoretically show that RDA maximizes the input-output mutual information. Our approach achieves promising performance in SSL under a variety of scenarios of mismatched distributions, as well as the conventional matched SSL setting. Our code is available at: <https://github.com/NJUyued/RDA4RobustSSL>.

Keywords: distribution alignment, mismatched distributions

1 Introduction

Semi-supervised learning (SSL) leverages the abundant unlabeled data to alleviate the lack of labeled data for machine learning [5,32,24]. Lately, *confidence-based pseudo-labeling* [23,21] and *distribution alignment* [4,2,21,10] have been introduced to SSL, boosting the performance to a new height. These techniques improve the label imputation for unlabeled data, which alleviates the confirmation bias [1]. In brief, pseudo-labeling aims to achieve entropy minimization [11] by producing hard labels. Recently, FixMatch [23] utilizes the confidence-based threshold to select more accurate pseudo-labels and proves the superiority of this technique. Despite this threshold preventing the model from risk of noisy pseudo-labels, since

* Corresponding author: Y. Shi (e-mail: syh@nju.edu.cn).

Y. Duan, Y. Shi are with the National Key Laboratory for Novel Software Technology and the National Institute of Healthcare Data Science, Nanjing University.

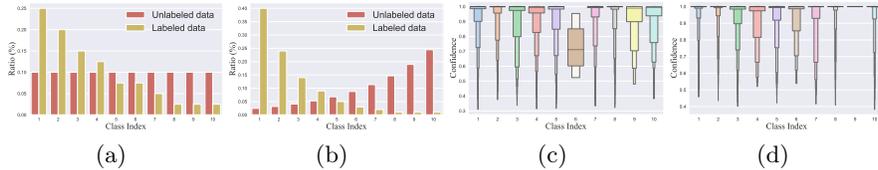


Fig. 1. Some examples of mismatched distributions in SSL. The x-axis represents the index of classes in CIFAR-10. In (a) and (b), the figures show the distributions of the labeled and unlabeled data. In (c) and (d), the figures show the confidences of FixMatch’s predictions on the unlabeled data. Letter-value plots [14] are displayed for multi-level quantile information. In (a) and (c), we show imbalanced labeled data and balanced unlabeled data with 40 labels $N_0 = 10$. In (b) and (d), the labeled and unlabeled data are mismatched and imbalanced with 100 labels, $N_0 = 40$ and $\gamma = 10$. More details about imbalance ratio N_0 and γ can be found in Sec. 4.2. In (c) and (d), we can see that the confidences of FixMatch’s predictions on the unlabeled data of different classes are totally irregular, which means it is difficult for us to adjust the confidence threshold to judge whether the prediction is correct. *i.e.*, confidence-based pseudo-labeling is also not suitable for the mismatched distributions.

the learning difficulties of different classes are different, a fixed threshold is not a “silver bullet” for all scenarios of SSL. Although [27,30] demonstrate the potential to dynamically adjust the threshold, the adjustment is complicated and the waste of unlabeled data with low confidence will become a latent limitation [9]. We try to ask — *is the confidence-based threshold really necessary for pseudo-labeling?*

Motivated by this, we rethink pseudo-labeling in a hyperparameter-free way while noticing that distribution alignment (DA) has been introduced to SSL [2,21,10]. DA scales the predictions on unlabeled data by prior information about labeled data distribution for strong regularization on the pseudo-labels, which can mitigate the confirmation bias. Inspired by this, we consider only using DA to improve the pseudo-labels without additional hyperparameters, *i.e.*, DA is enough for pseudo-labeling. Meanwhile, DA shows great potential in addressing the SSL under long-tailed distribution [26]. We expect that this technique can play a positive role in SSL in a more general scope. However, even though DA could help us improve pseudo-labeling by protecting SSL from noise, it is based on a strong assumption: “*labeled data and unlabeled data share the same distribution,*” *e.g.*, they are all balanced in CIFAR-10. The scenarios of *mismatched distributions* have not been widely discussed, *i.e.*, the distribution of labeled data doesn’t match that of unlabeled data, which is illustrated in Fig. 1. Some typical scenarios lead to mismatched distributions, such as biased sampling, label missing not at random [13] and so on. Mismatched distributions might cause biased pseudo-labels, significantly degrading the SSL model performance which is demonstrated by experimental results in Sec. 5.2. Under mismatched distributions, we cannot simply use the distribution of the labeled data to align predictions on unlabeled data with a very different distribution. This drives us to explore a more general distribution alignment to meet the above challenge of mismatched distributions.

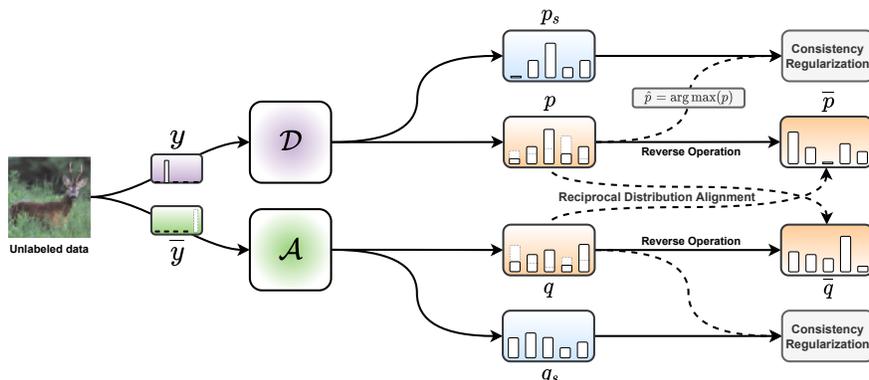


Fig. 2. Diagram of proposed Reciprocal Distribution Alignment (RDA). We use ground-truth label y and complementary label \bar{y} (dash line means \bar{y} is selected randomly from classes excluding ground-truth label) of labeled data to train Default Classifier \mathcal{D} and Auxiliary Classifier \mathcal{A} , respectively. Given an unlabeled sample u , \mathcal{D} predicts pseudo-label p and \mathcal{A} predicts complementary label q for its weakly-augmented version. RDA is applied on p and q by reciprocally scaling each other to the distributions of their reversed versions obtained by *Reverse Operation* (Proposition 1). We then enforce consistency regularization on the aligned pseudo-label and complementary label against corresponding predictions for strongly-augmented u , *i.e.*, p_s (from \mathcal{D}) and q_s (from \mathcal{A}).

Given motivations mentioned above, we propose Reciprocal Distribution Alignment (**RDA**) to establish a promising semi-supervised learning paradigm, which provides an integrated scheme to handle both the matched and mismatched scenarios in SSL. To relax the assumption about the class distribution of unlabeled data, we consider starting from the model itself to tap the potential guidance of distribution by regularizing the predictions from complementary perspectives. Inspired by [16,18,22], we consider simultaneously predict the class labels and their complementary labels (*i.e.*, indicating what class a sample is not), and utilize their distributions to regularize each other. Thus, we introduce two classifiers to RDA, one is Default Classifier (**DC**) and the other is Auxiliary Classifier (**AC**). Specifically, DC and AC are used to predict pseudo-labels and complementary labels for unlabeled data, respectively. The pseudo-labels and the complementary labels could be transformed into each other through their reversed version using the *Reverse Operation* defined in Proposition 1 in Sec. 3.3. Then a reciprocal alignment is employed to adjust the distributions of DC’s predictions and AC’s predictions by scaling them according to their corresponding reversed versions. We prove that RDA produces a “high-entropy” form of prediction distribution, which lead to maximizing the objective of input-output mutual information [4,2]. With the aligned pseudo-labels and complementary labels, the commonly used consistency regularization is further applied on DC and AC, respectively, which helps the model remain unchanged prediction on perturbed data. RDA could be applied to help the model improve pseudo-labels without suffering from

the threat of mismatched distributions since no prior information about class distribution of data is used. A diagram of RDA is shown in Fig. 2.

Despite its simplicity, our method shows superior performance in various settings, *e.g.*, on widely-used SSL benchmark CIFAR-10, RDA achieves an accuracy of $92.03 \pm 2.01\%$ with only 20 labels in the conventional setting, and in mismatched distributions, outperforms CoMatch [21], a recently-proposed algorithm for SSL, by up to a 52.09% gain on accuracy. Besides the significant performance improvement, our contributions can be presented as follows:

- We propose Reciprocal Distribution Alignment (RDA), a novel SSL algorithm, which can improve pseudo-labels in a hyperparameter-free way.
- RDA can be safely applied to SSL in both the conventional setting and the scenarios of mismatched distributions.
- We theoretical show that RDA could optimize the objective of mutual information between input data and predictions [4,2] under the premise of rational use of class distribution guidance information.

2 Related Work

Pseudo-labeling Based Entropy Minimization. Entropy minimization is a significant idea in recent SSL methods, which is closely related to pseudo-labeling (*i.e.*, convert model’s predictions to hard labels to reduce entropy) [20,23,21,28]. In another word, pseudo-labeling results in a form of entropy minimization [11]. This idea argues that model should ensure classes are well-separated while utilizing unlabeled data, which can be achieved by encouraging the model output prediction with low entropy [11]. Recent SSL algorithms like [23,21,27,31] set a confidence-based threshold to refine the pseudo-labels and obtain outstanding performance. However, the existence of confidence threshold leads to a waste of unlabeled samples with low confidence because they were filtered out. Moreover, it will lead to a significant increase in the cost of dynamic adjustment on confidence threshold like [27,30]. Meanwhile, under mismatched distributions, it is not reasonable to use a fixed threshold for all classes to filter pseudo-labels, because the model will also be affected by the unlabeled data with a potential risk of unknown distribution. In this work, we use distribution alignment to improve pseudo-labeling in a hyperparameter-free way which can achieve a better performance than algorithms introducing confidence threshold.

Distribution Alignment in SSL. Distribution alignment is proposed in [4] and originally applied to SSL in [2]. Briefly, [2] integrates it into pseudo-label inference step without additional loss terms or hyper-parameters. The main idea is the marginal distribution of predictions on unlabeled data and the marginal distribution of ground-truth labels should be consistent. This alleviates the confirmation bias [1] by improving pseudo-labels with the help of distributional guidance information. For class-imbalanced semi-supervised learning, [26] improves this technique by replacing the distribution of ground-truth labels with a smoothed form, resulting in superior performance in this setting. This improved distribution alignment in [26] helps the model benefit from rebalancing distribution.

In short, the objective of distribution alignment is to maximize the mutual information between the predictions and input data, *i.e.*, input-output mutual information [4,2]. Denoting the input data as x , the class prediction for x as y , and the predicted class distribution as $P(y|x)$, we can formalize this objective as:

$$\mathcal{I}(y; x) = \mathcal{H}(\mathbb{E}_x[P(y|x)]) - \mathbb{E}_x[\mathcal{H}(P(y|x))], \quad (1)$$

where $\mathcal{H}(\cdot)$ refers to entropy. For specific, distribution alignment aims to maximize the term $\mathcal{H}(\mathbb{E}_x[P(y|x)])$. However, the implementation of this technique in both [2] and [26] is based on an idealized assumption: “*labeled and unlabeled data fall in the same distribution.*” More realistically, we cannot guarantee that the distribution of labeled data matches that of unlabeled data. Such mismatched distributions can cause the distribution alignment in [2,26] to fail and is even detrimental to the model’s predictions on unlabeled data. In this work, we propose Reciprocal Distribution Alignment without the assumption of matched distributions and any prior information about the labeled data distribution.

3 Method

In this section, we discuss the setting of mismatched distributions in SSL and propose a novel SSL algorithm called Reciprocal Distribution Alignment (**RDA**) without additional hyper-parameters to improve pseudo-labeling in various scenarios of SSL. Moreover, we theoretically analyze the effectiveness of our method.

3.1 Matched and Mismatched Distributions in SSL

In semi-supervised learning, we have a training set divided into labeled portion \mathcal{X} and unlabeled portion \mathcal{U} . We denote class distribution of \mathcal{X} as \mathcal{C}_x and class distribution of \mathcal{U} as \mathcal{C}_u . Note that \mathcal{C}_u is inaccessible in training. Given $x \in \mathcal{X}$ with corresponding label y and unlabeled data $u \in \mathcal{U}$, we can review the SSL algorithms as the following optimization task:

$$\min \mathcal{L} = \mathcal{L}_{sup}(x, y; \theta) + \mathcal{L}_{unsup}(u; \theta), \quad (2)$$

where θ is the parameters of the model, \mathcal{L}_{sup} is supervised loss for the labeled data and \mathcal{L}_{unsup} is unsupervised loss for the unlabeled data. Recent pseudo-labeling based SSL methods try to impute the unknown label of u for \mathcal{L}_{unsup} . Therefore, the accuracy of pseudo-labels has become the top priority. In the traditional SSL setting, we assume $\mathcal{C}_x \approx \mathcal{C}_u$. Under this assumption, we can use \mathcal{C}_x to guide the prediction for u by distribution alignment [2,21], which can improve the performance of consistency-based or pseudo-labeling based methods [2,23,21,10]. Unfortunately, this assumption is too impractical and idealistic. More in line with the actual situation is $\mathcal{C}_x \not\approx \mathcal{C}_u$, which is called *mismatched distributions* in SSL. Unlike the conventional SSL, in mismatched distributions, the model learns a distribution from \mathcal{C}_x that differs from \mathcal{C}_u , so it cannot correctly predict the pseudo-labels. In other words, the distribution gap

caused by mismatch leads to strong confirmation bias [1], which could affect the performance of the model. It is worth noting that the distribution alignment used in [26] to solve the SSL under long-tail distribution also cannot be applied to the mismatched scenarios because [26] still depends on the assumption of matched distributions. To design a method that can tackle mismatched scenarios in SSL, we must face to $\mathcal{C}_x \not\approx \mathcal{C}_u$, and abandon prior of \mathcal{C}_x used in previous method [2,26].

3.2 Overview

We introduce two classifiers for our method. One is called Default Classifier (DC) \mathcal{D} and the other is called Auxiliary Classifier (AC) \mathcal{A} . In a nut shell, for an unlabeled image, \mathcal{D} is used to predict pseudo-label and \mathcal{A} is used to predict complementary label. We obtain labeled data $\mathcal{X} = \{(x_b, y_b)\}_{b=1}^B$ consisting of B images and unlabeled data $\mathcal{U} = \{(u_b)\}_{b=1}^{\mu B}$ consisting of μB images in a batch of data. At first, we construct complementary label \bar{y} for every labeled data by their ground-truth. Complementary label [15,16] represents which class the sample does not belong to. Denoting $y \in \mathcal{Y} = \{1, \dots, n\}$ as the ground-truth label of x where n is the number of classes, following [18], the complementary label of x is randomly selected from $\mathcal{Y} \setminus \{y\}$, which is denoted as \bar{y} .

Following [23], we integrate *consistency regularization* into RDA. Weak and strong augmentations are performed on images then we enforce consistency regularization on \mathcal{D} and \mathcal{A} . Denoting u_w as the weakly-augmented image and u_s as the strongly-augmented image for the same unlabeled data u , let y_c be the class prediction for input image. $P_G(y_c|\cdot)$ refers to the predicted class distribution outputted by classifier G for input. We can obtain pseud-labels $p = P_{\mathcal{D}}(y_c|u_w)$, $p_s = P_{\mathcal{D}}(y_c|u_s)$, and complementary labels $q = P_{\mathcal{A}}(y_c|u_w)$, $q_s = P_{\mathcal{A}}(y_c|u_s)$ respectively. Note that p, q are n -dimensional vectors of class probability where n is the number of classes. p_i, q_i represent the probability of belonging to the i -th class in the predictions. Then, dual consistency regularization can be achieved by minimizing the default consistency loss \mathcal{L}_{cd} and auxiliary consistency loss \mathcal{L}_{ca} :

$$\mathcal{L}_{cd} = \frac{1}{\mu B} \sum_{n=1}^{\mu B} H(\hat{p}_n, p_{s,n}), \quad (3)$$

$$\mathcal{L}_{ca} = \frac{1}{\mu B} \sum_{n=1}^{\mu B} H(q_n, q_{s,n}), \quad (4)$$

where $H(\cdot, \cdot)$ refers to the cross-entropy loss and $\hat{p} = \arg \max(p)$, which means we use hard labels for consistency regularization on \mathcal{D} . Differently, soft labels are used for \mathcal{A} instead. RDA exploits all unlabeled data for training, whereas previous consistency-based methods waste low-confidence data [23,21,27].

In addition, we enforce cross-entropy loss on \mathcal{D} between weakly-augmented version of x (denoted as x_w) and y , and on \mathcal{A} between x_w and \bar{y} respectively:

$$\mathcal{L}_{sd} = \frac{1}{B} \sum_{n=1}^B H(y_n, P_{\mathcal{D}}(y_c|x_{w,n})), \quad (5)$$

$$\mathcal{L}_{sa} = \frac{1}{B} \sum_{n=1}^B H(\bar{y}_n, P_{\mathcal{A}}(y_c|x_{w,n})), \quad (6)$$

where \mathcal{L}_{sd} is default supervised loss for \mathcal{D} and \mathcal{L}_{sa} is auxiliary supervised loss for \mathcal{A} . To sum up, RDA jointly optimizes four losses mentioned above:

$$\mathcal{L} = \mathcal{L}_{sd} + \lambda_a \mathcal{L}_{sa} + \lambda_{cd} \mathcal{L}_{cd} + \lambda_{ca} \mathcal{L}_{ca}, \quad (7)$$

where λ_a , λ_{cd} and λ_{ca} are trade-off coefficients and are all set to 1 for simplicity.

Previous entropy minimization based methods like [23,21,27] achieve superior performance in SSL by pseudo-labeling. Their key to success is the confidence threshold set to control the selection of pseudo-labels. To eliminate this hyper-parameter that becomes cumbersome in mismatched distributions, we consider a way to improve pseudo-labels using only distribution alignment. According to Eq. (1), we can formalize the objective of distribution alignment for \mathcal{D} as:

$$\max_{\mathcal{D}} \mathcal{H}[\mathbb{E}_u(P_{\mathcal{D}}(y_c|u_w))], \quad (8)$$

where $\mathcal{H}(\cdot)$ refers to the entropy. Likewise, we formalize the objective of distribution alignment for \mathcal{A} as:

$$\max_{\mathcal{A}} \mathcal{H}[\mathbb{E}_u(P_{\mathcal{A}}(y_c|u_w))]. \quad (9)$$

This two objectives encourage model to make predictions with equal frequency but these are not necessarily useful when dataset’s class distribution of ground-truth is not uniform. We use Reciprocal Distribution Alignment described in next paragraph to incorporate these two objectives.

3.3 Reciprocal Distribution Alignment

Following [2], we notice that making one distribution approach to another (distribution of labeled data is used in [2]) can achieve the purpose of maximizing Eq. (1). In this way, a form of “high entropy” could be achieved for the objective described by Eqs. (8) and (9). In brief, we define the objective over \mathcal{D} and \mathcal{A} as:

$$\max_{\mathcal{D}, \mathcal{A}} h(\mathcal{D}, \mathcal{A}) = \mathcal{H}[\mathbb{E}_u(p)] + \mathcal{H}[\mathbb{E}_u(q)]. \quad (10)$$

However, due to the existence of mismatched scenarios, the class distribution of labeled data cannot be directly used for alignment like [2]. So, next we will use the distribution of class predictions (*i.e.*, $\mathbb{E}_u(p)$) and the distribution of complementary class predictions (*i.e.*, $\mathbb{E}_u(q)$) to build a reciprocal alignment. Considering there is no strong correlation between the distribution of class predictions and that of complementary class predictions, we assume that \mathcal{A} is used to predict pseudo-label \bar{q} (a “reversed” version of q), so that the “reversed” version of $\mathbb{E}_u(q)$ (*i.e.*, $\mathbb{E}_u(\bar{q})$) can be used to align $\mathbb{E}_u(p)$.

Proposition 1 (Reverse Operation). *In the case of using \mathcal{A} to predict pseudo-labels, we have $\bar{q} = \text{Norm}(\mathbb{1} - q)$, where $\mathbb{1}$ is all-one vector and $\text{Norm}(x)$ is the normalized operation defined as $x'_i = x_i / \sum_{j=1}^n x_j$, $i \in (1, \dots, n)$.*

Proof. Assuming we use \mathcal{A} to predict pseudo-label \bar{q} , ideally, the probability of one class (*i.e.*, q_i) should randomly fall on a class which is different from the class predicted currently (*i.e.*, \bar{q}_j where $j \neq i$). Thus, for any $\bar{q}_j \in \bar{q}$, its value is the sum of the values randomly assigned to it by all q_i :

$$\bar{q}_j = \sum_{i=1, i \neq j}^n \frac{q_i}{n-1} = \frac{1-q_j}{n-1}. \quad (11)$$

Rewriting it we obtain:

$$\begin{aligned} \bar{q}_j &= \frac{1-q_j}{n - \sum_{k=1}^n q_k} = \frac{1-q_j}{(1-q_1) + \dots + (1-q_n)} \\ &= \frac{1-q_j}{\sum_{k=1}^n (1-q_k)} = \text{Norm}(1-q_j). \end{aligned} \quad (12)$$

Now, $\bar{q} = \text{Norm}(1-q)$ follows by combining the similar proof for any $q_i \in q$. \square

Likewise, if we use \mathcal{D} to predict complementary label \bar{p} , it can be calculated as $\bar{p} = \text{Norm}(\mathbb{1} - p)$. By Eq. (11), we notice that Reverse Operation does not change the relative relationship between classes in the class distribution, but just reverses the order, which allows us to still obtain helpful guidance information from the pseudo-label and complementary label perspectives.

Then, distribution alignment is conducted on $\mathbb{E}_u(p)$ by scaling it to $\mathbb{E}_u(\bar{q})$. Reciprocally, we align $\mathbb{E}_u(q)$ by scaling it to $\mathbb{E}_u(\bar{p})$. Following [2], we also integrate distribution alignment into RDA without hyper-parameters. We compute the moving average $\Psi(\cdot)$ of p , q , and their reversed version \bar{p} , \bar{q} over last 128 batches, which can respectively serve as the estimation of $\mathbb{E}_u(p)$, $\mathbb{E}_u(q)$, $\mathbb{E}_u(\bar{p})$ and $\mathbb{E}_u(\bar{q})$. Given an unlabeled image u , we scale the prediction of \mathcal{D} , *i.e.*, pseudo-label p by:

$$\tilde{p} = \text{Norm}\left(p \times \frac{\Psi(\bar{q})}{\Psi(p)}\right), \quad (13)$$

where \tilde{p} is an aligned probability distribution. Then, $\hat{p} = \arg \max \tilde{p}$ is used as hard pseudo-label for default consistency loss \mathcal{L}_{cd} . Meanwhile, we scale the prediction of \mathcal{A} , *i.e.*, complementary label q by:

$$\tilde{q} = \text{Norm}\left(q \times \frac{\Psi(\bar{p})}{\Psi(q)}\right), \quad (14)$$

where \tilde{q} is an aligned probability distribution. Then \tilde{q} is used as soft complementary label for auxiliary consistency loss \mathcal{L}_{ca} . The following theorem shows why RDA results in maximizing the objective Eq. (10). In this way, the input-output mutual information could be maximized, boosting the model's performance [4,2].

Theorem 1. For pseudo-label p and the reversed pseudo-label \bar{p} obtained by *Reverse Operation*, we show that the entropy of \bar{p} is larger than that of p :

$$\mathcal{H}(\bar{p}) \geq \mathcal{H}(p). \quad (15)$$

Proof. We sort the sequence p_1, \dots, p_n in descending order and denote the sorted sequence as $p_1 \geq \dots \geq p_n$ for simplicity. Considering the case where $p_1 < \frac{1}{2}$ firstly, we prove a equivalent form of Theorem 1:

$$\sum_{i=1}^n [p_i \log p_i - \frac{(1-p_i)}{n-1} \log \frac{(1-p_i)}{n-1}] \geq 0. \quad (16)$$

We define the function as

$$f(x) = x \log x - \frac{1-x}{n-1} \log \frac{1-x}{n-1}, \quad (17)$$

where $x \in [0, \frac{1}{2}]$ by $\frac{1}{2} \geq p_1 \geq \dots \geq p_n$. The second derivative of this function is

$$f''(x) = \frac{1}{x} - \frac{1}{(n-1)(1-x)} = \frac{(n-1) - nx}{x(n-1)(1-x)} \quad (18)$$

Let $f''(x) \geq 0$, we obtain $x \leq \frac{n-1}{n}$. Considering $n \geq 2$, the minimum of the term $\frac{n-1}{n}$ is $\frac{1}{2}$. By $x \leq \frac{1}{2}$, $f''(x) \geq 0$ holds, which means the $f(x)$ is a convex function. Thus, by Jensen's Inequality, we have

$$\frac{1}{n} \sum_{i=1}^n f(x_i) \geq f\left(\frac{1}{n} \sum_{i=1}^n x_i\right) \quad (19)$$

Substituting in $x_i = p_i$, by Eq. (19), we obtain

$$\frac{1}{n} \sum_{i=1}^n (p_i \log p_i - \frac{1-p_i}{n-1} \log \frac{1-p_i}{n-1}) \geq \frac{1}{n} \log \frac{1}{n} - \frac{1-\frac{1}{n}}{n-1} \log \frac{1-\frac{1}{n}}{n-1} = 0 \quad (20)$$

□

The proofs of this theorem when $p_1 \geq \frac{1}{2}$ and the version of complementary label q are provided in Sec. A of Supplementary Material. Given the above proof, \mathcal{D} and \mathcal{A} are optimized to output predictions \bar{p} and \bar{q} with larger entropy, *i.e.*,

$$\mathcal{H}[\mathbb{E}_u(p)] + \mathcal{H}[\mathbb{E}_u(q)] \leq \mathcal{H}[\mathbb{E}_u(\bar{p})] + \mathcal{H}[\mathbb{E}_u(\bar{q})]. \quad (21)$$

RDA maximizes the objective Eq. (10) by aligning $\mathbb{E}_u(p)$ to $\mathbb{E}_u(\bar{q})$ and aligning $\mathbb{E}_u(q)$ to $\mathbb{E}_u(\bar{p})$ reciprocally, so as the input-output mutual information objective Eq. (1) could be maximized.

With *Reverse Operation*, we can apply distribution alignment while ensuring that the relative relationship between classes in the class distribution can be utilized, so as RDA could achieve a more reasonable form of “high entropy” for the objective of distribution alignment without using prior about C_x . So far, we construct hyperparameter-free Reciprocal Distribution Alignment (**RDA**), which is robust to SSL under both mismatched distributions and the conventional setting. The whole algorithm is presented in Sec. B of Supplementary Material.

4 Experimental Setup

We evaluate RDA on various standard benchmarks of SSL image classification task under diverse settings, including mismatched distributions (*i.e.*, $C_x \not\approx C_u$) and the conventional SSL setting (*i.e.*, $C_x \approx C_u$ and they are all balanced). Experiments show that RDA outperforms significantly over current state-of-the-art (SOTA) SSL methods under most settings. We also conduct further ablation studies on the effectiveness of each components in our method.

4.1 Datasets

RDA is evaluated on four datasets used in SSL widely: CIFAR-10/100 [19], STL-10 [6] and mini-ImageNet [25]. CIFAR-10/100, are composed of 60,000 images from 10/100 classes. Both of them are divided into training set with 50,000 images and test set with 10,000 images. STL-10 is composed of 5,000 labeled images and 100,000 unlabeled images which extracted from a broader distribution. mini-ImageNet is a subset of ImageNet [8] consisting of 100 classes, and each class has 600 images.

4.2 Settings of C_x and C_u

In addition to the conventional matched setting (*i.e.*, both C_x and C_u are balanced), we verify the efficacy of our method in more realistic mismatched scenarios, as discussed in Sec. 3.1. In view of the complexity of this problem, we mainly use the following three scenarios to summary our experimental protocol:

- Training with imbalanced C_x and balanced C_u . We are interested in the impact of mismatched distributions resulting from this simple setting. A graphical explanation of this setting is shown in Fig. 1(a).
- Training with mismatched and imbalanced C_x, C_u , which is shown in Fig. 1(b). This challenging setting can fully test the robustness of RDA.
- Training with balanced C_x and imbalanced C_u .

For experiments in above scenarios, we randomly select samples from dataset to construct imbalanced C_x and C_u . For C_x , the number of labeled data N_i in each class is fixed by N_0 . N_i is calculated as $N_i = N_0 \times \gamma_x^{-\frac{i-1}{n-1}}$, where n is the number of classes and $i \in (1, \dots, n)$. For fairness, we hold N_0 and search a proper γ_x for each N_i to keep the total number of labeled data consistent with we set. Details on searching for γ_x are shown in Sec. C.2 of Supplementary Material.

Specially, C_u is constructed in a form similar to reversely ordered C_x for more challenging setting. After a random selection of unlabeled data from dataset, the remaining data is seen as unlabeled data. The number of unlabeled data M_i of each class is fixed by: $M_i = M_0 \times \gamma^{-\frac{n-i}{n-1}}$, where $M_0 = 5000$ in CIFAR-10, $M_0 = 500$ in mini-ImageNet. In this way, we construct C_u as a “reversed” version of C_x as shown in Fig. 1(b). Likewise, DARP’s protocol [17] also produces datasets with mismatched distributions from CIFAR-10 and STL-10. So we also make a fair comparison with DARP under this protocol. More details about DARP’s protocol can be found in Sec. C.1 of Supplementary Material.

4.3 Baselines

We compare RDA mainly with three recent state-of-the-art SSL methods: (1) FixMatch [23], combining consistency regularization and entropy minimization; (2) FixMatch with distribution alignment [2]; (3) CoMatch [21], combining graph-based contrastive learning and consistency regularization. Moreover, we provide more comparisons with MixMatch [3], AlphaMatch [10], and DARP [17].

4.4 Implementation Details

Unless noted otherwise, we adopt Wide ResNet [29] and Resnet-18 [12] as the backbone for experiments. For specific, WRN-28-2 is used for CIFAR-10, WRN-28-8 is used for CIFAR-100 and Resnet-18 is used for STL-10/mini-ImageNet. Following [23], RandAugment [7] is used for strong augmentation. For simplicity, we train models using SGD with a momentum of 0.9 and a weight decay of 0.0005 in all experiments. In addition, we use a learning rate of 0.03 with cosine decay schedule to train the models for 1024 epochs. For hyper-parameters, we set $\mu = 7$, $B = 64$, $\lambda_a = \lambda_{cd} = \lambda_{ca} = 1$ for all experiments. Particularly, we report the results averaged on five folds and the standard deviation is calculated.

5 Results and Analysis

5.1 Conventional Setting (Matched Distributions)

For a fair comparison with baseline SSL methods, we conduct experiments in the conventional setting, *i.e.*, both C_x and C_w are balanced. We test the accuracy of RDA on CIFAR-10, mini-ImageNet, and STL-10 by varying the number of labeled data. Tab. 1 shows that the performance of RDA is compatible to (if not better than) that of the conventional SSL methods under matched class distributions. This results also confirm our view that with our design, the distribution alignment alone is enough for pseudo-labeling. RDA outperforms CoMatch by 3.60% when labels are scarce (with 20 labels). Moreover, on datasets with more classes, our method consistently achieves improvement on accuracy than the best baseline, *e.g.*, 46.91% (ours) vs 43.72% (CoMatch) on mini-ImageNet with 1000 labels. The superior performance benefits from RDA, which improves pseudo-labels with the co-regularization of complementary class distribution and utilizes the entire unlabeled data, whereas low-confidence samples are filtered out in [23,21].

5.2 Mismatched Distributions

Imbalanced C_x and Balanced C_u . We keep balanced distribution in the unlabeled data and vary N_0 to change the imbalance degree of C_x while the total number of labeled data remains unchanged in the way described in Sec. 4.2. Tab. 2 shows the results on CIFAR-10, CIFAR-100, and mini-ImageNet. RDA outperforms all baseline methods by a large margin. *e.g.*, on CIFAR-10, with 100 labels and $N_0 = 80$, RDA outperforms FixMatch by 7.43% and CoMatch

Table 1. Results of accuracy (%) in the conventional matched SSL setting. Results with * are copied from CoMatch [21] and with † are copied from AlphaMatch [10]. Results of other baselines are based on our reimplementation.

Method	CIFAR-10				mini-ImageNet	STL-10
	20 labels	40 labels	80 labels	100 labels	1000 labels	1000 labels
MixMatch*	27.84±10.63	51.90±11.76	80.79±1.28	-	-	38.02±8.29
AlphaMatch†	-	91.35±3.38	-	-	-	-
FixMatch	84.97±10.37	89.18±1.54	91.99±0.71	93.14±0.76	39.03±0.66	65.38±0.42*
CoMatch	88.43±7.22	93.21±1.55	94.08±0.31	94.55±0.27	43.72±0.58	79.80±0.38*
RDA	92.03±2.01	94.13±1.22	94.24±0.42	94.35±0.25	46.91±1.16	82.63±0.54

Table 2. Results of accuracy (%) in the mismatched scenario with imbalanced C_x (*i.e.*, alter N_0) and balanced C_u . Experiments are conducted on CIFAR-10, CIFAR-100 and mini-ImageNet varying the number of labels and N_0 . Baseline methods are using our reimplementation. Results with **DA** are achieved by combining the original *distribution alignment* in [2]. **Note that CoMatch [21] also integrates DA technique.**

Method	CIFAR-10				CIFAR-100		mini-ImageNet	
	40 labels		100 labels		400 labels	1000 labels	1000 labels	
	$N_0 = 10$	20	40	80	40	80	40	80
FixMatch	85.72±0.93	76.53±3.03	93.01±0.72	71.57±1.88	25.66±0.46	40.22±1.00	36.20±0.36	28.33±0.41
FixMatch w. DA	71.23±1.25	47.85±1.99	56.78±1.28	34.18±0.86	22.66±1.53	31.06±0.51	33.87±0.40	23.53±0.72
CoMatch	60.27±3.22	39.48±2.20	52.82±2.03	26.91±0.75	23.97±0.62	28.35±1.20	30.24±1.37	21.47±0.86
RDA	92.57±0.53	81.78±6.44	94.23±0.36	79.00±2.67	30.86±0.78	41.29±0.43	42.73±0.84	36.73±1.01

by 52.09%. We witness that mismatched C_x and C_u significantly decrease the models’ performance. Notably, the traditional distribution alignment, assuming the labeled and unlabeled data share the same distribution, significantly degrades the performance of model when the distributions mismatch, whereas our method improves this situation by utilizing guidance of distribution information without any prior. As shown in Figs. 3(a) and 3(c), RDA resists the impact of imbalanced C_x and computes a more balanced pseudo-label distribution than FixMatch, demonstrating the effectiveness of RDA in this mismatched distributions scenario. Additionally, Figs. 3(b) and 3(d) show that the predictions of RDA are not necessarily more confident than that of FixMatch, but RDA reduces the overfitting on false pseudo-labels, *i.e.*, RDA is not as overconfident as FixMatch on pseudo-labels that may be wrong. Thanks to no requirement of prior about the labeled data distribution, RDA can be safely applied to this scenario without being overwhelmingly affected by distribution gap, thus exhibiting robust performance. **Mismatched and Imbalanced C_x, C_u .** Results of the more challenging setting are summarized in Tab. 3. While FixMatch and CoMatch fail to correct the severely biased prediction on unlabeled data caused by reversely ordered labeled data, RDA shows its superior performance in this setting and outperforms baseline methods significantly once again. As shown in Figs. 3(e) and 3(g), while imbalanced and mismatched C_x, C_u lead to strong bias on FixMatch’s predictions, RDA shows extraordinary robustness to this scenario. In contrast to FixMatch, RDA prevents overfitting of false pseudo-labels, as shown in Figs. 3(f) and 3(h).

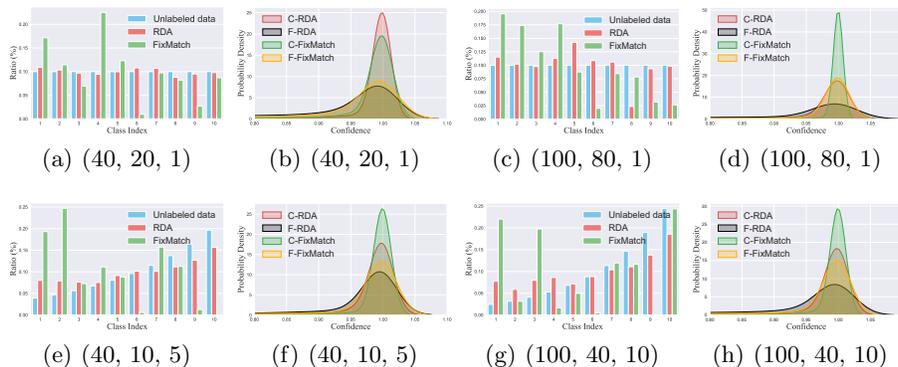


Fig. 3. In the caption, (x,y,z) denotes $(\text{labels}, N_0, \gamma)$. In (a), (b), (c) and (d), C_x is imbalanced and C_u is balanced. In (e), (f), (g) and (h), C_x and C_u are imbalanced and they mismatch. In (a), (c), (e) and (g), the x-axis represents the index of classes in CIFAR-10 and the y-axis represents the ratio of label to the total. *RDA/FixMatch* in figures indicates the class predictions from RDA/FixMatch and *Unlabeled data* indicates the ground-truth label of unlabeled samples. In (b), (d), (f) and (h), the x-axis represents the confidence of prediction from RDA/FixMatch and the y-axis represents the probability density of confidence estimated by kernel density estimation (KDE). *C-X* and *F-X* indicate the correct and false class predictions of X , respectively.

Table 3. Results of accuracy (%) with mismatched and imbalanced C_x , C_u (*i.e.*, alter both N_0 and γ at the same time). Baseline methods are based on our reimplement. We omit the results of baselines that combine DA considering their poor performance.

Method	CIFAR-10				mini-ImageNet
	40 labels, $N_0 = 10$		100 labels, $N_0 = 40$		1000 labels, $N_0 = 40$
	$\gamma = 2$	5	5	10	10
FixMatch	74.97±5.80	64.62±6.13	58.72±3.61	57.49±4.56	21.40±0.53
RDA	88.58±4.05	79.90±2.80	79.33±1.37	70.93±2.91	25.99±0.19

Balanced C_x and Imbalanced C_u . As shown in Tab. 4, RDA shows the compatibility to this scenario and also outperforms baselines combining distribution alignment. Mismatched distributions caused by balanced C_x and imbalanced C_u also lead to poor performance of methods with original distribution alignment.

Other Mismatched Settings. We also show results of RDA within the DARP’s protocol averaged on all five runs. As shown in Tab. 5, RDA consistently outperforms current class-imbalanced SSL method DARP [17] and shows the largest gains in all settings with mismatched C_x and C_u . More discussions on generalized settings of mismatched distributions can be found in Sec. D of Supplementary Material.

Table 4. Accuracy (%) on CIFAR-10 with balanced C_x and imbalanced C_u (*i.e.*, alter γ).

Method	40 labels, $\gamma = 200$
FixMatch w. DA	41.37±1.22
CoMatch	38.85±2.19
RDA	46.50±1.07

Table 5. Accuracy (%) under DARP’s protocol (see Sec. C.1 of Supplementary Material for more details and baselines). WRN-28-2 is adopted as the backbone for all datasets.

Method	CIFAR-10 ($\gamma_l = 100$)				STL-10 ($\gamma_l \neq \gamma_u$)	
	$\gamma_u = 1$	50	150	100 (reversed)	$\gamma_l = 10$	20
FixMatch	68.90±1.95	73.90±0.25	69.60±0.60	65.50±0.05	72.90±0.09	63.40±0.21
DARP	85.40±0.55	77.30±0.17	72.90±0.24	74.90±0.51	77.80±0.33	69.90±0.40
RDA	93.35±0.24	79.77±0.06	74.48±0.24	79.25±0.52	87.21±0.44	83.21±0.52

Table 6. Accuracy (%) of ablation studies on CIFAR-10 with two alternative alignment strategies. “/” represents the conventional setting and $\gamma = 1$ represents balanced C_u .

Method	40 labels			100 labels		
	$N_0, \gamma = /$	20, 1	10, 5	/	80, 1	40, 10
$\mathbb{E}_u(p) \Rightarrow \mathbb{E}_u(\bar{q})$	91.88±1.46	73.54±3.44	74.83±2.99	94.14±0.52	54.88±11.79	62.96±3.43
$\mathbb{E}_u(q) \Rightarrow \mathbb{E}_u(\bar{p})$	93.35±0.12	58.90±3.50	57.38±3.63	94.60±0.08	54.26±4.34	55.39±14.14
RDA	94.13±1.22	81.78±6.44	79.90±2.88	94.35±0.25	79.00±2.67	70.93±2.91

5.3 Ablation Study

To prove the effectiveness of each component in RDA, we conduct ablation studies on CIFAR-10 using consistent experimental setup with Sec. 4.4. We mainly conduct experiments in three settings described in Sec. 4.2 and change the strategy performing distribution alignment from each direction as follows:

$\mathbb{E}_u(p) \Rightarrow \mathbb{E}_u(\bar{q})$. We keep Eq. (13) and discard Eq. (14). *i.e.*, we align distribution of class predictions to “reversed” distribution of complementary predictions.

$\mathbb{E}_u(q) \Rightarrow \mathbb{E}_u(\bar{p})$. We keep Eq. (14) and discard Eq. (13). *i.e.*, we align distribution of complementary predictions to “reversed” distribution of class predictions.

As shown in Tab. 6, the performance of default RDA in mismatched distributions is dominant. RDA helps the model better maximize the objective Eq. (10) while obtaining helpful guidance information of class distribution without prior.

6 Conclusion

In this work, we propose a semi-supervised learning approach which is robust to both the conventional SSL and SSL in mismatched distributions. First, we describe a scenario that has not been discussed extensively by recently-proposed SSL work: mismatched distributions. Second, we improve distribution alignment by proposed RDA so that this technique could be applied into mismatched scenario safely. Then we show RDA results in a form of maximizing the input-out mutual information without any prior information. Finally, we demonstrate that our method outperforms existing baselines significantly under various scenarios.

Acknowledgements. This work is supported by projects from NSFC Major Program (62192783), CAAI-Huawei MindSpore (CAAIXSJLJJ-2021-042A), China Postdoctoral Science Foundation (2021M690609), Jiangsu NSF (BK20210224), and CCF-Lenovo Bule Ocean. Thanks to Prof. Penghui Yao’s helpful discussions.

References

1. Arazo, E., Ortego, D., Albert, P., O’Connor, N.E., McGuinness, K.: Pseudo-labeling and confirmation bias in deep semi-supervised learning. In: International Joint Conference on Neural Networks (2020)
2. Berthelot, D., Carlini, N., Cubuk, E.D., Kurakin, A., Sohn, K., Zhang, H., Raffel, C.: Remixmatch: Semi-supervised learning with distribution matching and augmentation anchoring. In: International Conference on Learning Representations (2020)
3. Berthelot, D., Carlini, N., Goodfellow, I., Papernot, N., Oliver, A., Raffel, C.A.: Mixmatch: A holistic approach to semi-supervised learning. In: Advances in Neural Information Processing Systems (2019)
4. Bridle, J.S., Heading, A.J., MacKay, D.J.: Unsupervised classifiers, mutual information and ‘phantom targets’. In: Advances in Neural Information Processing Systems (1992)
5. Chapelle, O., Scholkopf, B., Zien, A.: Semi-supervised learning. *IEEE Transactions on Neural Networks* **20**(3), 542–542 (2009)
6. Coates, A., Ng, A.Y., Lee, H.: An analysis of single-layer networks in unsupervised feature learning. In: International Conference on Artificial Intelligence and Statistics (2011)
7. Cubuk, E.D., Zoph, B., Shlens, J., Le, Q.: Randaugment: Practical automated data augmentation with a reduced search space. In: Advances in Neural Information Processing Systems (2020)
8. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2009)
9. Duan, Y., Zhao, Z., Qi, L., Wang, L., Zhou, L., Shi, Y., Gao, Y.: Mutexmatch: Semi-supervised learning with mutex-based consistency regularization. *arXiv preprint arXiv:2203.14316* (2022)
10. Gong, C., Wang, D., Liu, Q.: Alphamatch: Improving consistency for semi-supervised learning with alpha-divergence. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2021)
11. Grandvalet, Y., Bengio, Y.: Semi-supervised learning by entropy minimization. In: *Advances in Neural Information Processing Systems* (2005)
12. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2016)
13. Hernán, M.A., Robins, J.M.: *Causal inference: What if*. Boca Raton: Chapman & Hall/CRC (2020)
14. Hofmann, H., Kafadar, K., Wickham, H.: Letter-value plots: Boxplots for large data. *Tech. rep.* (2011)
15. Ishida, T., Niu, G., Hu, W., Sugiyama, M.: Learning from complementary labels. In: *Advances in Neural Information Processing Systems* (2017)
16. Ishida, T., Niu, G., Menon, A.K., Sugiyama, M.: Complementary-label learning for arbitrary losses and models. In: *International Conference on Machine Learning* (2018)
17. Kim, J., Hur, Y., Park, S., Yang, E., Hwang, S.J., Shin, J.: Distribution aligning refinery of pseudo-label for imbalanced semi-supervised learning. In: *Advances in Neural Information Processing Systems* (2020)
18. Kim, Y., Yim, J., Yun, J., Kim, J.: Nlnl: Negative learning for noisy labels. In: *IEEE/CVF International Conference on Computer Vision* (2019)

19. Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images. Tech. rep., University of Toronto (2009)
20. Lee, D.H., et al.: Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In: Workshop on challenges in representation learning, International Conference on Machine Learning (2013)
21. Li, J., Xiong, C., Hoi, S.C.: Comatch: Semi-supervised learning with contrastive graph regularization. In: IEEE/CVF International Conference on Computer Vision (2021)
22. Rizve, M.N., Duarte, K., Rawat, Y.S., Shah, M.: In defense of pseudo-labeling: An uncertainty-aware pseudo-label selection framework for semi-supervised learning. In: International Conference on Learning Representations (2021)
23. Sohn, K., Berthelot, D., Li, C.L., Zhang, Z., Carlini, N., Cubuk, E.D., Kurakin, A., Zhang, H., Raffel, C.: Fixmatch: Simplifying semi-supervised learning with consistency and confidence. In: Advances in Neural Information Processing Systems (2020)
24. Van Engelen, J.E., Hoos, H.H.: A survey on semi-supervised learning. *Machine Learning* **109**(2), 373–440 (2020)
25. Vinyals, O., Blundell, C., Lillicrap, T., Wierstra, D., et al.: Matching networks for one shot learning. In: Advances in Neural Information Processing Systems (2016)
26. Wei, C., Sohn, K., Mellina, C., Yuille, A., Yang, F.: Crest: A class-rebalancing self-training framework for imbalanced semi-supervised learning. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (2021)
27. Xu, Y., Shang, L., Ye, J., Qian, Q., Li, Y.F., Sun, B., Li, H., Jin, R.: Dash: Semi-supervised learning with dynamic thresholding. In: International Conference on Machine Learning (2021)
28. Yang, L., Zhuo, W., Qi, L., Shi, Y., Gao, Y.: Mining latent classes for few-shot segmentation. In: IEEE/CVF International Conference on Computer Vision. pp. 8721–8730 (2021)
29. Zagoruyko, S., Komodakis, N.: Wide residual networks. In: British Machine Vision Conference (2016)
30. Zhang, B., Wang, Y., Hou, W., Wu, H., Wang, J., Okumura, M., Shinozaki, T.: Flexmatch: Boosting semi-supervised learning with curriculum pseudo labeling. In: Advances in Neural Information Processing Systems (2021)
31. Zhao, Z., Zhou, L., Wang, L., Shi, Y., Gao, Y.: Lassl: Label-guided self-training for semi-supervised learning. In: AAAI Conference on Artificial Intelligence (2022)
32. Zhu, X.: Semi-supervised learning. *Encyclopedia of Machine Learning and Data Mining* pp. 1142–1147 (2017)