

# Towards Semi-supervised Learning with Non-random Missing Labels

Yue Duan<sup>1</sup> Zhen Zhao<sup>2</sup> Lei Qi<sup>3</sup> Luping Zhou<sup>2</sup> Lei Wang<sup>4</sup> Yinghuan Shi<sup>1\*</sup>  
<sup>1</sup>Nanjing University <sup>2</sup>University of Sydney <sup>3</sup>Southeast University <sup>4</sup>University of Wollongong

## Abstract

*Semi-supervised learning (SSL) tackles the label missing problem by enabling the effective usage of unlabeled data. While existing SSL methods focus on the traditional setting, a practical and challenging scenario called label Missing Not At Random (MNAR) is usually ignored. In MNAR, the labeled and unlabeled data fall into different class distributions resulting in biased label imputation, which deteriorates the performance of SSL models. In this work, class transition tracking based Pseudo-Rectifying Guidance (PRG) is devised for MNAR. We explore the class-level guidance information obtained by the Markov random walk, which is modeled on a dynamically created graph built over the class tracking matrix. PRG unifies the historical information of class distribution and class transitions caused by the pseudo-rectifying procedure to maintain the model's unbiased enthusiasm towards assigning pseudo-labels to all classes, so as the quality of pseudo-labels on both popular classes and rare classes in MNAR could be improved. Finally, we show the superior performance of PRG across a variety of MNAR scenarios, outperforming the latest SSL approaches combining bias removal solutions by a large margin. Code and model weights are available at <https://github.com/NJUyued/PRG4SSL-MNAR>.*

## 1. Introduction

Semi-supervised learning (SSL), which is in the ascendant, yields promising results in solving the shortage of large-scale labeled data [4, 29]. Current prevailing SSL methods [2, 27, 38, 7, 11] utilize the model trained on the labeled data to impute pseudo-labels for the unlabeled data, thereby boosting the model performance. Although these methods have made exciting advances in SSL, they only work well in the conventional setting, *i.e.*, the labeled and unlabeled data fall into the same (balanced) class distribution. Once this setting is not guaranteed, the gap between the class distributions of the labeled and unlabeled data will lead to a significant accuracy drop of the pseudo-labels, resulting in strong confirmation bias [1] which ultimately corrupts the performance




Figure 1: An example of the MNAR scenarios on CIFAR-10 (see Sec. 4 for details). The class distribution of total data is balanced whereas labeled data is unevenly distributed across classes. For better illustration, the y-axis has different scaling for labeled (blue) and unlabeled data (green).

of SSL models. [14] originally terms the scenario of the labeled and unlabeled data belonging to mismatched class distributions as label Missing Not At Random (MNAR) and proposes an unified doubly robust framework to train an unbiased SSL model in MNAR. During the same period, [39, 6] also independently explored the issue of mismatched distributions. For example, a typical MNAR scenario is shown in Fig. 1, in which the popular classes of labeled data cause the model to ignore the rare classes, increasingly magnifying the bias in label imputation on the unlabeled data. It is worth noting that although some recent SSL methods [15, 32] are proposed to deal with the class imbalance, they are still built upon the assumption of the matched class distributions between the labeled and unlabeled data, and their performance inevitably declines in MNAR.

MNAR is a more realistic scenario than the conventional SSL setting. In the practical labeling process, labeling all classes uniformly is usually not affordable because some classes are more difficult to recognize [26, 23]. Meanwhile, most automatic data collection methods also have difficulty in ensuring that the collected labeled data is balanced [22, 14]. In a nutshell, MNAR is almost inevitable in SSL. In MNAR, the tricky troublemaker is the mismatched class distributions between the labeled and unlabeled data. Training under MNAR, the model increasingly favors some classes, seriously affecting the *pseudo-rectifying* procedure. Pseudo-rectifying is defined as the change of the label assignment decision made by the SSL model for the same sample according to the knowledge learned at each new epoch. This

\*Corresponding author: syh@nju.edu.cn.

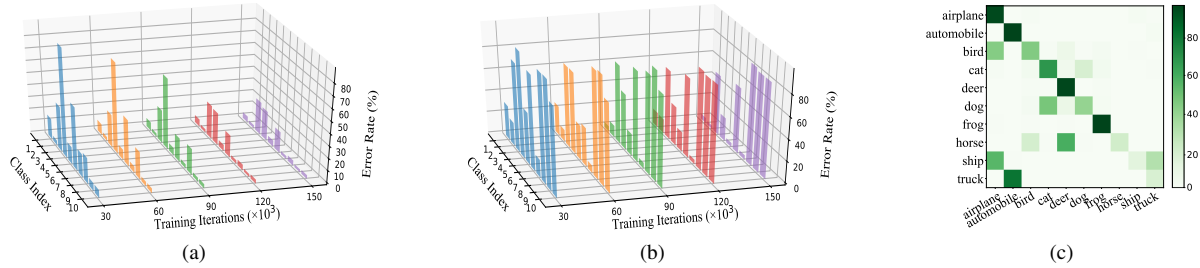


Figure 2: Results of FixMatch [27] in MNAR and the conventional SSL setting (*i.e.*, balanced labeled and unlabeled data). The models are trained on CIFAR-10 with WRN-28-2 backbone [37]. (a) and (b): Class-wise pseudo-label error rate. (c): Confusion matrix of pseudo-labels. In (b) and (c), experiments are conducted with the setting of Fig. 1, whereas in (a) with the conventional setting. The label amount used in (a) is the same as that in (b) and (c).

process may cause *class transition*, *i.e.*, given a sample, its class prediction at the current epoch is different from that at the last epoch. In the self-training process of the SSL model driven by the labeled data, the model is expected to gradually rectify the pseudo-labels mispredicted for the unlabeled data in last epochs. With pseudo-rectifying, the model trapped in the learning of extremely noisy pseudo-labels will be rescued due to its ability to correct these labels.

Unfortunately, the pseudo-rectifying ability of the SSL model could be severely perturbed in MNAR. Taking the setting in Fig. 1 for example, the model’s “confidence” in predicting the pseudo-labels into the labeled rare classes is attenuated by over-learning the samples of the labeled popular classes. Thus, the model fails to rectify those pseudo-labels mispredicted as the popular classes to the correct rare classes (even if the class distribution is balanced in unlabeled data). As shown in Fig. 2b, compared with FixMatch [27] trained in the conventional setting (Fig. 2a), FixMatch trained in MNAR (Fig. 1) significantly deteriorates its pseudo-rectifying ability. Even after many iterations, the error rates of the pseudo-labels predicted for labeled rare classes remain high. This phenomenon hints the necessity to provide additional guidance to the rectifying procedure to address MNAR. Meanwhile, as observed in Fig. 2c, we notice that the mispredicted pseudo-labels for each class are often concentrated in a few classes, rather than scattered across all other classes. Intuitively, a class can easily be confused with the classes similar to it. For example, as shown in Fig. 2c, the “automobile” samples are massively mispredicted as the most similar class: “truck”. Inspired by this, we argue that it is feasible to guide pseudo-rectifying from the *class level*, *i.e.*, pointing out the latent direction of class transition based on its current class prediction only. For instance, given a sample classified as “truck”, the model could be given a chance to classify it as “automobile” sometimes, and vice versa. Notably, our approach *does not* require predefined semantically similar classes. We believe that two classes are conceptually similar only if they are frequently misclassified to each other

by the classifier. In this sense, we develop a novel definition of the similarity of two classes, which is directly determined by model’s output. Even if there are no semantically similar classes, as long as the model makes incorrect prediction during the training, this still leads to class transitions which has seldom been investigated before. Our intuition could be regarded as perturbations on some confident class predictions to preserve the pseudo-rectifying ability of the model. Such a strategy does not rely on the matched class distributions assumption and therefore is amenable to MNAR.

Given the motivations above, we propose class transition tracking based Pseudo-Rectifying Guidance (**PRG**) to address SSL in MNAR, which is shown in Fig. 3. Our main idea can be presented as dynamically tracking the class transitions caused by pseudo-rectifying procedures at previous epoch to provide the class-level guidance for pseudo-rectifying at next epoch. We argue that every class transition of each pseudo-label could become the cure for the deterioration of the pseudo-rectifying ability of the traditional SSL methods in MNAR. A graph is first built on the class tracking matrix recording each pseudo-label’s class transitions occurring in pseudo-rectifying procedure. Then we propose to model the class transition by the Markov random walk, which brings information about the difference in the propensity to rectify pseudo-labels of one class into various other classes. Specifically, we guide the class transitions of each pseudo-label during the rectifying process according to the transition probability corresponding to the current class prediction. The probability is obtained by the transition matrix of Markov random walk, and it is also rescaled based on the class distribution of assigned pseudo-labels to better provide class-level guidance. PRG recalls classes that are easily overlooked but appear in class transition history. They are deemed as similar to the ground-truth and have more chance to be assigned rather than simply letting the model assign the classes it favors without hesitation. In sum, PRG perturbs some confident class predictions to preserve the pseudo-rectifying ability of the model with the usage of tran-

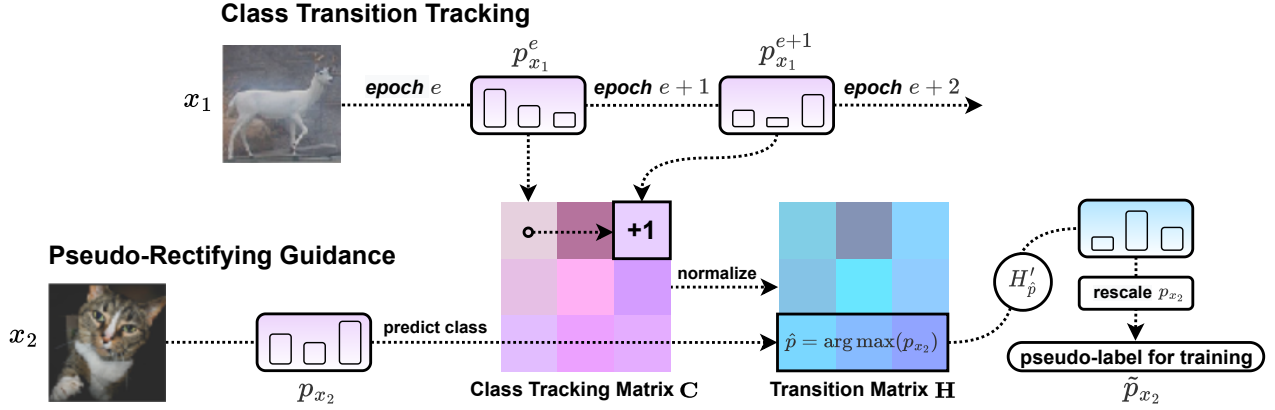


Figure 3: Overview of PRG. Class tracking matrix  $\mathbf{C}$  is obtained by tracking the class transitions of pseudo-labels (e.g.,  $p_{x_1}$  for sample  $x_1$ ) between epoch  $e$  and epoch  $e + 1$  caused by pseudo-rectifying procedure (Eq. (5)). The Markov random walk defined by transition matrix  $\mathbf{H}$  (each row  $H_i$  represents the transition probability vector corresponding to class  $i$ ) is modeled on the graph constructed over  $\mathbf{C}$ . Generally, given a pseudo-label, e.g.,  $p_{x_2}$  for sample  $x_2$ , class- and batch-rescaled  $\mathbf{H}$  (i.e.,  $\mathbf{H}'$ ) is utilized to provide the class-level pseudo-rectifying guidance for  $p_{x_2}$  according to its class prediction  $\hat{p} = \arg \max(p_{x_2})$  (Eqs. (6)~(7)). Finally, the rescaled pseudo-label  $\tilde{p}_{x_2}$  is used for the training.

sition history and class distribution information, which could help improve the quality of pseudo-labels suffered from biased label imputation caused by MNAR. We evaluate PRG on several widely-used SSL benchmarks, demonstrating its effectiveness in coping with SSL in MNAR.

- **What is the novelty and contribution?** Towards addressing SSL in MNAR, we propose transition tracking based Pseudo-Rectifying Guidance (**PRG**) to mitigate the adverse effects of mismatched distributions via combining information from the class transition history. We propose that the pseudo-rectifying guidance can be carried out from the class level, by modeling the class transition of the pseudo-label as a Markov random walk on the graph.
- **How about the performance improvement?** Our solution is computation and memory friendly without introducing additional network components. PRG achieves superior performance in MNAR under various protocols and, e.g., it outperforms CADR [14], a newly-proposed method for addressing MNAR, by up to 15.11% on CIFAR-10 and 15.21% on mini-ImageNet in accuracy.

## 2. Related Work

Addressing the issue of supervised learning requiring a large amount of labeled data has always been a focal point [10, 31, 21, 35]. *Semi-supervised learning* (SSL) is a promising paradigm to address this problem by effectively utilizing the unlabeled data. Given an input  $x$  (labeled or unlabeled data), our objective in SSL can be described as the learning of a predictor for generating label  $y$  for it. In the conventional SSL [19, 3, 33, 28, 25, 40], underlying most of them

is the assumption: the distributions of labeled and unlabeled data are all balanced. Some more practical scenarios for SSL are now extensively discussed. Recently, some work has focused on addressing the class-imbalanced issue in SSL [15, 32]. [15] refines the pseudo-labels softly by formulating a convex optimization. [32] proposes class-rebalancing self-training combining *distribution alignment*. However, these existing methods still underestimate the complexity of practical scenarios of SSL, e.g., [32] works based on strong assumptions: the labeled data and unlabeled data fall in the same distribution (i.e., their distributions match). Furthermore, [14, 39, 6] propose a novel and realistic setting: SSL with mismatched distributions, which pops up in various fields such as social analysis, medical sciences and so on [8, 13]. Typically, [14] designs a class-aware doubly robust (CADR) estimator to remove the bias on label imputation caused by the mismatched distributions. Differently, our method alleviates the bias from another perspective, that is to guide the pseudo-rectifying direction based on the historical information of class transitions.

## 3. Method

Formally, we denote the input space as  $\mathcal{X}$  and the label space as  $\mathcal{Y} = \{1, \dots, k\}$  over  $k$  classes. SSL can be reviewed as a label missing problem, and following [14], *label missing indicator* set is defined as  $\mathcal{M}$  with  $m \in \{0, 1\}$  (only used to define the MNAR setting), where  $m = 1$  indicates label is missing and  $m = 0$  is the otherwise. Given the training dataset in SSL, we obtain a set of labeled data:  $D_L \subseteq \mathcal{X} \times \mathcal{Y} \times \mathcal{M}$  and a set of unlabeled data:  $D_U \subseteq \mathcal{X} \times \hat{\mathcal{Y}} \times \mathcal{M}$ . Since the ground-truth  $y_U \in \hat{\mathcal{Y}}$




Figure 4: Visualization of  $\mathbf{C}$  obtained in training process of FixMatch [27] on CIFAR-10 with the same setting as in Figs. 2c and 2b (results with PRG are in Supplementary Material). The darker the color, the more frequent the class transitions. Overall, the number of class transitions decreases as the training progresses. Class transitions occur intensively between the popular classes, and class transitions between the rare classes gradually disappear (e.g., between “ship” and “truck”).

of unlabeled data  $x_U$  is inaccessible in SSL, prevailing self-training based SSL methods impute  $y_U$  with pseudo-label  $p$ .  $p = f(x_U; \theta)$  is predicted by the model which is parametrized by  $\theta$  and trained on the labeled data. Let  $(x_L^{(i)}, y_L^{(i)}, m_L^{(i)}) \in D_L, i \in \{1, \dots, n_L\}$  be the labeled data pairs consisting of the sample with corresponding ground-truth label (i.e.,  $m^{(i)} = 0$ ), and  $(x_U^{(i)}, y_U^{(i)}, m_U^{(i)}) \in D_U, i \in \{n_L + 1, \dots, n_T\}$  be the unlabeled data missing labels (i.e.,  $m^{(i)} = 1$ ), where  $n_L$  and  $n_T$  refer to the number of labeled data and total training data respectively. Hereafter, the SSL dataset can be defined as  $D = D_L \cup D_U$ . and we can review the conventional SSL as a optimization task for loss  $\mathcal{L}$ :

$$\min_{\theta} \sum_{(x,y,m) \in D} \mathcal{L}(x,y;\theta), \quad (1)$$

where  $D$  is a dataset with independent  $\mathcal{Y}$  and  $\mathcal{M}$ . In this sense, the model trained on  $D_L$  can easily impute unbiased pseudo-labels for unlabeled data  $x_U$  [14]. Conversely, the scenario where  $\mathcal{M}$  is dependent with  $\mathcal{Y}$ , namely label Missing Not At Random (MNAR), will make the model produce strong bias on label imputation, which causes the ability of pseudo-rectifying suffer greatly. Take the current most popular SSL method FixMatch [27] as an example. In FixMatch, the term  $\mathcal{L}(x,y;\theta)$  in Eq. (1) can be decomposed into two loss terms  $\mathcal{L}_L$  and  $\mathcal{L}_U$  with a confidence threshold  $\tau$

$$\mathcal{L}(x,y;\theta) = \mathcal{L}_L(x_L, y_L; \theta) + \lambda_U \mathbb{1}(\max(p) \geq \tau) \mathcal{L}_U(x_U, \arg \max(p); \theta), \quad (2)$$

where  $\lambda_L$  is the unlabeled loss weight and  $\mathbb{1}(\cdot)$  is the indicator function. Training with MNAR setting in Fig. 1, FixMatch is gradually seduced by samples predicted to be the labeled popular classes with confidence above  $\tau$  (even though most of them are wrong), while samples predicted to be the rare class with confidence below  $\tau$  do not participate into training, resulting in biased propensity on label imputation. In this work, we propose class transition tracking based Pseudo-Rectifying Guidance (PRG) to help model better self-correct pseudo-labels with additional guidance.

### 3.1. Pseudo-Rectifying Guidance

Firstly, we formally describe the pseudo-rectifying process in SSL. In this paper, label assignment is considered as a procedure for generating soft labels. We denote the  $i$ -th component of vector  $x$  as  $x_i$ . Let  $p \in \mathbb{R}_+^k$  be the soft label vector assigned to unlabeled data  $x_U$ , where  $\mathbb{R}_+$  is the set of nonnegative real numbers and  $\sum_{i=1}^k p_i = 1$ . Denoting  $x$  at epoch  $e$  as  $x^e$ , the pseudo-rectifying process can be described as the change on  $p$  by the next epoch:  $p^{e+1} = g_{\theta}(p^e)$ , where  $g_{\theta}(p^e)$  is a mapping from  $p^e$  to  $p^{e+1}$  determined by the knowledge learned from the model parametrized by  $\theta$  at epoch  $e + 1$ . In MNAR, take imbalanced  $D_L$  and balanced  $D_U$  as an example, as the training progresses, the model’s confidence is gradually slashed and unexpectedly grows on the rare and popular classes in  $D_L$  respectively. To address this issue, it is necessary to provide more guidance to assist the model in pseudo-rectifying. In general, the Pseudo-Rectifying Guidance (PRG) can be described as

$$\tilde{p}^{e+1} = \text{Normalize}(\eta \circ g_{\theta}(p^e)), \quad (3)$$

where  $\circ$  is Hadamard product, scaling weight vector  $\eta \in \mathbb{R}_+^k$  and  $\text{Normalize}(x)_i = x_i / \sum_{j=1}^k x_j$ .

We can review the technical contributions of some popular self-training works as obtaining more effective  $\eta$  for pseudo-rectifying. For example, pseudo-labeling based methods [19, 27, 20, 34, 38] set  $\eta_i = 1/p_i^{e+1}, i \in \{i \mid i = \arg \max(p^{e+1}) \wedge p_i^{e+1} \geq \tau\}$  and  $\eta_j = 0, j \in \{j \mid j \in (1, \dots, k) \wedge j \neq i\}$  and, i.e., using a confidence threshold to filter low-confidence samples. However, it is difficult to set an apposite  $\eta$  at the sample-level (e.g., for simplicity, [27] fixes  $\tau$  to determine  $\eta$  for all samples and the value of  $\tau$  is usually set based on experience) to guide pseudo-rectifying, especially in the MNAR settings. In addition, some variants of class-balancing algorithms [2, 20, 9] can be integrated into pseudo-rectifying framework. These methods utilize *distribution alignment* to make the class distribution of predictions close to the prior distribution (e.g., the distribution of labeled data). This process can be summa-

---

**Algorithm 1: PRG: Class Transition Tracking Based Pseudo-Rectifying Guidance**


---

**Input:** class tracking matrices  $\mathcal{C} = \{\mathbf{C}^{(i)}; i \in (1, \dots, N_B)\}$ , labeled training dataset  $D_L$ , unlabeled training dataset  $D_U$ , model  $\theta$ , label bank  $\{l^{(i)}; i \in (1, \dots, n_T - n_L)\}$

```

1 for  $n = 1$  to MaxIteration do
2   From  $D_L$ , draw a mini-batch  $\mathcal{B}_L = \{(x_L^{(b)}, y_L^{(b)}); b \in (1, \dots, B)\}$ 
3   From  $D_U$ , draw a mini-batch  $\mathcal{B}_U = \{(x_U^{(b)}); b \in (1, \dots, B_U)\}$ 
4    $\mathbf{H} = \text{RowWiseNormalize}(\text{Average}(\mathcal{C}))$  // Construct transition matrix
5    $H'_{ij} = \frac{H_{ij}}{\frac{L_j}{\sum_{d=1}^k L_d}}$  // Rescale  $\mathbf{H}$  at class-level
6   for  $b = 1$  to  $B_U$  do
7      $p^{(b)} = f_\theta(x_U^{(b)})$  // Compute model prediction
8      $\text{idx} = \text{Index}(x_U^{(b)})$  // Obtain the index of  $x_U^{(b)}$  in  $D_U$ 
9      $\hat{p}^{(b)} = \arg \max(p^{(b)})$  // Compute class prediction
10    if  $l^{(\text{idx})} \neq \hat{p}^{(b)}$  then
11       $C_{l^{(\text{idx})}\hat{p}^{(b)}}^{(n)} = C_{l^{(\text{idx})}\hat{p}^{(b)}}^{(n)} + 1$  // Perform class transition tracking
12       $l^{(\text{idx})} = \hat{p}^{(b)}$ 
13    end
14     $\tilde{p}^{(b)} = \text{Normalize}(H'_{\tilde{p}^{(b)}} \circ p^{(b)})$  // Perform pseudo-rectifying guidance
15  end
16   $\mathcal{L}_L, \mathcal{L}_U = \text{FixMatch}(\mathcal{B}_L, \mathcal{B}_U, \{\tilde{p}^{(b)}; b \in (1, \dots, B_U)\})$  // Run FixMatch
17   $\theta = \text{SGD}(\mathcal{L}_L + \mathcal{L}_U, \theta)$  // Update model parameters  $\theta$ 
18 end

```

---

alized as dataset-level pseudo-rectifying guidance by setting  $\eta$  as the ratio of the current class distribution of predictions to the prior distribution, *i.e.*, the fixed  $\eta$  are used for all samples. Performing pseudo-rectifying guidance in this way strongly relies on an ideal assumptions: the labeled data and unlabeled data share the same class distribution, *i.e.*, in  $D$ ,  $\mathcal{Y}$  is independent with  $\mathcal{M}$ . Thus, these approaches fail miserably in the MNAR scenarios, which can be demonstrated in Sec. D.1 of Supplementary Material. As we discussed in Sec. 1, it is also feasible to guide pseudo-rectifying at class level. Hence, we define rectifying weight matrix as  $\mathbf{A} \in \mathbb{R}_+^{k \times k}$ , where each row  $A_i$  is representing the rectifying weight vector corresponding to class  $i$ . Denoting the class prediction as  $\hat{p} = \arg \max(p)$ , the class-level pseudo-rectifying guidance can be conducted by plugging  $A_{\hat{p}^{e+1}}$  into  $\eta$  in Eq. (3):

$$\tilde{p}^{e+1} = \text{Normalize}(A_{\hat{p}^{e+1}} \circ g_\theta(p^e)). \quad (4)$$

Next, we will introduce a simple and feasible way to obtain an effective  $\mathbf{A}$  for PRG to improve the pseudo-labels.

### 3.2. Class Transition Tracking

Firstly, we consider building a fully connected graph  $G$  in class space  $\mathcal{Y}$ . This graph is constructed by adjacency matrix  $\mathbf{C} \in \mathbb{R}_+^{k \times k}$  (dubbed as class tracking matrix), where each element  $C_{ij}$  represents the frequency of class transitions that occur from class  $i$  to class  $j$  (*i.e.*, an edge directed from vertex  $i$  to vertex  $j$  on  $G$ ).  $C_{ij}$  is parametrized by *class*

*transition tracking* on last  $N_B$  batches with unlabeled data batch size  $B_U$ :  $C_{ij} = \sum_{n=1}^{N_B} C_{ij}^{(n)} / N_B$ , *i.e.*,

$$C_{ij}^{(n)} = \left| \left\{ \hat{p}^{(b)} \mid \hat{p}^{(b),e} = i, \hat{p}^{(b),e+1} = j, i \neq j, b \in \{1, \dots, B_U\} \right\} \right|, \quad (5)$$

where  $n \in \{1, \dots, N_B\}$  and  $C_{ii}^{(n)} = 0$ . Hereafter, we define the Markov random walk along the nodes of  $G$ , which is characterized by its transition matrix  $\mathbf{H} \in \mathbb{R}_+^{k \times k}$  and  $H_{ij}$  represents the transition probability for the class prediction  $\hat{p}$  transits from class  $i$  at epoch  $e$  to class  $j$  at epoch  $e + 1$ .  $\mathbf{H}$  is computed by conducting row-wise normalization on  $\mathbf{C}$ . The above designs are desirable for the following reasons.

(1) In the self-training process of the model, the historical information of pseudo-rectifying contains the relationship between classes, which is often ignored in previous methods and can be utilized to help the model assign labels at a new epoch. We can record the class transition trend in pseudo-rectifying by Eq. (5), which corresponds to the transition probability represented by  $H_{ij}$ , *i.e.*, for a sample  $x$ , when its class prediction  $\hat{p}$  is in the state of class  $i$ , if a rectifying procedure resulting in a class transition occurs, what probability will it transit to class  $j$ . Intuitively, given  $p$  with  $\hat{p} = i$ , the model prefers to rectify it to another class similar to class  $i$  in one class transition, *i.e.*, the preference of class transitions can also be regarded as the similarity between

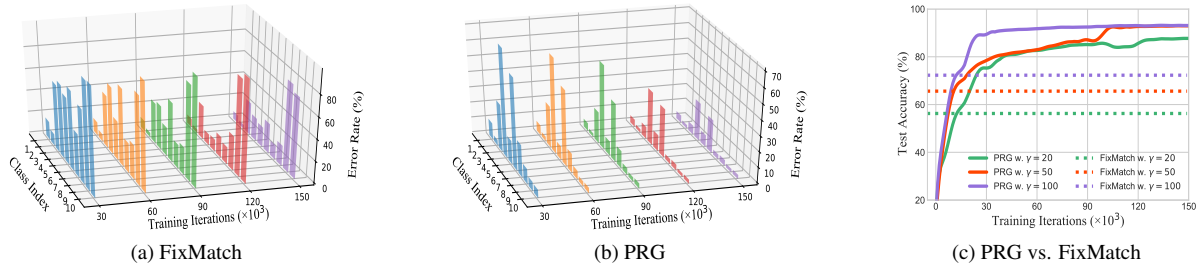


Figure 5: Results on CIFAR-10 under CADR’s protocol. (a) and (b): Class-wise pseudo-label error rate with  $\gamma = 50$ . (c): Learning curve of PRG. We mark the final results of FixMatch as dash lines.

Table 1: Accuracy (%) in MNAR under CADR’s protocol. The results of \* are derived from CADR [14]. The larger  $\gamma$ , the more imbalanced the labeled data. Our accuracies are averaged on 3 runs while the standard deviations ( $\pm$ Std.) and the performance difference ( $\uparrow$ ,  $\downarrow$ , Diff.) compared to original baseline methods are reported (**bold** indicates the best results).

Method	CIFAR-10			CIFAR-100			mini-ImageNet	
	$\gamma = 20$	50	100	50	100	200	50	100
II Model*	21.59	27.54	30.39	24.95	29.93	33.91	11.77	15.30
MixMatch*	26.63	31.28	28.02	37.82	41.32	42.92	13.12	18.30
ReMixMatch*	41.84	38.44	38.20	42.45	39.71	39.22	22.64	23.50
FixMatch*	56.26	65.61	72.28	50.51	48.82	50.62	23.56	26.57
+ Crest*	51.10 $\downarrow$ <sup>5.16</sup>	55.40 $\downarrow$ <sup>10.21</sup>	63.60 $\downarrow$ <sup>8.68</sup>	40.30 $\downarrow$ <sup>10.21</sup>	46.30 $\downarrow$ <sup>2.52</sup>	49.60 $\downarrow$ <sup>1.02</sup>	–	–
+ DARP*	63.14 $\uparrow$ <sup>6.88</sup>	70.44 $\uparrow$ <sup>4.83</sup>	74.74 $\uparrow$ <sup>2.46</sup>	38.87 $\downarrow$ <sup>11.64</sup>	40.49 $\downarrow$ <sup>8.33</sup>	44.15 $\downarrow$ <sup>6.47</sup>	–	–
+ CADR*	79.63 $\uparrow$ <sup>23.37</sup>	93.79 $\uparrow$ <sup>23.37</sup>	93.97 $\uparrow$ <sup>21.69</sup>	59.53 $\uparrow$ <sup>9.02</sup>	60.88 $\uparrow$ <sup>12.06</sup>	63.30 $\uparrow$ <sup>12.68</sup>	29.07 $\uparrow$ <sup>5.51</sup>	32.78 $\uparrow$ <sup>6.21</sup>
+ PRG (Ours)	<b>94.04</b> $\uparrow$ <sup>37.78</sup> $\pm 0.18$	<b>94.09</b> $\uparrow$ <sup>28.48</sup> $\pm 0.18$	94.28 $\uparrow$ <sup>22.00</sup> $\pm 0.22$	59.11 $\uparrow$ <sup>8.60</sup> $\pm 0.54$	61.84 $\uparrow$ <sup>13.02</sup> $\pm 0.45$	63.41 $\uparrow$ <sup>12.79</sup> $\pm 4.08$	44.28 $\uparrow$ <sup>20.72</sup> $\pm 0.54$	44.99 $\uparrow$ <sup>18.42</sup> $\pm 1.25$
+ PRG <sup>Last</sup> (Ours)	93.81 $\uparrow$ <sup>37.55</sup> $\pm 0.98$	93.44 $\uparrow$ <sup>27.83</sup> $\pm 1.05$	93.48 $\uparrow$ <sup>21.20</sup> $\pm 0.79$	59.54 $\uparrow$ <sup>9.03</sup> $\pm 0.99$	62.36 $\uparrow$ <sup>13.54</sup> $\pm 0.23$	60.56 $\uparrow$ <sup>11.86</sup> $\pm 1.86$	40.73 $\uparrow$ <sup>17.17</sup> $\pm 1.27$	43.89 $\uparrow$ <sup>17.32</sup> $\pm 0.14$
SimMatch	83.45 $\pm 2.32$	86.77 $\pm 2.15$	90.12 $\pm 1.90$	60.06 $\pm 1.17$	60.35 $\pm 0.59$	61.14 $\pm 0.24$	39.49 $\pm 1.04$	40.37 $\pm 0.96$
+ PRG (Ours)	86.87 $\uparrow$ <sup>3.42</sup> $\pm 2.38$	91.68 $\uparrow$ <sup>4.91</sup> $\pm 1.59$	<b>94.59</b> $\uparrow$ <sup>4.47</sup> $\pm 0.32$	<b>65.65</b> $\uparrow$ <sup>5.59</sup> $\pm 0.61$	<b>65.89</b> $\uparrow$ <sup>5.54</sup> $\pm 0.30$	66.50 $\uparrow$ <sup>5.36</sup> $\pm 0.17$	<b>44.61</b> $\uparrow$ <sup>5.12</sup> $\pm 0.24$	<b>46.48</b> $\uparrow$ <sup>6.11</sup> $\pm 0.15$
+ PRG <sup>Last</sup> (Ours)	86.46 $\uparrow$ <sup>3.01</sup> $\pm 1.75$	90.48 $\uparrow$ <sup>3.71</sup> $\pm 0.96$	94.22 $\uparrow$ <sup>4.10</sup> $\pm 0.14$	65.10 $\uparrow$ <sup>5.04</sup> $\pm 0.40$	65.52 $\uparrow$ <sup>5.17</sup> $\pm 0.29$	<b>66.62</b> $\uparrow$ <sup>5.48</sup> $\pm 0.19$	42.06 $\uparrow$ <sup>2.57</sup> $\pm 1.81$	44.86 $\uparrow$ <sup>9.95</sup> $\pm 4.49$

classes and the more similar two classes are, the more likely they are to be misclassified as each other’s classes. The label is more likely to oscillate between the two classes, resulting in more swinging class transitions. As shown in Fig. 4, in the “dog” class predictions, the predictions transitioning to the “cat” class are significantly more than to other classes, and vice versa in the “cat” labels. We can observe that  $\mathbf{C}$  behaves like a symmetric matrix, reflecting the symmetric nature of class similarity. Consequently, this similarity between classes can be utilized to provide information for our class-level pseudo-rectifying guidance.

(2) In the MNAR settings, the tricky problem is that the mismatched distributions lead to biased label imputation for unlabeled data. The feedback loop of self-reinforcing errors is not achieved overnight. Empirically, as the training progresses, the model becomes more and more confident in the popular classes (in labeled or unlabeled data), which leads to misclassify the samples that it initially thought might be the rare classes to the popular classes later. As shown in Fig. 4, the lower left corner and upper right corner of the heatmap (*i.e.*, the class transitions between the popular

classes and rare classes) is getting lighter and always lighter than the upper left corner (*i.e.*, the class transitions among the popular classes), which means the model is increasingly reluctant to transfer the class prediction to the rare classes during the pseudo-rectifying process. If we only focus on what the model has learned at present, the model’s past efforts to recognize the rare classes will be buried. The latent relational information between classes is hidden in the pseudo-rectifying process producing class transitions. The history of class transitions can point the way for bias removal on label imputation with an abnormal propensity on different classes caused by mismatched distributions in MNAR.

With obtained  $\mathbf{H}$ , some preparations are done for plugging it into Eq. (4) to replace  $\mathbf{A}$ . We’re only modeling the pseudo-rectifying process resulting in class transition (*i.e.*,  $C_{ii} = 0$ ), which means  $H_{ii} = 0$ , *i.e.*,  $\eta_{\hat{p}^{e+1}}$  is set to 0 in Eq. (3). This will encourage the class prediction to transition to other classes during each pseudo-rectifying process, which is unreasonable for training a robust classifier. Hence, we control the probability that does not transition class by setting  $H_{ii} = \frac{\alpha}{k-1}$ , where  $\frac{1}{k-1}$  is the average of the transi-

Table 2: Accuracy (%) in MNAR under our protocol with the varying labeled data sizes  $n_L$  and imbalanced ratios  $N_1$ . Baseline methods are based on our reimplemention.

Method	CIFAR-10 ( $n_L = 40$ )		CIFAR-10 ( $n_L = 250$ )		CIFAR-100 ( $n_L = 2500$ )		mini-ImageNet ( $n_L = 1000$ )	
	$N_1 = 10$	20	100	200	100	200	40	80
FixMatch	85.72 $\pm$ 0.93	76.53 $\pm$ 3.03	69.76 $\pm$ 5.57	46.53 $\pm$ 8.12	61.31 $\pm$ 3.67	41.38 $\pm$ 2.84	36.20 $\pm$ 0.36	28.33 $\pm$ 0.41
+ CADR	85.54 $\pm$ 2.07	75.11 $\pm$ 3.41	92.25 $\pm$ 1.61	63.92 $\pm$ 5.47	<b>61.62</b> $\pm$ 0.31	46.16 $\pm$ 1.45	36.08 $\pm$ 0.84	30.52 $\pm$ 0.99
+ PRG (Ours)	<b>91.87</b> $\pm$ 1.05	77.44 $\pm$ 15.96	<b>93.93</b> $\pm$ 0.16	<b>67.86</b> $\pm$ 16.98	61.49 $\pm$ 3.93	<b>49.84</b> $\pm$ 1.37	<b>39.99</b> $\pm$ 0.76	<b>35.39</b> $\pm$ 0.47
+ PRG <sup>Last</sup> (Ours)	85.66 $\pm$ 5.93	<b>77.85</b> $\pm$ 1.86	92.80 $\pm$ 1.44	64.00 $\pm$ 5.02	60.41 $\pm$ 1.01	43.80 $\pm$ 1.71	39.84 $\pm$ 0.05	33.17 $\pm$ 0.52

tion probabilities in each row of  $\mathbf{H}$  and  $\alpha$  is a pre-defined hyper-parameter. In addition, to better provide class-level guidance, we scale each element in  $\mathbf{H}$  by

$$H'_{ij} = \frac{H_{ij}}{\sum_{d=1}^k L_d} \quad (6)$$

where  $L \in \mathbb{R}_+^k$  and  $L_i$  records the number of class predictions belonging to class  $i$  averaged on last  $N_B$  batches. To sum up, we are trying to fight against MNAR by yielding the effect of adjusting the class distribution of pseudo-labels (*i.e.*,  $\frac{L_c}{\sum_{d=1}^k L_d}$ ). After all, the trouble MNAR brings us is biased label imputation caused by mismatched class distribution. We encourage the model to perform the pseudo-rectifying process, which leads to more labels being transition to classes with too few assigned labels, rather than ignoring the rare classes due to over-learning of popular classes. More explanations and alternatives of Eq. (6) can be found in Sec. B of Supplementary Material. Hereafter, we plug  $\mathbf{H}'$  into  $\mathbf{A}$  in Eq. (4) for pseudo-rectifying guidance framework:

$$\tilde{p}^{e+1} = \text{Normalize}(H'_{\tilde{p}^{e+1}} \circ g_\theta(p^e)), \quad (7)$$

where  $H'_{\tilde{p}^{e+1}}$  can be regarded as the class prediction for one sample randomly walks along the nodes of  $C$  at the current epoch, *i.e.*, drive a possible class transition in the pseudo-rectifying for bias removal on label imputation propensity due to MNAR. It is also feasible to use the class transition driven by  $\tilde{p}^e$  to revise  $p^{e+1}$  (what is the class prediction after a class transition), *i.e.*, replace  $H'_{\tilde{p}^{e+1}}$  in Eq. (7) with  $H'_{\tilde{p}^e}$ , which is dubbed as **PRG<sup>Last</sup>**. The algorithms of PRG and PRG<sup>Last</sup> are presented in Algorithm 1 and Algorithm 2 in Supplementary Material, respectively.

**Why does our method work for MNAR?** In the setting of MNAR, being aware of rare class plays a key role, PRG enhances the model to preserve a certain probability to generate class transition to rare classes when assigning pseudo-labels. This form of probability based on class transition history produces effective results, because we do not spare any attempt of the model to identify the rare class by class transition tracking (such attempts would be slowly buried due to overlearning of the popular classes). Thereby, PRG helps the model to still try to identify rare classes with

a certain probability while combines the class distribution information of pseudo-labels so that the model can assign labels to rare classes with a clear purpose.

## 4. Experiment

**Dataset and Baselines.** We evaluate PRG on three widely used SSL benchmarks, including CIFAR-10, CIFAR-100 [16] and mini-ImageNet [30] (a subset of ImageNet [5] composed of 100 classes). Following [14], we mainly report the mean accuracy of PRG in both conventional SSL settings and various MNAR scenarios. Multiple baseline methods are compared, including conventional SSL algorithms:  $\Pi$  Model [24], MixMatch [3], ReMixMatch [2], FixMatch [27] and SimMatch [41]. More importantly, we provide comparisons with the recent label bias removal methods for imbalanced SSL: DARP [15], Crest [32], and the latest approach designed for addressing SSL in MNAR: CADR [14].

**MNAR Settings.** Following [14], the MNAR scenarios are mimicked by constructing the class-imbalanced subset of the original dataset for either  $D_L$  or  $D_U$ . Let  $\gamma$  denote the imbalanced ratio,  $N_i$  and  $M_i$  respectively refer to the number of the labeled and the unlabeled data in class  $i$  from  $k$  classes. Three MNAR protocols are used for the evaluations on PRG: (1) CADR’s protocol [14].  $N_i = \gamma^{\frac{k-i}{k-1}}$ , in which  $N_1 = \gamma$  is the maximum number of labeled data in all classes, and the larger the value of  $\gamma$ , the more imbalanced  $D_L$ . For example, Fig. 1 shows CIFAR-10 with  $\gamma = 20$ . (2) Our protocol. Because the total number of labeled data  $n_L$  in the CADR’s protocol varies with  $\gamma$ , which violates the principle of controlling variables,  $n_L$  is fixed by users in our protocol.  $N_1$  is altered for different scales of imbalance, *i.e.*,  $N_i = N_1 \times \gamma^{-\frac{i-1}{k-1}}$  while  $\gamma$  is calculated by the constraint  $\sum_{i=1}^k N_i = n_L$ . We further consider the MNAR settings where  $D_U$  is also imbalanced, *i.e.*,  $M_i = M_1 \times \gamma_u^{-\frac{k-i}{k-1}}$  (implying inversely imbalanced distribution compared with  $D_L$ ), where  $M_1 = 5000$  in CIFAR-10. (3) DARP’s protocol [15]:  $N_i = N_1 \times \gamma_l^{-\frac{i-1}{k-1}}$ ,  $M_i = M_1 \times \gamma_u^{-\frac{i-1}{k-1}}$ , where  $N_1 = 1500$  and  $M_1 = 3000$  in CIFAR-10, where  $\gamma_l$  and  $\gamma_u$  are varied for  $D_L$  and  $D_U$  respectively, *i.e.*, the distributions of  $D_L$  and  $D_U$  are mismatched and imbalanced.




Figure 6: Results on CIFAR-10 under two protocols. The imbalanced distributions of labeled and unlabeled data are in reverse order of each other in (a) and the case of  $\gamma_u$  marked with “R” in (b).

Table 3: Geometric mean scores (GM) in MNAR on CIFAR-10 under CADR’s protocol.

Method	$\gamma = 20$	$\gamma = 50$	$\gamma = 100$
FixMatch	41.90 $\pm$ 8.55	53.61 $\pm$ 2.29	60.35 $\pm$ 1.84
+ CADR	75.25 $\uparrow$ <sup>33.35</sup> $\pm$ 1.55	92.98 $\uparrow$ <sup>39.37</sup> $\pm$ 0.43	93.15 $\uparrow$ <sup>32.8</sup> $\pm$ 0.36
+ PRG (Ours)	<b>93.53</b> $\uparrow$ <sup>18.63</sup> $\pm$ 0.39	<b>93.70</b> $\uparrow$ <sup>40.19</sup> $\pm$ 0.20	<b>93.94</b> $\uparrow$ <sup>33.59</sup> $\pm$ 0.35
+ PRG <sup>Last</sup> (Ours)	93.35 $\uparrow$ <sup>14.45</sup> $\pm$ 1.10	92.99 $\uparrow$ <sup>39.38</sup> $\pm$ 1.17	93.25 $\uparrow$ <sup>32.90</sup> $\pm$ 0.97

**Implementation Details.** In this section, PRG is mainly implemented as a plugin to FixMatch [27] and SimMatch [41]. Thus, we keep the same hyper-parameters as their original paper, whereas the class invariance coefficient  $\alpha = 1$  and the tracked batch number  $N_B = 128$  are set for PRG. The complete list of hyper-parameters can be found in Sec. C of Supplementary Material. Following [27], our models are trained for  $2^{20}$  iterations, respectively using the backbone of WideResNet-28-2 (WRN) [37] for CIFAR-10, WRN-28-8 for CIFAR-100 and ResNet-18 [12] for mini-Imagenet.

**Experimental Results List (SM refers to Supplementary Material).** (1) Imbalanced  $D_L$  and balanced  $D_U$  / Mismatched imbalanced  $D_L$  and  $D_U$ : *Main Results / More MNAR Settings* in Sec. 4.1. (2) Balanced  $D_L$  and imbalanced  $D_U$ : Sec. D.3.1 of SM. (3) Balanced  $D_L$  and  $D_U$ : Sec. D.3.5 of SM. (4) *More Application Scenarios* in Sec. 4.1. (5) Ablation studies: Sec. 4.2. (6) Results with distribution alignment: Sec. D.1. (7) More evaluation metrics (*e.g.*, precision and recall): Sec. D.3.2 of SM. (8) PRG built on other SSL learners: Sec. D.3.3 of SM.

#### 4.1. Results in MNAR Settings

**Main Results.** The experimental results under CADR’s and our protocol with various levels of imbalance are summarized in Tabs. 1 and 2. PRG consistently wins baseline methods across most of the settings, benefiting from the information offered by class transition tracking. As shown in Figs. 5a and 5b, the pseudo-rectifying ability of PRG is significantly improved compared with the original FixMatch, *i.e.*, as the training progresses, the error rates of both the popular classes and the rare classes of the labeled data are greatly reduced, eventually yielding improvements in test

accuracy shown in Fig. 5c. Meanwhile, in Tab. 3 we further provide *geometric mean scores* (GM, a metric often used for imbalanced dataset [17, 15]), which is defined by the geometric mean over class-wise sensitivity for evaluate the classification performance of models trained in MNAR.

Our main competitors include three categories. (1) State-of-The-Art (SOTA) SSL methods: FixMatch [27] and SimMatch [41]. As shown in Tabs. 1 and 2, these methods show poor performance under MNAR. FixMatch almost can’t cope with MNAR, whereas with our method, the performance is significantly improved by more than 10% in most cases. Likewise, SimMatch’s performance is also improved by a large margin. (2) Imbalanced SSL methods: DARP [15] and Crest [32]. These two SOTA methods addressing long-tailed distribution in SSL emphasize the bias removal in matched distribution (*i.e.*, the unlabeled data is equally imbalanced as the labeled data), showing very limited capacity in handling MNAR. (3) SSL solutions devised for the MNAR scenarios: CADR [14]. Our method outperforms CADR under its proposed protocol across the board, demonstrating PRG is more effective for bias removal on label imputation than it. With extremely few labels, the class-aware propensity estimation in CADR is not reliable whereas our method still works well, yielding a performance gap of up to 14.41%.

**More MNAR Settings.** More MNAR scenarios are considered for evaluation. In our protocol, we alter  $N_1$  and  $\gamma_u$  to mimic the case where the distributions of the labeled and unlabeled data are imbalanced and mismatched, *i.e.*, the two distributions are different. Likewise, DARP’s protocol produces similar mismatched distributions. As shown in Fig. 6, PRG achieves promising results in all the comparisons with the baseline methods. Our method boosts the accuracy of FixMatch by up to 35.51% and 24.33% in our and DARP’s protocols respectively. The activated class transitions make the model less prone to over-learning unexpected classes so that the negative effect of MNAR can be mitigated.

**More Application Scenarios.** To explore the broader image recognition applications of PRG, we further apply it to tabular MNIST (handwritten digital image from 10 classes [18] interpreting as tabular data with 784 features) by plugging it into VIME [36] (see Sec. D.3.4 of Supplementary Material for details). As shown in Tab. 4, PRG outperforms



Table 4: Accuracy (%) on tabular MNIST.  $\gamma$  is varied for CADR’s protocol whereas  $n_L$  and  $N_1$  are varied for our protocol.

Method	$\gamma = 20$	50	100	$n_L, N_1 = 40, 10$	40, 20	250, 100	250, 200
VIME	63.38±4.42	63.75±6.10	64.80±2.76	50.13±7.56	30.73±8.69	60.58±2.68	21.44±0.58
+ PRG (Ours)	59.41±14.45	65.92±13.90	<b>66.60±12.58</b>	49.28±11.09	<b>34.08±16.05</b>	<b>66.14±11.88</b>	<b>24.51±9.56</b>
+ PRG <sup>Last</sup> (Ours)	<b>63.49±10.73</b>	<b>66.19±14.22</b>	66.21±10.24	<b>53.17±8.84</b>	32.45±10.10	65.25±9.46	23.62±11.39

Table 5: Ablation studies on re-weighting scheme in Eq. (6). We report the results (accuracy (%) / GM) on CIFAR-10 under CADR’s protocol.

Method	$\gamma = 20$	$\gamma = 50$	$\gamma = 100$
PRG wo. Eq. (6)	88.97 / 87.37	91.73 / 91.28	92.72 / 92.55
PRG	94.04 / 93.53	94.09 / 93.70	94.28 / 93.94

baselines in the most of settings and its upper performance limits consistently exceed VIME by a large margin. Performance fluctuations can be alleviated by adjusting  $\alpha$  (here we keep consistent setting as previous experiments). The scheme of class-transition-based pseudo-rectifying guidance is high-level and general for classification task, so it shows promising potential in a broader MNAR scenario.

## 4.2. Ablation Studies

**Re-weighting scheme on H.** As shown in Tab. 5, the re-weighting scheme can effectively boost the performance of PRG in MNAR because it better provide class-level guidance by involving class distribution information and controls the intensity of class transition. Additionally, for the utilization of  $\mathbf{H}'$  in Eq. (7), we consider taking  $k$  steps, *i.e.*, multiply by  $\mathbf{H}'^k$  instead of  $\mathbf{H}'$  to uncover more complex patterns of misclassification than simple pairwise class relations. However, as shown in Tab. 6, we can observe that the performance is inversely proportional to  $k$ . The advantage of PRG is that  $\mathbf{H}'$  is updated in each iteration, which means that the value of  $\mathbf{H}'$  is dynamic. As the model learns new knowledge, the past  $\mathbf{H}'$  may not be suitable for the pseudo-rectifying process anymore. If  $\mathbf{H}'^k$  is used, this means that we are using the same  $\mathbf{H}'$  multiple times for a given sample, which wastes the advantage of dynamic  $\mathbf{H}'$ .  $\mathbf{H}'^k$  using a suitable  $k$  or a dynamic selection of  $k$  might yield better performance, but it is complicated to determine the value of  $k$ . Therefore, PRG is designed for simplicity and exhibits superior performance.

**Hyper-parameters.** We investigate the effect of the class invariance coefficient  $\alpha$  and the tracked batch number  $N_B$  on PRG, which is shown in Fig. 7. Choosing an appropriate  $\alpha$  to control the degree of class invariance in pseudo-rectifying is important for PRG, which ensures stability of supervision information and training. Meanwhile, we note that too small  $N_B$  is not sufficient to estimate the underlying distribution of class transitions, where  $N_B = 128$  is a sensible choice for both memory overhead and performance.

Table 6: Ablation studies on step  $k$  of driving possible class transitions. We respectively report accuracy (%) and GM on CIFAR-10 under CADR’s protocol with  $\gamma = 20$ .

$k$	1 (default PRG)	2	5	10
Accuracy	94.04	91.33	87.76	82.60
GM	93.53	90.79	85.80	80.24




Figure 7: Ablation studies with  $\alpha$  and  $N_B$  on CIFAR-10 under CADR’s protocol with  $\gamma = 20$ .

## 5. Conclusion

This paper can be concluded as proposing a effective SSL framework called class transition based Pseudo-Rectifying Guidance (PRG) to address SSL in the MNAR scenarios. Firstly, we argue that the history of class transition caused by pseudo-rectifying can be utilized to offer informative guidance for future label assignment. Thus, we model the class transition as a Markov random walk along the nodes of the graph constructed on the class tracking matrix. Finally, we propose to utilize the class prediction information at current epoch (or last epoch) to guide the class transition for pseudo-rectifying so that the bias of label imputation can be alleviated. Except for MNAR, we believe PRG can be used for robust semi-supervised learning in broader scenarios.

## Acknowledgement

Yue Duan and Yinghuan Shi are with the National Key Laboratory for Novel Software Technology and the National Institute of Healthcare Data Science, Nanjing University. This work is supported by the Natural Science Foundation of China Program (62222604, 62206052, 62192783), China Postdoctoral Science Foundation Project (2023T160100), Jiangsu Natural Science Foundation Project (BK20210224), and CCF-Lenovo Bule Ocean Research Fund.

## References

- [1] Eric Arazo, Diego Ortego, Paul Albert, Noel E. O’Connor, and Kevin McGuinness. Pseudo-labeling and confirmation bias in deep semi-supervised learning. *arXiv preprint arXiv:1908.02983*, 2019. 1
- [2] David Berthelot, Nicholas Carlini, Ekin D. Cubuk, Alex Kurakin, Kihyuk Sohn, Han Zhang, and Colin Raffel. Remixmatch: Semi-supervised learning with distribution matching and augmentation anchoring. In *International Conference on Learning Representations*, 2020. 1, 4, 7, 2, 3
- [3] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A. Raffel. Mixmatch: A holistic approach to semi-supervised learning. In *Advances in Neural Information Processing Systems*, 2019. 3, 7
- [4] Olivier Chapelle, Bernhard Scholkopf, and Alexander Zien. Semi-supervised learning. *IEEE Transactions on Neural Networks*, 20(3):542–542, 2009. 1
- [5] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2009. 7
- [6] Yue Duan, Lei Qi, Lei Wang, Luping Zhou, and Yinghuan Shi. Rda: Reciprocal distribution alignment for robust semi-supervised learning. In *European Conference on Computer Vision*, 2022. 1, 3
- [7] Yue Duan, Zhen Zhao, Lei Qi, Lei Wang, Luping Zhou, Yinghuan Shi, and Yang Gao. Mutexmatch: semi-supervised learning with mutex-based consistency regularization. *IEEE Transactions on Neural Networks and Learning Systems*, 2022. 1
- [8] Craig K Enders. *Applied missing data analysis*. Guilford press, 2010. 3
- [9] Chengyue Gong, Dilin Wang, and Qiang Liu. Alphamatch: Improving consistency for semi-supervised learning with alpha-divergence. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021. 4, 2
- [10] Zhangxuan Gu, Siyuan Zhou, Li Niu, Zihan Zhao, and Liqing Zhang. Context-aware feature generation for zero-shot semantic segmentation. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 1921–1929, 2020. 3
- [11] Guan Gui, Zhen Zhao, Lei Qi, Luping Zhou, Lei Wang, and Yinghuan Shi. Improving barely supervised learning by discriminating unlabeled samples with super-class. In *Advances in Neural Information Processing Systems*, 2022. 1
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2016. 8
- [13] James J Heckman. Sample selection bias as a specification error (with an application to the estimation of labor supply functions). Technical report, National Bureau of Economic Research, 1977. 3
- [14] Xinting Hu, Yulei Niu, Chunyan Miao, Xian-Sheng Hua, and Hanwang Zhang. On non-random missing labels in semi-supervised learning. In *International Conference on Learning Representations*, 2022. 1, 3, 4, 6, 7, 8, 5
- [15] Jaehyung Kim, Youngbum Hur, Sejun Park, Eunho Yang, Sung Ju Hwang, and Jinwoo Shin. Distribution aligning refinery of pseudo-label for imbalanced semi-supervised learning. In *Advances in Neural Information Processing Systems*, 2020. 1, 3, 7, 8
- [16] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009. 7
- [17] Miroslav Kubat, Stan Matwin, et al. Addressing the curse of imbalanced training sets: one-sided selection. In *International Conference on Machine Learning*, 1997. 8
- [18] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. 8
- [19] Dong-Hyun Lee 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. 3, 4
- [20] Junnan Li, Caiming Xiong, and Steven CH Hoi. Comatch: Semi-supervised learning with contrastive graph regularization. In *IEEE/CVF International Conference on Computer Vision*, 2021. 4, 2, 3
- [21] Shumeng Li, Heng Cai, Lei Qi, Qian Yu, Yinghuan Shi, and Yang Gao. Pln: Parasitic-like network for barely supervised medical image segmentation. *IEEE Transactions on Medical Imaging*, 42(3):582–593, 2022. 3
- [22] Dhruv Mahajan, Ross Girshick, Vignesh Ramanathan, Kaiming He, Manohar Paluri, Yixuan Li, Ashwin Bharambe, and Laurens Van Der Maaten. Exploring the limits of weakly supervised pretraining. In *European Conference on Computer Vision*, 2018. 1
- [23] Ishan Misra, C Lawrence Zitnick, Margaret Mitchell, and Ross Girshick. Seeing through the human reporting bias: Visual classifiers from noisy human-centric labels. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2016. 1
- [24] Antti Rasmus, Harri Valpola, Mikko Honkala, Mathias Berglund, and Tapani Raiko. Semi-supervised learning with ladder networks. In *International Conference on Neural Information Processing Systems*, 2015. 7
- [25] Mamshad Nayeem Rizve, Kevin Duarte, Yogesh S Rawat, and Mubarak Shah. In defense of pseudo-labeling: An uncertainty-aware pseudo-label selection framework for semi-supervised learning. In *International Conference on Learning Representations*, 2021. 3, 4
- [26] Saharon Rosset, Ji Zhu, Hui Zou, and Trevor Hastie. A method for inferring label sampling mechanisms in semi-supervised learning. In *Advances in Neural Information Processing Systems*, 2005. 1
- [27] Kihyuk Sohn, David Berthelot, Chun-Liang Li, Zizhao Zhang, Nicholas Carlini, Ekin D. Cubuk, Alex Kurakin, Han Zhang, and Colin Raffel. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. In *Advances in Neural Information Processing Systems*, 2020. 1, 2, 4, 7, 8, 3
- [28] Kai Sheng Tai, Peter Bailis, and Gregory Valiant. Sinkhorn label allocation: Semi-supervised classification via annealed

- self-training. In *International Conference on Machine Learning*, 2021. 3
- [29] Jesper E Van Engelen and Holger H Hoos. A survey on semi-supervised learning. *Machine Learning*, 109(2):373–440, 2020. 1
- [30] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. In *Advances in Neural Information Processing Systems*, 2016. 7
- [31] Xiyue Wang, Jinxi Xiang, Jun Zhang, Sen Yang, Zhongyi Yang, Ming-Hui Wang, Jing Zhang, Wei Yang, Junzhou Huang, and Xiao Han. Scl-wc: Cross-slide contrastive learning for weakly-supervised whole-slide image classification. In *Advances in Neural Information Processing Systems*, 2022. 3
- [32] Chen Wei, Kihyuk Sohn, Clayton Mellina, Alan Yuille, and Fan Yang. Crest: A class-rebalancing self-training framework for imbalanced semi-supervised learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021. 1, 3, 7, 8
- [33] Qizhe Xie, Zihang Dai, Eduard Hovy, Thang Luong, and Quoc Le. Unsupervised data augmentation for consistency training. In *Advances in Neural Information Processing Systems*, 2020. 3
- [34] Yi Xu, Lei Shang, Jinxing Ye, Qi Qian, Yu-Feng Li, Baigui Sun, Hao Li, and Rong Jin. Dash: Semi-supervised learning with dynamic thresholding. In *International Conference on Machine Learning*, 2021. 4
- [35] Fan Yang, Kai Wu, Shuyi Zhang, Guannan Jiang, Yong Liu, Feng Zheng, Wei Zhang, Chengjie Wang, and Long Zeng. Class-aware contrastive semi-supervised learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 3
- [36] Jinsung Yoon, Yao Zhang, James Jordon, and Mihaela van der Schaar. Vime: Extending the success of self-and semi-supervised learning to tabular domain. In *Advances in Neural Information Processing Systems*, 2020. 8, 4, 5
- [37] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In *British Machine Vision Conference*, 2016. 2, 8
- [38] Bowen Zhang, Yidong Wang, Wenxin Hou, Hao Wu, Jindong Wang, Manabu Okumura, and Takahiro Shinozaki. Flexmatch: Boosting semi-supervised learning with curriculum pseudo labeling. In *Advances in Neural Information Processing Systems*, 2021. 1, 4
- [39] Zhen Zhao, Luping Zhou, Yue Duan, Lei Wang, Lei Qi, and Yinghuan Shi. Dc-ssl: Addressing mismatched class distribution in semi-supervised learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 1, 3
- [40] Zhen Zhao, Luping Zhou, Lei Wang, Yinghuan Shi, and Yang Gao. Lassl: Label-guided self-training for semi-supervised learning. In *AAAI Conference on Artificial Intelligence*, 2022. 3
- [41] Mingkai Zheng, Shan You, Lang Huang, Fei Wang, Chen Qian, and Chang Xu. Simmatch: Semi-supervised learning with similarity matching. In *IEEE/CVF Conference on Computer Vision and Pattern Recogn*, 2022. 7, 8