

CroSel: Cross Selection of Confident Pseudo Labels for Partial-Label Learning

Shiyu Tian¹ Hongxin Wei² Yiqun Wang¹ Lei Feng^{3*}

¹Chongqing University

²Southern University of Science and Technology

³Singapore University of Technology and Design

tiansy7@foxmail.com, weihx@sustech.edu.cn, yiqun.wang@cqu.edu.cn, lfengqqaq@gmail.com

Abstract

Partial-label learning (PLL) is an important weakly supervised learning problem, which allows each training example to have a candidate label set instead of a single ground-truth label. Identification-based methods have been widely explored to tackle label ambiguity issues in PLL, which regard the true label as a latent variable to be identified. However, identifying the true labels accurately and completely remains challenging, causing noise in pseudo labels during model training. In this paper, we propose a new method called CroSel, which leverages historical predictions from the model to identify true labels for most training examples. First, we introduce a cross selection strategy, which enables two deep models to select true labels of partially labeled data for each other. Besides, we propose a novel consistency regularization term called co-mix to avoid sample waste and tiny noise caused by false selection. In this way, CroSel can pick out the true labels of most examples with high precision. Extensive experiments demonstrate the superiority of CroSel, which consistently outperforms previous state-of-the-art methods on benchmark datasets. Additionally, our method achieves over 90% accuracy and quantity for selecting true labels on CIFAR-type datasets under various settings.

1. Introduction

The past few years have seen an increased interest in deep learning due to its outstanding performance in various application domains, including image processing [3], automatic driving [38], and medical diagnosis [23]. The success of deep learning heavily relies on a massive amount of fully labeled data. However, it is challenging to obtain a large-scale dataset with completely accurate annotations in the real world. To address this challenge, many

researchers have explored a promising weakly supervised learning problem called partial-label learning (PLL) [5, 7, 12, 26, 27, 29, 35, 36, 39, 41], where each training example to have a set of candidate labels that includes the true label. This problem arises in many real-world tasks such as automatic image annotation [2] and facial age estimation [22].

As PLL focuses on multi-class classification, there is only one ground-truth label for each training example, and other labels in the candidate label set are actually wrong (false positive) labels, which would have a negative impact on model training. Therefore, there exists the challenge of *label ambiguity* in PLL. To address this challenge, the current mainstream solution is to disambiguate the candidate labels so as to figure out the true label for each training instance [7, 14, 17, 19, 43]. However, most of the existing disambiguation methods normally leverage simple heuristics to iteratively update the labeling confidences or pseudo labels [6, 7, 19, 29, 36], which could not achieve convincing performance in identifying the true label during the training phase. Generally, if more true labels of training instances can be identified, we can train a better model. This serves as our primary motivation for identifying as many true labels as possible for training instances, ultimately resulting in the creation of a desired model.

In this paper, we propose a method called **CroSel** (**Cross Selection of Confident Pseudo Labels**), which leverages historical prediction from deep neural networks to accurately identify true labels for most training examples. Our selecting criteria are based on the assumption that if a model consistently predicts the same label for an input image with high confidence and low volatility, then that label has a high probability of being the true label for that example. Using the cross selection strategy, the true labels of the vast majority of training examples can be accurately identified, with only negligible noise. Moreover, in order to avoid sample waste and tiny noise resulting from the selection, we also propose a co-mix consistency regularization to gener-

*Corresponding author.

ate trainable targets for all examples. This regulation term serves as an essential complement to our method, which can further enhance the scope and accuracy of our selection of “true” labels. The algorithm details are shown in Section 3.

Our main contribution can be summarized as follows:

- We propose a cross selection strategy to select confident pseudo labels in the candidate label set based on historical prediction. Our approach demonstrates a notable combination of high precision and ratio in selecting “true” labels within candidate label sets.
- We introduce a new consistency regularization term that can leverage MixUp [46] to enhance the data and generate trainable targets for all examples, serving as an important supplement to our method.
- We experimentally show that CroSel achieves state-of-the-art performance on common benchmark datasets. We also provide extensive ablation studies to examine the effect of the different components of CroSel.

2. Related Work

Partial-label learning. This setting allows each training example to be annotated with a set of candidate labels for which the ground truth label is guaranteed to be included. However, *label ambiguity* can pose a significant challenge in PLL. Early methods use an averaging strategy, which tends to treat each candidate label equally. [5, 47] But such methods are easily affected by negative labels in the candidate label set, thus forming wrong classification boundaries. Afterwards, identification-based methods [14, 17, 43] have received more attention from the community, which regard the ground-truth label as a latent variable, and will maintain a confidence score for each candidate label. For instance, Yu *et al.* [43] introduced the maximum margin constraint to PLL, trying to optimize the margin between the model outputs from candidate labels and other negative labels.

Recently, partial-label learning has been combined with deep networks, leading to significant improvements in performance. Feng *et al.* [7] assumed that partial labels come from a uniform generation model and provided a mathematical formulation, which is adopted by most of the algorithms proposed later [11, 29, 35, 40]. Based on this, they also propose an algorithm for classification consistency and risk consistency. PRODEN [19] assumes the true label should be the one with the smallest loss among the candidate labels, and improves the classification risk algorithm accordingly. Wen *et al.* [35] proposed a risk-consistent leveraged weighted loss with label-specific candidate label sampling. PiCO [29] innovatively introduces contrastive learning [21] to the field and provides a solid theoretical analysis based on the EM (Expectation-Maximum) algorithm. POP [40] updates the model and purifies each candidate label set progressively in every epoch and eventually approximates the Bayes optimal classifier with mild assumptions.

Sample selection. Sample selection is a popular technique in deep learning, especially for learning with noisy labels [4, 18, 25, 32, 33, 37, 42, 48]. Then an obvious idea is to separate the clean samples and noisy samples in the mixed dataset. To address this issue, many existing works adopt the small loss criterion [10, 13, 31], which assumes that clean samples tend to have a smaller loss than noisy samples during training. MentorNet [13] is a representative work that makes the teacher model pick up clean samples for the student model. Co-teaching [10] constructs a double branches network to select clean samples for each branch, which is different from the teacher-student approach since none of the models supervise the other but rather help each other out. This idea was improved by some research later [30, 44]. Curriculum learning is also applied to this field [9], which considers clean labeled data as an easy task, while noisily labeled data as a harder task. Guo *et al.* [8] splitted data into subgroups according to their complexities, in order to optimize the training objectives in the early stage of course learning. OpenMatch [28] trains n OVA classifiers to select the in-distribution samples under the open set setting. Generative models such as Beta Mixture Model [1] and Gaussian Mixture Model [16] are also used to fit loss functions to distinguish clean labels from noisy labels. Recently, the fluctuation magnitude of the output [34, 45] and normalized entropy [11] are also considered as an important credential to judge whether the label is clean.

3. Our Proposed Method

In this section, we provide a detailed explanation of how our algorithm works. Our method is composed of two main components: a cross selection strategy that utilizes two models to select confident pseudo labels for each other, and a consistency regularization term that is applied across different data augmentation versions. The latter part not only addresses the issue of label waste resulting from the selection process but also enhances the quantity and accuracy of selections. The pseudo-code for our algorithm is presented in Algorithm 1.

3.1. Problem Setting

Suppose the feature space is $\mathcal{X} \in \mathbb{R}^d$ with d dimensions and the label space is $\mathcal{Y} = \{1, 2, \dots, k\}$ with k classes. We are given a dataset $\mathcal{D} = \{(\mathbf{x}_i, S_i)\}_{i=1}^n$ with n examples, where the instance $\mathbf{x}_i \in \mathcal{X}$ and corresponding candidate label set $S_i \subset \mathcal{Y}$. Same as previous studies, we assume that the true label $y_i \in \mathcal{Y}$ of each input \mathbf{x}_i is concealed in S_i .

Our aim is to train a multi-class classifier $f : \mathbb{R}^d \rightarrow \mathcal{Y}$ that minimizes the classification risk on the given dataset. For our classifier f , we use $f(\mathbf{x})$ to represent the output of classifier f on given input \mathbf{x} . And we use $\hat{y} = \operatorname{argmax}_{y \in \mathcal{Y}} f_y(\mathbf{x})$ to denote the prediction of our classifier, where $f_y(\mathbf{x})$ is the y -th coordinate of $f(\mathbf{x})$.

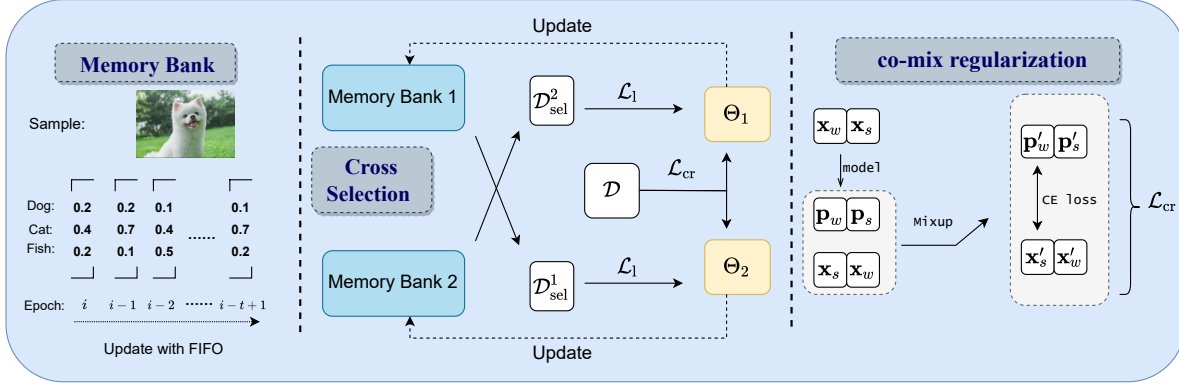


Figure 1. The left side of the figure is a brief example of our memory bank(MB) that stores the softmax output of the model for the last t epochs, which is updated by the FIFO (First In First Out) principle; the middle is the cross selection strategy: within each epoch, data subsets \mathcal{D}_{sel} with confident pseudo-labels are selected from the MB of each network, which produce loss function \mathcal{L}_1 to the training process for the other network; the right side illustrates our co-mix regularization term and the corresponding loss function \mathcal{L}_{cr} .

Algorithm 1 Pseudo-code of CroSel

Input: Training dataset $\mathcal{D} = \{(x_i, S_i)\}_{i=1}^n$, consistency regularization parameter λ_{cr} , sharpen parameter T , confidence threshold γ , memory bank $\text{MB}^{(1)}$, $\text{MB}^{(2)}$, network $\Theta^{(1)}$, $\Theta^{(2)}$, epoch E , iteration I .

Procedure:

$\text{MB}^{(1)}, \text{MB}^{(2)}, \Theta^{(1)}, \Theta^{(2)} = \text{WarmUp}(\text{MB}^{(1)}, \text{MB}^{(2)}, \Theta^{(1)}, \Theta^{(2)}, \mathcal{D})$. // CC algorithm

for $e = 1$ **to** E **do**

 Select labeled dataset $\mathcal{D}_{\text{sel}}^1$ through $\text{MB}^{(1)}$.

 Select labeled dataset $\mathcal{D}_{\text{sel}}^2$ through $\text{MB}^{(2)}$. // Eq. (4)

for $k = 1, 2$ **do**

for $i = 1$ **to** I **do**

 Fetch a labeled batch \hat{B}_i from the opposite selected dataset $\mathcal{D}_{\text{sel}}^k$.

 Fetch a batch B_i from the training dataset \mathcal{D} .

 Calculate the loss L among the two batches \hat{B}_i and B_i through Eq. (13).

 Update the weight of $\Theta^{(k)}$ by optimizer.

end for

 Update the memory bank $\text{MB}^{(k)}$ through the FIFO principle.

end for

end for

3.2. Selection Strategy

In the current partial-label learning task, the large number of candidate labels can confuse the classifier, making it difficult for the classifier to capture specific features belonging to a certain label. Therefore, our goal is to identify the most likely true label among the candidate labels and eliminate the interference of other negative labels during model training. By selecting these “true” labels, we can train the data with a supervised learning approach.

Warm up. Before selecting, we warm up the network using the entire training set. The goal of this stage is to reduce the classification risk of the input x to the whole candidate label set S , and obtain some historical information that can be used for selection. Therefore, here we use CC

algorithm [7] to warm up models for 10 epochs. At the same time, we will update the value of memory bank MB.

Selecting criteria. We have three criteria for selecting the confident pseudo labels from the candidate label set. Depending on the setup of our problem, the true label of each example must be in its candidate label set. This is our first criterion. In addition, we believe that an example’s predicted label is likely to be the true label if the model predicts it with high confidence and low fluctuation. To determine the latter, we maintain a memory bank MB to store historical prediction of this neural network.

The size of the memory bank is $t \times n \times k$, where t denotes the length of time it stores, n is the length of dataset, and k is the number of categories in the classification. In other words, MB stores the output after softmax of the model in

the last t epochs. MB is structured as a queue, with each element being a k -dimensional vector \mathbf{q} representing the output of an example in a particular epoch. We update MB with the FIFO principle. These selection criteria can be summarized as follows:

$$\beta_1 = \mathbb{I}(\operatorname{argmax}(\mathbf{q}^i) \in S), \quad (1)$$

$$\beta_2 = \mathbb{I}(\operatorname{argmax}(\mathbf{q}^i) = \operatorname{argmax}(\mathbf{q}^{i+1})), \quad (2)$$

$$\beta_3 = \mathbb{I}\left(\frac{1}{t} \sum_{i=1}^t \max(\mathbf{q}^i) > \gamma\right), \quad (3)$$

where γ is the confidence threshold of selection, $\mathbb{I}(\cdot)$ is the indicator function. β_1 lets our selected label be in the candidate label set. β_2 limits that the label we picked has not flipped in the past t epochs, which is a volatility consideration, and β_3 ensures that the label we selected has a high confidence level. The final selected dataset is \mathcal{D}_{sel} :

$$\mathcal{D}_{\text{sel}} = \{(\mathbf{x}_i, \operatorname{argmax}(\mathbf{q}_i^t)) | (\beta_1^i \wedge \beta_2^i \wedge \beta_3^i) = 1, \mathbf{x}_i \in \mathcal{D}\}, \quad (4)$$

Selected label loss. After selecting the high-confidence examples, we obtain a dynamically updated data subset $\mathcal{D}_{\text{sel}} = (X, \hat{Y})$. Each tuple of the subset has an instance \mathbf{x} and a selected label $\hat{y} = \operatorname{argmax}(\mathbf{q}^t)$. In this configuration, we can use the basic cross-entropy loss to deal with this part of the samples:

$$\mathcal{L}_1 = \frac{1}{|\mathcal{D}_{\text{sel}}|} \sum_{\mathbf{x} \in \mathcal{D}_{\text{sel}}} \mathcal{L}_{\text{CE}}(f(\mathbf{x}_w), \hat{y}), \quad (5)$$

where $\mathcal{L}_{\text{CE}}(\cdot, \cdot)$ denotes the softmax cross entropy loss, \mathbf{x}_w denotes the weak augmented version of \mathbf{x} .

Cross selection. However, In the process of selecting labels, it is difficult to guarantee that the selected labels are 100% accurate. In order to maximize the accuracy of selection, we propose a cross selection framework based on the idea of ensemble learning. Specifically, we train two identical models $\Theta^{(1)}$ and $\Theta^{(2)}$ through the same training and label selection process. By forming different decision boundaries, the two models can adaptively correct most of the errors even if there is noise in the selected confident pseudo labels.

To maximize the accuracy of the selected labels, we employ a cross-supervised training process. This involves using the selected dataset $\mathcal{D}_{\text{sel}}^1$ from MB⁽¹⁾ to train $\Theta^{(2)}$, and vice versa, using the selected dataset $\mathcal{D}_{\text{sel}}^2$ from MB⁽²⁾ to train $\Theta^{(1)}$. By doing so, the two models can learn from each other and further improve their ability to select true labels. In the test process, we will average the output of the two models to get the final prediction:

$$f'(\mathbf{x}) = \frac{1}{2}(f^1(\mathbf{x}) + f^2(\mathbf{x})), \quad (6)$$

where $f'(\mathbf{x})$ denotes the final output of our method in the test process, $f^1(\mathbf{x})$ denotes the output of $\Theta^{(1)}$, $f^2(\mathbf{x})$ denotes the output of $\Theta^{(2)}$.

3.3. Co-mix Consistency Regulation

Motivation. When dealing with complex partial-label learning tasks, our label selection strategy may not accurately select the true labels for all examples. If we only use the selected examples with their corresponding labels, it will result in a significant amount of wasted data, which contradicts our goal of utilizing as many examples as possible. Therefore, we aim to provide a trainable target for the remaining examples that are not selected. However, our setting differs from traditional semi-supervised learning as the proportion of unlabeled examples is relatively small. So it is unreasonable to directly transfer the existing semi-supervised learning tools to unlabeled data like other weakly supervised learning methods.

Motivated by this, we hope to propose a regularization term that can serve as an important supplement to our method and help us select examples. We proposed the co-mix regularization term, which employs two widely used data augmentation methods: weak augmentation and strong augmentation, to generate pseudo-labels as training targets for consistency regularization. We further employ MixUp [46] to further enhance the data. It is worth mentioning that the term ‘pseudo labels’ in this part specifically refers to the soft labels generated from different data augmentation versions, rather than the confident hard labels selected during the selection process.

Pseudo label generating. Consistency loss is a simple but effective idea in weakly supervised learning, whose key point is to reduce the gap between the output of two perturbed examples after passing through the model. As discussed in the previous section, we used two widely used data augmentations: ‘weak’ and ‘strong’, and crossed them to generate pseudo labels. Specifically, as stated in Figure 1, the pseudo label corresponding to weak augmented example is generated by strong augmented example, while the pseudo label corresponding to strong augmented example is generated by weak augmented example.

To generate these pseudo labels, we fix the parameters of the neural network, and pass the augmented images through the model to get the logits output. Then we perform two operations on the logits, sharpening and normalization. For sharpening operation, we use a hyper-parameter T , the more T goes to zero, the more logits tend to become a one-hot distribution. These two operations can be summarized by the following formula.

$$\mathbf{p}_i = \begin{cases} \frac{\exp(f_i(\mathbf{x})^{\frac{1}{T}})}{\sum_{i \in S} \exp(f_i(\mathbf{x})^{\frac{1}{T}})}, & i \in S, \\ 0, & i \notin S, \end{cases} \quad (7)$$

where \mathbf{p}_i denotes the i -th coordinate of pseudo label.

MixUp. After generating the pseudo labels of each example, we end up with two datasets that can be trained:

(X_w, P_s) and (X_s, P_w) . The subscripts w and s indicate the type of data augmentation used, i.e., weak or strong. Then, We spliced the two datasets together to further enhance the data with MixUp [46]. For a pair of two examples with their corresponding pseudo labels $(\mathbf{x}_1, \mathbf{p}_1), (\mathbf{x}_2, \mathbf{p}_2)$, we compute $(\mathbf{x}', \mathbf{p}')$ by the following formula:

$$\lambda \sim \text{Beta}(\alpha, \alpha), \quad (8)$$

$$\lambda' = \max(\lambda, 1 - \lambda), \quad (9)$$

$$\mathbf{x}' = \lambda' \mathbf{x}_1 + (1 - \lambda') \mathbf{x}_2, \quad (10)$$

$$\mathbf{p}' = \lambda' \mathbf{p}_1 + (1 - \lambda') \mathbf{p}_2, \quad (11)$$

After MixUp, we finally obtain $2n$ trainable example pair $(\mathbf{x}', \mathbf{p}')$. Then we can use the typical cross-entropy loss on every example, the consistency regulation loss will be:

$$\mathcal{L}_{\text{cr}} = \frac{1}{2n} \sum_{i=1}^{2n} \mathcal{L}_{\text{CE}}(f(\mathbf{x}'_i), \mathbf{p}'_i), \quad (12)$$

where $\mathcal{L}_{\text{CE}}(\cdot, \cdot)$ denotes the softmax cross entropy loss, n denotes the length of a dataset.

3.4. Algorithm overview

Overall loss. In the formal training phase, our loss function will be composed of two parts, the supervised loss \mathcal{L}_1 in the selected label set \mathcal{D}_{sel} and the consistency regularization item loss \mathcal{L}_{cr} . The two will be dynamically combined into the final loss function by a hyperparameter:

$$\mathcal{L}_{\text{all}} = \mathcal{L}_1 + \lambda_d * \mathcal{L}_{\text{cr}}, \quad (13)$$

where λ_d is a dynamically changing parameter, \mathcal{L}_1 and \mathcal{L}_{cr} can be calculated by Eq. (5) and Eq. (12), respectively. The use of MixUp can cause significant changes to the original feature space, making it necessary to adjust the weight of the regularization term in the loss function as the number of selected samples increases. To achieve this, we proposed a gradually decreasing λ_d with the increase of selected samples. We can control the magnitude of λ_d with a set hyperparameter λ_{cr} . The updated rules of λ_d are as follows:

$$\lambda_d = (1 - r_s) * \lambda_{\text{cr}}, \quad (14)$$

where r_s denotes the percentage of labeled data that we picked out, λ_{cr} is a hyperparameter that collaboratively adjusts the ratio of two loss items.

4. Experiments

4.1. Experimental Setup

Datasets. We used three widely used benchmarks in this field: SVHN [20], CIFAR-10 [15] and CIFAR-100 [15]. The way we generate partial labels is by flipping the negative labels $\bar{y} \neq y$ of the example with a set probability $q = P(\bar{y} \in S | \bar{y} \neq y)$. With the increase of q , the noise

of the dataset increases gradually. Following PiCO [29], we consider $q = \{0.01, 0.05, 0.1\}$ for CIFAR-100 and $q = \{0.1, 0.3, 0.5\}$ for other datasets.

Compared methods. We choose six well-performed partial-label learning algorithms to compare:

- PoP [40], an algorithm that updates the learning model and purifies each candidate label set progressively in every epoch.
- CRDPLL [36], an algorithm that takes non-candidate labels as supervision information and proposes a new consistency loss term between augmented images.
- PiCO [29], a theoretical solid framework that combines contrastive learning and prototype-based label disambiguation algorithm.
- LWS [35], an algorithm that wants to balance the risk error between the candidate label set and the non-candidate label set.
- PRODEN [19], a self-training algorithm that dynamically updates the confidence of candidate labels.
- CC [7], an algorithm that wants to minimize the classification error of the whole candidate label sets.

Implementations. Our implementation is based on PyTorch [24]. We use WRN-34-10 (short for Wide-ResNet-34-10) as the backbone model with a weight decay of 0.0001 on all the datasets for all compared methods. We present the mean and standard deviation in each case based on three independent runs with different random seeds. More detailed hyper-parameter setting can be found in Supplementary Materials.

4.2. Main Empirical Results

As shown in Table 1, our methods achieve state-of-the-art results on all the settings except CIFAR-10 with $q = 0.1$. Notably, on the complex dataset CIFAR-100, our method significantly improves performance.

However, a counter-intuitive phenomenon appears in our experiment, that is, the performance of our method does not necessarily decline strictly with the increase of noise magnitude q ; on the contrary, it may perform best in the case of moderate noise. We believe the possible reason for this phenomenon is that: On the one hand, when generating pseudo labels, we normalize them according to Eq. (7). That is to say, as the number of candidate labels increases, the ingredients involved in MixUp will also increase, leading to better-enhanced data interpolation. On the other hand, the increase of candidate labels also represents the increase of noise, which can have a negative impact on the algorithm's performance. Therefore, CroSel performs well in cases of intermediate noise, representing an optimal solution found in such a trade-off problem.

Table 2 presents the selection ratio and accuracy of CroSel in selecting the true labels for the subdataset \mathcal{D}_{sel} . It is evident that, CroSel can accurately select the true labels

Table 1. Accuracy (mean±std) comparisons on benchmark datasets.

Dataset	q	Ours	PoP	CRDPLL	PiCO	PRODEN	LWS	CC
CIFAR-10	0.1	97.31±.04%	97.17±.01%	97.41±.06%	96.10±.06%	95.66±.08%	91.20±.07%	90.73±.10%
	0.3	97.50±.05%	97.08±.01%	97.38±.04%	95.74±.10%	95.21±.07%	89.20±.09%	88.04±.06%
	0.5	97.34±.05%	96.66±.03%	96.76±.05%	95.32±.12%	94.55±.13%	80.23±.21%	81.01±.38%
SVHN	0.1	97.71±.05%	97.55±.06%	97.63±.06%	96.58±.04%	96.20±.07%	96.42±.09%	96.99±.17%
	0.3	97.96±.05%	97.50±.03%	97.65±.07%	96.32±.09%	96.11±.05%	96.15±.08%	96.67±.20%
	0.5	97.86±.06%	97.31±.01%	97.70±.05%	95.78±.05%	95.97±.03%	95.79±.05%	95.83±.23%
CIFAR-100	0.01	84.24±.09%	83.03±.04%	82.95±.10%	74.89±.11%	72.24±.12%	62.03±.21%	66.91±.24%
	0.05	83.92±.24%	82.79±.02%	82.38±.09%	73.26±.09%	70.03±.18%	57.10±.17%	64.51±.37%
	0.10	84.07±.16%	82.39±.04%	82.15±.20%	70.03±.10%	69.82±.11%	52.60±.54%	61.50±.36%

Table 2. Selection ratio and selection accuracy (mean±std) on benchmark datasets. S-ratio represents the selection ratio and S-acc represents selection accuracy in \mathcal{D}_{sel} .

Datasets	Setting	Index	Performance
CIFAR-10	$q = 0.1$	S-ratio	99.09±.07%
		S-acc	99.79±.05%
	$q = 0.3$	S-ratio	98.10±.10%
		S-acc	99.55±.03%
	$q = 0.5$	S-ratio	96.25±.12%
		S-acc	99.44±.06%
SVHN	$q = 0.1$	S-ratio	97.25±.14%
		S-acc	99.84±.06%
	$q = 0.3$	S-ratio	76.42±.21%
		S-acc	99.77±.06%
	$q = 0.5$	S-ratio	73.21±.15%
		S-acc	99.34±.02%
CIFAR-100	$q = 0.01$	S-ratio	96.58±.13%
		S-acc	99.71±.06%
	$q = 0.05$	S-ratio	95.45±.21%
		S-acc	98.29±.15%
	$q = 0.10$	S-ratio	93.61±.12%
		S-acc	97.93±.11%

of most examples in the dataset under various noise conditions. Notably, for CIFAR-10 and CIFAR-100, the selection ratio and accuracy are over 90% both. However, in SVHN, the selection ratio drops significantly with the increase of noise. This could be attributed to the fact that digital images in SVHN have relatively simple shape features, and MixUp may significantly disturb the feature space.

More challenging experimental settings. We conducted fine-grained settings on CIFAR-100 by limiting the candidate labels to the super-class of the true label, which is closer to reality. We set $q = 0.5$, indicating that half of the remaining superclass labels for each example may be converted into candidate labels. Our method still exhibits superior performance in this scenario. The results can be found in Table 3.

Table 3. Results on CIFAR-100 in fine-grained settings.

Method	Accuracy	Method	Accuracy
PoP	82.04%	CRDPLL	81.53%
PiCO	73.38%	PRODEN	71.16%
LWS	54.08%	CC	64.91%
Crosel (ours)		83.34%	

4.3. Ablation Studies

In this section, our primary objective is to demonstrate the collective effectiveness of all components within our method. Subsequently, we will delve into illustrating the consequences of either the absence or adjustment of certain components on the results.

All the components matter. Our algorithm comprises several components, including a selection strategy with three selection criteria and a regularization term. We conducted comprehensive ablation studies on CIFAR-100 with $q = 0.1$ to further prove that every component of our algorithm matters. The results are shown in Table 4. Notably, the absence of any component leads to a significant performance decline. The Co-mix regularization term serves as a training target for the examples that are not selected which can improve the selection accuracy, selection number, and final test accuracy when utilizing the same selection criteria.

Compared with using a combination of three criteria, employing individual selection criteria implies a relaxation of selection standards. When regularization terms are employed conventionally, this relaxation inevitably leads to decreased accuracy. While it enables the selection of more samples, it also introduces additional noise that can significantly misguide the model’s training process. In scenarios where no regularization terms are utilized, the imposition of more stringent selection criteria poses challenges in acquiring sufficient data for training. In such cases, relaxing the selection criteria serves to expand the model’s training data, thereby enhancing the overall performance of the model and subsequently improving its test performance.

The influence on the scope of the consistency regulariza-

Table 4. Results of thorough ablation experiments.

cr1	cr2	cr3	\mathcal{L}_{cr}	Acc	S-acc	S-ratio
✓				73.12%	95.15%	90.32%
✓	✓			72.00%	93.51%	85.09%
✓	✓	✓		70.68%	96.22%	78.65%
✓			✓	77.04%	91.02%	97.27%
✓	✓		✓	79.90%	94.66%	95.51%
✓	✓	✓	✓	84.07%	97.93%	93.61%

Table 5. Results for ablation studies on the scope of regularization.

Setting	Scope	Index	Performance
CIFAR-10 $q = 0.5$	All data	Acc	97.34%
		S-ratio	96.25%
		S-acc	99.44%
	Unselected data	Acc	90.32%
		S-ratio	93.27%
		S-acc	95.72%
None	Acc	81.01%	
	S-ratio	90.23%	
	S-acc	89.72%	
CIFAR-100 $q = 0.1$	All data	Acc	84.07%
		S-ratio	93.61%
		S-acc	97.93%
	Unselected data	Acc	77.61%
		S-ratio	90.12%
		S-acc	97.63%
None	Acc	70.68%	
	S-ratio	78.65%	
	S-acc	96.22%	

tion term. As described in Section 3, our co-mix regularization is designed to avoid sample waste. As such, a natural idea is to apply the regularization term to the unselected samples, as in traditional semi-supervised learning. However, our setting differs from traditional semi-supervised learning in that our unselected data only constitutes a small portion of all the data. So we conducted experiments by trying three cases: no regularization term, using the regularization term only for the unselected dataset, and using the regularization term for all examples.

The results in Table 5 suggest that with the expansion of the application scope of co-mix regularization term, the model’s performance steadily improves, and the number of selected labeled examples also gradually increases. This also shows that our co-mix regularization term and selection strategy can achieve a mutually beneficial effect.

Impact of other parameters on selection criteria. Except for the number of selection rules, the two parameters t and γ in Eq. (2) and Eq. (3) determine the strictness of our selection criteria. t represents the length of historical prediction stored in MB, while γ represents the select threshold for the average prediction confidence of the model for the example prediction in the past t epochs. A larger t and a higher γ represent a stricter selection criterion, resulting in

Table 6. Accuracy for ablation study on selection criteria.

Setting	t	Accuracy	γ	Accuracy
CIFAR-10 $q = 0.3$	$t = 2$	97.03%	$\gamma = 0.80$	96.24%
	$t = 3$	97.50%	$\gamma = 0.90$	97.50%
	$t = 4$	96.15%	$\gamma = 0.95$	97.38%
CIFAR-100 $q = 0.1$	$t = 2$	82.74%	$\gamma = 0.80$	80.20%
	$t = 3$	84.07%	$\gamma = 0.90$	84.07%
	$t = 4$	83.56%	$\gamma = 0.95$	83.56%

a smaller \mathcal{D}_{sel} size but higher precision, which can affect model training. During the early stages, it may be difficult to select enough examples under the condition of using stricter selection criteria. Therefore, in this experiment, we set the label flipping probability q to 0.3 on CIFAR-10. However, in the later stages, the impact of selection criteria on the final accuracy rate is not significant if the initial stage is passed smoothly. Our experiment shows that $t = 3$ and $\gamma = 0.9$ are suitable values that can be applied to most experimental environments. The experimental results are presented in Table 6 and Supplementary Materials.

Parameter test on λ_{cr} . In this experiment, we test the parameter λ_d , which weights the contribution of the consistency regularization term to the training loss. As mentioned in Eq. (14), the parameter λ_d is directly influenced by the hyperparameter λ_{cr} . Therefore, we test $\lambda_{cr} = \{1, 2, 4\}$ on CIFAR-10 ($q = 0.5$) and CIFAR-100 ($q = 0.1$). At the same time, we also try to fix the parameter λ_d , that is, its value is not related to the selection ratio r_s . This setting we denote by $\lambda_{cr}(\text{fix})$, we test $\lambda_{cr}(\text{fix}) = \{0.5, 1, 2\}$. The results are visualized in Figure 2, and detailed data can be found in Supplementary Material.

As mentioned in Table 2 above, CroSel achieves a very high selection ratio. When using a dynamically changing λ_d , the contribution made by the regularization term would be quite small in the later stages as the learning rate decays. In contrast, with a fixed value of λ_d , the regularization term would still have a significant contribution in the later stages. Although the accuracy rate on the test set does not show a significant difference between different parameters, the selection effect is critical. A too large contribution of the regularization item will slightly improve the selection ratio at the cost of a decrease in the selection accuracy, which goes against the original intention of our algorithm design. Therefore, we finally decided to use dynamically changing parameters. In other words, we want the algorithm to return to a supervised learning setting with minimal noise as much as possible at the end of the training process.

The influence on double model. As recognized by the community, dual models tend to achieve better performance than single models. We are curious about how effective our selection criteria would be without the adaptive error correction capability of cross selection. Figure 3 visualizes the gap in the selection ratio between dual-model and single-

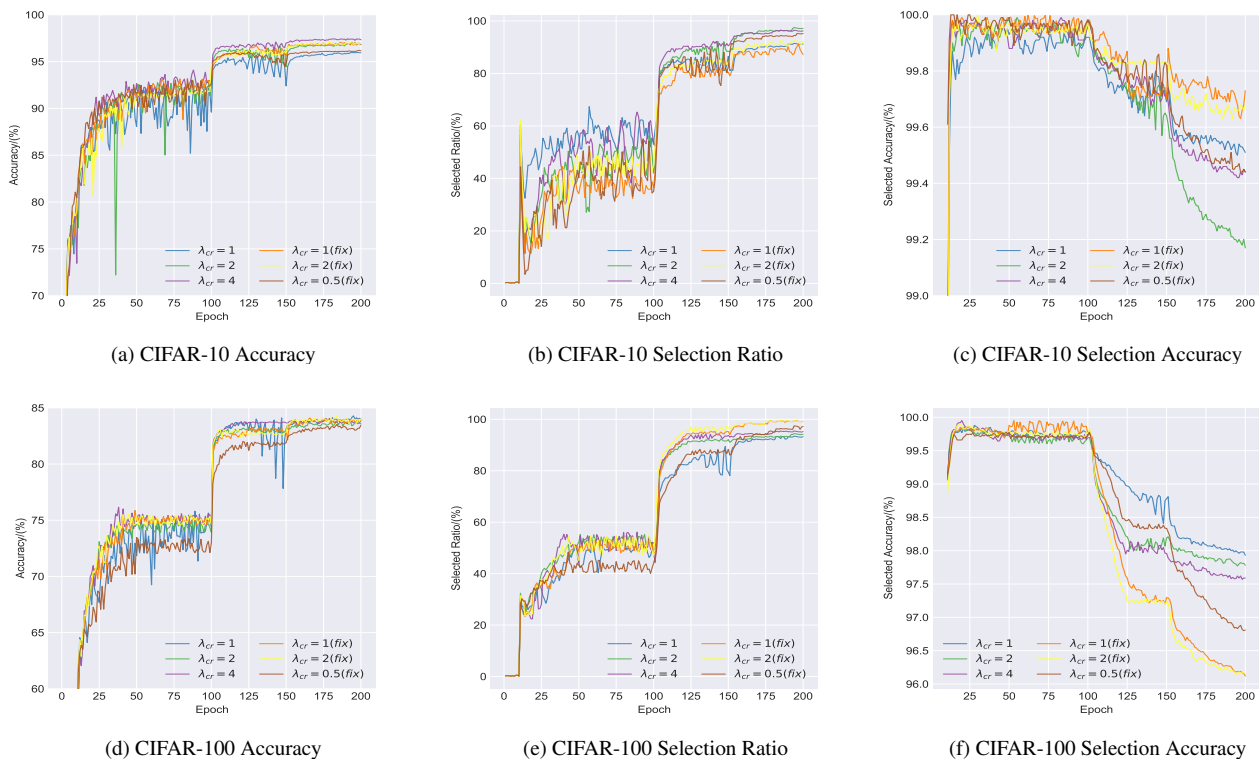


Figure 2. Parameter analysis of λ_{cr} on CIFAR-10 and CIFAR-100.

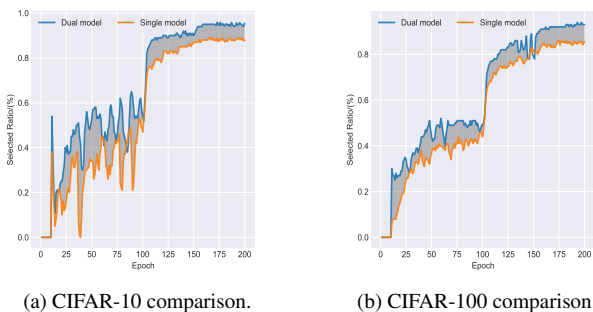


Figure 3. Selection ratio comparison between dual model and single model on CIFAR-10 and CIFAR-100.

model training. It shows that our cross selection strategy can select samples more comprehensively, and on average, about 10% more training examples can be selected on each dataset. However, the effectiveness of our algorithm is not solely due to the dual model. Even when using a single model with our selection criteria, the accuracy rate on the test set only decreases by 0.83% and 2.68% on CIFAR-10 and CIFAR-100 respectively. Furthermore, the high precision of selection is reflected in both settings. Detailed results can be found in Supplementary Material.

5. Conclusion

This work introduces CroSel, a novel partial-label learning method that leverages historical prediction to select confident pseudo labels from candidate label sets. The proposed method consists of two parts: a cross selection strategy that enables two models to select “true” labels for each other, and a consistency regularization term co-mix that avoids sample waste and tiny noise caused by false selection. Empirically, extensive experiments demonstrate the superiority of CroSel, which consistently outperforms previous state-of-the-art methods on multiple benchmark datasets.

Acknowledgements

Lei Feng is supported by National Natural Science Foundation of China (Grant No. 62106028) and Chongqing Overseas Chinese Entrepreneurship and Innovation Support Program. Hongxin Wei is supported by Shenzhen Fundamental Research Program JCYJ20230807091809020 and the Center for Computational Science and Engineering at the Southern University of Science and Technology. Yiqun Wang is supported by the Open Projects Program of State Key Laboratory of Multimodal Artificial Intelligence Systems and the NSFC No. 62202076.

References

- [1] Eric Arazo, Diego Ortego, Paul Albert, Noel O’Connor, and Kevin McGuinness. Unsupervised label noise modeling and loss correction. In *ICML*, pages 312–321, 2019. 2
- [2] Ching-Hui Chen, Vishal M Patel, and Rama Chellappa. Learning from ambiguously labeled face images. In *T-PAMI*, pages 1653–1667, 2017. 1
- [3] Hanting Chen, Yunhe Wang, Tianyu Guo, Chang Xu, Yiping Deng, Zhenhua Liu, Siwei Ma, Chunjing Xu, Chao Xu, and Wen Gao. Pre-trained image processing transformer. In *CVPR*, pages 12299–12310, 2021. 1
- [4] Hao Cheng, Zhaowei Zhu, Xing Sun, and Yang Liu. Mitigating memorization of noisy labels via regularization between representations. In *ICLR*, 2023. 2
- [5] Timothee Cour, Ben Sapp, and Ben Taskar. Learning from partial labels. In *JMLR*, pages 1501–1536. *JMLR.org*, 2011. 1, 2
- [6] Lei Feng and Bo An. Leveraging latent label distributions for partial label learning. In *IJCAI*, pages 2107–2113, 2018. 1
- [7] Lei Feng, Jiaqi Lv, Bo Han, Miao Xu, Gang Niu, Xin Geng, Bo An, and Masashi Sugiyama. Provably consistent partial-label learning. In *NeurIPS*, pages 10948–10960, 2020. 1, 2, 3, 5
- [8] Sheng Guo, Weilin Huang, Haozhi Zhang, Chenfan Zhuang, Dengke Dong, Matthew R Scott, and Dinglong Huang. Curriculumnet: Weakly supervised learning from large-scale web images. In *ECCV*, pages 135–150, 2018. 2
- [9] Bo Han, Ivor W Tsang, Ling Chen, P Yu Celina, and Sai-Fu Fung. Progressive stochastic learning for noisy labels. In *NeurIPS*, pages 5136–5148, 2018. 2
- [10] Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor Tsang, and Masashi Sugiyama. Co-teaching: Robust training of deep neural networks with extremely noisy labels. In *NeurIPS*, 2018. 2
- [11] Shuo He, Guowu Yang, and Lei Feng. Candidate-aware selective disambiguation based on normalized entropy for instance-dependent partial-label learning. In *ICCV*, pages 1792–1801, 2023. 2
- [12] Feng Hong, Jiangchao Yao, Zhihan Zhou, Ya Zhang, and Yanfeng Wang. Long-tailed partial label learning via dynamic rebalancing. In *ICLR*, 2023. 1
- [13] Lu Jiang, Zhengyuan Zhou, Thomas Leung, Li-Jia Li, and Li Fei-Fei. Mentornet: Learning data-driven curriculum for very deep neural networks on corrupted labels. In *ICML*, pages 2304–2313, 2018. 2
- [14] Rong Jin and Zoubin Ghahramani. Learning with multiple labels. In *NeurIPS*, 2002. 1, 2
- [15] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. In *Technique Report*, 2009. 5
- [16] Junnan Li, Richard Socher, and Steven CH Hoi. Dividemix: Learning with noisy labels as semi-supervised learning. In *arXiv preprint arXiv:2002.07394*, 2020. 2
- [17] Liping Liu and Thomas Dietterich. A conditional multinomial mixture model for superset label learning. In *NeurIPS*, 2012. 1, 2
- [18] Tongliang Liu and Dacheng Tao. Classification with noisy labels by importance reweighting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(3):447–461, 2015. 2
- [19] Jiaqi Lv, Miao Xu, Lei Feng, Gang Niu, Xin Geng, and Masashi Sugiyama. Progressive identification of true labels for partial-label learning. In *ICML*, pages 6500–6510, 2020. 1, 2, 5
- [20] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bisacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. In *Workshop on Deep Learning and Unsupervised Feature Learning*, 2011. 5
- [21] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. In *arXiv preprint arXiv:1807.03748*, 2018. 2
- [22] Gabriel Panis and Andreas Lanitis. An overview of research activities in facial age estimation using the fg-net aging database. In *ECCV*, 2015. 1
- [23] Seong Ho Park and Kyunghwa Han. Methodologic guide for evaluating clinical performance and effect of artificial intelligence technology for medical diagnosis and prediction. In *Radiology*, pages 800–809, 2018. 1
- [24] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *NeurIPS*, 2019. 5
- [25] Giorgio Patrini, Alessandro Rozza, Aditya Krishna Menon, Richard Nock, and Lizhen Qu. Making deep neural networks robust to label noise: A loss correction approach. In *CVPR*, pages 1944–1952, 2017. 2
- [26] Congyu Qiao, Ning Xu, and Xin Geng. Decompositional generation process for instance-dependent partial label learning. In *ICML*, 2023. 1
- [27] Congyu Qiao, Ning Xu, Jiaqi Lv, Yi Ren, and Xin Geng. Fre-dis: A fusion framework of refinement and disambiguation for unreliable partial label learning. In *ICML*, pages 28321–28336, 2023. 1
- [28] Kuniaki Saito, Donghyun Kim, and Kate Saenko. Open-match: Open-set semi-supervised learning with open-set consistency regularization. In *NeurIPS*, pages 25956–25967, 2021. 2
- [29] Haobo Wang, Ruixuan Xiao, Yixuan Li, Lei Feng, Gang Niu, Gang Chen, and Junbo Zhao. Pico: Contrastive label disambiguation for partial label learning. In *ICLR*, 2022. 1, 2, 5
- [30] Xiaobo Wang, Shuo Wang, Jun Wang, Hailin Shi, and Tao Mei. Co-mining: Deep face recognition with noisy labels. In *CVPR*, pages 9358–9367, 2019. 2
- [31] Hongxin Wei, Lei Feng, Xiangyu Chen, and Bo An. Combating noisy labels by agreement: A joint training method with co-regularization. In *CVPR*, pages 13726–13735, 2020. 2
- [32] Jiaheng Wei, Hangyu Liu, Tongliang Liu, Gang Niu, Masashi Sugiyama, and Yang Liu. To smooth or not? when label smoothing meets noisy labels. In *ICML*, 2022. 2

- [33] Jiaheng Wei, Zhaowei Zhu, Hao Cheng, Tongliang Liu, Gang Niu, and Yang Liu. Learning with noisy labels revisited: A study using real-world human annotations. In *ICLR*, 2022. [2](#)
- [34] Qi Wei, Haoliang Sun, Xiankai Lu, and Yilong Yin. Self-filtering: A noise-aware sample selection for label noise with confidence penalization. In *ECCV*, pages 516–532. Springer, 2022. [2](#)
- [35] Hongwei Wen, Jingyi Cui, Hanyuan Hang, Jiabin Liu, Yisen Wang, and Zhouchen Lin. Leveraged weighted loss for partial label learning. In *ICML*, pages 11091–11100, 2021. [1](#), [2](#), [5](#)
- [36] Dong-Dong Wu, Deng-Bao Wang, and Min-Ling Zhang. Revisiting consistency regularization for deep partial label learning. In *ICML*, pages 24212–24225, 2022. [1](#), [5](#)
- [37] Xiaobo Xia, Tongliang Liu, Nannan Wang, Bo Han, Chen Gong, Gang Niu, and Masashi Sugiyama. Are anchor points really indispensable in label-noise learning? In *NeurIPS*, 2019. [2](#)
- [38] Zuobin Xiong, Wei Li, Qilong Han, and Zhipeng Cai. Privacy-preserving auto-driving: a gan-based approach to protect vehicular camera data. In *ICDM*, pages 668–677, 2019. [1](#)
- [39] Ning Xu, Congyu Qiao, Xin Geng, and Min-Ling Zhang. Instance-dependent partial label learning. In *NeurIPS*, pages 27119–27130, 2021. [1](#)
- [40] Ning Xu, Biao Liu, Jiaqi Lv, Congyu Qiao, and Xin Geng. Progressive purification for instance-dependent partial label learning. In *ICML*, pages 38551–38565. PMLR, 2023. [2](#), [5](#)
- [41] Ning Xu, Biao Liu, Jiaqi Lv, Congyu Qiao, and Xin Geng. Progressive purification for instance-dependent partial label learning. In *ICML*, pages 38551–38565, 2023. [1](#)
- [42] Yu Yao, Tongliang Liu, Bo Han, Mingming Gong, Jiankang Deng, Gang Niu, and Masashi Sugiyama. Dual t: Reducing estimation error for transition matrix in label-noise learning. In *NeurIPS*, pages 7260–7271, 2020. [2](#)
- [43] Fei Yu and Min-Ling Zhang. Maximum margin partial label learning. In *ACML*, pages 96–111, 2016. [1](#), [2](#)
- [44] Xingrui Yu, Bo Han, Jiangchao Yao, Gang Niu, Ivor Tsang, and Masashi Sugiyama. How does disagreement help generalization against label corruption? In *ICML*, pages 7164–7173, 2019. [2](#)
- [45] Suqin Yuan, Lei Feng, and Tongliang Liu. Late stopping: Avoiding confidently learning from mislabeled examples. In *ICCV*, pages 16079–16088, 2023. [2](#)
- [46] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *arXiv preprint arXiv:1710.09412*, 2017. [2](#), [4](#), [5](#)
- [47] Min-Ling Zhang, Fei Yu, and Cai-Zhi Tang. Disambiguation-free partial label learning. *IEEE Transactions on Knowledge and Data Engineering*, pages 2155–2167, 2017. [2](#)
- [48] Zhilu Zhang and Mert Sabuncu. Generalized cross entropy loss for training deep neural networks with noisy labels. In *NeurIPS*, 2018. [2](#)