



Self-paced resistance learning against overfitting on noisy labels

Xiaoshuang Shi^a, Zhenhua Guo^b, Kang Li^c, Yun Liang^d, Xiaofeng Zhu^{a,*}

^a School of Computer Science and Technology, University of Electronic Science and Technology of China, Chengdu, Sichuan, China

^b Smart Transportation, Suzhou, China

^c West China Biomedical Big Data Center, West China Hospital, Sichuan University, Chengdu, Sichuan, China; Med-X Center for Informatics, Sichuan University, Chengdu, Sichuan, China; Shanghai Artificial Intelligence Laboratory, Shanghai, China

^d J. Crayton Pruitt Family Department of Biomedical Engineering, University of Florida, Gainesville, FL, USA



ARTICLE INFO

Article history:

Received 14 May 2022

Revised 21 July 2022

Accepted 20 September 2022

Available online 29 September 2022

Keywords:

Convolutional neural networks

Self-paced resistance

Model overfitting

Noisy labels

ABSTRACT

Noisy labels composed of correct and corrupted ones are pervasive in practice. They might significantly deteriorate the performance of convolutional neural networks (CNNs), because CNNs are easily overfitted on corrupted labels. To address this issue, inspired by an observation, deep neural networks might first memorize the probably correct-label data and then corrupt-label samples, we propose a novel yet simple self-paced resistance framework to resist corrupted labels, without using any clean validation data. The proposed framework first utilizes the memorization effect of CNNs to learn a curriculum, which contains confident samples and provides meaningful supervision for other training samples. Then it adopts selected confident samples and a proposed resistance loss to update model parameters; the resistance loss tends to smooth model parameters' update or attain equivalent prediction over each class, thereby resisting model overfitting on corrupted labels. Finally, we unify these two modules into a single loss function and optimize it in an alternative learning. Extensive experiments demonstrate the significantly superior performance of the proposed framework over recent state-of-the-art methods on noisy-label data. Source codes of the proposed method are available on <https://github.com/xssh2015/Self-paced-Resistance-Learning>.

© 2022 Elsevier Ltd. All rights reserved.

1. Introduction

Recently, convolutional neural networks (CNNs) have achieved tremendous success on various different tasks, such as image classification [1–4], retrieval [5,6], detection [7] and segmentation [8]. However, most CNNs usually require large-scale high-quality labels to obtain desired accuracy, because deep CNNs are capable of memorizing the entire training data even with completely random labels, i.e., deep neural networks (DNNs) will eventually memorize the wrongly given labels [9]. This infers that noisy labels might significantly deteriorate the performance of CNNs during training. Unfortunately, noisy labels are pervasive in practice and it is expensive to obtain accurate labeled data [10].

To tackle noisy labels for effectively and robustly training CNNs, some methods [11,12] utilize regularization terms for label correction to alleviate the deterioration of deep networks during training, but they often fail to attain the optimal accuracy. Another popular way is to estimate a label transition matrix without using regularizations for loss correction [13]. However,

it is usually difficult to accurately estimate the label transition matrix, especially for a large number of classes. The third promising direction is to select confident samples based on small-loss distances in order to update networks robustly, without estimating the label transition matrix. MentorNet [14] and Co-teaching [15] are two representative methods. When no clean validation data is available, self-paced MentorNet learns a neural network to approximate a predefined curriculum to provide meaningful supervision for StudentNet, so that it can focus on the samples with probably correct labels. Self-paced MentorNet is similar to the self-training method [16], and it inherits the same inferiority of accumulated errors generated by sample-selection bias. To address the issue, Co-teaching utilizes the memorization effect of DNNs and symmetrically trains two networks, each of which filters corrupted labels and selects the samples with small-loss to update the peer network, where the memorization effect of DNNs denotes that DNNs first memorize easy samples and then hard samples during training [17], and when noisy labels exist, DNNs eventually memorize the corrupted labels [9]. This phenomenon exists on different training optimizations, network architectures or datasets (please see Figure A10 in the supplemental material). However, with the increasing number of training epochs, the two networks will gradually form consensus predictions and Co-teaching will

* Corresponding author.

E-mail address: seanzhuxf@gmail.com (X. Zhu).

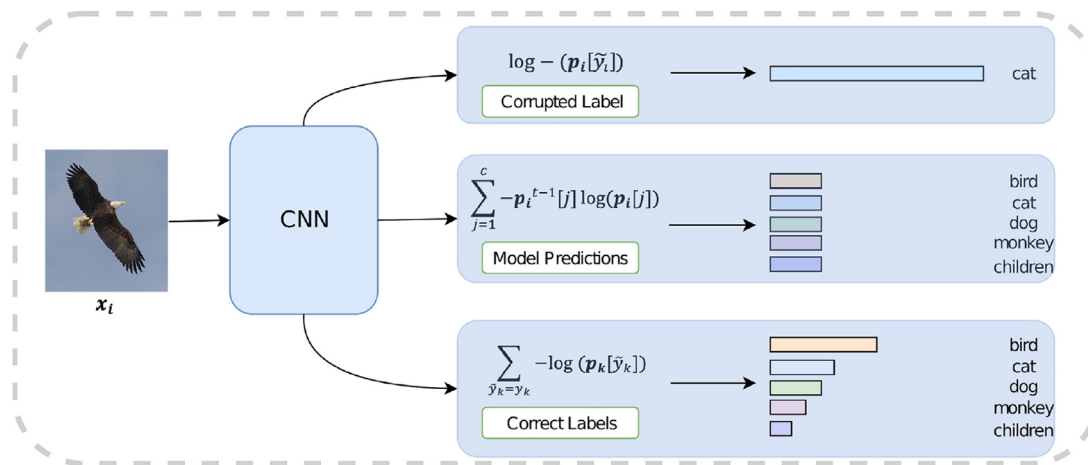


Fig. 1. The difference between the resistance loss and the traditional cross-entropy loss. The middle is the resistance loss tending to make the prediction over each class be equivalent. The top and bottom cross-entropy losses utilize a corrupted or correct label to update model parameters, respectively. (p_i^{t-1} and p_i are the prediction of the i th training sample x_i at the $t-1$ th and t th epochs, respectively, t is the current number of training epochs, c is the number of classes, y_k denotes the correct label of the k th ($1 \leq k \leq n$) sample and n is the total number of training samples.).

functionally deteriorate to self-paced MentorNet. Although the strategy of “Update by Disagreement” [18] can slow down the two networks of Co-teaching to form consensus predictions, it still cannot prevent the effect of sample-selection bias in many cases [19]. Additionally, when training data is with extremely noisy labels, MentorNet and Co-teaching easily select the corrupt-label data as confident samples, so that the networks are overfitted on corrupted labels, thereby decreasing their performance (please see Figure A11 in the supplemental material). Moreover, Co-teaching aims to filter corrupt-label training samples and thus might fail to explore their correct semantic information.

To address the performance deterioration of CNNs generated by model overfitting on corrupted labels, and meanwhile explore the correct semantic information of training samples with corrupted labels, in this paper, we propose a novel self-paced resistance framework using a resistance loss to robustly train CNNs on noisy labels, without using any clean validation data. The proposed framework is mainly inspired by: (i) DNNs might first memorize the probably correct-label data and then samples with corrupted labels or outliers [9,15,17] (see Figure A10); (ii) A curriculum consisting of confident samples can provide meaningful supervision for other training data [14]; (iii) A resisting model overfitted on corrupted labels can reduce the deterioration of model performance. We summarize three major contributions as follows:

- We propose a novel resistance loss to significantly alleviate model overfitting on corrupted labels, by smoothing model parameters’ update or attaining equivalent prediction on each class. For clarity, we present the difference between the resistance loss and the traditional cross-entropy loss in Fig. 1.
- We propose a novel yet simple framework, self-paced resistance learning (SPRL), by effectively using the memorization effect of DNNs, curriculum learning (CL) and a resistance loss to robustly train CNNs on noisy labels.
- Extensive experiments on four image datasets demonstrate that (i) The proposed framework can prevent the accuracy deterioration of CNNs on noisy labels, leading to superior classification accuracy over recent state-of-the-art methods on multiple types of label noise; (ii) With clean training data only, the proposed method usually obtains better results than standard networks.

The rest of the paper is organized as follows. Section 2 briefly reviews some popular methods to tackle noisy labels; Section 3 introduces the preliminaries on curriculum learning; Section 4 presents the proposed framework, SPRL; Section 5 shows

and analyzes experimental results of various methods; Finally, Section 6 concludes this paper and points out the future work.

2. Related work

Here, we briefly review some popular statistical learning methods for tackling noisy labels and DNNs with noisy labels.

Statistical learning methods. There are numerous statistical learning algorithms to handle noisy labels [20]. They can be roughly categorized into three groups: probabilistic modeling, surrogate losses and noise rate estimation. One popular probabilistic modeling method is [21], which proposes a two-coin model to handle noisy labels provided by multiple annotators. For surrogate losses based methods, [22] proposes an unbiased estimator to provide the noise corrected loss and then presents a weighted loss function for handling class-dependent noisy labels; [23] introduces a robust non-convex loss for tackling the contamination of data with outliers and a boosting algorithm, SavageBoost, to minimize the loss; [24] presents a convex loss modified from the hinged loss and proves its robustness to symmetric label noise. In the noise rate estimation category, [25] designs consistent estimators for classification with asymmetric (class-dependent) label noise; [26] utilizes kernel embeddings onto reproducing kernel Hilbert space for mixture proportion estimation; [27] estimates class proportions when the distributions of training and test samples are different; [27] and [28] introduce class-probability estimators using order statistics on the range of scores. Most of these statistical learning methods are proposed for traditional algorithms on relatively small datasets. Thus they usually fail to obtain promising performance on real applications, especially large datasets.

Deep neural networks with noisy labels. Because DNNs are sensitive to noisy labels, a few methods have been proposed to handle noisy labels for robust network training [29]. Mnih and Hinton [30] proposes two robust loss functions for binary classification of aerial image patches to handle omission and wrong location of training labels. Ghosh et al. [31] Ghosh et al. [32] Manwani and Sastry [33] investigate noise-tolerant of loss functions under risk minimization. Reed et al. [11] Laine and Aila [12] Shi et al. [34] consider the prediction consistency via adding a regularization term for robustly training DNNs. This strategy cannot prevent the performance deterioration of CNNs in many cases and it usually fails to obtain optimal accuracy. Sukhbaatar et al. [35] and Patrini et al. [13] estimate a label transition matrix, which summarizes the probability of one class being flipped into another, to

correct loss functions, and [36] employs a dimensionality-driven learning strategy to estimate the correct labels of samples during training and adapt the loss function. However, it is difficult to accurately estimate the label transition matrix or the labels of training samples. Wang et al. [37] proposes an iterative learning framework to handle open-set noisy labels. Han et al. [38] designs a self-learning framework to train the network in an end-to-end manner with iteratively correcting labels. Sun et al. [39] proposes a warped probabilistic inference for robust learning with noisy labels under the meta-learning scenario. Xiao et al. [40], Li et al. [41], Veit et al. [42] and Vahdat [43] adopt a small clean dataset to leverage samples with noisy labels; [44] adopts a small clean dataset to assign weights for training samples based on their gradient directions to reduce the effect of corrupted labels. These methods usually require an additional clean dataset to alleviate the overfitting of CNNs on noisy labels. Northcutt et al. [45] and Zhang and Sabuncu [46] adopt the confident samples for training by cleaning up corrupted labels, and thus they fail to exploit the semantic information of the samples with corrupted labels. [18] introduces a strategy, "Update by Disagreement", that updates the parameters of two networks by using the samples with different predictions. This strategy cannot handle noisy labels explicitly, because the disagreement predictions usually contain corrupted labels. Han et al. [47] integrates the learning regime of CL with the update process of vanilla SGD for robust model training on noisy labels. MentorNet [14] and Co-teaching [15] are two popular learn-to-teach methods to handle noisy labels. They select confident samples based on small-loss distances to teach the student or other network. [19] extends Co-teaching to alleviate the performance deterioration of DNNs. However, these learn-to-teach methods easily select corrupt-label samples as confident ones and then make CNNs be overfitted on corrupted labels in many cases, especially on extremely noisy labels (please refer to Figure A11 in the supplemental material), thereby deteriorating and decreasing the accuracy of CNNs during training. Yao et al. [48] formulates the sample selection from noisy labels as a function approximation problem, and proposes a novel Newton algorithm to solve the problem. However, its selection performance is still far from satisfying on extremely noisy labels.

Similar to previous learn-to-teach methods, the proposed method utilizes the memorization effect of DNNs to select confident samples as a curriculum to provide supervision of other training samples. However, unlike previous learn-to-teach methods that are very likely to deteriorate with the increasing number of training epochs, the proposed method can prevent the performance degradation during training. This is because the proposed resistance loss can significantly reduce the effect of corrupted labels by alleviating model overfitting. Additionally, the proposed framework does not require the noise rate and only trains a single network, differing from Co-teaching [15] and its variant [19] that need to know or estimate the rate of label noise and train two networks in a symmetric way. Overall, the proposed method is easy to utilize and can obtain good performance for image classification.

3. Preliminaries on curriculum learning

CL [49] is a training strategy inspired by the learning process of humans and animals that gradually proceeds easy to difficult samples. CL predetermines the curriculum based on the prior knowledge so that training data is ranked in a meaningful order to facilitate learning. In the following, we briefly introduce three major variants of CL that are related to our proposed method.

Self-paced learning (SPL) [50]: CL heavily relies on the prior knowledge and ignores the feedback of the learner (model); to address this issue, SPL dynamically determines the curriculum based on the learner's abilities. Given training data $\mathbf{X} = \{\mathbf{x}_i\}_{i=1}^n$ and the

corresponding labels $\mathbf{y} = \{y_i\}_{i=1}^n$, where \mathbf{x}_i and y_i denote the i th sample and its correct label, respectively. Let $f(\cdot)$ represent a classifier and \mathbf{w} be its model parameters. SPL simultaneously selects easy samples and learns model parameters in each iteration by solving the following problem:

$$\min_{\mathbf{w}, \mathbf{v}} E(\mathbf{w}, \mathbf{v}; \lambda) = \sum_{i=1}^n v_i L(y_i, f(\mathbf{x}_i, \mathbf{w})) - \lambda \sum_{i=1}^n v_i, \quad \text{s.t. } \mathbf{v} \in \{0, 1\}^n, \quad (1)$$

where $L(y_i, f(\mathbf{x}_i, \mathbf{w}))$ denotes the loss function that calculates the cost between the ground truth label y_i and the estimated label $f(\mathbf{x}_i, \mathbf{w})$, \mathbf{v} is a binary vector to indicate which ones are easy samples, and λ is a parameter to control the learning pace. Eq. (1) is usually solved by an alternative minimization strategy: with fixing \mathbf{w} , calculating \mathbf{v} by $\mathbf{v} = \begin{cases} 1 & L(\mathbf{x}_i, f(\mathbf{x}_i, \mathbf{w})) < \lambda, \\ 0 & \text{otherwise.} \end{cases}$, and then with fixing \mathbf{v} , updating \mathbf{w} by using selected easy samples to train the classifier $f(\cdot)$.

Self-paced curriculum learning (SPCL) [51]: Although SPL can dynamically learn the curriculum, it does not take into account the prior knowledge. Let Ψ be a feasible region encoding the information of a predetermined curriculum. To connect CL with SPL, SPCL [51] employs both the predetermined curriculum obtained by the prior knowledge before training and the learned curriculum during training with the following model:

$$\min_{\mathbf{w}, \mathbf{v}} E(\mathbf{w}, \mathbf{v}; \lambda) = \sum_{i=1}^n v_i L(y_i, f(\mathbf{x}_i, \mathbf{w})) + G(\mathbf{v}, \lambda), \quad \text{s.t. } \mathbf{v} \in [0, 1]^n, \mathbf{v} \in \Psi, \quad (2)$$

where \mathbf{v} is a weight vector to reflect the significance of samples, and $G(\cdot)$ is a self-paced function to control the learning scheme. For example, in SPL, $G(\mathbf{v}, \lambda) = -\lambda \sum_{i=1}^n v_i$. Similar to Eqs. (1), (2) can also be solved by using an alternative minimization method.

Self-paced MentorNet [14]: Because the learning procedure of DNNs is very complicated, it is difficult to be accurately modeled by the predefined curriculum. To tackle this issue, [14] employs two neural networks, one network called MentorNet $f_m(\cdot)$ and the other called StudentNet $f_s(\cdot)$. MentorNet is to approximate a predefined curriculum in order to compute time-varying weights $f_m(\mathbf{z}_i; \Theta^*) \in [0, 1]$ for each training sample, where Θ^* denotes the optimal parameters in $f_m(\cdot)$, $\mathbf{z}_i = \phi(\mathbf{x}_i, \tilde{y}_i, \mathbf{w})$ represents the input feature to MentorNet of the i th sample \mathbf{x}_i , \tilde{y}_i is the noisy label of \mathbf{x}_i and \mathbf{w} is the parameter of StudentNet $f_s(\cdot)$, which will utilize the learned weights $f_m(\mathbf{z}_i; \Theta^*)$ to update \mathbf{w} . To learn a Θ^* , MentorNet minimizes the following function:

$$\operatorname{argmin}_{\Theta} \sum_{i=1}^n f_m(\mathbf{z}_i; \Theta) l_i + G(f_m(\mathbf{z}_i; \Theta); \lambda), \quad (3)$$

where l_i is the loss between one hot vector $\tilde{\mathbf{y}}_i$ of the noisy label \tilde{y}_i and a predicting class probability vector $f_s(\mathbf{x}_i, \mathbf{w})$, which is a discriminative function of StudentNet. Similar to SPCL, $G(f_m(\mathbf{z}_i; \Theta); \lambda)$ is a self-paced function.

4. Self-paced resistance learning (SPRL)

Although MentorNet can boost model robustness when no clean validation data is used, it easily selects corrupt-label samples as confident ones and then overfits a model on them. To address this problem, we propose a novel training strategy, SPRL. It employs the memorization effect of DNNs to approximate a predefined curriculum in order to provide meaningful supervision for other training samples, and adopts a resistance loss to resist the effect of corrupted labels on the network. For clarity, we present the proposed SPRL framework in Fig. 2.

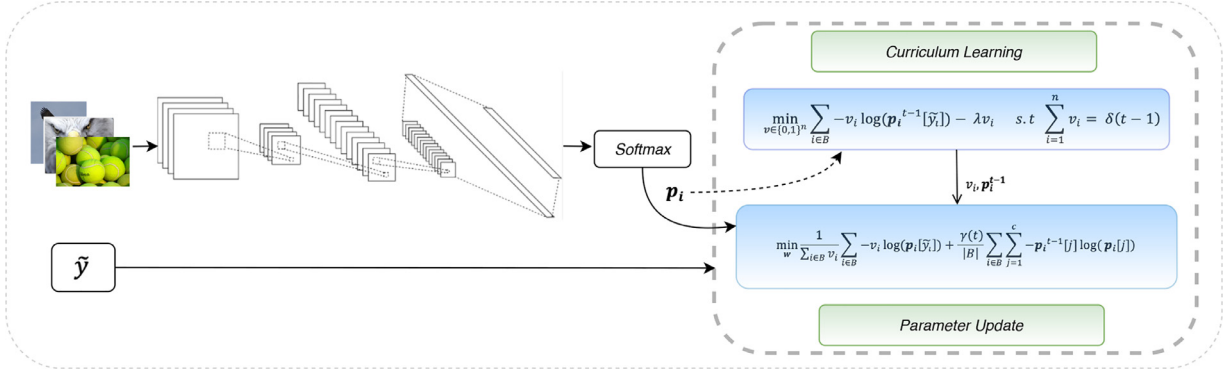


Fig. 2. The flowchart of the proposed SPRL, which alternatively learns a curriculum \mathbf{v} and updates model parameters \mathbf{w} during training, i.e., SPRL first conducts curriculum learning by using the model prediction \mathbf{p}_i^{t-1} to obtain v_i , and then feed v_i and \mathbf{p}_i^{t-1} into the model to update parameters \mathbf{w} . Note that \mathbf{p}_i with the dotted line denotes the \mathbf{p}_i^{t-1} and \mathbf{p}_i with the solid line represents the model prediction in the t th epoch, and $\tilde{\mathbf{y}}$ is used for both curriculum learning and parameter update.

4.1. Curriculum learning using the memorization effect

Given n training samples $\mathbf{X} = \{\mathbf{x}_i\}_{i=1}^n$, $\tilde{\mathbf{y}} = \{\tilde{y}_i\}_{i=1}^n$ denotes their corresponding noisy labels, where \mathbf{x}_i is the i th training sample, $\tilde{y}_i \in \{1, \dots, c\}$ is its label and c is the number of classes. To avoid the abuse of symbols, we utilize $f(\cdot)$ to represent an L -layer convolutional neural network and \mathbf{w} to denote model parameters. Let $\mathbf{P} = \{\mathbf{p}_i\}_{i=1}^n$ be label predictions of training samples and B represent the index set of selected training data in each mini-batch, where $\mathbf{p}_i = f(\mathbf{x}_i, \mathbf{w}) \in \mathbb{R}^c$ is the label prediction of the sample \mathbf{x}_i . To update model parameters, we adopt the cross-entropy loss function as follows:

$$\min_{\mathbf{w}} \frac{1}{|B|} \sum_{i \in B} -\log(\mathbf{p}_i[\tilde{y}_i]), \quad (4)$$

where $|B|$ denotes the length of the index set B .

Suppose that we train the network for T epochs in total. When we only utilize Eq. (4) to update model parameters during training, the model performance usually deteriorates after a few epochs, because the network might first memorize the correct and easy samples at initial epochs and then it will eventually overfit on the corrupted labels or outliers [15]. Based on this memorization of deep networks, we first run the model T_1 epochs and then select m samples based on the small-loss distances to construct a predefined curriculum, which contains the probably correct data. Afterwards, we gradually add a number of samples into the curriculum every a few epochs for training. For clarity, we formulate this procedure as the following model:

$$\begin{aligned} \min_{\mathbf{w}, \mathbf{v}} \frac{1}{\sum_{i \in B} v_i} \sum_{i \in B} -v_i \log(\mathbf{p}_i[\tilde{y}_i]) - \lambda v_i, \\ \text{s.t. } \mathbf{v} \in \{0, 1\}^n, \sum_{i=1}^n v_i = \delta(t), \end{aligned} \quad (5)$$

where $\delta(t)$ is a piecewise linear function to determine how many training samples are added into the curriculum and t is the current number of training epochs. Note that there are many possibilities to set $\delta(t)$. To make confident samples play a better role in model training, we define it as:

$$\delta(t) = \begin{cases} n & t \leq T_1 \\ \min(m + \left\lfloor \frac{t - T_1}{\frac{T - T_1}{K - \frac{m}{K} + 1}} \right\rfloor \left\lfloor \frac{n}{K} \right\rfloor, n) & T_1 < t \leq T \end{cases} \quad (6)$$

where $K \in \mathbb{Z}$ is the number of subsets, each of which contains some training samples. During the first T_1 epochs, we utilize all training data for model training so as to obtain correct-label data as many as possible, because the model firstly memorizes the probably correct-label data from the training data. Additionally, after T_1 epochs, because the number of samples is usually much

larger than the number of epochs, we gradually add the subset into the curriculum in each epoch based on the prediction confidence of samples, so that all training samples can be added in the curriculum during the training process in order to leverage more information. Eq. (6) suggests that Eq. (5) is equivalent to Eq. (4) when $t \leq T_1$.

4.2. Resistance loss

Directly using Eq. (5) to train CNNs is similar to self-paced learning, and the model will gradually overfit on corrupted labels, with the increasing number of training epochs. To address this problem, we propose a resistance loss using the cross entropy between model predictions of previous and current training epochs in each mini-batch. Because knowledge distillation methods [52–55] using model predictions as the teacher can also alleviate model overfitting, we present the core idea of the proposed resistance loss and their differences in Fig. 3, where the resistance loss (Fig. 3a) contains c weighted cross-entropy losses (see Eq. (7)), \mathbf{p}_i^{t-1} is the weight, \mathbf{p}_i in the bottom row is the model prediction, and $\mathbf{p}_i[j] = 1$ ($1 \leq j \leq c$) in the top row is the target; Knowledge distillation (Fig. 3b) employs the prediction \mathbf{p}_i^{t-1} of previous training epoch or a peer model as a teacher (see Eq. (15) in Section 5.2).

Suppose that $\mathbf{p}_i[j]$ is the label prediction of the sample \mathbf{x}_i belonging to the j th class in the t th training epoch, and $\mathbf{p}_i^{t-1}[j] \in \mathbf{p}_i^{t-1}$ is the label prediction before using \mathbf{x}_i to update model parameters in the $t - 1$ th training epoch. We propose the resistance loss as follows:

$$\min_{\mathbf{w}} \frac{1}{|B|} \sum_{i \in B} \sum_{j=1}^c -\mathbf{p}_i^{t-1}[j] \log(\mathbf{p}_i[j]). \quad (7)$$

Eq. (7) is used to resist model overfitting of CNNs on corrupted labels for boosting model robustness. It is mainly inspired by: (i) Eq. (7) might smooth the update of model parameters; (ii) Eq. (7) tends to make $\mathbf{p}_i[j] \rightarrow \frac{1}{c}$ ($1 \leq j \leq c$). Let \mathbf{p}_i^t be the label prediction of \mathbf{x}_i before using it to update model parameters in the t th training epoch, to better illustrate these two motivations, we present Proposition 1 and show its proof in the following.

Proposition 1. Suppose that solving the problem in Eq. (7) with gradient descent, for any two entries $\mathbf{p}_i[j], \mathbf{p}_i[k] \in \mathbf{p}_i$ and $\mathbf{p}_i^t[j] > \mathbf{p}_i^t[k]$. There are only three cases between $\frac{\mathbf{p}_i^t[j]}{\mathbf{p}_i^t[k]}$ and $\frac{\mathbf{p}_i^{t-1}[j]}{\mathbf{p}_i^{t-1}[k]}$: (i) $\frac{\mathbf{p}_i^{t-1}[j]}{\mathbf{p}_i^{t-1}[k]} < \frac{\mathbf{p}_i^t[j]}{\mathbf{p}_i^t[k]}$; (ii) $\frac{\mathbf{p}_i^{t-1}[j]}{\mathbf{p}_i^{t-1}[k]} > (\frac{\mathbf{p}_i^t[j]}{\mathbf{p}_i^t[k]})^2$; (iii) $\frac{\mathbf{p}_i^t[j]}{\mathbf{p}_i^t[k]} \leq \frac{\mathbf{p}_i^{t-1}[j]}{\mathbf{p}_i^{t-1}[k]} \leq (\frac{\mathbf{p}_i^t[j]}{\mathbf{p}_i^t[k]})^2$. For case (i) and (ii), there exists $\frac{\mathbf{p}_i^t[j]}{\mathbf{p}_i^t[k]} < \frac{\mathbf{p}_i^{t-1}[j]}{\mathbf{p}_i^{t-1}[k]}$ and $\frac{\mathbf{p}_i^{t-1}[j]}{\mathbf{p}_i^{t-1}[k]} > \frac{\mathbf{p}_i^t[j]}{\mathbf{p}_i^t[k]} > \frac{\mathbf{p}_i^{t-1}[j]}{\mathbf{p}_i^{t-1}[k]}$, respectively, thereby smoothing the update of model parameters; for case

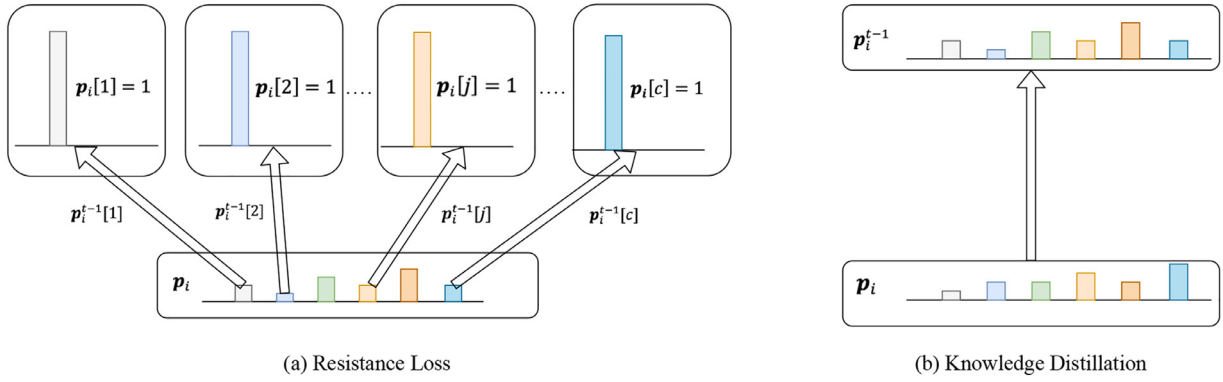


Fig. 3. The core idea of the proposed resistance loss and its difference from knowledge distillation.

(iii), there exists $\frac{\mathbf{p}_i[j]}{\mathbf{p}_i[k]} \leq \frac{\mathbf{p}_i^{t-1}[j]}{\mathbf{p}_i^{t-1}[k]} \leq \frac{\mathbf{p}_i^{t-1}[j]}{\mathbf{p}_i^{t-1}[k]}$, upon which each entry in \mathbf{p}_i tends to be gradually equivalent, i.e. $\mathbf{p}_i[j] = \mathbf{p}_i[k] = \frac{1}{c}, \forall 1 \leq j, k \leq c$.

Proof: Let $E(\mathbf{p}_i) = \sum_{j=1}^c -\mathbf{p}_i^{t-1}[j] \log(\mathbf{p}_i[j])$, taking its derivative with respect to (w.r.t) $\mathbf{p}_i[j]$, we have:

$$\frac{\partial E(\mathbf{p}_i)}{\partial \mathbf{p}_i[j]} = -\frac{\mathbf{p}_i^{t-1}[j]}{\mathbf{p}_i[j]}, \quad (8)$$

which means that $\nabla E(\mathbf{p}_i^{t-1}[j]) = -\frac{\mathbf{p}_i^{t-1}[j]}{\mathbf{p}_i^{t-1}[j]}$. Here, the entries in \mathbf{p}_i are independent, because Eq. (7) is used as c weighted cross-entropy losses (please refer to Fig. 3a). Note that in this paper, log utilizes e as its base.

If $c = 1$, then $\mathbf{p}_i[j] = \mathbf{p}_i^{t-1}[j] + \eta \frac{\mathbf{p}_i^{t-1}[j]}{\mathbf{p}_i^{t-1}[j]}$, where η denotes the learning rate. Because $\mathbf{p}_i^{t-1}[j] > 0, \mathbf{p}_i^{t-1}[j] > 0$ and $\eta > 0$, $\mathbf{p}_i[j]$ will gradually approximate to 1, i.e. $-\mathbf{p}_i^{t-1}[j] \log(\mathbf{p}_i[j]) \rightarrow 0$.

If $c > 1$, for any two entries $\mathbf{p}_i^{t-1}[j] > \mathbf{p}_i^{t-1}[k], 1 \leq j, k \leq c$, then when $\eta > 0$, there exists:

$$\frac{\mathbf{p}_i[j]}{\mathbf{p}_i[k]} = \frac{\mathbf{p}_i^{t-1}[j] + \eta \frac{\mathbf{p}_i^{t-1}[j]}{\mathbf{p}_i^{t-1}[j]}}{\mathbf{p}_i^{t-1}[k] + \eta \frac{\mathbf{p}_i^{t-1}[k]}{\mathbf{p}_i^{t-1}[k]}}. \quad (9)$$

Based on Eq. (9), when $\frac{\mathbf{p}_i^{t-1}[j]}{\mathbf{p}_i^{t-1}[k]} < \frac{\mathbf{p}_i^{t-1}[j]}{\mathbf{p}_i^{t-1}[k]}$, there exists $\frac{\mathbf{p}_i^{t-1}[j]}{\mathbf{p}_i^{t-1}[k]} < 1$,

leading to $\frac{\mathbf{p}_i[j]}{\mathbf{p}_i[k]} < \frac{\mathbf{p}_i^{t-1}[j]}{\mathbf{p}_i^{t-1}[k]}$, thereby smoothing the update of model parameters. In addition, if $\frac{\mathbf{p}_i^{t-1}[j]}{\mathbf{p}_i^{t-1}[k]} < 1$, there exists $\frac{\mathbf{p}_i^{t-1}[j]}{\mathbf{p}_i^{t-1}[k]} < \frac{\mathbf{p}_i[j]}{\mathbf{p}_i[k]} < \frac{\mathbf{p}_i^{t-1}[j]}{\mathbf{p}_i^{t-1}[k]}$.

When $\frac{\mathbf{p}_i^{t-1}[j]}{\mathbf{p}_i^{t-1}[k]} > (\frac{\mathbf{p}_i^{t-1}[j]}{\mathbf{p}_i^{t-1}[k]})^2$, i.e., $\frac{\mathbf{p}_i^{t-1}[j]}{\mathbf{p}_i^{t-1}[k]} > \frac{\mathbf{p}_i^{t-1}[j]}{\mathbf{p}_i^{t-1}[k]}$, it has $\frac{\mathbf{p}_i[j]}{\mathbf{p}_i[k]} > \frac{\mathbf{p}_i^{t-1}[j]}{\mathbf{p}_i^{t-1}[k]}$.

Eq. (9) equals $\frac{\mathbf{p}_i[j]}{\mathbf{p}_i[k]} = \frac{(\mathbf{p}_i^{t-1}[j])^2 + \eta \mathbf{p}_i^{t-1}[j]}{(\mathbf{p}_i^{t-1}[k])^2 + \eta \mathbf{p}_i^{t-1}[k]} \cdot \frac{\mathbf{p}_i^{t-1}[k]}{\mathbf{p}_i^{t-1}[j]}$ and $\frac{\mathbf{p}_i^{t-1}[k]}{\mathbf{p}_i^{t-1}[j]} < 1$, so they suggest $\frac{\mathbf{p}_i[j]}{\mathbf{p}_i[k]} < \frac{(\mathbf{p}_i^{t-1}[j])^2 + \eta \mathbf{p}_i^{t-1}[j]}{(\mathbf{p}_i^{t-1}[k])^2 + \eta \mathbf{p}_i^{t-1}[k]} < \frac{\mathbf{p}_i^{t-1}[j]}{\mathbf{p}_i^{t-1}[k]}$. Thus, $\frac{\mathbf{p}_i^{t-1}[j]}{\mathbf{p}_i^{t-1}[k]} > \frac{\mathbf{p}_i[j]}{\mathbf{p}_i[k]} > \frac{\mathbf{p}_i^{t-1}[j]}{\mathbf{p}_i^{t-1}[k]}$, which means model parameters' update would be smoothed.

When $\frac{\mathbf{p}_i^{t-1}[j]}{\mathbf{p}_i^{t-1}[k]} \leq (\frac{\mathbf{p}_i^{t-1}[j]}{\mathbf{p}_i^{t-1}[k]})^2$, i.e., $\frac{\mathbf{p}_i^{t-1}[j]}{\mathbf{p}_i^{t-1}[k]} \leq \frac{\mathbf{p}_i^{t-1}[j]}{\mathbf{p}_i^{t-1}[k]}$, it has $\frac{\mathbf{p}_i[j]}{\mathbf{p}_i[k]} \leq \frac{\mathbf{p}_i^{t-1}[j]}{\mathbf{p}_i^{t-1}[k]}$.

With an additional constraint $\frac{\mathbf{p}_i^{t-1}[j]}{\mathbf{p}_i^{t-1}[k]} \geq \frac{\mathbf{p}_i^{t-1}[j]}{\mathbf{p}_i^{t-1}[k]}$, it means $\frac{\mathbf{p}_i[j]}{\mathbf{p}_i[k]} \leq \frac{\mathbf{p}_i^{t-1}[j]}{\mathbf{p}_i^{t-1}[k]}$. In this case, each entry in \mathbf{p}_i will gradually becomes equiv-

alent, i.e. $\mathbf{p}_i[j] = \mathbf{p}_i[k]$. With a constraint $\sum_{j=1}^c \mathbf{p}_i^{t-1}[j] = 1$, there will be $\mathbf{p}_i[j] = \mathbf{p}_i[k] \rightarrow \frac{1}{c}$.

Therefore, Proposition 1 is proved.

In practice, because cases (i) and (ii) in Proposition 1 smoothly update model parameters, the relationship between $\frac{\mathbf{p}_i^{t-1}[j]}{\mathbf{p}_i^{t-1}[k]}$ and

$\frac{\mathbf{p}_i^{t-1}[j]}{\mathbf{p}_i^{t-1}[k]}$ might gradually satisfy the case (iii), thereby causing each

entry in \mathbf{p}_i gradually to be equivalent, where $\mathbf{p}_i[j]$ and $\mathbf{p}_i[k]$ ($\forall 1 \leq j, k \leq c$) are gradually equivalent, i.e. $\mathbf{p}_i[j] = \mathbf{p}_i[k] \rightarrow \frac{1}{c}$, so that

the objective of Eq. (7) becomes larger. This infers the case (iii) in Proposition 1, i.e. Eq. (7) can gradually make each entry of probability ratios be equivalent.

For clarity, Fig. 4 presents two examples to show the change of the objective in Eq. (7) from 101 to 300 epochs during training, with randomly selecting 1000 digits from '0' to '9' in MNIST [56] and 1000 images belonging to 10 categories from CIFAR-10 [57], respectively. Specifically, we first utilize Eq. (4) to train ResNet18 [58] with 100 epochs, and then adopt Eq. (7) to train the network for the subsequent 200 epochs. The loss gradually approximates to $\sum_{j=1}^{10} -0.1 * \log(0.1) = 2.303$.

Moreover, Fig. 5 shows an example to display the change of prediction probability of training data on their true classes when using Eqs. (4), (5) and (10) (Eqs. (5)+(7)), with randomly selecting training data from CIFAR-10 [57] and flip their labels using Eq. (14b). Figure 5 presents that the curve in Fig. 5c is smoother than that in Fig. 5a-b. This suggests that Eq. (7) can smooth the update of model parameters.

Figure 5c also illustrates that SPRL using Eq. (7) can resist model overfitting on corrupted labels. Note that SPRL aims to smooth all samples via the resistance loss during training, thereby causing the decrease of prediction probability on correct labels in Fig. 5c. However, correct-label data, which usually obtain high prediction confidence and their number is larger than that of corrupt-label ones in each class, play a significant role in guiding the model to attain correct predictions.

Figure 5c also illustrates that SPRL using Eq. (7) can resist model overfitting on corrupted labels. Note that SPRL aims to smooth all samples via the resistance loss during training, thereby causing the decrease of prediction probability on correct labels in Fig. 5c. However, correct-label data, which usually obtain high prediction confidence and their number is larger than that of corrupt-label ones in each class, play a significant role in guiding the model to attain correct predictions.

4.3. Self-paced resistance loss

Based on the learned curriculum and the proposed resistance loss, we can obtain the loss function of the proposed framework. Specifically, combining Eqs. (5) with (7), we have:

$$\begin{aligned} \min_{\mathbf{w}, \mathbf{v}} E(\mathbf{w}, \mathbf{v}; \lambda) &= \frac{1}{\sum_{i \in B} v_i} \sum_{i \in B} -v_i (\log(\mathbf{p}_i[\tilde{y}_i]) + \lambda) \\ &+ \frac{\gamma(t)}{|B|} \sum_{i \in B} \sum_{j=1}^c -\mathbf{p}_i^{t-1}[j] \log(\mathbf{p}_i[j]), \end{aligned} \quad (10)$$

s.t. $\mathbf{v} \in \{0, 1\}^n, \sum_{i=1}^n v_i = \delta(t)$,

where $\gamma(t)$ is a time-dependent weighting function to gradually enhance the weight of model predictions with the increasing number of epochs, so that Eq. (7) is mainly used to prevent model overfitting on corrupted labels. Because DNNs might first memo-




Fig. 4. The loss of Eqs. (4) and (7) at different numbers of training epochs.




Fig. 5. The prediction probability of training data on their true classes with respect to training epochs, with training ResNet18 [58] on noisy-label data by using three methods: (a) Standard (Eq. (4)), (b) CL (Eq. (5)), (c) SPRL (Eq. (10)).

size the correct-label data and then corrupt-label samples, and the noise rate of selected samples usually increases eventually.

There are many choices for $\gamma(t)$. Similar to the popular ramp-up function in [12], we utilize the following function:

$$\gamma(t) = \begin{cases} 0 & t \leq T_1 \\ \gamma_{\max} e^{-5\|1-\mu\|_F^2} & T_1 < t \leq T, \end{cases} \quad (11)$$

where μ linearly ramps up from 0 to 1 during $T - T_1$ epochs, γ_{\max} is the maximum of $\gamma(t)$ depending on m , e.g. $\gamma_{\max} = \gamma_d(10 - \lceil \frac{m}{0.1n} \rceil)$. This is because a larger γ_{\max} is required for a larger noise rate.

The optimization of Eq. (10) is similar to that of Eqs. (1) and (2), and thus we solve it by utilizing an alternative minimization strategy [50,51]. Specifically, it can be divided into two sub-problems:

$$\begin{aligned} & \min_{\mathbf{v}} \sum_{i \in B} -v_i \log(\mathbf{p}_i^{t-1}[\tilde{y}_i]) - \lambda v_i, \\ & \text{s.t. } \mathbf{v} \in \{0, 1\}^n, \sum_{i=1}^n v_i = \delta(t-1). \end{aligned} \quad (12a)$$

$$\begin{aligned} & \min_{\mathbf{w}} \frac{1}{\sum_{i \in B} v_i} \sum_{i \in B} -v_i \log(\mathbf{p}_i[\tilde{y}_i]) \\ & + \frac{\gamma(t)}{|B|} \sum_{i \in B} \sum_{j=1}^c -\mathbf{p}_i^{t-1}[j] \log(\mathbf{p}_i[j]). \end{aligned} \quad (12b)$$

Eq. (12a) is a \mathbf{v} -subproblem, in which the model parameter \mathbf{w} is known, and it aims to learn a curriculum consisting of confident samples; Eq. (12b) is a \mathbf{w} -subproblem, which consists of a cross-entropy loss to utilize selected confident samples to update model parameters, and a resistance loss to resist model overfitting of CNNs on corrupted labels. We alternatively solve Eqs. (12a) and (12b), i.e. fixing \mathbf{w} , based on Eq. (12a), we can calculate \mathbf{v} as follows:

$$v_i^* = \begin{cases} 1 & \text{if } -\log(\mathbf{p}_i^{t-1}[\tilde{y}_i]) < \lambda \\ 0 & \text{otherwise.} \end{cases} \quad (13)$$

Note that in each epoch we might need to adjust λ so that $\sum_{i=1}^n v_i = \delta(t-1)$. Then with a fixed \mathbf{v} , we can update the model parameter \mathbf{w} by solving Eq. (12b) via any optimizer, e.g. Adam [59]. In summary, we present the detailed procedure to solve Eq. (10) in Algorithm 1.

5. Experiments

To evaluate the proposed SPRL, we conduct experiments on four large-scale benchmark datasets: MNIST, CIFAR-10, CIFAR-100 and Mini-ImageNet. We briefly introduce them in the following.

MNIST [56] consists of 70K images with handwritten digits from '0' to '9'. There are 60K training and 10K testing images, each of which has a size of 28×28 .

CIFAR-10 [57] contains 60K color images belonging to 10 classes, each of which consists of 6K images. There are 50K train-

Algorithm 1 SPRL.

Input: Training data $\mathbf{X} = \{\mathbf{x}_i\}_{i=1}^n$, noisy labels $\tilde{\mathbf{y}} = \{\tilde{y}_i\}_{i=1}^n$, number of training epochs: T_1, T , parameters λ, K, γ_d piecewise linear function $\delta(t)$, stochastic neural network with parameters \mathbf{w} : $f(\cdot)$, stochastic input augmentation function: $h(\cdot)$

Output: Parameters \mathbf{w}

1. **for** t **in** $[1, T_1]$ **do**
2. **for** each mini-batch B **do**
3. $\mathbf{p}_{i \in B} \leftarrow f(h(\mathbf{x}_{i \in B}))$
4. loss \leftarrow Eq. (4)
5. updating \mathbf{w} using optimizers, e.g. Adam
6. **end for**
7. **end for**
8. **for** t **in** $[T_1 + 1, T]$ **do**
9. $\mathbf{v} \leftarrow$ Eq. (13) ▷ Adjust λ to make $\sum_{i=1}^n v_i = \delta(t - 1)$
10. **for** each mini-batch B **do**
11. $\mathbf{p}_{i \in B} \leftarrow f(h(\mathbf{x}_{i \in B}))$
12. loss \leftarrow Eq. (12b)
13. updating \mathbf{w} using optimizers, e.g. Adam
14. $\mathbf{p}_{i \in B}^{t-1} \leftarrow \mathbf{p}_{i \in B}$
15. **end for**
16. **end for**

ing and 10K testing images. Each one is aligned and cropped to 32×32 pixels.

CIFAR-100 [57] has 60K color images in 100 classes, with 600 images per class. There are also 50K training and 10K testing images. Each image has a size of 32×32 .

Mini-ImageNet [60] is more complex than CIFAR-100. It is composed of 60K color images selected from the ImageNet dataset [61]. These images belong to 100 classes, with 600 images per class. We divide them into a training set with 50K images and a testing set containing 10K images, and resize each image to 32×32 .

The images in CIFAR-10, CIFAR-100 and Mini-ImageNet datasets are with the popular augmentation: random translations ($\{\Delta x, \Delta y\} \sim [-4, 4]$) and horizontal flip ($p = 0.5$), and each image in MNIST is only augmented by the random translation ($\{\Delta x, \Delta y\} \sim [-2, 2]$).

For these four datasets, their noisy labels are generated by symmetry and pair flipping. Specifically, following [13,15], we corrupt the four datasets manually via a label transition matrix \mathbf{Q} that is calculated by $q_{ij} = \Pr(\tilde{y} = j | y = i)$, where the noisy label \tilde{y} is flipped from the correct label y . Similar to Han et al. [15], here \mathbf{Q} has two representative structures: symmetric flipping (class-independent noise) and pair flipping (class-dependent noise). For clarity, we present the definition of \mathbf{Q} with symmetric and pair flipping structures in Eqs. (14a) and (14b), respectively. It is worth noting that for symmetric flipping, the noise rate ϵ should be smaller than $\frac{c-1}{c}$, i.e. $\epsilon < \frac{c-1}{c}$; for pair flipping, $\epsilon < 0.5$ so that more than half of labels are correct. Note that the noise rate ϵ denotes the ratio of corrupted labels in the whole training data.

$$\mathbf{Q} = \begin{bmatrix} 1 - \epsilon & \frac{\epsilon}{c-1} & \cdots & \frac{\epsilon}{c-1} & \frac{\epsilon}{c-1} \\ \frac{\epsilon}{c-1} & 1 - \epsilon & \frac{\epsilon}{c-1} & \cdots & \frac{\epsilon}{c-1} \\ \vdots & & \ddots & & \vdots \\ \frac{\epsilon}{c-1} & \cdots & \frac{\epsilon}{c-1} & 1 - \epsilon & \frac{\epsilon}{c-1} \\ \frac{\epsilon}{c-1} & \frac{\epsilon}{c-1} & \cdots & \frac{\epsilon}{c-1} & 1 - \epsilon \end{bmatrix} \quad (14a)$$

$$\mathbf{Q} = \begin{bmatrix} 1 - \epsilon & \epsilon & 0 & \cdots & 0 \\ 0 & 1 - \epsilon & \epsilon & & 0 \\ \vdots & & \ddots & & \vdots \\ 0 & & & 1 - \epsilon & \epsilon \\ \epsilon & 0 & \cdots & 0 & 1 - \epsilon \end{bmatrix} \quad (14b)$$

In addition to symmetry and pair flipping labels, we also conduct experiments on noisy labels generated by CNNs with using the datasets CIFAR-10 and CIFAR-100. Specifically, we uniformly select 4K and 10K images from the training set of CIFAR-10 and CIFAR-100 as labeled data, and view the remaining images of training sets as unlabeled ones. Next, we only utilize labeled data to train models, and then employ the well-trained model to predict the whole training data to generate noisy labels.

Moreover, we conduct experiments on real-world noisy labels from the datasets Food101 and Clothing1M, respectively. Specifically,

Food101 [62] contains 101K images belonging to 101 food categories, with 750 training and 250 testing images per category. Training images are with noisy labels, while testing images have clean labels.

Clothing1M [40] consists of 1 million clothing images in 14 classes. We utilize training images with noisy labels for model training and 10K testing images with clean labels for testing.

5.1. Implementation details

We implement SPRL with the PyTorch framework and employ a 13-layer convolutional neural network (ConvNet) [12,63] or ResNet18 [58] as the backbone network. We adopt the optimizer, Adam [59], to update the network parameters, with initializing the momentum parameters $\beta_1 = 0.9$ and $\beta_2 = 0.999$. By default, we follow [15] to set the maximum learning rate η to be 0.001, run the network for $T = 200$ epochs and set the batch size to be 128. When using ResNet18 on MNIST, we choose $\eta = 0.0001$ to avoid exploding gradient. After the first 80 epochs, β_1 becomes 0.1 and the learning rate linearly decreases to 0 over the following 120 epochs. T_1 can be obtained through a validation set. Specifically, we randomly select 10% noisy training data to construct a validation set. T_1 is the epoch number, at which the network attains the best validation accuracy, in order to obtain the best model predictions. When the noise rate ϵ is not known, m is the maximum number of training data whose prediction $\mathbf{p}_i[y_i] \geq 0.5$ ($1 \leq i \leq n$) during the first T_1 epochs; when ϵ is known, we can empirically choose m within the range of $[0.5(1 - \epsilon)n, 0.8(1 - \epsilon)n]$. Additionally, m should satisfy $m \in [0.1n, 0.5n]$, because the noise rate ϵ is usually smaller than 0.9 and a large m might reduce the effect of curriculum learning. There are many choices for K , we set $K = 10$. γ_d can be estimated with cross-validation on noisy validation sets. For clarity, we present the detailed parameter settings (T_1 and γ_d) of each experiment in the supplemental materials (Please refer to Tables A3-A4).

5.2. Experimental settings

We compare the proposed SPRL with seven state-of-the-art algorithms. We briefly introduce them as follows:

Standard: the standard DNNs trained on noisy datasets.

Bootstrap [11]: which corrects the label by using the weighted combination of predicted and original labels. We adopt hard labels in our experiments because they usually perform better than soft ones.

F-correction [13]: which utilizes a label transition matrix to correct model predictions. We employ the forward strategy, which

Table 1

The best testing accuracy (%) of eight different methods on MNIST, CIFAR-10, CIFAR-100 and Mini-ImageNet with clean training data ($\epsilon = 0$). We bold the best accuracy and its similar results (within 0.5%).

Method	ResNet18			
	MNIST	CIFAR-10	CIFAR-100	Mini-ImageNet
Standard	99.63	93.06	72.35	55.26
Bootstrap	99.65	94.25	73.03	58.27
F-correction	99.65	94.08	72.92	58.03
Decoupling	99.68	92.10	69.87	47.74
MentorNet	99.65	93.87	70.46	53.48
Co-teaching	99.58	92.88	72.68	57.04
Co-teaching+	99.58	93.17	70.48	55.88
SPRL	99.67	94.20	73.88	63.04
Method	ConvNet			
	MNIST	CIFAR-10	CIFAR-100	Mini-ImageNet
Standard	99.67	92.66	71.04	53.82
Bootstrap	99.69	93.57	72.34	56.36
F-correction	99.69	93.76	73.01	58.66
Decoupling	99.48	92.35	70.03	47.68
MentorNet	99.71	92.15	69.23	55.40
Co-teaching	99.74	93.60	72.04	58.57
Co-teaching+	99.69	92.74	70.96	59.10
SPRL	99.69	92.76	72.23	60.27

usually yields better performance, and utilize a validation set to estimate the label transition matrix.

Decoupling [18]: which updates model parameters using the samples with different predictions of two classifiers.

MentorNet [14]: which adopts an additional network to learn an approximate predefined curriculum and employs another network, StudentNet, for classification. We utilize self-paced MentorNet, which is used for the case that no clean validation data is known.

Co-teaching [15]: which trains two networks in a symmetric way and each network selects the samples with the small-loss distance as the confident data for the other one.

Co-teaching+ [19]: which is based on Co-teaching but using the strategy of "Update by Disagreement" [18].

Here, we suppose that the noise rate is known in Co-teaching and Co-teaching+, but the noise rate is unknown in the proposed SPRL. For fairness, we re-implement all the seven state-of-the-art algorithms with the PyTorch framework based on their provided public codes and utilize their default parameter settings. Additionally, they adopt the same backbone networks and training procedure as SPRL.

5.3. Experiments on labels with symmetry and pair flipping


To better illustrate the strength of the proposed SPRL, we first run all the eight methods with clean training data of the four datasets, and then present their best testing accuracy in Table 1. As we can see, SPRL can achieve better or very competitive testing accuracy to the best competitors when using clean training data, and it consistently outperforms Standard, especially for more difficult datasets CIFAR-100 and Mini-ImageNet. A main possible reason is that the proposed method could reduce overfitting caused by outliers. This finding is very important. It shows that the proposed method has a wide range of applications.

Table 2 shows the average of testing accuracy of the proposed SPRL and seven compared algorithms on MNIST, CIFAR-10, CIFAR-100 and Mini-ImageNet over the last ten epochs. It illustrates that SPRL significantly outperforms the other seven algorithms on the four datasets, especially on extremely noisy labels. For example, when using ResNet18, for symmetric flipping with $\epsilon = 0.8$, the average accuracy of SPRL is 22.43%, 15.48%, 9.36% and 13.54% higher than the best competitors on the four datasets, respectively; for pair flipping with $\epsilon = 0.45$, its accuracy is 11.86%, 12.80%, 14.10%

Table 2


Average of testing accuracy (%) on MNIST, CIFAR-10, CIFAR-100 and Mini-ImageNet over the last ten epochs. We bold the best results and highlight the second best ones via underlines.

Method	ResNet18			
	Symmetry		Pair	
	$\epsilon = 0.2$	$\epsilon = 0.5$	$\epsilon = 0.8$	$\epsilon = 0.45$
MNIST				
Standard	92.69 ± 0.18	65.49 ± 0.33	24.59 ± 0.18	58.50 ± 0.34
Bootstrap	93.89 ± 0.08	66.48 ± 0.63	24.38 ± 0.35	59.92 ± 0.52
F-correction	97.08 ± 0.11	92.86 ± 0.14	40.93 ± 0.30	10.32 ± 0.01
Decoupling	95.70 ± 0.64	72.60 ± 4.17	27.00 ± 0.52	71.58 ± 2.35
MentorNet	93.51 ± 0.01	83.10 ± 0.01	24.96 ± 0.01	82.51 ± 0.01
Co-teaching	96.89 ± 0.11	91.01 ± 0.13	75.92 ± 0.47	87.44 ± 0.33
Co-teaching+	<u>99.04 ± 0.02</u>	<u>94.69 ± 0.15</u>	38.34 ± 1.23	87.36 ± 0.39
SPRL	99.58 ± 0.01	99.53 ± 0.01	98.35 ± 0.16	99.30 ± 0.01
CIFAR-10				
Standard	79.47 ± 0.24	45.37 ± 0.53	10.01 ± 0.19	51.48 ± 0.80
Bootstrap	83.39 ± 0.26	57.24 ± 0.48	17.96 ± 0.27	51.95 ± 0.57
F-correction	80.11 ± 0.16	45.92 ± 0.65	6.51 ± 0.19	52.35 ± 0.48
Decoupling	76.60 ± 1.25	53.63 ± 1.21	16.18 ± 0.17	50.22 ± 2.58
MentorNet	80.04 ± 0.21	53.20 ± 0.16	<u>42.02 ± 0.19</u>	49.93 ± 0.11
Co-teaching	88.34 ± 0.23	81.64 ± 0.19	<u>32.39 ± 0.26</u>	<u>79.09 ± 0.43</u>
Co-teaching+	90.87 ± 0.10	<u>83.52 ± 0.10</u>	23.18 ± 0.10	60.07 ± 0.56
SPRL	92.68 ± 0.03	88.25 ± 0.06	57.50 ± 0.11	91.89 ± 0.06
CIFAR-100				
Standard	54.14 ± 0.21	28.73 ± 0.19	7.03 ± 0.10	35.24 ± 0.08
Bootstrap	55.94 ± 0.29	31.37 ± 0.22	7.48 ± 0.10	37.07 ± 0.21
F-correction	56.32 ± 0.12	37.73 ± 0.08	9.09 ± 0.08	37.79 ± 0.19
Decoupling	54.56 ± 0.61	30.51 ± 0.36	7.37 ± 0.09	36.74 ± 0.28
MentorNet	52.11 ± 0.16	26.71 ± 0.16	12.76 ± 0.08	33.92 ± 0.14
Co-teaching	62.71 ± 0.13	48.14 ± 0.15	15.94 ± 0.10	<u>39.49 ± 0.23</u>
Co-teaching+	<u>66.41 ± 0.12</u>	<u>51.65 ± 0.13</u>	<u>19.17 ± 0.09</u>	34.96 ± 0.28
SPRL	70.93 ± 0.06	59.31 ± 0.07	28.53 ± 0.10	53.59 ± 0.06
Mini-ImageNet				
Standard	34.07 ± 0.18	16.17 ± 0.21	3.55 ± 0.10	22.78 ± 0.16
Bootstrap	34.91 ± 0.39	18.22 ± 0.19	4.14 ± 0.18	24.28 ± 0.27
F-correction	31.81 ± 0.14	12.29 ± 0.10	2.13 ± 0.04	6.13 ± 0.07
Decoupling	33.38 ± 0.17	16.63 ± 0.12	4.26 ± 0.06	22.78 ± 0.19
MentorNet	29.19 ± 0.01	14.14 ± 0.01	1.06 ± 0.01	21.89 ± 0.01
Co-teaching	48.84 ± 0.09	36.98 ± 0.17	5.86 ± 0.11	<u>29.21 ± 0.11</u>
Co-teaching+	51.13 ± 0.14	36.86 ± 0.23	<u>7.23 ± 0.05</u>	27.46 ± 0.07
SPRL	57.24 ± 0.09	47.66 ± 0.11	20.77 ± 0.09	39.53 ± 0.07
ConvNet				
	$\epsilon = 0.2$	$\epsilon = 0.5$	$\epsilon = 0.8$	$\epsilon = 0.45$
MNIST				
Standard	86.84 ± 0.27	60.80 ± 0.59	24.80 ± 0.51	57.14 ± 0.56
Bootstrap	91.48 ± 0.16	61.05 ± 0.69	21.11 ± 0.32	55.51 ± 0.72
F-correction	87.12 ± 0.19	66.36 ± 0.45	58.17 ± 0.54	57.70 ± 0.64
Decoupling	96.26 ± 0.24	88.93 ± 0.40	71.02 ± 0.39	61.22 ± 2.34
MentorNet	95.78 ± 0.01	92.44 ± 0.01	47.55 ± 0.01	73.67 ± 0.01
Co-teaching	98.91 ± 0.04	96.55 ± 0.07	<u>89.54 ± 0.26</u>	93.64 ± 0.14
Co-teaching+	99.56 ± 0.01	<u>99.15 ± 0.02</u>	<u>77.77 ± 0.03</u>	<u>97.25 ± 0.17</u>
SPRL	99.56 ± 0.01	99.43 ± 0.01	97.52 ± 0.04	99.28 ± 0.01
CIFAR-10				
Standard	77.82 ± 0.27	48.11 ± 0.42	22.31 ± 0.34	50.73 ± 0.62
Bootstrap	75.34 ± 0.84	47.37 ± 0.74	18.00 ± 0.53	51.11 ± 0.74
F-correction	84.26 ± 0.21	62.90 ± 0.42	11.58 ± 0.16	61.98 ± 0.40
Decoupling	83.49 ± 0.19	68.73 ± 0.27	<u>40.16 ± 0.30</u>	50.61 ± 3.12
MentorNet	81.80 ± 0.01	73.62 ± 0.14	27.90 ± 0.05	52.96 ± 0.02
Co-teaching	86.40 ± 2.58	83.06 ± 0.15	28.39 ± 0.31	<u>80.21 ± 0.58</u>
Co-teaching+	90.43 ± 0.10	85.87 ± 0.08	20.41 ± 0.04	77.51 ± 0.22
SPRL	90.47 ± 0.06	85.99 ± 0.05	60.42 ± 0.12	83.69 ± 0.12
CIFAR-100				
Standard	23.87 ± 0.15	9.37 ± 0.11	34.83 ± 0.30	
Bootstrap	50.51 ± 0.25	25.21 ± 0.18	9.66 ± 0.15	34.43 ± 0.21
F-correction	54.42 ± 0.13	33.19 ± 0.11	5.54 ± 0.08	38.70 ± 0.34
Decoupling	53.99 ± 0.17	32.84 ± 0.10	14.81 ± 0.09	37.31 ± 0.23
MentorNet	52.70 ± 0.01	38.75 ± 0.02	11.02 ± 0.01	31.87 ± 0.01
Co-teaching	66.30 ± 0.43	57.29 ± 0.13	<u>19.96 ± 0.17</u>	37.50 ± 0.19
Co-teaching+	69.00 ± 0.12	<u>59.79 ± 0.13</u>	11.57 ± 0.07	<u>43.21 ± 0.21</u>
SPRL	<u>67.65 ± 0.10</u>	59.81 ± 0.12	35.82 ± 0.14	47.26 ± 0.11
Mini-ImageNet				
Standard	38.06 ± 0.31	19.61 ± 0.25	7.96 ± 0.11	26.69 ± 0.16
Bootstrap	38.73 ± 0.38	19.12 ± 0.15	5.75 ± 0.11	27.60 ± 0.22
F-correction	33.45 ± 0.18	26.96 ± 0.08	2.24 ± 0.03	5.04 ± 0.06
Decoupling	30.37 ± 0.15	15.46 ± 0.17	6.21 ± 0.05	24.12 ± 0.11
MentorNet	43.47 ± 0.02	31.09 ± 0.01	1.80 ± 0.01	27.01 ± 0.01
Co-teaching	53.62 ± 0.13	43.54 ± 0.15	5.51 ± 0.05	30.68 ± 0.16
Co-teaching+	<u>54.99 ± 0.13</u>	<u>45.02 ± 0.23</u>	6.06 ± 0.06	<u>33.94 ± 0.15</u>
SPRL	55.32 ± 0.08	46.32 ± 0.13	24.40 ± 0.08	37.78 ± 0.12



(a) Symmetry $\epsilon = 0.2$ (b) Symmetry $\epsilon = 0.5$ (c) Symmetry $\epsilon = 0.8$ (d) Pair $\epsilon = 0.45$

Fig. 6. Testing accuracy of SPRL with different values of γ_d on MNIST when $T_1 = 15$. ‘Average’ means the average of testing accuracy over the last ten epochs, and ‘Max’ denotes the maximum of testing accuracy among all training epochs.



(a) Symmetry $\epsilon = 0.2$ (b) Symmetry $\epsilon = 0.5$ (c) Symmetry $\epsilon = 0.8$ (d) Pair $\epsilon = 0.45$

Fig. 7. Testing accuracy of SPRL with different values of γ_d on CIFAR-100 when $T_1 = 40$. ‘Average’ means the average of testing accuracy over the last ten epochs, and ‘Max’ denotes the maximum of testing accuracy among all training epochs.

and 10.32% higher than the best competitors on the four datasets, respectively. The superior accuracy of SPRL over the others can also be observed when using ConvNet. Note that, the implementation results of Co-teaching with ConvNet are significantly better than the reported ones in [15]. Because we utilize the data augmentation, which boosts the model performance. Moreover, we present testing accuracy of the eight methods at different numbers of training epochs on the four datasets in the supplemental materials (please see Figures. A4–A7), which further illustrate that SRL can obtain the best accuracy among all methods on two different network architectures, and its accuracy is much smoother than that of the others during training.

5.3.1. Parameter analysis

The proposed SPRL has three essential parameters γ_d , K , and T_1 , where γ_d and K determine $\gamma(t)$ and $\delta(t)$, respectively, and T_1 determines m . Here, we evaluate them by utilizing an easy dataset MNIST, a complex dataset CIFAR-100 and the network ConvNet. Specifically, Figures 6–7 show testing accuracy of SPRL with different values of γ_d on symmetric or pair flipping label noise, including $\gamma_d \in \{0, 1, 5, 10, 50, 100, 300, 500, 1000\}$ on MNIST and $\gamma_d \in \{0, 1, 3, 5, 10, 30, 50, 100, 300\}$ on CIFAR-100. Additionally, we present the testing accuracy of SPRL with different K and T_1 in Figures. A1–A2, and display the effect of different noise rates on the loss function Eq. (12b) during training in Figure A3 in the supplemental material.

Figures 6–7 suggest that a large weight of the resistance loss (Eq. (7)) can prevent the performance degradation of CNNs on symmetric or pair flipping label noise. Additionally, Fig. 7 also suggests that Eq. (7) with a large weight can boost the model accuracy. When $\gamma_d \geq 50$, SPRL obtains the best or sub-optimal accuracy on MNIST; when $\gamma_d \in [1, 10]$, SPRL obtains the best or sub-optimal accuracy on CIFAR-100 with symmetric label noise, and when $\gamma_d \in [10, 50]$, it achieves the best or sub-optimal accuracy on pair flipping label noise. However, Fig. 7 illustrates that if γ_d is too large, the model accuracy will decrease on CIFAR-100, probably be-

cause the resistance loss with model predictions tends to make the prediction on each class be equivalent. Furthermore, $\gamma_d = 0$ means removing the resistance loss (Eq. (7)) from the proposed loss function (Eq. (10)), as shown in Figures 6–7, the proposed resistance loss is very helpful to improve performance.

5.3.2. Comparison with knowledge distillation and label smooth regularization

Knowledge distillation [52] and label smooth [2] are two popular methods for boosting the model generalization. Here, we utilize Eq. (15) to distill knowledge from previous training epochs and Eq. (16) to smooth labels, and replace Eq. (7) with them in Eq. (10), respectively. They are:

$$\min_{\mathbf{w}} \frac{1}{|B|} \sum_{i \in B} \mathbf{p}_i^{t-1} \log \left(\frac{\mathbf{p}_i^{t-1}}{\mathbf{p}_i} \right), \tag{15}$$

$$\min_{\mathbf{w}} \frac{1}{|B|} \sum_{i \in B} \mathbf{u}_i \log \left(\frac{\mathbf{u}_i}{\mathbf{p}_i} \right), \tag{16}$$

where \mathbf{w} denotes model parameters and $\mathbf{u}_i = \left\{ \frac{1}{c}, \frac{1}{c}, \dots, \frac{1}{c} \right\} \in \mathbb{R}^c$.

Figure 8 shows their performance using ResNet18 as the backbone network on CIFAR-10 and CIFAR-100 with symmetric label noise. It demonstrates the superior performance of Eq. (7) over Eqs. (15) and (16).

5.4. Experiments on noisy labels generated by CNNs

In practice, labels might be not only symmetric or pair flipping. To further illustrate the strength of the proposed SPRL, we conduct experiments on noisy labels that are generated by CNNs. Table 3 presents the accuracy of trained models on training and testing sets of CIFAR-10 and CIFAR-100, by only employing 4K and 10K labeled data for model training, respectively. Then we apply trained models on the whole training set and utilize predicting labels as



(a) CIFAR-10@Symmetry $\epsilon = 0.5$ (b) CIFAR-100@Symmetry $\epsilon = 0.5$

Fig. 8. Testing accuracy of the proposed framework using ResNet18 with Eqs. (15), (16) and Eq. (7), respectively.

Table 3

Accuracy (%) of ResNet18 and ConvNet trained by partially labeled data on training and testing sets of CIFAR-10 and CIFAR-100 datasets (4K for CIFAR-10 and 10K for CIFAR-100).

Network	CIFAR-10		CIFAR-100	
	Training	Testing	Training	Testing
ResNet18	81.97	80.93	63.16	54.64
ConvNet	82.02	80.52	64.23	54.95

Table 4

Average of testing accuracy (%) on CIFAR-10 and CIFAR-100 over the last ten epochs by CNN generated noisy labels. We bold the best results and highlight the second best ones via underlines.

Method	CIFAR-10		CIFAR-100	
	ResNet18	ConvNet	ResNet18	ConvNet
Standard	81.94	81.88	54.86	54.52
Bootstrap	81.24	81.71	54.59	55.18
F-correction	83.40	81.28	54.20	55.25
Decoupling	79.31	78.46	49.80	50.79
MentorNet	81.27	80.24	52.39	54.06
Co-teaching	<u>82.80</u>	<u>82.80</u>	<u>55.44</u>	<u>55.56</u>
Co-teaching+	82.26	81.73	54.66	55.16
SPRL	85.63	84.00	62.08	58.73

noisy labels. Finally, we run the eight methods by utilizing training data with noisy labels to train models.

Table 4 shows the average accuracy of the eight methods on test sets of CIFAR-10 and CIFAR-100 over the last ten epochs. As shown in Tables 3–4, both SPRL and Co-teaching with noisy labels can consistently outperform ResNet18 and ConvNet with only partially labeled data. However, SPRL always achieves better average accuracy than the best competitor, Co-teaching, on two different deep architectures and datasets, especially on heavy noisy labels, e.g. labels (36.84% noise rate) generated by ResNet18, which is trained with only partially labeled data of CIFAR-100. We also present testing accuracy of the eight methods at different numbers of training epochs in the supplemental materials (please see Figure A8).

Table 5

Average accuracy (%) of four methods over the last ten epochs on real-world noisy labels.

Network	Food101		Clothing1M	
	ResNet18	ConvNet	ResNet18	ConvNet
Standard	71.01	73.56	66.59	68.15
Co-teaching	<u>71.36</u>	<u>74.34</u>	<u>69.78</u>	69.94
Co-teaching+	69.84	71.10	67.93	<u>70.08</u>
SPRL	76.14	74.61	71.63	71.81

5.5. Experiments on real-world noisy labels

To be further demonstrate the strength of the proposed SPRL on boosting model robustness, we conduct experiments on real-world noisy labels from the datasets Food101 and Clothing1M, respectively.

Table 5 displays the average accuracy of Standard, Co-teaching, Co-teaching+ and SPRL on Food101 and Clothing1M over the last ten epochs. It illustrates that SPRL consistently outperforms the best competitors Co-teaching and Co-teaching+ on real-world noisy labels, especially using ResNet18 as the backbone network. Additionally, We show their testing accuracy at different numbers of training epochs in the supplemental materials (please refer to Figure A9).

6. Conclusion and future work

In this paper, we propose a novel framework, SPRL, to alleviate model overfitting for robustly training CNNs on noisy labels. The proposed framework contains two major modules: curriculum learning, which utilizes the memorization skill of DNNs to learn a curriculum to provide meaningful supervision for other training samples; parameters update, which leverages the selected confident samples and a resistance loss to simultaneously update model parameters and significantly reduce the effect of corrupted labels. Experiments on multiple large-scale benchmark datasets demonstrate that SPRL can significantly reduce the effects of various types of corrupted labels by using the resistance loss to alleviate model overfitting. Additionally, experiments on noisy labels generated by CNNs suggest that SPRL can be potentially utilized to further improve the performance of semi-supervised and unsupervised deep methods.

Although SPRL has achieved robust and better accuracy than many state-of-the-art methods, SPRL cannot be directly applied on multi-label datasets with noisy labels [64], because it calculates the class probability of each sample by using the softmax function, which usually performs poorly on multi-label classification tasks. However, SPRL might be extended to handle multi-label tasks by replacing the softmax function with a sigmoid function. In the future, SPRL might be further improved based on the following two potential directions: (i) Introducing a small amount of clean validation data for training [65], instead of training models with only noisy training data without using any clean validation data. (ii) Employing model predictions of SPRL to generate labels to further boost the model performance (like Section 5.4), or distinguishing and changing the possibly corrupted labels by using the other popular methods [66–68].

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.patcog.2022.109080.

References

- [1] A. Krizhevsky, I. Sutskever, G.E. Hinton, Imagenet classification with deep convolutional neural networks, in: *Advances in Neural Information Processing Systems*, 2012, pp. 1097–1105.
- [2] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, Z. Wojna, Rethinking the inception architecture for computer vision, in: *IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 2818–2826.
- [3] Y. Feng, Z. Zhang, X. Zhao, R. Ji, Y. Gao, GVCNN: group-view convolutional neural networks for 3d shape recognition, in: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [4] Y. Feng, H. You, Z. Zhang, R. Ji, Y. Gao, Hypergraph neural networks, in: *AAAI Conference on Artificial Intelligence*, vol. 33, 2019, pp. 3558–3565.
- [5] X. Shi, M. Sapkota, F. Xing, F. Liu, L. Cui, L. Yang, Pairwise based deep ranking hashing for histopathology image classification and retrieval, *Pattern Recognit.* 81 (2018) 14–22.
- [6] X. Shi, Z. Guo, F. Xing, Y. Liang, L. Yang, Anchor-based self-ensembling for semi-supervised deep pairwise hashing, *Int. J. Comput. Vis.* (2020) 1–18.
- [7] R. Girshick, J. Donahue, T. Darrell, J. Malik, Rich feature hierarchies for accurate object detection and semantic segmentation, in: *IEEE Conference on Computer Vision and Pattern Recognition*, 2014, pp. 580–587.
- [8] J. Long, E. Shelhamer, T. Darrell, Fully convolutional networks for semantic segmentation, in: *IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 3431–3440.
- [9] C. Zhang, S. Bengio, M. Hardt, B. Recht, O. Vinyals, Understanding deep learning requires rethinking generalization, in: *International Conference on Learning Representations*, 2017.
- [10] D. Zheng, J. Xiao, K. Chen, X. Huang, L. Chen, Y. Zhao, Soft pseudo-label shrinkage for unsupervised domain adaptive person re-identification, *Pattern Recognit.* 127 (2022) 108615.
- [11] S. Reed, H. Lee, D. Anguelov, C. Szegedy, D. Erhan, A. Rabinovich, Training deep neural networks on noisy labels with bootstrapping, in: *Workshop of International Conference on Learning Representations*, 2015.
- [12] S. Laine, T. Aila, Temporal ensembling for semi-supervised learning, in: *International Conference on Learning Representations*, 2016.
- [13] G. Patrini, A. Rozza, A. Krishna Menon, R. Nock, L. Qu, Making deep neural networks robust to label noise: a loss correction approach, in: *IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 1944–1952.
- [14] L. Jiang, Z. Zhou, T. Leung, L.-J. Li, L. Fei-Fei, MentorNet: learning data-driven curriculum for very deep neural networks on corrupted labels, in: *International Conference on Machine Learning*, 2018.
- [15] B. Han, Q. Yao, et al., Co-teaching: robust training of deep neural networks with extremely noisy labels, in: *Advances in Neural Information Processing Systems*, 2018, pp. 8527–8537.
- [16] O. Chapelle, B. Scholkopf, A. Zien, et al., *Semi-supervised learning* (Chapelle, O. et al., Eds.; 2006)[Book Reviews], *IEEE Trans. Neural Netw.* 20 (3) (2009), 542–542.
- [17] D. Arpit, S. J. et al., A closer look at memorization in deep networks, in: *International Conference on Machine Learning*, 2017, pp. 233–242.
- [18] E. Malach, S. Shalev-Shwartz, Decoupling when to “update” from “how to update”, in: *Advances in Neural Information Processing Systems*, 2017, pp. 960–970.
- [19] X. Yu, B. Han, J. Yao, G. Niu, I. Tsang, M. Sugiyama, How does disagreement help generalization against label corruption? in: *International Conference on Machine Learning*, 2019, pp. 7164–7173.
- [20] B. Fréna, M. Verleysen, Classification in the presence of label noise: a survey, *IEEE Trans. Neural Netw. Learn. Syst.* 25 (5) (2013) 845–869.
- [21] V.C. Raykar, S. Yu, et al., Learning from crowds, *J. Mach. Learn. Res.* 11 (Apr) (2010) 1297–1322.
- [22] N. Natarajan, I.S. Dhillon, P.K. Ravikumar, A. Tewari, Learning with noisy labels, in: *Advances in Neural Information Processing Systems*, 2013, pp. 1196–1204.
- [23] H. Masnadi-Shirazi, N. Vasconcelos, On the design of loss functions for classification: theory, robustness to outliers, and savageboost, in: *Advances in Neural Information Processing Systems*, 2009, pp. 1049–1056.
- [24] B. Van Rooyen, A. Menon, R.C. Williamson, Learning with symmetric label noise: the importance of being unhinged, in: *Advances in Neural Information Processing Systems*, 2015, pp. 10–18.
- [25] C. Scott, G. Blanchard, G. Handy, Classification with asymmetric label noise: consistency and maximal denoising, in: *Conference On Learning Theory*, 2013, pp. 489–511.
- [26] H. Ramaswamy, C. Scott, A. Tewari, Mixture proportion estimation via kernel embeddings of distributions, in: *International Conference on Machine Learning*, 2016, pp. 2052–2060.
- [27] T. Sanderson, C. Scott, Class proportion estimation with application to multiclass anomaly rejection, in: *Artificial Intelligence and Statistics*, 2014, pp. 850–858.
- [28] T. Liu, D. Tao, Classification with noisy labels by importance reweighting, *IEEE Trans. Pattern Anal. Mach. Intell.* 38 (3) (2015) 447–461.
- [29] H. Song, M. Kim, D. Park, Y. Shin, J.-G. Lee, Learning from noisy labels with deep neural networks: a survey, *IEEE Trans. Neural Netw. Learn. Syst.* (2022).
- [30] V. Mnih, G.E. Hinton, Learning to label aerial images from noisy data, in: *International Conference on Machine Learning*, 2012, pp. 567–574.
- [31] A. Ghosh, H. Kumar, P.S. Sastry, Robust loss functions under label noise for deep neural networks, in: *AAAI Conference on Artificial Intelligence*, vol. 31, 2017.
- [32] A. Ghosh, N. Manwani, P.S. Sastry, Making risk minimization tolerant to label noise, *Neurocomputing* 160 (2015) 93–107.
- [33] N. Manwani, P.S. Sastry, Noise tolerance under risk minimization, *IEEE Trans. Cybern.* 43 (3) (2013) 1146–1151.
- [34] X. Shi, H. Su, et al., Graph temporal ensembling based semi-supervised convolutional neural network with noisy labels for histopathology image analysis, *Med. Image Anal.* 60 (2020) 101624.
- [35] S. Sukhbaatar, J. Bruna, M. Paluri, L. Bourdev, R. Fergus, Training convolutional networks with noisy labels, in: *Workshop of International Conference on Learning Representations*, 2015.
- [36] X. Ma, Y. Wang, et al., Dimensionality-driven learning with noisy labels, in: *International Conference on Machine Learning*, 2018.
- [37] Y. Wang, W. Liu, X. Ma, J. Bailey, H. Zha, L. Song, S.-T. Xia, Iterative learning with open-set noisy labels, in: *IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 8688–8696.
- [38] J. Han, P. Luo, X. Wang, Deep self-learning from noisy labels, in: *IEEE International Conference on Computer Vision*, 2019, pp. 5138–5147.
- [39] H. Sun, C. Guo, Q. Wei, Z. Han, Y. Yin, Learning to rectify for robust learning with noisy labels, *Pattern Recognit.* 124 (2022) 108467.
- [40] T. Xiao, T. Xia, Y. Yang, C. Huang, X. Wang, Learning from massive noisy labeled data for image classification, in: *IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 2691–2699.
- [41] Y. Li, J. Yang, Y. Song, L. Cao, J. Luo, L.-J. Li, Learning from noisy labels with distillation, in: *IEEE International Conference on Computer Vision*, 2017, pp. 1910–1918.
- [42] A. Veit, N. Alldrin, et al., Learning from noisy large-scale datasets with minimal supervision, in: *IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 839–847.
- [43] A. Vahdat, Toward robustness against label noise in training deep discriminative neural networks, in: *Advances in Neural Information Processing Systems*, 2017, pp. 5596–5605.
- [44] M. Ren, W. Zeng, B. Yang, R. Urtasun, Learning to reweight examples for robust deep learning, in: *International Conference on Machine Learning*, 2018.
- [45] C.G. Northcutt, T. Wu, I.L. Chuang, Learning with confident examples: Rank pruning for robust classification with noisy labels, *Uncertainty in Artificial Intelligence*, 2017.
- [46] Z. Zhang, M. Sabuncu, Generalized cross entropy loss for training deep neural networks with noisy labels, in: *Advances in Neural Information Processing Systems*, 2018, pp. 8778–8788.
- [47] B. Han, I.W. Tsang, L. Chen, P.Y. Celina, S.-F. Fung, Progressive stochastic learning for noisy labels, *IEEE Trans. Neural Netw. Learn. Syst.* 29 (10) (2018) 5136–5148.
- [48] Q. Yao, H. Yang, B. Han, G. Niu, J.T.-Y. Kwok, Searching to exploit memorization effect in learning with noisy labels, in: *International Conference on Machine Learning*, 2020, pp. 10789–10798.
- [49] Y. Bengio, J. Louradour, R. Collobert, J. Weston, Curriculum learning, in: *International Conference on Machine Learning*, ACM, 2009, pp. 41–48.
- [50] M.P. Kumar, B. Packer, D. Koller, Self-paced learning for latent variable models, in: *Advances in Neural Information Processing Systems*, 2010, pp. 1189–1197.

- [51] L. Jiang, D. Meng, Q. Zhao, S. Shan, A.G. Hauptmann, Self-paced curriculum learning, in: AAAI Conference on Artificial Intelligence, 2015.
- [52] G. Hinton, O. Vinyals, J. Dean, Distilling the knowledge in a neural network, *Stat* 1050 (2015) 9.
- [53] A. Polino, R. Pascanu, D. Alistarh, Model compression via distillation and quantization, in: International Conference on Learning Representations, 2018.
- [54] K. Lee, K. Lee, J. Shin, H. Lee, Overcoming catastrophic forgetting with unlabeled data in the wild, in: IEEE International Conference on Computer Vision, 2019, pp. 312–321.
- [55] K. Kim, B. Ji, D. Yoon, S. Hwang, Self-knowledge distillation with progressive refinement of targets, in: IEEE International Conference on Computer Vision, 2021, pp. 6567–6576.
- [56] Y. LeCun, L. Bottou, Y. Bengio, P. Haffner, et al., Gradient-based learning applied to document recognition, *Proc. IEEE* 86 (11) (1998) 2278–2324.
- [57] A. Krizhevsky, G. Hinton, Learning Multiple Layers of Features from Tiny Images, Technical Report, Citeseer, 2009.
- [58] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: IEEE Conference on Computer Vision and Pattern Recognition, 2016, pp. 770–778.
- [59] D.P. Kingma, J. Ba, Adam: a method for stochastic optimization, in: International Conference on Learning Representations, 2015.
- [60] O. Vinyals, C. Blundell, et al., Matching networks for one shot learning, in: Advances in Neural Information Processing Systems, 2016, pp. 3630–3638.
- [61] O. Russakovsky, J. Deng, et al., ImageNet large scale visual recognition challenge, *Int. J. Comput. Vis.* 115 (3) (2015) 211–252.
- [62] L. Bossard, M. Guillaumin, L. Van Gool, Food-101 – mining discriminative components with random forests, in: European Conference on Computer Vision, 2014.
- [63] A. Rasmus, M. Berglund, M. Honkala, H. Valpola, T. Raiko, Semi-supervised learning with ladder networks, in: Advances in Neural Information Processing Systems, 2015, pp. 3546–3554.
- [64] K.O. Mikalsen, C. Soguero-Ruiz, F.M. Bianchi, R. Jenssen, Noisy multi-label semi-supervised dimensionality reduction, *Pattern Recognit.* 90 (2019) 257–270.
- [65] Z. Zhang, H. Zhang, S.O. Arik, H. Lee, T. Pfister, Distilling effective supervision from severe label noise, in: IEEE Conference on Computer Vision and Pattern Recognition, 2020, pp. 9294–9303.
- [66] E. Arazo, D. Ortego, et al., Unsupervised label noise modeling and loss correction, in: International Conference on Machine Learning, 2019, pp. 312–321.
- [67] D. Berthelot, N. Carlini, I. Goodfellow, N. Papernot, A. Oliver, C.A. Raffel, Mix-Match: a holistic approach to semi-supervised learning, *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [68] J. Li, R. Socher, S.C.H. Hoi, DivideMix: learning with noisy labels as semi-supervised learning, in: International Conference on Learning Representations, 2019.

Xiaoshuang Shi is a faculty member in the Department of Computer Science and Engineering at the University of Electronic Science and Technology of China (UESTC). He obtained his PhD degree (2019) from University of Florida, Master degree (2013) from Tsinghua University, and Bachelor degree (2009) from Northwestern Polytechnical University. Before joining UESTC, he worked as a Postdoctoral fellow at the National Institutes of Health (NIH) (2020.01–2021.04), and as a research assistant at Tsinghua University (2013.09–2015.04). His major research interests include large-scale data retrieval, deep learning, medical image analysis.

Zhenhua Guo received the MS and PhD degrees in computer science from the Harbin Institute of Technology and The Hong Kong Polytechnic University, in 2004 and 2010, respectively. Since 2010, he has been with the Graduate School at Shenzhen, Tsinghua University. His research interests include pattern recognition, texture classification, biometrics, and video surveillance.

Kang Li received the PhD degree in Mechanical Engineering from University of Illinois at Urbana Champaign, Champaign, IL, USA, in 2009. He is now a full professor of the Biomedical Big Center at West China Hospital. Before joining West China Hospital, He was an associate professor with the Department of Orthopaedics, New Jersey Medical School (NJMS), Rutgers University, Newark, NJ, USA, and an assistant professor with Department of Industrial and Systems Engineering, Rutgers University. His research interests include AI in healthcare, musculoskeletal biomechanics, medical imaging, design and biorobotics, human reliability, and human factors/ergonomics.

Yun Liang is a PhD student in Biomedical Engineering at the University of Florida. She received the bachelor degree in Computing from the Hong Kong Polytechnic University and the master degree in Statistics from George Washington University. Her current research is in medical image processing and computer vision, especially in whole slide cytopathology image analysis.

Xiaofeng Zhu is a faculty member of University of Electronic Science and Technology of China, Chengdu, China. His current research interests include large-scale multimedia retrieval, feature selection, sparse learning, data preprocess, and medical image analysis.