

Co-learning: Learning from Noisy Labels with Self-supervision

Cheng Tan^{1,2}, Jun Xia^{1,2}, Lirong Wu^{1,2}, Stan Z. Li^{1,2} *

¹AI Lab, School of Engineering, Westlake University.

²Institute of Advanced Technology, Westlake Institute for Advanced Study

18 Shilongshan Road, Hangzhou 310024, Zhejiang Province, China

{tancheng,xiajun,wulirong,stan.zq.li}@westlake.edu.cn

ABSTRACT

Noisy labels, resulting from mistakes in manual labeling or webly data collecting for supervised learning, can cause neural networks to overfit the misleading information and degrade the generalization performance. Self-supervised learning works in the absence of labels and thus eliminates the negative impact of noisy labels. Motivated by co-training with both supervised learning view and self-supervised learning view, we propose a simple yet effective method called *Co-learning* for learning with noisy labels. Co-learning performs supervised learning and self-supervised learning in a cooperative way. The constraints of intrinsic similarity with the self-supervised module and the structural similarity with the noisily-supervised module are imposed on a shared common feature encoder to regularize the network to maximize the agreement between the two constraints. Co-learning is compared with peer methods on corrupted data from benchmark datasets fairly, and extensive results are provided which demonstrate that Co-learning is superior to many state-of-the-art approaches. The code is available at GitHub.

CCS CONCEPTS

• Computing methodologies → Computer vision.

KEYWORDS

label noise, self-supervised learning, computer vision

ACM Reference Format:

Cheng Tan^{1,2}, Jun Xia^{1,2}, Lirong Wu^{1,2}, Stan Z. Li^{1,2}. 2021. Co-learning: Learning from Noisy Labels with Self-supervision. In *Proceedings of the 29th ACM International Conference on Multimedia (MM '21)*, October 20–24, 2021, Virtual Event, China. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3474085.3475622>

1 INTRODUCTION

The success of deep learning relies on carefully labeled data in large-scale. However, precise annotations are extremely expensive and time-consuming. To alleviate the problem, inexpensive alternatives are often used, such as web crawling [35] or completing

annotations with crowd-sourcing and online queries. Unfortunately, these alternative methods can often leads to noisy labels. However, deep neural networks can easily overfit to noisy labels, as shown in recent research [18, 22, 34, 40], and this can dramatically degrade the generalization performance.




Figure 1: Comparison of leading methods in dealing with noisy labels. All the methods use two cooperative modules (A & B). As the training goes, the two become more and more in agreement with each other, but in different ways. *Decoupling* updates the parameters of the two networks with prediction-disagreed (!=) samples from a mini-batch. *Co-teaching* [11] uses small-loss samples of one network to teach its peer network for further training. *Co-teaching+* [39] first predicts each mini-batch with the two networks but uses disagreed (!=) samples only to compute the training loss. *JoCoR* [37] trains the two networks as a whole with a joint loss of weighted goals: making the two predictions agree with each other, and making the predictions stick to ground-true labels as far as possible. *Co-learning* trains a single shared encoder network with two heads (the self-supervised and the noisily-supervised) that constrain each other and maximizes the agreement between them in latent space.

Several studies have been conducted to investigate learning with noisy labels, where semi-supervised learning frameworks play crucial roles [1, 8, 17, 25, 38]. Most of these deploy unsupervised learning to obtain information about label-independent features and then do further training using noisy labels in a supervised learning manner. These methods commonly leverage labels through two views: (1) supervised learning taking advantage of direct supervision with labels and (2) unsupervised learning exploiting intrinsic features from data to combat label noise. The core challenge here is how to combine the two views effectively. Arazo et al. [1] design a beta mixture model as an unsupervised generative model of sample loss values and then adopt MixUp [41] augmentation to assure reliable convergence under extreme noise levels. Li et al. [17] models sample loss with Gaussian mixtures to divide training data into unlabeled data and labeled data, and then apply the state-of-the-art

*Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MM '21, October 20–24, 2021, Virtual Event, China

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8651-7/21/10...\$15.00

<https://doi.org/10.1145/3474085.3475622>

semi-supervised learning method MixMatch [2] to deal with the divided sets. However, the above methods are not end-to-end and the step-by-step learning manner increases complexity in both time and space.

Co-training and model ensembling have been shown to be beneficial in learning with noisy labels [11, 23, 37, 39]. Decoupling [23] trains two networks simultaneously and updates them using instances with different predictions. Co-teaching [11] also trains two networks simultaneously and selects small-loss data to teach the peer network during the training process. Co-teaching+ [39] follows a similar scheme as Co-teaching but with a scheme which selects small-loss data from disagreement data. JoCoR [37] maintains two networks but trains them as a whole with a joint loss to make predictions of them get closer. These methods are based on the assumption that two networks can provide two different views of the data, but the extra information gain they bring is very limited since the differences between two networks of the same architecture mainly come from random initialization. A comparison between Co-learning and other co-training-like methods is illustrated in Figure 1.

Motivated by semi-supervised learning and co-training, we propose a simple yet effective learning paradigm called "*Co-learning*" to combat the problem with noisy labels, in which self-supervised learning is introduced as a feature-dependent view to assist supervised learning. Different from those co-training-like methods, Co-learning has a shared feature encoder with two exclusive heads. While the projection head performs a self-supervised learning and exploits feature-dependent information through the intrinsic similarity, the classifier head performs a vanilla supervised learning and learns from label-dependent information. In addition, a structural similarity loss is utilized between the two exclusive heads to regularize the classifier and avoid bias caused by noisy labels. This is based on the constraint that output from the projection head and that from the classifier head should share similar structure pairwise. Unlike the other methods, Co-learning can be implemented conveniently without the need for such prior knowledge as noise rates, data distributions, and additional clean samples, thus avoiding the needed hyperparameters thereof.

We conducted experiments on simulated and real-world noisy datasets, including CIFAR-10, CIFAR-100, Animal-10N [32], and Food-101N [16] datasets. Among these datasets, Animal-10N is a noisy dataset that contains confusing images for manual annotations, and Food-101N is a webly dataset that directly collects data from the Web. Comparative results demonstrate that Co-learning is superior to state-of-the-art methods and robust to high-level noise in labels.

2 RELATED WORK

2.1 Learning with noisy labels

Various methods have been studied towards robust learning with noisy labels, and we briefly review existing works on learning with noisy labels in this section.

Noise transition matrix. Some of the existing methods estimate the label transition matrix. Forward-correction [29] performs forward correction by multiplying the estimated transition matrix with

the softmax outputs during the forward propagation step. Gold loss correction [14] leverages available trusted labels for loss correction and utilizes trusted labels' confusion matrices to obtain an accurate transition matrix. However, estimating the transition matrix is challenging in real-world settings and maybe infeasible for finding closed-form solutions.

Sample selection. Another direction of handling noisy labels without noise transition matrix is training on selected samples. These approaches try to select "clean" instances according to certain rules. In particular, MentorNet [15] pretrains a teacher network for selecting clean instances to guide the training of a student network. Co-teaching [11] chooses clean instances from a network to feed its peer network for updating parameters. Co-teaching+ [39] and JoCoR [37] also highlight the small-loss sample selection to explicitly assist the co-training procedures of two networks. INCV [5] divides noisy training data randomly to apply cross-validation and removes large-loss samples. This family of methods avoids the risk of error correction by simply excluding unreliable samples. Although these methods have achieved significant improvements, they neglect the potential supervisory signal of the large-loss training samples. Besides, either the noise rate or a clean validation set is required as prior knowledge.

Semi-supervised learning. Several recent studies transform the problem of learning with noisy labels into a semi-supervised learning task. There are mainly two directions in these methods: (1) divide noisy training data into two sets where reliable samples are treated as labeled data, and unreliable samples are treated as unlabeled data, then perform certain existing semi-supervised learning methods [1, 8, 17, 25, 38], (2) perform unsupervised learning on noisy training data to assist the supervised learning [12, 24]. Our method follows the second direction that utilizes unsupervised learning to implicitly help to learn from noisy labels.

2.2 Self-supervised learning

Self-supervised learning is a subset of unsupervised learning that designs pretext tasks to produce labels derived from the data itself. While solving pretext tasks, the model learns valuable representations; that is, the feature encoder of the model is able to extract useful features. There are various pretext tasks that can help the model learn representations, such as colorization [42], inpainting [28], rotation [9], jigsaw [26] and contrastive learning [4, 6, 7, 10, 13]. Among them, recent studies on contrastive learning have become a dominant component in self-supervised learning methods. Contrastive learning aims at grouping similar samples closer and diverse samples away from each other. A typical contrastive learning pipeline is shown in Figure 2. To measure the similarity between the feature representations, contrastive learning utilizes a metric that we call *intrinsic similarity* in this paper.

Current state-of-the-art contrastive learning methods mainly rely on a combination of intrinsic similarity and a set of image transformations. MoCo [13] builds a dynamic dictionary with a queue and a momentum encoder to obtain large and consistent dictionaries of visual representations. SimCLR [6] shows that the




Figure 2: Typical contrastive learning pipeline.

composition of data augmentations is crucial and a nonlinear projection head will further improve the quality of the learned representations. While contrastive learning may fail into collapsed solutions like outputting the same representations for all samples, BYOL [10] takes advantage of both the target network and the slow-moving average online network to avoid collapsing. To simplify the framework in BYOL, SimSiam [7] proposes Siamese networks with a stop-gradient operation. SwAV [4] introduces an online clustering strategy to further improve the performance without either a large memory bank or a momentum encoder.

2.3 Relations to other approaches

We compare Co-learning with related methods and illustrate the main differences in Table 1. Although Co-learning is motivated by Co-training and other Co-training-like methods [11, 23, 37, 39], the only similarity is that Co-learning tries to maximize agreements from different views.

	Co-teaching	Co-teaching+	JoCoR	Co-learning
Agreement	✓	✓	✓	✓
Small-loss	✓	✓	✓	✗
Double classifiers	✓	✓	✓	✗
Cross update	✓	✓	✗	✗

Table 1: Comparison with state-of-the-art methods in techniques used – “Agreement”: maximizing the agreement from different views by regularization; “Small-loss”: regarding small-loss samples as “Clean” samples; “Double classifiers”: training two classifiers simultaneously; “Cross update”: updating parameters in a cross manner instead of a parallel manner.

These methods treat the samples with small losses as reliable data and abandon the remaining samples. They cannot take full advantage of all the given data, hindering the potential of the model. This small-loss trick deliberately avoided the problem of difficult samples and noisy samples. Furthermore, their small-loss trick needs noisy rates, and other carefully chosen hyperparameters [11, 37, 39] lead to difficulties in real-world applications.

Our approach approves the assumption that different views can give more reliable information on noisy training data. However, the above methods simply apply two same classifiers that can provide little variety on views. The different views from networks with the same architecture rely on random initialization. Instead of maximizing agreements on models’ predictions, Co-learning tries to maximize agreements on the feature encode f so that both feature-dependent and label-dependent information can play roles during the training procedure. A shared feature encoder in Co-learning

not only ensures the soft agreements from different views but also reduces computation cost to a great extent.

3 METHODS

3.1 Preliminaries

Considering a classification problem with c classes, suppose $\mathcal{X} \subset \mathbb{R}^d$ is the input space and $\mathcal{Y} = \{0, 1\}^c$ is the ground-truth label space in an one-hot manner. The training set $\mathcal{D} = \{(x^{(i)}, y^{(i)})\}_{i=1}^N$ is obtained from a joint distribution over $\mathcal{X} \times \mathcal{Y}$. For the general supervised learning, the model learns a mapping function $\mathcal{F} : \mathcal{X} \rightarrow \mathcal{Y}$ to minimize the empirical risk $\mathcal{R}_{\mathcal{L}}(\mathcal{F})$ under a loss function \mathcal{L} ,

$$\mathcal{R}_{\mathcal{L}}(\mathcal{F}) = \mathbb{E}_{\mathcal{D}} [\mathcal{L}(\mathcal{F}(x), y)] = \frac{1}{|\mathcal{D}|} \sum_{(x, y) \in \mathcal{D}} \mathcal{L}(\mathcal{F}(x), y). \quad (1)$$

In our case, the mapping function \mathcal{F} is a neural network model parameterized by Θ . The optimal parameters Θ^* minimize the empirical risk subjected to network parameters,

$$\Theta^* = \arg \min_{\Theta} \mathcal{R}_{\mathcal{L}}(\mathcal{F}(\cdot; \Theta)) \quad (2)$$

A noisy training dataset $\widehat{\mathcal{D}} = \{(x^{(i)}, \hat{y}^{(i)})\}_{i=1}^N$ is provided from a noisy joint distribution over $\mathcal{X} \times \widehat{\mathcal{Y}}$ where \hat{y}_i is a noisy label. Under the standard training procedure, a mini-batch $\mathcal{B}_t = \{(x^{(i)}, \hat{y}^{(i)})\}_{i=1}^b$ is randomly selected from $\widehat{\mathcal{D}}$ at time t with batch size b . In this condition, parameters of the model may be mistakenly updated according to the gradient direction of the mini-batch \mathcal{B}_t ,

$$\Theta_{t+1} = \Theta_t - \eta \nabla \left(\frac{1}{|\mathcal{B}_t|} \sum_{(x, \hat{y}) \in \mathcal{B}_t} \mathcal{L}(\mathcal{F}(x; \Theta_t), \hat{y}) \right), \quad (3)$$

where η is the learning rate.

The optimization with noisy labels may easily lead to inaccurate direction, thus misguide and degenerate the model generalization performance on unseen data. The goal of learning with noisy labels is to mitigate the mischief of noisy labels.

3.2 Co-learning

The Co-learning framework is illustrated in Figure 3. The main learning task is to learn the encoder and the classifier, which is done with a vanilla cross-entropy loss defined on label-dependent information. A contrastive prediction-based pretext task is used to assist the main task with an *intrinsic similarity* loss defined on feature-dependent information, as motivated by the recent success of self-supervised learning [4, 6, 7, 10, 13]. A *structural similarity* loss is introduced in this work to maximize agreement on the feature encoder shared by the two tasks. We also hope this can help the classifier learn with noisy labels.

Given a mini-batch $\mathcal{B}_t = \{(x^{(i)}, \hat{y}^{(i)})\}_{i=1}^b$, we sample two independent strong transformations from the same distribution \mathcal{T} and one weak transformation from another distribution \mathcal{T}' to obtain the transformed data $\{\tilde{x}_k^{(i)}\}_{k=1,2,3}$, for which the definitions of the strong and weak transformations are illustrated in Appendix. A feature encoder f with parameters θ_1 takes the transformed images as inputs and produces representation vectors $u_k^{(i)} = f(\tilde{x}_k^{(i)}; \theta_1)$. The first representation $u_1^{(i)}$ is fed into the classifier g with parameters θ_2 and formulated as $\tilde{y}^{(i)} = g(u_1^{(i)}; \theta_2)$. The projection head




Figure 3: The architecture of Co-learning. The main task of supervised classifier learning and the auxiliary pretext task of self-supervised learning work cooperatively in the process of learning the encoder and the classifier. The subscripts 1, 2 and 3 index the three transformed versions of training example while the superscripts (i) and (j) index pairs of training examples.

h which is a θ_3 parameterized multi-layer perceptron then non-linearly transforms $u_2^{(i)}, u_3^{(i)}$ into projections, i.e., $v_2^{(i)} = h(v_2^{(i)}; \theta_3)$, $v_3^{(i)} = h(v_3^{(i)}; \theta_3)$.

The intrinsic similarity loss is calculated between $v_2^{(i)}$ and $v_3^{(i)}$ by taking $(v_2^{(i)}, v_3^{(i)})$ as a positive pair and $\{(v_2^{(i)}, v_2^{(j)}), (v_3^{(i)}, v_3^{(j)})\}_{j \neq i}$ as the negative pairs. This indicates the same data has intrinsic similarity through different transformations. The two different heads which perform contrastive prediction and supervised classification simultaneously can be regarded as two different pretext tasks. This indicates the same data has structural similarity through different pretext tasks.

Specifically, the complete loss function \mathcal{L} is a linear combination of the supervised loss \mathcal{L}_{sup} , intrinsic similarity loss \mathcal{L}_{int} , and structural similarity loss \mathcal{L}_{str} as follows:

$$\mathcal{L} = \mathcal{L}_{sup} + \mathcal{L}_{int} + \mathcal{L}_{str} \quad (4)$$

Supervised classifier loss. For multi-class classification problems, the cross-entropy loss is a common supervised loss to measure the relative entropy between the corrupted label \hat{y} and its corresponding prediction \tilde{y} .

$$\mathcal{L}_{ce}(\hat{y}, \tilde{y}) = - \sum_{i=1}^N \hat{y}^{(i)} \log(\tilde{y}^{(i)}) \quad (5)$$

Considering contrastive learning requires a long training period comparing to its supervised counterpart, Co-learning deploys MixUp [41] augmentation to slow down the convergence of the feature encoder f and avoid the supervised learning quickly overfitting to noisy labels. Then we create new training samples by linearly interpolating a sample with another sample (indexed by $m(i)$) randomly chosen from the same mini-batch:

$$\tilde{x}^{(i)} = \lambda \tilde{x}^{(i)} + (1 - \lambda) \tilde{x}^{(m(i))} \quad (6)$$

$$\tilde{y}^{(i)} = \lambda \hat{y}^{(i)} + (1 - \lambda) \hat{y}^{(m(i))} \quad (7)$$

where $\lambda \sim \text{Beta}(\alpha, \alpha)$.

Thus, the MixUp version of the supervised loss is formulated as:

$$\mathcal{L}_{sup} = - \sum_{i=1}^N \left[\lambda \tilde{y}^{(i)} \log(\hat{y}^{(i)}) + (1 - \lambda) \tilde{y}^{(i)} \log(\hat{y}^{(m(i))}) \right] \quad (8)$$

Intrinsic similarity loss. In the contrastive prediction task, projections of the same image from different views should be similar. The intrinsic similarity between $v_2^{(i)}$ and $v_3^{(i)}$ (the projections of the two strongly transformed versions $\tilde{x}_2^{(i)}$ and $\tilde{x}_3^{(i)}$ from $x^{(i)}$) is measured by the contrastive loss function as in InfoNCE [27]:

$$\ell(v_a^{(i)}, v_b^{(i)}) = - \log \frac{\exp(D(v_a^{(i)}, v_b^{(i)})/\tau)}{\sum_{j=1, i \neq j}^N \sum_{t_i, t_j \in \{2, 3\}} \exp(D(v_{t_i}^{(i)}, v_{t_j}^{(j)})/\tau)} \quad (9)$$

where $D(v_a, v_b) = \frac{v_a^T v_b}{\|v_a\| \|v_b\|}$ is the pairwise similarity between v_a and v_b , and τ is a temperature parameter. The intrinsic similarity loss is thus defined as:

$$\mathcal{L}_{int} = \sum_{i=1}^N \ell(v_2^{(i)}, v_3^{(i)}) + \ell(v_3^{(i)}, v_2^{(i)}) \quad (10)$$

Structural similarity loss. The classifier g and the projection head h share the same encoder network, so they jointly influence the process of the encoder learning. However, the classifier could be misled by noisy labels. To overcome this problem, we impose a structure-preserving constraint [19, 20] on pairs of features from $g(\cdot)$ and the corresponding ones from $h(\cdot)$, as illustrated in Figure 4. The structure similarity loss is defined as follows.




Figure 4: The corresponding outputs from g and h should have similar relational structures between pairs of points. Left: transformation $\tilde{y} \rightarrow v$; Right: The relationship between transformed points is similar to the previous.

First, we convert the Euclidean distances $d(v^{(i)}, v^{(j)}), d(\tilde{y}^{(i)}, \tilde{y}^{(j)})$ into the similarity metrics $p(d(v^{(i)}, v^{(j)})), p(d(\tilde{y}^{(i)}, \tilde{y}^{(j)})) \in [0, 1]$ that satisfy $\lim_{d \rightarrow +\infty} p(d) = 0$ and $\lim_{d \rightarrow 0} p(d) = 1$. The similarity metric is formulated as:

$$p(d) = C_\sigma \frac{1}{\sigma \sqrt{2\pi}} e^{-\frac{1}{2} \left(\frac{d-\mu}{\sigma}\right)^2} \quad (11)$$

where $C_\sigma = \sigma \sqrt{2\pi}$ is a constant normalizing the similarity metric in $[0, 1]$, and we set $\mu = 0, \sigma = 0.5$ empirically.

Then the structural similarity loss is defined as the KL-divergence between similarity metrics converted from the projection z and logits \tilde{y} :

$$\mathcal{L}_{str} = \sum_{i \neq j} p(d(v^{(i)}, v^{(j)})) \log \frac{p(d(v^{(i)}, v^{(j)}))}{p(d(\tilde{y}^{(i)}, \tilde{y}^{(j)}))} \quad (12)$$

The proposed approach is delineated in Algorithm 1.

Algorithm 1 Co-learning

Require: Encoder f with θ_1 , classifier g with θ_2 , projection head h with θ_3 , learning rate η , batch size B , epoch T_{\max}

- 1: **for** $t = 1, 2, \dots, T_{\max}$ **do**
- 2: Shuffle training set $\widehat{\mathcal{D}}$
- 3: **for** $n = 1, 2, \dots, \frac{|\widehat{\mathcal{D}}|}{B}$ **do**
- 4: Fetch n -th mini-batch $\widehat{\mathcal{D}}_n$ from $\widehat{\mathcal{D}}$
- 5: **for** $i = 1, 2, \dots, |B|$ **do**
- 6: Weak augment $\tilde{x}_1^{(i)} = t'(x^{(i)})$
- 7: Strong augment $\tilde{x}_2^{(i)} = t_1(x^{(i)})$, $\tilde{x}_3^{(i)} = t_2(x^{(i)})$
- 8: Obtain representations $u_k = f(\tilde{x}_k^{(i)}; \theta_1)$, $k \in \{1, 2, 3\}$
- 9: Obtain prediction $\tilde{y} = g(u_1; \theta_2)$
- 10: Obtain projections $v_k = h(u_k; \theta_3)$, $k \in \{2, 3\}$
- 11: **end for**
- 12: Calculate the loss \mathcal{L} by Eq. (4)
- 13: Update $\Theta = \Theta - \eta \nabla \mathcal{L}$
- 14: **end for**
- 15: **end for**
- 16: **return** $\theta_1, \theta_2, \theta_3$

4 EXPERIMENTS

4.1 Experimental settings

Datasets. Experiments are conducted on four benchmark datasets: CIFAR-10, CIFAR-100, Animal-10N [32], and Food-101N [16] which are popularly used for evaluating models in learning from noisy labels [1, 31].

Since CIFAR-10 and CIFAR-100 are clean datasets with true labels, we need to corrupt the original labels manually. Similar to previous works [29, 30, 33], we consider two types of synthetic label noise: *symmetric* noise and *asymmetric* noise. Following [36, 44], label noise is uniformly distributed among all categories with probability $\frac{p}{|C|}$, $p \in [0, 1]$ in the symmetric noise setup. In the asymmetric noise setup, labels are generated by flipping each class to a similar class or the next class circularly with probability p . We corrupt CIFAR-10 and CIFAR-100 datasets by the label transition matrix Q , where $Q_{ij} = \Pr[\hat{y} = j | y = i]$ given that noisy \hat{y} is flipped from clean y . Figure 6 shows an example of noise transition matrix on CIFAR-100.

Following previous works [37, 43], we evaluate our method on various noise rates η , with $\eta \in \{20\%, 50\%, 80\%\}$ for symmetric label noise and $\eta \in \{20\%, 30\%, 40\%\}$ for asymmetric label noise. We use no more than 40% because certain classes become theoretically indistinguishable for asymmetric noise larger than 50%. As for the asymmetric noise of CIFAR-10, following [29] the noisy labels were generated by mapping "truck" \rightarrow "automobile", "bird" \rightarrow "airplane", "deer" \rightarrow "horse" and "cat" \rightarrow "dog" with probability p . But for the asymmetric noise of CIFAR-100, we generate labels by flipping each class to the next class with probability p .

Besides, Animal-10N [32] is a real-world noisy dataset of human-labeled online images for ten confusing animals, with 50,000 training and 5,000 testing images. As people may label images with confusing patterns, this dataset has natural noisy data. Its noise rate is estimated at 8%. Food-101N is a webly noisy dataset that [16] collects 310k images from Google, Bing, Yelp, and TripAdvisor using the Food-101 [3] taxonomy. We train models on Food-101N and evaluate them on the test set of Food-101.

Baselines. We compare our Co-learning with the following state-of-the-art methods, and we implement all methods with their default settings by PyTorch based on their open-source code.

- (1) Standard cross-entropy, which refers to train on noisy datasets directly with standard cross-entropy loss. This method is applied as a simple baseline.
- (2) Decoupling [23], which updates the parameters only using samples that have different predictions from two classifiers.
- (3) Co-teaching [11], which trains two networks simultaneously and cross-updates parameters of their peer networks.
- (4) Co-teaching+ [39], which trains two networks simultaneously and cross-updates their peer networks with the samples with different predictions from two classifiers.
- (5) JoCoR [37], which trains two networks jointly and makes predictions from two classifiers similar by a Jensen-Shannon divergence.
- (6) APL [21], which proposes a framework that combines two robust loss functions that mutually boost each other.

Network Structure and Optimizer. We utilize ResNet-18 for CIFAR-10 and CIFAR-100 datasets. For Animal-10N and Food-101N datasets, we utilize ResNet-34.

For experiments on CIFAR-10 and CIFAR-100, Adam optimizer is applied with an initial learning rate of 0.001 and a momentum of 0.9. The batch size is fixed as 128. For each trial, we run 200 epochs, and the learning rate is linearly decayed to zero from 80 to 200 epochs. All the experiments are repeated five times with different random seeds, and we report the average test accuracy of five trials for each experiment. For experiments on the real-world dataset Animal-10N and Food-101N, Adam optimizer is also applied with an initial learning rate of 0.001 and a momentum of 0.9. The batch size is set to 64. We run each trial for 100 epochs.

Measurement. To evaluate the performance, the test accuracy is used, i.e., test accuracy = (# of correct predictions) / (# of the test). As the experiments of the synthetic noisy datasets are repeated five times, the MEAN and STD of the test accuracy are calculated. We have been highlighted the error bar for STD as a shade in each figure.

4.2 Comparison with the State-of-the-Arts

Results on CIFAR-10. Figure 5 shows the test accuracy on CIFAR-10. Co-learning performs best in all six cases. In these figures, test accuracy of the standard cross-entropy first reaches a very high level and then decreases gradually, which indicates the memorization effects of neural networks. The higher the noise rates, the more serious the problems become. A robust training method toward learning with noisy labels should avoid overfitting on noisy training data.




Figure 5: Results on CIFAR-10 dataset.

Flipping-Rate	Standard	Decoupling	Co-teaching	Co-teaching+	JoCoR	APL	Co-learning
Symmetric-20%	84.81±0.24	85.75±0.31	90.29±0.19	88.63±0.32	90.43±0.25	88.54±0.45	92.21±0.31
Symmetric-50%	61.49±0.58	61.93±0.82	63.45±3.89	76.27±2.80	66.00±0.53	76.51±1.73	84.49±0.34
Symmetric-80%	28.98±0.26	27.23±0.84	28.03±1.67	30.37±1.69	29.19±1.64	24.75±2.87	61.20±2.29
Asymmetric-20%	87.00±0.20	87.66±0.29	89.38±0.33	89.00±0.18	89.20±0.26	88.02±0.29	91.07±0.32
Asymmetric-30%	81.99±0.31	81.83±0.26	86.58±1.32	86.22±0.26	86.41±0.45	86.03±0.21	86.89±0.87
Asymmetric-40%	76.30±0.34	74.97±0.38	74.25±0.38	81.25±0.75	73.95±1.00	80.97±0.19	81.42±0.52

Table 2: Average test accuracy (%) on CIFAR-10 over the last 10 epochs.




Figure 6: Example of noise transition matrices (taking 5 classes and noise ratio 0.4 as an example)

The test accuracy of different algorithms is compared and reported in Table 2. In the easiest Symmetry-20% case, all approaches work well except standard cross-entropy, demonstrating their robustness to low-level label noise. When it comes to higher noisy rates like Symmetry-50% and Asymmetry-40%, Decoupling starts to overfit to noisy data as standard cross-entropy does. In the hardest Symmetry-80% case, Co-learning significantly outperforms other approaches and suffers less from noisy training data. Co-learning performs better than others in high noisy rates due to the feature-dependent information that comes from self-supervision.

Results on CIFAR-100. Table 3 shows results on CIFAR-100. As there are only 10 classes in CIFAR-10 but 100 classes in CIFAR-100, the overall accuracy is much lower than the previous. The memorization effects of neural networks are severe in standard cross-entropy as well as other approaches like Decoupling and Co-teaching+. As we can see, the observations are consistently the same as those for the CIFAR-10 dataset, which is, Co-learning still works much better than all other approaches. In the easiest Symmetry-20% case, only standard cross-entropy and Decoupling fail to combat noisy training data while others perform well. However, Co-learning works significantly better than other methods under high noisy rate cases such as Symmetry-50% and Symmetry-80%.

Figure 7 shows the test accuracy during the training process with different noise types. Co-learning performs more stable learning curves with much less overfitting than other methods.

Results on Animal-10N. Table 4 summarizes the test results on the real-world noisy labels using the Animal-10N dataset, where *best* denotes the optimal validation accuracy, and *last* denotes the final test accuracy. As Co-teaching, Co-teaching+, and JoCoR require a noisy rate estimation to assist their training, we roughly set it as 10%. Despite training with less prior knowledge, Co-learning consistently outperforms other methods.

Results on Food-101N. Table 5 illustrates the test results on Food-101N dataset, where *best* denotes the optimal validation accuracy,




Figure 7: Results on CIFAR-100 dataset.

Flipping-Rate	Standard	Decoupling	Co-teaching	Co-teaching+	JoCoR	APL	Co-learning
Symmetric-20%	57.79±0.44	56.18±0.32	64.28±0.32	55.40±0.71	62.29±0.71	59.21±0.50	66.58±0.15
Symmetric-50%	33.75±0.46	31.58±0.54	32.62±0.51	26.49±0.45	30.19±0.60	43.53±1.84	54.54±0.43
Symmetric-80%	8.64±0.22	7.71±0.23	6.65±0.71	8.57±1.55	6.84±0.92	13.97±0.53	35.45±0.79
Asymmetric-20%	59.36±0.36	57.97±0.24	59.76±0.53	56.11±0.60	58.58±0.51	58.89±0.40	65.26±0.76
Asymmetric-30%	51.06±0.44	49.86±0.54	49.53±0.79	47.12±0.73	49.04±0.91	51.46±0.15	56.97±1.22
Asymmetric-40%	42.49±0.23	41.51±0.67	40.62±0.79	38.98±0.54	39.72±0.76	41.96±0.92	47.62±0.79

Table 3: Average test accuracy (%) on CIFAR-100 over the last 10 epochs.

Methods	best	last
Standard CE	82.68	81.10
Decoupling	79.22	78.24
Co-teaching	82.43	81.52
Co-teaching+	50.66	48.52
JoCoR	82.82	81.06
Co-learning	82.95	82.18

Table 4: Classification accuracy (%) on Animal-10N dataset.

and *last* denotes the final test accuracy. Co-learning still outperforms other methods on the webly real-world noisy dataset, which demonstrates its robustness.

Methods	best	last
Standard CE	84.50	83.86
Decoupling	85.53	85.28
Co-teaching	61.91	61.86
Co-teaching+	81.61	81.24
JoCoR	77.94	77.86
Co-learning	87.57	86.56

Table 5: Classification accuracy (%) on Food-101N dataset.

4.3 Ablation study

To study the effect of our proposed components, the ablation study is conducted on CIFAR-10 with Symmetry-20% and Symmetry-50% noise, as shown in Figure 8.

Standard cross-entropy suffers from overfitting on noisy training data. Simply adding a self-supervision signal helps to a small extent but is not good enough. The main issue is the slow convergence of self-supervised learning, making the feature encoder f quickly overfits label-dependent information but ignores feature-dependent information.

In order to slow down the inaccurate convergence of supervised learning, a data augmentation method, MixUp, is utilized to avoid early overfitting. As we can see, a balanced training that prompts both label-dependent information and feature-dependent information play crucial roles during the training procedure. Though Co-learning without structural similarity loss can obtain a good performance under the low noisy rate Symmetry-20%, it still overfits noisy training data under the high noisy rate Symmetry-50%. Co-learning utilizes the structural similarity between different pre-text tasks to regularize supervised learning and keeps a robust noise-tolerant training procedure. These results show that self-supervised learning can help combat noisy labels, and a balanced training manner further regularizes the model from overfitting matters. Our proposed Co-learning is a more efficient paradigm to take advantage of feature-dependent information.

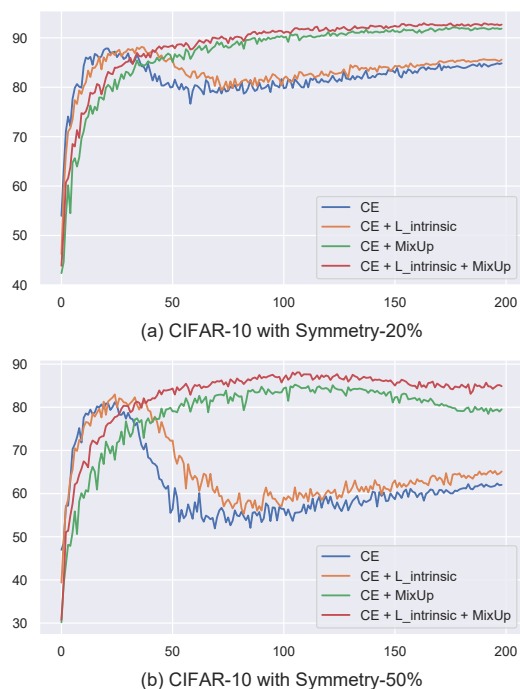


Figure 8: Results of ablation study on CIFAR-10 with Symmetry-20% and Symmetry-50% noise.

4.4 Discussion

Though these co-training-based methods for combating noisy labels are based on the assumption that different views of data should provide extra valuable information, networks with the same architecture provide limited variety caused by random initialization. In general, Decoupling firstly leverages the co-training method with a cross-update technique based on disagreements from two different views. Inspired by the small-loss trick from MentorNet, Co-teaching cross-updates samples selected by the small-loss. Co-teaching+ combines disagreement strategy and the small-loss trick, which means add a disagreement selection step in Co-teaching. JoCoR simplifies the above ideas and jointly updates two networks still based on the small-loss trick and an extra Jensen-Shannon divergence between predictions from two peer networks.

Co-learning further predigests these co-training methods through a shared feature encoder with two exclusive heads, which provides different views of data. Besides, Co-learning explores a reasonable scheme that leverages unsupervised learning to assist robust deep learning.

5 CONCLUSION

In this paper, we propose Co-learning for combating noisy labels by utilizing both label-dependent information from supervised learning and feature-dependent information from self-supervised learning. Instead of training two classifiers, Co-learning trains two heads with a shared feature encoder and regularizes the model

with both the intrinsic similarity and the structural similarity. Co-learning achieves robust training without trivial tricks, while other co-training-based methods need the small-loss trick and noisy rate estimation. Through extensive experiments on synthetic and real-world noisy datasets, we demonstrate that Co-learning outperforms other co-training-based methods and provides a new perspective to leverage self-supervised learning to assist robust deep learning. In the future, we will further explore the theoretical foundation and generalization analysis of Co-learning.

REFERENCES

- [1] Eric Arazo, Diego Ortego, Paul Albert, Noel O'Connor, and Kevin Mcguinness. 2019. Unsupervised Label Noise Modeling and Loss Correction. In *Proceedings of the 36th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 97)*, Kamalika Chaudhuri and Ruslan Salakhutdinov (Eds.). PMLR, Long Beach, 312–321. <http://proceedings.mlr.press/v97/arazo19a.html>
- [2] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A Raffel. 2019. MixMatch: A Holistic Approach to Semi-Supervised Learning. In *Advances in Neural Information Processing Systems*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Eds.), Vol. 32. Curran Associates, Inc., Vancouver. <https://proceedings.neurips.cc/paper/2019/file/1cd138d0499a68f4bb72bee04bbec2d7-Paper.pdf>
- [3] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. 2014. Food-101 – Mining Discriminative Components with Random Forests. In *Computer Vision – ECCV 2014*, David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars (Eds.). Springer International Publishing, Cham, 446–461.
- [4] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. 2020. Unsupervised learning of visual features by contrasting cluster assignments. *Advances in Neural Information Processing Systems* 33 (2020).
- [5] Pengfei Chen, Ben Ben Liao, Guangyong Chen, and Shengyu Zhang. 2019. Understanding and Utilizing Deep Neural Networks Trained with Noisy Labels. In *Proceedings of the 36th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 97)*, Kamalika Chaudhuri and Ruslan Salakhutdinov (Eds.). PMLR, Long Beach, 1062–1070. <http://proceedings.mlr.press/v97/chen19g.html>
- [6] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A Simple Framework for Contrastive Learning of Visual Representations. In *Proceedings of the 37th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 119)*, Hal Daumé III and Aarti Singh (Eds.). PMLR, Vienna, 1597–1607. <http://proceedings.mlr.press/v119/chen20j.html>
- [7] Xinlei Chen and Kaiming He. 2021. Exploring Simple Siamese Representation Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, Virtual, 15750–15758.
- [8] Yifan Ding, Liqiang Wang, Deliang Fan, and Boqing Gong. 2018. A Semi-Supervised Two-Stage Approach to Learning from Noisy Labels. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*. IEEE, Lake Tahoe, 1215–1224. <https://doi.org/10.1109/WACV.2018.00138>
- [9] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. 2018. Unsupervised Representation Learning by Predicting Image Rotations. In *International Conference on Learning Representations*. ICLR, Vancouver, 1–16. <https://openreview.net/forum?id=S1v4N2l0>
- [10] Jean-Bastien Grill, Florian Strub, Florent Alché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, Bilal Piot, koray kavukcuoglu, Remi Munos, and Michal Valko. 2020. Bootstrap Your Own Latent - A New Approach to Self-Supervised Learning. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., Virtual, 21271–21284. <https://proceedings.neurips.cc/paper/2020/file/f3ada80d5c4ee70142b17b8192b2958e-Paper.pdf>
- [11] Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor Tsang, and Masashi Sugiyama. 2018. Co-teaching: Robust training of deep neural networks with extremely noisy labels. In *Advances in Neural Information Processing Systems*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Eds.), Vol. 31. Curran Associates, Inc., Montreal. <https://proceedings.neurips.cc/paper/2018/file/a19744e268754fb0148b017647355b7b-Paper.pdf>
- [12] Jiangfan Han, Ping Luo, and Xiaogang Wang. 2019. Deep Self-Learning From Noisy Labels. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. IEEE, Seoul, 5138–5147.
- [13] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. 2020. Momentum Contrast for Unsupervised Visual Representation Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, Virtual, 9729–9738.
- [14] Dan Hendrycks, Mantas Mazeika, Duncan Wilson, and Kevin Gimpel. 2018. Using Trusted Data to Train Deep Networks on Labels Corrupted by Severe Noise.

- In *Advances in Neural Information Processing Systems*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Eds.), Vol. 31. Curran Associates, Inc., Montreal. <https://proceedings.neurips.cc/paper/2018/file/ad554d8c3b06d6b97ee76a2448bd7913-Paper.pdf>
- [15] Lu Jiang, Zhengyuan Zhou, Thomas Leung, Li-Jia Li, and Li Fei-Fei. 2018. MentorNet: Learning Data-Driven Curriculum for Very Deep Neural Networks on Corrupted Labels. In *Proceedings of the 35th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 80)*, Jennifer Dy and Andreas Krause (Eds.). PMLR, Stockholm, 2304–2313. <http://proceedings.mlr.press/v80/jiang18c.html>
- [16] Kuang-Huei Lee, Xiaodong He, Lei Zhang, and Linjun Yang. 2018. CleanNet: Transfer Learning for Scalable Image Classifier Training With Label Noise. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, Salt Lake City, 5447–5456.
- [17] Junnan Li, Richard Socher, and Steven C.H. Hoi. 2020. DivideMix: Learning with Noisy Labels as Semi-supervised Learning. In *International Conference on Learning Representations*. ICLR, New Orleans, 1–14. <https://openreview.net/forum?id=HJgExaVtwr>
- [18] Junnan Li, Yongkang Wong, Qi Zhao, and Mohan S. Kankanhalli. 2019. Learning to Learn From Noisy Labeled Data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, Long Beach, 5051–5059.
- [19] Stan Z Li, Zelin Zang, and Lirong Wu. 2020. Markov-Lipschitz deep learning. *arXiv preprint arXiv:2006.08256* (2020), 1–10.
- [20] Stan Z Li, Zelin Zang, and Lirong Wu. 2020. Deep Manifold Transformation for Dimension Reduction. *arXiv preprint arXiv:2010.14831* (2020), 1–10.
- [21] Xingjun Ma, Hanxun Huang, Yisen Wang, Simone Romano, Sarah Erfani, and James Bailey. 2020. Normalized Loss Functions for Deep Learning with Noisy Labels. In *Proceedings of the 37th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 119)*, Hal Daumé III and Aarti Singh (Eds.). PMLR, Vienna, 6543–6553. <http://proceedings.mlr.press/v119/ma20c.html>
- [22] Dhruv Mahajan, Ross Girshick, Vignesh Ramanathan, Kaiming He, Manohar Paluri, Yixuan Li, Ashwin Bharambe, and Laurens van der Maaten. 2018. Exploring the Limits of Weakly Supervised Pretraining. In *Proceedings of the European Conference on Computer Vision (ECCV)*. IEEE, Munich, 181–196.
- [23] Eran Malach and Shai Shalev-Shwartz. 2017. Decoupling “when to update” from “how to update”. In *Advances in Neural Information Processing Systems*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.), Vol. 30. Curran Associates, Inc., Long Beach, 960–970. <https://proceedings.neurips.cc/paper/2017/file/58d4d1e7b1e97b258c9ed0b37e02d087-Paper.pdf>
- [24] Devraj Mandal, Shrisha Bharadwaj, and Soma Biswas. 2020. A Novel Self-Supervised Re-labeling Approach for Training with Noisy Labels. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*. IEEE, Snowmass Village, 1381–1390.
- [25] Duc Tam Nguyen, Chaithanya Kumar Mummadi, Thi Phuong Nhung Ngo, Thi Hoai Phuong Nguyen, Laura Beggel, and Thomas Brox. 2020. SELF: Learning to Filter Noisy Labels with Self-Ensembling. In *International Conference on Learning Representations*. ICLR, Virtual, 1–15. <https://openreview.net/forum?id=HkgsPhNYPS>
- [26] Mehdi Noroozi and Paolo Favaro. 2016. Unsupervised Learning of Visual Representations by Solving Jigsaw Puzzles. In *Computer Vision – ECCV 2016*, Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling (Eds.). Springer International Publishing, Cham, 69–84.
- [27] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748* 1807, 03748 (2018), 1–13.
- [28] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A. Efros. 2016. Context Encoders: Feature Learning by Inpainting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, Caesars Palace, 2536–2544.
- [29] Giorgio Patrini, Alessandro Rozza, Aditya Krishna Menon, Richard Nock, and Lizhen Qu. 2017. Making deep neural networks robust to label noise: A loss correction approach. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, Hawaii, 1944–1952.
- [30] Scott E Reed, Honglak Lee, Dragomir Anguelov, Christian Szegedy, Dumitru Erhan, and Andrew Rabinovich. 2015. Training Deep Neural Networks on Noisy Labels with Bootstrapping. In *ICLR (Workshop)*. ICLR, San Diego, 1–11.
- [31] Jun Shu, Qi Xie, Lixuan Yi, Qian Zhao, Sanping Zhou, Zongben Xu, and Deyu Meng. 2019. Meta-Weight-Net: Learning an Explicit Mapping For Sample Weighting. In *Advances in Neural Information Processing Systems*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett (Eds.), Vol. 32. Curran Associates, Inc., Vancouver, 1919–1930. <https://proceedings.neurips.cc/paper/2019/file/e58cc5ca94270acaced13bc82dfed7-Paper.pdf>
- [32] Hwanjun Song, Minseok Kim, and Jae-Gil Lee. 2019. SELFIE: Refurbishing Unclean Samples for Robust Deep Learning. In *Proceedings of the 36th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 97)*, Kamalika Chaudhuri and Ruslan Salakhutdinov (Eds.). PMLR, Long Beach, 5907–5915. <http://proceedings.mlr.press/v97/song19b.html>
- [33] Daiki Tanaka, Daiki Ikami, Toshihiko Yamasaki, and Kiyoharu Aizawa. 2018. Joint optimization framework for learning with noisy labels. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, Salt Lake City, 5552–5560.
- [34] Ryutaro Tanno, Ardavan Saeedi, Swami Sankaranarayanan, Daniel C Alexander, and Nathan Silberman. 2019. Learning from noisy labels by regularized estimation of annotator confusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, Long Beach, 11244–11253.
- [35] Bart Thomee, David A Shamma, Gerald Friedland, Benjamin Elizalde, Karl Ni, Douglas Poland, Damian Borth, and Li-Jia Li. 2016. YFCC100M: The new data in multimedia research. *Commun. ACM* 59, 2 (2016), 64–73.
- [36] Yisen Wang, Xingjun Ma, Zaiyi Chen, Yuan Luo, Jinfeng Yi, and James Bailey. 2019. Symmetric cross entropy for robust learning with noisy labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, Long Beach, 322–330.
- [37] Hongxin Wei, Lei Feng, Xiangyu Chen, and Bo An. 2020. Combating noisy labels by agreement: A joint training method with co-regularization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, Washington, 13726–13735.
- [38] Yan Yan, Zhongwen Xu, Ivor Tsang, Guodong Long, and Yi Yang. 2016. Robust semi-supervised learning through label aggregation. In *Proceedings of the AAAI Conference on Artificial Intelligence*. AAAI Press, Phoenix, 1–7.
- [39] Xingrui Yu, Bo Han, Jiangchao Yao, Gang Niu, Ivor Tsang, and Masashi Sugiyama. 2019. How does Disagreement Help Generalization against Label Corruption?. In *Proceedings of the 36th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 97)*, Kamalika Chaudhuri and Ruslan Salakhutdinov (Eds.). PMLR, Long Beach, 7164–7173. <http://proceedings.mlr.press/v97/yy19b.html>
- [40] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. 2021. Understanding deep learning (still) requires rethinking generalization. *Commun. ACM* 64, 3 (2021), 107–115.
- [41] Hongyi Zhang, Moustapha Cisse, Yann N. Dauphin, and David Lopez-Paz. 2018. mixup: Beyond Empirical Risk Minimization. In *International Conference on Learning Representations*. ICLR, Vancouver, 1–13. <https://openreview.net/forum?id=riDdp1-Rb>
- [42] Richard Zhang, Phillip Isola, and Alexei A. Efros. 2016. Colorful Image Colorization. In *Computer Vision – ECCV 2016*, Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling (Eds.). Springer International Publishing, Cham, 649–666.
- [43] Yikai Zhang, Songzhu Zheng, Pengxiang Wu, Mayank Goswami, and Chao Chen. 2021. Learning with Feature-Dependent Label Noise: A Progressive Approach. In *International Conference on Learning Representations*. ICLR, Virtual, 1–13. <https://openreview.net/forum?id=ZPa2SyGcbwh>
- [44] Zhilu Zhang and Mert R Sabuncu. 2018. Generalized cross entropy loss for training deep neural networks with noisy labels. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*. Curran Associates, Inc., Montreal, 8792–8802.