

Estimating Noisy Class Posterior with Part-level Labels for Noisy Label Learning

Rui Zhao¹ Bin Shi^{1*} Jianfei Ruan¹ Tianze Pan² Bo Dong³

¹School of Computer Science and Technology, Xi'an Jiaotong University, China

²School of Physics, Xi'an Jiaotong University, China

³School of Distance Education, Xi'an Jiaotong University, China

rayn.z@stu.xjtu.edu.cn, shibin@xjtu.edu.cn

jianfei.ruan@hotmail.com, terrypztz@stu.xjtu.edu.cn, dong.bo@xjtu.edu.cn

Abstract

In noisy label learning, estimating noisy class posteriors plays a fundamental role for developing consistent classifiers, as it forms the basis for estimating clean class posteriors and the transition matrix. Existing methods typically learn noisy class posteriors by training a classification model with noisy labels. However, when labels are incorrect, these models may be misled to overemphasize the feature parts that do not reflect the instance characteristics, resulting in significant errors in estimating noisy class posteriors. To address this issue, this paper proposes to augment the supervised information with part-level labels, encouraging the model to focus on and integrate richer information from various parts. Specifically, our method first partitions features into distinct parts by cropping instances, yielding part-level labels associated with these various parts. Subsequently, we introduce a novel single-to-multiple transition matrix to model the relationship between the noisy and part-level labels, which incorporates part-level labels into a classifier-consistent framework. Utilizing this framework with part-level labels, we can learn the noisy class posteriors more precisely by guiding the model to integrate information from various parts, ultimately improving the classification performance. Our method is theoretically sound, while experiments show that it is empirically effective in synthetic and real-world noisy benchmarks.

1. Introduction

The increase in model capacity has enabled deep artificial neural networks to fit almost any data, which unfortunately makes them prone to memorizing even the mislabeled instances [45, 54]. Thus, to train an effective neural network, a large amount of accurately labeled data is required. However, obtaining accurate labeling in real-world tasks

typically involves manual labeling, which is often time-consuming and costly [46]. In contrast, massive amounts of noisy labels are readily available through web crawlers, questionnaires and crowdsourcing [22].

To reduce the negative impact of noisy labels, various heuristic strategies are applied to Noisy Label Learning (NLL), such as selecting reliable samples [6, 15, 17, 22, 43, 48, 53] and correcting labels [23, 40, 44]. Although these methods can train classifiers that perform well empirically, they do not guarantee a *consistent classifier*, meaning that the classifiers learned from noise-labeled data will not necessarily converge to the optimal one learned from clean data [46, 54].

To overcome this problem, various *classifier-consistent* methods have been proposed, with the most successful methods employing loss correction procedures [3, 13, 27, 37, 46, 47, 49, 52, 54]. The basic idea of these methods is that, the *clean class posterior* $P(\mathbf{Y}|\mathbf{x})$ can be inferred from *noisy class posterior* $P(\tilde{\mathbf{Y}}|\mathbf{x})$ and *transition matrix* $T(\mathbf{x})$, from the equation $P(\tilde{\mathbf{Y}}|\mathbf{x}) = T(\mathbf{x})^\top P(\mathbf{Y}|\mathbf{x})$ ¹. The noisy class posteriors estimation plays a fundamental role in this process, directly affecting the calculation of clean class posteriors while also being critical for the accurate transition matrix estimation [4].

Considering that the ground-truth noisy class posterior is typically unavailable, existing methods often learn it through the classification tasks supervised by noisy labels [3, 13, 27, 37, 46, 47, 49, 52, 54]. In classification, labels will encourage the model to focus on label-related feature part to minimize the classification loss. However, when labels are incorrect, models will be misled into overemphasizing the erroneous parts that do not reflect the instance characteristics. As illustrated in Figure 1a, for an instance of a "feathered monkey" labeled as "bird", the noisy la-

¹We define $P(\mathbf{Y}|\mathbf{x}) = [P(Y = 1|X = \mathbf{x}), \dots, P(Y = c|X = \mathbf{x})]^\top$ and $T_{ij}(\mathbf{x}) = P(\tilde{Y} = j|Y = i, X = \mathbf{x})$, where c denotes the number of classes, X and Y represents the variable for clean labels and instances, respectively.

*Corresponding author




Figure 1. Illustration of overemphasis that arises when learning noisy class posterior with classification loss (Figure 1a), as well as the framework we proposed to alleviate this overemphasis (Figure 1b). Class activation maps are used to visualize feature importance for estimation, where the highlighted areas (with stronger red intensity) represent the focus regions of the model.

bel "bird" will force the backbone network to focus on features related to birds, such as feathers, while ignoring other important features, such as monkey's facial features, to achieve a smaller loss through backpropagation. In such cases, the model's output tends to reflect the probability of 'birds' given "feathers" features rather than "birds" given this "feathered monkey" instance. This introduces significant errors in estimating the noisy class posteriors. To address this issue, a straightforward idea is to augment supervised information with part-related labels, thus encouraging the model to focus on disparate feature parts. This can enhance the model's ability to capture instance characteristics and aid in estimating noisy class posteriors. This idea is consistent with object recognition research, which reveals that feature parts are used by humans and machines to identify objects [47], drawing on evidence from psychological and physiological studies [30, 36, 42], as well as computational theories and learning algorithms [1, 2, 7, 14, 35, 41].

Motivated by this idea, we propose a Part-Level Multi-labeling (PLM) method, which generates multiple part-level labels to guide the model's focus. As illustrated in Figure 1b, by augmenting the supervised information, we attempt to capture distinct features of various parts, thereby rectifying the model's excessive focus on specific misleading parts. Specifically, to generate part-level labels, we design a multi-labeling approach based on instance crop-

ping. As displayed in the second row of Figure 1b, this approach partitions features into distinct parts by instance cropping, such that some of the parts do not contain excessively focused features. By means of this partitioning, we can employ a network trained on noisy data to assign labels to these parts, resulting in multiple part-level labels associated with various diverse focused regions. Subsequently, a single-to-multiple transition matrix is utilized to model the relationship between the single noisy label and the multiple part-level labels for each instance. With this matrix, we proposed a label joint training framework that incorporates both part-level and noisy supervised information into a classifier-consistent framework. As shown in the first row of Figure 1b, this framework guides the model to focus on more diverse regions, thereby obtaining representations that reflect instance characteristics, ultimately enhancing the learning of noisy class posteriors. Then, it can be used to estimate other critical metrics such as transition matrix or serve as a component of loss correction for NLL. We extensively evaluated our proposed method on synthetic datasets with instances-independent and -dependent noise, as well as on real-world datasets, offering empirical evidence of its efficacy in reducing errors in noisy class posterior estimation and improving performance on NLL tasks.

Our main contributions are summarized as follows:

- We focus on a novel problem of estimating noisy class

posteriors in noisy label learning, which forms the basis for building classifier-consistent algorithms.

- We are the first to note the misleading effect of incorrect labels on noisy class posteriors, where incorrect labels guide the model to overly focus on feature details that do not reflect instance characteristics. To counter this, we propose to incorporate part-related supervision by a consistent classifier, which guides the model to integrate information from various parts.
- Extensive experiments on a variety of synthetic and real-world noisy datasets have confirmed the effectiveness of proposed method. The method notably enhances the estimation of noisy class posteriors and can be integrated with various NLL methods that rely on noisy class posteriors to boost their classification performance.

2. Related work

The primary goal of NLL is to mitigate the impact of mislabeled data in the training, fostering more robust models. In this context, various data-centric [6, 8, 10, 11, 17, 22–25, 34, 38, 40, 43, 44, 53, 55] and loss-centric [3, 13, 26, 27, 29, 32, 37, 46, 47, 49, 51, 52, 54, 56] approaches have been proposed. The former involves data preprocessing to reduce the involvement of noisy data in training, while the latter constructs robust losses to diminish the strength of the supervised signals generated by noisy data.

The first category of methods often uses heuristic algorithms to reduce the side effects of error data, such as selecting reliable samples [5, 17, 22, 43, 53], correcting labels [18, 23, 40, 44], reweighting samples [38], smoothing labels [31] and noise reduction [45]. However, the use of heuristic algorithms makes these methods prone to overfitting noisy data. To address this issue, the second category of methods aims to introduce robust loss functions to weaken the supervised signals from mislabeled data, avoiding the reliance on heuristic algorithms. The most prominent approach within this category is *loss correction* methods [3, 13, 27, 37, 46, 47, 49, 52, 54]. These methods typically use a transition matrix to correct the loss to guarantee that the trained classifier will converge to the optimal classifier learned from the clean data (i.e., statistically consistent classifier). The critical task of this method is to estimate the transition matrix and thus infer the clean class posterior from the noisy data.

In these methods, the noisy class posterior is often utilized to estimate important parameters, such as calculating losses to obtain low-loss samples and estimating the transition probabilities to obtain a transition matrix, and even inferring clean class posteriors with a consistent classifier. Unfortunately, the estimation of these important metrics are all affected by the estimation errors of noisy class posteriors, which is considered a unresolved bottleneck of classification performance [4, 27]. While some label processing

methods [9, 31, 45] implicitly reduce overemphasis on specific parts by decreasing the model’s fitting to noisy labels, they usually apply the same treatment to all labels, lacking an effective guidance to help integrate information from other parts. Hence, these methods struggle to address the issue where incorrect labels mislead the network to focus on features that do not reflect the instance characteristics.

3. Method

3.1. Problem setting

Let D denote the joint probability distribution of a pair of random variables $(X, Y) \in \mathcal{X} \times \mathcal{C}$, where X and Y represent the random variables associated with instances and their corresponding clean labels, respectively. In this context, \mathcal{X} denotes the instance space, while $\mathcal{C} = \{1, \dots, c\}$ represents the label space with c denoting the number of classes. Given a labeled training dataset $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$ with size n , where each example (\mathbf{x}_i, y_i) is drawn independently from the probability distribution D , the classification task aims to learn a classifier $f : \mathcal{X} \rightarrow \mathcal{C}$ that maps each instance \mathbf{x}_i to its corresponding label y_i based on the training data \mathcal{D} . However, obtaining samples from the distribution D in real-world tasks presents a significant challenge, since the observed labels are often corrupted by noise. Let \tilde{Y} be the random variable of noisy labels and \tilde{D} be the distribution of a pair of random variables (X, \tilde{Y}) . The goal of NLL is to learn a robust classifier from noisy sample set $\tilde{\mathcal{D}} = \{(\mathbf{x}_i, \tilde{y}_i)\}_{i=1}^n$ independently drawn from \tilde{D} , which can assign the clean labels for the instances.

3.2. Preliminary: NLL with consistent classifiers

For building a consistent classifier, the mainstream methods using a transition matrix, which can relate the random variables of \tilde{Y} and Y [3, 6, 46], to infer the clean class posteriors from noisy class posteriors. Specifically, the noisy class posterior vector $P(\tilde{Y}|\mathbf{x}) = [P(\tilde{Y} = 1|\mathbf{x}), \dots, P(\tilde{Y} = c|\mathbf{x})]^\top$ can be transformed into the multiplication of the clean class posterior vector $P(Y|\mathbf{x}) = [P(Y = 1|\mathbf{x}), \dots, P(Y = c|\mathbf{x})]^\top$ and the noise transition matrix $T(\mathbf{x}) \in \mathbb{R}^{c \times c}$. Here, $T(\mathbf{x})$ denotes the transition matrix of instance \mathbf{x} , and the element of the i -th row and j -th column of $T(\mathbf{x})$ is defined by $T_{ij}(\mathbf{x}) = P(\tilde{Y} = j|Y = i, \mathbf{x})$. Specifically, according to Eq. (1), there is $P(\tilde{Y}|\mathbf{x}) = T(\mathbf{x})^\top P(Y|\mathbf{x})$. It means that the NLL task can be resolved by estimating the transition matrix $T(\mathbf{x})$ and posterior $P(\tilde{Y}|\mathbf{x})$ [28, 46, 49].

$$\begin{aligned} P(\tilde{Y} = j|\mathbf{x}) &= \sum_{i=1}^c P(\tilde{Y} = j|Y = i, \mathbf{x})P(Y = i|\mathbf{x}) \\ &= \sum_{i=1}^c T_{ij}(\mathbf{x})P(Y = i|\mathbf{x}). \end{aligned} \quad (1)$$

Thus, in the context of constructing a consistent classifier with a transition matrix, the accuracy of clean class posteriors directly relies on the estimation of noisy class posteriors. Besides, the estimation of the transition matrix, in many mainstream methods, also depends on noisy class posteriors [4]. Therefore, the estimation of noisy class posteriors is crucial in building classifier-consistent algorithms. To reduce errors in estimating noisy class posteriors, this paper introduces a training framework with a single-to-multiple transition matrix. It can serve as a component in loss correction methods to build a classifier-consistent NLL algorithm. Section 3.4 provides a detailed discussion on the label joint training framework.

3.3. Part-level multi-labeling

This section provides a detailed explanation of the process of generating part-level labels through instance cropping, as illustrated in the second row of Figure 1b. Specifically, for each instance, we obtain a set of cropped parts by applying K times crop, each capturing different features. Then, a noisy classifier trained on the raw noisy data is employed to assign labels to these parts, resulting in part-level labels that reflect diverse features. Since many parts do not contain the overemphasized features after cropping, the network can attend to other features when assigning labels to them, rather than being restricted by the features that are strongly correlated with the noisy labels. This leads to more informative labels that better reflect the comprehensive information of the instance.

Formally, for an instance \mathbf{x}_i , the goal of instance cropping is to generate a set of sub-instances $\mathcal{S}_i = \{\mathbf{s}_{i1}, \dots, \mathbf{s}_{iK}\}^\top$, where \mathbf{s}_{ij} signifies the j -th part of \mathbf{x}_i . Given a cropping number K , the results of instance cropping are influenced by the chosen cropping operation, which can be guided by different criteria, such as user-specified location, saliency, or attention maps. In the experiments conducted in this paper, we use image data as an example and crop five equally-sized parts uniformly from the four corners and the central position. After the instance cropping, we employ a labeling classifier network f^l trained from noisy data to label each part. Part-level labels are a form of multi-label, and thus can be decomposed into c independent binary labels, each corresponding to one of the possible labels in the label space. Thus, the part-level labels of instance \mathbf{x}_i can be represented as a vector $\mathbf{y}'_i = [y'_{i1}, \dots, y'_{ic}]$, where y'_{ij} denotes the part-level label associated with the j -th class. The value of y'_{ij} can be calculated as follows:

$$y'_{ij} = \begin{cases} 1 & \exists \mathbf{s} \in \mathcal{S}_i, f^l(\mathbf{s}) = j \\ 0 & \forall \mathbf{s} \in \mathcal{S}_i, f^l(\mathbf{s}) \neq j \end{cases}. \quad (2)$$

Through the labeling of distinct parts, the model is compelled to focus on and label each part independently. Then,

these labels can guide the model to focus on and integrate the corresponding parts to avoid misleading.

3.4. Label joint training framework

To leverage the part-level labels for estimating the noise class posterior, we propose a label joint training framework. This framework uses noisy labels and part-level labels jointly to estimate the noise class posterior, based on a single-to-multiple transition matrix that models how a single noisy label relates to multiple part-level labels.

As discussed in the previous section, the part-level labels of an instance can be interpreted as a multi-label. Therefore, the random variable associated with the part-level labels can be represented by a joint distribution (Y'_1, \dots, Y'_c) , where $Y'_j \in \{0, 1\}$ represents the random variable associated with the j -th label component of the part-level labels. Then, similar to the discussion in Section 3.2, the part-level labels can be linked to the noisy labels as follows:

$$P(Y'_j = 1 | X = \mathbf{x}) = \sum_{i=1}^c U_{ij}(\mathbf{x}) P(\tilde{Y} = i | X = \mathbf{x}), \quad (3)$$

where $U_{ij}(\mathbf{x}) = P(Y'_j = 1 | \tilde{Y} = i, X = \mathbf{x})$ is the ij -th entry of matrix $U(\mathbf{x}) \in \mathbb{R}^{c \times c}$, and $U(\mathbf{x})$ denotes the transition matrix of instance \mathbf{x} . In this way, we implemented the joint integration of part-level and noisy labels within a classifier-consistent framework.

Joint training with single-to-multiple matrix. For estimating single-to-multiple transition matrix and noisy class posterior, we construct a single-to-multiple transition matrix estimation network $g^u : \mathcal{X} \rightarrow \mathbb{R}^{c \times c}$ and a noisy class posterior estimation network $g^e : \mathcal{X} \rightarrow \mathbb{R}^c$. We perform a joint training of g^u and g^e by minimizing the following loss function:

$$\mathcal{L}(\mathbf{x}, \tilde{\mathbf{y}}, \mathbf{y}') = \frac{1}{2} (\ell_c(g^e(\mathbf{x}), \tilde{\mathbf{y}}) + \ell_m(g^p(\mathbf{x}), \mathbf{y}')), \quad (4)$$

where network $g^p(\mathbf{x}) = g^u(\mathbf{x})g^e(\mathbf{x})$, ℓ_c and ℓ_m represent a classification loss (e.g., the cross entropy loss) and a multi-label classification loss (e.g., the binary cross-entropy loss), respectively. By minimizing the loss function ℓ_c , the output of network g^e is adjusted to fit the noisy class posterior $\hat{P}(\tilde{\mathbf{Y}} | \mathbf{x})$. Furthermore, by minimizing ℓ_m , $g^u(\mathbf{x})g^e(\mathbf{x})$ will output part-level multi-labels. Therefore, based on Eq. (3), $g^u(\mathbf{x})$ is compelled to model the single-to-multiple transition matrix. As the estimation of g^u becomes more accurate, the noisy class posterior estimation of g^e also improves by minimizing ℓ_m . By employing the joint training approach, we ensure that the noisy class posterior estimation benefits from the supervision of noisy labels and partial-level labels as Eq. (3). The richer part contextual supervision prevents the overemphasis by g^e on specific part, thus avoiding overconfidence. Here, we introduce a novel matrix, and its identifiability is discussed in Appendix B.

Classification with consistent classifier. Beyond enhancing the efficiency of noise class posterior estimation to improve the estimation precision of other crucial metrics, our method can also function as a classifier-consistent component to be incorporated into a consistent classifier. From Eq. (1) and Eq. (3), we can formulate the relationship between the part-level labels and the clean labels as follows:

$$\begin{aligned}
 P(Y'_j = 1|\mathbf{x}) &= \sum_{i=1}^c U_{ij}(\mathbf{x})P(\tilde{Y} = i|\mathbf{x}) \\
 &= \sum_{i=1}^c U_{ij}(\mathbf{x}) \sum_{k=1}^c T_{ki}(\mathbf{x})P(Y = k|\mathbf{x}).
 \end{aligned}
 \tag{5}$$

The estimation of the transition matrix $T_{ki}(\mathbf{x})$ can be undertaken using existing methods. Subsequently, our framework can be exploited to amplify the effectiveness of NLL. The details of classifier training can be found in Appendix A.

4. Experiments

4.1. Experiment setup

Datasets. We evaluated the performance of our proposed model using both synthetic and real-world datasets. The synthetic datasets included MNIST [20], CIFAR-10 [19] and CIFAR-100 [19], which we generated by applying three types of noise: symmetry flipping [37], pair flipping [11] and instance-dependent label noise (IDN) [47]. The MNIST dataset consists of 60,000 training images and 10,000 test images, representing handwritten digits from 0 to 9. The CIFAR-10 and CIFAR-100 datasets consist of 50,000 training images and 10,000 test images, representing 10 and 100 object classes, respectively. We also used Clothing1M [50], a large-scale dataset of 1M real-world clothing images with noisy labels for training and 10k images with clean labels for testing. The dataset also provides 50k and 14k clean labels for training and validation, but we did not use them following the settings of VolMinNet [27]. For model selection, we randomly sampled 10% of the noisy training data as validation sets from each dataset.

Backbone and implementation details. For MNIST, we use a Lenet [21] as backbone network. The Stochastic Gradient Descent (SGD) optimizer is used to train the network, with an initial learning rate 5×10^{-2} , weight decay 10^{-4} and momentum 0.9. For CIFAR-10, we use a ResNet-18 [12] as backbone network. The SGD optimizer is used to train the network, with an initial learning rate 10^{-2} , weight decay 10^{-2} , momentum 0.9. We use the cosine learning rate decay strategy for MNIST and CIFAR-10. For CIFAR-100, we use a ResNet-32 [12] as backbone network. SGD optimizer is used to train the network, with an initial learning rate 5×10^{-2} , weight decay 10^{-3} , momentum 0.9. The learning rate is divided by 100 after the 40-th epoch.

We apply random horizontal flips and random crops of size 32×32 pixels after adding a 4-pixel padding on each side for both CIFAR-10 and CIFAR-100 datasets. For all synthetic datasets, the labeling network and estimation network are both trained for 50 epochs with batch size 128 with same parameters. For clothing1M, we use a ResNet-50 [12] pre-trained on ImageNet as backbone network. SGD optimizer is used to train the network, with an initial learning rate 10^{-2} , weight decay 10^{-3} , momentum 0.9. The learning rate is divided by 10 after the 10-th epoch and the 20-th epoch. We resize the image to 256×256 pixels and apply random crops of size 224×224 . The labeling network and estimation network are both trained 30 epochs with batch size 128, and we undersample the noisy labels for each class in every epoch to balance them like VolMinNet[27]. To generate sub-instances, we select five parts uniformly at the four corners and one central position. The part sizes are 22×22 pixels for MNIST, 25×25 pixels for CIFAR-10 and CIFAR-100, and 179×179 pixels for Clothing1M. Additionally, we discussed and conducted experiments on various instance cropping approach in Appendix D. The code are available at <https://github.com/RyanZhaoIc/PLM.git>.

4.2. Noisy class posterior estimation

We compared PLM with the CE approach in terms of the estimation error of noisy class posteriors across diverse datasets and noise rates. This comparison is pertinent as the previous loss correction methods [27, 37, 46, 49, 54] predominantly utilize the CE loss to estimate the posterior probability of noisy labels. For parity in comparison, we implemented both methods using the same backbone network as Section 4.1. The estimation error is calculated as the l_2 -distance between the ground-truth and the estimated noisy class posteriors of on the pseudo-anchor dataset, where the ground-truth originates from the noise generation process. Specifically, we first train a model with clean labels to select high-confidence samples as pseudo-anchor dataset. Then we assume that a pseudo-anchor sample (\mathbf{x}, i) satisfies $P(Y = i|X = \mathbf{x}) = 1$, thereby the ground-truth noisy class posterior of (\mathbf{x}) can be calculated as $P(\tilde{Y}|X = \mathbf{x}) = P(\tilde{Y}|Y = i, X = \mathbf{x})$. Thus, on the synthetic datasets, noise class posteriors can be obtained through the defined noise flipping probability $P(\tilde{Y}|Y = i, X = \mathbf{x})$. This setting are similar to the Dual-T [54].

We conducted these experiments across noise rates of [0.1, 0.2, 0.3, 0.4, 0.5]. As depicted in Figure 1, our method and the comparison methods demonstrated similar trends in the estimation error of noisy class posterior probabilities. However, the estimation error of PLM is consistently smaller than that of the CE at different noise rates, especially at higher ones. The results indicate that our method can provide estimations that are closer to the true noisy class posterior probabilities. The visualization in Appendix E can

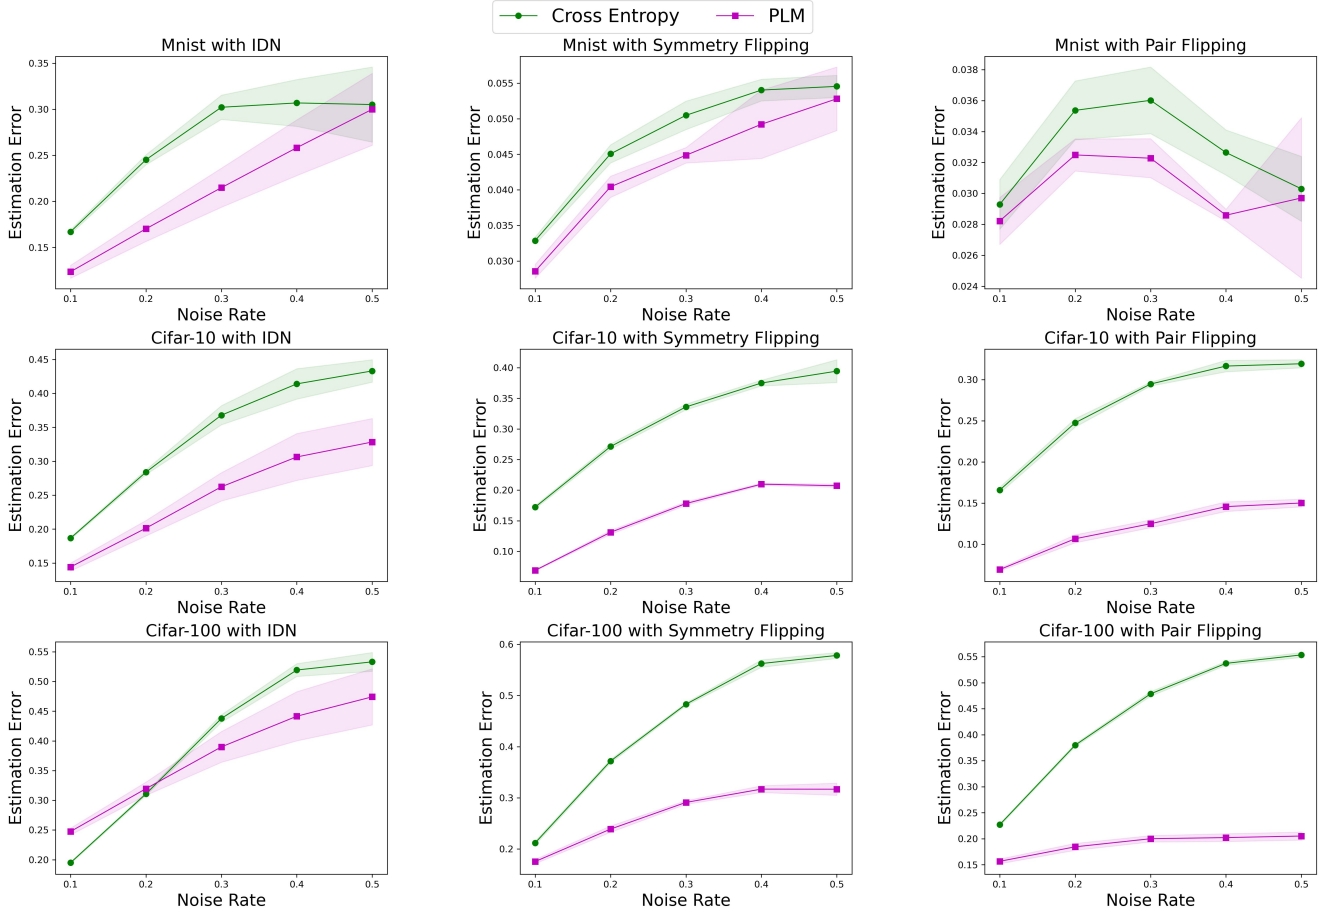


Figure 2. Mean estimation errors of noise class posterior for 5 trials on Mnist, CIFAR-10 and CIFAR-100. The error bars of standard deviation are shaded in each plot. The lower the better.

also serve as evidence for this result.

4.3. Comparison for classification

To compare the accuracy of the NLL, we first estimate the noise transition matrix using the same approach as Forward [37]. Subsequently, we train a single-to-multiple transition matrix estimation network using the method proposed in this paper. Afterward, we fix the transition matrix estimation network and train the classification network with a method like Forward, following the label transition relationship in Eq. (5) and the training objective in Eq. (4). Here, we use the average of the 10 estimated anchors' predicted probabilities as a row in the transition matrix, and train the classification network with the same parameters as the estimation network in Section 4.1. We compare the classification accuracy of PLM with the following baseline methods: (1) Decoupling [33], which trains two predictors and updates them only on examples that they disagree on; (2) Co-teaching [11], which trains two deep networks and exchanges the examples with small loss for network up-

dating; (3) Forward [37], which corrects the training loss by the class flipping transition matrix; (4) T-Revision [46], which estimates transition matrix with a slack variable to correct training loss; (5) Dual-T [54], which estimates transition matrix by factorizing the original transition matrix with an intermediate class; (6) DMI [51], which handles label noise by using information-theoretic loss based on the determinant of the joint distribution matrix; (7) VolMinNet [27], which estimates the matrix that incorporates the minimum volume constraint into the label-noise learning; (8) Class2Simi [45], which transforms noisy class labels into noisy similarity labels to reduce the noise rate. For a fair comparison, we implement all methods with the same backbone network and default parameters on each dataset by the PyTorch and run all experiments on the NVIDIA RTX 3090 GPUs. It should be noted that we did not compare our method with some state-of-the-art methods, such as DivideMix[22] and NCR[16]. These methods incorporate multiple advanced methods (e.g., semi-supervised learning, contrastive learning and complex data augmenta-

Table 1. The average classification accuracy and standard deviation (expressed in percentage) across five trials on the MNIST, CIFAR-10, and CIFAR-100 datasets under various synthetic noisy label settings. The best classification accuracy is indicated in **bold**.

	MNIST		CIFAR-10		CIFAR-100	
	Sym-20%	Sym-50%	Sym-20%	Sym-50%	Sym-20%	Sym-50%
Decoupling	97.58 ± 0.16	95.66 ± 0.17	87.11 ± 0.24	77.62 ± 0.30	69.71 ± 0.34	48.30 ± 0.68
Co-teaching	96.35 ± 0.10	92.01 ± 0.24	78.09 ± 0.91	67.16 ± 0.32	44.89 ± 0.63	32.19 ± 0.90
T-Revision	98.86 ± 0.04	98.41 ± 0.08	89.41 ± 0.18	83.68 ± 1.07	61.14 ± 0.69	40.01 ± 0.90
Dual-T	98.82 ± 0.17	98.24 ± 0.14	89.79 ± 0.40	77.97 ± 1.91	65.61 ± 0.39	50.07 ± 1.56
DMI	98.87 ± 0.09	96.96 ± 0.68	84.65 ± 0.83	70.62 ± 3.73	51.89 ± 2.24	32.02 ± 1.36
VolMinNet	98.82 ± 0.12	98.13 ± 0.15	90.19 ± 0.17	84.09 ± 0.68	67.70 ± 0.99	57.99 ± 0.40
Class2Simi	99.11 ± 0.06	98.33 ± 0.26	84.92 ± 0.50	70.02 ± 3.27	49.78 ± 1.05	32.02 ± 1.33
Forward	98.47 ± 0.19	97.78 ± 0.20	85.49 ± 1.12	74.33 ± 1.31	58.70 ± 0.73	39.82 ± 2.15
PLM	99.32 ± 0.02	98.96 ± 0.07	91.10 ± 0.24	85.08 ± 0.16	69.54 ± 0.23	60.44 ± 0.28
	Pair-20%	Pair-45%	Pair-20%	Pair-45%	Pair-20%	Pair-45%
Decoupling	97.37 ± 0.16	84.01 ± 1.90	88.93 ± 0.38	70.55 ± 2.29	66.65 ± 0.27	43.82 ± 1.75
Co-teaching	95.38 ± 0.29	89.17 ± 0.27	79.18 ± 0.44	68.58 ± 0.68	45.78 ± 1.02	28.98 ± 0.42
T-Revision	99.07 ± 0.10	96.42 ± 4.41	91.16 ± 0.16	81.88 ± 13.30	60.56 ± 0.66	49.44 ± 1.64
Dual-T	98.86 ± 0.09	92.47 ± 5.91	89.59 ± 1.34	74.39 ± 4.30	71.08 ± 0.27	52.25 ± 3.28
DMI	98.81 ± 0.13	96.45 ± 0.54	87.66 ± 0.69	80.30 ± 3.38	43.63 ± 0.86	32.22 ± 3.42
VolMinNet	99.11 ± 0.08	99.10 ± 0.08	91.55 ± 0.13	86.12 ± 1.26	71.65 ± 0.62	61.21 ± 2.98
Class2Simi	98.97 ± 0.12	98.33 ± 0.26	87.27 ± 0.31	75.47 ± 6.09	53.39 ± 0.95	36.04 ± 2.31
Forward	98.82 ± 0.15	91.02 ± 10.81	89.33 ± 1.09	69.99 ± 9.79	60.12 ± 0.33	37.99 ± 0.37
PLM	99.40 ± 0.03	99.39 ± 0.05	92.84 ± 0.17	91.80 ± 0.63	72.67 ± 0.32	61.90 ± 1.82

Table 2. The classification accuracy (expressed in percentage) on the Clothing1M dataset. Comparative baseline method results are quoted from PTD [47] and VolMinNet [27] since we employed the same experimental setup. The best classification accuracy is indicated in **bold**.

Co-teaching	Dual-T	T-Revision	DMI	VolMinNet	PTD	Forward	PLM
56.79	70.97	60.15	70.12	71.67	72.42	69.91	73.30

Table 3. The average classification accuracy and standard deviation (expressed in percentage) across five trials on the CIFAR-10 dataset with IDN. The results of the compared baseline methods are quoted from MEIDTM [3] since we used the same experimental setup. The best classification accuracy is indicated in **bold**.

	IDN-20%	IDN-30%	IDN-40%	IDN-50%
Co-teaching	88.43 ± 0.08	86.40 ± 0.41	80.85 ± 0.97	62.63 ± 1.51
DMI	89.99 ± 0.15	86.87 ± 0.34	80.74 ± 0.44	63.92 ± 3.92
PTD	89.33 ± 0.70	85.33 ± 1.86	80.59 ± 0.41	64.58 ± 2.86
TMDNN	88.14 ± 0.66	84.55 ± 0.48	79.71 ± 0.95	63.33 ± 2.75
MEIDTM	91.38 ± 0.34	87.68 ± 0.26	82.63 ± 0.24	72.17 ± 1.51
Forward	89.62 ± 0.14	86.93 ± 0.15	80.29 ± 0.27	65.91 ± 1.22
PLM	91.41 ± 0.17	88.60 ± 0.53	83.98 ± 2.53	76.87 ± 1.59

tion, etc.) to improve their performance under noisy labels, while PLM is solely focused on reducing the error of estimating noisy class posterior to benefit the NLL, negating the need for additional techniques. Thus, a direct comparison would not present a fair assessment.

In Table 1, we report the classification accuracy of PLM

and all baseline methods, under varying noise generation methods and rates. Here, we use Sym- ϵ and Pair- ϵ to denote the symmetry flipping and pair flipping methods with noise rates of ϵ , respectively. Compared to other baseline methods, PLM shows superior classification accuracy. We executed five experiments on synthetic noisy datasets with distinct random seeds, introducing a higher degree of label randomness. Despite varying data generations, our method demonstrates robust performance. Our method outperformed the comparative methods in all metrics. It is worth noting that although we did not combine PLM with the latest loss correction methods in this study, the PLM method still achieved competitive performance and significantly improved the classification performance of the model under high noise rates. At the same time, our method only exhibits a linear increase in computational time, and we provide relevant discussions in Appendix C. The results indicate that PLM can serve as an effective auxiliary component to improve the robustness of the model to label noise.

Table 4. The average classification accuracy and standard deviation (expressed in percentage) across five trials on the CIFAR-10 and CIFAR-100 datasets. The better classification accuracy is indicated in **bold**.

	CIFAR-10				CIFAR-100			
	Sym-20%	Sym-50%	Pair-20%	Pair-45%	Sym-20%	Sym-50%	Pair-20%	Pair-45%
Forward	85.62 ± 0.58	73.78 ± 0.99	89.20 ± 1.56	66.20 ± 10.05	58.70 ± 0.73	39.82 ± 2.15	60.12 ± 0.33	37.99 ± 0.37
PLM-F	89.45 ± 0.59	81.94 ± 0.62	91.29 ± 0.18	74.90 ± 1.37	68.19 ± 0.92	58.59 ± 0.92	70.94 ± 1.16	55.47 ± 1.27
Dual-T	89.79 ± 0.40	77.97 ± 1.91	89.59 ± 1.34	74.39 ± 4.30	65.62 ± 0.39	50.07 ± 1.56	71.08 ± 0.27	52.25 ± 3.28
PLM-D	90.95 ± 0.28	85.38 ± 0.49	93.15 ± 0.26	91.10 ± 2.17	69.44 ± 0.27	60.84 ± 0.38	71.90 ± 0.56	68.38 ± 1.11
T-Revision	88.68 ± 0.54	82.59 ± 1.82	91.10 ± 0.21	72.65 ± 16.11	61.14 ± 0.69	40.01 ± 0.90	60.56 ± 0.66	49.44 ± 1.64
PLM-R	91.33 ± 0.47	85.22 ± 0.54	92.69 ± 0.23	91.33 ± 0.60	67.52 ± 1.19	50.81 ± 3.31	70.58 ± 1.13	56.85 ± 1.73
VolMinNet	90.19 ± 0.17	84.09 ± 0.68	91.55 ± 0.13	86.12 ± 1.26	67.70 ± 0.99	57.99 ± 0.40	71.65 ± 0.62	61.21 ± 2.98
PLM-V	91.75 ± 0.19	84.10 ± 0.55	93.40 ± 0.23	86.91 ± 1.02	70.95 ± 0.49	62.37 ± 0.29	74.55 ± 0.12	64.05 ± 0.53

Additionally, we report the results on the Clothing1M dataset in Table 2, utilizing the same dataset settings as VolMinNet. Our proposed method demonstrates superior performance over the compared baseline methods on the Clothing1M dataset, reflecting a 3.91% enhancement relative to the Forward method. Additionally, we reported the performance on the real-world Animal-10N [39] dataset in Appendix H. These results of real-world datasets can be attributed to the PLM approach, which models instance information at the part-level, thereby bolstering the model’s robustness against instance-dependent noise. To further demonstrate this advantage, we conducted experiments in IDN synthetic noise in Table 3. Here, we adopted Forward’s approach to estimate the noise transition matrix to evaluate how PLM improves the robustness to IDN using this method, and employ a backbone network in alignment with the experimental and network settings of PTD [3]. We compare with the following baseline methods: (1) Co-teaching [11]; (2) DMI [51]; (3) Forward [37]; (4) PTD [47], which uses a part-based noise transition matrix estimation technique; (5) TMDNN [52], which uses deep neural networks to estimate the transition matrix by exploiting Bayes optimal labels; (6) MEIDTM [3], which estimates IDN transition matrix by formulating a manifold assumption. Among them, PTD, TMDNN, and MEIDTM are methods that are designed for the IDN modeling. The results demonstrates that PLM achieves competitive performance without utilizing IDN modeling techniques.

Learning with different noise transition matrices. To assess the efficacy of our proposed method in combination with various noise transition matrix estimation techniques, we present classification experiment results on the CIFAR-10 and CIFAR-100 datasets in Table 4. Specifically, we merge PLM with the matrix estimation techniques outlined in Forward [37], Dual-T [54], T-Revision [46], and VolMinNet [27] (denoted as PLM-F, PLM-D, PLM-R, and PLM-V). As both original Forward and T-Revision methods filtered the highest-confidence portions of samples, their methods may not be universally applicable [54]. Therefore,

to assess PLM’s performance in a more general context when combined with various transition matrices, we utilize the highest probability samples in each category as anchors for Forward, T-Revision, PLM-F, and PLM-R. For PLM-R, we also integrated the reweighting strategy outlined in T-Revision [46], detailed in Appendix A. Table 4 illustrates that our approach attains increased accuracy while significantly reducing biases. The results suggest that PLM can serve as a component to enhance the robustness of existing classifier-consistent methods. Furthermore, experiments for transition matrix in Appendix G indicate that PLM can help reduce the matrix estimation error.

5. Conclusion and limitation

Estimating the noisy class posterior accurately is critical for noisy label learning. In this paper, we introduces a part-level multi-labeling method aimed at augmenting supervised information, thereby reducing the estimation error of estimating noisy class posterior. By introducing a single-to-multiple transition matrix, we incorporate the part-level supervised information derived from cropped instances into a classifier-consistent framework, effectively mitigating overemphasis. Extensive experiments validate the robust performance of our method, both for estimating the noisy class posterior and for noisy label learning as a component of loss correction. One significant limitation of this study is that cropping original instances with a uniform criterion for labeling may be too simplistic. Future work will involve a detailed exploration of the theories and practices associated with instance representation learning, with the objective of formulating a more appropriate cropping criterion for noisy label learning.

Acknowledgments. This research was partially supported by the National Key Research and Development Project of China No.2021ZD0110700, the Key Research and Development Project in Shaanxi Province No.2023GXLH-024, the National Science Foundation of China under Grant No.62037001 and No.62250009.

References

- [1] Shivani Agarwal, Aatif Awan, and Dan Roth. Learning to detect objects in images via a sparse, part-based representation. *IEEE transactions on pattern analysis and machine intelligence*, 26(11):1475–1490, 2004. 2
- [2] Irving Biederman. Recognition-by-components: a theory of human image understanding. *Psychological review*, 94(2): 115, 1987. 2
- [3] De Cheng, Tongliang Liu, Yixiong Ning, Nannan Wang, Bo Han, Gang Niu, Xinbo Gao, and Masashi Sugiyama. Instance-dependent label-noise learning with manifold-regularized transition matrix estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16630–16639, 2022. 1, 3, 7, 8
- [4] De Cheng, Yixiong Ning, Nannan Wang, Xinbo Gao, Heng Yang, Yuxuan Du, Bo Han, and Tongliang Liu. Class-dependent label-noise learning with cycle-consistency regularization. *Advances in Neural Information Processing Systems*, 35:11104–11116, 2022. 1, 3, 4
- [5] Hao Cheng, Zhaowei Zhu, Xingyu Li, Yifei Gong, Xing Sun, and Yang Liu. Learning with instance-dependent label noise: A sample sieve approach. *arXiv preprint arXiv:2010.02347*, 2020. 3
- [6] Jiacheng Cheng, Tongliang Liu, Kotagiri Ramamohanarao, and Dacheng Tao. Learning with bounded instance and label-dependent label noise. In *International Conference on Machine Learning*, pages 1789–1799. PMLR, 2020. 1, 3
- [7] Thomas G Dietterich, Richard H Lathrop, and Tomás Lozano-Pérez. Solving the multiple instance problem with axis-parallel rectangles. *Artificial intelligence*, 89(1-2):31–71, 1997. 2
- [8] Arpit Garg, Cuong Nguyen, Rafael Felix, Thanh-Toan Do, and Gustavo Carneiro. Instance-dependent noisy label learning via graphical modelling. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 2288–2298, 2023. 3
- [9] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. In *International conference on machine learning*, pages 1321–1330. PMLR, 2017. 3
- [10] Bo Han, Jiangchao Yao, Gang Niu, Mingyuan Zhou, Ivor Tsang, Ya Zhang, and Masashi Sugiyama. Masking: A new perspective of noisy supervision. *Advances in neural information processing systems*, 31, 2018. 3
- [11] Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor Tsang, and Masashi Sugiyama. Co-teaching: Robust training of deep neural networks with extremely noisy labels. *Advances in neural information processing systems*, 31, 2018. 3, 5, 6, 8
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part IV 14*, pages 630–645. Springer, 2016. 5
- [13] Dan Hendrycks, Mantas Mazeika, Duncan Wilson, and Kevin Gimpel. Using trusted data to train deep networks on labels corrupted by severe noise. *Advances in neural information processing systems*, 31, 2018. 1, 3
- [14] Ehsan Hosseini-Asl, Jacek M Zurada, and Olfa Nasraoui. Deep learning of part-based representation of data using sparse autoencoders with nonnegativity constraints. *IEEE transactions on neural networks and learning systems*, 27(12):2486–2498, 2015. 2
- [15] Yingsong Huang, Bing Bai, Shengwei Zhao, Kun Bai, and Fei Wang. Uncertainty-aware learning against label noise on imbalanced datasets. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 6960–6969, 2022. 1
- [16] Ahmet Iscen, Jack Valmadre, Anurag Arnab, and Cordelia Schmid. Learning with neighbor consistency for noisy labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4672–4681, 2022. 6
- [17] Nazmul Karim, Mamshad Nayeem Rizve, Nazanin Rahnavard, Ajmal Mian, and Mubarak Shah. Unicon: Combating label noise through uniform selection and contrastive learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9676–9686, 2022. 1, 3
- [18] Jan Kremer, Fei Sha, and Christian Igel. Robust active label correction. In *International conference on artificial intelligence and statistics*, pages 308–316. PMLR, 2018. 3
- [19] Alex Krizhevsky et al. Learning multiple layers of features from tiny images. *Technical report*, 2009. 5
- [20] Yann LeCun. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998. 5
- [21] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. 5
- [22] Junnan Li, Richard Socher, and Steven CH Hoi. Dividemix: Learning with noisy labels as semi-supervised learning. In *International Conference on Learning Representations*, pages 1–13, 2019. 1, 3, 6
- [23] Jichang Li, Guanbin Li, Feng Liu, and Yizhou Yu. Neighborhood collective estimation for noisy label identification and correction. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXIV*, pages 128–145. Springer, 2022. 1, 3
- [24] Mingchen Li, Mahdi Soltanolkotabi, and Samet Oymak. Gradient descent with early stopping is provably robust to label noise for overparameterized neural networks. In *International conference on artificial intelligence and statistics*, pages 4313–4324. PMLR, 2020.
- [25] Shikun Li, Xiaobo Xia, Shiming Ge, and Tongliang Liu. Selective-supervised contrastive learning with noisy labels. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 316–325, 2022. 3
- [26] Shikun Li, Xiaobo Xia, Hansong Zhang, Yibing Zhan, Shiming Ge, and Tongliang Liu. Estimating noise transition matrix with label correlations for noisy multi-label learning. *Advances in Neural Information Processing Systems*, 35: 24184–24198, 2022. 3

- [27] Xuefeng Li, Tongliang Liu, Bo Han, Gang Niu, and Masashi Sugiyama. Provably end-to-end label-noise learning without anchor points. In *International Conference on Machine Learning*, pages 6403–6413. PMLR, 2021. 1, 3, 5, 6, 7, 8
- [28] Tongliang Liu and Dacheng Tao. Classification with noisy labels by importance reweighting. *IEEE Transactions on pattern analysis and machine intelligence*, 38(3):447–461, 2015. 3
- [29] Yang Liu and Hongyi Guo. Peer loss functions: Learning from noisy labels without knowing noise rates. In *International conference on machine learning*, pages 6226–6236. PMLR, 2020. 3
- [30] Nikos K Logothetis and David L Sheinberg. Visual object recognition. *Annual review of neuroscience*, 19(1):577–621, 1996. 2
- [31] Michal Lukasik, Srinadh Bhojanapalli, Aditya Menon, and Sanjiv Kumar. Does label smoothing mitigate label noise? In *International Conference on Machine Learning*, pages 6448–6458. PMLR, 2020. 3
- [32] Yueming Lyu and Ivor W Tsang. Curriculum loss: Robust learning and generalization against label corruption. *arXiv preprint arXiv:1905.10045*, 2019. 3
- [33] Eran Malach and Shai Shalev-Shwartz. “Decoupling” when to update” from” how to update”. *Advances in neural information processing systems*, 30, 2017. 6
- [34] Tam Nguyen, C Mummadi, T Ngo, L Beggel, and Thomas Brox. Self: learning to filter noisy labels with self-ensembling. In *International Conference on Learning Representations (ICLR)*, 2020. 3
- [35] Mohammad Norouzi, Tomas Mikolov, Samy Bengio, Yoram Singer, Jonathon Shlens, Andrea Frome, Greg S Corrado, and Jeffrey Dean. Zero-shot learning by convex combination of semantic embeddings. In *2nd International Conference on Learning Representations, ICLR 2014*, 2014. 2
- [36] Stephen E Palmer. Hierarchical structure in perceptual representation. *Cognitive psychology*, 9(4):441–474, 1977. 2
- [37] Giorgio Patrini, Alessandro Rozza, Aditya Krishna Menon, Richard Nock, and Lizhen Qu. Making deep neural networks robust to label noise: A loss correction approach. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1944–1952, 2017. 1, 3, 5, 6, 8
- [38] Mengye Ren, Wenyuan Zeng, Bin Yang, and Raquel Urtasun. Learning to reweight examples for robust deep learning. In *International conference on machine learning*, pages 4334–4343. PMLR, 2018. 3
- [39] Hwanjun Song, Minseok Kim, and Jae-Gil Lee. SELFIE: Refurbishing unclean samples for robust deep learning. In *ICML*, 2019. 8
- [40] Daiki Tanaka, Daiki Ikami, Toshihiko Yamasaki, and Kiyoharu Aizawa. Joint optimization framework for learning with noisy labels. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5552–5560, 2018. 1, 3
- [41] Shimon Ullman. High-level vision: Object recognition and visual cognition. *Cell*, 86:599–606, 1996. 2
- [42] E Wachsmuth, MW Oram, and DI Perrett. Recognition of objects and their component parts: responses of single units in the temporal cortex of the macaque. *Cerebral Cortex*, 4(5):509–522, 1994. 2
- [43] Yikai Wang, Xinwei Sun, and Yanwei Fu. Scalable penalized regression for noise detection in learning with noisy labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 346–355, 2022. 1, 3
- [44] Shijie Wu and Xun Gong. Boundaryface: A mining framework with noise label self-correction for face recognition. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XIII*, pages 91–106. Springer, 2022. 1, 3
- [45] Songhua Wu, Xiaobo Xia, Tongliang Liu, Bo Han, Mingming Gong, Nannan Wang, Haifeng Liu, and Gang Niu. Class2simi: A noise reduction perspective on learning with noisy labels. In *International Conference on Machine Learning*, pages 11285–11295. PMLR, 2021. 1, 3, 6
- [46] Xiaobo Xia, Tongliang Liu, Nannan Wang, Bo Han, Chen Gong, Gang Niu, and Masashi Sugiyama. Are anchor points really indispensable in label-noise learning? *Advances in Neural Information Processing Systems*, 32:1–12, 2019. 1, 3, 5, 6, 8
- [47] Xiaobo Xia, Tongliang Liu, Bo Han, Nannan Wang, Mingming Gong, Haifeng Liu, Gang Niu, Dacheng Tao, and Masashi Sugiyama. Part-dependent label noise: Towards instance-dependent label noise. *Advances in Neural Information Processing Systems*, 33:7597–7610, 2020. 1, 2, 3, 5, 7, 8
- [48] Xiaobo Xia, Tongliang Liu, Bo Han, Mingming Gong, Jun Yu, Gang Niu, and Masashi Sugiyama. Sample selection with uncertainty of losses for learning with noisy labels. *arXiv preprint arXiv:2106.00445*, 2021. 1
- [49] Xiaobo Xia, Bo Han, Nannan Wang, Jiankang Deng, Jiantong Li, Yinian Mao, and Tongliang Liu. Extended t: Learning with mixed closed-set and open-set noisy labels. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. 1, 3, 5
- [50] Tong Xiao, Tian Xia, Yi Yang, Chang Huang, and Xiaogang Wang. Learning from massive noisy labeled data for image classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2691–2699, 2015. 5
- [51] Yilun Xu, Peng Cao, Yuqing Kong, and Yizhou Wang. L_{dmi}: A novel information-theoretic loss function for training deep nets robust to label noise. *Advances in neural information processing systems*, 32, 2019. 3, 6, 8
- [52] Shuo Yang, Erkun Yang, Bo Han, Yang Liu, Min Xu, Gang Niu, and Tongliang Liu. Estimating instance-dependent bayes-label transition matrix using a deep neural network. In *International Conference on Machine Learning*, pages 25302–25312. PMLR, 2022. 1, 3, 8
- [53] Quanming Yao, Hansi Yang, Bo Han, Gang Niu, and James Tin-Yau Kwok. Searching to exploit memorization effect in learning with noisy labels. In *International Conference on Machine Learning*, pages 10789–10798. PMLR, 2020. 1, 3
- [54] Yu Yao, Tongliang Liu, Bo Han, Mingming Gong, Jiankang Deng, Gang Niu, and Masashi Sugiyama. Dual t: Reducing estimation error for transition matrix in label-noise learn-

ing. *Advances in neural information processing systems*, 33: 7260–7271, 2020. [1](#), [3](#), [5](#), [6](#), [8](#)

[55] Yu Yao, Tongliang Liu, Mingming Gong, Bo Han, Gang Niu, and Kun Zhang. Instance-dependent label-noise learning under a structural causal model. *Advances in Neural Information Processing Systems*, 34:4409–4420, 2021. [3](#)

[56] Zhilu Zhang and Mert Sabuncu. Generalized cross entropy loss for training deep neural networks with noisy labels. *Advances in neural information processing systems*, 31, 2018. [3](#)