

Combating Semantic Contamination in Learning with Label Noise

Wenxiao Fan, Kan Li[†]

School of Computer Science, Beijing Institute of Technology
{wenxiaofan, likan}@bit.edu.cn

Abstract

Noisy labels can negatively impact the performance of deep neural networks. One common solution is label refurbishment, which involves reconstructing noisy labels through predictions and distributions. However, these methods may introduce problematic semantic associations, a phenomenon that we identify as Semantic Contamination. Through an analysis of Robust LR, a representative label refurbishment method, we found that utilizing the logits of views for refurbishment does not adequately balance the semantic information of individual classes. Conversely, using the logits of models fails to maintain consistent semantic relationships across models, which explains why label refurbishment methods frequently encounter issues related to Semantic Contamination. To address this issue, we propose a novel method called Collaborative Cross Learning, which utilizes semi-supervised learning on refurbished labels to extract appropriate semantic associations from embeddings across views and models. Experimental results show that our method outperforms existing approaches on both synthetic and real-world noisy datasets, effectively mitigating the impact of label noise and Semantic Contamination.

Extended version — <https://arxiv.org/abs/2412.11620>

Introduction

In recent years, notable progress has been achieved in various fields through deep learning methodologies (Bochkovskiy, Wang, and Liao 2020; Marriott, Romdhani, and Chen 2021). The use of labeled datasets plays a pivotal role in achieving these notable outcomes. Nevertheless, as datasets continue to expand in size, the probability of encountering noisy labels also increases. Such corrupted knowledge can be assimilated by models, which consequently results in a noticeable decrease in their performance (Zhang et al. 2017; Arpit et al. 2017). This occurrence naturally prompts an urgent inquiry into how deep learning continues to succeed despite the presence of label noise.

State-of-the-art methods in Learning with Noisy Labels (LwNL) have notably enhanced noise robustness through label refurbishment methods (Malach and Shalev-Shwartz

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

[†]Corresponding author.

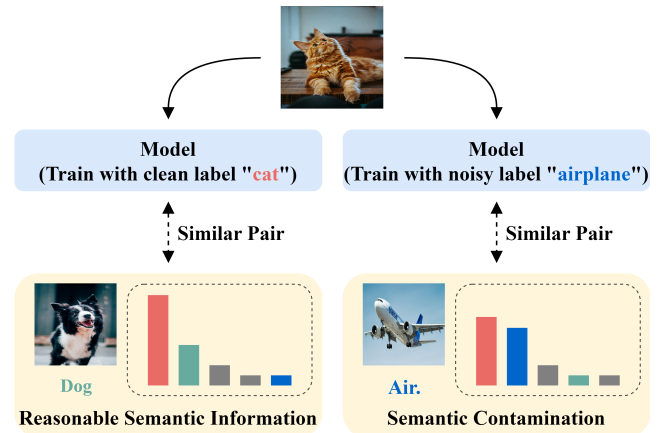


Figure 1: Illustration of Semantic Contamination. Air. is short for airplane. After being trained with noisy labels, models may learn the problematic semantic pairs, such as (cat, airplane) is more similar than (cat, dog). This issue may cause the model to learn incorrect feature spaces, impacting its performance. In this study, we mainly focus on how to enable the model to learn reasonable semantic information in order to overcome Semantic Contamination.

2017; Song, Kim, and Lee 2019; Chen et al. 2020, 2023). The core idea of label refurbishment methods is to transform problematic labels into new, informative labels for the model to learn. Previous researches (Tarvainen and Valpola 2017; Song, Kim, and Lee 2019; Chen et al. 2023) have shown that label refurbishment methods face self-reinforcing errors and confirmation bias, which can hinder performance.

Meanwhile, we have also discovered that, in addition to the two previously mentioned drawbacks, label refurbishment methods face another significant issue that can be more detrimental to the model: **Semantic Contamination (SC)**. This scenario pertains to the model’s inability to comprehend reasonable semantic associations, leading to the failure to acquire robust and consistent representations. For instance, if a cat is mislabeled in an airplane, the similarity between a cat and a dog could be smaller than the similarity between a cat and an airplane, as illustrated in Fig. 1. For clean datasets, reasonable semantic information can be eas-

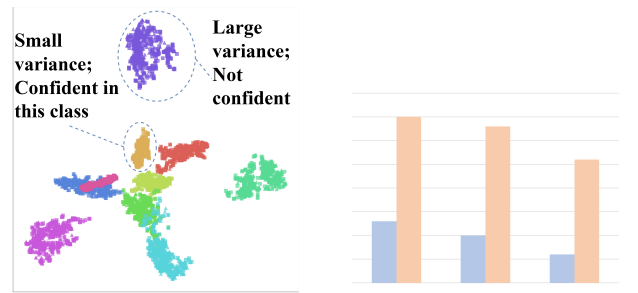
ily learned by the model. However, in LwNL, samples in the same class could be clustered into different categories, leading to inconsistent predictions. Hence, how to acquire the relevant semantic information remains an open question.

This study focuses on how label refurbishment methods can mitigate the impact of noisy labels and Semantic Contamination. We initially analyze *why the label refurbishment methods are prone to experiencing Semantic Contamination* and observe that directly aligning logits from different models, a common practice in label refurbishment, does not effectively align the embeddings. While aligning the logits of different views can cluster the embeddings, the unequal confidences for each class hinder the models from acquiring appropriate semantic information. Instead of relying solely on refurbishing with logits, we suggest mining latent relevancy in embeddings across views and models to learn semantic information and propose a novel method called Collaborative Cross Learning, which consists of two components: Cross-view learning and Cross-model learning. For Cross-view learning, we decouple the class label and the semantic concept and utilize self-supervised learning to prevent the incorporation of harmful semantic information. For Cross-model learning, we propose promoting the alignment of different models by using Collaborative Contrastive Learning Refurbished Labels (CCLRL) and theoretically establish that optimizing CCLRL enhances mutual information between the two models. The superiority of our method over state-of-the-art (SOTA) methods is demonstrated through validation on various synthetic and real-world benchmarks. Our contributions can be summarized as follows:

- We introduce a new challenge called Semantic Contamination in LwNL and analyze the reason why label refurbishment methods may be susceptible to SC from the perspective of views and models.
- We propose a novel method called Collaborative Cross Learning. By decoupling the semantic concept between views and mimicking the contrastive distributions between models, it successfully obtains robust and consistent representations while alleviating the damage of both Semantic Contamination and label noise.
- Experimental results show that our method advances state-of-the-art results on CIFAR with synthetic label noise, as well as on real-world noisy datasets.

Related Work

Label Refurbishment in LwNL. For label refurbishment, mainstream methods estimate refurbished labels through three ways: 1) Models: Decouple (Malach and Shalev-Shwartz 2017) updates two models with only disagreed samples between models. SEAL (Chen et al. 2020) retrains the model with the average predictions of the teacher model as refurbished labels. 2) Models and Views: RoLR (Chen et al. 2023) refurbishes noisy labels by aligning the predictions between models and views. 3) Historical predictions: SELFIE (Song, Kim, and Lee 2019) only includes samples with consistent predictions in recent epochs for refurbishment. However, the process of refurbishing labels poses



(a) T-SNE analyze for RoLR. (b) Entropy of variances of each class in CIFAR-10 under different symmetric noise (20%, 50% and 80%).

Figure 2: Illustration of Semantic Imbalance Among Classes. Fig. 2(b) shows that our method (orange) can learn more balance representations compared with RoLR (blue).

challenges in high-noise environments due to self-refining errors, which can impede the training of models.

Semi-supervised Learning. The field of semi-supervised learning (SSL) has experienced significant advances through the application of consistency regularization, which aims to minimize the disparity in model predictions between two views of the same sample or two predictions of the same sample using different models. MixMatch (Berthelot et al. 2019) initially aligned outputs from different views and models, but the enhancement strategies for different views were consistent. In contrast, FixMatch (Sohn et al. 2020) utilized both strong and weak transformations, in addition to confidently pseudo-labeling, which led to favorable results. Currently, SSL is also implemented in LwNL, for instance, DivideMix (Li, Socher, and Hoi 2020) incorporates MixMatch in LwNL, while RankMatch (Zhang et al. 2023) partially adopts FixMatch. RoLR (Chen et al. 2023) also employs distinct SSL augmentation strategies. However, SSL is susceptible to confirmation bias, and our experimental results indicate that SSL is still influenced by SC, both of which can impact the model’s performance.

Semantic Contamination in LwNL

As mentioned in the Introduction, Semantic Contamination refers to the phenomenon where, in the presence of latent semantic relationships among samples, a model fails to capture accurate semantic associations. The captured corrupted semantic information often corresponds to noisy labels. Such toxic information is commonly utilized in various label refurbishment methods as part of the pseudo-labels and can harm generalization and lead to poor robustness in LwNL. To address this issue, in this study, we first explore *why label refurbishment methods are susceptible to Semantic Contamination*. First, we take RoLR (Chen et al. 2023) as an example of label refurbishment methods and analyze the influence of different views and models through which RoLR may have learned potentially erroneous semantic information. We demonstrate that relying solely on views for refurbishing

bishing could lead to **semantic imbalance among classes**, ultimately impairing performance. Apart from that, we find that relying on models for refurbishing cannot even maintain the **semantic consistency across models**. Addressing these drawbacks is the focus of our method.

Preliminary. For convenience, notations within our work are clarified first. For a C -way image classification task with noisy labels, the training dataset is denoted by $D = \{(x, \hat{y})\}$, where x is the training sample and \hat{y} is the label which may be incorrectly annotated. We denote the models in the training stage as $\theta = g(f(x))$, where f is the feature extractor and g represents the classifier. p is the prediction and p_i is the probability of i -th class of input distribution p .

RoLR, which we used as the example, is a SOTA method that integrates pseudo-labeling and confidence estimation techniques to refurbish noisy labels. In pseudo-labeling stage, the pseudo-labels \tilde{y} is create by two averaged and sharpened models' predictions from different augmentation strategies, as depicted in Eq. (1).

$$\tilde{y}^m = \text{Sharpen}\left(\frac{p^w(y|x; \theta_m) + p^w(y|x; \theta_{(1-m)})}{2}\right) \quad (1)$$

$$\text{Sharpen}(p)_i = \frac{p_i^{\frac{1}{T}}}{\sum_{j=1}^C p_j^{\frac{1}{T}}}$$

where \tilde{y}^m is the pseudo-labels for model θ_m , $m \in \{0, 1\}$ means two models, p^w, p^s mean weak and strong augmentation strategies for p . C is the number of classes and T is the temperature. The loss of RoLR can be written by:

$$\begin{aligned} \mathcal{L}(\mathbf{p}_{\theta_m}^s, y^*) &= \mathcal{L}_c(\mathbf{p}_{\theta_m}^s, \omega \hat{y} + (1 - \omega) \tilde{y}^m) \\ &= \underbrace{\omega \mathcal{L}_c(\mathbf{p}_{\theta_m}^s, \hat{y})}_{\text{Correct guidance}} + \frac{1 - \omega}{2} \left(\underbrace{\mathcal{L}_c(\mathbf{p}_{\theta_m}^s, \mathbf{p}_{\theta_m}^w)}_{\text{Cross-view learning}} + \underbrace{\mathcal{L}_c(\mathbf{p}_{\theta_m}^s, \mathbf{p}_{\theta_v}^w)}_{\text{Cross-model learning}} \right) \end{aligned} \quad (2)$$

where y^* is the refurbished label, ω is the label confidence obtained by the confidence estimation stage and $\mathbf{p}_{\theta_m}^w$ is simplified form of $p^w(y|x; \theta_m)$, \mathcal{L}_c is the cross-entropy, $v = 1 - m$ for short. We omit the Sharpen (set $T = 1$) and decompose the loss into three parts in Eq. (2): Correct guidance, Cross-view learning and Cross-model learning. The last two terms directly impact the semantic information and may cause harmful drawbacks: semantic imbalance among classes and semantic inconsistency among, respectively.

Semantic imbalance among classes. Cross-view learning in Eq. (2) aligns different views of the same model. To evaluate its influence specifically on semantic information, we decouple it into two components: one related to the predicted class $y_{\text{pred}} = \max_i p_i^s$ and the other related to semantic concepts in Eq. (3).

$$\mathcal{L}_c(p^s, p^w) = \underbrace{\mathcal{L}_c(p_{y_{\text{pred}}}^s, p_{y_{\text{pred}}}^w)}_{\text{Prediction guidance}} + \underbrace{\sum_{i \neq y_{\text{pred}}} \mathcal{L}_c(p_i^s, p_i^w)}_{\text{Semantic smoothing}} \quad (3)$$

While Prediction guidance ensures alignment of views of the samples during training, Semantic smoothing aims

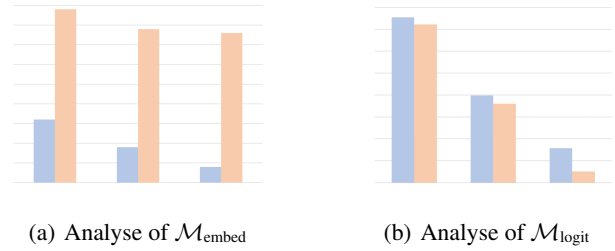


Figure 3: Evaluation results of semantic consistency across models between RoLR (blue) with our method (orange) on CIFAR-10 with different symmetric noise (20%, 50% and 80%).

to obtain consistent semantic representations, as illustrated in Fig. 2(a). However, Semantic smoothing can be highly vulnerable to noisy labels. Upon evaluating the entropy of variances of each class in Fig. 2(b), it is observed that the variances are significantly imbalanced due to the fail of Semantic smoothing. This phenomenon, named semantic imbalance among classes, leads to the model showing overconfidence in certain classes. Such an imbalance may introduce discrepancies in representations of different classes, which ultimately results in performance degradation.

Semantic inconsistency among models. Beside Cross-view learning in Eq. (2), Cross-model learning term offers an alternative way to gain semantic information across models by aligning the models in logits. In general, we aim to achieve semantic consistency among different models by learning a unified representation across models. However, the existence of noisy labels and the random initialization of different models impede the extraction of universal embeddings among models. Therefore, instead of pursuing this unattainable goal, our aim is to align the semantic relationships across models, ensuring semantic consistency across models. To evaluate the alignment of semantic relationships, we investigate the following two metrics:

- Embedding-wise metric $\mathcal{M}_{\text{embed}}$: We utilize Eq. (4) to assess the consistency of relationships between embeddings across models (Mikolov et al. 2013), such as $f_{\theta_1}(\text{dog}) = f_{\theta_1}(\text{cat}) - f_{\theta_0}(\text{cat}) + f_{\theta_0}(\text{dog})$.

$$\begin{aligned} \mathcal{M}_{\text{embed}}(D, \theta_0, \theta_1) &= \\ &= \mathbb{E}_{x_0, x_1 \in D} [\mathcal{M}_{\text{cos}}(f_{\theta_0}(x_0) - f_{\theta_0}(x_1); f_{\theta_1}(x_0) - f_{\theta_1}(x_1))] \end{aligned} \quad (4)$$

where \mathcal{M}_{cos} is the cosine similarity. Larger $\mathcal{M}_{\text{embed}}$ shows more greater semantic consistency across models.

- Logits-wise metric $\mathcal{M}_{\text{logit}}$: The Wasserstein Distance of logits between two models can show the level of disparity among models. Models with larger disparities exhibit less semantic consistency.

Fig. 3 displays the two metrics for various methods. It is observed that both metrics for RoLR are weaker than those for our method, indicating that RoLR does not maintain semantic consistency across models, leading to models being unable to learn consistent and robust representations.

We conclude the above analysis: *Cross-view learning in label refurbishment leads to semantic imbalance among classes, while Cross-model learning results in semantic inconsistency across models, which are why label refurbishment methods fail to obtain the appropriate semantic information.* This conclusion inspires us that simply refurbishing logits is vulnerable with noisy labels and insufficient for learning the right representations. It is necessary to incorporate independent modules that can extract reasonable semantic information, which we introduce next.

Methodology

The previous section concludes that refurbishments on logits are inadequate for accurate semantic information. To tackle this challenge, we propose a method, named Collaborative Cross Learning, that learns semantic relationships from embedding perspectives. Specifically, we propose two modules: Semantic-wise Decoupling with Confident Learning (SDCL), which can help in balancing the confidence among classes for Cross-view learning, and Embedding-based Interactive Alignment (EIA), which aids in aligning the models and maintaining semantic consistency across models for Cross-model learning. The overall pipeline is shown in Fig. 4 and the algorithm pseudocode is in Appendix.

Semantic-wise Decoupling with Confident Learning. Similar to Eq. (3), Semantic-wise Decoupling decouples the prediction into the predicted class and the semantic concepts. To avoid the effect of noisy labels, instead of using the logits in Semantic smoothing term, we employ an Augmentation-wise Contrastive Learning approach on embeddings. Strong augmentations are used as anchors to align with weak augmentations to reduce training difficulty, as depicted in Eq. (5).

$$\begin{aligned} \mathcal{L}_{\text{ACL}}(x, \theta) &= -\log q_{\theta}^{s \rightarrow w} \\ &= -\log \frac{\exp(f_{\theta}^s(x) \cdot f_{\theta}^w(x))}{\sum_{j=1}^N \exp(f_{\theta}^s(x) \cdot f_{\theta}^w(x_j))} \end{aligned} \quad (5)$$

where N is the batch number, $q_{\theta}^{s \rightarrow w}$ is the contrastive distribution of the model θ from strong augmentations s to weak augmentations w .

In addition to directly aligning the representations of different views, we also use **KL** divergence to align the latent relationships of samples between views, named as **View-wise Mimicry**:

$$\mathcal{L}_{\text{VM}}(x, \theta) = \text{KL}(q_{\theta}^{s \rightarrow w} || q_{\theta}^{w \rightarrow s}) \quad (6)$$

For **Predicted Guidance** term in Eq. (3), considering the influence of noise that the prediction may be corrupted, the sample should only be learned when it exhibits a high confidence in the prediction as **Confident Learning**:

$$\mathcal{L}_{\text{PG}}(x, \theta) = -\mathbb{1}[y_{\text{pred}} \geq c] \cdot y_{\text{pred}} \log \mathbf{p}_{\theta}^s \quad (7)$$

where c is the threshold for confidence and $\mathbb{1}[\cdot]$ is the characteristic function.

The loss of Cross-view Learning (CVL) can be calculated as follows:

$$\mathcal{L}_{\text{CVL}} = \mathcal{L}_{\text{PG}} + \mathcal{L}_{\text{ACL}} + \mathcal{L}_{\text{VM}} \quad (8)$$

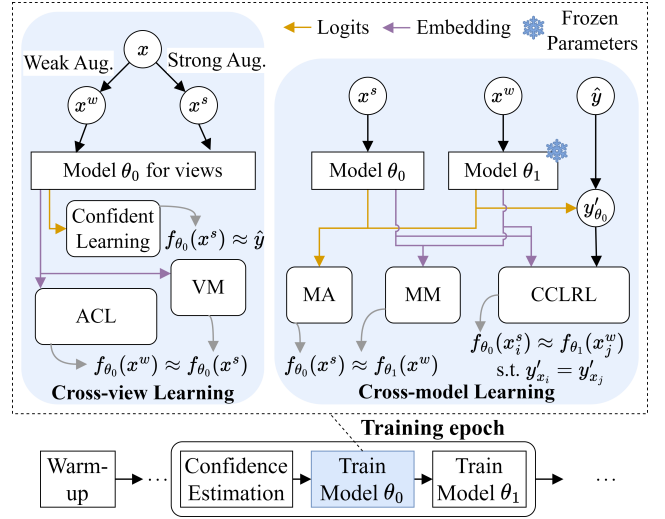


Figure 4: Pipeline of our method. The details of Warm-up and Confidence estimation by small-loss criterion and GMM can be found in Appendix.

\mathcal{L}_{PG} is related to the predicted class and $\mathcal{L}_{\text{ACL}} + \mathcal{L}_{\text{VM}}$ can acquire relevant semantic information, separately.

Embedding-based Interactive Alignment. In order to get semantic consistency across models, in addition to aligning the logits in Eq. (1), it is necessary to explore cross-model relationships in embeddings. To address this, we introduce a novel approach called Collaborative Contrastive Learning on Refurbished Labels (CCLRL) to fully exploit the information interaction among diverse peer models. Specifically, CCLRL considers same-class samples with strong transformations $f_{\theta_m}^s(x)$ from θ_m and weak transformations $f_{\theta_{(1-m)}}^w(x)$ from $\theta_{(1-m)}$ as positive sample pairs, while treating all other samples pairs as negative sample pairs, as shown in Eq. (9).

$$\begin{aligned} \mathcal{L}_{\text{CCLRL}}(x, \theta_m, \theta_{(1-m)}) &= -\log q_{\theta_m \rightarrow \theta_{(1-m)}}^{s \rightarrow w} \\ &= -\log \frac{\sum_{j=1, y'_x=y'_x_j}^N \exp(f_{\theta_m}^s(x) \cdot f_{\theta_{(1-m)}}^w(x_j))}{\sum_{k=1}^N \exp(f_{\theta_m}^s(x) \cdot f_{\theta_{(1-m)}}^w(x_k))} \end{aligned} \quad (9)$$

Given the presence of noisy labels, the identification of samples of the same class is based on refurbished collaborative labels y' . Furthermore, to avoid self-refining errors, the collaborative refurbished labels for the current model θ_m are obtained from another model $\theta_{(1-m)}$, as shown in Eq. (10).

$$y'_{\theta_m} = \omega \hat{y} + (1 - \omega) \mathbf{p}_{\theta_{(1-m)}}^w \quad (10)$$

Theoretical Analysis. We attribute the effectiveness of minimizing Eq. (9) to maximizing the lower bound on the mutual information $I(f_{\theta_m}(x), f_{\theta_{(1-m)}}(x))$ between θ_m and $\theta_{(1-m)}$, which can be formulated as:

$$I(f_{\theta_m}(x), f_{\theta_q}(x)) \geq \log(N) - \mathbb{E}[\mathcal{L}_{\text{CCLRL}}(x, \theta_m, \theta_v)] \quad (11)$$

Dataset	CIFAR-10					CIFAR-100				
	Noise Type		Sym	Pair	Ins	Sym		Pair	Ins	
Method / Noise Rate	20%	50%	80%	40%	40%	20%	50%	80%	40%	40%
Co-teaching (Han et al. 2018)	88.2	50.7	21.1	55.3	59.5	58.5	33.0	5.8	39.2	40.7
SCE (Wang et al. 2019)	89.9	78.5	31.8	58.8	77.8	55.9	40.2	12.9	39.9	42.4
JoCoR (Wei et al. 2020)	89.4	53.3	25.8	56.1	60.9	55.4	32.7	6.6	34.1	34.9
DivideMix (Li, Socher, and Hoi 2020)	95.7	94.4	92.9	92.1	95.1	76.9	74.2	59.6	52.3	76.1
Co-learning (Tan et al. 2021)	91.8	79.3	37.0	66.3	78.9	70.3	63.9	38.9	49.1	62.9
SELC+ (Lu and He 2022)	94.9	87.2	78.6	88.1	84.2	76.4	62.4	37.2	45.2	44.3
RoLR (Chen et al. 2023)	96.4	95.7	94.2	92.8	93.7	78.6	74.6	66.2	76.1	77.2
RankMatch (Zhang et al. 2023)	96.4	95.4	94.2	94.4	93.8	79.3	77.6	67.2	75.8	76.5
CrossSplit (Kim et al. 2023)	96.9	96.3	95.4	96.0	95.8	79.9	75.7	64.6	76.8	79.2
DISC (Li et al. 2023)	96.3	95.4	92.9	94.6	96.0	78.6	76.3	59.3	75.1	78.4
DMLP (Naive) (Tu et al. 2023)	94.2	94.0	93.2	93.9	93.2	72.3	70.1	63.2	71.8	72.2
DMLP (DivideMix) (Tu et al. 2023)	96.2	95.6	94.3	95.0	95.4	79.4	76.1	68.5	76.4	78.9
Ours	97.0	96.5	94.6	96.1	96.2	79.5	77.4	70.3	77.2	80.0

Table 1: Comparison with state-of-the-art methods on CIFAR-10/100 datasets under various types of noise. The results of other methods are from the published results of corresponding papers. The best results are indicated in bold.

Inspired by (Tian, Krishnan, and Isola 2020; Yang et al. 2022), a detailed proof is provided in Appendix. Intuitively, the mutual information $I(f_{\theta_m}(x), f_{\theta_{(1-m)}}(x))$ measures the uncertainty in contrastive embeddings from $\theta_{(1-m)}$ when the anchor embeddings from θ_m are known (Yang et al. 2022). Higher mutual information implies that θ_m gains additional contrastive knowledge from others, ensuring the preservation of semantic consistency across models.

Similar to View-wise Mimicry, **Model-wise Mimicry** aligns the contrastive distribution q of samples between models:

$$\mathcal{L}_{\text{MM}}(x, \theta_m, \theta_{(1-m)}) = \mathbf{KL}(q_{\theta_m \rightarrow \theta_{(1-m)}}^{s \rightarrow w} || q_{\theta_{(1-m)} \rightarrow \theta_m}^{w \rightarrow s}) \quad (12)$$

Combining with the model alignment with logits in Eq. (1), the loss for Cross-model learning (CML) can be formulated as:

$$\mathcal{L}_{\text{CML}} = \mathcal{L}_{\text{c}}(\mathbf{p}_{\theta_m}^s, \mathbf{p}_{\theta_{(1-m)}}^w) + \mathcal{L}_{\text{CCLRL}} + \mathcal{L}_{\text{MM}} \quad (13)$$

Overall Loss Function. Following (Li, Socher, and Hoi 2020; Chen et al. 2023; Zhang et al. 2023), we apply the regularization \mathcal{L}_{div} to increase the diversity of predictions:

$$\mathcal{L}_{\text{div}} = \sum_{i=1}^C \frac{1}{C} \log\left(\frac{1}{C} / \frac{\sum_{j=1}^N p_j^s[i]}{N}\right) \quad (14)$$

where $p_j^s[i]$ is the i -th class of the prediction of the strong augmentation of sample x_j .

The overall loss function for model θ_m optimization is as follows:

$$\mathcal{L} = \omega \mathcal{L}_{\text{c}}(\mathbf{p}_{\theta_m}^s, \hat{y}) + \frac{1-\omega}{2} (\mathcal{L}_{\text{CVL}} + \mathcal{L}_{\text{CML}}) + \mathcal{L}_{\text{div}} \quad (15)$$

Following (Li, Socher, and Hoi 2020; Chen et al. 2023), we warm up the models and estimate the label confidence ω using the small-loss criterion and Gaussian Mixture Model (GMM). The detailed algorithm is provided in the Appendix.

Experiments

Experimental Setup

Datasets. To verify the effectiveness of our method, we perform our method on classification tasks with six benchmarks: CIFAR-10 (Krizhevsky, Hinton et al. 2009), CIFAR-100 (Krizhevsky, Hinton et al. 2009), CIFAR-10N (Wei et al. 2022), CIFAR-100N (Wei et al. 2022), Animal-10N (Song, Kim, and Lee 2019) and WebVision (Li et al. 2017). The last four benchmarks are real-world noisy datasets.

Synthetic noise injection. Under such an assumption that the corruption process is conditionally independent of data features when the true label is given (Song et al. 2020; Zhang et al. 2017), we can construct the dataset containing noises by the noise transition matrix T (Song et al. 2020; Jiang et al. 2018; Malach and Shalev-Shwartz 2017; Song, Kim, and Lee 2019), where $T_{ij} \stackrel{\text{def}}{=} p(\hat{y} = j | y = i)$ is the probability of the clean label i being corrupted into a noisy label j . T can model two types of noises: (1) symmetric noise (Sym): $\forall i \neq j, T_{ij} = \frac{\tau_0}{C-1}$ and (2) pair noise (Pair): $\exists i \neq j, T_{ij} = \tau_0 \wedge \forall k \neq i, k \neq j, T_{ik} = 0$, which is also known as the asymmetric noise. To evaluate the performance on varying noise rates from light noises to heavy noises, we run our method and other methods on different noise rates $\tau_0 \in \{0.2, 0.4, 0.5, 0.8\}$.

We have also conducted experiments with another type of noise: instance-dependent noise (Ins) (Chen et al. 2020). Unlike the above two types of noise, Ins allows the label noise to depend mandatorily on the samples, and optionally on the labels.

Weak and strong augmentations in our paper are identical with (Chen et al. 2023). All the results from our runs are the average test accuracy over the last 10 epochs. We replicated experimental results that were missing from the corresponding papers. A detailed description of Comparison Methods

Dataset		CIFAR-10N					CIFAR-100N
Method	Noise Type Noise Rate	Aggre 9.0%	Rand1 17.2%	Rand2 18.12%	Rand3 17.64%	Worst 40.2%	Fine 40.2%
Decouple (Malach and Shalev-Shwartz 2017)		88.1	85.5	85.5	85.4	74.6	45.5
GCE (Zhang and Sabuncu 2018)		89.8	88.0	87.8	87.8	75.3	50.3
Co-teaching (Han et al. 2018)		89.9	87.8	87.2	87.4	62.3	40.5
DivideMix (Li, Socher, and Hoi 2020)		93.2	92.8	92.6	93.1	89.2	55.2
Co-learning (Tan et al. 2021)		92.4	91.3	91.2	91.4	81.0	47.9
Cores + LC (Wei et al. 2023)		92.1	90.9	91.2	91.1	82.8	59.2
RoLR (Chen et al. 2023)		95.4	94.9	94.7	95.2	92.3	62.3
RankMatch (Zhang et al. 2023)		95.6	94.8	95.1	95.3	92.8	65.2
Ours		96.4	96.0	95.8	96.1	93.1	65.5

Table 2: Comparison with state-of-the-art methods on CIFAR-N. The results are from (Wei et al. 2023) and our replication.

Method	WebVision		ILSVRC12	
	top-1	top-5	top-1	top-5
Co-teaching	63.6	85.2	61.5	84.7
Iterative-CV	65.2	85.3	61.6	85.0
DSOS	77.8	92.0	74.4	90.8
DivideMix	77.3	91.6	75.2	90.8
UNICON	77.6	93.4	75.3	93.7
RoLR	81.8	94.1	75.5	93.8
RankMatch	79.9	93.6	77.4	94.3
Ours	82.3	94.6	78.2	94.9

Table 3: Comparison with state-of-the-art methods on (mini) WebVision dataset. Numbers denote top-1 (top-5) accuracy (%) on the WebVision and ImageNet ILSVRC12 validation set. The results of other methods are from (Chen et al. 2023; Zhang et al. 2023).

and Implementation Details can be found in Appendix.

Experimental Results

Results on CIFAR with synthetic noise. Tab. 1 illuminates our method outperforms the state-of-the-art models across all noisy levels on CIFAR-10/100 with different types of synthetic noise. Benefiting from the advanced semantic mining mechanism, our method can boost the averaged test accuracy of the last 10 epochs from 94.2% to 94.6% on CIFAR-10 and 66.2% to 70.3% on CIFAR-100 under the extreme case of 80% noise compared with RoLR (Chen et al. 2023). Compared with sample selection methods such as RankMatch (Zhang et al. 2023), our method can achieve around 3.1% (70.3% v.s. 67.2%) on CIFAR-100 with 80% noisy labels. We surpass DMLP (Tu et al. 2023), a recent semi-supervised learning-based method, under all noise ratios, especially on the more challenging CIFAR-100 dataset with extreme noise.

Results on real-world datasets. Tab. 2, Tab. 3 and Tab. 4 demonstrate the results on CIFAR-N, Animal-10N and WebVision, respectively. Our method demonstrates superior performances compared to all other methods across vari-

SELFIE	PLC	NCT	RoLR	DISC	Ours
81.8	83.4	84.1	88.5	87.1	89.7

Table 4: Comparison with other methods on Animal-10N. The results of other methods are from (Chen et al. 2023).

ous noisy real-world datasets. In particular, compared with RoLR, our method achieves 1.2% performance gains in Animal-10N. Besides, compared with UNICON (Karim et al. 2022), another hybrid method that combines the advantages of semi-supervised learning and contrastive learning, our method surpasses the SOTA by over 3% in top-1 accuracy on both mini-WebVision and ILSVRC12 validation sets, ensuring the SOTA top-5 accuracy on WebVision and ILSVRC12, demonstrating the effectiveness of our method.

Results on Semi-supervised learning methods in LwNL.

In addition to label refurbishment methods, recent research has widely adopted semi-supervised learning mechanisms like MixMatch (Berthelot et al. 2019), such as DivideMix (Li, Socher, and Hoi 2020), to overcome the impact of noise. However, these methods also face the challenge of Semantic Contamination. To address this issue, we integrate (\mathcal{L}_{CVL} and \mathcal{L}_{CML}) into the existing loss function to further enhance the learning of better semantic information. From the results shown in Tab. 5, we can uncover the following empirical results: 1) Semi-supervised learning methods do not actually receive proper semantic information. By utilizing both (\mathcal{L}_{CVL} and \mathcal{L}_{CML}), the model can acquire better semantic information, resulting in performance enhancement. 2) Ensuring semantic consistency between the two models (\mathcal{L}_{CML}) can yield greater performance improvements compared to learning semantic information between views (\mathcal{L}_{CVL}). This also implies the importance of addressing semantic inconsistencies between two models.

Results for combating Semantic Contamination. To validate the effectiveness of our method in combating Semantic Contamination, we initially train our method and RoLR on CIFAR-100 with 80% symmetric noise. We pick up samples with accurate predictions but differing second-

Dataset	CIFAR-10		CIFAR-100	
	50%	80%	50%	80%
DivideMix	94.4	92.9	74.2	59.6
+ \mathcal{L}_{CVL}	95.2	93.3	74.8	62.3
+ \mathcal{L}_{CML}	95.9	93.7	75.3	64.2
+ ($\mathcal{L}_{CVL} + \mathcal{L}_{CML}$)	96.1	94.1	76.1	68.4

Table 5: Results on the combination of DivideMix and our method in terms of test accuracy (%) on CIFAR-10 and CIFAR-100 with symmetric noise.





					
	Truth Label	Auto.	Cat	Truck	Dog
	Noisy Label	Auto.	Ship	Horse	Dog
RoLR	Prediction	Auto.	Ship	Horse	Dog
	Similarity	-	0.79	0.56	0.42
Ours	Prediction	Auto.	Ship	Horse	Dog
	Similarity	-	0.52	0.78	0.49

Table 6: Results on Semantic Contamination on CIFAR-10 with 80% symmetric noise. Auto. is short for Automobile. Similarity refers to the cosine similarity between the embeddings of the first sample and the other three samples.

largest logits and then identify samples from the class associated with the second-largest logits in the same batch and compute the similarity among these samples. As illustrated in Tab. 6, we observe that: 1) Although the predictions are accurate, RoLR is still influenced by Semantic Contamination, which tends to assign higher similarity to pairs that lack semantic relevance (e.g. automobiles and cats). 2) in comparison to RoLR, our method assigns higher similarities to samples with semantic relationships (e.g. automobiles and trucks). This demonstrates that our method can capture more relevant semantic information, allowing the model to acquire a more continuous and consistent representation space.

Ablation Study

Effects of components of our method. We remove the corresponding components to study the effects of each component of our method: We remove \mathcal{L}_{CVL} and \mathcal{L}_{CML} to validate the effect of these two modules, respectively. Moreover, we remove both the mimicry in Eq. (6) and Eq. (12) to validate the effect of the mimicry. We also remove the threshold in Eq. (7) to validate the effect of the threshold for filtering the noise. We replace the loss on the refurbished labels in Eq. (9) with the vanilla contrastive learning. As shown in Tab. 7, the results validate the effectiveness of each component of our method. The mimicry, the threshold in \mathcal{L}_{PG} and refurbished labels in Eq. (9) are beneficial to our method.

The role of different augmentation strategies in LwNL.

To explore the role of different augmentation strategies in LwNL, we focus on the effects of different data augmentation strategies on model results. Specifically, we directly

Dataset	CIFAR-10		CIFAR-100	
	50%	80%	50%	80%
Ours	96.5	94.6	77.4	70.3
w/o Mimicry	95.7	93.8	76.7	68.2
w/o Threshold in \mathcal{L}_{PG}	95.8	92.3	75.2	67.9
w/o Refurbished labels	96.1	93.2	76.2	68.4
w/o \mathcal{L}_{CVL}	95.2	92.4	76.5	68.2
w/o \mathcal{L}_{CML}	94.8	91.4	75.1	66.2

Table 7: Ablation study results of test accuracy (%) on CIFAR-10 and CIFAR-100 with symmetric noise.

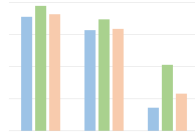
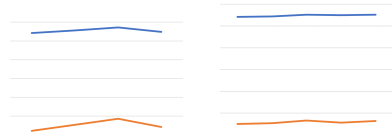


Figure 5: Results on different augmentation strategies on CIFAR-100 under 80% symmetric noises.



(a) Temperature τ from 0.05 to 0.9 (b) Threshold c from 0.90 to 0.99

Figure 6: Sensitivity to the variance of hyperparameters. Experiments are conducted on CIFAR-10 (blue) and CIFAR-100 (orange) under 80% symmetric noises.

train with the exception of different augmentation strategies, everything else remains the same. In Fig. 5, from left to right, the three data augmentation strategies are: both weak (blue), one weak and one strong (green) and both strong (orange), we find that using the same augmentation can affect performance. We speculate that the reason is that various augmentations can increase the diversity of training samples.

Sensitivity Analysis. Our method introduces the temperature τ in Eq. (5) and Eq. (9) and the threshold c in Eq. (7) as hyper-parameters. We vary τ from 0.05 to 0.9 and range c from 0.90 to 0.99. Fig. 6(a) shows that τ should be reasonable, as setting it too high or too low can both degrade the model’s performance. Fig. 6(b) shows that our method is robust against various choices for c .

Conclusion

In this paper, we study the problem of Semantic Contamination together with label noise. We analyze the drawbacks of label refurbishment methods and explain why these methods cannot overcome Semantic Contamination. The conclusion motivates us to propose our method that can learn the more reasonable semantic information through Semantic Decoupling with Confident Learning and Embedding-based Interactive Alignment. Experimental results illustrate that our method surpasses current methods in performance on synthetic and real-world noisy datasets, effectively reducing the influence of label noise and Semantic Contamination. In the future, we will further explore the theoretical foundation and generalization analysis of our method.

Acknowledgments

This work is supported by Beijing Natural Science Foundation (No.4222037, L181010).

References

- Arpit, D.; Jastrzebski, S.; Ballas, N.; Krueger, D.; Bengio, E.; Kanwal, M. S.; Maharaj, T.; Fischer, A.; Courville, A. C.; Bengio, Y.; and Lacoste-Julien, S. 2017. A Closer Look at Memorization in Deep Networks. 70: 233–242.
- Berthelot, D.; Carlini, N.; Goodfellow, I. J.; Papernot, N.; Oliver, A.; and Raffel, C. 2019. MixMatch: A Holistic Approach to Semi-Supervised Learning. In Wallach, H. M.; Larochelle, H.; Beygelzimer, A.; d’Alché-Buc, F.; Fox, E. B.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, 5050–5060.
- Bochkovskiy, A.; Wang, C.; and Liao, H. M. 2020. YOLOv4: Optimal Speed and Accuracy of Object Detection. *CoRR*, abs/2004.10934.
- Chen, M.; Cheng, H.; Du, Y.; Xu, M.; Jiang, W.; and Wang, C. 2023. Two Wrongs Don’t Make a Right: Combating Confirmation Bias in Learning with Label Noise. In Williams, B.; Chen, Y.; and Neville, J., eds., *Thirty-Seventh AAAI Conference on Artificial Intelligence, AAAI 2023, Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence, IAAI 2023, Thirteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2023, Washington, DC, USA, February 7-14, 2023*, 14765–14773. AAAI Press.
- Chen, P.; Ye, J.; Chen, G.; Zhao, J.; and Heng, P.-A. 2020. Beyond Class-Conditional Assumption: A Primary Attempt to Combat Instance-Dependent Label Noise. In *AAAI Conference on Artificial Intelligence*.
- Han, B.; Yao, Q.; Yu, X.; Niu, G.; Xu, M.; Hu, W.; Tsang, I. W.; and Sugiyama, M. 2018. Co-teaching: Robust training of deep neural networks with extremely noisy labels. *Advances in Neural Information Processing Systems*, 2018-Decem(NeurIPS): 8527–8537.
- Jiang, L.; Zhou, Z.; Leung, T.; Li, L. J.; and Fei-Fei, L. 2018. Mentornet: Learning data-driven curriculum for very deep neural networks on corrupted labels. *35th International Conference on Machine Learning, ICML 2018*, 5(6): 3601–3620.
- Karim, N.; Rizve, M. N.; Rahnavard, N.; Mian, A.; and Shah, M. 2022. UNICON: Combating Label Noise Through Uniform Selection and Contrastive Learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, 9666–9676. IEEE.
- Kim, J.; Baratin, A.; Zhang, Y.; and Lacoste-Julien, S. 2023. CrossSplit: Mitigating Label Noise Memorization through Data Splitting. In Krause, A.; Brunskill, E.; Cho, K.; Engelhardt, B.; Sabato, S.; and Scarlett, J., eds., *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, 16377–16392. PMLR.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images.
- Li, J.; Socher, R.; and Hoi, S. C. H. 2020. DivideMix: Learning with Noisy Labels as Semi-supervised Learning. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Li, W.; Wang, L.; Li, W.; Agustsson, E.; and Gool, L. V. 2017. WebVision Database: Visual Learning and Understanding from Web Data. *CoRR*, abs/1708.02862.
- Li, Y.; Han, H.; Shan, S.; and Chen, X. 2023. DISC: Learning from Noisy Labels via Dynamic Instance-Specific Selection and Correction. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, 24070–24079. IEEE.
- Lu, Y.; and He, W. 2022. SELC: Self-Ensemble Label Correction Improves Learning with Noisy Labels. In *International Joint Conference on Artificial Intelligence*.
- Malach, E.; and Shalev-Shwartz, S. 2017. Decoupling “when to update” from “how to update”. *Advances in Neural Information Processing Systems*, 2017-Decem: 961–971.
- Marriott, R. T.; Romdhani, S.; and Chen, L. 2021. A 3D GAN for Improved Large-Pose Facial Recognition. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, 13445–13455. Computer Vision Foundation / IEEE.
- Mikolov, T.; Chen, K.; Corrado, G.; and Dean, J. 2013. Efficient Estimation of Word Representations in Vector Space. In Bengio, Y.; and LeCun, Y., eds., *1st International Conference on Learning Representations, ICLR 2013, Scottsdale, Arizona, USA, May 2-4, 2013, Workshop Track Proceedings*.
- Sohn, K.; Berthelot, D.; Carlini, N.; Zhang, Z.; Zhang, H.; Raffel, C.; Cubuk, E. D.; Kurakin, A.; and Li, C. 2020. Fix-Match: Simplifying Semi-Supervised Learning with Consistency and Confidence. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Song, H.; Kim, M.; and Lee, J. 2019. SELFIE: Refurbishing Unclean Samples for Robust Deep Learning. volume 97 of *Proceedings of Machine Learning Research*, 5907–5915. PMLR.
- Song, H.; Kim, M.; Park, D.; and Lee, J. 2020. Learning from Noisy Labels with Deep Neural Networks: A Survey. *CoRR*, abs/2007.08199.
- Tan, C.; Xia, J.; Wu, L.; and Li, S. Z. 2021. Co-learning: Learning from Noisy Labels with Self-supervision. In Shen, H. T.; Zhuang, Y.; Smith, J. R.; Yang, Y.; Cesar, P.; Metzger, F.; and Prabhakaran, B., eds., *MM ’21: ACM Multimedia Conference, Virtual Event, China, October 20 - 24, 2021*, 1405–1413. ACM.
- Tarvainen, A.; and Valpola, H. 2017. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In Guyon, I.; von Luxburg, U.; Bengio, S.; Wallach, H. M.; Fergus, R.;

Vishwanathan, S. V. N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, 1195–1204.

Tian, Y.; Krishnan, D.; and Isola, P. 2020. Contrastive Multiview Coding. In Vedaldi, A.; Bischof, H.; Brox, T.; and Frahm, J., eds., *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part XI*, volume 12356 of *Lecture Notes in Computer Science*, 776–794. Springer.

Tu, Y.; Zhang, B.; Li, Y.; Liu, L.; Li, J.; Wang, Y.; Wang, C.; and Zhao, C. 2023. Learning from Noisy Labels with Decoupled Meta Label Purifier. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, 19934–19943. IEEE.

Wang, Y.; Ma, X.; Chen, Z.; Luo, Y.; Yi, J.; and Bailey, J. 2019. Symmetric Cross Entropy for Robust Learning With Noisy Labels. In *ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, 322–330. IEEE.

Wei, H.; Feng, L.; Chen, X.; and An, B. 2020. Combating Noisy Labels by Agreement: A Joint Training Method with Co-Regularization. In *CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, 13723–13732. Computer Vision Foundation / IEEE.

Wei, H.; Zhuang, H.; Xie, R.; Feng, L.; Niu, G.; An, B.; and Li, Y. 2023. Mitigating Memorization of Noisy Labels by Clipping the Model Prediction. In *ICML 2023*.

Wei, J.; Zhu, Z.; Cheng, H.; Liu, T.; Niu, G.; and Liu, Y. 2022. Learning with Noisy Labels Revisited: A Study Using Real-World Human Annotations. In *ICLR 2022*. OpenReview.net.

Yang, C.; An, Z.; Cai, L.; and Xu, Y. 2022. Mutual Contrastive Learning for Visual Representation Learning. In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelfth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022*, 3045–3053. AAAI Press.

Zhang, C.; Bengio, S.; Hardt, M.; Recht, B.; and Vinyals, O. 2017. Understanding deep learning requires rethinking generalization. In *ICLR 2017*. OpenReview.net.

Zhang, Z.; Chen, W.; Fang, C.; Li, Z.; Chen, L.; Lin, L.; and Li, G. 2023. RankMatch: Fostering Confidence and Consistency in Learning with Noisy Labels. In *IEEE/CVF International Conference on Computer Vision, ICCV 2023, Paris, France, October 1-6, 2023*, 1644–1654. IEEE.

Zhang, Z.; and Sabuncu, M. R. 2018. Generalized Cross Entropy Loss for Training Deep Neural Networks with Noisy Labels. In Bengio, S.; Wallach, H. M.; Larochelle, H.; Grauman, K.; Cesa-Bianchi, N.; and Garnett, R., eds., *NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, 8792–8802.