CARE: Class-Adaptive Expert Consensus for Reliable Learning with Long-Tailed Noisy Labels

Mengke Li [1] **Haiquan Ling** [2 1] **Lihao Chen** [2 1] **Yang Lu** [3] **Yiqun Zhang** [4] **Hui Huang*** [1]

Abstract

Learning from real-world data is frequently hindered by the compound challenge of long-tailed class distributions and noisy annotations. Existing methods partially address these issues but typically ignore the non-uniform impact of label noise across classes, resulting in ineffective correction for tail classes and over-regularization for head classes. To address this issue, we propose Class-Adaptive Rectification with Experts (CARE), a parameter-efficient framework that leverages three complementary supervision sources from vision-language models (VLM): observed noisy labels, VLM text embeddings, and visual features. CARE introduces a class-adaptive expert consensus mechanism that enforces stricter agreement for tail classes and more permissive agreement for head classes based on class frequency. By aggregating high-confidence predictions across these sources, CARE filters unreliable signals and recalibrates class distributions, yielding more reliable rectification under long-tailed distributions. Extensive experiments on both synthetic and real-world benchmarks demonstrate that CARE consistently outperforms state-of-the-art methods, achieving up to 3.0% performance gains. The source code is available at https://github.com/qwq123-study/CARE.

1. Introduction

Deep learning models rely heavily on large-scale datasets that are not only accurately annotated but also approximately class-balanced to ensure robust generalization (Menon et al., 2021). The rise of large-scale foundation models in recent years has further amplified this demand for high-quality labeled data (Dong et al., 2023; Li et al., 2024a), as insufficient label quality may introduce unintended prior biases during training. However, such ideal conditions are rarely met in real-world scenarios. Accurate labeling is often prohibitively time-consuming and labor-intensive, while data acquisition pipelines inherently introduce imbalance due to varying collection difficulty across classes (De Angeli et al., 2022; Hoyer et al., 2022). Rare or under-represented categories, often referred to as tail classes, are particularly difficult to collect, leading to extreme data skew (Wang et al., 2017). Consequently, real-world datasets are inherently challenged by the co-occurrence of pervasive annotation errors (noisy labels (Xiao et al., 2015)) and extreme class imbalance (long-tailed distributions (Zhang et al., 2023b)), which synergistically degrade model generalization and severely impair recognition performance, especially on tail classes. When tackled separately, long-tail learning methods assume clean labels and thus may amplify mislabeling in scarce tail classes, while noisy-label learning methods assume roughly balanced classes and tend to discard noisy samples, thereby losing valuable information in under-represented categories (Zhang et al., 2024).

To address the practical and pervasive challenge of learning from data with both label noise and long-tailed distributions, the concept of long-tailed noisy label (LTNL) learning has recently been introduced (Lu et al., 2023; Zhang et al., 2023a). Several recent studies have explored this joint setting, with some methods incorporating long-tail characteristics into the noise detection process. A common strategy involves first identifying noisy samples, followed by applying long-tailed learning techniques (Yi et al., 2022; Zhang et al., 2023a). However, such sequential pipelines critically depend on the performance of the specifically designed noise detectors, which are typically tailored to either noise robustness or long-tailed data, limiting their effectiveness in addressing both issues simultaneously. For example, (Chen et al., 2025) propose a Robust Logit Adjustment (RLA), which first corrects the noisy labels and then leverages the refurbished label distribution for logit adjustment, yet still relies on a class-agnostic label correction strategy. Moreover, most existing methods simplify the problem by

[1]College of Computer Science and Software Engineering, Shenzhen University, Shenzhen, China. [2]Guangming Laboratory, Shenzhen, China. [3]School of Informatics, Xiamen University, Xiamen, China. [4]School of Computer Science and Technology, Guangdong University of Technology, Guangzhou, China. Correspondence to: Hui Huang <hhzhiyan@gmail.com>.

Proceedings of the 43rd International Conference on Machine Learning, Seoul, South Korea. PMLR 306, 2026. Copyright 2026 by the author(s).

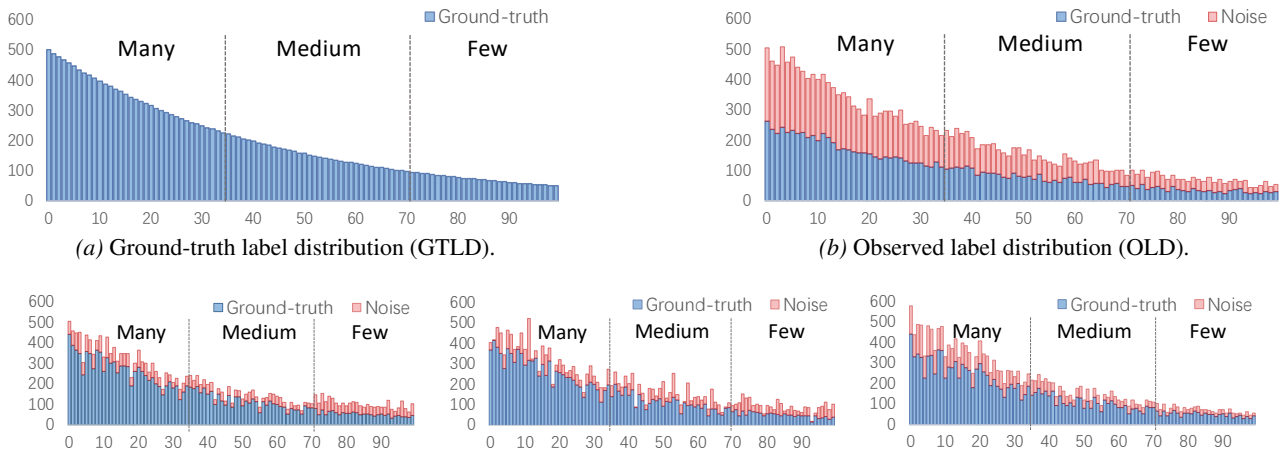

(a) Ground-truth label distribution (GTLD). *(b)* Observed label distribution (OLD).

(c) RLD 1 by class-agnostic LC. *(d)* RLD 2 by class-agnostic LC. *(e)* RLD 3 by class-aware LC.

Figure 1. Class-wise label distribution comparison on CIFAR-100. (a) The ground-truth distribution with an imbalance factor of 10. (b) The distribution under 50% joint noise. (c-e) Rectified label distributions (RLDs) obtained by different label correction (LC) methods.

Table 1. Acc. (%) comparison across different rectified labels.

Lable distribution	NR	Head	Med.	Tail	All
GTLD (Figure 1a)	0	86.2	84.9	82.9	84.8
OLD (Figure 1b)	50%	79.6	79.9	78.8	79.5
RLD 1 (Figure 1c)	24%	80.6	79.7	78.2	79.6
RLD 2 (Figure 1d)	25%	77.0	76.5	77.0	76.8
RLD 3 (Figure 1e)	28%	81.4	80.9	80.2	80.9

assuming a uniform noise rate across classes, ignoring the more realistic scenario where tail classes tend to exhibit higher noise levels. (Lu et al., 2023) take a more flexible view by considering discrepancies between the observed and intrinsic class distributions, highlighting the need for more flexible and class-aware solutions.

Although class-agnostic label refurbishment methods show promise in reducing overall noise, their effectiveness under long-tailed scenarios is limited and can even degrade performance in certain cases. In particular, we preliminarily adopted the label correction strategy from (Chen et al., 2025), generating two rectified label distributions under different hyperparameters, as illustrated in Figure 1c and Figure 1d. For comparison, we also present the ground-truth label distribution, the initial 50% joint noise distribution, and the CARE-refined class-aware label distribution in Figure 1a, Figure 1b, and Figure 1e, respectively. Fine-tuning CLIP with these three types of noisy-labeled data yields the results shown in Table 1. Despite achieving similar overall noise rates and significantly reducing label noise compared to the original noisy labels, the results reveal that insufficient correction of tail-class labels (e.g., Figure 1c) yields limited performance gains (79.6% vs. 79.5%), or even leads to degradation, as seen in Figure 1d vs. Figure 1b (total accuracy 76.8% vs. 79.5%). The decline stems from inaccurate regularization introduced during long-tailed learning. This observation shows that merely lowering the noise ratio is not sufficient, and class-aware rectification is essential. In contrast, by emphasizing accurate tail label correction, RLD 3 (shown in Figure 1e) ensures that the strong calibration induced by tail classes is meaningful, while the weaker regularization from uncorrected head labels avoids over-regularization and/or over-calibration. This class-aware strategy leads to more balanced supervision and improves overall robustness.

To this end, we propose a novel framework named Class-Adaptive Rectification with Experts (CARE) for robust learning under long-tailed noisy label (LTNL) settings. CARE leverages three complementary sources of information, namely, observed labels, Vision-Language Model (VLM)-based text embeddings, and VLM-based image features, to derive class-adaptive expert consensus for robust label rectification. Without introducing additional parameters, CARE identifies low-noise candidates through adaptive class-wise Top-K voting, where larger K values are used for head classes and smaller K for tail classes to reflect their respective data abundances and reduce confirmation bias in underrepresented categories. The expert-consensus semantics are then dynamically integrated into class frequency estimation, while low-confidence predictions are filtered through a strict hierarchical filtering mechanism that enforces consistency across classes. This enables CARE to produce reliable label distributions and perform effective noise correction. Our main contributions are summarized as follows: 1) We identify that ignoring tail label noise, while significantly reducing overall noise rate, fails to improve overall performance under long-tailed distributions. 2) We propose a parameter-free CARE framework that unifies noisy label correction and long-tailed calibration for improved generalization. 3) We design a class-aware consensus mechanism that dynamically adjusts correction confi-

dence based on class frequency and inter-expert agreement. 4) Extensive experiments on both synthetic and real-world LTNL benchmarks demonstrate that CARE consistently outperforms state-of-the-art methods under various noise and imbalance settings.

2. Proposed Method

2.1. Preliminaries

Basic Notations. Given a C-class classification problem, where the training set is denoted as $\mathcal{D}_{train} = \{(x_i, \tilde{y}_i)\}_{i=1}^{N}$, with x_i representing an input image and $\tilde{y}_i \in \tilde{\mathcal{Y}}$ being the corresponding observed label. Due to annotation noise, the observed label set $\tilde{\mathcal{Y}}$ may differ from the unknown ground-truth label set $y \in \mathcal{Y}$, and such discrepancies are regarded as noisy labels. Without loss of generality, classes are indexed in descending order of sample frequency n_i, i.e., $n_1 > n_2 > \ldots > n_C$, with $n_1 \gg n_C$, characterizing the long-tailed nature of the data distribution. Throughout this paper, subscripts are used to indicate sample or class indices.

Problem Analysis. In long-tailed noisy-label learning, the combination of sample scarcity and annotation noise severely degrades tail-class performance (Sheng et al., 2024b). Existing methods often apply class-agnostic corrections that fail to distinguish intrinsic uncertainty (from few samples) and actual noise, leading to over-filtering of tail data and biased learning (Shu et al., 2025; Zhang et al., 2023a; Lu et al., 2023). Moreover, long-tail rebalancing techniques may amplify the impact of mislabeled tail samples (Menon et al., 2021; Li et al., 2023). These issues hinder both label refinement and distribution calibration (see Appendix A for detailed analysis). We propose **CARE** that refines noisy labels by computing high-confidence consensus across multiple experts. By explicitly enhancing tail-class label reliability and disentangling class uncertainty from annotation noise, CARE enables more accurate and robust long-tail distribution calibration.

2.2. Method Overview.

CARE integrates three complementary sources of information to estimate reliable class distributions: 1) *Text Expert (TE).* A CLIP-based text encoder (Radford et al., 2021) predicts class confidences from the semantic description of each label (no additional fine-tuning). 2) *Image Expert (IE).* A CLIP-based image encoder generates visual-semantic class confidences for each input image, fine-tuned via AdaptFormer (Chen et al., 2022) for downstream tasks. 3) *Observed-label expert* is referred to as Base Expert (BE). BE utilizes the original (noisy) annotations, which nonetheless retain valuable class-indicative information. CARE aggregates class-wise confidence scores from all three experts using a relevance-weighted strategy to compute a consensus distribution. This yields a refined label distribution that simultaneously 1) down-weights unlikely noisy assignments in tail classes, 2) preserves high-confidence head-class labels, and 3) approximates the true underlying class frequencies. The resulting corrected labels exhibit substantially reduced noise rates and a distribution closer to the intrinsic (clean) data distribution. Finally, standard long-tail calibration methods, for example, logit adjustment (Menon et al., 2021) can be applied.

2.3. Composition of Parameter-Free Experts

This approach integrates signals from three distinct experts (text, image, and observed-label), leveraging their complementary strengths to identify high-confidence candidate classes, rather than relying on a single source.

Text Expert (TE). We leverage the text encoder of a pretrained VLM, such as CLIP (Radford et al., 2021), to embed each class label into a shared visual-semantic space and obtain $\{\mathbf{t}_c\}_{c=1}^{C}$. This encoder captures rich prior knowledge from large-scale pre-training and aligns textual semantics with visual representations. By computing the similarity between embedded label texts and image features, the model can estimate class-wise confidence scores without requiring additional fine-tuning. This allows TE to provide robust semantic guidance and highlight potentially correct classes, unaffected by the corrected observed labels. For an input image x[1], TE provides the confidence vector $\mathbf{p}^{\text{TE}} = \{p_c^{\text{TE}}\}$ based on the cosine similarity between the general image feature $\hat{\mathbf{f}}$ (extracted from the pretrained image encoder) and the corresponding text features $\mathbf{t}_c$:

$$p_c^{TE} = \frac{\exp(z_c^{\text{TE}})}{\sum_{j=1}^{C} \exp(z_j^{\text{TE}})}, \text{ with } z_c^{\text{TE}} = s \cdot \frac{\mathbf{t}_c^\top \hat{\mathbf{f}}}{\|\mathbf{t}_c\| \|\hat{\mathbf{f}}\|}, \quad (1)$$

where s is the scale parameter. The softmax-normalized similarity yields a semantic prior over class labels based on text-image alignment, without additional fine-tuning.

Image Expert (IE). The image expert is derived from the fine-tuned CLIP image encoder, which produces task-adapted visual representations $\mathbf{f}$ for each input image x. To enhance adaptability to the downstream tasks while maintaining efficiency, we adopt AdaptFormer (Chen et al., 2022) for parameter-efficient fine-tuning. It is worth noting that this fine-tuning is performed solely for the final classification purposes, and the resulting adapted image encoder is reused for label correction without introducing additional training objectives. IE provides the confidence vector $\mathbf{p}^{\text{IE}} = \{p_c^{\text{IE}}\}$ by the similarity between $\mathbf{f}$ and the class-specific weights $\mathbf{w}_c$ of the linear classifier $\mathbf{W}$:

$$p_c^{\text{IE}} = \frac{\exp(z_c^{\text{IE}})}{\sum_{j=1}^{C} \exp(z_j^{\text{IE}})}, \quad \text{where} \quad z_c^{\text{IE}} = s \cdot \frac{\mathbf{w}_c^\top \mathbf{f}}{\|\mathbf{w}_c\| \|\mathbf{f}\|}. \quad (2)$$

[1]We omit subscript notation for clarity.

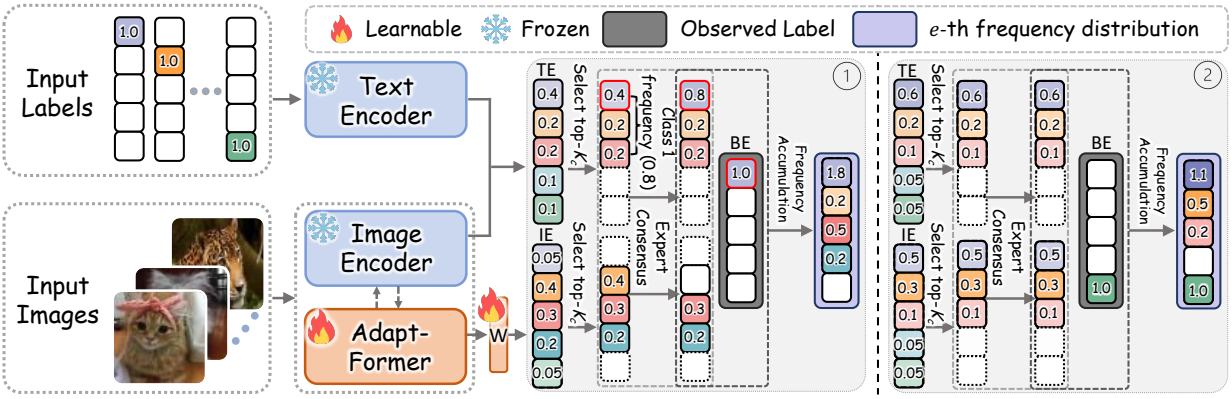

Figure 2. Overview of CARE. Different colors in the label denote classes, and intensity reflects confidence. ① illustrates expert consensus (TE and BE agree on Top-1 of TE). The Top-1 confidence of TE that is refined by removing unlikely classes is assigned to the consensus class (blue). Top-2 and Top-3 confidences, lacking consensus, are individually accumulated. ② depicts a no-consensus case, where each expert contributes confidence to their top predictions.

IE leverages the fine-tuned visual representation tailored to the target domain and serves as a byproduct of the main training pipeline. It remains parameter-free during the expert consensus phase.

Observed-Label Expert (BE) The observed-label expert, namely the base expert, relies directly on the original (potentially noisy) label annotations. Despite their unreliability, these labels still contain valuable supervision, especially for head classes that contain a large number of accurate labels. BE treats the observed label as an indicator of class assignment and serves as a complementary weak signal in the consensus-building process. The observed label $\hat{y}$ is represented as a one-hot confidence vector $\mathbf{p}^{BE} = \{p_c^{BE}\}$, where only the $\hat{y}$-th entry is set to 1.

2.4. Class Aware Expert Consensus Accumulation

To address the limitations of class-agnostic label refurbishment, particularly its inability to differentiate between head and tail class noise, we introduce a class-aware expert consensus mechanism. We adopt a consensus strategy that accumulates class-wise support across experts, effectively refining the label distribution with an awareness of class frequency.

Class Aware Expert Consensus. To evaluate the reliability of the observed label, we first identify the top-K most confident predictions from both TE and IE:

$$\mathcal{T}_K^* = \text{Top-K}(\mathbf{p}^*), \tag{3}$$

where $* \in \{\text{TE}, \text{IE}\}$, $\mathbf{p}^*$ denotes the confidence scores from the respective expert. To improve robustness against noisy labels and long-tailed class distributions, we introduce a class-adaptive Top-K consensus mechanism. The number of top predictions considered, K is adaptively determined based on the frequency of each class. Specifically, classes with higher sample counts (head classes) are assigned larger K, while those with fewer samples (tail classes) are assigned smaller K, allowing the model to allocate trust more appropriately across the class spectrum. Formally, the class-specific K is defined to be proportional to the class frequency: $K_c \propto n_c^e$, where n_c^e denotes the sample count for class c at epoch e. At the initial stage ($e = 0$), n_c^0 is computed based on the observed labels. For $e > 0$, n_c^e is updated using the corrected labels from the previous training epoch.

Class Frequency Accumulation. For each sample x and class c, the accumulated frequency is updated by:

$$F_c^{(e)} = F_c^{(e-1)} + \sum_{m \in \{\text{TE, IE, BE}\}} \alpha_m(x)\, g_m(c), \tag{4}$$

where m indexes the experts (TE, IE, BE). The class-wise contribution of expert m is:

$$g_m(c) = \mathbb{I}[c \in \mathcal{T}_K^m]\; p_c^m, \tag{5}$$

meaning that expert m contributes confidence only for classes in its Top-K list. The reliability weight assigned to expert m for sample x is:

$$\alpha_m(x) = \begin{cases} \sum_{j \in \mathcal{T}_K^m} p_j^m, & \text{if } \tilde{y} \in \mathcal{T}_K^m, \\ 1, & \text{otherwise.} \end{cases} \tag{6}$$

If the observed label $\tilde{y}$ appears in expert m's Top-K predictions, the expert is treated as reliable and its confidence reinforces $\tilde{y}$. Otherwise, only its Top-K predictions are retained to avoid reinforcing a potentially noisy label.

Noisy Label Correction and Long-Tail Calibration. In the dynamic accumulation process, the dynamically rectified label of sample x is obtained from the accumulated class frequency distribution in each round and is used for model supervision. Let $y^{r,(e)}$ be the corrected label at epoch e:

$$y^{r,(e)} = \arg\max_c \mathbf{F}_c^{(e)}. \tag{7}$$

Based on the rectified labels, we can recount the distribution of labels $n_c^{r,(e)}$ for class c that are relatively closer to the ground-truth label distribution:

$$n_c^{r,(e)} = \sum_{i=1}^{N} \mathbb{I}(y^{r,(e)} = c), \quad c = 1, 2, \ldots, C, \tag{8}$$

where $\mathbb{I}$ is the indicator function. The rectified label distribution can be used in conjunction with various long-tail learning methods, for example, the logit adjustment (LA) loss (Menon et al., 2021) for optimization:

$$\mathcal{L}_{\text{LA}}(x; y^{r,(e)} = c) = \\ -\log\left(\frac{\exp\left(z_c^{\text{IE}} + \log \widehat{\mathrm{P}}(y^{r,(e)} = c)\right)}{\sum_{j=1}^{C} \exp\left(z_j^{\text{IE}} + \log \widehat{\mathrm{P}}(y^{r,(e)} = j)\right)}\right) \tag{9}$$

where $\widehat{\mathrm{P}}(y^{r,(e)} = c)$ denotes the rectified class prior probability that is derived based on Equation (8).

The label correction for a single input image in CARE, which is performed jointly during training, is illustrated in Figure 2. The overall training algorithm is summarized in Algorithm 1 in Appendix B.

2.5. Rationale Analysis

Noise Suppression via Expert Consensus. Multi-expert Top-K consensus naturally suppresses noisy labels by leveraging agreement among experts. Since the true labels are more likely to consistently appear in Top-K predictions across each expert, while random noisy labels are not, this mechanism acts as a robust filter, similar in spirit to ensemble-based noise suppression (Lee & Chung, 2020). Formally, we state:

Theorem 1 (Reliability Amplification via Class-aware Expert Consensus)**.** *Let $\mathbf{p}^{TE}$ and $\mathbf{p}^{IE}$ be conditionally independent, and assume the expected confidence on the true class y is higher than any other class:*

$$\mathbb{E}[p_y^{TE}] > \mathbb{E}[p_c^{TE}], \quad \mathbb{E}[p_y^{IE}] > \mathbb{E}[p_c^{IE}], \quad \forall c \neq y. \tag{10}$$

Then, for a class-specific Top-K_c consensus set, the probability of both experts including y is significantly higher than for any noisy label $c \neq y$:

$$\Pr(y \in \mathcal{T}_{K_c}^{TE} \cap \mathcal{T}_{K_c}^{IE}) \gg \Pr(c \in \mathcal{T}_{K_c}^{TE} > \mathcal{T}_{K_c}^{IE}). \tag{11}$$

Proof Sketch. Since both experts assign higher expected confidence to the true label y than any noisy label $c \neq y$, the chance of y being jointly included in the Top-K_c sets of both experts is much higher than for any incorrect class. This amplifies the reliability of true labels through consensus.

A formal proof is provided in Appendix C. Theorem 1 yields the consensus-driven denoising effect:

Corollary 2 (Consensus-Based Label Refinement)**.** *By aggregating the consensus votes over training, the method constructs a frequency matrix aligned with clean labels. As training progresses (e increases), correct consensus increasingly dominates $\mathbf{F}_y^{(e)}$ in Equation (4), enabling a bootstrapped denoising that iteratively improves label quality.*

The above corollary leads to the following implications.

Remark 3. (Consensus-Based Label Refinement).

- Independent expert consensus suppresses high-confidence noisy supervisions.
- Aggregated over time, consensus acts as a robust voting that progressively corrects noisy labels.

Class-Adaptive Top-K Balances Head-Tail Confidence. CARE adaptively adjusts the number of Top-K candidates per class based on estimated class frequency. Tail classes, with limited data and unstable predictions, require stricter consensus to suppress noise-induced interference. In contrast, head classes benefit from relaxed consensus due to concentrated confidence and sufficient supervision. This design introduces the following theoretical advantage.

Proposition 4 (Tail-Class Consensus Robustness)**.** *For any tail class c with low n_t^e, selecting a smaller K_t improves the expected consensus precision compared to a global K:*

$$\mathbb{E}[\text{Precision@}K_t] > \mathbb{E}[\text{Precision@}K], \quad \textit{for } K > K_t. \tag{12}$$

The proof is provided in Appendix D. Theorem 4 yields the following practical implications.

Remark 5. (Implications of Class-Aware Consensus).

- The class-aware consensus mechanism offers theoretical robustness for tail-class prediction under noisy and imbalanced conditions.
- The adaptive K_c accounts for class imbalance, more flexibility for head classes and conservative consensus for tail classes while preventing spurious head-class candidates from dominating tail-class decisions, ensuring a discriminative and class-sensitive consensus.
- As training progresses and labels are refined, y^r becomes increasingly aligned with the true class distribution, reinforcing supervision and long-tailed bias calibration, especially for rare categories.

3. Experiment

3.1. Experimental Setup and Baseline Methods

Datasets. *Synthetic Datasets:* We evaluate CARE on both synthetic and real-world long-tailed noisy label datasets following Lu et al. (2023) and Zhang et al. (2023a). For the synthetic dataset, we construct a long-tailed noisy-label version of CIFAR-100 (Krizhevsky et al., 2009) (CIFAR-100-LTN) by subsampling its training set to follow an exponential class distribution with a predefined imbalance factor

(IF), and then injecting symmetric, asymmetric, or joint noise at a predefined noise rate (NR), following the protocol in (Zhang et al., 2023b). *Real-world Noise:* The WebVision-50 (Li et al., 2017) with inherent long-tailed distributions and noisy labels is adopted as real-world noise. Following SSBL_2 (Zhang et al., 2024), we subsample Food101N (Lee et al., 2018) into three different long-tailed noisy datasets by the IF of 20, 50, and 100 with inherent noise as real-world noise. For mini-ImageNet (Jiang et al., 2020), we create a long-tailed version via exponential subsampling and inject real-world web noise, forming ImageNet-LTN^r (abbreviated as Img-LTN^r).

Implementation Details. We adopt CLIP (Radford et al., 2021) ViT-B/16 as the backbone and Adaptformer (Chen et al., 2022) as the fine-tuning strategy. The model is trained for 20 epochs on all datasets with a batch size of 128 using SGD (initial learning rate 0.01, momentum 0.9, weight decay 5×10^{-4}). The computer resources for the experiments are provided in Appendix Q.

Compared Methods. We compare CARE with three groups of methods: *Long-tail (LT) learning*: LDAM (Cao et al., 2019), NCM (Kang et al., 2020), MiSLAS (Liu et al., 2022), LA (Menon et al., 2021), and influence-balanced loss (IB) (Park et al., 2021). *Noisy label (NL) learning*: Co-teaching (CT) (Han et al., 2018), MentorNet (Lu et al., 2018), CDR (Xia et al., 2020), ELR+ (Liu et al., 2020), MoPro (Li et al., 2021), NGC (Wu et al., 2021), Sel-CL (Li et al., 2022e) DivideMix (Li et al., 2020), and UNICON (Karim et al., 2022). *Long-tailed noisy label (LTNL) learning*: MW-Net (Shu et al., 2019), ROLT (Wei et al., 2026), HAR (Yi et al., 2022), ULC (Huang et al., 2022), H2E (Yi et al., 2022), TABASCO (Lu et al., 2023), RCAL (Zhang et al., 2023a), ECBS (Li et al., 2024c), SSBL_2 (Zhang et al., 2024) and RLA (Chen et al., 2025).

3.2. Main Comparison Results[2]

Comparison Results on Synthetic LTNL Datasets. We evaluate CARE on CIFAR-100-LTN under joint, symmetric, and asymmetric noise (see Tables 2 and 3). Standard LT or NL learning methods outperform CE, as they address class imbalance or label noise, respectively. LT methods benefit tail class learning by leveraging estimated class priors from observed labels, while NL methods mitigate noise through instance selection or loss correction. LTNL methods, tailored to address both challenges, achieve better overall results. Recent methods like RCAL, TABASCO, and RLA yield robust performance, but degrade under noise rates and severe class imbalance.

For example, RCAL drops to 30.3% under joint noise with $IF = 100$ and $NR = 50\%$, while TABASCO reaches only 50.5% under asymmetric noise ($IF = 10$, $NR = 40\%$). In contrast, CARE consistently outperforms both LT/LN methods and prior LTNL methods, especially under challenging settings. For example, on CIFAR-100-LTN, our method achieves 76.7% under joint noise ($IF = 100$, $NR = 50\%$) and 79.2% under symmetric noise ($IF = 10$, $NR = 60\%$), surpassing CLIP+LA by 1.4% and 3.2%, respectively. Compared to CLIP+RLA, CARE further improves accuracy by 10.0% and 2.2% under the same settings. It is also worth noting that while RLA is designed to enhance LA in noisy and imbalanced scenarios, it may occasionally underperform LA, particularly under severe noise conditions, possibly due to class-agnostic label refurbishment.

Comparison Results on Real-world Datasets. Table 4, Table 5, and Table 6 report the comparative results on three real-world datasets. As shown in Table 4, our method achieves the highest top-1 accuracy on Img-LTN^r under both moderate ($IF = 10$) and severe ($IF = 100$) imbalance. Notably, CARE consistently improves performance when integrated with CE or LA, surpassing existing methods such as DivideMix, ECBS, and CLIP+RLA. Besides, Appendix E analyses why CE loss outperforms the LA loss in an imbalance factor of 10. In Table 5, CARE achieves the highest accuracy (85.3%) on the WebVision test set and maintains competitive performance on ImageNet, demonstrating strong robustness to real-world noise and large-scale data. On Food101N (Table 6), our method achieves superior accuracy across all IF. In particular, it outperforms CLIP+RLA by a large margin under severe imbalance (e.g., +6.9% at $IF = 100$) and further improves upon CLIP+LA by up to +0.4%, indicating its complementary effect in refining label adjustment in noisy long-tailed scenarios. These results verify the effectiveness of CARE in enhancing noise tolerance and representation learning, especially under realistic and challenging long-tailed distributions. CARE's effectiveness on real-world noisy data (e.g., WebVision-50, Food101N) is primarily driven by two fundamental mechanisms. First, Noise-Distribution-Agnostic Rectification: Rather than estimating complex transition matrices, CARE leverages a multimodal consensus across text semantics, image features, and observed labels to naturally filter unknown noise without prior assumptions. Second, Class-Aware Thresholding: Since real-world noise disproportionately harms data-scarce tail classes, CARE applies a stricter, adaptive Top-K consensus for them. This prevents noisy head samples from dominating tail predictions and avoids over-filtering intrinsic tail uncertainty.

3.3. Further Analysis

Evolution of Class-Level Noise During Training. As illustrated in Figure 3, the class-level noise rates on the CIFAR-100-LTN dataset (50% joint noise, $IF = 10$) decrease markedly as training progresses. After the second epoch, both overall and noise rates across classes of vary-

[2]Entries marked with * denote our re-implementation with CLIP backbone; best results are highlighted in **bold**.

Table 2. Top-1 acc. (%) on CIFAR-100-LTN with joint noise. Res32 and Res18 are abbreviations for ResNet-32 and PreAct ResNet18, respectively.

Dataset	CIFAR-100-LTN									
Imbalance Factor	10					100				
Noise Ratio	0.1	0.2	0.3	0.4	0.5	0.1	0.2	0.3	0.4	0.5
CE	48.5	43.3	37.4	32.9	26.2	31.8	26.2	21.8	17.9	14.2
LDAM-DRW (Cao et al., 2019)	54.0	50.4	45.1	39.4	32.2	37.2	32.3	27.6	21.2	15.2
NCM (Kang et al., 2020)	50.8	45.2	41.3	35.4	29.3	34.9	29.5	24.7	21.8	16.8
MiSLAS(Liu et al., 2022)	57.7	53.7	50.0	46.1	40.6	41.0	37.4	32.8	27.0	21.8
Co-teaching (Han et al., 2018)	45.6	41.3	36.1	32.1	25.3	30.6	25.7	22.0	16.2	13.5
CDR (Xia et al., 2020)	47.0	40.6	35.4	30.9	24.9	27.2	25.5	22.0	17.3	13.6
Sel-CL+ (Li et al., 2022e)	55.7	53.5	50.9	47.6	44.9	37.5	36.8	35.1	32.0	28.6
RoLT (Wei et al., 2026)	54.1	51.0	47.4	44.6	38.6	35.2	31.0	27.6	24.7	20.1
RoLT-DRW (Wei et al., 2026)	55.4	52.4	49.3	46.3	40.9	37.6	32.7	30.2	26.6	21.1
HAR-DRW(Yi et al., 2022)	51.0	46.2	41.2	37.4	31.3	33.2	26.3	22.6	19.0	14.8
RCAL (Zhang et al., 2023a)	57.5	54.9	51.7	48.9	44.4	41.7	39.9	36.6	33.4	30.3
ECBS-Res32 (Li et al., 2024c)	58.4	55.6	53.8	52.8	51.2	41.6	40.7	38.5	37.1	35.5
ECBS-Res18 (Li et al., 2024c)	63.6	62.3	60.1	58.2	55.3	45.6	44.8	43.0	39.9	39.1
CLIP (zero-shot)	64.4	64.4	64.4	64.4	64.4	64.4	64.4	64.4	64.4	64.4
CLIP+CE	82.1	81.3	80.1	78.8	77.2	71.0	69.2	67.8	66.5	63.8
CLIP+CE w. CARE (ours)	82.4	81.7	81.0	79.8	78.6	71.5	70.5	69.4	67.9	65.4
CLIP+LA (Shi et al., 2024)	**84.0**	83.0	82.4	80.8	79.5	80.5	79.1	77.4	76.6	75.3
CLIP+RLA* (Chen et al., 2025)	83.9	83.0	82.5	81.4	**80.7**	74.0	72.3	71.0	69.2	66.7
CLIP+LA w. CARE (ours)	83.9	**83.5**	**83.0**	**81.9**	**80.7**	**80.7**	**79.8**	**78.5**	**77.5**	**76.7**

Table 3. Top-1 Acc. (%) on CIFAR-100-LTN with $IF = 10$ under symmetric and asymmetric noise.

Dataset	CIFAR-100-LTN			
Noise Type	Symmetric		Asymmetric	
Noise Ratio	0.4	0.6	0.2	0.4
CE	34.5	23.6	44.5	32.1
LDAM (Cao et al., 2019)	31.3	23.1	40.1	33.3
LA (Menon et al., 2021)	29.1	23.2	39.3	28.5
IB (Park et al., 2021)	32.4	25.8	45.0	35.3
DivdeMix (Li et al., 2020)	54.7	45.0	58.1	42.0
UNICON (Karim et al., 2022)	52.3	45.9	56.0	44.7
MW-Net (Shu et al., 2019)	32.0	21.7	42.5	30.4
RoLT (Wei et al., 2026)	42.0	32.6	48.2	39.3
HAR (Yi et al., 2022)	38.2	26.1	48.5	33.2
ULC (Huang et al., 2022)	54.9	44.7	54.5	43.2
TABASCO (Lu et al., 2023)	56.5	46.0	59.4	50.5
ECBS (Li et al., 2024c)	56.7	48.1	60.5	52.1
CLIP zero-shot	64.4	64.4	64.4	64.4
CLIP+CE	78.8	74.5	79.3	64.4
CLIP+CE w. CARE (ours)	80.3	77.2	80.3	67.5
CLIP+LA (Shi et al., 2024)	80.5	76.0	80.1	68.2
CLIP+RLA* (Chen et al., 2025)	80.9	77.0	80.6	69.3
CLIP+LA w. CARE (ours)	**81.5**	**79.2**	**80.8**	**70.5**

ing sample sizes drop sharply and then stabilize, offering increasingly reliable supervision. This trend provides increasingly accurate supervision for model training. Remarkably, despite having fewer clean samples, tail classes benefit greatly from CARE and achieve noise rates comparable to head classes, showing its effectiveness in correcting noisy labels even under severe imbalance.

Effectiveness of Expert Collaboration. The ablation results in Table 7 verify the individual contributions of each component (BE, TE, IE), where full integration yields a substantial performance gain (+1.6%) and significantly lowers the noise ratio from 50.0% to 27.8%. Removing either TE or IE leads to no improvement in accuracy or noise reduction, indicating that collaborative consensus among the experts is crucial. In our framework, CARE is integrated with a logit-adjustment-based long-tailed calibration loss function. Once the labels are corrected, the improved label quality makes the calibration more precise, which can further amplify this effect on head classes. Further breakdown in Table 8 shows that CARE not only reduces noise more effectively across all class splits (especially tail: from 72.8% to 53.3%), but also achieves higher accuracy for tail classes (+5.0%) compared to the baseline expert (BE). These results validate the design of CARE and its ability to enhance both label quality and generalization in long-tailed noisy settings. Interestingly, while CARE improves tail performance, head class accuracy slightly decreases. This trade-off arises from the stronger regularization effect imposed by LA on head classes, which is further reinforced by the improved label reliability of tail classes under the proposed consensus mechanism. Appendix F-O provides more complementary analysis for CARE.

Generalization of CARE across Different Backbones. For supplementary analysis, we evaluate the setting where the image expert is implemented with a ResNet backbone, and the text expert is based on GloVe embeddings (Pennington et al., 2014). The experimental results of ResNet + GloVe (Pennington et al., 2014) embeddings are presented as Table 9. Despite using less powerful models, CARE still surpasses the base model, demonstrating its generalization ability. CARE corrects noisy labels through a consensus mechanism that leverages the *discrepancy* among the predicted label distributions of different experts. Rather than

Table 4. Top-1 Acc. (%) on Img-LTN^r with NR of 0.4.

Dataset	Img-LTN^r	
Imbalance Factor	10	100
CE	31.5	31.5
DivdeMix	49.0	34.7
UNICON	41.6	31.1
MW-Net	40.3	31.1
RoLT	24.2	16.9
HAR	38.7	31.3
ULC	47.1	34.8
TABASCO	49.7	37.1
ECBS	50.8	36.9
CLIP zero-shot	54.0	54.0
CLIP+CE	83.4	80.5
CLIP+CE w. CARE	**84.6**	81.5
CLIP+LA	83.7	80.8
CLIP+RLA*	84.0	79.9
CLIP+LA w. CARE	84.4	**81.9**

Table 5. Top-1 Acc. (%) on WebVision-50 (WV50).

Train	Webvision-50	
Test	WV50	IMG12
CE	62.5	58.5
Co-teaching	63.6	61.5
MentorNet	63.0	57.8
ELR+	77.8	70.3
MoPro	77.6	76.3
NGC	79.2	74.4
Sel-CL+	80.0	76.8
RCAL+	79.6	76.3
ECBS	80.0	76.1
CLIP zero-shot	62.0	62.0
CLIP+CE	84.2	83.5
CLIP+CE w. CARE	84.7	83.6
CLIP+LA	85.1	**84.3**
CLIP+RLA*	85.0	84.2
CLIP+LA w. CARE	**85.3**	83.8

Table 6. Top-1 Acc. (%) on Food101N with raw noise.

Dataset	Food101N		
Imbalance Factor	20	50	100
LA	64.2	48.7	46.4
ELR+	58.6	49.0	44.1
DivideMix	64.0	54.7	52.6
DivideMix-LA	71.5	63.7	59.0
DivideMix-DRW	64.2	54.0	53.7
MW-Net	56.7	47.4	41.9
HAR-DRW	54.1	48.3	41.8
H2E	70.4	63.7	58.7
$SSBL_2$	74.4	69.1	63.7
CLIP zero-shot	69.0	69.0	69.0
CLIP+CE	81.3	76.9	72.6
CLIP+CE w. CARE	81.6	79.0	75.2
CLIP+LA	84.8	84.4	83.7
CLIP+RLA*	84.1	82.2	77.2
CLIP+LA w. CARE	**85.3**	**84.6**	**84.1**

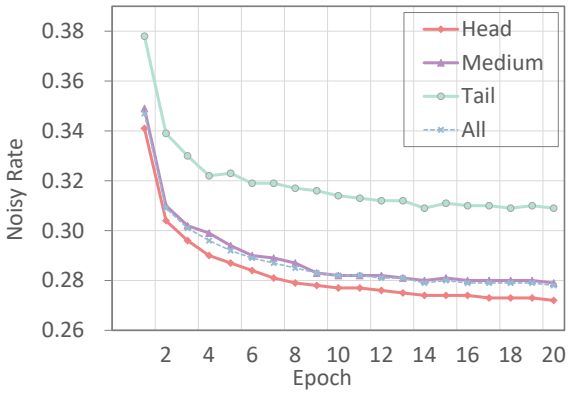

Figure 3. Noise rate dynamics during training of class splits.

Table 7. Ablation study on CIFAR-100-LTN

BE	TE	IE	NR (%)	ACC (%)
✓			50.0	78.7
✓	✓		50.0	78.7
✓		✓	50.0	78.7
✓	✓	✓	27.8	80.3

Table 8. Comparison between CARE and BE

		Head	Med.	Tail	All
BE	NR ↓ (%)	36.0	56.9	72.8	50.0
CARE	NR ↓ (%)	18.1	33.3	53.3	27.8
BE	ACC ↑ (%)	83.4	79.1	72.3	78.7
CARE	ACC ↑ (%)	81.8	81.1	77.3	80.3

depending on the absolute performance of IE or TE, CARE benefits from the diversity in their predictions, ensuring that even when the experts are relatively low-performing, as long as their outputs differ meaningfully from BE, the consensus mechanism remains effective. This robustness is also demonstrated in Table 4 to Table 6, where the zero-shot performance of CLIP (used as IE in our method) varies significantly, by up to 15% (54.0% vs. 69.0%) across two real-world datasets. Despite this large performance gap, CARE consistently achieves strong results on both datasets, maintaining accuracy above 80%, which highlights its ability to leverage expert disagreement rather than relying on expert strength alone.

In addition, to demonstrate that our framework remains highly effective using purely visual backbones, we conduct experiments with TABASCO (Lu et al., 2023), a long-tailed noisy label learning method that uses ResNet as backbone, and MLP-Mixer (Tolstikhin et al., 2021), an all-MLP architecture for vision, respectively. As shown in Table 10 and Table 11, CARE improves the performance of TABASCO and MLP-Mixer under different noise ratios. These results indicate that the performance gains are not solely attributable to the backbone but arise from the core design of CARE, which effectively integrates complementary expert signals to perform robust label correction and tail-aware calibration.

Table 12 reports the performance of the CLIP ResNet backbone in our experiments. These improvements are achieved within just 20 training epochs on noisy label data. The effectiveness of CARE does not rely on the absolute accuracy of a single expert. Instead, it is driven by the *discrepancy* among expert predictions. By comparing the cumulative frequency of predicted labels across different experts, CARE identifies and corrects potential noisy labels when there is sufficient disagreement with the base expert. Therefore, even when the auxiliary experts provide relatively inaccurate predictions, their diversity provides meaningful signals for correction.

Table 13 reports the performance of the BLIP (Li et al., 2022b) in our experiments. As shown in Table 13, applying our CARE framework consistently improves upon the BLIP baseline (w/o CARE) across various noise settings on CIFAR-100-LTN. These new results clearly demonstrate that CARE is a highly flexible, plug-and-play framework.

Table 9. Accuracy (%) Comparison of CARE (with ResNet + GloVe) on CIFAR100-LTN.

Datasets/Methods	Noise type	w/o CARE	w/ CARE
CIFAR100_IF100_NR50	Joint	45.4	47.0
CIFAR100_IF10_NR40	Symmetric	61.0	63.4
CIFAR100_IF10_NR50	Joint	58.9	61.1
CIFAR100_IF10_NR60	Symmetric	52.3	56.2

Table 10. Accuracy (%) Comparison of CARE (with TABASCO) on CIFAR100-LTN.

Datasets/Methods	Noise type	w/o CARE	w/ CARE
CIFAR100_IF10_NR20	Asymmetric	59.39	60.32
CIFAR100_IF10_NR40	Symmetric	56.52	57.20

Table 11. Accuracy (%) Comparison of CARE (with MLP-Mixer) on CIFAR100-LTN.

Datasets/Methods	Noise type	w/o CARE	w/ CARE
CIFAR100_IF10_NR40	Symmetric	76.0	77.8
CIFAR100_IF100_NR50	Joint	63.8	65.1

Table 12. Accuracy (%) Comparison of CARE (with CLIP ResNet backbone) on CIFAR100-LTN.

Datasets/Methods	Noise type	w/o CARE	w/ CARE
CIFAR100_IF100_NR50	Joint	53.3	53.8
CIFAR100_IF10_NR20	Asymmetric	63.7	64.1
CIFAR100_IF10_NR40	Asymmetric	52.4	53.2
CIFAR100_IF10_NR60	Symmetric	56.5	57.5

Table 13. Accuracy (%) Comparison of CARE (with BLIP) on CIFAR100-LTN.

Datasets/Methods	Noise type	w/o CARE	w/ CARE
CIFAR100_IF10_NR40	Symmetric	81.4	82.0
CIFAR100_IF10_NR40	Asymmetric	68.0	68.9
CIFAR100_IF10_NR60	Symmetric	77.7	79.2

Its class-aware multi-expert consensus mechanism is not restricted to CLIP, but can effectively integrate with and enhance other advanced VLMs to robustly tackle long-tailed noisy labels.

4. Related Work

Noisy Label Learning. Noisy label learning can generally be categorized into three groups: noisy sample separation, label correction, and noise regularization. To separate noisy samples, various metrics have been explored, including loss values (Lu et al., 2018; Han et al., 2018; Li et al., 2020), logit values (Pleiss et al., 2020), Jensen-Shannon divergence (Yao et al., 2021; Karim et al., 2022), and the distance between sample features and class prototypes (Ji et al., 2024). For label correction, many methods aim to revise noisy annotations. For example, partial label learning (Sheng et al., 2024a; Wang et al., 2022; Feng & An, 2019) assumes each training instance is associated with a candidate label set that contains the true label. PLS (Albert et al., 2023) further incorporates the confidence of the corrected labels. Other methods estimate the noise transition matrix (Yao et al., 2020; Cheng et al., 2022; Lin et al., 2024) to infer the label most likely to be correct. In terms of noise regularization, several strategies have been proposed to reduce the influence of noisy labels. These include reweighting training examples (Ren et al., 2018; Liu & Tao, 2015) and designing noise-robust loss functions, such as backward and forward loss correction (Patrini et al., 2017), gold loss correction (Hendrycks et al., 2018), MW-Net (Shu et al., 2019), and Dual-T (Yao et al., 2020).

Long-Tailed Noisy Label Learning. Several recent studies have been proposed to address the compounded challenges arising from the coexistence of noisy labels and unbalanced/long-tailed data. A common strategy is to identify and separate clean and noisy samples. For example, CNLCU (Xia et al., 2022) enhances the small loss method (Han et al., 2018) by designating a subset of samples with large losses as clean. TABASCO (Lu et al., 2023) addresses a complex scenario where noisy labels may cause an intrinsic tail class to appear as a head class, proposing a bi-dimensional separation metric to adapt to different cases. SSBL_2 (Zhang et al., 2024) separates noisy samples based on model prediction confidence and loss values. Another line of work focuses on emphasizing the importance of different samples by reweighting or regularization (Ren et al., 2018; Shu et al., 2019; Cao et al., 2021; Wei et al., 2022; Jiang et al., 2022; Chen et al., 2025). Meanwhile, improving representation learning (Zhou et al., 2022; Yi et al., 2022; Zhang et al., 2023a; Li et al., 2024c) explores intrinsic feature-space structures to mitigate noise propagation. Despite their merits, these methods underutilize the observed labels that implicitly contain weak but valuable information about the underlying ground-truth. Additional discussion of long-tailed learning is provided in Appendix I.

A discussion towards limitations of CARE and directions for future work is provided in Appendix P.

5. Conclusion

In this paper, we have demonstrated that simply rectifying noisy labels in a class-agnostic way does not guarantee improved learning for tail classes, nor does it necessarily enhance overall model performance. The challenge of correcting noisy labels in tail classes is even more pronounced due to their inherent sparsity and underrepresentation. To address this, we propose Class-Adaptive Rectification with Experts (CARE), which leverages a class-aware expert-consensus strategy to effectively rectify noisy labels across all classes. By ensuring more reliable supervision, CARE enables the model to better capture representative features for each category, ultimately boosting overall performance.

Acknowledgements

This work was supported in part by National Key R&D Program of China (2024YFB3908500, 2024YFB3908502), NSFC (62306181), Guangdong Basic and Applied Basic Research Foundation (2024A1515010163), Shenzhen S&T Program (RCBS20231211090659101), and National Key Lab of Radar Signal Processing (JKW202403).

Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. Our research focuses on enhancing the robustness of models under long-tailed and noisy label distributions. By enabling effective learning from imperfect, real-world data, this work has the potential to improve the reliability of visual recognition systems in practical settings where high-quality data curation is expensive or infeasible. We do not foresee any specific negative societal consequences or ethical issues arising directly from this work, but we encourage responsible use of robust learning frameworks in safety-critical applications.

References

Albert, P., Arazo, E., Krishna, T., O'Connor, N. E., and McGuinness, K. Is your noise correction noisy? PLS: robustness to label noise with two stage detection. In *WACV*, pp. 118–127, 2023.

Baik, J. S., Yoon, I. Y., Kim, K. H., and Choi, J. W. Distribution-aware robust learning from long-tailed data with noisy labels. In *ECCV*, volume 15072, pp. 160–177, 2024.

Cai, J., Wang, Y., and Hwang, J.-N. Ace: Ally complementary experts for solving long-tailed recognition in one-shot. In *ICCV*, pp. 112–121, 2021.

Cao, K., Wei, C., Gaidon, A., Arechiga, N., and Ma, T. Learning imbalanced datasets with label-distribution-aware margin loss. In *NeurIPS*, pp. 1567–1578, 2019.

Cao, K., Chen, Y., Lu, J., Aréchiga, N., Gaidon, A., and Ma, T. Heteroskedastic and imbalanced deep learning with adaptive regularization. In *ICLR*, 2021.

Chen, M., Du, Y., Jiang, W., Zhang, B., Feng, S., Xin, Y., and Wang, C. Robust logit adjustment for learning with long-tailed noisy data. In *AAAI*, pp. 15830–15838, 2025.

Chen, S., Ge, C., Tong, Z., Wang, J., Song, Y., Wang, J., and Luo, P. Adaptformer: Adapting vision transformers for scalable visual recognition. In *NeurIPS*, volume 35, pp. 16664–16678, 2022.

Cheng, D., Ning, Y., Wang, N., Gao, X., Yang, H., Du, Y., Han, B., and Liu, T. Class-dependent label-noise learning with cycle-consistency regularization. In *NeurIPS*, 2022.

Cubuk, E. D., Zoph, B., Mané, D., Vasudevan, V., and Le, Q. V. Autoaugment: Learning augmentation strategies from data. In *CVPR*, pp. 113–123, 2019.

Cubuk, E. D., Zoph, B., Shlens, J., and Le, Q. V. Randaugment: Practical automated data augmentation with a reduced search space. In *CVPRW*, pp. 3008–3017, 2020.

Cui, J., Liu, S., Tian, Z., Zhong, Z., and Jia, J. Reslt: Residual learning for long-tailed recognition. *IEEE TPAMI*, 45 (3):3695–3706, 2023.

Cui, Y., Jia, M., Lin, T.-Y., Song, Y., and Belongie, S. Class-balanced loss based on effective number of samples. In *CVPR*, pp. 9268–9277, 2019.

De Angeli, K., Gao, S., Danciu, I., Durbin, E. B., Wu, X.-C., et al. Class imbalance in out-of-distribution datasets: Improving the robustness of the textcnn for the classification of rare cancer types. *JBI*, 125:103957, 2022.

Dong, B., Zhou, P., Yan, S., and Zuo, W. LPT: long-tailed prompt tuning for image classification. In *ICLR*, 2023.

Feng, L. and An, B. Partial label learning with self-guided retraining. In *AAAI*, pp. 3542–3549, 2019.

Foret, P., Kleiner, A., Mobahi, H., and Neyshabur, B. Sharpness-aware minimization for efficiently improving generalization. In *ICLR*, 2021.

Han, B., Yao, Q., Yu, X., Niu, G., Xu, M., Hu, W., Tsang, I., and Sugiyama, M. Co-teaching: Robust training of deep neural networks with extremely noisy labels. In *NeurIPS*, volume 31, 2018.

Hendrycks, D., Mazeika, M., Wilson, D., and Gimpel, K. Using trusted data to train deep networks on labels corrupted by severe noise. In *NeurIPS*, volume 31, 2018.

Hoyer, L., Dai, D., and Gool, L. V. Daformer: Improving network architectures and training strategies for domain-adaptive semantic segmentation. In *CVPR*, pp. 9914–9925, 2022.

Huang, Y., Bai, B., Zhao, S., Bai, K., and Wang, F. Uncertainty-aware learning against label noise on imbalanced datasets. In *AAAI*, volume 36, pp. 6960–6969, 2022.

Ji, X., Zhu, Z., Xi, W., Gadyatskaya, O., Song, Z., Cai, Y., and Liu, Y. Fedfixer: Mitigating heterogeneous label noise in federated learning. In *AAAI*, pp. 12830–12838, 2024.

Jiang, L., Huang, D., Liu, M., and Yang, W. Beyond synthetic noise: Deep learning on controlled noisy labels. In *ICML*, volume 119, pp. 4804–4815, 2020.

Jiang, S., Li, J., Wang, Y., Huang, B., Zhang, Z., and Xu, T. Delving into sample loss curve to embrace noisy and imbalanced data. *AAAI*, 36:7024–7032, 2022.

Kang, B., Xie, S., Rohrbach, M., Yan, Z., Gordo, A., Feng, J., and Kalantidis, Y. Decoupling representation and classifier for long-tailed recognition. In *ICLR*, 2020.

Karim, N., Rizve, M., Rahnavard, N., Mian, A., and Shah, M. Unicon: Combating label noise through uniform selection and contrastive learning. In *CVPR*, pp. 9666–9676, 2022.

Krizhevsky, A., Hinton, G., et al. Learning multiple layers of features from tiny images. *Technical Report*, 2009.

Lee, J. and Chung, S. Robust training with ensemble consensus. In *ICLR*, 2020.

Lee, K., He, X., Zhang, L., and Yang, L. Cleannet: Transfer learning for scalable image classifier training with label noise. In *CVPR*, pp. 5447–5456, 2018.

Li, B., Han, Z., Li, H., Fu, H., and Zhang, C. Trustworthy long-tailed classification. In *CVPR*, pp. 6970–6979, 2022a.

Li, J., Socher, R., and Hoi, S. C. Dividemix: Learning with noisy labels as semi-supervised learning. In *ICLR*, 2020.

Li, J., Xiong, C., and Hoi, S. C. H. Mopro: Webly supervised learning with momentum prototypes. In *ICLR*, 2021.

Li, J., Li, D., Xiong, C., and Hoi, S. C. H. BLIP: bootstrapping language-image pre-training for unified vision-language understanding and generation. In *ICML*, pp. 12888–12900, 2022b.

Li, J., Tan, Z., Wan, J., Lei, Z., and Guo, G. Nested collaborative learning for long-tailed visual recognition. In *CVPR*, pp. 6949–6958, 2022c.

Li, M. *Advances in Long-Tailed Visual Recognition*. PhD thesis, Hong Kong Baptist University, 2022.

Li, M., Cheung, Y.-m., and Lu, Y. Long-tailed visual recognition via gaussian clouded logit adjustment. In *CVPR*, pp. 6929–6938, June 2022d.

Li, M., Cheung, Y.-m., and Hu, Z. Key point sensitive loss for long-tailed visual recognition. *IEEE TPAMI*, 45(4): 4812–4825, 2023.

Li, M., Liu, Y., Lu, Y., Zhang, Y., Cheung, Y.-m., and Huang, H. Improving visual prompt tuning by gaussian neighborhood minimization for long-tailed visual recognition. In *NeurIPS*, volume 37, pp. 103985–104009, 2024a.

Li, M., Zhikai, H., Lu, Y., Lan, W., Cheung, Y.-m., and Huang, H. Feature fusion from head to tail for long-tailed visual recognition. In *AAAI*, volume 38, pp. 13581–13589, 2024b.

Li, S., Xia, X., Ge, S., and Liu, T. Selective-supervised contrastive learning with noisy labels. In *CVPR*, pp. 316–325, 2022e.

Li, W., Wang, L., Li, W., Agustsson, E., and Van Gool, L. Webvision database: Visual learning and understanding from web data. *arXiv preprint arXiv:1708.02862*, 2017.

Li, Y., Li, Z., and Jia, Y. Boosting class representation via semantically related instances for robust long-tailed learning with noisy labels. In *ICCV*, pp. 1516–1525, 2025.

Li, Z., Zhao, H., Li, Z., Liu, T., Guo, D., and Wan, X. Extracting clean and balanced subset for noisy long-tailed classification. *arXiv*, 2024c.

Lienen, J. and Hüllermeier, E. Mitigating label noise through data ambiguation. In *AAAI*, pp. 13799–13807, 2024.

Lin, Y., Yao, Y., and Liu, T. Learning the latent causal structure for modeling label noise. In *NeurIPS*, volume 37, pp. 120549–120577, 2024.

Liu, S., Niles-Weed, J., Razavian, N., and Fernandez-Granda, C. Early-learning regularization prevents memorization of noisy labels. In *NeurIPS*, volume 33, pp. 20331–20342, 2020.

Liu, T. and Tao, D. Classification with noisy labels by importance reweighting. *IEEE TPAMI*, 38(3):447–461, 2015.

Liu, Y., Cao, B., and Fan, J. Improving the accuracy of learning example weights for imbalance classification. In *ICLR*, 2022.

Lu, J., Zhou, Z., Leung, T., Li, L.-J., and Li, F.-F. Mentornet: Learning data-driven curriculum for very deep neural networks on corrupted labels. In *ICML*, pp. 2304–2313, 2018.

Lu, Y., Zhang, Y., Han, B., Cheung, Y.-m., and Wang, H. Label-noise learning with intrinsically long-tailed data. In *ICCV*, pp. 1369–1378, 2023.

Menon, A. K., Jayasumana, S., Rawat, A. S., Jain, H., Veit, A., and Kumar, S. Long-tail learning via logit adjustment. In *ICLR*, 2021.

Park, S., Lim, J., Jeon, Y., and Choi, J. Y. Influence-balanced loss for imbalanced visual classification. In *ICCV*, pp. 735–744, 2021.

Patrini, G., Rozza, A., Menon, A. K., Nock, R., and Qu, L. Making deep neural networks robust to label noise: A loss correction approach. In *CVPR*, pp. 2233–2241, 2017.

Pennington, J., Socher, R., and Manning, C. D. Glove: Global vectors for word representation. In *EMNLP*, pp. 1532–1543, 2014.

Pleiss, G., Zhang, T., Elenberg, E., and Weinberger, K. Q. Identifying mislabeled data using the area under the margin ranking. In *NeurIPS*, volume 33, pp. 17044–17056, 2020.

Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al. Learning transferable visual models from natural language supervision. In *ICLR*, pp. 8748–8763, 2021.

Ren, J., Yu, C., sheng, s., Ma, X., Zhao, H., Yi, S., and Li, h. Balanced meta-softmax for long-tailed visual recognition. In *NeurIPS*, volume 33, pp. 4175–4186, 2020.

Ren, M., Zeng, W., Yang, B., and Urtasun, R. Learning to reweight examples for robust deep learning. In *ICML*, volume 80, pp. 4331–4340, 2018.

Sheng, M., Sun, Z., Cai, Z., Chen, T., Zhou, Y., and Yao, Y. Adaptive integration of partial label learning and negative learning for enhanced noisy label learning. In *AAAI*, pp. 4820–4828, 2024a.

Sheng, M., Sun, Z., Chen, T., Pang, S., Wang, Y., and Yao, Y. Foster adaptivity and balance in learning with noisy labels. In *ECCV*, pp. 217–235. Springer, 2024b.

Shi, J.-X., Wei, T., Zhou, Z., Shao, J.-J., Han, X.-Y., and Li, Y.-F. Long-tail learning with foundation model: Heavy fine-tuning hurts. In *ICML*, 2024.

Shu, C., Li, M., Zhang, Y., Lu, Y., Han, B., Cheung, Y.-m., and Wang, H. Classifying long-tailed and label-noise data via disentangling and unlearning. *arXiv preprint arXiv:2503.11414*, 2025.

Shu, J., Xie, Q., Yi, L., Zhao, Q., Zhou, S., Xu, Z., and Meng, D. Meta-weight-net: Learning an explicit mapping for sample weighting. In *NeurIPS*, pp. 1917–1928, 2019.

Tolstikhin, I. O., Houlsby, N., Kolesnikov, A., Beyer, L., Zhai, X., Unterthiner, T., Yung, J., Steiner, A., Keysers, D., Uszkoreit, J., Lucic, M., and Dosovitskiy, A. Mlp-mixer: An all-mlp architecture for vision. In *NeurIPS*, pp. 24261–24272, 2021.

Wang, B., Wang, P., Xu, W., Wang, X., Zhang, Y., Wang, K., and Wang, Y. Kill two birds with one stone: Rethinking data augmentation for deep long-tailed learning. In *ICLR*, 2024a.

Wang, D., Zhang, M., and Li, L. Adaptive graph guided disambiguation for partial label learning. *IEEE TPAMI*, 44(12):8796–8811, 2022.

Wang, P., Zhao, Z., Wen, H., Wang, F., Wang, B., Zhang, Q., and Wang, Y. Llm-autoda: Large language model-driven automatic data augmentation for long-tailed problems. In *NeurIPS*, 2024b.

Wang, X., Lian, L., Miao, Z., Liu, Z., and Yu, S. X. Long-tailed recognition by routing diverse distribution-aware experts. In *ICLR*, 2021.

Wang, Y.-X., Ramanan, D., and Hebert, M. Learning to model the tail. In *NeurIPS*, volume 30, 2017.

Wei, T., Shi, J.-X., Li, Y.-F., and Zhang, M.-L. Prototypical classifier for robust class-imbalanced learning. In *PAKDD*, pp. 44–57, 2022.

Wei, T., Shi, J.-X., Tu, W.-W., and Li, Y.-F. Robust long-tailed learning under label noise. *Front. Comput. Sci.*, 20 (1):2001321, 2026.

Wu, Z.-F., Wei, T., Jiang, J., Mao, C., Tang, M., and Li, Y. Ngc: A unified framework for learning with open-world noisy data. In *ICCV*, pp. 62–71, 2021.

Xia, X., Liu, T., Han, B., Gong, C., Wang, N., Ge, Z., and Chang, Y. Robust early-learning: Hindering the memorization of noisy labels. In *ICLR*, 2020.

Xia, X., Liu, T., Han, B., Gong, M., Yu, J., Niu, G., and Sugiyama, M. Sample selection with uncertainty of losses for learning with noisy labels. In *ICLR*, 2022.

Xiao, T., Xia, T., Yang, Y., Huang, C., and Wang, X. Learning from massive noisy labeled data for image classification. In *CVPR*, pp. 2691–2699, 2015.

Yao, Y., Liu, T., Han, B., Gong, M., Deng, J., Niu, G., and Sugiyama, M. Dual t: reducing estimation error for transition matrix in label-noise learning. In *NeurIPS*, pp. 7260–7271, 2020.

Yao, Y., Sun, Z., Zhang, C., Shen, F., Wu, Q., Zhang, J., and Tang, Z. Jo-SRC: A contrastive approach for combating noisy labels. In *CVPR*, pp. 5188–5197, 2021.

Yi, X., Tang, K., Hua, X.-S., Lim, J.-H., and Zhang, H. Identifying hard noise in long-tailed sample distribution. In *ECCV*, pp. 739–756, 2022.

Zhang, L., Tian, Z., Zhou, W., and Wang, W. Learning from long-tailed noisy data with sample selection and balanced loss. In *IJCAI*, pp. 5471–5480, 2024.

Zhang, M., Zhao, X., Yao, J., Yuan, C., and Huang, W. When noisy labels meet long tail dilemmas: A representation calibration method. In *ICCV*, pp. 15844–15854, 2023a.

Zhang, S., Li, Z., Yan, S., He, X., and Sun, J. Distribution alignment: A unified framework for long-tail visual recognition. In *CVPR*, pp. 2361–2370, 2021.

Zhang, Y., Kang, B., Hooi, B., Yan, S., and Feng, J. Deep long-tailed learning: A survey. *IEEE TPAMI*, 45(9): 10795–10816, 2023b.

Zhao, Z., Wang, P., Wen, H., Xu, W., Lai, S., Zhang, Q., and Wang, Y. Two fists, one heart: Multi-objective optimization based strategy fusion for long-tailed learning. In *ICML*, pp. 61040–61071, 2024.

Zhao, Z., Wen, H., Wang, P., Wang, S., Wang, Z., Zhang, Q., and Wang, Y. Balancing model efficiency and performance: Adaptive pruner for long-tailed data. In *ICML*, pp. 77723–77739, 2025.

Zhong, Z., Cui, J., Liu, S., and Jia, J. Improving calibration for long-tailed recognition. In *CVPR*, pp. 16489–16498, 2021.

Zhou, B., Cui, Q., Wei, X.-S., and Chen, Z.-M. BBN: Bilateral-branch network with cumulative learning for long-tailed visual recognition. In *CVPR*, pp. 9719–9728, 2020.

Zhou, X., Liu, X., Zhai, D., Jiang, J., Gao, X., and Ji, X. Prototype-anchored learning for learning with imperfect annotations. In *ICML*, volume 162, pp. 27245–27267, 2022.

Zhou, Y., Qu, Y., Xu, X., and Shen, H. Imbsam: A closer look at sharpness-aware minimization in class-imbalanced recognition. In *ICCV*, pp. 11345–11355, 2023a.

Zhou, Z., Li, L., Zhao, P., Heng, P.-A., and Gong, W. Class-conditional sharpness-aware minimization for deep long-tailed recognition. In *CVPR*, pp. 3499–3509, 2023b.

Appendix: Class-Adaptive Expert Consensus for Reliable Learning with Long-Tailed Noisy Labels

A. Detailed Problem Analysis for Long-Tailed Noisy-Label Learning

In long-tailed noisy-label learning, performance degradation on tail classes arises from the compounding effects of *intrinsic uncertainty* due to sample scarcity and *extrinsic uncertainty* caused by label noise. We denote the observed label $\tilde{y}$ as a noisy version of the ground-truth label y, sampled from a corruption distribution $\eta(y \mid \tilde{y})$. Due to long-tailed distribution, the number of samples per class n_c varies greatly, with $n_c \ll n_h$ for tail ($c \in \mathcal{T}$) and head ($h \in \mathcal{H}$) classes, respectively.

We summarize the core challenges as follows:

Class-Agnostic Label Correction. Take the label correction methods (Han et al., 2018; Li et al., 2020; Lienen & Hüllermeier, 2024; Baik et al., 2024) as an example, noise-robust methods rely on a global confidence threshold τ:

$$\mathcal{D}_{\text{clean}} = \{(x_i, \tilde{y}_i) \mid \max_c p(c \mid x_i) > \tau\}. \tag{13}$$

However, for tail classes ($c \in \mathcal{T}$), the estimated confidence $p(c \mid x)$ tends to be lower due to fewer samples:

$$\mathbb{E}[p(c \mid x)] \propto n_c. \tag{14}$$

Thus, using a class-agnostic τ results in the disproportionate rejection of valid tail samples.

Confusion Between Uncertainty and Noise. Let $\sigma_c^2 = \text{Var}[p(c \mid x)]$ denote prediction variance for class c. Tail classes often have higher σ_c^2, yet this uncertainty is mistaken as noise:

$$\text{Uncertainty}(x) = 1 - \max_c p(c \mid x). \tag{15}$$

High Uncertainty(x) does not always imply label corruption for tail samples. This leads to erroneous filtering:

$$\Pr[\tilde{y}_i = y_i \mid \text{Uncertainty}(x_i) > \delta] \gg 0, \quad \text{for } c = \tilde{y}_i \in \mathcal{T}. \tag{16}$$

Amplified Mislabeling Impact via Rebalancing. Reweighting or logit adjustment methods amplify loss from under-represented classes:

$$\mathcal{L}(x_i, \tilde{y}_i) = \alpha_c \cdot \ell(f(x_i), \tilde{y}_i), \quad \alpha_c \propto \frac{1}{n_c}. \tag{17}$$

However, since noisy label rate ρ_c is typically higher for tail classes:

$$\rho_c = \Pr(\tilde{y} \neq y \mid y = c), \quad \rho_c \uparrow \text{ as } n_c \downarrow, \tag{18}$$

which leads to over-calibration driven by incorrect labels from tail classes.

Summary. Due to the imbalance in confidence, variance, and noise rate, class-agnostic correction and rebalancing strategies result in:

1) Over-filtering of clean tail samples,
2) Retention of confident noisy head labels,
3) Amplification of noise-induced gradient bias in tail classes.

To address this, our proposed CARE computes a class-specific confidence consensus that leverages semantic priors and adjusts to category frequency, thereby disentangling class uncertainty from annotation noise and enabling effective distribution refinement.

B. Algorithm for CARE

The training procedure for CARE is outlined in Algorithm 1.

Algorithm 1 Class-Aware Rectification with Experts

Require: Training set $\mathcal{D} = \{x_i, \tilde{y}_i\}$, pre-trained model
Ensure: Fine-tuned model
1: Initialize $\{y_i^{r,(0)}\}_{i=1}^N \leftarrow \{\tilde{y}_i\}_{i=1}^N$, candidate count vector $\mathbf{F} \in \mathbb{R}^{N\times C} \leftarrow \{\mathbf{p}_i^{BE}\}_{i=1}^N$, class frequency $n_c^{r,(0)} = \sum_{i=1}^N \delta(\tilde{y}_i = c)$, fine-turning module ϕ.
2: **for** $e = 1$ to E **do**
3: $K_c \leftarrow f(n_c^e)$, where $f(*)$ is monotonically increasing function , for example, $f(x) = \lfloor (x)^{1/4} \rfloor$.
4: $\mathcal{T}_i^{TE,(e)} \leftarrow$ Top-K$(\mathbf{p}_i^{\text{IE}}|K_c)$, $\mathcal{T}_i^{IE,(e)} \leftarrow$ Top-K$(\mathbf{p}_i^{\text{TE}}|K_c)$
5: Update $\mathbf{F}_i^{(e)}$ for x_i according to Equation (4)
6: $y_i^{r,(e)} \leftarrow \arg\max_c \mathbf{F}_i^{(e)}$
7: Update $n_c^{r,(e)}$ according to Equation (8)
8: Compute the final loss $\mathcal{L}$ using $\{x_i, y_i^{r,(e)}\}$ and $\{n_c^{r,(e)}\}_{c=1}^C$.
9: Update ϕ by $\phi^{e+1} = \phi^e - \alpha \cdot \nabla\mathcal{L}$
10: **end for**

C. Proof of Theorem 1

Proof. We decompose the probability of consensus on class y as follows:

$$\Pr(y \in \mathcal{T}_{K_c}^{\text{TE}} \cap \mathcal{T}_{K_c}^{\text{IE}}) = \Pr(y \in \mathcal{T}_{K_c}^{\text{TE}}) \cdot \Pr(y \in \mathcal{T}_{K_c}^{\text{IE}}). \tag{19}$$

Assuming $\mathbf{p}^{\text{TE}}$ and $\mathbf{p}^{\text{IE}}$ are conditionally independent. Under the assumption that y has the highest expected confidence:

$$\mathbb{E}[\text{rank}(y)] < \mathbb{E}[\text{rank}(c)], \quad \forall c \neq y. \tag{20}$$

Since $K_c \propto n_c^e$, frequent (head) classes are allowed higher K_c, which increases:

$$\Pr(y \in \mathcal{T}_{K_c}^{*}) = \sum_{k=1}^{K_c} \Pr[\text{rank}(y) = k]. \tag{21}$$

Top-K inclusion probability increases with K_c. For a noisy label $c \neq y$, the probability that both experts place it in Top-K_c is:

$$\Pr(c \in \mathcal{T}_{K_c}^{\text{TE}}) \cdot \Pr(c \in \mathcal{T}_{K_c}^{\text{IE}}) \ll \Pr(y \in \mathcal{T}_{K_c}^{\text{TE}}) \cdot \Pr(y \in \mathcal{T}_{K_c}^{\text{IE}}), \tag{22}$$

as c is not favored in expectation. Consequently,

$$\Pr(y \in \mathcal{T}_{K_c}^{\text{TE}} \cap \mathcal{T}_{K_c}^{\text{IE}}) \gg \Pr(c \in \mathcal{T}_{K_c}^{\text{TE}} \cap \mathcal{T}_{K_c}^{\text{IE}}). \tag{23}$$

□

D. Proof of Proposition 4

Proof. We aim to show that for any tail class t with low sample frequency n_t^e at epoch e, using a smaller class-specific Top-K value K_t yields higher consensus precision compared to using a larger K.

Notation

- $\mathcal{D}_t$: the set of all samples belonging to tail class t.

- E: the set of experts.

- $\mathcal{R}_m(x, K)$: Top-K classes predicted by expert m for input x, i.e.,

$$\mathcal{R}_m(x, K) := \texttt{TopK}(\mathbf{z}_m).$$

Define

- $\mathcal{P}_t(K)$: Total number of times class t appears in the Top-K predictions across all experts and all samples.

- $\mathcal{T}_t(K)$: Total number of correct Top-K predictions for class t, i.e., predictions on samples from $\mathcal{D}_t$ where $t \in \mathcal{R}_m(x, K)$.

Then the **consensus precision** for class t under Top-K is defined as:

$$\text{ConsensusPrecision}_t(K) := \frac{\mathcal{T}_t(K)}{\mathcal{P}_t(K)}.$$

1. Larger K increases false positives for tail class t.

Let $n_t^e = |\mathcal{D}_t|$ denote the number of samples of class t at epoch e. Then,

$$\mathcal{T}_t(K) = \sum_{x \in \mathcal{D}_t} \sum_{m \in E} \mathbf{1}[t \in \mathcal{R}_m(x, K)].$$

For each sample from class t, at most $|E|$ experts can correctly include t in their Top-K, as a result:

$$\mathcal{T}_t(K) \leq |E| \cdot |\mathcal{D}_t|.$$

This term is upper-bounded and cannot increase indefinitely. For $\mathcal{P}_t(K)$,

$$\begin{aligned}
\mathcal{P}_t(K) &= \sum_{x \in \mathcal{D}_t} \sum_{m \in E} \mathbf{1}\big[t \in \mathcal{R}_m(x, K)\big] \\
&= \sum_{x \in \mathcal{D}_t} \sum_{m \in E} \mathbf{1}[t \in \mathcal{R}_m(x, K)] + \sum_{x \notin \mathcal{D}_t} \sum_{m \in E} \mathbf{1}[t \in \mathcal{R}_m(x, K)] \\
&= \mathcal{T}_t(K) + \sum_{x \notin \mathcal{D}_t} \sum_{m \in E} \mathbf{1}[t \in \mathcal{R}_m(x, K)].
\end{aligned}$$

The second term accounts for false positives contributed by non-t samples. For a tail class t, the number of non-t samples is much larger than that of t-samples, i.e., $|\mathcal{D} \setminus \mathcal{D}_t| \gg |\mathcal{D}_t|$. The probability that t appears in the Top-K of a non-t sample increases with K. Consequently, as K grows, the contribution of false positives $\sum_{x \notin \mathcal{D}_t} \sum_{m \in E} \mathbb{P}[t \in \mathcal{R}_m(x, K)]$ increases, while the number of correct predictions $\mathcal{T}_t(K)$ remains bounded by $|E| \cdot n_t^e$. Then we have

$$\frac{d\mathcal{P}_t(K)}{dK} = \sum_{x \notin \mathcal{D}_t} \sum_{m \in E} \frac{\mathbf{1}[t \in \mathcal{R}_m(x, K)]}{dK} > \frac{d\mathcal{T}_t(K)}{dK} = 0,$$

Therefore, for tail classes with small n_t^e, increasing K inflates $\mathcal{P}_t(K)$ faster than $\mathcal{T}_t(K)$.

Additional Observation: Saturation of correct predictions vs. growth of false positives.

In contrast, as K increases, the number of samples from other classes that mistakenly include t in their Top-K grows significantly, leading to an increase in: $\mathcal{P}_t(K) - \mathcal{T}_t(K)$.

2. Smaller K suppresses the inclusion of t on non-t samples.

Reducing K limits the number of classes per prediction, which reduces the chance of incorrectly including rare class t in the predictions for other classes. This slows the increase of $\mathcal{P}_t(K)$ while preserving most of $\mathcal{T}_t(K)$ for confident cases, thus improving:

$$\text{ConsensusPrecision}_t(K) = \frac{\mathcal{T}_t(K)}{\mathcal{P}_t(K)}.$$

3. For tail classes, $\mathcal{P}_t(K)$ contains more noise under large K.

Let $\epsilon_t(K)$ denote the expected proportion of false positives in $\mathcal{P}_t(K)$. For tail classes, as K increases:

$$\epsilon_t(K) := \frac{\mathcal{P}_t(K) - \mathcal{T}_t(K)}{\mathcal{P}_t(K)} \uparrow,$$

indicating more of the Top-K inclusions of t are incorrect. Thus:

$$\text{ConsensusPrecision}_t(K) = 1 - \epsilon_t(K) \downarrow .$$

Conclusion: To mitigate noisy inclusions and preserve precision for rare classes, a smaller class-specific K_t should be adopted for tail class t:

$$\text{ConsensusPrecision}_t(K_t) > \text{ConsensusPrecision}_t(K), \quad \text{for } K_t < K.$$

□

E. Analysis of CE and LA loss in the results on mini-ImageNet

Table 14. Analysis of CE and LA loss in the results on mini-ImageNet.

Loss	Class Group	
	Head	Tail
CE	87.3	81.2
LA	81.7	85.2

In the results on mini-ImageNet with an imbalance factor of 10, CE loss outperforms the LA loss. An imbalance factor of 10 means that the sample size of the head classes is ten times that of the tail classes. However, mini-ImageNet is a relatively balanced and moderately sized dataset, and even under this level of imbalance, the tail classes still retain a reasonable number of training examples, typically in one hundred. This relatively sufficient amount of data for tail classes allows the model to generalize reasonably well across both head and tail classes when trained with the CE loss. Consequently, the disparity in classification performance between head and tail classes is not particularly severe, and the CE loss can maintain strong overall performance without introducing bias toward any specific subset of classes. Besides, LA explicitly modifies the logits during training to favor underrepresented classes by incorporating class prior distributions into the loss function. While this reweighting helps improve the performance on tail classes, it does so by effectively downscaling the contributions of head classes. In settings where the imbalance is mild and tail classes already have adequate data to learn from, this shift in focus can lead to overcompensation. As a result, the performance of head classes degrades significantly, overshadowing the marginal gains seen in tail-class accuracy. This trade-off ultimately leads to a decline in overall performance when using LA compared to CE. For the results on mini-ImageNet with an imbalance factor of 10, we provide the accuracy of head and tail classes as Table 14.

F. Complementary Analysis of CARE Components

Figure 4 complements Table 7 (in Section 3.3) by illustrating how the label noise ratio evolves across epochs. It further demonstrates that a single expert alone is insufficient to alter the base expert (BE). Since the confidence scores from the image expert (IE) and text expert (TE) remain consistently below 1. Consequently, their individual contributions to the frequency accumulation vector **F** (defined in Equation (4)) are insufficient to surpass the accumulated support from BE. The single expert cannot accumulate enough frequency to override BE and induce label correction Collaborative voting among multiple experts for robust label refinement is necessary.

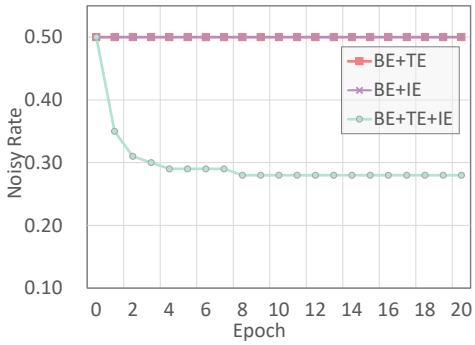

Figure 4. Noise Rate Evolution across CARE Components.

G. Analysis of Cumulative Frequency-Based Correction with Varying BE Weight

To further clarify the core mechanism of our correction strategy, which relies on cumulative frequency comparison, we analyze the case where the BE weight is set to 1. In this setting, correcting noisy labels in BE requires the cumulative frequency of predictions from IE and TE to exceed that of the noisy labels in BE. As shown in Table 7 (in Section 3.3),

Table 15. Expert Combinations Effect on Noisy Label Correction (BE Weight = 0.5, Symmetric Noise). NR: Noise Rate after correction.

Dataset	**BE + TE**		**BE + IE**	
	NR (↓)	Acc. (↑)	NR (↓)	Acc. (↑)
CIFAR100-IF10-NR50	31.8	79.8	29.3	79.7

neither IE nor TE alone can accumulate sufficient clean signals to achieve this correction. We then conducted additional experiments with the BE weight reduced to 0.5. The results in Table 15 demonstrate that, under this setting, either IE or TE alone is sufficient to correct noisy labels and improve model performance.

H. Ablation on the Form of Top-K Predictions

Table 16. Ablation Study on K (CIFAR-100-LTN, IF=100, 20%, Symmetric Noise).

K	Acc. (%)
CLIP+LA (baseline)	76.0
Step form (`H:8 M:4 T:1`)	76.5
Step form (`H:8 M:3 T:2`)	76.3
Step form (`H:7 M:4 T:2`)	76.4
Step form (`H:8 M:4 T:2`)	76.7
Step form (`H:9 M:4 T:2`)	76.9
Step form (`H:8 M:5 T:2`)	76.3
Step form (`H:8 M:4 T:3`)	76.6
Exponential form ($K_i = n_i^{0.25}$)	77.3
Logarithmic form ($K_i = \log(n_i)$)	77.3
Linear form	77.5

The K value for each class is proportional to the number of samples in that class. Specifically, it is computed as the one-fourth power of the number of samples in that class. We empirically adopt this power-law form as a practical choice. We also conducted ablation experiments on the calculation form of the hyperparameter K, and the results are shown as Table 16. "`H`", "`T`" and "`T`" refer to Head, Medium, and Tail classes, and n_i denotes the number of samples in class i. The linear form is computed as $K_i := \frac{n_i - n_{min}}{(n_{max} - n_{min})} * (K_{max} - K_{min}) + K_{min}$. n_{max} and n_{min} represent the maximum and minimum sample counts, which are set to 9 and 1, respectively. The experimental results demonstrate that different forms of K, particularly those proportional to class frequency, yield stable and improved performance.

Table 17. Comparison of class-wise K and global K.

Datasets / K	CIFAR-100-LTN ($IF = 10$, $NR = 20$, AN)				ImageNet-LTN^r (IF = 100, NR=40)			
	ALL	Head	Med	Tail	ALL	Head	Med	Tail
Global $K = 4$	80.7	81.0	82.7	77.9	80.7	81.1	96.0	73.6
Global $K = 8$	80.3	80.3	82.7	77.3	80.2	80.9	96.0	74.4
Class-wise K_c	80.8	80.9	82.5	78.6	81.9	83.0	94.5	74.8

Global vs. Class-wise K (Table 17): Using a global K for all classes improves overall accuracy compared to baseline (equivalent to $K = 1$), but the class-wise K_c consistently achieves higher performance, particularly benefiting tail classes.

I. Related Work of Long-Tail Learning

Long-tailed learning methods typically assume that datasets are correctly labeled (Cui et al., 2019), and adopt class-wise strategies that can be broadly categorized into four levels (Li, 2022; Shi et al., 2024). Input level: These approaches manipulate the training data to enhance learning, including re-weighting, re-sampling (Cui et al., 2019), and data augmentation techniques (Cubuk et al., 2019; 2020; Wang et al., 2024a;b). Representation level: Methods at this level modify model training to decouple representation learning from classifier learning, such as in decoupled training (Kang et al., 2020; Zhong et al., 2021) and BBN-based approaches (Zhou et al., 2020; Zhang et al., 2021), where features are learned on imbalanced data, followed by classifier retraining with class-balanced or reversed samples. Ensemble strategies are also used, including

redundant ensembles (Wang et al., 2021; Li et al., 2022a;c; Cai et al., 2021) and complementary ensembling (Zhou et al., 2020; Cui et al., 2023; Zhao et al., 2024), which leverage diverse expert outputs or data partitions. Output level: These methods directly adjust model outputs to model generalization and classifier performance. Logit adjustment (Menon et al., 2021; Ren et al., 2020) re-scales predictions, while re-margining techniques (Cao et al., 2019; Li et al., 2023; Menon et al., 2021; Li et al., 2022d; Zhao et al., 2025) assign class-dependent margins to promote better tail-class separation. Gradient level: CCSAM (Zhou et al., 2023b) and ImbSAM (Zhou et al., 2023a) adapt Sharpness-aware Minimization (Foret et al., 2021) to strengthen tail-class learning by adjusting gradient contributions based on class frequency. GNM (Li et al., 2024a) introduces parameter-independent perturbations to reduce head-class dominance during training.

J. Generalization analysis of CARE

We evaluated two additional noise types across multiple datasets:

1. Standard Symmetric Noise (Uniform Noise): Its transition matrix is defined as

$$T_{ii} = 1 - \eta \quad \text{(probability of keeping the correct label)},$$

$$T_{ij} = \frac{\eta}{\text{num_classes} - 1}, \quad i \neq j \quad \text{(probability of flipping from class } i \text{ to class } j\text{)}.$$

2. Asymmetric Pair-Flip Noise: Instead of flipping to any random class, each class is flipped to one specific, visually similar class, modeling realistic inter-class confusion.

The experimental results are summarized in Table 18. As shown, CARE consistently improves performance under both standard symmetric noise and asymmetric pair-flip noise. Specifically, for standard symmetric noise (NR40 and NR60), CARE achieves performance gains of 3.3 and 4.0 points, respectively, demonstrating its effectiveness in mitigating randomly distributed label corruption. Under the more challenging pair-flip noise (NR20 and NR40), which arises from class similarity, CARE still provides improvements of 1.5-3.7 points. Overall, these results show that CARE generalizes well across different noise types, enhancing robustness to both random and structured label noise.

Table 18. Experimental results (%) of two additional noise types.

Datasets/Methods	Noise type	w/o CARE	w/ CARE
CIFAR100_IF100_NR40	Uniform Noise	70.8	74.1
CIFAR100_IF100_NR60	Uniform Noise	65.3	69.3
CIFAR100_IF100_NR20	Pair-Flip Noise	77.9	79.4
CIFAR100_IF100_NR40	Pair-Flip Noise	67.0	70.7

K. Analysis of Macro F1-score

Table 19. The Macro F1-score (%) across multiple datasets.

Datasets/Methods	Noise type	w/o CARE	w/ CARE
CIFAR100_IF100_NR30	Joint Noise	77.2	78.3
CIFAR100_IF100_NR50	Joint Noise	75.0	76.4
CIFAR100_IF10_NR40	Symmetric Noise	80.3	81.4
CIFAR100_IF10_NR60	Symmetric Noise	75.7	78.9
CIFAR100_IF10_NR20	Asymmetric Noise	80.0	80.7
CIFAR100_IF10_NR40	Asymmetric Noise	67.6	69.7

Macro F1-score is the average of the F1-score computed independently for each class, giving equal weight to all classes regardless of their size. By combining precision and recall, Macro F1-score captures how well a model balances false positives and false negatives across classes. This makes it sensitive to class imbalance, as poor performance on minority classes cannot be offset by strong performance on majority classes.

We evaluated our method in terms of Macro F1-score(%) across multiple datasets, as shown in Table 19. CARE consistently improves Macro F1-score under all noise settings. For joint noise (NR30/NR50), it achieves gains of approximately

1-1.4 points. For symmetric noise (NR40/NR60), CARE achieves consistent improvements, including a substantial 3.2-point increase at $NR = 60$. Under the more challenging asymmetric noise scenario (NR20/NR40), CARE continues to demonstrate measurable performance gains. Overall, CARE reliably enhances class-balanced performance across different noise types and severities.

L. Experimental results (%) of Clothing1M

Table 20. Experimental results (%) of Clothing1M.

Datasets/Methods	w/o CARE	w/ CARE
Clothing1M	71.2	71.4

The experimental results of Clothing1M are presented as in Table 20. CARE yields a modest improvement on Clothing1M (71.2 vs. 71.4), which can be attributed to the dataset's abundant training samples and atypical class imbalance. Specifically, even the least frequent classes contain 19,743 samples, which is substantially higher than the tail classes in typical long-tailed datasets, indicating that these classes are already well represented. Consequently, the impact of methods designed to address tail-data scarcity is limited. Furthermore, in Clothing1M, some classes that are rare in the test set actually have a relatively large number of training samples, which is fundamentally different from real-world long-tailed distributions. In typical long-tailed scenarios (Menon et al., 2021; Zhou et al., 2020; Li et al., 2024b), each class is considered equally important, and tail classes are consistently underrepresented both in training and testing, leading to a relatively balanced test set that reflects the importance of all classes. This contrasts with Clothing1M, where the training set abundance for some rare test classes reduces the effect of tail scarcity. Despite these factors, CARE still achieves a measurable gain, showing that the method remains effective even when the dataset deviates from the specific long-tailed, noisy conditions it is primarily designed for.

M. Effect of Training Duration

Table 21. Accuracy (%) across epochs.

Epochs	15	20	30
Acc. (%)	80.8	81.5	82.0

We conducted additional experiments with 15, 20, and 30 epochs on CIFAR-100-LTN (Symmetric Noise, IF=10, NR=40%). As shown in Table 21, CARE is highly robust to training duration: it achieves a strong 80.8% accuracy even at 15 epochs, and smoothly improves to 82.0% at 30 epochs. We adopted 20 epochs (81.5%) as our default to optimally balance computational efficiency and state-of-the-art performance. By leveraging strong vision-language priors, CARE's rectification and representation learning converge significantly faster than from-scratch training. These results confirm that our method is not overly sensitive to exact epoch counts, making 20 epochs a highly reliable and efficient default setting.

N. More Recent Methods

Table 22. Accuracy (%) Comparison of IBC on CIFAR100-LTN.

Datasets/Methods	Noise type	CLIP+IBC	CLIP+CARE
CIFAR100_IF10_NR40	Symmetric	79.3	81.5
CIFAR100_IF10_NR40	Asymmetric	70.2	70.5
CIFAR100_IF10_NR40	Joint	70.2	76.7

We conduct a comparison with IBC (Li et al., 2025) (ICCV 2025). As shown in Table 22, CLIP+CARE outperforms CLIP+IBC across all noise types. Even under the most challenging joint noise setting, CARE achieves a significant +6.5% improvement over IBC (76.7% vs. 70.2%).

Table 23. Noise rates (%) for different classes in RLD 1, 2, and 3.

Methods/Classes	Head	Med.	Tail
RLD1	17.4	21.6	70.2
RLD2	17.0	30.5	59.2
RLD3	27.9	27.8	31.2

O. Analysis of Noise Rates for Different Classes in RLD

The noise rates for samples of different frequencies in RLD 1, 2, and 3 (as shown in Table 1) are presented in Table 23. Compared with RLD1 and RLD2, RLD3 exhibits a substantially more balanced noise-rate distribution across the head, medium, and tail categories. As shown in the table, the noise rates of RLD1 and RLD2 display a clear long-tailed pattern, where the tail classes suffer from significantly higher noise (e.g., 70.2% in RLD1 and 59.2% in RLD2), while the head classes remain comparatively clean. In contrast, RLD3 reduces this discrepancy: the noise rates of the three frequency groups (27.9% for head, 27.8% for medium, and 31.2% for tail) are closely aligned. This indicates that RLD3 mitigates the frequency-dependent bias of label noise, leading to a more uniform noise distribution. Such balance is beneficial for learning algorithms, as it prevents models from being disproportionately affected by high-noise tail classes.

P. Limitation Analysis & Future Work

While CARE achieves consistent improvements, it requires maintaining class frequency distributions during training, introducing extra memory overhead. Our future work will focus on optimizing memory usage to reduce this cost while preserving efficiency. In addition, CARE has the potential to integrate additional expert sources, such as image-aware experts from contrastive learning, to further enhance label rectification and model performance. We also leave this as an avenue for future work.

Q. Computational Resources

To ensure a fair comparison, all the experiments are executed on the following hardware: Intel (R) Xeon (R) Gold 5220, operating at 2.20GHz, equipped with 251GB RAM, and a single NVIDIA GeForce RTX 3090 GPU with 24 GB of memory.