

000 001 002 003 004 005 SCAD: SUPER-CLASS-AWARE DEBIASING FOR LONG- 006 TAILED SEMI-SUPERVISED LEARNING 007 008 009

010 **Anonymous authors**
011 Paper under double-blind review
012
013
014
015
016
017
018
019
020
021
022
023
024
025
026

ABSTRACT

027 In long-tailed semi-supervised learning (LTSSL), pseudo-labeling often creates a
028 vicious cycle of bias amplification, a problem that recent state-of-the-art methods
029 attempt to mitigate using logit adjustment (LA). However, their adjustment
030 schemes, inherited from LA, remain inherently hierarchy-agnostic, failing to ac-
031 count for the semantic relationships between classes. In this regard, we identify
032 a critical yet overlooked problem of *intra-super-class imbalance*, where a toxic
033 combination of high semantic similarity and severe local imbalance within each
034 super-class hinders effective LTSSL. This problem causes the model to reinforce
035 on its errors, leading to representation overshadowing. To break this cycle, we
036 propose Super-Class-Aware Debiasing (SCAD), a new framework that performs a
037 dynamic, super-class-aware logit adjustment. SCAD leverages the latent semantic
038 structure between classes to focus its corrective power on the most confusable
039 groups, effectively resolving the local imbalances. Our extensive experiments
040 validate that SCAD achieves new state-of-the-art performance, demonstrating the
041 necessity of a super-class-aware approach for robust debiasing.
042
043

1 INTRODUCTION

044 Deep learning models have achieved remarkable success on large, annotated datasets, but creating
045 them is expensive (Deng et al., 2009; Lin et al., 2014). Semi-supervised learning (SSL) is a powerful
046 approach to reduce this cost by using a small labeled dataset with a large unlabeled one (Sohn et al.,
047 2020; Berthelot et al., 2020). However, the performance of SSL is often hindered by the long-tailed
048 distribution of real-world data (Figure 1a), where a few majority classes vastly outnumber many
049 minority classes (Kang et al., 2020; Cui et al., 2019). This imbalance especially creates a vicious
050 cycle in long-tailed semi-supervised learning (LTSSL): a biased model generates biased pseudo-
051 labels, reinforcing the bias and causing performance to drop severely for minority classes (Kim
052 et al., 2020; Wei et al., 2021).

053 To break this vicious cycle, logit adjustment (LA) (Menon et al., 2021), a powerful method from sup-
054ervised long-tail learning, has become the *de facto standard* for pseudo-label debiasing in LTSSL.
055 LA applies a static, corrective offset to each class’s logits, calculated based on the global class fre-
056quencies. The appeal of this approach lies in its theoretical guarantee: LA provides a cost-free,
057 *Fisher-consistent* correction for any given class prior, making it an ideal foundational block for
058 more complex scenarios. Building upon this foundation, recent state-of-the-art methods have made
059 remarkable progress by deriving more sophisticated estimates of the true prior to leverage within
060 the LA framework, such as by estimating the distribution mismatch (Wei & Gan, 2023) or measur-
061ing the classifier’s intrinsic bias (Lee & Kim, 2024). However, while these methods innovate on
062 how to estimate the global prior, the adjustment scheme itself remains that of LA: a single, uniform
063 correction that is inherently blind to the semantic relationships between classes.

064 In this paper, we first uncover that this blindness to semantic relationships among classes leads to a
065 critical challenge we term *intra-super-class imbalance*. Specifically, we identify that this problem
066 arises from a toxic combination of two factors, best understood through the example of *truck* and
067 *automobile* on CIFAR10-LT. First, for these classes that have **high semantic similarity**, adjustment
068 mechanisms based on LA offer no solution, as its corrective offset is derived solely from class
069 frequencies and is agnostic to the inter-class similarities that cause model confusion. Second, this
070 challenge is dramatically amplified by the **extreme local imbalance** within their shared *vehicle*




Figure 1: **Illustration of the *intra-super-class imbalance* problem on CIFAR10-LT.** (a) The overall dataset exhibits a long-tailed distribution. (b) The problem arises from a toxic combination of high semantic similarity and extreme local imbalance, as seen within the *vehicle* super-class. (c) This leads to the failure of LA (Menon et al., 2021), which struggles to distinguish minority classes from their majority neighbors. (d) In contrast, SCAD **mitigates** these critical misclassifications.

super-class, where *automobile* vastly outnumbers *truck* (Figure 1b). LA’s globally derived correction is fundamentally mismatched for this local conflict. Its adjustment for the *truck* class is determined by its global rarity across the entire dataset, not by the specific, intense competition it faces from *automobile*. This globally calibrated boost is ill-suited to counteract the overwhelming pressure from its direct and dominant competitor. In the context of pseudo-labeling, this dual failure of LA triggers a self-reinforcing loop of bias amplification, leading to representation overshadowing and the systematic misclassification shown in Figure 1c.

To tackle this problem, we propose **Super-Class-Aware Debiasing (SCAD)**, a new framework designed to break this cycle by performing a dynamic, super-class-aware logit adjustment. Our approach is driven by a key insight: the debiasing mechanism must be adaptive, dynamically tailoring the corrective force to a sample’s predicted semantic group. SCAD implements this by first discovering a latent super-class structure from class names using pre-trained text encoders (e.g., CLIP (Radford et al., 2021)), then training a super-class classifier to infer each unlabeled sample’s semantic context, and finally refining standard LA with a super-class-aware correction that uses the super-class posterior as weights for a targeted local correction on top of LA’s global adjustment. This allows SCAD to mitigate systematic misclassifications on CIFAR10-LT (Figure 1d), and, as we further show on CIFAR100-LT (Figure 4 and Appendix A.2), to consistently alleviate analogous *intra-super-class* errors on more complex long-tailed benchmarks.

To validate our approach, we conduct experiments on diverse LTSSL benchmarks, ranging from CIFAR10/100-LT (Krizhevsky et al., 2009) to large-scale, fine-grained datasets such as ImageNet-127(Fan et al., 2022) and Food101-LT(Bossard et al., 2014). SCAD proves to be a flexible, pluggable framework that consistently enhances state-of-the-art methods under distribution mismatch.

2 RELATED WORK

Long-tailed Semi-supervised Learning LTSSL addresses the realistic scenario where labeled data is both scarce and imbalanced. Early methods focused on directly manipulating pseudo-labels, for instance, by iterative re-balancing Wei et al. (2021) or refining their distributions Kim et al. (2020); Wang et al. (2022). Subsequently, logit adjustment (LA) Menon et al. (2021) emerged as the *de facto standard* for debiasing, offering a more fundamental solution by directly correcting the model’s output logits. The effectiveness and simplicity of LA have made it the foundational block for recent state-of-the-art methods that tackle complex scenarios like distribution mismatch. These advanced methods, such as ACR Wei & Gan (2023) and CDMAD Lee & Kim (2024), innovate by dynamically modulating the strength of the LA correction. ACR achieves this by estimating the mismatch between labeled and unlabeled distributions, while CDMAD does so by measuring the classifier’s intrinsic bias. Despite these sophisticated advances, their adjustment schemes remain fundamentally that of LA. Consequently, they are all inherently hierarchy-agnostic and sub-optimal for alleviating the localized confusion between semantic neighbors.




Figure 2: **Illustration of Super-Class-Aware Debiasing (SCAD) mechanism.** For an unlabeled *truck* image, the class classifier initially produces biased logits, incorrectly favoring *automobile*. Standard LA applies a uniform, global correction, which is insufficient to change the outcome. In contrast, our mechanism first uses a reliable super-class classifier to infer the sample’s context (*vehicle*). In this context, $p(k|u_j)$, is then used to compute a **targeted local adjustment**. When this local adjustment is added to the standard LA, it successfully overcomes the initial bias and boosts the logit for the correct class.

Hierarchical and Super-class learning Leveraging class hierarchies to improve recognition is well established. In supervised learning, methods like HD-CNN Yan et al. (2015) and Deep-RTC Wu et al. (2020) use manually defined hierarchies, while SuperDisco Du et al. (2023) discovers super-classes from image features via a graph neural network. However, these approaches are less practical for LTSSL, as they often require pre-defined taxonomies or complex, computationally expensive training. In contrast, SCAD is designed to be lightweight and practical: it constructs super-classes by applying a pre-trained text encoder to class names—without manual hierarchies or pre-trained vision models—and serves as a simple, pluggable debiasing module for existing LTSSL methods.

3 METHOD

We propose Super-Class-Aware Debiasing (SCAD), a multi-stage framework to address *intra-super-class imbalance* in LTSSL. SCAD consists of three components: (i) super-class discovery from semantic information, (ii) a dedicated training procedure for the feature extractor and two classifiers, and (iii) a dynamic logit adjustment mechanism that corrects biased predictions at inference time. We detail each component below, after introducing notation and problem setup.

3.1 PRELIMINARIES

3.1.1 PROBLEM SETUP

Let $\mathcal{D}^l = \{(x_i, y_i)\}_{i=1}^N$ be a labeled dataset and $\mathcal{D}^u = \{u_j\}_{j=1}^M$ an unlabeled dataset, where $x_i, u_j \in \mathbb{R}^d$ denote the d -dimensional inputs. We consider C classes, indexed by $c \in [C] \triangleq \{1, \dots, C\}$. Let N_c be the number of labeled samples in class c and M_c the (generally unknown) number of unlabeled samples in class c . The labeled set is long-tailed with $N_1 \geq N_2 \geq \dots \geq N_C$ and imbalance ratio $\gamma_l = N_1/N_C$. We similarly define the unlabeled imbalance ratio $\gamma_u = \max_c M_c / \min_c M_c$.

Our model consists of three main components: a feature extractor $f_\theta : \mathbb{R}^d \rightarrow \mathbb{R}^p$, a primary class classifier $g_c : \mathbb{R}^p \rightarrow \mathbb{R}^C$, and an auxiliary super-class classifier $g_s : \mathbb{R}^p \rightarrow \mathbb{R}^K$, where p is the dimensionality of the feature representation and K denotes the number of discovered super-classes. For a given input x , we denote the extracted features as $z = f_\theta(x)$, the class logits as $\ell^c = g_c(z)$, and the super-class logits as $\ell^s = g_s(z)$.

162 3.1.2 SEMI-SUPERVISED LEARNING.
163164 Our method builds upon the consistency regularization paradigm, as popularized by FixMatch (Sohn
165 et al., 2020). The total loss for the primary SSL is

166
$$L = L_s + L_u, \quad (1)$$

167 where L_s is the supervised loss on labeled data and L_u is the unsupervised consistency loss on
168 unlabeled data.170 **Supervised loss.** For a mini-batch $\mathcal{B}_l \subset \mathcal{D}^l$, we apply a weak augmentation \mathcal{A}_w to x_i and compute

171
$$L_s = \frac{1}{|\mathcal{B}_l|} \sum_{(x_i, y_i) \in \mathcal{B}_l} \text{CE}(\text{softmax}(g_c(f_\theta(\mathcal{A}_w(x_i)))), y_i), \quad (2)$$

174 where $\text{CE}(\cdot, \cdot)$ denotes the standard cross-entropy with the label y_i treated as a class index.175 **Unsupervised consistency loss.** For a mini-batch $\mathcal{B}_u \subset \mathcal{D}^u$, we obtain a pseudo-label from a
176 weak view and enforce consistency on a strong view. Let $q_j = \text{softmax}(g_c(f_\theta(\mathcal{A}_w(u_j))))$ and
177 $\hat{y}_j = \arg \max_{c \in [C]} (q_j)_c$. We apply a strong augmentation \mathcal{A}_s to the same sample and compute

179
$$L_u = \frac{1}{|\mathcal{B}_u|} \sum_{u_j \in \mathcal{B}_u} \mathbb{I}(\max_c (q_j)_c \geq \tau) \text{CE}(\text{softmax}(g_c(f_\theta(\mathcal{A}_s(u_j)))), \hat{y}_j), \quad (3)$$

181 where $\tau \in (0, 1)$ is a confidence threshold. The mask $\mathbb{I}(\max_c (q_j)_c \geq \tau)$ prevents low-confidence
182 pseudo-labels from contributing to the loss. In practice, \hat{y}_j is treated with stop-gradient.184 3.2 SUPER-CLASS GENERATION
185186 The foundation of our method is a structural prior that groups fine-grained classes into meaningful
187 super-classes. We generate this prior automatically from class name semantics, requiring no manual
188 annotation or pre-trained vision models. [More details are explained in A.9](#).190 **Generation Process.** The generation process involves two steps: (1) We convert the names of
191 all C classes into semantic vectors using a pre-trained text encoder (e.g., SBERT (Devlin et al.,
192 2019) or CLIP’s text encoder (Radford et al., 2021)). (2) We then apply agglomerative clustering on
193 these vectors. This bottom-up technique builds a dendrogram representing the class hierarchy. By
194 cutting this dendrogram at a level that yields a pre-specified number of clusters (K), we partition
195 the C classes into K super-classes. This process yields a deterministic mapping $\mathcal{M} : \{1, \dots, C\} \rightarrow$
196 $\{1, \dots, K\}$, where $\mathcal{M}(c)$ is the super-class index for class c .197 **Justification.** Our approach of generating super-classes is grounded in the observation that se-
198 mantic relationships captured by pre-trained text embeddings often correlate with visual tax-
199 onomies (Radford et al., 2021). This strategy is particularly advantageous as it sidesteps the need
200 for manual hierarchy creation, which is often infeasible for large-scale or specialized datasets. It is
201 important to clarify that the novelty of our work does not lie in this generation process itself, which
202 is a straightforward application of existing tools. Rather, we employ this process as a simple and
203 effective means to obtain the structural prior required for our main contribution: a new, dynamic
204 debiasing mechanism. Our goal is to show that even an approximate hierarchy can [alleviate](#) critical
205 issues when leveraged by a powerful debiasing algorithm.206 Crucially, the effectiveness of our framework does not hinge on any single text encoder. This ro-
207 bustness is a key finding, as it proves our method’s principle is general and not reliant on a specific,
208 and potentially expensive, text model. SCAD consistently improves performance regardless of the
209 choice of text encoder, from lightweight GloVe (Pennington et al., 2014) to large-scale CLIP. This
210 indicates that SCAD’s benefit stems from the debiasing mechanism rather than encoder capacity, a
211 conclusion further supported by our ablations in Section 4.5.212 In this work, we obtain super-classes by applying text encoders to class names, but SCAD itself is
213 agnostic to the hierarchy source: it only requires a mapping from classes to super-classes. The same
214 interface can be instantiated using existing taxonomies or manually defined hierarchies. In extremely
215 rare or highly domain-specific settings where class names provide limited semantic signal, ontology-
or expert-defined super-classes offer a natural alternative to text-based grouping.

216 3.3 TRAINING PROCEDURE
217

218 The feature extractor f_θ , class classifier g_c , and super-class classifier g_s are trained jointly. The goal
219 of this stage is not only to train the main classifier, but also to train a reliable super-class classifier
220 g_s that can provide a trustworthy signal for our debiasing mechanism.

221 The rationale for this joint training hinges on the observation that the auxiliary task of super-class
222 classification is inherently simpler and more robust to data imbalance. This robustness stems from
223 two factors: (1) it is a coarse-grained problem with significantly fewer classes ($K \ll C$), and (2)
224 grouping classes often mitigates the label distribution imbalance, as noted in prior work (Du et al.,
225 2023). Consequently, g_s can learn to produce reliable predictions even for minority class samples,
226 providing a stable signal that is crucial for our debiasing mechanism. As we show in Appendix A.4,
227 g_s indeed yields a more balanced and reliable prediction distribution on real-world unlabeled data.

228 To achieve this, we introduce an auxiliary SSL loss for the super-class task, L_{super} , which mirrors
229 the structure of the primary loss:

$$230 \quad L_{\text{super}} = L_s^{\text{super}} + L_u^{\text{super}}. \quad (4)$$

231 Here, the supervised component, L_s^{super} , is the standard cross-entropy loss computed on the labeled
232 batch \mathcal{B}_l using super-class targets $y_i^{\text{super}} = \mathcal{M}(y_i)$, where \mathcal{M} maps each fine-grained class to a super-
233 class (either from a dataset-provided hierarchy or from our automatically discovered grouping). The
234 unsupervised component, L_u^{super} , is a consistency loss computed on the unlabeled batch \mathcal{B}_u , enforcing
235 that the model’s super-class predictions are consistent across strong and weak augmentations.

236 The final training objective is a weighted sum of the primary and auxiliary losses:

$$238 \quad L_{\text{total}} = L + \lambda L_{\text{super}}, \quad (5)$$

239 where λ is a hyperparameter balancing the two tasks. This objective encourages f_θ to learn repre-
240 sentations that are discriminative at both fine-grained and coarse-grained levels.

242 3.4 SUPER-CLASS–AWARE LOGIT ADJUSTMENT
243

244 This section introduces our framework’s core contribution: a novel logit adjustment method, Super-
245 Class–Aware logit Adjustment (SCAD), that addresses a critical limitation of the standard LA. Stan-
246 dard LA applies a static, global correction to the logits based on the class prior distribution π :

$$248 \quad \ell_j^{\text{LA}} = \ell_j^c - \log \pi, \quad (6)$$

249 While effective against global class imbalance, this super-class-aware correction is ill-equipped to
250 mitigate local ambiguities, particularly conflicts arising between classes within the same super-class.

251 SCAD overcomes this limitation with a dynamic, sample-aware adjustment mechanism, illustrated
252 in Figure 2. The process begins by inferring the coarse-grained context for each unlabeled sample
253 u_j . We utilize the super-class classifier, g_s , to obtain a posterior probability distribution over super-
254 classes, $p(k|u_j) = \text{softmax}(\ell_j^s)_k$. This distribution, $p(k|u_j)$, quantifies the estimated likelihood of
255 the sample belonging to each super-class k . For instance, an image of a *truck* would yield a high
256 probability for the *vehicle* super-class, i.e., $p(\text{vehicle}|u_j) \approx 1$.

257 This inferred context is then used to apply a tailored, super-class-aware adjustment. We pre-calculate
258 an adjustment vector $\Delta_k \in \mathbb{R}^C$ for each super-class k , specifically designed to counteract the *intra-*
259 *super-class imbalance* within that group. Let $\mathcal{C}_k = \{c \mid \mathcal{M}(c) = k\}$ be the set of all classes
260 belonging to super-class k . The c -th components of Δ_k are defined as:

$$262 \quad (\Delta_k)_c = \begin{cases} \beta_{k,c}, & \text{if } c \in \mathcal{C}_k, \\ \max_{c' \in \mathcal{C}_k} \beta_{k,c'}, & \text{if } c \notin \mathcal{C}_k, \end{cases} \quad (7)$$

265 where $\beta_{k,c} = n_{k,c}/(\max_{c' \in \mathcal{C}_k} n_{k,c'})$ is the relative dominance score of class c within its super-class
266 k , based on the sample count or estimated frequency $n_{k,c}$ for class c within super-class k . from a
267 mixture of labeled and high-confidence pseudo-labeled samples: whenever a labeled example with
268 ground-truth label c , or an unlabeled example whose pseudo-label is c with confidence above the
269 FixMatch threshold τ , is assigned to super-class k , we increment $n_{k,c}$ by one. These counts are
recomputed periodically (once per epoch in our implementation) using the current classifier. The

270 Table 1: Results on CIFAR10-LT and CIFAR100-LT datasets for various algorithms and settings.
271

272	273	CIFAR10-LT								CIFAR100-LT								
		$\gamma = \gamma_l = \gamma_u = 100$				$\gamma = \gamma_l = \gamma_u = 150$				$\gamma = \gamma_l = \gamma_u = 10$				$\gamma = \gamma_l = \gamma_u = 20$				
		$N_1 = 500$	$N_1 = 1500$	$M_1 = 4000$	$M_1 = 3000$	$N_1 = 500$	$N_1 = 1500$	$M_1 = 4000$	$M_1 = 3000$	$N_1 = 50$	$N_1 = 150$	$M_1 = 400$	$M_1 = 300$	$N_1 = 50$	$N_1 = 150$	$M_1 = 400$	$M_1 = 300$	
274	Algorithm	Supervised w/ LA	47.3 \pm 0.95	61.9 \pm 0.41	53.3 \pm 0.44	70.6 \pm 0.21	44.2 \pm 0.33	58.2 \pm 0.29	49.5 \pm 0.40	67.1 \pm 0.78	29.6 \pm 0.57	46.9 \pm 0.22	30.2 \pm 0.44	48.7 \pm 0.89	25.1 \pm 1.14	41.2 \pm 0.15		
275		FixMatch w/ DARP	67.8 \pm 1.13	77.5 \pm 1.32	74.5 \pm 0.78	77.8 \pm 0.63	62.9 \pm 0.36	72.4 \pm 1.03	67.2 \pm 0.32	73.6 \pm 0.73	45.2 \pm 0.55	56.5 \pm 0.06	49.4 \pm 0.20	58.1 \pm 0.44	43.4 \pm 0.87	52.2 \pm 0.66		
276		w/ CReST+	76.3 \pm 0.86	78.1 \pm 0.42	76.0 \pm 0.37	79.1 \pm 0.75	67.5 \pm 0.45	73.7 \pm 0.34	70.1 \pm 1.81	75.1 \pm 0.77	44.5 \pm 0.94	57.4 \pm 0.18	44.0 \pm 0.21	57.1 \pm 0.55	40.1 \pm 1.28	52.1 \pm 0.21		
277		w/ DASO	76.0 \pm 0.37	79.1 \pm 0.75	75.1 \pm 0.10	79.2 \pm 1.60	69.0 \pm 0.61	75.8 \pm 0.33	74.5 \pm 1.21	82.2 \pm 0.14	49.8 \pm 0.24	59.2 \pm 0.35	50.0 \pm 0.22	58.7 \pm 0.28	43.6 \pm 0.09	52.9 \pm 0.42		
278		w/ DASO + Ours	75.1 \pm 0.10	79.2 \pm 1.60	75.3 \pm 2.45	82.0 \pm 0.36	67.0 \pm 2.49	78.0 \pm 0.91	68.2 \pm 0.94	76.7 \pm 1.13	47.3 \pm 0.42	58.6 \pm 0.36	50.5 \pm 0.78	59.9 \pm 0.32	44.4 \pm 0.65	53.8 \pm 0.43		
279		FixMatch + LA	76.6 \pm 0.92	80.8 \pm 0.62	76.7 \pm 1.13	81.1 \pm 0.57	70.9 \pm 1.18	77.9 \pm 0.71	70.1 \pm 1.68	79.0 \pm 2.23	50.7 \pm 0.51	60.6 \pm 0.71	44.0 \pm 0.21	57.1 \pm 0.55	40.6 \pm 0.55	52.3 \pm 0.20		
280		w/ DARP	76.7 \pm 1.13	81.1 \pm 0.57	77.9 \pm 0.88	82.5 \pm 0.08	74.5 \pm 1.21	82.2 \pm 0.14	74.5 \pm 1.21	82.2 \pm 0.14	51.8 \pm 0.28	60.5 \pm 0.21	52.7 \pm 0.11	61.8 \pm 0.21	45.7 \pm 0.56	55.7 \pm 0.57		
281		w/ CReST+	81.6 \pm 0.22	84.0 \pm 0.99	81.6 \pm 0.22	84.0 \pm 0.99	77.0 \pm 1.19	80.9 \pm 0.22	78.6 \pm 0.56	83.3 \pm 0.20	51.3 \pm 0.48	61.1 \pm 0.11	44.8 \pm 0.21	55.9 \pm 0.31				
282		w/ DASO	81.6 \pm 0.22	84.0 \pm 0.99	83.5 \pm 0.16	85.5 \pm 0.03	78.6 \pm 0.56	83.3 \pm 0.20			52.7 \pm 0.11	61.8 \pm 0.21	45.8 \pm 0.20	56.4 \pm 0.10				
283		FixMatch + ACR	81.6 \pm 0.19	84.1 \pm 0.39	83.5 \pm 0.16	85.5 \pm 0.03												
284		w/ Ours																

286 Table 2: Test accuracy of previous LTSSL algorithms and ours under inconsistent class distributions,
287 i.e., $\gamma_l \neq \gamma_u$, on CIFAR-10-LT and STL10-LT datasets. The γ_l is fixed to 100 for CIFAR-10-LT,
288 while it is set to 10 and 20 for STL10-LT dataset. The best results are in **bold**.
289

290	291	CIFAR10-LT ($\gamma_l \neq \gamma_u$)								STL10-LT ($\gamma_u = N/A$)								
		$\gamma_u = 1$ (uniform)				$\gamma_u = 1/100$ (reversed)				$\gamma_l = 10$				$\gamma_l = 20$				
		$N_1 = 500$	$N_1 = 1500$	$M_1 = 4000$	$M_1 = 3000$	$N_1 = 500$	$N_1 = 1500$	$M_1 = 4000$	$M_1 = 3000$	$N_1 = 150$	$N_1 = 450$	$M_1 = 100k$	$M_1 = 100k$	$N_1 = 150$	$N_1 = 450$	$M_1 = 100k$	$M_1 = 100k$	
292	Algorithm	FixMatch	73.0 \pm 3.81	81.5 \pm 1.15	82.5 \pm 0.75	84.6 \pm 0.34	62.5 \pm 0.94	71.8 \pm 1.70	70.1 \pm 0.22	80.0 \pm 0.93	56.1 \pm 2.32	72.4 \pm 0.71	66.9 \pm 1.66	75.6 \pm 0.45	47.6 \pm 4.87	64.0 \pm 2.27		
293		w/ DARP	83.2 \pm 1.67	87.1 \pm 0.28	82.2 \pm 1.53	86.4 \pm 0.42	70.7 \pm 2.02	80.8 \pm 0.39	62.9 \pm 1.39	72.9 \pm 2.00	61.7 \pm 2.51	71.6 \pm 1.17	61.2 \pm 1.27	71.5 \pm 0.96	57.1 \pm 3.67	68.6 \pm 0.88		
294		w/ CReST	86.6 \pm 0.84	88.8 \pm 0.59	86.1 \pm 0.13	93.4 \pm 0.56	71.0 \pm 0.95	80.3 \pm 0.65	86.1 \pm 0.10	89.8 \pm 0.18	70.0 \pm 1.19	78.4 \pm 0.80	77.1 \pm 1.24	83.0 \pm 0.32	75.1 \pm 0.70	81.5 \pm 0.25		
295		w/ CReST+	92.1 \pm 0.18	93.5 \pm 0.11	93.0 \pm 0.13	93.4 \pm 0.56	85.0 \pm 0.09	89.5 \pm 0.17	86.1 \pm 0.10	89.8 \pm 0.18	77.8 \pm 0.45	83.6 \pm 0.45			75.8 \pm 0.10	82.0 \pm 0.35		
296		w/ DASO																
297		FixMatch + ACR																
298		w/ Ours																

300 normalization in Eq. 7 then yields $\beta_{k,c} \in [0, 1]$, making $(\Delta_k)_c$ a *relative*, rather than absolute, penalty and rendering the dynamic term insensitive to the overall dataset size. Since the SCAD correction term $\sum_k p(k \mid u_j) \Delta_k$ is a convex combination of these bounded adjustments, its components are themselves bounded in $[0, 1]$ and thus act as a moderate, local correction on top of the global LA term $-\log \pi$, whose magnitude can be much larger on highly imbalanced datasets. We empirically found this construction to be numerically stable across all benchmarks. For the *vehicle* super-class, this formulation assigns a high penalty (up to 1) to the dominant *automobile* class and a smaller penalty to the minority *truck* class. For all classes outside this super-class, a fixed maximum penalty (e.g. set to the strongest suppression level) is applied, preventing potential confusion. **This max penalty design is motivated by prior work (Tao et al., 2023)**

301 The final SCAD adjustment is formulated by augmenting the standard LA correction with our
302 sample-aware term. Specifically, we use the inferred super-class posterior $p(k \mid u_j)$ as a dynamic
303 weight to create a convex combination of the pre-computed adjustment vectors, Δ_k . This yields a
304 single, sample-specific adjustment vector. The complete SCAD formulation is:

$$315 \quad \ell_j^{\text{SCAD}} = \ell_j^c - \left(\log \pi + \sum_{k=1}^K p(k \mid u_j) \Delta_k \right), \quad (8)$$

316 This dynamic formulation is the key to SCAD's effectiveness. As illustrated in Figure 2, for an
317 image of a truck, the model's high confidence in the *vehicle* super-class ($p(\text{vehicle} \mid u_j) \approx 1$) ensures
318 that final adjustment is dominated by Δ_{vehicle} . This vector applies a large penalty to the competing
319 majority class (*automobile*) and a smaller penalty to the true class (*truck*), resolving the local
320 ambiguity. In stark contrast, standard LA applies a fixed, global correction that is blind to this local
321 context, making it ill-equipped to handle such *intra-super-class* conflicts.

324
325
326 Table 3: Test accuracy on CIFAR100-LT under *uniform* and *reversed* distributions.
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377

CIFAR100-LT ($\gamma_l \neq \gamma_u$)				
	$\gamma_u = 1$ (uniform)		$\gamma_u = 1/10$ (reversed)	
Algorithm	$N_1 = 50$	$N_1 = 150$	$N_1 = 50$	$N_1 = 150$
	$M_1 = 400$	$M_1 = 300$	$M_1 = 400$	$M_1 = 300$
FixMatch	45.5 ± 0.71	58.1 ± 0.72	44.2 ± 0.43	57.3 ± 0.19
w/ DARP	43.5 ± 0.95	55.9 ± 0.32	36.9 ± 0.48	51.8 ± 0.92
w/ CReST	43.5 ± 0.30	59.2 ± 0.25	39.0 ± 1.11	56.4 ± 0.62
w/ CReST+	43.6 ± 1.60	58.7 ± 0.16	39.1 ± 0.77	56.4 ± 0.78
w/ DASO	53.9 ± 0.66	61.8 ± 0.98	51.0 ± 0.19	60.0 ± 0.31
w/ Ours	54.0 ± 0.77	62.4 ± 0.98	48.6 ± 0.19	60.4 ± 0.10
FixMatch + ACR	57.2 ± 0.19	66.7 ± 0.30	51.6 ± 0.12	62.9 ± 0.25
w/ Ours	59.1 ± 0.25	66.8 ± 0.22	53.4 ± 0.11	63.3 ± 0.12

340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377

Table 4: Comparison with other LTSSL baselines on CIFAR10/100-LT.

Algorithm	CIFAR-10-LT				CIFAR-100-LT
	$\gamma_l = \gamma_u = 100$	$\gamma_l = 100, \gamma_u = 1$	$\gamma_l = 100, \gamma_u = 1/100$	$\gamma_l = \gamma_u = 20$	
FixMatch + SAW	77.5 ± 0.65	81.2 ± 0.68	72.3 ± 0.65	50.1 ± 0.10	
w/ Ours	79.4 ± 0.70	81.1 ± 0.30	75.7 ± 0.57	53.1 ± 0.16	
FixMatch + ABC	81.1 ± 1.14	82.7 ± 0.40	68.9 ± 0.61	53.3 ± 0.79	
w/ Ours	82.0 ± 0.30	83.0 ± 0.42	73.2 ± 1.50	55.1 ± 0.19	
FixMatch + CoSSL	83.1 ± 0.45	88.8 ± 0.42	85.1 ± 0.58	53.9 ± 0.78	
w/ Ours	83.3 ± 0.32	89.0 ± 0.27	84.8 ± 0.47	55.0 ± 0.21	
FixMatch + CDMAD	83.6 ± 0.46	87.5 ± 0.46	77.6 ± 0.70	54.3 ± 0.44	
w/ Ours	85.1 ± 0.11	87.9 ± 0.10	79.4 ± 1.15	54.9 ± 0.76	

4 EXPERIMENT

353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377

We conducted experiments used datasets in LTSSL, including CIFAR10-LT (Krizhevsky et al., 2009), CIFAR100-LT (Krizhevsky et al., 2009), and STL10-LT (Coates et al., 2011). Additionally, we performed experiments on the ImageNet-127, Food101-LT (Fan et al., 2022) dataset.

378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431
432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
10010
10011
10012
10013
10014
10015
10016
10017
10018
10019
10020
10021
10022
10023
10024
10025
10026
10027
10028
10029
10030
10031
10032
10033
10034
10035
10036
10037
10038
10039
10040
10041
10042
10043
10044
10045
10046
10047
10048
10049
10050
10051
10052
10053
10054
10055
10056
10057
10058
10059
10060
10061
10062
10063
10064
10065
10066
10067
10068
10069
10070
10071
10072
10073
10074
10075
10076
10077
10078
10079
10080
10081
10082
10083
10084
10085
10086
10087
10088
10089
10090
10091
10092
10093
10094
10095
10096
10097
10098
10099
100100
100101
100102
100103
100104
100105
100106
100107
100108
100109
100110
100111
100112
100113
100114
100115
100116
100117
100118
100119
100120
100121
100122
100123
100124
100125
100126
100127
100128
100129
100130
100131
100132
100133
100134
100135
100136
100137
100138
100139
100140
100141
100142
100143
100144
100145
100146
100147
100148
100149
100150
100151
100152
100153
100154
100155
100156
100157
100158
100159
100160
100161
100162
100163
100164
100165
100166
100167
100168
100169
100170
100171
100172
100173
100174
100175
100176
100177
100178
100179
100180
100181
100182
100183
100184
100185
100186
100187
100188
100189
100190
100191
100192
100193
100194
100195
100196
100197
100198
100199
100200
100201
100202
100203
100204
100205
100206
100207
100208
100209
100210
100211
100212
100213
100214
100215
100216
100217
100218
100219
100220
100221
100222
100223
100224
100225
100226
100227
100228
100229
100230
100231
100232
100233
100234
100235
100236
100237
100238
100239
100240
100241
100242
100243
100244
100245
100246
100247
100248
100249
100250
100251
100252
100253
100254
100255
100256
100257
100258
100259
100260
100261
100262
100263
100264
100265
100266
100267
100268
100269
100270
100271
100272
100273
100274
100275
100276
100277
100278
100279
100280
100281
100282
100283
100284
100285
100286
100287
100288
100289
100290
100291
100292
100293
100294
100295
100296
100297
100298
100299
100300
100301
100302
100303
100304
100305
100306
100307
100308
100309
100310
100311
100312
100313
100314
100315
100316
100317
100318
100319
100320
100321
100322
100323
100324
100325
100326
100327
100328
100329
100330
100331
100332
100333
100334
100335
100336
100337
100338
100339
100340
100341
100342
100343
100344
100345
100346
100347
100348
100349
100350
100351
100352
100353
100354
100355
100356
100357
100358
100359
100360
100361
100362
100363
100364
100365
100366
100367
100368
100369
100370
100371
100372
100373
100374
100375
100376
100377
100378
100379
100380
100381
100382
100383
100384
100385
100386
100387
100388
100389
100390
100391
100392
100393
100394
100395
100396
100397
100398
100399
100400
100401
100402
100403
100404
100405
100406
100407
100408
100409
100410
100411
100412
100413
100414
100415
100416
100417
100418
100419
100420
100421
100422
100423
100424
100425
100426
100427
100428
100429
100430
100431
100432
100433
100434
100435
100436
100437
100438
100439
100440
100441
100442
100443
100444
100445
100446
100447
100448
100449
100450
100451
100452
100453
100454
100455
100456
100457
100458
100459
100460
100461
100462
100463
100464
100465
100466
100467
100468
100469
100470
100471
100472
100473
100474
100475
100476
100477
100478
100479
100480
100481
100482
100483
100484
100485
100486
100487
100488
100489
100490
100491
100492
100493
100494
100495
100496
100497
100498
100499
100500
100501
100502
100503
100504
100505
100506
100507
100508
100509
100510
100511
100512
100513
100514
100515
100516
100517
100518
100519
100520
100521
100522
100523
100524
100525
100526
100527
100528
100529
100530
100531
100532
100533
100534
100535
100536
100537
100538
100539
100540
100541
100542
100543
100544
100545
100546
100547
100548
100549
100550
100551
100552
100553
100554
100555
100556
100557
100558
100559
100560
100561
100562
100563
100564
100565
1005

378 Table 5: Test accuracy on ImageNet-127. The best results are in **bold**.
379

Algorithm / ImageNet-127	32 × 32	64 × 64
FixMatch	29.7	42.3
w/ DARP	30.5	42.5
w/ DARP+cRT	39.7	51.0
w/ CReST+	32.5	44.7
w/ CReST++LA	40.9	55.9
w/ CoSSL	43.7	53.9
w/ ACR	57.2	63.6
w/ Ours	60.1	66.7
w/ ACR + Ours	60.5	67.0

391 4.2 RESULTS ON CIFAR10/100-LT AND STL10-LT
392


393 **In case of $\gamma_\ell = \gamma_u$** : To demonstrate the effectiveness of our proposed framework, we conducted
394 experiments by additionally incorporating FixMatch + DASO and FixMatch + ACR for comparison.
395 As shown in Table 1, SCAD enhances performance across almost all settings. It shows that SCAD
396 integrates effectively with various existing LTSSL algorithms, yielding a complementary effect. In
397 particular, the CIFAR100-LT performance was impressive with the ACR, where the settings $N_1 =$
398 50 and $M_1 = 400$ achieved state-of-the-art performance improvements of 2.7% and 2.2% when
399 $\gamma_\ell = \gamma_u = 10$ and $\gamma_\ell = \gamma_u = 20$.

400 **In case of $\gamma_\ell \neq \gamma_u$** : In the real world, there can be distributional differences between unlabeled
401 and labeled datasets. Therefore, we conducted experiments with both uniform and reverse settings
402 to explore these scenarios.
403

404 We observed overall performance improvements when SCAD was applied to ACR in the uniform
405 and reverse settings. Specifically, as shown in Table 3, SCAD marginally outperformed DASO on
406 CIFAR100-LT in both the uniform and reverse settings. Furthermore, combining ACR with SCAD
407 led to performance gains of up to 2% and 3.4% in the uniform and reverse settings, respectively.
408

409 4.3 RESULTS ON IMAGENET-127 AND FOOD101-LT
410

411 On ImageNet-127 (Fan et al., 2022), a challenging 127-class subset of ImageNet with a
412 severe long-tail distribution ($\gamma_u \approx 286$), our method shows a distinct advantage. Following
413 the standard setup with 10% labeled data, our
414 approach outperforms existing baselines, as
415 detailed in Table 5. Notably, it surpasses the
416 strong ACR baseline by a substantial margin
417 of 5.0% at 32×32 resolution and 4.0% at
418 64×64 resolution. This improvement is partic-
419 ularly noteworthy as ACR represents a highly
420 competitive recent method, highlighting the
421 efficacy of our proposed hierarchical debiasing.
422 Furthermore, when our method is combined
423 with ACR, it achieves new state-of-the-art
424 results, underscoring its high compatibility
425 as a plug-in module. These significant gains,
426 consistent with results on Food101-LT (see
427 Appendix A.6), validate that our approach
428 effectively **alleviates** the critical problem of
429 *intra-super-class imbalance*—a fine-grained
430 challenge that existing methods have struggled
431 to address.



432 Figure 3: **Left:** figure displays the average con-
433 fidence scores for minority-class samples (8, 9) of
434 the super-class *vehicle*. Our method effectively
435 reduces confidence in confusing majority classes
436 (0, 1) within the same super-class while increas-
437 ing confidence in the correct minority classes
438 (8, 9). **Right:** illustration demonstrates that our
439 approach consistently improves accuracy across
440 various LTSSL frameworks, especially benefit-
441 ing minority classes.

432 Table 6: Ablation studies of SCAD on CIFAR10-LT and STL10-LT datasets.
433

Ablations	CIFAR10-LT	STL10-LT
FixMatch	67.8	56.1
+ Super-class learning	69.2	69.0
+ Logit-Adjustment (LA)	76.9	70.4
+ Super-class-Aware Debiasing (SCAD)	78.7	71.3

440 Table 7: Comparison of SCAD with ground truth and various text encoders on CIFAR100-LT.
441

Algorithms	CIFAR100-LT
FixMatch + LA	47.3
w/ Ours with ground truth	50.3
w/ Ours with GloVe	49.7
w/ Ours with SBERT	50.1
w/ Ours with CLIP text encoder	50.4
w/ Ours with text-embedding-ada-002	50.5

449
450 4.4 RESULTS ON OTHER LTSSL BASELINES
451

452 Table 4 shows that our framework can also be applied effectively to other experiment settings
453 with state-of-the-art methods, we conducted experiments on the three settings of CIFAR10-LT and
454 consistent setting of CIFAR100-LT within the CDMAD (Lee & Kim, 2024), CoSSL (Fan et al.,
455 2022), ABC (Lee et al., 2021) and SAW (Lai et al., 2022) framework. According to CDMAD,
456 when we ran experiments with ACR under CDMAD settings, we observed a performance drop
457 from 84.1 to 81.8 at $N_1 = 1500$ and $M_1 = 3000$. Our framework remains robust when combined
458 with existing LTSSL in different settings. According to Table 4, when combining CDMAD with
459 our framework, which is another state-of-the-art method, there was a 1.6% improvement in the
460 consistent setting, a 0.3% increase in the uniform setting, and a 2.3% improvement in the reverse
461 setting. These results demonstrate that our proposed SCAD framework effectively integrates with
462 diverse LTSSL baselines, further validating its utility and applicability.

463
464 4.5 COMPREHENSIVE ANALYSIS OF OUR METHOD
465

466 **Ablation study** According to Table 6, we evaluated the performance improvements achieved by
467 incrementally adding each component to the algorithms on CIFAR10-LT and STL10-LT under a
468 consistent setting. First, adding the super-class learning component resulted in significant perfor-
469 mance gains, with an increase of 2.0% on CIFAR10-LT and 23% on STL10-LT. Finally, adding
470 SCAD yielded additional gains of 9.2% on CIFAR10-LT and 1.1% on STL10-LT. The limited im-
471 provement observed from adding SCAD on STL10-LT is likely due to the significant difference in
472 the number of labeled and unlabeled datasets.

473 **Comparison with ground truth super-class labels and various text embedding models** As
474 shown in Table 7, there is no performance difference between using our super-class labels generated
475 from text-embeddings-ada-002 (Ryan et al., 2022) or the CLIP text encoder (Radford et al., 2021),
476 and using the ground-truth super-class labels provided in CIFAR100. We further evaluate SCAD on
477 CIFAR100-LT using weaker embeddings such as GloVe (Pennington et al., 2014), SBERT (Nils &
478 Iryna, 2019), and still outperform the baseline. This indicates that our approach remains effective
479 even without large-scale models.

480 **Minority class accuracy comparison** As shown in Figure 3, Our method improves the average
481 confidence of minority classes (8,9) over the LA baseline, while correspondingly reducing the confi-
482 dence of majority classes (0,1) within the same super-class. This result yields consistent gains across
483 diverse LTSSL frameworks, highlighting its general applicability. Moreover, since our approach
484 merely appends a super-class classifier trained in a joint manner, it entails negligible computational
485 overhead ([more detailed in Appendix A.7](#)).

486 **Hyperparameter settings** we define the confidence threshold τ for class-level predictions as 0.95,
 487 following FixMatch. Since super-class learning is expected to converge faster than class-level learn-
 488 ing, we set the super-class threshold τ_s to the same value, i.e., $\tau_s = 0.95$, for simplicity and conve-
 489 nience. We determined the number of super-classes K by setting $K = \lceil \frac{C}{4} \rceil$. For a detailed analysis,
 490 choosing K and ensuring super-class balance, please refer to Appendix A.4.
 491

492 5 CONCLUSION

493
 494 In this paper, we introduce Super-Class-Aware Debiasing (SCAD), a novel framework designed
 495 to tackle the fundamental challenge of *intra-super-class imbalance* in long-tailed semi-supervised
 496 learning (LTSSL). Unlike previous approaches that rely solely on global class priors for logit ad-
 497 justment, SCAD dynamically adapts the corrective force to the semantic context of each sample.
 498 To validate the effectiveness of our approach, we conduct extensive experiments on a wide range
 499 of LTSSL benchmarks. The results demonstrate that SCAD substantially improves the reliability of
 500 pseudo-labels, enhances the learning dynamics of minority classes, and consistently narrows the per-
 501 formance gap caused by long-tailed distributions. Moreover, SCAD achieves state-of-the-art results
 502 across diverse datasets, outperforming strong baselines and existing debiasing methods. These find-
 503 ings highlight that incorporating semantic structure and local class relationships into the debiasing
 504 process is a powerful strategy, and establish SCAD as a new standard for robust and generalizable
 505 learning in real-world imbalanced settings.

506 6 LIMITATIONS AND FUTURE WORKS

507
 508 Our study mainly evaluates SCAD within LTSSL. Its effectiveness for other tasks such as detection
 509 or segmentation remains an open question and would require further validation. A promising di-
 510 rection for future work is to extend SCAD beyond classification. In particular, applying our idea to
 511 more complex tasks such as object detection or semantic segmentation could further demonstrate its
 512 utility, as these tasks often suffer from both long-tailed distributions and semantic overlap among
 513 categories.

514
 515
 516
 517
 518
 519
 520
 521
 522
 523
 524
 525
 526
 527
 528
 529
 530
 531
 532
 533
 534
 535
 536
 537
 538
 539

540 REFERENCES
541

- 542 David Berthelot, Nicholas Carlini, Ekin D Cubuk, Alexey Kurakin, Kihyuk Sohn, Han Zhang, and
543 Colin A Raffel. Remixmatch: Semi-supervised learning with distribution matching and augmen-
544 tation anchoring. In *International Conference on Learning Representations (ICLR)*, 2020.
- 545 Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101—mining discriminative compo-
546 nents with random forests. In *European conference on computer vision*, pp. 446–461. Springer,
547 2014.
- 548 Adam Coates, Andrew Ng, and Honglak Lee. An analysis of single-layer networks in unsupervised
549 feature learning. in proceedings of the fourteenth international conference on artificial intelligence
550 and statistics. *JMLR Workshop and Conference Proceedings*, 2011.
- 551 Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical automated
552 data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF conference on*
553 *computer vision and pattern recognition workshops*, pp. 702–703, 2020.
- 554 Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. Class-balanced loss based on
555 effective number of samples. In *Proceedings of the IEEE/CVF Conference on Computer Vision*
556 and *Pattern Recognition (CVPR)*, pp. 9268–9277, 2019.
- 557 Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hier-
558 archical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*
559 (*CVPR*), pp. 248–255, 2009.
- 560 Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep
561 bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of*
562 *the North American Chapter of the Association for Computational Linguistics: Human Language*
563 *Technologies, Volume 1 (Long and Short Papers)*, pp. 4171–4186, 2019.
- 564 Terrance DeVries and Graham W Taylor. Improved regularization of convolutional neural networks
565 with cutout. *arXiv preprint arXiv:1708.04552*, 2017.
- 566 Yingjun Du, Jiayi Shen, Xiantong Zhen, and Cees G. M. Snoek. Superdisco: Super-class discovery
567 improves visual recognition for the long-tail. In *Proceedings of the IEEE/CVF Conference on*
568 *Computer Vision and Pattern Recognition (CVPR)*, pp. 19944–19954, 2023.
- 569 Yue Fan, Dengxin Dai, Anna Kukleva, and Bernt Schiele. CossL: Co-learning of representation and
570 classifier for imbalanced semi-supervised learning. In *Proceedings of the IEEE/CVF Conference*
571 *on Computer Vision and Pattern Recognition (CVPR)*, pp. 14574–14584, 2022.
- 572 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recog-
573 nition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp.
574 770–778, 2016.
- 575 Bingyi Kang, Saining Xie, Marcus Rohrbach, Zhicheng Yan, Albert Gordo, Jiashi Feng, and Yannis
576 Kalantidis. Decoupling representation and classifier for long-tailed recognition. In *International*
577 *Conference on Learning Representations (ICLR)*, 2020.
- 578 Jaehyung Kim, Youngbum Hur, Sejun Park, Eunho Yang, Sung Ju Hwang, and Jinwoo Shin. Dis-
579 tribution aligning refinery of pseudo-label for imbalanced semi-supervised learning. In *Advances*
580 *in Neural Information Processing Systems (NeurIPS)*, pp. 14567–14579, 2020.
- 581 Alex Krizhevsky et al. Learning multiple layers of features from tiny images. *arXiv preprint*
582 *arXiv:0912.1234*, 2009.
- 583 Lai, Wang, et al. Smoothed adaptive weighting for imbalanced semi-supervised learning: Im-
584 prove reliability against unknown distribution data. In *ICML*, 2022. URL <https://api.semanticscholar.org/CorpusID:250341007>.
- 585 Hyuck Lee and Heeyoung Kim. Cdmad: Class-distribution-mismatch-aware debiasing for class-
586 imbalanced semi-supervised learning. In *Proceedings of the IEEE/CVF Conference on Computer*
587 *Vision and Pattern Recognition (CVPR)*, pp. 23891–23900, 2024.

- 594 Hyuck Lee, Seungjae Shin, and Heeyoung Kim. Abc: Auxiliary balanced classifier for class-
 595 imbalanced semi-supervised learning. In *Advances in Neural Information Processing Systems*
 596 (*NeurIPS*), pp. 7082–7094, 2021.
- 597 Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr
 598 Dollár, and C. Lawrence Zitnick. Microsoft COCO: Common objects in context. In *European*
 599 *Conference on Computer Vision (ECCV)*, pp. 740–755, 2014.
- 600 Aditya Krishna Menon, Sadeep Jayasumana, Ankit Singh Rawat, Himanshu Jain, Andreas Veit, and
 601 Sanjiv Kumar. Long-tail learning via logit adjustment. In *International Conference on Learning*
 602 *Representations (ICLR)*, 2021.
- 603 Nils and Iryna. Sentence-bert: Sentence embeddings using siamese bert-networks, 2019. URL
 604 <https://arxiv.org/abs/1908.10084>.
- 605 Youngtaek Oh, Dong-Jin Kim, and In So Kweon. Daso: Distribution-aware semantics-oriented
 606 pseudo-label for imbalanced semi-supervised learning. In *CVPR*, 2022.
- 607 Jeffrey Pennington, Richard Socher, and Christopher D Manning. Glove: Global vectors for word
 608 representation. In *Proceedings of the 2014 conference on empirical methods in natural language*
 609 *processing (EMNLP)*, pp. 1532–1543, 2014.
- 610 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,
 611 Girish Sastry, Amanda Askell, Pam Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever.
 612 Learning transferable visual models from natural language supervision. In *Proceedings of the*
 613 *International Conference on Machine Learning (ICML)*, 2021.
- 614 Ryan, Ted, et al. New and improved embedding model. *OpenAI*, 2022. URL <https://openai.com/blog/new-and-improved-embedding-model>.
- 615 Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin A Raffel, Ekin D
 616 Cubuk, Alexey Kurakin, and Chun-Liang Li. Fixmatch: Simplifying semi-supervised learning
 617 with consistency and confidence. In *Advances in Neural Information Processing Systems*
 618 (*NeurIPS*), pp. 596–608, 2020.
- 619 Yingfan Tao, Jingna Sun, Hao Yang, Li Chen, Xu Wang, Wenming Yang, Daniel Du, and Min Zheng.
 620 Local and global logit adjustments for long-tailed learning. In *Proceedings of the IEEE/CVF*
 621 *international conference on computer vision*, pp. 11783–11792, 2023.
- 622 Xudong Wang, Zhirong Wu, Long Lian, and Stella X Yu. Debiased learning from naturally im-
 623 balanced pseudo-labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and*
 624 *Pattern Recognition (CVPR)*, pp. 14647–14657, 2022.
- 625 Chen Wei, Kihyuk Sohn, Clayton Mellina, Alan Yuille, and Fan Yang. Crest: A class-
 626 rebalancing self-training framework for imbalanced semi-supervised learning. In *Proceedings*
 627 *of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 10857–
 628 10866, 2021.
- 629 Tong Wei and Kai Gan. Towards realistic long-tailed semi-supervised learning: Consistency is
 630 all you need. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern*
 631 *Recognition (CVPR)*, pp. 3469–3478, 2023.
- 632 Tz-Ying Wu, Pedro Morgado, Pei Wang, Chih-Hui Ho, and Nuno Vasconcelos. Solving long-
 633 tailed recognition with deep realistic taxonomic classifier. In *Proceedings of the 16th European*
 634 *Conference on Computer Vision (ECCV)*, volume 12353 of *Lecture Notes in Computer Science*,
 635 pp. 171–189. Springer, 2020. doi: 10.1007/978-3-030-58598-3_11.
- 636 Zhicheng Yan, Hao Zhang, Robinson Piramuthu, Vignesh Jagadeesh, Dennis DeCoste, Wei Di, and
 637 Yizhou Yu. HD-CNN: Hierarchical deep convolutional neural networks for large scale visual
 638 recognition. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*,
 639 pp. 2740–2748. IEEE Computer Society, 2015.
- 640 Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv preprint*
 641 *arXiv:1605.07146*, 2016.

648 **A APPENDIX**
649650 **A.1 THE ASSISTANCE OF A LARGE LANGUAGE MODEL (LLM)**
651652 This work used a large language model (LLM) solely for language editing (grammar, wording, and
653 minor typographical corrections) during manuscript preparation.
654655 **A.2 EXPERIMENT SETTING**
656658 For a fair comparison, we integrate our framework with the DASO (Oh et al., 2022), ACR (Wei
659 & Gan, 2023), SAW (Lai et al., 2022), ABC (Lai et al., 2022), CoSSL (Fan et al., 2022), and
660 CDMAD (Lee & Kim, 2024) algorithms using their official code repositories. For example, we
661 use the code for DASO from <https://github.com/ytaek-oh/daso.git>, for ACR from
662 <https://github.com/Gank0078/ACR.git>, and for CDMAD from <https://github.com/LeeHyuck/CDMAD.git>.
663664 When combining FixMatch (Sohn et al., 2020) with our framework, we adopt the ACR settings. It
665 is worth noting that the settings for each algorithm exhibit slight differences. For instance, ACR and
666 CDMAD utilize distinct configurations. The ACR setting is specifically tailored to generate weakly
667 and strongly augmented versions of input images for use within the FixMatch framework.
668

- 669
- **Weak augmentation:** This pipeline applies random horizontal flipping and random crop-
670 ping with reflective padding to preserve the image structure. These transformations are
671 designed to introduce minimal changes to the input image while maintaining its semantic
672 content.
 - **Strong augmentation:** In addition to the weak augmentations, this pipeline uses the Ran-
673 dAugment (Cubuk et al., 2020) method, which applies 2 transformations with 10 magnitude
674 with Cutout (DeVries & Taylor., 2017) operation.
675

676 In CDMAD setting, Weak augmentation is same however, there are slight difference in Strong aug-
677mentation. CDMAD uses RandAugment with 3 transformations of magnitude 4, followed by the
678 Cutout operation.
679680 The hyperparameter settings differ significantly between ACR and CDMAD. In ACR, the labeled
681 dataset batch size is set to 64, the ratio of the unlabeled batch size to the labeled batch size is set to
682 1, and the optimizer’s learning rate is adjusted using a cosine decay schedule. The learning rate at
683 each step is calculated as $\alpha \cos\left(\frac{7\pi t}{16T}\right)$, where α is the initial learning rate set to 3×10^{-3} , t denotes
684 the current iteration and T is the total number of iterations.
685686 In contrast, CDMAD uses a labeled dataset batch size of 32, the ratio of the unlabeled batch size to
687 the labeled batch size is set to 2, and a learning rate of 2×10^{-3} . All experiments are conducted
688 using one NVIDIA RTX 3090 GPU. The software environment includes PyTorch 2.1.2, TorchVision
0.16.2, and CUDA 12.1.
689690 All settings, We performed the experiments on CIFAR10-LT (Krizhevsky et al., 2009), CIFAR100-
691 LT (Krizhevsky et al., 2009), and STL-10 (Coates et al., 2011) using the WideResNet-28-2
692 (Zagoruyko & Komodakis, 2016) architecture, and on ImageNet-127 (Fan et al., 2022), Food101-
693 LT(Fan et al., 2022) using the ResNet-50 (He et al., 2016) architecture.
694695 **A.3 INTRA-SUPERCLASS-IMBALANCE PROBLEM ON CIFAR100-LT**
696698 We extended our investigation to CIFAR100-LT to prove the prevalence of intra-super-class imbal-
699 ance. Our results (see Figure 4) show that even on CIFAR-100-LT, substantial imbalance persists
700 within super-classes when using standard methods such as LA: minority classes inside each semantic
701 group remain underserved despite global prior correction. In contrast, SCAD consistently reduces
these intra-super-class disparities by explicitly targeting the local bias. These additional findings




Figure 4: Illustration of the *intra-super-class imbalance* problem on CIFAR100-LT. Within the Figure 4a, 4b super-class (e.g., camel and elephant), the model frequently confuses semantically related classes, and SCAD effectively alleviates these errors. A similar pattern appears in the Figure 4c, 4d super-class (e.g., porcupine and possum), where SCAD reduces misclassifications by addressing the local dominance within the group.




Figure 5: Illustration for balancing super-classes and determining the optimal value of K on CIFAR10-LT

suggest that intra-super-class imbalance is not a CIFAR10-specific artifact and that our solution is robust across different long-tailed datasets.

A.4 ANALYSIS OF SUPER-CLASS LEVEL DISTRIBUTION

According to Figure 5a, we set $K = \lceil \frac{C}{4} \rceil$. In addition, in Figure 5b we compare the class-level and superclass-level imbalance ratios of unlabeled data under the consistent, uniform, and reverse settings. Note that in the reverse setting, the reciprocal values are not employed. We observe that the imbalance ratio drops sharply in the consistent and reverse settings, while it slightly increases in the uniform setting but still remains nearly identical. Furthermore, owing to this more balanced distribution structure, Figure 5 reports class and super-class level accuracies across training epochs with our method, consistently achieves higher performance at all epochs. These results indicate that the super-class classifier is more balanced and reliable than the class-level counterpart. Experiments on CIFAR100-LT show a similar trend: According to Figure 6a, the performance reaches its maximum around $K = \lceil \frac{C}{4} \rceil$, and when the original unlabeled imbalance ratio γ_u is 10 or lower, the performance difference becomes negligible. (See Figure 6b)

A.5 VISUALIZATION ON SCAD

To visually demonstrate that our proposed method is particularly beneficial for minority classes, we further compare the confusion matrices of FixMatch, FixMatch + Ours, ACR and ACR+Ours. Specifically, as shown in Figure 7, our method significantly enhances classification performance for minority classes under consistent settings, highlighting the effectiveness of SCAD in alleviating class imbalance.




Figure 6: Illustration for balancing super-classes and determining the optimal value of K on CIFAR100-LT




Figure 7: Confusion matrices of the test set for FixMatch, FixMatch+Ours, ACR, and ACR+Ours under consistent settings.

A.6 FOOD101-LT RESULTS

Food101-LT, we sample 250 labeled and 500 unlabeled images per class, corresponding to imbalance ratios of $\gamma_l = \gamma_u = 50, 100$, followed by CoSSL. According to Table 8, Our method also improves the accuracy on Food101-LT about 3% and 7%, respectively.

Table 8: Test accuracy on Food101-LT with $\gamma_l = \gamma_u = 50, 100$. The best results are in **bold**.

Algorithm	$\gamma = 50$	$\gamma = 100$
FixMatch	42.6	35.3
w/ DARP	42.0	34.2
w/ DARP + cRT	41.5	34.4
w/ CReST+	43.8	31.2
w/ CReST+ + LA	47.7	36.1
w/ CoSSL	49.0	40.4
w/ Ours	51.0	43.5

A.7 CALIBRATION ANALYSIS ON SCAD

We further evaluate the calibration properties of SCAD compared to FixMatch and LA baselines. We measure calibration using Expected Calibration Error (ECE), which quantifies the discrepancy between predicted confidence and empirical accuracy. FixMatch exhibits severe miscalibration in tail classes, with ECE values reaching up to 0.59, while SCAD reduces these errors substantially. This indicates that SCAD alleviates underconfidence in tail classes and achieves a more balanced calibration across head and tail. We further evaluated calibration by integrating SCAD into ACR and observed consistent improvements across all frequency groups. The ECE drops from 0.1018 to 0.0517 (Head), 0.1339 to 0.1228 (Medium), and from 0.1654 to 0.0926 (Tail). These results show




Figure 8: ECE analysis on CIFAR10-LT

Table 9: ECE comparison between ACR and ACR+Ours

Method	Head	Medium	Tail
ACR	0.1018	0.1339	0.1654
ACR+Ours	0.0517	0.1228	0.0926

that SCAD not only preserves performance but also substantially enhances prediction reliability, especially for tail classes.

A.8 TRAINING TIME ANALYSIS

As presented in Table 10, the reported time represents the training time per epoch. The difference between our method and the baseline is approximately 1 second, which is negligible in practice.

A.9 DETAILED INFORMATION ABOUT THE SUPER-CLASS

We retrieved descriptions for each class name using the Wikipedia API, and for convenience, we only used the first sentence of the returned summary. If no corresponding Wikipedia entry exists, we simply use the class name itself. Agglomerative clustering is performed using the clustering module from scikit-learn. As previously mentioned, we set $K = \lceil \frac{C}{4} \rceil$. Specifically, for CIFAR10-LT (Krizhevsky et al., 2009) and STL10-LT (Coates et al., 2011), $K = 3$, for CIFAR100-LT (Krizhevsky et al., 2009), $K = 25$ and for ImageNet-127 (Fan et al., 2022), $K = 32$. The groupings for CIFAR10-LT, STL10-LT, CIFAR100-LT and ImageNet-127 are detailed in Table 11 and Table 13. These groupings are determined using the text-embedding-ada-002 model (Ryan et al., 2022).

864
865
866
867

Table 10: **Training time comparison (seconds) per epoch**

868
869
870
871
872
873
874
875
876
877
878
879

Method	FixMatch	FixMatch+Ours	ACR	ACR+Ours
Time (s)	28	29	35	36

Table 11: Super-class table on CIFAR10-LT. We use the Wikipedia API to retrieve a summary for a given query and perform agglomerative clustering using the text-embedding-ada-002 model, setting the number of clusters to 3.

880
881
882
883
884
885
886
887
888
889

Cluster	Classes
Cluster 0	0: airplane, 1: automobile, 8: ship, 9: truck
Cluster 1	2: bird, 6: frog
Cluster 2	3: cat, 4: deer, 5: dog, 7: horse

Table 12: Super-class table on CIFAR100-LT. We use the Wikipedia API to retrieve a summary for a given query and perform agglomerative clustering using the text-embedding-ada-002 model, setting the number of clusters to 25.

890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917

Cluster	Classes
Cluster 0	1: aquarium_fish, 26: crab, 45: lobster, 77: snail, 99: worm
Cluster 1	27: crocodile, 29: dinosaur, 44: lizard, 73: shark, 78: snake, 93: turtle
Cluster 2	21: chimpanzee, 42: leopard, 43: lion, 88: tiger
Cluster 3	54: orchid, 62: poppy, 70: rose, 82: sunflower, 92: tulip
Cluster 4	58: pickup_truck, 69: rocket, 85: tank, 89: tractor
Cluster 5	2: baby, 11: boy, 35: girl, 46: man, 98: woman
Cluster 6	22: clock, 86: telephone, 87: television
Cluster 7	6: bee, 7: beetle, 14: butterfly, 18: caterpillar, 24: cockroach, 79: spider
Cluster 8	36: hamster, 50: mouse, 65: rabbit, 74: shrew, 75: skunk, 80: squirrel
Cluster 9	0: apple, 51: mushroom, 53: orange, 57: pear, 83: sweet_pepper
Cluster 10	5: bed, 20: chair, 25: couch, 94: wardrobe
Cluster 11	23: cloud, 71: sea
Cluster 12	33: forest, 49: mountain, 60: plain
Cluster 13	9: bottle, 10: bowl, 28: cup
Cluster 14	13: bus, 81: streetcar, 90: train
Cluster 15	15: camel, 19: cattle, 31: elephant
Cluster 16	17: castle, 37: house, 76: skyscraper
Cluster 17	16: can, 39: keyboard, 40: lamp, 61: plate, 84: table
Cluster 18	32: flatfish, 67: ray, 91: trout
Cluster 19	30: dolphin, 95: whale, 72: seal
Cluster 20	8: bicycle, 41: lawn_mower, 48: motorcycle
Cluster 21	3: bear, 4: beaver, 34: fox, 55: otter, 97: wolf
Cluster 22	12: bridge, 68: road
Cluster 23	38: kangaroo, 63: porcupine, 66: raccoon, 64: possum
Cluster 24	47: maple_tree, 52: oak_tree, 56: palm_tree, 59: pine_tree, 96: willow_tree

918
 919
 920
 921
 922
 923
 924
 925
 926
 927
 928
 929
 930
 931
 932

933 Table 13: Super-class table on ImageNet-127. We use the Wikipedia API to retrieve a summary
 934 for a given query and perform agglomerative clustering using the text-embedding-ada-002 model,
 935 setting the number of clusters to 32.

936

937 Cluster	938 Classes
Cluster 0	17: electronic_equipment, 43: light, 60: camera, 99: magazine, 125: monitor
Cluster 1	16: furniture, 58: litter, 118: cushion
Cluster 2	2: jewelry, 22: hairpiece, 36: protective_garment, 44: garment, 45: nightwear, 50: uniform, 55: glove, 84: hosiery, 90: gown, 113: dress
Cluster 3	27: pot, 28: dish, 30: rod, 42: home_appliance, 71: toiletry, 80: pan, 94: cleaning_implement, 106: cooker, 109: turner
Cluster 4	29: athlete, 34: game_equipment, 41: golf_equipment, 47: gymnastic_apparatus, 53: sports_implement, 110: participant, 115: diver
Cluster 5	31: frozen_desert, 37: sauce, 39: dip, 63: pudding, 74: entree, 77: cracker, 111: concoction
Cluster 6	9: platform, 32: weight, 59: system, 76: structure, 105: sign, 121: drill_rig
Cluster 7	19: plaything
Cluster 8	75: ceratopsian, 86: mammal, 104: bird
Cluster 9	25: bus, 46: fare, 56: sled, 98: passenger_train
Cluster 10	38: ray, 79: plectognath, 107: shark, 108: ganoid, 124: food_fish, 126: soft_finned_fish
Cluster 11	72: wine, 81: punch, 114: coffee
Cluster 12	8: globule, 21: stick, 66: tissue, 83: covering
Cluster 13	11: fungus, 14: coelenterate, 54: echinoderm, 65: arthropod, 68: mollusk
Cluster 14	88: rescue_equipment, 91: safety_belt
Cluster 15	52: spacecraft, 64: rocket, 67: heavier_than_air_craft, 112: lighter_than_air_craft
Cluster 16	18: sailing_vessel, 35: ship, 93: oar, 119: boat
Cluster 17	6: anguid_lizard, 24: teiid_lizard, 49: lacertid_lizard, 78: gecko, 87: agamid, 102: venomous_lizard, 116: iguanid, 122: chameleon
Cluster 18	20: board, 120: bar
Cluster 19	95: bath_linen, 103: fabric
Cluster 20	1: vegetable, 4: fruit, 40: fodder, 51: flower
Cluster 21	70: pen, 101: eraser
Cluster 22	3: ligament
Cluster 23	13: frog, 89: snake, 96: turtle, 97: salamander, 100: worm
Cluster 24	69: cistern, 117: container
Cluster 25	0: loaf_of_bread, 73: bun
Cluster 26	5: cap, 26: helmet, 123: hat
Cluster 27	15: alligator, 62: crocodile
Cluster 28	12: device, 23: tool, 48: sharpener, 57: kit, 85: key, 92: duplicator
Cluster 29	10: military_vehicle
Cluster 30	33: geologicalFormation
Cluster 31	7: damselfish, 61: butterfly_fish, 82: scorpaenid

957

958

959

960

961

962

963

964

965

966

967

968

969

970

971