

# Flexible Distribution Alignment: Towards Long-tailed Semi-supervised Learning with Proper Calibration

Emanuel Sanchez Aimar<sup>✉</sup>, Nathaniel Helgesen<sup>✉</sup>,  
Yonghao Xu<sup>✉</sup>, Marco Kuhlmann<sup>✉</sup>, Michael Felsberg<sup>✉</sup>

Linköping University, Sweden

{emanuel.sanchez.aimar,nathaniel.helgesen,yonghao.xu,marco.kuhlmann,michael.felsberg}@liu.se


**Abstract.** Long-tailed semi-supervised learning (LTSSL) represents a practical scenario for semi-supervised applications, challenged by skewed labeled distributions that bias classifiers. This problem is often aggravated by discrepancies between labeled and unlabeled class distributions, leading to biased pseudo-labels, neglect of rare classes, and poorly calibrated probabilities. To address these issues, we introduce Flexible Distribution Alignment (FlexDA), a novel adaptive logit-adjusted loss framework designed to dynamically estimate and align predictions with the actual distribution of unlabeled data and achieve a balanced classifier by the end of training. FlexDA is further enhanced by a distillation-based consistency loss, promoting fair data usage across classes and effectively leveraging under-confident samples. This method, encapsulated in ADELLO (Align and Distill Everything All at Once), proves robust against label shift, significantly improves model calibration in LTSSL contexts, and surpasses previous state-of-art approaches across multiple benchmarks, including CIFAR100-LT, STL10-LT, and ImageNet127, addressing class imbalance challenges in semi-supervised learning. Our code is available at <https://github.com/emasa/ADELLO-LTSSL>.

**Keywords:** Distribution Alignment · Confidence Calibration · Long-tailed · Semi-supervised Learning

## 1 Introduction

Solving computer vision tasks with limited labeled data is a challenging problem that has motivated the development of semi-supervised learning (SSL) [10]. Training models on a mix of labeled and unlabeled data allows costly labeling to be circumvented, though unlabeled data has been shown to complicate training when the distribution of classes is highly imbalanced or follows a long-tailed distribution [53], as depicted in Fig. 1a. Notably, common SSL techniques, namely pseudo-labelling [38] and high-confidence thresholding [58], can lead to imbalanced pseudo-label distributions even in balanced settings [64], producing classifiers that are biased towards head classes [64, 67].

Distribution alignment (DA) [3, 37] aims to mitigate these issues by aligning pseudo-label distributions with actual label priors in balanced [3, 64] and, more generally, long-tailed (LT) settings [64, 68]. Despite recent progress in long-tailed semi-supervised learning (LTSSL) [16, 29, 35, 51], most DA methods assume that labeled and unlabeled data follow the same distribution [3, 64, 68], though in practical or low-label applications, the unlabeled class distribution is more likely to be unknown and distinct from



**Fig. 1:** Long-tailed semi-supervised learning considers a challenging scenario, where a labeled dataset with a **skewed class distribution**,  $\mathcal{P}_L(y)$ , see (a), can bias the model towards frequent classes. This challenge is exacerbated by the use of a larger unlabeled dataset with an **unknown class distribution**,  $\mathcal{Q}(y)$ , see (b), risking the reinforcement of data biases, which in turn leads to **uncalibrated probabilities**. Our evaluation focuses on misclassification error (complement of accuracy) and expected calibration error to test generalization and calibration. Our approach shows consistent improvements in both respects, as shown in (c) and (d).

the labeled distribution, as illustrated in Fig. 1b. This mismatch between labeled and unlabeled data, or *label distribution shift* [24], can lead to *low data utilization*, where models fail to effectively leverage the breadth of available data, especially samples from minority classes [29, 68]. Therefore, incorporating the correct data prior is critical for improving model performance [24, 29, 80]. Previous DA approaches either require access to the ground truth prior or an approximation from balanced holdout data in advance for SSL training [3, 29, 37], conditions that are often challenging to fulfill. Anchor distributions [46, 69] aim to sidestep these requirements, but challenges remain due to implicit assumptions about the class prior and the increased complexity of these approaches.

Model calibration, which guarantees that model predictions align with actual outcomes [20], has gained recent attention for both enabling effective pseudo-labeling in SSL [44] and facilitating distribution alignment in LT fully-supervised scenarios [1, 48, 74]. Beyond mitigating confirmation bias [44], calibration is crucial when using logit adjustment to address label shift [1]. This highlights the importance of further investigating the influence of proper calibration on distribution alignment within LTSSL.

In this work, we tackle distribution misalignment in LTSSL, focusing on the shift between labeled and unlabeled class distributions. We present the following contributions:

- **Flexible Distribution Alignment (FlexDA):** introduces novel supervised and consistency logit-adjusted losses to dynamically align the model with the class distribution of unlabeled data during training, boosting performance across diverse unlabeled class distributions. It integrates a progressive scheduler to adjust the target prior towards a balanced classifier gradually, crucial to accurate debiasing at inference.
- **Complementary Consistency Regularization:** augments FlexDA by distilling underconfident samples, optimizing the utilization of data often ignored due to low-confidence pseudo-labels.
- **Study of Model Calibration in LTSSL:** We delve into the relationship between calibration and generalization in LTSSL, corroborating that improved calibration is highly correlated with enhanced model generalization across various datasets and different degrees of label shift.

The effectiveness of our methodology is showcased in settings with both controlled and unknown label shifts, such as CIFAR100-LT and STL10-LT, where it excels in LTSSL contexts, see Fig. 1d and Fig. 1c. Under matching distributions, it surpasses state-of-the-art (SOTA) approaches in the large-scale dataset ImageNet127. Finally, ADELLO presents significant advances in calibration compared to previous LTSSL approaches, see Fig. 1c. These contributions form the foundation of *Align and Distill Everything All at Once* (ADELLO), outlined in Fig. 2, a versatile framework adept at handling distribution shifts, efficiently overcoming LTSSL challenges, and enhancing model calibration.

## 2 Related Work

**Semi-supervised learning.** Semi-supervised learning is a mature field, with significant advancements made in recent decades [10, 19, 38, 47, 57, 58]. Successful techniques such as pseudo-labeling [2, 38], consistency-regularization [49, 60], and their combinations [2–4, 34] have contributed to this progress. Encouraging consistency between weakly-perturbed data views has shown improvements [36, 56], and further progress has been made with the combination of weak and strong data augmentation [58, 71].

Confidence-based methods combine high-confidence thresholding and pseudo-labeling (hard [58] or temperature-scaled [71]) to minimize confirmation bias [31]. However, a fixed threshold can limit the utilization of unlabeled samples [12, 45], especially in low-labeled regimes. In response, later approaches introduced progressive thresholds [73] and class-specific thresholds [78]. Recently, SoftMatch [12] integrates adaptive thresholding and confidence-based weighting for better sample utilization. Contrasting with these methods is the *knowledge distillation* (KD) approach, characterized by the use of softened output logits to distill information about class similarities [23]. In contexts with minimal labeled data, self-supervised prototypes, acquired via online deep-clustering [8], are distilled to exploit predictions below an adaptive threshold [45]. Diverging from [12, 45], our method uniquely distills soft pseudo-labels that fall below a threshold, this process being steered by a distribution alignment loss.

To mitigate the bias induced by imbalanced pseudo-labeled distributions, some techniques aim to align the model distribution with the labeled prior, often uniform, based on maximum mean-entropy regularization [2, 31, 66]. Additionally, other distribution alignment approaches involve correcting the prior directly on pseudo-labels [3, 5, 64] and using margin-based losses to debias the classifier [64].

**Long-tailed recognition.** In real-world datasets, LT distributions are common, where a few classes dominate with numerous examples, and most have significantly fewer [61, 82]. Addressing this imbalance, methods based on data resampling [11, 27, 33, 63], loss reweighting [14, 25, 42, 59], and margin modifications [7, 50, 72] have been developed. Theoretically-grounded logit adjustments (LA) mitigate the LT bias [48], yielding balanced [48, 55] and well-calibrated classifiers [1, 75]. These adjustments can be applied during the optimization [48, 55] or as post-hoc bias correction [24, 48]. Additionally, expert-based models are specialized in handling either single or multiple target distributions [1, 40, 80] and have the capability to adjust the test target prior by leveraging unlabeled data transductively [80]. Lastly, KD [23] can also be applied for transferring knowledge from head to tail classes [22, 40, 70].

**Long-tailed semi-supervised learning.** SSL research has recently focused on long-tailed scenarios by relaxing the uniform assumption [29, 68]. DARP [29] refines pseudo-labels via convex optimization, while CReST [68] reduces class imbalance by expanding the labeled dataset with unlabeled data across multiple generations. Class-dependent approaches weight losses based on class difficulty [35] and varying pseudo-label thresholds based on relative class-frequencies [21]. Several approaches introduce auxiliary balanced classifiers [16, 39, 51, 69] or feature regularization [16] to address LT issues.

CReST+ [68] introduces a schedule for progressively aligning the distribution of pseudo-labels, transitioning from long-tailed to more balanced distributions, a strategy later incorporated as part of the training loss [37]. However, most DA approaches assume that labeled and unlabeled data share similar marginal distributions, which may not always hold, and often requires additional supervised pre-training to estimate distribution mismatch [29, 37]. Diverging from this, some studies have adopted re-weighting strategies across fixed anchor distributions [46, 69]. However, these methods might inadvertently rely on privileged information, as they constrain the adjustment to the family of label distributions typically used in evaluation benchmarks. In contrast, our approach utilizes distribution alignment losses, dynamically adjusting to the class distribution of unlabeled data. This reduces label bias and aims for an unbiased, balanced classifier during inference, without being confined to a specific family of prior distributions. Additionally, our method leverages all data samples for regularization, enhancing both prior estimation and distribution alignment, further detailed in Sections 4.1 and 4.2.

**Confidence calibration.** Over-parameterized networks tend to yield uncalibrated and overly confident predictions, particularly in the presence of out-of-distribution data [20, 41], a problem exacerbated by class imbalances [81]. While Mixup [79] and its variants [62, 76] improve calibration in fully-supervised settings and adaptations for long-tailed contexts exist [52, 74], integrating these approaches with threshold-based methods, such as FixMatch [58], remains challenging. Nevertheless, the crucial link between calibration and model performance in balanced SSL setups [44] underscores the importance of calibration. Finally, fully supervised LT distribution alignment methods critically hinge on well-calibrated probabilities [1, 24, 48], underscoring the urgency of addressing calibration within LTSSL frameworks.

### 3 Preliminaries

**Problem formulation.** In long-tailed semi-supervised learning for classification tasks, we address a scenario with a limited labeled dataset  $D_L = (X_L, Y_L)$  and a larger unlabeled dataset  $D_U = (X_U, \cdot)$ . The labeled dataset comprises  $N$  samples  $(x_i, y_i)$ , where  $y_i \in \{1, \dots, K\}$  denotes the class labels and  $K$  is the total number of classes. The unlabeled dataset contains  $M$   $(u_i, \cdot)$  samples, following an unknown class distribution  $Q(y)$ . Classes are sorted by descending order of labeled sample size, i.e.,  $N_1 \geq N_2 \geq \dots \geq N_K$ , where  $N_k$  represents the number of labeled samples for class  $k$ . The imbalance ratio for the labeled set,  $\gamma_l$ , is defined as  $N_1/N_K$ . Similarly, the (unknown) imbalance ratio for the unlabeled set,  $\gamma_u$ , is  $M_1/M_K$ , where  $M_k$  is the count of unlabeled samples in class  $k$ . The objective is to train a classifier that minimizes the *balanced error rate* (BER) [6, 9] on the test distribution, thus ensuring fair treatment of minority classes.

For labeled data, the standard supervised loss [58, 71], denoted by  $\mathcal{L}_s$ , is the average cross-entropy,  $\mathcal{H}(\cdot, \cdot)$ , between the true labels  $y$  and the model predictions  $p(y|x)$ ,

$$\mathcal{L}_s = \frac{1}{B} \sum_{b=1}^B \mathcal{H}(y_b, p(y|\omega(x_b))), \quad (1)$$

where  $p(y|x) = \sigma(f(x))$  denotes the prediction produced by a neural network  $f$ , normalized by the *softmax* function,  $\sigma(\cdot)$ ;  $\omega(\cdot)$  denotes a weak data augmentation procedure and  $B$  denotes the batch size.

For unlabeled data, we are interested in threshold-based consistency regularization approaches [58, 66, 71, 78]. Following [58, 71], we employ weak and strong data augmentations, denoted by  $\omega(\cdot)$  and  $\Omega(\cdot)$ , respectively. The unsupervised loss is defined as

$$\mathcal{L}_u = \frac{1}{\mu B} \sum_{b=1}^{\mu B} \mathcal{M}(u_b) \cdot \mathcal{H}(\hat{y}_b, p(y|\Omega(u_b))), \quad (2)$$

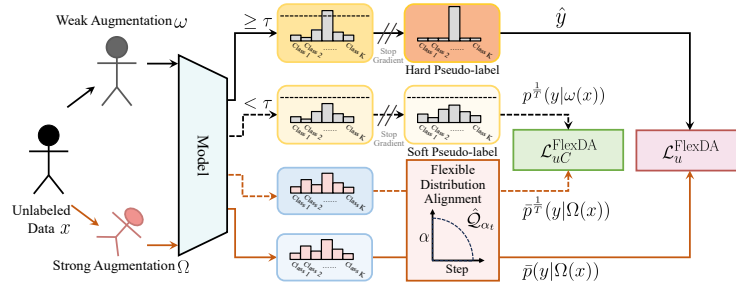
where  $\mathcal{M}(u_b) = \mathbb{1}(\max(p(y|\omega(u_b))) \geq \tau)$  denotes a sample mask relative to a threshold  $\tau$ ,  $\hat{y}_b = \arg \max p(y|\omega(u_b))$  is a one-hot pseudo-label [58].  $\mu$  determines the relative sizes of labeled and unlabeled data in a batch.

**Distribution alignment in SSL.** Table 1 contrasts various methods of distribution alignment tailored for SSL in scenarios with class imbalance. Originally introduced for the balanced setting, ReMixMatch [3] aligns the marginal distribution of predictions on unlabeled data, estimated by an exponential moving average (EMA),  $\hat{Q}(y) \approx \mathbb{E}_{u \sim X_U} [p(y|\omega(u))]$ , with the labeled prior distribution  $\mathcal{P}_L(y)$ , via pseudo-label adjustment (PL). CReST+ [68] adopts a more flexible multi-generational training approach for LT scenarios, controlling the rate and extent of PL debiasing to preserve the model precision and recall for unlabeled data. However, using only pseudo-label correction in long-tailed data scenarios can bias the classifier towards head classes, even with correct pseudo-labels, as noted in fully-supervised settings [28, 48].

In the fully-supervised case, logit-adjusted losses [48, 55] correct the long-tailed label bias by aligning the distribution towards a uniform prior. DebiasPL [64] combines pseudo-label generation with a margin-based unsupervised loss, controlled by a static hyper-parameter. UDAL [37], inspired by CReST+, progressively adjusts its losses to target a smooth long-tailed prior  $\mathcal{P}_{\alpha_t}(y)$ , achieving a more effective alignment. These methods typically presuppose that the unlabeled marginal distribution  $Q(y)$  is similar to the labeled distribution  $\mathcal{P}_L(y)$ . In situations where the target distribution  $Q(y)$  is unknown, previous research [29, 37] has proposed modifying the consistency loss to include a predefined target distribution. However, this approach can be challenging, particularly in the absence of prior knowledge about the true distribution or when a balanced, labeled hold-out dataset for estimating the prior [43] is not available.

**Table 1: Comparative overview of distribution alignment methods under class imbalance.** “PL” denotes pseudo-label adjustment; “S loss” refers to supervised logit-adjusted loss; “U loss” indicates unsupervised logit-adjusted loss; “LT” denotes long-tailed labeled prior; “Progressive” describes gradual smoothing of the target prior throughout training.

Method	Strategy	Target Prior	Progressive
ReMixMatch [3]	PL	LT	✗
CReST+ [68]	PL	LT	✓
LA [48, 55]	S loss	Uniform	✗
DebiasPL [64]	PL & U loss	Uniform	✗
UDAL [37]	S & U losses	LT	✓
ADELLO (ours)	S & U losses	<b>Adaptive</b>	✓



**Fig. 2: Method overview:** Our *flexible distribution alignment* (FlexDA) aligns the classifier with the correct prior, dynamically estimated from unlabeled data. This approach extends FixMatch with a bias-adjusted supervised loss ((4), Sec. 4.1) and a bias-adjusted consistency loss ((5), Sec. 4.1) to debias high-confidence *hard pseudo-labels*. We also introduce a bias-adjusted *complementary consistency* loss to learn from low-confidence *soft pseudo-labels* (Sec. 4.2). A progressive scheduler steadily smooths the target prior,  $\hat{Q}_{\alpha_t}$ , leading to a balanced classifier by the conclusion of training.

Aiming to resolve these challenges, ADELLO aligns the target distribution to the unknown unlabeled data and fosters robust consistency by distilling even low-confident predictions simultaneously to counter the class imbalance.

## 4 ADELLO Framework

As illustrated in Fig. 2, the key idea of ADELLO is to conduct progressive distribution alignment on unlabeled data via consistency regularization using high-confidence and low-confidence pseudo-labels. We propose the following objective,

$$\mathcal{L} = \mathcal{L}_s^{\text{FlexDA}} + \mathcal{L}_u^{\text{FlexDA}} + \mathcal{L}_{uC}^{\text{FlexDA}}, \quad (3)$$

where  $\mathcal{L}_s^{\text{FlexDA}}$  denotes the supervised loss (4),  $\mathcal{L}_u^{\text{FlexDA}}$  denotes the unsupervised consistency loss (5), and  $\mathcal{L}_{uC}^{\text{FlexDA}}$  denotes the complementary consistency loss (Sec. 4.2) within FlexDA. This framework facilitates model alignment with an adaptive target distribution and effective use of labeled and unlabeled data, ensuring robust model training through comprehensive data utilization. We will now provide a detailed description of the losses introduced in ADELLO.

### 4.1 Flexible Distribution Alignment

In the training of modern SSL frameworks, we typically observe two learning phases: 1) a supervised phase where the model is trained on a small labeled dataset using weak data augmentation, and 2) a phase where pseudo-labeling and consistency regularization are employed to learn from strongly-augmented unlabeled data.

Let us assume that the first phase yields a scorer function  $g_L(x)$  that perfectly fits the labeled distribution, represented as  $g_L(x) \propto \mathcal{P}_L(y|x)$ . In the second phase, we want to find the classifier that maximizes the number of correct pseudo-labels. The Bayes-optimal

classifier emerges as a solution:  $\hat{y} = \arg \max_y \mathcal{Q}(y|x) = \arg \max_y \mathcal{Q}(x|y) \cdot \mathcal{Q}(y)$ . Under *label shift* [24], where priors might differ, i.e.  $\mathcal{P}_L(y) \neq \mathcal{Q}(y)$ , yet the likelihood remains the same, i.e.  $\mathcal{P}_L(x|y) = \mathcal{Q}(x|y)$ , we can define an adjusted scorer for pseudo-labeling:  $g_U(x) = g_L(x) \cdot \frac{\mathcal{Q}(y)}{\mathcal{P}_L(y)}$ , and show that  $g_U(x) \propto \mathcal{Q}(y|x)$ . Furthermore, this indicates that  $g_U(x)$  is the best scorer for the unlabeled data in terms of Bayes optimality.

The definition of  $g_U(x)$  suggests that, in practice, obtaining a good classifier involves two challenges: neutralizing bias from the skewed labeled data and adjusting for the marginal distribution  $\mathcal{Q}(y)$ . However, for inference, our goal is to achieve a *balanced classifier* that treats all classes equally. This classifier should minimize the balanced error, which is independent of training or test priors:  $\hat{y} = \arg \max_y \mathcal{P}_{\text{bal}}(y|x) = \arg \max_y \mathcal{P}_{\text{bal}}(x|y) \cdot \mathcal{P}_{\text{bal}}(y)$ . Here,  $\mathcal{P}_{\text{bal}}(y) = \frac{1}{K}$  denotes the uniform prior. A balanced scorer can thus be defined as  $g_B(x) = g_U(x) \cdot \frac{\mathcal{P}_{\text{bal}}(y)}{\mathcal{Q}(y)} \propto \mathcal{P}_{\text{bal}}(y|x)$ , aligning with the Bayes-optimal rule for minimizing the BER [48].

Reconciling training demands with inference realities presents a challenge, as  $\mathcal{Q}(y)$  is often unknown. Even under consistent scenarios,  $\mathcal{P}_L(y)$  may be a biased estimate of the correct distribution, particularly with limited labeled samples. To address these challenges, we introduce the **Flexible Distribution Alignment** (FlexDA) approach. FlexDA dynamically adapts to the characteristics of unlabeled data by aligning the model with a target distribution based on the EMA of the pseudo-labels during optimization. It employs a smoothed target prior,  $\hat{\mathcal{Q}}_{\alpha_t}(y) = \frac{\hat{\mathcal{Q}}(y)^{\alpha_t}}{\sum_j \hat{\mathcal{Q}}(j)^{\alpha_t}}$ , with a time-updated decay factor  $\alpha_t$ . Our approach progressively smooths the target prior, starting from the unlabeled prior  $\hat{\mathcal{Q}}_1(y) = \hat{\mathcal{Q}}(y)$ , and gradually transitioning to a (near) balanced prior  $\hat{\mathcal{Q}}_0(y) = \frac{1}{K}$  by the end of the training process.

In the FlexDA approach, our proposed logit-adjusted supervised loss is defined as


$$\mathcal{L}_s^{\text{FlexDA}} = \frac{1}{B} \sum_{b=1}^B \mathcal{H}(y_b, \sigma(f(\omega(x_b)) + \log \frac{\mathcal{P}_L}{\hat{\mathcal{Q}}_{\alpha_t}})), \quad (4)$$

while our unsupervised consistency loss is defined as

$$\mathcal{L}_u^{\text{FlexDA}} = \frac{1}{\mu B} \sum_{b=1}^{\mu B} \mathcal{M}(u_b) \cdot \mathcal{H}(\hat{y}_b, \sigma(f(\Omega(u_b)) + \log \frac{\hat{\mathcal{Q}}}{\hat{\mathcal{Q}}_{\alpha_t}})), \quad (5)$$

where  $\alpha_t = 1.0 - (1.0 - \alpha_{\min}) \left(\frac{t}{t_{\text{total}}}\right)^d$  defines a schedule. Here,  $t$  is the current training step,  $t_{\text{total}}$  the total number of training steps, and  $d$  and  $\alpha_{\min}$  are hyper-parameters controlling the speed of the debiasing schedule and the minimum bias allowed, respectively.

**Statistical perspective.** For simplicity, let us assume  $\alpha_{\min} = 0$ , i.e. a uniform target prior by the end of training. When  $\alpha_t \rightarrow 1$ , our supervised loss compensates for the distribution shift between labeled and unlabeled data. While the unadjusted scorer might effectively classify labeled data, as in  $\sigma(f(x) + \log \frac{\mathcal{P}_L}{\hat{\mathcal{Q}}_{\alpha_t}}) \stackrel{\alpha_t \rightarrow 1}{\approx} g_L(x)$ , the model (adjusted scorer) aims to emulate the Bayes-optimal classifier for unlabeled data, i.e.,  $\sigma(f(x)) \approx g_L(x) \cdot \frac{\hat{\mathcal{Q}}_{\alpha_t}}{\mathcal{P}_L} \stackrel{\alpha_t \rightarrow 1}{\approx} g_U(x)$ . As Fig. 3 shows, the estimated model prior  $\hat{\mathcal{Q}}$  aligns closely with the ground-truth unlabeled prior  $\mathcal{Q}$  from the beginning of training under various levels of label shift, as shown by the KL divergence between these distributions.



(a) KL-div for *forward* case (b) KL-div for *balanced* case (c) KL-div for *reversed* case (d)  $\alpha_t$  ( $d = 2$ ,  $\alpha_{\min} = 0.1$ )

**Fig. 3: Prior estimation under label shift.** A comparison of KL divergence shows 1) a small difference between the estimated prior,  $\hat{Q}$ , and the ground-truth prior,  $Q$ , during most of the training (blue curve), and 2) a larger disparity between  $\hat{Q}$  and the uniform prior,  $\mathcal{P}_{\text{bal}}$ , (orange curve). The progression of a quadratic scheduler ( $d = 2$ ) is shown in (d) (green curve). Label shift settings: (a) forward, (b) balanced, and (c) reversed long-tailed, computed for CIFAR10-LT100.

Throughout the training, both supervised and unsupervised losses progressively reduce the bias induced by the unlabeled data, modulated by  $\alpha_t$  (refer to Fig. 3d). By the end of the training, as  $\alpha_t \rightarrow 0$  and  $\hat{Q}_{\alpha_t} \rightarrow \mathcal{P}_{\text{bal}}(y)$ , FlexDA steers the model towards a (more) balanced distribution, leading to  $\sigma(f(x)) \approx g_U(x) \cdot \frac{\hat{Q}_{\alpha_t}}{\hat{Q}(y)} \stackrel{\alpha_t \rightarrow 0}{\approx} g_B(x)$ . This dual loss structure effectively counters the bias introduced by labeled and unlabeled data samplers and adapts to the dynamic target prior. It provides a holistic approach for training semi-supervised models in the presence of class imbalance.

## 4.2 Complementary Consistency Regularization

The lack of ground-truth labels for tail classes complicates the generation of accurate pseudo-labels early in training, leading to a reduced supervisory signal when high-confidence thresholds are applied [45]. Furthermore, (progressive) distribution alignment can decrease the confidence of pseudo-labels as training progresses [68], further weakening the supervisory signal. To utilize all available unlabeled data, we enhance the conventional consistency objective (2) with a complementary consistency regularization (CCR) technique. We implement a masked-distillation loss that makes use of soft pseudo-labels that fall below the confidence threshold:

$$\mathcal{L}_{uC} = \frac{1}{\mu B} \sum_{b=1}^{\mu B} \mathcal{M}^C(u_b) \cdot \mathcal{H}(p^{\frac{1}{T}}(y|\omega(u_b)), p^{\frac{1}{T}}(y|\Omega(u_b))), \quad (6)$$

where  $\mathcal{M}^C(u_b) = 1 - \mathcal{M}(u_b)$  denotes the complementary mask,  $T$  is the temperature scaling factor, and  $p^{\frac{1}{T}}(y|x) = \sigma(\frac{1}{T}f(x))$  represents the temperature-scaled predictions.

When combining our distribution alignment with complementary consistency, we obtain  $\mathcal{L}_{uC}^{\text{FlexDA}}$  which is defined the same as  $\mathcal{L}_{uC}$  but replacing  $p^{\frac{1}{T}}(y|\Omega(u_b))$  with  $\bar{p}^{\frac{1}{T}}(y|\Omega(u_b)) = \sigma(\frac{1}{T}(f(\Omega(u_b)) + \log \frac{\hat{Q}}{\hat{Q}_{\alpha_t}}))$ , where  $\bar{p}^{\frac{1}{T}}$  denotes temperature-scale unadjusted predictions. For simplicity, we denote  $\bar{p}(y|x) = \bar{p}^{\frac{1}{T}}(y|x)$  when  $T = 1$ .

Contrasting with recent methods that sharpen (hence low temperature) low-confidence, less reliable pseudo-labels [3, 12, 78], increasing the risk of confirmation bias, we apply bias-corrected distillation-based consistency (hence high temperature) to uncertain samples, boosting model accuracy and reliability for LTSSL, as observed in Fig. 1.



**Imbalance-aware temperature selection.** Determining the optimal temperature for distillation can be critical, particularly when confronted with LT settings. While a temperature of  $T=1$  is shown to be effective for balanced (unlabeled) datasets, see Fig. 6, this may not hold for imbalanced ones, as observed in fully-supervised settings [22]. In response to this, we calculate the temperature to accommodate the class imbalance:

$$T = \exp(\text{KL}(\mathcal{P}_{\text{bal}} \parallel \hat{Q})), \quad (7)$$

by initiating distillation after a warm-up period. This temperature, once set, is kept constant throughout the remainder of the training. This strategy ensures that our distillation process is fine-tuned to the unlabeled data from the start and remains stable against minor changes in the pseudo-label distribution. We validate its effectiveness in Sec. 5.4.

## 5 Experiments

### 5.1 Setup

**Datasets.** We evaluate the performance of our approach on multiple benchmarks, assessing its robustness under varying degrees of class imbalance, labeled data availability, and class distribution mismatch.

**CIFAR10-LT** and **CIFAR100-LT** are based on CIFAR10 and CIFAR100 [32], each originally containing 60k  $32 \times 32$  color images across 10 and 100 classes, split into 50k for training and 10k for testing. Following the standard protocol [29], we sample these datasets to create their long-tailed versions, using a head class size  $N_1$  and an imbalance ratio  $\gamma_l$ .  $\gamma_u$  denotes the imbalanced ratio for unlabeled data. The number of images per class for labeled and unlabeled data is determined by  $N_k = N_1 \cdot \gamma_l^{-\kappa}$  and  $M_k = M_1 \cdot \gamma_u^{-\kappa}$ , respectively, where  $\kappa = (k - 1)/(K - 1)$ . We use two labeled data settings with 1/3 and 1/9 of the total data [51].

**STL10-LT** is derived by downsampling the labeled portion of the STL10 dataset [13], akin to the procedure used for the CIFAR long-tailed variants. STL10 consists of 5k training and 8k test images, each  $96 \times 96$  in resolution, spread across 10 classes. This is augmented by an extra 100k unlabeled images that include both in-distribution and related out-of-distribution (OOD) classes from the ImageNet [15] taxonomy.

**ImageNet127** [26], a naturally imbalanced large-scale dataset with  $\gamma_l \approx \gamma_u \approx 286$ , groups the 1k classes of ImageNet [15] into 127 classes based on the WordNet hierarchy. We evaluate under  $32 \times 32$  and  $64 \times 64$  image resolution using 10% of labeled data [16].

**Training and evaluation.** Our framework is based on FixMatch [58] with a confidence threshold of 0.95 [58]. Our experiments use Wide-ResNet-28-2 [77] for CIFAR10-LT, CIFAR100-LT, and STL10-LT, while ResNet-50 is used for ImageNet-127, following [16, 29, 51]. For CIFAR{10,100}-LT and STL10-LT, we train for 256 epochs of 1024 steps each, using SGD, Nesterov momentum of 0.9, and weight decay of  $5e-4$  [51]. The base learning rate (LR) is set to 0.03. ImageNet-127 experiments use Adam [30] with a base LR of 0.002 for 500 epochs of 500 steps each, following [16]. Batch sizes are 64 for labeled and 128 for unlabeled data. We set  $\alpha_{\min}$  to 0.1 and  $d$  to 2, following [37]. A warm-up period of 50k steps is used with CIFAR{10,100}-LT, while STL10-LT and ImageNet127 skip it. Following common practices [29, 37, 58, 68], we define equally

**Table 2:** Test accuracy (%) on CIFAR10-LT and CIFAR100-LT **under label shift**. †: labeled prior as target. ‡: results from prior work [51]. Best scores **bold**, second-best underlined.

	CIFAR10-LT			CIFAR100-LT			Friedman Final	
	$\gamma_l \rightarrow$	$\gamma_u \rightarrow$	$N_1 \rightarrow$	$M_1 \rightarrow$	Rank	Rank		
Supervised	63.8±0.3	63.8±0.3	63.8±0.3	36.3±0.3	36.3±0.3	36.3±0.3	-	-
FixMatch [58]	75.5±1.1	86.1±1.1	81.0±4.2	44.4±0.6	48.6±1.0	45.4±1.6	8.67	9
+DARP† [29]	76.6±1.0	68.8±0.7	63.3±1.3	44.7±0.4	43.3±0.6	40.4±0.7	9.50	10
+CReST+ [68]	78.1±0.6	<b>92.6±0.2</b>	68.5±0.5	44.9±0.2	<u>56.5±0.5</u>	40.5±0.4	6.17	7
+ABC [39]	82.3±0.7	89.0±0.2	<b>87.0±0.4</b>	47.2±0.6	52.4±1.5	48.7±2.0	4.00	3
+DASO [51]	79.1±0.7†	88.8±0.6‡	80.3±0.6†	44.7±0.2	51.7±2.0	48.5±2.1	7.00	8
+DebiasPL [64]	80.5±0.1	88.6±0.2	83.8±0.2	46.8±0.3	52.5±0.8	50.8±1.9	5.00	6
+CoSSL [16]	<b>84.6±0.1</b>	88.8±0.6	84.2±0.2	47.6±0.8	50.4±1.2	46.8±0.6	4.67	5
+UDAL† [37]	83.0±0.3	89.1±0.2	80.9±0.7	<u>48.6±0.5</u>	52.6±1.0	48.7±1.3	4.00	3
+ADELLO (ours)	<u>83.8±0.3</u>	<u>91.9±0.3</u>	<u>86.1±0.4</u>	<b>49.2±0.6</b>	<b>57.5±1.3</b>	<b>53.0±0.9</b>	<b>1.50</b>	<b>1</b>
SoftMatch [12]	79.6±0.2	89.6±0.4	83.0±0.8	46.4±0.9	<b>57.5±0.8</b>	<u>51.2±1.2</u>	<u>3.83</u>	<u>2</u>

weighted losses. An ablation study in Appendix A supports this choice. Appendix E includes pseudo-code for ADELLO and Appendix F details our hyperparameter settings.

To assess our method, an EMA network updates parameters at each step with a decay of 0.999 [3, 51]. We report the average of the *test balanced accuracy* over the final 20 epochs [16]. We provide the mean and standard deviation from three independent runs. Friedman ranking [17, 18] is used to fairly assess algorithms across different settings, subsequently determining the final ranking from the Friedman scores, following the methodology in [65]. All experiments were conducted on a single Nvidia V100-32GB GPU within a local cluster. A discussion on running times is deferred to Appendix B.

## 5.2 Main Results

In this section, we present extensive experiments to evaluate the performance of our approach against several SOTA approaches. These methods include a supervised baseline (using only labeled data), FixMatch [58] (SSL baseline), SoftMatch [12] (stronger SSL baseline), as well as representative LTSSL algorithms, including DARP [29], CReST+ [68], ABC [39], DASO [51], DebiasPL [64], CoSSL [16], and UDAL [37]. To demonstrate the effectiveness of our method across diverse scenarios, we assess its performance under varying levels of label shift in Table 2, different degrees of class imbalance under low-label regimes in Table 3, and an exceptionally challenging scenario on ImageNet127 in Table 4. Furthermore, we investigate model calibration in Section 5.3.

**Results under (unknown) label shift.** To address scenarios where the unlabeled class distribution differs from or is unknown relative to the labeled prior, we vary the unlabeled class distribution for CIFAR{10,100}-LT to obtain three evaluation settings: forward long-tailed ( $\gamma_u = \gamma_l$ ), balanced ( $\gamma_u = \frac{1}{K}$ ), and reversed long-tailed ( $\gamma_u = \frac{1}{\gamma_l}$ ).

Table 2 demonstrates the robustness of ADELLO in handling unknown distribution mismatches, particularly evident on CIFAR100-LT (see also Fig. 1d), which contains a larger number of classes. The rightmost columns of the table show the Friedman

**Table 3:** Test accuracy (%) on CIFAR{10,100}-LT and STL10-LT under low-label regimes. †: labeled prior as target. ‡: results from prior work [51]. Best scores **bold**, second-best underlined.

	CIFAR10-LT		CIFAR100-LT		STL10-LT		Friedman Rank	Final Rank
	$\gamma_l \rightarrow$	$\gamma_u \rightarrow$	$N_l \rightarrow$	$M_l \rightarrow$	$\gamma_l \rightarrow$	$\gamma_u \rightarrow$		
Supervised	46.6±0.9	43.4±1.9	27.7±1.8	25.1±1.1	46.4±0.6	40.8±0.6	-	-
FixMatch [58]	69.8±1.6	65.6±1.5	47.0±0.9	42.2±0.6	64.1±2.3	54.5±4.3	9.67	10
+DARP† [29]	72.9±1.3	67.2±2.0	47.7±0.7	42.8±1.2	62.1±1.4	54.7±2.6	9.00	9
+CRcST+ [68]	77.6±0.2	72.1±2.9	46.6±0.6	43.2±1.0	66.9±1.0	62.6±2.6	7.33	8
+ABC [39]	78.9±0.9	72.0±2.4	49.7±1.3	44.1±0.3	71.2±1.0	65.7±2.3	4.83	6
+DASO‡ [51]	80.1±1.2	70.6±0.8	49.8±0.2	43.6±0.1	70.0±1.2	65.7±1.8	5.83	7
+DebiasPL [64]	76.4±4.3	72.0±1.8	50.3±1.1	<u>45.4±0.5</u>	70.1±0.8	66.6±2.1	4.50	4
+CoSSL [16]	<u>80.8±0.5</u>	<b>76.8±0.7</b>	48.8±1.0	44.4±0.7	70.6±0.5	66.0±1.4	3.67	3
+UDAL† [37]	80.8±0.5	76.4±2.6	50.4±1.1	<b>46.5±0.1</b>	69.8±1.1	65.0±2.3	<u>3.50</u>	2
+ADELLO (ours)	<b>81.3±0.4</b>	76.0±1.7	<b>51.8±0.7</b>	<b>46.5±0.2</b>	<b>75.7±0.7</b>	<b>74.6±0.4</b>	<b>1.33</b>	<b>1</b>
SoftMatch [12]	77.1±0.8	71.0±1.4	50.2±0.7	43.8±0.5	<u>72.6±0.3</u>	<u>70.6±0.4</u>	4.67	5

**Table 4: Large-scale datasets.** Test balanced accuracy (%) on ImageNet127 at  $32 \times 32$  and  $64 \times 64$  image resolution. †: results from prior work [16]. Best scores **bold**, second-best underlined.


Method	$32 \times 32$	$64 \times 64$
FixMatch [58]†	29.7	42.3
+DARP [29]†	30.5	42.5
+DARP+cRT [29]†	39.7	51.0
+CRcST+ [68]†	32.5	44.7
+CRcST+ +LA [68]†	40.9	<u>55.9</u>
+CoSSL [16]†	43.7	53.8
+UDAL ( $\alpha_{\min}=0.55$ ) [37]	40.2	49.4
+UDAL ( $\alpha_{\min}=0.1$ ) [37]	<u>44.1</u>	52.3
+ADELLO (ours)	<b>47.5</b>	<b>58.0</b>

scoring and the final rank over test accuracies for each method. ADELLO secures the top position, showcasing its superior performance. Consistently ranking first or second in all settings, it demonstrates remarkable adaptability to degrees of label shift. The effectiveness of ADELLO becomes evident in both forward and reversed LT scenarios on CIFAR100, outperforming SoftMatch significantly. Distinctly outperforming previous SOTA approaches like ABC, DASO, and CoSSL, ADELLO delivers robust LTSSL performance without depending on auxiliary classifiers or data re-sampling.


We also compare the classification performance of ADELLO with ACR [69], a recent LTSSL approach. Fig. 4 shows that ADELLO outperforms ACR under label shift, with the performance gap widening as the distribution mismatch increases.

**Results with limited labeled data.** Table 3 highlights the effectiveness of ADELLO on CIFAR{10,100}-LT and STL10-LT, particularly in scenarios with limited labeled data and significant class imbalance. In STL10-LT, where only 150 labels are available for the head class amid a range of OOD data, ADELLO shows marked improvements. It notably surpasses CoSSL with a +8.0 gain in accuracy and ABC with a +4.5 increase at imbalance ratios of 20 and 10, respectively, while outperforming SoftMatch, a strong SSL baseline. Oddly, SoftMatch mistakenly classifies OOD data as known classes using hard PLs with targeted weights. Conversely, our method uses CCR to predict soft PLs on potential OOD samples, enhancing robustness.

Under consistent CIFAR{10,100}-LT settings, the performance of ADELLO matches or exceeds that observed in established methods like CoSSL and UDAL, reinforcing the effectiveness of correct prior estimation even with a low amount of labels and without reliance on strong assumptions. Notably, ADELLO outperforms SoftMatch by a large margin as imbalance ratios increase without using any adaptive thresholding technique.



**Fig. 4:** Varying label shift.



**Fig. 5: Trade-off between Generalization and Calibration Performance.** We report misclassification error (%) vs. expected calibration error (%), computed on the test split. The first row contrasts different datasets, and the second row examines various degrees of label shift.

**Results on ImageNet127.** The performance of various methods on the ImageNet127 dataset, a challenging variant of standard ImageNet featuring 127 classes and an imbalance ratio of 286, is summarized in Table 4. Due to its extensive sample size (1.28M), ImageNet127 serves as a unique testbed for assessing large-scale imbalanced datasets. ADELLO notably excels in balanced accuracy, surpassing the previous state-of-the-art, CoSSL, with gains of +3.8 at  $32 \times 32$  resolution and +4.2 at  $64 \times 64$  resolution. Compared to UDAL, accuracy increases of +3.4 and +5.7 are observed, respectively, at these resolutions. The advantage of ADELLO over UDAL and CReST+, both using consistent priors for distribution alignment, highlights the benefits of marrying FlexDA with complementary consistency for superior performance in large-scale settings.

### 5.3 Study of Model Calibration

Model calibration [20, 44], vital for accurately reflecting predictive uncertainty in SSL, is a focal point of our study [44]. We examine the impact of calibration on LTSSL by comparing expected calibration error (ECE) with misclassification error. Our findings align with [44], showing that better-calibrated models improve SSL performance, even with class imbalances, as Fig. 5 illustrates. ADELLO demonstrates a superior trade-off in reducing misclassification and calibration errors compared to other LTSSL methods, across CIFAR{10,100}-LT and STL10-LT datasets. Fig. 5 (first row) reveals that while most LTSSL methods surpass FixMatch, ADELLO further improves generalization and, particularly, calibration in scenarios where distributions align. The superiority of ADELLO is significantly highlighted in the STL10-LT benchmark, characterized by an unknown label shift and a substantial presence of unlabeled OOD samples. Section 5.4 attributes these improvements to our flexible distribution alignment and the significant role of complementary consistency regularization in such contexts.

In scenarios with controlled label shift, as illustrated in Fig. 5 (second row), certain LTSSL methods, such as DARP, CoSSL, and CReST+ to a degree, face difficulties with large label shifts where the bias in labeled data fails to accurately represent the charac-

**Table 5:** Influence of ADELLO components on model generalization (Test accuracy).

Components	CIFAR100-LT50 $\uparrow$			STL10-LT20 $\uparrow$
	$\gamma_u \rightarrow$	50	1	0.02
FixMatch		44.4 $\pm$ 0.6	48.6 $\pm$ 1.0	45.4 $\pm$ 1.6
+FlexDA		48.6 $\pm$ 0.7	53.6 $\pm$ 1.0	51.2 $\pm$ 0.9
+CCR		44.7 $\pm$ 0.7	51.6 $\pm$ 1.6	47.2 $\pm$ 2.1
<b>+FlexDA+CCR</b>		<b>49.2<math>\pm</math>0.6</b>	<b>57.5<math>\pm</math>1.3</b>	<b>53.0<math>\pm</math>0.9</b>
+FlexDA+KD		49.1 $\pm$ 0.6	<b>58.2<math>\pm</math>1.1</b>	52.8 $\pm$ 1.1

**Table 7:** Ablation of scheduler speed ( $d$ ).

$d$	$\gamma_u = 50$
0	46.6 $\pm$ 0.6
1	49.1 $\pm$ 0.4
2	49.2 $\pm$ 0.5
3	49.1 $\pm$ 0.8

**Table 8:** Ablation of minimum bias ( $\alpha_{\min}$ ).

$\alpha_{\min}$	$\gamma_u = 50$
0.0	49.1 $\pm$ 0.7
0.1	49.2 $\pm$ 0.5
0.2	49.1 $\pm$ 1.2
0.3	48.7 $\pm$ 0.7

**Table 6:** Influence of ADELLO components on model calibration (ECE/MCE).

Components	CIFAR100-LT50 $\downarrow$		STL10-LT20 $\downarrow$	
	$\gamma_u \rightarrow$	50	1	N/A
FixMatch		37.4 $\pm$ 0.4 / 57.3 $\pm$ 1.1	37.8 $\pm$ 4.5 / 55.1 $\pm$ 4.9	
+FlexDA		31.4 $\pm$ 0.4 / 52.0 $\pm$ 2.4	23.6 $\pm$ 1.4 / 49.5 $\pm$ 4.5	
+CCR		36.3 $\pm$ 0.7 / 56.8 $\pm$ 1.8	22.2 $\pm$ 2.3 / 38.5 $\pm$ 4.2	
<b>+FlexDA+CCR</b>		<b>26.1<math>\pm</math>0.9 / 46.2<math>\pm</math>0.6</b>	<b>6.9<math>\pm</math>0.3 / 25.9<math>\pm</math>1.0</b>	
+FlexDA+KD		33.4 $\pm$ 0.6 / 57.5 $\pm$ 2.0	10.0 $\pm$ 0.5 / 31.3 $\pm$ 7.5	

**Table 9:** Ablation of warm-up period.

#steps	$\gamma_u \rightarrow$	50	1	0.02
no warm-up		45.6 $\pm$ 0.6	58.4 $\pm$ 1.3	51.3 $\pm$ 1.9
25k steps		49.2 $\pm$ 0.7	57.7 $\pm$ 1.3	52.6 $\pm$ 1.0
50k steps		49.2 $\pm$ 0.6	57.5 $\pm$ 1.3	53.0 $\pm$ 0.9
100k steps		49.1 $\pm$ 0.4	57.8 $\pm$ 1.2	52.8 $\pm$ 1.0
no distillation		48.6 $\pm$ 0.7	53.6 $\pm$ 1.0	51.2 $\pm$ 0.9

teristics of the unlabeled distribution, a problem that ADELLO overcomes. Although the auxiliary balanced classifier in ABC demonstrates acceptable calibration, ADELLO showcases greater flexibility and robustness in generalization performance compared to ABC. Appendix C presents similar trends for the maximum calibration error (MCE).

## 5.4 Ablation Study


**Importance of the proposed losses.** Our ablation studies on CIFAR100-LT50, shown in Table 5, evaluate ADELLO objectives across imbalance ratios ( $\gamma_u$ ) of 50, 1, and 0.02. The FlexDA component in ADELLO significantly outperforms the baseline, FixMatch, with gains of +4.2, +5.0, and +5.8 points for these  $\gamma_u$  values, underscoring its effectiveness against class imbalance and distribution mismatch. Further, the studies indicate that complementary consistency enhances performance, highlighting its value in SSL. However, the synergy of FlexDA and CCR within ADELLO results in the most substantial improvements, with increases of +4.8, +8.9, and +7.6 points across the 50, 1, and 0.02 imbalance ratios, respectively. While FlexDA sees advantages from indiscriminate KD of all samples (FlexDA+KD), using masked distillation (FlexDA+CCR) more often results in enhanced generalization in imbalanced scenarios (see  $\gamma_u \in \{50, 0.02\}$ ).

**Are all components in ADELLO necessary for proper calibration?** Table 6 shows that both FlexDA and CCR boost calibration independently. We observe that their synergy markedly surpasses the baseline, by correcting the label bias on the whole data distribution, akin to fully-supervised approaches [1, 74]. Significantly, the key to enhanced calibration in LTSSL contexts lies not just in the naive distillation of all samples (FlexDA+KD), but in the strategic combination of soft pseudo-labels for underconfident samples and hard pseudo-labels for those with high confidence, as depicted in Fig. 2.

**Do we need a progressive scheduler?** The setup for FlexDA, as outlined in Section 5.2, adheres to configurations proposed by Lazarow et al. [37]. Within the ADELLO framework, this analysis investigates the effect of the speed of the scheduler ( $d$ ) and the minimum bias hyperparameters ( $\alpha_{\min}$ ) on model performance. Table 7 suggests that a moderate, yet progressive, scheduler, i.e.  $d \in (1, 3)$ , leads to optimal accuracy, while an aggressive debiasing rate ( $d = 0$ ) proves detrimental. Performance peaks when the minimum bias ( $\alpha_{\min}$ ) is near zero as observed in Table 8, suggesting that this configuration minimizes the balanced error, aligning with findings of fully-supervised approaches [48]. Generally, there is minimal sensitivity to the precise settings of these hyperparameters. We hypothesize that the union of FlexDA with complementary consistency is key to the effectiveness of ADELLO compared to other DA approaches, namely CReST+ and DebiasPL, which do not engage in full classifier debiasing to retain high data utilization.

**How robust is the inferred temperature  $T$ ?** We calibrate  $T$  based on the class imbalance of unlabeled data, using (7). Fig. 6 shows this strategy is nearly as effective as custom-tuning the temperature for each dataset. A  $T$  near one is preferred for balanced data, facilitating distillation from an increasingly balanced classifier and leading to marked performance gains. Under more imbalanced cases, our method takes a more cautious approach by opting for a higher temperature.

**Effect of warm-up period.** Table 9 shows that beginning complementary consistency after a warm-up stage using our temperature-selection procedure boosts performance compared to using an uninformative setting or not distilling at all, while showing robustness to the exact starting point for CCR.



**Fig. 6:** Inferred vs. tuned temperature.

## 6 Conclusion

We proposed a two-faceted framework for greatly improving the performance of LTSSL under label shift. First, our flexible distribution alignment (FlexDA) reduces the bias caused by differing labeled and unlabeled class-distribution marginals, and subsequently, the head-class bias intrinsic to imbalanced data. These reductions are achieved by aligning the model prior first to a dynamic estimate of the unlabeled marginal and gradually towards a more balanced distribution. Second, our complementary consistency regularization leverages the soft output signals of below-threshold pseudo-labels toward improving data utilization of minority classes. We demonstrate that this framework is state-of-the-art when unlabeled and labeled marginal distributions are mismatched, competitive when they are matched, and achieves better calibration than its competitors.

**Limitations.** All LTSSL benchmarks we are aware of focus on the *closed-world assumption* [54], where every class is labeled and known during inference. For the STL10-LT, which includes near-out-of-distribution unlabeled samples, ADELLO yields promising results, suggesting potential in handling "unknown" classes not present in the labeled set. Furthermore, while the proposed framework is developed for the classification task, it may also be beneficial to address the class imbalance in more complex visual tasks, such as object detection, instance segmentation, or tracking.

## Acknowledgements

This work was supported by the Wallenberg Artificial Intelligence, Autonomous Systems and Software Program (WASP), funded by the Knut and Alice Wallenberg Foundation. The computational resources were provided by the National Academic Infrastructure for Supercomputing in Sweden (NAISS), partially funded by the Swedish Research Council through grant agreement no. 2022-06725, and by the Berzelius resource, provided by the Knut and Alice Wallenberg Foundation at the National Supercomputer Centre.

## References

1. Aimar, E.S., Jonnarth, A., Felsberg, M., Kuhlmann, M.: Balanced product of calibrated experts for long-tailed recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 19967–19977 (2023)
2. Arazo, E., Ortego, D., Albert, P., O’Connor, N.E., McGuinness, K.: Pseudo-labeling and confirmation bias in deep semi-supervised learning. In: International Joint Conference on Neural Networks (IJCNN). IEEE (2020)
3. Berthelot, D., Carlini, N., Cubuk, E.D., Kurakin, A., Sohn, K., Zhang, H., Raffel, C.: Remix-match: Semi-supervised learning with distribution matching and augmentation anchoring. In: 8th International Conference on Learning Representations (2020)
4. Berthelot, D., Carlini, N., Goodfellow, I., Papernot, N., Oliver, A., Raffel, C.A.: Mixmatch: A holistic approach to semi-supervised learning. In: Advances in Neural Information Processing Systems (2019)
5. Berthelot, D., Roelofs, R., Sohn, K., Carlini, N., Kurakin, A.: Adamatch: A unified approach to semi-supervised learning and domain adaptation (2021)
6. Brodersen, K.H., Ong, C.S., Stephan, K.E., Buhmann, J.M.: The balanced accuracy and its posterior distribution. In: 2010 20th international conference on pattern recognition. pp. 3121–3124. IEEE (2010)
7. Cao, K., Wei, C., Gaidon, A., Arechiga, N., Ma, T.: Learning imbalanced datasets with label-distribution-aware margin loss. In: Advances in Neural Information Processing Systems (2019)
8. Caron, M., Misra, I., Mairal, J., Goyal, P., Bojanowski, P., Joulin, A.: Unsupervised learning of visual features by contrasting cluster assignments. *Advances in neural information processing systems* **33**, 9912–9924 (2020)
9. Chan, P.K., Stolfo, S.J.: Learning with non-uniform class and cost distributions: Effects and a distributed multi-classifier approach. In: In Workshop Notes KDD-98 Workshop on Distributed Data Mining (1998)
10. Chapelle, O., Schölkopf, B., Zien, A.: Semi-supervised learning (chapelle, o. et al., eds.; 2006). *IEEE Transactions on Neural Networks* **20**(3) (2009)
11. Chawla, N.V., Bowyer, K.W., Hall, L.O., Kegelmeyer, W.P.: Smote: synthetic minority over-sampling technique. *Journal of artificial intelligence research* (2002)
12. Chen, H., Tao, R., Fan, Y., Wang, Y., Savvides, M., Wang, J., Raj, B., Xie, X., Schiele, B.: Softmatch: Addressing the quantity-quality tradeoff in semi-supervised learning. In: Eleventh International Conference on Learning Representations. OpenReview. net (2023)
13. Coates, A., Ng, A., Lee, H.: An analysis of single-layer networks in unsupervised feature learning. In: Proceedings of the fourteenth international conference on artificial intelligence and statistics (2011)

14. Cui, Y., Jia, M., Lin, T.Y., Song, Y., Belongie, S.: Class-balanced loss based on effective number of samples. In: Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition (2019)
15. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: IEEE conference on Computer Vision and Pattern Recognition (2009)
16. Fan, Y., Dai, D., Schiele, B.: Cossl: Co-learning of representation and classifier for imbalanced semi-supervised learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2022)
17. Friedman, M.: The use of ranks to avoid the assumption of normality implicit in the analysis of variance. *Journal of the American Statistical Association* **32**, 675–701 (1937), <https://api.semanticscholar.org/CorpusID:120581754>
18. Friedman, M.: A comparison of alternative tests of significance for the problem of  $\$m\$$  rankings. *Annals of Mathematical Statistics* **11**, 86–92 (1940), <https://api.semanticscholar.org/CorpusID:121778036>
19. Grandvalet, Y., Bengio, Y.: Semi-supervised learning by entropy minimization. In: Advances in Neural Information Processing Systems (2005)
20. Guo, C., Pleiss, G., Sun, Y., Weinberger, K.Q.: On calibration of modern neural networks. In: International conference on machine learning. pp. 1321–1330. PMLR (2017)
21. Guo, L.Z., Li, Y.F.: Class-imbalanced semi-supervised learning with adaptive thresholding. In: International Conference on Machine Learning. pp. 8082–8094. PMLR (2022)
22. He, Y.Y., Wu, J., Wei, X.S.: Distilling virtual examples for long-tailed recognition. 2021 IEEE/CVF International Conference on Computer Vision (ICCV) (Oct 2021). <https://doi.org/10.1109/iccv48922.2021.00030>, <http://dx.doi.org/10.1109/ICCV48922.2021.00030>
23. Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network. arXiv preprint arXiv:1503.02531 (2015)
24. Hong, Y., Han, S., Choi, K., Seo, S., Kim, B., Chang, B.: Disentangling label distribution for long-tailed visual recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2021)
25. Huang, C., Li, Y., Loy, C.C., Tang, X.: Learning deep representation for imbalanced classification. In: Proceedings of the IEEE conference on Computer Vision and Pattern Recognition (2016)
26. Huh, M., Agrawal, P., Efros, A.A.: What makes imagenet good for transfer learning? arXiv preprint arXiv:1608.08614 (2016)
27. Hyun, M., Jeong, J., Kwak, N.: Class-imbalanced semi-supervised learning. arXiv preprint arXiv:2002.06815 (2020)
28. Kang, B., Xie, S., Rohrbach, M., Yan, Z., Gordo, A., Feng, J., Kalantidis, Y.: Decoupling representation and classifier for long-tailed recognition. In: International conference on learning representations (2020)
29. Kim, J., Hur, Y., Park, S., Yang, E., Hwang, S., Shin, J.: Distribution aligning refinery of pseudo-label for imbalanced semi-supervised learning. In: Advances in neural information processing systems (2020)
30. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: International conference on learning representations (2015)
31. Krause, A., Perona, P., Gomes, R.: Discriminative clustering by regularized information maximization. *Advances in neural information processing systems* **23** (2010)
32. Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images (2009), technical report
33. Kubat, M., Matwin, S., et al.: Addressing the curse of imbalanced training sets: one-sided selection. In: *Icml*. vol. 97, p. 179 (1997)



34. Kuo, C.W., Ma, C.Y., Huang, J.B., Kira, Z.: Featmatch: Feature-based augmentation for semi-supervised learning. In: European Conference on Computer Vision. pp. 479–495. Springer (2020)
35. Lai, Z., Wang, C., Gunawan, H., Cheung, S.C.S., Chuah, C.N.: Smoothed adaptive weighting for imbalanced semi-supervised learning: Improve reliability against unknown distribution data. In: International Conference on Machine Learning. pp. 11828–11843. PMLR (2022)
36. Laine, S., Aila, T.: Temporal ensembling for semi-supervised learning. In: 5th International Conference on Learning Representations (2017)
37. Lazarow, J., Sohn, K., Lee, C.Y., Li, C.L., Zhang, Z., Pfister, T.: Unifying distribution alignment as a loss for imbalanced semi-supervised learning. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 5644–5653 (2023)
38. Lee, D.H.: Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In: Workshop on challenges in representation learning, ICML (2013)
39. Lee, H., Shin, S., Kim, H.: Abc: Auxiliary balanced classifier for class-imbalanced semi-supervised learning. *Advances in Neural Information Processing Systems* **34**, 7082–7094 (2021)
40. Li, J., Tan, Z., Wan, J., Lei, Z., Guo, G.: Nested collaborative learning for long-tailed visual recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6949–6958 (2022)
41. Li, Z., Hoiem, D.: Improving confidence estimates for unfamiliar examples. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2686–2695 (2020)
42. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollar, P.: Focal loss for dense object detection. 2017 IEEE International Conference on Computer Vision (ICCV) (Oct 2017). <https://doi.org/10.1109/iccv.2017.324>, <http://dx.doi.org/10.1109/ICCV.2017.324>
43. Lipton, Z., Wang, Y.X., Smola, A.: Detecting and correcting for label shift with black box predictors. In: International conference on machine learning. pp. 3122–3130. PMLR (2018)
44. Loh, C., Dangovski, R., Sudalairaj, S., Han, S., Han, L., Karlinsky, L., Soljagic, M., Srivastava, A.: On the importance of calibration in semi-supervised learning. arXiv preprint arXiv:2210.04783 (2022)
45. Lucas, T., Weinzaepfel, P., Rogez, G.: Barely-supervised learning: Semi-supervised learning with very few labeled images. *Proceedings of the AAAI Conference on Artificial Intelligence* **36**(2), 1881–1889 (Jun 2022). <https://doi.org/10.1609/aaai.v36i2.20082>, <http://dx.doi.org/10.1609/aaai.v36i2.20082>
46. Ma, C., Elezi, I., Deng, J., Dong, W., Xu, C.: Three heads are better than one: Complementary experts for long-tailed semi-supervised learning. arXiv preprint arXiv:2312.15702 (2023)
47. McLachlan, G.J.: Iterative reclassification procedure for constructing an asymptotically optimal rule of allocation in discriminant analysis. *Journal of the American Statistical Association* **70**(350), 365–369 (1975)
48. Menon, A.K., Jayasumana, S., Rawat, A.S., Jain, H., Veit, A., Kumar, S.: Long-tail learning via logit adjustment. In: International Conference on Learning Representations (2021), <https://openreview.net/forum?id=37nvvqkCo5>
49. Miyato, T., Maeda, S.i., Koyama, M., Ishii, S.: Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE transactions on Pattern Analysis and Machine Intelligence* **41**(8) (2018)
50. Morik, K., Brockhausen, P., Joachims, T.: Combining statistical learning with a knowledge-based approach: a case study in intensive care monitoring. Tech. rep., Technical Report (1999)

51. Oh, Y., Kim, D.J., Kweon, I.S.: Daso: Distribution-aware semantics-oriented pseudo-label for imbalanced semi-supervised learning. 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (Jun 2022). <https://doi.org/10.1109/cvpr52688.2022.00956>, <http://dx.doi.org/10.1109/CVPR52688.2022.00956>
52. Park, S., Hong, Y., Heo, B., Yun, S., Choi, J.Y.: The majority can help the minority: Context-rich minority oversampling for long-tailed classification. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6887–6896 (2022)
53. Powers, D.M.: Applications and explanations of zipf’s law. In: New methods in language processing and computational natural language learning (1998)
54. Reiter, R.: On closed world data bases, p. 300–310. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA (1987)
55. Ren, J., Yu, C., Ma, X., Zhao, H., Yi, S., et al.: Balanced meta-softmax for long-tailed visual recognition. *Advances in neural information processing systems* **33**, 4175–4186 (2020)
56. Sajjadi, M., Javanmardi, M., Tasdizen, T.: Regularization with stochastic transformations and perturbations for deep semi-supervised learning. In: *Advances in neural information processing systems* (2016)
57. Scudder, H.: Probability of error of some adaptive pattern-recognition machines. *IEEE Transactions on Information Theory* (1965)
58. Sohn, K., Berthelot, D., Li, C.L., Zhang, Z., Carlini, N., Cubuk, E.D., Kurakin, A., Zhang, H., Raffel, C.: Fixmatch: Simplifying semi-supervised learning with consistency and confidence. In: *Advances in Neural Information Processing Systems* (2020)
59. Tan, J., Wang, C., Li, B., Li, Q., Ouyang, W., Yin, C., Yan, J.: Equalization loss for long-tailed object recognition. 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (Jun 2020). <https://doi.org/10.1109/cvpr42600.2020.01168>, <http://dx.doi.org/10.1109/cvpr42600.2020.01168>
60. Tarvainen, A., Valpola, H.: Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In: *Advances in Neural Information Processing Systems* (2017)
61. Van Horn, G., Perona, P.: The devil is in the tails: Fine-grained classification in the wild. *arXiv preprint arXiv:1709.01450* (2017)
62. Verma, V., Lamb, A., Beckham, C., Najafi, A., Mitliagkas, I., Lopez-Paz, D., Bengio, Y.: Manifold mixup: Better representations by interpolating hidden states. In: *International Conference on Machine Learning*. pp. 6438–6447. PMLR (2019)
63. Wallace, B.C., Small, K., Brodley, C.E., Trikalinos, T.A.: Class imbalance, redux. In: 2011 IEEE 11th international conference on data mining. pp. 754–763. Ieee (2011)
64. Wang, X., Wu, Z., Lian, L., Yu, S.X.: Debaised learning from naturally imbalanced pseudo-labels. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14647–14657 (2022)
65. Wang, Y., Chen, H., Fan, Y., Sun, W., Tao, R., Hou, W., Wang, R., Yang, L., Zhou, Z., Guo, L.Z., Qi, H., Wu, Z., Li, Y.F., Nakamura, S., Ye, W., Savvides, M., Raj, B., Shinozaki, T., Schiele, B., Wang, J., Xie, X., Zhang, Y.: Usb: A unified semi-supervised learning benchmark for classification (2022)
66. Wang, Y., Chen, H., Heng, Q., Hou, W., Savvides, M., Shinozaki, T., Raj, B., Wu, Z., Wang, J.: Freematch: Self-adaptive thresholding for semi-supervised learning. *arXiv preprint arXiv:2205.07246* (2022)
67. Wang, Y.X., Ramanan, D., Hebert, M.: Learning to model the tail. In: Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*. vol. 30. Curran Associates, Inc. (2017), [https://proceedings.neurips.cc/paper\\_files/paper/2017/file/147ebe637038ca50a1265abac8dea181-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2017/file/147ebe637038ca50a1265abac8dea181-Paper.pdf)

68. Wei, C., Sohn, K., Mellina, C., Yuille, A., Yang, F.: Crest: A class-rebalancing self-training framework for imbalanced semi-supervised learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2021)
69. Wei, T., Gan, K.: Towards realistic long-tailed semi-supervised learning: Consistency is all you need. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3469–3478 (2023)
70. Xiang, L., Ding, G., Han, J.: Learning from multiple experts: Self-paced knowledge distillation for long-tailed classification. In: Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part V 16. pp. 247–263. Springer (2020)
71. Xie, Q., Dai, Z., Hovy, E., Luong, M.T., Le, Q.V.: Unsupervised data augmentation for consistency training. In: Advances in Neural Information Processing Systems (2019)
72. Xie, Y., Manski, C.F.: The logit model and response-based samples. *Sociological Methods & Research* **17**(3), 283–302 (1989)
73. Xu, Y., Shang, L., Ye, J., Qian, Q., Li, Y.F., Sun, B., Li, H., Jin, R.: Dash: Semi-supervised learning with dynamic thresholding. In: International Conference on Machine Learning. pp. 11525–11536. PMLR (2021)
74. Xu, Z., Chai, Z., Yuan, C.: Towards calibrated model for long-tailed visual recognition from prior perspective (2021)
75. Xu, Z., Chai, Z., Yuan, C.: Towards calibrated model for long-tailed visual recognition from prior perspective. *Advances in Neural Information Processing Systems* **34**, 7139–7152 (2021)
76. Yun, S., Han, D., Oh, S.J., Chun, S., Choe, J., Yoo, Y.: Cutmix: Regularization strategy to train strong classifiers with localizable features. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 6023–6032 (2019)
77. Zagoruyko, S., Komodakis, N.: Wide residual networks (2017)
78. Zhang, B., Wang, Y., Hou, W., Wu, H., Wang, J., Okumura, M., Shinozaki, T.: Flexmatch: Boosting semi-supervised learning with curriculum pseudo labeling. *Advances in Neural Information Processing Systems* **34** (2021)
79. Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: Beyond empirical risk minimization. In: 6th International Conference on Learning Representations, ICLR (2018)
80. Zhang, Y., Hooi, B., Hong, L., Feng, J.: Self-supervised aggregation of diverse experts for test-agnostic long-tailed recognition. *Advances in Neural Information Processing Systems* **35**, 34077–34090 (2022)
81. Zhong, Z., Cui, J., Liu, S., Jia, J.: Improving calibration for long-tailed recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16489–16498 (2021)
82. Zhu, X., Anguelov, D., Ramanan, D.: Capturing long-tail distributions of object subcategories. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2014)