
Rethinking the Value of Labels for Improving Class-Imbalanced Learning

Yuzhe Yang
EECS

Massachusetts Institute of Technology
yuzhe@mit.edu

Zhi Xu
EECS

Massachusetts Institute of Technology
zhixu@mit.edu

Abstract

Real-world data often exhibits long-tailed distributions with heavy class imbalance, posing great challenges for deep recognition models. We identify a persisting dilemma on the *value of labels* in the context of imbalanced learning: on the one hand, supervision from labels typically leads to better results than its unsupervised counterparts; on the other hand, heavily imbalanced data naturally incurs “label bias” in the classifier, where the decision boundary can be drastically altered by the majority classes. In this work, we systematically investigate these two facets of labels. We demonstrate, theoretically and empirically, that class-imbalanced learning can significantly benefit in both semi-supervised and self-supervised manners. Specifically, we confirm that (1) positively, imbalanced labels are valuable: given more unlabeled data, the original labels can be leveraged with the extra data to reduce label bias in a semi-supervised manner, which greatly improves the final classifier; (2) negatively however, we argue that imbalanced labels are not useful always: classifiers that are first pre-trained in a self-supervised manner consistently outperform their corresponding baselines. Extensive experiments on large-scale imbalanced datasets verify our theoretically grounded strategies, showing superior performance over previous state-of-the-arts. Our intriguing findings highlight the need to rethink the usage of imbalanced labels in realistic long-tailed tasks. Code is available at <https://github.com/YyzHarry/imbalanced-semi-self>.

1 Introduction

Imbalanced data is ubiquitous in the real world, where large-scale datasets often exhibit long-tailed label distributions [1, 5, 24, 33]. In particular, for critical applications related to safety or health, such as autonomous driving and medical diagnosis, the data are by their nature heavily imbalanced. This poses a major challenge for modern deep learning frameworks [2, 5, 10, 20, 53], where even with specialized techniques such as data re-sampling approaches [2, 5, 41] or class-balanced losses [7, 13, 26], significant performance drops still remain under extreme class imbalance. In order to further tackle the challenge, it is hence vital to understand the different characteristics incurred by class-imbalanced learning.

Yet, distinct from balanced data, the labels in the context of imbalanced learning play a surprisingly controversial role, which leads to a persisting dilemma on the *value of labels*: (1) On the one hand, learning algorithms with supervision from labels typically result in more accurate classifiers than their unsupervised counterparts, demonstrating the *positive* value of labels; (2) On the other hand, however, imbalanced labels naturally impose “label bias” during learning, where the decision boundary can be significantly driven by the majority classes, demonstrating the *negative* impact of labels. Hence, the imbalanced label is seemingly a double-edged sword. Naturally, a fundamental question arises:

How to maximally exploit the value of labels to improve class-imbalanced learning?

In this work, we take initiatives to *systematically* decompose and analyze the two facets of imbalanced labels. As our key contributions, we demonstrate that the positive *and* the negative viewpoint of the

dilemma are indeed both enlightening: they can be effectively exploited, in semi-supervised and self-supervised manners respectively, to significantly improve the state-of-the-art.

On the positive viewpoint, we argue that imbalanced labels are indeed valuable. Theoretically, via a simple Gaussian model, we show that extra unlabeled data benefits imbalanced learning: we obtain a close estimate with high probability that increases exponentially in the amount of unlabeled data, even when unlabeled data are also (highly) imbalanced. Inspired by this, we confirm empirically that by leveraging the label information, class-imbalanced learning can be substantially improved by employing a simple pseudo-labeling strategy, which alleviates the label bias with extra data in a *semi-supervised* manner. Regardless of the imbalanceness on both the labeled and unlabeled data, superior performance is established consistently across various benchmarks, signifying the valuable supervision from the imbalanced labels that leads to substantially better classifiers.

On the negative viewpoint, we demonstrate that imbalanced labels are not advantageous all the time. Theoretically, via a high-dimensional Gaussian model, we show that if given informative representations learned without using labels, with high probability depending on the imbalanceness, we obtain classifier with exponentially small error probability, while the raw classifier always has constant error. Motivated by this, we verify empirically that by abandoning label at the beginning, classifiers that are first pre-trained in a *self-supervised* manner consistently outperform their corresponding baselines, regardless of settings and base training techniques. Significant improvements on large-scale datasets reveal that the biased label information can be greatly compensated through natural self-supervision.

Overall, our intriguing findings highlight the need to rethink the usage of imbalanced labels in realistic imbalanced tasks. With strong performance gains, we establish that not only the positive viewpoint but also the negative one are both promising directions for improving class-imbalanced learning.

Contributions. (i) We are the first to systematically analyze imbalanced learning through two facets of imbalanced label, validating and exploiting its value in novel semi- and self-supervised manners. (ii) We demonstrate, theoretically and empirically, that using unlabeled data can substantially boost imbalanced learning through semi-supervised strategies. (iii) Further, we introduce self-supervised pre-training for class-imbalanced learning without using any extra data, exhibiting both appealing theoretical interpretations and new state-of-the-art on large-scale imbalanced benchmarks.

2 Imbalanced Learning with Unlabeled Data

As motivated, we explore the positive value of labels. Naturally, we consider scenarios where extra unlabeled data is available and hence, the limited labeling information is critical. Through a simple theoretical model, we first build intuitions on how different ingredients of the originally imbalanced data and the extra data affect the overall learning process. With a clearer picture in mind, we design comprehensive experiments to confirm the efficacy of this direction on boosting imbalanced learning.

2.1 Theoretical Motivation

Consider a binary classification problem with the data generating distribution P_{XY} being a mixture of two Gaussians. In particular, the label Y is either positive (+1) or negative (-1) with equal probability (i.e., 0.5). Condition on $Y = +1$, $X|Y = +1 \sim \mathcal{N}(\mu_1, \sigma^2)$ and similarly, $X|Y = -1 \sim \mathcal{N}(\mu_2, \sigma^2)$. Without loss of generality, let $\mu_1 > \mu_2$. It is straightforward to verify that the optimal Bayes's classifier is $f(x) = \text{sign}(x - \frac{\mu_1 + \mu_2}{2})$, i.e., classify x as +1 if $x > (\mu_1 + \mu_2)/2$. Therefore, in the following, we measure our ability to learn $(\mu_1 + \mu_2)/2$ as a proxy for performance.

Suppose that a base classifier f_B , trained on imbalanced training data, is given. We consider the case where extra unlabeled data $\{\tilde{X}_i\}_{i=1}^{\tilde{n}}$ (potentially also imbalanced) from P_{XY} are available, and study how this affects our performance with the label information from f_B . Precisely, we create pseudo-label for $\{\tilde{X}_i\}_{i=1}^{\tilde{n}}$ using f_B . Let $\{\tilde{X}_i^+\}_{i=1}^{\tilde{n}_+}$ be the set of unlabeled data whose pseudo-label is +1; similarly let $\{\tilde{X}_i^-\}_{i=1}^{\tilde{n}_-}$ be the negative set. Naturally, when the training data is imbalanced, f_B is likely to exhibit different accuracy for different class. We model this as follows. Consider the case where pseudo-label is +1 and let $\{I_i^+\}_{i=1}^{\tilde{n}_+}$ be the indicator that the i -th pseudo-label is correct, i.e., if $I_i^+ = 1$, then $\tilde{X}_i^+ \sim \mathcal{N}(\mu_1, \sigma^2)$ and otherwise $\tilde{X}_i^+ \sim \mathcal{N}(\mu_2, \sigma^2)$. We assume $I_i^+ \sim \text{Bernoulli}(p)$, which means f_B has an accuracy of p for the positive class. Analogously, we define $\{I_i^-\}_{i=1}^{\tilde{n}_-}$, i.e., if $I_i^- = 1$, then $\tilde{X}_i^- \sim \mathcal{N}(\mu_2, \sigma^2)$ and otherwise $\tilde{X}_i^- \sim \mathcal{N}(\mu_1, \sigma^2)$. Let $I_i^- \sim \text{Bernoulli}(q)$, which

means f_B has an accuracy of q for the negative class. Denote by $\Delta \triangleq p - q$ the imbalance in accuracy. As mentioned, we aim to learn $(\mu_1 + \mu_2)/2$ with the above setup, via the extra unlabeled data. It is natural to construct our estimate as $\hat{\theta} = \frac{1}{2}(\sum_{i=1}^{\tilde{n}_+} \tilde{X}_i^+ / \tilde{n}_+ + \sum_{i=1}^{\tilde{n}_-} \tilde{X}_i^- / \tilde{n}_-)$. Then, we have:

Theorem 1. Consider the above setup. For any $\delta > 0$, with probability at least $1 - 2e^{-\frac{2\delta^2}{9\sigma^2} \cdot \frac{\tilde{n}_+ \tilde{n}_-}{\tilde{n}_+ + \tilde{n}_-}} - 2e^{-\frac{8\tilde{n}_+ \delta^2}{9(\mu_1 - \mu_2)^2}} - 2e^{-\frac{8\tilde{n}_- \delta^2}{9(\mu_1 - \mu_2)^2}}$ our estimates $\hat{\theta}$ satisfies

$$|\hat{\theta} - (\mu_1 + \mu_2)/2 - \Delta(\mu_1 - \mu_2)/2| \leq \delta.$$

Interpretation. The result illustrates several interesting aspects. (1) *Training data imbalance affects the accuracy of our estimation.* For heavily imbalanced training data, we expect the base classifier to have a large difference in accuracy between major and minor classes. That is, the more imbalanced the data is, the larger the gap Δ would be, which influences the closeness between our estimate and desired value $(\mu_1 + \mu_2)/2$. (2) *Unlabeled data imbalance affects the probability of obtaining such a good estimation.* For a reasonably good base classifier, we can roughly view \tilde{n}_+ and \tilde{n}_- as approximations for the number of actually positive and negative data in unlabeled set. For term $2 \exp(-\frac{2\delta^2}{9\sigma^2} \cdot \frac{\tilde{n}_+ \tilde{n}_-}{\tilde{n}_+ + \tilde{n}_-})$, note that $\frac{\tilde{n}_+ \tilde{n}_-}{\tilde{n}_+ + \tilde{n}_+}$ is maximized when $\tilde{n}_+ = \tilde{n}_-$, i.e., balanced unlabeled data. For terms $2 \exp(-\frac{8\tilde{n}_+ \delta^2}{9(\mu_1 - \mu_2)^2})$ and $2 \exp(-\frac{8\tilde{n}_- \delta^2}{9(\mu_1 - \mu_2)^2})$, if the unlabeled data is heavily imbalanced, then the term corresponding to the minor class dominates and can be moderately large. Our probability of success would be higher with balanced data, but in any case, more unlabeled data is *always* helpful.

2.2 Semi-Supervised Imbalanced Learning Framework

Our theoretical findings show that pseudo-label (and hence the label information in training data) can be helpful in imbalanced learning. The degree to which this is useful is affected by the imbalanceness of the data. Inspired by these, we systematically probe the effectiveness of unlabeled data and study how it can improve realistic imbalanced task, especially with varying degree of imbalanceness.

Semi-Supervised Imbalanced Learning. To harness the unlabeled data for alleviating the inherent imbalance, we propose to adopt the classic *self-training* framework, which performs semi-supervised learning (SSL) by generating *pseudo-labels* for unlabeled data. Precisely, we obtain an intermediate classifier $f_{\hat{\theta}}$ using the original imbalanced dataset \mathcal{D}_L , and apply it to generate pseudo-labels \hat{y} for unlabeled data \mathcal{D}_U . The data and pseudo-labels are combined to learn a final model $f_{\hat{\theta}_t}$ by minimizing a loss function as $\mathcal{L}(\mathcal{D}_L, \theta) + \omega \mathcal{L}(\mathcal{D}_U, \theta)$, where ω is the unlabeled weight. This procedure seeks to remodel the class distribution with \mathcal{D}_U , obtaining better class boundaries especially for tail classes.

We remark that besides self-training, more advanced SSL techniques can be easily incorporated into our framework by modifying only the loss function, which we will study later. As we do not specify the learning strategy of $f_{\hat{\theta}}$ and $f_{\hat{\theta}_t}$, the semi-supervised framework is also compatible with existing class-imbalanced learning methods. Accordingly, we demonstrate the value of unlabeled data — a simple self-training procedure can lead to substantially better performance for imbalanced learning.

Experimental Setup. We conduct thorough experiments on artificially created long-tailed versions of CIFAR-10 [7] and SVHN [36], which naturally have their unlabeled part with similar distributions: 80 Million Tiny Images [48] for CIFAR-10, and SVHN’s own extra set [36] with labels removed for SVHN. Following [7, 11], the class *imbalance ratio* ρ is defined as the sample size of the most frequent (head) class divided by that of the least frequent (tail) class. Similarly for \mathcal{D}_U , we define the *unlabeled imbalance ratio* ρ_U in the same way. More details of datasets are reported in Appendix D.

For long-tailed dataset with a fixed ρ , we augment it with 5 times more unlabeled data, denoted as $\mathcal{D}_U @ 5x$. As we seek to study the effect of unlabeled imbalance ratio, the total size of $\mathcal{D}_U @ 5x$ is fixed, where we vary ρ_U to obtain corresponding imbalanced \mathcal{D}_U . We select standard cross-entropy (CE) training, and a recently proposed state-of-the-art imbalanced learning method LDAM-DRW [7] as baseline methods. We follow [7, 25, 33] to evaluate models on corresponding balanced test datasets.

2.2.1 Main Results

CIFAR-10-LT & SVHN-LT. Table 1 summarizes the results on two long-tailed datasets. For each ρ , we vary the type of class imbalance in \mathcal{D}_U to be uniform ($\rho_U = 1$), half as labeled ($\rho_U = \rho/2$), same ($\rho_U = \rho$), and doubled ($\rho_U = 2\rho$). As shown in the table, the SSL scheme can consistently and

Table 1: Top-1 test errors (%) of ResNet-32 on long-tailed CIFAR-10 and SVHN. We compare SSL using 5x unlabeled data ($\mathcal{D}_U@5x$) with corresponding supervised baselines. Imbalanced learning can be drastically improved with unlabeled data, which is consistent across different ρ_U and learning strategies.

(a) CIFAR-10-LT												
Imbalance Ratio (ρ)	100				50				10			
\mathcal{D}_U Imbalance Ratio (ρ_U)	1	$\rho/2$	ρ	2ρ	1	$\rho/2$	ρ	2ρ	1	$\rho/2$	ρ	2ρ
CE	29.64				25.19				13.61			
CE + $\mathcal{D}_U@5x$	17.48	<u>18.42</u>	18.74	20.06	16.79	<u>16.88</u>	18.36	19.94	10.22	<u>10.48</u>	10.86	11.04
LDAM-DRW [7]	22.97				19.06				11.84			
LDAM-DRW + $\mathcal{D}_U@5x$	14.96	<u>15.18</u>	15.33	15.55	14.33	<u>14.70</u>	14.93	15.24	8.72	8.24	<u>8.68</u>	8.97

(b) SVHN-LT												
Imbalance Ratio (ρ)	100				50				10			
\mathcal{D}_U Imbalance Ratio (ρ_U)	1	$\rho/2$	ρ	2ρ	1	$\rho/2$	ρ	2ρ	1	$\rho/2$	ρ	2ρ
CE	19.98				17.50				11.46			
CE + $\mathcal{D}_U@5x$	13.02	<u>13.73</u>	14.65	15.04	13.07	<u>13.36</u>	<u>13.16</u>	14.54	10.01	<u>10.20</u>	<u>10.06</u>	10.71
LDAM-DRW [7]	16.66				14.59				10.27			
LDAM-DRW + $\mathcal{D}_U@5x$	11.32	<u>11.70</u>	11.92	12.78	10.98	<u>11.14</u>	11.26	11.51	<u>8.94</u>	9.08	8.70	9.35

substantially improve existing techniques across different ρ . Notably, under extreme class imbalance ($\rho = 100$), using unlabeled data can lead to +10% on CIFAR-10-LT, and +6% on SVHN-LT.

Imbalanced Distribution in Unlabeled Data. As indicated by Theorem 1, unlabeled data imbalance affects the learning of final classifier. We observe in Table 1 that gains indeed vary under different ρ_U , with smaller ρ_U (i.e., more balanced \mathcal{D}_U) leading to larger gains. Interestingly however, as the original dataset becomes more balanced, the benefits from \mathcal{D}_U tend to be similar across different ρ_U .

Qualitative Results. To further understand the effect of unlabeled data, we visualize representations learned with vanilla CE (Fig. 1a) and with SSL (Fig. 1b) using t-SNE [34]. The figures show that imbalanced training set can lead to poor class separation, particularly for tail classes, which results in mixed class boundary during class-balanced inference. In contrast, by leveraging unlabeled data, the boundary of tail classes can be better shaped, leading to clearer separation and better performance.

Summary. Consistently across various settings, class-imbalanced learning tasks benefit *greatly* from additional unlabeled data. The superior performance obtained demonstrates the positive value of imbalanced labels as being capable of exploiting the unlabeled data for extra benefits.

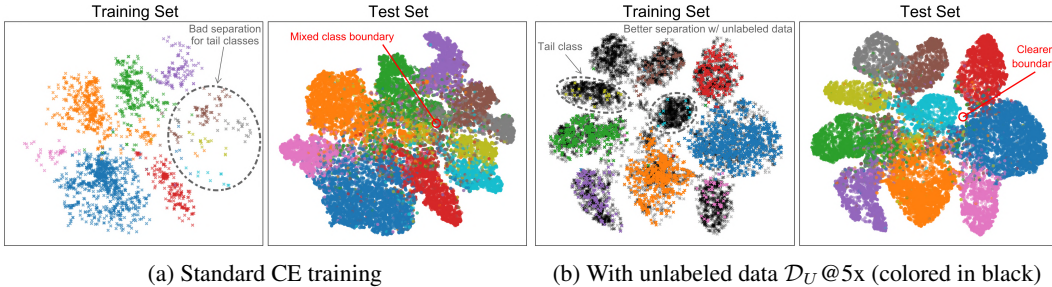


Figure 1: t-SNE visualization of training & test set on SVHN-LT. Using unlabeled data helps to shape clearer class boundaries and results in better class separation, especially for the tail classes.

2.2.2 Further Analysis and Ablation Studies

Different Semi-Supervised Methods (Appendix E.1). In addition to the simple pseudo-label, we select more advanced SSL techniques [35, 46] and explore the effect of different methods under the

imbalanced settings. In short, all SSL methods can achieve remarkable performance gains over the supervised baselines, with more advanced SSL method enjoying larger improvements in general.

Generalization on Minority Classes (Appendix E.2). Besides the reported top-1 accuracy, we further study generalization on each class with and without unlabeled data. We show that while all classes can obtain certain improvements, the minority classes tend to exhibit larger gains.

Unlabeled & Labeled Data Amount (Appendix E.3 & E.4). Following [39], we investigate how different amounts of \mathcal{D}_U as well as \mathcal{D}_L can affect our SSL approach in imbalanced learning. We find that larger \mathcal{D}_U or \mathcal{D}_L often brings higher gains, with gains gradually diminish as data amount grows.

3 A Closer Look at Unlabeled Data under Class Imbalance

With significant boost in performance, we confirm the value of imbalanced labels with extra unlabeled data. Such success naturally motivates us to dig deeper into the techniques and investigate whether SSL is the solution to practical imbalanced data. Indeed, in the balanced case, SSL is known to have issues in certain scenarios when the unlabeled data is not ideally constructed. Techniques are often sensitive to the relevance of unlabeled data, and performance can even degrade if the unlabeled data is largely mismatched [39]. The situation becomes even more complicated when imbalance comes into the picture. The relevant unlabeled data could also exhibit long-tailed distributions. With this, we aim to further provide an informed message on the utility of semi-supervised techniques.

Data Relevance under Imbalance. We construct sets of unlabeled data with the same imbalance ratio as the training data but varying relevance. Specifically, we mix the original unlabeled dataset with irrelevant data, and create unlabeled datasets with varying data relevance ratios (detailed setup can be found in Appendix D.2). Fig. 2 shows that in imbalanced learning, adding unlabeled data from mismatched classes can actually *hurt* performance. The relevance has to be as high as 60% to be effective, while better results could be obtained without unlabeled data at all when it is only moderately relevant. The observations are consistent with the balanced cases [39].

Varying ρ_U under Sufficient Data Relevance. Furthermore, even with enough relevance, what will happen if the relevant unlabeled data is (heavily) long-tailed? As presented in Fig. 3, for a fixed relevance, the higher ρ_U of the relevant data is, the higher the test error. In this case, to be helpful, ρ_U cannot be larger than 50 (which is the imbalance ratio of the training data). This highlights that unlike traditional setting, the imbalance of the unlabeled data imposes an additional challenge.

Why Do These Matter. The observations signify that semi-supervised techniques should be applied with care in practice for imbalanced learning. When it is readily to obtain relevant unlabeled data of each class, they are particularly powerful as we demonstrate. However, certain practical applications, especially those extremely imbalanced, are situated at the worst end of the spectrum. For example, in medical diagnosis, positive samples are always scarce; even with access to more “unlabeled” medical records, positive samples remain sparse, and the confounding issues (e.g., other disease or symptoms) undoubtedly hurt relevance. Therefore, one would expect the imbalance ratio of the unlabeled data to be higher, if not lower, than the training data in these applications.

To summarize, unlabeled data are useful. However, semi-supervised learning alone is *not* sufficient to solve the imbalanced problem. Other techniques are needed in case the application does not allow constructing meaningful unlabeled data, and this, exactly motivates our subsequent studies.




Figure 2: Test errors of different unlabeled data relevance ratios on CIFAR-10-LT with $\rho = 50$. We fix $\rho_U = 50$ for the relevant unlabeled data.




Figure 3: Test errors of different ρ_U of relevant unlabeled data on CIFAR-10-LT with $\rho = 50$. We fix the unlabeled data relevance ratio as 60%.

4 Imbalanced Learning from Self-Supervision

The previous studies naturally motivate our next quest: can the negative viewpoint of the dilemma, i.e., the imbalanced labels introduce bias and hence are “unnecessary”, be successfully exploited as well to advance imbalanced learning? In answering this, our goal is to seek techniques that can be broadly applied without extra data. Through a theoretical model, we first justify the usage of self-supervision in the context of imbalanced learning. Extensive experiments are then designed to verify its effectiveness, proving that thinking through the negative viewpoint is indeed promising as well.

4.1 Theoretical Motivation

We start with another inspiring model to study how imbalanced learning benefits from self-supervision. Consider d -dimensional binary classification with data generating distribution P_{XY} being a mixture of Gaussians. In particular, the label $Y = +1$ with probability p_+ while $Y = -1$ with probability $p_- = 1 - p_+$. Let $p_- \geq 0.5$, i.e., major class is negative. Condition on $Y = +1$, X is a d -dimensional isotropic Gaussian, i.e., $X|Y = +1 \sim \mathcal{N}(0, \sigma_1^2 \mathbf{I}_d)$. Similarly, $X|Y = -1 \sim \mathcal{N}(0, \beta \sigma_1^2 \mathbf{I}_d)$ for some constant $\beta > 3$, i.e., the negative samples have larger variance. The training data, $\{(X_i, Y_i)\}_{i=1}^N$, could be highly imbalanced, and we denote by N_+ & N_- as number of positive & negative samples.

To develop our intuition, we consider learning a linear classifier with and without self-supervision. In particular, consider the class of linear classifiers $f(x) = \text{sign}(\langle \theta, \text{feature} \rangle + b)$, where feature would be the raw input X in standard training, and for the self-supervised learning, feature would be $Z = \psi(X)$ for some representation ψ learned through a self-supervised task. For convenience, we consider the case where the intercept $b \geq 0$. We assume a properly designed black-box self-supervised task so that the learned representation is $Z = k_1 \|X\|_2^2 + k_2$, where $k_1, k_2 > 0$. Precisely, this means that we have access to the new features Z_i for the i -th data after the black-box self-supervised step, without knowing explicitly what the transformation ψ is. Finally, we measure the performance of a classifier f using the standard error probability: $\text{err}_f = \mathbb{P}_{(X,Y) \sim P_{XY}}(f(X) \neq Y)$.

Theorem 2. *Let Φ be the CDF of $\mathcal{N}(0, 1)$. For any linear classifier of the form $f(X) = \text{sign}(\langle \theta, X \rangle + b)$ where $b > 0$, the error probability satisfies: $\text{err}_f = p_+ \Phi\left(-\frac{b}{\|\theta\|_2 \sigma_1}\right) + p_- \Phi\left(\frac{b}{\|\theta\|_2 \sqrt{\beta} \sigma_1}\right) \geq \frac{1}{4}$.*

Theorem 2 states that for standard training, regardless of whether the training data is imbalanced or not, the linear classifier cannot have an accuracy $\geq 3/4$. This is rather discouraging for such a simple case. However, we show that self-supervision and training on the resulting Z provides a better classifier. Consider the same linear class $f(x) = \text{sign}(\langle \theta, \text{feature} \rangle + b)$, $b > 0$ and following explicit classifier with feature $Z = \psi(X)$: $f_{ss}(X) = \text{sign}(-Z + b)$, $b = \frac{1}{2} \left(\frac{\sum_{i=1}^{N_+} \mathbf{1}_{\{Y_i=+1\}} Z_i}{N_+} + \frac{\sum_{i=1}^{N_-} \mathbf{1}_{\{Y_i=-1\}} Z_i}{N_-} \right)$. The next theorem shows a high probability error bound for the performance of this linear classifier.

Theorem 3. *Consider the linear classifier with self-supervised learning, f_{ss} . For any $\delta \in (0, \frac{\beta-1}{\beta+1})$, we have that with probability at least $1 - 2e^{-N_- d \delta^2 / 8} - 2e^{-N_+ d \delta^2 / 8}$, the classifier satisfies*

$$\text{err}_{f_{ss}} \leq \begin{cases} p_+ e^{-d \cdot \frac{(\beta-1-(1+\beta)\delta)^2}{32}} + p_- e^{-d \cdot \frac{(\beta-1-(1+\beta)\delta)^2}{32\beta^2}}, & \text{if } \delta \in [\frac{\beta-3}{\beta+1}, \frac{\beta-1}{\beta+1}); \\ p_+ e^{-d \cdot \frac{(\beta-1-(1+\beta)\delta)}{16}} + p_- e^{-d \cdot \frac{(\beta-1-(1+\beta)\delta)^2}{32\beta^2}}, & \text{if } \delta \in (0, \frac{\beta-3}{\beta+1}). \end{cases}$$

Interpretation. Theorem 3 implies the following interesting observations. By first abandoning imbalanced labels and learning an informative representation via self-supervision, (1) *With high probability, we obtain a satisfying classifier f_{ss} , whose error probability decays exponentially on the dimension d .* The probability of obtaining such a classifier also depends exponentially on d and the number of data. These are rather appealing as modern data is of extremely high dimension. That is, even for imbalanced data, one could obtain a good classifier with proper self-supervised training; (2) *Training data imbalance affects our probability of obtaining such a satisfying classifier.* Precisely, given N data, if it is highly imbalanced with an extremely small N_+ , then the term $2 \exp(-N_+ d \delta^2 / 8)$ could be moderate and dominate $2 \exp(-N_- d \delta^2 / 8)$. With more or less balanced data (or just more data), our probability of success increases. Nevertheless, as the dependence is exponential, even for imbalanced training data, self-supervised learning can still help to obtain a satisfying classifier.

Table 2: Top-1 test error rates (%) of ResNet-32 on long-tailed CIFAR-10 and CIFAR-100. Using SSP, we consistently improve different imbalanced learning techniques, and achieve the best performance.

Dataset	CIFAR-10-LT			CIFAR-100-LT		
Imbalance Ratio (ρ)	100	50	10	100	50	10
CE	29.64	25.19	13.61	61.68	56.15	44.29
CB-CE [11]	27.63	21.95	13.23	61.44	55.45	42.88
CB-CE + <i>SSP</i>	23.47	19.60	11.57	56.94	52.91	41.94
Focal [32]	29.62	23.29	13.34	61.59	55.68	44.22
CB-Focal [11]	25.43	20.73	12.90	60.40	54.83	42.01
CB-Focal + <i>SSP</i>	22.90	18.74	11.75	57.03	53.12	41.16
CE-DRW [7]	24.94	21.10	13.57	59.49	55.31	43.78
CE-DRS [7]	25.53	21.39	13.73	59.62	55.46	43.95
CE-DRW + <i>SSP</i>	23.04	19.93	12.66	57.21	53.57	41.77
LDAM [7]	26.65	23.18	13.04	60.40	55.03	43.09
LDAM-DRW [7]	22.97	19.06	11.84	57.96	53.85	41.29
LDAM-DRW + <i>SSP</i>	22.17	17.87	11.47	56.57	52.89	41.09

4.2 Self-Supervised Imbalanced Learning Framework

Motivated by our theoretical results, we again seek to systematically study how self-supervision can help and improve class-imbalanced tasks in realistic settings.

Self-Supervised Imbalanced Learning. To utilize self-supervision for overcoming the intrinsic label bias, we propose to, in the first stage of learning, abandon the label information and perform *self-supervised pre-training* (SSP). This procedure aims to learn better initialization that is more label-agnostic from the imbalanced dataset. After the first stage of learning with self-supervision, we can then perform any standard training approach to learn the final model initialized by the pre-trained network. Since the pre-training is independent of the learning approach applied in the normal training stage, such strategy is compatible with any existing imbalanced learning techniques.

Once the self-supervision yields good initialization, the network can benefit from the pre-training tasks and finally learn more generalizable representations. Since SSP can be easily embedded with existing techniques, we would expect that any base classifiers can be consistently improved using SSP. To this end, we empirically evaluate SSP and show that it leads to consistent and substantial improvements in class-imbalanced learning, across various large-scale long-tailed benchmarks.

Experimental Setup. We perform extensive experiments on benchmark CIFAR-10-LT and CIFAR-100-LT, as well as large-scale long-tailed datasets including ImageNet-LT [33] and real-world dataset iNaturalist 2018 [24]. We again evaluate models on the corresponding balanced test datasets [7, 25, 33]. We use Rotation [16] as SSP method on CIFAR-LT, and MoCo [19] on ImageNet-LT and iNaturalist. In the classifier learning stage, we follow [7, 25] to train all models for 200 epochs on CIFAR-LT, and 90 epochs on ImageNet-LT and iNaturalist. Other implementation details are in Appendix D.3.

4.2.1 Main Results

CIFAR-10-LT & CIFAR-100-LT. We present imbalanced classification on long-tailed CIFAR in Table 2. We select standard cross-entropy (CE) loss, Focal loss [32], class-balanced (CB) loss [11], re-weighting or re-sampling training schedule [7], and recently proposed LDAM-DRW [7] as state-of-the-art methods. We group the competing methods into four sessions according to which basic loss or learning strategies they use. As Table 2 reports, in each session across different ρ , adding SSP consistently outperforms the competing ones by notable margins. Further, the benefits of SSP become more significant as ρ increases, demonstrating the value of self-supervision under class imbalance.

ImageNet-LT & iNaturalist 2018. Besides standard and balanced CE training, we also select other baselines including OLTR [33] and recently proposed classifier re-training (cRT) [25] which achieves state-of-the-art on large-scale datasets. Table 3 and 4 present results on two datasets, respectively.

Table 3: Top-1 test error rates (%) on ImageNet-LT. † denotes results reproduced with authors’ code.

Method	ResNet-10	ResNet-50
CE (Uniform)	65.2	61.6
CE (Uniform) + <i>SSP</i>	64.1	54.4
CE (Balanced)	62.9	59.7
CE (Balanced) + <i>SSP</i>	61.6	52.4
OLTR [33]	64.4	62.6 [†]
OLTR + <i>SSP</i>	62.3	53.9
cRT [25]	58.2	52.7
cRT + <i>SSP</i>	56.8	48.7

Table 4: Top-1 test error rates (%) on iNaturalist 2018. † denotes results reproduced with authors’ code.

Method	ResNet-50
CE (Uniform)	39.3
CE (Uniform) + <i>SSP</i>	35.6
CE (Balanced)	36.5
CE (Balanced) + <i>SSP</i>	34.1
LDAM-DRW [7]	35.4 [†]
LDAM-DRW + <i>SSP</i>	33.7
cRT [25]	34.8
cRT + <i>SSP</i>	31.9

On both datasets, adding SSP sets new state-of-the-arts, substantially improving current techniques with 4% absolute performance gains. The consistent results confirm the success of applying SSP in realistic large-scale imbalanced learning scenarios.

Qualitative Results. To gain additional insight, we look at the t-SNE projection of learnt representations for both vanilla CE training (Fig. 4a) and with SSP (Fig. 4b). For each method, the projection is performed over both training and test data, thus providing the same decision boundary for better visualization. The figures show that the decision boundary of vanilla CE can be greatly altered by the head classes, which results in the large leakage of tail classes during (balanced) inference. In contrast, using SSP sustains clear separation with less leakage, especially between adjacent head and tail class.

Summary. Regardless of the settings and the base training techniques, adding our self-supervision framework in the first stage of learning can uniformly boost the final performance. This highlights that the negative, “unnecessary” viewpoint of the imbalanced labels is also valuable and effective in improving the state-of-the-art imbalanced learning approaches.

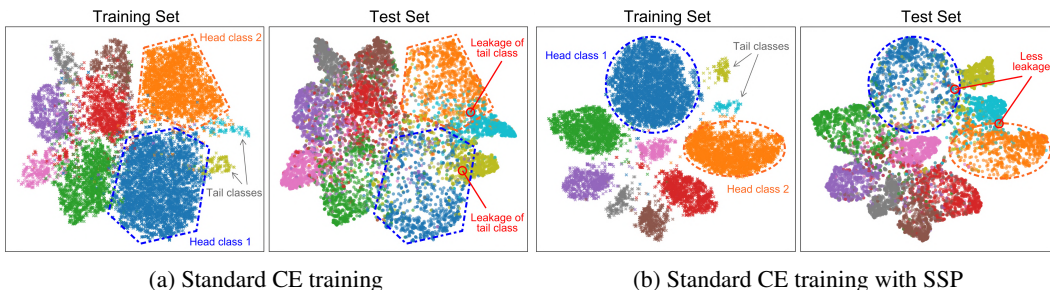


Figure 4: t-SNE visualization of training & test set on CIFAR-10-LT. Using SSP helps mitigate the tail classes leakage during testing, which results in better learned boundaries and representations.

4.2.2 Further Analysis and Ablation Studies

Different Self-Supervised Methods (Appendix F.1). We select four different SSP techniques, and evaluate them across four benchmark datasets. In general, all SSP methods can lead to notable gains compared to the baseline, while interestingly the gain varies across methods. We find that MoCo [19] performs better on large-scale datasets, while Rotation [16] achieves better results on smaller ones.

Generalization on Minority Classes (Appendix F.2). In addition to the top-1 accuracy, we further study the generalization on each specific class. On both CIFAR-10-LT and ImageNet-LT, we observe that SSP can lead to consistent gains across all classes, where trends are more evident for tail classes.

Imbalance Type (Appendix F.3). While the main paper is focused on the long-tailed imbalance distribution which is the most common type of imbalance, we remark that other imbalance types are also suggested in literature [5]. We present ablation study on another type of imbalance, i.e., step imbalance [5], where consistent improvements and conclusions are verified when adding SSP.

5 Related Work

Imbalanced Learning & Long-tailed Recognition. Literature is rich on learning long-tailed imbalanced data. Classical methods have been focusing on designing data re-sampling strategies, including over-sampling the minority classes [2,41,44] as well as under-sampling the frequent classes [5,31]. In addition, cost-sensitive re-weighting schemes [6,22,23,27,52] are proposed to (dynamically) adjust weights during training for different classes or even different samples. For imbalanced classification problems, another line of work develops class-balanced losses [7,11,13,26,32] by taking intra- or inter-class properties into account. Other learning paradigms, including transfer learning [33,54], metric learning [55,58], and meta-learning [1,45], have also been explored. Recent studies [25,59] also find that decoupling the representation and classifier leads to better long-tailed learning results. In contrast to existing works, we provide systematic strategies through two viewpoints of imbalanced labels, which boost imbalanced learning in both semi-supervised and self-supervised manners.

Semi-Supervised Learning. Semi-supervised learning is concerned with learning from both unlabeled and labeled samples, where typical methods ranging from entropy minimization [18], pseudo-labeling [30], to generative models [15,17,29,42,56]. Recently, a line of work that proposes to use consistency-based regularization approaches has shown promising performance in semi-supervised tasks, where consistency loss is integrated to push the decision boundary to low-density areas using unlabeled data [3,28,35,43,46,50]. The common evaluation protocol assumes the unlabeled data comes from the same or similar distributions as labeled data, while authors in [39] argue it may not reflect realistic settings. In our work, we consider the data imbalance in both labeled and unlabeled datasets, as well as the data relevance for the unlabeled data, which altogether provides a principled setup on semi-supervised learning for imbalanced learning tasks.

Self-Supervised Learning. Learning with self-supervision has recently attracted increasing interests, where early approaches mainly rely on pretext tasks, including exemplar classification [14], predicting relative locations of image patches [12], image colorization [57], solving jigsaw puzzles of image patches [37], object counting [38], clustering [9], and predicting the rotation of images [16]. More recently, a line of work based on contrastive losses [4,19,21,40,47] shows great success in self-supervised representation learning, where similar embeddings are learned for different views of the same training example, and dissimilar embeddings for different training examples. Our work investigates self-supervised pre-training in class-imbalanced context, revealing surprising yet intriguing findings on how the self-supervision can help alleviate the biased label effect in imbalanced learning.

6 Conclusion

We systematically study the value of labels in class-imbalanced learning, and propose two theoretically grounded strategies to understand, validate, and leverage such imbalanced labels in both semi- and self-supervised manners. On both sides, sound theoretical guarantees as well as superior performance on large-scale imbalanced datasets are demonstrated, confirming the significance of the proposed schemes. Our findings open up new avenues for learning imbalanced data, highlighting the need to rethink what is the best way to leverage the inherently biased labels to improve imbalanced learning.

Broader Impact

Real-world data often exhibits skewed distributions with a long tail, rather than the ideal uniform distributions over each class. We tackle this important problem through two novel perspectives: (1) using unlabeled data without depending on additional human labeling; (2) explore intrinsic properties from data itself with self-supervision. These simple yet effective strategies introduce new frameworks for improving generic imbalanced learning tasks, which we believe will broadly benefit practitioners dealing with heavily imbalanced data in realistic applications.

On the other hand, however, we only extensively test our strategies on academic datasets. In many real-world applications such as autonomous driving, medical diagnosis, and healthcare, beyond being naturally imbalanced, the data may impose additional constraints on learning process and final models, e.g., being fair or private. We focus on standard accuracy as our measure and largely ignore other ethical issues in imbalanced data, especially in minor classes. As such, the risk of producing unfair or biased outputs reminds us to carry rigorous validations in critical, high-stakes applications.

References

- [1] Muhammad Abdullah Jamal, Matthew Brown, Ming-Hsuan Yang, Liqiang Wang, and Boqing Gong. Rethinking class-balanced methods for long-tailed visual recognition from a domain adaptation perspective. *arXiv preprint arXiv:2003.10780*, 2020.
- [2] Shin Ando and Chun Yuan Huang. Deep over-sampling framework for classifying imbalanced data. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 770–785. Springer, 2017.
- [3] Philip Bachman, Ouais Alsharif, and Doina Precup. Learning with pseudo-ensembles. In *Advances in Neural Information Processing Systems*, pages 3365–3373, 2014.
- [4] Philip Bachman, R Devon Hjelm, and William Buchwalter. Learning representations by maximizing mutual information across views. In *NeurIPS*, 2019.
- [5] Mateusz Buda, Atsuto Maki, and Maciej A Mazurowski. A systematic study of the class imbalance problem in convolutional neural networks. *Neural Networks*, 106:249–259, 2018.
- [6] Jonathon Byrd and Zachary Lipton. What is the effect of importance weighting in deep learning? In *International Conference on Machine Learning*, pages 872–881, 2019.
- [7] Kaidi Cao, Colin Wei, Adrien Gaidon, Nikos Arechiga, and Tengyu Ma. Learning imbalanced datasets with label-distribution-aware margin loss. In *NeurIPS*, 2019.
- [8] Yair Carmon, Aditi Raghunathan, Ludwig Schmidt, John C Duchi, and Percy S Liang. Unlabeled data improves adversarial robustness. In *NeurIPS*, 2019.
- [9] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In *ECCV*, 2018.
- [10] Ronan Collobert and Jason Weston. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *Proceedings of the 25th international conference on Machine learning*, pages 160–167, 2008.
- [11] Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. Class-balanced loss based on effective number of samples. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9268–9277, 2019.
- [12] Carl Doersch, Abhinav Gupta, and Alexei A Efros. Unsupervised visual representation learning by context prediction. In *ICCV*, pages 1422–1430, 2015.
- [13] Qi Dong, Shaogang Gong, and Xiatian Zhu. Imbalanced deep learning by minority class incremental rectification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(6):1367–1381, Jun 2019.
- [14] Alexey Dosovitskiy, Jost Tobias Springenberg, Martin Riedmiller, and Thomas Brox. Discriminative unsupervised feature learning with convolutional neural networks. In *NIPS*, pages 766–774, 2014.
- [15] Vincent Dumoulin, Ishmael Belghazi, Ben Poole, Olivier Mastropietro, Alex Lamb, Martin Arjovsky, and Aaron Courville. Adversarially learned inference. *arXiv:1606.00704*, 2016.
- [16] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. *arXiv preprint arXiv:1803.07728*, 2018.
- [17] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- [18] Yves Grandvalet and Yoshua Bengio. Semi-supervised learning by entropy minimization. In *Advances in neural information processing systems*, pages 529–536, 2005.
- [19] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. *arXiv preprint arXiv:1911.05722*, 2019.
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.
- [21] Olivier J Hénaff, Ali Razavi, Carl Doersch, SM Eslami, and Aaron van den Oord. Data-efficient image recognition with contrastive predictive coding. *arXiv:1905.09272*, 2019.

- [22] Chen Huang, Yining Li, Chen Change Loy, and Xiaoou Tang. Learning deep representation for imbalanced classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5375–5384, 2016.
- [23] Chen Huang, Yining Li, Change Loy Chen, and Xiaoou Tang. Deep imbalanced learning for face recognition and attribute prediction. *IEEE transactions on pattern analysis and machine intelligence*, 2019.
- [24] iNaturalist. The inaturalist 2018 competition dataset. https://github.com/visipedia/inat_comp/tree/master/2018, 2018.
- [25] Bingyi Kang, Saining Xie, Marcus Rohrbach, Zhicheng Yan, Albert Gordo, Jiashi Feng, and Yannis Kalantidis. Decoupling representation and classifier for long-tailed recognition. *arXiv:1910.09217*, 2019.
- [26] Salman Khan, Munawar Hayat, Syed Waqas Zamir, Jianbing Shen, and Ling Shao. Striking the right balance with uncertainty. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [27] Salman H Khan, Munawar Hayat, Mohammed Bennamoun, Ferdous A Sohel, and Roberto Togneri. Cost-sensitive learning of deep feature representations from imbalanced data. *IEEE transactions on neural networks and learning systems*, 29(8):3573–3587, 2017.
- [28] Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. *arXiv preprint arXiv:1610.02242*, 2016.
- [29] Bruno Lecouat, Chuan-Sheng Foo, Houssam Zenati, and Vijay Chandrasekhar. Manifold regularization with gans for semi-supervised learning. *arXiv preprint arXiv:1807.04307*, 2018.
- [30] Dong-Hyun Lee. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *ICML Workshop on Challenges in Representation Learning*, 2013.
- [31] Hansang Lee, Minseok Park, and Junmo Kim. Plankton classification on imbalanced large scale database via convolutional neural networks with transfer learning. In *IEEE international conference on image processing (ICIP)*, pages 3713–3717, 2016.
- [32] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *ICCV*, pages 2980–2988, 2017.
- [33] Ziwei Liu, Zhongqi Miao, Xiaohang Zhan, Jiayun Wang, Boqing Gong, and Stella X Yu. Large-scale long-tailed recognition in an open world. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2537–2546, 2019.
- [34] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9:2579–2605, 2008.
- [35] Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE transactions on pattern analysis and machine intelligence*, 41(8):1979–1993, 2018.
- [36] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. *NIPS Workshop on Deep Learning and Unsupervised Feature Learning*, 2011.
- [37] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *European Conference on Computer Vision*, pages 69–84. Springer, 2016.
- [38] Mehdi Noroozi, Hamed Pirsiavash, and Paolo Favaro. Representation learning by learning to count. In *ICCV*, 2017.
- [39] Avital Oliver, Augustus Odena, Colin A Raffel, Ekin Dogus Cubuk, and Ian Goodfellow. Realistic evaluation of deep semi-supervised learning algorithms. In *NeurIPS*, 2018.
- [40] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv:1807.03748*, 2018.
- [41] Samira Pouyanfar, Yudong Tao, Anup Mohan, et al. Dynamic sampling in convolutional neural networks for imbalanced data classification. In *IEEE conference on multimedia information processing and retrieval (MIPR)*, 2018.
- [42] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.

- [43] Mehdi Sajjadi, Mehran Javanmardi, and Tolga Tasdizen. Regularization with stochastic transformations and perturbations for deep semi-supervised learning. In *Advances in Neural Information Processing Systems*, pages 1163–1171, 2016.
- [44] Li Shen, Zhouchen Lin, and Qingming Huang. Relay backpropagation for effective learning of deep convolutional neural networks. In *European conference on computer vision*, pages 467–482. Springer, 2016.
- [45] Jun Shu, Qi Xie, Lixuan Yi, Qian Zhao, Sanping Zhou, Zongben Xu, and Deyu Meng. Meta-weight-net: Learning an explicit mapping for sample weighting. *arXiv preprint arXiv:1902.07379*, 2019.
- [46] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *Advances in neural information processing systems*, pages 1195–1204, 2017.
- [47] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. *arXiv:1906.05849*, 2019.
- [48] Antonio Torralba, Rob Fergus, and William T Freeman. 80 million tiny images: A large data set for nonparametric object and scene recognition. *IEEE transactions on pattern analysis and machine intelligence*, 30(11):1958–1970, 2008.
- [49] Trieu H Trinh, Minh-Thang Luong, and Quoc V Le. Selfie: Self-supervised pretraining for image embedding. *arXiv preprint arXiv:1906.02940*, 2019.
- [50] Vikas Verma, Alex Lamb, Juho Kannala, Yoshua Bengio, and David Lopez-Paz. Interpolation consistency training for semi-supervised learning. *arXiv preprint arXiv:1903.03825*, 2019.
- [51] Martin J Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press, 2019.
- [52] Yu-Xiong Wang, Deva Ramanan, and Martial Hebert. Learning to model the tail. In *Advances in Neural Information Processing Systems*, pages 7029–7039, 2017.
- [53] Yuzhe Yang, Guo Zhang, Dina Katabi, and Zhi Xu. ME-Net: Towards effective adversarial robustness with matrix estimation. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, 2019.
- [54] Xi Yin, Xiang Yu, Kihyuk Sohn, Xiaoming Liu, and Manmohan Chandraker. Feature transfer learning for face recognition with under-represented data. In *In Proceeding of IEEE Computer Vision and Pattern Recognition*, Long Beach, CA, June 2019.
- [55] Chong You, Chi Li, Daniel P. Robinson, and René Vidal. A scalable exemplar-based subspace clustering algorithm for class-imbalanced data. In *ECCV*, 2018.
- [56] Xiaohua Zhai, Avital Oliver, Alexander Kolesnikov, and Lucas Beyer. S4l: Self-supervised semi-supervised learning. *arXiv preprint arXiv:1905.03670*, 2019.
- [57] Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful image colorization. In *European conference on computer vision*, pages 649–666. Springer, 2016.
- [58] Xiao Zhang, Zhiyuan Fang, Yandong Wen, Zhifeng Li, and Yu Qiao. Range loss for deep face recognition with long-tailed training data. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5409–5418, 2017.
- [59] Boyan Zhou, Quan Cui, Xiu-Shen Wei, and Zhao-Min Chen. Bbn: Bilateral-branch network with cumulative learning for long-tailed visual recognition. *arXiv:1912.02413*, 2019.

Appendices

A Proof of Theorem 1

Given the pseudo-labels, we note that we can rewrite $\tilde{X}_i^+ = (1 - I_i^+)(\mu_2 - \mu_1) + Z_i^+$, where $Z_i^+ \sim \mathcal{N}(\mu_1, \sigma^2)$ and $I_i^+ \sim \text{Bernoulli}(p)$. That is, if the pseudo-label is correct, then $\tilde{X}_i^+ \sim Z_i^+$ and otherwise, $\tilde{X}_i^+ \sim (\mu_2 - \mu_1) + Z_i^+$. Similarly, $\tilde{X}_i^- = (1 - I_i^-)(\mu_1 - \mu_2) + Z_i^-$, where $Z_i^- \sim \mathcal{N}(\mu_2, \sigma^2)$ and $I_i^- \sim \text{Bernoulli}(q)$. Now,

$$\begin{aligned} \hat{\theta} &= \frac{1}{2} \left(\frac{\sum_{i=1}^{\tilde{n}_+} \tilde{X}_i^+}{\tilde{n}_+} + \frac{\sum_{i=1}^{\tilde{n}_-} \tilde{X}_i^-}{\tilde{n}_-} \right) \\ &= \frac{1}{2} \left\{ \frac{\sum_{i=1}^{\tilde{n}_+} (1 - I_i^+)(\mu_2 - \mu_1) + Z_i^+}{\tilde{n}_+} + \frac{\sum_{i=1}^{\tilde{n}_-} (1 - I_i^-)(\mu_1 - \mu_2) + Z_i^-}{\tilde{n}_-} \right\} \\ &= \frac{1}{2} \left\{ \frac{\sum_{i=1}^{\tilde{n}_+} I_i^+(\mu_1 - \mu_2)}{\tilde{n}_+} + \frac{\sum_{i=1}^{\tilde{n}_-} I_i^-(\mu_2 - \mu_1)}{\tilde{n}_-} + \left(\frac{\sum_{i=1}^{\tilde{n}_+} Z_i^+}{\tilde{n}_+} + \frac{\sum_{i=1}^{\tilde{n}_-} Z_i^-}{\tilde{n}_-} \right) \right\}. \end{aligned} \quad (1)$$

We bound each term in Eq. (1) separately. First, note that

$$\left(\frac{\sum_{i=1}^{\tilde{n}_+} Z_i^+}{\tilde{n}_+} + \frac{\sum_{i=1}^{\tilde{n}_-} Z_i^-}{\tilde{n}_-} \right) \sim \mathcal{N} \left(\mu_1 + \mu_2, \sigma^2 \left(\frac{1}{\tilde{n}_+} + \frac{1}{\tilde{n}_-} \right) \right).$$

By standard Gaussian concentration, we have

$$\mathbb{P} \left(\left| \frac{\sum_{i=1}^{\tilde{n}_+} Z_i^+}{\tilde{n}_+} + \frac{\sum_{i=1}^{\tilde{n}_-} Z_i^-}{\tilde{n}_-} - (\mu_1 + \mu_2) \right| > t \right) \leq 2e^{-\frac{t^2}{2\sigma^2} \cdot \frac{1}{\frac{1}{\tilde{n}_+} + \frac{1}{\tilde{n}_-}}}. \quad (2)$$

Next, we bound the term $\frac{\sum_{i=1}^{\tilde{n}_+} I_i^+}{\tilde{n}_+}$. Since $I_i^+ \sim \text{Bernoulli}(p)$, applying Hoeffding inequality

$$\mathbb{P} \left(\left| \frac{\sum_{i=1}^{\tilde{n}_+} I_i^+}{\tilde{n}_+} - p \right| > t \right) \leq 2e^{-2\tilde{n}_+ t^2}. \quad (3)$$

Similarly, we have

$$\mathbb{P} \left(\left| \frac{\sum_{i=1}^{\tilde{n}_-} I_i^-}{\tilde{n}_-} - q \right| > t \right) \leq 2e^{-2\tilde{n}_- t^2}. \quad (4)$$

Note that by triangle inequality,

$$\begin{aligned} & \left| \hat{\theta} - \frac{\Delta(\mu_1 - \mu_2)}{2} - \frac{(\mu_1 + \mu_2)}{2} \right| \\ &= \left| \frac{1}{2} \left(\frac{\sum_{i=1}^{\tilde{n}_+} \tilde{X}_i^+}{\tilde{n}_+} + \frac{\sum_{i=1}^{\tilde{n}_-} \tilde{X}_i^-}{\tilde{n}_-} \right) - \frac{p(\mu_1 - \mu_2)}{2} - \frac{q(\mu_2 - \mu_1)}{2} - \frac{(\mu_1 + \mu_2)}{2} \right| \\ &\leq \left| \frac{1}{2} \frac{\sum_{i=1}^{\tilde{n}_+} I_i^+}{\tilde{n}_+} - \frac{p}{2} \right| |\mu_1 - \mu_2| + \left| \frac{1}{2} \frac{\sum_{i=1}^{\tilde{n}_-} I_i^-}{\tilde{n}_-} - \frac{q}{2} \right| |\mu_1 - \mu_2| \\ &\quad + \left| \frac{1}{2} \left(\frac{\sum_{i=1}^{\tilde{n}_+} Z_i^+}{\tilde{n}_+} + \frac{\sum_{i=1}^{\tilde{n}_-} Z_i^-}{\tilde{n}_-} \right) - \frac{\mu_1 + \mu_2}{2} \right|. \end{aligned}$$

Given $\delta > 0$, consider the event that

$$E = \left\{ \left| \frac{1}{2} \frac{\sum_{i=1}^{\tilde{n}_+} I_i^+}{\tilde{n}_+} - \frac{p}{2} \right| |\mu_1 - \mu_2| \leq \frac{\delta}{3} \text{ and } \left| \frac{1}{2} \frac{\sum_{i=1}^{\tilde{n}_-} I_i^-}{\tilde{n}_-} - \frac{q}{2} \right| |\mu_1 - \mu_2| \leq \frac{\delta}{3} \text{ and } \left| \frac{1}{2} \left(\frac{\sum_{i=1}^{\tilde{n}_+} Z_i^+}{\tilde{n}_+} + \frac{\sum_{i=1}^{\tilde{n}_-} Z_i^-}{\tilde{n}_-} \right) - \frac{\mu_1 + \mu_2}{2} \right| \leq \frac{\delta}{3} \right\}.$$

Using union bound and the concentration inequalities Eqs. (2), (3) and (4), we obtain the following lower bound on the probability of event E :

$$\mathbb{P}(E) \geq 1 - 2e^{-\frac{2\delta^2}{9\sigma^2} \cdot \frac{1}{\bar{n}_+ + \bar{n}_-}} - 2e^{-\frac{8\bar{n}_+ \delta^2}{9(\mu_1 - \mu_2)^2}} - 2e^{-\frac{8\bar{n}_- \delta^2}{9(\mu_1 - \mu_2)^2}}.$$

Finally, the above equation and the triangle inequality implies that

$$\mathbb{P}\left(\left|\hat{\theta} - \frac{\Delta(\mu_1 - \mu_2)}{2} - \frac{(\mu_1 + \mu_2)}{2}\right| \leq \delta\right) \geq 1 - 2e^{-\frac{2\delta^2}{9\sigma^2} \cdot \frac{1}{\bar{n}_+ + \bar{n}_-}} - 2e^{-\frac{8\bar{n}_+ \delta^2}{9(\mu_1 - \mu_2)^2}} - 2e^{-\frac{8\bar{n}_- \delta^2}{9(\mu_1 - \mu_2)^2}}.$$

This completes the proof.

B Proof of Theorem 2

Note that if $X \sim \mathcal{N}(0, \sigma^2 \mathbf{I}_d)$, then for a given θ , $\theta^T X \sim \mathcal{N}(0, \|\theta\|_2^2 \sigma^2)$. Based on the form of the linear classifier, we know that

$$\begin{aligned} \text{err}_f &= \mathbb{P}_{(X,Y) \sim P_{XY}}(y(\theta^T X + b) < 0) \\ &= \mathbb{P}_{(X,Y) \sim P_{XY}}(y(\theta^T X + b) < 0 | Y = +1) \mathbb{P}(Y = +1) \\ &\quad + \mathbb{P}_{(X,Y) \sim P_{XY}}(y(\theta^T X + b) < 0 | Y = -1) \mathbb{P}(Y = -1) \\ &= p_+ \mathbb{P}(\mathcal{N}(0, \|\theta\|_2^2 \sigma_1^2) < -b) + p_- \mathbb{P}(\mathcal{N}(0, \|\theta\|_2^2 \beta \sigma_1^2) > -b) \\ &= p_+ \mathbb{P}\left(\mathcal{N}(0, 1) < -\frac{b}{\|\theta\|_2 \sigma_1}\right) + p_- \mathbb{P}\left(\mathcal{N}(0, 1) < \frac{b}{\|\theta\|_2 \sqrt{\beta} \sigma_1}\right) \\ &= p_+ \Phi\left(-\frac{b}{\|\theta\|_2 \sigma_1}\right) + p_- \Phi\left(\frac{b}{\|\theta\|_2 \sqrt{\beta} \sigma_1}\right). \end{aligned}$$

Finally, note that for $b > 0$, $\Phi\left(\frac{b}{\|\theta\|_2 \sqrt{\beta} \sigma_1}\right) \geq 1/2$. Since p_- is assumed to be at least $1/2$, the error probability of the classifier is at least $1/4$.

C Proof of Theorem 3

We recall the following standard concentration inequality for sub-exponential random variables [51]. Suppose that W_1, W_2, \dots, W_n are i.i.d. sub-exponential random variables with parameters (ν, α) . Then,

$$\mathbb{P}\left(\frac{1}{n} \sum_{i=1}^n (W_i - E[X_i]) \geq \delta\right) \leq \begin{cases} e^{-\frac{n\delta^2}{2\nu^2}}, & \text{for } 0 \leq \delta \leq \frac{\nu^2}{\alpha}; \\ e^{-\frac{n\delta}{2\alpha}}, & \text{for } \delta > \frac{\nu^2}{\alpha}. \end{cases}$$

We remark that a similar two-sided tail bounds also hold. For our purpose, note that if $W \sim \mathcal{N}(0, 1)$, then W^2 is sub-exponential with parameter $(2, 4)$. Therefore, we have the standard χ^2 -concentration for a χ^2 random variable with degree n as follows:

$$\mathbb{P}\left(\left|\frac{1}{n} \sum_{i=1}^n W_i^2 - 1\right| \geq \delta\right) \leq 2e^{-n\delta^2/8}, \quad \forall \delta \in (0, 1).$$

Let us now analyze the classifier $f_{ss}(X) = \text{sign}(-Z + b)$ obtained by self-supervised training. Recall that we assume $Z = \psi(X) = k_1 \|X\|_2^2 + k_2$. For the negative training data, we note that $\frac{Z_i - k_2}{k_1 \beta \sigma_1^2} \sim \chi_d^2$ (the χ^2 distribution with d degree of freedom). Using the previous χ^2 concentration, we have that for $\delta \in (0, 1)$,

$$\mathbb{P}\left(\left|\frac{1}{N-d} \sum_{i=1}^N \mathbf{1}_{\{Y_i = -1\}} \frac{Z_i - k_2}{k_1 \beta \sigma_1^2} - 1\right| \geq \delta\right) \leq 2e^{-N-d\delta^2/8}.$$

Rearrange the terms, we obtain

$$\mathbb{P}\left(\left|\frac{1}{N_-}\sum_{i=1}^N \mathbf{1}_{\{Y_i=-1\}}Z_i - (dk_1\beta\sigma_1^2 + k_2)\right| \geq \delta dk_1\beta\sigma_1^2\right) \leq 2e^{-N_-d\delta^2/8}.$$

Similarly, for the positive training data, we have

$$\mathbb{P}\left(\left|\frac{1}{N_+}\sum_{i=1}^N \mathbf{1}_{\{Y_i=+1\}}Z_i - (dk_1\sigma_1^2 + k_2)\right| \geq \delta dk_1\sigma_1^2\right) \leq 2e^{-N_+d\delta^2/8}.$$

Therefore, by an application of union bound, with probability at least $1 - 2e^{-N_-d\delta^2/8} - 2e^{-N_+d\delta^2/8}$, the intercept b of the self-supervised classifier, i.e., $b = \frac{1}{2}\left(\frac{\sum_{i=1}^N \mathbf{1}_{\{Y_i=+1\}}Z_i}{N_+} + \frac{\sum_{i=1}^N \mathbf{1}_{\{Y_i=-1\}}Z_i}{N_-}\right) > 0$, satisfies

$$\left|b - \frac{dk_1(\beta+1)\sigma_1^2}{2} - k_2\right| \leq \frac{\delta dk_1(\beta+1)\sigma_1^2}{2}. \quad (5)$$

In the following, we condition on the event that b satisfies the bound in Eq. (5). Then,

$$\begin{aligned} \text{err}_{f_{ss}} &= \mathbb{P}_{(X,Y)\sim P_{XY}}(y(-Z+b) < 0) \\ &= \mathbb{P}_{(X,Y)\sim P_{XY}}(y(-Z+b) < 0|Y=+1)\mathbb{P}(Y=+1) \\ &\quad + \mathbb{P}_{(X,Y)\sim P_{XY}}(y(-Z+b) < 0|Y=-1)\mathbb{P}(Y=-1) \\ &= p_+\mathbb{P}_{(X,Y)\sim P_{XY}}(Z > b|Y=+1) + p_-\mathbb{P}_{(X,Y)\sim P_{XY}}(Z < b|Y=-1) \\ &\leq p_+\mathbb{P}_{(X,Y)\sim P_{XY}}\left(Z > \frac{dk_1(\beta+1)\sigma_1^2}{2} + k_2 - \frac{\delta dk_1(\beta+1)\sigma_1^2}{2} \middle| Y=+1\right) \\ &\quad + p_-\mathbb{P}_{(X,Y)\sim P_{XY}}\left(Z < \frac{dk_1(\beta+1)\sigma_1^2}{2} + k_2 + \frac{\delta dk_1(\beta+1)\sigma_1^2}{2} \middle| Y=-1\right). \end{aligned} \quad (6)$$

We analyze the two terms in the bound Eq. (6). First, condition on $Y = -1$, again note that $\frac{Z-k_2}{k_1\beta\sigma_1^2} \sim \chi_d^2$. Therefore,

$$\begin{aligned} &\mathbb{P}\left(Z < \frac{dk_1(\beta+1)\sigma_1^2}{2} + k_2 + \frac{\delta dk_1(\beta+1)\sigma_1^2}{2} \middle| Y=-1\right) \\ &= \mathbb{P}\left(Z - (dk_1\beta\sigma_1^2 + k_2) < -\frac{(\beta-1-\delta(\beta+1))}{2\beta}dk_1\beta\sigma_1^2 \middle| Y=-1\right) \\ &\leq \exp\left(-d \cdot \frac{(\beta-1-\delta(\beta+1))^2}{32\beta^2}\right), \end{aligned}$$

where the last inequality follows from the concentration inequality and note that $0 < (\beta-1-\delta(\beta+1))/2\beta < 1$. In a similar manner, for $Y = +1$, note that $\frac{Z-k_2}{k_1\sigma_1^2} \sim \chi_d^2$. Hence,

$$\begin{aligned} &\mathbb{P}\left(Z > \frac{dk_1(\beta+1)\sigma_1^2}{2} + k_2 - \frac{\delta dk_1(\beta+1)\sigma_1^2}{2} \middle| Y=+1\right) \\ &= \mathbb{P}\left(Z - (dk_1\sigma_1^2 + k_2) > \frac{(\beta-1-\delta(\beta+1))}{2}dk_1\sigma_1^2 \middle| Y=+1\right) \\ &\leq \begin{cases} \exp\left(-d \cdot \frac{(\beta-1-\delta(\beta+1))^2}{32}\right), & \text{if } \delta \in \left[\frac{\beta-3}{\beta+1}, \frac{\beta-1}{\beta+1}\right]; \\ \exp\left(-d \cdot \frac{(\beta-1-\delta(\beta+1))}{16}\right), & \text{if } \delta \in \left(0, \frac{\beta-3}{\beta+1}\right). \end{cases} \end{aligned}$$

For the last inequality, we note that if $\delta \in \left[\frac{\beta-3}{\beta+1}, \frac{\beta-1}{\beta+1}\right]$, $0 < \frac{(\beta-1-\delta(\beta+1))}{2} \leq 1$; otherwise, $\frac{(\beta-1-\delta(\beta+1))}{2} > 1$. Substituting the above inequalities into the error probability $\text{err}_{f_{ss}}$ (Eq. (6)) completes the proof.

D Experimental Details

D.1 Imbalanced Dataset Details

In this section, we provide the detailed information of five long-tailed imbalanced datasets we use in our experiments. Table 5 provides an overview of the five datasets.

Table 5: Overview of the five imbalanced datasets used in our experiments. ρ indicates the imbalance ratio.

Dataset	# Class	ρ	Head class size	Tail class size	# Training set	# Val. set	# Test set
CIFAR-10-LT	10	10 ~ 100	5,000	500 ~ 50	20,431 ~ 12,406	–	10,000
CIFAR-100-LT	100	10 ~ 100	500	50 ~ 5	19,573 ~ 10,847	–	10,000
SVHN-LT	10	10 ~ 100	1,000	100 ~ 10	4,084 ~ 2,478	–	26,032
ImageNet-LT	1,000	256	1,280	5	115,846	20,000	50,000
iNaturalist 2018	8,142	500	1,000	2	437,513	24,426	–

CIFAR-10-LT and CIFAR-100-LT. The original versions of CIFAR-10 and CIFAR-100 contain 50,000 images for training and 10,000 images for testing with class number of 10 and 100, respectively. We create the long-tailed CIFAR versions following [7, 11] with controllable degrees of data imbalance, and keep the test set unchanged. We vary the class imbalance ratio ρ from 10 to 100.

SVHN-LT. The original SVHN dataset contains 73,257 images for training and 26,032 images for testing with 10 classes. Similarly to CIFAR-LT, we create SVHN-LT dataset with maximally 1,000 images per class (head class), and vary ρ from 10 to 100 for different long-tailed versions.

ImageNet-LT. ImageNet-LT [33] is artificially truncated from its balanced version, with sample size in the training set following an exponential decay across different classes. ImageNet-LT has 1,000 classes and 115.8K training images, with number of images ranging from 1,280 to 5 images per class.

iNaturalist 2018. iNaturalist 2018 [24] is a real-world fine-grained visual recognition dataset that naturally exhibits long-tailed class distributions, consisting of 435,713 samples from 8,142 species.

D.2 Unlabeled Data Details

We provide additional information on the unlabeled data we use in the semi-supervised settings, i.e., the unlabeled data sourcing, and how we create unlabeled sets with different imbalanced distributions.

Unlabeled Data Sourcing. To obtain the unlabeled data needed for our semi-supervised setup, we follow [8] to mine the 80 Million Tiny Images (80M) dataset [48] to source unlabeled and uncuration data for CIFAR-10. In particular, CIFAR-10 is a human-labeled subset of 80M, which is manually restricted to 10 classes. Accordingly, most images in 80M do not belong to any image categories in CIFAR-10. To select unlabeled data that exhibit similar distributions as labeled ones, we follow the procedure as in [8], where an 11-class classification model is trained to distinguish CIFAR-10 classes and a “non-CIFAR” class. For each class, we then rank the images based on the prediction confidence, and construct the unlabeled (imbalanced) dataset \mathcal{D}_U according to our settings.

For SVHN, since its own dataset contains an extra part [36] with 531.1K additional (labeled) samples, we directly use these additional data to simulate the unlabeled dataset, which exhibits similar data distribution as the main dataset. Specifically, the ground truth labels are used only for preparing \mathcal{D}_U , and are abandoned throughout experiments (i.e., before performing pseudo-labeling).

Relevant (and Irrelevant) Unlabeled Data. To analyze the data relevance of unlabeled data with class imbalance (cf. Sec. 3), we again employ the 11-way classifier to select samples with prediction scores that are high for the extra class, and use them as proxy for irrelevant data. We then mix the irrelevant dataset and our main unlabeled dataset with different proportions, thus creating a sequence of unlabeled datasets with different degrees of data relevance.

Unlabeled Data with Class Imbalance. With the sourced unlabeled data, we construct the demanded unlabeled dataset \mathcal{D}_U also with class imbalance. In Fig. 5, we show an example of data distributions of both original labeled imbalanced dataset \mathcal{D}_L and \mathcal{D}_U with different unlabeled imbalance ratio. Specifically, Fig. 5a presents the training and test set of CIFAR-10-LT with $\rho = 100$, where a long




Figure 5: An illustration of labeled dataset (\mathcal{D}_L) as well as its corresponding unlabeled dataset ($\mathcal{D}_U@5x$) under different unlabeled imbalance ratio ρ_U . Given \mathcal{D}_L with a fixed ρ , the total amount of \mathcal{D}_U is fixed, while different ρ_U will lead to different class distributions of the unlabeled data.

tail can be observed for the training set, while the test set is balanced across classes. With labeled data on hand, we create different degrees of class imbalance in \mathcal{D}_U to be (1) uniform ($\rho_U = 1$, Fig. 5b), (2) half imbalanced as labeled set ($\rho_U = \rho/2$, Fig. 5c), (3) same imbalanced ($\rho_U = \rho$, Fig. 5d), and (4) double imbalanced ($\rho_U = 2\rho$, Fig. 5e). Note that the total data amount of \mathcal{D}_U is fixed (e.g., 5x as labeled set) given \mathcal{D}_L , while different ρ_U will result in different unlabeled data distributions.

D.3 Implementation Details

CIFAR-10-LT and CIFAR-100-LT. Following [1, 7, 11], we use ResNet-32 [20] for all CIFAR-LT experiments. The data augmentation follows [20] to use zero-padding with 4 pixels on each side and then random crop back to the original image size, after which a random horizontal flip is performed. We train all models for 200 epochs, and remain all other hyper-parameters the same as [7]. In the semi-supervised settings, we fix the unlabeled weight $\omega = 1$ for all experiments.

SVHN-LT. Similarly to CIFAR-LT, we use ResNet-32 model for all SVHN-LT experiments, and fix the same hyper-parameters as in CIFAR-LT experiments throughout training.

ImageNet-LT. We follow [25, 33] to report results with ResNet-10 and ResNet-50 models. Since [33] only employs ResNet-10 model, we reproduce the results with ResNet-50 using the public code from the authors for fair comparison. During the classifier training stage, we train all models for 90 epochs, and keep all other hyper-parameters identical to those in [25]. During the self-supervised pre-training stage, we leave the hyper-parameters unchanged as in [19], but only use samples from ImageNet-LT.

iNaturalist 2018. We follow [1, 7, 11, 25] to use ResNet-50 model. Similar to ImageNet-LT, we train all models for 90 epochs in the classifier training stage, and other hyper-parameters are kept the same as in [25]. The self-supervised pre-training stage is remained the same as that on ImageNet-LT. We reproduce the results for [7] on iNaturalist 2018 using the authors’ code.

E Additional Results for Semi-Supervised Imbalanced Learning

E.1 Different Semi-Supervised Learning Methods

We study the effect of different advanced semi-supervised learning methods, in addition to the simple pseudo-label strategy we apply in the main text. We select the following two methods for analysis.

Virtual Adversarial Training. Virtual adversarial training (VAT) [35] is one of the state-of-the-art semi-supervised learning methods, which aims to make the predicted labels robust around input data point against local perturbation. It approximates a tiny perturbation ϵ_{adv} to add to the (unlabeled) inputs which would most significantly affect the outputs of the model. Note that the implementation difference between VAT and the pseudo-label is the loss term on the unlabeled data, where VAT exhibits a consistency regularization loss rather than supervised loss, resulting in a loss function as $\mathcal{L}(\mathcal{D}_L, \theta) + \omega \mathcal{L}_{con}(\mathcal{D}_U, \theta)$. We add an additional entropy regularization term following [35].

Mean Teacher. The mean teacher (MT) [46] method is also a representative algorithm using consistency regularization, where a teacher model and a student model are maintained and a consistency cost between the student’s and the teacher’s outputs is introduced. The teacher weights are updated through an exponential moving average (EMA) of the student weights.

Similar to pseudo-label, the two semi-supervised methods can be seamlessly incorporated with our imbalanced learning framework. We present the results with these methods in Table 6. For each run, we construct $\mathcal{D}_U@5x$ with the same imbalance ratio as the labeled set (i.e., $\rho_U = \rho$). As

Table 6: Ablation study of different semi-supervised learning methods on CIFAR-10-LT and SVHN-LT. We fix $\rho_U = \rho$ for each specific setting. Best results of each column are in **bold** and the second best are underlined.

Dataset	CIFAR-10-LT			SVHN-LT		
Imbalance Ratio (ρ)	100	50	10	100	50	10
Vanilla CE	29.64	25.19	13.61	19.98	17.50	11.46
$\mathcal{D}_U@5x$ + Pseudo-label [30]	18.74	18.36	10.86	14.65	13.16	10.06
$\mathcal{D}_U@5x$ + VAT [35]	<u>17.93</u>	<u>16.53</u>	9.44	<u>13.07</u>	<u>12.27</u>	<u>9.29</u>
$\mathcal{D}_U@5x$ + MT [46]	16.52	15.79	<u>9.53</u>	12.34	11.12	8.62

Table 6 reports, across different datasets and imbalance ratios, adding unlabeled data can consistently benefit imbalanced learning via semi-supervised learning. Moreover, by using more advanced SSL techniques, larger improvements can be obtained in general.

E.2 Class-wise Generalization Results

In the main paper, we report the top-1 test errors as the final performance metric. To gain additional insights on how unlabeled data helps imbalanced tasks, we further look at the generalization results in each class, especially on the minority (tail) classes.

Generalization on Minority Classes. In Fig. 6 we plot the test error on each class on CIFAR-10-LT and SVHN-LT with $\rho = 50$. As the figure shows, regardless of the base training technique, using unlabeled data can consistently and substantially improve the generalization on tail classes.




Figure 6: Class-wise top-1 error rates. C0 stands for the head class, and C9 stands for the tail class. Using unlabeled data leads to better generalization on tail classes while keeping the performance on head classes almost unaffected, and can consistently boost different training techniques. Results are averaged across 5 runs.

Confusion Matrix. We further show the confusion matrices on CIFAR-10-LT with and without \mathcal{D}_U . Fig. 7 presents the results, where for the vanilla CE training, predictions for tail classes are biased towards the head classes significantly. In contrast, by using unlabeled data, the leakage from tail classes to head classes can be largely eliminated.

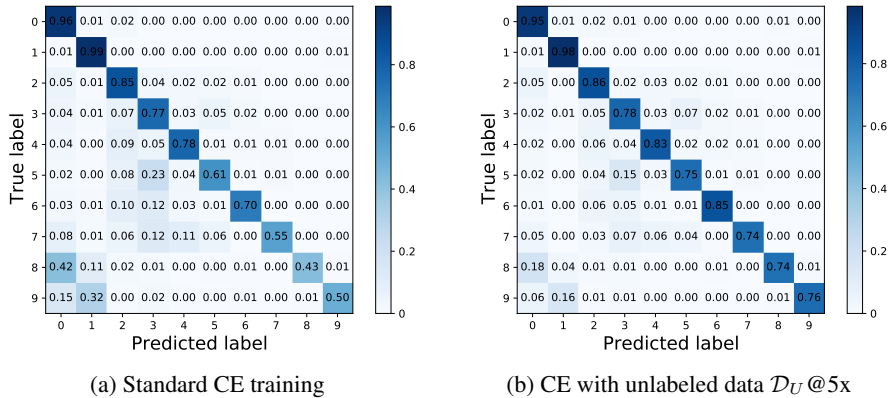


Figure 7: Confusion matrices of standard CE training and using $\mathcal{D}_U@5x$ on CIFAR-10-LT with $\rho = 100$.

E.3 Effect of Unlabeled Data Amount

We study the effect of the size of the unlabeled dataset on our SSL approach in imbalanced learning. We first fix the labeled dataset \mathcal{D}_L with $\rho = 50$, the unlabeled imbalance ratio to be $\rho_U = \rho$, and then vary the amount of \mathcal{D}_U to be $\{0.5x, 1x, 5x, 10x\}$ of the size of \mathcal{D}_L . Table 7 reports the results, where we can observe that larger \mathcal{D}_U consistently leads to higher gains. Furthermore, even with only 0.5x more unlabeled data, the performance can be boosted largely compared to that without unlabeled data. Interestingly however, as the size of \mathcal{D}_U becomes larger, the gains gradually diminish.

Table 7: Ablation study of how unlabeled data amount affects SSL in imbalanced learning. We fix the imbalance ratios as $\rho = \rho_U = 50$. We vary the amount of \mathcal{D}_U with respect to labeled data amount (e.g., 0.5x means the size of \mathcal{D}_U is half of \mathcal{D}_L). Best results of each part are in **bold** and the second best are underlined.

Dataset	CIFAR-10-LT				SVHN-LT			
\mathcal{D}_U Size (w.r.t. \mathcal{D}_L)	0.5x	1x	5x	10x	0.5x	1x	5x	10x
CE	25.19				17.50			
CE + \mathcal{D}_U	21.75	20.35	<u>18.36</u>	16.88	14.96	14.13	<u>13.16</u>	13.02
LDAM-DRW [7]	19.06				14.59			
LDAM-DRW + \mathcal{D}_U	17.43	16.59	<u>14.93</u>	13.91	13.93	13.07	<u>11.26</u>	11.09

E.4 Effect of Labeled Data Amount

Following [39], we further study how the labeled data amount affects SSL in imbalanced learning. We fix the imbalance ratios of \mathcal{D}_L and \mathcal{D}_U as $\rho = \rho_U = 50$, and also fix the size of \mathcal{D}_U to be 5x of \mathcal{D}_L . We vary \mathcal{D}_L amount to be $\{0.5x, 0.75x, 1x\}$ with respect to the original labeled data amount. As Table 8 shows, with smaller size of labeled data, the test errors of vanilla CE training increases largely, while adding unlabeled data can maintain sufficiently low errors. Interestingly, when unlabeled data is added, using only 50% of labeled data can already surpass the fully-supervised baseline on both datasets, demonstrating the power of unlabeled data in the context of imbalanced learning.

Table 8: Ablation study of how labeled data amount affects SSL in imbalanced learning. We fix the imbalance ratios as $\rho = \rho_U = 50$, and fix unlabeled data amount to be 5x of labeled data used. We vary the amount of \mathcal{D}_L with respect to their original labeled data amount (e.g., 0.5x means only half of the initial labeled data is used).

Dataset	CIFAR-10-LT			SVHN-LT		
\mathcal{D}_L Size	0.5x	0.75x	1x	0.5x	0.75x	1x
CE	33.35	28.65	25.19	23.19	19.73	17.50
CE + $\mathcal{D}_U@5x$	20.77	<u>18.67</u>	18.36	14.80	<u>13.51</u>	13.16

F Additional Results for Self-Supervised Imbalanced Learning

F.1 Different Self-Supervised Pre-Training Methods

In this section, we investigate the effect of different SSP methods on imbalanced learning tasks. We select four different SSP approaches, ranging from pretext tasks to recent contrastive methods.

Solving Jigsaw Puzzles. Jigsaw [37] is a classical method based on pretext tasks, where an image is divided into patches, and a classifier is trained to predict the correct permutation of these patches.

Rotation Prediction. Predicting rotation [16] is another simple yet effective method, where an image is rotated by a random multiple of 90 degrees, constructing a 4-way classification problem; a classifier is then trained to determine the degree of rotation applied to an input image.

Selfie. Selfie [49] works by masking out select patches in an image, and then constructs a classification problem to determine the correct patch to be filled in the masked location.

Momentum Contrast. The momentum contrast (MoCo) [19] method is one of the recently proposed contrastive techniques, where contrastive losses [19] are applied in a representation space to measure the similarities of positive and negative sample pairs, and a momentum-updated encoder is employed.

We conduct controlled experiments over four benchmark imbalanced datasets, and report the results in Table 9. As the table reveals, all self-supervised pre-training methods can benefit the imbalanced learning, consistently across different datasets. Interestingly however, the performance gain varies across SSP techniques. Specifically, on datasets with smaller scale, i.e., CIFAR-10-LT and CIFAR-100-LT, methods using pretext tasks are generally better than using contrastive learning, with *Rotation* performs the best. In contrast, on larger datasets, i.e., ImageNet-LT and iNaturalist, *MoCo* outperforms other SSP methods by a notable margin. We hypothesize that since MoCo needs a large number of (negative) samples to be effective, the smaller yet imbalanced datasets thus may not benefit much from MoCo, compared to those with larger size and more samples.

Table 9: Ablation study of different self-supervised pre-training methods. We set imbalance ratio of $\rho = 50$ for CIFAR-LT. Best results of each column are in **bold** and the second best are underlined.

Dataset	CIFAR-10-LT	CIFAR-100-LT	ImageNet-LT	iNaturalist 2018
Vanilla CE	25.19	56.15	61.6	39.3
+ Jigsaw [37]	24.68	55.89	60.2	38.2
+ Selfie [49]	<u>22.75</u>	<u>55.31</u>	58.3	36.9
+ Rotation [16]	21.80	54.96	<u>55.7</u>	<u>36.5</u>
+ MoCo [19]	24.18	55.83	54.4	35.6

F.2 Class-wise Generalization Results

Similar to the semi-supervised setting, we again take a closer look at generalization on each class to gain further insights, in addition to the top-1 test error rates reported in the main text.

Generalization on Minority Classes. We plot the class-wise top-1 errors on CIFAR-10-LT (Fig. 8a) and ImageNet-LT (Fig. 8b), respectively. For ImageNet-LT, we follow [25, 33] to split the test set into three subsets for evaluating shot-wise accuracy, namely *Many-shot* (classes with more than 100 images), *Medium-shot* (20 ~ 100 images), and *Few-shot* (less than 20 images). On both datasets, we can observe that regardless of training techniques for the base classifier, using SSP can consistently and substantially improve the generalization on tail (few-shot) classes, while maintaining or slightly improving the performance on head (many-shot) classes. The consistent gains demonstrate the effectiveness of SSP in the context of imbalanced learning, especially for the tail classes.




Figure 8: Class-wise top-1 error rates on CIFAR-10-LT and ImageNet-LT. C0 stands for the head class, and C9 stands for the tail class. On ImageNet-LT we follow [33] to report test error on three splits of the set of classes: Many-shot, Medium-shot, and Few-shot. Using SSP leads to better generalization on both head and tail classes, and can consistently boost different training techniques. Results are averaged across 5 runs.

Confusion Matrix. To further understand how self-supervision helps imbalanced learning, we again plot the confusion matrices on CIFAR-10-LT with and without SSP, respectively. As illustrated in Fig. 9, the prediction results of vanilla CE training suffers from the large leakage from tail classes to head classes, leading to low accuracy on minority categories. In contrast, by using self-supervised pre-training, the tail-to-head class leakages are greatly compensated, resulting in better performance and consistent improvements across the tail classes.




Figure 9: Confusion matrices of standard CE training and using SSP on CIFAR-10-LT with $\rho = 100$.

F.3 Effect of Imbalance Type

Finally, we conduct ablation study on another type of imbalance. The majority of the literature [1, 11, 25, 33, 59] focused on the long-tailed imbalance distribution, which is also the typical scenario for large-scale real-world datasets [24, 33]. Yet, few other manually designed imbalance types are also investigated by researchers [5, 7] to provide a comprehensive picture. For completeness, we study the performance of SSP under another imbalance type, i.e., the step imbalance [5], where the training instances of half of the classes (i.e., the minority classes) are reduced to a fixed size. The minority classes are defined to have the same sample size, and so do all frequent classes. The imbalance ratio ρ is the same as in long-tailed setting, i.e., $\rho = \max_i\{n_i\} / \min_i\{n_i\}$.

Table 10 presents the results, where we confirm that SSP can bring in consistent benefits across different imbalanced learning techniques on various datasets. Furthermore, when the dataset is more imbalanced (i.e., with higher ρ), the performance gains from SSP tend to be even larger, demonstrating the value of self-supervision under extreme class imbalance.

Table 10: Top-1 test errors (%) of ResNet-32 on CIFAR-10 and CIFAR-100 with step imbalance [5]. Using SSP, we can consistently and substantially improve different imbalanced learning techniques across various datasets, and achieve the best performance. [†] denotes results that reported in [7].

Dataset	Imbalanced CIFAR-10		Imbalanced CIFAR-100	
Imbalance Ratio (ρ)	100	10	100	10
CE	36.70	17.50	61.45	45.37
CB-CE [11] [†]	38.06	16.20	78.69	47.52
CE + SSP	27.27	12.04	55.57	42.90
Focal [32]	36.09	16.36	61.43	46.54
CB-Focal [11] [†]	39.73	16.54	80.24	49.98
Focal + SSP	27.00	12.07	55.12	42.93
LDAM [7] [†]	33.42	15.00	60.42	43.73
LDAM-DRW [7] [†]	23.08	12.19	54.64	40.54
LDAM-DRW + SSP	22.95	11.83	54.28	40.33