Pruning the Unlabeled Data to Improve Semi-Supervised Learning

Guy Hacohen,[1,2] **Daphna Weinshall**[1]

[1] School of Computer Science and Engineering, The Hebrew University of Jerusalem, Israel
[2] Edmond & Lily Safra Center for Brain Sciences, The Hebrew University of Jerusalem, Israel
{guy.hacohen,daphna@mail.huji.ac.il}

Abstract

In the domain of semi-supervised learning (SSL), the conventional approach involves training a learner with a limited amount of labeled data alongside a substantial volume of unlabeled data, both drawn from the same underlying distribution. However, for deep learning models, this standard practice may not yield optimal results. In this research, we propose an alternative perspective, suggesting that distributions that are more readily separable could offer superior benefits to the learner as compared to the original distribution. To achieve this, we present *PruneSSL*, a practical technique for selectively removing examples from the original unlabeled dataset to enhance its separability. We present an empirical study, showing that although *PruneSSL* reduces the quantity of available training data for the learner, it significantly improves the performance of various competitive SSL algorithms, thereby achieving state-of-the-art results across several image classification tasks.

1 Introduction

In recent years, extensive research has centered around the domain of deep semi-supervised learning (SSL), showcasing remarkable effectiveness across various domains. This success largely arises from the abundance of unlabeled data, in contrast with labeled data whose collection usually demands costly human annotation. Accordingly, most of the previous works focused on the creation and improvement of optimization algorithms that can utilize both labeled and unlabeled data. Differently, we focus on improving SSL performance by manipulating the unlabeled data directly.

Traditionally, SSL paradigms assume that both labeled and unlabeled data stem from the same underlying distribution. This is a reasonable assumption – when there is a significant disparity between the unlabeled and labeled data, the learning model may learn the wrong dependencies across these datasets. Consequently, this misalignment can lead to errors in generalization, ultimately mitigating the advantages that could otherwise be gained from incorporating unlabeled data.

Our investigation indicates that reliance on this assumption may not be optimal in the context of SSL. In particular, we demonstrate that modifying the distribution of unlabeled examples to enhance their distinctiveness boosts the performances of numerous deep SSL algorithms. With this objective, we present *PruneSSL*, a systematic approach designed to identify and prune unlabeled instances that undermine separability within the unlabeled dataset, thereby boosting overall performance.

More specifically, instances are considered to undermine the separability of the dataset if, given some meaningful embedding of the dataset and its corresponding labeling function, simple classifiers fail to accurately classify these instances when trained on the complete dataset. Clearly, executing this process directly on the unlabeled dataset is unfeasible due to the absence of labels. Consequently, the challenge lies in devising a technique capable of identifying these instances without relying on a labeling function.

Our proposed method, called *PruneSSL*, starts by employing a deep representation task over the unlabeled data, resulting in a meaningful embedding space. Subsequently, it calculates pseudo-labels for the unlabeled data using a deep-clustering algorithm or k-means within the embedding space. Equipped with both the embedding space and the pseudo labels, *PruneSSL* trains a simple classifier on the data and prunes instances that elicit the lowest levels of confidence from the classifier.

To illustrate this concept, we present a step-by-step visualization of each stage of *PruneSSL* in Fig. 1, utilizing a binary subset extracted from CIFAR-10. The figure shows a comparison between random instance pruning from the unlabeled dataset and the *PruneSSL* approach. For visualization purposes, the data is projected into a 2-dimensional space using t-SNE. Note that this projection serves exclusively for visualization purposes, as all computations occur within the original embedding space. Evidently, even in the projected space, the resulting data exhibits increased separation, as the 2 clusters in this case are more distinct. The exact details of this experiment are described in Section 3.

The inherent modularity of *PruneSSL* presents clear advantages as well as potential drawbacks. On the positive side, this modularity facilitates seamless integration of upcoming developments in both self-representation and pseudo-labeling tasks. Moreover, the framework readily adapts to diverse domain-specific challenges by tailoring the representation task accordingly. However, a notable drawback of *PruneSSL* lies in its reliance on the effectiveness of both the self-representation task and the pseudo-labeling process.

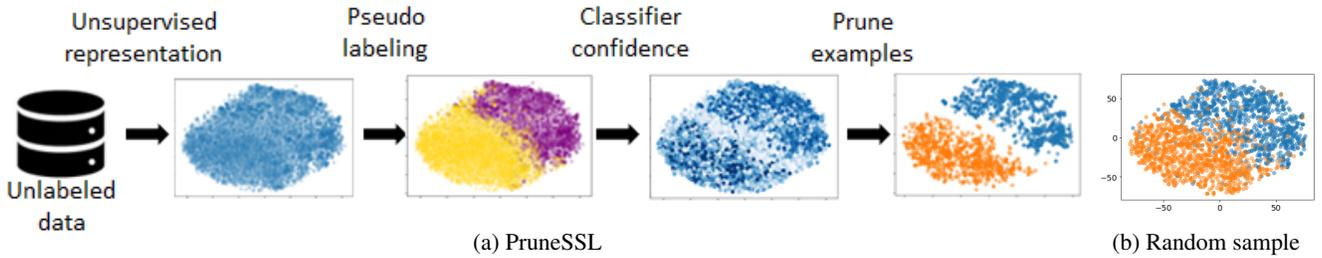In Section 3 we present the outcome of our empirical investigation, demonstrating the effectiveness of *PruneSSL*.

Figure 1: Flowchart of *PruneSSL*. In (a), we provide a step-by-step visual representation of how *PruneSSL* operates on the binary CIFAR-10 dataset. To aid visualization, each step is projected onto a 2-dimensional space using t-SNE. *PruneSSL* starts with an unsupervised representation task, followed by pseudo-labeling denoted in yellow and purple. Subsequently, a simple classifier is trained on the data, assigning confidence levels to each example – darker shades of blue indicate higher confidence. Finally, examples with low confidence are pruned from the data, with the orange-blue color scheme corresponding to actual labels. In contrast, (b) displays pruning using a random function. Evidently, *PruneSSL* results in a distribution that is more separable.

Despite its reduction of the unlabeled dataset's size, we find that a diverse array of SSL algorithms gain significant performance boosts when trained with the pruned unlabeled set, compared to training with the entire unlabeled dataset. Notably, these advantages are even more pronounced when replacing the pseudo-labeling process with an oracle possessing knowledge of the actual unlabeled dataset labels. These results demonstrate the validity of the idea behind *PruneSSL*.

In Section 3.1, we describe the results of an ablation study, designed to assess the significance of each component within the *PruneSSL* algorithm, as well as the varying influence of several hyperparameters. Our findings reveal a consistent trend: *PruneSSL* consistently produces comparable qualitative outcomes across a diverse range of embedding spaces and pseudo-labeling techniques. Additionally, our analysis highlights an interesting discovery — while an optimal value exists for the number of instances to be pruned, this parameter displays robustness, enabling the use of an extensive array of values with comparable effectiveness. Finally, we report that the benefits of *PruneSSL* increase as the size of the labeled set decreases.

Our approach bears a close connection to the fundamental cluster assumption of SSL (Chapelle et al., 2006). This assumption states that in order for any semi-supervised learning framework to work, even outside the realm of deep learning, the algorithm must assume that the data have inherent cluster structure, and thus, instances falling into the same cluster have the same class label. A direct result of this assumption is that the decision boundary of SSL algorithms should avoid intersecting high-density regions of the data. Instead, it should reside within low-density regions, thereby preserving the data's cluster structure. By enhancing the separability of unlabeled data, our method guides various SSL algorithms away from solutions that might traverse high-density regions. This ultimately enhances the performance of deep SSL techniques in a broader context.

Related work

In classical machine learning, training with unlabeled data could potentially lead to performance degradation Chawla

and Karakoulas (2005); Yang and Priebe (2011); Li and Zhou (2014). This phenomenon was primarily attributed to the manifold and cluster assumptions: unlabeled data is expected to be helpful only if it lies on a low-dimension manifold, and if similar classes are clustered together Chapelle et al. (2006); Singh et al. (2008). In contrast, the use of unlabeled data in deep learning is generally regarded as beneficial across most scenarios (Yang et al., 2022). This divergence could be attributed to deep learning's exceptional capacity (Johnson et al., 2016) to map data into spaces where both the manifold and cluster assumptions hold. This paper demonstrates the mutual benefits that can be harnessed between classical machine learning insights and deep learning. While unlabeled data generally helps deep learning, forcing it to better uphold the cluster assumption can further help the performance of deep models.

Our study diverges from previous art by focusing on altering the unlabeled data directly, rather than modifying the algorithms that use it. While numerous works draw inspiration from the cluster assumption, they often drive the separation boundary to reside in less dense regions of the unlabeled data via algorithmic optimizations (Chapelle and Zien, 2005; Ruiz et al., 2010; Verma et al., 2022). In contrast, our approach modifies the unlabeled data itself, implicitly encouraging SSL algorithms to adhere to the cluster assumption more closely. It's important to highlight that many methods assume a shared distribution between labeled and unlabeled sets (Ouali et al., 2020; Oliver et al., 2018; Berthelot et al., 2019; Li et al., 2020). Therefore, our method can be used in combination with these other methods, to boost their performance.

Recently, the recognition that specific instances within the unlabeled dataset can detrimentally impact the learning process has gained prominence, influencing several semi-supervised learning (SSL) strategies (Ren et al., 2020). For instance, FixMatch (Sohn et al., 2020) introduced the concept of integrating an unsupervised loss relying on pseudo-labels assigned exclusively to unlabeled instances demonstrating high model confidence. This approach effectively steers the model to learn from a selective set of examples in the unlabeled set, which dynamically evolves throughout the train-

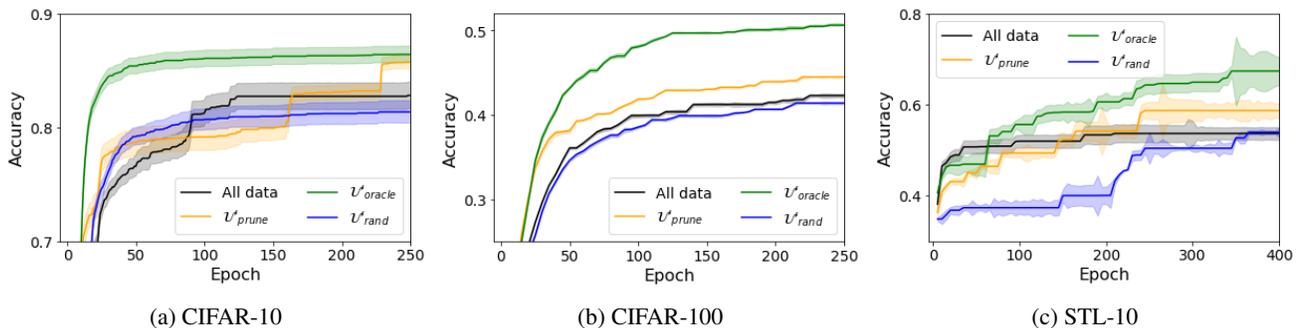|		
(a) CIFAR-10	(b) CIFAR-100	(c) STL-10

Figure 2: Comparison of *PruneSSL* across various classification tasks. Each graph illustrates the average accuracy over 5 repetitions of FreeMatch during the entire learning process. The shaded region signifies the standard error for each plot. Notably, it is observed that employing *PruneSSL*'s data pruning approach (depicted in orange) yields superior performance compared to training with the original unlabeled data (illustrated in black). The green line represents *PruneSSL*'s performance when using an oracle instead of the pseudo-labeling function. The blue line represents training with random pruning of the unlabeled dataset.

ing process. After the introduction of FixMatch, numerous competitive methods have integrated akin selective strategies (Zhang et al., 2021; Berthelot et al., 2022; Li et al., 2021; Zheng et al., 2022; Fan et al., 2023; Jiang et al., 2023).

The primary distinction between the methods discussed above and our approach lies in the utilization of examples from the unlabeled set. In the reviewed methods, all unlabeled examples have the potential to be employed, contingent upon the confidence of the trained model. In contrast, since our approach does not interfere with the optimization process of SSL, examples that are pruned become unavailable to the model irrespective of the model's confidence in them.

Another line of relevant works draws inspiration from adversarial attacks. In the context of attacking semi-supervised methods, Carlini (2021); Shejwalkar et al. (2022) demonstrated that introducing unlabeled examples that impair data separability can negatively impact the performance of several SSL algorithms. This corroborates our findings, underscoring the advantageous nature of removing analogous examples from the unlabeled set.

2   Method

We consider a K-class classification scenario within a semi-supervised learning (SSL) framework. Here, $\mathcal{X}$ is the set of all possible examples, and $\mathcal{Y}$ is their corresponding labels. A semi-supervised learner denoted as $f : \mathcal{X} \to \mathcal{Y}$ is trained using a set of labeled examples $\mathcal{L} \subseteq \mathcal{X} \times \mathcal{Y}$ and unlabeled set of examples $\mathcal{U} \subseteq \mathcal{X}$. As labeled examples are often more expensive to obtain than unlabeled examples, traditionally the labeled pool $\mathcal{L}$ is significantly smaller than the unlabeled pool $\mathcal{U}$. The empirical generalization error of f is its performance on a labeled test set distinct from both $\mathcal{L}$ and $\mathcal{U}$.

In many SSL scenarios, the labeled pool, unlabeled pool, and the test set all originate from the same underlying data distribution. In our experimental setup, we alter the distribution from which $\mathcal{U}$ is sampled. We observe that drawing $\mathcal{U}$ from the same distribution as $\mathcal{L}$ may not be ideal for learning. Specifically, our findings indicate that increasing the separability of $\mathcal{U}$, achieved by pruning hard examples from it, significantly improves the learning of f.

Let $\mathcal{U}' \subseteq \mathcal{U}$ denote a subset of the unlabeled pool $\mathcal{U}$, obtained by applying some pruning method. In our experimental setup, we explore three distinct unsupervised pool types: (i) $\mathcal{U}'_{rand}$ is obtained by uniformly and randomly removing examples from $\mathcal{U}$ while preserving its original distribution. (ii) $\mathcal{U}'_{prune}$ is obtained by removing points as suggested by *PruneSSL*, without relying on any labels of $\mathcal{U}$ to guide it. (iii) $\mathcal{U}'_{oracle}$ is similar to $\mathcal{U}'_{prune}$, but here the removal of points by *PruneSSL* is guided by the true labels of $\mathcal{U}$, rather than its using an inferred pseudo-labeling function.

When constructing $\mathcal{U}'_{oracle}$ and $\mathcal{U}'_{prune}$, our objective is to generate distributions that exhibit enhanced separability as compared to the original unlabeled pool $\mathcal{U}$. To achieve this, we outline a 4-step general protocol for the creation of these unlabeled datasets:

1. Conduct a deep-representation task on $\mathcal{U}$, resulting in an embedding space in which the problem is more linearly separated than in the original pixel space.

2. Obtain labels for the unsupervised pool $\mathcal{U}$, using the real labels for $\mathcal{U}'_{oracle}$ and pseudo-labels for $\mathcal{U}'_{prune}$.

3. Train a simple classifier using the representation from step (1) and the labels obtained in step (2).

4. Use the classifier's confidence to determine which examples are hardest to classify, and prune them from the unlabeled pool.

For a detailed algorithmic representation, please refer to Alg. 1. Additionally, an illustrative visualization of this process can be found in Fig. 1.
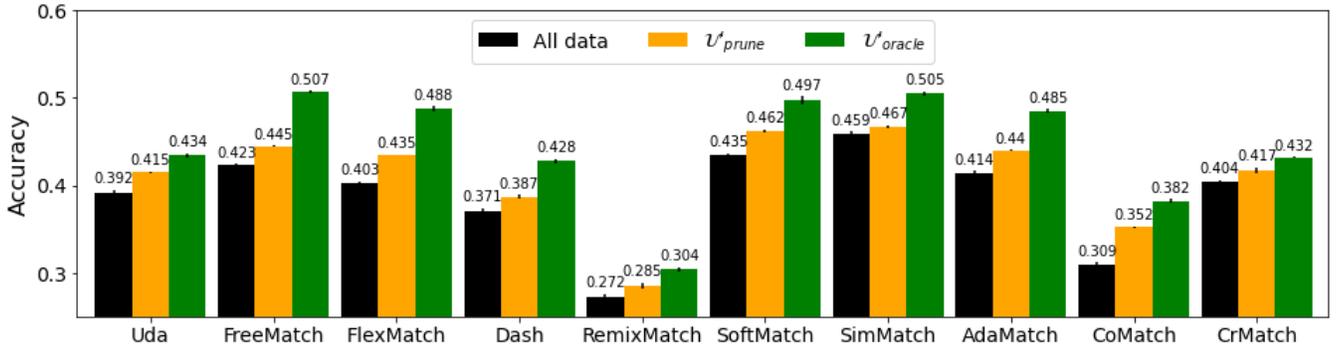
Figure 3: Different SSL algorithms trained with *PruneSSL*. Each group of bars depicts the mean final accuracy of 3 WRN networks, trained with different SSL algorithms, on CIFAR-100, with 300 labeled examples (3 per class). *PruneSSL* improves all the SSL methods, despite using a smaller unlabeled dataset.

Algorithm 1: PruneSSL

1: **Input:** Unlabeled set $\mathcal{U}$, $n = \#$ examples in final unlabeled set

2: **Output:** $\mathcal{U}'_{prune} \subseteq \mathcal{U}$
3: embedding_X $\leftarrow$ self_representation_task($\mathcal{U}$)
4: pseudo_labels $\leftarrow$ pseudo_labeling_task($\mathcal{U}$)
5: clf $\leftarrow$ simple_classifier.fit(embedding, pseudo_labels)
6: $score \leftarrow$ clf.confidence(embedding)
7: $\mathcal{U}'_{prune} \leftarrow n$ examples from $\mathcal{U}$ with highest $score$
8: **return** $\mathcal{U}'_{prune}$

In our work, we explored various alternatives for each step outlined above, yielding consistent qualitative outcomes, as elaborated in the subsequent section. Concerning embedding methods, we employed self-representation tasks like SimCLR (Chen et al., 2020) and transfer learning based on the Inception network (Szegedy et al., 2015) trained on ImageNet (Deng et al., 2009). In terms of the simple classifier, we evaluated choices including a linear SVM, SVM with an RBF kernel, or a small fully connected deep network. The SVM's confidence score was calculated based on the examples' distance from the separating hyperplane, while for the neural network, confidence was determined by the logit corresponding to the example's label. These diverse approaches consistently produced the same qualitative results, as shown in the following section.

2.1 Technical details

In the experiments outlined in Section 3, our evaluation considered a range of SSL algorithms. To ensure a fair comparison, we adopt the SSL evaluation environment crafted by (Wang et al., 2022). This repository offers a diverse selection of SSL methods, all evaluated on the same underlying architectures and datasets. We consider the following recent SSL methods: Dash (Xu et al., 2021), FlexMatch (Zhang et al., 2021), FreeMatch (Wang et al., 2023), RemixMatch (Berthelot et al., 2020), SoftMatch (Chen et al., 2023), Uda (Xie et al., 2020), SimMatch (Zheng et al., 2022), AdaMatch (Berthelot et al., 2022), CoMatch (Li et al., 2021) and Cr-Match (Fan et al., 2023). The specific architectures and hyper-

parameters used for each method are detailed below.

Datasets In the experiments below, we considered 4 datasets: CIFAR-10, CIFAR-100, STL-10 (Coates et al., 2011), and a binary subset of CIFAR-10. When using STL-10, we omitted the unlabeled split due to its inclusion of out-of-distribution examples. The binary subset of CIFAR-10 contained the examples from CIFAR-10 that belonged to the cats and the airplanes classes. This dataset was used due to its simplicity – while still containing real images, SSL algorithms could learn with a small labeled set, even when a small neural architecture was used, drastically reducing its computational cost. This was needed, especially in the ablation study, as state-of-the-art SSL methods are often computationally demanding.

Architectures and hyper-parameters When training SSL methods on CIFAR-10, CIFAR-100, and STL-10 datasets, we employed the Wide-ResNet-28 (WRN) architecture (Zagoruyko and Komodakis, 2016) as the underlying architecture, using 2 width factor, stochastic gradient descent optimizer, 64 batch size, 0.03 learning rate, and 0.9 momentum. We used 0.001 weight decay for CIFAR-10 and 5e-4 for CIFAR-100 and STL-10. When training SSL methods on the binary CIFAR-10, we used a small Vision-Transformer (ViT) (Dosovitskiy et al., 2021), termed *small-ViT*. *small-ViT* has 6 depth, 3 attention heads, 96 width, 5e-4 learning rate, 0.9 momentum, 8 batch size, and was trained using AdamW.

When training SCAN, we used a ResNet-18 architecture with Adam optimizer, learning rate of 1e-4, 0.9 momentum, 1e-4 weight decay, and 128 batch size. When training Sim-CLR, we used ResNet-18 architecture, with Adam optimzier, 3e-4 learning rate, 256 batch size, and 1e-4 weight decay.

Empirically, we observed that while small-ViT achieves significantly worse performance than SOTA architectures, the qualitative results of every experiment remain the same when small-ViT is replaced by WRN. All the experiments were performed using Nvidia GeForce RTX 2080 GPUs.

2.2 Implementation choices for PruneSSL

Unless explicitly stated otherwise, the implementation of *PruneSSL* on CIFAR-10, CIFAR-100, or STL-10, employed SimCLR as the feature space. Subsequently, pseudo-labels were derived using SCAN, followed by RBF-SVM as the simple classifier. When running *PruneSSL* on the binary CIFAR-10, we used k-means instead of SCAN to obtain pseudo-labels. In all datasets, *PruneSSL* pruned 40% of the data, keeping in the unlabeled set 60% of the examples. For CIFAR-10 and STL-10, we used 40 labels for $\mathcal{L}$ and 300 examples For CIFAR-100. In all cases, all classes had the same number of examples in the labeled pool $\mathcal{L}$.

3 Empirical Evaluation

PruneSSL enhances SSL algorithms on different datasets. Fig. 2 illustrates a performance comparison of FreeMatch using different sets of unlabeled data: $\mathcal{U}'_{prune}$, $\mathcal{U}'_{oracle}$, $\mathcal{U}'_{rand}$, and the complete unlabeled set $\mathcal{U}$. This evaluation is conducted on CIFAR-10, CIFAR-100, and STL-10 datasets, utilizing a competitive WRN architecture. Across all datasets, it is evident that while $\mathcal{U}'_{rand}$ exhibits slightly inferior performance as compared to the entire unlabeled dataset $\mathcal{U}$, both $\mathcal{U}'_{prune}$ and $\mathcal{U}'_{oracle}$ showcase significantly improved results.

Despite the fact that $\mathcal{U}$, $\mathcal{U}'_{rand}$, and $\mathcal{U}'_{oracle}$ maintain class balance, this balance is not maintained in $\mathcal{U}'_{prune}$ due to the absence of labels during its construction. Remarkably, even with this class imbalance, $\mathcal{U}'_{prune}$ still surpasses the performance of $\mathcal{U}'_{rand}$. As expected, $\mathcal{U}'_{oracle}$ outperforms $\mathcal{U}'_{prune}$, indicating that potential improvements in the pseudo-labeling function hold promise for advancing *PruneSSL*'s performance in the future.

PruneSSL improves different SSL methods. The results depicted in Fig. 2 extend beyond the FreeMatch algorithm. In the experiments depicted in Fig. 3, we focus on CIFAR-100 and perform the same experiment as Fig. 2, but with a wide variety of SSL algorithms. For each algorithm, we plot the mean final accuracy of 3 repetitions using $\mathcal{U}'_{oracle}$, $\mathcal{U}'_{prune}$ and the entire unlabeled set $\mathcal{U}$. Analogous to the findings observed with FreeMatch, all the considered SSL algorithms show the same qualitative results – training with *PruneSSL* significantly improves the learning, despite its using fewer unlabeled datapoints. Since the methods are vastly different from each other, such results indicate that potential improvements in the SSL algorithms in the future can also advance the performance of *PruneSSL*.

3.1 Ablation study

In this section, we check the significance of the specific selections done by *PruneSSL*. Our investigation reveals a notable degree of robustness in the results. Diverse choices exhibit analogous qualitative trends and provide similar benefits.

Different deep representation feature spaces can achieve comparable qualitative results. In Fig. 4a, we train FreeMatch on CIFAR-100 and contrast multiple variations of *PruneSSL* against training with the complete unlabeled dataset. These *PruneSSL* variations diverge in the embedding part of the algorithm. We explore the following deep-representation spaces: (i) SimCLR (as in Fig. 2); (ii) the penultimate layer of an Inception network, pre-trained on ImageNet; and (iii) the pixel-space of the images, without any embedding.

The deep representation spaces being compared have different characteristics: SimCLR is a contrastive-learning-based self-representation task, which is optimized specifically to get a meaningful representation space. On the other hand, the penultimate layer of Inception is initially tailored for a stand-alone classification task, later employed here in a manner analogous to transfer learning. Despite these differences, both representation spaces outperform training with the entire unlabeled set $\mathcal{U}$ by a significant margin. This result demonstrates the robustness of *PruneSSL* in incorporating a wide range of representation spaces.

It's worth noting that while various deep representation spaces can be effective, not all representations are inherently suitable for *PruneSSL*. As depicted in Fig. 4a, we see that employing the pixel space as a representation yields worse results than training with the entire unlabeled dataset.

Different pseudo labeling methods can achieve similar qualitative results. Illustrated in Fig. 4b, we conduct FreeMatch training on the binary CIFAR-10 and contrast two variants of *PruneSSL* against training with the complete unlabeled dataset. These *PruneSSL* variants diverge in the pseudo-labeling method. We consider using either SCAN or k-means with $k = 2$ on the feature space. Both variants outperform training with the entire unlabeled set $\mathcal{U}$ by a significant margin, suggesting that *PruneSSL* can incorporate different pseudo-labeling techniques.

Changing the classifier. *PruneSSL* relies on the confidence of a simple classifier to determine which examples should be pruned. In all the aforementioned experiments, we employed an SVM with an RBF kernel due to its simplicity and relatively good performance. Fig. 4c explores *PruneSSL* variants integrated with diverse classifiers, each exhibiting the same qualitative behavior. We conducted FreeMatch training on a binary subset of CIFAR-10, applying *PruneSSL* with linear SVM, RBF-kernel SVM, and a small fully connected neural network. Notably, while the RBF-kernel SVM yields optimal performance, all three classifiers outperform training with the entire unlabeled set $\mathcal{U}$, suggesting that *PruneSSL* is compatible with a diverse array of classifiers.

Manipulating the size of the labeled set $\mathcal{L}$. The benefit of pruning examples from $\mathcal{U}$ becomes more pronounced when the size of the labeled set $\mathcal{L}$ is small. In Fig. 5, we present the outcomes of FreeMatch training on the binary CIFAR-10, employing varying sizes for the labeled set $\mathcal{L}$. Evidently, as the size of $\mathcal{L}$ increases, the advantage of *PruneSSL* decreases. This result implies that pruning examples from $\mathcal{U}$ is more effective when the labeled pool is smaller.

Manipulating the size of the unlabeled set $\mathcal{U}$. While an optimal number of examples for pruning exists, a broad spectrum of pruned example quantities can enhance SSL algorithms. Highlighted in Fig. 6, we showcase the outcomes of FreeMatch training on the binary CIFAR-10, incorporating varying sizes for the unlabeled set $\mathcal{U}$. Strikingly, across all selections, a consistent qualitative trend emerges – $\mathcal{U}'_{oracle}$

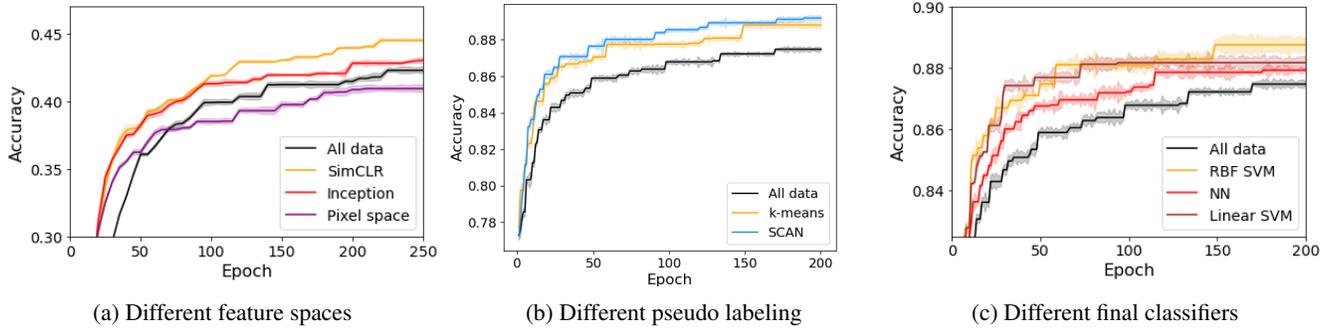| (a) Different feature spaces | (b) Different pseudo labeling | (c) Different final classifiers |

Figure 4: Evaluation of diverse choices at each stage of the *PruneSSL* Algorithm. *PruneSSL* begins by training a representation space, followed by a pseudo-labeling function, and culminates by training a simple classifier that effectively prunes examples with low confidence. In panel (a), we train FreeMatch on CIFAR-100, with *PruneSSL* variants based on either the representation space achieved by training SimCLR or when using the penultimate layer of Inception pre-trained on ImageNet. In panel (b) we train FreeMatch on the binary CIFAR-10 dataset, with *PruneSSL* variants based on either SCAN or k-means for the pseudo-labeling function. In panel (c) we train FreeMatch on the binary CIFAR-10 dataset, with *PruneSSL* variants based on either linear SVM, RBF kerneled SVM, or a small fully connected neural network. Across all three panels, we plot the mean of 5 repetitions, where the shaded area denotes the standard error. We see that whatever method is eventually used in each step, training with the data achieved by *PruneSSL*, despite being significantly smaller, is better than training with the entire unlabeled set. These results demonstrate the adaptability of *PruneSSL* and its effectiveness.
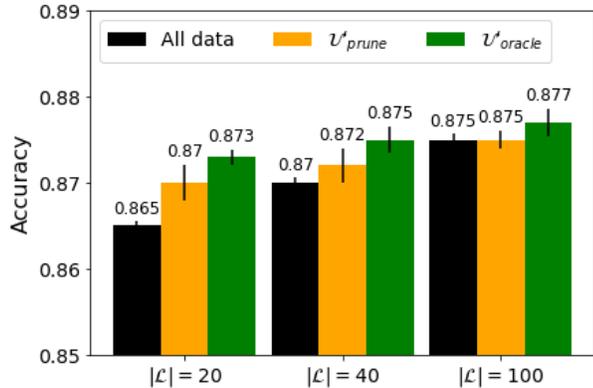

Figure 5: Performance analysis of *PruneSSL* across various sizes of the labeled set. The mean final accuracy of 5 networks trained with FreeMatch on the binary CIFAR-10 dataset is plotted in bar groups. Standard error bars are included. Notably, *PruneSSL* demonstrates more significant improvements with smaller labeled datasets.
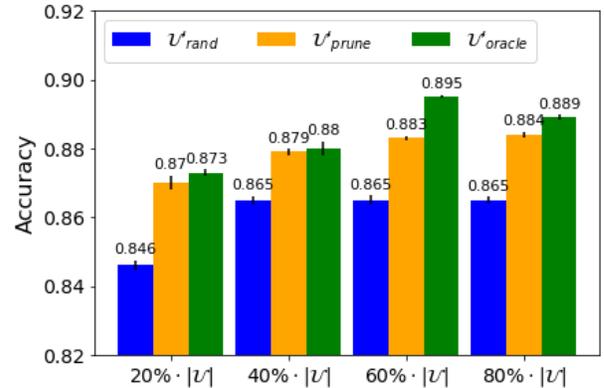

Figure 6: Performance analysis of *PruneSSL* across various sizes of the unlabeled set. The mean final accuracy of 5 networks trained with FreeMatch on the binary CIFAR-10 dataset is plotted in bar groups. Standard error bars are included. Remarkably, while *PruneSSL* achieves peak performance at 60% data retention, diverse pruning levels can consistently improve the performance of the SSL algorithm.

surpasses $\mathcal{U}'_{prune}$, which, in turn, outperforms $\mathcal{U}'_{rand}$. Notably, the best results are achieved when retaining 60% of the data. Nevertheless, we note that a large array of other pruning values also enhance the SSL algorithm performance.

Adding the pruned examples back to the data. A possible explanation for *PruneSSL*'s success might be drawn from the realm of curriculum learning (Bengio et al., 2009; Hacohen and Weinshall, 2019). In curriculum learning, a learner is progressively trained on tasks of increasing complexity. The underlying concept is that mastering simpler

tasks facilitates the acquisition of more complex ones. This parallel might hold here: pruning data could arguably make the problem simpler, given the enhanced separability of the unlabeled data. Once mastery over the easier version of the task is achieved, reintroducing the entire unlabeled dataset might prove advantageous.

We conducted the following experiment: we pruned the binary CIFAR-10 data according to *PruneSSL* and subjected it to 100 epochs of FreeMatch training. Subsequently, the pruned examples were reintegrated into the unlabeled dataset,

followed by an additional 100 epochs of training. We then compared these results to those obtained when training the model over 200 epochs using the complete unlabeled dataset, without any pruning. The results are shown in Fig. 7.

Inspecting Fig. 7, we observe that the reintroduction of pruned examples after several iterations does not yield learning benefits. Instead, upon reintegrating the pruned examples, performance drops, converging to the same levels as training without any pruning. These results imply that the pruned examples indeed have a negative impact on learning, indicative of more than just increased difficulty. Conceivably, these examples lead SSL algorithms to choose suboptimal separation boundaries.
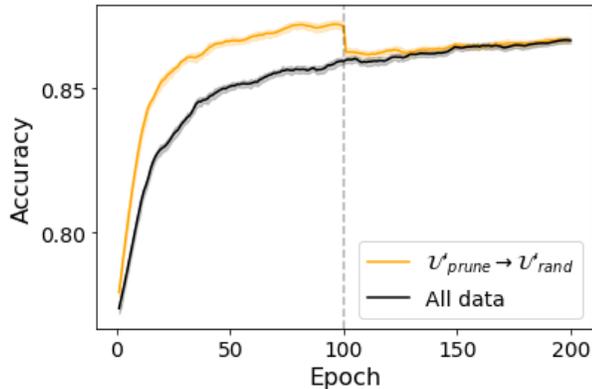

Figure 7: Impact of reintroducing the pruned examples back to the unlabeled set. We plot the mean accuracy of 5 networks, trained with FreeMatch on the binary CIFAR-10 dataset. The black line shows a 200-epoch training on full unlabeled dataset, while the orange line shows *PruneSSL* (100 epochs) followed by the full unlabeled dataset (100 epochs). The dashed vertical line marks the reintroduction of pruned examples. The shaded area represents the standard error. We see that reintroducing the pruned examples harms the performance, suggesting that the pruned examples indeed have a negative impact on SSL algorithms.

Discriminability vs coverage   *PruneSSL* is designed to amplify discriminability within unlabeled data, a tactic that inadvertently leads to a more constrained coverage of the unlabeled dataset, as certain parts of the distribution are pruned completely. In this section, we compare *PruneSSL* and a pruning technique that focuses on covering all parts of the unlabeled distribution.

Coverage emerges as a pivotal concept in the active learning domain (Ren et al., 2021). In active learning, the learner has access to a small set of labeled data and a large set of unlabeled examples. The goal is to pick examples from the unlabeled pool to be annotated so that the resulting labeled set would be optimal for the learner's performance. Several active learning approaches underscore the merit of annotating examples that cover the entire unlabeled dataset for the learners (Sener and Savarese, 2018; Kirsch et al., 2019; Mahmood et al., 2022).

To cultivate a coverage-based pruning, we draw inspiration from the work of Sener and Savarese (2018), replacing *PruneSSL*'s pseudo-labeling with k-means clustering with a large k to cover the data. This approach prunes instances situated farthest from the cluster centroids, allowing the advantage of keeping examples from diverse parts of the distribution, while also preserving representativeness, grounded in their proximity to respective centroids.

In Fig. 8, we plot the results of applying this coverage-oriented pruning technique to the binary CIFAR-10 dataset. A comparison between the removal of instances selected via the aforementioned method and random example removal reveals diminished performance. This outcome aligns with our initial hypothesis, as increasing the coverage of the unlabeled data decreases its discriminability.
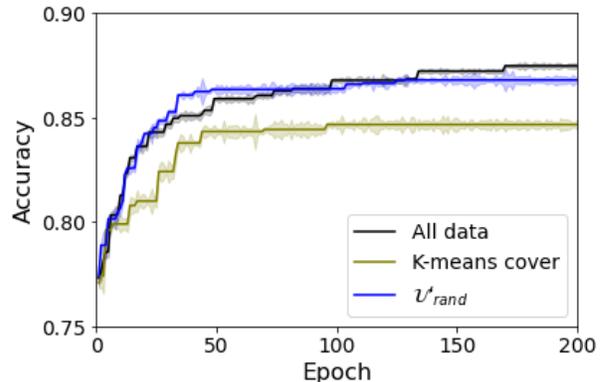

Figure 8: Coverage-focused pruning evaluation versus discriminability. We plot the mean accuracy of 5 networks, trained with FreeMatch on the binary CIFAR-10 dataset. The shaded area represents the standard error. The green line illustrates coverage-driven pruning using k-means centroids (see Section 3.1). Notably, compared to random pruning (blue), k-means pruning exhibits inferior outcomes. This discrepancy may stem from the coverage goal to represent all the areas in the distribution, disproportionately representing obstructive areas, thus negatively affecting data discriminability. Results are shown for $k = 50$; diverse k values yield similar trends.

4   Summary and Discussion

In this paper, we propose a way to improve the unlabeled dataset used in SSL algorithms by making it more separable. Our practical approach called *PruneSSL*, focuses on pruning examples that hinder the separability of the unlabeled data, thus highlighting the inherent cluster structure of the data. The paper presents a comprehensive empirical investigation, demonstrating that this pruning technique notably enhances the performance of various competitive SSL algorithms across a diverse range of image classification tasks.

The structure of *PruneSSL* involves a sequence of steps: a self-representation task followed by pseudo-labeling, and finally, a simple classifier. Through an in-depth analysis presented in the paper, we observe that each of these individual components are adaptable to a different method, suggesting that *PruneSSL* could accommodate future improvements in

each respective field. Testing the limits of its effectiveness, we find that *PruneSSL* yields better results when the labeled dataset is smaller and the data itself is more challenging.

Acknowledgement

References

Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. 2009. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*. pages 41–48.

David Berthelot, Nicholas Carlini, Ekin D. Cubuk, Alex Kurakin, Kihyuk Sohn, Han Zhang, and Colin Raffel. 2020. Remixmatch: Semi-supervised learning with distribution matching and augmentation anchoring. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net. https://openreview.net/forum?id=HklkeR4KPB.

David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A Raffel. 2019. Mixmatch: A holistic approach to semi-supervised learning. *Advances in neural information processing systems* 32.

David Berthelot, Rebecca Roelofs, Kihyuk Sohn, Nicholas Carlini, and Alexey Kurakin. 2022. Adamatch: A unified approach to semi-supervised learning and domain adaptation. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net. https://openreview.net/forum?id=Q5uh1Nvv5dm.

Nicholas Carlini. 2021. Poisoning the unlabeled dataset of {Semi-Supervised} learning. In *30th USENIX Security Symposium (USENIX Security 21)*. pages 1577–1592.

Olivier Chapelle, Bernhard Scholkopf, and Alexander Zien. 2006. Semi-supervised learning. 2006. *Cambridge, Massachusettes: The MIT Press View Article* 2.

Olivier Chapelle and Alexander Zien. 2005. Semi-supervised classification by low density separation. In *International workshop on artificial intelligence and statistics*. PMLR, pages 57–64.

Nitesh V Chawla and Grigoris Karakoulas. 2005. Learning from labeled and unlabeled data: An empirical study across techniques and domains. *Journal of Artificial Intelligence Research* 23:331–366.

Hao Chen, Ran Tao, Yue Fan, Yidong Wang, Jindong Wang, Bernt Schiele, Xing Xie, Bhiksha Raj, and Marios Savvides. 2023. Softmatch: Addressing the quantity-quality trade-off in semi-supervised learning. *CoRR* abs/2301.10921. https://doi.org/10.48550/arXiv.2301.10921.

Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*. PMLR, pages 1597–1607.

Adam Coates, Andrew Ng, and Honglak Lee. 2011. An analysis of single-layer networks in unsupervised feature learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*. JMLR Workshop and Conference Proceedings, pages 215–223.

Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*. Ieee, pages 248–255.

Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2021. An image is worth 16x16 words: Transformers for image recognition at scale. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net. https://openreview.net/forum?id=YicbFdNTTy.

Yue Fan, Anna Kukleva, Dengxin Dai, and Bernt Schiele. 2023. Revisiting consistency regularization for semi-supervised learning. *International Journal of Computer Vision* 131(3):626–643.

Guy Hacohen and Daphna Weinshall. 2019. On the power of curriculum learning in training deep networks. In *International conference on machine learning*. PMLR, pages 2535–2544.

Tao Jiang, Luyao Chen, Wanqing Chen, Wenjuan Meng, and Peihan Qi. 2023. Reliamatch: Semi-supervised classification with reliable match. *Applied Sciences* 13(15):8856.

Justin Johnson, Alexandre Alahi, and Li Fei-Fei. 2016. Perceptual losses for real-time style transfer and super-resolution. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part II 14*. Springer, pages 694–711.

Andreas Kirsch, Joost Van Amersfoort, and Yarin Gal. 2019. Batchbald: Efficient and diverse batch acquisition for deep bayesian active learning. *Advances in neural information processing systems* 32.

Junnan Li, Richard Socher, and Steven C. H. Hoi. 2020. Dividemix: Learning with noisy labels as semi-supervised learning. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net. https://openreview.net/forum?id=HJgExaVtwr.

Junnan Li, Caiming Xiong, and Steven CH Hoi. 2021. Comatch: Semi-supervised learning with contrastive graph regularization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pages 9475–9484.

Yu-Feng Li and Zhi-Hua Zhou. 2014. Towards making unlabeled data never hurt. *IEEE transactions on pattern analysis and machine intelligence* 37(1):175–188.

Rafid Mahmood, Sanja Fidler, and Marc T. Law. 2022. Low-budget active learning via wasserstein distance: An integer programming approach. In *The Tenth International Conference on Learning Representations, ICLR*

2022, Virtual Event, April 25-29, 2022. OpenReview.net. https://openreview.net/forum?id=v8OlxjGn23S.

Avital Oliver, Augustus Odena, Colin A Raffel, Ekin Dogus Cubuk, and Ian Goodfellow. 2018. Realistic evaluation of deep semi-supervised learning algorithms. *Advances in neural information processing systems* 31.

Yassine Ouali, Céline Hudelot, and Myriam Tami. 2020. An overview of deep semi-supervised learning. *arXiv preprint arXiv:2006.05278* .

Pengzhen Ren, Yun Xiao, Xiaojun Chang, Po-Yao Huang, Zhihui Li, Brij B Gupta, Xiaojiang Chen, and Xin Wang. 2021. A survey of deep active learning. *ACM computing surveys (CSUR)* 54(9):1–40.

Zhongzheng Ren, Raymond Yeh, and Alexander Schwing. 2020. Not all unlabeled data are equal: Learning to weight data in semi-supervised learning. *Advances in Neural Information Processing Systems* 33:21786–21797.

Carlos Ruiz, Myra Spiliopoulou, and Ernestina Menasalvas. 2010. Density-based semi-supervised clustering. *Data mining and knowledge discovery* 21(3):345–370.

Ozan Sener and Silvio Savarese. 2018. Active learning for convolutional neural networks: A core-set approach. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net. https://openreview.net/forum?id=H1aIuk-RW.

Virat Shejwalkar, Lingjuan Lyu, and Amir Houmansadr. 2022. The perils of learning from unlabeled data: Backdoor attacks on semi-supervised learning. *arXiv preprint arXiv:2211.00453* .

Aarti Singh, Robert Nowak, and Jerry Zhu. 2008. Unlabeled data: Now it helps, now it doesn't. *Advances in neural information processing systems* 21.

Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin A Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. 2020. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *Advances in neural information processing systems* 33:596–608.

Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. 2015. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. pages 1–9.

Vikas Verma, Kenji Kawaguchi, Alex Lamb, Juho Kannala, Arno Solin, Yoshua Bengio, and David Lopez-Paz. 2022. Interpolation consistency training for semi-supervised learning. *Neural Networks* 145:90–106.

Yidong Wang, Hao Chen, Yue Fan, Wang Sun, Ran Tao, Wenxin Hou, Renjie Wang, Linyi Yang, Zhi Zhou, Lan-Zhe Guo, Heli Qi, Zhen Wu, Yu-Feng Li, Satoshi Nakamura, Wei Ye, Marios Savvides, Bhiksha Raj, Takahiro Shinozaki, Bernt Schiele, Jindong Wang, Xing Xie, and Yue Zhang. 2022. Usb: A unified semi-supervised learning benchmark for classification. In *Thirty-sixth Conference on Neural Informa-

tion Processing Systems Datasets and Benchmarks Track*. https://doi.org/10.48550/ARXIV.2208.07204.

Yidong Wang, Hao Chen, Qiang Heng, Wenxin Hou, Yue Fan, Zhen Wu, Jindong Wang, Marios Savvides, Takahiro Shinozaki, Bhiksha Raj, Bernt Schiele, and Xing Xie. 2023. Freematch: Self-adaptive thresholding for semi-supervised learning. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net. https://openreview.net/pdf?id=PDrUPTXJI_A.

Qizhe Xie, Zihang Dai, Eduard Hovy, Thang Luong, and Quoc Le. 2020. Unsupervised data augmentation for consistency training. *Advances in neural information processing systems* 33:6256–6268.

Yi Xu, Lei Shang, Jinxing Ye, Qi Qian, Yu-Feng Li, Baigui Sun, Hao Li, and Rong Jin. 2021. Dash: Semi-supervised learning with dynamic thresholding. In *International Conference on Machine Learning*. PMLR, pages 11525–11536.

Ting Yang and Carey E Priebe. 2011. The effect of model misspecification on semi-supervised classification. *IEEE transactions on pattern analysis and machine intelligence* 33(10):2093–2103.

Xiangli Yang, Zixing Song, Irwin King, and Zenglin Xu. 2022. A survey on deep semi-supervised learning. *IEEE Transactions on Knowledge and Data Engineering* .

Sergey Zagoruyko and Nikos Komodakis. 2016. Wide residual networks. *arXiv preprint arXiv:1605.07146* .

Bowen Zhang, Yidong Wang, Wenxin Hou, Hao Wu, Jindong Wang, Manabu Okumura, and Takahiro Shinozaki. 2021. Flexmatch: Boosting semi-supervised learning with curriculum pseudo labeling. In Marc'Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan, editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*. pages 18408–18419. https://proceedings.neurips.cc/paper/2021/hash/995693c15f439e3d189b0 Abstract.html.

Mingkai Zheng, Shan You, Lang Huang, Fei Wang, Chen Qian, and Chang Xu. 2022. Simmatch: Semi-supervised learning with similarity matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pages 14471–14481.