

Roll with the Punches: Expansion and Shrinkage of Soft Label Selection for Semi-supervised Fine-Grained Learning

Yue Duan^{1,2}, Zhen Zhao³, Lei Qi⁴, Luping Zhou³, Lei Wang⁵, Yinghuan Shi^{1,2,*}

¹National Key Laboratory for Novel Software Technology, Nanjing University, China

²National Institute of Healthcare Data Science, Nanjing University, China

³School of Electrical and Information Engineering, The University of Sydney, Sydney

⁴School of Computer Science and Engineering, Southeast University, China

⁵School of Computing and Information Technology, University of Wollongong, Australia

yueduan@smail.nju.edu.cn, {zhen.zhao, luping.zhou}@sydney.edu.au, qilei@seu.edu.cn, leiw@uow.edu.au, syh@nju.edu.cn

Abstract

While semi-supervised learning (SSL) has yielded promising results, the more realistic SSL scenario remains to be explored, in which the unlabeled data exhibits extremely high recognition difficulty, *e.g.*, fine-grained visual classification in the context of SSL (SS-FGVC). The increased recognition difficulty on fine-grained unlabeled data spells disaster for pseudo-labeling accuracy, resulting in poor performance of the SSL model. To tackle this challenge, we propose Soft Label Selection with Confidence-Aware Clustering based on Class Transition Tracking (SoC) by reconstructing the pseudo-label selection process by jointly optimizing Expansion Objective and Shrinkage Objective, which is based on a soft label manner. Respectively, the former objective encourages soft labels to absorb more candidate classes to ensure the attendance of ground-truth class, while the latter encourages soft labels to reject more noisy classes, which is theoretically proved to be equivalent to entropy minimization. In comparisons with various state-of-the-art methods, our approach demonstrates its superior performance in SS-FGVC. Checkpoints and source code are available at <https://github.com/NJUyued/SoC4SS-FGVC>.

1 Introduction

Semi-supervised learning (SSL) aims to leverage a pool of unlabeled data to alleviate the dependence of deep models on labeled data (Zhu and Goldberg 2009). However, current SSL approaches achieve promising performance with clean and ordinary data, but become unstuck against the indiscernible unlabeled data. A typical and worth discussing example is the semi-supervised fine-grained visual classification (SS-FGVC) (Su, Cheng, and Maji 2021), where the unlabeled data faced by the SSL model is no longer like the “house” and “bird” that are easy to distinguish, but like the dizzying “*Streptopelia chinensi*” and “*Streptopelia orientalis*”, which are difficult to distinguish accurately even for ornithologists. The mentioned scenario also hints at the practicality of SS-FGVC, *i.e.*, it is resource-consuming to label the fine-grained data for supervised learning.

*Corresponding author.

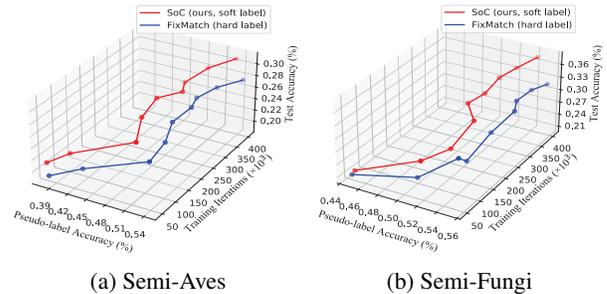


Figure 1: Experimental results on Semi-Aves and Semi-Fungi (see Sec. 4.1 for details of the dataset). Our SoC (soft label) consistently outperforms FixMatch (hard label) even when using pseudo-labels with lower accuracy.

SS-FGVC scenario poses a more problematic problem for pseudo-labeling-based SSL approaches. In general, pseudo-labeling (Lee et al. 2013; Iscen et al. 2019; Arazo et al. 2020) based methods assign pseudo-labels to the unlabeled data online for self-training loops, where these methods rely heavily on the accuracy of pseudo-labels (Xie et al. 2020; Sohn et al. 2020; Li, Xiong, and Hoi 2021; Zhang et al. 2021; Zhao et al. 2022). Unfortunately, fine-grained data may severely affect the quality of pseudo-labels and consequently pull down the model performance. As an example, FixMatch (Sohn et al. 2020), currently the most popular SSL method, selects high-confidence pseudo-labels by a threshold, which are more likely correct in intuition. Then the selected pseudo-labels are converted into one-hot labels for training. If applying to SS-FGVC, the low accuracy of pseudo-labels caused by fine-grained data will pose a great risk to the model training, resulting in further error accumulation in subsequent iterations. Therefore, considering the difficulty of accurately predicting the pseudo-labels in the SS-FGVC scenario, why not step back on a relaxed manner to generate pseudo-labels? A natural idea is to use the soft label. Our motivation is that training with proper soft labels is more robust to noisy pseudo-labels compared to hard labels in SS-FGVC. As shown in Fig. 1, the proposed SoC (as described below) unexpectedly maintains a perfor-

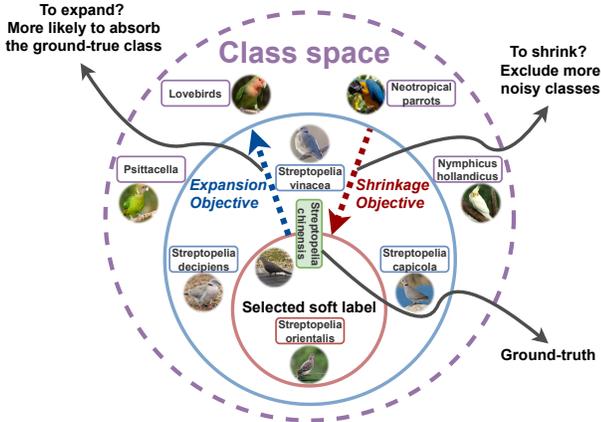


Figure 2: Illustration for proposed approach of soft label selection. We divide the class space of the SS-FGVC scenario into clusters with different granularity, where each cluster (e.g., the blue circle and red circle) contains classes that are more similar to each other. We encourage the soft pseudo-label to select the classes in a cluster with smaller granularity and higher probability of containing ground-truth (i.e., “streptopelia chinensis”), by optimizing *Expansion Objective* (to absorb more candidate classes) and *Shrinkage Objective* (to shrink the cluster for rejecting noisy classes).

mance advantage during training, even when using soft labels with lower pseudo-label accuracy. The bad effects of incorrect hard label on the model have overridden its low entropy advantage while soft label can benefit the model because it could still provide useful information although it is wrong. However, using the traditional soft label may robbing Peter to pay Paul, because it contains all the classes, which obviously introduces too much noise into the learning. We therefore revisit the soft label with the purpose of reducing the penalty to wrong pseudo-labels while preventing the noisy classes in the label from confusing the model. Based on this, we target to find a subset of the class space for each soft pseudo-label, which contains the ground-truth classes as much as possible and the noisy classes as few as possible. Thus, as illustrated in Fig. 2, we propose the soft label selection with a coupled optimization goal to include the classes that are more likely to be ground-truth, and to exclude the classes that are more likely to be noise. Specifically, for *Expansion Objective* (Obj. 1), the classes that are most similar to the current class prediction should be treated as the candidates for the ground-truth; For *Shrinkage Objective* (Obj. 2), the more confident the model is in its prediction, the more it should shrink the range of candidate classes.

Given aforementioned discussions, we propose **Soft Label Selection with Confidence-Aware Clustering based on Class Transition Tracking (SoC)** to establish a high-performance SS-FGVC framework, which can be decomposed into two parts: (1) To optimize *Expansion Objective*: Overall, we utilize a k -medoids-like clustering algorithm (Kaufman and Rousseeuw 2009) to serve as a class selector for the soft pseudo-label generation. In order to make soft labels more

likely to contain the ground-truth class, when the model gives a class prediction for an unlabeled sample, other classes similar to it, i.e., all classes in a cluster, are taken into account. The key of clustering on the class space is the distance metric between classes. However, Euclidean distances usually used in classic k -medoids algorithm will have difficulty in directly representing class-to-class distances. Thus, we innovatively introduce Class Transition Tracking (CTT) technique (Duan et al. 2023) to measure the similarity (rather than distance) between classes for clustering (k -medoids algorithm can directly accept similarity as input). For specific, we statistically track the transitions of the model’s class predictions for each sample, e.g., for the same unlabeled data, the model predicts “Streptopelia chinensis” for it in one epoch and “Streptopelia orien” for it in the next epoch. This oscillation of predictions indicates the degree of similarity between classes, i.e., the more frequent the class transition, the more similar the two classes are and the closer they are in the class space. With CTT, the k -medoids clustering can be performed to select the candidates of soft label for the unlabeled data. (2) To optimize *Shrinkage Objective*: In a nutshell, we shrink the obtained clusters based on the confidence scores of predictions on the unlabeled data. We theoretically prove that in SoC, shrinking the soft labels could lead to entropy minimization (Grandvalet and Bengio 2005). Thus, we use the number of clusters (called k) to control the granularity of clustering, where k is mapped by the confidence, i.e., certain predictions should correspond to a smaller extension range (i.e., selecting smaller k), and conversely, the soft label should be further extended for uncertain predictions (i.e., selecting larger k). The overview of SoC is shown in Fig. 3.

In summary, our contributions are as follows: (1) We propose a coupled optimization goal for SS-FGVC, which is composed of *Expansion Objective* and *Shrinkage Objective*; (2) We propose soft label selection framework with optimizing the above objectives by Class Transition Tracking based k -medoids clustering and Confidence-Aware k Selection; (3) Finally, SoC shows promising performance in SS-FGVC under both the conventional setting and more realistic setting, outperforming various state-of-the-art methods.

2 Related Work

Semi-supervised learning (SSL) relies on leveraging the unlabeled samples to boost the performance of model trained on the labeled data (Yang et al. 2019; Van Engelen and Hoos 2020; Li et al. 2022; Xing et al. 2022, 2023; Zhao et al. 2023; Duan et al. 2023; Yang et al. 2023). For semi-supervised image classification, recently proposed SSL approaches unify *pseudo-labeling* technique to assign pseudo-labels to unlabeled data, enabling their use in the training. The pseudo-label can be classified into two classes: the “hard” label (one-hot label) used in Lee et al. (2013); Sohn et al. (2020); Zhang et al. (2021) and the “soft” label used in Xie et al. (2020); Berthelot et al. (2020); Li, Xiong, and Hoi (2021); Zheng et al. (2022), while the training on pseudo-labels also can be classified into two classes: training multiple models to exploit the disagreement among them in Qiao et al. (2018); Dong, Gong, and Zhu (2018); and training on the la-

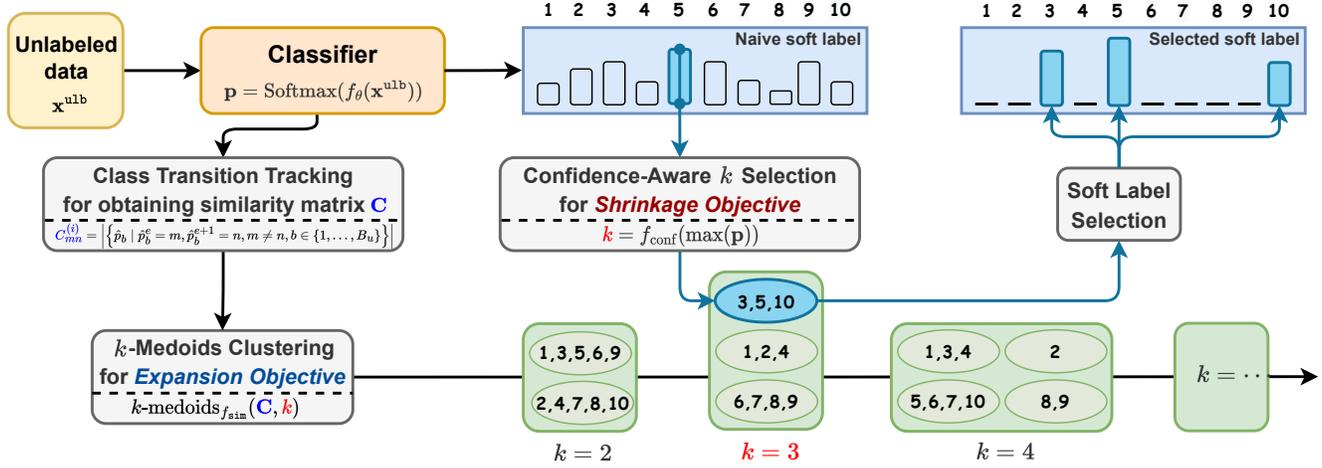


Figure 3: Overview of SoC. We select a subset of the class space to serve as the selection range of candidate classes, encouraging the attendance of ground-truth class as much as possible (Obj. 1) while rejecting noisy class as much as possible (Obj. 2). Given the unlabeled samples, we perform Class Transition Tracking (Eq. (7)) to count the transitions of class predictions to obtain the similarity between classes (see Sec. 3.2). With obtained similarity, we perform k -medoids clustering on the class space to obtain the clusters of candidate classes, which is used to select the soft pseudo-labels (Eq. (8) Eq. (9)). For different samples, we decide the selection of k based on the confidence scores of their class predictions (Eq. (12)), where higher confidence corresponds to a larger k , *i.e.*, a smaller selection range of candidate classes.

beled data to obtain pseudo-labels on the unlabeled data for subsequent training (Lee et al. 2013; Sohn et al. 2020; Li, Xiong, and Hoi 2021; Zhang et al. 2021; Duan et al. 2022).

In the conventional SSL, the datasets is coarse-grained (*e.g.*, “automobile”, “cat” in CIFAR-10 (Krizhevsky, Hinton et al. 2009)) while datasets in the real world usually contain fine-grained data. Although *fine-grained visual classification* (FGVC) is widely discussed (Biederman et al. 1999; Dubey et al. 2018; Zhang, Tang, and Jia 2018; Yang et al. 2018; Wei et al. 2021; Nassar et al. 2021), FGVC in the context of SSL (SS-FGVC) is always ignored. Recently, Su, Cheng, and Maji (2021) takes the lead in evaluating the SS-FGVC setting. In this work, we propose that the pseudo-labeling scheme could be performed in a relaxed manner with a coupled optimization goal, which reaches superior performance cross a variety of SS-FGVC scenarios.

3 Methodology

3.1 A Coupled Optimization Goal for SS-FGVC

SSL involves a training dataset \mathcal{D} divided into two portions: the set of labeled data $\mathcal{D}^{\text{lb}} = \{(\mathbf{x}_i^{\text{lb}}, \mathbf{y}_i)\}_{i=1}^N$ and the set of unlabeled data $\mathcal{D}^{\text{ulb}} = \{\mathbf{x}_i^{\text{ulb}}\}_{i=1}^M$, where N and M are the amount of samples from K classes in \mathcal{D}^{lb} and \mathcal{D}^{ulb} respectively. Formally, \mathbf{x} is sampled from distribution \mathcal{X} while $\mathbf{y} \in \mathcal{Y} = \{1, \dots, K\}$. SSL algorithms aim to learn a predictor $f_\theta: \mathcal{X} \rightarrow \mathcal{Y}$ to correctly classify the test data, where θ is the model parameters. In SS-FGVC, \mathcal{X} presents fine-grained similarity between classes, which poses new challenges of identifying sub-classes that belong to a general class.

Current conventional SSL methods incorporate pseudo-labeling to achieve entropy minimization for the predic-

tions on unlabeled data $\mathbf{x}_i^{\text{ulb}}$. Precisely, given $\mathbf{x}_i^{\text{ulb}}$ in the self-training loop, the model computes its output probability $\mathbf{p}_i = \text{Softmax}(f_\theta(\mathbf{x}_i^{\text{ulb}}))$ and utilizes the hard pseudo-label $\hat{p}_i = \arg \max(\mathbf{p}_i)$ to augment the training set by $\mathcal{D}' = (\mathcal{D} \setminus \{\mathbf{x}_i^{\text{ulb}}\}) \cup \{(\mathbf{x}_i^{\text{ulb}}, \hat{p}_i)\}$. In this way, the model is able to leverage both labeled and unlabeled data to produce a robust learning scheme to improve the performance of SSL. Mathematically, the common goal of pseudo-labeling based methods is to minimize the following objective:

$$\mathcal{L} = \frac{1}{N} \sum_{i=1}^N H(f_\theta(\mathbf{x}_i^{\text{lb}}), \mathbf{y}_i^{\text{lb}}) + \lambda \mathbb{1}(\max(\mathbf{p}_i) \geq \tau) \frac{1}{M} \sum_{i=1}^M H(f_\theta(\mathbf{x}_i^{\text{ulb}}), \hat{p}_i) \quad (1)$$

where $H(\cdot, \cdot)$ is the standard cross entropy, λ is the trade-off coefficient, $\mathbb{1}(\cdot)$ is the indicator function and τ is a pre-defined threshold to filter out low-confidence samples. Pseudo-labeling achieves remarkable performance in conventional SSL by reducing the entropy of predictions, *i.e.*,

$$\min_{\theta} \mathbb{E}_{\mathbf{x}_i^{\text{ulb}} \in \mathcal{D}^{\text{ulb}}} [\mathcal{H}(\mathbf{p}_i)], \quad (2)$$

where $\mathcal{H}(\cdot)$ refers to the entropy. However, there are potential risks in using this technique in SS-FGVC. Given that the difficulty of the model to correctly identify fine-grained images is not comparable to the difficulty of correctly identifying regular images, the resulting hard labels are mixed with too much noise. The damage caused to the model by a wrong hard label is nonnegligible, which means that pseudo-labeling could not work well when the model can generate pseudo-labels with low accuracy. Unfortunately, it is difficult to distinguish fine-grained data easily in SS-FGVC. Therefore, it is natural to consider the use of soft labels in

SS-FGVC to participate in the self-training process. Meanwhile, we still tend to achieve entropy minimization to some extent, which is involved in Obj. 2.

The core idea of SoC is that the selected soft label for each $\mathbf{x}_i^{\text{unlabeled}}$ correspondingly contains only a proper subset \mathcal{C}_i of \mathcal{Y} , aiming at reducing the noise present in pseudo-label to yield a superior performance. Formally, we denote the i -th component of vector v as $v_{(i)}$. Hereafter, *label selection indicator* is defined as $\mathbf{g}_i \subset \{0, 1\}^K$, which is a binary vector representing whether a class is selected or not for the pseudo-label \mathbf{p} on $\mathbf{x}_i^{\text{unlabeled}}$, i.e., $g_{i,(c)} = 1$ indicates $p_{i,(c)}$ is selected while $g_{i,(c)} = 0$ indicates $p_{i,(c)}$ is not selected. Specifically, we compute $g_{i,(c)}$ by

$$g_{i,(c)} = \mathbb{1}(c \in \mathcal{C}_i), \quad (3)$$

where we will describe later how \mathcal{C}_i is obtained. With obtained \mathbf{g}_i , we select pseudo-label \mathbf{p}_i by

$$\tilde{\mathbf{p}}_i = \text{Normalize}(\mathbf{g}_i \circ \mathbf{p}_i), \quad (4)$$

where $\text{Normalize}(x)_i = x_i / \sum_{j=1}^K x_j$ is normalization operation and \circ is Hadamard product. Then $\tilde{\mathbf{p}}_i$ is used for training with optimizing our *coupled goal for SS-FGVC*.

Objective 1 (Expansion Objective). *Encourage the pseudo-label to contain the ground-truth class as much as possible:*

$$\max_{\theta} \mathbb{E}_{\mathbf{x}_i^{\text{unlabeled}} \in \mathcal{D}^{\text{unlabeled}}} [\mathbb{1}(y_i^* \in \{c \mid g_{i,(c)} = 1\}) p_{i,(y_i^*)}], \quad (5)$$

where $y_i^* \in \mathcal{Y}$ is the ground-truth label of $\mathbf{x}_i^{\text{unlabeled}}$. For simplicity, we denote $p_{i,(y_i^*)} \mathbb{1}(y_i^* \in \{c \mid g_{i,(c)} = 1\})$ as z_i^{obj1} .

Intuitively, given that the soft label generated by SoC does not contain all classes (i.e., $\mathcal{C}_i \subsetneq \mathcal{Y}$), a necessary prerequisite for $\tilde{\mathbf{p}}_i$ to be useful for training is that y_i^* is selected into the pseudo-label. This objective corresponds to our more relaxed soft-label training scheme compared to the pseudo-labeling approaches. An extreme way to optimize this objective is to incorporate as many classes as possible into the pseudo-label, but this introduces more noisy information. Another extreme way to optimize this objective is to only select the semantic class of $\mathbf{x}_i^{\text{unlabeled}}$ (i.e., the correct hard pseudo-label), but it is very difficult for SS-FGVC to reach this point. Thus, we should balance the advantages and disadvantages to optimize the first objective, i.e., we suppress the tendency to make soft labels contain all classes.

Objective 2 (Shrinkage Objective). *Encourage the pseudo-label to contain as few classes as possible:*

$$\min_{\theta} \mathbb{E}_{\mathbf{x}_i^{\text{unlabeled}} \in \mathcal{D}^{\text{unlabeled}}} \sum_{c=1}^K g_{i,(c)}, \quad (6)$$

where $\sum_{c=1}^K g_{i,(c)}$ is denoted as z_i^{obj2} for simplicity.

Briefly, we optimize the first goal while minimizing the number of classes in soft labels, i.e., shrinking the range where the ground-truth class may exist to exclude noisy classes. In SoC, this objective corresponds to a mathematical interpretation: the familiar entropy minimization objective (Grandvalet and Bengio 2005), i.e., Eq. (2) (detailed proof is shown in Sec. 3.3).

Given y_i^* is unseen in SSL, we first introduce how to optimize Obj. 1 heuristically in the following subsection.

3.2 Class Transition Tracking Based Clustering

Notably, the key to optimizing Obj. 1 is to obtain a suitable \mathcal{C}_i that contains the semantic class of the unlabeled data as much as possible. The starting point of our solution is: *similar classes are more likely to be misclassified, and the correct answer is hidden among them with a high probability*. Thus, given $\mathbf{x}_i^{\text{unlabeled}}$, when the model considers it to be a “streptopelia chinensis”, those classes that are most similar to “streptopelia chinensis” are also likely to be the ground-truth class (e.g., “streptopelia orientalis”). A reliable solution is to cluster around “streptopelia chinensis”, and then we add the classes in the cluster to \mathcal{C}_i . For simplicity, we adopt k -medoids clustering algorithm (Kaufman and Rousseeuw 2009) in SoC to obtain \mathcal{C}_i . It is well known that k -medoids clustering can accept distance or similarity as input. Because we aim to cluster the classes instead of the samples, it is not convenient to use Euclidean distance, which is most commonly used in k -medoids clustering, to achieve our goal. Hence, we innovatively propose class transition tracking (CTT) (Duan et al. 2023) to simply and effectively model the similarity between fine-grained classes.

In the training, the class prediction of the SSL model for a given sample $\mathbf{x}_i^{\text{unlabeled}}$ is not constant. As new knowledge is learned, the class predictions output by the model may transit from one class to another, i.e., $\hat{p}_i^e = m$ at epoch e transits to $\hat{p}_i^{e+1} = n$ at epoch $e + 1$, where m and n are different classes. First, a adjacency matrix $\mathbf{C} \in \mathbb{R}_+^{K \times K}$ is constructed, where each element C_{mn} represents the frequency of class transitions that occur from class m to class n . C_{mn} is parametrized by the following CTT averaged on last N_B batches with unlabeled data batch size B_u , i.e., $C_{mn} = \sum_{i=1}^{N_B} C_{mn}^{(i)} / N_B$, where

$$C_{mn}^{(i)} = |\{\hat{p}_b^e \mid \hat{p}_b^e = m, \hat{p}_b^{e+1} = n, m \neq n, b \in \{1, \dots, B_u\}\}|. \quad (7)$$

Our core idea is: *the more frequent the transition between two classes, the greater the similarity between the two classes and the closer they are* (see Sec. D.1 of supplementary¹ for more empirical evaluations on this CTT-based similarity measure properties). Thus, we treat $f_{\text{sim}}(m, n) = \frac{C_{mn} + C_{nm}}{2}$ as a similarity function, i.e., the larger the $f_{\text{sim}}(m, n)$, the more similar the two classes m and n are. Finally, we plug $f_{\text{sim}}(m, n)$ into a k -medoids-like clustering process to obtain the class clusters:

$$\mathcal{S} = k\text{-medoids}_{f_{\text{sim}}}(\mathbf{C}, k), \quad (8)$$

where $\mathcal{S} = \{S_c \mid c \in \{1, \dots, k\} \wedge S_c \subsetneq \mathcal{Y}\}$ is obtained k clusters. Algorithmic presentation and computational cost analysis can be respectively found in Algorithm 1 and Sec. C.2 of supplementary. Then, we construct \mathcal{C}_i by

$$\mathcal{C}_i = S_s, s = \arg \max_{c \in \{1, \dots, k\}} \mathbb{1}(\hat{p}_i \in S_c) \quad (9)$$

which means we select the cluster containing the most confident class prediction outputted by the model as \mathcal{C}_i . In this way, although we cannot explicitly optimize Obj. 1 (since we can never know the ground-truth labels of the unlabeled

¹See <https://arxiv.org/abs/2312.12237> for supplementary.

data), we can optimize it heuristically, since the clusters generated based on CTT contain families of classes that are most likely to be misclassified by the classifier, and in which are more likely to contain the ground-truth class.

3.3 Confidence-Aware k Selection

As mentioned in Sec. 3.1, we first show that optimizing Obj. 2 is equivalent to entropy minimization (*i.e.*, Eq. (2)), which is a widely used objective for effectively improving the quality of pseudo-labels (Miyato et al. 2019; Sohn et al. 2020). Denoting \mathcal{C}_i at epoch e as $\mathcal{C}_i^{(e)}$, the following holds:

Theorem 1. *In SoC, minimizing $|\mathcal{C}_i^{(1)}|$ to $|\mathcal{C}_i^{(m)}|$: $\mathcal{C}_i^{(m)} \subsetneq \dots \subsetneq \mathcal{C}_i^{(2)} \subsetneq \mathcal{C}_i^{(1)}$ (*i.e.*, $\sum_{c=1}^K \mathbf{g}_i^{(m)} < \dots < \sum_{c=1}^K \mathbf{g}_i^{(2)} < \sum_{c=1}^K \mathbf{g}_i^{(1)}$), we show that the entropy of \mathbf{p}_i is minimizing:*

$$\mathcal{H}(\tilde{\mathbf{p}}_i^{(m)}) \leq \mathcal{H}(\tilde{\mathbf{p}}_i^{(m-1)}) \leq \dots \leq \mathcal{H}(\tilde{\mathbf{p}}_i^{(2)}) \leq \mathcal{H}(\tilde{\mathbf{p}}_i^{(1)}), \quad (10)$$

where $\mathcal{H}(\cdot)$ refers to the entropy.

To prove this theorem, we first give the following lemma.

Lemma 1. *Given $\tilde{\mathbf{p}}_i$ in SoC (implying $\mathcal{C}_i \subsetneq \mathcal{Y}$ and $\tilde{\mathbf{p}}_i \in \mathcal{C}_i$), we show that the entropy of $\tilde{\mathbf{p}}_i$ is smaller than that of \mathbf{p}_i :*

$$\mathcal{H}(\tilde{\mathbf{p}}_i) \leq \mathcal{H}(\mathbf{p}_i), \quad (11)$$

where $\mathcal{H}(\cdot)$ refers to the entropy.

See Sec. B of supplementary for proof. By Lemma 1, we can simply obtain the proof for Theorem 1.

Proof for Theorem 1. Given $\mathcal{C}_i^{(2)} \subsetneq \mathcal{C}_i^{(1)}$, we can treat $\tilde{\mathbf{p}}_i^{(1)}$ as \mathbf{p}_i and $\tilde{\mathbf{p}}_i^{(2)}$ as $\tilde{\mathbf{p}}_i$ in Lemma 1, *i.e.*, treat the previously obtained $\tilde{\mathbf{p}}_i$ as the naive soft pseudo-label. Thus, $\mathcal{H}(\tilde{\mathbf{p}}_i^{(2)}) \leq \mathcal{H}(\tilde{\mathbf{p}}_i^{(1)})$ holds. For any $\mathcal{C}_i^{(j)} \subsetneq \mathcal{C}_i^{(j-1)}$, we can repeat the above proof to obtain $\mathcal{H}(\tilde{\mathbf{p}}_i^{(j)}) \leq \mathcal{H}(\tilde{\mathbf{p}}_i^{(j-1)})$. Thus, we can obtain $\mathcal{H}(\tilde{\mathbf{p}}_i^{(m)}) \leq \dots \leq \mathcal{H}(\tilde{\mathbf{p}}_i^{(1)})$. \square

By Theorem 1, we propose to “shrink” \mathcal{C}_i to optimize Obj. 2. Since \mathcal{C}_i is obtained by Eq. (9), a feasible way to “shrink” \mathcal{C}_i is to decrease the number of classes in the clusters obtained by k -medoids clustering, *i.e.*, enlarging k . An extreme case is setting $k = K$, which means each cluster only contains one class, *i.e.*, $|\mathcal{C}_i| = 1$. Meanwhile, with \mathcal{C}_i obtained in this way, there is only one component in \mathbf{g}_i that is not zero, which means the pseudo-label \mathbf{p}_i generated by SoC will also degrade to one-hot labels (by Eq. (4)). With training, the model will become more confident in $\mathbf{x}_i^{\text{ulb}}$, and therefore we should shrink the range of candidate classes (*i.e.*, $|\mathcal{C}_i|$) for $\mathbf{x}_i^{\text{ulb}}$ to optimize Obj. 2, *i.e.*, the larger the $\max(\mathbf{p}_i)$, the finer the clustering granularity, which implies the larger k . Thus, we propose Confidence-Aware k Selection: given $\mathbf{x}_i^{\text{ulb}}$, k_i is calculated as

$$k_i = f_{\text{conf}}(\max(\mathbf{p}_i)), \quad (12)$$

where $f_{\text{conf}}(\cdot)$ is a monotonically increasing function that maps confidence scores to k . In SoC, the linear function is adopted as f_{conf} for simplicity:

$$k_i = \lceil \left[\left(\frac{\max(\mathbf{p}_i)}{\alpha} + \frac{2}{K} \right) \times K - \frac{1}{2} \right] \rceil, \quad (13)$$

where $\alpha \geq \frac{K}{K-2}$ is a pre-defined parameter. By this, we control $k_i \in \{2, \dots, K\}$. More discussion on other adopted functions for f_{conf} can be found in Sec. 4.3.

Additionally, the dynamic selection of k corresponds to another intuitive motivation. It is worth noting that since the identification difficulty of different samples is different, it is not reasonable that all samples’ \mathcal{C} are obtained from clusters of the same granularity (*i.e.*, k is fixed). Taking the above considerations into account, we believe that the clustering granularity should be determined at the sample level. The simpler samples should be from clusters with finer granularity (*i.e.*, larger k). Accordingly, the harder samples should be from clusters with coarser granularity (*i.e.*, smaller k), which implies optimizing Obj. 1 since a larger $|\mathcal{C}_i|$ will have a larger probability to contain the ground-truth class of $\mathbf{x}_i^{\text{ulb}}$. For simplicity, we regard the prediction’s confidence score $\max(\mathbf{p}_i)$ as an estimation of the learning difficulty of $\mathbf{x}_i^{\text{ulb}}$. This also corresponds to our k selection scheme: the larger the confidence, the larger the k .

3.4 Putting It All Together

Following prevailing SSL methods (Xie et al. 2020; Sohn et al. 2020; Li, Xiong, and Hoi 2021; Zhang et al. 2021; Tai, Bailis, and Valiant 2021), consistency regularization technique is integrated into SoC, *i.e.*, weak augmentation $\text{Aug}_w(\cdot)$ is applied on \mathbf{x}_i^{lb} and $\mathbf{x}_i^{\text{ulb}}$ while strong augmentation $\text{Aug}_s(\cdot)$ is only applied on $\mathbf{x}_i^{\text{ulb}}$. In a training iteration, we obtain a batch of B labeled data $\{(\mathbf{x}_i^{\text{lb}}, \mathbf{y}_i)\}_{i=1}^B$ and a batch of μB unlabeled data $\{\mathbf{x}_i^{\text{ulb}}\}_{i=1}^{\mu B}$, where μ determines the relative size of labeled batch to unlabeled batch. First, the supervised loss \mathcal{L}_{sup} is defined as

$$\mathcal{L}_{\text{sup}} = \frac{1}{B} \sum_{n=1}^B H(\mathbf{y}_i, f_{\theta}(\text{Aug}_w(\mathbf{x}_i^{\text{ulb}}))), \quad (14)$$

where $H(\cdot, \cdot)$ denotes the cross-entropy loss. Then, the whole algorithm proceeds by three steps:

Step 1. The soft pseudo-label for weakly-augmented $\mathbf{x}_i^{\text{ulb}}$ is computed as $\mathbf{p}_i^w = \text{Softmax}(f_{\theta}(\text{Aug}_w(\mathbf{x}_i^{\text{ulb}})))$. With obtained \mathbf{p}_i^w , the k_i^w for clustering is selected by Eq. (13).

Step 2. Class transition matrix \mathbf{C} is calculated by Eq. (7) in the current iteration. By Eq. (8), the CTT based k -medoids clustering is performed with k_i^w and \mathbf{C} to obtain \mathcal{C}_i^w .

Step 3. \mathbf{g}_i^w is computed by Eq. (3) with \mathcal{C}_i^w and then the selected pseudo-label $\tilde{\mathbf{p}}_i^w$ is computed by Eq. (4) with \mathbf{g}_i^w .

Finally, the consistency regularization is achieved by minimizing the consistency loss \mathcal{L}_{cos} :

$$\mathcal{L}_{\text{cos}} = \frac{1}{\mu B} \sum_{i=1}^{\mu B} H(\tilde{\mathbf{p}}_i^w, f_{\theta}(\text{Aug}_s(\mathbf{x}_i^{\text{ulb}}))). \quad (15)$$

Unlike previous prevailing pseudo-labeling based methods (*e.g.*, Sohn et al. (2020); Li, Xiong, and Hoi (2021)) that waste data with low confidence (as shown in Eq. (1)), SoC exploits all unlabeled data for training. The total loss function is given by:

$$\mathcal{L} = \mathcal{L}_{\text{sup}} + \lambda_{\text{cos}} \mathcal{L}_{\text{cos}}, \quad (16)$$

Dataset	Pseudo-label	Method	Year	from scratch		from ImageNet		from iNat	
				Top1	Top5	Top1	Top5	Top1	Top5
Semi-Aves	—	Supervised Oracle	—	57.4±0.3	79.2±0.1	68.5±1.4	88.5±0.4	69.9±0.5	89.8±0.7
		MoCo	CVPR' 20	28.2±0.3	53.0±0.1	52.7±0.1	78.7±0.2	68.6±0.1	87.7±0.1
	Hard label	Pseudo-Label	ICML' 13	16.7±0.2	36.5±0.8	54.4±0.3	78.8±0.3	65.8±0.2	86.5±0.2
		Curriculum Pseudo-Label	AAAI' 21	20.5±0.5	41.7±0.5	53.4±0.8	78.3±0.5	69.1±0.3	87.8±0.1
		FixMatch	NIPS' 20	28.1±0.1	51.8±0.6	57.4±0.8	78.5±0.5	70.2±0.6	87.0±0.1
		FlexMatch*	NIPS' 21	27.3±0.5	49.7±0.8	53.4±0.2	77.9±0.3	67.6±0.5	87.0±0.2
		MoCo + FlexMatch*	NIPS' 21	35.0±1.2	58.5±1.0	53.4±0.4	77.0±0.2	68.9±0.3	87.7±0.2
	Soft label	KD-Self-Training	CVPR' 21	22.4±0.4	44.1±0.1	55.5±0.1	79.8±0.1	67.7±0.2	87.5±0.2
		MoCo + KD-Self-Training	CVPR' 21	31.9±0.1	56.8±0.1	55.9±0.2	80.3±0.1	70.1±0.2	88.1±0.1
		SimMatch*	CVPR' 22	24.8±0.5	48.1±0.6	53.3±0.5	77.9±0.8	65.4±0.2	86.9±0.3
		MoCo + SimMatch*	CVPR' 22	32.9±0.4	57.9±0.3	53.7±0.2	78.8±0.5	65.7±0.3	87.1±0.2
		SoC	Ours	<i>31.3±0.8</i>	<i>55.3±0.7</i>	57.8±0.5	80.8±0.5	<i>71.3±0.3</i>	<i>88.8±0.2</i>
		MoCo + SoC	Ours	39.3±0.2	62.4±0.4	58.0±0.4	81.7±0.4	70.8±0.4	88.9±0.5
	Semi-Fungi	—	Supervised oracle	—	60.2±0.8	83.3±0.9	73.3±0.1	92.5±0.3	73.8±0.3
MoCo			CVPR' 20	33.6±0.2	59.4±0.3	55.2±0.2	82.9±0.2	52.5±0.4	79.5±0.2
Hard label		Pseudo-Label	ICML' 13	19.4±0.4	43.2±1.5	51.5±1.2	81.2±0.2	49.5±0.4	78.5±0.2
		Curriculum Pseudo-Label	AAAI' 21	31.4±0.6	55.0±0.6	53.7±0.2	80.2±0.1	53.3±0.5	80.0±0.5
		FixMatch	NIPS' 20	32.2±1.0	57.0±1.2	56.3±0.5	80.4±0.5	58.7±0.7	81.7±0.2
		FlexMatch*	NIPS' 21	36.0±0.9	59.9±1.1	59.6±0.5	82.4±0.5	60.1±0.6	82.2±0.5
		MoCo + FlexMatch*	NIPS' 21	44.2±0.6	67.0±0.8	59.9±0.8	82.8±0.7	61.4±0.6	83.2±0.4
Soft label		KD-Self-Training	CVPR' 21	32.7±0.2	56.9±0.2	56.9±0.3	81.7±0.2	55.7±0.3	82.3±0.2
		MoCo + KD-Self-Training	CVPR' 21	39.4±0.3	64.4±0.5	58.2±0.5	84.4±0.2	55.2±0.5	82.9±0.2
		SimMatch*	CVPR' 22	36.5±0.9	61.7±1.0	56.6±0.4	81.8±0.6	56.7±0.3	80.9±0.4
		MoCo + SimMatch*	CVPR' 22	42.2±0.5	67.0±0.4	56.5±0.2	82.5±0.3	57.4±0.2	81.3±0.4
		SoC	Ours	<i>39.4±2.3</i>	<i>62.5±1.1</i>	<i>61.4±0.4</i>	<i>83.9±0.6</i>	<i>62.4±0.2</i>	<i>85.1±0.2</i>
		MoCo + SoC	Ours	47.2±0.5	71.3±0.2	61.9±0.3	85.8±0.2	62.5±0.4	84.7±0.2

Table 1: Accuracy (%) on Semi-Aves and Semi-Fungi. We provide comparisons with multiple baseline methods reported in Su, Cheng, and Maji (2021) and the state-of-the-art SSL methods based on our re-implementation (marked as *). The models are trained from scratch, or ImageNet/iNat pre-trained or the model initialized with MoCo learning on the unlabeled data. Our results are averaged on 3 runs while the standard deviations \pm Std. are reported. Meanwhile, we mark out the best SSL results in italics and the best MoCo results in bold.

where λ_{cos} is a hyper-parameter to weight the importance of the consistency loss. So far, we have established the framework of **Soft** Label Selection with CTT based k -medoids Clustering (**SoC**) which is boosted by Confidence-Aware k Selection to address SS-FGVC. The whole algorithm is presented in Algorithm 2 of supplementary.

4 Experiment

4.1 Experimental Setup

Baselines. We use various SSL baseline methods (Pseudo-Labeling (Berg and Belhumeur 2013), Curriculum Pseudo-Labeling (Cascante-Bonilla et al. 2021), KD-Self-Training (Su, Cheng, and Maji 2021), FixMatch (Sohn et al. 2020), FlexMatch (Zhang et al. 2021) and SimMatch (Zheng et al. 2022)) and self-supervised learning baseline methods (MoCo (He et al. 2020) and MoCo + X) for comparisons. See Sec. C.1 of supplementary for more details. Specially,

supervised oracle means the ground-truth labels of unlabeled data are included for training.

Datasets. We evaluate SoC on two datasets that have been recently introduced into SS-FGVC: Semi-Aves and Semi-Fungi (Su and Maji 2021; Su, Cheng, and Maji 2021). Both of them are consisting of 200 fine-grained classes and exhibit heavy class imbalance. Semi-Aves is divided into the training set and validation set with a total of 5,959 labeled images and 26,640 unlabeled images, and the test set with 8,000 images; Semi-Fungi is divided into the training set and validation set with a total of 4,141 labeled images and 13,166 unlabeled images, and the test set with 4,000 images. Additionally, Semi-Aves and Semi-Fungi also contain 800 and 1194 out-of-distribution (OOD) classes respectively, including 122,208 and 64,871 unlabeled images. SoC is mainly tested on Semi-Aves and Semi-Fungi with in-distribution data, while the two datasets with OOD data are also used for a comprehensive evaluation.

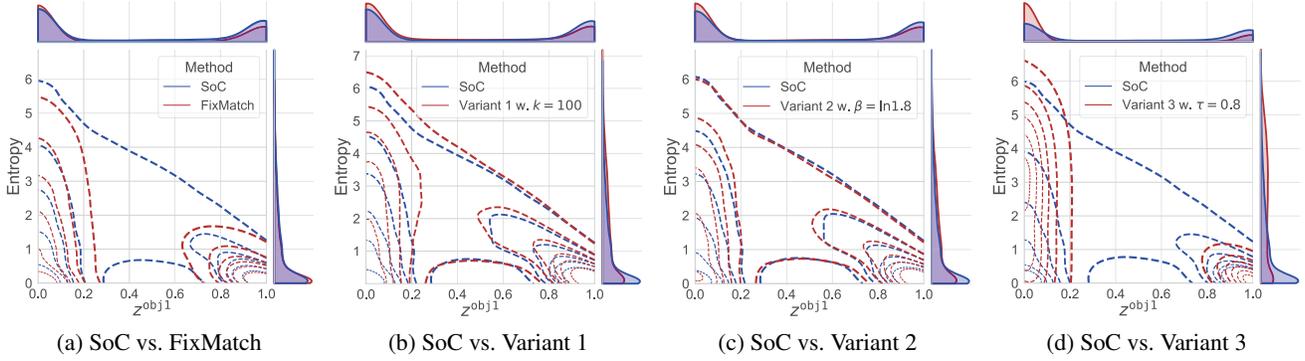


Figure 4: Kernel density estimation for the bivariate distribution of $z^{\text{obj}1}$ and entropy. The x-axis represents $z^{\text{obj}1}$, *i.e.*, the optimization effect of Obj. 1. The y-axis represents the entropy of predictions on the unlabeled data, *i.e.*, the optimization effect of Obj. 2. In general, the greater the density in the lower right corner of the region, the better the optimization effect. The experiments are conducted on Semi-Fungi training from scratch with the same experimental setting as Sec. 4.1. In (a), we show the comparison between SoC and FixMatch. In (b), (c) and (d), we show the comparisons between SoC and the results of Variant 1, Variant 2 and Variant 3 respectively.

Implementation Details. For Semi-Aves and Semi-Fungi, all images are resized to 224×224 . Following Su, Cheng, and Maji (2021), we adopt ResNet-50 (He et al. 2016) as the backbone for all experiments including model training from scratch and pre-trained models on ImageNet (Deng et al. 2009) and iNaturalist 2018 (iNat) (Van Horn et al. 2018), where iNat is a large-scale fine-grained dataset containing some overlapping classes with Semi-Aves (but there are no overlapping images). For consistency regularization, following Sohn et al. (2020), crop-and-flip is used for weak augmentation and RandAugment (Cubuk et al. 2020) is used for strong augmentation. The learning rate with cosine decay schedule is set to 0.01 for training from scratch and 0.001 for training from pre-trained models. For optimization, SGD is used with a momentum of 0.9 and a weight decay of 0.0005. For training from scratch, we train our models for 400k iterations (200k if MoCo is used), whereas our models are trained for 50k iterations when training from pre-trained models. For hyper-parameters, we set $N_b = 5120$, $B = 32$, $\mu = 5$, $\lambda_{\text{cos}} = 1$ and $\alpha = 5$. We report our results averaged on 3 runs with standard variances.

4.2 Results and Analysis

The comparisons of performance are summarized in Tab. 1. As shown in this table, SoC shows the most superior results on both Semi-Aves and Semi-Fungi. Without unsupervised learning (*i.e.*, MoCo (He et al. 2020)), SoC outperforms all baseline SSL methods by a tangible margin. Notably, from the perspective of hard pseudo-label based methods, FixMatch and FlexMatch underperform in SS-FGVC without pre-training. This confirms our claim that carefully selected soft labels are more suitable for SS-FGVC, because we can hardly guarantee the accuracy of pseudo-labels. Meanwhile, from the perspective of soft pseudo-label based methods, SimMatch’s performance is also consistently weaker than SoC, which proves the effectiveness of our soft label selection scheme. The core enhancement of SoC is to optimize

proposed *coupled optimization goal*, which enables to generate more effective soft labels for SS-FGVC. As shown in Fig. 4a, although SoC’s predictions has a lower probability density than FixMatch on the low entropy region, it has higher probability density on high $z^{\text{obj}1}$ regions, *i.e.*, SoC ensures the attendance of ground-truth class in the pseudo-label at the slight expense of minimizing entropy, which leads to a substantial performance gain. With MoCo, SoC also consistently wins all baselines.

On the other hand, the performance of both baselines and SoC is boosted by pre-training, since the SSL model produces high-quality pseudo-labels at the outset. Even so, our method still shows the great advantage of addressing FGVC. Additional results on Semi-Aves and Semi-Fungi with OOD data can be found in Sec. D.2 of supplementary.

4.3 Ablation Study

To explore the effectiveness of each components in SoC, we mainly conduct experiments on three variants of SoC. More ablation studies on class transition tacking (CTT) based clustering can be found Sec. D.3 of supplementary.

Variant 1 - Ablation of k selection: We retain the original CTT-based k -medoids clustering but use a fixed k for it (*i.e.*, ablating Confidence-Aware k Selection). As show in Tab. 2, no matter what value of k is set, the default SoC achieves the best performance. In Fig. 4a, we can observe that SoC achieves better results on our coupled optimization goal, *i.e.*, compared to Variant 1, SoC’s predictions are more clustered in regions with high $z^{\text{obj}1}$ (corresponding to Obj. 1) and low entropy (corresponding to Obj. 2).

Variant 2 - Ablation of f_{conf} : First, we alter α for the default f_{conf} used in Eq. (13) and the results are summarized in Tab. 2. Then, we use the exponential function to replace the liner function used for f_{conf} in default Confidence-Aware k Selection, *i.e.*, we rewrite Eq. (13) as

$$k_i = \lceil (\exp(\beta \times \max(\mathbf{p}_i)) - 1 + \frac{2}{K}) \times K - \frac{1}{2} \rceil, \quad (17)$$

k	25	50	100	150
Top-1 / Top-5	36.5 / 61.0	35.8 / 59.7	36.0 / 60.1	36.6 / 60.9
α	$K/(K-2)$	2	5	10
Top-1 / Top-5	36.6 / 60.8	37.8 / 61.2	39.4 / 62.5	35.8 / 58.7
β	$\ln 1.2$	$\ln 1.4$	$\ln 1.8$	$\ln(2 - \frac{2}{K})$
Top-1 / Top-5	34.9 / 58.4	36.7 / 60.2	37.2 / 60.3	36.9 / 60.2
τ	0.2	0.4	0.8	0.95
Top-1 / Top-5	34.7 / 58.2	34.3 / 57.6	30.3 / 54.9	26.15 / 46.9

Table 2: Accuracy (%) on Semi-Fungi ($K = 200$) training from scratch. k , α , β and τ are altered for Variant 1, 2, 2 and 3 respectively. We mark the results of default SoC as bold.

where $\beta \leq \ln(2 - \frac{2}{K})$ (to keep $k_i \in \{2, \dots, K\}$). In Tab. 2 and Fig. 4b, we observe that the default SoC achieves better performance and optimization result, which proves that simple linear function is competent enough as f_{conf} .

Variant 3 - Ablation of soft label selection: We discard all components in SoC and use the most simple approach to utilize ordinary soft pseudo-label (*i.e.*, \mathbf{g}_i is set to an all-one vector) for learning in SS-FGVC, *i.e.*, we replace the hard label used in FixMatch with soft label. As show in Tab. 2, with various confidence threshold τ in FixMatch, the performance of default SoC demonstrates the effectiveness of our proposed soft label selection. In addition, as shown in Fig. 4c, we can observe that although the ordinary soft label contains all classes, Variant 3 still does not outperform SoC in optimizing $z^{\text{obj}1}$. Meanwhile, the dynamic selection of k has an extraordinary effect on entropy minimization, *i.e.*, SoC’s predictions are clustered in the low entropy region.

Additionally, the comparison between SoC and FixMatch in Fig. 4d also demonstrated that a certain degree of sacrificing entropy minimization to make progress on optimizing Obj. 1 is more suitable for SS-FGVC. After all, it is necessary to output a “certain” prediction only after first ensuring that there is a correct answer in the pseudo-label.

5 Conclusion

The paper presents a new approach for semi-supervised fine-grained visual classification (SS-FGVC), which faces the challenge of working with unlabeled data that has a high recognition difficulty. We propose **Soft Label Selection with Confidence-Aware Clustering** based on Class Transition Tracking (SoC) to tackle the SS-FGVC scenario. SoC optimizes both *Extension Objective* and *Shrinkage Objective* in the tradeoff to improve the soft label selection for SS-FGVC. These objectives respectively encourage the inclusion of candidate classes that may contain the ground truth, while excluding noisy classes. Comprehensive experiments show that SoC consistently achieves significant improvements over the current baseline methods in SS-FGVC. In the future, we believe our method can also be borrowed for more complex and realistic scenarios in SSL.

Acknowledgements

This work is supported by the Science and Technology Innovation 2030 New Generation Artificial Intelligence Major Projects (SQ2023AAA010051), NSFC Program (62222604, 62206052, 62192783), Jiangsu Natural Science Foundation Project (BK20210224). In addition, we sincerely appreciate the discussion and helpful suggestions from Associate Professor Penghui Yao of Nanjing University.

References

- Arazo, E.; Ortego, D.; Albert, P.; O’Connor, N. E.; and McGuinness, K. 2020. Pseudo-labeling and confirmation bias in deep semi-supervised learning. In *International Joint Conference on Neural Networks*.
- Berg, T.; and Belhumeur, P. N. 2013. Poof: Part-based one-vs.-one features for fine-grained categorization, face verification, and attribute estimation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Berthelot, D.; Carlini, N.; Cubuk, E. D.; Kurakin, A.; Sohn, K.; Zhang, H.; and Raffel, C. 2020. ReMixMatch: Semi-Supervised Learning with Distribution Matching and Augmentation Anchoring. In *International Conference on Learning Representations*.
- Biederman, I.; Subramaniam, S.; Bar, M.; Kalocsai, P.; and Fiser, J. 1999. Subordinate-level object classification reexamined. *Psychological research*, 62(2): 131–153.
- Cascante-Bonilla, P.; Tan, F.; Qi, Y.; and Ordonez, V. 2021. Curriculum labeling: Revisiting pseudo-labeling for semi-supervised learning. In *AAAI Conference on Artificial Intelligence*.
- Cubuk, E. D.; Zoph, B.; Shlens, J.; and Le, Q. 2020. RandAugment: Practical Automated Data Augmentation with a Reduced Search Space. In *Advances in Neural Information Processing Systems*.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Dong, Q.; Gong, S.; and Zhu, X. 2018. Imbalanced deep learning by minority class incremental rectification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(6): 1367–1381.
- Duan, Y.; Zhao, Z.; Qi, L.; Wang, L.; Zhou, L.; Shi, Y.; and Gao, Y. 2022. Mutexmatch: semi-supervised learning with mutex-based consistency regularization. *IEEE Transactions on Neural Networks and Learning Systems*.
- Duan, Y.; Zhao, Z.; Qi, L.; Zhou, L.; Wang, L.; and Shi, Y. 2023. Towards Semi-supervised Learning with Non-random Missing Labels. In *IEEE/CVF International Conference on Computer Vision*.
- Dubey, A.; Gupta, O.; Guo, P.; Raskar, R.; Farrell, R.; and Naik, N. 2018. Pairwise confusion for fine-grained visual classification. In *European Conference on Computer Vision*.
- Grandvalet, Y.; and Bengio, Y. 2005. Semi-supervised learning by entropy minimization. In *Advances in Neural Information Processing Systems*.

- He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. 2020. Momentum contrast for unsupervised visual representation learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Iscen, A.; Tolias, G.; Avrithis, Y.; and Chum, O. 2019. Label propagation for deep semi-supervised learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Kaufman, L.; and Rousseeuw, P. J. 2009. *Finding groups in data: an introduction to cluster analysis*. John Wiley & Sons.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images. Technical report, University of Toronto.
- Lee, D.-H.; et al. 2013. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, International Conference on Machine Learning*.
- Li, J.; Xiong, C.; and Hoi, S. C. 2021. Comatch: Semi-supervised learning with contrastive graph regularization. In *IEEE/CVF International Conference on Computer Vision*.
- Li, S.; Cai, H.; Qi, L.; Yu, Q.; Shi, Y.; and Gao, Y. 2022. PLN: Parasitic-like network for barely supervised medical image segmentation. *IEEE Transactions on Medical Imaging*, 42(3): 582–593.
- Miyato, T.; Maeda, S.-I.; Koyama, M.; and Ishii, S. 2019. Virtual Adversarial Training: A Regularization Method for Supervised and Semi-Supervised Learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(8): 1979–1993.
- Nassar, I.; Herath, S.; Abbasnejad, E.; Buntine, W.; and Hafari, G. 2021. All labels are not created equal: Enhancing semi-supervision via label grouping and co-training. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Qiao, S.; Shen, W.; Zhang, Z.; Wang, B.; and Yuille, A. 2018. Deep co-training for semi-supervised image recognition. In *European Conference on Computer Vision*.
- Sohn, K.; Berthelot, D.; Li, C.-L.; Zhang, Z.; Carlini, N.; Cubuk, E. D.; Kurakin, A.; Zhang, H.; and Raffel, C. 2020. FixMatch: Simplifying Semi-Supervised Learning with Consistency and Confidence. In *Advances in Neural Information Processing Systems*.
- Su, J.-C.; Cheng, Z.; and Maji, S. 2021. A realistic evaluation of semi-supervised learning for fine-grained classification. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Su, J.-C.; and Maji, S. 2021. The semi-supervised inaturalist-aves challenge at fgvc7 workshop. *arXiv preprint arXiv:2103.06937*.
- Tai, K. S.; Bailis, P.; and Valiant, G. 2021. Sinkhorn Label Allocation: Semi-Supervised Classification via Annealed Self-Training. In *International Conference on Machine Learning*.
- Van Engelen, J. E.; and Hoos, H. H. 2020. A survey on semi-supervised learning. *Machine Learning*, 109(2): 373–440.
- Van Horn, G.; Mac Aodha, O.; Song, Y.; Cui, Y.; Sun, C.; Shepard, A.; Adam, H.; Perona, P.; and Belongie, S. 2018. The inaturalist species classification and detection dataset. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Wei, X.-S.; Song, Y.-Z.; Mac Aodha, O.; Wu, J.; Peng, Y.; Tang, J.; Yang, J.; and Belongie, S. 2021. Fine-grained image analysis with deep learning: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Xie, Q.; Dai, Z.; Hovy, E.; Luong, T.; and Le, Q. 2020. Unsupervised Data Augmentation for Consistency Training. In *Advances in Neural Information Processing Systems*.
- Xing, Z.; Dai, Q.; Hu, H.; Chen, J.; Wu, Z.; and Jiang, Y.-G. 2023. SVFormer: Semi-Supervised Video Transformer for Action Recognition. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Xing, Z.; Li, H.; Wu, Z.; and Jiang, Y.-G. 2022. Semi-supervised single-view 3d reconstruction via prototype shape priors. In *European Conference on Computer Vision*.
- Yang, L.; Zhao, Z.; Qi, L.; Qiao, Y.; Shi, Y.; and Zhao, H. 2023. Shrinking class space for enhanced certainty in semi-supervised learning. In *IEEE/CVF International Conference on Computer Vision*.
- Yang, Y.; Zhan, D.-C.; Wu, Y.-F.; Liu, Z.-B.; Xiong, H.; and Jiang, Y. 2019. Semi-supervised multi-modal clustering and classification with incomplete modalities. *IEEE Transactions on Knowledge and Data Engineering*, 33(2): 682–695.
- Yang, Z.; Luo, T.; Wang, D.; Hu, Z.; Gao, J.; and Wang, L. 2018. Learning to navigate for fine-grained classification. In *European Conference on Computer Vision*.
- Zhang, B.; Wang, Y.; Hou, W.; Wu, H.; Wang, J.; Okumura, M.; and Shinozaki, T. 2021. Flexmatch: Boosting semi-supervised learning with curriculum pseudo labeling. In *Advances in Neural Information Processing Systems*.
- Zhang, Y.; Tang, H.; and Jia, K. 2018. Fine-grained visual categorization using meta-learning optimization with sample selection of auxiliary data. In *European Conference on Computer Vision*.
- Zhao, Z.; Long, S.; Pi, J.; Wang, J.; and Zhou, L. 2023. Instance-specific and Model-adaptive Supervision for Semi-supervised Semantic Segmentation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Zhao, Z.; Zhou, L.; Duan, Y.; Wang, L.; Qi, L.; and Shi, Y. 2022. Dc-ssl: Addressing mismatched class distribution in semi-supervised learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Zheng, M.; You, S.; Huang, L.; Wang, F.; Qian, C.; and Xu, C. 2022. SimMatch: Semi-supervised Learning with Similarity Matching. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Zhu, X.; and Goldberg, A. B. 2009. Introduction to semi-supervised learning. *Synthesis lectures on artificial intelligence and machine learning*, 3(1): 1–130.