

# FATE: A Prompt-Tuning-Based Semi-Supervised Learning Framework for Extremely Limited Labeled Data

Hezhao Liu  
23020231154209@stu.xmu.edu.cn  
Key Laboratory of Multimedia  
Trusted Perception and Efficient  
Computing, Ministry of Education of  
China, Xiamen University  
Xiamen, China

Yiqun Zhang  
yqzhang@gdut.edu.cn  
School of Computer Science and  
Technology, Guangdong University of  
Technology  
Guangzhou, China

Yang Lu\*  
luyang@xmu.edu.cn  
Key Laboratory of Multimedia  
Trusted Perception and Efficient  
Computing, Ministry of Education of  
China, Xiamen University  
Xiamen, China

Shreyank N Gowda  
shreyank.narayanagowda@  
nottingham.ac.uk  
University of Nottingham  
Nottingham, UK

Hanzi Wang  
Hanzi.Wang@xmu.edu.cn  
Key Laboratory of Multimedia  
Trusted Perception and Efficient  
Computing, Ministry of Education of  
China, Xiamen University  
Xiamen, China

Mengke Li  
mengkejiajia@hotmail.com  
College of Computer Science and  
Software Engineering, Shenzhen  
University  
Shenzhen, China

Chen Gong  
chen.gong@sjtu.edu.cn  
School of Automation and Intelligent  
Sensing, Shanghai Jiao Tong  
University.  
Shanghai, China

## Abstract

Semi-supervised learning (SSL) has achieved significant progress by leveraging both labeled data and unlabeled data. Existing SSL methods overlook a common real-world scenario when labeled data is extremely scarce, potentially as limited as a single labeled sample in the dataset. General SSL approaches struggle to train effectively from scratch under such constraints, while methods utilizing pre-trained models often fail to find an optimal balance between leveraging limited labeled data and abundant unlabeled data. To address this challenge, we propose Firstly Adapt, Then categorize (FATE), a novel SSL framework tailored for scenarios with extremely limited labeled data. At its core, the two-stage prompt tuning paradigm FATE exploits unlabeled data to compensate for scarce supervision signals, then transfers to downstream tasks. Concretely, FATE first adapts a pre-trained model to the feature distribution of downstream data using volumes of unlabeled samples in an unsupervised manner. It then applies an SSL method

specifically designed for pre-trained models to complete the final classification task. FATE is designed to be compatible with both vision and vision-language pre-trained models. Extensive experiments demonstrate that FATE effectively mitigates challenges arising from the scarcity of labeled samples in SSL, achieving an average performance improvement of 33.74% across seven benchmarks compared to state-of-the-art SSL methods. Code is available at <https://github.com/ganchi-huanggua/FATE.git>.

## CCS Concepts

• **Computing methodologies** → **Computer vision problems**;  
**Computer vision**.

## Keywords

Semi-Supervised Learning, Prompt Tuning, Pre-trained Model, Few-shot Learning, CLIP, Parameter-Efficient Fine-Tuning

## ACM Reference Format:

Hezhao Liu, Yang Lu, Mengke Li, Yiqun Zhang, Shreyank N Gowda, Chen Gong, and Hanzi Wang. 2025. FATE: A Prompt-Tuning-Based Semi-Supervised Learning Framework for Extremely Limited Labeled Data. In *Proceedings of the 33rd ACM International Conference on Multimedia (MM '25)*, October 27–31, 2025, Dublin, Ireland. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3746027.3754998>

## 1 Introduction

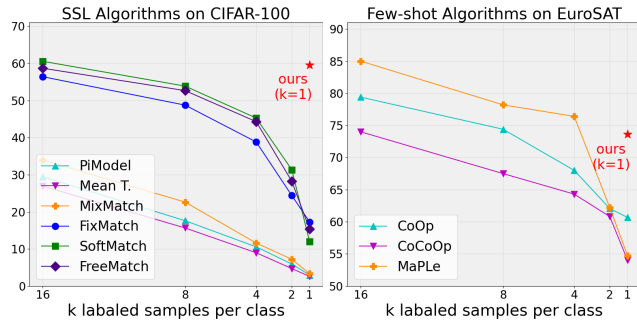
Semi-supervised learning (SSL) is a machine learning paradigm that aims to leverage both labeled and unlabeled data to improve

\*Corresponding author: Yang Lu (luyang@xmu.edu.cn)

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

MM '25, Dublin, Ireland

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.  
ACM ISBN 979-8-4007-2035-2/2025/10  
<https://doi.org/10.1145/3746027.3754998>



**Figure 1: Performance of SSL and few-shot algorithms under different amounts of labeled data  $k$  for each class. Generally, the performance degrades as the number of labeled data decreases.**

model performance. Compared to supervised learning, which relies solely on labeled data, SSL can enhance learning efficiency by utilizing additional unlabeled instances [41, 49]. Due to its ability to reduce reliance on large-scale labeled datasets, SSL has been successfully applied in various real-world scenarios like medical imaging, natural language processing, and autonomous driving [11, 20, 44], making it a crucial research area in the machine learning community.

Existing SSL methods [6, 35, 40, 46] typically assume the availability of a sufficient amount of labeled data, even if the full annotation is not required. However, obtaining even a small fraction of labeled data can be challenging in real-world scenarios. Consider tasks such as medical image annotation or satellite image labeling, which demand substantial human and financial resources. In such cases, datasets may contain only one or two labeled samples while the vast majority remain unlabeled, posing a significant challenge for SSL methods that rely on training a strong feature extractor from scratch. As shown in Fig. 1, we evaluate six classical and state-of-the-art (SOTA) SSL algorithms with varying numbers of labeled samples. The results indicate that as labeled data decreases, the performance of these SSL algorithms degrades sharply. For instance, SoftMatch [6] exhibits nearly a 50% performance drop when reducing labeled samples per class from 25 to just one. The core challenge lies in training a powerful feature extractor from scratch with extremely limited labeled data. On the other hand, with the advent of pre-trained models, this challenge can be mitigated by leveraging a robust backbone, thereby eliminating the need for training from scratch. A few labeled samples can be used to efficiently fine-tune the pre-trained model, and several methods [7, 19, 21, 58] have achieved promising results. These methods are collectively referred to as Parameter-Efficient Fine-Tuning (PEFT) [25]. As further demonstrated in Fig. 1, our evaluation of three SOTA few-shot PEFT algorithms [22, 57, 58] shows that their performance degrades significantly as the number of labeled samples decreases, highlighting their susceptibility to labeled data scarcity. Moreover, these methods lack effective mechanisms to leverage unlabeled data, resulting in substantial information waste.

To fully leverage abundant unlabeled data and mitigate the challenge of extremely scarce labeled data in SSL, we propose Firstly Adapt, Then catEgorize (FATE), a novel prompt-tuning-based SSL

framework. The key idea of FATE is to first adapt a pre-trained model to the downstream feature distribution using unsupervised learning, then refine it with limited labeled samples to enhance classification performance. Specifically, FATE leverages unlabeled data for unsupervised adaptation, aligning the model with downstream distributions before incorporating labeled data for final classification. This ensures that the model benefits from the vast amount of unlabeled data while overcoming the limitations of scarce supervision. FATE introduces a flexible two-stage approach that applies to both vision and vision-language pre-trained models, incorporating tailored prompt tuning strategies. We validate the effectiveness of applying FATE on the foundational models of the vision model [14] and the vision language model [33]. Extensive experiments demonstrate that FATE outperforms existing SOTA SSL and PEFT methods. Our contributions can be summarized as follows.

- We propose FATE, a novel two-stage prompt tuning framework for pre-trained models, specifically designed to address SSL challenges under extremely limited labeled data.
- FATE addresses the challenge of reconciling prompt tuning strategies between vision and vision-language pre-trained models, enabling flexible cross-modality adaptation.
- FATE effectively alleviates the challenge of SSL under minimal labels, demonstrated by a 33.74% average improvement over SOTA methods across seven benchmarks.

## 2 Preliminaries

### 2.1 Semi-Supervised Learning

SSL uses a large amount of unlabeled data and a small amount of labeled data to build models. Formally, to define a  $Y$ -class SSL classification problem, we let  $\mathcal{X} = \{(x_b, y_b) \mid b \in (1, \dots, B)\}$  be a batch of  $B$  labeled sample, where  $x_b$  is the training sample with its one-hot label  $y_b$ . Let  $\mathcal{U} = \{u_b \mid b \in (1, \dots, \mu B)\}$  be a batch of  $\mu B$  unlabeled samples, where  $\mu$  is a fixed parameter that determines the ratio of  $\mathcal{X}$  to  $\mathcal{U}$ .

### 2.2 Prompt Learning for Pre-trained Model

Prompt tuning requires adapting the pre-trained model to downstream tasks without updating its parameters. For a ViT model denoted as  $V$ , an input image  $I$  is cut into  $m$  fixed-sized patches. Each patch is embedded into  $d$ -dimensional latent space with positional encoding. The process of cutting and embedding can be denoted as  $\mathcal{E}$ . The collection of input image patch embeddings  $E$  is:  $E = \mathcal{E}(I)$ , where  $E \in \mathbb{R}^{m \times d}$ . Together with an extra frozen classification token and a group of learnable tokens  $P \in \mathbb{R}^{n \times d}$ , the whole ViT is formulated as:

$$[x_{cls}^{out}; E^{out}; P^{out}] = V([x_{cls}^{in}; E^{in}; P^{in}]), \quad (1)$$

where  $x_{cls}$  denotes the classification token,  $n$  denotes length of the group of prompts. Only  $P$  is updated while  $V$  and  $x_{cls}$  are kept frozen. For brevity, the symbols *in* and *out* below will be omitted. Concretely, The original CLIP [33] works in the following way: for a textual encoder denoted as  $T$  of Transformer [42], a fixed sequence of words, such as "a photo of a [CLASS]." will be converted with byte pair encoding representation [36] and encompassed at a fixed length  $l$  to get input embedding vectors  $S \in \mathbb{R}^{l \times d}$ . The process can be denoted as  $f = T(S)$ , where  $f \in \mathbb{R}^d$  represents the text feature

of S. CoOp [58] replaces the embedding vectors of fixed prompt “a photo of a” with a set of learnable tokens  $P \in \mathbb{R}^{n \times d}$  to get new input embedding vectors  $S'$  and new text feature  $f'$ . Only  $P$  are updated during training, and  $P$  are fixed during inference.

### 3 Method

#### 3.1 Motivation

As classical SSL algorithms struggle to train an effective backbone from scratch under conditions of extreme scarcity of labeled data, we leverage pre-trained models to obtain a powerful feature extractor and fine-tune it, thereby avoiding training from scratch. However, directly fine-tuning the pre-trained model with only labeled data will lead to serious information waste because there is no mechanism to utilize unlabeled data. We aim to develop a method with a mechanism that effectively leverages unlabeled data while retaining some pre-trained knowledge from the pre-trained model. Compared to other PEFT approaches [18, 19], prompt tuning does not require modifications to the backbone network and achieves high computational efficiency by optimizing only the input-related parameters to learn task-specific knowledge. Therefore, the prompt tuning method aligns perfectly with our goal of utilizing unlabeled data to adapt the pre-trained model while preserving some pre-training knowledge.

In Section 3.2, we provide a detailed introduction to the proposed FATE framework, followed by its implementation on the vision pre-trained model in Section 3.3 and the vision-language pre-trained model in Section 3.3, respectively.

#### 3.2 FATE Framework

To implement the aforementioned idea, we propose a two-stage prompt tuning framework for pre-trained models named FATE. In the first stage, we leverage a large amount of unlabeled data to adapt the model to the downstream data’s feature distribution, referred to as the *Adaptation Stage*. In the second stage, we utilize the entire downstream dataset to perform the final SSL classification task, referred to as the *Classification Stage*.

*Adaptation Stage.* Before tuning a pre-trained model with a few labeled data, FATE explores the feature distribution of a large amount of unlabeled data and stores this distribution information into the prompt tokens through prompt learning. Formally, for a frozen pre-trained model  $M$ , we first use unlabeled samples  $\mathcal{U}$  to update a group of prompt tokens through unsupervised learning methods with  $\mathcal{L}_u$ . This set of prompts is designed to learn the overall feature distribution of downstream data, so we refer to the prompts as Distribution-adaptive Prompts (DP), denoted as  $P_d$ . The process can be defined as:

$$P_d^* = \arg \min_{P_d} \frac{1}{\mu B} \sum_{b=1}^{\mu B} \mathcal{L}_u(M(u_b; P_d)), \quad (2)$$

where  $P_d^*$  is the optimized DP. The model’s understanding of downstream data is improved by incorporating  $P_d^*$  during prediction.  $P_d^*$  remains fixed and is used during both the Classification Stage and final inference.

*Classification Stage.* Since unlabeled data was only used to help the model adapt to the feature of downstream data, generally rather than directly for the classification task, the vast amount of feature information provided by unlabeled data is not directly exploited for classification. Therefore, we again leverage unlabeled data along with the supervised signal from a small amount of labeled data to complete the classification task instead of simply fine-tuning the model with labeled samples. This approach allows unlabeled data to fully contribute to adapting the model at a coarse level and improving its classification performance at a fine-grained level.

Formally, we freeze the DP just learned and optimize a new set of prompts named Classification Prompt (CP), denoted as  $P_c$ , to perform the final classification task. We further modify FixMatch [37] through prompt learning to adapt it to the pre-trained model architecture. For each unlabeled sample  $u_b$  in  $\mathcal{U}$ , frozen DP is incorporated when the weak augmentation view of  $u_b$  is input into the model. During the forward propagation of data features, incorporating DP enhances the model’s understanding of downstream data. Thus, the model’s predicted class distributions for both the strong augmented view  $q_s$  and weak augmented view  $q_w$  can be denoted as:

$$q_s = M(\mathcal{A}(u_b); P_c), \quad (3)$$

$$q_w = M(\mathcal{W}(u_b); [P_d^*, P_c]), \quad (4)$$

where  $\mathcal{A}(\cdot)$  and  $\mathcal{W}(\cdot)$  represent the strong and the weak augmentation operation to the sample, respectively. The unsupervised loss of unlabeled data can be defined as:

$$\mathcal{L}_u = \frac{1}{\mu B} \sum_{b=1}^{\mu B} 1_{[\max(q_w) \geq \theta]} \mathcal{H}(\hat{y}_b, q_s), \quad (5)$$

where  $\hat{y}_b = \arg \max(q_w)$  represents the one-hot pseudo-label of the unlabeled sample  $u_b$ , and  $\hat{y}_b$  is generated by the predicted class distribution of  $\mathcal{W}(u_b)$ .  $\theta$  is a hyperparameter representing the threshold above to retain a pseudo-label, and  $\mathcal{H}(\cdot, \cdot)$  represents the cross-entropy loss function. For each labeled sample  $x_b$  in  $\mathcal{X}$ , DP is incorporated again when  $\mathcal{W}(x_b)$  is input into the model. The corresponding predicted class distribution and the supervised loss of labeled data are:

$$q_l = M(\mathcal{W}(x_b); P_d^*, P_c), \quad (6)$$

$$\mathcal{L}_s = \frac{1}{B} \sum_{b=1}^B \mathcal{H}(y_b, q_l). \quad (7)$$

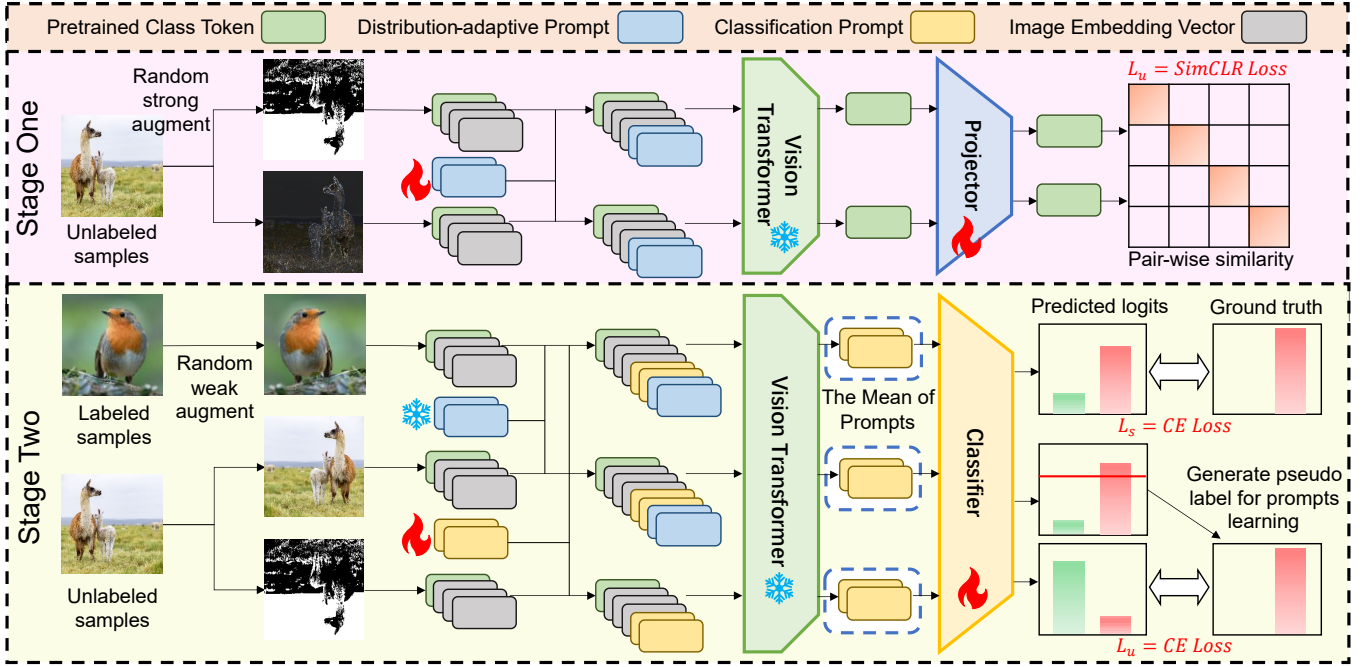
This process of FATE can be defined as:

$$P_c^* = \arg \min_{P_c} (\mathcal{L}_s + \lambda \mathcal{L}_u), \quad (8)$$

where  $\lambda$  is a hyperparameter denoting the relative weight of the unsupervised loss and  $P_c^*$  represents the optimized CP. During inference, both sets of prompts DP and CP are fixed and concatenated with the input embedding vectors to obtain the prediction value.

#### 3.3 Implementation

**Vision Model** For vision models, the focus of FATE implementation should be on the feature level of the data.



**Figure 2: Implementation of FATE on the vision model.** Firstly, we concatenate the DP to the embedding vectors of the unlabeled samples and optimize it with contrastive learning loss. Then we modify FixMatch by first concatenating learnable CP to all embedding vectors, fixing the DP just learned, and concatenating it to the branches of the weak augmented views of both labeled and unlabeled samples. Finally, the mean of the CP is used as the classification feature.

*Adaptation Stage.* For the pre-trained vision model like ViT that requires a classifier for classification, we can use contrastive learning to enable DP to recognize different augmented views of the same sample. This can achieve the purpose of utilizing the unlabeled data to let the model learn the overall feature information of the downstream data mentioned in Section 3.2. For each unlabeled sample  $u_b$  in  $\mathcal{U}$ , as shown in Fig. 2, we make two random strong augmented views  $\mathcal{A}(u_b)_1$  and  $\mathcal{A}(u_b)_2$  where  $\mathcal{A}(u_b)_1 \neq \mathcal{A}(u_b)_2$ , then cut and map them into embedding vectors:  $E_s^1 = \mathcal{E}(\mathcal{A}(u_b)_1)$  and  $E_s^2 = \mathcal{E}(\mathcal{A}(u_b)_2)$ . We concatenate DP on the two sets of input embedding vectors  $E_s^1$  and  $E_s^2$  and then add a frozen pre-trained class token  $x_{cls}$ . Two final input vector groups from  $\mathcal{A}(u_b)_1$  and  $\mathcal{A}(u_b)_2$  and their corresponding output vectors are:

$$[x_{cls}^1; P_d^1; E_s^1] = V([x_{cls}; P_d; E_s^1]), \quad (9)$$

$$[x_{cls}^2; P_d^2; E_s^2] = V([x_{cls}; P_d; E_s^2]), \quad (10)$$

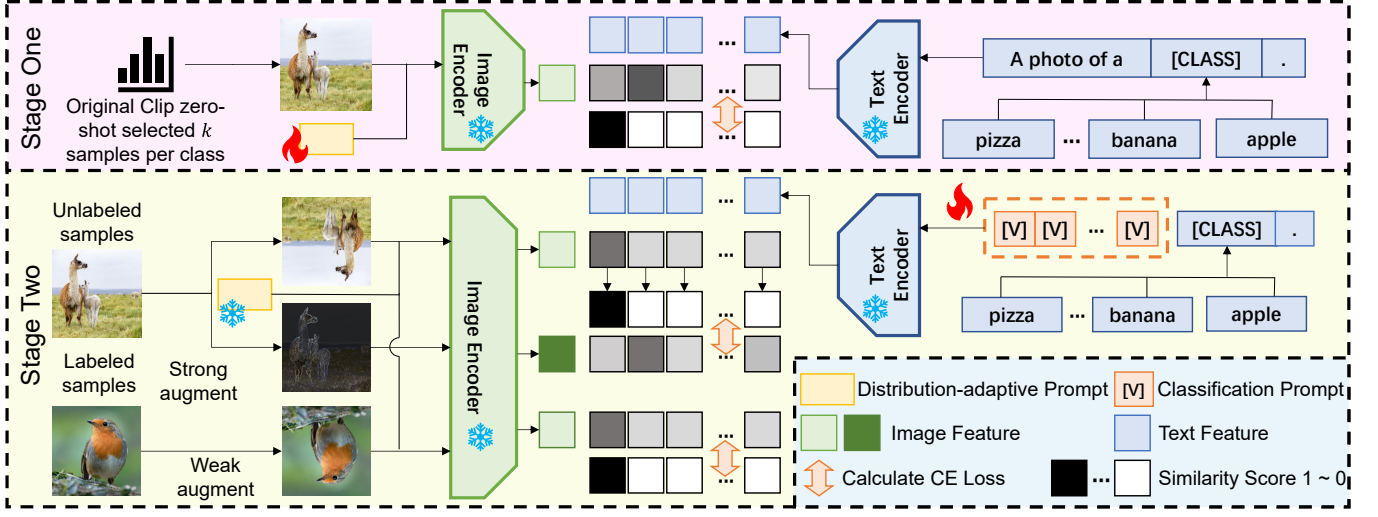
Note that the class token and DP in the two input vector sets are the same, but their values differ after ViT mapping due to the distinct embedding vector groups of the strong augmented views ( $E_s^1 \neq E_s^2$ ). We take the two mapped  $x_{cls}^1$  and  $x_{cls}^2$  and project them into a low-dimensional space using a projector  $J$  to get two projected class tokens  $x_{cls}^1$  and  $x_{cls}^2$ . For batch unlabeled samples, each unlabeled sample  $u_i$  in  $\mathcal{U}$  is corresponding two projected class tokens  $x_{cls}^{i,1}, x_{cls}^{i,2}$  where  $i \in \mu B$ . The distance between the two class tokens  $x_{cls}^{i,1}$  and  $x_{cls}^{i,2}$  obtained by augmenting the two views with the same sample  $u_i$  should be shortened, and the distance between the class tokens

$x_{cls}^{i,*}, x_{cls}^{j,*}$  of each view augmented with different samples should be further, where  $i, j \in \mu B, i \neq j$  and  $*$  represents any of the two augmented views. The contrastive loss of the batch of unlabeled data, also serving as the unsupervised loss Eq. (2) used to optimize DP mentioned in Section 3.2 can be denoted as:

$$\mathcal{L}_u = - \sum_{i=1}^{\mu B} \log \frac{\exp(\langle x_{cls}^{i,1}, x_{cls}^{i,2} \rangle / \tau)}{\sum_{k=1}^{2\mu B} 1_{[k \neq i]} \exp(\langle x_{cls}^{k,*}, x_{cls}^{i,*} \rangle / \tau)}, \quad (11)$$

where  $\tau$  is a scalar temperature hyperparameter adjusting the distribution of similarity scores and  $\langle \cdot, \cdot \rangle$  represents the cosine similarity. Only  $P_d$  and  $J$  are updated in this process.

*Classification Stage.* After the model has completed the adaptation to the downstream data by optimizing DP, we leverage all data to complete the classification task. In contrast to the Adaptation Stage, the pre-trained class token is treated as a feature vector rather than the classification feature, as  $x_{cls}$  has been used for feature contrastive learning instead of classifying. For each unlabeled sample or labeled sample, we can obtain its embedding vectors and corresponding input vector groups based on the method proposed in Section 3.2. After the input vector groups are mapped by ViT, the mean of  $P_c$  will serve as the high-dimensional feature for classification, which is then fed into a classifier  $C$  to obtain  $q_s, q_w$  and  $q_l$  based on Eq. (3), Eq. (4) and Eq. (6). The supervised loss  $\mathcal{L}_s$  and the unsupervised loss  $\mathcal{L}_u$  can be calculated using Eq. (7) and Eq. (5). Then update  $P_c$  and  $C$  using Eq. (8).



**Figure 3: Implementation of FATE on the vision-language model.** We use the original CLIP’s zero-shot capability to pseudo-label the unlabeled data, selecting the top- $k$  samples with the highest predicted values for each class and training the DP for the visual encoder with the pseudo-labeled samples. Then we fix the DP just learned and design CP at the textual encoder side, optimizing it using Eq. (8).

**Vision-Language Model Implementation** For vision-language models, we should prioritize the extraction of semantic features from the data.

*Adaptation Stage.* For the vision-language model, such as CLIP, we can utilize its strong zero-shot capability to generate pseudo-labels for the unlabeled samples. Following [27], we then select samples with high predicted logits to represent the feature distribution of the downstream data because the features of these samples are quite distinct and easy to classify. Based on the idea mentioned in Section 3.2, we can let DP learn the feature distribution of the downstream data with these pseudo-labeled samples. Specifically, as shown in Fig. 3, we first pseudo-label all unlabeled data  $\mathcal{U}$  using the original CLIP, and then divide the unlabeled samples according to the pseudo-labeled classes. We take the  $k$  samples with the highest predicted logits (similarity score) in each class to get a set of samples with high confidence pseudo labels:  $\hat{\mathcal{X}} = \{(\hat{u}_c, \hat{y}_c) \mid c \in (1, \dots, Y \times k)\}$ .  $\hat{\mathcal{X}}$  can be recognized as a representative of the entire downstream unlabeled data’s feature distribution. DP is still done on the visual encoder  $V$ . We use each sample  $\hat{u}_c$  in  $\hat{\mathcal{X}}$  to patch and embed to form embedding vectors  $E_v = \mathcal{E}(\hat{u}_c)$ , and splice the DP into  $E_v$  and input them together into the frozen visual encoder  $V$ . At the same time, the class name prompts of the samples ( $Y$  sentences in the format of “a photo of a [CLASS]”) are also embedded into vectors  $S$ , and  $S$  are input into the frozen textual encoder  $T$ . The process of extracting image and text features can be defined as:

$$[x_{cls}; P_d; E_v] = V([x_{cls}; P_d; E_v]), \quad (12)$$

$$f = T(S), S \in \mathbb{R}^{Y \times l \times d}, f \in \mathbb{R}^{Y \times d}. \quad (13)$$

Then calculate the cosine similarity between text features  $f$  and image features  $x_{cls}$  to get the classification distribution, and calculate the unsupervised loss Eq. (2) mentioned in Section 3.2 with the

one-hot pseudo-label  $\hat{y}_c$  of the sample  $u_c$  to update  $P_d$ :

$$\hat{\mathcal{L}}_u = \frac{1}{k \times Y} \sum_{c=1}^{k \times Y} \mathcal{H}(\hat{y}_c, \langle x_{cls}, f \rangle). \quad (14)$$

*Classification Stage.* Following CLIP’s adaptation, we replicated our implementation on ViT. However, the difference is that our DP is only modified on the visual side and is not aligned with the textual side. To ensure the consistency between the two modalities encoders, our CP is designed on the textual side according to the prompt design method of CoOp [58] and  $P_c$  replaces the embedding vectors of “a photo of a”. The final input embedding vectors  $S'$  and the text features  $f'$  can be defined as:

$$S' = \{[P_c; [\text{CLASS}]_i] \mid i \in (1, \dots, Y)\}, \quad (15)$$

$$f' = T(S'), S' \in \mathbb{R}^{Y \times l \times d}, f' \in \mathbb{R}^{Y \times d}. \quad (16)$$

For each unlabeled sample or labeled sample, we can get its corresponding class tokens mapped by ViT. The operation is based on Section 3.2 and follows the approach we used on ViT without CP. We calculate the similarity between  $x_{cls}$  and  $f'$  to obtain the logits  $q_s, q_w$  and  $q_l$  based on Eq. (3), Eq. (4) and Eq. (6).  $\mathcal{L}_s$  and  $\mathcal{L}_u$  can be calculated using Eq. (7) and Eq. (5), respectively. Finally, we only update  $P_c$  using Eq. (8).

## 4 Experiment

### 4.1 Datasets and Experimental Settings

**Basic Settings:** All experiments are conducted through PyTorch [31] framework on Ubuntu-20.04.6 with one single GeForce RTX 3090 GPU. To construct training datasets for experiments, we randomly sample **one labeled instance** from each class in the original dataset, and the collection of these samples is used as a labeled training dataset. The remaining samples in the original dataset are

**Table 1: Accuracies on seven datasets for SOTA methods in SSL. The backbone used in the upper part is Wide-Resnet [52], and the backbone used in the lower part is ViT. Bold indicates the best performance.**

Method	CIFAR-10	CIFAR-100	ImageNet100	Semi-iNat	SVHN	SUN397	CUB200	Avg.
Pseudo-label [1]	19.07	3.43	2.98	1.87	18.69	1.43	2.37	7.12
FixMatch [37]	17.36	14.59	5.06	1.76	11.05	2.48	3.53	7.98
MixMatch [2]	25.37	3.90	2.10	1.95	20.36	1.85	1.94	8.21
SoftMatch [6]	49.62	9.30	6.44	3.26	16.56	6.05	4.52	13.68
FullMatch [8]	26.25	20.77	5.96	1.88	12.35	2.76	5.73	10.81
FreeMatch [46]	75.65	15.69	7.04	2.81	18.25	2.94	5.17	18.22
Pseudo-label [1]	19.00	4.97	19.42	2.24	12.07	3.24	2.45	9.06
FixMatch [37]	26.59	10.02	12.56	2.44	9.72	4.58	13.78	11.38
SoftMatch [6]	29.15	5.85	6.22	2.47	10.11	1.86	3.23	8.41
Fine-Tuning Classifier (sup)	52.48	34.17	74.72	20.62	16.22	31.93	41.11	38.75
VPT (sup) [21]	51.54	34.69	74.20	20.70	15.90	30.97	40.59	38.37
FATE (Ours)	<b>85.89</b>	<b>59.55</b>	<b>76.48</b>	<b>23.86</b>	<b>20.97</b>	<b>40.03</b>	<b>56.96</b>	<b>51.96</b>

**Table 2: The results for the ablation study on the vision model. If CP is not selected, the class token is used as the classification feature.**

DP	CP	CIFAR-100	SVHN	SUN397	CUB200
×	×	47.61	17.09	35.22	47.78
√	×	58.35	17.71	<b>40.29</b>	54.41
×	√	48.14	20.21	37.39	50.24
√	√	<b>59.55</b>	<b>20.97</b>	40.03	<b>56.96</b>

used as unlabeled examples. In all implementations, strong augmentation strategies for unlabeled samples are Random Augmentation [10] including Cutout [13]. The relative weight of the unsupervised loss  $\lambda$  is set to 1. During training, we randomly sample batches of data from the labeled and unlabeled sets. The average results over three runs with different random seeds are reported.

**Vision Model Setting:** We use the vit-base-patch16-224 pre-trained model published on hugging face from TIMM deep learning library [14, 38, 47] as the feature backbone and evaluate the implementation of FATE on the vision model on seven visual classification benchmarks: (1) CIFAR-10 [24], (2) CIFAR-100, (3) ImageNet100 [12], (4) Semi-iNat [39], (5) SVHN [28], (6) SUN397 [48], (7) CUB200 [43]. We set  $\mu = 1$ ,  $B = 32$ ,  $\tau = 0.5$  and  $\theta = 0.95$  for training. The training process is done with SGD [34] and an initial learning rate of 0.03, decayed by the cosine annealing rule. Both the length of DP and CP are fixed to 12. The training epoch is set to 10 in the Adaptation Stage and 50 in the Classification Stage. Experiments are implemented based on the SSL PyTorch framework USB [45].

**Vision-Language Model Setting:** We use the ViT-B/16 CLIP pre-trained model published by OpenAI as the feature backbone and evaluate the implementation of FATE on the vision-language model on ten visual classification benchmarks: (1) StanfordCars [23], (2) SUN397 [48], (3) Food101 [3], (4) OxfordPets [30], (5) DTD [9], (6) OxfordFlowers [29], (7) FGVC Aircraft [26], (8) Caltech101 [15], (9) EuroSAT [17], (10) UCF101 [32]. We set  $\mu = 16$ ,  $B = 4$ ,  $\tau = 0.5$  and  $\theta = 0.95$  for training. Training is done with SGD [34] and an initial learning rate of 0.1 when adapting and 0.0025 when classifying. The length of DP is fixed to 12 while CP is 16. The

training epoch is set to 20 in the Adaptation Stage, and 20 in the Classification Stage. The results when  $k = 16$  are shown, and the influence of  $k$  will be discussed in Section 4.4. Experiments are implemented based on the PyTorch framework Dassel [56, 59].

## 4.2 Main Results

We compared the performance of FATE on the vision model with the current SOTA SSL methods, and the results are shown in Table 1. It can be seen that FATE achieves SOTA performance on all seven benchmarks, and FATE can improve the test accuracy by 4.41% to 50.25% compared with SoftMatch. FATE’s accuracy improved by 3.16% to 34.35% compared to directly fine-tuning the vision model with labeled data. This is because the semantic information from only one labeled sample is insufficient to train a strong backbone from scratch. Directly replacing these methods with the vision model and fine-tuning is not feasible. After all, the vision model needs more label data to fine-tune from scratch [14]. Directly fixing the vision model and only fine-tuning a classifier with one labeled sample also has poor performance due to overfitting.

## 4.3 Ablation Study on Vision Model

To prove the role of DP and CP, we conducted an ablation experiment on the vision model. We take Fixmatch with VPT (use Fixmatch to optimize VPT parameters) as the basic method and gradually add modules. The ablation results are shown in Table 2, and we can see that DP plays an important role in improving the performance of FATE while the inclusion of CP further improved the experimental results. Furthermore, to visually observe the distribution of classification features, we construct 2D t-SNE visualizations of the classification features of the CIFAR-10 test set, adding DP/CP or not. As shown in Fig. 4, the original data feature points would be highly scattered without DP and CP. The above experimental results demonstrate that the DP, optimized through contrastive learning, captures the feature distribution of downstream data. When the model incorporates the DP for classification tasks, it enhances its understanding of downstream data. Meanwhile, CP refines class boundaries, resulting in compact and distinct clusters, further enhancing the model’s performance.




Figure 4: The t-SNE visualization of FATE’s implementation on the vision model for the CIFAR-10 test set. With the inclusion of DP/CP, the features of data points belonging to the same class cluster together, whereas the original features without DP/CP remain dispersed.




Figure 5: Test accuracy of FATE implemented on the vision language model on four datasets with different  $k$ . We performed ablation experiments on DP and CP. The model performance with DP is already higher than the original CLIP zero-shot, and with CP, the model performance can be further improved. As  $k$  increases to a certain extent, the model’s performance tends to converge or decrease.

Table 3: The results for fine-tuning a classifier (F C) on the vision model with/without DP using one labeled data.

	CIFAR-100	SVHN	SUN397	CUB200
F C w.o. DP	34.17	16.22	31.93	41.11
F C w. noisy DP	33.68	14.37	26.76	40.08
F C w. DP	<b>40.61</b>	<b>17.09</b>	<b>35.25</b>	<b>46.92</b>

#### 4.4 Influence on the Value of $k$

This section will discuss the influence of  $k$  on the implementation of the vision-language model. We set  $k = 1, 2, 4, 8, 16, 32$  and use these pseudo-labeled data sets to train our DP. Furthermore, we conducted ablation experiments on DP and CP, examining the impact of DP and CP on the final predictive performance as  $k$  varies. As shown Fig. 5, the prediction effect will be stronger after using DP than the original zero-shot CLIP. Moreover, incorporating DP during the Classification Stage to optimize CP further enhances the model’s performance. This shows that when DP is used for the Classification Stage, with the prior knowledge from unlabeled data provided by the DP, the model can perform better in the Classification Stage. When  $k$  is very small, the trained DP cannot play a good role, which also indirectly shows that a certain amount of unlabeled data is needed to support the training of DP. When  $k$  reaches a certain level, due to the decline in the quality of pseudo-labels, the trained

DP will be affected by noise and show a trend of performance convergence or decline.

#### 4.5 Effectiveness of Fine-tuning with DP

We conducted more experiments on the vision model to prove the effectiveness of DP as shown in Table 3. The quality of pseudo-labeling will be improved when the DP is added to pseudo-label unlabeled data in the Classification Stage. We fixed the DP after training it in the Adaptation Stage, and then directly used a labeled sample to fine-tune a classifier and evaluate the model’s performance. It was found that the final model performance obtained when adding DP was better than the performance without DP. This shows that our DP learned the overall feature distribution of downstream data and smoothed the semantic information from only one labeled data point to alleviate the overfitting phenomenon. Furthermore, we add DP without updating it to illustrate that DP has indeed learned the feature distribution of downstream data, rather than simply smoothing the classification boundaries of the model by adding noise.

#### 4.6 Comparison with Other PEFT Strategies

As for FATE’s implementation on the vision-language model, we compare it with a wider range of PEFT strategies in one-shot learning scenarios, including adapter-based approaches: TIP-Adapter-F [55], CLIP-Adapter [16]; alternative prompt tuning methods on

**Table 4: Accuracies on ten datasets for SOTA methods in other PEFT strategies and FATE’s implementation on the Vision-Language model. Underlined values indicate suboptimal performance. Aside from GLoRA, CoOp, CoCoOp, MaPLe, and our experiments, other results are taken from [53].**

Method	Cars	SUN	Food	Pets	DTD	Flowers	Aircraft	Caltech	EuroSAT	UCF	Avg.
CLIP [33]	58.5	62.5	85.9	89.0	44.1	70.7	24.8	93.3	48.3	67.6	64.47
CoOp [58]	67.5	66.8	84.2	90.4	49.7	78.6	16.5	93.3	60.7	72.0	67.96
CoCoOp [57]	68.3	<u>68.6</u>	85.5	91.9	51.5	74.7	27.6	94.2	54.0	71.2	68.75
MaPLe [22]	67.8	68.8	85.8	<b>92.2</b>	50.3	76.1	28.2	94.2	54.7	71.7	68.96
TIP-Adapter-F [55]	67.1	67.2	85.8	90.6	51.6	83.8	28.8	94.0	<u>67.8</u>	73.4	71.01
CLIP-Adapter [16]	65.7	65.4	86.1	89.0	44.2	71.3	25.2	92.0	49.3	66.9	65.51
PLOT++ [5]	<u>68.8</u>	66.8	<u>86.2</u>	91.9	<b>54.6</b>	80.5	28.6	<u>94.3</u>	65.4	<b>74.3</b>	<u>71.14</u>
KgCoOp [50]	66.7	68.4	<b>86.4</b>	<u>92.1</u>	52.7	74.7	26.8	94.2	61.9	72.8	69.67
ProGrad [60]	68.2	67.0	84.9	91.4	52.8	80.9	28.8	93.5	57.0	73.3	69.78
GLoRA [4]	7.9	39.5	30.9	68.5	2.6	<b>87.3</b>	9.2	77.2	12.8	41.0	37.70
TaskRes [51]	<u>68.8</u>	68.1	84.6	90.2	53.8	81.7	<b>31.3</b>	93.6	65.4	71.7	70.92
FATE (Ours)	<b>68.9</b>	<b>68.9</b>	83.1	91.2	<u>53.9</u>	<u>84.9</u>	<u>28.9</u>	<b>94.8</b>	<b>73.6</b>	<u>73.8</u>	<b>72.20</b>

**Table 5: Comparison of our method with other SSL methods fine-tuned on CLIP across four benchmarks. -V/-T means applying visual/textual prompts as a tuning strategy.**

Method	DTD	EuroSAT	Flowers	Aircraft	Avg.
CPL-V [54]	<u>56.90</u>	<b>76.96</b>	70.85	17.84	55.64
FPL-V [54]	51.15	66.73	69.61	19.76	51.81
GRIP-V [27]	54.57	63.48	67.95	19.43	51.36
CPL-T [54]	<b>57.78</b>	71.81	<b>84.99</b>	20.59	<u>58.79</u>
FPL-T [54]	51.54	63.08	74.87	<u>20.80</u>	52.57
GRIP-T [27]	56.07	58.66	83.60	16.98	53.83
FATE (Ours)	53.93	<u>73.60</u>	<u>84.93</u>	<b>28.87</b>	<b>60.33</b>

CLIP: PLOT++ [5], KgCoOp [50], ProGrad [60], CoOp [58], CoCoOp [57], MaPLe [22]; improved LoRA [19] variants adapted for vision tasks: GLoRA [4]; and processes that tune CLIP’s text embeddings by introducing residual parameters: TaskRes [51]. Experimental results shown in Table 4 demonstrate that FATE achieves optimal performance on four benchmarks and suboptimal results on four others among ten evaluated benchmarks, surpassing all other PEFT methods in overall performance. This superior performance indicates that our method exploits additional information from unlabeled data, yielding better results than approaches relying solely on supervised fine-tuning signals.

#### 4.7 Comparison with Other CLIP-Based SSL Methods

Furthermore, we compare FATE’s implementation on vision-language models against alternative SSL approaches employing prompt tuning on CLIP. These methods [27, 54] operate through an iterative process of screening high-confidence samples via CLIP’s prompt tuning mechanism and retraining new prompts with the selected samples. While using identical dataset configurations, we maintained all other experimental parameters at their default values as provided in the authors’ original implementations. As shown in Table 5, FATE obtains one optimal and two suboptimal results on four

**Table 6: The results for adding DP to the strong augment branch of the implementation on the vision model.**

DP	CIFAR-100	Semi-iNat	SVHN	SUN397	CUB200
√	48.63	23.46	17.10	37.02	46.36
×	<b>59.55</b>	<b>23.86</b>	<b>20.97</b>	<b>40.03</b>	<b>56.96</b>

benchmarks. Notably, our method outperforms the second-best approach by an average margin of nearly 2 points. This performance gap highlights FATE’s superior robustness in extremely label-scarce scenarios, whereas other SSL methods fail to fully leverage this advantage despite utilizing pre-trained models.

#### 4.8 Discussion on the Position of DP

We experimented by adding DP to the embedding vector group of the strong augmented view of unlabeled data in the Classification Stage, and found that the performance shown in Table 6 was worse than the final result we provided. It is not necessary to add DP to the branch of the strong augmented view, because the significance of the strong augmented view is to let the model recognize different samples of the same class and reduce the inductive bias of the model. Without DP, the strong augmented view undergoes an additional augmentation at the embedding vector level compared to when DP is included. The reason why the DP is involved in the predicted class distribution mapping of the weak augmented view of the unlabeled data and the labeled data is that we need the predicted class distributions of both to be relatively accurate.

### 5 Conclusion

This paper proposes a two-stage prompt tuning framework called FATE for pre-trained models to address the SSL problem when labeled samples are extremely scarce. FATE is adaptable to both vision and vision-language models. It first uses a large amount of unlabeled data to help the model adapt to the feature distribution of the downstream data, then adapts SSL algorithms to suit pre-trained models, thereby completing the classification task. We demonstrate the effectiveness of FATE through extensive experiments.

## Acknowledgments

This study was supported in part by the National Key Research and Development Program of China under Grant 2022ZD0160402; in part by the National Natural Science Foundation of China under Grants 62376233, 62476063, 62336003, and 12371510; in part by the Guangdong Basic and Applied Basic Research Foundation under Grants 2024A1515010163 and 2025A1515011293; in part by the Shenzhen Science and Technology Program under Grant RCBS20231211-090659101; in part by the National Key Laboratory of Radar Signal Processing under Grant JKW202403.

## References

- [1] Eric Arazo, Diego Ortego, Paul Albert, Noel E O'Connor, and Kevin McGuinness. 2020. Pseudo-labeling and confirmation bias in deep semi-supervised learning. In *2020 International Joint Conference on Neural Networks*. IEEE, 1–8.
- [2] 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 (2019).
- [3] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. 2014. Food-101 – Mining Discriminative Components with Random Forests. In *European Conference on Computer Vision*.
- [4] Arnab Chavan, Zhuang Liu, Deepak Gupta, Eric Xing, and Zhiqiang Shen. 2023. One-for-all: Generalized lora for parameter-efficient fine-tuning. *arXiv preprint arXiv:2306.07967* (2023).
- [5] Guangyi Chen, Weiran Yao, Xiangchen Song, Xinyue Li, Yongming Rao, and Kun Zhang. 2022. Prompt learning with optimal transport for vision-language models. (2022).
- [6] 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. *arXiv preprint arXiv:2301.10921* (2023).
- [7] Shoufa Chen, Chongjian Ge, Zhan Tong, Jiangliu Wang, Yibing Song, Jue Wang, and Ping Luo. 2022. Adaptformer: Adapting vision transformers for scalable visual recognition. *Advances in Neural Information Processing Systems* 35 (2022), 16664–16678.
- [8] Yuhao Chen, Xin Tan, Borui Zhao, Zhaowei Chen, Renjie Song, Jiajun Liang, and Xuequan Lu. 2023. Boosting Semi-Supervised Learning by Exploiting All Unlabeled Data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 7548–7557.
- [9] M. Cimpoi, S. Maji, I. Kokkinos, S. Mohamed, and A. Vedaldi. 2014. Describing Textures in the Wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- [10] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. 2020. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*. 702–703.
- [11] Andrew M Dai and Quoc V Le. 2015. Semi-supervised sequence learning. *Advances in Neural Information Processing Systems* 28 (2015).
- [12] 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, 248–255.
- [13] Terrance DeVries. 2017. Improved Regularization of Convolutional Neural Networks with Cutout. *arXiv preprint arXiv:1708.04552* (2017).
- [14] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiuhua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* (2020).
- [15] Li Fei-Fei, Rob Fergus, and Pietro Perona. 2007. Learning generative visual models from few training examples: An incremental Bayesian approach tested on 101 object categories. *Computer Vision and Image Understanding* 106, 1 (2007), 59–70.
- [16] Peng Gao, Shijie Geng, Renrui Zhang, Teli Ma, Rongyao Fang, Yongfeng Zhang, Hongsheng Li, and Yu Qiao. 2024. Clip-adapter: Better vision-language models with feature adapters. *International Journal of Computer Vision* 132, 2 (2024), 581–595.
- [17] Patrick Helber, Benjamin Bischke, Andreas Dengel, and Damian Borth. 2019. Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 12, 7 (2019), 2217–2226.
- [18] Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. Parameter-efficient transfer learning for NLP. In *International Conference on Machine Learning*. PMLR, 2790–2799.
- [19] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. 2022. Lora: Low-rank adaptation of large language models. *International Conference on Learning Representations* 1, 2 (2022), 3.
- [20] Shu Hu, Chun-Hao Liu, Jayanta Dutta, Ming-Ching Chang, Siwei Lyu, and Naveen Ramakrishnan. 2022. Pseudoprop: Robust pseudo-label generation for semi-supervised object detection in autonomous driving systems. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 4390–4398.
- [21] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and Ser-Nam Lim. 2022. Visual prompt tuning. In *European Conference on Computer Vision*. Springer, 709–727.
- [22] Muhammad Uzair Khattak, Hanoona Rasheed, Muhammad Maaz, Salman Khan, and Fahad Shahbaz Khan. 2023. Maple: Multi-modal prompt learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 19113–19122.
- [23] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 2013. 3D Object Representations for Fine-Grained Categorization. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*.
- [24] Alex Krizhevsky, Geoffrey Hinton, et al. 2009. Learning multiple layers of features from tiny images. (2009).
- [25] Vladislav Lialin, Vijeta Deshpande, and A Rumshisky. 2023. Scaling down to scale up: A guide to parameter-efficient fine-tuning (2023). *arXiv preprint arXiv:2303.15647* (2023).
- [26] Subhansu Maji, Esa Rahtu, Juho Kannala, Matthew Blaschko, and Andrea Vedaldi. 2013. Fine-grained visual classification of aircraft. *arXiv preprint arXiv:1306.5151* (2013).
- [27] Cristina Menghini, Andrew Delworth, and Stephen Bach. 2023. Enhancing clip with clip: Exploring pseudolabeling for limited-label prompt tuning. *Advances in Neural Information Processing Systems* 36 (2023), 60984–61007.
- [28] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Baolin Wu, Andrew Y Ng, et al. 2011. Reading digits in natural images with unsupervised feature learning. In *NIPS workshop on deep learning and unsupervised feature learning*, Vol. 2011. Granada, 4.
- [29] Maria-Elena Nilsback and Andrew Zisserman. 2008. Automated Flower Classification over a Large Number of Classes. In *Indian Conference on Computer Vision, Graphics and Image Processing*.
- [30] Omkar M. Parkhi, Andrea Vedaldi, Andrew Zisserman, and C. V. Jawahar. 2012. Cats and Dogs. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [31] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems* 32 (2019).
- [32] Yuxin Peng, Yunzhen Zhao, and Junchao Zhang. 2018. Two-stream collaborative learning with spatial-temporal attention for video classification. *IEEE Transactions on Circuits and Systems for Video Technology* 29, 3 (2018), 773–786.
- [33] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*. PMLR, 8748–8763.
- [34] Sebastian Ruder. 2016. An overview of gradient descent optimization algorithms. *arXiv preprint arXiv:1609.04747* (2016).
- [35] Mehdi Sajjadi, Mehran Javanmardi, and Tolga Tasdizen. 2016. Regularization with stochastic transformations and perturbations for deep semi-supervised learning. *Advances in Neural Information Processing Systems* 29 (2016).
- [36] Rico Sennrich. 2015. Neural machine translation of rare words with subword units. *arXiv preprint arXiv:1508.07909* (2015).
- [37] 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 (2020), 596–608.
- [38] Andreas Steiner, Alexander Kolesnikov, Xiaohua Zhai, Ross Wightman, Jakob Uszkoreit, and Lucas Beyer. 2021. How to train your vit? data, augmentation, and regularization in vision transformers. *arXiv preprint arXiv:2106.10270* (2021).
- [39] Jong-Chyi Su and Subhansu Maji. 2021. The semi-supervised inaturalist challenge at the fgvc8 workshop. *arXiv preprint arXiv:2106.01364* (2021).
- [40] Antti Tarvainen. 2017. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Advances in Neural Information Processing Systems* 30 (2017).
- [41] Jesper E Van Engelen and Holger H Hoos. 2020. A survey on semi-supervised learning. *Machine Learning* 109, 2 (2020), 373–440.
- [42] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in Neural Information Processing Systems* 30 (2017).
- [43] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. 2011. *The Caltech-UCSD Birds-200-2011 Dataset*. Technical Report CNS-TR-2011-001. California Institute of Technology.
- [44] Dong Wang, Yuan Zhang, Kexin Zhang, and Liwei Wang. 2020. Focalmix: Semi-supervised learning for 3d medical image detection. In *Proceedings of the IEEE/CVF*

- Conference on Computer Vision and Pattern Recognition*. 3951–3960.
- [45] Yidong Wang, Hao Chen, Yue Fan, Wang Sun, Ran Tao, Wenxin Hou, Renjie Wang, Linyi Yang, Zhi Zhou, Lan-Zhe Guo, et al. 2022. Usb: A unified semi-supervised learning benchmark for classification. *Advances in Neural Information Processing Systems* 35 (2022), 3938–3961.
- [46] Yidong Wang, Hao Chen, Qiang Heng, Wenxin Hou, Yue Fan, Zhen Wu, Jindong Wang, Marios Savvides, Takahiro Shinozaki, Bhiksha Raj, et al. 2022. Freematch: Self-adaptive thresholding for semi-supervised learning. *arXiv preprint arXiv:2205.07246* (2022).
- [47] Ross Wightman. 2019. PyTorch Image Models. <https://github.com/huggingface/pytorch-image-models>. doi:10.5281/zenodo.4414861
- [48] Jianxiong Xiao, James Hays, Krista A Ehinger, Aude Oliva, and Antonio Torralba. 2010. Sun database: Large-scale scene recognition from abbey to zoo. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. IEEE, 3485–3492.
- [49] Xiangli Yang, Zixing Song, Irwin King, and Zenglin Xu. 2022. A survey on deep semi-supervised learning. *IEEE Transactions on Knowledge and Data Engineering* 35, 9 (2022), 8934–8954.
- [50] Hantao Yao, Rui Zhang, and Changsheng Xu. 2023. Visual-language prompt tuning with knowledge-guided context optimization. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*. 6757–6767.
- [51] Tao Yu, Zhihe Lu, Xin Jin, Zhibo Chen, and Xinchao Wang. 2023. Task residual for tuning vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 10899–10909.
- [52] Sergey Zagoruyko. 2016. Wide residual networks. *arXiv preprint arXiv:1605.07146* (2016).
- [53] Maxime Zanella and Ismail Ben Ayed. 2024. Low-rank few-shot adaptation of vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 1593–1603.
- [54] Jiahua Zhang, Qi Wei, Feng Liu, and Lei Feng. 2024. Candidate pseudolabel learning: Enhancing vision-language models by prompt tuning with unlabeled data. *arXiv preprint arXiv:2406.10502* (2024).
- [55] Renrui Zhang, Wei Zhang, Rongyao Fang, Peng Gao, Kunchang Li, Jifeng Dai, Yu Qiao, and Hongsheng Li. 2022. Tip-adapter: Training-free adaption of clip for few-shot classification. In *European Conference on Computer Vision*. Springer, 493–510.
- [56] Kaiyang Zhou, Ziwei Liu, Yu Qiao, Tao Xiang, and Chen Change Loy. 2022. Domain generalization: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2022).
- [57] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. 2022. Conditional prompt learning for vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 16816–16825.
- [58] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. 2022. Learning to prompt for vision-language models. *International Journal of Computer Vision* 130, 9 (2022), 2337–2348.
- [59] Kaiyang Zhou, Yongxin Yang, Yu Qiao, and Tao Xiang. 2021. Domain adaptive ensemble learning. *IEEE Transactions on Image Processing* 30 (2021), 8008–8018.
- [60] Beier Zhu, Yulei Niu, Yucheng Han, Yue Wu, and Hanwang Zhang. 2023. Prompt-aligned gradient for prompt tuning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 15659–15669.