

EfficientFSL: Enhancing Few-Shot Classification via Query-Only Tuning in Vision Transformers

Wenwen Liao¹, Hang Ruan¹, Jianbo Yu^{2*}, Bing Song³, Yuansong Wang⁴, Xiaofeng Yang²

¹College of Intelligent Robotics and Advance Manufacturing, Fudan University, Shanghai China

²School of Microelectronics, Fudan University, Shanghai China

³School of Information Science and Engineering, East China University of Science and Technology, Shanghai China

⁴Tsinghua Shenzhen International Graduate School, Tsinghua University, Shenzhen China

Abstract

Large models such as Vision Transformers (ViTs) have demonstrated remarkable superiority over smaller architectures like ResNet in few-shot classification, owing to their powerful representational capacity. However, fine-tuning such large models demands extensive GPU memory and prolonged training time, making them impractical for many real-world low-resource scenarios. To bridge this gap, we propose **EfficientFSL**, a query-only fine-tuning framework tailored specifically for few-shot classification with ViT, which achieves competitive performance while significantly reducing computational overhead. EfficientFSL fully leverages the knowledge embedded in the pre-trained model and its strong comprehension ability, achieving high classification accuracy with an extremely small number of tunable parameters. Specifically, we introduce a lightweight trainable Forward Block to synthesize task-specific queries that extract informative features from the intermediate representations of the pre-trained model in a query-only manner. We further propose a Combine Block to fuse multi-layer outputs, enhancing the depth and robustness of feature representations. Finally, a Support-Query Attention Block mitigates distribution shift by adjusting prototypes to align with the query set distribution. With minimal trainable parameters, EfficientFSL achieves state-of-the-art performance on four in-domain few-shot datasets and six cross-domain datasets, demonstrating its effectiveness in real-world applications.

Introduction

Despite the remarkable progress of deep learning in computer vision, it typically relies on large amounts of labeled data. In contrast, Few-shot Learning (FSL) (Finn, Abbeel, and Levine 2017; Munkhdalai et al. 2018; Antoniou, Edwards, and Storkey 2018) enables the construction of models that can recognize and classify previously unseen categories using only a very limited number of samples. Since the model must learn discriminative features from minimal labeled data (Zhang et al. 2022; Afrasiyabi et al. 2022), the representation capability of the model becomes crucial.

To enhance this capability, many recent approaches (Li et al. 2024; Zhang et al. 2024; Li, Wang, and Li 2024) have replaced CNN-based ResNet models with Vision Transformers (ViT) (Dosovitskiy et al. 2020) for feature extraction, achieving significant performance gains. A common strategy involves pre-training on a general, large-scale dataset, followed by adaptation to few-shot tasks. The most straightforward way to adapt is to fully fine-tune all parameters of the pre-trained model. However, this requires storing a complete set of model parameters for each individual task, leading to considerable storage overhead (Zhai et al. 2022).

To improve storage efficiency, researchers have explored Parameter-Efficient Transfer Learning (PETL) methods (Chen et al. 2022a; Hu et al. 2022; Jia et al. 2022; Jie and Deng 2023; Jie, Wang, and Deng 2023), which adapt large pre-trained models to downstream tasks by tuning only a small portion of parameters. Typically, these methods insert small modules into the frozen pre-trained model and only update the parameters of these added components. However, these methods inevitably modify the feature flow and remain coupled with the backbone weights, risking overfitting and reduced generalization, especially in data-limited settings like FSL. To overcome this, we propose a novel query-only paradigm **EfficientFSL**, a fine-tuning method tailored specifically for few-shot classification with ViT. The core distinction from existing methods lies in EfficientFSL’s approach: instead of modifying the backbone, it freezes it entirely and introduces a lightweight querying module to selectively extract task-relevant information.

As illustrated in Figure 1, EfficientFSL takes the intermediate representations from a pre-trained model as input and integrates them with task-specific knowledge through a stack of modular components. First, EfficientFSL follows a decoupling design, employing a lightweight Forward Block that consists of a trainable Active Block and a Frozen Block. The Active Block focuses on learning task-specific knowledge by generating adaptive queries. These queries guide the model to attend to relevant information. The following Frozen Block provides general knowledge by reusing and freezing the intermediate representations from the pre-trained backbone, effectively preserving prior knowledge while avoiding unnecessary parameter updates. Next, building on the rich hierarchical features from the Frozen Block, we introduce the Combine Block, which adaptively ag-

*Corresponding to jb_yu@fudan.edu.cn

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

gregates multi-layer features using the final layer’s output as guidance. Throughout the architecture, Bottleneck-Structured Fully Connected Layers are utilized to reduce the number of trainable parameters.

Importantly, to address the distribution shift between the support set and the query set, Support-Query Attention Block (SQ Attention Block) is proposed. This component adjusts the position of prototypes to better align with the center of the corresponding query distribution. EfficientFSL outperforms recent approaches across all benchmarks.

The contributions are summarized as follows:

- **A Novel PEFT Framework for FSL.** We propose **EfficientFSL**, a parameter-efficient fine-tuning framework for few-shot classification. It effectively integrates pre-trained representations with task-specific knowledge while significantly reducing the number of trainable parameters.
- **Enhance Support-Query Alignment.** To mitigate the distribution shift between the support and query sets and improve generalization in few-shot scenarios, we propose the SQ Attention Block, which dynamically adjusts prototype positions by aligning them with the query set distribution.
- **Outstanding Performance.** Extensive experiments on 4 in-domain few-shot datasets and 6 cross-domain datasets, validate the effectiveness of our method and its individual components.

Related Works

Few-Shot Learning

Few-shot learning (FSL) aims to learn classifiers for new categories using only a few labeled examples. Studies on the few-shot classification problem can be roughly divided into the following categories: (i) Optimization-based methods leverage meta-learning to acquire well-initialized parameters that enable adaptation to new classes with minimal optimization steps (Finn, Abbeel, and Levine 2017; Ravi and Larochelle 2017; Rusu et al. 2019; Baik, Hong, and Lee 2020). (ii) Data augmentation-based approaches adaptively predict classifier weights for new categories based on the feature embeddings of new examples (Wang et al. 2018; Zhang, Zhang, and Koniusz 2019; Wang et al. 2020). (iii) Metric-based methods focus on learning a metric space in which a query example can be classified using a nearest neighbor approach (Snell, Swersky, and Zemel 2017; Vinyals et al. 2016; Koch, Zemel, and Salakhutdinov 2015; Chen et al. 2020).

Among these, metric-based methods have emerged as particularly appealing due to their simplicity, efficiency in low-data regimes. These approaches aim to learn an embedding space where samples from the same category are close together, while samples from different categories are farther apart (Fu et al. 2022). Typically, this involves projecting all input instances into vectors of fixed dimensionality, with predictions for query samples made using methods such as nearest neighbor classifiers (Vinyals et al. 2016), robust category weights (Tang et al. 2022), parameterized metrics

(Sung et al. 2018), or other task-specific distance metrics (Liu et al. 2020).

Parameter-Efficient Transfer Learning

Updating all parameters of large-scale pre-trained models can lead to substantial computational overhead, thereby hindering the efficiency of rapid model adaptation and deployment. To address this issue, researchers have proposed Parameter-Efficient Transfer Learning (PETL), which enables effective adaptation to downstream tasks by updating only a small subset of parameters while keeping the majority of the model frozen.

When applied to large vision models such as ViT, PETL offers an efficient and cost-effective solution. Common PETL techniques include the use of soft prompts (Li and Liang 2021; Jia et al. 2022), lightweight adapter modules (Houlsby et al. 2019; Chen et al. 2022a; He et al. 2021), and low-rank matrix decomposition (Zhang et al. 2023; Valipour et al. 2022; Hu et al. 2022). Notably, recent studies have demonstrated that in few-shot settings with large language models, parameter-efficient fine-tuning (PEFT) can achieve performance comparable to or even surpassing that of full fine-tuning (Liu et al. 2022; Yang, Liu, and Wei 2024). This reveals a natural synergy between PETL strategies and FSL.

Preliminaries

Problem Setting

In few-shot classification, the dataset is divided into a train set \mathcal{D}^{train} , and a test set, \mathcal{D}^{test} , where there is no overlap between the classes of the train and test sets. Each set is further split into two subsets. One is support set $\mathcal{S} = \{(\mathbf{x}_i, y_i)\}_{i=0}^{N \times K}$, containing a few labeled examples for each class used for forming the prototypes of categories. The other subset, the query set $\mathcal{Q} = \{(\mathbf{x}_i, y_i)\}_{i=0}^{N \times Q}$, consists of unlabeled examples used to compute their distances to the prototypes for classification. Here, N denotes the number of classes per episode, and K represents the number of support examples per class, and Q the number of query examples per class. This setting is commonly known as the N -way K -shot classification problem, where \mathbf{x}_i is the input image and y_i its label.

Prototypical Network

A representative and widely adopted metric-based method is the Prototypical Network (PN) (Snell, Swersky, and Zemel 2017), which has gained popularity for its ability to learn discriminative and compact class representations in few-shot classification tasks.

We adopt PN as the classification head. The feature extractor f is trained to construct an embedding space where samples from the same class are mapped close to one another, while those from different classes are pushed farther apart. For each class c , the prototype \mathbf{s}^c is computed as the mean of the feature representations of its support samples, as defined by:

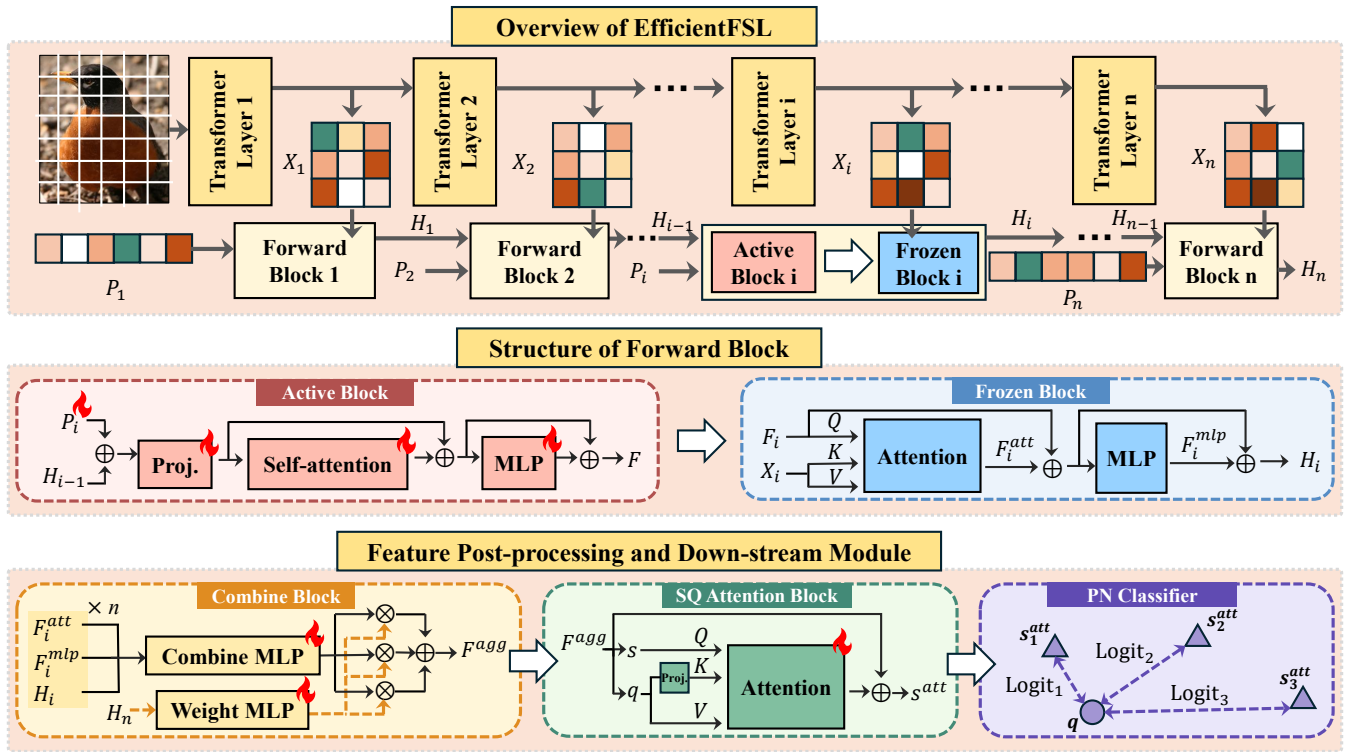


Figure 1: An overview of our pipeline. EfficientFSL takes intermediate representations from a pre-trained model as input and integrates task-specific knowledge through the Forward Block, Combine Block, and SQ Attention Block, followed by classification using a PN classifier.

$$s^c = \frac{1}{|\mathcal{D}^c|} \sum_{(x_i, y_i) \in \mathcal{D}^c} f(x_i), \quad (1)$$

where \mathcal{D}^c denotes the set of support examples belonging to class c . Classification is then performed by measuring the distance between each query sample and the prototypes of all classes, enabling efficient and effective prediction.

Methodology

Overview

From the upper part of Figure 1, it can be observed that EfficientFSL takes the intermediate features from each layer of the pre-trained ViT as the input for each block in EfficientFSL.

Each block consists of an active sub-block and a frozen sub-block: The Active Block is a lightweight module designed to learn dataset-specific knowledge, adapt to specific downstream tasks, and particularly enhance few-shot classification performance. The Frozen Block utilizes the tokens generated by the Active Block as queries, while using the intermediate features from the pre-trained ViT as keys and values. By reusing the original model’s weights, it preserves the inherent capability of ViT. Afterward, a Combine Block is employed to adaptively fuse the rich features extracted from each Frozen Block using adaptive weights, resulting in a unified feature output.

To reduce the distribution bias between the support set and query set, the SQ Attention Block is introduced, which refines the distribution of prototypes. Finally, PN serves as the classification head to perform classification.

Forward Block

Active Block. The Active Block is a lightweight network used to adapt to downstream tasks. The output from the previous layer H_{i-1} is added with a trainable prompt P_i , and then feed it into a bottleneck-structured projection layer to obtain the output Z_i . These per-layer added prompts and the projection layer $\text{Proj}(\cdot)$ enhance the Active Block’s ability to capture task-specific knowledge.

$$Z_i = \text{Proj}(H_{i-1} + P_i). \quad (2)$$

Inspired by the Transformer structure, Z_i interacts through a self-attention layer and is further passed into a Multilayer Perceptron (MLP). Notably, to maintain parameter efficiency, we extensively use bottleneck-structured fully connected layers here to reduce the number of parameters. In the attention layer $\text{Att}(\cdot)$, Z_i passes through three bottleneck-structured projection layers to produce the query Q_i^A , key K_i^A , and value V_i^A . Then, the computation in the attention layer is as follows, controlled by a scale factor ξ :

$$Z_i' = \xi \cdot \text{Att}(Q_i^A, K_i^A, V_i^A) + Z_i. \quad (3)$$

Subsequently, Z'_i is passed through a layer normalization $\text{LN}(\cdot)$ and a fully connected layer $\text{MLP}(\cdot)$ to obtain F_i , controlled by another scale factor ζ :

$$F_i = \zeta \cdot \text{MLP}(\text{LN}(Z'_i)) + Z'_i. \quad (4)$$

Since the Active Block is trained from scratch on the downstream dataset, its output F_i captures task-specific knowledge. In the next section, F_i will be used in a query-only manner.

Frozen Block. The Frozen Block reuses and freezes the pre-trained parameters for feature interaction, requiring no additional training, which ensures the efficiency of our training process. Its input consists of F_i obtained from the Active Block, as well as the outputs from each layer in the ViT forward propagation process $\{X_1, X_2, \dots, X_n\}$, where n is the number of layers in the pre-trained ViT model.

In the attention layer $\text{Att}'(\cdot)$, we use F_i to obtain the query Q_i^F , and X_i to obtain the key K_i^F and value V_i^F . The subsequent attention computation is as follows:

$$F_i^{att} = \text{Att}'(Q_i^F, K_i^F, V_i^F) + F_i. \quad (5)$$

Next, F_i^{att} passes through a layer normalization $\text{LN}(\cdot)$ and a fully connected layer $\text{MLP}'(\cdot)$ to obtain F_i^{mlp} :

$$F_i^{mlp} = \text{MLP}'(\text{LN}(F_i^{att})). \quad (6)$$

Finally, the feature H_i is obtained as:

$$H_i = F_i^{mlp} + F_i^{att}. \quad (7)$$

To maximize the potential of the Frozen Block, in addition to the final output H_i , we also leverage the features F_i^{att} and F_i^{mlp} from the attention module and the fully connected layer for classification.

Combine Block

Feature Alignment using a Shared Combine MLP. Through the Forward Block, hierarchical features are obtained as follows:

$$\{F_1^{att}, F_1^{mlp}, H_1, F_2^{att}, F_2^{mlp}, H_2, \dots, F_n^{att}, F_n^{mlp}, H_n\}.$$

First, we align all features using a shared combine MLP. Specifically, we employ a bottleneck-structured fully connected layer and share the projection layers across all features to improve parameter efficiency. The transformation for each feature is given by:

$$\begin{cases} \hat{F}_i^{att} = \text{MLP}_{\text{Combine}}^{\text{shared}}(F_i^{att}), \\ \hat{F}_i^{mlp} = \text{MLP}_{\text{Combine}}^{\text{shared}}(F_i^{mlp}), \\ \hat{H}_i = \text{MLP}_{\text{Combine}}^{\text{shared}}(H_i), \end{cases} \quad (8)$$

where \hat{F}_i^{att} , \hat{F}_i^{mlp} , and \hat{H}_i are the projected representations of the original features.

Adaptive Feature Aggregation. To compute feature weights, we leverage an adaptive weighting mechanism based on the final-layer hidden state H_n . We use a Weight MLP to extract the weights from H_n , and then apply them to all features through a weighted sum. The weighting process is defined as:

$$w_i^{att}, w_i^{mlp}, w_i^H = \sigma(\text{MLP}_{\text{weight}}(H_n)), \forall i \in \{1, \dots, n\} \quad (9)$$

$$F^{\text{agg}} = \sum_{i=1}^n \left(w_i^{att} \cdot \hat{F}_i^{att} + w_i^{mlp} \cdot \hat{F}_i^{mlp} + w_i^H \cdot \hat{H}_i \right), \quad (10)$$

where $w_i^{att}, w_i^{mlp}, w_i^H$ are the learned adaptive weights for each type of feature, and F^{agg} is the aggregated feature representation used for classification; σ denotes the softmax function.

Support-Query Attention Block

Support and query samples are passed through the above blocks to extract their features and compute class prototypes following the PN method (Eq. (1)). We denote the prototype of a batch's support set as s and the feature of a query sample as q . Due to slight variations between the images in the support set and the query set (such as differences in background, lighting, or shooting angles), the prototypes generated from the support set may have a certain distance from the corresponding query set of the same class. This distribution bias leads to classification inaccuracy.

To mitigate this distribution bias, we propose the SQ Attention Block to adjust the position of the prototypes, making them closer to the distribution center of the corresponding query set, as shown in Figure 3. Besides, before computing the correlation between prototype s and query q , we apply a learnable projection $\text{Proj}(\cdot)$ to q for adaptive, class-aware alignment, enhancing robustness under distribution shift. The formulation is as follows:

$$s^{att} = \alpha(s \cdot \text{Proj}(q)^T) \cdot q + (1 - \alpha) \cdot s. \quad (11)$$

Thus, the updated prototype s^{att} is obtained. Finally, the cosine similarity between each query sample and s^{att} is computed to produce the predicted labels.

Experiments

Datasets

Our work conducts experiments on four common FSL benchmarks, i.e., miniImageNet (Vinyals et al. 2016), tiered-ImageNet (Ren et al. 2018), CIFAR (Krizhevsky, Hinton et al. 2009) and FC100 (Oreshkin, Rodríguez López, and Lacoste 2018), in which miniImageNet and tieredImageNet are derivatives of the ImageNet-1K dataset (Russakovsky and et al. 2015). And we also conduct cross-domain few shot learning test on six datasets CUB (Wah et al. 2011), Stanford Cars (Krause et al. 2013), Places (Zhou et al. 2017), Plantae (Van Horn et al. 2018), EuroSAT (Helber et al. 2019) and CropDiseases (Mohanty, Hughes, and Salathé 2016).

Setting

We utilize three backbone models: ViT-S/16 and ViT-B/16 pre-trained on the ImageNet-1K dataset, and ViT-B/16 pre-trained on the ImageNet-21K dataset. The ξ , ζ , and α are tuned within the range $\{0.1, 1\}$. The hidden size of all bottleneck projections is set to 48, while the hidden size of the attention bottleneck projections for the query, key, and value in Active Block is set to 8. Following common practice (Jie and Deng 2023), we adopt AdamW (Loshchilov and Hutter 2017) as the optimizer and use a cosine learning rate scheduler. We set the batch size to 64 and optimized the model using a learning rate of 0.0001 for 5 epochs. For evaluation, we tested the model on 320 batches and reported the average performance along with the 95% confidence interval. For data augmentation, we first resize the input image to 256×256 and then apply a center crop to 224×224 . All experiments are conducted on a single NVIDIA V100 GPU.

Baseline

State-of-the-art (SOTA) FSL methods are applied to the few-shot classification task to facilitate a thorough comparison, including FewTURE(Hiller et al. 2022), MetaFormer-A (Yang, Yao, and Wei 2024), FewVS (Li, Wang, and Li 2024), SCAM-Net(Di et al. 2025), SP(Chen et al. 2023), KTPP (Li et al. 2024), SemFew (Zhang et al. 2024) as well as cross-domain SOTA methods like LDPNet(Zhou et al. 2023), FLoR (Zou et al. 2024), StyleAdv (Fu et al. 2023), AMT-FT (Yang, Liu, and Wei 2024). To further demonstrate the effectiveness of EfficientFSL under the PETL setting, we also compare against several commonly used PETL baselines Adapter (Chen et al. 2022b), AdaptFormer (Chen et al. 2022a), and LoRA (Hu et al. 2022).

Experimental Results and Analysis

Comparison with SOTA Methods. In Table 1, we evaluate our method on miniImageNet, tieredImageNet, CIFAR, and FC100 datasets, covering both the 5-way 1-shot and 5-shot settings. Our method demonstrates consistently superior performance across all settings and backbones, while using significantly fewer parameters.

Specifically, when using ViT-S, EfficientFSL surpasses the fully fine-tuned ViT-S and all recent works using the same pre-trained model. In particular, EfficientFSL only uses 1.25M parameters, which is substantially fewer than other ViT-based approaches, highlighting its efficiency. With ViT-B as the backbone and pre-trained on ImageNet-21K, EfficientFSL further improves the performance.

To compare with SOTAs on cross-domain few-shot benchmarks, we evaluate the performance on six cross-domain few-shot benchmarks under the 5-way 1-shot and 5-way 5-shot scenario, using meta-learning on miniImageNet with a pre-trained ViT-S/16 model on ImageNet1K. As shown in Figure 2, our approach outperforms existing SOTAs across all benchmarks, demonstrating its generalization capabilities in various domains.

Overall, these results underscore the effectiveness of EfficientFSL in various challenging FSL tasks, showing its superior generalization capabilities compared to recent SOTAs.




Figure 2: Cross-domain Results under Different Shots.

Comparison with mainstream PETL approaches. To ensure a fair and rigorous comparison, we used the same backbone network on FC100 for all methods and carefully tuned their hyperparameters to maintain a comparable number of trainable parameters.

As Table 2 shows, EfficientFSL achieves significantly higher accuracy than the other compared methods on both ViT-S and ViT-B, demonstrating its superior performance. Concurrently, it also leads comprehensively in efficiency, achieving faster training and inference speeds and highly competitive memory usage.

Impact of Removing Training Modules on Parameter Consumption and Performance. In Table 3, removing different components from the Active Block and Combine Block has a significant impact on both parameter consumption and performance.

When all modules are retained, the model achieves the best performance. Removing the projection layer (Proj.) reduces the parameter count but leads to a sharp drop in performance, indicating the essential role of the projection layer in ensuring high-quality feature representations. Removing the attention and MLP modules (Att & MLP) in the Active Block reduces the parameter count by more than half, and performance drops. This shows that the attention mechanism and MLP modules are critical for capturing task-specific information and enhancing the model’s performance. On the other hand, removing the Combine Block modules results in a smaller reduction in parameters, and the accuracy drop is relatively minor.

These results highlight the importance of the projection layer and attention mechanism in the Active Block, as well as the crucial role of the Combine Block in efficient feature fusion, indicating its effectiveness in integrating multi-scale features.

Ablation Study on Active and Combine Block Components. Table 4 analyzes the impact of independently

Models	Backbone	Params	miniImageNet		tieredImageNet		CIFAR		FC100	
			5w-1s	5w-5s	5w-1s	5w-5s	5w-1s	5w-5s	5w-1s	5w-5s
Full Fine-Tuning	ViT-S	21.7M	89.41 ± 0.88	95.59 ± 0.36	77.87 ± 1.35	89.13 ± 0.74	84.86 ± 1.11	92.30 ± 0.62	52.93 ± 1.29	66.70 ± 1.01
FewTURE (NeurIPS'22)	ViT-S	21.7M	68.02 ± 0.88	84.51 ± 0.53	72.96 ± 0.92	86.43 ± 0.67	72.80 ± 0.88	86.14 ± 0.64	46.20 ± 0.79	63.14 ± 0.73
MetaF. (ICML'24)	ViT-S	24.5M	84.78 ± 0.79	91.39 ± 0.42	88.38 ± 0.78	93.37 ± 0.45	88.34 ± 0.76	92.21 ± 0.59	58.04 ± 0.99	70.80 ± 0.76
FewVS (MM'24)	ViT-S	21.7M	86.80 ± 0.28	90.32 ± 0.22	87.87 ± 0.36	92.27 ± 0.26	85.63 ± 0.37	90.73 ± 0.33	61.01 ± 0.40	70.37 ± 0.39
SCAM-Net (arXiv'25)	ViT-S	21.7M	75.93 ± 0.45	89.75 ± 0.22	78.89 ± 0.52	88.67 ± 0.65	79.45 ± 0.53	90.90 ± 0.45	47.48 ± 0.32	65.85 ± 0.75
Ours(1K)	ViT-S	1.25M	97.40 ± 0.47	99.05 ± 0.14	89.72 ± 0.96	95.41 ± 0.48	88.82 ± 0.86	94.60 ± 0.42	69.94 ± 1.16	81.68 ± 0.95
Full Fine-Tuning	ViT-B	85.8M	90.51 ± 0.82	95.01 ± 0.37	80.17 ± 1.16	91.15 ± 0.58	85.83 ± 1.18	92.88 ± 0.59	52.86 ± 1.28	65.03 ± 0.99
SP (CVPR'23)	Visformer-T	10.3M	72.31 ± 0.40	83.42 ± 0.30	78.03 ± 0.46	88.55 ± 0.32	82.18 ± 0.40	88.24 ± 0.32	48.53 ± 0.38	61.55 ± 0.41
KTPP (MM'24)	Visformer-T	11.1M	76.71 ± 0.37	86.46 ± 0.27	80.80 ± 0.43	90.01 ± 0.29	83.63 ± 0.57	90.19 ± 0.30	51.59 ± 0.40	65.18 ± 0.40
SemFew (CVPR'24)	Swin-T	88.0M	78.94 ± 0.66	86.49 ± 0.50	82.37 ± 0.77	89.89 ± 0.52	84.34 ± 0.67	89.11 ± 0.54	54.27 ± 0.77	65.02 ± 0.72
Ours(1K)	ViT-B	2.48M	97.57 ± 0.48	98.96 ± 0.14	87.63 ± 1.08	93.13 ± 0.62	85.25 ± 1.21	92.64 ± 0.54	72.60 ± 1.48	80.74 ± 1.04
Ours(21K)	ViT-B	2.48M	98.34 ± 0.30	99.12 ± 0.13	93.27 ± 0.75	96.78 ± 0.33	93.25 ± 0.69	97.28 ± 0.31	80.13 ± 1.26	88.81 ± 0.81

Table 1: Comparison of our method with SOTAs on miniImageNet, tieredImageNet, CIFAR, and FC100 datasets. The ‘‘Params’’ column indicates the number of parameters (in millions). Ours (1K) refers to the full results with pre-trained ViT on ImageNet1K, while Ours (21K) uses ImageNet21K.

Method	Acc ↑	TT ↓	PM ↓	IS ↑
Backbone: ViT-S				
Adapter	62.80	26.68	0.50	380.46
AdaptFormer	64.13	38.22	0.57	274.53
LoRA	69.04	34.02	0.49	298.14
EfficientFSL	69.94	23.61	0.49	391.57
Backbone: ViT-B				
Adapter	67.48	74.19	1.09	134.21
AdaptFormer	68.83	76.87	1.19	129.98
LoRA	74.19	76.35	1.10	131.84
EfficientFSL	80.13	59.62	1.08	134.29

Table 2: Comparison of different methods with ViT-S and ViT-B. Acc = Accuracy (%), TT = Train Time (s/epoch), PM = Peak Memory (GB), IS = Inference Speed (images/s).

removing key components from the Active and Combine Blocks on parameter count and performance.

For the Active Block, removing p_i causes a significant performance drop with negligible parameter reduction, confirming its critical role in guiding learning. Ablating Att slightly reduces accuracy alongside parameter savings, while removing MLP leads to a larger performance decrease, underscoring its importance in refining features. Overall, all three modules are essential for efficient feature learning.

For the Combine Block, ablating F_i^{att} has limited effect, mainly on 1-shot tasks; removing F_i^{mlp} yields a larger drop, highlighting its role in decision adjustment; removing H_i causes the largest performance loss, showing its key contribution to knowledge retention. While parameter reduction is minimal, multi-scale outputs remain important for performance. These highlight the contribution of each component and provide guidance for future model optimization.

Proj.	Att & MLP	Combine	Params (M)	1-shot	5-shot
✗	✓	✓	1.58	51.15	68.95
✓	✗	✓	1.05	72.15	87.81
✓	✓	✗	2.37	75.55	88.60
✓	✓	✓	2.48	80.13	88.81

Table 3: Impact of Removing Training Modules on Parameter Consumption and Performance.

Module	Setting	Params (M)	1-shot	5-shot
Active Block	w/o p_i	2.44	39.36	60.99
	w/o Att	1.95	75.98	87.83
	w/o MLP	1.56	55.09	69.99
Combine Block	w/o F_i^{att}	2.47	78.31	89.71
	w/o F_i^{mlp}	2.47	73.26	87.52
	w/o H_i	2.47	77.77	89.20
EfficientFSL	Complete Model	2.48	80.13	88.81

Table 4: Ablation Study on the Internal Components of Active Block and Combine Block.

Impact of SQ Attention Block. As shown in Table 5, removing the SQ Attention Block leads to a decrease in accuracy, demonstrating the substantial impact of the SQ Attention Block on performance. Moreover, removing projection on q results in a slight yet consistent performance drop. This confirms that adjusting q with a class-aware projection layer prior to the correlation calculation with s alleviates the adverse effects of inter-class cluster overlap.

Figure 3 shows the t-SNE visualization of one batch of test samples from the experiment. The figure clearly demonstrates that SQ Attention Block, through the attention mechanism, brings the prototypes closer to the distribution center of the corresponding class’s query samples. This highlights that the SQ Attention Block significantly reduces the distri-

Setting	Params (M)	Mini		Tiered		CIFAR		FC100	
		1shot	5shot	1shot	5shot	1shot	5shot	1shot	5shot
w/o SQA	2.48	95.95	98.97	89.78	96.61	90.15	96.99	75.42	87.82
w/o proj	2.48	98.34	99.12	93.27	96.78	93.25	97.28	80.13	88.81
Ours	2.55	98.49	99.23	93.85	96.89	93.34	97.28	80.20	89.53

Table 5: Ablation Analysis of the SQ Attention Block.




Figure 3: t-SNE visualization of a 5-way 1-shot 15-query task on the FC100 dataset. Squares represent prototypes computed by simply averaging support features (s^c); triangles represent prototypes computed by SQ Attention Block (s^{att}); circles denote query samples.

bution shift between the support set and the query set.

Impact of Different Feature Aggregation Methods on Parameter Size and Few-Shot Performance. Table 6 compares three feature aggregation strategies in the Combine Block and their impact on parameter size and few-shot performance. Simple averaging treats all layers equally, while fixed weights introduce input-invariant learnable parameters with slight improvements. In contrast, conditional weights dynamically generate aggregation weights based on the final-layer output, achieving the best performance with minimal parameter increase (2.48M).

These results show that adaptive, condition-aware weighting significantly improves generalization in FSL, especially under low-data settings. With minimal parameter overhead, it provides an efficient yet effective alternative to simpler aggregation strategies.

From Chaos to Order: Evolution of the Representation Space through Module Integration. Figure 4 shows how the feature space is progressively refined as different modules are incrementally introduced. Starting from the pre-trained ViT-B/16 without fine-tuning (Figure 4a), query features are highly entangled with vague class boundaries. Adding trainable prompts (P_i) (Figure 4b) induces coarse clustering, showing initial task adaptation but with noticeable inter-class overlap. Introducing the projection layer ($\text{Proj}(\cdot)$) (Figure 4c) further improves separability and cluster compactness. With the integration of self-attention and MLP into an Active Block (Figure 4d), clusters become tighter and class boundaries clearer. Adding the Combine Block (Figure 4e) enhances intra-class cohesion and inter-class separation. Finally, the SQ Attention Block dynamically adjusts prototypes, yielding the full EfficientFSL

Combine Block	Params (M)	1shot	5shot
Simple Average	2.45	69.58 \pm 1.22	85.38 \pm 0.83
Fixed Weights	2.45	69.97 \pm 1.22	85.49 \pm 0.84
Conditional Weights	2.48	80.13 \pm 1.26	88.81 \pm 0.81

Table 6: Ablation of Different Feature Fusion Methods.

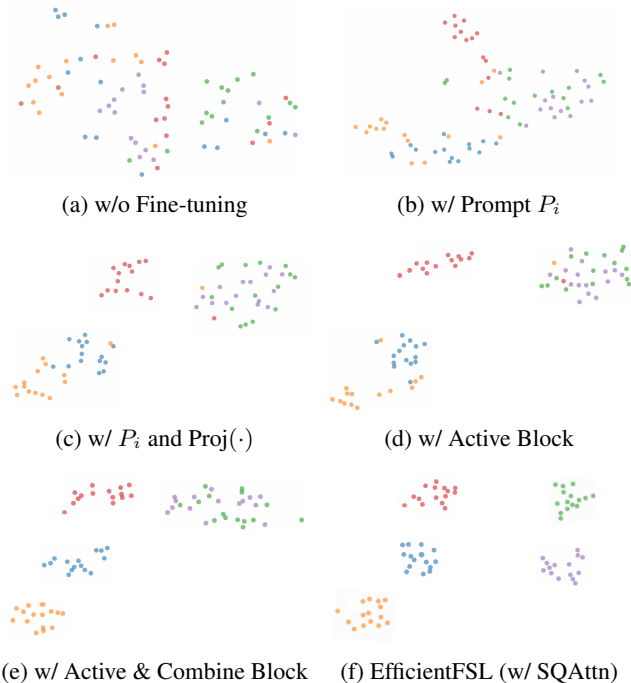


Figure 4: t-SNE visualizations of test query samples under six settings. (b)-(f) incrementally add components to the previous setting.

framework (Figure 4f), where features form clear, well-separated clusters. These results validate the effectiveness of our modular design in improving representation quality for FSL.

Conclusion

To bridge the gap between powerful pre-trained representations and the limited resources in FSL, we propose **EfficientFSL**, a parameter-efficient fine-tuning framework for few-shot classification. By introducing the Forward Block, Combine Block, and SQ Attention Block, EfficientFSL enables task-specific adaptation with minimal trainable parameters. Extensive experiments on both in-domain and cross-domain benchmarks show that EfficientFSL consistently outperforms existing methods, exhibiting strong performance and generalization capabilities. In particular, the SQ Attention Block effectively mitigates support-query distribution shift, thereby enhancing the effectiveness of prototype-based classification. Overall, EfficientFSL strikes a strong balance between efficiency, scalability, and performance, offering an effective solution for FSL applications.

Acknowledgments

This work is supported by the National Natural Science Foundation of China (Grant Number: 62203118) and the Fundamental Research Funds for the Central Universities (Grant Number: 2025SMECP012).

References

- Afrasiyabi, A.; Larochelle, H.; Lalonde, J.-F.; and Gagné, C. 2022. Matching feature sets for few-shot image classification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9014–9024.
- Antoniou, A.; Edwards, H.; and Storkey, A. 2018. How to train your MAML. *arXiv preprint arXiv:1810.09502*.
- Baik, S.; Hong, S.; and Lee, K. M. 2020. Learning to Forget for Meta-Learning. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2376–2384.
- Chen, J.; Zhan, L.-M.; Wu, X.-M.; and lai Chung, F. 2020. Variational metric scaling for metric-based meta-learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 3478–3485.
- Chen, S.; Ge, C.; Tong, Z.; Wang, J.; Song, Y.; Wang, J.; and Luo, P. 2022a. Adaptformer: Adapting vision transformers for scalable visual recognition. *Advances in Neural Information Processing Systems*, 35: 16664–16678.
- Chen, W.; Si, C.; Zhang, Z.; Wang, L.; Wang, Z.; and Tan, T. 2023. Semantic prompt for few-shot image recognition. *arXiv preprint arXiv:2303.14123*.
- Chen, Z.; Duan, Y.; Wang, W.; He, J.; Lu, T.; Dai, J.; and Qiao, Y. 2022b. Vision transformer adapter for dense predictions. *arXiv preprint arXiv:2205.08534*.
- Di, K.; Li, X.; Han, Y.; Li, Z.; Li, Q.; and Wu, X. 2025. Brain Inspired Adaptive Memory Dual-Net for Few-Shot Image Classification. *arXiv preprint arXiv:2503.07396*.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Finn, C.; Abbeel, P.; and Levine, S. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, 1126–1135. PMLR.
- Fu, L.; et al. 2022. Learning Robust Discriminant Subspace Based on Joint L2,p and L2,s-Norm Distance Metrics. *IEEE Transactions on Neural Networks and Learning Systems*, 33(1): 130–144.
- Fu, Y.; Xie, Y.; Fu, Y.; and Jiang, Y.-G. 2023. StyleAdv: Meta Style Adversarial Training for Cross-Domain Few-Shot Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 24575–24584.
- He, J.; Zhou, C.; Ma, X.; Berg-Kirkpatrick, T.; and Neubig, G. 2021. Towards a unified view of parameter-efficient transfer learning. *arXiv preprint arXiv:2110.04366*.
- Helber, P.; Bischke, B.; Dengel, A.; and Borth, D. 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): 2217–2226.
- Hiller, M.; Ma, R.; Harandi, M.; and Drummond, T. 2022. Rethinking generalization in few-shot classification. *Advances in neural information processing systems*, 35: 3582–3595.
- Houlsby, N.; Giurgiu, A.; Jastrzebski, S.; Morrone, B.; De Laroussilhe, Q.; Gesmundo, A.; Attariyan, M.; and Gelly, S. 2019. Parameter-efficient transfer learning for NLP. In *International conference on machine learning*, 2790–2799. PMLR.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W.; et al. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2): 3.
- Jia, M.; Tang, L.; Chen, B.-C.; Cardie, C.; Belongie, S.; Hariharan, B.; and Lim, S.-N. 2022. Visual prompt tuning. In *European conference on computer vision*, 709–727. Springer.
- Jie, S.; and Deng, Z.-H. 2023. Fact: Factor-tuning for lightweight adaptation on vision transformer. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, 1060–1068.
- Jie, S.; Wang, H.; and Deng, Z.-H. 2023. Revisiting the parameter efficiency of adapters from the perspective of precision redundancy. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 17217–17226.
- Koch, G.; Zemel, R.; and Salakhutdinov, R. 2015. Siamese neural networks for oneshot image recognition. In *ICML deep learning workshop*, volume 2. Lille.
- Krause, J.; Stark, M.; Deng, J.; and Fei-Fei, L. 2013. 3d object representations for fine-grained categorization. In *Proceedings of the IEEE international conference on computer vision workshops*, 554–561.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images.
- Li, W.; Wang, Q.; Zhao, P.; and Yin, Y. 2024. KNN Transformer with Pyramid Prompts for Few-Shot Learning. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 1082–1091.
- Li, X. L.; and Liang, P. 2021. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv preprint arXiv:2101.00190*.
- Li, Z.; Wang, Y.; and Li, K. 2024. FewVS: A Vision-Semantics Integration Framework for Few-Shot Image Classification. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 1341–1350.
- Liu, B.; et al. 2020. Negative Margin Matters: Understanding Margin in Few-Shot Classification. In *European Conference on Computer Vision (ECCV)*, 438–455.
- Liu, H.; Tam, D.; Muqeeth, M.; Mohta, J.; Huang, T.; Bansal, M.; and Raffel, C. A. 2022. Few-shot parameter-efficient fine-tuning is better and cheaper than in-context learning. *Advances in Neural Information Processing Systems*, 35: 1950–1965.

- Loshchilov, I.; and Hutter, F. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Mohanty, S. P.; Hughes, D. P.; and Salathé, M. 2016. Using deep learning for image-based plant disease detection. *Frontiers in Plant Science*, 7: 215232.
- Munkhdalai, T.; Yuan, X.; Mehri, S.; and Trischler, A. 2018. Rapid adaptation with conditionally shifted neurons. In *International conference on machine learning*, 3664–3673. PMLR.
- Oreshkin, B.; Rodríguez López, P.; and Lacoste, A. 2018. Tadam: Task dependent adaptive metric for improved few-shot learning. *Advances in neural information processing systems*, 31.
- Ravi, S.; and Larochelle, H. 2017. Optimization as a Model for Few-Shot Learning. In *International Conference on Learning Representations (ICLR)*.
- Ren, M.; Triantafillou, E.; Ravi, S.; Snell, J.; Swersky, K.; Tenenbaum, J. B.; Larochelle, H.; and Zemel, R. S. 2018. Meta-learning for semi-supervised few-shot classification. In *International Conference on Learning Representations (ICLR)*.
- Russakovsky, O.; and et al. 2015. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision*, 115(3): 211–252.
- Rusu, A. A.; Rao, D.; Sygnowski, J.; Vinyals, O.; Pascanu, R.; Osindero, S.; and Hadsell, R. 2019. Meta-Learning with Latent Embedding Optimization. In *International Conference on Learning Representations (ICLR)*.
- Snell, J.; Swersky, K.; and Zemel, R. 2017. Prototypical Networks for Few-Shot Learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 4077–4087.
- Sung, F.; Yang, Y.; Zhang, L.; Xiang, T.; Torr, P. H. S.; and Hospedales, T. M. 2018. Learning to Compare: Relation Network for Few-Shot Learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 1199–1208.
- Tang, H.; Yuan, C.; Li, Z.; and Tang, J. 2022. Learning Attention-Guided Pyramidal Features for Few-Shot Fine-Grained Recognition. *Pattern Recognition*, 130: 108792.
- Valipour, M.; Rezagholizadeh, M.; Kobzyev, I.; and Ghodsi, A. 2022. Dylora: Parameter efficient tuning of pre-trained models using dynamic search-free low-rank adaptation. *arXiv preprint arXiv:2210.07558*.
- 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 *Proceedings of the IEEE conference on computer vision and pattern recognition*, 8769–8778.
- Vinyals, O.; Blundell, C.; Lillicrap, T.; and Wierstra, D. 2016. Matching Networks for One Shot Learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 3630–3638.
- Wah, C.; Branson, S.; Welinder, P.; Perona, P.; and Belongie, S. 2011. The Caltech-UCSD Birds-200-2011 dataset.
- Wang, H.; Zhang, Y.; Ji, Z.; Pang, Y.; and Ma, L. 2020. Consensus-aware Visual-Semantic Embedding for Image-Text Matching. In *European Conference on Computer Vision (ECCV)*, volume 12369, 18–34.
- Wang, Y.-X.; Girshick, R. B.; Hebert, M.; and Hariharan, B. 2018. Low-shot Learning from Imaginary Data. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 7278–7286.
- Yang, X.; Liu, C.; and Wei, Y. 2024. Mixture of Adversarial LoRAs: Boosting Robust Generalization in Meta-Tuning. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Yang, X.; Yao, H.; and Wei, Y. 2024. One meta-tuned transformer is what you need for few-shot learning. In *Forty-first International Conference on Machine Learning*.
- Zhai, X.; Kolesnikov, A.; Houlsby, N.; and Beyer, L. 2022. Scaling vision transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 12104–12113.
- Zhang, C.; Cai, Y.; Lin, G.; and Shen, C. 2022. Deepemd: Differentiable earth mover’s distance for few-shot learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(5): 5632–5648.
- Zhang, H.; Xu, J.; Jiang, S.; and He, Z. 2024. Simple semantic-aided few-shot learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 28588–28597.
- Zhang, H.; Zhang, J.; and Koniusz, P. 2019. Few-shot Learning via Saliency-guided Hallucination of Samples. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2770–2779.
- Zhang, Q.; Chen, M.; Bukharin, A.; Karampatziakis, N.; He, P.; Cheng, Y.; Chen, W.; and Zhao, T. 2023. Adalora: Adaptive budget allocation for parameter-efficient fine-tuning. *arXiv preprint arXiv:2303.10512*.
- Zhou, B.; Lapedriza, A.; Khosla, A.; Oliva, A.; and Torralba, A. 2017. Places: A 10 million image database for scene recognition. *IEEE transactions on pattern analysis and machine intelligence*, 40(6): 1452–1464.
- Zhou, F.; Wang, P.; Zhang, L.; Wei, W.; and Zhang, Y. 2023. Revisiting prototypical network for cross domain few-shot learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 20061–20070.
- Zou, Y.; Liu, Y.; Hu, Y.; Li, Y.; and Li, R. 2024. Flatten long-range loss landscapes for cross-domain few-shot learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 23575–23584.