

MPA: Multimodal Prototype Augmentation for Few-Shot Learning

Liwen Wu^{1,2}, Wei Wang¹, Lei Zhao³, Zhan Gao³, Qika Lin⁴, Shaowen Yao¹, Zuozhu Liu⁵, Bin Pu^{3*}

¹Yunnan University, Kunming, China

²Engineering Research Center of Cyberspace, Kunming, China

³Hunan University, Changsha, China

⁴National University of Singapore, Singapore

⁵Zhejiang University, Hangzhou, China

{lwwu,yaosw}@ynu.edu.cn, {zhaolei,gzhan,pubin}@hnu.edu.cn, weiwang@stu.ynu.edu.cn, zuozhuliu@intl.zju.edu.cn

Abstract

Recently, few-shot learning (FSL) has become a popular task that aims to recognize new classes from only a few labeled examples and has been widely applied in fields such as natural science, remote sensing, and medical images. However, most existing methods focus only on the visual modality and compute prototypes directly from raw support images, which lack comprehensive and rich multimodal information. To address these limitations, we propose a novel Multimodal Prototype Augmentation FSL framework called **MPA**, including LLM-based Multi-Variant Semantic Enhancement (LMSE), Hierarchical Multi-View Augmentation (HMA), and an Adaptive Uncertain Class Absorber (AUCA). LMSE leverages large language models to generate diverse paraphrased category descriptions, enriching the support set with additional semantic cues. HMA exploits both natural and multi-view augmentations to enhance feature diversity (e.g., changes in viewing distance, camera angles, and lighting conditions). AUCA models uncertainty by introducing uncertain classes via interpolation and Gaussian sampling, effectively absorbing uncertain samples. Extensive experiments on four single-domain and six cross-domain FSL benchmarks demonstrate that MPA achieves superior performance compared to existing state-of-the-art methods across most settings. Notably, MPA surpasses the second-best method by 12.29% and 24.56% in the single-domain and cross-domain setting, respectively, in the 5-way 1-shot setting.

Code — <https://github.com/ww36user/MPA>

Introduction

In recent years, few-shot learning (FSL) has attracted widespread attention for its ability to learn efficiently with limited labeled data (Song, Wang, and Zhong 2024; Liu et al. 2025b; Zou et al. 2024; Zhu et al. 2025; Li et al. 2024), and has been widely applied in many fields such as image recognition (Su, He, and Li 2025; Cai et al. 2024), crop disease diagnosis (Zhang et al. 2024b), and remote sensing (Ma et al. 2025). Prototype-based metric methods have been a popular direction in FSL due to their simplicity, efficiency, and generalization ability (Zhang et al. 2021a). These methods

*Corresponding author.

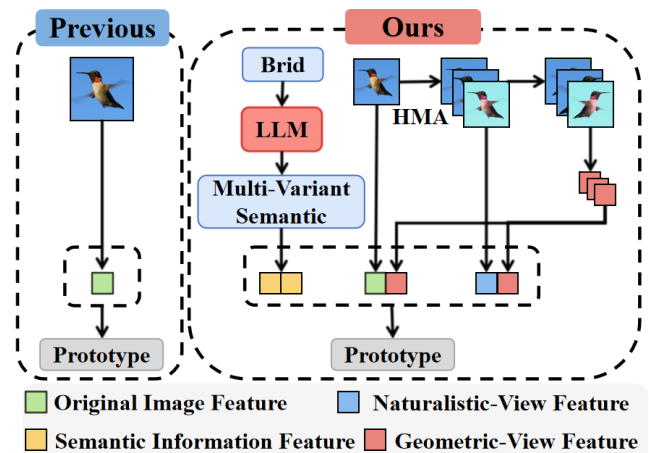


Figure 1: Unlike existing methods that typically focus on original image feature-based prototype, MPA integrates multimodal feature-based prototype (i.e., LLM-based multi-variant semantic features and hierarchical Multi-View features), significantly improving the model’s generalization performance and robustness across diverse tasks.

typically focus on single-modal, single-view image features to construct prototypes and classify based on the distance between query samples and prototypes.

For example, FSL-PRS utilizes a self-attention mechanism to jointly leverage support sets and query sets, dynamically adjusting visual prototypes to enhance classification discrimination capabilities (Zhao et al. 2024). However, due to its focus solely on the visual modality, its performance remains limited.

Similarly, TEFSL proposes a few-shot learning framework that enhances spectral-spatial embedding, improving classification performance through attention feature embedding and graph iterative prototype optimization. However, it still focuses solely on the visual modality, neglecting the importance of semantic information (Xi et al. 2025; Wu et al. 2025). A key limitation remains: *existing methods often overlook the potential of multimodal prototypes and multi-view representations in few-shot scenarios.*

As shown in Figure 1, single-modal/single-view-based

methods may fail to adequately capture the diverse characteristics of limited-support set image samples.

For example, consider a bird image where different species share very similar visual appearances—similar colors, patterns, and shapes—making it difficult for single-modal methods to distinguish between them. However, incorporating semantic information such as species descriptions (“small, with a sharp, curved beak and a bold stripe across its wings”) can provide crucial clues that complement the visual data, enabling more accurate recognition. **(1) This highlights the limitation of relying solely on visual modality and the advantage of leveraging semantic information, since a bird’s representation is inherently multi-dimensional and cannot be fully captured through visual information alone.** Especially by leveraging large language models to fully harness the semantic representation capabilities of language-based multimodal information, the semantic richness of samples can be significantly enhanced.

Similarly, while few-shot images offer limited expressive information on their own, leveraging multiple views can fully capture their representations. **(2) In other words, natural variations in viewing distance, camera angles, and lighting conditions provide diverse perspectives that enhance feature diversity.** In short, incorporating multimodal and multi-view augmentations can enrich sample representations and significantly enhance both discrimination and generalization.

Based on the above analysis, we propose a novel multimodal prototype augmentation FSL framework called MPA, which integrates LLM-based multi-variant semantic enhancement (LMSE), hierarchical multi-view augmentation (HMA), and adaptive uncertain class absorption (AUCA), effectively overcoming these limitations and providing a more comprehensive and robust solution to improving prototype representativeness in few-shot learning. Specifically, LMSE is designed to capture high-quality semantic features during prototype construction, thereby breaking the dependence on image features. Then, HMA is proposed to leverage natural variations and geometric views to generate diverse representations, effectively enriching the support set and enhancing feature diversity. Finally, AUCA combines interpolation with normal distribution dynamic sampling to generate an adaptive uncertain class for absorbing uncertain samples and mitigating inter-class interference. Throughout the process, we utilized the CLIP model to extract multimodal features from multi-variant semantic information, and multi-view visual information. Finally, we trained a logistic regression classifier using these features, whose effectiveness was thoroughly validated across multiple evaluation tasks. The major contributions of this study are outlined as:

- We introduce an LLM-based Multi-Variant Semantic Enhancement module that generates high-quality semantic features from class names and descriptions, reducing reliance on image-only features in prototype construction.
- We design a Hierarchical Multi-View Augmentation module that leverages natural data augmentations and geometric views to generate diverse feature representations, thereby significantly enriching feature diversity and en-

hancing prototype robustness.

- We propose an Adaptive Uncertain Class Absorber that generates uncertain classes via interpolation and normal distribution sampling, effectively absorbing uncertain samples and reducing inter-class interference.
- Extensive experiments on both single-domain and cross-domain FSL datasets, validate the effectiveness of MPA in addressing the challenges of limited data and complex labeling in healthcare, achieving outstanding performance across diverse settings.

Related Work

Few-shot Learning

Few-shot learning aims to generalize to novel classes using only a limited number of labeled samples. Existing approaches can be grouped into metric-based (Su, He, and Li 2025; Li et al. 2022; Zhang et al. 2024b), optimization-based (Ma et al. 2025; Song, Wang, and Zhong 2024; Zou et al. 2024), transfer-based (Fu et al. 2023; Liu et al. 2025b), and data augmentation-based methods. Metric-based methods aim to construct an embedding space where samples from the same class are closely clustered, while those from different classes are well separated. For instance, PrototypeFormer (Su, He, and Li 2025) introduces a transformer-based Prototype Extraction Module that leverages learnable prototype tokens to capture highly discriminative class representations. Additionally, it employs a Prototype Contrastive Loss to improve the robustness and generalization of prototypes, achieving state-of-the-art performance on few-shot learning benchmarks. Similarly, SEVPro (Cai et al. 2024) improves prototype quality by mapping semantic cues to boost class separation, making it easier to distinguish similar categories. This method significantly improves performance across various prototype-based FSL tasks.

Cross-Domain Few-shot Learning

Despite its progress, FSL models often suffer performance degradation when applied to classes with significant domain shifts. To tackle domain shift, various approaches have been proposed (Zhou et al. 2023; Wang et al. 2023, 2025; Cai et al. 2024; Zhang et al. 2024a; Liu et al. 2025b,a). For example, Zhou et al. (Zhou et al. 2023) proposed a local-global distillation framework, representing the first effort to combine both local and global distillation strategies to tackle domain shifts and improve cross-domain generalization. Similarly, SPM (Song, Wang, and Zhong 2024) introduces semantic information as prompts to adaptively guide the feature extractor, enhancing class-specific representations and boosting generalization under domain shift. However, previous prototype-based methods often lack semantic richness and multi-view representations, which limits their ability to fully capture feature distributions and reduces their robustness in few-shot scenarios. However, although SPC (Wu et al. 2025) focuses on local semantic cues to refine models, it still ignores the importance of multi-view diversity and rich textual semantics. The method mainly relies on intra-image spatial relations, lacking complementary visual views

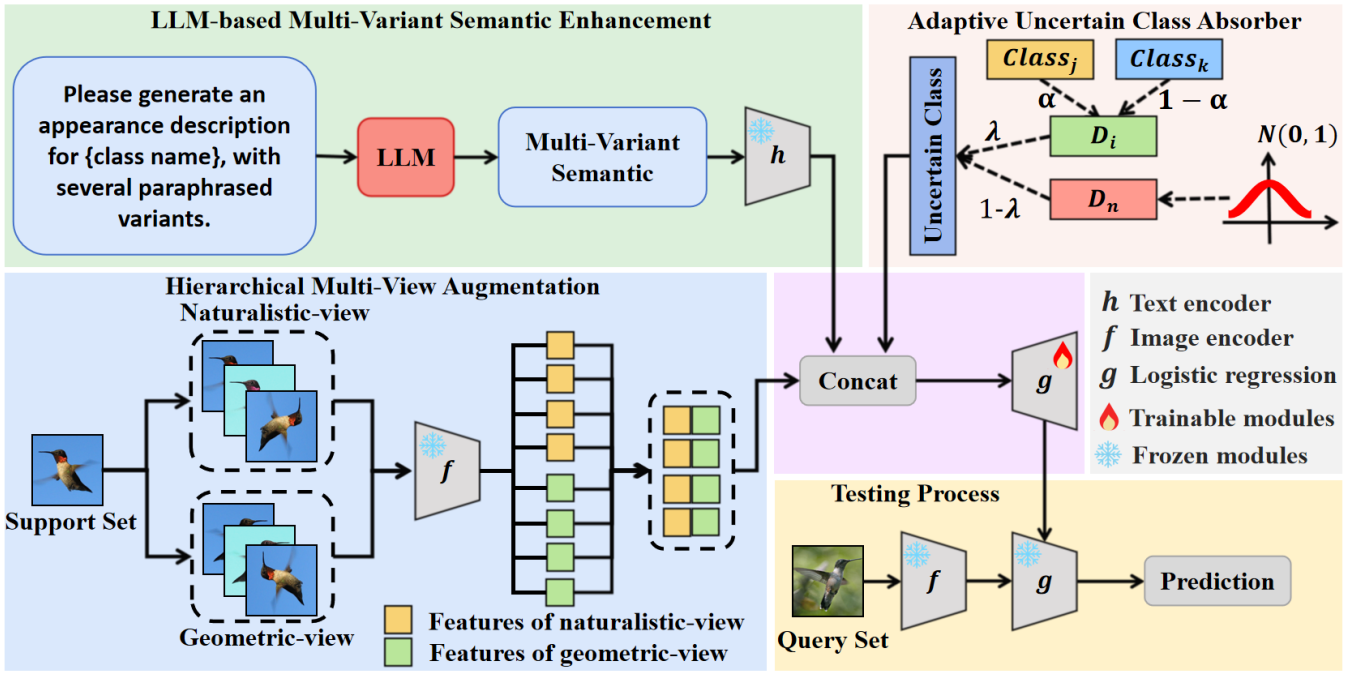


Figure 2: Overview of MPA. Our framework includes three components: LLM-based Multi-Variant Semantic Enhancement (LMSE), Hierarchical Multi-View Augmentation (HMA), and Adaptive Uncertain Class Absorber (AUCA). LMSE produces high-quality semantic features by leveraging LLM, reducing reliance on visual inputs for prototype construction. HMA improves data diversity and feature robustness through multi-view and feature-level augmentations. AUCA mitigates sample bias by interpolating between prototypes and sampling from a normal distribution, where the interpolation weight λ is adaptively determined based on prototype differences. Finally, logistic regression is applied to optimized features for classification.

and textual guidance. As a result, its representations are biased toward limited visual contexts, reducing adaptability and generalization to unseen domains or ambiguous classes.

Methodology

Overview

As shown in Figure 2, we propose a multimodal support set enhancement framework to improve prototype quality and generalization in few-shot learning. Specifically, we first introduce an LLM-based multi-variant semantic enhancement method, which provides high-quality class-level semantic features to complement visual features and alleviate reliance on limited image data. Then, a hierarchical multi-view augmentation strategy is employed, combining naturally transformed naturalistic-views and geometric-views to enrich the diversity of the support set and enhance the expressiveness of feature representations. To further enhance robustness, we introduce an adaptive uncertain class absorber that adaptively constructs uncertain classes via interpolation and normal distribution sampling, helping to absorb uncertain samples. Together, these components form a unified approach that strengthens prototype representations, mitigates noise interference, and significantly improves model generalization in FSL scenarios.

LLM-based Multi-Variant Semantic Enhancement

To enrich class representations, we generate diverse semantic variants using an language model, encode them, and integrate them into the support set. Unlike images with noisy backgrounds, semantic descriptions focus on key attributes, thus improving feature diversity. The support image features are extracted using the CLIP image encoder (Radford et al. 2021), and can be expressed as:

$$F_i = f(\{I_n\}_{n=1}^N) = \{f(I_1), f(I_2), \dots, f(I_N)\}, \quad (1)$$

where $\{I_n\}_{n=1}^N$ represents the set of all N images in the support set, and $f(\cdot)$ is the CLIP image encoder function applied to each image individually.

To obtain richer and more diverse semantic representations, we use a large language model (LLM, e.g., GPT-4.0) to generate multi-variant semantics for each class with the prompt: “Please generate an appearance description for {class name}, with four paraphrased variants.” Specifically, we define a generation function \mathcal{G}_{LLM} that takes a class label c as input and produces a set of semantically enriched information $\{t_m\}_{m=1}^M$ (including the original description and four paraphrased variants) to capture diverse and complementary semantic features. These semantic representations are then projected into the embedding space using the CLIP text encoder $h(\cdot)$, expressed as:

$$F_t = h(\mathcal{G}_{\text{LLM}}(c)), \quad (2)$$

where $\mathcal{G}_{\text{LLM}}(c) = \{t_m\}_{m=1}^M$ represents the set of M multi-variant semantic informations generated for class c , and $h(\cdot)$ represents the CLIP text encoder that projects the textual variants into a shared semantic space. As a result, $F_t \in \mathbb{R}^d$ is the resulting semantic embedding that captures the enriched textual information. By incorporating LLM-generated variants, we introduce contextual diversity and latent knowledge into the textual representation, which significantly enhances generalization under few-shot conditions and alleviates the limitations of fixed or manually written class names.

Hierarchical Multi-View Augmentation

The hierarchical multi-view augmentation module aims to enrich the support set and improve generalization in few-shot scenarios by introducing diverse visual perspectives. We first apply transformations such as central cropping, rotation, and color perturbation to simulate natural variations. Visually diverse yet label-consistent, these augmented samples help reduce dependence on scarce data.

Specifically, given a support image I , the naturalistic-view is generated as:

$$I_a = \{\tau_n(I) \mid \tau_n \in \mathcal{T}\}, \quad (3)$$

where \mathcal{T} represents the set of individual augmentation functions applied independently to I , and \mathcal{I}_a is the resulting set of augmented images generated by applying each transformation τ_n separately. In HMA, to enhance the diversity of visual inputs, we apply hierarchical augmentations including central crops (120, 170, 200 pixels), rotations (45°, 90°, 180°, 270°, 315°), and color jitter (brightness = 0.5, contrast = 0.5, saturation = 0.5, hue = 0.2). The corresponding features are then extracted by an image encoder:

$$F_a = f(\{I_a^{(m)}\}_{m=1}^M) = \{f(I_a^{(1)}), f(I_a^{(2)}), \dots, f(I_a^{(M)})\}, \quad (4)$$

where $\{I_a^{(m)}\}_{m=1}^M$ represents the set of M naturalistic views generated from a raw support image I , and $f(\cdot)$ is the CLIP image encoder function applied to each naturalistic-view individually, resulting in a set of feature vectors $F_a \in \mathbb{R}^{M \times d}$.

In parallel, complementary views are generated via horizontal reflection of support images, and their features are combined with the originals to enrich representation and boost few-shot generalization.

Adaptive Uncertain Class Absorber

To enhance model generalizability, we introduce an adaptive uncertain class absorber by leveraging interpolation and sampling from a normal distribution. The interpolation-based data can be formulated as:

$$D_i = [\alpha \quad 1 - \alpha] \cdot \begin{bmatrix} F_j \\ F_k \end{bmatrix}, \quad (5)$$

where the interpolation coefficient α ranges from 0.2 to 0.8 in our experiments, and F_j, F_k denote features from different classes. Next, the normal distribution is expressed as:

$$D_n \sim \mathcal{N}(0, 1), \quad (6)$$

where $\text{Data}_n \in \mathbb{R}^{1 \times 768}$. Then, the cosine similarity $S_{j,k}$ between the prototypes of class j and class k is computed as:

$$S_{j,k} = \frac{C_j \cdot C_k}{\|C_j\| \|C_k\|}, \quad j \leq k, \quad (7)$$

where $\|\cdot\|$ represents the Euclidean norm of a vector. Based on the cosine similarity, construct a symmetric matrix \mathbf{S} , where the lower triangular part is set to 0:

$$\mathbf{S} = \begin{bmatrix} S_{1,1} & S_{1,2} & \cdots & S_{1,N} \\ 0 & S_{2,2} & \cdots & S_{2,N} \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & S_{N,N} \end{bmatrix}, \quad (8)$$

where N is the total number of classes. Normalize the matrix \mathbf{S} globally to map the cosine similarity to the range $[0, 1]$:

$$S'_{j,k} = \arg \min_{\lambda \in \mathbb{R}} \|S_{j,k} - ((1 - \beta)S_{\min} + \beta S_{\max})\|^2, \quad (9)$$

where S_{\min} and S_{\max} are the minimum and maximum values of the matrix \mathbf{S} , respectively.

Finally, compute average of the normalized similarities for all class pairs and convert it into a difference measure:

$$\lambda = 1 - \frac{2}{\binom{C}{2}} \sum_{j=1}^{C-1} \sum_{k=j+1}^C S'_{j,k}, \quad (10)$$

where $\binom{C}{2}$ is the combination number. As a result, the uncertain class data can be expressed as:

$$\mathbb{E}[D_u] = (1 - \lambda) \cdot D_n + \lambda \cdot D_i, \\ D_u \sim \begin{cases} D_n, & P(D_u = D_n) = 1 - \lambda \\ D_i, & P(D_u = D_i) = \lambda \end{cases}, \quad (11)$$

where λ is a dynamic factor determined by the inter-prototype differences within each batch. D_u , D_n , and D_i denote uncertain-class, Gaussian-sampled, and interpolated data, respectively.

In cross-domain scenarios, where features are more clustered and harder to distinguish, λ is typically smaller; in contrast, it tends to be larger in single-domain datasets with more separable features. MPA injects diversity via interpolation and Gaussian sampling to foster robust decision boundaries and flexible overlap with real classes.

Experiments

Datasets and Evaluation

We conducted extensive experiments on ten datasets, including four single-domain and six cross-domain datasets. The publicly available datasets used in our experiments are divided into two groups. The single-domain datasets consist of miniImageNet (Vinyals et al. 2016), tieredImageNet (Hariharan and Girshick 2017), CIFAR-FS (Bertinetto et al. 2018), and FC100 (Oreshkin, Rodríguez López, and Lacoste 2018). The cross-domain datasets include fine-grained classification benchmarks CUB (Welinder et al. 2010) and Cars (Krause et al. 2013), which present subtle inter-class differences, as well as more general domain datasets such as Places (Wang et al. 2020), Plantae (Van Horn et al. 2018), EuroSAT (Helber et al. 2019), and CropDisease (Mohanty, Hughes, and Salathé 2016).

Method	Marker	PT	miniImageNet	tieredImageNet	CIFAR-FS	FC100	Average
Meta-UAFS	IJCAI-21	×	64.22	69.13	74.08	41.99	62.36
HCTransformers	CVPR-22	×	74.62	79.57	78.88	48.15	70.31
CPEA	ICCV-23	×	71.97	76.93	77.82	47.24	68.49
Diff-ResNet	TPAMI-24	×	73.47	79.74	-	-	-
FedFSL-CFRD	AAAI-25	×	56.66	60.52	63.00	-	-
SP-CLIP	CVPR-23	✓	72.31	78.03	82.18	48.53	70.26
CAML	ICLR-24	✓	96.20	95.40	70.80	-	-
SEVPro	IJCAI-24	✓	71.81	72.77	80.36	-	-
SemFew-Trans	CVPR-24	✓	78.94	82.37	84.34	54.27	74.98
SPM	AAAI-24	✓	93.70	88.79	82.40	<u>68.35</u>	<u>83.31</u>
ECER-FSL	AAAI-25	✓	81.14	81.81	86.01	57.34	76.58
MLVLM	AAAI-25	✓	<u>98.24</u>	<u>98.06</u>	<u>95.02</u>	-	-
MPA (Ours)	-	✓	98.87	98.57	97.47	87.47	95.60

Table 1: Performance comparison of state-of-the-art methods on 5-way 1-shot single-domain datasets. **PT** represents a public pre-trained model. The best results are in **bold**, and the second-best are underlined.

Method	Mark	PT	CUB	Cars	Places	Plantae	EuroSAT	CropDisease	Average
TPN+ATA	IJCAI-21	×	51.89	38.07	57.26	40.75	70.84	82.47	56.88
RDC	CVPR-22	×	50.09	39.94	61.17	41.30	70.51	85.79	58.13
LDP-net	CVPR-23	×	55.94	37.44	62.21	41.04	73.25	81.24	58.52
TPN+FAP	IJCAI-24	×	50.56	34.39	57.34	37.44	62.62	76.11	53.08
PRM-Net	NIPS-24	×	59.48	38.86	62.90	44.06	69.56	84.01	59.81
StyleAdv-FT	CVPR-23	✓	84.01	40.48	72.64	55.52	74.93	84.11	68.62
FloR	CVPR-24	✓	55.94	40.01	61.27	41.70	71.38	86.30	59.43
SPM	AAAI-24	✓	84.39	41.71	72.35	53.85	74.97	84.43	68.62
SVasP	AAAI-25	✓	85.56	40.51	<u>75.93</u>	<u>56.25</u>	75.51	83.98	<u>69.62</u>
DAMIM	AAAI-25	✓	-	-	-	-	<u>77.23</u>	86.74	-
ECER-FSL	AAAI-25	✓	-	-	-	-	74.13	82.13	-
MLVLM	AAAI-25	✓	<u>96.64</u>	99.74	-	-	-	-	-
MPA (Ours)	-	✓	98.95	<u>98.51</u>	93.55	91.73	87.05	95.28	94.18

Table 2: Comparison with state-of-the-art methods on 5-way 1-shot classification across cross-domain datasets.

Baseline Methods

We compare our model with a broad range of state-of-the-art baselines to validate the effectiveness of MPA. For cross-domain FSL tasks, we evaluate against: TPN+ATA (Wang and Deng 2021), RDC (Li et al. 2022), LDP-Net (Zhou et al. 2023), TPN+FAP (Zhang et al. 2024b), PRM-Net (Zhou et al. 2024), StyleAdv-FT (Fu et al. 2023), FloR (Zou et al. 2024), SVasP (Li, Fang, and Xue 2025), and DAMIM (Ma et al. 2025). For standard FSL tasks, we additionally include comparisons with: Meta-UAFS (Zhang et al. 2021b), HCTransformers (He et al. 2022), CPEA (Hao et al. 2023), Diff-ResNet (Wang et al. 2023), FedFSL-CFRD (Wang et al. 2025), SP-CLIP (Chen et al. 2023), CAML (Fifty et al. 2024), SEVPro (Cai et al. 2024), SemFew-Trans (Zhang et al. 2024a) and DAMIM (Ma et al. 2025). Notably, SPM (Song, Wang, and Zhong 2024), ECER-FSL (Liu et al. 2025b), and MLVLM (Liu et al. 2025a) are evaluated under both single-domain and cross-domain FSL settings.

Implementation Details

Our model is implemented using the PyTorch framework, employing the CLIP model with the ViT-L/14 backbone as the feature extractor. In each epoch, we randomly sample 100 episodes from the target domain. For each task, the support set consists of 5 classes, with each class containing either 5 or 1 sample, while the query set includes 15 images per class. All experiments were conducted on a workstation equipped with an NVIDIA RTX 4090 GPU, Intel(R) Xeon(R) Silver 4310 CPU, and 32GB RAM.

Comparison with State-of-the-arts

Overall Results. As shown in Tables 1, 2, 3, and 4, MPA exhibits outstanding classification performance and broad applicability across all experimental conditions. Notably, our method performs effectively not only in natural scenarios (e.g., miniImageNet and CIFAR-FS datasets) but also achieves strong results in fine-grained tasks (CUB and Cars datasets). Moreover, MPA demonstrates

Method	Mark	PT	miniImageNet	tieredImageNet	CIFAR-FS	FC100	Average
Meta-UAFS	IJCAI-21	×	79.99	84.33	85.92	57.43	76.92
HCTransformers	CVPR-22	×	89.19	91.72	90.50	66.42	84.46
CPEA	CVPR-23	×	87.26	90.12	88.98	65.02	82.85
Diff-ResNet	TPAMI-24	×	83.86	87.10	-	-	-
FedFSL-CFRD	AAAI-25	×	72.15	78.95	74.58	-	-
SP-CLIP	CVPR-23	✓	83.42	88.55	88.24	61.55	80.44
CAML	ICLR-24	✓	<u>98.60</u>	<u>98.10</u>	85.50	-	-
SEVPro	IJCAI-24	✓	78.88	84.04	86.12	-	-
SemFew-Trans	CVPR-24	✓	86.49	89.89	89.11	65.02	82.63
SPM	AAAI-24	✓	98.30	96.20	<u>93.10</u>	<u>83.78</u>	<u>92.85</u>
MPA (Ours)	-	✓	99.07	98.79	97.64	90.57	96.52

Table 3: Performance comparison of state-of-the-art methods on 5-way 5-shot single-domain datasets.

Method	Mark	PT	CUB	Cars	Places	Plantae	EuroSAT	CropDisease	Average
TPN+ATA	IJCAI-21	×	70.14	55.23	73.87	59.02	85.47	93.56	72.88
RDC	CVPR-22	×	67.23	53.49	74.91	57.47	84.29	93.30	71.78
LDP-Net	CVPR-23	×	73.34	53.06	75.47	59.64	82.01	89.40	72.15
TPN+FAP	IJCAI-24	×	64.17	47.38	72.05	53.58	80.24	88.34	67.63
PRM-Net	NeurIPS-24	×	76.68	55.44	76.98	63.08	83.22	93.09	74.75
StyleAdv-FT	CVPR-23	✓	95.82	66.02	88.33	78.01	90.12	95.99	85.72
FloR	CVPR-24	✓	74.06	57.98	74.25	61.70	83.76	93.60	74.23
SPM	AAAI-24	✓	95.95	62.89	88.79	74.52	89.72	96.11	84.66
SVasP	AAAI-25	✓	<u>95.95</u>	<u>66.47</u>	<u>89.19</u>	<u>78.67</u>	90.55	96.17	<u>86.17</u>
DAMIM	AAAI-25	✓	-	-	-	-	<u>91.08</u>	<u>96.49</u>	-
MPA (Ours)	-	✓	99.32	99.63	93.37	95.37	91.77	98.21	96.28

Table 4: Comparison with state-of-the-art methods on 5-way 5-shot classification across cross-domain datasets.

robust capabilities in single-domain few-shot learning (miniImageNet, tieredImageNet, CIFAR-FS and FC100 datasets), cross-domain few-shot learning (EuroSAT and CropDisease), and fine-grained recognition tasks (Cars, CUB, Places and Plantae datasets). These results highlight

Method	CUB		Cars	
	1-shot	5-shot	1-shot	5-shot
TPN+ATA	51.89	70.14	38.07	55.23
RDC	50.09	67.23	39.94	53.49
LDP-net	55.94	73.34	37.44	53.06
TPN+FAP	50.56	64.17	34.39	47.38
PRM-Net	59.48	76.68	38.86	55.44
StyleAdv-FT	84.01	95.82	40.48	66.02
FloR	55.94	74.06	40.01	57.98
SPM	84.39	95.95	41.71	62.89
SVasP	85.56	<u>95.95</u>	40.51	<u>66.47</u>
MLVLM	96.64	-	99.74	-
MPA (Ours)	98.95	99.32	<u>98.51</u>	99.63

Table 5: Comparison results on fine-grained datasets (CUB and Cars) under 5-way 1-shot and 5-shot settings.

MPA’s potential for practical use across a wide range of real-world scenarios.

Results on single-domain FSL datasets. As shown in Table 1, under the 5-way 1-shot setting, the average classification accuracy of MPA surpasses the second-best method by 12.29%. Specifically, MPA improves the average accuracy by 5.17%, 9.78%, 15.07%, and 19.12% compared to the second-best method on miniImageNet, tieredImageNet, CIFAR-FS and FC100 datasets, respectively. As shown in Table 3, under the 5-way 5-shot setting, the average classification accuracy of MPA surpasses the second-best method (SPM) by 3.67%. Specifically, MPA improves the average classification accuracy by 0.77%, 2.59%, 4.54%, and 6.79% compared to the second-best method on miniImageNet, tieredImageNet, CIFAR-FS and FC100 datasets, respectively.

Results on cross-domain FSL datasets. As shown in Table 2, MPA achieves superior performance on all six datasets under the 5-way 1-shot setting. Specifically, MPA achieves the best performance on CUB, Places, Plantae, EuroSAT, and CropDisease datasets and ranks second-best on Cars,

[†]The appendices and supplementary materials are available at <https://github.com/ww36user/MPA>.

with an average improvement of 24.56%. As shown in Table 4, under the 5-way 5-shot setting, MPA achieves the best performance across all six datasets. Specifically, MPA improves the average accuracy by 3.37%, 33.16%, 4.18%, 16.7%, 1.22% and 2.04% compared to the second-best method on CUB, Cars, Places, Plantae, EuroSAT and CropDisease datasets, respectively.

Results on fine-grained FSL datasets. As shown in Table 5, MPA performs exceptionally well on the CUB and Cars datasets in the 5-way 1-shot and 5-way 5-shot settings, achieving the best results in three metrics and ranking second in the other. Specifically, in the 5-way 5-shot classification task on the Cars dataset, MPA achieves a 33.16% improvement in accuracy compared to the second-place model. These results demonstrate that MPA possesses significant advantages and strong generalization capabilities in low-sample fine-grained classification tasks.

Ablation Study

Effect of MPA Components. As presented in Table 6, to comprehensively evaluate the contributions of each component in our MPA framework, we conduct an ablation study across three datasets. Without introducing HMA and AUCA, using LMSE alone can significantly improve performance on EuroSAT and CIFAR-FS. After introducing HMA, performance is further improved on Places and CIFAR-FS, indicating that multi-view augmentation helps model generalization. Finally, when the three modules are used together, the best performance is achieved on all datasets, showing the effectiveness and complementarity of each module.

LMSE	HMA	AUCA	EuroSAT	Places	CIFAR-FS
×	×	×	76.41	87.24	93.69
✓	×	×	83.03	93.43	95.36
×	✓	×	79.44	84.71	94.17
✓	✓	×	85.69	92.64	96.32
✓	✓	✓	87.05	93.55	97.47

Table 6: Ablation study on the effects of the three components across three datasets in 5-way 1-shot setting.

Effect of Backbone. To evaluate the generalizability of MPA under the challenging 5-way 1-shot setting, we conduct experiments on multiple CLIP backbone variants, including ViT-L/14, ViT-B/32, ViT-B/16, and ResNet101. As shown in Table 7, MPA consistently outperforms the baseline across all backbones and datasets, demonstrating strong robustness and adaptability in few-shot scenarios.

Extended Evaluation and Analysis. To further validate the generality and robustness of our approach, we provide several extended experiments and analyses. Specifically, Table 8 examines the impact of different large language models on semantic enhancement, and GPT-4.0 achieves the best accuracy. *Appendix. A2[†]* evaluates the effectiveness of AUCA under extreme conditions with uncertain samples. A statistical analysis of the dynamic factor λ is presented in *Ap-*

Method	Backbone	EuroSAT	Places	FC100
Baseline	ViT-L/14	76.41	87.24	77.18
Ours	ViT-L/14	87.05	93.55	87.47
Baseline	ViT-B/32	65.47	82.32	60.20
Ours	ViT-B/32	74.15	91.08	75.95
Baseline	ViT-B/16	71.19	83.92	67.95
Ours	ViT-B/16	78.92	92.00	79.73
Baseline	ResNet101	57.88	78.71	46.40
Ours	ResNet101	65.99	83.99	51.61

Table 7: Performance comparison across different CLIP backbones on three datasets under the 5-way 1-shot setting.

pendix. A3[†] to demonstrate the adaptability of AUCA across various domains. UMAP visualizations of MPA features on five target-domain datasets are provided in *Appendix. A4[†]*, showing that MPA clearly separates category boundaries compared with the baseline. Additional feature visualizations and analysis are presented in *Appendix. A5[†]*.

These findings consistently support the effectiveness of our method in learning transferable representations and addressing diverse data distributions.

LLM	tieredImageNet	CIFAR-FS
GPT-3.5	97.88	97.20
GPT-4.0	98.57	97.47
DeepSeek-V3	97.79	97.01
DeepSeek-R1	97.79	97.25
ChatGLM-3	97.99	97.19
Claude-4	98.17	97.39
Gemini-2.5	98.17	96.53
InternVL3	98.16	97.15
Qwen2.5	98.15	97.36
ERNIE X1	98.03	97.32

Table 8: Performance comparison of LMSE using different LLMs, evaluated on 5-way 1-shot tasks across two datasets.

Conclusion

In this paper, we propose MPA, a novel multimodal few-shot learning framework called MPA, which includes LLM-based Multi-Variant Semantic Enhancement (LMSE), Hierarchical Multi-View Augmentation (HMA), and Adaptive Uncertain Class Absorber (AUCA) for improving generalization and stability. LMSE is designed to generate high-quality semantic features, reducing reliance on image data. HMA enriches support representations by combining natural augmentations and geometric views. AUCA dynamically generates adaptive uncertain classes through interpolation and normal distribution sampling to address uncertain samples and reduce class overlap. Experiments on four single-domain datasets and six cross-domain datasets show state-of-the-art performance, demonstrating strong generalization and effectiveness in various FSL scenarios.

Acknowledgments

This study is supported by the Yunnan Fundamental Research Projects (Nos. 202301AU070194 and 202501AT070233), the Yunnan Province Special Project (No. 202403AP140021), the Practical Innovation Project of Postgraduate Students in the Professional Degree of Yunnan University (No. ZC-24248950), the Yunnan Province expert workstations (No. 202305AF150078), and the National Natural Science Foundation of China (Nos. 62162067 and 62562061).

References

- Bertinetto, L.; Henriques, J. F.; Torr, P. H.; and Vedaldi, A. 2018. Meta-learning with differentiable closed-form solvers. *arXiv preprint arXiv:1805.08136*.
- Cai, H.; Liu, Y.; Huang, S.; and Lv, J. 2024. With a little help from language: semantic enhanced visual prototype framework for few-shot learning. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*, 3751–3759.
- Chen, W.; Si, C.; Zhang, Z.; Wang, L.; Wang, Z.; and Tan, T. 2023. Semantic Prompt for Few-Shot Image Recognition. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 23581–23591. IEEE.
- Fifty, C.; Duan, D.; Junkins, R. G.; Amid, E.; Leskovec, J.; Re, C.; and Thrun, S. 2024. Context-Aware Meta-Learning. *arXiv:2310.10971*.
- 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.
- Hao, F.; He, F.; Liu, L.; Wu, F.; Tao, D.; and Cheng, J. 2023. Class-aware patch embedding adaptation for few-shot image classification. In *Proceedings of the IEEE/CVF international conference on computer vision*, 18905–18915.
- Hariharan, B.; and Girshick, R. 2017. Low-shot visual recognition by shrinking and hallucinating features. In *Proceedings of the IEEE international conference on computer vision*, 3018–3027.
- He, Y.; Liang, W.; Zhao, D.; Zhou, H.-Y.; Ge, W.; Yu, Y.; and Zhang, W. 2022. Attribute surrogates learning and spectral tokens pooling in transformers for few-shot learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9119–9129.
- 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.
- 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.
- Li, P.; Gong, S.; Wang, C.; and Fu, Y. 2022. Ranking distance calibration for cross-domain few-shot learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9099–9108.
- Li, W.; Fang, P.; and Xue, H. 2025. Svasp: Self-versatility adversarial style perturbation for cross-domain few-shot learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 15275–15283.
- Li, X.; Tan, Y.; Liang, B.; Pu, B.; Yang, J.; Zhao, L.; Kong, Y.; Yang, L.; Zhang, R.; Li, H.; et al. 2024. Tkr-fsod: Fetal anatomical structure few-shot detection utilizing topological knowledge reasoning. *IEEE Journal of Biomedical and Health Informatics*, 29(1): 547–557.
- Liu, F.; Cai, W.; Huo, J.; Zhang, C.; Chen, D.; and Zhou, J. 2025a. Making Large Vision Language Models to be Good Few-shot Learners. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 5415–5423.
- Liu, M.; Wu, F.; Li, B.; Lu, Z.; Yu, Y.; and Li, X. 2025b. Envisioning class entity reasoning by large language models for few-shot learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 18906–18914.
- Ma, R.; Zou, Y.; Li, Y.; and Li, R. 2025. Reconstruction Target Matters in Masked Image Modeling for Cross-Domain Few-Shot Learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 19305–19313.
- Mohanty, S. P.; Hughes, D. P.; and Salathé, M. 2016. Using deep learning for image-based plant disease detection. *Frontiers in plant science*, 7: 1419.
- 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.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PmLR.
- Song, M.; Wang, H.; and Zhong, G. 2024. Self-prompt mechanism for few-shot image recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 4934–4942.
- Su, M.; He, F.; and Li, F. 2025. PrototypeFormer: Learning to Explore Prototype Relationships for Few-shot Image Classification. *arXiv:2310.03517*.
- 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.; Wierstra, D.; et al. 2016. Matching networks for one shot learning. *Advances in Neural Information Processing Systems*, 29.
- Wang, H.; and Deng, Z.-H. 2021. Cross-Domain Few-Shot Classification via Adversarial Task Augmentation. In Zhou, Z.-H., ed., *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, 1075–1081.
- Wang, H.; Wu, X.; Huang, Z.; and Xing, E. P. 2020. High-frequency component helps explain the generalization of convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8684–8694.

Wang, S.; Li, J.; Liu, Z.; Zhang, Y.; and Gong, M. 2025. FedFSL-CFRD: Personalized Federated Few-Shot Learning with Collaborative Feature Representation Disentanglement. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 21243–21251.

Wang, T.; Dou, Z.; Bao, C.; and Shi, Z. 2023. Diffusion mechanism in residual neural network: theory and applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(2): 667–680.

Welinder, P.; Branson, S.; Mita, T.; Wah, C.; Schroff, F.; Belongie, S.; and Perona, P. 2010. Caltech-UCSD birds 200.

Wu, L.; Wang, W.; Fang, J.; Zhao, L.; Gao, Z.; Li, N.; Liang, P.; Pu, B.; and Yao, S. 2025. SPC: Selective Patches Consistency for Cross-Domain Few-Shot Learning. *IEEE Transactions on Consumer Electronics*.

Xi, B.; Zhang, Y.; Li, J.; Huang, Y.; Li, Y.; Li, Z.; and Chanussot, J. 2025. Transductive Few-Shot Learning With Enhanced Spectral-Spatial Embedding for Hyperspectral Image Classification. *IEEE Transactions on Image Processing*, 34: 854–868.

Zhang, B.; Li, X.; Ye, Y.; Huang, Z.; and Zhang, L. 2021a. Prototype completion with primitive knowledge for few-shot learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 3754–3762.

Zhang, H.; Xu, J.; Jiang, S.; and He, Z. 2024a. Simple semantic-aided few-shot learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 28588–28597.

Zhang, T.; Cai, Q.; Gao, F.; Qi, L.; and Dong, J. 2024b. Exploring Cross-Domain Few-Shot Classification via Frequency-Aware Prompting. *arXiv preprint arXiv:2406.16422*.

Zhang, Z.; Lan, C.; Zeng, W.; Chen, Z.; and Chang, S.-F. 2021b. Uncertainty-Aware Few-Shot Image Classification. In *International Joint Conference on Artificial Intelligence*.

Zhao, P.; Wang, L.; Zhao, X.; Liu, H.; and Ji, X. 2024. Few-shot learning based on prototype rectification with a self-attention mechanism. *Expert Systems with Applications*, 249: 123586.

Zhou, F.; Wang, P.; Zhang, L.; Chen, Z.; Wei, W.; Ding, C.; Lin, G.; and Zhang, Y. 2024. Meta-Exploiting Frequency Prior for Cross-Domain Few-Shot Learning. *Advances in Neural Information Processing Systems*, 37: 116783–116814.

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.

Zhu, Y.; Liang, B.; Li, N.; Zhao, L.; Li, X.; Li, H.; Yang, F.; and Pu, B. 2025. Anatomical Structure Few-Shot Detection Utilizing Enhanced Human Anatomy Knowledge in Ultrasound Images. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 35–45. Springer.

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.