

Delving into Multimodal Prompting for Fine-Grained Visual Classification

Xin Jiang^{1*}, Hao Tang^{1*}, Junyao Gao², Xiaoyu Du¹, Shengfeng He³, Zechao Li^{1†}

¹Nanjing University of Science and Technology, China

²Tongji University, China

³Singapore Management University, Singapore

{xinjiang, tanghao0918, duxy, zechao.li}@njjust.edu.cn, junyao.gao@tongji.edu.cn, shengfenghe7@gmail.com

Abstract

Fine-grained visual classification (FGVC) involves categorizing fine subdivisions within a broader category, which poses challenges due to subtle inter-class discrepancies and large intra-class variations. However, prevailing approaches primarily focus on uni-modal visual concepts. Recent advancements in pre-trained vision-language models have demonstrated remarkable performance in various high-level vision tasks, yet the applicability of such models to FGVC tasks remains uncertain. In this paper, we aim to fully exploit the capabilities of cross-modal description to tackle FGVC tasks and propose a novel multimodal prompting solution, denoted as MP-FGVC, based on the contrastive language-image pertaining (CLIP) model. Our MP-FGVC comprises a multimodal prompts scheme and a multimodal adaptation scheme. The former includes Subcategory-specific Vision Prompt (SSVP) and Discrepancy-aware Text Prompt (DATP), which explicitly highlights the subcategory-specific discrepancies from the perspectives of both vision and language. The latter aligns the vision and text prompting elements in a common semantic space, facilitating cross-modal collaborative reasoning through a Vision-Language Fusion Module (VLFM) for further improvement on FGVC. Moreover, we tailor a two-stage optimization strategy for MP-FGVC to fully leverage the pre-trained CLIP model and expedite efficient adaptation for FGVC. Extensive experiments conducted on four FGVC datasets demonstrate the effectiveness of our MP-FGVC.

Introduction

Fine-grained visual classification (FGVC) is a significant task in the field of computer vision that aims to categorize distinct subcategories within a given supercategory. Unlike traditional classification tasks, FGVC faces the challenges of subtle inter-class variations caused by similar subordinate categories, as well as large intra-class variations due to factors such as location, scale, and deformation. Recognizing these subtle differences between subcategories, such as distinguishing bird subcategories based on characteristics like eye shape, flipper appearance, and coloration of the tail region, is a key aspect of human visual perception. To address

these challenges, deep learning methods (Yu et al. 2018; Xu et al. 2021; Gu and Feng 2023; Tang et al. 2023; Gu et al. 2023; Shen et al. 2023) have been proposed for FGVC, falling into two main categories: feature-encoding methods and part-localization methods. Part-localization methods, in particular, have garnered significant interest due to their consistency with human recognition intuition, as fine-grained subcategory-specific discrepancies (Dong et al. 2023) often reside in the unique properties of object parts.

Although CNN-based models have achieved satisfactory performance for FGVC in a weakly-supervised manner, a prominent drawback persists: CNNs tend to focus solely on a small, irrelevant region within an image, resulting in the captured subcategory-specific discrepancies that lack robustness and discriminability. Recently, vision transformers like ViT (Dosovitskiy et al. 2021a), have recorded remarkable performance in the image classification tasks and are actively being employed in FGVC (He et al. 2022; Xu et al. 2023). Compared to CNNs, ViT’s patch-by-patch processing lends itself well to FGVC as each image patch can be regarded as a local part. Additionally, the inherent self-attention mechanism in ViT enables the modeling of long-range dependencies in the entire image. However, due to the relatively limited size of FGVC datasets, these CNN-based and ViT-based methods often exhibit erratic behavior during optimization and heavily rely on visual representations from a predefined closed set during inference, which hampers the full exploitation of these models’ potential for FGVC.

The human visual perception system possesses a unique mechanism of top-down **cognitive penetrability** (Maier and Abdel Rahman 2019) in which prior linguistic knowledge (Li et al. 2023), such as appearance descriptions, is instinctively used to adaptively adjust ongoing visual perceptual processing to stimulus features relevant to object categories, thereby facilitating the recognition of fine-grained objects. Recently, large-scale vision-language (VL) models have been developed to connect visual representation with high-level language descriptions, showing potential for adaptation to various vision tasks. Thus, it raises the question of whether the cross-modal description ability in the VL models can be effectively utilized for FGVC tasks. Recent studies (Bahng et al. 2022; Jia et al. 2022) have already confirmed the possibility of utilizing the knowledge from the CLIP model via prompt techniques (Zhou et al. 2022b,a)

*Equal contribution.

†Corresponding author.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

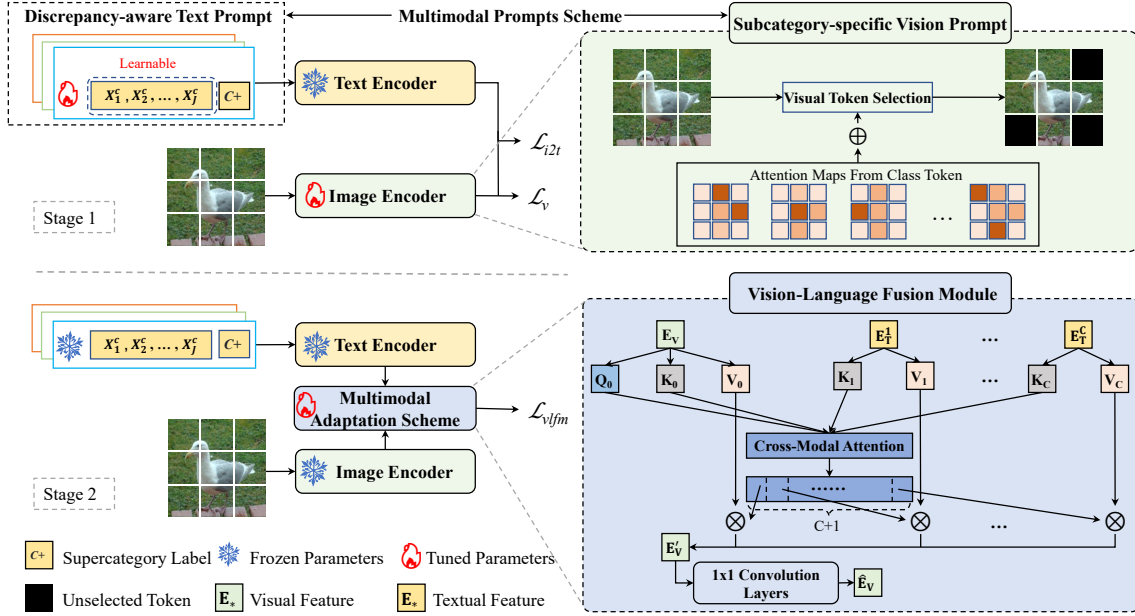


Figure 1: The pipeline of our framework, which consists of a multimodal prompts scheme and a multimodal adaptation scheme.

for classification tasks. However, existing prompt strategies are primarily designed to capture high-level category semantics rather than the more detailed subcategory-level discrepancies. Furthermore, the lack of specific annotations in existing FGVC datasets hampers the widespread adoption of VL models in FGVC. Fortunately, recent works (Li, Sun, and Li 2023; Wang et al. 2023) have introduced large-scale contrastive language-image pretraining (CLIP) (Radford et al. 2021) model to the domain of fine-grained understanding. However, these methods aim to associate a broader range of visual concepts, including both pre-defined and unknown categories (Tang et al. 2020), with their textual descriptions (Yan et al. 2023), addressing the problem of fine-grained image retrieval in open-set scenarios. Therefore, it is worth investigating how to effectively make the pre-trained CLIP model sensitive to subcategory-specific discrepancies according to the semantic prompts to enhance the performance of FGVC tasks in close-set scenarios.

To this end, leveraging the capabilities of the recently introduced CLIP model, we present a multimodal prompting solution, called MP-FGVC, for FGVC. The primary objective of MP-FGVC is to align the FGVC task with the classification task performed by the pre-trained CLIP model. Fig. 1 illustrates the proposed MP-FGVC framework, which consists of a multimodal prompts scheme and a multimodal adaptation scheme. The multimodal prompts scheme, comprising Subcategory-specific Vision Prompt (SSVP) and Discrepancy-aware Text Prompt (DATP), explicitly emphasizes the subcategory-specific discrepancies from the input perspective. In particular, SSVP identifies subcategory-specific discrepancies by selecting highly salient patches from a set of generated image patches, employing an importance ranking mechanism. On the other hand, DATP adaptively transforms learnable text tokens using supercat-

egory names to generate subcategory-specific discrepancy descriptions, which serve as additional discriminative cues. To maintain semantic coherence, a multimodal adaptation scheme is proposed to align the vision and text prompt into a common semantic space, enabling cross-modal collaborative reasoning through a Vision-Language Fusion Module (VLFM). This further improves the performance of multimodal prompting on FGVC tasks. However, due to the absence of specific annotations for fine-grained images and the availability of only subcategory labels or indexes, there is limited supervision for optimizing MP-FGVC. Thereby, a two-stage training strategy is proposed. In the first stage, we keep the text encoder of the pre-trained CLIP fixed and optimize the image encoder with the multimodal prompts scheme. In the second stage, the image encoder, text encoder, and multimodal prompts scheme remain static. Together, they provide subcategory-specific visual embeddings and discrepancy-aware textual embeddings, facilitating the multimodal adaptation scheme to effectively exploit multimodal prompting for FGVC.

In summary, the primary contributions of this work are as follows: (1) We propose MP-FGVC, a novel multimodal prompting solution for fine-grained visual classification. To the best of our knowledge, this is the first approach that efficiently adapts the image-text pre-trained CLIP model to FGVC tasks in close-set scenarios. (2) Our MP-FGVC encompasses a multimodal prompts scheme and a multimodal adaptation scheme. The former explicitly highlights the subcategory-specific discrepancies from both the vision and language perspectives, while the latter achieves cross-modal collaborative reasoning based on the former to further enhance multimodal prompting in FGVC. (3) MP-FGVC employs a two-stage training strategy to fully leverage the cross-modal description ability of the CLIP model, resulting

in improved FGVC with a simple modification. Experimental evaluations conducted on four FGVC datasets demonstrate the remarkable performance achieved by our proposed MP-FGVC.

Related Work

Fine-Grained Visual Classification

Existing CNN-based approaches for FGVC can generally be categorized into two groups. The first group, known as feature-encoding methods (Yu et al. 2018; Gao et al. 2020; Xu et al. 2021), focuses on learning more comprehensive features by integrating higher-order information or exploring various training strategies. The second group, referred to as part-localization methods (Ge, Lin, and Yu 2019; Liu et al. 2020; Tang et al. 2022), aims to identify multiple local regions for capturing more discriminative features. Recently, vision Transformers (ViTs) (Dosovitskiy et al. 2021b) have exhibited impressive performance in image classification on large-scale datasets and are currently being actively employed in FGVC (Hu et al. 2021; Sun, He, and Peng 2022; Xu et al. 2023). By virtue of their patch-based processing and powerful self-attention mechanism, ViT-based methods can outperform CNN-based methods (Zhuang, Wang, and Qiao 2020; Luo et al. 2019; Min et al. 2023) by solely relying on a pure Transformer encoder architecture along with diverse patch selection mechanisms. However, these methods exclusively rely on uni-modal visual representation from a pre-defined closed set and completely disregard corresponding high-level language descriptions that encompass rich semantic information that can be generated by vision-language models. In this paper, we are the first to exploit the cross-modal description capabilities of CLIP to establish associations between subcategory-specific visual disparities in images and their corresponding textual descriptions. This simple but effective approach yields significant improvements on FGVC.

Vision-Language Learning

Vision-Language (VL) models, comprising image and text encoders, are trained in a contrastive manner using large-scale image-text pairs to learn a common feature space for images and textual labels. Recent advancements in contrastive training with large-scale Internet data, such as CLIP (Radford et al. 2021), ALIGN (Jia et al. 2021), and FILIP (Yao et al. 2022), have further scaled up VL models. Specifically, CLIP (Radford et al. 2021) has achieved remarkable success in zero-shot tasks by leveraging text prompts to transfer visual concepts to downstream tasks. Consequently, numerous subsequent approaches (Khattak et al. 2023; Lee et al. 2023; Yang et al. 2023) have been proposed to exploit pre-trained VL models for diverse downstream tasks. Drawing inspiration from recent developments in NLP, prompt techniques have begun to be explored in multimodal computer vision problems for transferring knowledge from large-scale VL models to downstream tasks. In the case of ViTs, some methods (Yu et al. 2023; Wasim et al. 2023) employ trainable prompts to guide pre-trained ViTs. Conversely, certain methods (Zhou et al.

2022b) introduce learnable vectors into the text encoder of CLIP for transfer learning in image recognition tasks. For fine-grained understanding, some approaches propose utilizing pre-trained CLIP to address the problem of fine-grained image retrieval in open-set scenarios. However, to the best of our knowledge, no related work has specifically focused on how to develop a pre-trained CLIP model to solve FGVC tasks in close-set scenarios. The primary challenge lies in existing prompt strategies tailored towards capturing supercategory-level semantics rather than more detailed subcategory-level discrepancies. In response, we propose a multimodal prompts scheme and a multimodal adaptation scheme to enhance the sensitivity of the pre-trained CLIP model for solving FGVC tasks.

Preliminary

Review of Vision Transformer (ViT)

For the input image \mathbf{I} , ViT architecture initially split it into $N = N_h \times N_w$ non-overlapping patches of size $P \times P$. Here, N_h and N_w represent the number of patches in each column and row of the input image, respectively. Subsequently, these patches are transformed into embedding tokens $\mathbf{E} = [\mathbf{E}^1, \mathbf{E}^2, \dots, \mathbf{E}^N] \in \mathbb{R}^{N \times D}$ using a learnable linear projection $\mathbf{P}_{emb} \in \mathbb{R}^{P^2 \times D}$, where D denotes the output dimension for each token. Lastly, the embedded tokens \mathbf{E} and the class token $\mathbf{E}^{class} \in \mathbb{R}^D$ are concatenated and combined with an additional position embedding $\mathbf{E}_{pos} \in \mathbb{R}^{(N+1) \times D}$ to form the initial input token sequence as $\mathbf{E}_0 = [\mathbf{E}^{class}, \mathbf{E}^1, \mathbf{E}^2, \dots, \mathbf{E}^N] + \mathbf{E}_{pos}$. The ViT encoder comprises L transformer layers, each consisting of multi-head self-attention (MHSA) and multi-layer perception (MLP) blocks. Thus, when provided with the input \mathbf{E}_{i-1} , the output of the i -th layer can be expressed as:

$$\begin{aligned} \mathbf{E}'_i &= \text{MHSA}(\text{LN}(\mathbf{E}_{i-1})) + \mathbf{E}_{i-1}, \\ \mathbf{E}_i &= \text{MLP}(\text{LN}(\mathbf{E}'_i)) + \mathbf{E}'_i. \end{aligned} \quad (1)$$

Here, $i = \{1, 2, \dots, L\}$, and $\text{LN}(\cdot)$ denotes the layer normalization. The first token of the last layer, \mathbf{E}_L^0 , in ViT serves as the global feature representation and is forwarded to the classifier head to generate the final predictions.

Methodology

The overall pipeline of our proposed MP-FGVC is depicted in Fig. 1. Obviously, our framework comprises two key components: the **multimodal prompts scheme** and the **multimodal adaptation scheme**, both of which are constructed using a pre-trained CLIP model through a two-stage training process. In the first stage, we refine the multimodal prompts scheme to empower the pre-trained CLIP model in recognizing subcategory-specific semantics in the vision prompt, as well as generating discrepancy-aware descriptions in the text prompt. Subsequently, in the second stage, we optimize the multimodal adaptation scheme to fully exploit the benefits derived from the multimodal prompts scheme specifically for FGVC.

Multimodal Prompts Scheme

Subtle yet discriminative discrepancies are widely acknowledged as being crucial for FGVC. However, the conventional ViTs employed in existing CLIP-like vision-language models are originally designed to represent common visual concepts and recognize them into coarse categories, rather than effectively capturing subtle differences to predict subordinate categories of a given object. To address this, we propose a multimodal prompts scheme that encompasses two components: subcategory-specific vision prompt and discrepancy-aware text prompt.

Subcategory-specific Vision Prompt. An important challenge in FGVC is accurately localizing discriminative regions that capture subcategory-specific discrepancies among confusing subcategories. However, the ViT treats all patch-divided image regions equally, resulting in the inclusion of much redundant patch information that can negatively impact the final decision. This limitation restricts the applicability of the pre-trained CLIP model in FGVC tasks. To address this limitation, we propose a parameter-free operation called Subcategory-specific Vision Prompt (SSVP), which can be applied to the penultimate layer of the ViT to modify its input to the last layer. In SSVP, the importance scores generated by M heads in MHSA are considered to select discriminative tokens that reflect the attention given to each region. Let $\mathbf{A} = [\mathbf{A}^1, \mathbf{A}^2, \dots, \mathbf{A}^m, \dots, \mathbf{A}^M]$ denote the attention scores of the class token in the penultimate layer, where $\mathbf{A}^m \in \mathbb{R}^N$ represents the attention score of the m -th head and reflects the attention degree of class prediction towards each region. As shown in Fig. 1, we first aggregate the attention scores of all heads and then select the top- k valuable tokens with high activation values as $\text{id} = \text{TopK}(\sum_{i=1}^M \mathbf{A}^i) \in \mathbb{Z}^K$, where $\text{TopK}(\cdot)$ returns the index id of the first k tokens based on the aggregated attention scores, sorted in descending order. Instead of treating all tokens equally, we select k tokens by id and concatenate them along with the class token to form the input sequence of the L -th transformer layer, denoted as:

$$\mathbf{E}_{L-1} = [\mathbf{E}_{L-1}^{\text{class}}, \mathbf{E}_{L-1}^{\text{id}(1)}, \mathbf{E}_{L-1}^{\text{id}(2)}, \dots, \mathbf{E}_{L-1}^{\text{id}(k)}]. \quad (2)$$

By replacing the original entire input sequence of the last transformer layer with tokens that solely cover the discriminative regions of an object, the proposed SSVP encourages the image encoder to focus more on the subtle differences between different subcategories while disregarding less discriminative regions such as background or common features, thereby facilitating fine-grained understanding.

Discrepancy-aware Text Prompt. The prompt strategies employed by existing CLIP-like VL models primarily focus on capturing supercategory-level semantics rather than detailed visual discrepancies necessary for distinguishing fine-grained objects. Consequently, the acquisition of subcategory-specific visual discrepancies in FGVC tasks is constrained, despite the provision of concrete subcategory names for text prompts. To maintain semantic coherence with the subcategory-specific vision prompt, we propose a Discrepancy-aware Text Prompt (DATP) operation

that supplements the lack of textual information by pre-training a set of learnable text tokens. These tokens automatically generate appropriate text descriptions about more detailed visual discrepancies among subcategories. Specifically, the text prompt template fed into $\mathcal{T}(\cdot)$ is designed as $\mathbf{P}_{\text{class}}^c = [\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_j, \dots, \mathbf{X}_J, \text{class++}]$, where $c \in \{1, \dots, C\}$ and C denotes the number of subcategories. Each $\mathbf{X}_j \in \mathbb{R}^D$ represents the j -th learnable text token, and D represents the dimension of the token vector. Notably, $\text{class++} \in \mathbb{R}^D$ corresponds to the word embedding generated by the **supercategory** label (e.g., bird/dog/food), rather than specific subcategory labels. To ensure semantic coherence with the subcategory-specific vision prompt, each text prompt $\mathbf{P}_{\text{class}}^c$ serves as a discrepancy-aware textual description and shares the same supercategory name class++ within a dataset, thereby improving the computation efficiency. Ultimately, these learnable prompt vectors, in conjunction with the supercategory names, construct a discrepancy-aware text prompt that the text encoder of the CLIP model can understand. Consequently, this mechanism generates subcategory-specific discrepancy descriptions, which serve as additional discriminative cues.

Multimodal Adaptation Scheme

The objective of the multimodal prompts scheme is to enhance the sensitivity of the pre-trained CLIP model to subcategory-specific discrepancies from the perspectives of both vision and text prompts. To enhance the performance of multimodal prompting on FGVC, we additionally introduce a multimodal adaptation scheme by incorporating a Vision-Language Fusion Module (VLFM) to align the vision and text prompts in a shared semantic space, thereby facilitating collaborative reasoning across modalities.

Vision-Language Fusion Module. As shown in Fig. 1, VLFM first mutually aligns the vision prompt and text prompts in a semantic space using a cross-modal attention mechanism. Specifically, we consider the class token of the ViT as the subcategory-specific visual embedding, denoted as \mathbf{E}_V , and the textual features generated by the learned text tokens $\mathbf{P}_{\text{class}} = \{\mathbf{P}_{\text{class}}^1, \mathbf{P}_{\text{class}}^2, \dots, \mathbf{P}_{\text{class}}^C\}$ as the textual embeddings that are aware of the subcategory-specific discrepancies, denoted as \mathbf{E}_T . With the multimodal input $[\mathbf{E}_V, \mathbf{E}_T]$, we generate three visual embeddings, namely the query (\mathbf{Q}_0), key (\mathbf{K}_0), and value (\mathbf{V}_0), using a 1×1 convolution layer. Similarly, a key-value textual embedding pair ($\mathbf{K}_c - \mathbf{V}_c$) is generated for each textual embedding \mathbf{E}_T^c ($c \in \{1, \dots, C\}$) using the same process. Thus, the cross-modal attention vector \mathbf{S}_i is computed between the visual query vector \mathbf{Q}_0 and all key vectors \mathbf{K}_i ($i = \{0, 1, \dots, C\}$) as $\mathbf{S}_i = \text{Softmax}(\frac{\mathbf{Q}_0 \mathbf{K}_i^T}{\sqrt{D}})$, where D denotes the embedding dimension and T represents the transpose operation. Utilizing the cross-modal attention \mathbf{S}_i , the subcategory-specific descriptions from the text prompts are transformed into subcategory-specific visual discrepancies, given by $\mathbf{E}'_V = \sum_{i=0}^C \mathbf{S}_i \otimes \mathbf{V}_i$, which aims to recover the missing discrepancies and form a discriminative and generalized visual embedding. Furthermore, to incorporate more complete discrepancies into the VLFM, we introduce two

1×1 convolution layers that constitute a translation layer denoted as $\text{Trans}(\cdot)$. This layer is applied to \mathbf{E}'_V to obtain $\hat{\mathbf{E}}_V = \text{Trans}(\mathbf{E}'_V)$. The first convolution layer, which has an expansion factor of 4, serves as an expansion layer, while the second layer maps the feature dimension back to the original input dimension.

Optimization

In this paper, we aim to fully exploit multimodal prompting for FGVC by utilizing the recently introduced CLIP model. However, in most FGVC tasks, only subcategory labels or indexes are available, which lack specific textual descriptions for fine-grained understanding. As a result, the adoption of CLIP-like VL models in FGVC has been limited. To address this issue, we propose a two-stage training strategy illustrated to enhance the sensitivity of the proposed MP-FGVC to subcategory-specific discrepancies, thereby facilitating efficient adaptation for FGVC tasks.

The first training stage. In this stage, we fix the parameters of $\mathcal{T}(\cdot)$ and optimize only the parameters of the image encoder $\mathcal{I}(\cdot)$ and the learned tokens in DATP are optimized. Similar to CLIP, we employ the image encoder $\mathcal{I}(\cdot)$ and the text encoder $\mathcal{T}(\cdot)$ to obtain the visual embeddings $\mathbf{V}_i \in \mathbb{R}^D$ and the textual embeddings $\mathbf{T}_i \in \mathbb{R}^{C \times D}$. However, instead of using the original text as input to $\mathcal{T}(\cdot)$, we replace it with \mathbf{P}_{class} , where $\mathbf{P}_{class} = [\mathbf{P}_{class}^1, \mathbf{P}_{class}^2, \dots, \mathbf{P}_{class}^C]$. Here, we use \mathbf{T}_i^c to indicate the textual feature of subcategory c . To optimize $\mathcal{I}(\cdot)$, we simply utilize \mathbf{V}_i to generate subcategory predictions \hat{y}_1 and compute the cross-entropy loss \mathcal{L}_v as follows:

$$\mathcal{L}_v = \text{CrossEntropy}(\hat{y}_1, y), \quad (3)$$

where y is the one-hot label. By minimizing \mathcal{L}_v , we aim to project the semantic features into a new space where subcategory-specific discrepancies are highlighted. Furthermore, to generate discrepancy-aware textual features, we employ the image-to-text contrastive loss \mathcal{L}_{i2t} learnable tokens in DATP. This loss is computed as:

$$\mathcal{L}_{i2t} = - \sum_{c=1}^C y_c \cdot \log \frac{\exp(\cos(\mathbf{V}_i, \mathbf{T}_i^c)/\tau)}{\sum_{k=1}^C \exp(\cos(\mathbf{V}_i, \mathbf{T}_i^k)/\tau)}, \quad (4)$$

where τ is the temperature parameter and is set to 0.07. By minimizing both \mathcal{L}_v and \mathcal{L}_{i2t} , the gradients are back-propagated through the unfixed $\mathcal{I}(\cdot)$ and fixed $\mathcal{T}(\cdot)$ to identify subcategory-specific semantics in the vision prompt and generate discrepancy-aware descriptions in the text prompt. Thus, the overall loss function of the first stage is formulated as:

$$\mathcal{L}_{stage1} = \mathcal{L}_v + \mathcal{L}_{i2t}. \quad (5)$$

The second training stage. In this stage, we fix the parameters of $\mathcal{I}(\cdot)$ and $\mathcal{T}(\cdot)$, and only optimize VLFM to align the vision prompt and text prompts in the semantic space, in order to fully exploit the advantages of the multimodal prompts scheme for final classification. We utilize the visual features $\hat{\mathbf{E}}_V$ output by VLFM to predict the subcategory label \hat{y}_2 and calculate the cross-entropy loss \mathcal{L}_{vlfm} as in Eq. 3. Hence, the loss for the second training stage is:

$$\mathcal{L}_{stage2} = \mathcal{L}_{vlfm} = \text{CrossEntropy}(\hat{y}_2, y). \quad (6)$$

Method	Backbone	CUB	Dogs
MRDMN	CNN	88.8	89.1
FDL	CNN	89.1	84.9
MSHQP	CNN	89.0	90.4
API-Net	CNN	90.0	90.3
PRIS	CNN	90.0	90.7
CAL	CNN	90.6	88.7
TransFG*	ViT-B-16	91.1	90.5
IELT*	ViT-B-16	90.9	90.6
SIM-Trans*	ViT-B-16	91.3	89.2
MP-FGVC (Ours)	ViT-B-16	91.8	91.0

Table 1: Comparisons with competitive methods on CUB-200-2011 and Stanford Dogs, * indicates the result reproduced according to the source code.

Experiments

Experiments Setup

Dataset. **CUB-200-2011** (Wah et al. 2011) comprises 11,788 bird images from 200 bird species, which are officially split into 5,994 training images and 5,794 test images.

Stanford Dogs (Khosla et al. 2011) consists of 20,580 images depicting 120 dog variants, with 12,000 images allocated for training and 8,580 images designated for testing.

NABirds (Horn et al. 2015) contains 48,562 images showcasing North American birds across 555 subcategories. It is split into 23,929 training images and 24,633 test images.

Food101 (Bossard, Guillaumin, and Gool 2014) encompasses 101 different kinds of foods, totaling 101,000 images. Within each class, 250 images are for testing, while the remaining 750 images are for training.

Implementation details. We employ the ViT-B-16 (Dosovitskiy et al. 2021a), as our image encoder $\mathcal{I}(\cdot)$. The text encoder $\mathcal{T}(\cdot)$, is a pre-trained Transformer model from CLIP (Radford et al. 2021). All input images are resized to a resolution of 448×448 . In the first training stage, we initialize the learning rate as $3e-2$, except for Stanford Dogs where it is initialized as $3e-3$. The number of training epochs is set to 30 for both CUB-200-2011 and Stanford Dogs, while the remaining datasets are trained for 10 epochs. For the second training stage, we initialize the learning rate as $1e-3$, except for Stanford Dogs where it is initialized as $1e-4$. The number of training epochs is set to 10 for both CUB-200-2011 and Stanford Dogs, while the other datasets are trained for 3 epochs. Both training stages utilize the SGD optimizer and employ cosine annealing as the optimization scheduler.

Comparison with State-of-the-Art Methods

The comparison results for CUB-200-2011 and Stanford Dogs datasets are presented in Table 1, while the results for the NABirds dataset can be found in Table 2. Furthermore, the comparison results for the Food101 dataset are

Method	Backbone	NABirds
Cross-X	CNN	86.4
DSTL	CNN	87.9
GHRD	CNN	88.0
API-Net	CNN	88.1
PRIS	CNN	88.4
MGE-CNN	CNN	88.6
TransFG*	ViT-B-16	<u>90.5</u>
IELT*	ViT-B-16	90.1
SIM-Trans*	ViT-B-16	90.3
MP-FGVC (Ours)	ViT-B-16	91.0

Table 2: Comparisons with competitive methods on NABirds.

Method	Backbone	Food101
DCL	CNN	88.9
SGLANet	CNN	89.7
IG-CMAN	CNN	90.4
MSMVFA	CNN	90.6
PRENet	CNN	91.1
TransFG*	ViT-B-16	<u>92.1</u>
IELT*	ViT-B-16	91.4
SIM-Trans*	ViT-B-16	91.9
MP-FGVC (Ours)	ViT-B-16	93.0

Table 3: Comparisons with competitive methods on Food101.

provided in Table 3. From these tables, it can be observed that our proposed method outperforms other state-of-the-art methods across all datasets. Specifically, our method achieves an improvement of 0.5%, 0.3%, 0.5%, and 0.9% in terms of the Top-1 Accuracy metric compared to the second-best results on the CUB-200-2011, Stanford Dogs, NABirds, and Food101 datasets, respectively. Notably, our method exhibits substantial improvements over the ViT-based methods, which solely rely on visual training. These findings serve as compelling evidence of the effectiveness of our proposed MP-FGVC.

Ablation Studies and Analysis

Efficacy of various components. The proposed MP-FGVC comprises three essential components: SsVP and DATP from multimodal prompts scheme, and VLFM from multimodal adaptation scheme. We conducted ablation experiments on these components, and the results are reported in Table 5, with the Baseline representing the pure ViT. The introduction of DATP led to a 0.5% improvement in multimodal prompting performance on the CUB-200-2011 dataset and a 0.4% improvement on the NABirds dataset. By utilizing VLFM to enhance FGVC with language modality, further performance gains were achieved on both CUB-200-2011 and NABirds datasets. This observation indicates that DATP incorporates subcategory-specific discrepancy

Text Prompt Templates	CUB	NABirds
“a photo of a <code>class</code> ”	91.5	90.7
“a photo of $\mathbf{X}_1, \dots, \mathbf{X}_J, \text{class}$ ”	91.7	90.9
“a photo of $\mathbf{X}_1, \dots, \mathbf{X}_J, \text{class++}$ ”	91.8	91.0
“ $\mathbf{X}_1, \dots, \mathbf{X}_J, \text{class++}$ ”	91.8	91.0

Table 4: Comparisons with different text prompt design of accuracy (%) on CUB-200-2011 and NABirds.

Method	CUB	NABirds
Baseline	90.8	90.0
Baseline + DATP	91.3	90.4
Baseline + DATP + VLFM	91.5	90.6
Baseline + SsVP	91.2	90.3
Baseline + SsVP + DATP	91.5	90.8
Baseline + SsVP + DATP + VLFM	91.8	91.0

Table 5: The accuracy (%) results of component ablation study on CUB-200-2011 and NABirds.

Strategies	CUB	NABirds
One stage	91.1	90.3
Two stage	91.8	91.0

Table 6: Comparison of different training strategies on CUB-200-2011 and NABirds.

descriptions that complement the image content, thereby improving the performance. Moreover, the inclusion of SsVP in all model variations improved their performance on the CUB-200-2011 and NABirds datasets, demonstrating the effectiveness of selecting discriminative regions to identify subcategory-specific discrepancies.

Comparison of various text prompt templates. We conduct a comparative experiment to analyze the impact of different text prompt templates in DATP. Specifically, `class` denotes the concrete subcategory name and `class++` denotes the ambiguous supercategory name. The results in Table 4 demonstrate that learnable text tokens for the text prompt outperform handcrafted prompt templates such as “a photo of a `class`”. This indicates that our DATP can learn a set of words that describe the subcategory-specific discrepancies between subcategories, leading to improved performance on FGVC. Moreover, the performance is not significantly affected by whether “a photo of a” is used as a prefix or if the subcategory name is used as a suffix.

Number of selected visual tokens in SsVP. The performance comparison of different values of k , representing the number of selected visual tokens in SsVP, is presented in Figure 2. From the analysis, it is observed that performance degrades when k exceeds 14. This performance drop may be attributed to the fact that as the number of selected tokens increases, the vision prompt may highlight the locations of objects or parts instead of the subcategory-specific differences, rendering it ineffective for FGVC. Consequently, we




Figure 2: Analyses of hyper-parameter k on CUB-200-2011 and NABirds.




Figure 3: Analyses of hyper-parameter J on CUB-200-2011 and NABirds.

have chosen to set k to 14 for all datasets.

Number of learnable text tokens in DATP. What is the appropriate number of learnable tokens in DATP for accurately summarizing subcategory-specific discrepancies? We conduct analyses on the length J of the text prompt template, and present evaluation results in Fig. 3. The performance gradually increases with the number of learnable tokens but decreases when the prompt length exceeds 32. One potential explanation for this observation is that as the length of the text prompt reaches 32, the model tends to prioritize optimizing DATP itself, hindering the optimization process for the subcategory-specific vision prompt. Thus, we set the token number in DATP to 16 for all datasets.

Necessity of two-stage training strategy. Our proposed MP-FGVC employs a two-stage training strategy to align and fuse embeddings from text and image domains. It is worth noting that integrating the learning of vision prompt, text prompt, and multimodal adaptation into single-stage training, similar to CLIP-like VL models, is possible. To de-

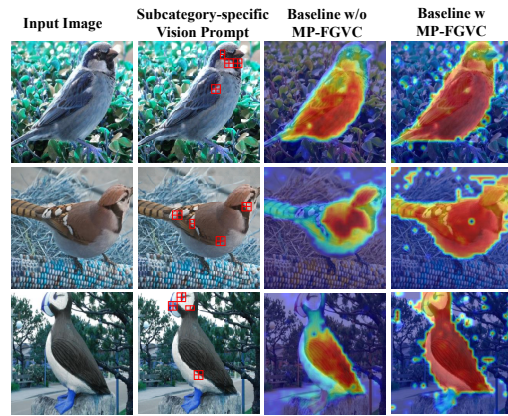


Figure 4: Visualization of MP-FGVC on CUB-200-2011.

termine the effectiveness of each strategy, we conduct comparison experiments between the two-stage training and the one-stage training. Table 6 demonstrates that the one-stage training strategy performs less effectively. The inadequate description of subcategory-specific discrepancies by the randomly initialized text tokens during the early training stage hinders the overall optimization of the model. Therefore, we adopt the two-stage training to retain the hidden states of the pre-trained image encoder and text encoder, enabling CLIP to maintain its advantages on FGVC.

Visualization of MP-FGVC. We conduct visualization experiments to illustrate the effectiveness of MP-FGVC, as shown in Fig. 4. The input images are depicted in the first column, followed by the second column which represents the selected vision tokens ($k=14$) in our SSVP. The third and fourth columns display the class activation maps computed using GradCAM for the baseline model with and without our proposed MP-FGVC, respectively. The visualization results demonstrate that MP-FGVC enhances the baseline model’s ability to focus on a more comprehensive region, particularly in capturing more detailed discrepancies specific to subcategories, such as the head, back, and abdomen.

Conclusion

In this paper, we introduce MP-FGVC, a novel multimodal prompting solution designed to improve fine-grained visual classification through the utilization of the cross-modal capabilities of the pre-trained CLIP model. By employing a multimodal prompts scheme and a multimodal adaptation scheme, MP-FGVC effectively highlights subcategory-specific discrepancies and achieves cross-modal reasoning, ultimately leading to improved FGVC performance with minimal modifications. Experimental evaluations conducted on four FGVC datasets validate the effectiveness of MP-FGVC. As the pioneer effort to extend the applicability of large-scale vision-language models to FGVC, MP-FGVC incorporates novel prompting and adaptation techniques. It is our aspiration that our findings serve as inspiration and facilitate further research on efficient adaptation methods of large-scale vision-language models for FGVC tasks.

Acknowledgements

This work was partially supported by the National Key Research and Development Program of China under Grant 2022ZD0118802, the National Natural Science Foundation of China (Grant No. U20B2064 and 62172226), Singapore MOE Tier 1 Funds (MSS23C002), AISG Research Grant (No. AISG3-GV-2023-011), and the Postgraduate Research & Practice Innovation Program of Jiangsu Province (KYCX230491).

References

- Bahng, H.; Jahanian, A.; Sankaranarayanan, S.; and Isola, P. 2022. Visual Prompting: Modifying Pixel Space to Adapt Pre-trained Models. *CoRR*, abs/2203.17274.
- Bossard, L.; Guillaumin, M.; and Gool, L. V. 2014. Food-101 - Mining Discriminative Components with Random Forests. In *ECCV*, volume 8694, 446–461.
- Dong, N.; Zhang, L.; Yan, S.; Tang, H.; and Tang, J. 2023. Erasing, transforming, and noising defense network for occluded person re-identification. *IEEE Transactions on Circuits and Systems for Video Technology*.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Houlsby, N. 2021a. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *ICLR*.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Houlsby, N. 2021b. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *ICLR*.
- Gao, Y.; Han, X.; Wang, X.; Huang, W.; and Scott, M. R. 2020. Channel Interaction Networks for Fine-Grained Image Categorization. In *AAAI*, 10818–10825.
- Ge, W.; Lin, X.; and Yu, Y. 2019. Weakly Supervised Complementary Parts Models for Fine-Grained Image Classification From the Bottom Up. In *CVPR*, 3034–3043.
- Gu, Z.; and Feng, S. 2023. Individuality Meets Commonality: A Unified Graph Learning Framework for Multi-View Clustering. *ACM Transactions on Knowledge Discovery from Data*, 17(1): 7:1–7:21.
- Gu, Z.; Feng, S.; Hu, R.; and Lyu, G. 2023. ONION: Joint Unsupervised Feature Selection and Robust Subspace Extraction for Graph-based Multi-View Clustering. *ACM Transactions on Knowledge Discovery from Data*, 17(5): 70:1–70:23.
- He, J.; Chen, J.; Liu, S.; Kortylewski, A.; Yang, C.; Bai, Y.; and Wang, C. 2022. TransFG: A Transformer Architecture for Fine-Grained Recognition. In *AAAI*, 852–860.
- Horn, G. V.; Branson, S.; Farrell, R.; Haber, S.; Barry, J.; Ipeirotis, P.; Perona, P.; and Belongie, S. J. 2015. Building a bird recognition app and large scale dataset with citizen scientists: The fine print in fine-grained dataset collection. In *CVPR*, 595–604.
- Hu, Y.; Jin, X.; Zhang, Y.; Hong, H.; Zhang, J.; He, Y.; and Xue, H. 2021. RAMS-Trans: Recurrent Attention Multi-scale Transformer for Fine-grained Image Recognition. In *ACM Multimedia*, 4239–4248.
- Jia, C.; Yang, Y.; Xia, Y.; Chen, Y.; Parekh, Z.; Pham, H.; Le, Q. V.; Sung, Y.; Li, Z.; and Duerig, T. 2021. Scaling Up Visual and Vision-Language Representation Learning With Noisy Text Supervision. In *ICML*, volume 139, 4904–4916.
- Jia, M.; Tang, L.; Chen, B.; Cardie, C.; Belongie, S. J.; Hariharan, B.; and Lim, S. 2022. Visual Prompt Tuning. In *ECCV*, volume 13693, 709–727.
- Khattak, M. U.; Rasheed, H.; Maaz, M.; Khan, S.; and Khan, F. S. 2023. Maple: Multi-modal prompt learning. In *CVPR*, 19113–19122.
- Khosla, A.; Jayadevaprakash, N.; Yao, B.; and Fei-Fei, L. 2011. Novel Dataset for Fine-Grained Image Categorization. In *CVPR Workshop*, volume 2.
- Lee, Y.; Tsai, Y.; Chiu, W.; and Lee, C. 2023. Multimodal Prompting with Missing Modalities for Visual Recognition. In *CVPR*, 14943–14952.
- Li, S.; Sun, L.; and Li, Q. 2023. CLIP-ReID: Exploiting Vision-Language Model for Image Re-identification without Concrete Text Labels. In *AAAI*, volume 37, 1405–1413.
- Li, Z.; Tang, H.; Peng, Z.; Qi, G.-J.; and Tang, J. 2023. Knowledge-guided semantic transfer network for few-shot image recognition. *IEEE Transactions on Neural Networks and Learning Systems*.
- Liu, C.; Xie, H.; Zha, Z.; Ma, L.; Yu, L.; and Zhang, Y. 2020. Filtration and Distillation: Enhancing Region Attention for Fine-Grained Visual Categorization. In *AAAI*, 11555–11562.
- Luo, W.; Yang, X.; Mo, X.; Lu, Y.; Davis, L.; Li, J.; Yang, J.; and Lim, S. 2019. Cross-X Learning for Fine-Grained Visual Categorization. In *ICCV*, 8241–8250.
- Maier, M.; and Abdel Rahman, R. 2019. No matter how: Top-down effects of verbal and semantic category knowledge on early visual perception. *Cognitive, Affective, & Behavioral Neuroscience*, 19: 859–876.
- Min, W.; Wang, Z.; Liu, Y.; Luo, M.; Kang, L.; Wei, X.; Wei, X.; and Jiang, S. 2023. Large Scale Visual Food Recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(8): 9932–9949.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; Krueger, G.; and Sutskever, I. 2021. Learning Transferable Visual Models From Natural Language Supervision. In *ICML*, volume 139, 8748–8763.
- Shen, F.; Xie, Y.; Zhu, J.; Zhu, X.; and Zeng, H. 2023. Git: Graph interactive transformer for vehicle re-identification. *IEEE Transactions on Image Processing*.
- Sun, H.; He, X.; and Peng, Y. 2022. SIM-Trans: Structure Information Modeling Transformer for Fine-grained Visual Categorization. In *ACM Multimedia*, 5853–5861.
- Tang, H.; Li, Z.; Peng, Z.; and Tang, J. 2020. Block-Mix: Meta Regularization and Self-Calibrated Inference for

Metric-Based Meta-Learning. In *ACM Multimedia*, 610–618.

Tang, H.; Liu, J.; Yan, S.; Yan, R.; Li, Z.; and Tang, J. 2023. M3Net: Multi-view Encoding, Matching, and Fusion for Few-shot Fine-grained Action Recognition. In *ACM Multimedia*, 1719–1728.

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.

Wah, C.; Branson, S.; Welinder, P.; Perona, P.; and Belongie, S. 2011. The caltech-ucsd birds-200-2011 dataset.

Wang, S.; Chang, J.; Li, H.; Wang, Z.; Ouyang, W.; and Tian, Q. 2023. Open-Set Fine-Grained Retrieval via Prompting Vision-Language Evaluator. In *CVPR*, 19381–19391.

Wasim, S. T.; Naseer, M.; Khan, S. H.; Khan, F. S.; and Shah, M. 2023. Vita-CLIP: Video and text adaptive CLIP via Multimodal Prompting. In *CVPR*, 23034–23044.

Xu, K.; Lai, R.; Gu, L.; and Li, Y. 2021. Multiresolution Discriminative Mixup Network for Fine-Grained Visual Categorization. *IEEE Transactions on Neural Networks and Learning Systems*, 1–13.

Xu, Q.; Wang, J.; Jiang, B.; and Luo, B. 2023. Fine-Grained Visual Classification Via Internal Ensemble Learning Transformer. *IEEE Transactions on Multimedia*, 1–14.

Yan, S.; Tang, H.; Zhang, L.; and Tang, J. 2023. Image-specific information suppression and implicit local alignment for text-based person search. *IEEE Transactions on Neural Networks and Learning Systems*.

Yang, L.; Wang, Y.; Li, X.; Wang, X.; and Yang, J. 2023. Fine-Grained Visual Prompting. *CoRR*, abs/2306.04356.

Yao, L.; Huang, R.; Hou, L.; Lu, G.; Niu, M.; Xu, H.; Liang, X.; Li, Z.; Jiang, X.; and Xu, C. 2022. FILIP: Fine-grained Interactive Language-Image Pre-Training. In *ICLR*.

Yu, C.; Zhao, X.; Zheng, Q.; Zhang, P.; and You, X. 2018. Hierarchical Bilinear Pooling for Fine-Grained Visual Recognition. In *ECCV*, volume 11220, 595–610.

Yu, Y.; Chung, J.; Yun, H.; Hessel, J.; Park, J. S.; Lu, X.; Zellers, R.; Ammanabrolu, P.; Bras, R. L.; Kim, G.; and Choi, Y. 2023. Fusing Pre-Trained Language Models with Multimodal Prompts through Reinforcement Learning. In *CVPR*, 10845–10856.

Zhou, K.; Yang, J.; Loy, C. C.; and Liu, Z. 2022a. Conditional Prompt Learning for Vision-Language Models. In *CVPR*, 16795–16804.

Zhou, K.; Yang, J.; Loy, C. C.; and Liu, Z. 2022b. Learning to Prompt for Vision-Language Models. *International Journal of Computer Vision*, 130(9): 2337–2348.

Zhuang, P.; Wang, Y.; and Qiao, Y. 2020. Learning Attentive Pairwise Interaction for Fine-Grained Classification. In *AAAI*, 13130–13137.