

Domain Prompt Learning with Quaternion Networks

Qinglong Cao^{1,2}, Zhengqin Xu¹, Yuntian Chen^{2*}, Chao Ma^{1*}, and Xiaokang Yang¹

¹ MoE Key Lab of Artificial Intelligence, AI Institute, Shanghai Jiao Tong University, Shanghai, China

² Ningbo Institute of Digital Twin, Eastern Institute of Technology, Ningbo, China

{caoql2022, fate311}@sjtu.edu.cn, ychen@eitech.edu.cn, {chaoma, xkyang}@sjtu.edu.cn

Abstract

Prompt learning has emerged as a potent and resource-efficient technique in large Vision-Language Models (VLMs). However, its application in adapting VLMs to specialized domains like remote sensing and medical imaging, termed domain prompt learning, remains relatively unexplored. Although large-scale domain-specific foundation models offer a potential solution, their focus on a singular vision level presents challenges in prompting both vision and language modalities. To address this limitation, we propose leveraging domain-specific knowledge from these foundation models to transfer the robust recognition abilities of VLMs from generalized to specialized domains, employing quaternion networks. Our method entails utilizing domain-specific vision features from domain-specific foundation models to guide the transformation of generalized contextual embeddings from the language branch into a specialized space within quaternion networks. Furthermore, we introduce a hierarchical approach that derives vision prompt features by analyzing intermodal relationships between hierarchical language prompt features and domain-specific vision features. Through this mechanism, quaternion networks can effectively explore intermodal relationships in specific domains, facilitating domain-specific vision-language contrastive learning. Extensive experiments conducted on domain-specific datasets demonstrate that our proposed method achieves new state-of-the-art results in prompt learning. Codes are available at <https://github.com/caoql98/DPLQ>.

1. Introduction

Supervised learning using large-scale training samples has shown remarkable success in various visual understanding tasks [2, 12, 30, 36, 37, 44]. However, conventional supervised learning often requires a significant amount of data. To address this issue, large-scale Vision-Language Models




Figure 1. Performance comparison of the novel class generalization. Our method achieves state-of-the-art performance on various remote sensing image datasets in terms of harmonic mean.

(VLMs) [15, 28, 32] have recently emerged, making prompt learning a more popular and efficient paradigm for various vision tasks. Prompt learning [3, 22, 29, 38, 46, 48] focuses on adapting VLMs effectively for downstream tasks with limited training data.

Current VLMs have achieved remarkable progress in contrastive learning with extensive image-text pairs. Among these models, Contrastive Language-Image Pre-training (CLIP) [28] is widely recognized for its exceptional zero-shot generalization capabilities. However, the reliance on fixed prompts challenges the adaptability of CLIP to downstream tasks. CoOp [48] introduces context optimization by using learnable context vectors as language prompts to address this limitation, which enhances the adaptability of CLIP for visual recognition tasks. However, challenges related to class shifting remain. Thus, CoCoOp [47] further introduces input-conditional prompts for the language branch to refine the alignment of both vision and language branches. MaPLe [17] goes further by introducing a multi-modal prompt learning approach. MaPLe fine-tunes vision and language branches concurrently, leading to more harmonious alignment and improved performance.

* Corresponding authors.

It is worth noting that these prompt learning techniques primarily focus on natural images, and these models without domain-specific knowledge tend to misinterpret specific domain images in natural image patterns. Consequently, the existing adaptation of VLMs to specific domains such as remote sensing and medical imaging would be compromised and further result in a significant performance gap. Large-scale domain-specific foundation model [23, 33, 35] could be leveraged to address this challenge. However, existing domain-specific foundation models are only pre-trained at the vision level, lacking inherent support for prompting vision-language pairs in a contrastive learning framework.

We draw inspiration from quaternion networks [26], renowned for their effective modeling of orthogonal relations and adept exploration of inter- and intra-correlations within the quaternion hidden space. Leveraging domain-specific knowledge from foundation models, we transfer the robust recognition abilities of VLMs from generalized to specialized domains using quaternion networks. Our approach involves propagating domain-specific vision features from foundation models and generalized contextual embeddings from the language branch into the quaternion hidden space. Within this space, the quaternion network extracts domain-specific cross-modal knowledge for the language branch and projects the generalized contextual embeddings into the specialized domain. To mitigate potential overfitting concerns [47] during prompt learning process, we introduce random noise into the quaternion hidden space, thereby enhancing the robustness of learning process.

Pre-trained VLMs establish a strong and consistent vision-language matching relationship. Thus, we can easily forward the domain-specific information from the specialized language branch into the vision branch. We utilize the learnable language prompt feature as input and the orthogonal domain-specific vision features as guidance to mine the intermodal relations in each vision-language layer, which hierarchically provides vision prompt features. As shown in Figure 1, both the vision and language branches become domain-specific, and the domain-specific contrastive learning minimizes the domain gap, leading to uncompromised domain-specific recognition performance.

In summary, we make the following contributions:

- To our knowledge, we first introduce the quaternion concept into prompt learning for specific domains. It successfully transfers the strong recognition ability of VLMs from the generalized domain to specialized fields such as remote sensing and medical images.
- We forward domain-specific info of pre-trained VLMs hierarchically to the vision branch via quaternion hidden space, enabling better recognition performance.
- We extensively validate our method on large-scale domain-specific datasets, and our method achieves state-of-the-art performance.

2. Related Work

Vision Language Models. Vision Language Models (VLMs) aim to establish a stable connection between visual content and textual descriptions by creating a unified embedding space encompassing visual and linguistic modalities. These models can generally be classified into three main groups: models with contrastive objectives [19, 24, 28], generative objectives [1, 18], and alignment objectives [9, 31]. Among these models, CLIP [28] pioneered a symmetrical image-language contrastive loss, enabling robust zero-shot prediction capabilities for various downstream tasks. However, its effectiveness heavily relies on the availability of large and expensive image-text paired datasets. To mitigate this dependency, ALIGN [15] leverages large-scale noisy image-text pairs, achieving comparable performance through a noise-robust contrastive learning approach. Subsequently, various researchers have explored more efficient VL model pre-training with fewer image-text pairs. For example, OTTER [39] uses optimal transport distillation to establish a soft image-text correspondence, enabling efficient zero-shot recognition with significantly reduced training data. DeCLIP [19] harnesses self-supervision, multi-view supervision, and nearest-neighbor supervision to extract valuable information from limited data. ZeroVL [6] successfully employs debiased data sampling and coin flipping mixup for efficient contrastive learning. In contrast to models primarily designed for recognition tasks, there is a separate category of VL models [4, 14, 45] oriented toward captioning tasks. For instance, COCA [45] employs an encoder-decoder architecture to align embeddings between images and their corresponding captions, yielding impressive results. Through utilizing these advanced VL models, many traditional vision tasks like object detection [10], semantic segmentation [41], and caption [34], have achieved great progress. Yet, these works are still limited to natural images. In this paper, we creatively transfer the VLMs into the specialized domains with quaternion networks.

Prompt Learning. VLMs offer adaptability to downstream tasks through full fine-tuning or linear probing methods. However, full fine-tuning can be computationally intensive and put the established cross-modal representations at risk. On the other hand, linear probing has limitations, particularly concerning models with zero-shot capabilities like CLIP. In the field of natural language processing [20], prompt learning techniques [17, 47, 48] have been proposed, which are efficient for VL models. For instance, CoOp [48] proposes an effective prompt-learning approach where language prompts are modeled with learnable embeddings while preserving the pre-trained parameters. However, the learned context in CoOp might lack generalizability to unseen classes within the same dataset.

To overcome this challenge, CoCoOp [47] introduces input-conditional tokens for each image, which act as dynamic prompts to mitigate class shift issues and enhance generalization performance. To prevent prompt tuning from causing a loss of general knowledge, ProGrad [50] selectively updates prompts based on their gradients aligning with the pre-defined prompt predictions. Despite the promise, existing prompt learning methods mostly focus on the language branch, with limited consideration for the image branch and its depth. To fill this gap, MaPLe [17] proposes an effective prompting approach where both vision and language branches are aligned simultaneously. However, these prompt learning methods still focus on natural images and few consider the domain-specific adaption problem. To address this issue, we proposed to leverage the domain-specific knowledge from the domain-specific foundation model to achieve domain prompt learning with quaternion networks.

3. Preliminaries

Quaternion Networks. In four-dimensional space, a quaternion Q extends a complex number and can be expressed as follows:

$$Q = r1 + xi + yj + zk, \quad (1)$$

where $r, x, y,$ and z are real numbers, and $1, \mathbf{i}, \mathbf{j},$ and \mathbf{k} are the quaternion unit basis. The real part of Q is denoted by r , while $x\mathbf{i} + y\mathbf{j} + z\mathbf{k}$ is the imaginary or vector part. Quaternions are useful for describing spatial rotations and other applications because they contain embedded information that can be represented by a matrix of real numbers:

$$Q = \begin{bmatrix} r & -x & -y & -z \\ x & r & -z & y \\ y & z & r & -x \\ z & -y & x & r \end{bmatrix}. \quad (2)$$

A quaternion neural network can be defined as:

$$Q_{out} = \alpha(W \otimes Q), \quad (3)$$

where W represents the learnable parameters of the quaternion neural networks, \otimes denotes the Hadamard product, and α is the activation function defined as:

$$\alpha(Q) = f(r)1 + f(x)\mathbf{i} + f(y)\mathbf{j} + f(z)\mathbf{k}, \quad (4)$$

where f is any standard activation function. We suggest using quaternion networks to discover orthogonal intermodal relationships between domain-specific vision features from the domain-specific foundation models and contextual embeddings from the language branch. This approach is inspired by the unique feature processing pattern of the quaternion networks.

Prompt Learning. The CLIP model consists of a visual encoder and a text encoder that generate image embeddings and corresponding text embeddings, respectively. The CLIP model is based on the vision transformer and follows the same setting as previous methods [17, 47]. During training, the CLIP model maximizes the cosine similarity between the image and its matched text and minimizes the cosine similarity between the image and its unmatched text. This allows the CLIP model to perform zero-shot classification. To perform zero-shot classification, the text embedding ω_i is generated from a hand-crafted prompt, such as “a photo of *category*”, where *category* is the i -th class name. If there are C categories and the visual embedding of the image is x , the probability of the image belonging to the i -th class name is given by:

$$p(y_i|x) = \frac{\exp(\text{sim}(x, \omega_i)/\tau)}{\sum_{i=1}^C \exp(\text{sim}(x, \omega_i)/\tau)}, \quad (5)$$

where $\text{sim}()$ is the cosine similarity function and τ is the temperature hyperparameter.

4. Proposed Method

4.1. Overview

Aims to prompt VLMs efficiently from a generalized domain to specific domains like remote sensing and medical images, domain prompt learning with quaternion networks is proposed to facilitate the integration of domain-specific knowledge from large-scale foundation models into VLMs. Illustrated in Figure 2, the quaternion network enables the identification of cross-modal relationships between domain-specific vision features from the foundation model and generalized contextual embeddings from the language branch. Subsequently, this information is utilized to map the generalized contextual embeddings into the specialized domain. Furthermore, well-aligned vision-language relationships in pre-trained VLMs are leveraged to propagate domain-specific information from the specialized language branch into the vision branch. Consequently, the proposed domain-specific prompt learning significantly enhances recognition performance.

4.2. Domain-Specific Foundation Model

The emergence of large-scale domain-specific foundation models [23, 33, 35] has notably enhanced representation quality for downstream vision tasks, particularly in remote sensing and medical imaging domains. Building upon this progress, we integrate the large-scale remote sensing foundation model [35] and MedSAM [23] to offer crucial domain-specific knowledge for remote sensing and medical images, respectively. The remote sensing foundation model primarily employs architectures such as ViT [8] and Vi-TAE [42], training these networks with millions of remote

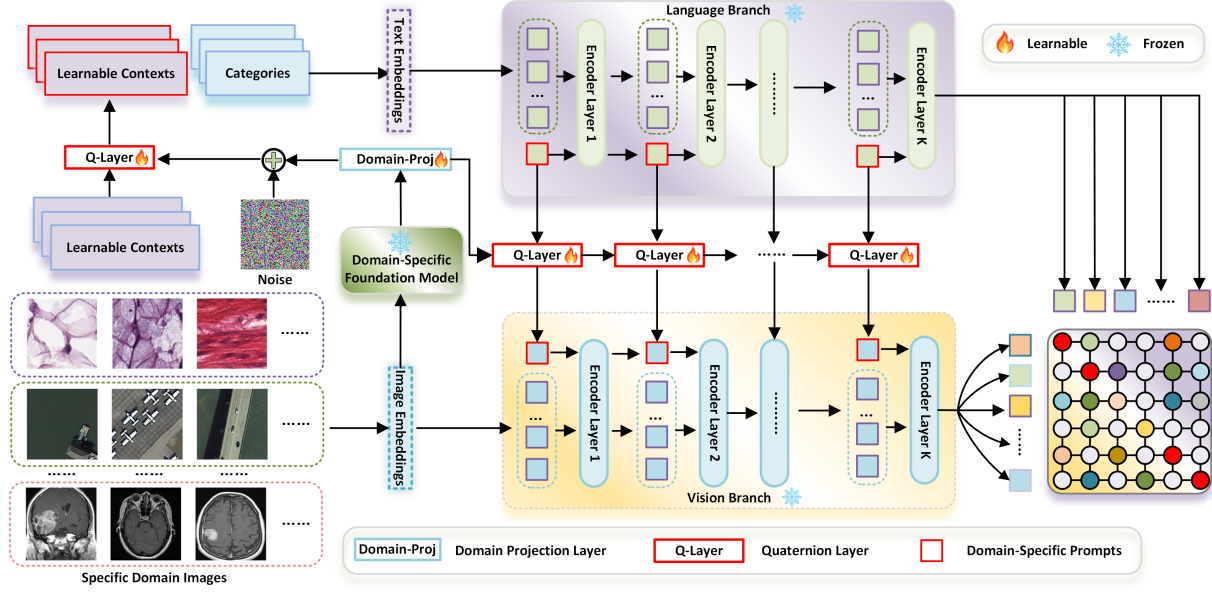


Figure 2. Overview of our proposed Domain Prompt Learning. We use the large-scale domain-specific foundation model as guidance, and exploit quaternion networks to mine the intermodal relationships between domain-specific vision features from the domain-specific foundation model and contextual embeddings from the language branch. Based on the stable vision-language matching relationships in pre-trained VLMs, the domain-specific information is hierarchically forwarded from the language branch to the vision branch.

sensing images via a masked autoencoder approach [13]. This framework reconstructs masked images, prioritizing the retrieval of visible portions within an encoder-decoder architecture. Network optimization involves minimizing the loss between the reconstructed regions and corresponding ground-truth masked regions. In the medical image processing domain, MedSAM [23] utilizes a meticulously curated dataset of over one million medical images for pre-training. To effectively guide the visual and language branches into the specific domain, we introduce domain-specific foundation models to furnish domain-specific knowledge.

4.3. Prompting Language Branch

To obtain domain-specific vision features F_d , we propagate the image patch embeddings into the domain-specific foundation model. To prepare the domain-specific information in the quaternion hidden space, we apply a domain-projection layer L_d consisting of two linear layers on F_d , resulting in the projected domain-specific vision features \hat{F}_d :

$$\hat{F}_d = L_d(F_d), \quad (6)$$

To mine the critical orthogonal intermodal relationship, we model the domain-specific vision features \hat{F}_d and the learnable context embeddings T_c in two orthogonal axes in the quaternion hidden space:

$$Q_l = T_c + \hat{F}_d \mathbf{i} + 0\mathbf{j} + 0\mathbf{k}, \quad (7)$$

To support the quaternion projection, we construct a zero tensor Z_0 with the same size of T_c . However, as mentioned

in CoCoOp [47], prompt learning is prone to the problem of overfitting due to the limited amount of data involved. To address this issue, we add some random Gaussian noise N_G into the quaternion hidden space to implicitly enhance the robustness of the learning process. The noise is scaled by the mean of the domain-specific vision features \hat{F}_d :

$$N_G = Mean(\hat{F}_d)N_\theta, \quad (8)$$

where N_θ denotes the standard Gaussian noise. Given quaternion layer Q_t , the generation of domain-specific context embedding T_d is computed as follows:

$$T_d = Q_t([\hat{F}_d + T_c + N_G, Z_0]), \quad (9)$$

In this way, the orthogonal intermodal relationship between domain-specific vision features and contextual embeddings is well-mined in the quaternion hidden space, and the generalized contextual embeddings are projected into the specialized space. To prompt the language encoder with the domain-specific information, we leverage domain-specific context embedding T_d and add a set of learnable language prompt features $[P_l^1, P_l^2, \dots, P_l^m]$ with the setting depth of k into encoder layers $[L_l^1, L_l^2, \dots, L_l^m]$ of the language branch, where m denotes the total layer numbers. Given the fixed text embedding C_t for categories, we first concatenate T_d with C_t to acquire the complete domain-specific text embeddings, i.e., $W_1 = [T_d, C_t]$. The domain-specific information propagation could then be computed as follows:

$$[W_i, -] = L_l^i([W_{i-1}, P_l^{i-1}]) \quad i = 1, \dots, k, \quad (10)$$

$$[W_j, P_l^j] = L_l^j([W_{j-1}, P_l^{j-1}]) \quad j = k+1, \dots, m, \quad (11)$$

4.4. Prompting Vision Branch

The pre-trained VLMs establish a solid vision-language matching relationship, which allows us to easily transfer domain-specific information from the language branch to the vision branch. To achieve this, we set a group of learnable vision prompt features, denoted by $[P_v^1, P_v^2, \dots, P_v^m]$, with a depth of k , which correspond to the encoder layers $[L_v^1, L_v^2, \dots, L_v^m]$ of the vision branch. Next, we introduce a group of quaternion layers $[Q_v^1, Q_v^2, \dots, Q_v^m]$ with a depth of k , which are responsible for providing cross-modal domain-specific information. In a similar way to how the quaternion computation is performed in the language branch, domain-specific vision features \hat{F}_d and language prompt features P_l^i are modeled in two orthogonal axes in the quaternion hidden space for the vision branch. We use the following equation to perform this computation:

$$Q_v^i = P_l^i + \hat{F}_d \mathbf{i} + 0\mathbf{j} + 0\mathbf{k}, \quad (12)$$

As the vision-language relationship is well-matched, random noise is no longer required. Therefore, we can compute the vision prompt features as follows:

$$P_v^i = Q_v^i([\hat{F}_d + P_l^i, Z_0]) \quad i = 1, 2, \dots, k, \quad (13)$$

By leveraging the well-matched vision-language relations, we propagate the domain-specific knowledge into the vision prompt features within the quaternion hidden space. To propagate domain-specific information for the vision branch, we start with the original image embeddings E_1 and class token c_1 , and compute the propagation as follows:

$$[E_i, c_i, -] = L_v^i([E_{i-1}, c_{i-1}, P_v^{i-1}]) \quad i = 1, \dots, k, \quad (14)$$

We then compute the following equation to propagate the domain-specific information further:

$$[E_j, c_i, P_v^j] = L_v^j([E_{j-1}, c_{j-1}, P_v^{j-1}]) \quad j = k+1, \dots, m, \quad (15)$$

Finally, given vision embeddings E_m of last vision layer and language embeddings $[W_m^1, W_m^2, \dots, W_m^C]$ for C categories of the last language layer, we compute the probability of an image belonging to a specific category as follows:

$$p(y_i|x) = \frac{\exp(\text{sim}(E_m, W_m^i)/\tau)}{\sum_{j=1}^C \exp(\text{sim}(E_m, W_m^j)/\tau)}, \quad (16)$$

where $p(y_i|x)$ is the probability of the image belonging to the i -th category, sim denotes the similarity function, and τ is a temperature parameter.

5. Experiments

In order to assess the effectiveness of our proposed quaternion learning approach, we conducted extensive experiments on remote sensing and medical images. Our experiments covered three distinct problem settings: 1) generalization from known to unknown classes within a dataset, 2) transfer between datasets, and 3) generalization across domains. This section provides a detailed description of the datasets, evaluation metrics, and experimental implementations. We also present a thorough performance analysis, as well as ablation experiments to clarify the effectiveness of our proposed designs.

5.1. Datasets and Evaluation Metrics

We evaluate the proposed method on 8 remote sensing datasets, namely MLRSNet [27], PatternNet [49], RSSCN7 [51], AID [40], RSICD [21], UCM [43], WHURS19 [7], and NWPU [5]. Additionally, we evaluated the proposed method on 3 medical datasets, including BTMRI [25], CCBTM [11], and CHMNIST [16]. As with previous methods [17], we use accuracy and Harmonic Mean (HM) as evaluation metrics:

$$HM = \frac{2 \times Acc_{base} \times Acc_{novel}}{Acc_{base} + Acc_{novel}}, \quad (17)$$

where Acc_{base} denotes the accuracy of base categories, and Acc_{novel} denotes the accuracy of novel categories. We report the experimental results as an average of three independent runs for better statistical reliability. For the base-to-novel generalization scenario, we conduct experiments on all eight remote sensing datasets and three medical datasets. In addition, we evaluate the proposed method on cross-dataset generalization and domain generalization, where MLRSNet is used as the source dataset, and other remote sensing datasets are used as target datasets.

5.2. Implementation Details

For a fair comparison, we use similar training settings in MaPLe. All experiments are conducted using a few-shot training strategy with 16 shots, randomly sampled for each class. We use the pre-trained ViT-B/16 CLIP model for prompt tuning. The network is trained with a batch size of 4 and a learning rate of 0.0035 using the SGD optimizer on a single NVIDIA A100 GPU. We use the template ‘‘a photo of *category*’’ for the word embeddings. To ensure fair and reliable comparisons, we maintain consistent hyperparameters across all datasets, except for the domain-project layer, designed quaternion layer, and learnable contexts, which had their corresponding dimensions specified for each layer. The domain-project layer follows a Linear-ReLU-Linear architecture, while the designed quaternion layer adheres to the standard quaternion layer structure outlined in [26].

Table 1. Comparison with SOTA methods in base-to-novel generalization on 8 remote sensing recognition datasets. Our method consistently performs well over the SOTA approaches. We use **red** and **blue** to highlight the first and second best scores.

(a) Average over 8 datasets				(b) MLRSNet				(c) PatternNet			
	Base	Novel	HM		Base	Novel	HM		Base	Novel	HM
CLIP [28]	71.19	71.33	70.63	CLIP [28]	64.50	60.30	62.33	CLIP [28]	70.60	62.60	66.36
CoOp [48]	87.61	70.84	78.03	CoOp [48]	79.37	58.90	67.62	CoOp [48]	87.30	64.20	73.99
CoCoOp [47]	91.82	68.98	78.43	CoCoOp [47]	83.30	59.50	69.42	CoCoOp [47]	93.70	59.90	73.08
MaPLe [17]	93.12	71.71	80.42	MaPLe [17]	85.23	59.60	70.15	MaPLe [17]	95.30	57.90	72.03
Ours (ViTAE)	94.28	73.43	82.05	Ours (ViTAE)	88.96	57.10	69.56	Ours (ViTAE)	97.07	62.37	75.94
Ours (ViT)	94.08	75.06	83.50	Ours (ViT)	87.07	59.00	70.34	Ours (ViT)	95.80	66.20	78.30
(d) RSSCN7				(e) AID				(f) RSICD			
	Base	Novel	HM		Base	Novel	HM		Base	Novel	HM
CLIP [28]	66.70	95.30	78.48	CLIP [28]	73.50	70.40	71.92	CLIP [28]	71.50	60.20	65.37
CoOp [48]	84.80	89.13	86.91	CoOp [48]	87.63	70.37	78.06	CoOp [48]	88.43	60.20	71.63
CoCoOp [47]	90.97	90.00	90.48	CoCoOp [47]	92.63	65.73	76.89	CoCoOp [47]	92.37	58.80	71.86
MaPLe [17]		91.67	93.70	MaPLe [17]	92.73	74.57	82.66	MaPLe [17]	93.93	56.27	70.38
Ours (ViTAE)	91.53	94.53	93.01	Ours (ViTAE)	94.03	74.97	83.43	Ours (ViTAE)	94.57	65.20	77.19
Ours (ViT)	91.20	95.57	93.33	Ours (ViT)	94.50	75.77	84.10	Ours (ViT)	95.67	64.83	77.29
(g) UCM				(h) WHURS19				(i) NWPU			
	Base	Novel	HM		Base	Novel	HM		Base	Novel	HM
CLIP [28]	80.60	68.00	73.77	CLIP [28]	73.10	90.80	80.99	CLIP [28]	69.00	63.00	65.87
CoOp [48]	93.60	74.53	82.98	CoOp [48]	95.20	82.40	88.34	CoOp [48]	84.53	66.97	74.73
CoCoOp [47]	95.23	71.57	81.72	CoCoOp [47]	97.10	77.00	85.89	CoCoOp [47]	89.27	69.37	78.07
MaPLe [17]		97.70	70.90	MaPLe [17]	97.70	88.03	92.61	MaPLe [17]	90.70	72.70	80.71
Ours (ViTAE)	97.10	72.10	82.75	Ours (ViTAE)	99.40	89.90	94.41	Ours (ViTAE)	91.60	71.23	80.14
Ours (ViT)	97.90	73.30	83.83	Ours (ViT)	98.80	90.80	94.63	Ours (ViT)	91.70	75.03	82.53

5.3. Generalization from Base-to-Novel Classes

Prompt learning targets effective transfer of pre-trained VLMs to downstream tasks, facing the challenge of robust generalization across base and novel classes. To rigorously assess our method’s generalization from base to novel classes, we experiment on remote sensing and medical images. Utilizing the large-scale remote sensing foundation model [35], which provides ViTAE and ViT backbones, we conduct comprehensive experiments and compare results with advanced methods. Detailed experimental comparisons are presented in Table 1. With the ViT backbone, our method achieves state-of-the-art (SOTA) performance, with a 3.08% average HM improvement. Specifically, it enhances base accuracy from 93.12% to 94.08% and novel accuracy by 3.35%. On the RSICD dataset, our approach significantly boosts base accuracy from 93.93% to 95.67% and novel accuracy from 60.20% to 64.83%, yielding an HM improvement of nearly 6%. These findings highlight the advantage of domain-specific information for RSICD. However, on the challenging MLRSNet dataset, our method achieves only a modest 0.19% HM improvement, primarily due to limited knowledge of the domain-specific foundation model. Additionally, while deeper ViTAE performs better under some settings, it does not outperform ViT overall. These results suggest that deeper architectures may

not be optimal for extracting domain-specific information, as redundant parameters could hinder performance. Furthermore, we experiment with medical images using a specialized medical domain model [23] to further evaluate our prompt learning method’s effectiveness. Results and comparisons with advanced techniques are detailed in Table 2. Overall, our method achieves the highest accuracy across datasets. On average, it achieves a base category accuracy of 74.36% and enhances novel category accuracy from 44.40% to 44.74%. Regarding HM, our approach brings about nearly a 4% improvement. Interestingly, although our method excels in base category accuracy and HM across datasets, it does not rank highest for novel categories. This observation suggests that existing domain-specific medical image models, while significantly improving base category accuracy through domain-specific knowledge, may not fully exploit their potential to enhance generalization capabilities. This limitation could be a subject for future research.

5.4. Cross-Dataset Evaluation

To evaluate the dataset-generalization capabilities of our proposed method, we implement our approach on the MLRSNet and test the trained model directly on the remaining datasets. The experiments and comparison results with other advanced prompt learning methods are presented

Table 2. Comparison between our method and SOTA methods for base-to-novel generalization on medical image classification datasets. Our method performs well over the compared methods. We use **red** and **blue** to indicate the first and second best scores.

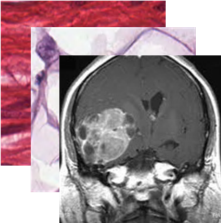
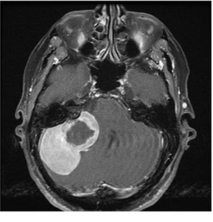
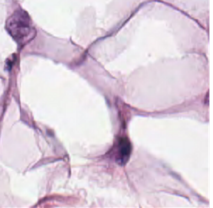
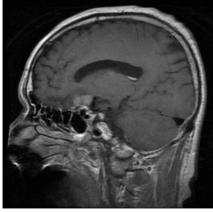
(a) Average over datasets				(b) All datasets				(c) BTMRI				(d) BTMRI dataset			
	Base	Novel	HM						Base	Novel	HM				
CLIP [28]	49.83	41.83	45.18		CLIP [28]	50.60	51.20	50.89		CLIP [28]	50.60	51.20	50.89		
CoOp [48]	51.59	43.77	46.81			CoOp [48]	48.93	53.30			51.02	CoOp [48]	48.93	53.30	51.02
CoCoOp [47]	64.45	43.16	49.45			CoCoOp [47]	52.37	52.80			52.58	CoCoOp [47]	52.37	52.80	52.58
MaPLe [17]	62.39	44.40	49.01			MaPLe [17]	53.67	61.60			57.36	MaPLe [17]	53.67	61.60	57.36
Ours	74.36	44.74	53.36			Ours	60.97	56.30			58.54	Ours	60.97	56.30	58.54
(e) CHMNIST				(f) CHMNIST dataset				(g) CCBTM				(h) CCBTM dataset			
	Base	Novel	HM						Base	Novel	HM				
CLIP [28]	31.60	27.40	29.35		CLIP [28]	67.30	46.90	55.28		CLIP [28]	67.30	46.90	55.28		
CoOp [48]	41.70	25.67	31.78			CoOp [48]	64.13	52.33			57.63	CoOp [48]	64.13	52.33	57.63
CoCoOp [47]	74.30	25.30	37.74			CoCoOp [47]	66.67	51.37			58.03	CoCoOp [47]	66.67	51.37	58.03
MaPLe [17]	74.03	25.10	37.49			MaPLe [17]	59.47	46.50			52.19	MaPLe [17]	59.47	46.50	52.19
Ours	87.80	26.60	40.83			Ours	74.30	51.33			60.72	Ours	74.30	51.33	60.72

Table 3. Comparisons with SOTA methods for cross-dataset generalization with the MLRSNet dataset as the source domain and remaining remote sensing datasets as the target domains. Our method achieves better performance than the compared methods. We use **red** and **blue** to highlight the first and second best scores.

	Source			Target					
	MLRSNet	PatternNet	RSSCN7	AID	RSICD	UCM	WHURS19	NWPU	Average
CoOp [48]	72.53	66.97	69.03	67.30	63.50	77.57	85.47	70.43	71.60
CoCoOp [47]	71.70	65.67	68.80	66.63	62.57	76.40	85.33	70.30	70.92
MaPLe [17]	76.83	68.53	71.43	65.13	59.53	79.90	85.23	72.80	72.42
Ours	78.73	68.17	70.60	66.70	62.27	79.93	91.07	73.13	73.83

Table 4. Comparisons with SOTA methods for single-source multi-target domain generalization with the MLRSNet dataset as the source domain and remaining datasets as the target domains. Our method achieves better performance than the compared methods. We use **red** and **blue** to highlight the first and second best scores.

	Source			Target					
	MLRSNet	PatternNetv2	RSSCN7v2	AIDv2	RSICDv2	UCMv2	WHURS19v2	NWPUv2	Average
CoOp [48]	72.53	66.97	69.07	67.13	64.27	77.40	85.20	71.17	71.72
CoCoOp [47]	71.70	65.57	69.37	67.13	62.73	75.70	84.83	70.97	71.00
MaPLe [17]	76.83	68.03	72.50	64.90	59.73	78.93	83.07	73.17	72.15
Ours	78.73	67.63	71.33	66.87	62.33	78.33	89.90	73.67	73.60

in Table 3. Our proposed method achieves superior performance on the MLRSNet with a 1.90% performance improvement over the state-of-the-art. Notably, our method achieves the highest performance, reaching a 91.07% accuracy, on the WHURS19 dataset, suggesting a significant overlap in domain-specific information between MLRSNet and WHURS19. While our method does not attain top performance in every dataset, it ranks first in terms of average performance with a 1.41% accuracy improvement. These results affirm the enhanced cross-dataset generalization ability of our method.

5.5. Domain Generalization

To further validate the generalization capabilities, we conduct evaluations under the domain generalization setting like the previous methods. The experiments and comparison results with advanced algorithms are shown in Table 4. While the proposed method does not achieve top performance on each dataset, it excels with an average accuracy of 73.60%, corresponding to a 1.35% improvement over MaPLe. Notably, the best performance is attained on the WHURS19V2 dataset, where our method exhibits a remarkable 4.70% performance boost. These re-

Table 5. The utilization of quaternion network (QN). Our method performs well over the compared methods due to successful cross-modal mining with QN.

Methods	Base	Novel	HM
Baseline	93.93	56.27	70.38
Ours w/o QN	94.17	63.53	75.87
Ours	95.67	64.83	77.29

Table 6. The prompting of vision and language branches. PL: prompting language branch; PV: prompting vision branch. Our method performs better due to simultaneous PL and PV.

Methods	Base	Novel	HM
Baseline	93.93	56.27	70.38
PL	95.23	63.73	76.36
PV	95.60	64.23	76.84
Ours (PL+PV)	95.67	64.83	77.29

sults demonstrate that our proposed domain prompt learning effectively enhances the generalization and robustness of vision-language models, offering promising capabilities for domain generalization across datasets.

5.6. Ablation Studies

Quaternion Network. In order to examine the impact of different components of the proposed domain prompt learning, we carry out a series of ablation experiments. Initially, we study the effect of the quaternion neural network and its impact. Table 5 shows that even without the quaternion network, our method still significantly improves the performance. This suggests that incorporating the domain-specific knowledge to transfer the VLMs from generalized to specialized is an efficient approach, and the quaternion network efficiently models the orthogonal cross-modal relationships, leading to further improved performance.

Prompting Vision and Language Branches. Table 6 presents the experimental results of studying the impact of prompting vision and language branches. The results show that prompting either the language or vision branch alone leads to better performance due to improved domain-specific vision-language contrastive learning. By prompting both the vision and language branches complementarily, our proposed method successfully propagates domain-specific information into contrastive learning and significantly enhances recognition performance.

Prompting Depth. In domain prompt learning, the depth of prompting plays a crucial role, and we conduct experiments to study its impact. The results in Table 7 show that performance improves as the depth increases until it reaches its highest point at a depth of 9. However, beyond that point, a further increase in depth leads to a decline in performance due to redundancy. Therefore, a depth of 9 is the appropriate choice for optimal performance.

Adding Noise. In order to mitigate potential overfitting

Table 7. Ablation studies of prompting depth. Our method performs better due to appropriate prompting depth.

Depth	Base	Novel	HM
3	94.73	60.7	73.99
6	95.63	61.83	75.10
9	95.67	64.83	77.29
12	95.57	64.67	77.14

Table 8. Ablation studies on the addition of noise. VB denotes the vision branch. LB denotes the language branch. Our method performs better due to correctly adding noise.

Depth	Base	Novel	HM
Baseline	93.93	56.27	70.38
w/o Noise	95.43	63.03	75.92
VB + Noise	95.38	62.95	75.84
Ours (LB+ Noise)	95.67	64.83	77.29

problems, we devise an approach that introduces noise into the language branch. We assess the impact of this technique, and the results are presented in Table 8. Interestingly, our method still shows some performance improvements even in the absence of added noise. However, we have observed a decrease in accuracy when we add noise in the vision branch. We attribute this decline to the possibility of noise disrupting well-matched vision-language relationships. Our approach achieves the best performance, indicating that adding noise in the language branch is a beneficial strategy for addressing overfitting.

6. Conclusion

Most advanced prompt learning algorithms, tailored for natural images, often face difficulties in adapting to specific domains. To tackle these, we leverage domain-specific insights from foundation models and employ quaternion networks to extend the robust recognition abilities of VLMs into specialized domains. By examining cross-modal connections between domain-specific vision features and contextual embeddings, the quaternion network smoothly incorporates generalized contextual embeddings into specialized realms. Furthermore, our method integrates domain-specific knowledge into the vision branch, facilitating effective vision-language contrastive learning. Extensive experiments on domain-specific datasets showcase the superiority of our approach with compelling results.

Acknowledgements

This work was supported by the National Natural Science Foundation of China (Grant No. 62106116, 62322113, and 62376156), China Meteorological Administration under Grant QBZ202316, Natural Science Foundation of Ningbo of China (No. 2023J027), as well as by the High Performance Computing Centers at Eastern Institute of Technology, Ningbo, and Ningbo Institute of Digital Twin.

References

- [1] Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei. Beit: Bert pre-training of image transformers. *arXiv:2106.08254*, 2021. **2**
- [2] Qinglong Cao, Yuntian Chen, Chao Ma, and Xiaokang Yang. Few-shot rotation-invariant aerial image semantic segmentation. *IEEE Transactions on Geoscience and Remote Sensing*, 62:1–13, 2023. **1**
- [3] Qinglong Cao, Zhengqin Xu, Yuntian Chen, Chao Ma, and Xiaokang Yang. Domain-controlled prompt learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 936–944, 2024. **1**
- [4] Xi Chen, Xiao Wang, Soravit Changpinyo, AJ Piergiovanni, Piotr Padlewski, Daniel Salz, Sebastian Goodman, Adam Grycner, Basil Mustafa, Lucas Beyer, et al. Pali: A jointly-scaled multilingual language-image model. *arXiv:2209.06794*, 2022. **2**
- [5] Gong Cheng, Junwei Han, and Xiaoqiang Lu. Remote sensing image scene classification: Benchmark and state of the art. *Proceedings of the IEEE*, 105(10):1865–1883, 2017. **5**
- [6] Quan Cui, Boyan Zhou, Yu Guo, Weidong Yin, Hao Wu, Osamu Yoshie, and Yubo Chen. Contrastive vision-language pre-training with limited resources. In *Proceedings of European Conference on Computer Vision*, 2022. **2**
- [7] Dengxin Dai and Wen Yang. Satellite image classification via two-layer sparse coding with biased image representation. *IEEE Transactions on Geoscience and Remote Sensing*, 8(1):173–176, 2011. **5**
- [8] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv:2010.11929*, 2020. **3**
- [9] Zi-Yi Dou, Aishwarya Kamath, Zhe Gan, Pengchuan Zhang, Jianfeng Wang, Linjie Li, Zicheng Liu, Ce Liu, Yann LeCun, Nanyun Peng, et al. Coarse-to-fine vision-language pre-training with fusion in the backbone. *Advances in Neural Information Processing Systems*, 2022. **2**
- [10] Yu Du, Fangyun Wei, Ziheng Zhang, Miaoqing Shi, Yue Gao, and Guoqi Li. Learning to prompt for open-vocabulary object detection with vision-language model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. **2**
- [11] Seyed Mohammad Hossein Hashemi. Crystal clean: Brain tumors mri dataset, 2023. **5**
- [12] Kaifeng He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016. **1**
- [13] Kaifeng He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. **4**
- [14] Runhui Huang, Yanxin Long, Jianhua Han, Hang Xu, Xiwen Liang, Chunjing Xu, and Xiaodan Liang. Nlip: Noise-robust language-image pre-training. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2023. **2**
- [15] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International Conference on Machine Learning*, 2021. **1, 2**
- [16] JN Kather, CA Weis, F Bianconi, SM Melchers, LR Schad, T Gaiser, A Marx, and Zollner F. Multi-class texture analysis in colorectal cancer histology. *Scientific Reports*, 2016. **5**
- [17] Muhammad Uzair Khattak, Hanoona Rasheed, Muhammad Maaz, Salman Khan, and Fahad Shahbaz Khan. Maple: Multi-modal prompt learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023. **1, 2, 3, 5, 6, 7**
- [18] Byungsoo Ko and Geonmo Gu. Large-scale bilingual language-image contrastive learning. *arXiv preprint arXiv:2203.14463*, 2022. **2**
- [19] Yangguang Li, Feng Liang, Lichen Zhao, Yufeng Cui, Wanli Ouyang, Jing Shao, Fengwei Yu, and Junjie Yan. Supervision exists everywhere: A data efficient contrastive language-image pre-training paradigm. *arXiv:2110.05208*, 2021. **2**
- [20] Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM Computing Surveys*, 55(9), 2023. **2**
- [21] Xiaoqiang Lu, Binqiang Wang, Xiangtao Zheng, and Xuelong Li. Exploring models and data for remote sensing image caption generation. *IEEE Transactions on Geoscience and Remote Sensing*, 56(4):2183–2195, 2017. **5**
- [22] Timo Lüddecke and Alexander Ecker. Image segmentation using text and image prompts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7086–7096, 2022. **1**
- [23] Jun Ma and Bo Wang. Segment anything in medical images. *arXiv:2304.12306*, 2023. **2, 3, 4, 6**
- [24] Norman Mu, Alexander Kirillov, David Wagner, and Saining Xie. Slip: Self-supervision meets language-image pre-training. In *Proceedings of European Conference on Computer Vision*, 2022. **2**
- [25] Msoud Nickparvar. Brain tumor mri dataset, 2021. **5**
- [26] Titouan Parcollet, Mirco Ravanelli, Mohamed Morchid, Georges Linarès, Chiheb Trabelsi, Renato De Mori, and Yoshua Bengio. Quaternion recurrent neural networks. *arXiv:1806.04418*, 2018. **2, 5**
- [27] Xiaoman Qi, Panpan Zhu, Yuebin Wang, Liqiang Zhang, Junhuan Peng, Mengfan Wu, Jialong Chen, Xudong Zhao, Ning Zang, and P Takis Mathiopoulos. Mlrsnet: A multi-label high spatial resolution remote sensing dataset for semantic scene understanding. *ISPRS Journal of Photogrammetry and Remote Sensing*, 169:337–350, 2020. **5**
- [28] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language super-

- vision. In *International Conference on Machine Learning*, 2021. 1, 2, 6, 7
- [29] Yongming Rao, Wenliang Zhao, Guangyi Chen, Yansong Tang, Zheng Zhu, Guan Huang, Jie Zhou, and Jiwen Lu. Densclip: Language-guided dense prediction with context-aware prompting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 1
- [30] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv:1409.1556*, 2014. 1
- [31] Amanpreet Singh, Ronghang Hu, Vedanuj Goswami, Guillaume Couairon, Wojciech Galuba, Marcus Rohrbach, and Douwe Kiela. Flava: A foundational language and vision alignment model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 2
- [32] Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. Vi-bert: Pre-training of generic visual-linguistic representations. *arXiv:1908.08530*, 2019. 1
- [33] Xian Sun, Peijin Wang, Wanxuan Lu, Zicong Zhu, Xiaonan Lu, Qibin He, Junxi Li, Xuee Rong, Zhujun Yang, Hao Chang, et al. Ringmo: A remote sensing foundation model with masked image modeling. *IEEE Transactions on Geoscience and Remote Sensing*, 2022. 2, 3
- [34] Mingkan Tang, Zhanyu Wang, Zhenhua Liu, Fengyun Rao, Dian Li, and Xiu Li. Clip4caption: Clip for video caption. In *Proceedings of the 29th ACM International Conference on Multimedia*, 2021. 2
- [35] Di Wang, Qiming Zhang, Yufei Xu, Jing Zhang, Bo Du, Dacheng Tao, and Liangpei Zhang. Advancing plain vision transformer towards remote sensing foundation model. *IEEE Transactions on Geoscience and Remote Sensing*, 61:1–15, 2022. 2, 3, 6
- [36] Yimu Wang, Bo Xue, Quan Cheng, Yuhui Chen, and Lijun Zhang. Deep unified cross-modality hashing by pairwise data alignment. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence*, 2021. 1
- [37] Yimu Wang, Xiangru Jian, and Bo Xue. Balance act: Mitigating hubness in cross-modal retrieval with query and gallery banks. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 2023. 1
- [38] Zifeng Wang, Zizhao Zhang, Chen-Yu Lee, Han Zhang, Ruoxi Sun, Xiaoqi Ren, Guolong Su, Vincent Perot, Jennifer Dy, and Tomas Pfister. Learning to prompt for continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 1
- [39] Bichen Wu, Ruizhe Cheng, Peizhao Zhang, Peter Vajda, and Joseph E Gonzalez. Data efficient language-supervised zero-shot recognition with optimal transport distillation. *arXiv:2112.09445*, 2021. 2
- [40] Gui-Song Xia, Jingwen Hu, Fan Hu, Baoguang Shi, Xiang Bai, Yanfei Zhong, Liangpei Zhang, and Xiaoqiang Lu. Aid: A benchmark data set for performance evaluation of aerial scene classification. *IEEE Transactions on Geoscience and Remote Sensing*, 55:3965–3981, 2017. 5
- [41] Mengde Xu, Zheng Zhang, Fangyun Wei, Yutong Lin, Yue Cao, Han Hu, and Xiang Bai. A simple baseline for open-vocabulary semantic segmentation with pre-trained vision-language model. In *Proceedings of European Conference on Computer Vision*, 2022. 2
- [42] Yufei Xu, Qiming Zhang, Jing Zhang, and Dacheng Tao. Vi-tae: Vision transformer advanced by exploring intrinsic inductive bias. *Advances in Neural Information Processing Systems*, 2021. 3
- [43] Yi Yang and Shawn Newsam. Bag-of-visual-words and spatial extensions for land-use classification. In *Proceedings of the 18th SIGSPATIAL International Conference on Advances in Geographic Information Systems*, 2010. 5
- [44] Xiwen Yao, Qinglong Cao, Xiaoxu Feng, Gong Cheng, and Junwei Han. Learning to assess image quality like an observer. *IEEE Transactions on Neural Networks and Learning Systems*, 2022. 1
- [45] Jiahui Yu, Zirui Wang, Vijay Vasudevan, Legg Yeung, Mojtaba Seyedhosseini, and Yonghui Wu. Coca: Contrastive captioners are image-text foundation models. *arXiv:2205.01917*, 2022. 2
- [46] Wangbo Zhao, Jiasheng Tang, Yizeng Han, Yibing Song, Kai Wang, Gao Huang, Fan Wang, and Yang You. Dynamic tuning towards parameter and inference efficiency for vit adaptation. *arXiv:2403.11808*, 2024. 1
- [47] Kaiyang Zhou, Jingkan Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 1, 2, 3, 4, 6, 7
- [48] Kaiyang Zhou, Jingkan Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *International Journal of Computer Vision*, 130(9):2337–2348, 2022. 1, 2, 6, 7
- [49] Weixun Zhou, Shawn Newsam, Congmin Li, and Zhenfeng Shao. Patternnet: A benchmark dataset for performance evaluation of remote sensing image retrieval. *ISPRS Journal of Photogrammetry and Remote Sensing*, 145:197–209, 2018. 5
- [50] Beier Zhu, Yulei Niu, Yucheng Han, Yue Wu, and Hanwang Zhang. Prompt-aligned gradient for prompt tuning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023. 3
- [51] Qin Zou, Lihao Ni, Tong Zhang, and Qian Wang. Deep learning based feature selection for remote sensing scene classification. *IEEE Geoscience and Remote Sensing Letters*, 12(11):2321–2325, 2015. 5