

# Towards Generative Class Prompt Learning for Fine-grained Visual Recognition

Soumitri Chattopadhyay<sup>1</sup>

soumitri@cs.unc.edu

Sanket Biswas<sup>2</sup>

sbiswas@cvc.uab.es

Emanuele Vivoli<sup>2, 3</sup>

evivoli@cvc.uab.es

Josep Lladós<sup>2</sup>

josep@cvc.uab.es

<sup>1</sup> Department of Computer Science,  
University of North Carolina at Chapel  
Hill, USA

<sup>2</sup> Computer Vision Center & Computer  
Science Department,  
Universitat Autònoma de Barcelona,  
Spain

<sup>3</sup> MICC, University of Florence, Italy

---

## Abstract

Although foundational vision-language models (VLMs) have proven to be very successful for various semantic discrimination tasks, they still struggle to perform faithfully for fine-grained categorization. Moreover, foundational models trained on one domain do not generalize well on a different domain without fine-tuning. We attribute these to the limitations of the VLM's semantic representations and attempt to improve their fine-grained visual awareness using generative modeling. Specifically, we propose two novel methods: *Generative Class Prompt Learning* (GCPL) and *Contrastive Multi-class Prompt Learning* (CoMPLe). Utilizing text-to-image diffusion models, GCPL significantly improves the visio-linguistic synergy in class embeddings by conditioning on few-shot exemplars with learnable class prompts. CoMPLe builds on this foundation by introducing a contrastive learning component that encourages inter-class separation during the generative optimization process. Our empirical results demonstrate that such a generative class prompt learning approach substantially outperform existing methods, offering a better alternative to few shot image recognition challenges. The source code will be made available at: <https://github.com/soumitri2001/GCPL>.

## 1 Introduction

Foundational vision-language models (VLMs) [22, 45, 56] have emerged as powerful open-world learners due to their large-scale web-based pre-training, covering a broad range of multi-modal concepts. Models like CLIP [45] and ALIGN [24], trained with a contrastive learning objective to align image-text pairs, show promising generalization to downstream tasks [55], often in a zero-shot manner [50, 72]. CLIP-based zero-shot classification, for example, uses a handcrafted prompt ("A photo of a [CLASS]") to represent a class, predicting the class by computing the similarity between text and image embeddings [45, 70].

However, the performance of zero-shot CLIP baseline is often suboptimal, especially for fine-grained image recognition tasks [51, 41, 53]. This limitation is largely due to the quality of handcrafted prompts, which lack discriminative task-specific knowledge [77, 59, 70].

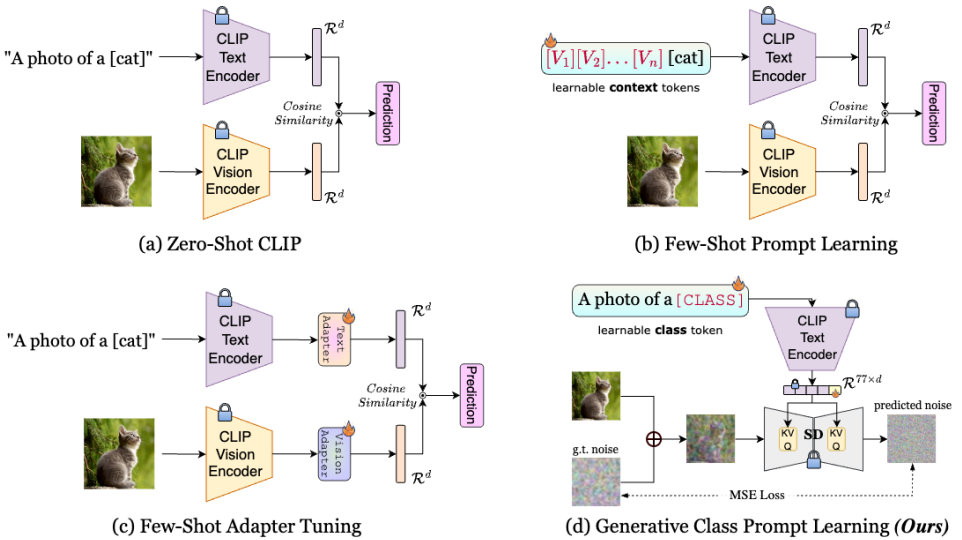


Figure 1: Overview of our approach compared to existing VLM adaptation methods. (a) Zero-shot inference with CLIP; (b) Contextual prompt token learning; (c) Adapter-based tuning with handcrafted prompts on frozen CLIP representations; (d) **Our setup (GCPL)**: generatively learning the [CLASS] token by prompting a frozen text-to-image LDM [46].

To mitigate this issue, techniques such as *adapter learning* [43, 58] and *prompt learning* [27, 54, 70, 73] have emerged in the VLM space. While the former introduces extra layers on top of frozen vision and language representations for alignment [43], prompt learning, inspired by [0], transforms the previously handcrafted prompts (e.g., "A photo of a") into soft, learnable prompts, with the [CLASS] token fixed. This enables the prompts to capture task-specific contextual information, optimized over few-shot support sets. Subsequent work [9, 27, 54, 73] has further refined this approach, enhancing contextual learning and improving visiolinguistic synergy [9, 27].

Despite improvements, several issues persist in CLIP’s transferability to specialized downstream tasks, particularly in fine-grained classification. Firstly, fine-grained category names are often highly dataset-specific or domain-specific, lacking the visual semantic knowledge that could assist in accurate recognition. For example, consider the FGVCaircraft [58] dataset that has aircraft model names like Tu-154, DR-400, MD-80 etc., or StanfordCars [51] with fine-grained vehicular models such as "2012 Acura TL Sedan," "2011 Audi S6 Sedan", etc. Such customized categories *do not* contain sufficient relevant visual cues and would naturally lead to spurious representations when fed through CLIP. Although multimodal prompt learning [27, 28] seeks to enhance visio-linguistic synergy, it depends on frozen CLIP embeddings, causing their suboptimality to transfer to learnable tokens via self-attention [46, 57]. Secondly, CLIP’s training on web-scraped data focuses on natural visual content, limiting its ability to generalize to specialized domains such as medical imaging [26, 59], affecting multimodal prompting [27] due to a lack of domain-specific knowledge in both visual and textual branches. Thirdly, certain visual concepts are *hard to describe by text* (e.g., abstract fractals [25]), making CLIP’s class embeddings less effective. Thus, we can single out the core underlying problem: suboptimality of raw CLIP representations, which often lack fine-grained visual semantic awareness.

In this work, we challenge the existing paradigm of learning contextual prompts [27, 62, 70] and instead advocate for learning stronger class embeddings with enhanced visual semantic knowledge. In our attempt to tackle these aforementioned limitations head-on, we take a radically different route from conventional approaches [27, 45, 62, 70]; leveraging text-to-image diffusion models [46] for few-shot class prompt learning. Two primary factors motivate our line of exploration: the immense potential unearthed from zero-shot generative classifiers [23, 63] and their emergent properties such as strong shape bias and human-level out-of-distribution accuracy [23]; and the explicit *parameterization of textual conditions* for controllable generation [19, 49]. Text-to-image latent diffusion models such as Stable Diffusion (SD) [46, 57] rely on a text conditional prompt that interacts with a U-Net denoiser via cross-attention [46] during the denoising process [21], creating strong visio-linguistic synergy. In our few-shot learning approach with an exemplar support set and a predefined prompt (e.g., "A photo of [CLASS]"), we make the [CLASS] token *learnable* while keeping the contextual token fixed (unlike traditional prompt tuning [59]). This learnable prompt embedding can be utilized as the text control for the diffusion model [46]. The class embedding is *learned* via optimisation of the mean-squared error loss between estimated and the added noises [21, 46] on the exemplar images of that particular class. This approach, coined as Generative Class Prompt Learning (**GCPL**), is derived from the so-called textual inversion paradigm [19], with however, a *different goal*: to learn rich pseudo-text-like representations that are enhanced with fine-grained visual semantics, to aid fine-grained discriminative tasks. Empirical results (Section 5) highlight the high robustness and effective fine-grained feature learning achieved via GCPL as compared to existing works [33, 45, 68, 69].

Therefore, in this way we present a very simple yet effective means to adapt a frozen SD model [46] to learn class prompts for few-shot classification. However, a key weakness is that our method learns class embeddings directly from exemplar images without learning any inter-class discrimination, potentially leading to poor separation of latent class representations and reduced performance at a fine-grained level. To address this, we additionally propose **Contrastive Multi-class Prompt Learning (CoMPLe)**, which optimizes all class prompts simultaneously using a contrastive learning approach. For each noisy image latent, the denoised output should be closest to its ground-truth noise and distinct from those of other classes. This involves optimizing two loss terms: minimizing for positive (same class) noise pairs and maximizing for negative (different class) pairs. Empirical results show that CoMPLe enhances GCPL performance, especially on out-of-domain tasks like medical imaging datasets [26, 39], highlighting the benefits of inter-class awareness.

Summing up, our main contributions are: *(i)* We identify potential limitations in CLIP’s visio-linguistic knowledge bed with practical cases of failure; *(ii)* we propose a generative class prompt learning baseline, leveraging pre-trained diffusion models to mitigate those limitations and advocating for learning stronger vision-induced textual representations with inter-class discriminative knowledge; *(iii)* via thorough experimentations and comparisons with prior approaches, we demonstrate the potential and robustness of our proposal.

Being a first-of-its-kind attempt to introduce generation-aided prompting for few-shot V/L tasks, we hope our work would serve as a baseline to stir up further interest in advancing the study of generative classifiers and parameterization in diffusion models in general.

## 2 Related Work

**Few-shot Image Classification:** Few-shot learning (FSL) focuses on training models with a limited support set, typically formulated as an  $N$ -way  $K$ -shot classification task, where  $N$

is the number of classes and  $K$  is the number of examples per class [10, 82, 69]. Traditional FSL methods use meta-learning frameworks [10, 12, 82, 66], involving meta-training on base classes and meta-testing on novel classes through simulated episodes [12]. These methods are categorized into metric-based [62, 64, 68], alignment-based [60, 61, 62, 65], and optimization-based approaches [17, 18, 40]. More recently, large-scale multimodal models such as CLIP [45] have enabled few-shot learning across modalities by leveraging broad world knowledge [45, 64, 68]. Unlike these discriminative methods, our work introduces a *generative modeling-based* framework for few-shot classification.

**Foundational Vision-Language Models (VLMs):** Using language as a supervisory signal with visual content has led to the development of robust multi-modal representations in large-scale VLMs like CLIP [45], ALIGN [24], and Flamingo [0]. These models, trained on vast amounts of internet data (e.g., 400M image-text pairs for CLIP), employ contrastive learning [10, 42] to capture extensive open-world visio-linguistic knowledge [66]. This makes them highly effective for zero-shot tasks such as classification [61, 65], cross-modal retrieval [60, 60], anomaly detection [64, 70], and segmentation [8, 15, 72]. However, despite their strong representations, models like CLIP struggle with fine-grained categorization [61, 70] and domain-specific tasks [26, 69]. To address this, methods such as linear probing [22, 45, 68], prompting [27, 28, 64, 69, 70, 73], and adapters [43, 68] have been introduced to adapt VLMs to specialized tasks using few-shot support sets, enhancing their performance [65, 63]. Inspired by these adaptation techniques, we propose a *generation-aided* class prompt learning method that combines CLIP [45] and Stable Diffusion [46] to learn strong few-shot exemplar representations.

**Prompt Learning in VLMs:** Originating from NLP [0], prompt learning is a parameter-efficient method to adapt foundational models like CLIP [45] for downstream tasks in few-shot settings [63]. The standard handcrafted prompt “a photo of a [CLASS]” used in zero-shot CLIP is suboptimal [64, 70] as it lacks task-specific information. To address this, prompt learning replaces fixed context words with learnable soft prompts that encode task-specific knowledge. Methods like CoOp [70] and CoCoOp [69] optimize continuous context vectors in CLIP’s text branch, while MaPLe [72] applies prompt learning to both text and vision branches for better visio-linguistic synergy. Further advances include techniques such as optimal transport [9], self-regularization [28], and consistency regularization [48]. All these methods keep the [CLASS] embedding frozen while learning context tokens. Our approach differs by optimizing the [CLASS] embedding while using handcrafted context words, learning class prototypes via prompting with a frozen text-to-image diffusion model [46] (details in Section 3).

**Diffusion models for discriminative tasks:** With recent advances, diffusion models [21, 46] have become a standard for visual content generation [19, 46, 49], relying on an iterative denoising process [21] that enables *controllable* image generation through proper conditioning [9, 20, 46, 57]. This same process can be adapted for discriminative tasks, such as classification [8, 13, 63], image retrieval [29], and segmentation [8, 14], by parameterizing and learning the conditions via the diffusion model’s reconstruction objective [46]. Many of these approaches use Stable Diffusion, similar to CLIP [45], to leverage large-scale, open-world knowledge for zero-shot inference. Zero-shot adaptations of SD for classification [13, 23, 63] benefit from the unique properties of generative classifiers [23]. However, like zero-shot CLIP, naive use of SD faces similar challenges, which we address by introducing a few-shot generative classifier that learns class-specific prototypes as conditional inputs to SD [19, 46].




Figure 2: **Contrastive multi-class prompt learning (CoMPLE)** framework. Our proposed CoMPLE learns class prompts by optimizing the LDM loss for the trainable class token, minimizing noise reconstruction for ground truth noise while maximizing it for other class noises. Red arrows show "maximize," and blue arrows show "minimize." The diffusion classifier [63] uses our few-shot learned [CLASS] embeddings for inference.

### 3 Proposed Methodology

**Overview:** We propose a class prototype learning framework using a pre-trained text-to-image diffusion model, using few-shot exemplar sets per class. Specifically, we derive from [19] and parameterize a handcrafted prompt (e.g. “A photo of a [CLASS]”), where the [CLASS] token is made *learnable* keeping the rest of the prompt fixed. This prompt is then fed through CLIP’s text encoder [45] to obtain a continuous text embedding, which is then used to condition the denoising process of a frozen Stable Diffusion model [46] that aims to reconstruct the few-shot images provided for that particular class (Section 3.2). This is our vanilla few-shot learning approach that learns each [CLASS] embedding independently; we further extend this to **CoMPLE**, a **contrastive multi-class prompt learning** setup where all the class prompts are optimized simultaneously using a contrastive reconstruction loss objective that minimizes the noise prediction error for a sample’s true noise while maximizing the error for other (i.e. negative) samples (Section 3.3). For classification, we use Diffusion Classifier [63], replacing CLIP class embeddings (zero-shot) with our *few-shot learned embeddings*.

#### 3.1 Revisiting Text-conditional Latent Diffusion Models

**Preliminaries of LDM.** The working principle of diffusion models [21] is based on two complementary iterative processes of adding noise (forward) and removing noise (reverse) to a given clean image. Both processes are spaced over  $T$  timesteps. Latent diffusion models (LDM) [46] shifted the diffusion process i.e. noise addition and removal processes from pixel space to the latent space, projected by a pre-trained VAE encoder  $\mathcal{E}(\cdot)$  and eventually decoded back to the image space by a pre-trained VAE decoder  $\mathcal{D}(\cdot)$ , such that  $\mathcal{D}(\mathcal{E}(x)) \approx x$ . Formally, a clean image latent  $x_0$  undergoes iterative Gaussian noise addition across  $T$

timesteps, such that at a given timestep  $t$  the noisy image  $x_t = \sqrt{\bar{\alpha}_t}x_0 + \sqrt{1 - \bar{\alpha}_t}\varepsilon$ , where  $\varepsilon \sim \mathcal{N}(0, I)$  is the noise added,  $\{\alpha_t\}_1^T$  refers to a noise schedule where  $\bar{\alpha}_t = \prod_{i=1}^t \alpha_i$  [49]. A denoising U-Net parameterizes the noise removal process [44]  $\Phi_\theta$ , taking the noisy latent as input and aims to estimate the added noise  $\varepsilon \approx \Phi_\theta(x_t, t)$ , trained using a reconstruction (i.e. MSE loss). Having learned how to effectively denoise samples, the denoiser network  $\Phi_\theta$  can now start from pure Gaussian noise and employing  $T$  timesteps, it can produce cleaner images  $\{x_{t-1}, x_{t-2} \dots x_0\}$ , mimicking in the original distribution of  $x_0$  [20].

**Text-conditioned LDM.** The unconditional denoising diffusion model [20] can be made “conditional” by fusing auxiliary control signals with the denoiser U-Net  $\Phi_\theta$ . The control signals may be of various modalities, e.g. textual prompts [11, 19, 46, 67], semantic maps [49, 67], 2D pose [67] etc. We take particular focus in text-conditioned LDM, where a pre-trained text encoder like CLIP [45] is used to obtain tokenized textual embeddings from raw text prompts, which interact with the denoiser network through cross-attention (shown in Figure 2). This brings out the visio-linguistic synergical modeling in text-conditioned LDMs (in particular, we use Stable Diffusion [46]), which we exploit in our work.

## 3.2 Generative Class Prompt Learning

Given a few-shot support set and a corresponding prompt containing the class name, we aim to learn a visually enriched representation for the [CLASS] token present in our prompt embedding. Taking a leaf out of [19], we first tokenize the prompt to unique word embeddings [45], and make the class-corresponding token *learnable*. To do so, we introduce a placeholder token,  $p_c$ , into the vocabulary of CLIP [45], and make the parameterization trainable. Now we leverage the frozen SD model [46] to perform conditional denoising over the given set of images, with the textual condition containing the learnable token  $p_c$ . The optimization of SD follows the standard  $l_2$  loss between the predicted and ground truth noise added to the respective latents. The optimization follows the following equation:

$$\mathcal{L}_{GCPL} = \mathbb{E}_{x \sim \mathcal{E}(x), p_c, \varepsilon \sim \mathcal{N}(0,1), t} \left[ \|\varepsilon_c - \varepsilon_c^\theta(x_t, t, c_\theta(p_c))\|_2^2 \right] \quad (1)$$

where  $\varepsilon_c$  is the ground truth noise added to the image latent of class ‘ $c$ ’,  $\varepsilon_c^\theta$  being the noise predicted by the denoiser parameterized by  $\theta$ , and  $c_\theta$  maps the learnable class prompt  $p_c$  to a conditioning embedding. Essentially,  $p_c$  is learned via:

$$p_c^* = \arg \min_{p_c} \mathbb{E}_{x \sim \mathcal{E}(x), p_c, \varepsilon \sim \mathcal{N}(0,1), t} \left[ \|\varepsilon_c - \varepsilon_c^\theta(x_t, t, c_\theta(p_c))\|_2^2 \right] \quad (2)$$

This generative optimization induces a synergy between the class prompts and the visual exemplars fed into the LDM, resulting in visually enhanced class prototypes (refer to Figure 2).

## 3.3 CoMPLE: Contrastive Multi-Class Prompt Learning

As discussed previously, mere learning of class prompts, without any discriminative knowledge of other classes, can potentially lead to suboptimal class embeddings, given our primary goal is downstream discriminative tasks. For this, we extend GCPL to a multi-class setting where all class prompts  $\{p_{c_i}\}$  are initialised together and simultaneously optimized via the above-discussed objective. We identify a simple contrastive constraint: in addition to aligning the predicted noises with our ground truth added ones, we also *enforce divergence of predicted noises against the ground truths noises of other classes*. As shown in Figure 2, we apply two loss terms: a minimization objective over true noise pairs (i.e. belonging to same class), and a maximization objective over negative pairs of noises. Together, the loss function across a batch of  $B$  samples containing  $c_{\{1..B\}}$  classes can be written as:



$$\begin{aligned} \mathcal{L}_{\text{CoMPLe}} = & \frac{1}{B} \sum_{i=j}^B \mathbb{E}_{x \sim \mathcal{E}(x), p_{c_j}, \varepsilon \sim \mathcal{N}(0,1), t} \left[ \|\varepsilon_{c_i} - \varepsilon_{c_j}^\theta(x_t^j, t, c_\theta(p_{c_j}))\|_2^2 \right] \\ & - \lambda \cdot \frac{1}{B(B-1)} \sum_{i \neq j}^B \mathbb{E}_{x \sim \mathcal{E}(x), p_{c_j}, \varepsilon \sim \mathcal{N}(0,1), t} \left[ \|\varepsilon_{c_i} - \varepsilon_{c_j}^\theta(x_t^j, t, c_\theta(p_{c_j}))\|_2^2 \right] \end{aligned} \quad (3)$$

where  $\lambda$  is a weighting hyperparameter that balances the two contrasting MSE objectives. Learning of  $\{p_{c_j}\}$  follows the optimization similar to Equation 2. This allows us to imbibe inter-class discriminative knowledge within our generatively trained class prompting setup, boosting fine-grained representation learning.

### 3.4 Few-shot Diffusion Classifier

We adopt the recently proposed diffusion classifier [63] and extend it to our few-shot inference setup with minimal changes, which ensures fair comparisons and robustness.

**Preliminaries.** As discussed in preceding sections, an LDM [46] parameterizes the reverse process  $p_\theta(x_{t-1} | x_t, c)$ , i.e. denoising  $x_t$  to  $x_{t-1}$ , where  $c$  is the conditioning signal. Now, using  $x_t = \sqrt{\bar{\alpha}_t}x_0 + \sqrt{1 - \bar{\alpha}_t}\varepsilon$  (refer to Section 3.1), diffusion model learns the noise estimation network  $\varepsilon_\theta(x_t, c)$ , and hence following [63] the variational lower bound (ELBO) may be rewritten as  $-\mathbb{E}_\varepsilon \left[ \sum_{t=2}^T w_t \|\varepsilon - \varepsilon_\theta(x_t, c)\|_2^2 - \log p_\theta(x_0 | x_1, c) \right] + C$ , which can be further simplified to  $-\mathbb{E}_{\varepsilon, t} \left[ \|\varepsilon - \varepsilon_\theta(x_t, c)\|_2^2 \right] + C$ , assuming  $w_t = 1$  and  $\log p_\theta(x_0 | x_1, c) \approx 0$  for large  $T = 1000$  [6, 63].

**Classification.** For classification [6, 63], we leverage Bayes’ Theorem on model predictions  $p_\theta(x | c_i)$  over labels  $\{c_i\}$  and obtain the inverse probability:  $p_\theta(c_i | x) = \frac{p(c_i)p_\theta(x|c_i)}{\sum_j p(c_j)p_\theta(x|c_j)}$ . Simplifying  $p_\theta(x | c)$  with the ELBO derived above, we can rewrite the formulation as:

$$p_\theta(c_i | x) = \frac{\exp\{-\mathbb{E}_{\varepsilon, t} [\|\varepsilon - \varepsilon_\theta(x_t, c_i)\|_2^2]\}}{\sum_j \exp\{-\mathbb{E}_{\varepsilon, t} [\|\varepsilon - \varepsilon_\theta(x_t, c_j)\|_2^2]\}} \quad (4)$$

Following [63], an unbiased Monte Carlo estimate is calculated for each expectation by sampling  $N(t_i, \varepsilon_i)$  pairs and computing  $\frac{1}{N} \sum_{i=1}^N \|\varepsilon_i - \varepsilon_\theta(\sqrt{\bar{\alpha}_{t_i}}x + \sqrt{1 - \bar{\alpha}_{t_i}}\varepsilon_i, c_j)\|_2^2$ . Plugging this quantity in Equation 4 gives us the diffusion classifier.

**Adapting to our few-shot setting.** In our framework, the conditioning signal  $c$  is derived from the learned class prompts, i.e.  $c_i = c_\theta(p_{c_i})$ . Using this value, along with a further approximation of Equation 4 noted by [63], leads to the following equation:

$$p_\theta(c_i | x) = \frac{1}{\sum_j \exp\{\mathbb{E}_{\varepsilon, t} [\|\varepsilon - \varepsilon_\theta(x_t, c_\theta(p_{c_i}))\|_2^2 - \|\varepsilon - \varepsilon_\theta(x_t, c_\theta(p_{c_j}))\|_2^2]\}} \quad (5)$$

Thus, given a test image sample  $x$ , we can predict its class label  $c_i$  via conditional denoising across all class prompts  $\{p_{c_i}\}$  and estimating the relative differences in prediction errors across each condition (i.e.  $c_\theta(p_{c_i})$ ). Note that the class prompts  $\{p_{c_i}\}$  have been learned via GCPL (or CoMPLe) in a few-shot manner, hence we call it a *few-shot diffusion classifier*.

## 4 Experiments

### 4.1 Datasets

We conduct experiments on six diverse image classification datasets with fine-grained categories. span across multiple domains of computer vision, making our empirical observations

generalizable and robust.

**Natural images:** We use StanfordCars [51] comprising 196 fine-grained car model categories, Flowers102 [40] consisting of 102 floral subspecies and Cornseeds [53] with 4 seed image variants.

**Medical imaging:** We use the CRC5k [26], ISIC2018 [39] and LC25000 [8] datasets. While CRC5k and LC25000 are histopathology datasets comprising 8 colorectal and 5 lung-colon tissue categories respectively, ISIC2018 is a dermatology image corpus spanning 7 types of skin lesions. These can be considered *fine-grained* image datasets, as all image categories belong to the same anatomical region and only differ in minute attributes.

**Hard-to-describe visual concepts:** We use the Fractals [25] dataset, comprising 60 categories of intricate fine-grained visual patterns. We conjecture these abstract images are hard to describe via a hard-coded word in a text vocabulary [45] (intuitively supported by empirical findings), which drives the motivation of prompt learning class prototypes using a generative approach that imbibes visual attributes into the learned embeddings.

## 4.2 Implementation

Our model is implemented in PyTorch [44] accelerated by Nvidia RTX A6000 48GB GPU. We use Stable Diffusion v1.4 [46] as the frozen text-to-image diffusion model that encapsulates a CLIP [45] text encoder with embedding dimension  $d = 768$ . For class prompting, hand-crafted prompt templates are used that contains dataset-specific context [53] words (shown in Table 1) as well as the learnable class token. For GCPL, we follow [19] and set the following training hyperparameters: AdamW [36] optimizer ( $\beta_1 = 0.9, \beta_2 = 0.999, \omega_{decay} = 0.01$ ) with learning rate of  $5e - 4$  and trained for 2000 epochs. For CoMPL, the learning rate was set to  $1e - 3$ , batch size of 4, and trained for 4000 epochs; the contrastive weighting factor  $\lambda$  being set to 0.001 (based on empirical findings). For downstream inference, we followed the protocol outlined in [53], replacing the CLIP [45] class embeddings with *our prompt learned embeddings* and using them for the conditional denoising process [53, 46].

Dataset	Visual concept	Prompt template	Initializer word
StanfordCars [51]	Vehicular variants	“A photo of [CLASS], a type of car.”	car
Cornseeds [53]	Natural images, agriculture	“A photo of [CLASS] corn seed.”	seed
CRC5k [26]	Histopathology	“[CLASS] tissue.”	tissue
ISIC2018 [39]	Dermatology	“[CLASS] skin lesion.”	skin
LC25000 [8]	Histopathology	“[CLASS] tissue.”	tissue
Fractals [25]	Abstract imagery	“[CLASS] fractal.”	fractal

Table 1: Handcrafted prompt templates used for different datasets, where [CLASS] denotes the learnable class prompt token.

## 4.3 Competitors

**Zero-shot approaches.** First, we compare our approach against zero-shot baselines such as CLIP [45] and Diffusion Classifier [53]. Both of these baselines quantify the innate representational knowledge embedded in their respective foundational models [45, 46], with one being discriminative and the other generative. Since the results reported in [53] were obtained using SD v2.0 (i.e. different checkpoint), we rerun their codes using SD v1.4 and report the results.



Method	StanfordCars [68]	Cornseeds [68]	Flowers102 [68]	CRC5k [68]	ISIC2018 [68]	LC25000 [6]	Fractals [68]
<b>Zero-Shot</b>							
CLIP [45]	65.56	18.47	70.73	21.49	14.43	25.40	9.25
Diffusion Classifier [68]	76.77	17.77	54.21	24.16	10.41	17.29	6.25
<b>Adapter</b>							
Tip-Adapter [68]	65.82 ± 0.51	34.27 ± 3.97	89.28 ± 0.55	59.90 ± 2.18	33.88 ± 7.26	80.48 ± 1.93	81.49 ± 1.22
Tip-Adapter-F [68]	75.14 ± 0.35	39.61 ± 2.88	94.25 ± 0.43	71.44 ± 2.46	40.32 ± 5.19	86.02 ± 1.59	86.16 ± 0.54
<b>Prompt learning</b>							
CoCoOp [69]	71.57 ± 0.76	36.56 ± 5.42	87.84 ± 0.48	60.91 ± 2.98	24.67 ± 6.54	73.86 ± 4.19	67.89 ± 1.29
KgCoOp [64]	78.76 ± 0.61	38.45 ± 4.84	91.97 ± 0.44	59.90 ± 5.17	29.16 ± 6.82	75.87 ± 3.88	72.84 ± 0.93
MaPLe [67]	74.39 ± 0.43	34.37 ± 15.44	93.96 ± 0.61	40.56 ± 16.12	30.33 ± 13.67	71.96 ± 5.22	76.91 ± 6.55
PromptSRC [68]	83.33 ± 0.35	33.69 ± 4.55	<b>97.06 ± 0.27</b>	56.45 ± 18.28	44.18 ± 7.02	77.54 ± 1.51	<b>93.45 ± 0.52</b>
<b>Ours</b>							
Ours-GCPL	<b>88.47 ± 0.27</b>	43.42 ± 2.84	93.45 ± 1.39	74.76 ± 1.94	48.84 ± 2.13	93.44 ± 0.78	90.76 ± 2.23
Ours-CoMPLe	87.69 ± 1.47	<b>45.79 ± 2.12</b>	90.73 ± 1.05	<b>76.36 ± 1.82</b>	<b>49.27 ± 2.59</b>	<b>94.83 ± 0.28</b>	88.83 ± 1.57

Table 2: Comparison of our proposed approaches against SoTA VLM adaptation (16-shot).

**Few-shot VLM adaptation methods.** As discussed earlier, two prominent directions of few-shot VLM adaptation include use of learnable adapters [68] and prompt learning [27, 64, 69]. We compare our approach with (a) Tip-Adapter and Tip-Adapter-F [68], the latter being a fine-tuning variant of the otherwise training-free version; and (b) SoTA few-shot prompt learning methods: CoCoOp [69], KgCoOp [64] MaPLe [27] and PromptSRC [68]. As discussed in Section 2, these methods use the fixed class embedding initialised from CLIP and optimizes the context tokens of the prompt, which is *orthogonal* to our generatively-learned approach that directly optimizes the [CLASS] embedding in a handcrafted prompt.

## 5 Performance Analysis

### 5.1 Few-shot Classification

We compare against the above-mentioned methods under few-shot settings, and report the 16-shot comparisons in Table 2. We observe several intriguing patterns. Firstly, zero-shot models severely fail to generalise to newer domains (medical imaging or even abstract visual patterns like Fractal). This is in accordance to our discussion early-on regarding the inherent limitations of CLIP [45]. Secondly, here we also notice the weaknesses in prompt learning approaches [27, 68, 69], especially under domain-shift and fine-grained categories. The high standard deviation scores for prompt learning methods on medical imaging show their noisy representations, conjecturing the importance of class token learning. Overall, our generation-aided class prompting leads to favourable performances across most datasets, and GCPL despite having no inter-class discriminatory knowledge emerges as a strong feature learning paradigm. It is interesting to note however that, though CoMPLe is supposedly an improved version of GCPL, it deteriorates on datasets with larger number of classes [25, 61, 61], though it shows superior performance on smaller datasets [6, 26, 69]. We note this a very intriguing observation and hope future research on this direction would possibly explore this anomaly.

### 5.2 Varying the number of shots

We further experiment across different number of shots per class – 1, 4, 8, 16 across all of the datasets to see how our method fares in comparison to others under severely low-data settings. The results are graphically depicted in Figure 3. We observe that both GCPL and CoMPLe show robust trends across all datasets, whereas many prompt learning approaches [27, 69] show spurious behavior on out-of-domain datasets. This convincingly proves that our class prompting paradigm is a strong few-shot learner.




Figure 3: Few-shot performance across varying number of shots per class, over the 7 datasets used in this work. Note that the first sub-figure depicts the **mean** few-shot performance over all datasets (following prior works [27, 69]).

## 6 Limitations

While our methods, GCPL and CoMPL, significantly improve the fine-grained image recognition capabilities of VLMs, they also face certain limitations, which we feel are important to discuss. Diffusion models [46], in particular, are resource-intensive and pose a barrier to widespread adoption. GCPL’s class-specific training process is inherently slow, while CoMPL’s higher GPU memory requirements, due to multiple denoising steps per batch, further increase the computational burden. This stands in contrast to prompt learning [27] and adapter tuning [68], which are less resource-intensive. Furthermore, the need for small batch sizes in CoMPL, driven by memory constraints, may affect performance, as larger batches typically improve contrastive learning outcomes [10]. Moreover, the slow inference speed of the diffusion classifier, a common issue across generative classifiers [23, 53], limits its practical deployment. Addressing these limitations will lead to more powerful and practical generation-aided discriminative models.

## 7 Conclusion

In this work, we proposed a radically different paradigm for fine-grained visio-linguistic representation learning with limited data, in the form of two generative class prompt learning frameworks, GCPL and CoMPL. We identified the potential pitfalls of CLIP representations for fine-grained and out-of-domain categories, which we were able to overcome significantly via generative modeling. Despite showing great promise, our methods require high computational resources, limiting their practical usage. Future work should focus on optimizing these methods for greater computational efficiency and scalability. As generative classifiers continue to garner interest due to their emergent properties [23], we anticipate that forthcoming research will build on our work, overcoming current challenges and unlocking new capabilities in the intersection of generative and discriminative representation learning.

## Acknowledgements

This work acknowledges the Spanish projects GRAIL PID2021-126808OB-I00, DocAI 2021-SGR-01559, the CERCA Program / Generalitat de Catalunya, and PhD Scholarship from AGAUR 2023 FI-3- 00223. The authors also acknowledge Gedas Bertasius and Feng Cheng of UNC Chapel Hill for constructive discussions and hardware resources.

## References

- [1] Yuval Alaluf, Elad Richardson, Gal Metzer, and Daniel Cohen-Or. A neural space-time representation for text-to-image personalization. In *SIGGRAPH*, 2023.
- [2] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. In *NeurIPS*, 2022.
- [3] Sidra Aleem, Fangyijie Wang, Mayug Maniparambil, Eric Arazo, Julia Dietlmeier, Kathleen Curran, Noel E O'Connor, and Suzanne Little. Test-time adaptation with salip: A cascade of sam and clip for zero shot medical image segmentation. *arXiv preprint arXiv:2404.06362*, 2024.
- [4] Arpit Bansal, Hong-Min Chu, Avi Schwarzschild, Soumyadip Sengupta, Micah Goldblum, Jonas Geiping, and Tom Goldstein. Universal guidance for diffusion models. In *ICLR*, 2024.
- [5] Moinak Bhattacharya and Prateek Prasanna. Gazediff: A radiologist visual attention guided diffusion model for zero-shot disease classification. In *Medical Imaging with Deep Learning*, 2024.
- [6] Andrew A Borkowski, Marilyn M Bui, L Brannon Thomas, Catherine P Wilson, Lauren A DeLand, and Stephen M Mastorides. Lung and colon cancer histopathological image dataset (lc25000). *arXiv preprint arXiv:1912.12142*, 2019.
- [7] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. In *NeurIPS*, 2020.
- [8] Ryan Burgert, Kanchana Ranasinghe, Xiang Li, and Michael S Ryoo. Peekaboo: Text to image diffusion models are zero-shot segmentors. *arXiv preprint arXiv:2211.13224*, 2022.
- [9] Guangyi Chen, Weiran Yao, Xiangchen Song, Xinyue Li, Yongming Rao, and Kun Zhang. Plot: Prompt learning with optimal transport for vision-language models. In *ICLR*, 2023.
- [10] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, 2020.
- [11] Wei-Yu Chen, Yen-Cheng Liu, Zsolt Kira, Yu-Chiang Frank Wang, and Jia-Bin Huang. A closer look at few-shot classification. In *ICLR*, 2019.

- [12] Yinbo Chen, Zhuang Liu, Huijuan Xu, Trevor Darrell, and Xiaolong Wang. Meta-baseline: Exploring simple meta-learning for few-shot learning. In *ICCV*, 2021.
- [13] Kevin Clark and Priyank Jaini. Text-to-image diffusion models are zero shot classifiers. *NeurIPS*, 2023.
- [14] Barbara Toniella Corradini, Mustafa Shukor, Paul Couairon, Guillaume Couairon, Franco Scarselli, and Matthieu Cord. Freeseq-diff: Training-free open-vocabulary segmentation with diffusion models. *arXiv preprint arXiv:2403.20105*, 2024.
- [15] Jian Ding, Nan Xue, Gui-Song Xia, and Dengxin Dai. Decoupling zero-shot semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11583–11592, 2022.
- [16] 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. In *ICLR*, 2021.
- [17] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, 2017.
- [18] Chelsea Finn, Kelvin Xu, and Sergey Levine. Probabilistic model-agnostic meta-learning. In *NeurIPS*, 2018.
- [19] Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. An image is worth one word: Personalizing text-to-image generation using textual inversion. In *ICLR*, 2023.
- [20] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022.
- [21] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020.
- [22] Yunshi Huang, Fereshteh Shakeri, Jose Dolz, Malik Boudiaf, Houda Bahig, and Ismail Ben Ayed. Lp++: A surprisingly strong linear probe for few-shot clip. *arXiv preprint arXiv:2404.02285*, 2024.
- [23] Priyank Jaini, Kevin Clark, and Robert Geirhos. Intriguing properties of generative classifiers. In *ICLR*, 2024.
- [24] 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 *ICML*, 2021.
- [25] Hirokatsu Kataoka, Kazushige Okayasu, Asato Matsumoto, Eisuke Yamagata, Ryosuke Yamada, Nakamasa Inoue, Akio Nakamura, and Yutaka Satoh. Pre-training without natural images. In *ACCV*, 2020.
- [26] Jakob Nikolas Kather, Johannes Krisam, Pornpimol Charoentong, Tom Luedde, Esther Herpel, Cleo-Aron Weis, Timo Gaiser, Alexander Marx, Nektarios A Valous, Dyke Ferber, et al. Predicting survival from colorectal cancer histology slides using deep learning: A retrospective multicenter study. *PLoS medicine*, 16(1):e1002730, 2019.

- [27] Muhammad Uzair Khattak, Hanoona Rasheed, Muhammad Maaz, Salman Khan, and Fahad Shahbaz Khan. Maple: Multi-modal prompt learning. In *CVPR*, 2023.
- [28] Muhammad Uzair Khattak, Syed Talal Wasim, Muzammal Naseer, Salman Khan, Ming-Hsuan Yang, and Fahad Shahbaz Khan. Self-regulating prompts: Foundational model adaptation without forgetting. In *ICCV*, 2023.
- [29] Subhadeep Koley, Ayan Kumar Bhunia, Aneeshan Sain, Pinaki Nath Chowdhury, Tao Xiang, and Yi-Zhe Song. Text-to-image diffusion models are great sketch-photo match-makers. *arXiv preprint arXiv:2403.07214*, 2024.
- [30] Subhadeep Koley, Ayan Kumar Bhunia, Aneeshan Sain, Pinaki Nath Chowdhury, Tao Xiang, and Yi-Zhe Song. You’ll never walk alone: A sketch and text duet for fine-grained image retrieval. *arXiv preprint arXiv:2403.07222*, 2024.
- [31] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *ICCV Workshops*, 2013.
- [32] Steinar Laenen and Luca Bertinetto. On episodes, prototypical networks, and few-shot learning. In *NeurIPS*, 2021.
- [33] Alexander C Li, Mihir Prabhudesai, Shivam Duggal, Ellis Brown, and Deepak Pathak. Your diffusion model is secretly a zero-shot classifier. In *ICCV*, 2023.
- [34] Yiting Li, Adam Goodge, Fayao Liu, and Chuan-Sheng Foo. Promptad: Zero-shot anomaly detection using text prompts. In *WACV*, 2024.
- [35] Fan Liu, Tianshu Zhang, Wenwen Dai, Wenwen Cai, Xiaocong Zhou, and DeLong Chen. Few-shot adaptation of multi-modal foundation models: A survey. *arXiv preprint arXiv:2401.01736*, 2024.
- [36] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2018.
- [37] Xu Luo, Hao Wu, Ji Zhang, Lianli Gao, Jing Xu, and Jingkuan Song. A closer look at few-shot classification again. In *ICML*, 2023.
- [38] Subhansu Maji, Esa Rahtu, Juho Kannala, Matthew Blaschko, and Andrea Vedaldi. Fine-grained visual classification of aircraft. *arXiv preprint arXiv:1306.5151*, 2013.
- [39] Md Ashraful Alam Milton. Automated skin lesion classification using ensemble of deep neural networks in isic 2018: Skin lesion analysis towards melanoma detection challenge. *arXiv preprint arXiv:1901.10802*, 2019.
- [40] Alex Nichol and John Schulman. Reptile: a scalable metalearning algorithm. *arXiv preprint arXiv:1803.02999*, 2018.
- [41] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *ICVGIP*, 2008.
- [42] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.

- [43] Omiros Pantazis, Gabriel Brostow, Kate Jones, and Oisín Mac Aodha. Svl-adapter: Self-supervised adapter for vision-language pretrained models. *arXiv preprint arXiv:2210.03794*, 2022.
- [44] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. 2019.
- [45] 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 supervision. In *ICML*, 2021.
- [46] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022.
- [47] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, 2015.
- [48] Shuvendu Roy and Ali Etemad. Consistency-guided prompt learning for vision-language models. In *ICLR*, 2024.
- [49] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *CVPR*, 2023.
- [50] Aneeshan Sain, Ayan Kumar Bhunia, Pinaki Nath Chowdhury, Subhadeep Koley, Tao Xiang, and Yi-Zhe Song. Clip for all things zero-shot sketch-based image retrieval, fine-grained or not. In *CVPR*, 2023.
- [51] Shoaib Ahmed Siddiqui, Andreas Dengel, and Sheraz Ahmed. Analyzing the potential of zero-shot recognition for document image classification. In *ICDAR*, 2021.
- [52] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *NeurIPS*, 2017.
- [53] M Sobhana, Pranathi Dabbara, Girija Ravulapalli, and Krishna Sahithi Kakunuri. Seed quality prediction using computer vision and convolutional neural networks. In *IEEE ICIRCA*, 2022.
- [54] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip HS Torr, and Timothy M Hospedales. Learning to compare: Relation network for few-shot learning. In *CVPR*, 2018.
- [55] Bowen Tang, Jing Zhang, Long Yan, Qian Yu, Lu Sheng, and Dong Xu. Data-free generalized zero-shot learning. In *AAAI*, 2024.
- [56] Yonglong Tian, Yue Wang, Dilip Krishnan, Joshua B Tenenbaum, and Phillip Isola. Rethinking few-shot image classification: a good embedding is all you need? In *ECCV*, 2020.
- [57] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017.



- [58] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. In *NeurIPS*, 2016.
- [59] Yaqing Wang, Quanming Yao, James T Kwok, and Lionel M Ni. Generalizing from a few examples: A survey on few-shot learning. *ACM computing surveys (CSUR)*, 53(3):1–34, 2020.
- [60] Davis Wertheimer, Luming Tang, and Bharath Hariharan. Few-shot classification with feature map reconstruction networks. In *CVPR*, 2021.
- [61] Jijie Wu, Dongliang Chang, Aneeshan Sain, Xiaoxu Li, Zhanyu Ma, Jie Cao, Jun Guo, and Yi-Zhe Song. Bi-directional feature reconstruction network for fine-grained few-shot image classification. In *AAAI*, 2023.
- [62] Ziyang Wu, Yuwei Li, Lihua Guo, and Kui Jia. Parn: Position-aware relation networks for few-shot learning. In *ICCV*, 2019.
- [63] JiaLu Xing, JianPing Liu, Jian Wang, LuLu Sun, Xi Chen, XunXun Gu, and YingFei Wang. A survey of efficient fine-tuning methods for vision-language models—prompt and adapter. *Computers & Graphics*, 2024.
- [64] Hantao Yao, Rui Zhang, and Changsheng Xu. Visual-language prompt tuning with knowledge-guided context optimization. In *CVPR*, 2023.
- [65] Chi Zhang, Yujun Cai, Guosheng Lin, and Chunhua Shen. Deepemd: Few-shot image classification with differentiable earth mover’s distance and structured classifiers. In *CVPR*, 2020.
- [66] Jingyi Zhang, Jiaying Huang, Sheng Jin, and Shijian Lu. Vision-language models for vision tasks: A survey. *IEEE TPAMI*, 2024.
- [67] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *ICCV*, 2023.
- [68] Renrui Zhang, Rongyao Fang, Wei Zhang, Peng Gao, Kunchang Li, Jifeng Dai, Yu Qiao, and Hongsheng Li. Tip-adapter: Training-free clip-adapter for better vision-language modeling. In *ECCV*, 2022.
- [69] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. In *CVPR*, 2022.
- [70] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *IJCV*, 2022.
- [71] Qihang Zhou, Guansong Pang, Yu Tian, Shibo He, and Jiming Chen. Anomaly-clip: Object-agnostic prompt learning for zero-shot anomaly detection. *arXiv preprint arXiv:2310.18961*, 2023.
- [72] Ziqin Zhou, Yinjie Lei, Bowen Zhang, Lingqiao Liu, and Yifan Liu. Zegclip: Towards adapting clip for zero-shot semantic segmentation. In *CVPR*, 2023.
- [73] Beier Zhu, Yulei Niu, Yucheng Han, Yue Wu, and Hanwang Zhang. Prompt-aligned gradient for prompt tuning. In *ICCV*, 2023.