

LiteEmbed: Adapting CLIP to Rare Classes

Aishwarya Agarwal^{1,2} Srikrishna Karanam² Vineet Gandhi¹
 aishagar@adobe.com skaranam@adobe.com v Gandhi@iiit.ac.in

¹CVIT, Kohli Centre for Intelligent Systems, IIIT Hyderabad, India

²Adobe Research, Bengaluru, India

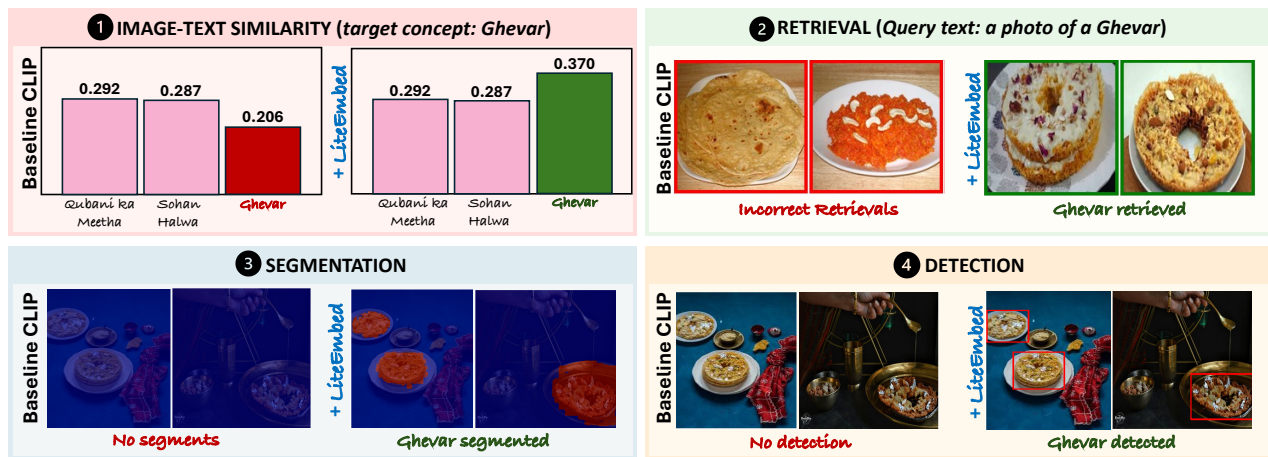


Figure 1. Given a target concept, we introduce LiteEmbed, which refines CLIP’s text embedding through Subspace-Guided Optimization using only 3–5 reference images. The optimized embedding plugs directly into CLIP-based models, here applied to classification, retrieval, segmentation [24], and detection [7], and consistently improves performance without task-specific retraining.

Abstract

Large-scale vision–language models such as CLIP achieve strong zero-shot recognition but struggle with classes rarely seen in pretraining, including newly emerging entities and culturally specific classes. We introduce LiteEmbed, a lightweight framework for few-shot personalization of CLIP that enables new classes to be added without retraining its encoders. LiteEmbed performs subspace-guided optimization of text embeddings within CLIP’s vocabulary, leveraging a PCA-based decomposition that disentangles coarse semantic directions from fine-grained variations. Two complementary objectives, coarse alignment and fine separation, jointly preserve global semantic consistency while enhancing discriminability among visually similar classes. Once optimized, the embeddings are plug-and-play, seamlessly substituting CLIP’s original text features across classification, retrieval, segmentation, and detection. Extensive experiments demonstrate substantial gains over prior methods, establishing LiteEmbed as an effective approach for adapting CLIP to underrepresented, rare or unseen classes.

1. Introduction

Large-scale Vision–Language Models (VLMs) like CLIP [27] have brought about an era of powerful zero-shot visual recognition. Yet, their knowledge is limited to training domains: newly emerging classes (e.g., a recently prominent actor) or culturally specific items (e.g., regional Indian dishes) remain poorly represented. Fig. 2 reports CLIP’s zero-shot classification accuracy, showing that while common Food-101 [6] and Indian Actor classes [4] score high, 24% of regional dishes and 14% of emerging actors receive zero recognition, with most others below 40%. This highlights systematic failure on underrepresented classes and a critical gap in real-world use-cases.

To bridge this gap, we explore extending CLIP to new classes using only 3–5 images, without retraining its large encoders or relying on prior data. Our goal is to make new classes immediately usable across any task or pipeline that leverages CLIP’s text embeddings, from classification




Figure 2. Performance disparity in CLIP’s zero-shot classification.

and retrieval to generation, segmentation, and detection. Achieving this requires preserving CLIP’s core strengths: (i) alignment, maintaining a coherent cross-modal representation; (ii) discrimination, preserving clear separability among visually similar fine-grained classes; and (iii) compositionality, enabling flexible use of new classes within natural prompts such as “a boy eating Ghevar”. In essence, the adapted model should operate as if these classes had been part of CLIP’s pretraining, generalizing seamlessly while retaining full downstream utility.

Prior efforts to adapt CLIP have shown strong task-specific gains but are not directly suited for few-shot, extensible class addition. Adapter-based approaches [13, 42] insert small trainable modules into CLIP’s frozen backbone to specialize on tasks such as classification. While efficient, these modules modify the shared representation space, making the adapted features task-specific and less generalizable to other applications such as retrieval, scoring, or generation. Prompt tuning methods [1, 10, 18, 38, 46, 47] learn contextual tokens to improve alignment across known classes but rely on large-scale supervision and shared context embeddings rather than class-specific ones, limiting discrimination and compositionality in few-shot, single-class settings. Moreover, these methods assume static training distributions and cannot extend CLIP incrementally. Together, these limitations highlight the need for approaches that can flexibly introduce new classes into CLIP’s embedding space, while preserving its core properties of alignment, discrimination, and compositionality.

Building on these insights, we introduce **LiteEmbed**, a lightweight optimization framework for adapting CLIP to rare classes. Each new class is represented by a learnable token within CLIP’s vocabulary, individually optimized using few-shot examples. We observe that directly aligning a new embedding with the few-shot examples by maximizing similarity often leads to two key failures: (a) the embedding may diverge from its semantic neighborhood (e.g., Golden Retriever separating from other dog-related classes), or (b) collapse toward visually similar cat-

egories (e.g., Labrador Retriever), thereby reducing fine-grained discriminability. We posit that semantic alignment and discrimination operate at distinct scales: alignment should preserve coarse global semantics (e.g., recognizing a Golden Retriever as a dog, not a cat), while discrimination should capture subtle distinctions among related breeds. Interestingly, a PCA analysis of CLIP’s text embeddings supports this hypothesis, revealing that high-variance directions encode coarse categorical structure, whereas low-variance directions capture fine-grained variations (see Sec. 3.2).

To this end, LiteEmbed employs subspace-guided optimization, which leverages CLIP’s semantic structure through two complementary loss terms. For each novel class (e.g., Golden Retriever), we assemble a semantic neighborhood of related coarse classes (e.g., dog, canine) and a visually confusable set of fine-grained negatives (e.g., Labrador Retriever, Irish Setter). PCA on their combined text embeddings defines a local subspace that disentangles coarse, high-variance directions from fine-grained, low-variance ones. We then introduce (i) a **Coarse Alignment Loss** that aligns the embedding with the semantic neighborhood along high-variance components, and (ii) a **Fine Separation Loss** that enforces separation from confusable classes along low-variance components. These subspace-aware objectives yield a lightweight, training-free adaptation of CLIP to new classes, requiring only minimal optimization over few-shot examples at test time.

LiteEmbed produces embeddings that are fully plug-and-play: once optimized, they can seamlessly substitute CLIP’s original text embeddings across diverse downstream tasks. This includes classification, retrieval, segmentation, and detection, all without task-specific retraining. Fig. 1 demonstrates substantial gains using LiteEmbed in retrieving accurate images, generating faithful masks, detecting correct objects, and producing compositional generations.

We comprehensively evaluate LiteEmbed across classification and retrieval tasks on three publicly available datasets: TV100 [44], Korean Celebrities [30], and Indian Food Images [3]. To facilitate a more diverse and rigorous evaluation of adaptation to novel and underrepresented classes, we further introduce **NOVA** (New Object and Visual Alignment), a benchmark comprising five meticulously curated datasets. NOVA encompasses game characters, emerging landmarks, fashion ensembles, and recently prominent public figures, thereby establishing a realistic and challenging testbed for few-shot generalization. In addition, LiteEmbed is assessed on UECFood100 [6] for segmentation and detection tasks, and on CustomConcept101 [19] for text-to-image generation. Formally, our key contributions are as follows:

- We propose LiteEmbed, a Subspace-Guided Optimiza-

tion framework that adapts CLIP to rare classes from a few visual exemplars. PCA is employed to disentangle CLIP’s text-embedding manifold into global and local semantic subspaces, within which novel-class token embeddings are optimized through two complementary losses that jointly preserve coarse alignment and enhance fine-grained discriminability.

- We present comprehensive quantitative, qualitative, and ablative experiments across ten datasets and five tasks, demonstrating that LiteEmbed enables plug-and-play use of learned embeddings, directly substituting CLIP text features without retraining and yields consistent, pronounced gains over existing methods.

2. Related Works

Large-scale vision–language models [15, 20, 21, 27, 32, 35] have established strong cross-modal alignment but remain limited to their pretraining domains. A large body of work has explored adapting CLIP to new distributions through prompt or adapter tuning. Prompt-based methods [18, 23, 46–48] learn context tokens to better align textual and visual features, while adapter-based methods [13, 42] insert lightweight modules for efficient fine-tuning. Despite strong performance, these approaches depend on multiclass supervision and often modify CLIP’s shared embedding space, limiting compositionality and incremental class addition. To reduce supervision, test-time [10, 29, 36, 38, 48] and continual adaptation strategies [8, 11, 14, 16, 17, 25, 31, 39, 45] dynamically update CLIP. However, these approaches still couple adaptation with specific training or inference contexts and often suffer from degraded cross-task generalization.

A related line of work on generative personalization [2, 12, 19, 26, 28, 33, 43] learns new tokens within CLIP’s vocabulary to represent novel or user-specific entities. However, these embeddings are optimized for image synthesis and fail to generalize to the broad range of downstream tasks where CLIP is applied. In contrast, we adapt CLIP’s text embeddings via subspace-guided objectives that balance semantic consistency and visual separability, yielding a unified, few-shot, plug-and-play framework for enhancing CLIP across diverse tasks [7, 9, 22, 24, 34, 37, 40, 41].

3. Methodology

LiteEmbed, summarized in Fig. 3, learns a class-specific text representation by optimizing only the embedding e_c of the placeholder token “*” corresponding to a new class c , while keeping both CLIP encoders f_{img} and f_{text} fixed. Given an exemplar set $I_c = \{x_i\}_{i=1}^N$ corresponding to the class, the prompt p_c (e.g., “a photo of *”) is encoded into $z_c = f_{\text{text}}(p_c)$, initialized from the base CLIP text embedding z_c^0 . Adaptation is governed by




Figure 3. Overview of LiteEmbed. Given a target class (e.g., *anarsa*) and a few reference images, the learnable token e_c replaces the placeholder “*” in the prompt and is optimized while all CLIP encoders remain frozen. Subspace-Guided Optimization balances three objectives: image–text alignment ($\mathcal{L}_{\text{img-align}}$), coarse semantic anchoring ($\mathcal{L}_{\text{coarse}}$), and fine-grained separation ($\mathcal{L}_{\text{fine}}$), promoting discriminative yet semantically consistent embeddings.

three losses: a similarity-alignment loss encouraging high $\text{sim}(f_{\text{img}}(x_i), z_c)$ over exemplars, and two additional loss terms that exploit the coarse-to-fine PCA decomposition of CLIP’s text space. These terms push z_c to remain anchored to its broad semantic neighborhood along coarse components, while enforcing separation from visually confusable negatives $\mathcal{N} = \{n_k\}_{k=1}^K$ via their embeddings z_{n_k} , along fine components. We now motivate the design of our approach and delineate each component in detail.

3.1. Motivation

We begin by examining how CLIP behaves for under-represented classes. Consider the example of *anarsa*, a traditional Indian sweet from the Indian Food Images dataset. Its base CLIP text embedding z_c^0 exhibits a very low average image-text similarity (0.19) and fails completely in classification (0% accuracy), indicating poor alignment with the visual feature space (Fig. 4 (a)).

To rectify this misalignment, we first refine the learnable token embedding e_c for class c (*anarsa*) to improve its alignment with the exemplar set I_c , leveraging image-text correspondence as CLIP’s principal supervisory signal. Formally, we optimize the alignment loss:

$$\mathcal{L}_{\text{img-align}} = -\frac{1}{N} \sum_{i=1}^N \text{sim}(f_{\text{img}}(x_i), f_{\text{text}}(p_c(e_c))),$$

and update the embedding via gradient descent as $e_c \leftarrow e_c - \eta \nabla_{e_c} \mathcal{L}_{\text{img-align}}$, where η is the learning rate.

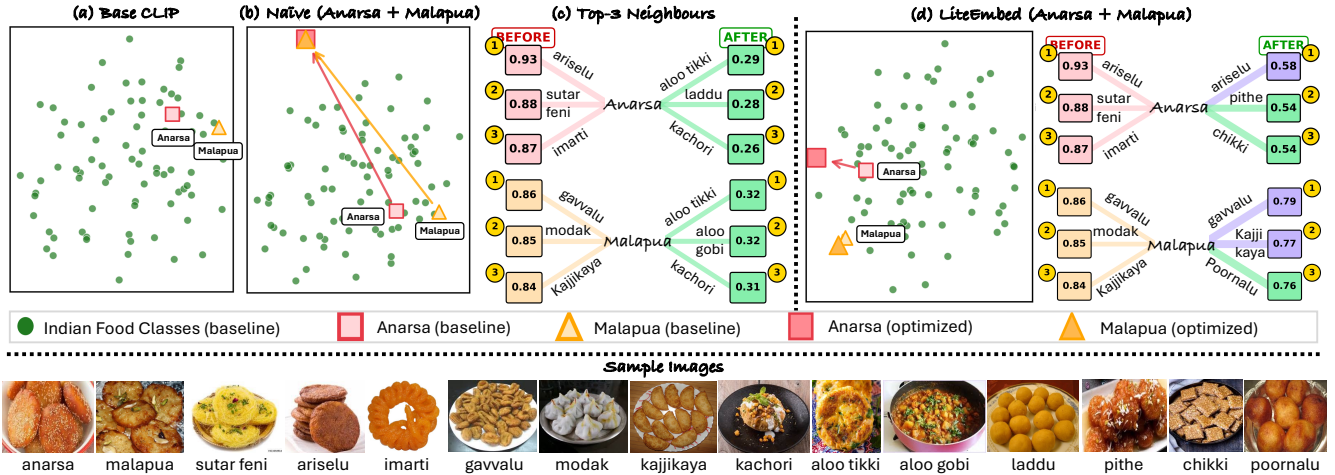


Figure 4. t-SNE visualization of CLIP’s text embeddings for *anarsa* and *malapua*. (a) Base CLIP embeddings before any adaptation. (b) Naïve image-only optimization leads to discriminative collapse, bringing the two classes too close. (c) Their nearest neighbors shift from related sweets to visually similar but semantically unrelated snacks, indicating semantic drift. (d) Our Subspace-Guided Optimization restores separation while preserving semantic relationships. Sample class images are shown below.

While such alignment-only optimization boosts performance for the adapted class, with *anarsa* achieving 100% accuracy when evaluated for classification against baseline CLIP embeddings, it quickly exposes two critical issues when multiple classes are adapted sequentially:

1. **Discriminative collapse:** When we subsequently optimize a new embedding (from scratch) for another visually similar sweet, *malapua*, the classification accuracy of *anarsa* drops sharply to 30%. The t-SNE visualization in Fig. 4 (b) reveals that the optimized embeddings for both classes almost completely overlap, indicating that the optimization has captured shared visual cues (golden-brown circular texture) but failed to encode discriminative detail.
2. **Semantic drift:** The optimized embedding for *anarsa* no longer resides within its original semantic neighborhood. In the baseline CLIP space, its nearest neighbors are related Indian sweets such as *ariselu*, *sutar feni*, and *imarti* (Fig. 4(c)). After naive optimization, these shift to visually similar but semantically unrelated snacks like *kachori* and *aloo tikki*, which are not sweets. Quantitatively, the maximum cosine similarity to any CLIP text embedding drops from about 0.93 to 0.29, confirming a substantial loss of semantic connectivity in the text manifold.

This shows that alignment-only optimization, while visually faithful, distorts the embedding’s semantic neighborhood and fails to enhance discriminative power. An effective adaptation strategy must therefore preserve **global semantic alignment**, while avoiding **discriminative collapse** with related fine-grained classes.

Accordingly, we pose a central question: *can CLIP’s embedding space itself reveal how a novel class such as *anarsa* should be encoded, retaining its global semantics (e.g., Indian sweets) while enhancing separation from visually confusable classes like *malapua*?* This motivates our analysis of CLIP’s latent structure, presented next.

3.2. Understanding CLIP’s Embedding Space

To examine the semantic structure of CLIP’s text embedding space, we conduct PCA over embeddings spanning diverse categories such as animals, vehicles, and food items. Fig. 5 visualizes the separations along the first (PC_1) and fourth (PC_4) principal components. In (a), PC_1 captures coarse semantic variance: cross-category pairs (e.g., British Shorthair vs SUV) show markedly higher separation distance than fine-grained pairs (e.g., Maine Coon vs Bengal), indicating that high-variance directions encode coarse semantic structure. The accompanying t-SNE plot reveals two distinct clusters corresponding to dog and cat breeds, consistent with this coarse organization. In contrast, (b) shows PC_4 , where fine-grained distinctions dominate and cross-category distances shrink, suggesting that lower-variance directions capture localized, discriminative semantics. The corresponding t-SNE projection exhibits mixed clusters with finer intra-breed separation, reinforcing this observation.

Overall, CLIP’s text space demonstrates a hierarchical organization: early components preserve broad semantic alignment, while later ones capture fine-grained variability, an insight that underpins our subspace-guided optimization framework for balanced alignment and discrimination.

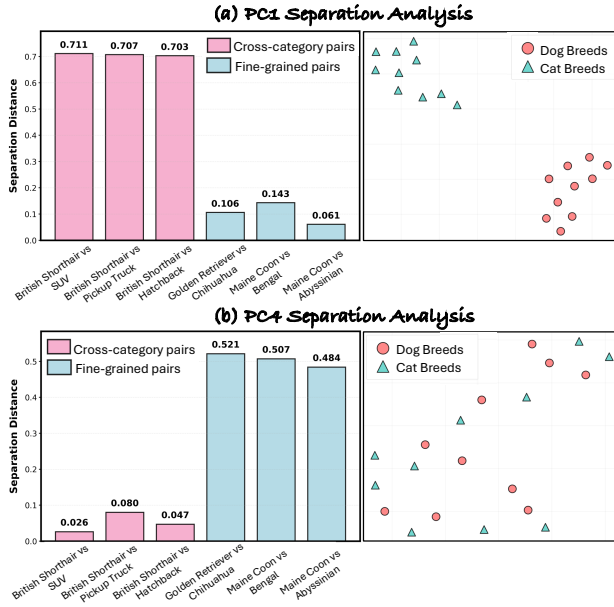


Figure 5. PCA Analysis of CLIP’s Text Embedding Space.

3.3. Subspace-Guided Optimization

Given a new class c , our goal is to learn a text embedding e_c that can be directly inserted into CLIP’s text encoder and used across downstream tasks.

For each class, we define two reference neighborhoods: a **coarse set** \mathcal{C} , containing broader semantic groups suggested by an LLM that c should stay aligned with (e.g., “Indian sweet”, “dessert”, “snack”), and a **fine set** \mathcal{F} , containing visually similar yet distinct classes (e.g., *malapua* for *anarsa*). The fine set \mathcal{F} constitutes the repository of hard-negative exemplars and is systematically constructed through a two-stage pipeline. In the first stage, a large language model (LLM) enumerates visually proximate candidates for each class. Subsequently, a CLIP-based filtering mechanism prunes this set by retaining candidates that exhibit the high similarity between their text embeddings and the image embeddings of the exemplars, thereby preserving only the most visually confusable instances.

We obtain CLIP text embeddings for all classes in these sets, each of dimensionality $d \in \{512, 768\}$ depending on the backbone. Performing PCA on the union $\mathcal{P} = \mathcal{C} \cup \mathcal{F}$ yields an orthogonal basis $U = [u_1, u_2, \dots, u_D]$, where $u_i \in \mathbb{R}^d$ and $D \leq d$. Guided by our PCA analysis in Fig. 5, the top-variance components capture coarse semantics (broad category distinctions), while the low-variance ones encode fine-grained variations. We thus split U into $U_{\text{coarse}} = [u_1, \dots, u_k]$ and $U_{\text{fine}} = [u_{k+1}, \dots, u_D]$, which serve as complementary subspaces for optimization.

The **coarse alignment loss** anchors e_c to its semantic

neighborhood along high-variance directions,

$$\mathcal{L}_{\text{coarse}} = \frac{1}{|\mathcal{C}|} \sum_{p \in \mathcal{C}} (1 - \text{sim}(U_{\text{coarse}}^\top z_c, U_{\text{coarse}}^\top z_p)),$$

while the **fine separation loss** reduces similarity to visually confusable classes along low-variance directions,

$$\mathcal{L}_{\text{fine}} = \frac{1}{|\mathcal{F}|} \sum_{n \in \mathcal{F}} \text{sim}(U_{\text{fine}}^\top z_c, U_{\text{fine}}^\top z_n).$$

The final objective combines these with the standard image–text alignment loss:

$$\mathcal{L} = \mathcal{L}_{\text{img-align}} + \lambda_1 \mathcal{L}_{\text{coarse}} + \lambda_2 \mathcal{L}_{\text{fine}}.$$

This formulation anchors e_c to its coarse semantic neighborhood while maintaining discrimination against visually similar classes, enabling few-shot, training-free adaptation without retraining CLIP, preserving compositionality and plug-and-play use across downstream tasks.

3.4. Putting It All Together

When applied back to the motivating example of *anarsa* and *malapua* (Sec. 3.1), our subspace-guided optimization effectively resolves both issues observed with naïve alignment. As shown in Fig. 4 (d), the optimized embeddings for the two classes are now well-separated, preventing the overlap that previously caused **discriminative collapse**. At the same time, **semantic drift** is mitigated. Nearest-neighbor analysis in Fig. 4 (d) reveals that the optimized *malapua* embedding remains surrounded by other Indian sweets such as *gavvalu*, *poornalu*, and *kajjikaya*, many of which were its original neighbors in CLIP’s base embedding space.

These results show that LiteEmbed attains both fine-grained visual discriminability and coarse semantic fidelity: fine-subspace movement separates confusable classes, while coarse-subspace anchoring preserves linguistic and semantic grounding.

4. Experiments

Datasets: We evaluate on 8 diverse classification datasets spanning both public and newly curated sources. The public ones include Indian Food Images [3], Korean Celebrities [30], and TV100 [44]. The remaining five namely Indian Singers, Indian Actors, Game Characters, Fashion Outfits, and Landmarks were newly collected using iCrawler as part of our proposed NOVA benchmark. For each of these, we first queried an LLM to identify classes that have emerged after 2023 within the respective domain, and then scraped 50 images per class. For downstream tasks, we use UECFood100 [5] for segmentation and detection, Indian Food Images for retrieval, and CustomConcept101 [19] for

text-to-image generation.

Experimental Setup: We evaluate our method across three complementary settings: few-shot learning, continual learning, and downstream transfer. In the few-shot setup, we follow the standard N -shot classification protocol with $N \in \{1, 2, 4, 8, 16\}$ labeled samples per class. In the continual learning setup, we adopt a class-incremental scenario where classes appear sequentially, each task providing four samples from a single class; cumulative accuracy is reported over all classes observed so far and averaged across five random class orders to reduce ordering bias. For downstream transfer, we evaluate the learned embeddings within CLIP-based pipelines for text-guided retrieval, semantic segmentation, object detection, and text-to-image generation, without additional finetuning, to assess compatibility and reuse across tasks. All experiments use CLIP ViT-B/16 as the base model, except detection, where YOLO-World [7] uses CLIP ViT-B/32 as its text encoder. Optimization is performed using Adam with a learning rate of 1×10^{-4} , a 1000-step warm-up, and 5000 total steps.

Compared Approaches: We compare our method against zero-shot CLIP and a broad range of adaptation techniques, including prompt-learning methods (CoOp [46], CoCoOp [47], MaPLe [18]), adapter-based approaches (CLIP-Adapter [13], TIP-Adapter-F [42]), and test-time adaptation methods (DiffTPT [10], C-TPT [38], DynaPrompt [36], PromptAlign [1]). For continual learning, we evaluate Continual-CLIP [31], ENGINE [45], MoE-Adapters4CL [39], and an adapted CoOp variant that relearns context tokens after each task. For downstream transfer, we benchmark against CLIP-based task models, including CLIPSeg [24] for segmentation, YOLO-World [7] for detection, Textual Inversion [12] for text-to-image generation, and standard CLIP for retrieval.

4.1. Main Results

By design, LiteEmbed is trained in a continual manner, sequentially incorporating new classes without requiring joint access to earlier samples. However, for fair comparison with other methods, we categorize experiments into (i) few-shot classification, where all class samples are provided simultaneously, allowing joint adaptation, and (ii) continual learning, where classes are introduced sequentially.

We first assess LiteEmbed under the standard few-shot classification setup with four samples per class provided simultaneously, and report the 4-shot accuracies in Tab. 1. LiteEmbed decisively outperforms existing approaches across all datasets, with an average improvement of 14% over the second-best approach (CoOp) and 35.6% over baseline CLIP zero-shot. The most notable improvements occur on low-alignment domains such as Fashion Outfits (+37.8%) and TV100 (+33.1%), where CLIP’s zero-

shot performance is especially weak. Extended 1/2/8/16-shot results appear in the supplementary material.

We next evaluate LiteEmbed in its native continual setup, shown in Fig. 6, where adaptation proceeds sequentially over new class groups. LiteEmbed consistently outperforms baselines, achieving 32% average higher final accuracy over the second-best approach (ENGINE) and significantly lower forgetting, whereas CoOp shows near-complete forgetting of previously learned context tokens. Minor performance drops across tasks, primarily arise from increased candidate classes (rather than forgetting), highlighting LiteEmbed’s ability to integrate classes progressively while maintaining stable performance.

4.2. Evaluation on Downstream Tasks

To assess the generality of LiteEmbed beyond few-shot classification, we evaluate it on downstream tasks that rely on CLIP text embeddings, including retrieval, segmentation, detection, and text-to-image generation. Unlike prior methods optimized solely for classification, LiteEmbed requires no task-specific retraining; the learned embedding simply replaces the original CLIP text embedding wherever the text encoder is used. This illustrates its plug-and-play nature and broad applicability across vision–language tasks. We benchmark against task-specific CLIP-based models, including CLIPSeg for segmentation, YOLO-World for detection, and Textual Inversion for generative adaptation.

Retrieval: We evaluate LiteEmbed on text-based image retrieval using the Indian Food Images dataset. Replacing the original CLIP text embeddings with LiteEmbed embeddings, we observe substantial gains in retrieval precision, with Precision@5 improving from 26.2% to 68.7% across 80 food classes. For the 20 most underrepresented dishes, the gains are even more pronounced, with Precision@5 improving from 13.3% to 93.6%. Qualitatively, Fig. 7 (a) illustrates common failure modes of baseline CLIP: for the class *Pootharekulu*, CLIP retrieves images from unrelated categories, while for *Misi Roti*, it confuses visually similar breads such as *Bhatura* and *Chapati*. In contrast, LiteEmbed embeddings retrieve accurate images across the top-5 positions, highlighting their ability to capture distinctive visual characteristics for rare classes.

Segmentation: CLIPSeg consumes a text prompt and outputs a corresponding segmentation mask. On the Indian Food Dataset (IFD), which contains many underrepresented categories, baseline CLIPSeg fails to generate masks for 23 of the 80 classes (zero-coverage classes, where coverage denotes the percentage of segmented pixels), underscoring its limitations on rare classes. Although IFD lacks ground-truth masks and precise quantitative evaluation is not feasible, we observe that LiteEmbed produces nonzero segmentation coverage for all classes after adaptation, eliminating all zero-coverage cases. Qualitative results in Fig. 7

Method	Indian Food (80 classes)	Korean Celebrities (50 classes)	Indian Singers (50 classes)	Indian Actors (50 classes)	Game Characters (65 classes)	Fashion Outfits (55 classes)	TV100 (60 classes)	Landmarks (35 classes)	Average
CLIP Zero-Shot	39.1	27.9	27.7	21.6	23.2	7.6	2.6	35.6	23.2
<i>Prompt learning methods</i>									
CoOp [46] CVPR'22	<u>53.3</u>	<u>52.6</u>	<u>55.2</u>	<u>43.2</u>	<u>56.5</u>	<u>30.7</u>	1.1	58.2	44.0
CoCoOp [47] IJCV'22	46.9	44.5	37.2	26.5	54.9	12.3	1.2	57.1	35.1
MaPLe [18] CVPR'23	50.9	50.7	32.8	34.8	55.2	17.5	0.9	<u>65.7</u>	38.6
<i>Adapter-based methods</i>									
CLIP-Adapter [13] IJCV'24	43.8	37.5	30.6	21.2	36.0	11.4	0.7	49.3	28.9
TIP-Adapter-F [42] ECCV'22	52.3	40.2	39.0	33.0	40.8	11.1	<u>8.5</u>	57.9	35.4
<i>Test-time adaptation methods</i>									
DiffTPT [10] ICCV'23	42.3	27.6	20.6	17.7	23.6	8.1	1.7	38.6	22.4
PromptAlign [1] NeurIPS'23	23.8	26.2	17.8	16.1	23.4	6.1	2.4	33.9	18.7
C-TPT [38] ICLR'24	39.2	23.1	19.1	15.9	21.9	6.6	1.0	36.1	20.5
DynaPrompt [36] ICLR'25	41.6	24.0	19.6	21.0	30.4	6.8	2.6	45.7	24.0
LiteEmbed (Ours)	69.1	60.6	58.9	56.9	68.9	45.4	35.7	74.3	58.7

Table 1. Few-Shot Classification Accuracy Results (4-shot). LiteEmbed consistently outperforms all baselines across 8 diverse datasets.

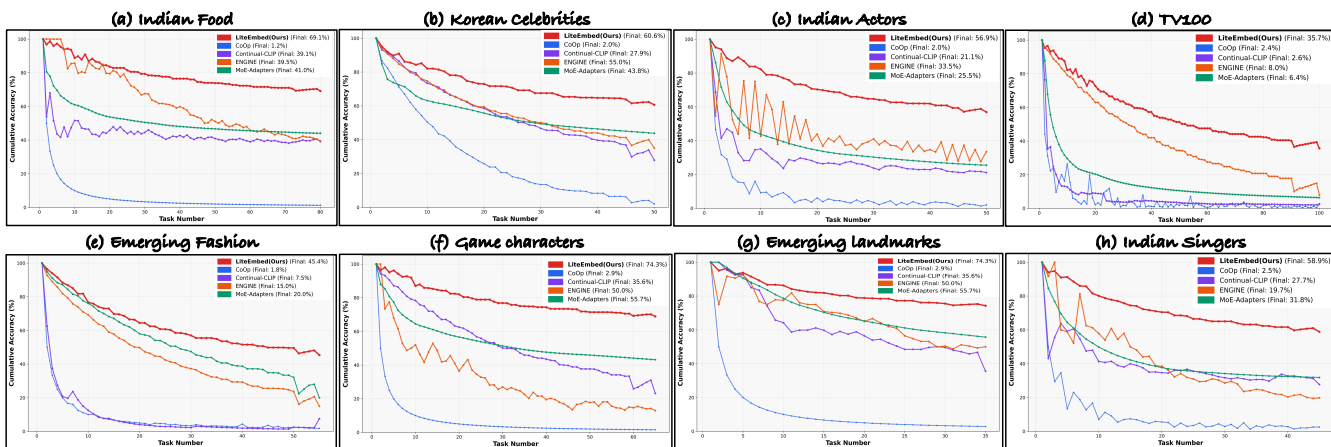


Figure 6. Comparison of 4-shot continual adaptation results across sequentially added classes. LiteEmbed in red, ENGINE in orange, CoOp in blue, Continual-CLIP in purple, and MoE-Adapters in green; best viewed in color.

illustrate this effect. Panel (b), first row, shows that baseline CLIPSeg fails to segment *adhirasam*, a failure that is fully recovered with LiteEmbed.

UECFood100 provides ground-truth segmentation masks and includes multi-object images, where CLIPSeg often lacks discrimination and segments incorrect regions. Quantitatively, LiteEmbed delivers a substantial IoU improvement, increasing overall IoU from 0.43 to 0.71. Qualitative examples in Fig. 7(b) illustrate this effect: in the third row, baseline CLIPSeg segments additional objects beyond the intended *toast*, whereas LiteEmbed correctly isolates the target; the fourth-row *Chinese soup* example shows a similar correction. These results confirm that LiteEmbed significantly enhances both coverage and mask fidelity for rare and visually confusable categories.

Detection: We evaluate LiteEmbed on UECFOOD100 open-vocabulary detection using YOLO-World, replacing class text embeddings with the optimized embeddings. We report mean IoU (mIoU) as the primary evaluation metric. With LiteEmbed, overall mIoU notably increases from

0.37 to 0.72, compared to using baseline CLIP embeddings. Qualitatively, Fig. 7 (c) illustrates typical failure modes of YOLO-World with baseline embeddings: in the first two rows, objects of interest (*Nanbanzuke* and *Miso soup*) are entirely missed, while in others, incorrect items are instead detected. Using LiteEmbed embeddings corrects these errors, enabling accurate detection of all target objects and improving localization precision.

Generation: We extend LiteEmbed to custom image generation via Textual Inversion (TI), which learns new class tokens by optimizing embeddings in CLIP’s text space. We replace LiteEmbed’s alignment loss \mathcal{L}_{align} with TI’s reconstruction loss, keeping all other aspects of training unchanged. This ensures the learned embedding generates images faithful to the class. Qualitative results Fig. 8 on examples from the CustomConcept101 dataset show that while baseline TI struggles with compositional prompts, LiteEmbed learned embeddings successfully generate images such as a custom cat wearing sunglasses, demonstrating improved editability without sacri-

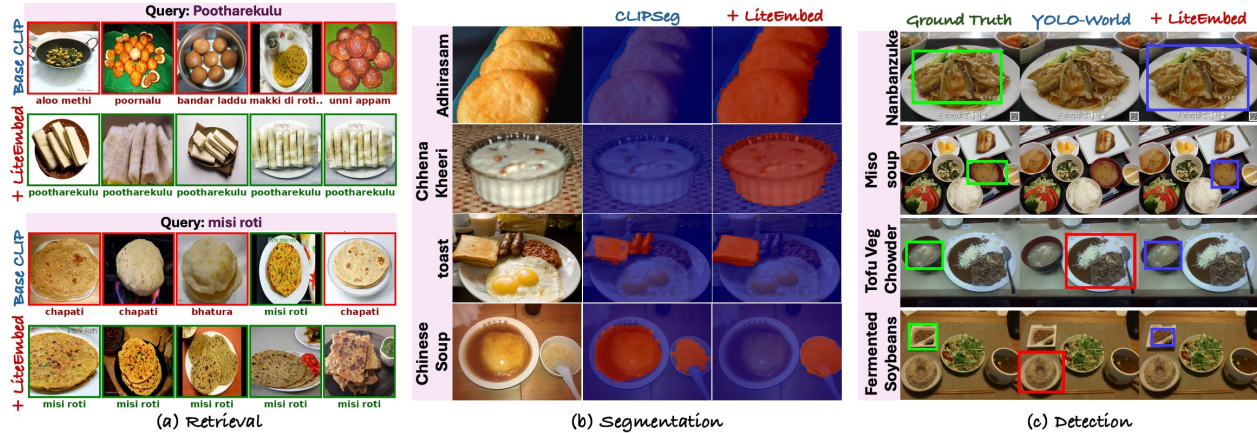


Figure 7. Qualitative results across downstream tasks. LiteEmbed consistently improves retrieval, segmentation, and detection quality.



Figure 8. Textual Inversion results showing that LiteEmbed enables more editable image generation compared to baseline TI. Here, $\langle sks \rangle$ in prompts denotes the learned class embedding.

ficing reconstruction fidelity.

4.3. Ablations

Ablation experiments on the Indian Food Images dataset evaluate the contribution of each loss component in our optimization objective. Tab. 2 reports classification accuracy and the average image–text cosine similarity between class images and their corresponding ground-truth text embeddings. Introducing the image–text alignment loss $\mathcal{L}_{\text{img-align}}$ improves alignment scores over base CLIP but does not substantially boost accuracy, indicating limited generalization when optimized alone. Adding the coarse alignment loss $\mathcal{L}_{\text{coarse}}$ leads to modest accuracy gains, likely due to reduced confusion between classes belonging to different semantic neighborhoods, while also preserving compositional

Table 2. Ablation study on the Indian Food Images dataset.

Method	Accuracy (%)	Avg. Image–Text Cosine Similarity
Base CLIP	39.1	0.28 (range: 0.18-0.35)
+ $\mathcal{L}_{\text{img-align}}$	45.3	0.65 (range: 0.56-0.71)
+ $\mathcal{L}_{\text{img-align}}$ + $\mathcal{L}_{\text{coarse}}$	52.6	0.48 (range: 0.40-0.53)
+ $\mathcal{L}_{\text{img-align}}$ + $\mathcal{L}_{\text{coarse}}$ + $\mathcal{L}_{\text{fine}}$	69.1	0.39 (range: 0.33-0.42)

structure within CLIP’s text space. Finally, incorporating the fine separation loss $\mathcal{L}_{\text{fine}}$ enhances discriminative power, yielding the highest overall accuracy. These results collectively demonstrate that the full subspace-guided objective achieves a balanced trade-off between preserving semantic consistency and enhancing visual discrimination, leading to the highest overall performance. We further provide ablations across different CLIP backbones and experiments analyzing the choice of k for subspace-guided optimization in the supplementary material.

5. Conclusion

In this work, we addressed the challenge of adapting CLIP to underrepresented and culturally diverse concepts, where standard text embeddings often fail to align with visual semantics. We introduced LiteEmbed, a plug-and-play framework that refines CLIP’s text embeddings through Subspace-Guided Optimization using only a few reference images. Our formulation balances coarse semantic anchoring with fine-grained discriminative adaptation, yielding embeddings that integrate seamlessly into CLIP-based models across tasks such as retrieval, segmentation, detection, and generation, all without task-specific retraining. Extensive experiments demonstrate consistent improvements over prior adaptation methods and validate the generality of our approach. Future work will investigate applying subspace-guided adaptation to other vision–language models, studying its generality across diverse multimodal backbones.

References

- [1] Jameel Abdul Samadh, Mohammad Hanan Gani, Noor Hussein, Muhammad Uzair Khattak, Muhammad Muzammal Naseer, Fahad Shahbaz Khan, and Salman H Khan. Align your prompts: Test-time prompting with distribution alignment for zero-shot generalization. *Advances in Neural Information Processing Systems (NeurIPS)*, 36:80396–80413, 2023. 2, 6, 7
- [2] Aishwarya Agarwal, Srikrishna Karanam, Tripti Shukla, and Balaji Vasani Srinivasan. An image is worth multiple words: Multi-attribute inversion for constrained text-to-image synthesis. In *Winter Conference on Applications of Computer Vision (WACV)*, pages 6053–6062. IEEE, 2025. 3
- [3] Sourav Banerjee. Indian food images dataset, 2022. 2, 5
- [4] Sourav Banerjee. Indian actor images dataset, 2022. 1
- [5] Elena Battini Sönmez, Sefer Memiş, Berker Arslan, and Okan Zafer Batur. The segmented uec food-100 dataset with benchmark experiment on food detection. *Multimedia Systems*, 29(4):2049–2057, 2023. 5
- [6] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101—mining discriminative components with random forests. In *European Conference on Computer Vision (ECCV)*, pages 446–461. Springer, 2014. 1, 2
- [7] Tianheng Cheng, Lin Song, Yixiao Ge, Wenyu Liu, Xingang Wang, and Ying Shan. Yolo-world: Real-time open-vocabulary object detection. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16901–16911, 2024. 1, 3, 6
- [8] Yuxuan Ding, Lingqiao Liu, Chunna Tian, Jingyuan Yang, and Haoxuan Ding. Don’t stop learning: Towards continual learning for the clip model. *arXiv preprint arXiv:2207.09248*, 2022. 3
- [9] Sepideh Esmailpour, Bing Liu, Eric Robertson, and Lei Shu. Zero-shot out-of-distribution detection based on the pre-trained model clip. In *Association for the Advancement of Artificial Intelligence (AAAI)*, pages 6568–6576, 2022. 3
- [10] Chun-Mei Feng, Kai Yu, Yong Liu, Salman Khan, and Wangmeng Zuo. Diverse data augmentation with diffusions for effective test-time prompt tuning. In *International Conference on Computer Vision (ICCV)*, pages 2704–2714, 2023. 2, 3, 6, 7
- [11] Hao Fu, Hanbin Zhao, Jiahua Dong, Henghui Ding, Chao Zhang, and Hui Qian. Iap: Improving continual learning of vision-language models via instance-aware prompting. *arXiv preprint arXiv:2503.20612*, 2025. 3
- [12] 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 *International Conference on Learning Representations (ICLR)*, 2023. 3, 6
- [13] Peng Gao, Shijie Geng, Renrui Zhang, Teli Ma, Rongyao Fang, Yongfeng Zhang, Hongsheng Li, and Yu Qiao. Clip-adapter: Better vision-language models with feature adapters. *International Journal of Computer Vision (IJCV)*, 132(2):581–595, 2024. 2, 3, 6, 7
- [14] Saurav Jha, Dong Gong, and Lina Yao. Clap4clip: Continual learning with probabilistic finetuning for vision-language models. *Advances in Neural Information Processing Systems (NeurIPS)*, 37:129146–129186, 2024. 3
- [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 (ICML)*, pages 4904–4916. PMLR, 2021. 3
- [16] Zhiqi Kang, Liyuan Wang, Xingxing Zhang, and Karteek Alahari. Advancing prompt-based methods for replay-independent general continual learning. *arXiv preprint arXiv:2503.00677*, 2025. 3
- [17] Muhammad Gul Zain Ali Khan, Muhammad Ferjad Naeem, Luc Van Gool, Didier Stricker, Federico Tombari, and Muhammad Zeshan Afzal. Introducing language guidance in prompt-based continual learning. In *International Conference on Computer Vision (ICCV)*, pages 11463–11473, 2023. 3
- [18] Muhammad Uzair Khattak, Hanoona Rasheed, Muhammad Maaz, Salman Khan, and Fahad Shahbaz Khan. Maple: Multi-modal prompt learning. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19113–19122, 2023. 2, 3, 6, 7
- [19] Nupur Kumari, Bingliang Zhang, Richard Zhang, Eli Shechtman, and Jun-Yan Zhu. Multi-concept customization of text-to-image diffusion. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1931–1941, 2023. 2, 3, 5
- [20] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International Conference on Machine Learning (ICML)*, pages 12888–12900. PMLR, 2022. 3
- [21] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International Conference on Machine Learning (ICML)*, pages 19730–19742. PMLR, 2023. 3
- [22] Yuqi Lin, Minghao Chen, Wenxiao Wang, Boxi Wu, Ke Li, Binbin Lin, Haifeng Liu, and Xiaoifei He. Clip is also an efficient segmenter: A text-driven approach for weakly supervised semantic segmentation. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15305–15314, 2023. 3
- [23] Yuning Lu, Jianzhuang Liu, Yonggang Zhang, Yajing Liu, and Xinmei Tian. Prompt distribution learning. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5206–5215, 2022. 3
- [24] Timo Lüddecke and Alexander Ecker. Image segmentation using text and image prompts. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7086–7096, 2022. 1, 3, 6
- [25] Aristeidis Panos, Rahaf Aljundi, Daniel Olmeda Reino, and Richard E Turner. Efficient few-shot continual learning in vision-language models. *arXiv preprint arXiv:2502.04098*, 2025. 3
- [26] Pramuditha Perera, Matthew Trager, Luca Zancato, Alessandro Achille, and Stefano Soatto. Discriminative-generative

- custom tokens for vision-language models. *arXiv preprint arXiv:2502.12095*, 2025. 3
- [27] 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 *International Conference on Machine Learning (ICML)*, pages 8748–8763. PmLR, 2021. 1, 3
- [28] 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 *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 22500–22510, 2023. 3
- [29] Liyana Sahir, Anwesha Banerjee, and Soma Biswas. Adaprompt: Prompt tuning with adaptive neighbours for generalized category discovery. In *International Conference on Image Processing (ICIP)*, pages 1133–1138. IEEE, 2024. 3
- [30] Jinwoo Seo, Soora Choi, Eungyeom Ha, Beomjune Kim, and Dongbin Na. New benchmarks for asian facial recognition tasks: Face classification with large foundation models. *arXiv preprint arXiv:2310.09756*, 2023. 2, 5
- [31] Vishal Thengane, Salman Khan, Munawar Hayat, and Fahad Khan. Clip model is an efficient continual learner. *arXiv preprint arXiv:2210.03114*, 2022. 3, 6
- [32] Michael Tschannen, Alexey Gritsenko, Xiao Wang, Muhammad Ferjad Naeem, Ibrahim Alabdulmohsin, Nikhil Parthasarathy, Talfan Evans, Lucas Beyer, Ye Xia, Basil Mustafa, et al. Siglip 2: Multilingual vision-language encoders with improved semantic understanding, localization, and dense features. *arXiv preprint arXiv:2502.14786*, 2025. 3
- [33] Kai Wang, Fei Yang, Bogdan Raducanu, and Joost van de Weijer. Multi-class textual-inversion secretly yields a semantic-agnostic classifier. In *Winter Conference on Applications of Computer Vision (WACV)*, pages 4400–4409. IEEE, 2025. 3
- [34] Zhaoqing Wang, Yu Lu, Qiang Li, Xunqiang Tao, Yandong Guo, Mingming Gong, and Tongliang Liu. Cris: Clip-driven referring image segmentation. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11686–11695, 2022. 3
- [35] Bin Xiao, Haiping Wu, Weijian Xu, Xiyang Dai, Houdong Hu, Yumao Lu, Michael Zeng, Ce Liu, and Lu Yuan. Florence-2: Advancing a unified representation for a variety of vision tasks. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4818–4829, 2024. 3
- [36] Zehao Xiao, Shilin Yan, Jack Hong, Jiayin Cai, Xiaolong Jiang, Yao Hu, Jiayi Shen, Qi Wang, and Cees GM Snoek. Dynaprompt: Dynamic test-time prompt tuning. In *International Conference on Learning Representations (ICLR)*, 2025. 3, 6, 7
- [37] Shuanglin Yan, Neng Dong, Liyan Zhang, and Jinhui Tang. Clip-driven fine-grained text-image person re-identification. *IEEE Transactions on Image Processing*, 32:6032–6046, 2023. 3
- [38] Hee Suk Yoon, Eunseop Yoon, Joshua Tian Jin Tee, Mark Hasegawa-Johnson, Yingzhen Li, and Chang D Yoo. C-tpt: Calibrated test-time prompt tuning for vision-language models via text feature dispersion. In *International Conference on Learning Representations (ICLR)*, 2024. 2, 3, 6, 7
- [39] Jiazuo Yu, Yunzhi Zhuge, Lu Zhang, Ping Hu, Dong Wang, Huchuan Lu, and You He. Boosting continual learning of vision-language models via mixture-of-experts adapters. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 23219–23230, 2024. 3, 6
- [40] Qihang Yu, Ju He, Xueqing Deng, Xiaohui Shen, and Liang-Chieh Chen. Convolutions die hard: Open-vocabulary segmentation with single frozen convolutional clip. *Advances in Neural Information Processing Systems (NeurIPS)*, 36:32215–32234, 2023. 3
- [41] Wenwen Yu, Yuliang Liu, Wei Hua, Deqiang Jiang, Bo Ren, and Xiang Bai. Turning a clip model into a scene text detector. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6978–6988, 2023. 3
- [42] Renrui Zhang, Wei Zhang, Rongyao Fang, Peng Gao, Kun-chang Li, Jifeng Dai, Yu Qiao, and Hongsheng Li. Tip-adapter: Training-free adaptation of clip for few-shot classification. In *European Conference on Computer Vision (ECCV)*, pages 493–510. Springer, 2022. 2, 3, 6, 7
- [43] Yuxin Zhang, Weiming Dong, Fan Tang, Nisha Huang, Haibin Huang, Chongyang Ma, Tong-Yee Lee, Oliver Deussen, and Changsheng Xu. Prospect: Prompt spectrum for attribute-aware personalization of diffusion models. *ACM Transactions on Graphics (TOG)*, 42(6):1–14, 2023. 3
- [44] Da-Wei Zhou, Zhi-Hong Qi, Han-Jia Ye, and De-Chuan Zhan. Tv100: a tv series dataset that pre-trained clip has not seen. *Frontiers of Computer Science*, 18(5):185349, 2024. 2, 5
- [45] Da-Wei Zhou, Kai-Wen Li, Jingyi Ning, Han-Jia Ye, Lijun Zhang, and De-Chuan Zhan. External knowledge injection for clip-based class-incremental learning. In *International Conference on Computer Vision (ICCV)*, 2025. 3, 6
- [46] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16816–16825, 2022. 2, 3, 6, 7
- [47] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *International Journal of Computer Vision (IJCV)*, 130(9):2337–2348, 2022. 2, 6, 7
- [48] Beier Zhu, Yulei Niu, Yucheng Han, Yue Wu, and Hanwang Zhang. Prompt-aligned gradient for prompt tuning. In *International Conference on Computer Vision (ICCV)*, pages 15659–15669, 2023. 3