

Towards Purified Multi-Label Test-Time Adaptation of Vision-Language Models

Yiwen Liang¹, Hui Chen^{1*}, Yizhe Xiong¹, Mengyao Lyu², Yuhan Cao³, Zijia Lin¹, Shuaicheng Niu², Sicheng Zhao¹, Jungong Han¹, and Guiguang Ding^{1*}

¹ Tsinghua University, Beijing, China

² Nanyang Technological University, Singapore

³ University of Washington, Seattle, WA, USA

{evenliang789, jichenhui2012}@gmail.com, dinggg@tsinghua.edu.cn

Abstract. Test-time adaptation (TTA) has been widely explored in single-label recognition, effectively mitigating distribution shifts, especially when combined with vision-language models. However, real-world images often contain multiple objects, while the more practical multi-label test-time adaptation (MLTTA) has received little attention so far. Recent cache-based TTA methods have shown promising efficiency and effectiveness, yet directly extending them to multi-label scenarios suffers from a one-to-many mapping problem: a shared global representation entangling co-occurring objects is stored as class-wise cache prototypes, inducing dominant-label bias and compromised cache calibration. While introducing region-level cues helps isolate class-specific evidence, such regional evidence can also be unreliable under distribution shifts, making its identification and utilization non-trivial. To address these issues, we introduce **PuRF**, a novel **PuRiF**ication-driven cache-based method for multi-label test-time adaptation of vision-language models. Specifically, PuRF first performs region purification to identify reliable regions, providing comprehensive regional cues for multi-label recognition and enabling fine-grained alignment. Based on these purified regions, PuRF conducts cache purification to enhance cache representation and adaptability, where episodic purification builds a discriminative region-based cache, and temporal refreshing further promotes long-term cache adaptability. Experiments demonstrate that PuRF consistently outperforms state-of-the-art methods, achieving a notable 4.05% mAP improvement on ViT-B/32 across five datasets. The code is available at: <https://github.com/EvelynIyWliang/PuRF-MLTTA>.

Keywords: Test-Time Adaptation · Vision-Language Models · Multi-Label Learning

1 Introduction

Vision-Language Models (VLMs) [20, 41] have exhibited remarkable zero-shot capability, successfully tackling a wide range of computer vision tasks [11, 26, 58, 64].

* Corresponding author.

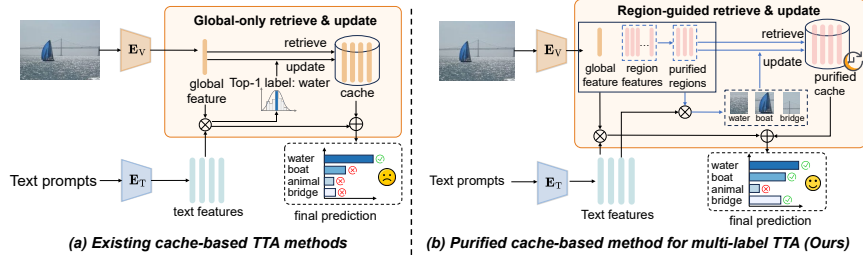


Fig. 1: Comparison between existing cache-based methods and ours. (a) Existing cache-based TTA methods [23, 29, 65] suffer from one-to-many representation coupling in multi-label TTA, where shared global features are dominated by salient classes (e.g., “water”), misleading class-wise cache calibration. (b) PuRF performs region purification and cache purification, enabling fine-grained and accurate multi-label adaptation.

However, these models often fail to generalize well when facing testing distributions that deviate from those seen during pre-training, known as distribution shift [34, 59, 60]. Such shifts commonly occur in real scenarios, such as images captured by different environments [25, 39] or sensing conditions [1, 16, 61].

Test-time adaptation (TTA) has emerged as a practical solution for handling distribution shifts by adapting models at inference using only unlabeled test data. Recent studies primarily leverage learnable prompts [13, 45, 66] for dynamic alignment and cache mechanisms [5, 18, 23, 29] for historical knowledge retrieval to enhance the adaptation capability of VLMs. However, most existing methods are developed for *single-label recognition*, whereas real-world images often contain multiple co-occurring objects, rendering *multi-label test-time adaptation* substantially more challenging and complex than single-label settings.

As the first work on multi-label TTA, ML-TTA [57] introduces a Bound Entropy Minimization objective with prompt optimization to enhance the confidence of top- k labels while mitigating over-suppression from naive entropy minimization. However, it still relies on intensive text retrieval and episodic prompt resetting, failing to exploit historical information and resulting in suboptimal performance. In contrast, cache-based methods [5, 29, 65, 69] have demonstrated strong performance in single-label TTA by accumulating and reusing historical information for prediction refinement via retrieval-based aggregation, but remain largely unexplored in multi-label settings, leaving their potential underexploited. When directly applied to multi-label images, these methods encounter a *one-vs-many dilemma* [22, 74], where a single spatially pooled global feature represents all co-occurring objects and simultaneously acts as a shared prototype for cache retrieval and storage. Such shared representations cannot effectively capture label-specific features [22] and tend to be dominated by salient objects [9, 19, 42], leading to biased predictions and weakened label-specific discriminability. Consequently, the cache fails to perform reliable class-specific calibration, thereby hampering adaptation performance, as illustrated in Fig. 1(a). Although prior works [36, 38, 74] incorporate region-level cues to improve class-specific representation in multi-label recognition, obtaining reliable region evidence and lever-

aging it in fully unsupervised test-time settings remains challenging, as naive region proposals are often noisy and include irrelevant context [47, 51, 75].

After revealing the above limitations, we propose **PuRF**, a **PuRiF**ication-driven cache-based method for multi-label Test-Time Adaptation, which builds upon the cache mechanism for its effectiveness with historical information while extending this paradigm to more practical multi-label scenario. Unlike ML-TTA [57] that tailors optimization objectives for multi-label settings, our key idea is to exploit purified regional evidence to enable finer-grained multi-label recognition and construct a more discriminative cache to better utilize historical information. To this end, PuRF proposes two main components that collectively enhance multi-label adaptation: 1) Region Purification (RP) via multi-granularity consistency identifies purified regions via relative activation-based selection and derives reliable pseudo supervision via global-local consistency, enabling fine-grained alignment and effective semantic optimization for accurate multi-label test-time recognition; (2) Cache Purification (CP) integrates episodic purification to build class-specific prototypes from purified high-confidence regions for more discriminative representations, and temporal refreshing to improve long-term cache adaptability under evolving distributions via time-aware entropy weighting. In Fig. 1(b), our purified regional evidence and purified cache improve class-specific discriminability for better multi-label recognition. Overall, PuRF advances multi-label test-time adaptation by extending the cache-based paradigm to multi-label settings with two novel purification strategies. Experimental results show that PuRF exhibits strong generalization across diverse backbones and prompt settings while maintaining high efficiency.

We summarize our contributions as follows:

- We propose PuRF, a novel purification-driven cache-based method tailored for multi-label test-time adaptation, enabling fine-grained recognition while enhancing label-specific cache calibration.
- PuRF integrates Region Purification to identify informative regions and leverage reliable and comprehensive regional cues for fine-grained alignment, Cache Purification for constructing discriminative class-specific prototypes from purified regions for reliable calibration, and temporal refreshing for improving long-term cache adaptability under streaming data.
- Extensive experiments across five standard multi-label benchmarks and four backbones show that PuRF consistently outperforms the strongest state-of-the-art method, demonstrating its superior effectiveness.

2 Related Work

Vision-Language Test-Time Adaptation. Test-time adaptation (TTA) improves model generalization under distribution shifts, and the rise of vision-language models (VLMs) further extends it to more challenging open-world shifts. Existing TTA methods for VLMs can be broadly categorized into prompt-based and cache-based methods. Prompt-based methods [32, 45, 63, 67] pioneer this field by optimizing textual prompts via marginal entropy minimization, but




Fig. 2: The overall framework of our proposed PuRF. We first adaptively identify purified regions for prediction aggregation and then derive reliable pseudo supervision by enforcing global-local prediction consistency. Guided by the pseudo labels, confident region-text pairs are stored as class-specific region prototypes for episodic cache purification, improving class-wise discriminability. Temporal refreshing further purifies the cache by incorporating newly reliable samples while maintaining long-term adaptivity of cache. $\mathcal{H}$ denotes the prediction entropy.

suffer from limited efficiency due to per-sample prompt resetting and heavy back-propagation. Cache-based methods [5, 18, 23, 70] offer a more efficient alternative by constructing a dynamic cache of test features and refining predictions through retrieval-based similarity aggregation. A line of work [23, 69, 70] is training-free and focuses on the design of multiple caches, where different caches serve distinct roles to enhance adaptation. Another line [29, 65] introduces lightweight prototype evolution through residual updates for better cross-modal alignment. Despite advancements, most current methods mainly focus on single-label classification. Motivated by this gap, we extend the cache-based paradigm to multi-label adaptation by addressing challenges in purified regional cue selection and utilization, while improving discriminative class-wise cache calibration.

Multi-Label Recognition with VLMs. Multi-label recognition (MLR) serves as a fundamental task in computer vision, where the key challenges lie in modeling label dependencies [8, 50] and identifying discriminative properties for each label [21, 35]. Earlier methods rely solely on visual features and capture label correlations via graph structures [44, 62], recurrent networks [53, 55], or attention mechanisms [43]. Recently, the emergence of VLMs has introduced a new paradigm for MLR, enabling strong zero-shot and open-vocabulary capabilities, particularly when combined with parameter-efficient tuning techniques such as prompt learning. Methods like DualCoOp [46], TaI-DPT [15], and PVP [56] introduce learnable prompts to capture label semantics, effectively modeling label co-occurrence with minimal additional parameters. However, they still rely on annotated samples or offline fine-tuning on batches for stable adaptation. In con-

trast, our work targets a fully unlabeled setting with online adaptation for each incoming sample, which remains challenging for precise multi-label recognition.

3 Methodology

The overall pipeline of PuRF is shown in Fig. 2. PuRF enhances multi-label test-time adaptation through purification at the region and cache levels. Sec. 3.2 proposes Region Purification (RP) via Multi-granularity Consistency, which selects informative regions using adaptive relative activation to provide comprehensive visual cues, and exploits global-local prediction consistency to derive reliable pseudo supervision for semantic residual optimization and fine-grained alignment. Sec. 3.3 presents Cache Purification (CP), comprising Episodic Purification (EP) that constructs region-aware class-specific prototypes from purified regions for stronger discriminability, and Temporal Refreshing (TR) that further enhances long-term cache adaptivity. Finally, Sec. 3.4 establishes a unified objective that enables joint optimization and feature alignment.

3.1 Preliminaries

Cache-based TTA for VLMs. Cache-based TTA methods [23, 65, 69] build on the key-value design of [68] that store low-entropy samples and continually update during inference. For image x and a class prompt $\mathcal{T}^c$, given the CLIP visual feature $\mathbf{F}_V = \mathbf{E}_V(x)$ and textual feature $\mathbf{F}_T^c = \mathbf{E}_T(\mathcal{T}^c)$ extracted by its visual and text encoders $\mathbf{E}_V$ and $\mathbf{E}_T$. The prediction logits and probabilities are obtained from the cosine similarities $\langle \cdot, \cdot \rangle$:

$$s_i^c = \langle \mathbf{F}_V, \mathbf{F}_T^c \rangle \quad (1)$$

The cache $\mathcal{M} = \{\mathcal{M}^c\}_{c=1}^C$ is composed of C class-wise subsets, each with a capacity of L entries containing image features and their entropy values. For test-time samples arriving on the fly, new items enter the corresponding cache slots based on their one-hot pseudo-labels. Once the cache reaches capacity, items with the highest entropy are replaced. At inference, test feature $\mathbf{F}_V$ retrieves relevant cache items, and the prediction combines CLIP logits with cache logits:

$$s_{\text{TTA}}^c = s_i^c + s_{\text{cache}}^c = s_i^c + \mathcal{A}(\langle \mathbf{F}_V, \mathbf{F}_{\text{cache}}^c \rangle) \mathbf{L}_p \quad (2)$$

where $\mathbf{F}_{\text{cache}}^c$ denotes cached prototypes of class c , and $\mathbf{L}_p$ is the one-hot pseudo-label from CLIP prediction. $\mathcal{A}(x) = \alpha e^{-\beta(1-x)}$ is a modulating function with amplitude α and sharpness β . Akin to [65] and [29], $\mathbf{F}_{\text{cache}}^c$ is obtained by averaging features in $\mathcal{M}^c$ for compact visual prototypes.

Some cache-based methods [5, 29, 65] introduce lightweight residual learning with entropy-driven objectives to further enhance adaptation without heavy computation. Recently, ReTA [29] proposes adjacent textual embeddings that extend each class prototype into multiple semantic experts, with learnable residuals for distribution calibration, enabling more fine-grained contextual semantics. Following this design, we implement textual residual learning as:

$$\hat{\mathbf{F}}_T^c = \mathbf{F}_T^c + \Delta \mathbf{F}_T^c \quad (3)$$

where $\Delta \mathbf{F}_T^c$ denotes the learnable residual and the original prototype is frozen. We incorporate adjacent text embeddings as in [29] to enhance semantic diversity, averaging them to form the textual prototype. Additionally, we also progressively update the shared semantic embedding via running averaging over streaming test samples, following [29, 65]:

$$\hat{\mathbf{F}}_T^c = \text{Norm}\left((l-1)\hat{\mathbf{F}}_T^c + \hat{\mathbf{F}}_T^{c(*)}\right) \quad (4)$$

where $\text{Norm}(\mathbf{x}) = \mathbf{x}/\|\mathbf{x}\|$ denotes ℓ_2 normalization, and l is the number of accumulated updates. However, we do not adopt the entropy reweighting strategy, as it is tailored for single-label consistency.

3.2 Region Purification via Multi-granularity Consistency

Conventional TTA methods [23, 52, 70] typically rely on a global representation. However, in multi-label scenarios, such representations entangle co-occurring objects and are dominated by salient labels occupying large space [22, 42], making them insufficient for modeling fine-grained label-specific features. An intuitive approach is to incorporate region-level cues via region cropping [14, 31, 54] or localization [2, 49], which can isolate object-related evidence and benefit multi-label recognition. However, region candidates often contain redundancy and noise, especially in unsupervised TTA settings. Naively aggregating all regions may dilute discriminative signals and introduce irrelevant context [47, 75], thereby limiting recognition performance. This motivates us to design a region purification mechanism to select informative regions and facilitate fine-grained alignment to optimize semantic representations.

Adaptive Region Purification. To support reliable regional evidence, we leverage both global and local cues. At the global level, we apply holistic augmentations (*e.g.*, flipping and color jitter) to form a multi-view set $\mathcal{V}$ following [23, 69]. Each augmented view x_i yields predictions $p(x_i)$, which are used to derive pseudo-label candidates to guide the purification and semantic alignment:

$$\mathcal{G} = \bigcap_{x_i \in \mathcal{V}} \{c \mid c \in \text{Top-}\kappa_g(p(x_i))\}. \quad (5)$$

where $\text{Top-}\kappa_g(\cdot)$ identifies the indices of the highest-confidence classes, with the threshold κ_g defined as a fixed proportion of the total class count $\mathcal{C}$.

At the local level, we extract Q local region features $\mathcal{F}_V^{\text{local}} = \{\mathbf{f}_V^{(1)}, \mathbf{f}_V^{(2)}, \dots, \mathbf{f}_V^{(Q)}\}$ by performing stochastic scaling and cropping on the input image x under predefined scale ranges, following [3, 54]. We aim to identify regions that provide more isolated and less entangled label-specific evidence, which helps compensate for missing spatial cues and improve multi-label discrimination [4], particularly for small or spatially localized objects [73]. Intuitively, a purified region should exhibit clear dominance of one class over the others. To quantify this dominance, we compute the relative confidence of each region via normalization across classes:

$$\hat{s}_q^c = \text{softmax}_c\left(\left\langle \mathbf{f}_V^{(q)}, \mathbf{F}_T^c \right\rangle\right). \quad (6)$$

We identify the dominant class for each region as $c^* = \arg \max_{c \in \mathcal{C}} \hat{s}_q^c$ and use the normalized score $\hat{s}_q^{c^*}$ as the purification confidence. Since confidence varies across classes due to class imbalance [28], a uniform global threshold becomes suboptimal for region selection. We therefore maintain an adaptive class-wise threshold μ_c via running averaging of incoming region confidence scores:

$$\mu_c \leftarrow \mu_c + \frac{1}{T}(\bar{s}_c - \mu_c), \quad \bar{s}_c = \frac{1}{Q} \sum_{q=1}^Q \hat{s}_q^c. \quad (7)$$

where T is the number of test steps. Regions whose dominant-class confidence surpasses the class-wise adaptive threshold are retained, forming the purified region feature set that provides reliable class-specific evidence:

$$\mathcal{F}_V^{\text{pur}} = \left\{ \mathbf{f}_V^{(q)} \mid \hat{s}_q^{c^*} \geq \mu_{c^*}, q \in [1, Q] \right\}. \quad (8)$$

Region-level pseudo-labels are obtained by aggregating the top-1 predictions from purified regions, serving as local counterparts to the global pseudo-labels:

$$\mathcal{R} = \left\{ c^* \mid c^* = \arg \max_{c \in \mathcal{C}} \langle \mathbf{f}_V^{(q)}, \mathbf{F}_T^c \rangle, \mathbf{f}_V^{(q)} \in \mathcal{F}_V^{\text{pur}} \right\}. \quad (9)$$

Thus, purified regions provide reliable local supervision for semantic optimization, while the final pseudo supervision is obtained via intersection: $\hat{\mathcal{Y}} = \mathcal{G} \cap \mathcal{R}$. In Supplementary Sec. 3, we provide a theoretical analysis that formally demonstrates the effectiveness of the proposed region purification mechanism.

Global-Local Aggregation and Alignment. We employ global-local aggregation to capture complementary information by extending the prediction in Eq. (2) to balance global and local contributions:

$$s_{\text{TTA}}^c = \underbrace{\frac{1}{2} \langle \mathbf{F}_V, \mathbf{F}_T^c \rangle}_{\text{Global}} + \underbrace{\frac{1}{2} \max_{\mathbf{f}_V^{(q)} \in \mathcal{F}_V^{\text{pur}}} \langle \mathbf{f}_V^{(q)}, \mathbf{F}_T^c \rangle}_{\text{Local}} + s_{\text{cache}}^c \quad (10)$$

where $\max(\cdot)$ preserves the strongest class-specific regional activation as [17, 54] to suppress noise. This aggregation integrates purified regional evidence for more fine-grained recognition. With reliable pseudo-labels derived above, we adopt the standard BCE loss for multi-label optimization:

$$\mathcal{L}_{\text{BCE}} = - \sum_{c=1}^{\mathcal{C}} [\tilde{y}^c \log(\sigma(s_{\text{TTA}}^c)) + (1 - \tilde{y}^c) \log(1 - \sigma(s_{\text{TTA}}^c))] \quad (11)$$

where $\sigma(\cdot)$ denotes the Sigmoid function. Overall, the proposed aggregation and alignment leverage complementary visual cues from purified regions for more precise adaptation, while enabling mutual refinement between region purification and semantic optimization.

3.3 Cache Purification with Episodic and Temporal Strategies

Cache-based methods typically rely on global features for retrieval and storage. In multi-label scenarios, these features are often dominated by salient classes and lack label purity, leading to cache contamination and degraded class-wise calibration. To alleviate this issue, we perform cache purification at two complementary




Fig. 3: Empirical studies of cache dynamics on COCO2014. (a): The cache turnover rate measures the frequency of cache updates, where $N_{\text{in}}^{(t)}$ and $N_{\text{out}}^{(t)}$ denote the numbers of samples inserted and removed at adaptation step t , respectively. (b): The evolution of model performance during adaptation.

timescales: *short-term* episodic purification and *long-term* temporal refreshing. The former ensures reliable cache entries for each incoming sample, while the latter enhances adaptation over time by refreshing outdated cache entries.

Episodic Purification. For each incoming test sample (episode), we obtain its purified regions as defined in Eq. (8). Given the reliable pseudo-label set $\tilde{\mathcal{Y}}$, we select the corresponding regions to form region-label cache entries. Specifically, for each label $\tilde{y}_i \in \tilde{\mathcal{Y}}$, we select the most confident region in $\mathcal{F}_V^{\text{pur}}$ as anchor:

$$r^* = \arg \min_{\mathbf{f}_V^{(q)} \in \mathcal{F}_V^{\text{pur}}} \mathcal{H} \left(\left\langle \mathbf{f}_V^{(q)}, \mathbf{F}_T^{\tilde{y}_i} \right\rangle \right) \quad (12)$$

where region r^* denotes the anchor region index for pseudo-label $\tilde{y}_i$. Each anchor feature $\mathbf{f}_V^{(r^*)}$ is paired with its entropy to form region-entropy cache candidates. This design yields cleaner cache representations and improves class-wise discriminability, as empirically demonstrated in Fig. 5. The cache score is computed on the same purified region set as Eq. (10):

$$s_{\text{cache}}^c = \mathcal{A} \left(\max_{\mathbf{f}_V^{(q)} \in \mathcal{F}_V^{\text{pur}}} \left\langle \mathbf{f}_V^{(q)}, \mathbf{F}_{\text{cache}}^c \right\rangle \right) \mathbf{L}_p. \quad (13)$$

Unlike typical entropy-based selection that stores a single cache entry per sample [23, 65], our episodic purification exploits purified regions to form region-label pairs, enabling multiple class-specific cache entries within each episode. Cache logits are computed through region-level max aggregation over the purified regions, improving prediction robustness and class-wise discriminability.

Temporal Refreshing. As mentioned above, each multi-label image could produce multiple cache pairs, which accelerates cache saturation under limited capacity compared with single-label images that typically provide only one candidate. As adaptation proceeds, the entropy-based admission threshold becomes increasingly stringent, making it difficult for later samples to enter the cache. Consequently, the cache is rarely updated in the later stages and may eventually solidify into a fixed state. We refer to this phenomenon as *cache saturation*, where overly strict admission suppresses the inclusion of later samples, causing the cache to be dominated by early entries and gradually lose temporal dynamics — its ability to update with samples appearing later in the data stream. As

shown in Fig. 3, the blue curve denotes the vanilla cache without any adjustment. Its turnover rate first increases and then declines, while the mAP curve plateaus later, indicating that the cache becomes progressively less adaptive over time. As the cache becomes nearly fixed, it struggles to incorporate later samples, thereby overlooking their contribution. Meanwhile, the progressive prototype update in Eq. (4) relies on accumulating semantic evidence over time. When later samples fail to enter the cache, it cannot benefit from the progressively refined semantic prototypes for reliable entry selection, ultimately limiting long-term adaptation.

To maintain cache purification and enhance long-term adaptability, we design a Temporal Refreshing mechanism with a time-decay strategy to adjust the priority of cached entries, mitigating early-stage bias that causes early samples to dominate the cache and enabling the integration of later ones. To be specific, we additionally record for each cached item its cumulative time in the cache, denoted as t_i , and each class-wise cache $\mathcal{M}^c$ consists of L feature entries represented as:

$$\mathcal{M}^c = \{(\mathbf{F}_{\text{cache}}, \mathcal{H}(\langle \mathbf{F}_{\text{cache}}, \mathbf{F}_T^c \rangle), t_i)\}_{i=1}^L \quad (14)$$

Inspired by the exponential decay strategy in learning rate scheduling [27, 33], we formulate a similar temporal decay strategy to gradually downweight the priority of long-retained cache entries by penalizing their entropy, mitigating early-sample dominance. Let δ denote a temporal decay constant that controls the retention time, beyond which the penalization progressively increases. The temporal weight and updated entropy are defined as:

$$w_i^{\text{time}} = \exp\left(\frac{t_i - \delta}{\delta}\right), \quad \mathcal{H}'_i = \mathcal{H}_i \cdot w_i^{\text{time}}. \quad (15)$$

The updated entropy $\mathcal{H}'_i$ replaces the original one in Eq. (14) to penalize long-retained entries. When $t_i > \delta$, the weight $w_i^{\text{time}} > 1$ amplifies their entropy, lowering priority and encouraging replacement by newer samples. This refreshing strategy enables the cache to evolve with streaming data, allowing later samples to leverage refined semantic representations to progressively improve class-specific representations, thereby enhancing overall adaptation (see Fig. 3).

3.4 Overall Objectives

Our optimization goal is to achieve accurate recognition in multi-label test-time adaptation, which is challenging for conventional TTA methods that mainly rely on vanilla entropy minimization. Based on the reliable pseudo-labels obtained from global-local consistency (Sec. 3.2), we incorporate the binary cross-entropy loss in Eq. (11) to provide additional supervision and stabilize optimization under distribution shifts. Following [29, 65], we adopt the standard entropy loss and an alignment loss as auxiliary objectives to optimize textual residuals and calibrate predictions. Given test sample x and its augmented set $\mathcal{V} = \{x_1, \dots, x_{N-1}\}$, the marginal entropy loss can be written as:

$$\mathcal{L}_{\text{ent}}(x) = \frac{1}{|\mathcal{V}_f|} \sum_{x_i \in \mathcal{V}_f} \mathcal{H}(p(x_i)) \quad (16)$$

where $\mathcal{H}(p) = -\sum_{c=1}^C p^c \log p^c$ is the prediction entropy, and $\mathcal{V}_f \subset \mathcal{V} \cup \{x\}$ denotes the subset of the most confident augmented views (typically the top 10%). Following [29, 65], we introduce an alignment loss to enforce cross-modal consistency between cached visual prototypes and text, formulated as:

$$\mathcal{L}_{\text{align}} = \sum_{c=1}^C \left[-\log \frac{\exp((\hat{\mathbf{F}}_T^c)^\top \mathbf{F}_{\text{cache}}^c)}{\sum_{j=1}^C \exp((\hat{\mathbf{F}}_T^c)^\top \mathbf{F}_{\text{cache}}^j)} \right] \quad (17)$$

where $\hat{\mathbf{F}}_T^c$ and $\mathbf{F}_{\text{cache}}^c$ denote the mean prototypes of the textual and visual prototypes, respectively. In summary, our final loss function is defined as:

$$\mathcal{L}^* = \mathcal{L}_{\text{ent}} + \lambda_1 \mathcal{L}_{\text{BCE}} + \lambda_2 \mathcal{L}_{\text{align}} \quad (18)$$

where λ_1 and λ_2 are trade-off coefficients balancing contributions of pseudo-label supervision and cross-modal alignment. The integration of multi-label supervision with test-time regularization via residual updates enables PuRF to capture label co-occurrence knowledge and mitigate test-time uncertainty.

4 Experiments

4.1 Experimental Settings

Datasets. Following [57], we evaluate on five widely used multi-label datasets: VOC [10] (versions 2007 and 2012), MS-COCO [30] (versions 2014 and 2017), and NUS-WIDE [6]. For the VOC datasets, the 2007 and 2012 versions contain 20 object categories with 4.9K and 5.8K test images, respectively. MS-COCO covers 80 object categories, and we use the validation sets of COCO2014/COCO2017 containing 40K/5K testing images. For NUS-WIDE, we evaluate on 81 manually annotated categories covering broader concepts, with 83K test images.

Implementation Details. We verify the generalization of our PuRF across diverse CLIP architectures (RN50, RN101, ViT-B/16, ViT-B/32) using the default prompt “*a photo of a*” and report results in mAP. We further assess its compatibility with different prompt initializations, including (1) pretrained prompt weights from MaPLe [24] and CoOp [72], and (2) dataset-specific templates in CuPL [40] manner. For all datasets, we set the number of generated regions to 50 and fix the temporal decay constant as $\delta = 1000$. For global filtering, we select the top- κ_g labels with $\kappa_g = 0.1$. We use 63 augmented visual views, consistent with common practice in prior baselines [23, 45, 65, 69]. The loss coefficients are set to $\lambda_1=0.2$ and $\lambda_2=0.5$. The cache size $L=3$ with three adjacent text embeddings follows [29, 65]. Please refer to Supplementary for more details.

4.2 Comparison with State-of-the-Art Methods

Different CLIP Visual Backbones. As shown in Table 1, PuRF consistently shows superior performance across all datasets and backbones. Notably, on ViT-B/32, PuRF improves mAP by 8.56% over its prompt-based counterpart ML-TTA [57] and 4.05% over ReTA [29] on average across datasets, demonstrating

Table 1: Comparison with state-of-the-art methods under different visual encoder architectures *w.r.t.* mAP (%). Bold indicates best results.

Methods	ResNet-50						ResNet-101					
	VOC07	VOC12	COCO14	COCO17	NUS	Avg.	VOC07	VOC12	COCO14	COCO17	NUS	Avg.
CLIP [41] _{ICML'21}	75.91	74.25	47.53	47.32	41.53	57.31	76.72	74.21	48.83	48.15	41.93	57.97
TPT [45] _{NeurIPS'22}	75.54	73.92	48.52	48.51	41.97	57.69	74.82	73.39	49.71	48.89	43.10	57.98
DiffTPT [13] _{CVPR'23}	75.89	74.13	48.56	48.67	41.33	57.72	74.98	74.31	49.45	49.19	42.93	58.17
RLCF [71] _{ICLR'24}	65.75	64.73	36.87	36.73	29.83	46.78	71.21	69.63	40.53	39.79	31.77	50.59
ML-TTA [57] _{ICLR'25}	78.62	76.63	51.58	51.39	42.53	60.15	78.72	78.13	52.92	52.24	43.62	61.13
TDA [23] _{CVPR'24}	76.64	75.12	48.91	49.11	42.34	58.42	78.12	77.13	50.19	49.78	43.13	59.67
DMN [70] _{CVPR'24}	74.87	74.13	44.54	44.18	41.32	55.81	76.82	75.32	46.28	45.44	42.71	57.31
DPE [65] _{NeurIPS'24}	82.51	80.64	54.38	54.74	45.96	63.65	82.55	81.21	54.92	54.53	46.64	63.97
BoostAdapter [69] _{NeurIPS'24}	79.53	78.00	52.07	52.92	44.20	61.34	81.13	80.06	53.90	54.33	45.47	62.98
ReTA [29] _{ACM MM'25}	83.56	81.98	55.09	55.40	46.90	64.59	83.69	82.40	56.71	56.39	47.02	65.24
PuRF (Ours)	86.35	84.38	61.47	61.56	47.57	68.27	86.77	85.43	62.60	61.93	47.91	68.93

Methods	ViT-B/16						ViT-B/32					
	VOC07	VOC12	COCO14	COCO17	NUS	Avg.	VOC07	VOC12	COCO14	COCO17	NUS	Avg.
CLIP [41] _{ICML'21}	79.58	79.25	54.42	54.13	45.65	62.61	77.18	76.85	50.31	50.15	42.90	59.48
TPT [45] _{NeurIPS'22}	77.54	77.39	53.32	54.20	46.15	61.72	74.21	71.93	48.12	48.63	43.63	57.30
DiffTPT [13] _{CVPR'23}	77.93	77.24	53.91	54.15	46.13	61.87	74.50	72.98	48.73	49.19	43.42	57.76
RLCF [71] _{ICLR'24}	79.29	79.26	54.21	54.43	43.18	62.07	77.12	76.83	50.28	49.59	43.29	59.42
ML-TTA [57] _{ICLR'25}	81.28	81.13	57.52	57.49	46.55	64.80	78.70	77.97	52.83	52.99	44.12	61.32
TDA [23] _{CVPR'24}	80.12	79.92	55.21	55.46	46.72	63.49	77.62	77.12	51.23	51.49	44.13	60.32
DMN [70] _{CVPR'24}	79.83	79.67	52.52	52.37	46.27	62.13	77.42	76.60	49.32	48.13	43.41	58.98
DPE [65] _{NeurIPS'24}	84.10	83.62	59.21	59.28	48.90	67.02	82.38	81.74	55.79	55.32	47.11	64.47
BoostAdapter [69] _{NeurIPS'24}	83.01	82.56	58.07	58.64	47.76	66.01	80.82	80.03	54.49	54.75	46.14	63.25
ReTA [29] _{ACM MM'25}	85.09	84.49	60.37	60.30	49.20	67.89	83.99	83.51	57.02	56.58	48.05	65.83
PuRF (Ours)	88.18	87.12	66.05	65.33	50.72	71.48	87.32	86.04	63.42	63.02	49.61	69.88

strong generalization under multi-label setting. By incorporating purified informative regions with multi-granularity aggregation and alignment, PuRF extracts reliable and complementary cues for more accurate multi-label recognition. In addition, the purified cache enables accurate class-wise calibration and maintains long-term cache adaptability under evolving distributions. Consequently, PuRF achieves a notable 6.09% mAP gain over ReTA on the large-scale COCO2014 dataset and consistent improvements on NUS-WIDE across backbones.

Different Prompt Initialization. In Table 2, PuRF still achieves the best performance under all prompt initialization methods. With MaPLe [24] initialization, PuRF achieves 73.79% mAP, outperforming ReTA (70.88%) and ML-TTA (70.38%), highlighting its strong adaptability to various prompt initializations. Templates constructed via CuPL [40] offer richer semantics that benefit cache-based methods, from which PuRF further benefits by exploiting purified regional cues for more reliable class-wise predictions, surpassing ReTA by 2.28% mAP. While prompt-based methods are incompatible with such template design and thus fail to effectively exploit its semantic richness.

Efficiency Analysis. Table 3 compares the efficiency and performance of different baselines on a single NVIDIA 3090 GPU. Prompt-based methods (*e.g.*, [45, 57]) exhibit high latency and memory due to repeated prompt re-initialization and optimization. DPE requires expensive visual prototype learning that leads to high memory consumption. In contrast, our PuRF adopts lightweight text residual learning with tolerable region-level computation, achieving a remarkable 8.87% mAP gain with low memory cost.

Table 2: Comparison with state-of-the-art methods under different prompt initializations on ViT-B/16 *w.r.t.* mAP (%). Bold indicates best results.

Methods	VOC2007	VOC2012	COCO2014	COCO2017	NUS-WIDE	Avg.
<i>CoOp</i>						
CoOp [72] _{IJCV'22}	79.14	77.85	56.12	56.35	46.74	63.24
TPT [45] _{NeurIPS'22}	79.72	77.85	55.35	55.23	47.27	63.08
DiffTPT [13] _{CVPR'23}	79.86	77.61	55.30	55.47	47.13	63.07
RLCF [71] _{ICLR'24}	80.15	78.24	56.72	56.18	47.62	63.78
ML-TTA [57] _{ICLR'25}	83.17	81.36	59.68	59.33	48.12	66.33
TDA [23] _{CVPR'24}	80.20	78.58	56.93	57.15	47.82	64.13
DPE [65] _{NeurIPS'24}	84.97	83.86	61.39	61.04	48.49	67.95
BoostAdapter [69] _{NeurIPS'24}	82.32	81.71	58.77	58.02	47.56	65.68
ReTA [29] _{ACM MM'25}	85.23	84.15	61.63	61.31	48.75	68.21
PuRF (Ours)	88.31	86.97	66.69	66.02	50.84	71.77
<i>MaPLe</i>						
MaPLe [24] _{CVPR'23}	85.34	84.79	62.18	62.35	48.42	68.62
TPT [45] _{NeurIPS'22}	85.04	83.92	63.36	63.75	48.90	69.01
DiffTPT [13] _{CVPR'23}	85.15	83.78	62.93	63.14	48.81	68.76
RLCF [71] _{ICLR'24}	85.35	85.28	62.84	62.90	49.37	69.15
ML-TTA [57] _{ICLR'25}	86.40	85.69	64.75	64.86	50.21	70.38
TDA [23] _{CVPR'24}	85.76	84.15	63.25	63.64	49.55	69.27
DPE [65] _{NeurIPS'24}	87.55	86.73	64.80	64.38	52.39	71.17
BoostAdapter [69] _{NeurIPS'24}	85.99	84.62	63.74	63.62	50.16	69.63
ReTA [29] _{ACM MM'25}	87.40	86.60	64.59	64.21	51.62	70.88
PuRF (Ours)	89.42	88.56	69.31	68.82	52.84	73.79
<i>Template</i>						
CLIP [41] _{ICML'21}	82.93	82.26	59.94	59.62	47.76	66.50
TDA [23] _{CVPR'24}	83.18	82.09	60.06	60.05	47.95	66.67
DPE [65] _{NeurIPS'24}	85.86	84.82	61.99	61.54	49.29	68.70
BoostAdapter [69] _{NeurIPS'24}	84.85	84.06	61.82	61.79	48.64	68.23
ReTA [29] _{ACM MM'25}	87.71	86.54	63.43	62.56	48.41	69.73
PuRF (Ours)	88.73	87.75	66.94	66.16	50.47	72.01

Table 3: Efficiency and performance over five multi-label datasets on a single NVIDIA 3090 GPU (ViT-B/16).

Method	Inference Speed (FPS) $\uparrow$	Memory (GB) $\downarrow$	mAP (%) $\uparrow$	Gain (%) $\uparrow$
CLIP [41]	106.25	0.35	62.61	-
TPT [45]	3.59	5.77	61.72	-0.89
ML-TTA [57]	3.05	5.77	64.80	2.19
DPE [65]	4.78	7.49	67.02	4.41
ReTA [29]	6.35	0.74	67.89	5.28
PuRF (Ours)	5.74	1.04	71.48	8.87

Generalization across Diverse VLM Backbones. To validate the generalization of PuRF beyond the original CLIP, we further evaluate it on three advanced VLM backbones, including EVA-02 [12] (ViT-B/16), SigLIP2 [48] (SO400M/14), and MetaCLIP2 [7] (ViT-H/14). As shown in Table 4, PuRF consistently achieves the best performance across all datasets compared with

Table 4: Comparison with state-of-the-art cache-based methods across diverse VLM backbones. Bold indicates best results.

	Method	VOC2007	VOC2012	COCO2017	Avg.
EVA-02 [12] (ViT-B/16)	No Adapt.	84.05	83.88	57.88	75.27
	DPE [65]	87.19	86.49	61.22	78.30
	ReTA [29]	86.76	86.23	61.15	78.05
	PuRF (Ours)	89.65	88.43	66.59	81.56
SigLIP2 [48] (SO400M/14)	No Adapt.	88.35	88.09	67.32	81.25
	DPE [65]	89.40	88.98	68.65	82.34
	ReTA [29]	88.84	88.47	67.92	81.74
	PuRF (Ours)	91.13	90.34	72.81	84.76
MetaCLIP2 [7] (ViT-H/14)	No Adapt.	84.16	83.91	61.16	76.41
	DPE [65]	85.64	84.57	62.87	77.69
	ReTA [29]	86.30	85.22	63.41	78.31
	PuRF (Ours)	89.59	88.74	70.29	82.87

Table 5: Ablation of key components on ViT-B/16.

	RP (region)	EP (cache)	TR (cache)	VOC2007	COCO2017
1				84.83	60.71
2	✓			86.67	63.35
3		✓		85.79	62.59
4	✓	✓		87.50	64.56
5	✓		✓	87.36	64.12
6		✓	✓	86.65	63.37
7	✓	✓	✓	88.18	65.35

Table 6: Ablation on Region Purification (COCO2017, ViT-B/16).

Region Selection	mAP (%)
Global-only	61.67
All regions	63.46
Abs-score selection	65.02
Rel-score selection (ours)	65.35

Table 7: Ablation on Cache Purification (COCO2017, ViT-B/16).

Cache setting	mAP (%)
Global-only cache	62.21
PuRF (global retrieve)	64.80
PuRF (min-ent) + avg	64.92
PuRF (min-ent) + max (ours)	65.35

two strong cache-based baselines, yielding average improvements of 3.56% mAP over DPE [65] and 3.69% mAP over ReTA [29]. These results demonstrate the effectiveness and generalization of PuRF across multiple VLM architectures.

4.3 Ablation Studies

Effectiveness of the Key Components. In Table 5, we analyze the effectiveness of each module in PuRF, including Region Purification (RP), Episodic Purification (EP) and Temporal Refreshing (TR) for cache purification. RP exploits informative regions through purification, providing fine-grained cues that enable reliable pseudo supervision via global-local consistency for semantic op-




Fig. 4: Impact of hyper-parameters. (a) Variation of κ_g in Eq. (5). (b) Variation of δ in Eq. (15). (c) Joint impact of λ_1 and λ_2 in Eq. (18) on COCO2017.




Fig. 5: t-SNE visualization of cached visual features on COCO2017. The capacity of all caches is set to 7, and the visualization is taken at the 2000-th adaptation step.

timization, boosting performance by 2.24% on average across two datasets. EP improves cache purity by constructing class-specific prototypes from purified regions, enabling more effective cache calibration. It yields a larger gain on COCO2017 (+1.88% mAP), where dense label co-occurrence requires finer class-wise discrimination. TR offers comparable improvements to EP by enhancing long-term cache adaptability and mitigating cache saturation and early-stage bias. Together, these components form a complementary design that balances performance and temporal adaptivity in multi-label test-time adaptation.

Impact of Region and Cache Purification. Table 6 and Table 7 evaluate the impact of purification at region and cache levels. In Eq. (8), relative activation-based selection measures the competitiveness of a region *w.r.t.* other class responses, favoring regions with strong class-specific dominance and yielding purer regional evidence. For cache purification, we compare global-only caching, purified-region caching with global retrieval, and entropy-based filtering with different aggregations. Caching purified regions improves performance by constructing class-specific prototypes from less ambiguous signals. Max aggregation further achieves the best results by preserving the most confident class-wise regional activation, preventing dominant responses from being diluted.

Hyperparameter Sensitivity Analysis. Here, we analyze four hyperparameters in PuRF: κ_g for selecting the top- κ_g classes from global predictions, δ for the refreshing onset and entropy penalization scale, and λ_1, λ_2 for balancing the trade-off between different loss terms. As shown in Fig. 4(a) and (b), smaller κ_g suppresses noisy candidate classes, while $\delta=500-1000$ strikes a balance between early and delayed refreshing, with stable performance across a wide range. In Fig. 4(c), $\lambda_1 = 0.2$ and $\lambda_2 = 0.5$ achieve the optimal performance.




Fig. 6: Visualization of class activation maps for region and cache purification.

4.4 Visualization

Fig. 5 presents the t-SNE [37] visualization of cached features for DPE [65], ReTA [29], and our PuRF. DPE and ReTA improve cache prototype quality for single-label recognition, but rely on global representations that entangle multiple labels, resulting in noisy and ambiguous cache entries in multi-label settings. In contrast, PuRF produces class-specific representations with well-separated clusters and strong inter-class discriminability. Fig. 6 illustrates the role of RP and CP. RP enhances fine-grained evidence aggregation by exploiting purified regions and enforcing global-local consistency for effective semantic optimization. EP leverages purified region-level representations to refine class-specific focus with more discriminative features, yielding improved cache calibration.

5 Conclusion

In this work, we focus on the largely unexplored problem of multi-label test-time adaptation. We propose PuRF, a purification-driven cache-based method that addresses the limitations of global representations in multi-label TTA, where label entanglement introduces bias and ambiguity in predictions and cache. PuRF performs region and cache purification to effectively exploit informative and reliable regions, improving recognition accuracy and discriminative cache calibration. Extensive experiments on five benchmarks verify its consistent superiority over existing methods, achieving state-of-the-art results and establishing it as a strong baseline for future multi-label adaptation research.

Acknowledgements

This work was supported by National Natural Science Foundation of China (Nos. 62271281, 62525103, 62441235) and Beijing Natural Science Foundation (L247026).

References

1. Barbu, A., Mayo, D., Alverio, J., Luo, W., Wang, C., Gutfreund, D., Tenenbaum, J., Katz, B.: Objectnet: A large-scale bias-controlled dataset for pushing the limits of object recognition models. In: Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*. vol. 32. Curran Associates, Inc. (2019)
2. Bilen, H., Vedaldi, A.: Weakly supervised deep detection networks. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2016)
3. Chen, S., Duan, B., Khan, S., Khan, F.S.: Interpretable zero-shot learning with locally-aligned vision-language model (2025), <https://arxiv.org/abs/2506.23822>
4. Chen, T., Wang, Z., Li, G., Lin, L.: Recurrent attentional reinforcement learning for multi-label image recognition. *Proceedings of the AAAI Conference on Artificial Intelligence* **32**(1) (Apr 2018). <https://doi.org/10.1609/aaai.v32i1.12281>
5. Chen, X., Zhai, H., Zhang, C., Shi, X., Li, R.: Multi-cache enhanced prototype learning for test-time generalization of vision-language models. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 2281–2291 (October 2025)
6. Chua, T.S., Tang, J., Hong, R., Li, H., Luo, Z., Zheng, Y.: Nus-wide: A real-world web image database from national university of singapore. In: *Proceedings of the ACM International Conference on Image and Video Retrieval. CIVR '09*, Association for Computing Machinery, New York, NY, USA (2009). <https://doi.org/10.1145/1646396.1646452>
7. Chuang, Y.S., Li, Y., Wang, D., Yeh, C.F., Lyu, K., Raghavendra, R., Glass, J., Huang, L., Weston, J., Zettlemoyer, L., et al.: Meta clip 2: A worldwide scaling recipe. *Advances in Neural Information Processing Systems* **38**, 48009–48036 (2026)
8. Ding, Z., Wang, A., Chen, H., Zhang, Q., Liu, P., Bao, Y., Yan, W., Han, J.: Exploring structured semantic prior for multi label recognition with incomplete labels. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 3398–3407 (June 2023)
9. Duan, S., Yang, X., Wang, N.: Multi-label prototype visual spatial search for weakly supervised semantic segmentation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 30241–30250 (June 2025)
10. Everingham, M., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The pascal visual object classes (voc) challenge. *International Journal of Computer Vision* **88**, 303–338 (2010)
11. Fan, X., Chen, X., Yang, L., Yap, C.H., Qureshi, R., Dou, Q., Yap, M.H., Shah, M.: Test-time retrieval-augmented adaptation for vision-language models. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 8810–8819 (October 2025)
12. Fang, Y., Sun, Q., Wang, X., Huang, T., Wang, X., Cao, Y.: Eva-02: A visual representation for neon genesis. *Image and Vision Computing* **149**, 105171 (2024)
13. Feng, C.M., Yu, K., Liu, Y., Khan, S., Zuo, W.: Diverse data augmentation with diffusions for effective test-time prompt tuning. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 2704–2714 (October 2023)

14. Gao, B.B., Zhou, H.Y.: Learning to discover multi-class attentional regions for multi-label image recognition. *IEEE Transactions on Image Processing* **30**, 5920–5932 (2021). <https://doi.org/10.1109/TIP.2021.3088605>
15. Guo, Z., Dong, B., Ji, Z., Bai, J., Guo, Y., Zuo, W.: Texts as images in prompt tuning for multi-label image recognition. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 2808–2817 (June 2023)
16. Hendrycks, D., Dietterich, T.: Benchmarking neural network robustness to common corruptions and perturbations. *arXiv preprint arXiv:1903.12261* (2019)
17. Hu, P., Sun, X., Sclaroff, S., Saenko, K.: Dualcoop++: Fast and effective adaptation to multi-label recognition with limited annotations. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **46**(5), 3450–3462 (2024). <https://doi.org/10.1109/TPAMI.2023.3346405>
18. Huang, F., Jiang, J., Jiang, Q., Li, H., Khan, F.N., Wang, Z.: Cosmic: Clique-oriented semantic multi-space integration for robust clip test-time adaptation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 9772–9781 (June 2025)
19. Huang, S., Yan, J., Liu, B., Liu, B., Hong, R.: Dual-view alignment learning with hierarchical-prompt for class-imbalance multi-label image classification. *IEEE Trans. Image Process.* **34**, 5989–6001 (2025). <https://doi.org/10.1109/TIP.2025.3609185>
20. Jia, C., Yang, Y., Xia, Y., Chen, Y.T., Parekh, Z., Pham, H., Le, Q., Sung, Y.H., Li, Z., Duerig, T.: Scaling up visual and vision-language representation learning with noisy text supervision. In: Meila, M., Zhang, T. (eds.) *Proceedings of the 38th International Conference on Machine Learning. Proceedings of Machine Learning Research*, vol. 139, pp. 4904–4916. PMLR (18–24 Jul 2021)
21. Jia, J., Gao, N., He, F., Chen, X., Huang, K.: Learning disentangled attribute representations for robust pedestrian attribute recognition. *Proceedings of the AAAI Conference on Artificial Intelligence* **36**(1), 1069–1077 (Jun 2022)
22. Jia, J., He, F., Gao, N., Chen, X., Huang, K.: Learning disentangled label representations for multi-label classification. *arXiv preprint arXiv:2212.01461* (2022)
23. Karmanov, A., Guan, D., Lu, S., El Saddik, A., Xing, E.: Efficient test-time adaptation of vision-language models. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 14162–14171 (June 2024)
24. Khattak, M.U., Rasheed, H., Maaz, M., Khan, S., Khan, F.S.: Maple: Multi-modal prompt learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 19113–19122 (June 2023)
25. Koh, P.W., Sagawa, S., Marklund, H., Xie, S.M., Zhang, M., Balasubramani, A., Hu, W., Yasunaga, M., Phillips, R.L., Gao, I., Lee, T., David, E., Stavness, I., Guo, W., Earnshaw, B., Haque, I., Beery, S.M., Leskovec, J., Kundahe, A., Pierson, E., Levine, S., Finn, C., Liang, P.: Wilds: A benchmark of in-the-wild distribution shifts. In: Meila, M., Zhang, T. (eds.) *Proceedings of the 38th International Conference on Machine Learning. Proceedings of Machine Learning Research*, vol. 139, pp. 5637–5664. PMLR (18–24 Jul 2021)
26. Lee, Y., Kim, D., Kang, J., Bang, J., Song, H., Lee, J.G.: RA-TTA: Retrieval-augmented test-time adaptation for vision-language models. In: *The Thirteenth International Conference on Learning Representations* (2025), <https://openreview.net/forum?id=V3zobHnS61>
27. Li, Z., Arora, S.: An exponential learning rate schedule for deep learning. *arXiv preprint arXiv:1910.07454* (2019)

28. Liang, Y., Chen, H., Lin, Z., Liu, P., Wang, J., Wang, G., Li, L., Zhao, S., Han, J., Ding, G.: Dar-prompt: Dynamic regulation in prompt tuning for multi-label zero-shot learning. *IEEE Transactions on Image Processing* **34**, 7697–7711 (2025). <https://doi.org/10.1109/TIP.2025.3626157>
29. Liang, Y., Chen, H., Xiong, Y., Zhou, Z., Lyu, M., Lin, Z., Niu, S., Zhao, S., Han, J., Ding, G.: Advancing reliable test-time adaptation of vision-language models under visual variations. In: *Proceedings of the 33rd ACM International Conference on Multimedia*. p. 4788–4797. MM '25, Association for Computing Machinery, New York, NY, USA (2025)
30. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: *Computer Vision – ECCV 2014*. pp. 740–755 (2014)
31. Lin, Y., Chen, M., Zhang, K., Li, H., Li, M., Yang, Z., Lv, D., Lin, B., Liu, H., Cai, D.: Tagclip: A local-to-global framework to enhance open-vocabulary multi-label classification of clip without training. *Proceedings of the AAAI Conference on Artificial Intelligence* **38**(4), 3513–3521 (Mar 2024)
32. Liu, Z., Sun, H., Peng, Y., Zhou, J.: Dart: Dual-modal adaptive online prompting and knowledge retention for test-time adaptation. *Proceedings of the AAAI Conference on Artificial Intelligence* **38**(13), 14106–14114 (Mar 2024). <https://doi.org/10.1609/aaai.v38i13.29320>
33. Loshchilov, I., Hutter, F.: Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983* (2016)
34. Lyu, M., Hao, T., Xu, X., Chen, H., Lin, Z., Han, J., Ding, G.: Learn from the learnt: Source-free active domain adaptation via contrastive sampling and visual persistence. In: *Computer Vision – ECCV 2024*. pp. 228–246. Springer Nature Switzerland, Cham (2025)
35. Ma, L.L., Xu, S., Xie, M.K., Wang, L., Sun, D., Zhao, H.: Correlative and discriminative label grouping for multi-label visual prompt tuning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 25434–25443 (June 2025)
36. Ma, L., Xie, H., Wang, L., Fu, Y., Sun, D., Zhao, H.: Text-region matching for multi-label image recognition with missing labels. In: *Proceedings of the 32nd ACM International Conference on Multimedia*. p. 6133–6142. MM '24, Association for Computing Machinery, New York, NY, USA (2024). <https://doi.org/10.1145/3664647.3680815>
37. Maaten, L.v.d., Hinton, G.: Visualizing data using t-sne. *Journal of machine learning research* **9**(Nov), 2579–2605 (2008)
38. Narayan, S., Gupta, A., Khan, S., Khan, F.S., Shao, L., Shah, M.: Discriminative region-based multi-label zero-shot learning. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 8731–8740 (October 2021)
39. Peng, X., Bai, Q., Xia, X., Huang, Z., Saenko, K., Wang, B.: Moment matching for multi-source domain adaptation. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (October 2019)
40. Pratt, S., Covert, I., Liu, R., Farhadi, A.: What does a platypus look like? generating customized prompts for zero-shot image classification. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 15691–15701 (October 2023)
41. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., Sutskever, I.: Learning transferable

- visual models from natural language supervision. In: Meila, M., Zhang, T. (eds.) Proceedings of the 38th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 139, pp. 8748–8763. PMLR (18–24 Jul 2021)
42. Rawlekar, S., Cai, Y., Wang, Y., Yang, M.H., Ahuja, N.: Efficiently disentangling clip for multi-object perception. arXiv preprint arXiv:2502.02977 (2025)
 43. Sarafianos, N., Xu, X., Kakadiaris, I.A.: Deep imbalanced attribute classification using visual attention aggregation. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 708–725 (September 2018)
 44. Shi, M., Tang, Y., Zhu, X., Liu, J.: Multi-label graph convolutional network representation learning. IEEE Transactions on Big Data **8**(5), 1169–1181 (2022). <https://doi.org/10.1109/TBDATA.2020.3019478>
 45. Shu, M., Nie, W., Huang, D.A., Yu, Z., Goldstein, T., Anandkumar, A., Xiao, C.: Test-time prompt tuning for zero-shot generalization in vision-language models. In: Koyejo, S., Mohamed, S., Agarwal, A., Belgrave, D., Cho, K., Oh, A. (eds.) Advances in Neural Information Processing Systems. vol. 35, pp. 14274–14289. Curran Associates, Inc. (2022)
 46. Sun, X., Hu, P., Saenko, K.: Dualcoop: Fast adaptation to multi-label recognition with limited annotations. In: Koyejo, S., Mohamed, S., Agarwal, A., Belgrave, D., Cho, K., Oh, A. (eds.) Advances in Neural Information Processing Systems. vol. 35, pp. 30569–30582. Curran Associates, Inc. (2022)
 47. Tan, H., Tan, Z., Li, J., Liu, A., Wan, J., Lei, Z.: Recover and match: Open-vocabulary multi-label recognition through knowledge-constrained optimal transport. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4650–4660 (June 2025)
 48. Tschannen, M., Gritsenko, A., Wang, X., Naeem, M.F., Alabdulmohsin, I., Parthasarathy, N., Evans, T., Beyer, L., Xia, Y., Mustafa, B., Hénaff, O., Harmsen, J., Steiner, A., Zhai, X.: Siglip 2: Multilingual vision-language encoders with improved semantic understanding, localization, and dense features. arXiv preprint arXiv:2502.14786 (2025)
 49. Wan, F., Liu, C., Ke, W., Ji, X., Jiao, J., Ye, Q.: C-mil: Continuation multiple instance learning for weakly supervised object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019)
 50. Wang, A., Chen, H., Lin, Z., Ding, Z., Liu, P., Bao, Y., Yan, W., Ding, G.: Hierarchical prompt learning using clip for multi-label classification with single positive labels. In: Proceedings of the 31st ACM International Conference on Multimedia. p. 5594–5604. MM '23, Association for Computing Machinery, New York, NY, USA (2023)
 51. Wang, A., Chen, H., Lin, Z., Han, J., Ding, G.: Lsnet: See large, focus small. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 9718–9729 (June 2025)
 52. Wang, D., Shelhamer, E., Liu, S., Olshausen, B., Darrell, T.: Tent: Fully test-time adaptation by entropy minimization (2021), <https://arxiv.org/abs/2006.10726>
 53. Wang, J., Yang, Y., Mao, J., Huang, Z., Huang, C., Xu, W.: Cnn-rnn: A unified framework for multi-label image classification. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2016)
 54. Wang, M., Luo, C., Hong, R., Tang, J., Feng, J.: Beyond object proposals: Random crop pooling for multi-label image recognition. IEEE Transactions on Image Processing **25**(12), 5678–5688 (2016). <https://doi.org/10.1109/TIP.2016.2612829>

55. Wang, Z., Chen, T., Li, G., Xu, R., Lin, L.: Multi-label image recognition by recurrently discovering attentional regions. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV) (Oct 2017)
56. Wu, X., Jiang, Q.Y., Yang, Y., Wu, Y.F., Chen, Q.G., Lu, J.: Tai++: Text as image for multi-label image classification by co-learning transferable prompt (2024), <https://arxiv.org/abs/2405.06926>
57. Wu, X., Yu, F., Yang, Y., Chen, Q.G., Lu, J.: Multi-label test-time adaptation with bound entropy minimization. In: The Thirteenth International Conference on Learning Representations (2025), <https://openreview.net/forum?id=75PhjtbBdr>
58. Xie, C.W., Sun, S., Xiong, X., Zheng, Y., Zhao, D., Zhou, J.: Ra-clip: Retrieval augmented contrastive language-image pre-training. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 19265–19274 (June 2023)
59. Xiong, Y., Chen, H., Hao, T., Lin, Z., Han, J., Zhang, Y., Wang, G., Bao, Y., Ding, G.: Pyra: Parallel yielding re-activation for training-inference efficient task adaptation. In: Computer Vision – ECCV 2024. pp. 455–473. Springer Nature Switzerland, Cham (2025)
60. Xiong, Y., Chen, H., Lin, Z., Zhao, S., Ding, G.: Confidence-based visual dispersal for few-shot unsupervised domain adaptation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 11621–11631 (October 2023)
61. Xiong, Y., Zhou, Z., Liang, Y., Chen, H., Lin, Z., Hao, T., Zhang, F., Han, J., Ding, G.: Neutralizing token aggregation via information augmentation for efficient test-time adaptation (2025), <https://arxiv.org/abs/2508.03388>
62. Ye, J., He, J., Peng, X., Wu, W., Qiao, Y.: Attention-driven dynamic graph convolutional network for multi-label image recognition. In: Computer Vision – ECCV 2020. pp. 649–665 (2020)
63. Yoon, H.S., Yoon, E., Tee, J.T.J., Hasegawa-Johnson, M., Li, Y., Yoo, C.D.: C-tp: Calibrated test-time prompt tuning for vision-language models via text feature dispersion (2024)
64. Zanella, M., Fuchs, C., De Vleeschouwer, C., Ben Ayed, I.: Realistic test-time adaptation of vision-language models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 25103–25112 (2025)
65. Zhang, C., Stepputtis, S., Sycara, K.P., Xie, Y.: Dual prototype evolving for test-time generalization of vision-language models. In: The Thirty-eighth Annual Conference on Neural Information Processing Systems (2024)
66. Zhang, C., Xu, K., Liu, Z., Peng, Y., Zhou, J.: Scap: Transductive test-time adaptation via supportive clique-based attribute prompting. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 30032–30041 (June 2025)
67. Zhang, J., Huang, J., Zhang, X., Shao, L., Lu, S.: Historical test-time prompt tuning for vision foundation models. In: Globerson, A., Mackey, L., Belgrave, D., Fan, A., Paquet, U., Tomczak, J., Zhang, C. (eds.) Advances in Neural Information Processing Systems. vol. 37, pp. 12872–12896. Curran Associates, Inc. (2024). <https://doi.org/10.52202/079017-0410>
68. Zhang, R., Zhang, W., Fang, R., Gao, P., Li, K., Dai, J., Qiao, Y., Li, H.: Tip-adapter: Training-free adaption of clip for few-shot classification. In: Computer Vision – ECCV 2022. pp. 493–510. Springer Nature Switzerland, Cham (2022)

69. Zhang, T., Wang, J., Guo, H., Dai, T., Chen, B., Xia, S.T.: Boostadapter: Improving vision-language test-time adaptation via regional bootstrapping. In: The Thirty-eighth Annual Conference on Neural Information Processing Systems (2024)
70. Zhang, Y., Zhu, W., Tang, H., Ma, Z., Zhou, K., Zhang, L.: Dual memory networks: A versatile adaptation approach for vision-language models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 28718–28728 (June 2024)
71. Zhao, S., Wang, X., Zhu, L., Yang, Y.: Test-time adaptation with CLIP reward for zero-shot generalization in vision-language models. In: The Twelfth International Conference on Learning Representations (2024)
72. Zhou, K., Yang, J., Loy, C.C., Liu, Z.: Learning to prompt for vision-language models. *International Journal of Computer Vision* **130**(9), 2337–2348 (2022)
73. Zhu, F., Li, H., Ouyang, W., Yu, N., Wang, X.: Learning spatial regularization with image-level supervisions for multi-label image classification. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (July 2017)
74. Zhu, K., Wu, J.: Residual attention: A simple but effective method for multi-label recognition. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 184–193 (October 2021)
75. Zhu, X., Liu, J., Tang, D., Ge, J., Liu, W., Liu, B., Cao, J.: Query-based knowledge sharing for open-vocabulary multi-label classification. *ACM Trans. Multimedia Comput. Commun. Appl.* **21**(12) (Nov 2025). <https://doi.org/10.1145/3762195>