

Is Less More? Exploring Token Condensation as Training-free Test-time Adaptation

Zixin Wang[†] Dong Gong[‡] Sen Wang[†] Zi Huang[†] Yadan Luo^{†*}

[†]The University of Queensland, Australia [‡]University of New South Wales, Australia

{zixin.wang, sen.wang, helen.huang, y.luo}@uq.edu.au, dong.gong@unsw.edu.au

Abstract

*Contrastive Language-Image Pretraining (CLIP) excels at learning generalizable image representations but often falls short in zero-shot inference on certain downstream datasets. Test-time adaptation (TTA) mitigates this issue by adjusting components like normalization layers or context prompts, yet it typically requires large batch sizes and extensive augmentations, leading to high computational costs. This raises a key question: **Can VLMs’ performance drop in specific test cases be mitigated through efficient, training-free approaches?** To explore the solution, we investigate token condensation (TC) techniques, originally designed to enhance vision transformer efficiency by refining token usage during inference. We observe that informative tokens improve visual-text alignment in VLMs like CLIP on unseen datasets. However, existing TC methods often fail to maintain in-distribution performance when reducing tokens, prompting us to ask: *How can we transform TC into an effective “free-lunch” adaptation strategy for VLMs?* To address this, we propose **Token Condensation as Adaptation (TCA)**, a training-free adaptation method that takes a step beyond standard TC. Rather than passively discarding tokens, TCA condenses token representation by introducing reservoir-based domain anchor tokens for information-preserving token reduction and logits correction. TCA achieves up to a 21.4% performance improvement over the strongest baseline on cross-dataset benchmark and the CIFAR-100-Corrupted dataset while reducing GFLOPs by 12.2% to 48.9%, with minimal hyperparameter dependency on both CLIP and SigLIP series. Code is available at <https://github.com/Jo-wang/TCA>.*

1. Introduction

Online test-time adaptation (TTA) [69] has emerged as a promising strategy to handle distribution shifts encountered during inference [10, 11, 33]. As an online source-

free extension of domain adaptation [8, 39, 68], TTA dynamically fine-tunes pretrained models on unlabeled data batches, enhancing generalization by aligning intermediate-layer batch statistics [46], optimizing for first-order flatness in the loss landscape [19], promoting self-supervised consistency across augmentations [82], or tracking model historical weights [32]. It is particularly useful in zero-shot scenarios [6, 7, 9, 12], where no labeled data from the target domain is available. Despite the success of traditional TTA methods, they often require computationally expensive tuning of the backbone’s parameters. This challenge is further exacerbated in CLIP [48], SigLIP [80], and SigLIP v2 [60], which rely on visual-text similarity for zero-shot prediction. With vast parameter sets and large batch sizes (e.g., 256) to stabilize adaptation [15], applying conventional TTA to VLMs becomes increasingly impractical.

To circumvent the need for full model tuning, prompting-based methods have been introduced [35], with test-time prompting (TPT) emerging as a more efficient alternative. By learning a small set of task-specific context prompts, TPT better aligns text features with visual representations, enabling lightweight adaptation. However, TPT primarily focuses on refining text inputs while largely overlooking the impact of visual distribution shifts. Besides, adapting to high-variance target images through prompts often relies on external source data [54] or extensive data augmentation [18] (e.g., 60× more AugMix or diffusion-based synthetic samples). In strict online TTA settings, where the batch size is constrained to one, this reliance on augmentation significantly inflates computational costs, leading to a 60× rise in GFLOPs compared to single-sample processing (i.e., 1108.61 vs. 17.59 GFLOPs). The need for gradient backpropagation during inference further increases the computation burden, making exiting TPT suboptimal for many resource-constrained applications.

The trade-off between the increased computation of TTA and the efficiency demands of testing raises a key question: Can VLMs achieve TTA while maintaining testing efficiency? In this paper, we explore training-free solutions to address this challenge. Given that many VLM visual

*Corresponding author.

encoders rely on Vision Transformers (ViTs), token pruning and merging accelerate inference by reducing redundancy while preserving essential tokens through adjusted token usage in the forward pass. However, efficiency-driven token adjustments often come at the cost of degraded *in-distribution* performance (e.g., ImageNet-1K validation) in exchange for lower GFLOPs [34]. In contrast, our analysis reveals that selectively condensing low-attentiveness tokens not only preserves performance but enhance it on certain *unseen* datasets (Tab. 1). This motivates our exploration of efficient, training-free TTA through token condensation.

To better understand which tokens contribute most to VLM adaptation, we conducted a preliminary analysis on token importance, investigating their role in visual-text alignment (Fig. 2a). This analysis highlights two key types of tokens that benefit from condensation: (1) class-irrelevant background tokens, which may mislead the model by emphasizing non-essential regions that deviate from pre-training data distributions, and (2) class-ambiguous object tokens, such as animal fur or textures, which overlap across categories and disperse visual embeddings. However, simply adapting to individual test samples in isolation is suboptimal, as it fails to capture domain-level trends that evolve over time. While textual features encode valuable class-specific information, their embedding space differs from visual representations, limiting direct alignment.

To address these challenges, we introduce Token Condensation as Adaptation (TCA), a training-free adaptation method that dynamically refines token selection based on evolving domain knowledge. Instead of indiscriminately reducing tokens, TCA tracks and maintains domain-representative tokens over time, allowing the model to be adapted to new domains. As shown in Fig. 1, TCA stores domain-aware anchor tokens (e.g., `<cls>` tokens in CLIP or pooled vectors in SigLIP) in a reservoir (DTR), which then guides adaptation by refining class-specific representations. At each time step, TCA aligns token usage with domain trends by selectively adjusting token attentiveness based on both the current test sample and accumulated domain information in DTR. Additionally, TCA utilizes domain anchor tokens to refine model logits, improving visual-text alignment without modifying model parameters.

To our knowledge, this is the first work to explore token condensation as a form of test-time adaptation. Unlike previous approaches, TCA offers a lightweight, scalable, and training-free solution that generalizes across diverse data domains and can be readily extended to SigLIP and SigLIP v2. The extensive evaluations on the cross-dataset benchmark and the CIFAR-100-Corrupted dataset demonstrate that TCA consistently outperforms traditional TTA, prompting, and test-time prompting methods, achieving up to a 21.4% improvement over the strongest baseline while reducing GFLOPs by 12.2% to 48.9%.

2. Related Work

Online Test-time Adaptation. To address performance degradation during test time, online test-time adaptation (TTA) has gained significant attention. Current TTA methods can be categorized into three main types [69]: optimization-, data-, and model-based approaches. Optimization-based methods focus on model updates and optimization objectives [21, 42, 82]. A prominent example is Tent [61], which adapts Batch Normalization layers [26] by entropy minimization. SAR [46] extends this approach to Layer Normalization [1] and Group Normalization [74] with sharpness-aware minimization [19]. Data-based methods include augmentations like selective [64] and adversarial augmentation [59], and memory banks [5, 20, 78]. Model-based approaches involve architectural modifications to enhance model adaptability during testing [27, 28, 37, 67]. However, they typically depend on large batch sizes and augmentations, which introduce significant latency for online prediction.

Recently, vision-language models like CLIP [48] have excelled beyond fixed label sets, rendering traditional TTA methods less suitable [15]. As a result, various online adaptation strategies have been proposed to improve zero-shot generalization. Test-time prompt tuning has emerged as a key approach in this context. TPT [57] optimizes learnable prompts using data augmentations and soft entropy minimization, Diff-TPT [18] enriches this with more diverse augmentations [52], while C-TPT [77] focusing on model calibration. Other methods like VTE [15] and DART [38] leverage prompt ensembles with DART further employing moving averages to boost performance. SwapPrompt [40] incorporates an EMA-updated target prompt. AdaPrompt [81] utilizes a class-balanced memory bank to enhance adaptability. SCP [66] builds on TPT with a teacher-student framework to prevent semantic drift, while RLCF [83] incorporates reinforcement learning strategy [71] to optimize the adaptation process. Beyond these, MTA [79] introduces a new objective based on test-time augmentation to optimize visual features in the semantic space. TDA [29] further improves CLIP’s zero-shot ability by incorporating positive and negative caches with a training-free adapter. However, it relies on a large number of hyperparameters and is highly sensitive to them, while incurring significant computational costs during inference. In contrast, our approach strikes a better balance between computational efficiency and performance, outperforming both training-required and training-free methods.

Token Condensation in Vision Transformers. Vision transformers have achieved notable success in image recognition tasks, but their deployment is often limited by resource-constrained environments. To address this, various token condensation methods [30, 44, 50, 53, 63, 76, 86] have been proposed to reduce the computational overhead,

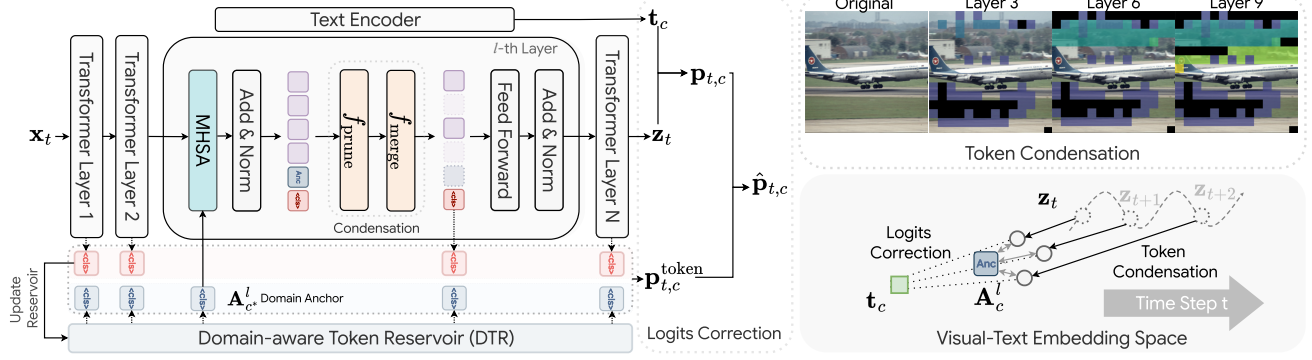


Figure 1. **Proposed Token Condensation as Adaptation (TCA)**. Using CLIP as an example, to adapt visual embeddings to text embeddings during test-time, TCA utilizes a domain-aware token reservoir (DTR) to retain historical `<cls>` tokens with the lowest uncertainty as domain anchor tokens. These anchor tokens assist in (1) condensing tokens with low attentiveness scores (*top-right*) and (2) acting as token-level classifiers to refine predictions through logits self-correction, moving visual embeddings \mathbf{z}_t toward text embeddings \mathbf{t}_c .

primarily through two strategies: token pruning and token merging. Token pruning eliminates less informative tokens to save computation, as seen in methods like EViT [34], which retains tokens based on their attentiveness to the `<cls>` tokens. ATS [16] introduces input-dependent token pruning to adapt to variability across inputs. Token merging, on the other hand, seeks to combine similar tokens to reduce redundancy. For instance, ToME [3] uses bipartite soft matching to merge neighboring tokens that exhibit similarity. Hybrid approaches have also emerged, such as TPS [70], which prunes tokens and transfers information to retained ones using nearest-neighbor matching, and PruMerge [56], which prunes inattentive tokens using interquartile range and merges via k-nearest neighbors. While previous works have focused on enhancing efficiency within pure ViT models, our approach goes far beyond a mere adaptation of these methods, addressing multimodal distribution shifts in VLMs. This shift remains underexplored, particularly in how to use semantic guidance to prune irrelevant visual tokens that introduce ambiguity. By condensing these tokens, we effectively reduce such distribution shifts, enhancing test-time performance while simultaneously lowering computational costs (See Tab. 1).

3. Token Condensation as Adaptation

Without loss of generality, we use CLIP as the representative model in the rest of the sections. We show our method can be effortlessly extended to other VLMs in Sec. 4.3.

Problem Set-up. We begin by revisiting online test-time adaptation. For a given downstream task \mathcal{D}_{tar} , the test data $\mathbf{x} = \{\mathbf{x}_t\}_{t=1}^T$ arrives sequentially at each time step t . The objective is to adapt the model on the fly to classify the incoming test samples into one of C classes, each represented by a textual prompt like “a photo of a `<classname>`”. CLIP embeds both visual and textual inputs into a shared space. The visual encoder E_v extracts

visual features $\mathbf{z}_t = E_v(\mathbf{V}_t) \in \mathbb{R}^D$ from image patches $\mathbf{V}_t = [\mathbf{v}_{\text{cls}}, \mathbf{v}_1, \dots, \mathbf{v}_N] \in \mathbb{R}^{(N+1) \times D_v}$ of dimension D_v , where \mathbf{v}_{cls} is a CLIP-only `<cls>` token appended to N patches. The text encoder E_t generates class embeddings $\mathbf{T} = \{\mathbf{t}_c\}_{c=1}^C$, where each $\mathbf{t}_c \in \mathbb{R}^D$ corresponding to a class prompt. Classification is performed by computing the cosine similarity between the visual embedding \mathbf{z}_t and each class embedding \mathbf{t}_c with the probabilities calculated as:

$$\mathbf{p}_{t,c}(\mathbf{z}_t, \mathbf{t}_c) = \frac{\exp(\cos(\mathbf{z}_t, \mathbf{t}_c)/\tau)}{\sum_{j=1}^C \exp(\cos(\mathbf{z}_t, \mathbf{t}_j)/\tau)}, \quad (1)$$

where τ denotes the temperature parameter controlling the sharpness of the output distribution.


For the visual encoder of ViT-based CLIP, given an L -layer ViT, the forward pass through the l -th Transformer block, where $l \in [1, 2, \dots, L]$, is formulated as:

$$\begin{aligned} \mathbf{V}^{l+1} &= \hat{\mathbf{V}}^l + \text{MLP}(\hat{\mathbf{V}}^l), \\ \hat{\mathbf{V}}^l &= \mathbf{V}^l + \frac{1}{H} \sum_{h=1}^H \text{Attn}(\mathbf{V}^l \mathbf{W}_Q^h, \mathbf{V}^l \mathbf{W}_K^h) \mathbf{V}^l \mathbf{W}_V^h, \end{aligned} \quad (2)$$

where $\mathbf{V}^l \in \mathbb{R}^{(N+1) \times D_v}$ is the token embeddings at layer l . The matrices $\mathbf{W}_Q^h, \mathbf{W}_K^h, \mathbf{W}_V^h \in \mathbb{R}^{D_v \times D_v}$ are the linear projection matrices for the query, key, and value vectors in the h -th attention head, respectively, within the total number of attention heads H .

3.1. Empirical Discussion

Pitfalls of TTA. In CLIP, since the target domain \mathcal{D}_{tar} is unseen during pre-training, the alignment between visual embeddings \mathbf{z}_t and the textual embeddings \mathbf{T} may be sub-optimal. Previous methods have attempted to address this by learning domain-specific prompts [77] or replacing classifier weights with visual centroids [27] to move \mathbf{T} closer to \mathbf{z}_t . However, the variability in CLIP’s visual embeddings



(a) Impact of token removal on alignment. The warmer color indicates higher attentiveness. (b) Alignment between the updated domain anchor token and text embedding over time.

Figure 2. Empirical studies of token influence and the strategy of caching domain anchor token (i.e., $\langle \text{cls} \rangle$ tokens in CLIP) to improve alignment.

is often much *higher* than in textual embeddings [48]. At the patch level, individual tokens within the visual embeddings can drift and vary significantly [48]. Thus, it becomes more urgent to derive methods that adjust \mathbf{z}_t towards \mathbf{T} for improved alignment.

How to Mine $\langle \text{cls} \rangle$ Tokens Aligned with Text? The $\langle \text{cls} \rangle$ token in CLIP visual encoder is trained for broad concept alignment, often extending beyond target classes. Due to the mismatch between textual and visual token spaces, a key challenge is finding a more representative visual $\langle \text{cls} \rangle$ token that better aligns with text embeddings. To address this, we track visual $\langle \text{cls} \rangle$ tokens from test samples with the lowest entropy as **domain anchors**, updating them dynamically at each time step t . We then measure the alignment between these stored $\langle \text{cls} \rangle$ tokens and text embeddings \mathbf{T} over time (Fig. 2b). Our results reveal a progressive correlation – as the stored domain anchors consistently favor lower-entropy samples, their $\langle \text{cls} \rangle$ tokens increasingly align with text embeddings, effectively bridging the gap between visual and textual representations.

3.2. Method Overview

Building on our empirical findings, we propose Token Condensation as Adaptation (TCA), a training-free online adaptation strategy that enhances VLMs by filtering out tokens that contribute to visual-text misalignment, ensuring stable and efficient adaptation.

To leverage VLMs’ existing knowledge, TCA introduces a domain-aware token reservoir (DTR) that retains representative anchor tokens for adaptation (Sec. 3.3). Note that in models like SigLIP [60] and SigLIP v2 [60], which lack a $\langle \text{cls} \rangle$ token, we use the pooled feature vector to align with their holistic aggregation strategy. These domain anchors provide a stable reference for adaptation. Guided by these anchors, we perform cross-head token condensation between multi-head self-attention and feed-forward layers [34], selectively merging or discarding less informative tokens to ensure that only relevant, domain-consistent information is retained (Sec. 3.4). Finally, stored domain anchor tokens aid in logits self-correction, refining model predic-

tions based on accumulated domain knowledge (Sec. 3.5).

3.3. Domain-aware Token Reservoir

To keep representative anchor tokens at domain level, we introduce domain-aware token reservoir $\mathfrak{R} = \{\mathfrak{R}_c\}_{c=1}^C$. Notably, both token reduction and logits correction are building on the top of these saved tokens. In \mathfrak{R} , each buffer $\mathfrak{R}_c = \{(\mathbf{H}_c(\mathbf{z}_i; \mathbf{t}_c), \mathbf{A}_{i,c}^{\text{cls}})\}_{i=1}^M$ is structured as a *priority queue* that retains the top M most reliable domain anchor tokens of target samples, which serve to implicitly distil semantic information from the corresponding text prompt \mathbf{t}_c to guide the visual adaptation. These domain anchor tokens are crucial *alignment proxies*: although the architectures of the text encoder E_t and the visual encoder E_v differ, the selected domain anchor tokens help determine which visual tokens best align with text features. The reliability of these domain anchor tokens is quantified by entropy scores,

$$\mathbf{H}_c(\mathbf{z}_t, \mathbf{t}_c) = -\mathbf{p}_{t,c}(\mathbf{z}_t, \mathbf{t}_c) \log \mathbf{p}_{t,c}(\mathbf{z}_t, \mathbf{t}_c), \quad (3)$$

which act as keys to update the reservoir \mathfrak{R}_c . At each time step t , for each visual embedding \mathbf{z}_t , the corresponding anchor embeddings from all L layers $\mathbf{A}_{t,c}^{\text{cls}} = [\mathbf{v}_{\text{cls}}^1, \dots, \mathbf{v}_{\text{cls}}^L] \in \mathbb{R}^{L \times D_v}$ will be stored in \mathfrak{R}_c if $\text{argmax}(\mathbf{p}_{t,c}) = c$, ensuring that only the most semantically consistent samples are retained:

$$\mathfrak{R}_c \leftarrow \text{update}(\mathfrak{R}_c, (\mathbf{H}_c(\mathbf{z}_t, \mathbf{t}_c), \mathbf{A}_{t,c}^{\text{cls}})). \quad (4)$$

If the priority queue \mathfrak{R}_c has reached its capacity M , the sample with the highest entropy score is discarded and replaced with the new sample. Strategies for *updating* the reservoir, such as first-in-first-out (FIFO) and similarity- or diversity-enforcing methods, are explored in Sec. 4.3.

3.4. Domain-aware Cross-head Token Reduction

Inspired by our preliminary studies as despite in Fig. 2a, we introduce token condensation for training-free adaptation. Prior token reduction methods [34] primarily discard patch tokens with lower averaged attention scores $\mathbf{S} \in \mathbb{R}^N$ relative to the $\langle \text{cls} \rangle$ token $\mathbf{v}_{\text{cls}}^l$ across all attention heads $\mathbf{S}_i = \frac{1}{H} \sum_{h=1}^H \text{Attn}(\mathbf{v}_{\text{cls}}^l \mathbf{W}_Q^h, \mathbf{v}_i^l \mathbf{W}_K^h)$. However, this approach faces two limitations when applied in TTA tasks: (1) The $\langle \text{cls} \rangle$ token is *universal* and may not be specifically aligned with the target class set. It may capture broad, unrelated semantics (e.g., “cat food”), leading to the retention of irrelevant tokens that mislead the model into making incorrect predictions of the target class (e.g., “cat”). (2) Averaging attention scores across all heads risks omitting important details, as each attention head tends to focus on distinct features (e.g., shape, color). Outliers in attention heads (highlighted by red circles in Fig. 3) may disproportionately dominate the overall score, overshadowing valuable information by other heads. To overcome these limitations, we

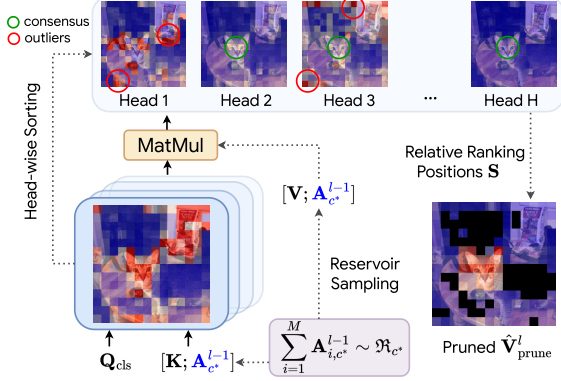


Figure 3. An overview of domain-aware cross-head token pruning.

propose a *domain-aware cross-head token reduction* that evaluates the token importance individually for each attention head with the consideration of saved domain anchor tokens and utilizes the averaged relative ranking positions to determine which tokens to prune and which to merge (see Fig. 3). This approach reaches a more robust cross-head consensus and mitigates the impact of outliers.

Domain-aware Token Evaluation. Firstly, we sample the domain anchor token for the $(l-1)$ -th layer from domain-aware class buffer \mathfrak{R}_{c^*} , where c^* is determined by the maximum cosine similarity between the current $\mathbf{v}_{\text{cls}}^l$ token embedding and the average stored domain anchor tokens in \mathfrak{R}_c for each class c . Here, the averaged domain anchor token is calculated by $\mathbf{A}_c^{l-1} = \frac{1}{M} \sum_{i \in [M]} \mathbf{A}_{i,c}^{l-1}$. $c^* = \arg \max_{c \in [C]} \cos(\mathbf{v}_{\text{cls}}^l, \mathbf{A}_c^{l-1})$. Subsequently, we refine the attention map by concatenating $\mathbf{v}_{\text{cls}}^l$ with the historical domain anchor token $\mathbf{A}_{c^*}^{l-1}$,

$$\text{Attn}([\mathbf{v}_{\text{cls}}^l; \mathbf{A}_{c^*}^{l-1}] \mathbf{W}_Q^h, [\mathbf{V}^l; \mathbf{A}_{c^*}^{l-1}] \mathbf{W}_K^h), \quad (5)$$

where $[\cdot; \cdot]$ indicates concatenation. This providing historical context that is *better* aligned with target semantics.

Cross-head Token Reduction. To better evaluate the attentiveness of the tokens, we compute the token reduction score $\mathbf{S}_i^{\text{head}} = \frac{1}{H} \sum_{h=1}^H \text{rank}_h(i)$ for the i -th token, where $\text{rank}_h(i)$ gives the relative ranking position of token i in head h based on the attention score in Eq. (5). This ensures that tokens receiving consistently high attention across individual heads are retained, thereby achieving greater robustness to outliers in specific attention heads.

We find out the class-irrelevant and class-ambiguous tokens by the cross-head score $\mathbf{S}_i^{\text{head}}$ and formulate token reduction in TCA as:

$$\hat{\mathbf{V}}^l = f_{\text{merge}} \circ f_{\text{prune}}(\mathbf{V}^l; \mathfrak{R}), \quad (6)$$

where $f_{\text{prune}}(\cdot; \mathfrak{R}) : \mathbb{R}^{N+1} \mapsto \mathbb{R}^{(\alpha \cdot R \cdot N)+1}$ and $f_{\text{merge}}(\cdot; \mathfrak{R}) : \mathbb{R}^{(\alpha \cdot R \cdot N)+1} \mapsto \mathbb{R}^{R \cdot N+1}$ are pruning and

merging respectively, responsible for reducing the number of tokens from $N+1$ (including the $\langle \text{cls} \rangle$ token) to $R \cdot N+1$, where R is the fraction of tokens to be preserved. α controls the extent of token pruning. We first *prune the class-irrelevant ones*:

$$\hat{\mathbf{V}}_{\text{prune}}^l \leftarrow \{\hat{\mathbf{v}}_i^l \mid \mathbf{S}_i^{\text{head}} \leq \theta_{\text{prune}}(\alpha, R), \forall i \in [N]\}, \quad (7)$$

where $\hat{\mathbf{V}}_{\text{prune}}^l$ is the set of tokens retained after pruning at layer l . The pruning threshold $\theta_{\text{prune}}(\alpha, R)$ allows only top-ranked $\alpha \cdot R \cdot N$ tokens are retained.

As depicted in Fig. 2a, the *class-ambiguous* tokens, although relevant to the target class, exhibit high uncertainty:

$$\Phi = \{i \mid \theta_{\text{merge}}(R) \leq \mathbf{S}_i^{\text{head}} \leq \theta_{\text{prune}}(\alpha, R), \forall i\}, \quad (8)$$

where $\theta_{\text{merge}}(R)$ denotes thresholds for token selection during merging. The selected tokens $\mathbf{V}_{\Phi}^l = \{\mathbf{v}_i^l\}_{i \in \Phi}$ can introduce variance or noise into latent representation \mathbf{z}_t and negatively impact the final classification decision. To address this, we propose a *domain-aware token merging strategy* to consolidate these tokens into more representative ones. Here, rather than truncating neighbored token pairs with bipartite soft matching [3], applying spectral clustering [2], or graph pooling [73], we adopt a more efficient *coreset selection* approach. See supplementary material.

3.5. Logits Self-correction

To counter the shifts on the *semantic* side after token condensation, we introduce a logits self-correction mechanism that leverages domain anchor tokens stored in \mathfrak{R} . In particular, the visual $\langle \text{cls} \rangle$ token of the current sample (if using CLIP) $\mathbf{V}_t^{\text{cls}} \in \mathbb{R}^{L \times D_v}$ is compared with the stored domain anchor tokens $\mathcal{A} = \{\mathbf{A}_{i,c}^{\text{cls}}\}_{i=1}^M$. The cosine similarity between these cross-layer tokens serves as a token-level classifier, which provides auxiliary information to adjust the predicted probability $\mathbf{p}_{t,c}$ from a *visual* perspective:

$$\begin{aligned} \tilde{\mathbf{p}}_{t,c} &= \mathbf{p}_{t,c} + \lambda \mathbf{p}_{t,c}^{\text{token}}, \\ \mathbf{p}_{t,c}^{\text{token}} &= \frac{1}{M} \sum_{i=1}^M \cos(\mathbf{V}_t^{\text{cls}}, \mathbf{A}_{i,c}^{\text{cls}}) \cdot \mathbf{P} \cdot \mathbb{1}_c, \end{aligned} \quad (9)$$

where λ is the logits correction weight. $\mathbb{1}_c \in \mathbb{R}^C$ the one-hot vector for the c -the class. The layer-specific exponential scaling coefficients are denoted as $\mathbf{P} = [\exp(\frac{l}{\beta})]_{l=1}^L \in \mathbb{R}^L$, where β controls the influence of different layers. We show that this correction temperature β provides semantic interpretability, as further discussed in Sec. 4.3. This self-correction mechanism ensures better alignment between the visual and semantic contexts, improving robustness in handling semantic shifts.

4. Experiments

4.1. Experimental Setup

Datasets. Following prior works, we evaluate our method on the cross-dataset (CD) benchmark, which measures model performance on unseen classes across 10 datasets: Aircraft [41], Caltech101 [17], Cars [31], DTD [13], EuroSAT [22], Flower102 [45], Food101 [4], Pets [47], SUN397 [75], and UCF101 [58]. Besides, we assess TCA’s robustness to distribution shifts using CIFAR-100-Corrupted (CIFAR-100-C) [23].

Baselines. We compare TCA with existing approaches across four categories: (1) *Prompt-tuning methods* like CoOp [85] and CoCoOp [84], which require multi-epoch adaptation; (2) *Conventional online test-time adaptation (TTA) methods* such as Tent [61] and SAR [46]. Tent updates batch normalization layers, while SAR further incorporates sharpness-aware minimization for reliable model updates. Following [15], we reran these experiments with adjusted batch sizes to align with our settings; (3) *Test-time prompting methods*, including TPT [57], C-TPT [77], and Diff-TPT [18], as well as TTA methods for CLIP such as MTA [79] and TDA [29]; and (4) *Token pruning and merging methods for ViTs*, such as EViT [34], ToMe [3], and ATS [16]. As ATS is an adaptive token pruning method with no fixed budget, we constrain its computational cost by an upper bound to ensure fair comparison.

Implementation Details. For CLIP, we use its official prompts. The batch size is set to 1, without data augmentations, to mimic realistic deployment scenarios. All experiments are conducted using pre-trained CLIP models with ViT-B/16 and ViT-L/14 architectures as the visual backbone. We set K to 2. For SigLIP and SigLIP v2, we select models with ViT-B/16 as the backbone. Since the SigLIP series learns visual embeddings differently (*i.e.*, via MAP rather than relying on the last layer’s $\langle \text{cls} \rangle$ token), we adapt our method accordingly. Specifically, we use pooled attention weights to obtain the domain anchor and the token attentiveness score (S_i^{head}) for cross-head token reduction and logits self-correction. Unless otherwise stated, all ablation studies are conducted on CLIP with ViT-B/16. Notably, our method is training-free, enabling rapid adaptation with minimal hyperparameter dependency. All experiments are performed on a single NVIDIA RTX A6000 GPU.

4.2. Main Results

Tab. 1 presents the results for fine-grained cross-dataset benchmark using the ViT-B/16 architecture on CLIP. As observed in Fig. 2a, the core idea behind TCA is that condensing inattentive tokens can effectively mitigate distribution shifts caused by visual-text misalignment. This concept is first validated by the improved performance of token pruning baselines over CLIP inference, where a condensed to-




Figure 4. Impact of reservoir size and GFLOPs budgets.

ken set yields a 0.9% increase in average accuracy when $R = 0.9$. TCA further enhances its performance by dealing with visual-text misalignment, moving visual features toward historical domain anchor tokens from DTR. As a result, TCA achieves an average accuracy of 68.69%, outperforming both train-required and training-free baselines without augmentation. Conventional TTA methods perform poorly on all datasets even with the requirement of fine-tuning a large amount of learnable parameters. In contrast, prompt-tuning methods, although requiring fewer learnable parameters, rely heavily on augmentation and struggle to effectively handle visual shifts. While TDA is a training-free method, it requires a large number of hyperparameters (a total of 10 for managing positive and negative caches) to achieve optimal performance. On the other hand, TCA uses significantly fewer hyperparameters and delivers a 1.72% improvement in average accuracy over TDA, with approximately 12.2% fewer GFLOPs.

4.3. Analysis and Ablations

Results on Various Severity Levels. To evaluate the effectiveness of TCA across different shift levels, we conduct experiments on five severity levels of three corruption types in the CIFAR-100-Corrupted dataset. Our results in Tab. 2 reveal a performance decline for EViT across most severity levels, particularly for snow and brightness corruptions. This may be due to the CLIP model embedding perturbed samples differently when token reduction is applied, leading to misinterpretations. In contrast, our proposed TCA consistently demonstrates improvements, highlighting the effectiveness of the DTR module in preserving informative and reliable token representations for logits correction.

Results on Multimodal VLMs. To evaluate the generalizability of our proposed method, we conduct additional experiments on SigLIP and SigLIP v2. As shown in Tab. 3, TCA consistently improves the performance of both SigLIP and SigLIP v2 without requiring model tuning. Notably, EuroSAT achieves a significant gain of 21.12% over SigLIP direct inference, aligning with trends observed in Tab. 1. On average, TCA yields improvements of 2.21% for SigLIP and 1.32% for SigLIP v2 across datasets, demonstrating its

Table 1. Results on the cross-dataset benchmark using CLIP ViT-B/16, including the number of learnable parameters (L-Param.) for learning-based TTA methods. * denotes the averaged GFLOPs across all datasets. The best performance (aug-free) is **bolded**.

Method	Aug-free	Aircraft	Caltech101	Cars	DTD	EuroSAT	Flower102	Food101	Pets	SUN397	UCF101	Average	GFLOPs	L-Param.
CLIP	✓	23.22	93.55	66.11	45.04	50.42	66.99	82.86	86.92	65.63	65.16	64.59	17.59	0
CoOp	✗	18.47	93.70	64.51	41.92	46.39	68.71	85.30	89.14	64.15	66.55	63.88	17.59	2048
CoCoOp	✗	22.29	93.79	64.90	45.45	39.23	70.85	83.97	90.46	66.89	68.44	64.63	17.59	34,816
Tent	✓	8.97	93.39	62.69	39.78	20.85	61.23	83.70	87.76	65.30	66.93	59.06	17.59	40,960
SAR	✓	21.09	91.85	61.15	44.68	46.19	63.54	81.43	87.95	59.74	65.58	62.32	17.59	31,744
TPT	✗	24.78	94.16	66.87	47.75	42.44	68.98	84.67	87.79	65.50	68.04	65.10	1108.61	2048
Diff-TPT	✗	25.60	92.49	67.01	47.00	43.13	70.10	87.23	88.22	65.74	62.67	65.47	-	-
C-TPT	✗	23.90	94.10	66.70	46.80	48.70	69.90	84.50	87.40	66.00	66.70	65.47	1108.61	2048
MTA	✗	25.32	94.21	68.47	45.90	45.36	68.06	85.00	88.24	66.67	68.11	65.53	-	-
TDA	✓	23.91	94.24	67.28	47.40	58.00	71.42	86.14	88.63	67.62	70.66	67.53	17.59	0
EViT _{R=0.9}	✓	24.12	92.25	64.57	45.09	48.41	70.24	84.99	88.96	64.58	68.46	65.17	15.41	0
ToME _{R=0.9}	✓	24.66	92.49	63.10	44.92	48.64	69.22	85.04	87.90	64.22	68.62	64.88	15.31	0
ATS _{R=0.9}	✓	22.86	92.21	57.90	40.96	40.62	67.52	80.16	85.34	61.53	67.22	61.63	11.15*	0
EViT _{R=0.7}	✓	23.31	91.20	58.44	43.32	43.26	67.11	79.70	85.77	61.41	66.69	62.02	11.62	0
ToME _{R=0.7}	✓	22.26	90.79	55.48	42.32	40.12	64.11	79.36	84.19	60.66	63.97	60.33	11.45	0
ATS _{R=0.7}	✓	17.28	85.40	33.65	36.52	27.79	52.62	55.97	72.94	48.82	56.44	48.74	8.76*	0
TCA _{R=0.9}	✓	24.87	93.63	65.33	46.16	70.43	73.33	85.31	89.53	65.92	72.38	68.69	15.45 <small>-12.2%</small>	0
TCA _{R=0.7}	✓	23.19	92.13	58.15	44.50	61.63	69.79	79.99	85.99	61.89	67.38	64.46	11.69 <small>-33.5%</small>	0

Table 2. Improvements over CLIP inference on CIFAR-100-C.

Severity	Contrast			Snow			Brightness		
	CLIP	EViT	Ours	CLIP	EViT	Ours	CLIP	EViT	Ours
1	31.90	-3.17%	+2.16%	35.34	-2.49%	+1.50%	41.00	-2.10%	+2.10%
2	20.67	+0.68%	+3.58%	29.72	-1.99%	+1.14%	41.44	-1.54%	+2.85%
3	15.05	-1.53%	+7.04%	29.14	-1.68%	+1.61%	41.83	-1.82%	+2.37%
4	8.85	-0.34%	+11.41%	27.04	-2.51%	+1.48%	41.12	-1.82%	+4.74%
5	2.69	+1.49%	+18.59%	24.85	-0.80%	+3.30%	38.10	-2.13%	+4.75%

potential. While minor degradations appear in Food101, overall performance remains stable or improved across most datasets, reinforcing TCA’s robustness.

Impact of GFLOPs Budget. We evaluate TCA under different GFLOPs budgets: $R = \{0.6, 0.7, 0.8, 0.9\}$, resulting in GFLOPs of 9.91, 11.68, 13.27, and 15.45, respectively, compared to the baseline ($R = 1$, 17.58 GFLOPs). As shown in Fig. 4b, condensing inattentive tokens can even enhance performance on certain datasets, notably Pets, and EuroSAT. Specifically for EuroSAT, when $R = 0.9$, the model’s adaptation performance is significantly improved, aligning with our findings in Fig. 2a. However, excessively aggressive pruning budgets (e.g., GFLOPs less than 13) lead to significant performance degradation across all datasets. This occurs since higher pruning rates may inadvertently remove informative tokens, causing irreversible harm in training-free scenarios where we lack the ability to update the model for further correction.

Impact of Reservoir Saving Strategy. In Tab. 4, we examine the performance changes across different reservoir saving strategies. We compare several approaches: *First-In-First-Out (FIFO)*; an *uncertainty-based* strategy, which discards the most uncertain sample when the reservoir reaches

capacity; a *similarity-enforced* strategy, where samples with high certainty and high cosine similarity to the saved samples are preferred; and a *diversity-enforced* strategy, which prioritizes saving prototypes that contain distinct tokens compared to those already stored. Our results show that the FIFO strategy performs poorly on Flower102 and EuroSAT, likely because CLIP’s low confidence leads to retaining misclassified samples. Conversely, Pets has high CLIP zero-shot accuracy (86.91% in Tab. 1), which makes FIFO acceptable. Among all strategies, the diversity-based approach consistently achieves the best performance. This is intuitive, as it maintains a representative set of features by capturing dataset diversity, whereas entropy-based methods may store homogenous features and overlook multiple class prototypes. By prioritizing diversity, our method ensures that a more representative set of features is maintained, leading to more robust performance across datasets.

Impact of Reservoir Size M . We assess the effectiveness of TCA across various reservoir sizes M on Pets, Flower102, and EuroSAT datasets, as illustrated in Fig. 4a. Remarkably, although the best performances are achieved at different reservoir sizes for different datasets, our TCA consistently maintains stable and high performance across a wide range of M values. This showcases the robustness and flexibility of TCA with respect to different reservoir budgets. Notably, even under extreme conditions with a minimal reservoir size (i.e., $M = 1$), our strategy significantly surpasses the strongest baseline method, TDA, by a large proportion on the EuroSAT dataset (14.2%).

Impact of Component. The impact of the domain anchor

Table 3. Improvements in Cross-Dataset benchmark over SigLIP and SigLIP v2 inference on ViT-B/16.

Method	Aircraft	Caltech101	Cars	DTD	EuroSAT	Flower102	Food101	Pets	SUN397	UCF101	Average
SigLIP	40.50	97.44	90.71	62.83	39.86	84.13	89.06	93.10	69.65	70.84	73.81
+TCA _{R=0.9}	+0.74%	+0.58%	+0.12%	+2.44%	+21.12%	+1.39%	-0.17%	+0.56%	+2.02%	+3.40%	+2.21%
SigLIP v2	50.05	97.61	93.36	62.41	42.09	85.30	90.14	94.88	73.34	72.09	76.13
+TCA _{R=0.9}	+0.42%	+0.25%	+0.11%	+2.76%	+11.17%	+0.39%	-0.04%	+0.02%	+0.64%	+3.22%	+1.32%

Table 4. Impact of reservoir updating strategy with CLIP.

	FIFO	Uncertainty	Similarity-enf	Diversity-enf
Pets	89.21	89.18	88.91	89.53
Flower102	70.65	73.28	72.15	73.33
EuroSAT	50.28	70.20	68.48	70.43

Table 5. Ablation study of the proposed components in TCA.

$A_{c^*}^l$	S^{head}	Food101	Pets	EuroSAT
		85.25	88.77	68.14
✓		85.31	88.96	67.69
	✓	85.25	89.23	67.14
✓	✓	85.31	89.53	70.43

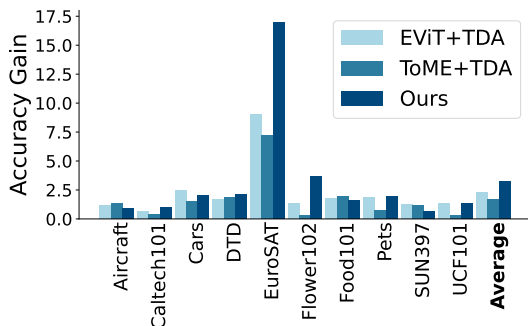


Figure 5. Accuracy gain of TCA and TDA + EViT/ToME over TDA combined with ATS.

tokens $A_{c^*}^l$ and the head-wise sorting score S^{head} (Sec. 3.4) is presented in Tab. 5. We observe that each component individually contributes to performance improvements. On the Food101 and Pets datasets, incorporating either component yields measurable gains in accuracy. By leveraging historical domain anchor tokens, the model acquires rich contextual information, enhancing the stability of token importance over time. Simultaneously, cross-head token sorting ensures that token pruning decisions are more robust by accounting for consensus across attention heads. An intriguing case arises with the EuroSAT dataset. Here, the baseline performance without any components is 68.14%. Applying either component alone results in a slight performance decrease. However, when both components are used together, performance significantly improves to 70.43%. This outcome emphasizes the necessity of combining historical domain anchor tokens and cross-head token sorting to fully realize the model’s potential.




Figure 6. Examples of our token condensation.

Comparison with TDA + Token Condensation. We compare TDA combined with token pruning and merging methods, and report the performance gains over TDA + ATS in Fig. 5. While TDA alone offers improvements, it depends heavily on negative caching and a large hyperparameter set. In contrast, TCA consistently outperforms TDA + EViT and TDA + ToME across multiple datasets and on average, despite using only minimal hyperparameter settings, demonstrating its superior adaptability.

Visualization. We visualize the pruned and merged tokens of different ViT layers in Fig. 6. The black mask indicates pruned regions while different colors are set for different merging clusters. As token condensation progresses, non-discriminative tokens are gradually removed, leading to better alignment with the text semantics.

5. Conclusion

We introduced Token Condensation as Adaptation (TCA), a novel training-free test-time adaptation method on VLMs such as CLIP. Our experiments across various VLMs demonstrated that token condensation significantly enhances visual-text alignment while also serving as an interpretation of visual semantics. Besides, TCA reduces GFLOPs as a beneficial byproduct, improving computational efficiency. Despite our empirical findings, a rigorous theory of CLIP’s generalizability remains underexplored. To bridge the gap, we provide a theoretical analysis, examine TCA’s applicability to other VLMs, and acknowledge its limitations in the supplementary material.

Acknowledgment

This work was partially supported by ARC DE240100105, DP240101814, DP230101196, DP230101753, BA24006, the ARC Training Centre for Information Resilience IC20 0100022.

References

- [1] Lei Jimmy Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. Layer normalization. *CoRR*, abs/1607.06450, 2016. 2
- [2] Filippo Maria Bianchi, Daniele Grattarola, and Cesare Alippi. Spectral clustering with graph neural networks for graph pooling. In *ICML*, pages 874–883. PMLR, 2020. 5
- [3] Daniel Bolya, Cheng-Yang Fu, Xiaoliang Dai, Peizhao Zhang, Christoph Feichtenhofer, and Judy Hoffman. Token merging: Your vit but faster. In *ICLR*. OpenReview.net, 2023. 3, 5, 6
- [4] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101 - mining discriminative components with random forests. In *ECCV*, pages 446–461. Springer, 2014. 6
- [5] Dian Chen, Dequan Wang, Trevor Darrell, and Sayna Ebrahimi. Contrastive test-time adaptation. In *CVPR*, pages 295–305. IEEE, 2022. 2
- [6] Zhi Chen, Yadan Luo, Ruihong Qiu, Sen Wang, Zi Huang, Jingjing Li, and Zheng Zhang. Semantics disentangling for generalized zero-shot learning. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 1
- [7] Zhi Chen, Yadan Luo, Sen Wang, Jingjing Li, and Zi Huang. Gsmflow: Generation shifts mitigating flow for generalized zero-shot learning. *IEEE Transactions on Multimedia*, 2022. 1
- [8] Zhuoxiao Chen, Yadan Luo, Zheng Wang, Mahsa Baktashmotlagh, and Zi Huang. Revisiting domain-adaptive 3d object detection by reliable, diverse and class-balanced pseudo-labeling. In *ICCV*, pages 3691–3703. IEEE, 2023. 1
- [9] Zhi Chen, Pengfei Zhang, Jingjing Li, Sen Wang, and Zi Huang. Zero-shot learning by harnessing adversarial samples. In *Proceedings of the 31th ACM International Conference on Multimedia 2023*, 2023. 1
- [10] Zhuoxiao Chen, Zixin Wang, Yadan Luo, Sen Wang, and Zi Huang. DPO: dual-perturbation optimization for test-time adaptation in 3d object detection. In *ACM MM*, pages 4138–4147. ACM, 2024. 1
- [11] Zhuoxiao Chen, Junjie Meng, Mahsa Baktashmotlagh, Yonggang Zhang, Zi Huang, and Yadan Luo. MOS: model synergy for test-time adaptation on lidar-based 3d object detection. In *ICLR*, 2025. 1
- [12] Zhi Chen, Zecheng Zhao, Jingcai Guo, Jingjing Li, and Zi Huang. Svip: Semantically contextualized visual patches for zero-shot learning. In *ICCV2025*, 2025. 1
- [13] Mircea Cimpoi, Subhransu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. In *CVPR*, pages 3606–3613. IEEE Computer Society, 2014. 6
- [14] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255. IEEE Computer Society, 2009. 3
- [15] Mario Döbler, Robert A. Marsden, Tobias Raichle, and Bin Yang. A lost opportunity for vision-language models: A comparative study of online test-time adaptation for vision-language models. *CoRR*, abs/2405.14977, 2024. 1, 2, 6
- [16] Mohsen Fayyaz, Soroush Abbasi Koohpayegani, Farnoush Rezaei Jafari, Sunando Sengupta, Hamid Reza Vaezi Joze, Eric Sommerlade, Hamed Pirsiavash, and Jürgen Gall. Adaptive token sampling for efficient vision transformers. In *ECCV*, pages 396–414. Springer, 2022. 3, 6
- [17] Li Fei-Fei, Robert Fergus, and Pietro Perona. Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. *Comput. Vis. Image Underst.*, 106(1):59–70, 2007. 6
- [18] Chun-Mei Feng, Kai Yu, Yong Liu, Salman Khan, and Wangmeng Zuo. Diverse data augmentation with diffusions for effective test-time prompt tuning. In *ICCV*, pages 2704–2714. IEEE, 2023. 1, 2, 6
- [19] Pierre Foret, Ariel Kleiner, Hossein Mobahi, and Behnam Neyshabur. Sharpness-aware minimization for efficiently improving generalization. In *ICLR*. OpenReview.net, 2021. 1, 2
- [20] Taesik Gong, Jongheon Jeong, Taewon Kim, Yewon Kim, Jinwoo Shin, and Sung-Ju Lee. NOTE: robust continual test-time adaptation against temporal correlation. In *NeurIPS*, 2022. 2
- [21] Sachin Goyal, Mingjie Sun, Aditi Raghunathan, and J. Zico Kolter. Test time adaptation via conjugate pseudo-labels. In *NeurIPS*, 2022. 2
- [22] Patrick Helber, Benjamin Bischke, Andreas Dengel, and Damian Borth. Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. *IEEE J. Sel. Top. Appl. Earth Obs. Remote. Sens.*, 12(7):2217–2226, 2019. 6
- [23] Dan Hendrycks and Thomas G. Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. In *ICLR*, 2019. 6
- [24] Dan Hendrycks, Kevin Zhao, Steven Basart, Jacob Steinhardt, and Dawn Song. Natural adversarial examples. *CoRR*, abs/1907.07174, 2019. 3
- [25] Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, Dawn Song, Jacob Steinhardt, and Justin Gilmer. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *ICCV*, pages 8320–8329. IEEE, 2021. 3
- [26] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, pages 448–456. JMLR.org, 2015. 2
- [27] Yusuke Iwasawa and Yutaka Matsuo. Test-time classifier adjustment module for model-agnostic domain generalization. In *NeurIPS*, pages 2427–2440, 2021. 2, 3
- [28] Minguk Jang, Sae-Young Chung, and Hye Won Chung. Test-time adaptation via self-training with nearest neighbor information. In *ICLR*. OpenReview.net, 2023. 2
- [29] Adilbek Karmanov, Dayan Guan, Shijian Lu, Abdulmotaleb El-Saddik, and Eric P. Xing. Efficient test-time adaptation of vision-language models. *CoRR*, abs/2403.18293, 2024. 2, 6

- [30] Zhenglun Kong, Peiyan Dong, Xiaolong Ma, Xin Meng, Wei Niu, Mengshu Sun, Xuan Shen, Geng Yuan, Bin Ren, Hao Tang, Minghai Qin, and Yanzhi Wang. Spvit: Enabling faster vision transformers via latency-aware soft token pruning. In *ECCV*, pages 620–640. Springer, 2022. 2
- [31] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *ICCV Workshops*, pages 554–561. IEEE Computer Society, 2013. 6
- [32] Jae-Hong Lee and Joon-Hyuk Chang. Stationary latent weight inference for unreliable observations from online test-time adaptation. In *ICML*. OpenReview.net, 2024. 1
- [33] Jian Liang, Ran He, and Tieniu Tan. A comprehensive survey on test-time adaptation under distribution shifts. *CoRR*, abs/2303.15361, 2023. 1
- [34] Youwei Liang, Chongjian Ge, Zhan Tong, Yibing Song, Jue Wang, and Pengtao Xie. Evit: Expediting vision transformers via token reorganizations. In *ICLR*. OpenReview.net, 2022. 2, 3, 4, 6
- [35] Jia Syuen Lim, Zhuoxiao Chen, Zhi Chen, Mahsa Baktashmotlagh, Xin Yu, Zi Huang, and Yadan Luo. Dipex: Dispersing prompt expansion for class-agnostic object detection. In *NeurIPS*, 2024. 1
- [36] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. In *NeurIPS*, 2023. 3
- [37] Jiaming Liu, Senqiao Yang, Peidong Jia, Renrui Zhang, Ming Lu, Yandong Guo, Wei Xue, and Shanghang Zhang. Vida: Homeostatic visual domain adapter for continual test time adaptation. In *ICLR*. OpenReview.net, 2024. 2
- [38] Zichen Liu, Hongbo Sun, Yuxin Peng, and Jiahuan Zhou. DART: dual-modal adaptive online prompting and knowledge retention for test-time adaptation. In *AAAI*, pages 14106–14114. AAAI Press, 2024. 2
- [39] Yadan Luo, Zijian Wang, Zhuoxiao Chen, Zi Huang, and Mahsa Baktashmotlagh. Source-free progressive graph learning for open-set domain adaptation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 45(9):11240–11255, 2023. 1
- [40] Xiaosong Ma, Jie Zhang, Song Guo, and Wenchao Xu. Swapprompt: Test-time prompt adaptation for vision-language models. In *NeurIPS*, 2023. 2
- [41] Subhransu Maji, Esa Rahtu, Juho Kannala, Matthew B. Blaschko, and Andrea Vedaldi. Fine-grained visual classification of aircraft. *CoRR*, abs/1306.5151, 2013. 6
- [42] Robert A. Marsden, Mario Döbler, and Bin Yang. Universal test-time adaptation through weight ensembling, diversity weighting, and prior correction. In *WACV*, pages 2543–2553. IEEE, 2024. 2
- [43] Prasanna Mayilvahanan, Thaddäus Wiedemer, Evgenia Rusak, Matthias Bethge, and Wieland Brendel. Does clip’s generalization performance mainly stem from high train-test similarity? In *ICLR*, 2024. 1
- [44] Lingchen Meng, Hengduo Li, Bor-Chun Chen, Shiyi Lan, Zuxuan Wu, Yu-Gang Jiang, and Ser-Nam Lim. Advit: Adaptive vision transformers for efficient image recognition. In *CVPR*, pages 12299–12308. IEEE, 2022. 2
- [45] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *ICVGIP*, pages 722–729. IEEE Computer Society, 2008. 6
- [46] Shuaicheng Niu, Jiayang Wu, Yifan Zhang, Zhiqian Wen, Yaofu Chen, Peilin Zhao, and Mingkui Tan. Towards stable test-time adaptation in dynamic wild world. In *ICLR*. OpenReview.net, 2023. 1, 2, 6
- [47] Omkar M. Parkhi, Andrea Vedaldi, Andrew Zisserman, and C. V. Jawahar. Cats and dogs. In *CVPR*, pages 3498–3505. IEEE Computer Society, 2012. 6
- [48] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763. PMLR, 2021. 1, 2, 4
- [49] Maithra Raghu, Thomas Unterthiner, Simon Kornblith, Chiyuan Zhang, and Alexey Dosovitskiy. Do vision transformers see like convolutional neural networks? In *NeurIPS*, pages 12116–12128, 2021. 2
- [50] Yongming Rao, Wenliang Zhao, Benlin Liu, Jiwen Lu, Jie Zhou, and Cho-Jui Hsieh. Dynamicvit: Efficient vision transformers with dynamic token sparsification. In *NeurIPS*, pages 13937–13949, 2021. 2
- [51] Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishaal Shankar. Do imagenet classifiers generalize to imagenet? In *ICML*, pages 5389–5400. PMLR, 2019. 3
- [52] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, pages 10674–10685. IEEE, 2022. 2
- [53] Michael S. Ryoo, A. J. Piergiovanni, Anurag Arnab, Mostafa Dehghani, and Anelia Angelova. Tokenlearner: Adaptive space-time tokenization for videos. In *NeurIPS*, pages 12786–12797, 2021. 2
- [54] Jameel Abdul Samadh, Hanan Gani, Noor Hussein, Muhammad Uzair Khattak, Muzammal Naseer, Fahad Shahbaz Khan, and Salman H. Khan. Align your prompts: Test-time prompting with distribution alignment for zero-shot generalization. In *NeurIPS*, 2023. 1
- [55] Ozan Sener and Silvio Savarese. Active learning for convolutional neural networks: A core-set approach. In *ICLR*. OpenReview.net, 2018. 1
- [56] Yuzhang Shang, Mu Cai, Bingxin Xu, Yong Jae Lee, and Yan Yan. Llava-prumerge: Adaptive token reduction for efficient large multimodal models. *CoRR*, abs/2403.15388, 2024. 3
- [57] Manli Shu, Weili Nie, De-An Huang, Zhiding Yu, Tom Goldstein, Anima Anandkumar, and Chaowei Xiao. Test-time prompt tuning for zero-shot generalization in vision-language models. In *NeurIPS*, 2022. 2, 6
- [58] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. UCF101: A dataset of 101 human actions classes from videos in the wild. *CoRR*, abs/1212.0402, 2012. 6
- [59] Devavrat Tomar, Guillaume Vray, Behzad Bozorgtabar, and Jean-Philippe Thiran. Tesla: Test-time self-learning with automatic adversarial augmentation. In *CVPR*, pages 20341–20350. IEEE, 2023. 2
- [60] Michael Tschannen, Alexey Gritsenko, Xiao Wang, Muhammad Ferjad Naeem, Ibrahim Alabdulmohsin, Nikhil Parthasarathy, Talfan Evans, Lucas Beyers, Ye Xia, Basil

- Mustafa, et al. Siglip 2: Multilingual vision-language encoders with improved semantic understanding, localization, and dense features. *arXiv:2502.14786*, 2025. 1, 4
- [61] Dequan Wang, Evan Shelhamer, Shaoteng Liu, Bruno A. Olshausen, and Trevor Darrell. Tent: Fully test-time adaptation by entropy minimization. In *ICLR*. OpenReview.net, 2021. 2, 6
- [62] Haohan Wang, Songwei Ge, Zachary C. Lipton, and Eric P. Xing. Learning robust global representations by penalizing local predictive power. In *NeurIPS*, pages 10506–10518, 2019. 3
- [63] Hongjie Wang, Bhishma Dedhia, and Niraj K. Jha. Zero-truncate: Zero-shot token pruning through leveraging of the attention graph in pre-trained transformers. *CoRR*, abs/2305.17328, 2023. 2
- [64] Qin Wang, Olga Fink, Luc Van Gool, and Dengxin Dai. Continual test-time domain adaptation. In *CVPR*, pages 7191–7201. IEEE, 2022. 2
- [65] Qizhou Wang, Yong Lin, Yongqiang Chen, Ludwig Schmidt, Bo Han, and Tong Zhang. A sober look at the robustness of clips to spurious features. In *NeurIPS*, 2024. 1
- [66] Ran Wang, Hua Zuo, Zhen Fang, and Jie Lu. Towards robustness prompt tuning with fully test-time adaptation for clip’s zero-shot generalization. In *ACM Multimedia 2024*, 2024. 2
- [67] Shuai Wang, Daoan Zhang, Zipei Yan, Jianguo Zhang, and Rui Li. Feature alignment and uniformity for test time adaptation. In *CVPR*, pages 20050–20060. IEEE, 2023. 2
- [68] Zixin Wang, Yadan Luo, Zhi Chen, Sen Wang, and Zi Huang. Cal-sfda: Source-free domain-adaptive semantic segmentation with differentiable expected calibration error. In *ACM MM*, pages 1167–1178. ACM, 2023. 1
- [69] Zixin Wang, Yadan Luo, Liang Zheng, Zhuoxiao Chen, Sen Wang, and Zi Huang. In search of lost online test-time adaptation: A survey. *CoRR*, abs/2310.20199, 2023. 1, 2
- [70] Siyuan Wei, Tianzhu Ye, Shen Zhang, Yao Tang, and Jiajun Liang. Joint token pruning and squeezing towards more aggressive compression of vision transformers. In *CVPR*, pages 2092–2101. IEEE, 2023. 3
- [71] Ronald J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Mach. Learn.*, 8:229–256, 1992. 2
- [72] Gert W. Wolf. Facility location: concepts, models, algorithms and case studies. series: Contributions to management science. *Int. J. Geogr. Inf. Sci.*, 25(2):331–333, 2011. 1
- [73] Junran Wu, Xueyuan Chen, Ke Xu, and Shangzhe Li. Structural entropy guided graph hierarchical pooling. In *ICML*, pages 24017–24030. PMLR, 2022. 5
- [74] Yuxin Wu and Kaiming He. Group normalization. In *ECCV*, pages 3–19. Springer, 2018. 2
- [75] Jianxiong Xiao, James Hays, Krista A. Ehinger, Aude Oliva, and Antonio Torralba. SUN database: Large-scale scene recognition from abbey to zoo. In *CVPR*, pages 3485–3492. IEEE Computer Society, 2010. 6
- [76] Yifan Xu, Zhijie Zhang, Mengdan Zhang, Kekai Sheng, Ke Li, Weiming Dong, Liqing Zhang, Changsheng Xu, and Xing Sun. Evo-vit: Slow-fast token evolution for dynamic vision transformer. In *AAAI*, pages 2964–2972. AAAI Press, 2022. 2
- [77] Hee Suk Yoon, Eunseop Yoon, Joshua Tian Jin Tee, Mark A. Hasegawa-Johnson, Yingzhen Li, and Chang D. Yoo. C-TPT: calibrated test-time prompt tuning for vision-language models via text feature dispersion. In *ICLR*. OpenReview.net, 2024. 2, 3, 6
- [78] Longhui Yuan, Binhui Xie, and Shuang Li. Robust test-time adaptation in dynamic scenarios. In *CVPR*, pages 15922–15932. IEEE, 2023. 2
- [79] Maxime Zanella and Ismail Ben Ayed. On the test-time zero-shot generalization of vision-language models: Do we really need prompt learning? *CoRR*, abs/2405.02266, 2024. 2, 6
- [80] Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language image pre-training. In *ICCV*, 2023. 1
- [81] Dingchu Zhang, Zhi Zhou, and Yufeng Li. Robust test-time adaptation for zero-shot prompt tuning. In *AAAI*, pages 16714–16722. AAAI Press, 2024. 2
- [82] Marvin Zhang, Sergey Levine, and Chelsea Finn. MEMO: test time robustness via adaptation and augmentation. In *NeurIPS*, 2022. 1, 2
- [83] Shuai Zhao, Xiaohan Wang, Linchao Zhu, and Yi Yang. Test-time adaptation with CLIP reward for zero-shot generalization in vision-language models. In *ICLR*. OpenReview.net, 2024. 2
- [84] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. In *CVPR*, pages 16795–16804. IEEE, 2022. 6
- [85] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *Int. J. Comput. Vis.*, 130(9):2337–2348, 2022. 6
- [86] Zhuofan Zong, Kunchang Li, Guanglu Song, Yali Wang, Yu Qiao, Biao Leng, and Yu Liu. Self-slimmed vision transformer. In *ECCV*, pages 432–448. Springer, 2022. 2