

Domain Aligned CLIP for Few-shot Classification

Muhammad Waleed Gondal, Jochen Gast, Inigo Alonso Ruiz, Richard Droste,
Tommaso Macri, Suren Kumar, Luitpold Staudigl
Amazon

{wgondal, jogast, inruiz, rdroste, tmacri, ssurkum, luitpold}@amazon.com

Abstract

Large vision-language representation learning models like CLIP have demonstrated impressive performance for zero-shot transfer to downstream tasks while largely benefiting from inter-modal (image-text) alignment via contrastive objectives. This downstream performance can further be enhanced by full-scale fine-tuning which is often compute intensive, requires large labelled data, and can reduce out-of-distribution (OOD) robustness. Furthermore, sole reliance on inter-modal alignment might overlook the rich information embedded within each individual modality. In this work, we introduce a sample-efficient domain adaptation strategy for CLIP, termed Domain Aligned CLIP (DAC), which improves both intra-modal (image-image) and inter-modal alignment on target distributions without fine-tuning the main model. For intra-modal alignment, we introduce a lightweight adapter that is specifically trained with an intra-modal contrastive objective. To improve inter-modal alignment, we introduce a simple framework to modulate the precomputed class text embeddings. The proposed few-shot fine-tuning framework is computationally efficient, robust to distribution shifts, and does not alter CLIP’s parameters. We study the effectiveness of DAC by benchmarking on 11 widely used image classification tasks with consistent improvements in 16-shot classification upon strong baselines by about 2.3% and demonstrate competitive performance on 4 OOD robustness benchmarks.

1. Introduction

The growing popularity of large-scale representation learning models for multi-modal data like CLIP [33], ALIGN [20] and Florence [50] has highlighted the need for efficiently adapting these models to various downstream tasks across multiple domains and applications. In CLIP, for instance, successful zero-shot transfer relies on adapting to both image and text modalities in new domains and aligning their representations (inter-modal alignment) in a shared representation space. However, the generalization

capability of CLIP is constrained by its pre-training distribution [8]. To enhance their performance, these models are often transferred to target distributions either through fine-tuning or employing various few-shot strategies. While fine-tuning can be resource intensive and prone to overfitting [32, 45], few-shot adaptation offers a training and sample efficient alternative.

In the few-shot setup, in addition to the class labels, we are also provided with a few labeled images from target distributions. These labeled images serve as data-specific priors which can be used to update CLIP’s existing inter-modal predictions. In this work, we present an approach called DAC (Domain Aligned CLIP) that effectively leverages this prior knowledge to adapt CLIP for new downstream tasks. We posit that a combination of improved intra-modal (image-image) and inter-modal (image-text) alignment in the target domain results in better few-shot transfer to downstream classification tasks. To achieve this, we split the overall classification task into an ensemble of intra- and inter-modal classifications. While the inter-modal classification leverages image-text similarity (as in CLIP), the intra-modal classification is performed by means of a visual cache that is composed of precomputed image embeddings of the few (seen) labelled images, *cf.* Fig. 1.

A similar ensembling framework was recently leveraged in Tip-Adapter [53] for few-shot CLIP adaptation. However, no explicit regularization is done to improve the intra-modal alignment. While keeping the inter-modal classification fixed, Tip-Adapter(-F) treats the visual cache as learnable parameters and optimizes them to learn the residual information required to improve the upstream classification performance. We show that such an optimization causes the visual cache to lose its diverse, rich visual information and deteriorates its discriminative capability as depicted by its intra-modal classification performance *cf.* Fig. 2. Hence, despite relying on an ensemble to exploit feature diversity, Tip-Adapter-F reduces this diversity and limits feature reuse which is crucial for transfer learning and robustness [27]. Furthermore, Tip-Adapter-F does not adapt textual features in the target domain which can be crucial as recent work

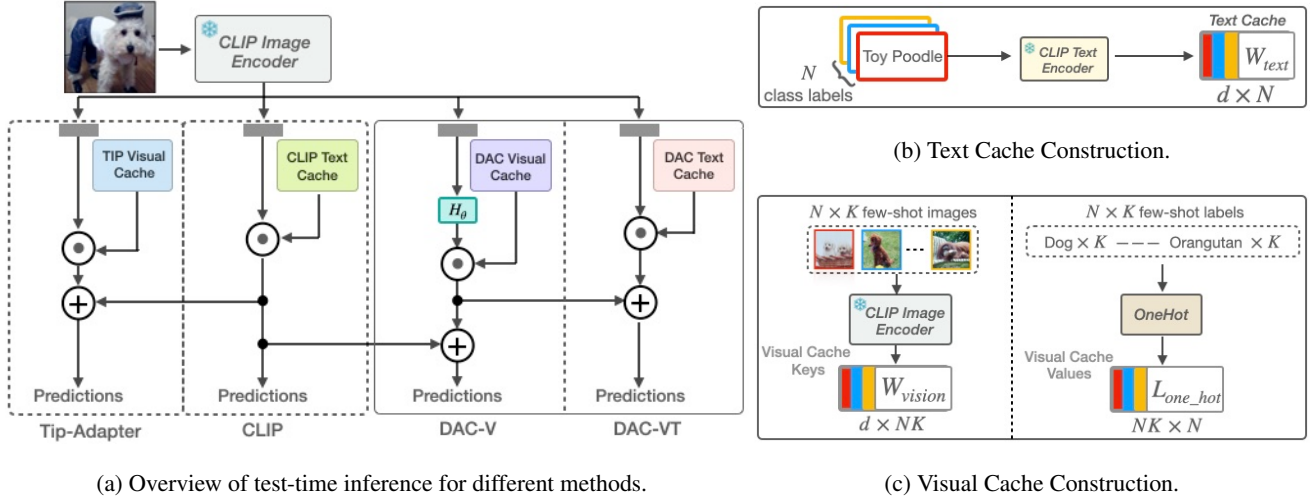


Figure 1. **Overview of CLIP and its few-shot adaptations.** (a) To make zero-shot predictions, CLIP uses precomputed text embeddings of class labels (termed as text cache) to determine inter-modal similarities between images and text embeddings. Tip-Adapter extends CLIP for few-shot prediction by compounding its inter-modal logits with intra-modal logits. The intra-modal logits stem from precomputed image embeddings of a few labeled images (termed as visual cache). DAC-V adapts image embeddings of the visual cache to the target distributions. In addition to the adaptation of image embeddings, DAC-VT adapts textual embeddings on the target distributions. The construction of text and visual caches from a few seen examples are shown in (b) and (c), respectively.

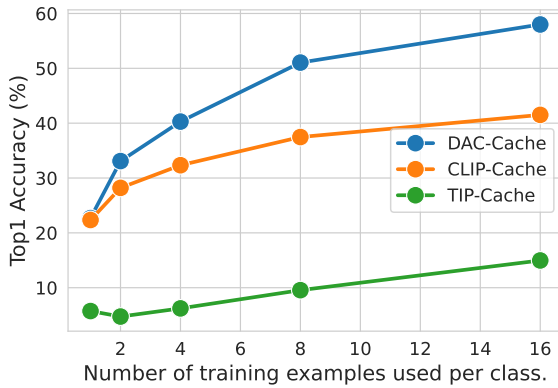


Figure 2. **Comparative analysis of different intra-modal classifiers using 16-shots visual-cache.** Here, TIP-Adapter’s cache loses the intra-modal classification advantage of CLIP. In contrast, the intra-modal classifier based on DAC visual features performs consistently better.

highlights the limitations of CLIP for inter-modal alignment [24] and how performance degrades as downstream vocabulary expands [35].

In this work, unlike Tip-Adapter(-F), we introduce a two stage adaptation strategy that focuses on individually improving both intra- and inter-modal classifiers. See Fig. 1 for an overview of methods. Our hypothesis is grounded in the widely acknowledged phenomenon that effective ensembles consist of models that are both accurate and make uncorrelated errors [12, 29, 44]. In the first stage of the

proposed setup, a linear adapter layer is exclusively trained by a (self-)supervised contrastive objective to contrast images sampled from different classes. The goal is to improve the affinity of the latent representations of images coming from the same class while keeping representations of images belonging to different classes distant. This results in an improved intra-modal classifier which forms the basis of DAC-V. In the second stage, we introduce a framework that directly optimizes CLIP’s text embeddings and improves inter-modal classifier performance while ensembling it with the frozen intra-modal classifier from the first stage. We call this overall framework DAC-VT where both visual and textual representations are adapted to the target distribution.

Our **primary contributions** are as follows:

- We present DAC, a novel framework for adapting CLIP for few-shot classification tasks that learns by explicitly aligning intra-modal and inter-modal representations on target distributions. To the best of our knowledge, this is the first work to leverage intra-modal regularization for few-shot adaptation of large vision-language models.
- We conduct comprehensive quantitative analysis on 11 widely used image classification benchmarks and show that our method outperforms competitive baselines, while maintaining reasonable robustness to distributions shifts (measured on 4 benchmarks).

2. Related Work

Learning rich representations of data that generalize well to multiple tasks is challenging but desirable [1]. Such representations not only enable sample-efficient transfer to downstream tasks but also simplify the hyper-parameter optimization [22]. To this end, contrastive learning [14, 42] based self-supervised methods have shown to be promising for learning transferable representations of images [4, 15, 39] and text [42]. More recently, these objectives have been extended to align data from different modalities in a joint representation space [7, 20, 33, 48, 50], achieving impressive zero-shot transfer learning performance on a number of downstream tasks. The performance of these models can further be improved by either fine-tuning them on the labeled target data [33, 45] or by adapting their feature space on the target distribution while maintaining feature re-useability [27]. The second strategy is practically more appealing, as it is sample efficient and only requires simple hyperparameter tuning [19]. For CLIP [33], such sample-efficient adaptation methods can be broadly classified into two categories; (1) methods that learn to optimize text prompts [25, 49, 52, 54, 55], and, (2) methods that introduce lightweight adapter layers to align image and text embeddings on target distributions [10, 13, 31, 53]. The latter setting offers more flexibility in adapting both visual and textual domains, where CLIP-Adapter [10] tunes adapter layers appended to CLIP’s frozen image and text encoders.

Differences to CLIP-Adapter: DAC differs from CLIP-Adapter in three ways: (1) DAC explicitly optimizes the intra-modal alignment of visual features, whereas optimization in CLIP-Adapter is solely geared towards inter-modal alignment. (2) DAC only uses a single linear-layer adapter for intra-modal alignment, whereas CLIP-Adapter uses two-layer MLPs on both visual and textual features for inter-modal alignment. (3) Inter-modal alignment in DAC is done by directly modulating the text cache which is much more efficient in comparison to using a separate adapter for text embeddings (as in CLIP-Adapter).

Differences to Tip-Adapter: Tip-Adapter [53] employs a visual caching structure to split the overall classification into an ensemble of intra-modal and inter-modal classifiers. However, note that the ensembling in Tip-Adapter yields inefficient sub-classifiers, limiting the full utilization of few-shot knowledge available. We propose two ways to improve it. First, we introduce an intra-modal contrastive learning framework to improve the visual alignment of features in the target domain. We extend the function contrastive objective used in [11] to approximate the proxy visual function-space of downstream classes. Second, we fine-tune CLIP’s precomputed textual embeddings to mitigate its limitations with unseen vocabulary [35] for class labels in the target domain. The framework is simpler than prompt tuning and does not require any additional parameters. Moreover, un-

like Tip-Adapter, our method eliminates the need for an extra sharpness parameter for tuning image similarity scores.

While efforts have been made to adapt CLIP features without fine-tuning at test-time at the cost of reduced in-distribution performance [36, 41], we specifically focus on few-shot fine-tuning for CLIP adaptation, noting that these methods complement our work. Additionally, recent work leverages pre-trained language models to generate additional category information [47, 51] and visual generative models to synthesize images for expanding few-shot training data [51]. Unlike these methods, DAC only uses the few-shot data provided with the task, making a comparison with these methods unfair.

3. Background

We start by explaining CLIP and how its zero-shot prediction can be formulated by means of a *text-cache*. We then expand on how this formulation is extended in Tip-Adapter to support few-shot classification tasks via a *visual-cache*.

Zero-shot Classification with a Text-Cache. CLIP is a vision-language representation learning model that aligns vision and text modalities in a joint embedding space by learning from image-text pairs (\mathbf{x}, \mathbf{t}) where \mathbf{x} are vectorized images and \mathbf{t} correspond to tokenized text inputs. At inference-time, CLIP encodes these input modalities into d -dimensional embeddings with separate encoders for image and text, *i.e.* $\mathbf{v}_{\text{image}} = \mathbf{E}_{\text{image}}(\mathbf{x})$ and $\mathbf{v}_{\text{text}} = \mathbf{E}_{\text{text}}(\mathbf{t})$. For brevity, we will refer to the L2-normalized embeddings as $\mathbf{z}_{\text{image}} = \mathbf{v}_{\text{image}} / \|\mathbf{v}_{\text{image}}\|$ and $\mathbf{z}_{\text{text}} = \mathbf{v}_{\text{text}} / \|\mathbf{v}_{\text{text}}\|$. Alignment between image and text embeddings is then computed via cosine similarity, *i.e.* $\text{sim}(\mathbf{v}_{\text{image}}, \mathbf{v}_{\text{text}}) := \mathbf{v}_{\text{image}}^T \mathbf{v}_{\text{text}} / (\|\mathbf{v}_{\text{image}}\| \|\mathbf{v}_{\text{text}}\|) = \mathbf{z}_{\text{image}}^T \mathbf{z}_{\text{text}}$. CLIP leverages this image-text alignment for zero-shot classification with novel inputs. Assume that a given task consists of N classification labels $\{y^{(i)}\}_i^N$. We first construct a precomputed weight matrix (or *text-cache*) by concatenating (normalized) text embeddings of all classification labels.

$$\mathbf{W}_{\text{text}} = \begin{pmatrix} \mathbf{z}_{\text{text}}^{(1)} & \mathbf{z}_{\text{text}}^{(2)} & \dots & \mathbf{z}_{\text{text}}^{(N)} \end{pmatrix} \in \mathbb{R}^{d \times N}, \quad (1)$$

which encapsulates pre-computed, textual knowledge associated with the task. Subsequently, the text-cache \mathbf{W}_{text} can be used to classify a new, unseen input image into N classes by computing *inter-modal* logits, *i.e.*

$$\text{logits}_{\text{CLIP}} = \mathbf{W}_{\text{text}}^T \mathbf{z}_{\text{image}} \in \mathbb{R}^{N \times 1}. \quad (2)$$

Note that \mathbf{W}_{text} only needs to be computed once per task.

Few-shot Classification with a Visual-Cache. Tip-Adapter [53] extends CLIP for few-shot classification. For each new task, it requires a few labeled training examples from a target distribution $\mathcal{D}_{\text{train}} = \{(x^{(i)}, y^{(i)})\}_i^{N \times K}$ where

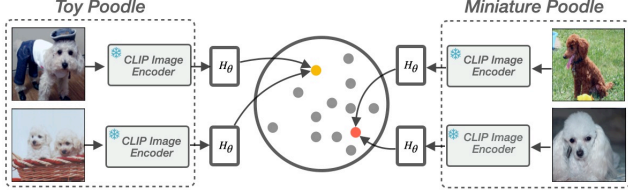


Figure 3. **Construction of supervised contrastive objective used to fine-tune the visual adapter layer in DAC.**

N is the number of classes and K is the number of examples (or shots) per class. Tip-Adapter encodes this few-shot knowledge into a precomputed *visual-cache* with separate cache keys and values. Akin to Eq. (1), cache keys are computed as

$$\mathbf{W}_{\text{image}} = \left(\mathbf{z}_{\text{image}}^{(1,1)} \quad \mathbf{z}_{\text{image}}^{(1,2)} \quad \dots \quad \mathbf{z}_{\text{image}}^{(K,N)} \right) \in \mathbb{R}^{d \times NK}, \quad (3)$$

where $\mathbf{W}_{\text{image}}$ is the concatenation of (sub-)weight matrices $\mathbf{z}_{\text{image}}^{(1,\cdot)} \mathbf{z}_{\text{image}}^{(2,\cdot)} \dots \mathbf{z}_{\text{image}}^{(K,\cdot)}$ per classification label horizontally. Corresponding cache values are then constructed as one-hot encodings $\mathbf{L}_{\text{one_hot}} \in \mathbb{R}^{NK \times N}$ of ground truth labels $\{y_i\}_i^{N \times K}$ by vertically concatenating one-hot encodings per shot, followed by horizontal concatenation per classification label. Note that such a key-value configuration effectively enables the visual-cache to retain all the available few-shot knowledge in $\mathcal{D}_{\text{train}}$. To update CLIP logits (Eq. 2) with the few-shot knowledge encoded in the visual-cache (Eq. 3), Tip-adapter introduces an affinity vector

$$\mathbf{w}_{\text{affinity}} = \exp \left(\beta \left(\mathbf{W}_{\text{image}}^T \mathbf{z}_{\text{image}} - 1 \right) \right) \in \mathbb{R}^{NK \times 1}, \quad (4)$$

where \exp denotes a pointwise exponential function and β modulates the sharpness of affinities. The affinity vector $\mathbf{w}_{\text{affinity}}$ retains the similarity (or compatibility) between a given image $\mathbf{z}_{\text{image}}$ and the images stored in the visual-cache $\mathbf{W}_{\text{image}}$. Tip-Adapter finally computes aggregated logits as

$$\text{logits}_{\text{TIP}} = \text{logits}_{\text{CLIP}} + \alpha \mathbf{L}_{\text{one_hot}}^T \mathbf{w}_{\text{affinity}}, \quad (5)$$

where the second term denotes *intra-modal* logits. Here, the few-shot knowledge in the visual-cache is used to update CLIP’s inter-modal predictions, *cf.* Fig. 1. Note that α trades off contributions of the visual and text-cache towards the final prediction. Furthermore, [53] proposes Tip-Adapter-F which improves upon Tip-Adapter by optimizing the visual-cache $\mathbf{W}_{\text{image}}$ w.r.t. $\mathcal{D}_{\text{train}}$ to learn the residual information required to increase the upstream classification performance in the target domain.

4. Domain Aligned CLIP

We now introduce Domain Aligned CLIP (DAC), a method that improves the few-shot domain adaptation of

CLIP in two stages. In the first stage, we tune a visual adapter layer to align CLIP’s visual representation in the target distribution, resulting in an improved intra-modal classifier. This intra-modal classifier later becomes the basis of DAC-V. In the second stage, we fine-tune CLIP’s textual representation for improved inter-modal alignment in the target distribution. This inter-modal classifier, together with the intra-modal classifier from the first stage results in DAC-VT. See Fig. 1 for an overview of both methods.

4.1. Adapting the Visual Domain

Tip-Adapter shows how an inter-modal classifier based on a text-cache can be improved when ensembled with a visual-cache based intra-modal classifier. However, the performance of the intra-modal classifier is inferior to the inter-modal classifier. This is due to CLIP being explicitly trained for inter-modal alignment between images and text where it is not explicitly encouraged to align embeddings of images sharing the same underlying concept or class. Moreover, recent work [7, 48] has shown the benefits of enforcing intra-modal alignment in pre-training of CLIP-like models.

In this work, we do not train models from scratch. Instead, we rely on a few labeled examples from the target domain to enhance the intra-modal alignment of pre-trained CLIP models without affecting their inter-modal alignment.

DAC Visual Adapter. To align the visual features of CLIP in the target domain, we introduce a linear layer \mathbf{H}_{θ} as an adapter that is appended to the frozen CLIP image encoder. During visual adaptation we only fine-tune the parameters θ . Unlike the two-layered adapter in [10], we found that a single linear layer is effective and avoids over-fitting. To allow the unimpeded passage of features at the beginning of fine-tuning, we initialize \mathbf{H}_{θ} as the identity, making it stable and efficient [19].

Visual Adapter Training. Next, we train the adapter layer \mathbf{H}_{θ} to minimize the distance between image embeddings of the same class in the latent space while pushing them apart for images of different classes. Similar to Sec. 3, we assume a few-shot setting with a novel target distribution $\mathcal{D}_{\text{train}}$ given N classes and K shots. To ensure that all K images of the same class are mapped to similar representations, we formulate a supervised contrastive objective, as illustrated in Fig. 3. Here, we consider K images as the context set \mathcal{C}_s of a parent class $s \in \{1, \dots, N\}$ which is further supplemented with M randomly augmented views of the given images, *i.e.* $\mathcal{C}_s = \{(\mathbf{x}_s^{(i)}, \mathbf{y}_s^{(i)})\}_i^{K \times M}$. In Sec. 6 we provide ablations for choosing an appropriate number of augmented views. To apply the visual adaptation, we linearly transform the image embeddings obtained from the frozen CLIP image encoder $\mathbf{v}_{\text{image}}^{\theta} = \mathbf{H}_{\theta} \mathbf{E}_{\text{image}}(\mathbf{x})$, followed by L2-normalization $\mathbf{g} = \mathbf{v}_{\text{image}}^{\theta} / \|\mathbf{v}_{\text{image}}^{\theta}\|$. Note that we drop the dependency on θ for brevity. We aim to find the

optimal transformation by minimizing the contrastive loss

$$\sum_{n=1}^N \sum_{1 \leq i < j \leq MK} \log \frac{\exp \left[(\mathbf{g}_j^{(n)})^T \mathbf{g}_i^{(n)} / \tau \right]}{\sum_{q=1}^N \exp \left[(\mathbf{g}_j^{(n)})^T \mathbf{g}_i^{(q)} / \tau \right]}, \quad (6)$$

where τ is a temperature to scale cosine similarities. Minimizing Eq. (6) aims to maximize the similarity between embeddings pairs coming from the same class (*positive* pairs), while maximizing dissimilarity between embedding pairs of different classes (*negative* pairs). Note that the summation $1 \leq i < j \leq MK$ in Eq. (6) considers all the positive pairs in \mathcal{C}_N , a total of $\binom{MK}{2}$ combinations. Our visual adaptation enforces structure onto the visual embeddings which is demonstrated by a much better cluster separation; cf. Fig. 4.

Constructing DAC-V. Leveraging the learned transformation \mathbf{H}_θ , we can improve Eq. (3) by an adapted visual-cache

$$\mathbf{W}_{\text{image}}^{\text{DAC}} = \left(\mathbf{g}^{(1,1)} \quad \mathbf{g}^{(1,2)} \quad \dots \quad \mathbf{g}^{(K,N)} \right) \in \mathbb{R}^{d \times NK}, \quad (7)$$

where we apply a horizontal concatenation as before. By inserting the improved visual-cache (Eq. 7) into Eq. (4), we obtain an optimized affinity vector that is visually adapted to the given task. However, note that the parameter β is subsumed into the learnable linear transformation. Hence, we have the optimized affinity vector

$$\mathbf{w}_{\text{affinity}}^{\text{DAC-V}} = \exp \left((\mathbf{W}_{\text{image}}^{\text{DAC}})^T \mathbf{g}_{\text{image}} - 1 \right) \in \mathbb{R}^{NK \times 1}. \quad (8)$$

While Eq. (4) and Eq. (8) share similarities, they differ in a crucial aspect. That is, by introducing a learned linear transformation in Eqs. (7) and (8), the intra-modal representation of DAC-V gets tailored towards the novel task, while Eq. (4) remains static and does not perform such domain adaptation. Similar to Eq. (5), we obtain the final logits as

$$\text{logits}_{\text{DAC-V}} = \text{logits}_{\text{CLIP}} + \alpha \mathbf{L}_{\text{one-hot}}^T \mathbf{w}_{\text{affinity}}^{\text{DAC-V}}. \quad (9)$$

In contrast to Eq. (5), the second term in Eq. (9) is composed of image features that are visually aligned in the target distribution. Our experiments demonstrate that DAC-V, on average, outperforms the fine-tuned Tip-Adapter-F on 11 image benchmarks by 0.83%, cf. Sec. 5.

4.2. Adapting the Textual Domain

Having optimized the intra-modal representations in DAC-V, we now look at enhancing inter-modal alignment between images and text features in the target distribution. Previous work [10, 55] demonstrates the benefits of optimizing text embeddings for few-shot classification through the optimizing of text prompts, while concurrently keeping the attached class names fixed. Such a prompt-tuning

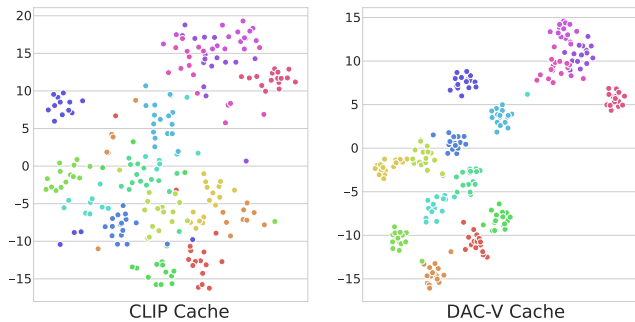


Figure 4. **2D tSNE projections of the CLIP's and DAC-V's cached image embeddings.** Our intra-modal contrastive objective enforces structure on the representations.

framework, however, significantly lacks the flexibility of fine-tuning the text embeddings. There are two severe issues: First, it does not address the adaptation of CLIP's vocabulary to the new class names from a target distribution as CLIP's learned vocabulary is shown to have limitations [35]. Second, there may be multiple visual concepts associated with different class names that can cause confusion among competing text features of CLIP [2, 40]; see Appendix C. To circumvent these challenges, we present a significantly simpler framework to align images and text labels in the target distribution.

DAC Textual Adapter. Unlike previous work, we do not introduce a new adapter module for improving inter-modal alignment. Instead, we directly fine-tune the text-cache \mathbf{W}_{text} . Note that the text embeddings are continuous vectors that encapsulate the concepts specified by class names from a target distribution. Therefore, modulating them influences the overall class description.

Constructing DAC-VT. To ensure smooth integration of the inter-modal alignment with the previously proposed intra-modal classifier in DAC-V, we optimize \mathbf{W}_{text} in the ensembled setting. More specifically, we convert \mathbf{W}_{text} into a learnable vector and freeze all the remaining components including the visual cache components and \mathbf{H}_θ . Thereafter, using the few-shot dataset, we optimize \mathbf{W}_{text} to align text embeddings with the visual embeddings while keeping the weighting parameter fixed, i.e. $\alpha = 1$. The optimized text weights $\mathbf{W}_{\text{text}}^{\text{DAC}}$ result in the DAC-VT classifier, i.e.

$$\text{logits}_{\text{DAC-VT}} = (\mathbf{W}_{\text{text}}^{\text{DAC}})^T \mathbf{z}_{\text{image}} + \alpha \mathbf{L}_{\text{one-hot}}^T \mathbf{w}_{\text{affinity}}^{\text{DAC-V}}. \quad (10)$$

Intuitively, optimizing the inter-modal alignment in this ensembled setting encourages $\mathbf{W}_{\text{text}}^{\text{DAC}}$ to assimilate the prior few-shot knowledge acquired by DAC-V. In Sec. 6, we ablate over other ways of constructing DAC-VT (including an end-to-end setting) that results into sub-optimal ensembles of intra-modal and inter-modal classifiers.

5. Experiments

In this section, we quantitatively evaluate our proposed method on 11 commonly used image classification tasks. We also study its robustness to distribution shifts. In Sec. 6, we ablate over DAC components and the design choices.

Datasets. For our experiments, we consider ImageNet [6], Caltech101 [9], FGVC Aircraft [26], UCF101 [38], EuroSAT [16], Flowers102 [28], StanfordCars [23], DTD [5], Food101 [3], OxfordPets [30], and, SUN397 [46].

Training and Evaluation Protocol. We follow the few-shot protocol by [33, 53] and fine-tune our models using 1, 2, 4, 8, and 16 shots per class, sampled from the training sets. Based on the validation sets, we then select the best fine-tuned adapters and the optimal α . Finally, we evaluate on the respective test sets. For ImageNet, like [53], we report results on the validation sets. Using Adam [21], we train the visual adapter for 500 epochs with learning rate 0.00003, temperature τ 0.008 and a batch size equal to the number of classes in the dataset. We set the number of randomly augmented views M to 7, and ablate over this parameter in Sec. 6. The training augmentations include random horizontal flips, random cropping, and, resizing to 224×224 pixels. To train the textual adapter, we follow the same data pre-processing protocol, but fine-tune for only 100 epochs with a learning rate of 0.00001. At inference time, we apply CLIP’s pre-processing (center cropping and resizing). On a single Nvidia V100 GPU, visual adapter training takes ~ 1 hour for the 16-shot setting. In contrast, textual adapter training only takes about 30 seconds. Note that for building the visual cache, following [53], we randomly augment each training image 10 times and use the mean embedding as a cache entry. For a fair comparison, we apply the prompt ensembling of [53] for ImageNet and a single prompt for the other datasets.

Baselines. We compare DAC with the strong existing few-shot adaptation methods for CLIP. This includes linear-probe CLIP [33], CoOp [55], CLIP-Adapter [10], and TIP-Adapter [53]. Note that we do not compare with [47, 51] as these recent works leverage language and vision generative models to generate bigger training sets from a few examples. We reproduce the results for TIP-Adapter using their official code¹. For other baselines, we provide officially reported scores for a fair comparison.

Results and Discussion. In Fig. 5, we compare the few-shot classification performances of DAC-V and DAC-VT with Tip-Adapter variants on all the datasets. It can be seen that DAC-V performs comparable to the strong baseline of Tip-Adapter-F. With an increasing number of training shots (*cf.* 4, 8, and 16-shot), the average performance of DAC-V surpasses that of Tip-Adapter-F by 0.8% or (performs better on 6 out of 11 datasets). Note that DAC-V is only

Models	RN50	RN101	V-B/32	V-B/16	V-L/14
Zero-shot CLIP	60.33	62.53	63.80	68.73	75.92
CoOp	62.95	66.60	66.85	71.92	-
CLIP-Adapter	63.59	65.39	66.19	71.13	-
SgVA-CLIP	65.70	68.51	68.26	73.30	-
Tip-Adapter	62.03	64.79	65.60	70.83	77.70
Tip-Adapter-F	65.47	68.53	68.74	73.70	79.43
DAC-V	64.89	67.38	67.77	72.98	79.62
DAC-VT	66.61	69.37	69.64	74.59	80.20

Table 1. **16-shot classification performance of different methods using different CLIP variants on ImageNet.** Here, *RN* refers to ResNet and *V* refers to ViT. For Eg., V-B/32 \rightarrow ViT-B/32.

optimized to align visual representations in the target domain and no explicit fine-tuning is done to increase its upstream few-shot performance. This clearly demonstrates the benefit of having a strong intra-modal classifier. With further optimization for inter-modal alignment on DAC-V, our proposed DAC-VT method surpasses all baselines by a significant margin. This strong result further illustrates the benefits of aligning both visual and textual domains on target distributions. The results in Fig. 5 correspond to the ResNet-50 variant of CLIP. In Tab. 1, we present few-shot adaptation results on ImageNet validation sets using different CLIP backbones. The results indicate a robust performance of DAC-VT across all CLIP variants, significantly outperforming the other baselines.

Distributional Robustness. Radford *et al.* [33] show that while fine-tuning improves the in-distribution performance, it reduces the overall robustness to shifts in distributions. So far we have observed that improving both visual and textual representations of CLIP in new domains consistently enhances its downstream performance in that domain. However, does it come at the cost of reduced robustness to natural shifts in distributions? In this section, we study the transfer of DAC models trained on ImageNet to four ImageNet variants *i.e.* ImageNet-V2 [34], ImageNet-Sketch [43], ImageNet-A [18] and ImageNet-R [17]. In Tab. 2, we conduct a cross-dataset evaluation and find that intra-modal alignment (DAC-V) results in better OOD performance, when compared to other methods that aim for inter-modal alignment. We conjecture that alignment of visual features is more robust to distribution shifts than inter-modal alignment [48]. Note that Tip-Adapter-F, optimized for inter-modal alignment, also does not outperform its untuned version in OOD setting.

Assaying Inter- and Intra-modal Classifiers in DAC.

The results in Fig. 5 and Tab. 1 empirically verify our main hypothesis that an ensemble of strong intra- and inter-modal classifiers leads to a better overall classifier. However, two questions arise: (1) How much does the visual cache benefit

¹<https://github.com/gaopengcuhk/Tip-Adapter>

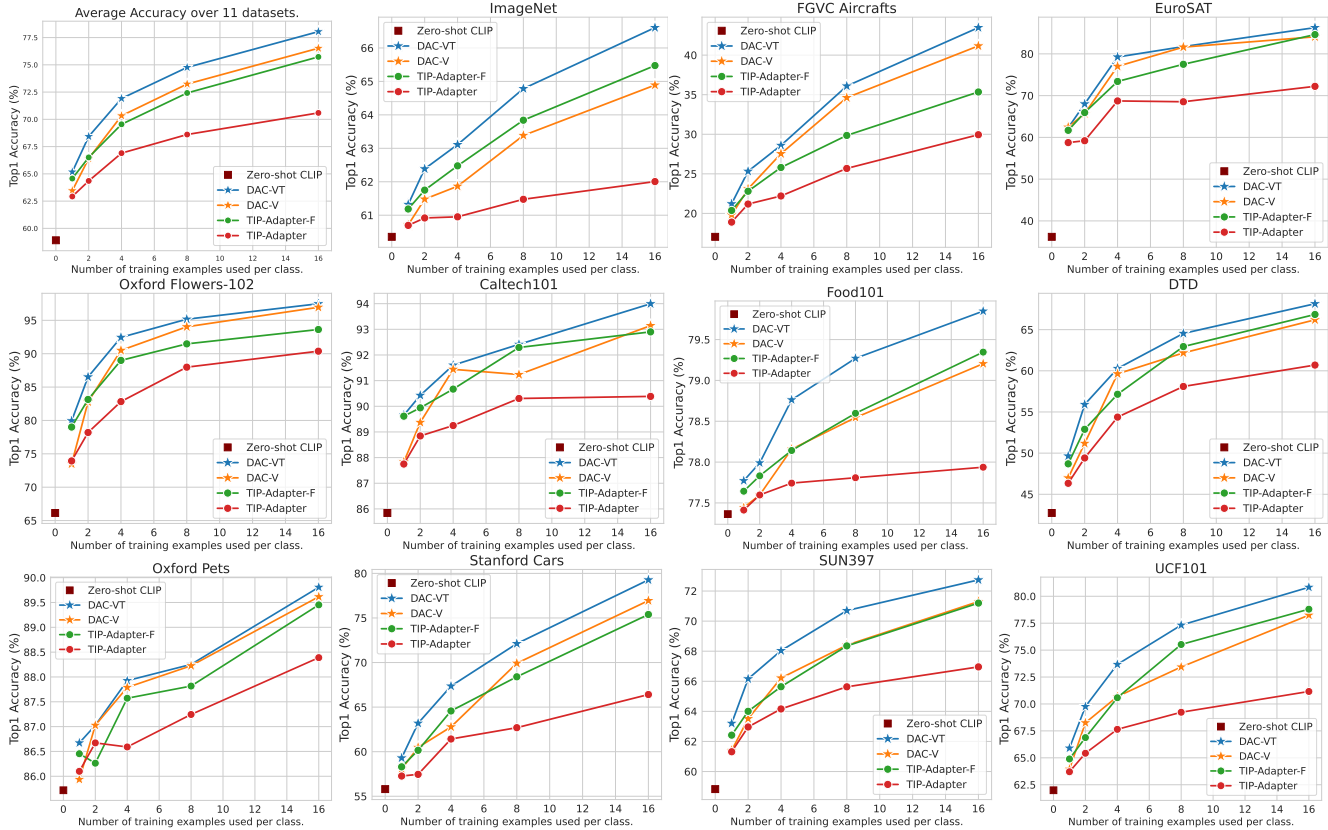


Figure 5. **A comparison of top1 accuracy (%) obtained by different few-shot CLIP adaptation methods over 11 datasets.** Here, the x-axis represents the number of training examples used per class from the target distribution. Our proposed methods, DAC-V which aligns only visual representations of CLIP on the target distribution and DAC-VT which aligns both visual and textual representations, perform comparable or better than the baselines. **(best viewed in color)**

	Source	Target Datasets			
		ImageNet	-V2	-A	-R
Linear-probe CLIP	56.13	45.61	12.71	34.86	19.13
CoOp	62.95	54.58	23.06	54.96	31.04
CoCoOp	62.81	55.72	23.32	57.74	34.48
CALIP-FS	65.81	55.98	23.42	56.74	35.37
TIP-Adapter	62.03	54.56	23.61	60.33	35.86
Tip-Adapter-F	65.47	56.79	20.93	58.48	34.62
DAC-V	64.89	56.56	23.92	60.52	36.27
DAC-VT	66.61	57.68	20.92	58.68	35.33

Table 2. **Robustness to Distributional Shifts.** We use CLIP ResNet-50 backbone for all the methods. Here, DAC-V demonstrates better performance than other baselines.

from intra-modal contrastive learning? And, (2) do inter- and intra-modal classifiers make sufficiently uncorrelated mistakes to justify ensembling? Figure 2 demonstrates that intra-modal contrastive fine-tuning enhances the discriminative capability of the visual cache, increasing DAC-Cache’s performance by 17% in the 16-shot setting. In contrast, Tip-Adapter-F fine-tuning reduces this performance by implicitly forcing the visual cache to learn the residual informa-

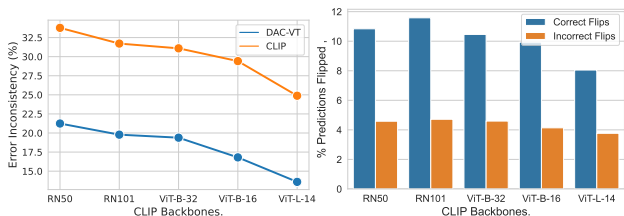


Figure 6. **(left)** Error inconsistencies between inter- and intra-modal classifiers using different CLIP backbones on ImageNet. DAC-VT significantly reduces this inconsistency. **(right)** Percentage of correct vs incorrect prediction flips in DAC-VT.

tion required to improve the upstream (inter-modal) classification, indicating that intra-modal adaptation is indeed beneficial. To analyze whether ensembling inter- and intra-modal classifiers would improve the overall performance, we plot their error inconsistencies in Fig. 6 (left). This reveals that, using pre-trained CLIP features, classifiers make highly uncorrelated mistakes, presenting an opportunity to flip the incorrect predictions via ensembling. DAC-VT reduces this error inconsistency, meaning that the predictions of both classifiers are flipped either correctly or incorrectly. In Fig. 6 (right) we show that the percentage of correct flips

V.A.	T.A.	α	Augmented Views	Top1 (%)
◇	-	-	-	41.51
✓	-	-	-	56.36
✓	-	-	✓	58.00
✓	◇	-	✓	62.01
✓	◇	✓	✓	64.89
-	✓	-	-	64.56
-	✓	-	✓	65.37
✓	✓	-	-	65.33
✓	✓	✓	-	65.97
✓	✓	-	✓	66.07
✓	✓	✓	✓	66.61

Table 3. **Ablation of DAC’s components on ImageNet using ResNet-50 in 16-shot setting.** Here, “-” indicates non-existence of the corresponding feature, ◇ indicates the usage of frozen CLIP features and ✓ indicates their adaptation for DAC-VT. Here V.A and T.A stands for visual and textual adaptation respectively.

is more than double than that of incorrect flips, indicating that the reduction in error inconsistency is due to correctly flipped predictions in ensembled setting. We refer readers to the ablations in Appendix B to see a similar behavior of error inconsistencies across different datasets.

6. Ablations

In this section, we ablate over all the components used in the construction of DAC-VT to justify our design choices. In Tab. 3, we see how the four components of DAC-VT interact on the ImageNet dataset using a CLIP ResNet50 backbone. Our textual adaptation alone provides a significant 5% boost in performance over zero-shot CLIP’s performance (i.e., 60.33%). Ensembling it with visual adaptation contributes to an additional 1.3% gain in accuracy. We also observe that the number of randomly augmented views M and the weighting parameter α play crucial roles in finding the optimal performance. We further ablate over the number of augmented views in Fig. 7. The accuracy increases monotonically up to $M = 7$ augmented views. Therefore, we select $M = 7$ for all of our experiments. To select the optimal weighting parameter α , we use grid search on the validation sets of each dataset (range [0.1, 10]). Note that [53] used the same strategy to find the optimal residual parameter. For the ImageNet 16-shot classification setting, we empirically find the values of 8.3 and 3.3 to be optimal for DAC-V and DAC-VT, respectively. We also ablate over the depth of the visual adapter layer \mathbf{H}_θ and find a single linear layer to be optimal (66.61% vs 65.58% of double layered adapter). We provide additional ablations on alpha values and adapter layers in Appendix A.

Ablating Alternative Ensembling Choices. In addition to the proposed DAC-VT framework, we also looked at other alternatives. The following experiments consider 16-shot

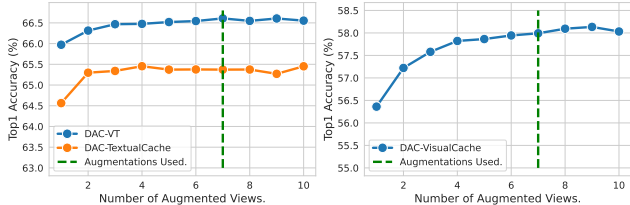


Figure 7. **Effects of varying number of augmented views used to train DAC inter-modal adapter (left) and intra-modal adapter (right).** The vertical green line indicates the number of augmented views used for all experiments in this paper.

classification on ImageNet using CLIP ResNet-50.

End-to-end Visual and Textual Domains Adaptation.

We tried to adapt both visual and textual representations together in an end-to-end fashion. The strategy is similar to CLIP-Adapter [10], however, we use a linear layer for visual features adaption as in DAC-V and adapted textual features as described in Sec. 4.2. The resulting model obtained 64.22% (vs 66.61% of DAC-VT). Noticeably, the intra-modal classification deteriorated similar to Tip-Adapter *i.e.* 28.44% (vs 41.51% of CLIP).

Building Visual Cache with Class Prototypes. Instead of retaining all image embeddings in the visual cache, we can reduce them to class prototypes [37] *i.e.* averaging image embeddings corresponding to all classes. This approach in DAC-VT’s style ensembling obtained 66.06% while it obtained 65.83% top1 accuracy in an end-to-end setting.

Cross-entropy Loss for Training Visual Adapter. The visual adapter can also be trained via cross-entropy instead of a contrastive objective. We ran an experiment where apart from the visual adapter’s training objective, all settings remained the same. The cross-entropy method gets 65.62% accuracy, while the contrastive objective achieves 66.61%.

7. Conclusion

We presented a sample-efficient framework, DAC, for adapting CLIP to downstream classification tasks. Using only a few labeled examples from a target distribution, DAC boosts the overall classification performance by improving both intra- and inter-modal representations of CLIP. Extensive experiments on 11 widely used image classification benchmarks show that DAC outperforms the competitive baselines while maintaining robustness to natural shifts in distributions. The performance improvements come at negligible additional computational cost during inference, as our framework requires only a linear layer for adaptation. Although the inference cost of DAC remains low, the two-stage adaptation increases the computational overhead for fine-tuning in comparison to the competitive baselines. We posit that there is room for further improving the ensembling of intra- and inter-modal classifiers, as both classifiers continue to exhibit uncorrelated errors *cf.* Fig. 6.

References

- [1] Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE T. Pattern Anal. Mach. Intell.*, 35(8):1798–1828, Aug. 2013. 3
- [2] Lucas Beyer, Olivier J. Hénaff, Alexander Kolesnikov, Xiaohua Zhai, and Aäron van den Oord. Are we done with ImageNet? *arXiv preprint arXiv:2006.07159*, 2020. 5
- [3] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101 – Mining discriminative components with random forests. In *ECCV*, volume 6, pages 446–461, 2014. 6
- [4] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, pages 1597–1607, 2020. 3
- [5] Mircea Cimpoi, Subhansu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. In *CVPR*, pages 3606–3613, 2014. 6
- [6] 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, 2009. 6
- [7] Jiali Duan, Liqun Chen, Son Tran, Jinyu Yang, Yi Xu, Belinda Zeng, and Trishul Chilimbi. Multi-modal alignment using representation codebook. In *CVPR*, pages 15651–15660, 2022. 3, 4
- [8] Alex Fang, Gabriel Ilharco, Mitchell Wortsman, Yuhao Wan, Vaishaal Shankar, Achal Dave, and Ludwig Schmidt. Data determines distributional robustness in contrastive language image pre-training (CLIP). In *ICML*, pages 6216–6234, 2022. 1
- [9] Li Fei-Fei, Rob Fergus, and Pietro Perona. Learning generative visual models from few training examples: An incremental Bayesian approach tested on 101 object categories. *Comput. Vis. Image Und.*, 106(1):59–70, Apr. 2007. 6
- [10] Peng Gao, Shijie Geng, Renrui Zhang, Teli Ma, Rongyao Fang, Yongfeng Zhang, Hongsheng Li, and Yu Qiao. CLIP-Adapter: Better vision-language models with feature adapters. *arXiv preprint arXiv:2110.04544*, 2021. 3, 4, 5, 6, 8
- [11] Muhammad Waleed Gondal, Shruti Joshi, Nasim Rahman, Stefan Bauer, Manuel Wuthrich, and Bernhard Schölkopf. Function contrastive learning of transferable meta-representations. In *ICML*, pages 3755–3765, 2021. 3
- [12] Raphael Gontijo-Lopes, Yann N. Dauphin, and Ekin D. Cubuk. No one representation to rule them all: Overlapping features of training methods. In *ICLR*, 2022. 2
- [13] Ziyu Guo, Renrui Zhang, Longtian Qiu, Xianzheng Ma, Xupeng Miao, Xuming He, and Bin Cui. CALIP: Zero-shot enhancement of CLIP with parameter-free attention. In *AAAI*, pages 746–754, 2023. 3
- [14] Michael Gutmann and Aapo Hyvärinen. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *AISTATS*, pages 297–304, 2010. 3
- [15] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross B. Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, pages 9729–9738, 2020. 3
- [16] 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, July 2019. 6
- [17] 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, 2021. 6
- [18] Dan Hendrycks, Kevin Zhao, Steven Basart, Jacob Steinhardt, and Dawn Song. Natural adversarial examples. In *CVPR*, pages 15262–15271, 2021. 6
- [19] Neil Houlsby, Andrei Giurgiu, Stanisław Jastrzębski, Bruna Morrone, Quentin de Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for NLP. In *ICML*, pages 2790–2799, 2019. 3, 4
- [20] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc V. Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *ICML*, pages 4904–4916, 2021. 1, 3
- [21] Diederik P. Kingma and Jimmy Lei Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 6
- [22] Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Joan Puigcerver, Jessica Yung, Sylvain Gelly, and Neil Houlsby. Big transfer (BIT): General visual representation learning. In *ECCV*, volume 5, pages 491–507, 2020. 3
- [23] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3D object representations for fine-grained categorization. In *ICCVW*, pages 554–561, 2013. 6
- [24] Victor Weixin Liang, Yuhui Zhang, Yongchan Kwon, Serena Yeung, and James Y Zou. Mind the gap: Understanding the modality gap in multi-modal contrastive representation learning. In *NeurIPS*, pages 17612–17625, 2022. 2, 12
- [25] Chengcheng Ma, Yang Liu, Jiankang Deng, Lingxi Xie, Weiming Dong, and Changsheng Xu. Understanding and mitigating overfitting in prompt tuning for vision-language models. *arXiv preprint arXiv:2211.02219*, 2022. 3
- [26] Subhansu Maji, Esa Rahtu, Juho Kannala, Matthew B. Blaschko, and Andrea Vedaldi. Fine-grained visual classification of aircraft. *arXiv preprint arXiv:1306.5151*, 2013. 6
- [27] Behnam Neyshabur, Hanie Sedghi, and Chiyuan Zhang. What is being transferred in transfer learning? In *NeurIPS*, pages 512–523, 2020. 1, 3
- [28] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *ICVGIP*, pages 722–729, 2008. 6
- [29] David W. Opitz and Richard Maclin. Popular ensemble methods: An empirical study. *J. Artif. Intell. Res.*, 11:169–198, Aug. 1999. 2
- [30] Omkar M. Parkhi, Andrea Vedaldi, Andrew Zisserman, and C. V. Jawahar. Cats and dogs. In *CVPR*, pages 3498–3505, 2012. 6

- [31] Fang Peng, Xiaoshan Yang, Linhui Xiao, Yaowei Wang, and Changsheng Xu. SgVA-CLIP: Semantic-guided visual adapting of vision-language models for few-shot image classification. *arXiv preprint arXiv:2211.16191*, 2022. 3
- [32] Hieu Pham, Zihang Dai, Golnaz Ghiasi, Kenji Kawaguchi, Hanxiao Liu, Adams Wei Yu, Jiahui Yu, Yi-Ting Chen, Minh-Thang Luong, Yonghui Wu, Mingxing Tan, and Quoc V. Le. Combined scaling for zero-shot transfer learning. *Neurocomputing*, 555:126658, Oct. 2023. 1
- [33] 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, 2021. 1, 3, 6
- [34] Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishaal Shankar. Do ImageNet classifiers generalize to ImageNet? In *ICML*, pages 5389–5400, 2019. 6
- [35] Shuhuai Ren, Lei Li, Xuancheng Ren, Guangxiang Zhao, and Xu Sun. Delving into the openness of CLIP. In *ACL*, pages 9587–9606, 2023. 2, 3, 5, 11
- [36] 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*, pages 14274–14289, 2022. 3
- [37] Jake Snell, Kevin Swersky, and Richard S. Zemel. Prototypical networks for few-shot learning. In *NIPS*, pages 4077–4087, 2017. 8
- [38] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. UCF101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012. 6
- [39] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. In *ECCV*, volume 11, pages 776–794, 2020. 3
- [40] Dimitris Tsipras, Shibani Santurkar, Logan Engstrom, Andrew Ilyas, and Aleksander Madry. From ImageNet to image classification: Contextualizing progress on benchmarks. In *ICML*, pages 9625–9635, 2020. 5, 11, 12
- [41] Vishal Udandarao, Ankush Gupta, and Samuel Albanie. SuS-X: Training-free name-only transfer of vision-language models. *arXiv preprint arXiv:2211.16198*, 2022. 3
- [42] Aäron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 3
- [43] 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. 6
- [44] Yeming Wen, Dustin Tran, and Jimmy Ba. BatchEnsemble: An alternative approach to efficient ensemble and lifelong learning. In *ICLR*, 2020. 2
- [45] Mitchell Wortsman, Gabriel Ilharco, Jong Wook Kim, Mike Li, Simon Kornblith, Rebecca Roelofs, Raphael Gontijo Lopes, Hannaneh Hajishirzi, Ali Farhadi, Hongseok Namkoong, and Ludwig Schmidt. Robust fine-tuning of zero-shot models. In *CVPR*, pages 7959–7971, 2022. 1, 3
- [46] 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, 2010. 6
- [47] Taihong Xiao, Zirui Wang, Liangliang Cao, Jiahui Yu, Shengyang Dai, and Ming-Hsuan Yang. Exploiting category names for few-shot classification with vision-language models. *arXiv preprint arXiv:2211.16594*, 2022. 3, 6
- [48] Jinyu Yang, Jiali Duan, Son Tran, Yi Xu, Sampath Chanda, Liqun Chen, Belinda Zeng, Trishul Chilimbi, and Junzhou Huang. Vision-language pre-training with triple contrastive learning. In *CVPR*, pages 15650–15659, 2022. 3, 4, 6
- [49] Yuan Yao, Ao Zhang, Zhengyan Zhang, Zhiyuan Liu, Tat-Seng Chua, and Maosong Sun. CPT: Colorful prompt tuning for pre-trained vision-language models. *arXiv preprint arXiv:2109.11797*, 2021. 3
- [50] Lu Yuan, Dongdong Chen, Yi-Ling Chen, Noel Codella, Xiyang Dai, Jianfeng Gao, Houdong Hu, Xuedong Huang, Boxin Li, Chunyuan Li, Ce Liu, Mengchen Liu, Zicheng Liu, Yumao Lu, Yu Shi, Lijuan Wang, Jianfeng Wang, Bin Xiao, Zhen Xiao, Jianwei Yang, Michael Zeng, Luwei Zhou, and Pengchuan Zhang. Florence: A new foundation model for computer vision. *arXiv preprint arXiv:2111.11432*, 2021. 1, 3
- [51] Renrui Zhang, Xiangfei Hu, Bohao Li, Siyuan Huang, Hanqiu Deng, Yu Qiao, Peng Gao, and Hongsheng Li. Prompt, generate, then cache: Cascade of foundation models makes strong few-shot learners. In *CVPR*, pages 15211–15222, 2023. 3, 6
- [52] Renrui Zhang, Longtian Qiu, Wei Zhang, and Ziyao Zeng. VT-CLIP: Enhancing vision-language models with visual-guided texts. *arXiv preprint arXiv:2112.02399*, 2021. 3
- [53] Renrui Zhang, Wei Zhang, Rongyao Fang, Peng Gao, Kun-chang Li, Jifeng Dai, Yu Qiao, and Hongsheng Li. Tip-Adapter: Training-free adaption of CLIP for few-shot classification. In *ECCV*, volume 35, pages 493–510, 2022. 1, 3, 4, 6, 8, 11
- [54] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. In *CVPR*, pages 16795–16804, 2022. 3
- [55] 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, July 2022. 3, 5, 6

A. Ablations

Architecture for Visual Adapter Layer. In the main paper, we use a linear layer for adapting visual features. To ablate the varying depth of \mathbf{H}_θ , we increase the number of layers (with ReLU activations in between). Table 4 shows the results on the ImageNet validation set using CLIP RN-50 on 16-shot classification.

Weighting Parameter α for different datasets. DAC improves the classification capability of both inter-modal and intra-modal classifiers. We use a scalar α to balance the contributions of each classifier towards the final accuracy. The value for α is selected based on the performance on the validation sets. A similar strategy was employed by Tip-Adapter [53]. However, in Tip-Adapter it is used to determine how much residual information should flow from intra-modal classifier to update the inter-modal predictions. To find the optimal value, we perform a grid search with a step size of 0.01, a search range in [0.1, 10], and the number of search steps being 10000. In this section, we present different values of α used to compute the final test performances of each dataset. In Fig. 8, we show how varying α influences the performance on 16-shot ImageNet classification. Table 5 lists our optimal values for α for all datasets (both DAC-V and DAC-VT). Since α is multiplied with the intra-modal logits, it can be seen that DAC-V consumes more information from the intra-modal classifier. Remember that in DAC-V, we only optimize the visual representations of CLIP without optimizing it for the upstream few-shot classification task. This further highlights the benefits of having better intra-modal representations in few-shot adaptation setting.

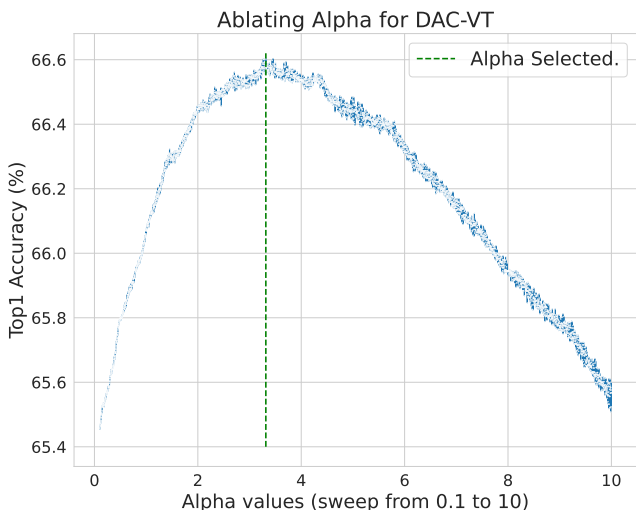


Figure 8. Ablating the α parameter for ImageNet using CLIP ResNet50.

B. Detailed Analysis on Error Inconsistencies

We analyze the error inconsistencies observed across various datasets in Fig. 9. This plot complements our analysis in Sec. 5 about the role of inter- and intra-modal classifiers in an ensembled setting, and further illustrates how DAC-VT reduces inconsistencies between intra and inter-modal classifiers. The consensus between the DAC-VT’s sub-classifiers is higher for some datasets (e.g., Flowers102, Caltech101), however, the inconsistencies for certain datasets (e.g., FGVCaircraft) are still high.

C. A Case for Aligning Textual Representations in Target Domain

We further elaborate on why it is important to align textual features on each downstream task. Previous work [35] has shown that CLIP’s zero-shot transfer is vulnerable to expansion of downstream vocabulary used for class labels. This becomes even more important when the visual concepts in the target domain get associated with different class labels, presented at different granularities. Figure 10 shows an example (taken from [40]) where multiple, different labels from ImageNet can be used to describe the same image. Such cases are particularly difficult for vision-language

Structure of \mathbf{H}_θ	DAC-V	DAC-VT
Linear Layer	64.89	66.61
2 Layer MLP	64.452	65.582
3 Layer MLP	64.08	65.274
4 Layer MLP	64.01	65.064

Table 4. Ablating structure of adapter layer \mathbf{H}_θ

Datasets	DAC-V α	DAC-VT α
UCF-101	3.78	1.16
Caltech101	2.40	1.33
ImageNet	8.32	3.31
SUN397	5.95	1.39
FGVCAircraft	8.2	6.91
StanfordCars	6.50	2.42
Flowers102	8.17	3.43
Food101	1.17	1.05
OxfordPets	1.07	0.73
DTD	3.05	1.11
EuroSAT	5.17	0.76

Table 5. Details of α used to weigh intra and inter-modal classifiers for different datasets in DAC-V and DAC-VT. In DAC-V the contribution from intra-modal features is weighted more which indicates that the adapted visual cache contains reliable information to update CLIP’s inter-modal knowledge.

models to generalize to in zero-shot manner, unless more context is given by either prompts or some domain-specific training data.

Note that the adaptation of textual representations introduced in *cf.* 4.2 aims to caters for such confusing examples as it modulates the overall textual embedding (including the class name). Such an optimization allows the textual cache to adapt the class description according to the visual concepts defined by a few observed images.

Few-shots	1	2	4	8	16
Linear-probe CLIP	22.17	31.98	41.20	49.52	56.13
CoOp	57.15	57.81	59.99	61.56	62.95
CLIP-Adapter	61.20	61.52	61.84	62.68	63.59
Tip-Adapter	60.70	60.92	60.95	61.48	62.00
Tip-Adapter-F	61.19	61.75	62.48	63.84	65.47
DAC-V	60.71	61.48	61.87	63.38	64.89
DAC-VT	61.32	62.39	63.11	64.78	66.61

Table 6. Top1 accuracy of different methods on ImageNet at different shots.

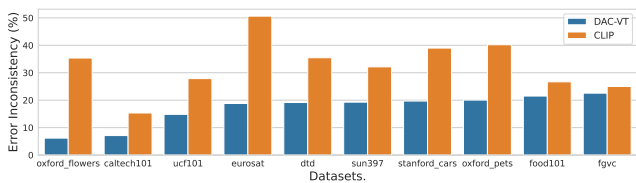


Figure 9. Comparative analysis of error inconsistencies between intra-modal and inter-modal classifiers of CLIP and DAC-VT on 10 different datasets (sorted by DAC-VT’s performance). We observe that DAC-VT significantly reduces the error inconsistencies, however, the performance gap reduces on certain datasets such Food101 and FGVAircrafts.




Figure 10. Examples of confusing labels in ImageNet [40]. The labels above appear to correctly describe the visual concepts, however, ImageNet assigns *acoustic_guitar* and *projectile* labels to the images, respectively.

D. Understanding Inter-Modal and Intra-Modal Representations Alignment

In this section, we delve into understanding how DAC-VT modulates the interactions between inter-modal and intra-modal representations. We look at them from the perspective of cone effects occurrences in representations distances that’s been extensively studied in [24]. In Fig. 11, we showcase the range of cosine similarities scores obtained by computing similarities between inter-modal and intra-modal representations. It can be seen that even after updating textual representations, DAC-VT maintains the same range of inter-modal similarity between images and text as in CLIP. The bigger shift is observed in intra-modal alignment where the visual representations tuned with DAC have a different support in comparison to TIP and CLIP based intra-modal alignments. We conjecture that this shift happens because the supervised contrastive objective used to tune visual representations introduce a different learning inductive bias than what was used to aligning image-text representations.

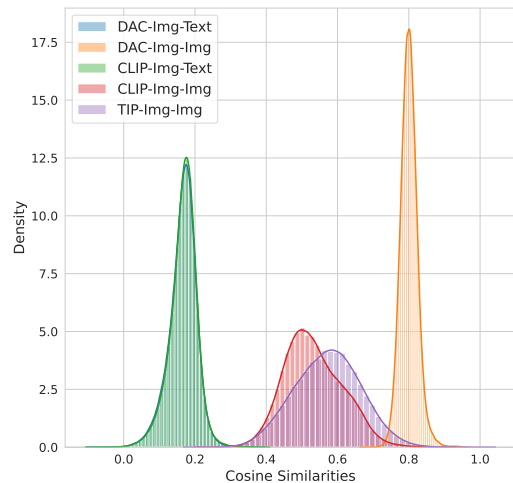


Figure 11. Pictorial depiction of modality gaps between intra-modal and inter-modal representations of different methods (illustrated by cosine similarities). It can be seen that the DAC-VT’s and CLIP image-text similarities remain within the same range.