

Reevaluating the Intra-Modal Misalignment Hypothesis in CLIP

Jonas Herzog
Zhejiang University
jherzog@zju.edu.cn

Yue Wang
Zhejiang University
wangyue@iipc.zju.edu.cn

Abstract

Recent research suggested that the embeddings produced by CLIP-like contrastive language-image training are sub-optimal for image-only tasks. The main theory is that the inter-modal (language-image) alignment loss ignores intra-modal (image-image) alignment, leading to poorly calibrated distances between images. In this study, we question this intra-modal misalignment hypothesis. We reexamine its foundational theoretical argument, the indicators used to support it, and the performance metrics affected. For the theoretical argument, we demonstrate that there are no such supposed degrees of freedom for image embedding distances. For the empirical measures, our findings reveal they yield similar results for language-image trained models (CLIP, SigLIP) and image-image trained models (DINO, SigLIP2). This indicates the observed phenomena do not stem from a misalignment specific to the former. Experiments on the commonly studied intra-modal tasks retrieval and few-shot classification confirm that addressing supposed misalignment is unnecessary for best performance. Project Page: <https://vision-kek.github.io/Is-CLIP-Really-Misaligned>.

1. Introduction

The success of contrastive language-image pretraining exemplified by CLIP [33] has significantly shaped the contemporary era of multimodal, foundational models. The most fundamental and still ubiquitous use-case is to embed images and texts in a shared space, and then compare them via cosine similarity. Following this simple approach, many long-standing computer vision problems from classification to detection [23, 24] to segmentation [19, 49] became addressable in the popular open-vocabulary zero-shot setting.

Nevertheless, recent works have raised concerns that these embeddings may not be ideal for image-only tasks. The so-called *intra-modal misalignment hypothesis* argues that contrastive language-image training focuses solely on inter-modal (image-text) alignment while neglecting intra-modal (image-image) alignment [2, 26, 40, 44]. As a result,

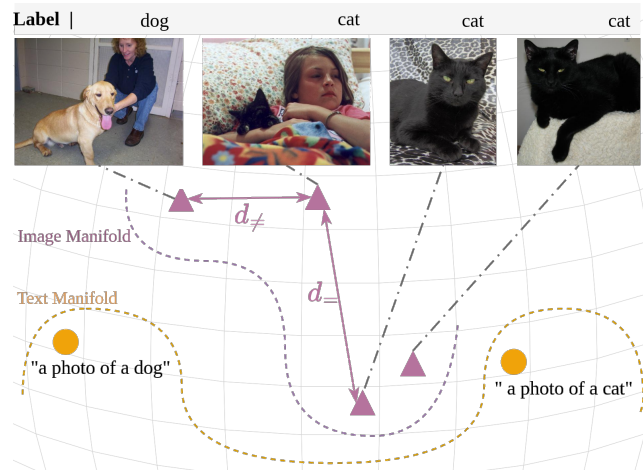


Figure 1. Previous work illustrated an intra-modal misalignment in CLIP space by showing there are cat images closer to a dog (d_{\neq}) than to another cat ($d_{=}$). We argue $d_{\neq} < d_{=}$ is no sign of misalignment. For a labeled downstream dataset, intra-class variance of open vocabulary models is *expected and desired* to capture semantics and style beyond the narrow dataset-specific labels. Classifying and retrieving with frozen CLIP image embeddings still works well when similarities are measured along the dataset-specific semantic axes. Here the horizontal axis captures dog/cat.

even though the model produces good image-text similarities, the same might not be true for image-image similarities. This effect is often illustrated as in Fig. 1.

The hypothesis relies on the following pillars: (i) a theoretical argument that misalignment emerges from unconstrained degrees of freedom in the embedding space, and (ii) empirical evidence through proposed misalignment indicators such as cosine similarity histograms, modality-gap magnitudes, and retrieval or few-shot classification performance metrics.

In this work, we critically re-examine these arguments.

For the theoretical model, we demonstrate that the degree of freedom argument breaks when considering a sufficiently large set of embeddings. Image-image similarities can in fact be recovered entirely from image-text similarities. This shows that image-image structure is not arbitrary

but a consequence of the learned image-text structure.

For the empirical indicators that were proposed to reveal the misalignment, we experiment with these measures and find they are no reliable indicators of intra-modal alignment quality. For example, we argue intra-class variance as shown in Fig. 1 is not a sign of misalignment. Instead, this is normal behavior because without fine-tuning, open-vocabulary models are expected to capture semantics and styles beyond the closed-vocabulary of a narrow downstream dataset. More crucially, the same indicators would also suggest “misalignment” in non-CLIP vision encoders such as DINO [5], which have never been trained with language supervision — revealing that these signals are artifacts of the metrics rather than of the training objective.

Based on our findings, we propose a simple alternative method that measures similarities in class-relevant axes. We conduct evaluations on retrieval and few-shot classification, which have been previously assumed to be affected by the misalignment, and provide a reinterpretation of the results. Finally we discuss the role of the modality gap [20] which is often cited [25, 26, 44] in the context of intra-modal misalignment. In summary:

- We critically re-examine the intra-modal misalignment hypothesis, which directly impacts all methods that use CLIP or SigLIP to compare image embeddings with each other.
- We point out the limitations of previous arguments and evidences for such misalignment and provide an alternative explanation for their findings.
- Consequently, we construct a simple alternative method that confirms that the best performance on the previously studied few-shot classification and retrieval tasks can be achieved without assuming a misalignment in CLIP.

2. The Intra-Modal Misalignment Hypothesis

We examine the following belief:

CLIP is not optimized for uni-modal scenarios [44]. This leads to intra-modal misalignment [26] / miscalibration [40], its performance in intra-modal tasks is not guaranteed [44]. Only exploiting the image encoder of such models is highly suboptimal for intra-modal tasks like image-to-image retrieval [26].

2.1. The Supporting Evidence

We review some evidence in favor of the above intra-modal misalignment hypothesis.

The consequence of having only a cross-modal loss term in CLIP has been **theoretically** explained [40, 44] by degrees of freedom in the embedding space. This idea has been illustrated similar to Fig. 4a-c.

Further, to measure and expose poorly calibrated similarities, several **indicators** were proposed. The main tool




Figure 2. Pairwise cosine similarity distributions. *Left*: Similarities between same class (blue) and opposite class (orange) image feature pairs. A high overlap ratio between the two colors was previously highlighted as an indicator for an intra-modal misalignment issue in CLIP. *Right*: Similarity distributions of image-text pairs (purple) versus image-image pairs (green). Because CLIP is only supervised on the former, the divergence has previously prompted concerns about whether the latter reflect true similarities. CLIP ViT-B/16. Dataset as in Tab. 1.

for this purpose are cosine similarity histograms, where the distribution of pairwise similarities between images are recorded. There are two types of similarities that are examined, by class and by modality, shown in Fig. 2. For the similarity distributions by classes [18, 26], the point is straightforward: Two images of the same class should have higher similarity than two images of different classes. As in Fig. 2a, if the two histograms that capture the similarities of same-class and different-class image pairs have a high overlap, it is hypothesized this indicates poor class separation and hence is a sign of misalignment. On the other hand, measuring the similarity distributions by modality [2, 38, 40] as in Fig. 2b reveals that image embeddings have a much higher similarity to each other than to text embeddings. This is considered problematic especially for few-shot classification in the vision-language model adaptation literature [2, 38, 40, 44] because these work combine inter-modal and intra-modal similarity scores, such that it is questioned if the two differently distributed scores can be treated equally.

This finding of higher intra-modal similarities is a consequence of the modality gap [20], which describes the effect that image embeddings cluster in a manifold clearly separated from the text embedding cluster. Distances across the cluster are therefore naturally higher on average than within. Some work [2, 26] has also attributed the misalignment to the modality gap, but the existing experiments [20, 26] documented that a shifting or fine-tuning approach to narrow the gap tends to worsen performance.

To mitigate the assumed misalignment, prior works then proposed methods for better **performance**. The most common idea is to avoid measuring image-image similarities in favor of text-image similarities. For few-shot classification, Tip-X [40] proposed to address uncalibrated image-image embedding distances in few-shot classification by using image-text similarities as a bridge. Instead of directly comparing image features, it constructs “signatures”




Figure 3. Motivated by the intra-modal misalignment hypothesis, previous work [26] posited it is necessary to convert image-image comparison (left) into image-text comparison (right).

for both test and few-shot images by measuring their affinity with text classifier weights. The final affinity between images is then calculated as the KL-divergence between these inter-modal probability distributions. With a similar motivation, Mistretta *et al.* [26] proposed to tune a text token that represents the image information, an approach illustrated in Fig. 3 and dubbed Optimization-based Textual Inversion (OTI) inspired by [13]. This proved to be more efficient on image-image retrieval than comparing image embeddings directly. In another work, Barbier *et al.* [2] found that simply normalizing the cosine similarity distributions as in Fig. 2a to equal mean and variance can bring improvements for few-shot object detection with CLIP-like VLMs. Chao *et al.* [44] also assumed wrong distances between images for intra-modal tasks, and therefore introduced to represent an image by its cross-modal distances to text generated with the help of large language models.

In summary, a substantial body of prior work has interpreted calibrated similarity statistics, modality-wise clustering, and improved performance from text-based surrogates as evidence of intra-modal misalignment.

2.2. The Prior Opposing Indicators that Call for Reassessment

While the preceding section reviewed evidence supporting the intra-modal misalignment hypothesis, there also exist several observations that lower the prior probability that the hypothesis is true before considering the evidence presented in our paper.

A first source of skepticism comes from some influential literature that has successfully employed CLIP and related vision-language models for uni-modal tasks such as image classification [14], few-shot adaptation [3, 42, 47], image generation [34], and video synthesis [11, 48]. Further, [25] report that image-to-image retrieval actually outperforms image-to-text retrieval in downstream VLM adaptation tasks. If CLIP embeddings were indeed severely misaligned within the visual modality, it would be difficult to reconcile this with the strong results reported in these works, particularly those experiments that exclusively rely on the image encoder. A second reason to question the mis-

Table 1. Repeating the demonstrative experiment in [26] on the simplistic Dogs vs Cats legacy dataset, where near-perfect results are expected. It was suggested in [26] that poor results with CLIP image-image similarities evidences a misalignment in the image-image space. This hypothesis gets no evidence when swapping model for the uni-modal DINO, a widely acknowledged state-of-the-art image embedder. CLIP scores highest, suggesting that the observed low metrics are not caused by a model weakness and hence neither by a misalignment. Instead, the performance gap between text-image (T-I) and image-image (I-I) can be attributed to the ambiguity [21] in the way the task is conveyed to the model: Two images with opposite labels might still share enough other concepts to be justifiably similar.

Model	Retrieval		Classification	
	<i>T-I</i>	<i>I-I</i>	<i>T-I</i> (0-shot)	<i>I-I</i> (1 16-shot)
CLIP ViT-B/16	99.3	87.1	99.6	84.2 99.7
DINOv2 ViT-B/14	-	81.8	-	76.2 97.3
DINOv3 ViT-L/16	-	84.3	-	80.2 97.8

alignment interpretation concerns the meaning of intra-class variance. An alternative explanation for the phenomena illustrated in Fig. 1 has been articulated in [53]: embeddings naturally contain both task-specific and task-irrelevant components. From this perspective, high intra-class variance may reflect the coexistence of other semantic factors rather than represent a flaw in the embedding space. The challenge, therefore, is not to “correct” image-image distances, but to identify task-relevant axes (classifier) given the limited samples [12, 21].

Taken together, Sec. 2.1 and Sec. 2.2 highlight a tension: while many works argue for intra-modal misalignment, several existing results in the broader literature are difficult to reconcile with the misalignment view. This inconsistency calls for a careful reassessment of the hypothesis.

2.3. A Preliminary Experiment

As a first step toward reassessing the misalignment hypothesis, we revisit a simple experiment in [26] which particularly inspired this work. In this experiment, poor CLIP retrieval performance on a very easy cats and dogs legacy dataset [9] was attributed to the intra-modal misalignment.

However, we consider it unlikely that CLIP fails on such dataset. A simple way to rule out a model weakness is comparison with uni-modal vision models such as DINO [29, 36]. Tab. 1 demonstrates how image-image metrics with DINO lack even more behind the text-image result, prompting the suggestion that the observed measures are *not a weakness of CLIP*.

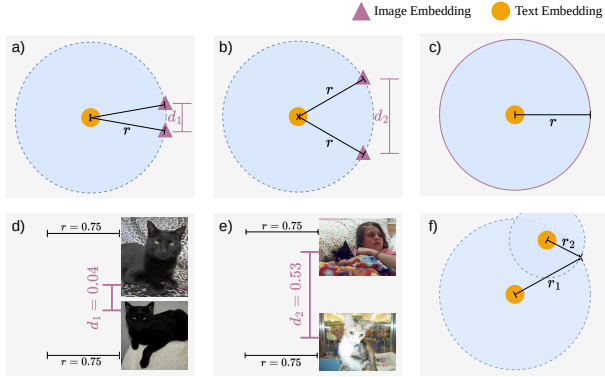


Figure 4. (a)-(c): **The previous intra-modal degree-of-freedom argument** in [40] illustrates that two image embeddings can be either close together (a) or far apart (b) while having the same text-image distance r , concluding that two image embeddings can lie on any two arbitrary points on the circumference (c), leaving a degree of freedom for image-image miscalibration. **Our interpretation (d)-(f):** The previous line of argumentation overlooks that each image embedding is bound to more than one text anchor (f). Moreover, the two different configurations in (a) and (b) are not arbitrary, but have a good reason to exist: Images in (a,d) and (b,e) have equal distance r to the “cat” text, but the two images in (a,d) are much more similar to each other than those in (b,e). Displayed distance values are real measurements.

3. Methodology

Following the structure established by the previous supporting evidence from Sec. 2.1, our methodology to re-evaluate the intra-modal misalignment hypothesis is threefold:

From the theoretical perspective, Sec. 3.1 re-examines if text-image alignment leaves degrees of freedom for image-image misalignment.

From the empirical perspective, Sec. 3.2 builds on the insight from the preliminary experiment that we can substitute models to allow for controlled study of indicators.

Lastly, to explain performance gains of previous image-to-text augmentation approaches, Sec. 3.3 proposes a simple alternative method.

3.1. Analyzing intra-modal degrees of freedom

We begin with re-examining of the degree of freedom argument in [40], where it was hypothesized that intra-modal misalignment emerges from the degrees of freedom that image embeddings have during contrastive language-image training. The idea how this leaves room for poor calibration is illustrated in Fig. 4a-c.

We argue the shortcoming of this hypothesis and its visual demonstration lies in its reduction to a single text embedding. If we extend their model to a larger number of embeddings, then the conclusion no longer holds.

The general acknowledged assumption is that text-image

distances are well calibrated. An argument for the misalignment states that given certain text-image distances, there is still room for miscalibrated image-image distances. We can test this argument, formulating the question:

In the pre-training dataset with n_T texts and n_I images, given fixed pair-wise *inter*-modal cosine similarities expressed through matrix $S_{inter} \in \mathbb{R}^{n_T \times n_I}$, is there any room for miscalibrated *intra*-modal similarities $S_{intra} \in \mathbb{R}^{n_I \times n_I}$?

If there were a way to express S_{intra} as an entangled consequence of S_{inter} , then we could conclude that there is no degree of freedom. The underlying image embeddings $X_I \in \mathbb{R}^{n_I \times d}$ are the unknown. Also the text embeddings $X_T \in \mathbb{R}^{n_T \times d}$ are in general not necessary to infer S_{intra} . We can show that in a d -dimensional space, with only d sampled text anchors represented by row indices $\mathcal{J} \subset \{1, \dots, n_T\}, |\mathcal{J}| = d$, we can set up a linear system that allows for unique recovery of all image embeddings $X_I \in \mathbb{R}^{n_I \times d}$. We know that:

$$S_{inter}[\mathcal{J}] = X_T[\mathcal{J}] \cdot X_I^\top. \quad (1)$$

Solving for X_I ,

$$X_I = (X_T[\mathcal{J}])^{-1} \cdot S_{inter}[\mathcal{J}]. \quad (2)$$

Sampling at least d text embeddings ensures $X_T[\mathcal{J}]$ is invertible, given that every row in X_T is a unique embedding unit vector. Typically $n_T, n_I \gg d$, so the number of required anchor points is comparatively small. The logic behind Eq. (2) is illustrated for $d = 2$ in Fig. 4f. From Eq. (2) follows that all image-image similarities are *well-defined without further degrees of freedom*:

$$S_{intra} = X_I X_I^\top. \quad (3)$$

The dot product equals cosine similarity because the rows in X_I are unit vectors. Eq. (2) produces unit vectors provided that the rows of X_T are normalized likewise.

While in this section we borrow the text anchor assumption from previous work for illustration, it should be noted that such fixed points do not exist in reality. Therefore, in Appx. A we show an alternative way to infer S_{intra} without sampled anchor embeddings.

3.2. Contrasting with non-CLIP

We seek to put the previous proposed indicators for a misalignment under test.

To isolate effects specific to the lack of an intra-modal training objective, we employ the method of comparing models solely trained with inter-modal objectives (CLIP, SigLIP [46]) with models trained with intra-modal objectives (DINO, SigLIP2 [39]).

As in [26] and Fig. 2a, we measure same- vs. opposite class cosine similarity distributions for paired and unpaired samples. Following [2, 40] and Fig. 2b, we further measure same- vs. opposite modality similarity (S_{intra} vs S_{intra}) distributions by substituting CLIP models with models that have additional image supervision as in the DINO objective. Finally, we conduct the same CLIP vs. non-CLIP comparison on performance metrics such as retrieval average precision and few-shot classification accuracy.

This way, if an experimental outcome reproduces for both model types, it cannot be attributed to a missing intra-modal term in CLIP’s loss function.

3.3. A simple alternative for more “classness” in similarity measure

When relying on similarity measures between images, many typical classification and retrieval tasks would benefit from only focusing on the image’s single most dominant concept. Ignoring all visual details and background objects would then remove information that does not correlate with the dataset’s labels.

We attribute previous improvements brought by the modality inversion [1] technique from Fig. 3 to this effect. This technique was used in [26] as a means to overcome intra-modal misalignment by converting an image to a text token. This text token v^* within the “a photo of a v^* ” prompt is tuned through backpropagation through the text encoder in order to approximate the image embedding.

We argue this is only efficient because it forces the information contained in the image to collapse to a single word-like token. The resulting text embedding is like “What single word describes the image best?”, which very often correlates strongly with the to-predict class in the test dataset.

To validate our claim, we propose to replace the entire Modality Inversion procedure through simply projecting the image embeddings on the subspace that explains the most variance of “a photo of x ” prompts, where x is a word-like class name. This is similar to [52, 54]. Irrespective of the downstream dataset, we choose ImageNet class names for x . Given n class names, we extract n text embeddings with dimension d . The first $d/2$ principal components are selected. All test images are projected onto the subspace spanned by the selected components. We then obtain results by measuring image-image cosine similarities in this subspace. We refer to this method as PCA^+ in the experiments.

4. Results

Sec. 4.2 evaluates the cosine similarity histogram indicators under the CLIP/non-CLIP substitution methodology from Sec. 3.2. Sec. 4.3 and Sec. 4.4 present results on few-shot classification and image-to-image retrieval, respectively, including the method from Sec. 3.3.

4.1. Datasets and Metrics

Following the line of works [47, 50, 51] on few-shot classification with CLIP, we evaluate on 11 image classification datasets. These datasets encompass a variety of image recognition tasks, such as generic object recognition with ImageNet [8] and Caltech101 [41], fine-grained recognition using OxfordPets [30], StanfordCars [17], Flowers102 [28], Food101 [4], and FGVCAircraft [22], satellite image classification with EuroSAT [15], action classification with UCF101 [37], texture classification with DTD [7], and scene recognition with SUN397 [43]. We measure the common classification accuracy metric.

We use these datasets also for image-to-image retrieval, following the work that hypothesized an image-image misalignment [26]. For retrieval, we additionally evaluate on the more commonly used dataset ROxford and RParis [32]. Additional ablations are conducted on the autonomous driving dataset BDD100k [45]. For retrieval evaluations, following [26], we measure mean average precision (mAP) over retrieval queries.

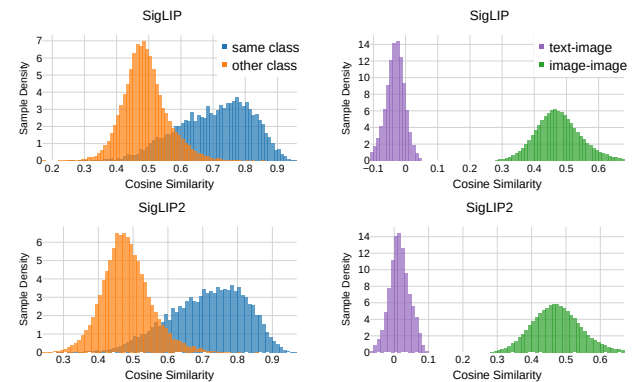


Figure 5. Cosine similarity histograms **by class** (left) and **by modality** (right). The distributions are almost identical for purely text-image trained SigLIP (first row) and SigLIP2 (second row) which includes an image-image self-supervised objective as in the DINO line of work. This indicates intra-class variation (left) and the gap between text and image embeddings (right) are no signs of misalignment brought by pure text-image training, but rather normal behavior. All models are ViT-B. Embeddings are sampled from ImageNet validation set.

4.2. Cosine Similarity Distributions

Setup. We analyze the two cosine similarity distribution indicators proposed as evidence for intra-modal misalignment (see Fig. 2 for a recap). To test whether these indicators are specific to models lacking an intra-modal objective, we contrast a model trained with only an inter-modal loss (SigLIP) with a model that includes an additional image-image objective (SigLIP2). If the indicators are caused by a missing intra-modal loss, they should be “fixed” in SigLIP2. Results

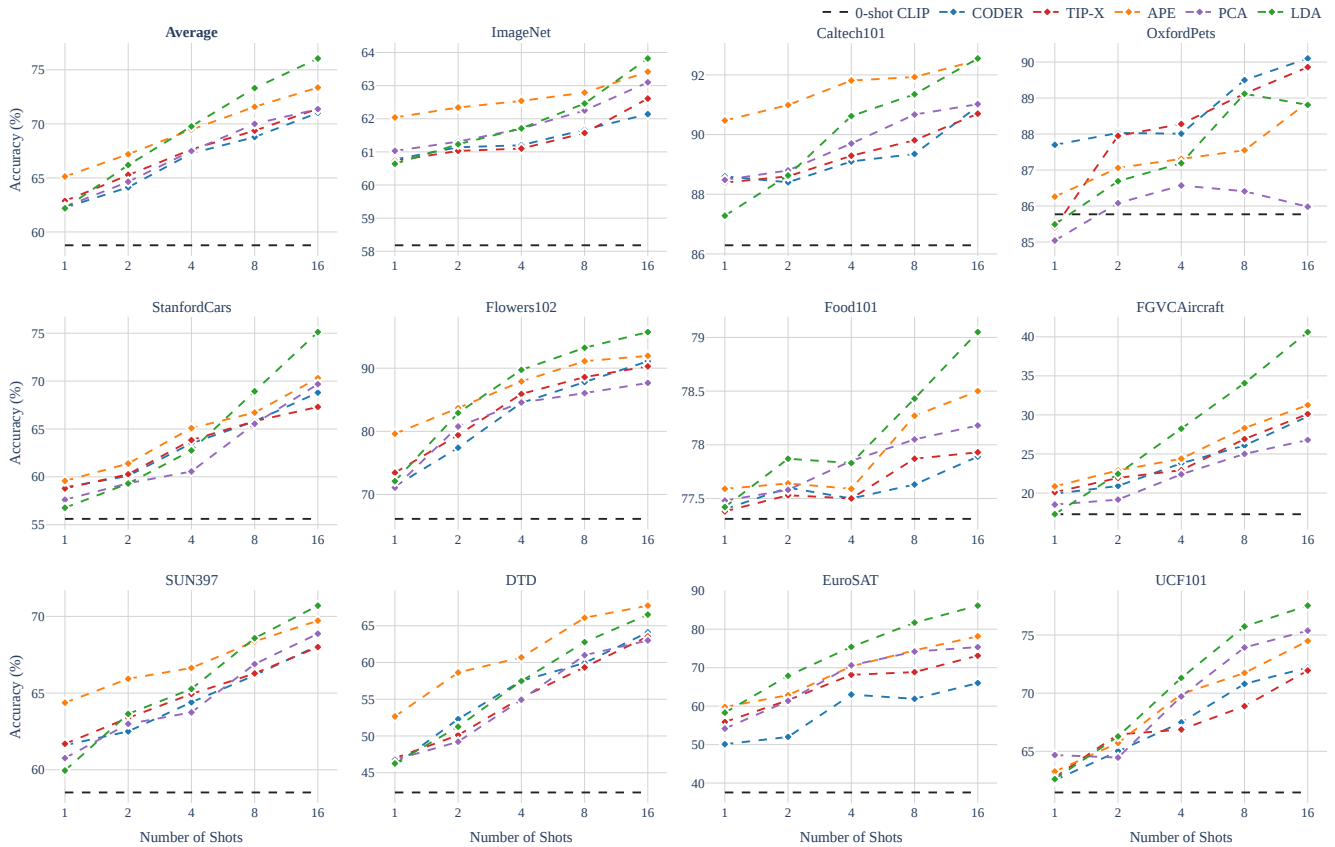


Figure 6. **Few-shot classification as in training-free VLM adaptation** literature. Approaches that attempt to address the intra-modal misalignment issue through avoiding use of image-image similarities (Tip-X, Coder) [40, 44] do not win against feature selection in APE [53] or a basic Linear Discriminant Analysis (LDA) [42] on image embeddings.

are shown in Fig. 5. The key insight is that both indicators are indistinguishable between the two models, with details discussed in the following.

Table 2. **Few-shot classification with no use of text prompts.** Accuracy is therefore only based on image-to-image similarities. This allows directly to test the intra-modal alignment quality of models with inter-modal loss only (CLIP, SigLIP), both inter- and intra-modal losses (SigLIP2) and image-only models (DINO). Results on PCA^{\leftarrow} -projected features are in (\cdot) parenthesis. Two best models per classifier are underlined>. SigLIP scores higher than DINO, indicating well-aligned image embeddings despite inter-modal training only. All models are ViT-B. Reported values are the average over the 11 datasets as in Fig. 6.

Model	Classifier	1-shot	2-shot	4-shot	8-shot	16-shot
CLIP	Prototype	43.4 (50.7)	55.3 (62.5)	63.8 (69.7)	69.6 (74.7)	73.5 (77.5)
	LDA	2.3 (2.3)	60.0 (62.5)	69.8 (70.8)	76.1 (76.5)	79.5 (79.6)
SigLIP	Prototype	57.3 (62.1)	<u>68.6</u> (71.9)	<u>76.3</u> (78.5)	<u>80.5</u> (81.8)	<u>82.5</u> (83.7)
	LDA	2.3 (2.3)	<u>71.0</u> (72.4)	<u>79.0</u> (79.4)	<u>83.3</u> (83.3)	<u>85.3</u> (85.3)
SigLIP2	Prototype	<u>58.0</u> (63.4)	<u>69.7</u> (73.2)	<u>77.0</u> (79.4)	<u>80.8</u> (82.6)	<u>83.0</u> (84.5)
	LDA	2.3 (2.3)	<u>73.2</u> (74.5)	<u>80.5</u> (81.0)	<u>84.5</u> (84.7)	<u>86.5</u> (86.5)
DINOv2	Prototype	<u>59.6</u>	67.1	71.8	76.0	78.2
	LDA	2.3	69.2	75.3	80.3	83.3

Class distances. Fig. 5 (left) shows that both SigLIP and SigLIP2 produce cosine similarity distributions where there are many image pairs from the same class less similar to each other than some image pairs from a different class, mirroring the introductory cat-cat cat-dog example in Fig. 1. The finding that also SigLIP2 manifests the high overlaps in the histograms further supports our alternative hypothesis that this is not a sign of misalignment, but an expected and desirable property of general-purpose pretrained vision encoders. Moreover, the result suggests that SigLIP2’s improvement over SigLIP is not due to a better class separation that can be visualized through the used histogram.

Modality gap. Fig. 5 (right) leads to the same conclusion. The divergence in similarity distributions between inter-modal (image-text) and intra-modal (image-image) pairs is nearly identical for both models. The fact that SigLIP2’s additional image-image training does not close this gap demonstrates that it is a misalignment introduced by a missing image-image objective. We do not see that this divergence is harmful for intra-modal tasks. In fact, the different range can be explained straightforward as a consequence of

the modality gap: Because image and text embeddings cluster in two distinct regions in the embedding space, distances measured within the cluster are naturally smaller than between clusters. We discuss this in context of findings from prior work in Sec. 5.

4.3. Comparison on the Few-Shot Classification Task

Setup. We study few-shot classification performance in two complementary settings: (1) the *standard VLM few-shot adaptation* setting [50, 51] where a text prompt provides a zero-shot classifier that is refined using few-shot labeled images; and (2) an *image-only few-shot* setting with no text prompts, designed to isolate the intrinsic alignment quality of the image embeddings. The former corresponds to Fig. 6, the latter to Tab. 2.

Few-Shot VLM Adaptation. Following prior work, CLIP ResNet-50 results are reported. Fig. 6 summarizes performance across 11 datasets. Methods such as TIP-X and Coder were explicitly designed to circumvent potential intra-modal misalignment by avoiding image-image similarity. However, the results show that these approaches do not outperform simple image-space discriminative baselines such as APE feature selection or Linear Discriminant Analysis (LDA).¹ The reported PCA^{\leftarrow} results simply use the class mean of the projected few-shot samples as classifier. It performs competitively with Tip-X and Coder, supporting the view that classification in the image space is sufficient, without requiring the avoidance of image-image similarity.

Image-Only Few-Shot Classification. We evaluate few-shot classification *without* any text prompts, more akin to a classical uni-modal few-shot learning setup [6]. Classification is performed measuring image-image cosine-similarity to prototypes (class means) or by estimating the classifier as in [42]. The two classifiers are tested with original and PCA^{\leftarrow} projected embeddings. CLIP ViT-B/16 is used.

Across 11 datasets, Tab. 2 evidences how SigLIP (inter-modal only) outperforms DINOv2 (image-only) in image-image similarity based classification. This ranking is consistent across classifiers. While comparing models that have been trained on different data does not allow for conclusions about the impact of the training objective (loss function), the experiment still reveals a notable result relevant to the misalignment hypothesis: pure contrastive language-image training can yield image embeddings that are aligned well enough to outperform strong pure vision encoders. In other words, the results lack any evidence of misalignment.

¹While in the original work the method was introduced as Gaussian Discriminant Analysis (GDA), we here refer to it as Linear Discriminant Analysis (LDA) because they assume a shared identical covariance matrix for all classes.

4.4. Comparison on the Retrieval Task

Setup. We compare six models that we classify in inter-modal and intra-modal, following our contrasting approach. Inter-modal models are those with no image self-supervised objective: CLIP B/32, L/14 and SigLIP. Intra-modal models are SLIP [27], SigLIP2 and DINO, that all have some self-supervised loss on the images. CLIP, SigLIP and SLIP were also used in the OTI work [26], SigLIP2 and DINO were added by us in order to cover the intra-modal models.

Results. Tab. 3 reports our image-to-image retrieval results. Our proposed PCA^{\leftarrow} alternative consistently outperforms OTI across all datasets and across all models.

We consider this critical as it proves two key points our paper makes. First, it shows comparing image embeddings works well, such that there is no reason to assume a misalignment and thus no need to invert images to pseudo-texts as in OTI and Fig. 3. Second, it reveals no experimental outcome is specific to pure text-image training as in CLIP. Specifically, the original models SigLIP and SigLIP2 perform comparable, with only 1.2 mAP gain on SigLIP2, which is an expected improvement given it is an evolution of SigLIP, but no large gap which would point towards a significant problem in SigLIP with pure text-image training. Also the degree to which the PCA^{\leftarrow} projection helps is comparable, with 5.6 and 5.8 mAP gains, respectively.

We include DINOv3 L/16 as a reference upper limit. Interestingly, it does not always score highest. On Stanford Cars, it scores 8.6 mAP lower than our SigLIP B/16 PCA^{\leftarrow} from the inter-modal category. On the other hand, DINO excels on the more standard retrieval datasets ROxford and RParis. We attribute this to dataset curation: ROxford/RParis are designed for unambiguous retrieval, while the other datasets (to our knowledge only used in [26]) suffer from query ambiguity akin to one-shot classification. Their retrieval mAPs we show here are relevant for the study of the misalignment hypothesis, but for the task ambiguity reason, we caution against viewing them as generally meaningful retrieval benchmarks in future work.

Ablation. To test our explanation why PCA^{\leftarrow} can improve performance metrics, we conduct an experiment on a task where we assume it is *not* helpful to focus on the dominant concept in the image. Classifying the weather and time of the day requires an image representation that captures more than the main class, as background information matters. We therefore measure retrieval metrics on “weather” and “time of day” classification on the BDD100k driving dataset. Unlike on all datasets in Tab. 3, we find a performance *drop* by 2.3% mAP on weather and 1.0% mAP on time of day compared to original CLIP (ViT-L/14). This confirms that PCA^{\leftarrow} ’s information removal only brings gains when the dataset labels correspond to the main semantic concept in the images (Fig. 1), but does not fix some misalignment in general.

Table 3. **Image-to-image retrieval** results (mAP) following [26]. Rows with $\langle I, I \rangle$ or $\langle I^{\leftarrow}, I^{\leftarrow} \rangle$ measure intra-modal similarities. We show the performance gain brought by OTI [26] - an image-to-text conversion attempt motivated by the hypothesis that only $\langle T, I \rangle$ is aligned - can be significantly surpassed through simply measuring $\langle I^{\leftarrow}, I^{\leftarrow} \rangle$ along the main axes of variance of ImageNet class names. This trend is no different between models without intra-modal and with intra-modal training objectives.

	ViT	Method	$\langle \cdot, \cdot \rangle$	$\mathcal{R}_{\text{Oxford}}$	$\mathcal{R}_{\text{Paris}}$	Cars	Pets	Flowers	Aircraft	DTD	EuroSAT	Food101	SUN397	Caltech	UCF101	ImageNet	Average
CLIP	B/32	Original	$\langle I, I \rangle$	42.4	74.0	24.9	31.2	62.5	14.5	28.3	49.3	33.6	34.6	77.7	46.2	21.6	41.6
		OTI	$\langle T, I \rangle$	43.0	70.3	28.0	37.5	62.6	14.4	31.9	47.2	34.7	36.3	79.9	48.6	23.8	42.9
		PCA^{\leftarrow}	$\langle I^{\leftarrow}, I^{\leftarrow} \rangle$	51.4	80.9	34.6	47.7	70.3	16.1	34.0	53.8	43.0	40.0	83.3	53.3	28.6	49.0
	L/14	Original	$\langle I, I \rangle$	57.1	77.8	43.8	47.2	84.2	25.8	33.9	57.8	55.0	39.2	83.8	59.5	33.0	53.7
		OTI	$\langle T, I \rangle$	62.4	77.1	50.5	56.0	86.0	27.1	37.7	56.3	55.9	43.5	87.3	62.8	38.2	57.0
		PCA^{\leftarrow}	$\langle I^{\leftarrow}, I^{\leftarrow} \rangle$	64.5	83.0	57.2	62.7	88.9	28.7	39.9	62.8	64.4	46.0	89.5	66.8	42.2	61.3
Sig LIP	B/16	Original	$\langle I, I \rangle$	50.6	73.1	65.7	56.4	87.5	37.9	39.8	53.3	56.3	42.8	87.2	56.9	35.8	57.2
		OTI	$\langle T, I \rangle$	55.2	79.1	71.8	64.2	89.7	37.6	43.3	52.9	59.0	43.6	88.9	54.9	38.8	60.0
		PCA^{\leftarrow}	$\langle I^{\leftarrow}, I^{\leftarrow} \rangle$	57.9	78.4	77.2	68.5	92.0	40.9	44.2	54.2	61.8	46.9	91.2	60.3	43.5	62.8
SLIP	B/16	Original	$\langle I, I \rangle$	36.5	79.2	4.9	17.8	65.7	9.1	29.7	53.7	19.5	26.1	65.4	40.2	15.4	35.6
		OTI	$\langle T, I \rangle$	36.4	79.3	5.0	19.3	65.1	9.0	30.5	50.6	20.0	26.4	67.6	40.6	14.8	35.7
		PCA^{\leftarrow}	$\langle I^{\leftarrow}, I^{\leftarrow} \rangle$	43.6	83.9	6.2	23.1	72.5	10.3	32.2	53.2	25.1	30.7	70.7	42.4	19.1	39.4
Sig LIP2	B/16	Original	$\langle I, I \rangle$	52.5	75.6	70.8	56.6	89.3	40.7	38.6	49.3	59.7	43.0	89.2	59.2	37.9	58.6
		PCA^{\leftarrow}	$\langle I^{\leftarrow}, I^{\leftarrow} \rangle$	59.4	78.6	80.2	67.8	93.2	46.3	44.1	51.1	65.3	48.9	93.0	63.0	46.5	64.4
	DINOv3	L/16	$\langle I, I \rangle$	91.4	93.5	68.6	82.3	99.4	39.8	44.9	59.9	71.0	50.1	90.7	70.6	57.6	70.8

5. Discussion

Is the modality gap actually a problem? Given that previous papers titled “Mind the gap” [20], “Mitigate the Gap” [10], “Cross the gap” [26], “Bridging the Gap” [2], the gap may appear as a problem. The resulting different distributions of cosine similarities (Fig. 2b, Fig. 5b) further made some [2, 40] believe something is intra-modally misaligned. However, we failed to identify any such misalignment that negatively impacts image-only tasks. Attempts of previous work to narrow the gap could also not bring consistent improvement [16, 20, 26]. Eslami *et al.* [10] suggest adding an intra-modal separation loss reduces the gap while helpful for downstream performance, but this result is confined to their from-scratch trained modified shared text-vision encoder architecture. Qian *et al.* [31] theoretically demonstrate that the gap cannot be reduced sufficiently by minimizing the contrastive loss in CLIP. Jiang *et al.* [16] further prove that exact modality alignment is suboptimal in general for downstream prediction tasks. We therefore see it as reasonable, not problematic, that image and text embeddings lie in two separate manifolds.

What about other tasks? E.g. segmentation, depth estimation, VQA? – We only evaluate on retrieval and few-shot classification because these were the tasks that were previously studied in the context of the intra-modal misalignment hypothesis. To our knowledge, literature on other image-image applications of CLIP has not mentioned the misalignment issue. If supposed misalignment turns out to be a non-issue for the previously studied tasks, we do not

expect other tasks to be impacted, either. As for the PCA^{\leftarrow} method, we do not believe it will be useful for these tasks.

Did we disprove the hypothesis? For the experimental results, strictly speaking, no. Tabs. 1 to 3 and Fig. 5 prove intra-modal misalignment was not the cause of the observed trends, but do not prove it is impossible to exist. Similar to [35], the refutation logic is to show that what previously was believed to be evidence for a hypothesis, can actually not serve as such. For the theoretical argument in Sec. 3 and Appx. A, yes it seeks to refute the belief in the underlying cause. But ultimately, the hypotheses in prior work vary to some extent, so we believe it is more useful and appropriate to consider the evidence for the individual arguments than to speak of holistic proof/disproof.

6. Conclusion

In this work, we reevaluate the intra-modal misalignment hypothesis in CLIP-style vision-language models. Through theoretical analysis, we demonstrate that intra-modal similarities are not unconstrained with arbitrary degrees of freedom, but determined by cross-modal similarities. Empirically, we reexamine previously proposed indicators and performance metrics, showing they are no reliable diagnostic of intra-modal misalignment. Comparison with a PCA-style projection method further supports our alternative hypothesis that task ambiguity in the few-shot setting, not misalignment, best explains the results. We hope this can guide future work to leverage the pretrained intra-modal geometry rather than circumvent it.

Acknowledgment This work was supported by the National Nature Science Foundation of China under Grant No. 62522317 and 62373322, and by Zhejiang Provincial Natural Science Foundation of China under Grant No. LD25F030001.

References

- [1] Alberto Baldrati, Lorenzo Agnolucci, Marco Bertini, and A. Bimbo. Zero-shot composed image retrieval with textual inversion. In *ICCV*, 2023. 5
- [2] Clément Barbier, Baptiste Abello, and Stéphane Herbin. Bridging the modality gap: Training-free adaptation of vision-language models for remote sensing via visual prototypes. In *CVPR Workshops*, 2025. 1, 2, 3, 5, 8
- [3] Yassir Bendou, Amine Ouasfi, Vincent Gripon, and Adnane Boukhayma. Proker: A kernel perspective on few-shot adaptation of large vision-language models. In *CVPR*, 2025. 3
- [4] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101—mining discriminative components with random forests. In *ECCV*, 2014. 5
- [5] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, 2021. 2
- [6] Wei-Yu Chen, Yen-Cheng Liu, Zsolt Kira, Y. Wang, and Jia-Bin Huang. A closer look at few-shot classification. In *ICLR*, 2019. 7
- [7] Mircea Cimpoi, Subhansu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. *CVPR*, 2014. 5
- [8] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 5
- [9] Jeremy Elson, John R. Douceur, Jon Howell, and Jared Saul. Asirra: a captcha that exploits interest-aligned manual image categorization. In *Conference on Computer and Communications Security*, 2007. 3
- [10] Sedigheh Eslami and Gerard de Melo. Mitigate the gap: Investigating approaches for improving cross-modal alignment in clip. In *ICLR*, 2025. 8
- [11] Patrick Esser, Johnathan Chiu, Parmida Atighehchian, Jonathan Granskog, and Anastasis Germanidis. Structure and content-guided video synthesis with diffusion models. In *ICCV*, 2023. 3
- [12] Matteo Farina, Massimiliano Mancini, Giovanni Iacca, and Elisa Ricci. Rethinking few-shot adaptation of vision-language models in two stages. In *CVPR*, 2025. 3
- [13] Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit H. Bermano, Gal Chechik, and Daniel Cohen-Or. An image is worth one word: Personalizing text-to-image generation using textual inversion. In *ICLR*, 2023. 3
- [14] Robert Geirhos, Priyank Jaini, Austin Stone, Sourabh Medapati, Xi Yi, George Toderici, Abhijit Ogale, and Jonathon Shlens. Towards flexible perception with visual memory. In *ICML*, 2025. 3
- [15] 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 Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 12(7), 2019. 5
- [16] Qian Jiang, Changyou Chen, Han Zhao, Liqun Chen, Q. Ping, Son Dinh Tran, Yi Xu, Belinda Zeng, and Trishul M. Chilimbi. Understanding and constructing latent modality structures in multi-modal representation learning. In *CVPR*, 2023. 8
- [17] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *CVPR workshops*, 2013. 5
- [18] Shuo Li, Fang Liu, Zehua Hao, Xinyi Wang, Lingling Li, Xu Liu, Puhua Chen, and Wenping Ma. Logits deconfusion with clip for few-shot learning. In *CVPR*, 2025. 2
- [19] Feng Liang, Bichen Wu, Xiaoliang Dai, Kunpeng Li, Yanan Zhao, Hang Zhang, Peizhao Zhang, Péter Vajda, and Diana Marculescu. Open-vocabulary semantic segmentation with mask-adapted clip. In *CVPR*, 2023. 1
- [20] 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*, 2022. 2, 8
- [21] Zhiqiu Lin, Samuel Yu, Zhiyi Kuang, Deepak Pathak, and Deva Ramana. Multimodality helps unimodality: Cross-modal few-shot learning with multimodal models. In *CVPR*, 2023. 3
- [22] Subhansu Maji, Esa Rahtu, Juho Kannala, Matthew Blaschko, and Andrea Vedaldi. Fine-grained visual classification of aircraft. *arXiv preprint arXiv:1306.5151*, 2013. 5
- [23] Matthias Minderer, Alexey A. Gritsenko, Austin Stone, Maxim Neumann, Dirk Weissenborn, Alexey Dosovitskiy, Aravindh Mahendran, Anurag Arnab, Mostafa Dehghani, Zhuoran Shen, Xiao Wang, Xiaohua Zhai, Thomas Kipf, and Neil Houlsby. Simple open-vocabulary object detection with vision transformers. In *ECCV*, 2022. 1
- [24] Matthias Minderer, Alexey A. Gritsenko, and Neil Houlsby. Scaling open-vocabulary object detection. In *NeurIPS*, 2023. 1
- [25] Yifei Ming and Yixuan Li. Understanding retrieval-augmented task adaptation for vision-language models. In *ICML*, 2024. 2, 3
- [26] Marco Mistretta, Alberto Baldrati, Lorenzo Agnolucci, Marco Bertini, and Andrew D. Bagdanov. Cross the gap: Exposing the intra-modal misalignment in clip via modality inversion. In *ICLR*, 2025. 1, 2, 3, 5, 7, 8
- [27] Norman Mu, Alexander Kirillov, David A. Wagner, and Saining Xie. Slip: Self-supervision meets language-image pre-training. In *ECCV*, 2022. 7
- [28] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *Indian conference on computer vision, graphics & image processing*, 2008. 5
- [29] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Q. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Mahmoud Assran, Nicolas Ballas, Wojciech Galuba, Russ

- Howes, Po-Yao (Bernie) Huang, Shang-Wen Li, Ishan Misra, Michael G. Rabbat, Vasu Sharma, Gabriel Synnaeve, Huijiao Xu, Hervé Jégou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. DINOv2: Learning robust visual features without supervision. *ArXiv arXiv:2304.0719*, 2023. 3
- [30] Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, and CV Jawahar. Cats and dogs. In *CVPR*, 2012. 5
- [31] Qi Qian, Yuanhong Xu, and Juhua Hu. Intra-modal proxy learning for zero-shot visual categorization with clip. In *NeurIPS*, 2023. 8
- [32] Filip Radenović, Ahmet Iscen, Giorgos Tolias, Yannis Avrithis, and Ondřej Chum. Revisiting oxford and paris: Large-scale image retrieval benchmarking. In *CVPR*, 2018. 5
- [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*, 2021. 1
- [34] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *CVPR*, 2023. 3
- [35] Rylan Schaeffer, Brando Miranda, and Oluwasanmi Koyejo. Are emergent abilities of large language models a mirage? In *NeurIPS*, 2023. 8
- [36] Oriane Siméoni, Huy V. Vo, Maximilian Seitzer, Federico Baldassarre, Maxime Oquab, Cijo Jose, Vasil Khalidov, Marc Szafraniec, Seungeun Yi, Michaël Ramamonjisoa, Francisco Massa, Daniel Haziza, Luca Wehrstedt, Jianyuan Wang, Timothée Darcet, Théo Moutakanni, Leonel Sentana, Claire Roberts, Andrea Vedaldi, Jamie Tolan, John Brandt, Camille Couprie, Julien Mairal, Hervé Jégou, Patrick Labatut, and Piotr Bojanowski. DINOv3. *arXiv preprint arXiv:2508.10104*, 2025. 3
- [37] 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. 5
- [38] Christoph Timmermann, Hyunse Lee, and Woojin Lee. Semobridge: Semantic modality bridge for efficient few-shot adaptation of clip. *arXiv preprint arXiv:2509.26036*, 2025. 2
- [39] Michael Tschannen, Alexey Gritsenko, Xiao Wang, Muhammad Ferjad Naeem, Ibrahim M. Alabdulmohsin, Nikhil Parthasarathy, Talfan Evans, Lucas Beyer, Ye Xia, Basil Mustafa, Olivier H'enaff, Jeremiah Harmsen, Andreas Steiner, and Xiao-Qi Zhai. Siglip 2: Multilingual vision-language encoders with improved semantic understanding, localization, and dense features. *arXiv-preprint arXiv:2502.14786*, 2025. 4
- [40] Vishaal Udandarao, Ankush Gupta, and Samuel Albanie. Sus-x: Training-free name-only transfer of vision-language models. In *CVPR*, 2023. 1, 2, 4, 5, 6, 8
- [41] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. 2011. 5
- [42] Zhengbo Wang, Jian Liang, Lijun Sheng, Ran He, Zilei Wang, and Tieniu Tan. A hard-to-beat baseline for training-free clip-based adaptation. In *ICLR*, 2024. 3, 6, 7
- [43] Jianxiong Xiao, James Hays, Krista A Ehinger, Aude Oliva, and Antonio Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *CVPR*, 2010. 5
- [44] Chao Yi, Lu Ren, De-Chuan Zhan, and Han-Jia Ye. Leveraging cross-modal neighbor representation for improved clip classification. In *CVPR*, 2024. 1, 2, 3, 6
- [45] Fisher Yu, Haofeng Chen, Xin Wang, Wenqi Xian, Yingying Chen, Fangchen Liu, Vashisht Madhavan, and Trevor Darrell. Bdd100k: A diverse driving dataset for heterogeneous multitask learning. In *CVPR*, 2018. 5
- [46] Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language image pre-training. In *ICCV*, 2023. 4
- [47] Renrui Zhang, Rongyao Fang, Peng Gao, Wei Zhang, Kunchang Li, Jifeng Dai, Yu Qiao, and Hongsheng Li. Tip-adapter: Training-free clip-adapter for better vision-language modeling. In *ECCV*, 2022. 3, 5
- [48] Zhixing Zhang, Bichen Wu, Xiaoyan Wang, Yaqiao Luo, Luxin Zhang, Yanan Zhao, Peter Vajda, Dimitris Metaxas, and Licheng Yu. Avid: Any-length video inpainting with diffusion model. In *CVPR*, 2024. 3
- [49] Chong Zhou, Chen Change Loy, and Bo Dai. Extract free dense labels from clip. In *ECCV*, 2022. 1
- [50] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. In *IJCV*, 2022. 5, 7
- [51] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. In *CVPR*, 2022. 5, 7
- [52] Beier Zhu, Jiequan Cui, Han Chao Zhang, and Chi Zhang. Project-probe-aggregate: Efficient fine-tuning for group robustness. In *CVPR*, 2025. 5
- [53] Xiangyang Zhu, Renrui Zhang, Bowei He, A-Long Zhou, Dong Wang, Bingyan Zhao, and Peng Gao. Not all features matter: Enhancing few-shot clip with adaptive prior refinement. In *ICCV*, 2023. 3, 6
- [54] Xingyu Zhu, Beier Zhu, Shuo Wang, Kesen Zhao, and Han Chao Zhang. Enhancing clip robustness via cross-modality alignment. In *NeurIPS*, 2025. 5