

CANONICALIZING MULTIMODAL CONTRASTIVE REPRESENTATION LEARNING

Sharut Gupta*
MIT CSAIL
sharut@mit.edu

Sanyam Kansal*
IIT Kanpur
sanyamka23@iitk.ac.in

Stefanie Jegelka
MIT CSAIL, TU Munich
stefje@mit.edu

Phillip Isola
MIT CSAIL
phillipi@mit.edu

Vikas Garg
Aalto University, YaiYai Ltd
vgarg@csail.mit.edu

ABSTRACT

As models and data scale, independently trained networks often induce analogous notions of similarity. Yet, similarity-based measures are weaker than precise correspondence maps between distinct models. This gap is even more consequential for multimodal models, where convergence must hold not only within each modality but also for the learned image–text coupling. We therefore ask whether, given two independently trained multimodal contrastive models, there exists a single transformation that simultaneously aligns both their image and text representations? In this work, we show that the answer is yes, and, even more, the joint transformation is a simple orthogonal map. Across contrastive multimodal families, an orthogonal map $Q \in O(d)$ fit using only a few paired data from a single modality (e.g., images) i.e., $f(x) \approx Qf(x)$, also aligns the other modality (text), $\tilde{g}(y) \approx Qg(y)$. We quantify this transfer in the target text space by large gains in pointwise cosine similarity and class-level nearest neighbour accuracy after transformation. Theoretically, we show that the agreement of the multimodal similarity kernel $\langle f(x), g(t) \rangle \approx \langle f(x), \tilde{g}(y) \rangle$ on a small, finite set of points forces a shared orthogonal map Q across modalities. Broadly, this finding enables backward-compatible model upgrades, avoiding costly re-embedding, and has implications for the privacy of learned representations.

1 INTRODUCTION




Figure 1: *A single orthogonal map aligns two independent contrastive models across modalities.* (a) Illustrative PCA schematic on synthetic embeddings from two models (A and B) shows that the spaces are a priori incomparable; (b) A single orthogonal map Q , fit only using image embeddings of CLIP (OpenAI) and CLIP (LAION), almost perfectly aligns image embeddings of the two models (c, left) and simultaneously aligns their text embeddings, as evidenced by a large gain in text-text pointwise cosine similarity (c, right).

A recurring question in modern representation learning is *convergence*: as models and data scale, do independently trained networks—across datasets, architectures, and training runs—recover similar internal representations? A growing body of evidence suggests they often do, in the sense that different

systems can induce similar notions of similarity over inputs or contain universal neurons and “circuits” (Huh et al., 2024; Merullo et al., 2022; Gupta et al., 2026; Chughtai et al., 2023; Gurnee et al., 2024; Zimmermann et al., 2021). This idea is central to the *Platonic Representation Hypothesis* (PRH), which posits that, at a sufficiently large scale, learned embeddings converge towards a shared representation that reflects the underlying structure of the world (Huh et al., 2024). Empirically, this convergence is commonly studied through representational similarity analyses such as SVCCA and CKA (Raghu et al., 2017; Kornblith et al., 2019; Huh et al., 2024) over unimodal co-occurrence kernels.

However, much of this literature abstracts away parameterization. SVCCA is invariant to affine transformations (Raghu et al., 2017), and CKA compares induced similarity structure rather than the precise geometric correspondence between them (Kornblith et al., 2019). Two models may agree on many tasks while their internal embedding spaces remain related only by complex, sample-dependent distortions. From a geometric standpoint, the stronger and far more consequential question is whether independently trained models recover representations that are equivalent up to simple transformations. This question becomes especially important for *multimodal* models, which couple image and text through a contrastive objective while keeping the two modalities at arm’s length in the learned space, a phenomenon often referred to as the *modality gap*. To make this precise, let $\mathcal{M} := (f, g)$ and $\tilde{\mathcal{M}} := (\tilde{f}, \tilde{g})$ denote two distinct multimodal models, with image encoders f, \tilde{f} and text encoders g, \tilde{g} mapping inputs to embedding spaces. Here, it no longer suffices to ask whether image representations f and \tilde{f} of \mathcal{M} and $\tilde{\mathcal{M}}$ converge in isolation (or text encoders g and \tilde{g} of the two models converge in isolation). Instead, the central question becomes:

Given two independently trained multimodal models, does a systematic geometric relationship exist between their embedding spaces? If so, what is its form, and how does it differ across modalities?

In this work, we study this question for multimodal *contrastive* models and show that two independently trained instances—with different embedding dimensions, training distributions and modeling choices—exhibit a remarkably rigid, *modality-invariant* geometric relationship. Concretely, across model families such as CLIP (Radford et al., 2021), SigLIP (Zhai et al., 2023), and FLAVA (Singh et al., 2022), we find that the inter-model relationship is well-approximated by a *single orthogonal map* (up to a global mean shift) i.e., there exists an orthogonal map Q where $Q^\top Q = I$ such that $\tilde{f}(x) \approx Qf(x)$ for paired images x .^{*} Moreover, the *same* Q simultaneously aligns text encoders: $\tilde{g}(y) \approx Qg(y)$ for texts y (as shown in Figure 1). This induces a commuting correspondence between encoders; once Q is learned, any embedding produced by one model—image or text—can be mapped into the other model’s coordinate system and back and compared meaningfully with any embedding there.

Empirically, despite never using text to estimate Q , applying this map to text substantially improves cross-model agreement in the target text space, as measured both by (i) the mean cosine similarity between matched text embeddings after mapping, and (ii) prompt retrieval, in which each mapped prompt is matched to its nearest neighbor among the target model’s class prompts and scored by whether it selects the correct class. In the same aligned space, nearest-neighbor image classification using mapped source embeddings matches the target model’s performance, indicating that Q preserves semantic details and task-relevant geometry. Moreover, this transfer is data-efficient, requiring only about $\sim 30\%$ data to learn Q reliably. Finally, Q learned on one dataset transfers to others without re-fitting, and is consistent under composition, consistent with a global geometric relationship rather than instance- or class-specific tuning.

We complement these findings by theoretically characterizing when this coupling is guaranteed. At the population level, we derive the optimal multimodal contrastive critic and show that, on a fixed target domain, agreement of cross-modal similarity kernel i.e. $\langle f, g \rangle = \langle \tilde{f}, \tilde{g} \rangle$ on a small set of anchor pairs, forces a shared orthogonal map across modalities: the map that aligns images simultaneously determines the induced alignment of text. We further move beyond the exact regime, proving stability bounds that quantify how approximate cross-modal kernel alignment $\langle f, g \rangle \approx \langle \tilde{f}, \tilde{g} \rangle$ translates into reliable alignment.

^{*}Equivalently, Q can be learned from either modality: estimating it from paired text aligns images as well.

To summarize, the key contributions of our work are:

- We show that independently trained multimodal contrastive models can be closely approximated by a *single orthogonal* map. Additionally, this map is shared across modalities, i.e., estimating the map from images alone aligns text, and vice versa.
- Theoretically, we prove that matching multimodal kernels on a small anchor set across two distinct models forces a shared orthogonal alignment across modalities and derive stability bounds in the approximate regime.
- We validate these claims across five benchmarks and multiple model pairs, with extensive ablations showing that this map transfers across datasets without re-fitting and remains consistent under composition, yielding the most reliable cross-model, cross-modal transfer.

2 RELATED WORK

For an extended discussion of related work, see [Appendix A](#).

Representational Convergence and Functional Interoperability. A central question in deep learning is whether independently trained models converge to identical representations. The Platonic Representation Hypothesis (Huh et al., 2024) argues that large models across different modalities are converging towards the same representations, given the vast amounts of training data that are used by these models. Standard embedding similarity tools such as CKA (Kornblith et al., 2019) or SVCCA (Raghu et al., 2017) measure this convergence up to broad equivalence classes (e.g., invertible linear maps), but do not provide an explicit coordinate mapping. A stronger operational test is *model stitching*, which connects representations via simple learnable transformations to enable zero-shot interchangeability (Lenc & Vedaldi, 2015; Bansal et al., 2021; Merullo et al., 2022). However, these approaches are limited to coarse transfer metrics (e.g., image captioning), whereas we measure a stronger notion of alignment—tight pointwise agreement via cosine similarity—while simultaneously verifying that semantic structure is preserved through retrieval performance. Prior alignment results largely operate on *unimodal* marginals, including word embeddings (Mikolov et al., 2013; Dev et al., 2021) and vision features (Maystre et al., 2025; Merullo et al., 2022). Aligning marginals, however, does not in general identify the *joint* distribution: multiple distinct joint geometries may be consistent with the same unimodal alignments. We therefore study a strictly stronger question—whether the *joint* image-text geometry of two multimodal models is identifiable up to a *single, rigid* isometry *shared across modalities*, rather than allowing separate, unconstrained maps.

The Modality Gap and Geometry of Contrastive Representations. Our study is grounded in the geometry of contrastive vision-language models (Radford et al., 2021; Jia et al., 2021). Empirically, these models exhibit a *modality gap*, where image and text embeddings cluster in distinct cones (Liang et al., 2022; Udandarao, 2022; Shi et al., 2023). This separation renders alignment non-trivial, as naive mappings can easily collapse the gap or distort the intra-modal structure. Theoretically, prior work has analyzed conditions under which contrastive objectives identify the underlying latent factors up to linear or affine transformations (Zimmermann et al., 2021; Roeder et al., 2021). While general linear maps include rotations, they also permit shear and anisotropic scaling, which are poorly constrained and undesirable for preserving semantic structure. In contrast, our setting restricts attention to *isometries*, which preserve angles, norms, and neighborhood relations. Even though images and text remain separated within each model, we prove that multimodal kernel agreement on a *small* anchor set suffices to recover a *shared* isometry Q that aligns the two models across *both* modalities. As a result, instead of retraining a model or learning complex transfer operators, one can anchor a new model to a reference model using a modest number of examples and obtain transfer to another modality *for free*. This yields a substantially more economical alternative to retraining, general linear transfer, or optimal-transport-based adaptations.

3 PROBLEM FORMULATION

We consider the standard dual-encoder framework where data consists of co-occurring pairs (x, y) (e.g., images and text). A contrastive model consists of two encoders, $f : \mathcal{X} \rightarrow \mathbb{S}^{d-1}$ and $g : \mathcal{Y} \rightarrow \mathbb{S}^{d-1}$, which map inputs to the unit hypersphere in \mathbb{R}^d . The training objective maximizes the cosine similarity $\langle f(x), g(y) \rangle$ for semantically matched pairs while minimizing it for mismatched ones.

Consider two contrastive models, trained in complete isolation on different datasets, with different architectures, initializations, and modeling choices: a *source* model $\mathcal{M} = (f, g)$ and a *target* model $\tilde{\mathcal{M}} = (\tilde{f}, \tilde{g})$ mapping to dimensions d and \tilde{d} respectively (without loss of generality, assume $d \leq \tilde{d}$). Due to optimization stochasticity and training differences, the embedding spaces of \mathcal{M} and $\tilde{\mathcal{M}}$ are a priori incomparable. Even with identical data, the contrastive objective depends only on within-model dot products $\langle f(x), g(y) \rangle$, so jointly rotating both embeddings by any orthogonal matrix leaves the loss and all within-model similarities unchanged. Architectural mismatch, finite-sample noise, and optimization effects further amplify this ambiguity (Saunshi et al., 2022; Robinson et al., 2021); under distribution shift, the models may not even share the same population optimum. As a result, cross-model dot products and nearest-neighbor queries between \mathcal{M} and $\tilde{\mathcal{M}}$ are ill-defined unless we first learn a map between their spaces. We therefore seek a map $\psi : \mathbb{R}^d \rightarrow \mathbb{R}^{\tilde{d}}$ that transports embeddings from \mathcal{M} into $\tilde{\mathcal{M}}$. The key question is whether a consistent geometric relationship exists between the two models, and how it depends on modality: must one learn separate maps for images and text, or does a single ψ align both modalities? Concretely, if ψ aligns images so that $\tilde{f} \approx \psi(f)$, does the same ψ also align text, i.e., $\tilde{g} \approx \psi(g)$? Further, if such a map exists, what is its functional form: nonlinear, linear, or orthogonal?

4 MODALITY GAP IN CONTRASTIVE REPRESENTATIONS

This question is nontrivial due to the intrinsic geometry of contrastive representations. If matched image-text pairs collapsed to approximately the same point on the hypersphere (i.e., $f(x) \approx g(y)$), then aligning the image manifold would be equivalent to aligning the text manifold. In practice, however, contrastive models exhibit a pronounced *modality gap* where image and text embeddings occupy largely disjoint regions of the sphere (Liang et al., 2022; Shi et al., 2023; Udandarao, 2022). Prior work also suggests that naively “closing” this gap can harm downstream performance and fairness (Liang et al., 2022), as shown in Figure 2.




Figure 2: (left) PCA visualization of embeddings from CLIP (OpenAI) on CIFAR 100, showing a large modality gap. (right) Applying a translation to close this gap distorts the embedding geometry and significantly degrades classification accuracy.

Since the gap reflects systematic modality-specific structure rather than noise (Schrodi et al., 2024), it is *not* obvious that a map ψ learned solely on images will extend correctly to text: In principle, infinitely many maps can agree on the image manifold yet behave arbitrarily on the text manifold. Our work specifically tests whether, despite this gap, the *relative* geometry remains stable enough to permit transfer.




Figure 3: (a) Across CLIP variants, the multimodal kernel $\langle f(x), g(y) \rangle$ (relative angles between image and text embeddings) is strongly preserved (dashed lines), unlike the unimodal kernel $\langle f(x), \tilde{f}(x') \rangle$; (b) CKA on multimodal kernels shows high alignment across models.

Despite the modality gap and disjoint supports of the two models, we argue that the alignment problem is indeed solvable because *relative* geometry is remarkably stable. As shown in Figure 3,

while the absolute coordinates of the embedding cones shift arbitrarily between models, the angular arrangement of the texts with respect to the images remains consistent. Mathematically, this means that the *multimodal kernels* are approximately preserved across models: $\langle f, g \rangle \approx \langle \tilde{f}, \tilde{g} \rangle$. This observation can be viewed as a multimodal analogue of the Platonic Representation Hypothesis (Huh et al., 2024) that posits that models converge to similar *unimodal* kernels ($\langle f, f \rangle \approx \langle \tilde{f}, \tilde{f} \rangle$). In the next section, we prove that this preservation of multimodal kernels is a sufficient condition to constrain the functional form of ψ , forcing it to be an isometry.

5 THEORETICAL PERSPECTIVES

In this section, we formalize the above intuition by characterizing population-level optima and showing that contrastive models, even when trained on different distributions, can recover the same multimodal kernels up to a constant (Section 5.2, as shown in Figure 3). We then show that this agreement on a minimal anchor set constrains the alignment map to be linear (Section 5.3), which, under the hypersphere constraint, collapses to an isometry (Section 5.4). Next, we relax these conditions, proving that the recovered isometry preserves zero-shot retrieval even when pointwise alignment is imperfect (Appendix C.5.4). Finally, we extend these guarantees to the approximate setting, showing stability under approximate kernel matching (Appendix C.6). For a schematic overview of our theoretical analysis, refer to Figure 8 in the appendix.

5.1 CONTRASTIVE REPRESENTATION LEARNING

Building on the dual-encoder framework defined in Section 3, we assume access to data characterized by co-occurring pairs $(x, y) \in \mathcal{X} \times \mathcal{Y}$ (e.g., images and text) drawn from a joint distribution p_{XY} with marginals p_X and p_Y . Here, the contrastive objective aims to learn functions, $f : \mathcal{X} \rightarrow \mathbb{S}^{d-1}$ and $g : \mathcal{Y} \rightarrow \mathbb{S}^{d-1}$, which map inputs to a shared unit hypersphere $\mathbb{S}^{d-1} \subset \mathbb{R}^d$ where similarity of a pair (x, y) is measured by the score function

$$s(x, y) := \frac{1}{\tau} \langle f(x), g(y) \rangle + b, \tag{1}$$

where $\tau > 0$ is a temperature parameter and $b \in \mathbb{R}$ is a scalar bias. This alignment is achieved by optimizing the symmetric InfoNCE objective (Oord et al., 2018). Specifically, for a batch of N pairs $\{(x_i, y_i)\}_{i=1}^N$ drawn i.i.d. from P_{XY} , the loss \mathcal{L}_N minimizes the cross-entropy of identifying the correct positive pair relative to $N - 1$ negatives in both directions:

$$\mathcal{L}_N(s) := \mathcal{L}_N^{(X)}(s) + \mathcal{L}_N^{(Y)}(s) = \mathbb{E} \left[- \sum_{i=1}^N \log \frac{e^{s(x_i, y_i)}}{\sum_{j=1}^N e^{s(x_i, y_j)}} - \sum_{i=1}^N \log \frac{e^{s(x_i, y_i)}}{\sum_{j=1}^N e^{s(x_j, y_i)}} \right]. \tag{2}$$

5.2 OPTIMAL CRITIC FOR FAMILY OF CONTRASTIVE LEARNERS

We first characterize the score function $s(x, y)$ that minimizes the population InfoNCE loss. Let $r(x, y) := \frac{p_{XY}(x, y)}{p_X(x)p_Y(y)}$ denote the pointwise density ratio, and $\log r(x, y)$ be the Pointwise Mutual Information (PMI). For a fixed training distribution P_{XY} , any global minimizer of the symmetric objective \mathcal{L}_N induces an optimal score $s^*(x, y) = \log r(x, y) + C$, i.e. the pointwise mutual information (PMI) up to a global constant. We note that this has been partially studied in prior work (Huh et al., 2024; Oord et al., 2018), with a detailed derivation provided in Appendix C.2 and Appendix C.3. Consequently, any two globally optimal models trained on the *same* distribution—or on distributions related to P_{XY} by a bijective reparameterization—induce kernels that differ only by a constant (Corollary C.6). Here, we go beyond the fixed-distribution setting and prove that this invariance persists even for models trained on *distinct* internet-scale corpora.

Platonic distribution and Dataset Curation. Let P_{XY}^* denote an underlying “reality” distribution with density p_{XY}^* and marginals p_X^*, p_Y^* (satisfying positivity on the domain of interest). We model each training corpus as a *curation* of P_{XY}^* i.e. for dataset $a \in \{1, 2\}$, there exist measurable weights $u_a : \mathcal{X} \rightarrow (0, \infty)$ and $v_a : \mathcal{Y} \rightarrow (0, \infty)$ such that

$$p_{XY}^{(a)}(x, y) \propto u_a(x) v_a(y) p_{XY}^*(x, y). \tag{3}$$

Here u_a and v_a capture modality-specific curation (e.g., image quality and text language/safety etc.) that can substantially change the marginals while preserving the underlying semantic information. We further assume that curation acts *independently across modalities* in the sense that the expected text acceptance does not depend on the image it is paired with, and vice versa i.e.

$$\begin{aligned}\mathbb{E}^*[v_a(Y) | X = x] &= \mathbb{E}^*[v_a(Y)] \quad \text{a.e. in } x, \\ \mathbb{E}^*[u_a(X) | Y = y] &= \mathbb{E}^*[u_a(X)] \quad \text{a.e. in } y.\end{aligned}\tag{4}$$

where \mathbb{E}^* corresponds to expectation under p^* and a.e. stands for almost everywhere, causing these expectation terms to appear in the PMI only as additive constants.

Theorem 5.1. *Under Equation (4), if s_1^* and s_2^* are Bayes-optimal scores for the contrastive models trained on two distinct distributions defined in Equation (3), then there exists a constant Δ' such that*

$$s_1^*(x, y) = s_2^*(x, y) + \Delta' \text{ for } P_X^* \otimes P_Y^* \text{-a.e. } (x, y).\tag{5}$$

Theorem 5.1 establishes that even when trained on different distributions, independently trained contrastive learners can converge to optimal similarity scores that agree up to an additive constant. We note that our result holds for several widely used contrastive objectives such as softmax InfoNCE (Oord et al., 2018) and pairwise sigmoid objectives in SigLIP (Zhai et al., 2023). Since contrastive models approximate this target using dot products, any two models converging to the same score must implicitly align their kernels: $\langle f(x), g(y) \rangle \approx \langle \tilde{f}(x), \tilde{g}(y) \rangle$.^{*} For detailed proofs and discussion, refer to Appendix C.4.

5.3 LINEAR ALIGNMENT OF CONTRASTIVE MODELS

In the previous section, we showed that independently trained contrastive models can induce multimodal kernels that agree up to an additive constant, despite differences in data and modeling choices. We now assume this kernel agreement holds on domains of interest $\Omega_X \subseteq \mathcal{X}$ and $\Omega_Y \subseteq \mathcal{Y}$ (e.g., a downstream image-text dataset) and analyze its geometric consequences. Our first main result shows that matching kernels on a small set of *anchor* points suffices to determine a linear map relating the two embedding spaces.

Let the contrastive model pairs (f, g) and (\tilde{f}, \tilde{g}) map inputs to $\mathbb{S}^{d-1} \subset \mathbb{R}^d$ and $\mathbb{S}^{\tilde{d}-1} \subset \mathbb{R}^{\tilde{d}}$ respectively, where $d \leq \tilde{d}$, without loss of generality. We fix a set of *image anchors* $\{\bar{x}_j\}_{j=1}^d \subset \Omega_X$ and *text anchors* $\{\bar{y}_i\}_{i=1}^{\tilde{d}} \subset \Omega_Y$ and collect their embeddings into the following matrices: $G := [g(\bar{y}_1) \cdots g(\bar{y}_{\tilde{d}})] \in \mathbb{R}^{d \times \tilde{d}}$, $\tilde{G} := [\tilde{g}(\bar{y}_1) \cdots \tilde{g}(\bar{y}_{\tilde{d}})] \in \mathbb{R}^{\tilde{d} \times \tilde{d}}$, $F := [f(\bar{x}_1) \cdots f(\bar{x}_d)] \in \mathbb{R}^{d \times d}$.

Assumption 5.2. The multimodal kernels coincide on the set of anchors:

$$\begin{aligned}\langle f(x), g(\bar{y}_i) \rangle &= \langle \tilde{f}(x), \tilde{g}(\bar{y}_i) \rangle \quad \forall x \in \Omega_X, \forall i \in \{1, \dots, \tilde{d}\}, \\ \langle f(\bar{x}_j), g(y) \rangle &= \langle \tilde{f}(\bar{x}_j), \tilde{g}(y) \rangle \quad \forall y \in \Omega_Y, \forall j \in \{1, \dots, d\}.\end{aligned}$$

Theorem 5.3. *(Linear Identifiability, proof in Appendix C.5.1). Under Assumption 5.2, suppose \tilde{G} and F are invertible. Then there exists a linear map A such that $\tilde{f}(x) = Af(x) \forall x \in \Omega_X$. Further, if A has full column rank, then for every $y \in \Omega_Y$, $\text{Proj}_{\text{Im}(A)} \tilde{g}(y) = A(A^\top A)^{-1}g(y)$. If $\tilde{d} = d$, then $\tilde{g}(y) = A^{-\top}g(y)$.*

5.4 ISOMETRIC ALIGNMENT OF CONTRASTIVE MODELS

Theorem 5.3 shows that kernel matching identifies the representation up to a linear map A . But, contrastive encoders normalize embeddings to the unit hypersphere \mathbb{S}^{d-1} , forcing $\|\tilde{f}(x)\|_2 = \|Af(x)\|_2 = 1$ everywhere. This forces A to be an isometry ($A^\top A = I_d$) only if the data is sufficiently diverse to probe the matrix in all directions. We formalize this diversity via the following condition.

^{*}Exact equality implies the target PMI respects the contrastive parameterization constraints. Since CLIP scores are bounded and low-rank, they act as a low-rank approximation of the generally unbounded, high-rank population target.

Definition 5.4. (*Sym(d)-spanning*) A set of vectors $S \subset \mathbb{R}^d$ is *Sym(d)-spanning* if the rank-one matrices $\{xx^\top : x \in S\}$ span the space of symmetric matrices $\text{Sym}(d)$. Equivalently: if $M \in \text{Sym}(d)$ and $x^\top Mx = 0$ for all $x \in S$, then $M = 0$. This equivalence follows from the identity $x^\top Mx = \langle xx^\top, M \rangle$.

Theorem 5.5. (*Orthogonal Identifiability, proof in Appendix C.5.2*). Assume the conditions of Theorem 5.3 hold. If the set of image embeddings $\{f(x) : x \in \Omega_X\}$ contains a *Sym(d)-spanning subset*, then the linear map A has orthonormal columns ($A^\top A = I_d$). Consequently, $\tilde{f}(x) = Qf(x) \quad \forall x \in \Omega_X$, where $Q := A$ satisfies $Q^\top Q = I_d$. Furthermore, for the other modality:

$$\text{Proj}_{\text{Im}(Q)} \tilde{g}(y) = Qg(y) \quad \forall y \in \Omega_Y.$$

If $\tilde{d} = d$, Q is orthogonal i.e. $Q \in O(d)$ and $\tilde{g}(y) = Qg(y)$.

In Appendix C.5.3, we extend Theorem 5.5 to the case where $\text{span}\{f(x) : x \in \Omega_X\}$ lies in low-dimensions.

Throughout the preceding analysis, we assumed exact kernel equality. In Appendix C.6, we further relax Assumption 5.2 to an ϵ -approximate bound: $|\langle f(x), g(y) \rangle - \langle \tilde{f}(x), \tilde{g}(y) \rangle| \leq \epsilon$ and prove that the map Q becomes an approximate isometry.

6 THE PROCRUSTES ALGORITHM

Guided by the theoretical guarantees in Section 5, we translate the alignment problem into an optimization procedure. We align the source and target manifolds of one modality (say images) using a set of N unlabelled anchor images $\{x_i\}_{i=1}^N$. Let $X \in \mathbb{R}^{d \times N}$ and $\tilde{X} \in \mathbb{R}^{\tilde{d} \times N}$ be the data matrices containing the centered, normalized embeddings $f(x_i)$ and $\tilde{f}(x_i)$ as columns. We solve for the optimal isometry \hat{Q} by minimizing the transport cost subject to an orthogonality constraint: $\hat{Q} = \text{argmin}_{Q^\top Q = I} \|\tilde{X} - QX\|_F^2$. This is the classic *Orthogonal Procrustes Problem*, which has a closed-form solution via the Singular Value Decomposition (SVD) of the cross-covariance matrix $M = \tilde{X}X^\top$. Let $U\Sigma V^\top = \text{SVD}(M)$, then $\hat{Q} = UI_{\tilde{d} \times d}V^\top$. Here, $I_{\tilde{d} \times d}$ is a rectangular identity matrix. This formulation naturally handles $d < \tilde{d}$, where \hat{Q} becomes a semi-orthogonal i.e. $\hat{Q}^\top \hat{Q} = I_d$. For additional details and pseudocode, refer to Appendix D.4.1

7 EXPERIMENTAL RESULTS

In this section, we empirically evaluate our central claim across various benchmarks and configurations, leading to three main takeaways: (i) a single orthogonal map Q accurately captures the relationship between independently trained contrastive models and, crucially, the same Q applies to both image and text representations (Section 7.2); (ii) Q is data-efficient and only a few examples suffice to estimate it; it generalizes to unseen classes and even to new downstream datasets without re-fitting (Sections 7.3 and 7.4); and (iii) although more expressive linear or non-linear maps can increase pointwise similarity on the fitted domain, they fail to transfer to the second modality and distort task-relevant geometry, degrading downstream retrieval; enforcing orthogonality in contrast yields the most reliable transfer across models and modalities (Section 7.5). Finally, we show that the learned orthogonal map approximately commutes with cross-modal retrieval across models i.e. direct image alignment and text-mediated alignment recover the same semantic neighborhoods across models (Appendix E.12).

7.1 TRAINING PROTOCOL

We report each metric both before alignment and after applying the learned orthogonal map Q , and average results over three random seeds. We describe the evaluated model pairs, datasets, and metrics below and in detail in Appendix D.4.

Models. We evaluate three independently trained vision-language pairs: (i) CLIP ViT-B/32 (OpenAI) and CLIP ViT-B/32 trained on LAION-400M; (ii) CLIP ViT-L/14 (OpenAI) and SIGLIP; and (iii) CLIP ViT-L/14 (OpenAI) and FLAVA (Radford et al., 2021; Schuhmann et al.,

2021; Zhai et al., 2023; Singh et al., 2022). We ℓ_2 -normalize all embeddings such that dot products equal cosine similarity.

Datasets. We report results on Oxford-IIIT Pets (Parkhi et al., 2012), CIFAR-100 (Krizhevsky et al., 2009), Caltech-101 (Fei-Fei et al., 2004), STL10 (Coates et al., 2011) and DTD (Cimpoi et al., 2014). For more information about creating text prompts, refer to Appendix D.3.1. We report results only for Oxford Pets in the main paper and defer results on the remaining datasets to Appendix E.

Training. We learn the alignment across models using the standard orthogonal Procrustes solution described in Section 6. In practice, the two models can differ by a constant offset in embedding space due to finite-sample effects and dataset mismatch. We therefore fit and apply Q on centered embeddings, and then re-add the target mean i.e. $z \mapsto Q(z - \mu_A^{(\cdot)}) + \mu_B^{(\cdot)}$, where $\mu_A^{(\cdot)}$ and $\mu_B^{(\cdot)}$ are modality-specific means of the source and target embeddings (image or text, as appropriate). Centering isolates the rotational relationship by removing this offset while preserving the orthogonal correspondence in the centered space* Theoretically, this mean offset vanishes when the two models are exactly related by an orthogonal map, as discussed in Remark C.15.

Performance Metrics. We report five evaluation metrics: (1) Paired-instance cosine similarity, measured between aligned and target embeddings for either images or texts; (2) top-1 retrieval across models, evaluated for both image-image and text-text retrieval by nearest-neighbor matching at the class level. Zero-shot classification for (3) aligned images against target text (aligned image-text), (4) target images against aligned text (image-aligned text), and (5) both images and text aligned (aligned image-aligned text). All metrics are computed using cosine similarity; full metric definitions are deferred to Appendix D.4.2.

7.2 INDEPENDENTLY TRAINED CONTRASTIVE MODELS DIFFER BY AN ORTHOGONAL MAP COMMON TO BOTH MODALITIES




Figure 4: Inter-model alignment on Oxford-Pets before and after fitting a single orthogonal map Q from paired images. (a) Image-image cosine similarity; (b) Text-text cosine similarity; (c) Aligned Image and Aligned text class retrieval accuracy. Here, Model A and Model B denote each model’s within-model image-to-text baseline. Q aligns images across models with a single orthogonal map, and the same Q learned only from image embeddings transfers to text, boosting text-text cosine from near-chance to near-oracle, all while preserving strong image classification accuracy.

Figure 4 summarizes our findings on Oxford-Pets across three independently trained pairs.

An Orthogonal Map Aligns Different Models. First, from Figure 4(a), we find that a single orthogonal map almost perfectly aligns image embeddings across distinct multimodal contrastive models, improving the image-image cosine similarity from near zero to $\sim 0.8 - 0.9$. We observe analogous findings for text embeddings (see Figure 34)(c), indicating that independently trained contrastive models are related by an approximately orthogonal map.

This Map Transfers Across Modalities. Second, and more importantly, this map is modality-invariant: Figure 4(b) shows that the same orthogonal map Q fit using paired images sharply improves text alignment, significantly improving text-text cosine similarity across model pairs. Finally, as shown in Figure 4(c), aligned-image-to-aligned-text retrieval remains high, showing

*We find that mean-centering has a negligible effect on class-level retrieval and decision geometry; it primarily changes pointwise cosine agreement (see Appendix E.10). Thus, a pure orthogonal map on raw embeddings suffices for semantic alignment and preserves decision geometry.

that Q preserves task-relevant geometry while eliminating any need to compute the second model’s text embeddings. Additionally, in some cases, Q effectively transfers model A’s stronger decision geometry into model B’s space, matching or even exceeding model B’s native performance. Results across additional datasets and metrics appear in [Appendix E.3](#).

We extend these findings to mismatched embedding dimensions in [Appendix E.8](#) and also show the reverse direction i.e. fitting Q on text to align images in [Appendix E.6](#). Finally, we ablate mean-centering and find that it has negligible effect on class-level retrieval and decision geometry, and mainly affects pointwise cosine agreement. Thus, a pure orthogonal map on raw embeddings suffices for semantic alignment and preserves decision geometry (see [Appendix E.10](#)).

7.3 ONLY A FEW DATA POINTS ARE NEEDED TO LEARN THE ORTHOGONAL MAP




Figure 5: Generalization of orthogonal alignment under limited supervision (Oxford Pets; CLIP (OpenAI) and FLAVA) as measured in aligned-image-to-aligned text accuracy. Q learned from a few classes transfers and generalizes to unseen classes.




Figure 6: Generalization of the orthogonal map across data distributions across CLIP (OpenAI) and FLAVA. (left) text-text cosine similarity; (right) image classification accuracy post-transformation using aligned images and aligned texts. A map Q fit on Oxford-Pets transfers to Caltech-101 (and vice versa).

[Theorem 5.5](#) states that if the multimodal kernels induced by two contrastive models agree on a sufficiently rich but *small finite* set of anchors, a single global orthogonal map aligns their representations across both modalities. As shown in [Figure 5](#), we empirically validate this on Oxford-Pets, fitting Q using paired images from only N classes and evaluating transfer on the remaining unseen classes. Here, Model A and Model B denote each model’s within-model image-to-text baseline. Performance on both seen and unseen classes improves quickly with just a few anchor classes and essentially saturates once N reaches a modest value (around 10-15 classes), after which additional anchors provide little benefit. Thus, practitioners can recover near-full cross-model transfer by fitting Q on a lightweight image-only calibration set, rather than curating large-scale cross-model supervision. For additional metrics and results across model pairs and datasets, refer to [Appendix E.4](#).

7.4 THE ORTHOGONAL MAP GENERALIZES BROADLY

The previous experiment shows that Q is identifiable from a few anchors and generalizes to unseen classes within the same dataset. We next ask a stronger version of this question: does the *same* Q transfer to a completely new downstream distribution *without* re-fitting? From [Figure 6](#) (left), a map learned on Oxford-Pets substantially increases text-text cosine similarity on Caltech-101 (and vice versa). Correspondingly, on the right, aligned-image-to-aligned-text classification remains strong under transfer—often closely matching or even exceeding an in-domain fit—indicating that Q generalizes beyond the calibration dataset.

7.5 EVALUATING ALTERNATIVE ALIGNMENT MAPS THAN THE ORTHOGONAL MAPPING

Here, we ablate the alignment design by comparing three maps of increasing expressiveness: (i) an orthogonal map Q , (ii) a linear map, and (iii) a non-linear MLP. As shown in [Figure 7](#), more expressive maps improve *pointwise* image-image cosine similarity. However, these maps transfer poorly to the text modality and distort the image-text geometry. In contrast, the orthogonal

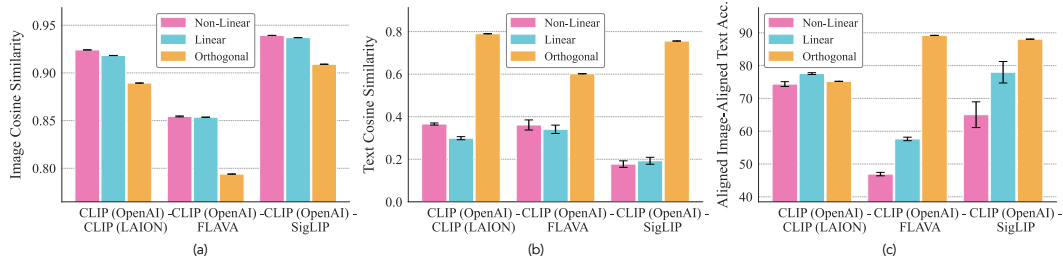


Figure 7: Comparison of alignment strategies on Oxford Pets across model pairs in terms of (a) image cosine similarity; (b) text cosine similarity, and (c) aligned-image-to-aligned-text accuracy. Unlike orthogonal maps, linear and non-linear maps distort task-relevant geometry and do not transfer to text, leading to poor downstream performance.

map consistently performs best on both pointwise text cosine similarity and geometry-sensitive downstream metrics. Extended results for additional datasets are provided in [Appendix E.9](#).

Additional Ablations. In [Appendix E.7](#), we show that \mathcal{Q} remains consistent under *cycle* and *composition*. When we hold the training data fixed and vary only the design choices, transfer is even stronger than under dataset shift (as shown in [Appendix E.11](#)). Finally, in [Appendix E.13](#) we show that \mathcal{Q} preserves fine-grained attributes (pose, etc.), beyond coarse class-level semantics.

8 DISCUSSION AND CONCLUSION

Conclusion. In this work, we show a rigid form of geometric convergence in multimodal contrastive models: across independently trained systems (with different data and modeling choices), a single orthogonal map learned in one modality can approximately align both image and text representations, inducing a shared coordinate system. Moreover, estimating this map requires only a small anchor set from a single modality (image or text). Theoretically, we characterize conditions under which agreement of the multimodal similarity kernel forces such a shared isometry and establish guarantees even under approximate agreement.

Discussion and Implications. Our results have several practical and scientific implications. In large embedding systems, switching models typically triggers full re-embedding, often infeasible at modern scale (billions of vectors) ([Jayaram Subramanya et al., 2019](#); [Johnson et al., 2019](#)) and costly in both time and compute ([OpenAI, 2024](#)). We show that a small anchor set can recover the orthogonal map that restores compatibility across models. Since it preserves inner products, it supports model upgrades without re-encoding while keeping the embedding geometry intact. Finally, models often specialize differently; one might have a stronger vision tower, while another has a stronger or multilingual text tower. Our approach lets practitioners swap and combine towers while preserving image-text geometry. For an extended discussion, refer to [Appendix B](#).

Limitations and Future Work. Our evaluation focuses on classification-style semantics; we do not establish gains for fine-grained retrieval or dense ranking. Although an orthogonal map preserves angles, we do not test whether fine-grained attributes remain easily decodable after alignment; a natural next step is to train lightweight decoders on the aligned space. Finally, we study image-text contrastive encoders; extending to other modalities (e.g., audio, video) is an important direction.

9 ACKNOWLEDGEMENTS

S.G. acknowledges the support of the MathWorks Engineering Fellowship. P.I. acknowledges support from a Packard Fellowship, the MIT-IBM Watson AI Lab, and the ONR MURI grant N00014-22-1-2740. S.J. acknowledges the support of the NSF AI Institute TILOS (NSF CCF2112665) and the Alexander von Humboldt Foundation. V.G. acknowledges the support from Saab-WASP (grant 411025), Academy of Finland (grant 342077), and the Jane and Aatos Erkkö Foundation (grant 7001703). We thank Kiril Bangachev for thorough and insightful discussions.

REFERENCES

- Samuel K Ainsworth, Jonathan Hayase, and Siddhartha Srinivasa. Git re-basin: Merging models modulo permutation symmetries. *arXiv preprint arXiv:2209.04836*, 2022.
- Sanjeev Arora, Hrishikesh Khandeparkar, Mikhail Khodak, Orestis Plevrakis, and Nikunj Saunshi. A theoretical analysis of contrastive unsupervised representation learning. *arXiv preprint arXiv:1902.09229*, 2019. URL <https://arxiv.org/abs/1902.09229>.
- Kiril Bangachev, Guy Bresler, Iliyas Noman, and Yury Polyanskiy. Global minimizers of sigmoid contrastive loss. *arXiv preprint arXiv:2509.18552*, 2025.
- Yamini Bansal, Preetum Nakkiran, and Boaz Barak. Revisiting model stitching to compare neural representations. In *Advances in Neural Information Processing Systems*, 2021. URL <https://arxiv.org/abs/2106.07682>.
- Bilal Chughtai, Lawrence Chan, and Neel Nanda. A toy model of universality: Reverse engineering how networks learn group operations. In *International Conference on Machine Learning*, pp. 6243–6267. PMLR, 2023.
- Mircea Cimpoi, Subhansu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3606–3613, 2014.
- Adam Coates, Andrew Ng, and Honglak Lee. An analysis of single-layer networks in unsupervised feature learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pp. 215–223. JMLR Workshop and Conference Proceedings, 2011.
- Corinna Cortes, Mehryar Mohri, and Afshin Rostamizadeh. Algorithms for learning kernels based on centered alignment. *Journal of Machine Learning Research*, 13(28):795–828, 2012. URL <https://jmlr.org/papers/v13/cortes12a.html>.
- Nello Cristianini, Andre Elisseeff, John Shawe-Taylor, and Jaz Kandola. On kernel-target alignment. In *Advances in Neural Information Processing Systems*, 2002. URL <https://papers.nips.cc/paper/1946-on-kernel-target-alignment.pdf>.
- Sunipa Dev, Safia Hassan, and Jeff M Phillips. Closed form word embedding alignment. *Knowledge and Information Systems*, 63(3):565–588, 2021.
- Felix Draxler, Kambis Veschgini, Manfred Salmhofer, and Fred Hamprecht. Essentially no barriers in neural network energy landscape. In *International conference on machine learning*, pp. 1309–1318. PMLR, 2018.
- Rahim Entezari, Hanie Sedghi, Olga Saukh, and Behnam Neyshabur. The role of permutation invariance in linear mode connectivity of neural networks. *arXiv preprint arXiv:2110.06296*, 2021.
- 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. In *2004 conference on computer vision and pattern recognition workshop*, pp. 178–178. IEEE, 2004.
- Timur Garipov, Pavel Izmailov, Dmitrii Podoprikin, Dmitry P Vetrov, and Andrew G Wilson. Loss surfaces, mode connectivity, and fast ensembling of dnns. *Advances in neural information processing systems*, 31, 2018.
- Sharut Gupta, Joshua Robinson, Derek Lim, Soledad Villar, and Stefanie Jegelka. Structuring representation geometry with rotationally equivariant contrastive learning. *International Conference on Learning Representations (ICLR)*, 2023.
- Sharut Gupta, Shobhita Sundaram, Chenyu Wang, Stefanie Jegelka, and Phillip Isola. Better together: Leveraging unpaired multimodal data for stronger unimodal models. *International Conference on Learning Representations (ICLR)*, 2026.
- Wes Gurnee, Theo Horsley, Zifan Carl Guo, Tara Rezaei Kheirkhah, Qinyi Sun, Will Hathaway, Neel Nanda, and Dimitris Bertsimas. Universal neurons in gpt2 language models. *arXiv preprint arXiv:2401.12181*, 2024.

Jeff Z HaoChen, Colin Wei, Adrien Gaidon, and Tengyu Ma. Provable guarantees for self-supervised deep learning with spectral contrastive loss. *Advances in neural information processing systems*, 34:5000–5011, 2021.

(First name unknown) Hu et al. Learning backward compatible embeddings. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2022.

Minyoung Huh, Brian Cheung, Tongzhou Wang, and Phillip Isola. The platonic representation hypothesis. *arXiv preprint arXiv:2405.07987*, 2024. URL <https://arxiv.org/abs/2405.07987>.

Young Kyun Jang and Ser-nam Lim. Towards cross-modal backward-compatible representation learning for vision-language models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2025. URL https://openaccess.thecvf.com/content/ICCV2025/papers/Jang_Towards_Cross-modal_Backward-compatible_Representation_Learning_for_Vision-Language_Models_ICCV_2025_paper.pdf.

Suhas Jayaram Subramanya, Fnu Devvrit, Harsha Vardhan Simhadri, Ravishankar Krishnawamy, and Rohan Kadekodi. Diskann: Fast accurate billion-point nearest neighbor search on a single node. *Advances in neural information processing systems*, 32, 2019.

Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 4904–4916. PMLR, 18–24 Jul 2021. URL <https://proceedings.mlr.press/v139/jia21b.html>.

Jeff Johnson, Matthijs Douze, and Hervé Jégou. Billion-scale similarity search with gpus. *IEEE Transactions on Big Data*, 7(3):535–547, 2019.

Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey Hinton. Similarity of neural network representations revisited. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 3519–3529. PMLR, 2019. URL <https://proceedings.mlr.press/v97/kornblith19a.html>.

Nikolaus Kriegeskorte, Marieke Mur, and Peter A. Bandettini. Representational similarity analysis—connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience*, 2:4, 2008. doi: 10.3389/neuro.06.004.2008. URL <https://www.frontiersin.org/articles/10.3389/neuro.06.004.2008/full>.

Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images.(2009), 2009.

Karel Lenc and Andrea Vedaldi. Understanding image representations by measuring their equivariance and equivalence. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. URL https://openaccess.thecvf.com/content_cvpr_2015/html/Lenc_Understanding_Image_Representations_2015_CVPR_paper.html.

Weixin Liang, Yuhui Zhang, Yongchan Kwon, Serena Yeung, and James Zou. Mind the gap: Understanding the modality gap in multi-modal contrastive representation learning. In *Advances in Neural Information Processing Systems*, 2022.

Lucas Maystre et al. When embedding models meet: Procrustes bounds and first-order equivalences for embedding interoperability. *arXiv preprint arXiv:2510.13406*, 2025. URL <https://arxiv.org/abs/2510.13406>.

Qiang Meng, Chixiang Zhang, Xiaoqiang Xu, and Feng Zhou. Learning compatible embeddings. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 9939–9948, 2021.

- Jack Merullo, Louis Castricato, Carsten Eickhoff, and Ellie Pavlick. Linearly mapping from image to text space. *arXiv preprint arXiv:2209.15162*, 2022.
- Tomas Mikolov, Quoc V. Le, and Ilya Sutskever. Exploiting similarities among languages for machine translation. *arXiv preprint arXiv:1309.4168*, 2013. URL <https://arxiv.org/abs/1309.4168>.
- Mazda Moayeri, Keivan Rezaei, Maziar Sanjabi, and Soheil Feizi. Text-to-concept (and back) via cross-model alignment. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 25037–25060. PMLR, 23–29 Jul 2023. URL <https://proceedings.mlr.press/v202/moayeri23a.html>.
- Ari S. Morcos, Maithra Raghu, and Samy Bengio. Insights on representational similarity in neural networks with canonical correlation. In *Advances in Neural Information Processing Systems*, 2018. URL <https://arxiv.org/abs/1806.05759>.
- Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. URL <https://arxiv.org/abs/1807.03748>.
- OpenAI. Api pricing. <https://openai.com/api/pricing/>, 2024. Accessed: January 2026.
- Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, and CV Jawahar. Cats and dogs. In *2012 IEEE conference on computer vision and pattern recognition*, pp. 3498–3505. IEEE, 2012.
- 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 *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 8748–8763. PMLR, 2021. URL <https://proceedings.mlr.press/v139/radford21a.html>.
- Maithra Raghu, Justin Gilmer, Jason Yosinski, and Jascha Sohl-Dickstein. Svcca: Singular vector canonical correlation analysis for deep learning dynamics and interpretability. In *Advances in Neural Information Processing Systems*, 2017. URL <https://arxiv.org/abs/1706.05806>.
- J. Robinson, L. Sun, K. Yu, K. Batmanghelich, S. Jegelka, and S. Sra. Can contrastive learning avoid shortcut solutions? In *Neural Information Processing Systems (NeurIPS)*, 2021.
- Geoffrey Roeder, Luke Metz, and Durk Kingma. On linear identifiability of learned representations. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 9030–9039. PMLR, 18–24 Jul 2021. URL <https://proceedings.mlr.press/v139/roeder21a.html>.
- N. Saunshi, J. Ash, S. Goel, D. Misra, C. Zhang, S. Arora, S. Kakade, and A. Krishnamurthy. Understanding contrastive learning requires incorporating inductive biases. In *International Conference on Machine Learning (ICML)*, 2022.
- Simon Schrodi, David T Hoffmann, Max Argus, Volker Fischer, and Thomas Brox. Two effects, one trigger: On the modality gap, object bias, and information imbalance in contrastive vision-language representation learning. *arXiv preprint arXiv:2404.07983*, 2024.
- Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. *arXiv preprint arXiv:2111.02114*, 2021. URL <https://arxiv.org/abs/2111.02114>.
- Yantao Shen, Yuanjun Xiong, Wei Xia, and Stefano Soatto. Towards backward-compatible representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6368–6377, 2020.

- Peiyang Shi, Michael C Welle, Mårten Björkman, and Danica Kragic. Towards understanding the modality gap in clip. In *ICLR 2023 workshop on multimodal representation learning: perks and pitfalls*, 2023.
- Amanpreet Singh et al. FLAVA: A foundational language and vision alignment model. *arXiv preprint arXiv:2112.04482*, 2022. URL <https://arxiv.org/abs/2112.04482>.
- Vishaal Udandarao. Understanding and fixing the modality gap in vision-language models. *Master's thesis, University of Cambridge*, 32, 2022.
- Tongzhou Wang and Phillip Isola. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In Hal Daumé III and Aarti Singh (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 9929–9939. PMLR, 13–18 Jul 2020. URL <https://proceedings.mlr.press/v119/wang20k.html>.
- Xiaohua Zhai et al. Sigmoid loss for language image pre-training. *arXiv preprint arXiv:2303.15343*, 2023. URL <https://arxiv.org/abs/2303.15343>.
- Roland S. Zimmermann, Yash Sharma, Steffen Schneider, Matthias Bethge, and Wieland Brendel. Contrastive learning inverts the data generating process. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 12979–12990. PMLR, 18–24 Jul 2021. URL <https://proceedings.mlr.press/v139/zimmermann21a.html>.
- Liu Ziyin and Isaac Chuang. Proof of a perfect platonic representation hypothesis. *arXiv preprint arXiv:2507.01098*, 2025.

APPENDIX

A Additional Related Works	17
B Additional Discussion and Implications	18
C Proof of Theoretical Results	18
C.1 Setup and Notations	18
C.2 Bayes Posterior and Optimal Critic for Unimodal Contrastive Learners	20
C.3 Optimal Critic for Multimodal Contrastive Learners	20
C.4 Bayes-Optimal Scores Are Identifiable Up to a Constant	21
C.5 Guarantees Under Exact Alignment	22
C.5.1 Linear Alignment of Independent Contrastive Models	22
C.5.2 Isometric Alignment of Independent Contrastive Models	23
C.5.3 Isometric Alignment of Independent Contrastive Models When One Modality Lies in Low Dimension	24
C.5.4 Alignment Upto Classification Boundaries	24
C.6 Guarantees Under Approximate Alignment	25
C.6.1 Approximate Linear Alignment of Independent Contrastive Models	26
C.6.2 Approximate Isometric Alignment of Independent Contrastive Models	26
C.6.3 Approximate Isometric Alignment of Independent Contrastive Models When One Modality Lies in Low Dimension	27
D Supplementary Experimental Details and Assets Disclosure	28
D.1 Assets	28
D.2 Hardware and setup	28
D.3 Datasets	28
D.3.1 Image Classification Benchmarks	28
D.3.2 Constructing Text Templates	28
D.4 Training and Evaluation Protocol	30
D.4.1 Learning An Orthogonal Transformation	30
D.4.2 Performance Metrics at Evaluation	31
E Additional Experiments	32
E.1 Modality Gap in Contrastive Models	32
E.2 Visualization of Multimodal Kernels	32
E.3 Independently Trained Contrastive Models Differ by an Orthogonal Map That Is Shared Across Modalities	33
E.4 Only a Few Data Points Are Needed to Learn the Orthogonal Map	38
E.5 The Learned Orthogonal Map Generalizes Broadly	45
E.6 Learning the Orthogonal Map from Text Instead of Images Transfers to Images	49

E.7 Cycle Consistency and Consistency Under Composition	52
E.8 Alignment Across Embedding Dimensions.	59
E.9 Evaluating Alternative Alignment Maps Than The Orthogonal Mapping	62
E.10 Orthogonal Alignment With and Without Centering	65
E.11 Isolating Architecture Effects Under Fixed Training Data	68
E.12 Qualitative Evidence of Commutativity of Image-Text Alignment Paths	70
E.13 Fine-grained Semantic Preservation Under Orthogonal Maps.	70

A ADDITIONAL RELATED WORKS

Representational Convergence. A long-standing theme in representation learning is whether independently trained networks converge to comparable internal representations. This question is frequently studied through representational similarity analyses that compare activations up to broad equivalence classes, including RSA (Kriegeskorte et al., 2008), CCA-based methods such as SVCCA (Raghu et al., 2017) and its refinements (Morcos et al., 2018), and kernel-based comparisons such as CKA (Kornblith et al., 2019). Earlier work also probed convergence by explicitly matching units or subspaces across independently trained networks (e.g., via neuron matching and sparse prediction), highlighting that models may learn similar *subspaces* even when individual coordinates do not align. These tools have been influential for documenting empirical convergence trends and motivating hypotheses such as the Platonic Representation Hypothesis (PRH) (Huh et al., 2024). PRH formalizes convergence at the level of induced similarity structure: two representations are considered aligned when their (co-occurrence) kernels agree on corresponding inputs. However, by construction, many similarity measures are invariant to broad classes of transformations (e.g., invertible linear maps) or compare induced kernels rather than producing an explicit coordinate-level correspondence. Our work targets this stronger notion of convergence: whether independently trained *multimodal embedding spaces* agree in (or can be brought into) a shared coordinate system via a simple global map.

Model Stitching and Functional Interoperability. Beyond similarity indices, several works test whether independently trained representations become *interchangeable* under simple transformations such as linear or orthogonal maps. Early studies introduced *stitching* layers to test equivalence between networks by swapping intermediate features (Lenc & Vedaldi, 2015), and subsequent work systematized stitching as a methodology for comparing learned representations and their compositionality across training regimes (Bansal et al., 2021). In parallel, (Mikolov et al., 2013) showed that independently trained word embedding spaces can often be aligned by a single linear map learned from limited supervision, suggesting that semantic structure is preserved up to a global change of basis. More recently, linear aligners have been used to translate representations between modern pretrained models, such as mapping vision features into CLIP space (Moayeri et al., 2023), and to study when embeddings from different models are mutually transferable via low-complexity maps (Maystre et al., 2025). These approaches provide evidence that two networks encode similar information, even when coordinates do not match directly.

Symmetries, Global Minima, and Landscape Connectivity. The existence of many apparently distinct solutions is also consistent with known symmetries of neural networks and the geometry of the loss landscape. Empirically, SGD solutions are often connected by low-loss paths (mode connectivity) (Garipov et al., 2018; Draxler et al., 2018), and accounting for permutation symmetries of hidden units can further reduce apparent barriers between independently trained models (Entezari et al., 2021). Recent work makes this operational by explicitly *rebasin*g one model to another via permutation alignment to enable weight-space merging (Ainsworth et al., 2022). In this context, our results can be viewed as a representation-space analogue: while weights admit large symmetry groups, we find that multimodal embedding spaces often differ primarily by an (approximately) orthogonal transform that is shared across modalities.

Vision-Language Contrastive Pretraining and the Modality Gap. Large-scale vision-language representation learning is dominated by dual-encoder contrastive objectives, as exemplified by CLIP (Radford et al., 2021) and ALIGN (Jia et al., 2021), with many variants exploring alternative losses, scaling strategies, and training recipes (e.g., Zhai et al., 2023; Singh et al., 2022). A recurring geometric observation in such models is the *modality gap*: image and text embeddings can occupy systematically shifted regions even after normalization, with consequences for optimization and downstream behavior (Liang et al., 2022; Udandarao, 2022; Shi et al., 2023). Our findings complement this line of work by showing that, despite the within-model modality gap, cross-model relationships exhibit a surprising rigidity: different models’ image and text spaces are often related by the *same* global orthogonal map, so that aligning one modality effectively determines the other.

Theory of Contrastive Learning and Kernel Alignment. Theoretical analyses of contrastive objectives have clarified what is identifiable from paired data and how representation geometry emerges. Key work includes contrastive predictive coding and InfoNCE-style objectives (Oord et al., 2018), general theoretical frameworks and guarantees for contrastive representation learning (Arora et al., 2019; HaoChen et al., 2021; Bangachev et al., 2025), and geometric decompositions into

alignment and uniformity (Wang & Isola, 2020). Gupta et al. (2023) show that minimizing a loss which preserves unimodal kernels across augmentations induces an orthogonally equivariant structure in the contrastive embedding space. Other analyses connect contrastive learning to inversion of latent generative structure under suitable assumptions (Zimmermann et al., 2021).

Separately, *kernel alignment* formalizes agreement between similarity functions and targets (Cristianini et al., 2002; Cortes et al., 2012), and underlies modern representational comparisons such as CKA (Kornblith et al., 2019). PRH also emphasizes kernel-level convergence, and recent theoretical work proves “perfect” PRH in simplified (deep linear) settings where representations align up to an orthogonal transformation (Ziyin & Chuang, 2025). Our theoretical results operate at this interface: we provide minimal conditions under which agreement of cross-modal similarity structure on a small anchor set forces the existence of a shared orthogonal transform coupling modalities across models, and we establish stability guarantees translating approximate kernel agreement into reliable retrieval transfer.

B ADDITIONAL DISCUSSION AND IMPLICATIONS

Our results have several practical and scientific implications. First, they point to a simple mechanism for backward-compatible upgrades in large embedding systems. Upgrading embedding models typically forces full re-embedding and ANN index rebuilds, which are prohibitively costly at modern scale (hundreds of millions to billions of vectors) (Jayaram Subramanya et al., 2019; Johnson et al., 2019) and often require multi-hour rebuilds and six-figure re-embedding budgets (OpenAI, 2024). This has motivated prior work on backward-compatible and interoperable embeddings, where new models are explicitly trained to preserve compatibility with deployed representations (Jang & Lim, 2025; Meng et al., 2021; Hu et al., 2022; Shen et al., 2020). Our findings show that a small anchor set can often restore cross-model comparability for multimodal contrastive systems. Because orthogonal maps preserve inner products, this can enable upgrades without re-encoding stored corpora and often without rebuilding indexes.

Second, a shared coordinate system enables mix-and-match multimodal pipelines. Different models often excel in different components, such as stronger vision encoders or stronger and more multilingual text encoders. When representations are aligned by a single orthogonal map, image and text towers from different models can be combined into a common space, enabling retrieval across heterogeneous encoders. This perspective complements recent work on representation compatibility and model stitching, which studies when independently trained networks can be connected via lightweight alignment layers (Bansal et al., 2021).

Finally, the ability to align text representations without accessing text has implications for data governance and security. In many deployments, raw text may be unavailable due to privacy, licensing, or retention constraints, even though embeddings are stored. Our results show that, in multimodal systems, text-space alignment can be recovered without using text, given a small anchor set from another modality. At the same time, easy cross-model transferability raises security considerations: if embeddings across models and modalities are easily transformable, then stored embeddings may encode more transferable semantic information than anticipated, reinforcing the need to treat embeddings as sensitive artifacts rather than model-specific byproducts.

C PROOF OF THEORETICAL RESULTS

For a schematic overview of our theoretical analysis, refer to Figure 8.

C.1 SETUP AND NOTATIONS

Let (X, Y) be random variables on measurable spaces \mathcal{X}, \mathcal{Y} with joint density p_{XY} and marginals p_X, p_Y .

Assumption C.1 (Positivity on the domain). On the domain of interest, assume $p_{XY}(x, y) > 0$ and $p_X(x)p_Y(y) > 0$.

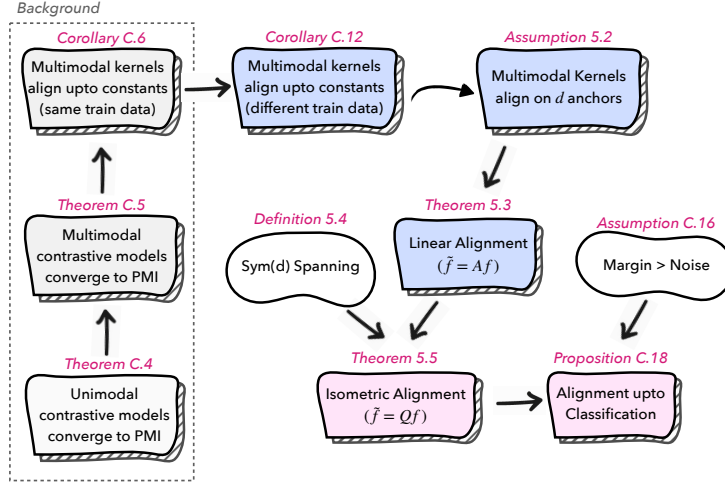


Figure 8: High-level overview of our theoretical results. Background results establish multimodal kernel alignment; additional assumptions progressively strengthen this to linear, then isometric, and finally task-level (classification) alignment.

Remark C.2. The above assumptions can be expressed without densities by replacing ratios of densities with Radon-Nikodym derivatives. Specifically, on the domain of interest, assume $P_{XY} \ll P_X \otimes P_Y$ (equivalently, $P_{Y|X=x} \ll P_Y$ for P_X -a.e. x), and let $r(x, y) := \frac{dP_{XY}}{d(P_X \otimes P_Y)}(x, y)$ denote the density ratio. Optionally, assume $r(x, y) > 0$ a.e. if $\log r$ is required to be finite. For clarity, we'll use densities.

Definition C.3 (Density ratio and Pointwise Mutual Information). Define the mutual density ratio $r : \mathcal{X} \times \mathcal{Y} \rightarrow (0, \infty)$ by

$$r(x, y) := \frac{p_{XY}(x, y)}{p_X(x)p_Y(y)}, \quad K^*(x, y) := \log r(x, y).$$

A contrastive model is a pair of measurable maps $f : \mathcal{X} \rightarrow \mathbb{S}^{d-1}, g : \mathcal{Y} \rightarrow \mathbb{S}^{d-1}$, where $\mathbb{S}^{d-1} := \{u \in \mathbb{R}^d : \|u\|_2 = 1\}$, and a temperature $\tau > 0$, defining a score

$$s(x, y) := \frac{1}{\tau} \langle f(x), g(y) \rangle + b, \quad b \in \mathbb{R}.$$

In practice, contrastive learning constructs a finite candidate set by pairing each query with one positive and several independent negatives; the InfoNCE loss is exactly the cross-entropy for identifying the positive within this list. Below, we formalize the standard finite-sample sampling procedure underlying InfoNCE (Oord et al., 2018) principle.

InfoNCE Sampling Model. Fix $N \geq 2$. Sample $X \sim P_X$, draw one *positive* $Y^+ \sim P_{Y|X}(\cdot | X)$, and draw $N - 1$ *negatives* $Y^{(1)}, \dots, Y^{(N-1)} \stackrel{\text{iid}}{\sim} P_Y$ independently of (X, Y^+) . Sample $J \sim \text{Unif}\{0, \dots, N - 1\}$ and place Y^+ in slot J to form the candidate list (Y_0, \dots, Y_{N-1}) . The learner observes $(X, Y_{0:N-1})$ but not J .

Given a measurable score $s : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$, define

$$q_s(j | x, y_{0:N-1}) := \frac{\exp(s(x, y_j))}{\sum_{k=0}^{N-1} \exp(s(x, y_k))}, \quad \mathcal{L}_N^{(X)}(s) := \mathbb{E}[-\log q_s(J | X, Y_{0:N-1})].$$

Define $\mathcal{L}_N^{(Y)}$ analogously by swapping the roles of X and Y , and set

$$\mathcal{L}_N(s) := \mathcal{L}_N^{(X)}(s) + \mathcal{L}_N^{(Y)}(s).$$

This gives us the objective of a multimodal contrastive learner. We now characterize the minimizers of \mathcal{L}_N .

C.2 BAYES POSTERIOR AND OPTIMAL CRITIC FOR UNIMODAL CONTRASTIVE LEARNERS

Theorem C.4 (Bayes posterior and characterization of minimizers). *Under Assumption C.1,*

1. For almost every $(x, y_{0:N-1})$ under the InfoNCE sampling model,

$$q^*(j \mid x, y_{0:N-1}) := \Pr(J = j \mid X = x, Y_{0:N-1} = y_{0:N-1}) = \frac{r(x, y_j)}{\sum_{k=0}^{N-1} r(x, y_k)}.$$

2. Any measurable s^* achieving the infimum of $\mathcal{L}_N^{(X)}$ must satisfy

$$s^*(x, y) = \log r(x, y) + c(x)$$

for some measurable $c : \mathcal{X} \rightarrow \mathbb{R}$, holding for almost every (x, y) with respect to $P_X \otimes P_Y$.

Proof. (i) Conditioning on $(J = j, X = x)$ gives

$$p(y_{0:N-1} \mid J = j, X = x) = p_{Y|X}(y_j \mid x) \prod_{i \neq j} p_Y(y_i).$$

Since J is uniform, Bayes' rule yields

$$\Pr(J = j \mid x, y_{0:N-1}) = \frac{p_{Y|X}(y_j \mid x) \prod_{i \neq j} p_Y(y_i)}{\sum_{k=0}^{N-1} p_{Y|X}(y_k \mid x) \prod_{i \neq k} p_Y(y_i)}.$$

Dividing numerator and denominator by $\prod_{i=0}^{N-1} p_Y(y_i)$ gives (i), since $\frac{p_{Y|X}(y_j \mid x)}{p_Y(y_j)} = r(x, y_j)$.

(ii) For each realization $(x, y_{0:N-1})$, the conditional risk is

$$\mathbb{E}[-\log q_s(J \mid x, y_{0:N-1}) \mid x, y_{0:N-1}] = H(q^*(\cdot \mid x, y_{0:N-1}), q_s(\cdot \mid x, y_{0:N-1})),$$

so $\mathcal{L}_N^{(X)}(s)$ is minimized iff $q_s(\cdot \mid x, y_{0:N-1}) = q^*(\cdot \mid x, y_{0:N-1})$ a.e. On this full-measure set, for any $j \neq k$,

$$\exp(s^*(x, y_j) - s^*(x, y_k)) = \frac{q_{s^*}(j \mid x, y_{0:N-1})}{q_{s^*}(k \mid x, y_{0:N-1})} = \frac{q^*(j \mid x, y_{0:N-1})}{q^*(k \mid x, y_{0:N-1})} = \frac{r(x, y_j)}{r(x, y_k)}.$$

Integrating out Y_2, \dots, Y_{N-1} (Fubini–Tonelli), gives $\exp(s^*(x, y) - s^*(x, y')) = r(x, y)/r(x, y')$ for $(P_X \otimes P_Y \otimes P_Y)$ -a.e. (x, y, y') . Thus $s^*(x, y) - \log r(x, y)$ is (for P_X -a.e. x) independent of y on a P_Y -full-measure set, i.e. $s^*(x, y) = \log r(x, y) + c(x)$ for some measurable c , holding $(P_X \otimes P_Y)$ -a.e. \square

C.3 OPTIMAL CRITIC FOR MULTIMODAL CONTRASTIVE LEARNERS

Theorem C.5. *Under the Assumption C.1, any measurable s^* achieving the infimum of \mathcal{L}_N must satisfy*

$$s^*(x, y) = \log r(x, y) + C$$

for some constant $C \in \mathbb{R}$, holding for almost every (x, y) with respect to $P_X \otimes P_Y$.

Proof. Let $\ell_X := \inf_s \mathcal{L}_N^{(X)}(s)$ and $\ell_Y := \inf_s \mathcal{L}_N^{(Y)}(s)$. By Theorem C.4(ii), the score $s_0(x, y) := \log r(x, y)$ simultaneously achieves ℓ_X and ℓ_Y (up to an additive constant), hence $\inf_s \mathcal{L}_N(s) = \ell_X + \ell_Y$.

If s^* achieves the infimum of \mathcal{L}_N , then

$$\mathcal{L}_N^{(X)}(s^*) + \mathcal{L}_N^{(Y)}(s^*) = \ell_X + \ell_Y.$$

Since each term is bounded below by its own infimum, we must have $\mathcal{L}_N^{(X)}(s^*) = \ell_X$ and $\mathcal{L}_N^{(Y)}(s^*) = \ell_Y$. Applying Theorem C.4(ii) in each direction yields

$$s^*(x, y) = \log r(x, y) + c(x) = \log r(x, y) + d(y)$$

$(P_X \otimes P_Y)$ -a.e. for measurable c and d . Thus $c(x) = d(y)$ on a full $(P_X \otimes P_Y)$ -measure set. Let $E \subseteq \mathcal{X} \times \mathcal{Y}$ be a full $(P_X \otimes P_Y)$ -measure set where $c(x) = d(y)$. By Fubini-Tonelli's theorem, there exists y_0 such that the section $E_{y_0} := \{x : (x, y_0) \in E\}$ has $P_X(E_{y_0}) = 1$. Then for $x \in E_{y_0}$, $c(x) = d(y_0)$, so c is P_X -a.e. constant. Call this constant C . Plugging back gives $s^*(x, y) = \log r(x, y) + C$ for $(P_X \otimes P_Y)$ -a.e. (x, y) . \square

Corollary C.6. *Let (f, g) and (\tilde{f}, \tilde{g}) be two contrastive models trained on the same distribution that achieve the global infimum of the symmetric objective \mathcal{L}_N . Then there exists a constant $\Delta \in \mathbb{R}$ such that*

$$\langle f(x), g(y) \rangle = \langle \tilde{f}(x), \tilde{g}(y) \rangle + \Delta$$

for almost every (x, y) with respect to $P_X \otimes P_Y$. If the temperatures differ, the relation becomes affine in $\langle \tilde{f}(x), \tilde{g}(y) \rangle$.

Remark C.7. On a finite domain $\{x_1, \dots, x_n\} \times \{y_1, \dots, y_m\}$, the CLIP score matrix $S \in \mathbb{R}^{n \times m}$ with entries $S_{ij} = \tau^{-1} \langle f(x_i), g(y_j) \rangle + b$ can be written as $S = \tau^{-1} F^\top G + b \mathbf{1}_n \mathbf{1}_m^\top$, where $F = [f(x_1) \cdots f(x_n)]$ and $G = [g(y_1) \cdots g(y_m)]$. Hence $\text{rank}(S) \leq d + 1$. Therefore, exact equality $S_{ij} = K^*(x_i, y_j) + C$ would require a strong low-rank structure of the PMI matrix on that domain. This is *not* assumed in [Theorem C.4](#) and [Theorem C.5](#) as it becomes relevant only when one tries to realize the Bayes-optimal critic within a dot-product parameterization.

C.4 BAYES-OPTIMAL SCORES ARE IDENTIFIABLE UP TO A CONSTANT

So far, [Theorem C.5](#) and [Corollary C.6](#) show that for a fixed training distribution P_{XY} , any two global minimizers of the symmetric InfoNCE objective induce the same kernel (up to a global additive constant). We now give a simple, formal setting in which this conclusion can *also* hold for two *different* training distributions.

Dataset Curation. Let P_{XY}^* be a ground truth distribution on $\mathcal{X} \times \mathcal{Y}$ with density p_{XY}^* and marginals p_X^*, p_Y^* , satisfying positivity on the domain of interest (cf. [Assumption C.1](#)). For each dataset $a \in \{1, 2\}$, let $P_{XY}^{(a)}$ be a training distribution with density $p_{XY}^{(a)}$ and marginals $p_X^{(a)}, p_Y^{(a)}$. Let r^*, K^* and $r^{(a)}, K^{(a)}$ denote the density ratios and PMIs defined earlier, applied to p_{XY}^* and $p_{XY}^{(a)}$, respectively.

Assumption C.8. For each dataset $a \in \{1, 2\}$, there exist measurable weights $u_a : \mathcal{X} \rightarrow (0, \infty)$ and $v_a : \mathcal{Y} \rightarrow (0, \infty)$ such that

$$Z_a := \mathbb{E}^*[u_a(X)v_a(Y)] < \infty, \quad p_{XY}^{(a)}(x, y) = \frac{u_a(x)v_a(y)}{Z_a} p_{XY}^*(x, y).$$

In [Assumption C.8](#), we view large-scale training corpora as distinct *curations* of a common underlying distribution, where each modality (image/text) is filtered by its own criteria (e.g., quality, safety, language) independently of the other.

Lemma C.9. *Under [Assumption C.8](#), for each $a \in \{1, 2\}$ and for $P_X^* \otimes P_Y^*$ -a.e. (x, y) ,*

$$r^{(a)}(x, y) = r^*(x, y) \cdot \frac{Z_a}{\mathbb{E}^*[v_a(Y) \mid X = x] \mathbb{E}^*[u_a(X) \mid Y = y]}.$$

Equivalently,

$$K^{(a)}(x, y) = K^*(x, y) + \log Z_a - \log \mathbb{E}^*[v_a(Y) \mid X = x] - \log \mathbb{E}^*[u_a(X) \mid Y = y].$$

Proof. By [Assumption C.8](#), $p_{XY}^{(a)}(x, y) = \frac{u_a(x)v_a(y)}{Z_a} p_{XY}^*(x, y)$. Thus

$$p_X^{(a)}(x) = \int p_{XY}^{(a)}(x, y) dy = \frac{u_a(x)p_X^*(x)}{Z_a} \mathbb{E}^*[v_a(Y) \mid X = x],$$

and similarly

$$p_Y^{(a)}(y) = \frac{v_a(y)p_Y^*(y)}{Z_a} \mathbb{E}^*[u_a(X) \mid Y = y].$$

Plugging into $r^{(a)}(x, y) = \frac{p_{XY}^{(a)}(x, y)}{p_X^{(a)}(x)p_Y^{(a)}(y)}$ gives the stated identity, and taking logs yields the PMI form. \square

Lemma C.9 shows that curation perturbs the density ratio by a multiplicative factor governed by the conditional expectations $\mathbb{E}^*[v(Y) | X = x]$ and $\mathbb{E}^*[u(X) | Y = y]$. Below, we impose a mild condition on these terms, requiring that dataset curation acts independently across modalities i.e the expected acceptance rate of texts does not depend on the image they are paired with, and vice versa.

Assumption C.10. For each dataset $a \in \{1, 2\}$,

$$\mathbb{E}^*[v_a(Y) | X = x] = \mathbb{E}^*[v_a(Y)] \quad \text{for } P_X^*\text{-a.e. } x, \quad \mathbb{E}^*[u_a(X) | Y = y] = \mathbb{E}^*[u_a(X)] \quad \text{for } P_Y^*\text{-a.e. } y.$$

Theorem C.11. Under [Assumption C.8](#) and [Assumption C.10](#), for each $a \in \{1, 2\}$ there exists a constant $C_a \in \mathbb{R}$ such that $K^{(a)}(x, y) = K^*(x, y) + C_a$ for $P_X^* \otimes P_Y^*$ -a.e. (x, y) . Consequently, there exists a constant $\Delta \in \mathbb{R}$ such that

$$K^{(1)}(x, y) = K^{(2)}(x, y) + \Delta$$

for $P_X^* \otimes P_Y^*$ -a.e. (x, y) .

Proof. Under [Assumption C.10](#), the conditional expectations in [Lemma C.9](#) are constants: $\mathbb{E}^*[v_a(Y) | X = x] = \mathbb{E}^*[v_a(Y)]$ and $\mathbb{E}^*[u_a(X) | Y = y] = \mathbb{E}^*[u_a(X)]$ a.e. Hence $r^{(a)}(x, y) = \alpha_a r^*(x, y)$ a.e., where

$$\alpha_a := \frac{Z_a}{\mathbb{E}^*[u_a(X)] \mathbb{E}^*[v_a(Y)]},$$

so $K^{(a)}(x, y) = K^*(x, y) + \log \alpha_a$ a.e. The second claim follows by subtraction with $\Delta = \log \alpha_1 - \log \alpha_2$. \square

Corollary C.12. Let s_1^* and s_2^* be measurable global minimizers of the symmetric objective \mathcal{L}_N when the InfoNCE sampling model is defined under $P_{XY}^{(1)}$ and $P_{XY}^{(2)}$, respectively (cf. [Appendix C.1](#)). Under [Theorem C.11](#), there exists $\Delta' \in \mathbb{R}$ such that

$$s_1^*(x, y) = s_2^*(x, y) + \Delta'$$

for $P_X^* \otimes P_Y^*$ -a.e. (x, y) .

Proof. By [Theorem C.5](#) applied to each dataset, $s_a^*(x, y) = K^{(a)}(x, y) + C'_a$ for constants C'_a (a.e.). By [Theorem C.11](#), $K^{(1)}(x, y) = K^{(2)}(x, y) + \Delta$ (a.e.), hence $s_1^* - s_2^*$ is constant. \square

Remark C.13 (Implication for contrastive dot-product kernels). If each critic is realized through (f, g) and (\tilde{f}, \tilde{g}) , then [Corollary C.12](#) implies

$$\langle f(x), g(y) \rangle = \langle \tilde{f}(x), \tilde{g}(y) \rangle + \tilde{\Delta}$$

for $P_X^* \otimes P_Y^*$ -a.e. (x, y) , for some constant $\tilde{\Delta} \in \mathbb{R}$. If the temperatures differ, the relation becomes affine in $\langle \tilde{f}(x), \tilde{g}(y) \rangle$.

C.5 GUARANTEES UNDER EXACT ALIGNMENT

C.5.1 LINEAR ALIGNMENT OF INDEPENDENT CONTRASTIVE MODELS

In this section, we prove our first main result showing that matching multimodal kernels on a small set of ‘‘anchor’’ points is sufficient to lock the two embedding spaces together up to a linear transformation.

Let $\Omega_X \subseteq \mathcal{X}$ and $\Omega_Y \subseteq \mathcal{Y}$ be subsets (the domain of interest, e.g. a downstream dataset of images and prompts). Let the contrastive model pairs (f, g) and (\tilde{f}, \tilde{g}) map inputs to $\mathbb{S}^{d-1} \subset \mathbb{R}^d$ and $\mathbb{S}^{\tilde{d}-1} \subset \mathbb{R}^{\tilde{d}}$ respectively, where $d \leq \tilde{d}$, without loss of generality. We fix a set of *image anchors* $\{\bar{x}_j\}_{j=1}^d \subset \Omega_X$ and *text anchors* $\{\bar{y}_i\}_{i=1}^{\tilde{d}} \subset \Omega_Y$ and collect their embeddings into the following matrices:

$$\begin{aligned} G &:= [g(\bar{y}_1) \cdots g(\bar{y}_{\tilde{d}})] \in \mathbb{R}^{d \times \tilde{d}}, \\ \tilde{G} &:= [\tilde{g}(\bar{y}_1) \cdots \tilde{g}(\bar{y}_{\tilde{d}})] \in \mathbb{R}^{\tilde{d} \times \tilde{d}}, \\ F &:= [f(\bar{x}_1) \cdots f(\bar{x}_d)] \in \mathbb{R}^{d \times d}. \end{aligned}$$

Assumption 5.2. The multimodal kernels coincide on the set of anchors:

$$\begin{aligned}\langle f(x), g(\bar{y}_i) \rangle &= \langle \tilde{f}(x), \tilde{g}(\bar{y}_i) \rangle \quad \forall x \in \Omega_X, \forall i \in \{1, \dots, \tilde{d}\}, \\ \langle f(\bar{x}_j), g(y) \rangle &= \langle \tilde{f}(\bar{x}_j), \tilde{g}(y) \rangle \quad \forall y \in \Omega_Y, \forall j \in \{1, \dots, d\}.\end{aligned}$$

Theorem 5.3. (Linear Identifiability, proof in [Appendix C.5.1](#)). Under [Assumption 5.2](#), suppose \tilde{G} and F are invertible. Then there exists a linear map A such that $\tilde{f}(x) = Af(x) \forall x \in \Omega_X$. Further, if A has full column rank, then for every $y \in \Omega_Y$, $\text{Proj}_{\text{Im}(A)} \tilde{g}(y) = A(A^\top A)^{-1}g(y)$. If $\tilde{d} = d$, then $\tilde{g}(y) = A^{-\top}g(y)$.

Proof. Fix $x \in \Omega_X$ and define

$$k(x) := (\langle f(x), g(\bar{y}_i) \rangle)_{i=1}^{\tilde{d}} = G^\top f(x), \quad \tilde{k}(x) := (\langle \tilde{f}(x), \tilde{g}(\bar{y}_i) \rangle)_{i=1}^{\tilde{d}} = \tilde{G}^\top \tilde{f}(x).$$

By [Assumption 5.2](#), $k(x) = \tilde{k}(x)$, hence $G^\top f(x) = \tilde{G}^\top \tilde{f}(x)$. Since \tilde{G} is invertible,

$$\tilde{f}(x) = \tilde{G}^{-\top} G^\top f(x) = Af(x),$$

for all $x \in \Omega_X$. Now fix $y \in \Omega_Y$. For each anchor \bar{x}_j ,

$$\langle f(\bar{x}_j), g(y) \rangle = \langle \tilde{f}(\bar{x}_j), \tilde{g}(y) \rangle = \langle Af(\bar{x}_j), \tilde{g}(y) \rangle = \langle f(\bar{x}_j), A^\top \tilde{g}(y) \rangle.$$

Thus $\langle f(\bar{x}_j), g(y) - A^\top \tilde{g}(y) \rangle = 0$ for all $j \in \{1, \dots, d\}$. Since $F = [f(\bar{x}_1) \cdots f(\bar{x}_d)]$ is invertible, $\{f(\bar{x}_j)\}_{j=1}^d$ spans \mathbb{R}^d , hence $g(y) = A^\top \tilde{g}(y)$ for all $y \in \Omega_Y$.

Finally, if A has full column rank then $P := A(A^\top A)^{-1}A^\top$ is the orthogonal projector onto $\text{Im}(A)$, and

$$\text{Proj}_{\text{Im}(A)} \tilde{g}(y) = P\tilde{g}(y) = A(A^\top A)^{-1}A^\top \tilde{g}(y) = A(A^\top A)^{-1}g(y).$$

If $\tilde{d} = d$ and A is invertible, then $A(A^\top A)^{-1} = A^{-\top}$, giving $\tilde{g}(y) = A^{-\top}g(y)$. \square

C.5.2 ISOMETRIC ALIGNMENT OF INDEPENDENT CONTRASTIVE MODELS

[Theorem 5.3](#) establishes that the representation is fixed up to a linear transformation A . However, standard contrastive encoders normalize embeddings to the unit hypersphere \mathbb{S}^{d-1} , forcing $\|\tilde{f}(x)\|_2 = \|Af(x)\|_2 = 1$ everywhere. This forces A to be an isometry ($A^\top A = I_d$) only if the data is sufficiently diverse to probe the matrix in all directions. We formalize this diversity via the following condition.

Definition 5.4. (Sym(d)-spanning) A set of vectors $S \subset \mathbb{R}^d$ is *Sym(d)-spanning* if the rank-one matrices $\{xx^\top : x \in S\}$ span the space of symmetric matrices $\text{Sym}(d)$. Equivalently: if $M \in \text{Sym}(d)$ and $x^\top Mx = 0$ for all $x \in S$, then $M = 0$. This equivalence follows from the identity $x^\top Mx = \langle xx^\top, M \rangle$.

Lemma C.14. Let $A \in \mathbb{R}^{\tilde{d} \times d}$ and let $U \subseteq \mathbb{S}^{d-1} \subset \mathbb{R}^d$. Assume $\|Au\|_2 = 1$ for all $u \in U$. If U contains a *Sym(d)-spanning subset*, then $A^\top A = I_d$.

Proof. Let $M := A^\top A - I_d \in \text{Sym}(d)$. For any $u \in \mathbb{S}^{d-1}$,

$$\|Au\|_2^2 = u^\top A^\top Au = u^\top (I_d + M)u = 1 + u^\top Mu.$$

Thus $\|Au\|_2 = 1$ implies $u^\top Mu = 0$ for all $u \in U$. Choose a *Sym(d)-spanning subset* $S = \{u_i\} \subseteq U$. Then $u_i^\top Mu_i = 0$ for all i implies $M = 0$ by definition of *Sym(d)-spanning*, hence $A^\top A = I_d$. \square

Theorem 5.5. (Orthogonal Identifiability, proof in [Appendix C.5.2](#)). Assume the conditions of [Theorem 5.3](#) hold. If the set of image embeddings $\{f(x) : x \in \Omega_X\}$ contains a *Sym(d)-spanning subset*, then the linear map A has orthonormal columns ($A^\top A = I_d$). Consequently, $\tilde{f}(x) = Qf(x) \quad \forall x \in \Omega_X$, where $Q := A$ satisfies $Q^\top Q = I_d$. Furthermore, for the other modality:

$$\text{Proj}_{\text{Im}(Q)} \tilde{g}(y) = Qg(y) \quad \forall y \in \Omega_Y.$$

If $\tilde{d} = d$, Q is orthogonal i.e. $Q \in O(d)$ and $\tilde{g}(y) = Qg(y)$.

Proof. By [Theorem 5.3](#), $\tilde{f}(x) = Af(x)$ for all $x \in \Omega_X$. Since $\|f(x)\|_2 = \|\tilde{f}(x)\|_2 = 1$, we have $\|Af(x)\|_2 = 1$ for all $x \in \Omega_X$. Apply [Lemma C.14](#) to conclude $A^\top A = I_d$. Set $Q := A$.

By [Theorem 5.3](#), $g(y) = Q^\top \tilde{g}(y)$ for all $y \in \Omega_Y$. Multiplying by Q gives

$$Qg(y) = QQ^\top \tilde{g}(y) = \text{Proj}_{\text{Im}(Q)} \tilde{g}(y),$$

since QQ^\top is the orthogonal projector onto $\text{Im}(Q)$ when $Q^\top Q = I_d$. If $\tilde{d} = d$, then $\text{Im}(Q) = \mathbb{R}^d$ and $\text{Proj}_{\text{Im}(Q)}$ is the identity. \square

Remark C.15. Suppose the exact regime of [Theorem 5.5](#) holds: $\tilde{f}(x) = Qf(x)$ for all $x \in \Omega_X$ (and similarly for g), and suppose the means are taken w.r.t. the same distribution on Ω_X : $\mu_f := \mathbb{E}[f(X)]$, $\mu_{\tilde{f}} := \mathbb{E}[\tilde{f}(X)]$. Then $\mu_{\tilde{f}} = Q\mu_f$, hence

$$\tilde{f}(x) - \mu_{\tilde{f}} = Q(f(x) - \mu_f).$$

Therefore the rigid map $z \mapsto Q(z - \mu_f) + \mu_{\tilde{f}}$ reduces exactly to the pure rotation $z \mapsto Qz$. In this sense, mean-centering is a finite-sample / distribution-mismatch correction that leaves the exact-identifiability statement unchanged.

C.5.3 ISOMETRIC ALIGNMENT OF INDEPENDENT CONTRASTIVE MODELS WHEN ONE MODALITY LIES IN LOW DIMENSION

The above analysis assumes that $\text{span}\{f(x) : x \in \Omega_X\} = \mathbb{R}^d$ since we assumed the existence of d anchor vectors from Ω_X that make F invertible. In other words, we assumed the existence of d linearly independent vectors $f(\bar{x}_1), \dots, f(\bar{x}_d)$, that span \mathbb{R}^d . We extend this analysis below in [Theorem C.16](#) to cases where these embeddings instead lie in a lower-dimensional subspace.

Let

$$U_X := \text{span}\{f(x) : x \in \Omega_X\} \subseteq \mathbb{R}^d, \quad \tilde{U}_X := \text{span}\{\tilde{f}(x) : x \in \Omega_X\} \subseteq \mathbb{R}^{\tilde{d}},$$

where U_X is a subspace of \mathbb{R}^d with rank $r := \dim(U_X)$. Choose anchors $\bar{x}_1, \dots, \bar{x}_r \in \Omega_X$ such that $\text{span}\{f(\bar{x}_1), \dots, f(\bar{x}_r)\} = U_X$.

Theorem C.16 (Orthogonal Identifiability). *Assume multimodal kernel alignment as in [Assumption 5.2](#) but on r anchors $\bar{x}_1, \dots, \bar{x}_r \in \Omega_X$ instead of d . Assume there exists a matrix $Q \in \mathbb{R}^{\tilde{d} \times d}$ with $Q^\top Q = I_d$ such that $\tilde{f}(x) = Qf(x)$ for all $x \in \Omega_X$. Then $\tilde{U}_X = QU_X$ and for every $y \in \Omega_Y$,*

$$\text{Proj}_{\tilde{U}_X} \tilde{g}(y) = Q \text{Proj}_{U_X} g(y).$$

In particular, if $U_X = \mathbb{R}^d$, then $\text{Proj}_{\text{Im}(Q)} \tilde{g}(y) = Qg(y)$ for all $y \in \Omega_Y$, and if additionally $\tilde{d} = d$ then $\tilde{g}(y) = Qg(y)$.

Proof. Since $\tilde{f}(x) = Qf(x)$ and $Q^\top Q = I_d$, for each anchor \bar{x}_j with $j \in \{1, \dots, r\}$ and $y \in \Omega_Y$,

$$\langle f(\bar{x}_j), g(y) \rangle = \langle \tilde{f}(\bar{x}_j), \tilde{g}(y) \rangle = \langle Qf(\bar{x}_j), \tilde{g}(y) \rangle = \langle f(\bar{x}_j), Q^\top \tilde{g}(y) \rangle.$$

Hence $\langle f(\bar{x}_j), g(y) - Q^\top \tilde{g}(y) \rangle = 0$ for all anchors \bar{x}_j . By linearity, $\langle u, g(y) - Q^\top \tilde{g}(y) \rangle = 0$ for all $u \in U_X$, so $g(y) - Q^\top \tilde{g}(y) \in U_X^\perp$. Equivalently, $\text{Proj}_{U_X} g(y) = \text{Proj}_{U_X} (Q^\top \tilde{g}(y))$. Applying Q and using $\tilde{U}_X = QU_X$ and $\text{Proj}_{\tilde{U}_X} = Q \text{Proj}_{U_X} Q^\top$ yields

$$\text{Proj}_{\tilde{U}_X} \tilde{g}(y) = Q \text{Proj}_{U_X} Q^\top \tilde{g}(y) = Q \text{Proj}_{U_X} g(y).$$

If $U_X = \mathbb{R}^d$, then $\tilde{U}_X = Q\mathbb{R}^d = \text{Im}(Q)$, giving $\text{Proj}_{\text{Im}(Q)} \tilde{g}(y) = Qg(y)$. If moreover $\tilde{d} = d$, then $\text{Im}(Q) = \mathbb{R}^d$ and the projection is the identity. \square

C.5.4 ALIGNMENT UPTO CLASSIFICATION BOUNDARIES

While the preceding theory establishes conditions for exact geometric alignment. In practice however, one often cares about alignment up to concepts or classification that are relevant for the downstream task. We now analyze this regime, where we seek to distinguish a finite family of prompts $\mathcal{Y}_{\text{cls}} =$

$\{y_c\}_{c=1}^K \subseteq \Omega_Y$ even when pointwise alignment is imprecise. Specifically, we care about the within-modality cross-model top-1 retrieval accuracy i.e.

$$\hat{c}(y_c) \in \arg \max_{c' \in \{1, \dots, K\}} \langle Qg(y_c), \tilde{g}(y_{c'}) \rangle.$$

Theorem C.16 is stated for all $y \in \Omega_Y$ and its proof is pointwise in y and uses the multimodal kernel equalities only for the pairs (\bar{x}_j, y) with $j \in \{1, \dots, r\}$. Therefore, if one only needs the conclusion for a subset $\mathcal{Y}_0 \subseteq \Omega_Y$ (e.g. $\mathcal{Y}_0 = \mathcal{Y}_{\text{cls}}$), it suffices to assume the anchor kernel equalities only on $\{\bar{x}_1, \dots, \bar{x}_r\} \times \mathcal{Y}_0$. Formally,

Assumption C.17 (Anchor kernel equalities only for class prompts). For the image anchors $\bar{x}_1, \dots, \bar{x}_r$ spanning U_X , assume that for all class prompts $y \in \mathcal{Y}_{\text{cls}}$,

$$\langle f(\bar{x}_j), g(y) \rangle = \langle \tilde{f}(\bar{x}_j), \tilde{g}(y) \rangle \quad \forall j \in \{1, \dots, r\}.$$

Now, for any class prompt $y \in \mathcal{Y}_{\text{cls}}$, decompose the embedding into an identifiable signal $u(y)$ and an unidentifiable residual $w(y)$:

$$g(y) = u(y) + w(y), \quad u(y) := \text{Proj}_{U_X} g(y), \quad w(y) \in U_X^\perp.$$

Then, from **Theorem C.16**, define analogously

$$\tilde{g}(y) = Qu(y) + \tilde{w}(y), \quad \tilde{w}(y) \in (QU_X)^\perp.$$

Definition C.18. Let $\mathcal{Y}_{\text{cls}} = \{y_c\}_{c=1}^K$. Define

$$\gamma := \min_c \left(\|u(y_c)\|_2^2 - \max_{c' \neq c} \langle u(y_c), u(y_{c'}) \rangle \right),$$

that measures class separability *within the image span* U_X . Define,

$$\eta := \max_{c, c'} |\langle Qw(y_c), \tilde{w}(y_{c'}) \rangle|.$$

that measures the worst-case cross-model interaction of the unidentifiable residuals across models.

Proposition C.19. Under **Assumption C.17**, if $\gamma > 2\eta$, then nearest-neighbor class retrieval is correct for every class prompt:

$$\hat{c}(y_c) = c \quad \forall c \in \{1, \dots, K\}.$$

Proof. The conclusion of **Theorem C.16** under **Assumption C.17** implies

$$\tilde{U}_X = QU_X \quad \text{and} \quad \text{Proj}_{\tilde{U}_X} \tilde{g}(y) = Q \text{Proj}_{U_X} g(y) \quad \forall y \in \mathcal{Y}_{\text{cls}}$$

Fix a query class c and a candidate c' . Using $g(y) = u(y) + w(y)$ and $\tilde{g}(y) = Qu(y) + \tilde{w}(y)$,

$$\langle Qg(y_c), \tilde{g}(y_{c'}) \rangle = \langle Q(u(y_c) + w(y_c)), Qu(y_{c'}) + \tilde{w}(y_{c'}) \rangle.$$

$Qu(y_c)$ is orthogonal to $\tilde{w}(y_{c'})$ and similarly $\langle Qw(y_c), Qu(y_{c'}) \rangle = 0$ since $w(y_c) \in U_X^\perp$. Also, $Q^\top Q = I_d$ implies $\langle Qu, Qu' \rangle = \langle u, u' \rangle$. Thus,

$$\langle Qg(y_c), \tilde{g}(y_{c'}) \rangle = \langle u(y_c), u(y_{c'}) \rangle + \langle Qw(y_c), \tilde{w}(y_{c'}) \rangle.$$

By the definition of γ , for any $c' \neq c$ we have $\langle u(y_c), u(y_{c'}) \rangle \leq \|u(y_c)\|_2^2 - \gamma$. Therefore, for all $c' \neq c$,

$$\langle Qg(y_c), \tilde{g}(y_{c'}) \rangle - \langle Qg(y_c), \tilde{g}(y_c) \rangle \geq (\|u(y_c)\|_2^2 - \eta) - (\|u(y_c)\|_2^2 - \gamma + \eta) = \gamma - 2\eta.$$

If $\gamma > 2\eta$, the right-hand side is strictly positive, so c is the unique maximizer and $\hat{c}(y_c) = c$. \square

Proposition C.19 shows that class retrieval depends on a *margin* condition inside U_X , and is stable as long as the residual terms in U_X^\perp do not overwhelm that margin.

C.6 GUARANTEES UNDER APPROXIMATE ALIGNMENT

We now relax the anchor equalities in **Appendix C.5** and allow small discrepancies on the anchor pairs. Like before, we assume the contrastive model pairs (f, g) and (\tilde{f}, \tilde{g}) to map inputs to $\mathbb{S}^{d-1} \subset \mathbb{R}^d$ and $\mathbb{S}^{\tilde{d}-1} \subset \mathbb{R}^{\tilde{d}}$ respectively, where $d \leq \tilde{d}$, without loss of generality.

C.6.1 APPROXIMATE LINEAR ALIGNMENT OF INDEPENDENT CONTRASTIVE MODELS

Similar to before, fix \tilde{d} text anchors $\tilde{y}_1, \dots, \tilde{y}_{\tilde{d}} \in \Omega_Y$ and define

$$G := [g(\tilde{y}_1) \cdots g(\tilde{y}_{\tilde{d}})] \in \mathbb{R}^{d \times \tilde{d}}, \quad \tilde{G} := [\tilde{g}(\tilde{y}_1) \cdots \tilde{g}(\tilde{y}_{\tilde{d}})] \in \mathbb{R}^{\tilde{d} \times \tilde{d}}.$$

Assumption C.20 (Approximate multimodal kernel equalities). There exists $\varepsilon \geq 0$ such that for all $x \in \Omega_X$ and all $i \in \{1, \dots, \tilde{d}\}$,

$$|\langle f(x), g(\tilde{y}_i) \rangle - \langle \tilde{f}(x), \tilde{g}(\tilde{y}_i) \rangle| \leq \varepsilon.$$

Theorem C.21. Assume \tilde{G} is invertible with smallest singular value $\sigma_{\min}(\tilde{G}) > 0$. Then under Assumption C.20, there exists a linear map A such that for every $x \in \Omega_X$,

$$\|\tilde{f}(x) - Af(x)\|_2 \leq \frac{\sqrt{\tilde{d}}}{\sigma_{\min}(\tilde{G})} \varepsilon.$$

Proof. Define $k(x) := G^\top f(x) \in \mathbb{R}^{\tilde{d}}$ and $\tilde{k}(x) := \tilde{G}^\top \tilde{f}(x) \in \mathbb{R}^{\tilde{d}}$. Set $A := \tilde{G}^{-\top} G^\top \in \mathbb{R}^{\tilde{d} \times d}$, then

$$\tilde{f}(x) - Af(x) = \tilde{f}(x) - \tilde{G}^{-\top} G^\top f(x) = \tilde{G}^{-\top} (\tilde{G}^\top \tilde{f}(x) - G^\top f(x)) = \tilde{G}^{-\top} (\tilde{k}(x) - k(x)).$$

Therefore

$$\|\tilde{f}(x) - Af(x)\|_2 \leq \|\tilde{G}^{-\top}\|_2 \|\tilde{k}(x) - k(x)\|_2 = \frac{1}{\sigma_{\min}(\tilde{G})} \|\tilde{k}(x) - k(x)\|_2.$$

By Assumption C.20, each coordinate of $\tilde{k}(x) - k(x)$ has magnitude at most ε , so $\|\tilde{k}(x) - k(x)\|_2 \leq \sqrt{\tilde{d}} \varepsilon$. \square

C.6.2 APPROXIMATE ISOMETRIC ALIGNMENT OF INDEPENDENT CONTRASTIVE MODELS

Assumption C.22 (Sym(d)-spanning). There exist $\bar{x}_1, \dots, \bar{x}_m \in \Omega_X$ such that $S = \{f(\bar{x}_1), \dots, f(\bar{x}_m)\} \subset \mathbb{S}^{d-1}$ is Sym(d)-spanning

Fix $A \in \mathbb{R}^{\tilde{d} \times d}$ and define the uniform deviation

$$\Delta := \sup_{x \in \Omega_X} \|\tilde{f}(x) - Af(x)\|_2.$$

Also define the constant

$$\kappa_S := \sup_{M \in \text{Sym}(d) \setminus \{0\}} \frac{\|M\|_2}{\max_{u \in S} |u^\top M u|} < \infty.$$

Theorem C.23 (Approximate Orthogonal Identifiability). Assume A has full column rank. Then under Assumption C.22, there exists $Q^\top Q = I_d$, such that for all $x \in \Omega_X$

$$\|\tilde{f}(x) - Qf(x)\|_2 \leq \Delta + \kappa_S (2 + \Delta) \Delta$$

Proof. Let $M := A^\top A - I_d \in \text{Sym}(d)$. For any $u = f(\bar{x}_i) \in S$, $\|\tilde{f}(\bar{x}_i)\|_2 = 1$ and $\|Au - \tilde{f}(\bar{x}_i)\|_2 \leq \Delta$ imply $\|Au\|_2 - 1 \leq \Delta$ and hence $\|Au\|_2^2 - 1 \leq (2 + \Delta)\Delta$. But $\|Au\|_2^2 - 1 = u^\top M u$, so $\max_{u \in S} |u^\top M u| \leq (2 + \Delta)\Delta$. By definition of κ_S , $\|M\|_2 \leq \kappa_S \max_{u \in S} |u^\top M u|$, showing that

$$\|A^\top A - I_d\|_2 \leq \kappa_S (2 + \Delta) \Delta.$$

Now, let $Q \in \mathbb{R}^{\tilde{d} \times d}$ denote the orthogonal factor in the polar decomposition of A i.e. $A = Q(A^\top A)^{1/2}$, $Q^\top Q = I_d$. Then $\|A - Q\|_2 = \|(A^\top A)^{1/2} - I_d\|_2 \leq \|A^\top A - I_d\|_2$. Finally, $\|\tilde{f}(x) - Qf(x)\|_2 \leq \|\tilde{f}(x) - Af(x)\|_2 + \|(A - Q)f(x)\|_2 \leq \Delta + \|A - Q\|_2$, proving the claim. \square

Now, fix d image anchors $\bar{x}_1, \dots, \bar{x}_d \in \Omega_X$ and define

$$F := [f(\bar{x}_1) \cdots f(\bar{x}_d)] \in \mathbb{R}^{d \times d}, \quad \sigma_{\min}(F) > 0.$$

Assumption C.24 (Approximate multimodal kernel equalities). There exists $\varepsilon' \geq 0$ such that for all $y \in \Omega_Y$ and all $j \in \{1, \dots, d\}$,

$$|\langle f(\bar{x}_j), g(y) \rangle - \langle \tilde{f}(\bar{x}_j), \tilde{g}(y) \rangle| \leq \varepsilon'.$$

Theorem C.25. Assume [Assumption C.24](#) and let $Q \in \mathbb{R}^{\tilde{d} \times d}$ satisfy $Q^\top Q = I_d$ and

$$\max_{j \in \{1, \dots, d\}} \left\| \tilde{f}(\bar{x}_j) - Qf(\bar{x}_j) \right\|_2 \leq \delta_f.$$

Then for every $y \in \Omega_Y$,

$$\left\| \text{Proj}_{\text{Im}(Q)} \tilde{g}(y) - Qg(y) \right\|_2 = \left\| Q^\top \tilde{g}(y) - g(y) \right\|_2 \leq \frac{\sqrt{\tilde{d}}}{\sigma_{\min}(F)} (\varepsilon' + \delta_f).$$

In particular, if $\tilde{d} = d$ then $\text{Im}(Q) = \mathbb{R}^d$ and the projection is redundant, giving a bound on $\|\tilde{g}(y) - Qg(y)\|_2$.

Proof. Fix $y \in \Omega_Y$ and set $v := g(y) - Q^\top \tilde{g}(y) \in \mathbb{R}^d$. For each anchor \bar{x}_j ,

$$\langle f(\bar{x}_j), v \rangle = \langle f(\bar{x}_j), g(y) \rangle - \langle Qf(\bar{x}_j), \tilde{g}(y) \rangle.$$

Thus

$$|\langle f(\bar{x}_j), v \rangle| \leq |\langle f(\bar{x}_j), g(y) \rangle - \langle \tilde{f}(\bar{x}_j), \tilde{g}(y) \rangle| + |\langle \tilde{f}(\bar{x}_j) - Qf(\bar{x}_j), \tilde{g}(y) \rangle| \leq \varepsilon' + \delta_f,$$

since $\|\tilde{g}(y)\|_2 = 1$. Hence $\|F^\top v\|_2 \leq \sqrt{\tilde{d}}(\varepsilon' + \delta_f)$ and, since F is invertible,

$$\|v\|_2 \leq \|F^{-\top}\|_2 \|F^\top v\|_2 = \frac{1}{\sigma_{\min}(F)} \|F^\top v\|_2 \leq \frac{\sqrt{\tilde{d}}}{\sigma_{\min}(F)} (\varepsilon' + \delta_f).$$

Finally, using $Q^\top Q = I_d$,

$$\left\| \text{Proj}_{\text{Im}(Q)} \tilde{g}(y) - Qg(y) \right\|_2 = \left\| Q(Q^\top \tilde{g}(y) - g(y)) \right\|_2 = \|v\|_2.$$

□

Remark C.26. [Assumption C.20](#) and [Assumption C.24](#) are ε -relaxations of [Assumption 5.2](#). [Theorem C.21](#) and [Theorem C.25](#) are the perturbation analogues of [Theorem 5.3](#) and [Theorem 5.5](#) respectively.

C.6.3 APPROXIMATE ISOMETRIC ALIGNMENT OF INDEPENDENT CONTRASTIVE MODELS WHEN ONE MODALITY LIES IN LOW DIMENSION

Similar to [Theorem C.16](#), we extend our analysis below to cases where the image embeddings lie in a lower-dimensional subspace. Let $U_X := \text{span}\{f(x) : x \in \Omega_X\} \subseteq \mathbb{R}^d$ where U_X is a subspace of \mathbb{R}^d of $\dim(U_X) = r$. Choose anchors $\bar{x}_1, \dots, \bar{x}_r \in \Omega_X$ such that $\text{span}\{f(\bar{x}_1), \dots, f(\bar{x}_r)\} = U_X$. Accordingly, denote

$$F := [f(\bar{x}_1) \cdots f(\bar{x}_r)] \in \mathbb{R}^{d \times r}, \quad \sigma_{\min}(F) > 0.$$

Assumption C.27. Assume there exists a matrix $Q \in \mathbb{R}^{\tilde{d} \times d}$ with $Q^\top Q = I_d$ and constants $\delta_f, \varepsilon \geq 0$ such that

$$\sup_{x \in \Omega_X} \left\| \tilde{f}(x) - Qf(x) \right\|_2 \leq \delta_f, \quad \sup_{x \in \Omega_X, y \in \Omega_Y} |\langle f(x), g(y) \rangle - \langle \tilde{f}(x), \tilde{g}(y) \rangle| \leq \varepsilon.$$

Theorem C.28. Under [Assumption C.27](#), for every $y \in \Omega_Y$,

$$\left\| \text{Proj}_{U_X} (g(y) - Q^\top \tilde{g}(y)) \right\|_2 \leq \rho, \quad \rho := \frac{\sqrt{r}}{\sigma_{\min}(F)} (\varepsilon + \delta_f).$$

Equivalently,

$$\left\| \text{Proj}_{QU_X} \tilde{g}(y) - Q \text{Proj}_{U_X} g(y) \right\|_2 \leq \rho.$$

Proof. Fix $y \in \Omega_Y$ and define $h := g(y) - Q^\top \tilde{g}(y)$. For any $x \in \Omega_X$,

$$\begin{aligned} |\langle f(x), h \rangle| &= |\langle f(x), g(y) \rangle - \langle Qf(x), \tilde{g}(y) \rangle| \\ &\leq |\langle f(x), g(y) \rangle - \langle \tilde{f}(x), \tilde{g}(y) \rangle| + |\langle \tilde{f}(x) - Qf(x), \tilde{g}(y) \rangle| \leq \varepsilon + \delta_f, \end{aligned}$$

since $\|\tilde{g}(y)\|_2 = 1$. Let $b \in \mathbb{R}^r$ have coordinates $b_j := \langle f(\tilde{x}_j), h \rangle$. Then $\|b\|_2 \leq \sqrt{r}(\varepsilon + \delta_f)$. Moreover, $F^\top \text{Proj}_{U_X} h = F^\top h = b$, and since F has full column rank,

$$\|\text{Proj}_{U_X} h\|_2 \leq \|(F^\top)^\dagger\|_2 \|b\|_2 = \frac{1}{\sigma_{\min}(F)} \|b\|_2 \leq \frac{\sqrt{r}}{\sigma_{\min}(F)} (\varepsilon + \delta_f).$$

The equivalent bound in \tilde{U}_X follows by left-multiplying by Q and using $\text{Proj}_{Q U_X} = Q \text{Proj}_{U_X} Q^\top$ (which holds whenever $Q^\top Q = I_d$).

□

D SUPPLEMENTARY EXPERIMENTAL DETAILS AND ASSETS DISCLOSURE

D.1 ASSETS

We do not introduce new data in the course of this work. Instead, we use publicly available, widely used image datasets for the purposes of benchmarking and comparison.

D.2 HARDWARE AND SETUP

Each experiment was conducted on 1 NVIDIA Tesla V100 GPUs, each with 32GB of accelerator RAM. The CPUs used were Intel Xeon E5-2698 v4 processors with 20 cores and 384GB of RAM. All experiments were implemented using the PyTorch deep learning framework. We provide our experimental code for reproducibility in the supplementary material.

D.3 DATASETS

D.3.1 IMAGE CLASSIFICATION BENCHMARKS

We evaluate cross-model alignment on three standard image classification benchmarks: Oxford Pets (Parkhi et al., 2012), CIFAR-100 (Krizhevsky et al., 2009), Caltech-101 (Fei-Fei et al., 2004), STL10 (Coates et al., 2011) and DTD (Cimpoi et al., 2014). More details about dataset statistics and splits are provided in Table 1.

Table 1: Detailed statistics of the 5 datasets for image classification.

Dataset	Classes	Train	Val	Test
Caltech-101 (Fei-Fei et al., 2004)	100	4,128	1,649	2,465
Oxford Pets (Parkhi et al., 2012)	37	2,944	736	3,669
CIFAR100 (Krizhevsky et al., 2009)	100	50,000	-	10,000
STL10 (Coates et al., 2011)	10	50,000	-	8,000
DTD (Cimpoi et al., 2014)	47	2,820	1,128	1,692

D.3.2 CONSTRUCTING TEXT TEMPLATES

We construct class-level text prototypes by instantiating natural-language templates with each class name and embedding the resulting strings using each model’s text encoder. All templates are taken directly from the official CLIP repository and prior CLIP evaluation code, ensuring that our protocol matches standard zero-shot classification practice. Table 2 summarizes the exact prompts used for each dataset. We embed all prompts for a class, aggregate them into a single prototype by averaging, and then ℓ_2 -normalizing. Concretely, for templates $\{\tau_k(c)\}_{k=1}^K$ for a class c ,

$$p_c = \text{norm}\left(\frac{1}{K} \sum_k g(\tau_k(c))\right) \tag{6}$$

When the orthogonal map Q is fit on images and evaluated on text, we use these averaged class prototypes p_c for text-side evaluation. In contrast, when fitting Q on text, averaging would collapse each class to a single vector and drastically reduce supervision; therefore, we fit Q on individual prompt embeddings $g(\tau_k(c))_{k=1}^K$, treating each template instantiation as a separate training sample.

Dataset	Text Templates (from official CLIP repo)
Oxford Pets	a photo of a {class}, a type of pet.
CIFAR-100	a photo of a {class}. a blurry photo of a {class}. a black and white photo of a {class}. a low contrast photo of a {class}. a high contrast photo of a {class}. a bad photo of a {class}. a good photo of a {class}. a photo of a small {class}. a photo of a big {class}. a photo of the {class}. a blurry photo of the {class}. a black and white photo of the {class}. a low contrast photo of the {class}. a high contrast photo of the {class}. a bad photo of the {class}. a good photo of the {class}. a photo of the small {class}. a photo of the big {class}.
Caltech-101	a photo of a {class}. a painting of a {class}. a plastic {class}. a sculpture of a {class}. a sketch of a {class}. a tattoo of a {class}. a toy {class}. a rendition of a {class}. an embroidered {class}. a cartoon {class}. a {class} in a video game. a plushie {class}. an origami {class}. art of a {class}. graffiti of a {class}. a drawing of a {class}. a doodle of a {class}. a photo of the {class}. a painting of the {class}. the plastic {class}. a sculpture of the {class}. a sketch of the {class}. a tattoo of the {class}. the toy {class}. a rendition of the {class}. the embroidered {class}. the cartoon {class}. the {class} in a video game. the plushie {class}. the origami {class}. art of the {class}. graffiti of the {class}. a drawing of the {class}. a doodle of the {class}.
STL-10	a photo of a {class}. a photo of the {class}.
DTD	a photo of a {class} texture. a photo of a {class} pattern. a photo of a {class} thing. a photo of a {class} object. a photo of the {class} texture. a photo of the {class} pattern. a photo of the {class} thing. a photo of the {class} object.

Table 2: Text templates used to construct class-level text prototypes for each dataset. All templates are taken from the official CLIP repository. For all reported results, we use a single template per class for evaluation

D.4 TRAINING AND EVALUATION PROTOCOL

D.4.1 LEARNING AN ORTHOGONAL TRANSFORMATION

We consider two independently trained multimodal contrastive models with image encoders f, \tilde{f} and text encoders g, \tilde{g} . Our goal is to learn a single orthogonal map $\mathcal{R} \in O(d)$ that aligns representations across models. All experiments use fixed train/validation/test splits with a fixed random seed. We report the mean and standard deviation averaged over three random seeds. We fit \mathcal{R} using paired embeddings from a single modality (images or text). Let $\{x_i\}_{i=1}^n$ denote paired inputs in the chosen modality, and define

$$X = \begin{bmatrix} f(x_1)^\top \\ \vdots \\ f(x_n)^\top \end{bmatrix}, \quad \tilde{X} = \begin{bmatrix} \tilde{f}(x_1)^\top \\ \vdots \\ \tilde{f}(x_n)^\top \end{bmatrix},$$

with f, \tilde{f} replaced by g, \tilde{g} when fitting on text. For numerical stability, we center the embeddings only during estimation i.e. $X \rightarrow X - \mu_X, \tilde{X} \rightarrow \tilde{X} - \mu_{\tilde{X}}$, and compute the closed form Orthogonal Procrustes solution

$$\mathcal{Q} = \arg \min_{Q \in O(d)} \|XQ - \tilde{X}\|_F = UV^\top,$$

where $U\Sigma V^\top = \text{SVD}(X^\top \tilde{X})$. At evaluation, we deploy \mathcal{Q} as a pure orthogonal transformation on the centered point clouds,

$$z \mapsto Q(z - \mu_A^{(\cdot)}) + \mu_B^{(\cdot)},$$

where $\mu_A^{(\cdot)}$ and $\mu_B^{(\cdot)}$ are modality-specific means of the source and target embeddings (image or text, as appropriate). Centering isolates the rotational relationship by removing this offset while preserving the orthogonal correspondence in the centered space. Note that, we also ablate using pure orthogonal transformation on the raw embedding point clouds and find that mean centering mainly improves *pointwise* cosine agreement, while class-level retrieval and decision geometry remain essentially unchanged even without mean-centering (see [Appendix E.10](#)).

For more details, refer to the pseudocode in [Algorithm 1](#).

Algorithm 1 Pseudocode for fitting an orthogonal cross-model map \mathcal{R} (Procrustes)

```

# f, g: source image/text encoders with outputs of unit-norm
# f_tilde, g_tilde: target image/text encoders with outputs of unit-norm
# pairs: paired samples from ONE modality (images x_i or texts y_i)
# fit_modality: "image" or "text"
# norm(v): l2-normalize v

def fit_procrustes(pairs, fit_modality = 'image'):
    X, X_tilde = [], []
    for s in pairs:
        if fit_modality == "image":
            X.append(f(s))
            X_tilde.append(f_tilde(s))
        else:
            X.append(g(s))
            X_tilde.append(g_tilde(s))

    X = stack_rows(X) # (n, d)
    X_tilde = stack_rows(X_tilde) # (n, d)

    # center for numerical stability during estimation
    Xc = X - mean_row(X)
    Xtc = X_tilde - mean_row(X_tilde)

    # orthogonal Procrustes
    M = transpose(Xc) @ Xtc
    U, S, Vt = svd(M)
    R = U @ Vt

# deployment (used for BOTH modalities at evaluation time)
def deploy(z, R, mu, mu_tilde):
    return R @ (z-mu) + mu_tilde

R = fit_procrustes(pairs, fit_modality)

```

D.4.2 PERFORMANCE METRICS AT EVALUATION

We evaluate alignment at two granularities: the instance level and the class level. For the former, we report paired-image cosine $\langle \mathcal{R}f(x), \tilde{f}(x) \rangle$ and paired-text cosine $\langle \mathcal{R}g(y), \tilde{g}(y) \rangle$ and the correspondingly the Euclidean distance between image embeddings ($\mathcal{R}f(x)$ and $\tilde{f}(x)$) and text embeddings ($\mathcal{R}g(y)$ and $\tilde{g}(y)$), measuring how well a single orthogonal map \mathcal{R} matches corresponding embeddings across models at the instance level. For the class level alignment, we report:

- **Intra-modal alignment across models.** We measure class-level image–image and text–text top-1 retrieval across models. For image–image, each query is an aligned image embedding $\mathcal{R}f(x)$, and we predict its class by nearest-neighbor search over the target image embedding set $\{\tilde{f}(x')\}_{x'}$, counting a hit if the retrieved neighbor shares the same class label. For text–text, each query is an aligned class text embedding $\mathcal{R}g(y_c)$, and we retrieve within the target class text embedding set $\{\tilde{g}(y_{c'})\}_{c'}$, counting a hit if the retrieved class equals c .
- **Cross-modal transfer across models (task accuracy).** We evaluate zero-shot classification under three retrieval settings, each probing a different notion of transfer: (a) aligned images with target text $\hat{c} = \arg \max_c \langle \mathcal{R}f(x), \tilde{g}(y_c) \rangle$ measures whether image–text semantics are preserved when only the image space is aligned; (b) target images with aligned text $\hat{c} = \arg \max_c \langle \tilde{f}(x), \mathcal{R}g(y_c) \rangle$ measures whether text semantics transfer under alignment without modifying the image space; (c) aligned images with aligned text $\hat{c} = \arg \max_c \langle \mathcal{R}f(x), \mathcal{R}g(y_c) \rangle$ measures whether a single map \mathcal{R} induces a coherent shared space across both modalities.

In all, we report five metrics: (1) paired-instance cosine (or euclidean distance), either image $\langle \mathcal{R}f(x), \tilde{f}(x) \rangle$ or text $\langle \mathcal{R}g(y), \tilde{g}(y) \rangle$ depending on the transfer direction; (2) class-level image-image top-1 retrieval or text-text top-1 retrieval; (3) zero-shot accuracy with aligned images and target text, $\hat{c} = \arg \max_c \langle \mathcal{R}f(x), \tilde{g}(y_c) \rangle$; (4) zero-shot accuracy with target images and aligned text, $\hat{c} = \arg \max_c \langle \tilde{f}(x), \mathcal{R}g(y_c) \rangle$; and (5) zero-shot accuracy with aligned images and aligned text, $\hat{c} = \arg \max_c \langle \mathcal{R}f(x), \mathcal{R}g(y_c) \rangle$. Throughout the evaluation, every inner product $\langle \cdot, \cdot \rangle$ is measured through *cosine similarity*, i.e., a dot product between ℓ_2 -normalized embeddings.

E ADDITIONAL EXPERIMENTS

E.1 MODALITY GAP IN CONTRASTIVE MODELS




Figure 9: PCA visualization in 2D of generated embeddings of CIFAR-100 from pre-trained models showing pervasive modality gap in multi-modal contrastive representation learning whereby image and text embeddings are located in two completely separate regions of the embedding space.

Figure 9 visualizes the modality gap in contrastive models via a 2D PCA projection of image and text embeddings. Across CLIP (OpenAI), CLIP (LAION), and FLAVA, image and text representations form two well-separated clusters, occupying distinct regions of the embedding space rather than a unified manifold. Gray lines connect matched image-text pairs, highlighting that although the pairs are semantically aligned, their embeddings remain geometrically distant.

E.2 VISUALIZATION OF MULTIMODAL KERNELS

In Figure 10, we visualize the PCA projection of the image/text geometry across models. We observe that while the absolute coordinates of the embedding cones shift arbitrarily between models, the angular arrangement of the texts with respect to the images remains consistent. Mathematically, this means that the *multimodal kernels* are approximately preserved across models: $\langle f, g \rangle \approx \langle \tilde{f}, \tilde{g} \rangle$. This observation can be viewed as a multimodal analogue of the Platonic Representation Hypothesis (Huh et al., 2024) that posits that models converge to similar *unimodal kernels* ($\langle f, f \rangle \approx \langle \tilde{f}, \tilde{f} \rangle$).




Figure 10: Across CLIP variants, the multimodal kernel $\langle f(x), g(y) \rangle$ (relative angles between image and text embeddings) is strongly preserved (dashed lines), unlike the unimodal kernel $\langle f(x), \tilde{f}(x') \rangle$;

E.3 INDEPENDENTLY TRAINED CONTRASTIVE MODELS DIFFER BY AN ORTHOGONAL MAP THAT IS SHARED ACROSS MODALITIES

In this section, we report complete results across five benchmarks—Oxford Pets (Figure 11), Caltech-101 (Figure 12), CIFAR-100 (Figure 13), STL10 (Figure 14), and DTD (Figure 15)—covering all five evaluation metrics used throughout the paper and explained in Appendix D.4.2. Two consistent conclusions emerge across datasets and model pairs. First, a single orthogonal map \mathcal{Q} accurately aligns image embeddings between independently trained contrastive models, yielding near-oracle image-image agreement. Second, the *same* map transfers across modalities: applying \mathcal{Q} learned from images substantially improves text-text pointwise alignment, cross-model retrieval, and zero-shot classification, without degrading image performance. Together, these results confirm that independently trained contrastive models are related by a shared, modality-invariant orthogonal reparameterization.

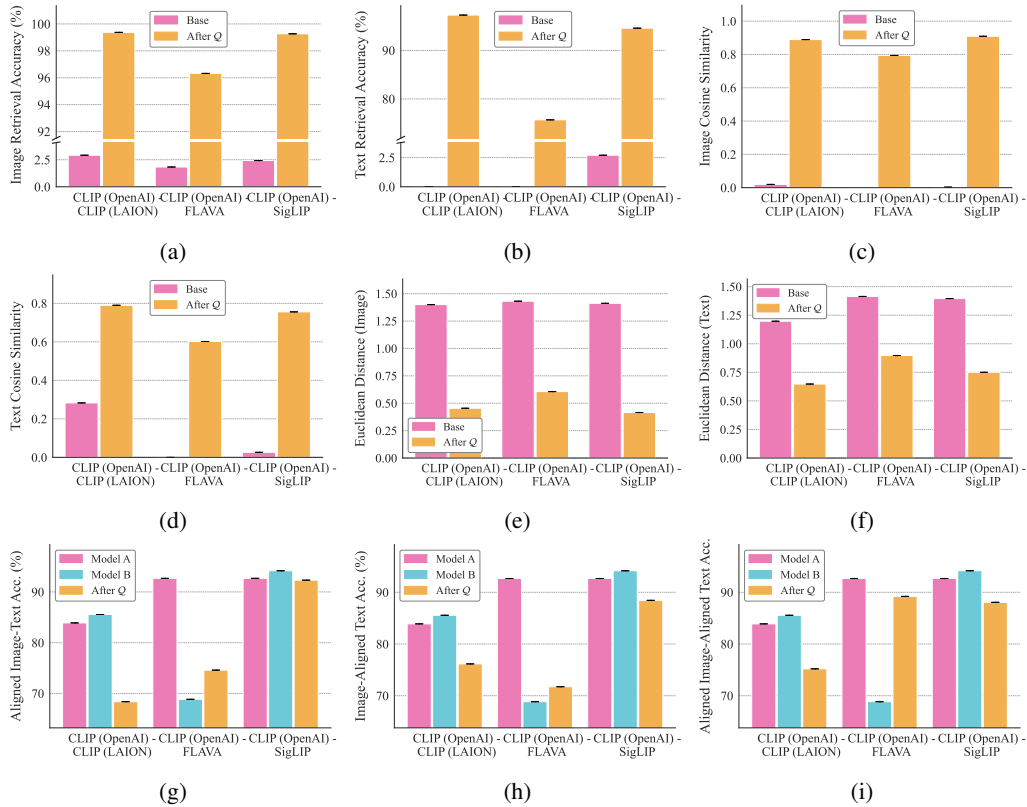


Figure 11: *Cross-model alignment on Oxford Pets across independently trained contrastive models before and after fitting a single orthogonal map \mathcal{R} .* (a) Image–image class retrieval and (b) text–text class retrieval (c) Mean image–image cosine similarity. (d) Mean text–text cosine similarity. (e) Euclidean distance (without normalization) between image embeddings. (f) Euclidean distance (without normalization) between text embeddings. (g) Image–text retrieval using aligned images from model A and text from model B. (h) Image–text retrieval using images from model B and aligned text from model A. (i) Image–text retrieval using aligned images and aligned text from model A. \mathcal{Q} aligns images across models with a single orthogonal map, and the same \mathcal{Q} learned only from image embeddings transfers to text, boosting text-text retrieval from near-chance to near-oracle, all while preserving strong image classification accuracy.

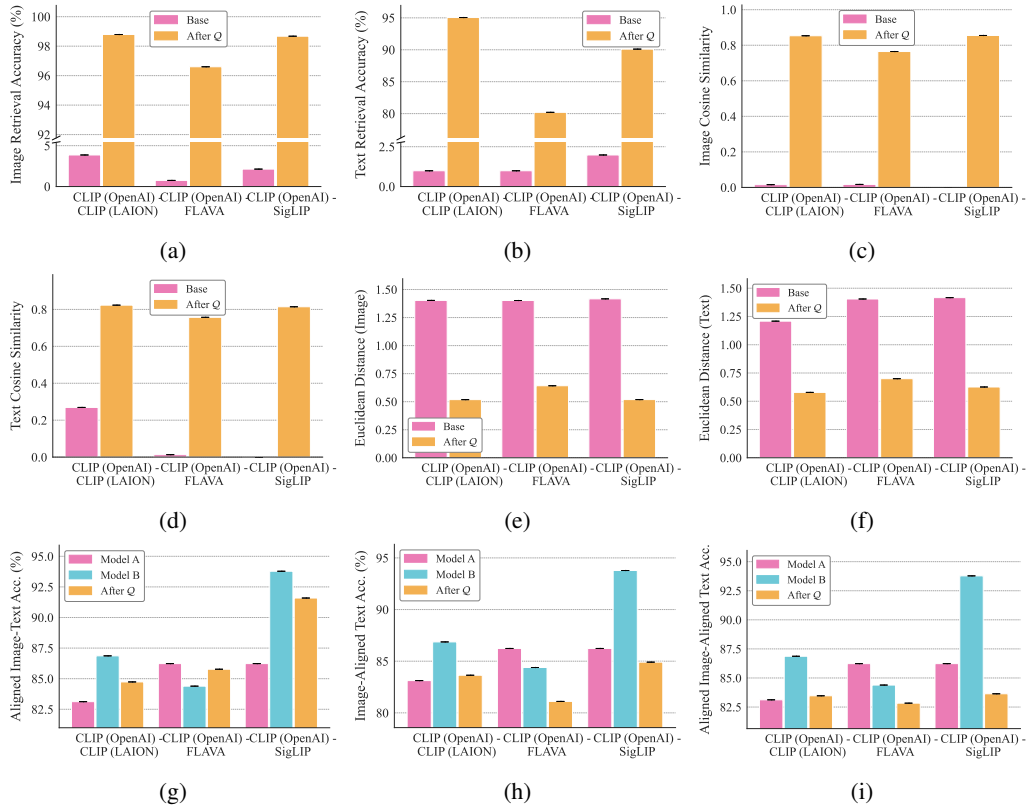


Figure 12: Cross-model alignment on Caltech-101 across independently trained contrastive models before and after fitting a single orthogonal map \mathcal{R} . (a) Image–image class retrieval and (b) text–text class retrieval (c) Mean image–image cosine similarity. (d) Mean text–text cosine similarity. (e) Euclidean distance (without normalization) between image embeddings. (f) Euclidean distance (without normalization) between text embeddings. (g) Image–text retrieval using aligned images from model A and text from model B. (h) Image–text retrieval using images from model B and aligned text from model A. (i) Image–text retrieval using aligned images and aligned text from model A. \mathcal{Q} aligns images across models with a single orthogonal map, and the same \mathcal{Q} learned only from image embeddings transfers to text, boosting text-text retrieval from near-chance to near-oracle, all while preserving strong image classification accuracy.

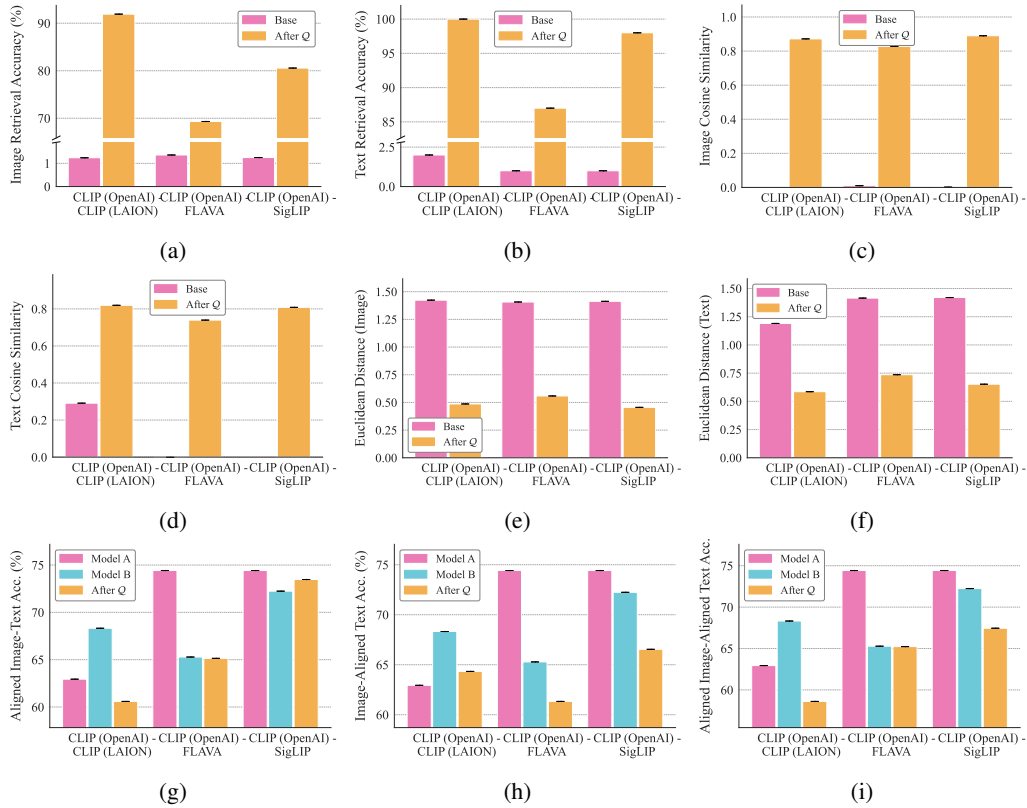


Figure 13: *Cross-model alignment on CIFAR-100 across independently trained contrastive models before and after fitting a single orthogonal map \mathcal{R} .* (a) Image–image class retrieval and (b) text–text class retrieval (c) Mean image–image cosine similarity. (d) Mean text–text cosine similarity. (e) Euclidean distance (without normalization) between image embeddings. (f) Euclidean distance (without normalization) between text embeddings. (g) Image–text retrieval using aligned images from model A and text from model B. (h) Image–text retrieval using images from model B and aligned text from model A. (i) Image–text retrieval using aligned images and aligned text from model A. \mathcal{Q} aligns images across models with a single orthogonal map, and the same \mathcal{Q} learned only from image embeddings transfers to text, boosting text-text retrieval from near-chance to near-oracle, all while preserving strong image classification accuracy.

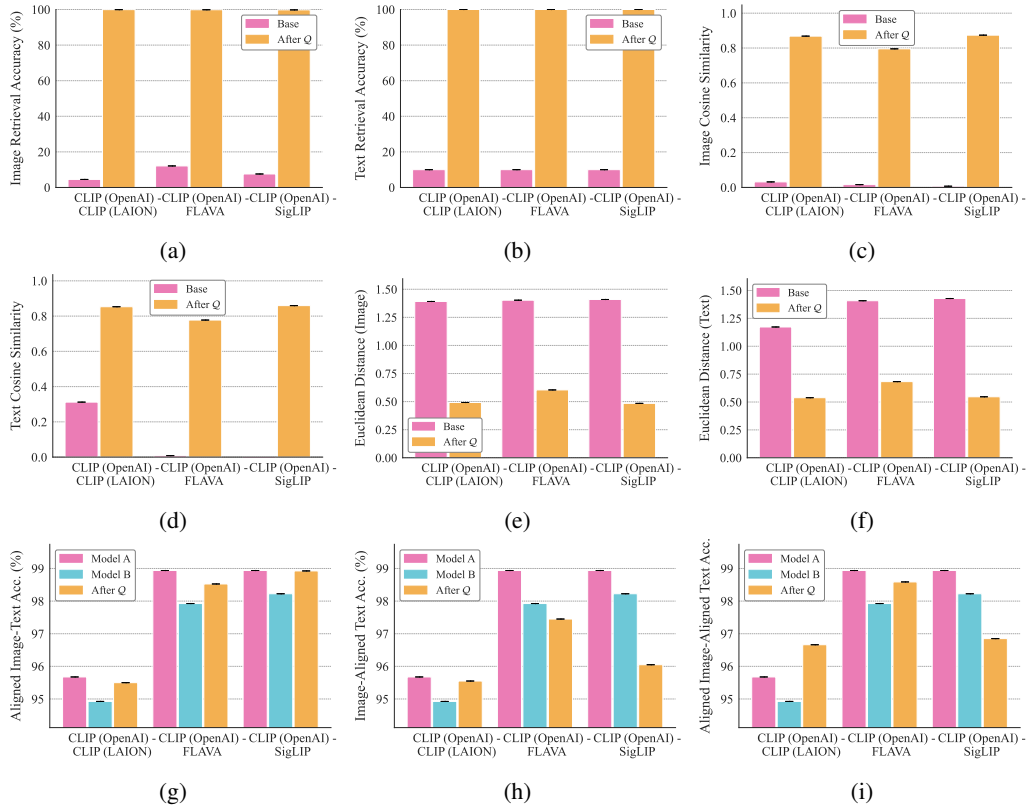


Figure 14: *Cross-model alignment on STL10 across independently trained contrastive models before and after fitting a single orthogonal map \mathcal{R} .* (a) Image–image class retrieval and (b) text–text class retrieval (c) Mean image–image cosine similarity. (d) Mean text–text cosine similarity. (e) Euclidean distance (without normalization) between image embeddings. (f) Euclidean distance (without normalization) between text embeddings. (g) Image–text retrieval using aligned images from model A and text from model B. (h) Image–text retrieval using images from model B and aligned text from model A. (i) Image–text retrieval using aligned images and aligned text from model A. \mathcal{Q} aligns images across models with a single orthogonal map, and the same \mathcal{Q} learned only from image embeddings transfers to text, boosting text-text retrieval from near-chance to near-oracle, all while preserving strong image classification accuracy.

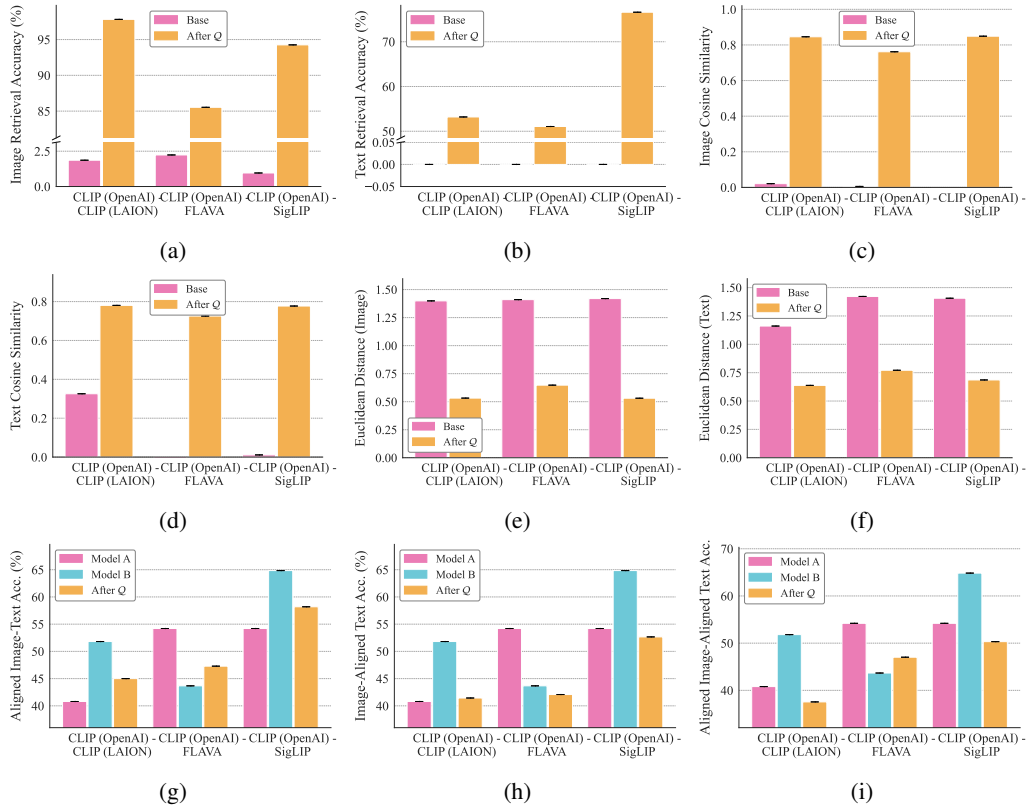


Figure 15: *Cross-model alignment on DTD across independently trained contrastive models before and after fitting a single orthogonal map \mathcal{R} .* (a) Image–image class retrieval and (b) text–text class retrieval (c) Mean image–image cosine similarity. (d) Mean text–text cosine similarity. (e) Euclidean distance (without normalization) between image embeddings. (f) Euclidean distance (without normalization) between text embeddings. (g) Image–text retrieval using aligned images from model A and text from model B. (h) Image–text retrieval using images from model B and aligned text from model A. (i) Image–text retrieval using aligned images and aligned text from model A. \mathcal{Q} aligns images across models with a single orthogonal map, and the same \mathcal{Q} learned only from image embeddings transfers to text, boosting text-text retrieval from near-chance to near-oracle, all while preserving strong image classification accuracy.

E.4 ONLY A FEW DATA POINTS ARE NEEDED TO LEARN THE ORTHOGONAL MAP

In Appendix C.5.2, we proved that if the multimodal kernels induced by two contrastive models agree on a sufficiently rich but *small finite* anchor set, then a single global orthogonal map aligns their representations across both modalities. We empirically validate this prediction by fitting Q using paired images from only N classes and evaluating transfer on the remaining unseen classes. We report this analysis across additional model pairs and datasets. The results for CLIP (OpenAI) aligned to CLIP (OpenAI), CLIP (LAION), FLAVA, and SIGLIP on Caltech-101 are shown in Figures 16 to 19. For CIFAR100 and Oxford Pets, the analogous results are reports in Figures 20 to 23 and Figures 24 to 27.

Across all settings, the same trend holds: performance on both seen and unseen classes improves quickly with just a few anchor classes and essentially saturates once N reaches a modest value, after which additional anchors provide little benefit. Thus, practitioners can recover near-full cross-model transfer by fitting Q on a lightweight image-only calibration set, rather than curating large-scale cross-model supervision.

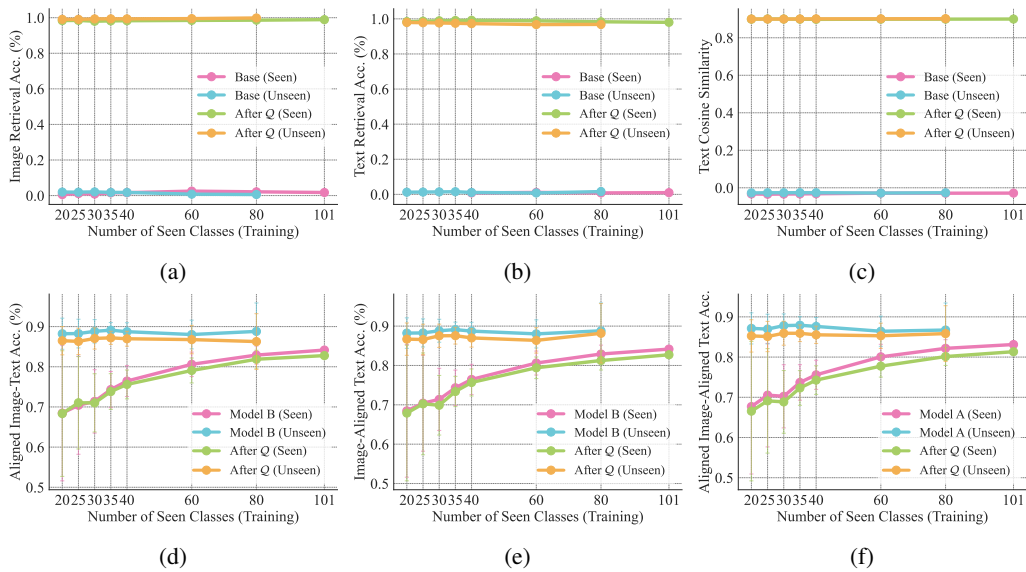


Figure 16: *Generalization of orthogonal alignment under limited supervision (Caltech-101; CLIP ViT-B/32 to CLIP ViT-B/16, OpenAI)*. (a) Image-image class retrieval, (b) text-text class retrieval, and (c) mean text-text cosine similarity, each reported on seen and unseen classes. (d-f) Downstream transfer: (d) aligned images from model A with text from model B, (e) images from model B with aligned text from model A, and (f) aligned images with aligned text from model A. \mathcal{R} learned from few classes transfers across modalities and generalizes to unseen classes, achieving near-oracle cross-model retrieval and classification.

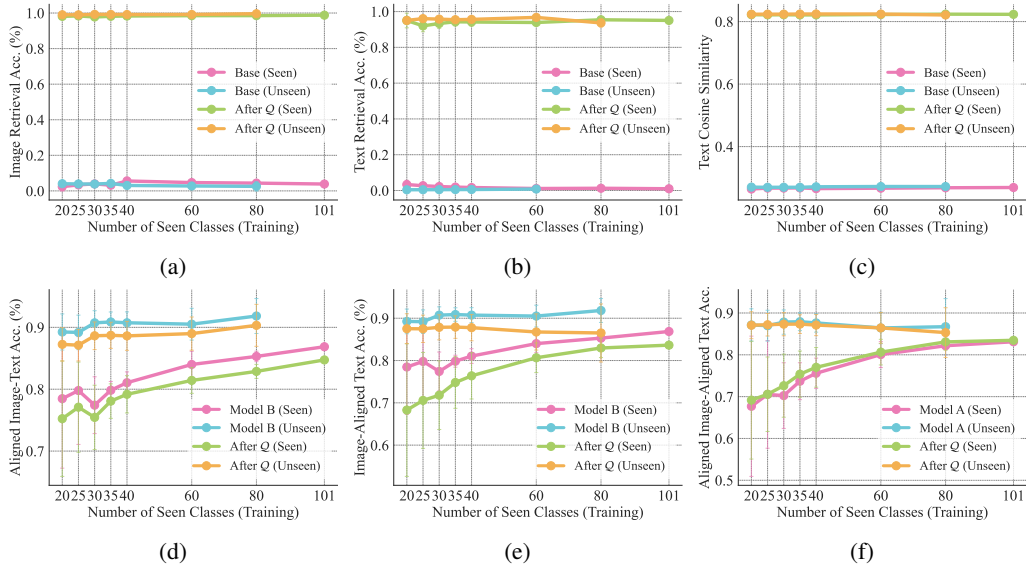


Figure 17: *Generalization of orthogonal alignment under limited supervision (Caltech-101; CLIP ViT-B/32 (OpenAI) to CLIP ViT-B/32 (LAION)).* (a) Image-image class retrieval, (b) text-text class retrieval, and (c) mean text-text cosine similarity, each reported on seen and unseen classes. (d-f) Downstream transfer: (d) aligned images from model A with text from model B, (e) images from model B with aligned text from model A, and (f) aligned images with aligned text from model A. \mathcal{R} learned from few classes transfers across modalities and generalizes to unseen classes, achieving near-oracle cross-model retrieval and classification.

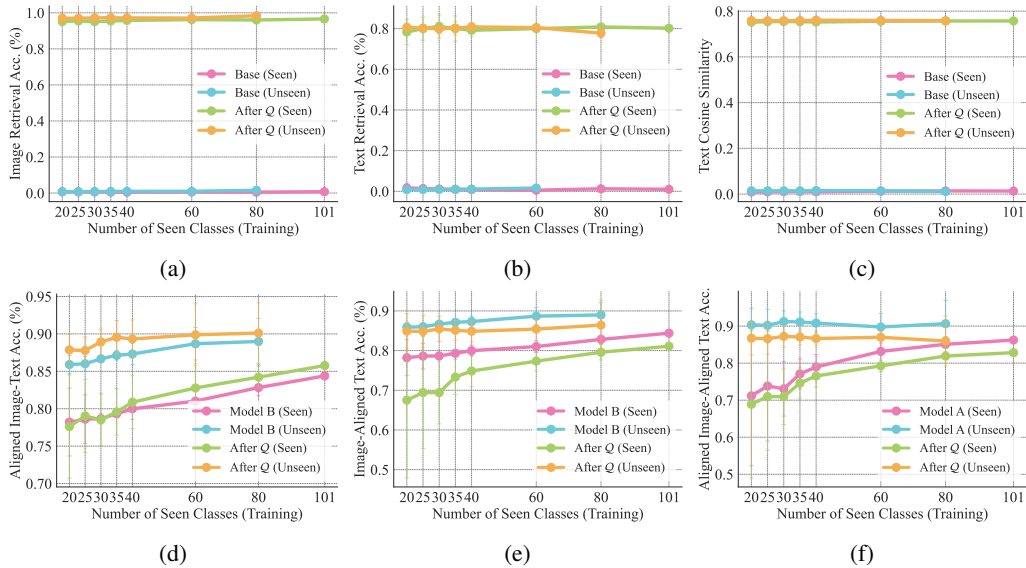


Figure 18: *Generalization of orthogonal alignment under limited supervision (Caltech-101; CLIP ViT-L/14 (OpenAI) to FLAVA)* (a) Image-image class retrieval, (b) text-text class retrieval, and (c) mean text-text cosine similarity, each reported on seen and unseen classes. (d-f) Downstream transfer: (d) aligned images from model A with text from model B, (e) images from model B with aligned text from model A, and (f) aligned images with aligned text from model A. \mathcal{R} learned from few classes transfers across modalities and generalizes to unseen classes, achieving near-oracle cross-model retrieval and classification.

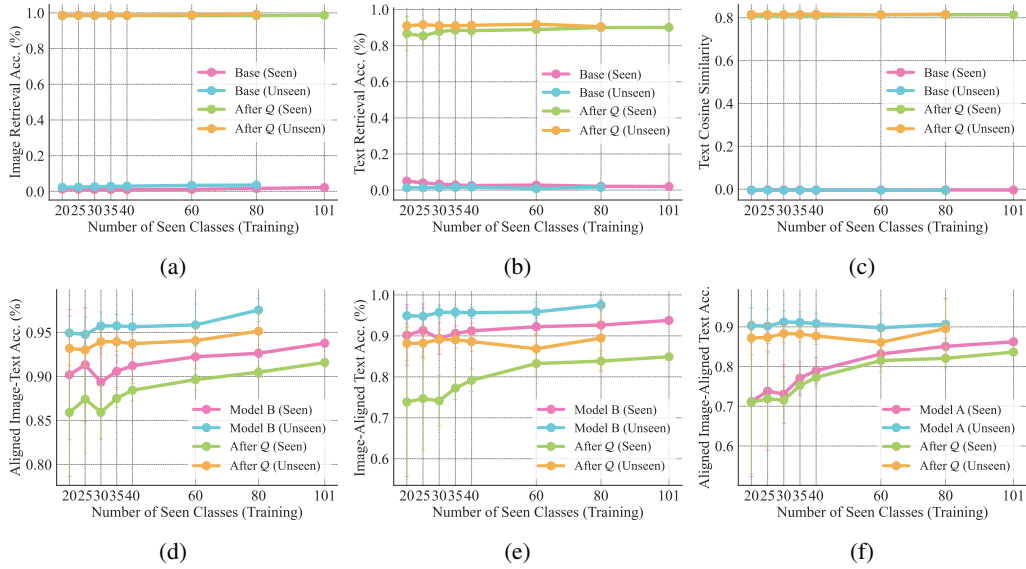


Figure 19: *Generalization of orthogonal alignment under limited supervision (Caltech-101; CLIP ViT-L/14 (OpenAI) to SigLIP.* (a) Image-image class retrieval, (b) text-text class retrieval, and (c) mean text-text cosine similarity, each reported on seen and unseen classes. (d-f) Downstream transfer: (d) aligned images from model A with text from model B, (e) images from model B with aligned text from model A, and (f) aligned images with aligned text from model A. \mathcal{R} learned from few classes transfers across modalities and generalizes to unseen classes, achieving near-oracle cross-model retrieval and classification.

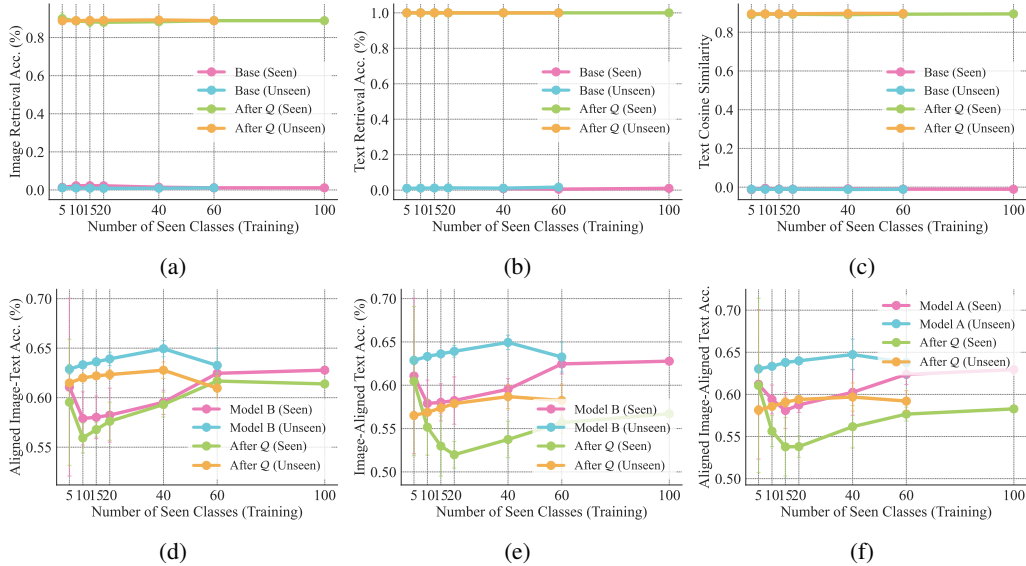


Figure 20: *Generalization of orthogonal alignment under limited supervision (CIFAR-100; CLIP ViT-B/32 to CLIP ViT-B/16, OpenAI).* (a) Image-image class retrieval, (b) text-text class retrieval, and (c) mean text-text cosine similarity, each reported on seen and unseen classes. (d-f) Downstream transfer: (d) aligned images from model A with text from model B, (e) images from model B with aligned text from model A, and (f) aligned images with aligned text from model A. \mathcal{R} learned from few classes transfers across modalities and generalizes to unseen classes, achieving near-oracle cross-model retrieval and classification.

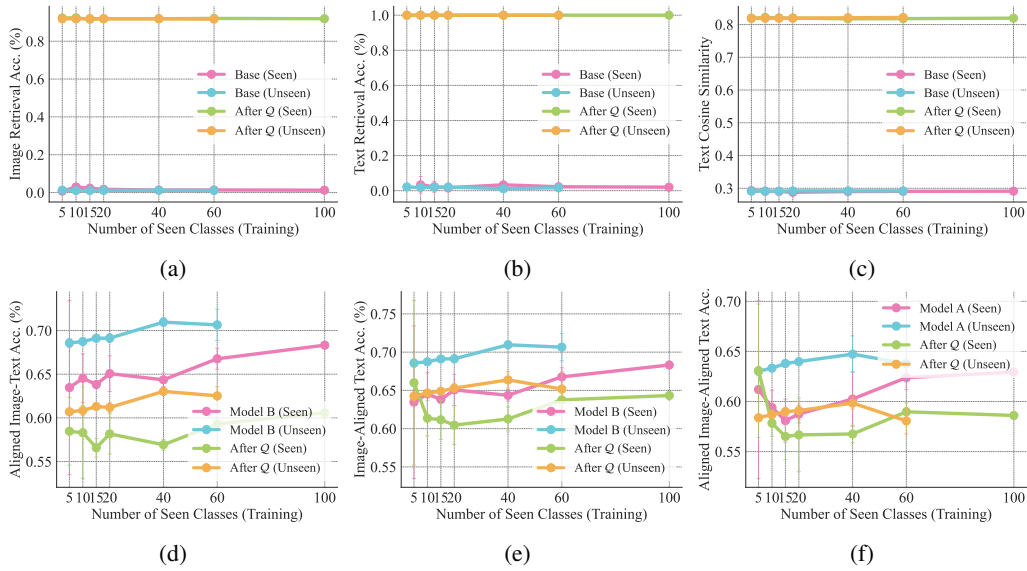


Figure 21: *Generalization of orthogonal alignment under limited supervision (CIFAR100; CLIP ViT-B/32 (OpenAI) to CLIP ViT-B/32 (LAION).* (a) Image-image class retrieval, (b) text-text class retrieval, and (c) mean text-text cosine similarity, each reported on seen and unseen classes. (d-f) Downstream transfer: (d) aligned images from model A with text from model B, (e) images from model B with aligned text from model A, and (f) aligned images with aligned text from model A. \mathcal{R} learned from few classes transfers across modalities and generalizes to unseen classes, achieving near-oracle cross-model retrieval and classification.

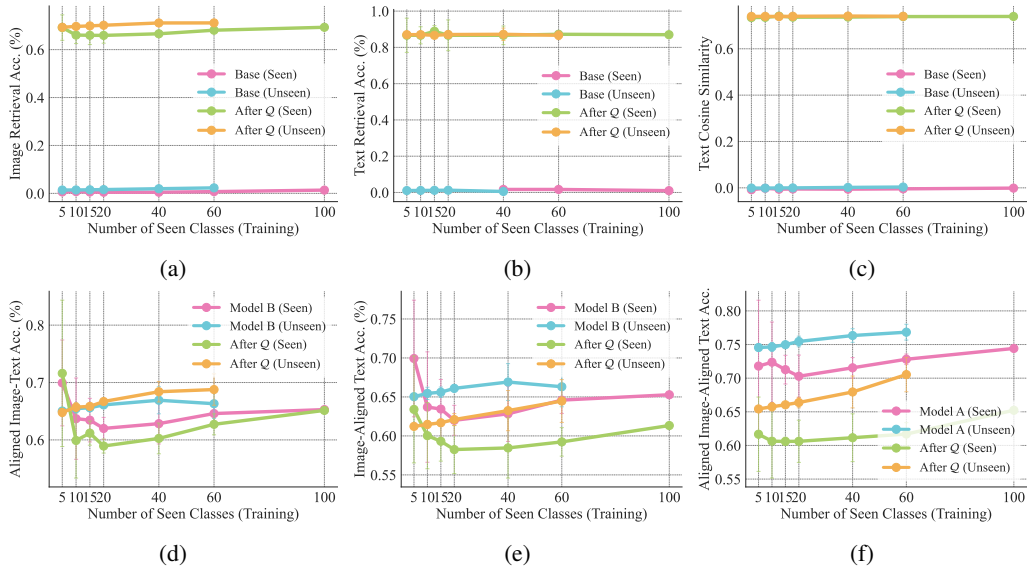


Figure 22: *Generalization of orthogonal alignment under limited supervision (CIFAR100; CLIP ViT-L/14 (OpenAI) to FLAVA.* (a) Image-image class retrieval, (b) text-text class retrieval, and (c) mean text-text cosine similarity, each reported on seen and unseen classes. (d-f) Downstream transfer: (d) aligned images from model A with text from model B, (e) images from model B with aligned text from model A, and (f) aligned images with aligned text from model A. \mathcal{R} learned from few classes transfers across modalities and generalizes to unseen classes, achieving near-oracle cross-model retrieval and classification.

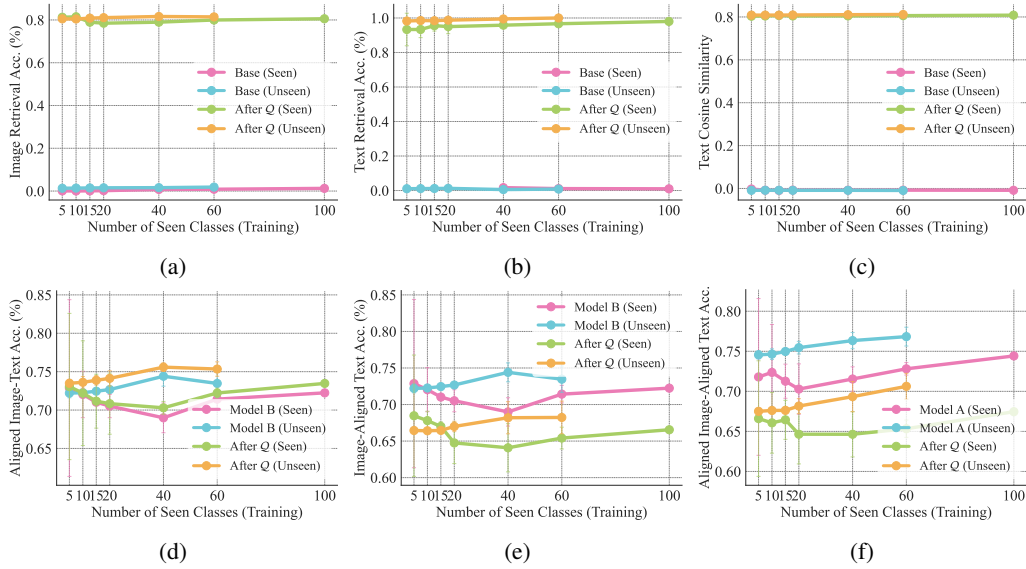


Figure 23: *Generalization of orthogonal alignment under limited supervision (CIFAR100; CLIP ViT-L/14 (OpenAI) to SigLIP.* (a) Image-image class retrieval, (b) text-text class retrieval, and (c) mean text-text cosine similarity, each reported on seen and unseen classes. (d-f) Downstream transfer: (d) aligned images from model A with text from model B, (e) images from model B with aligned text from model A, and (f) aligned images with aligned text from model A. \mathcal{R} learned from few classes transfers across modalities and generalizes to unseen classes, achieving near-oracle cross-model retrieval and classification.

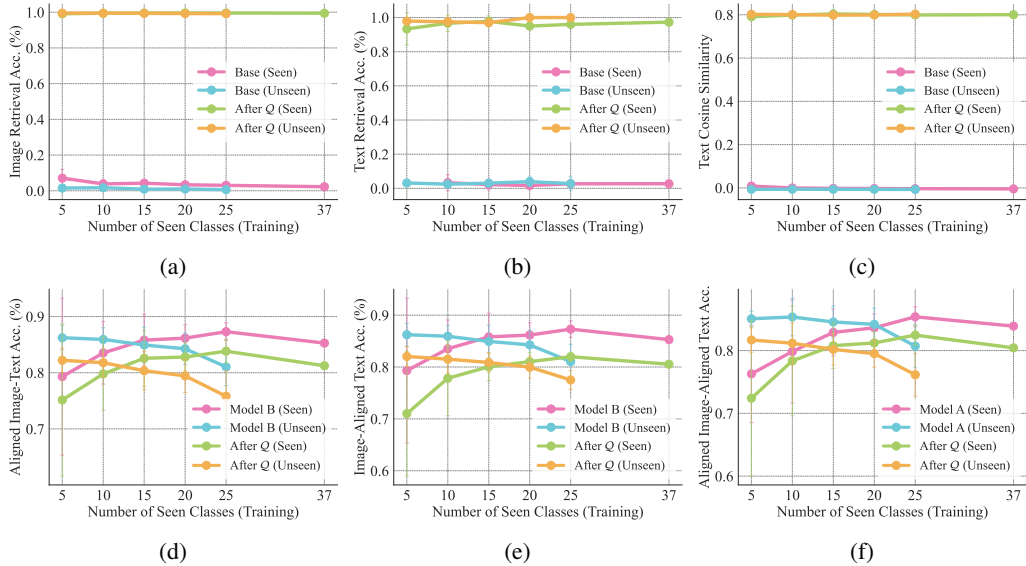


Figure 24: *Generalization of orthogonal alignment under limited supervision (Oxford Pets; CLIP ViT-B/32 to CLIP ViT-B/16, OpenAI).* (a) Image-image class retrieval, (b) text-text class retrieval, and (c) mean text-text cosine similarity, each reported on seen and unseen classes. (d-f) Downstream transfer: (d) aligned images from model A with text from model B, (e) images from model B with aligned text from model A, and (f) aligned images with aligned text from model A. \mathcal{R} learned from few classes transfers across modalities and generalizes to unseen classes, achieving near-oracle cross-model retrieval and classification.

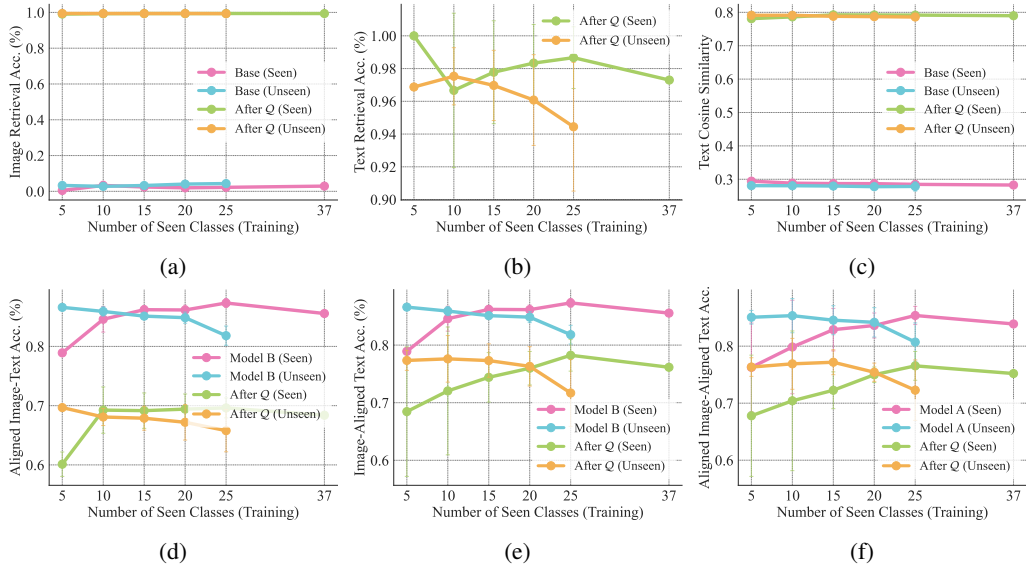


Figure 25: *Generalization of orthogonal alignment under limited supervision (Oxford Pets; CLIP ViT-B/32 (OpenAI) to CLIP ViT-B/32 (LAION)).* (a) Image-image class retrieval, (b) text-text class retrieval, and (c) mean text-text cosine similarity, each reported on seen and unseen classes. (d-f) Downstream transfer: (d) aligned images from model A with text from model B, (e) images from model B with aligned text from model A, and (f) aligned images with aligned text from model A. \mathcal{R} learned from few classes transfers across modalities and generalizes to unseen classes, achieving near-oracle cross-model retrieval and classification.

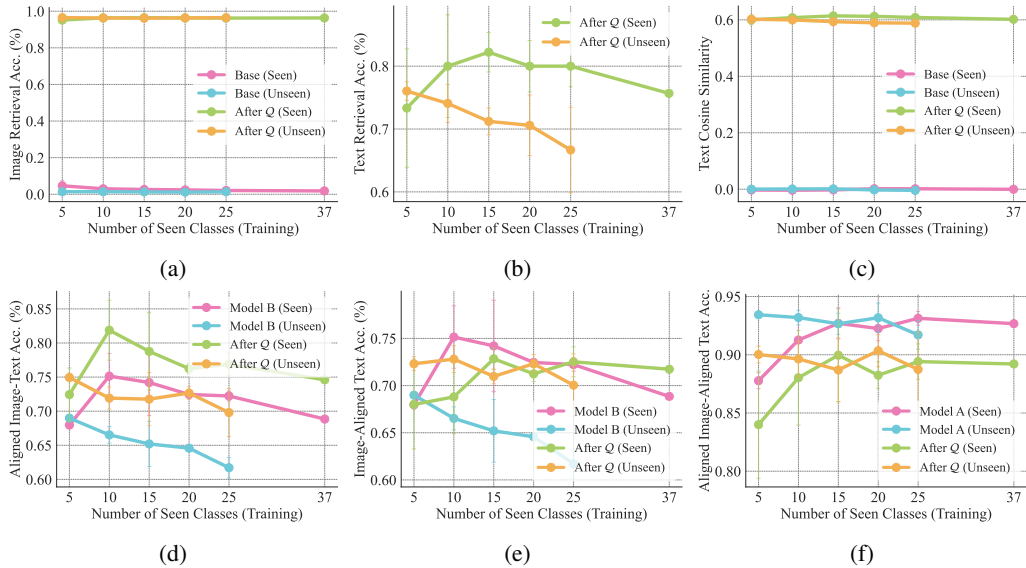


Figure 26: *Generalization of orthogonal alignment under limited supervision (Oxford Pets; CLIP ViT-L/14 (OpenAI) to FLAVA)* (a) Image-image class retrieval, (b) text-text class retrieval, and (c) mean text-text cosine similarity, each reported on seen and unseen classes. (d-f) Downstream transfer: (d) aligned images from model A with text from model B, (e) images from model B with aligned text from model A, and (f) aligned images with aligned text from model A. \mathcal{R} learned from few classes transfers across modalities and generalizes to unseen classes, achieving near-oracle cross-model retrieval and classification.

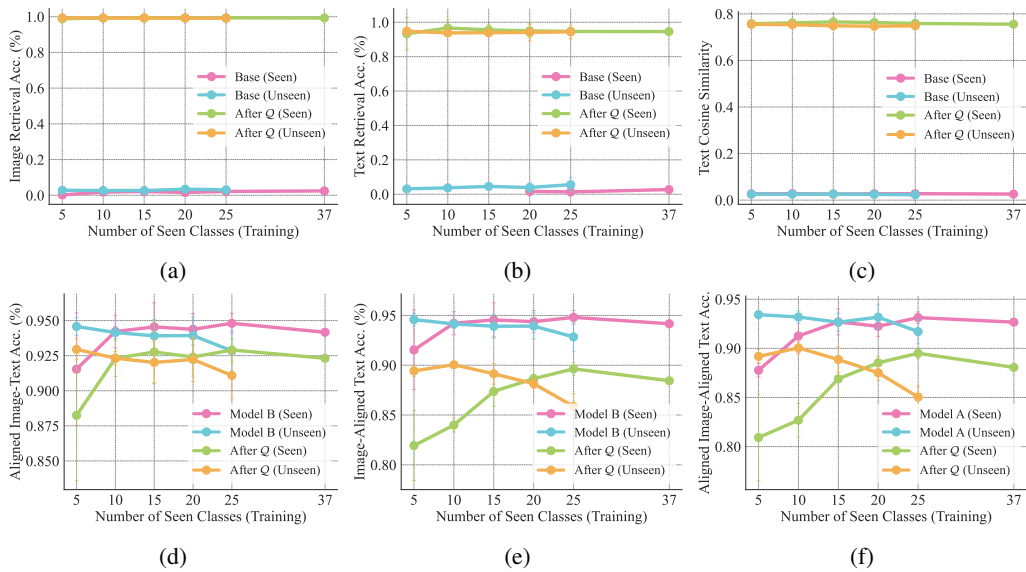


Figure 27: *Generalization of orthogonal alignment under limited supervision (Oxford Pets; CLIP ViT-L/14 (OpenAI) to SigLIP. (a) Image-image class retrieval, (b) text-text class retrieval, and (c) mean text-text cosine similarity, each reported on seen and unseen classes. (d-f) Downstream transfer: (d) aligned images from model A with text from model B, (e) images from model B with aligned text from model A, and (f) aligned images with aligned text from model A. \mathcal{R} learned from few classes transfers across modalities and generalizes to unseen classes, achieving near-oracle cross-model retrieval and classification.*

E.5 THE LEARNED ORTHOGONAL MAP GENERALIZES BROADLY

The previous experiment shows that Q is identifiable from few anchors within a dataset. We next ask whether the same Q depends on the downstream distribution used to estimate it. We fit Q using paired images from one dataset (Oxford Pets) and evaluate on a different dataset (Caltech-101). Results for transfer from CLIP (OpenAI) to CLIP (OpenAI), CLIP (LAION), FLAVA, and SIGLIP are shown in Figures 28 to 31 respectively.

Across all figures, the orthogonal map learned on images from one dataset transfers to the text from another dataset, as shown by strong text-text cosine similarity (subplot c). Further, downstream image-text retrieval metrics (subplots d,e,f) remain strong under transfer—often closely matching or even exceeding an in-domain fit—indicating that Q generalizes beyond the calibration dataset.

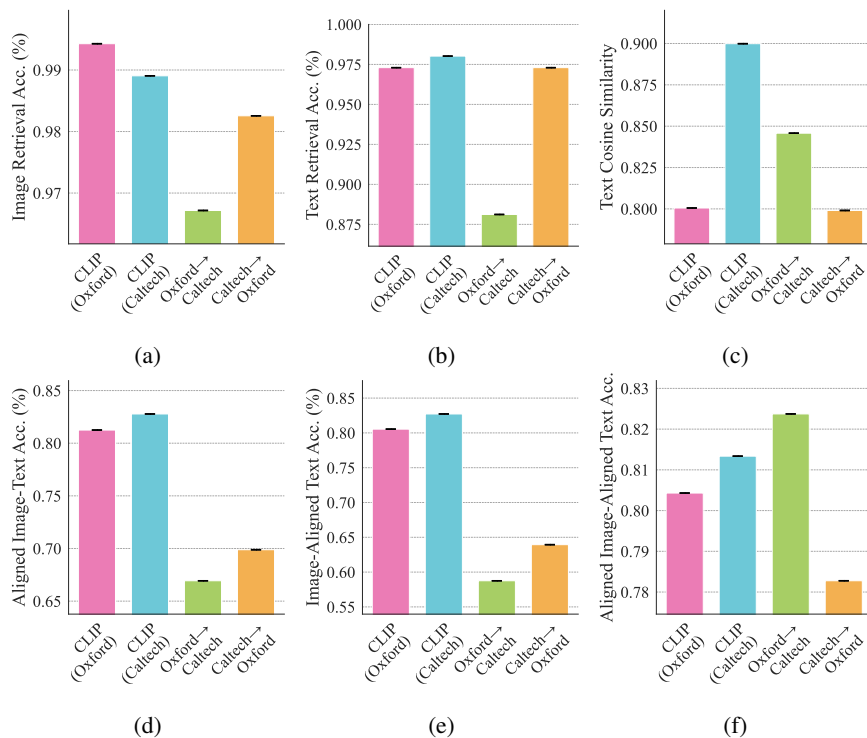


Figure 28: Transfer of Q learnt on Oxford Pets to Caltech-101 and vice-versa across CLIP ViT-B/32 and ViT-B/16 (OpenAI). (a) Image-image class retrieval and (b) text-text class retrieval (c) Mean text-text cosine similarity. (d) Image-text retrieval using aligned images from model A and text from model B. (e) Image-text retrieval using images from model B and aligned text from model A. (f) Image-text retrieval using aligned images and aligned text from model A.

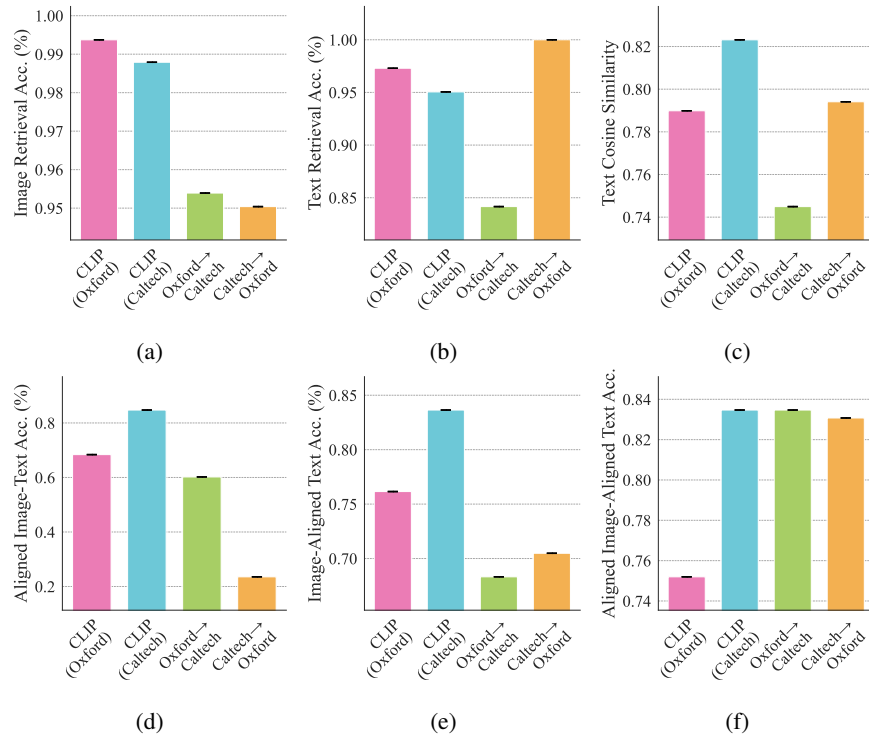


Figure 29: Transfer of \mathcal{Q} learnt on Oxford Pets to Caltech-101 and vice-versa across CLIP ViT-B/32 (OpenAI) and ViT-B/32 (LAION) (a) Image–image class retrieval and (b) text–text class retrieval (c) Mean text–text cosine similarity. (d) Image–text retrieval using aligned images from model A and text from model B. (e) Image–text retrieval using images from model B and aligned text from model A. (f) Image–text retrieval using aligned images and aligned text from model A.

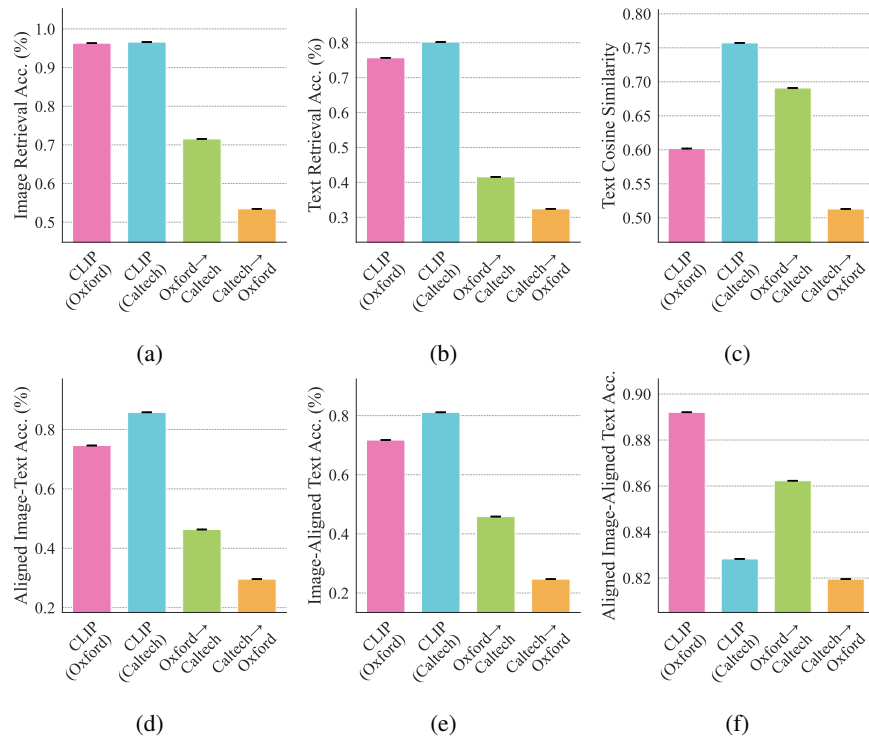


Figure 30: Transfer of \mathcal{Q} learnt on Oxford Pets to Caltech-101 and vice-versa across CLIP ViT-L/14 (OpenAI) and FLAVA (a) Image-image class retrieval and (b) text-text class retrieval (c) Mean text-text cosine similarity. (d) Image-text retrieval using aligned images from model A and text from model B. (e) Image-text retrieval using images from model B and aligned text from model A. (f) Image-text retrieval using aligned images and aligned text from model A.

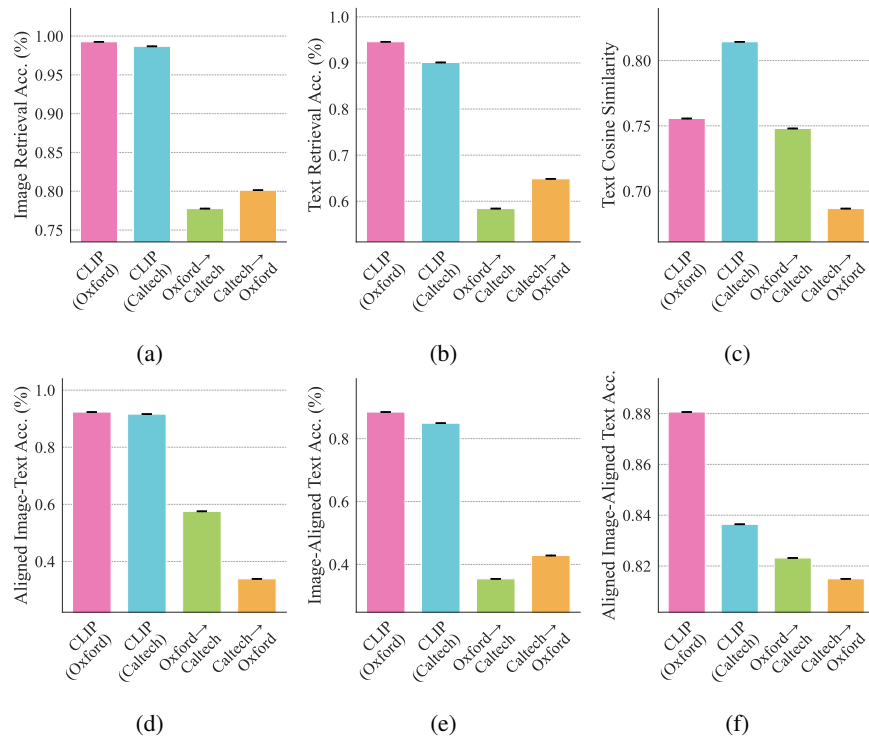


Figure 31: Transfer of \mathcal{Q} learnt on Oxford Pets to Caltech-101 and vice-versa across CLIP ViT-L/14 (OpenAI) and SigLIP. (a) Image-image class retrieval and (b) text-text class retrieval (c) Mean text-text cosine similarity. (d) Image-text retrieval using aligned images from model A and text from model B. (e) Image-text retrieval using images from model B and aligned text from model A. (f) Image-text retrieval using aligned images and aligned text from model A.

E.6 LEARNING THE ORTHOGONAL MAP FROM TEXT INSTEAD OF IMAGES TRANSFERS TO IMAGES

So far, we fit Q using paired images and evaluated whether it transfers to the text modality. We now ask the converse: can we fit Q using only text and recover the same cross-model transform that governs images? Figures 32 to 34, shows these results for Caltech-101, CIFAR-100 and Oxford Pets, respectively. As observed from these figures, across datasets and model pairs, this text-trained map substantially improves downstream task accuracy after Q (image-text retrieval), recovering a large fraction of the stronger model’s performance. However, compared to image-trained maps, text-trained maps yield weaker visual transfer (image-image retrieval and paired-image cosine), indicating that dense image supervision better constrains the shared orthogonal map than class-level text anchors.

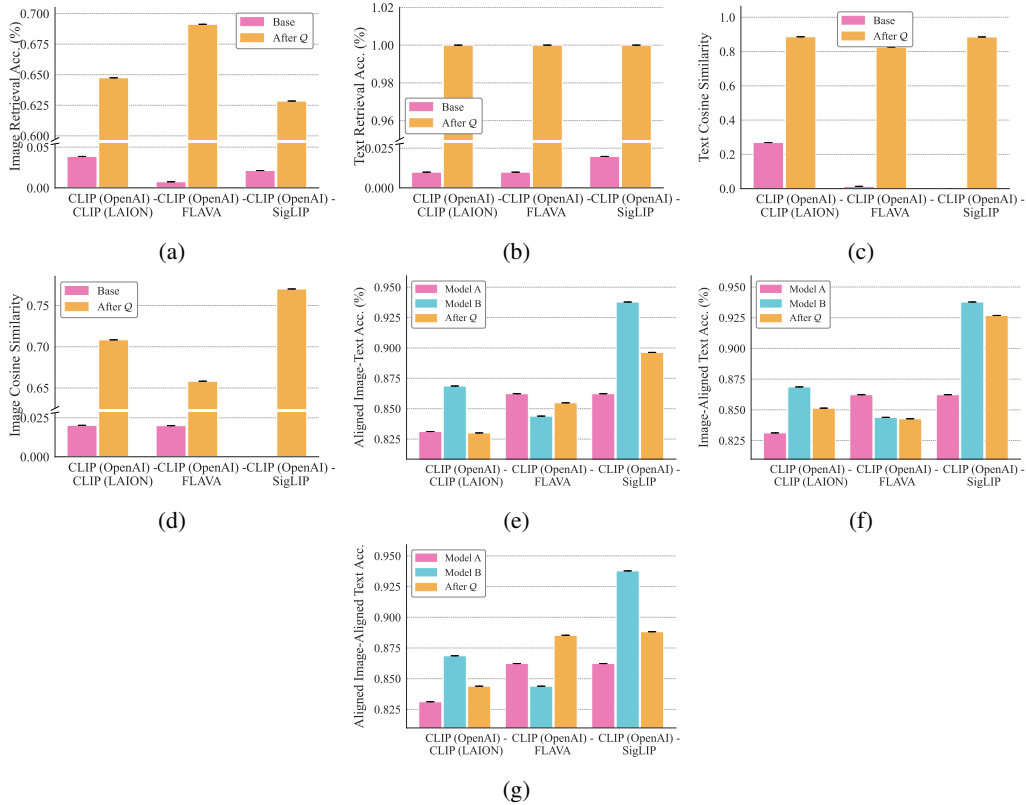


Figure 32: Cross-model alignment on Caltech-101 learned from text prototypes instead of images, before and after fitting a single orthogonal map Q . (a) Image–image class retrieval and (b) text–text class retrieval (c) Mean text–text cosine similarity. (d) Mean image–image cosine similarity. (e) Image–text retrieval using aligned images from model A and text from model B. (f) Image–text retrieval using images from model B and aligned text from model A. (g) Image–text retrieval using aligned images and aligned text from model A.

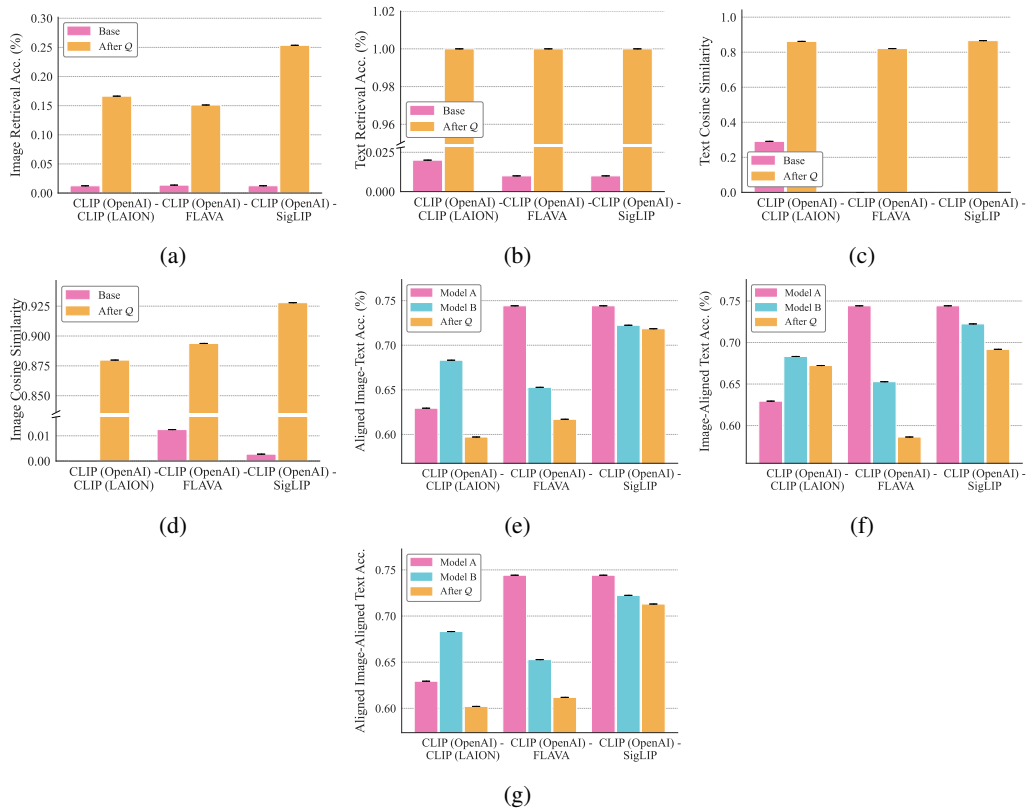


Figure 33: Cross-model alignment on CIFAR-100 learned from text prototypes alone before and after fitting a single orthogonal map Q . (a) Image-image class retrieval and (b) text-text class retrieval (c) Mean text-text cosine similarity. (d) Mean image-image cosine similarity. (e) Image-text retrieval using aligned images from model A and text from model B. (f) Image-text retrieval using images from model B and aligned text from model A. (g) Image-text retrieval using aligned images and aligned text from model A.

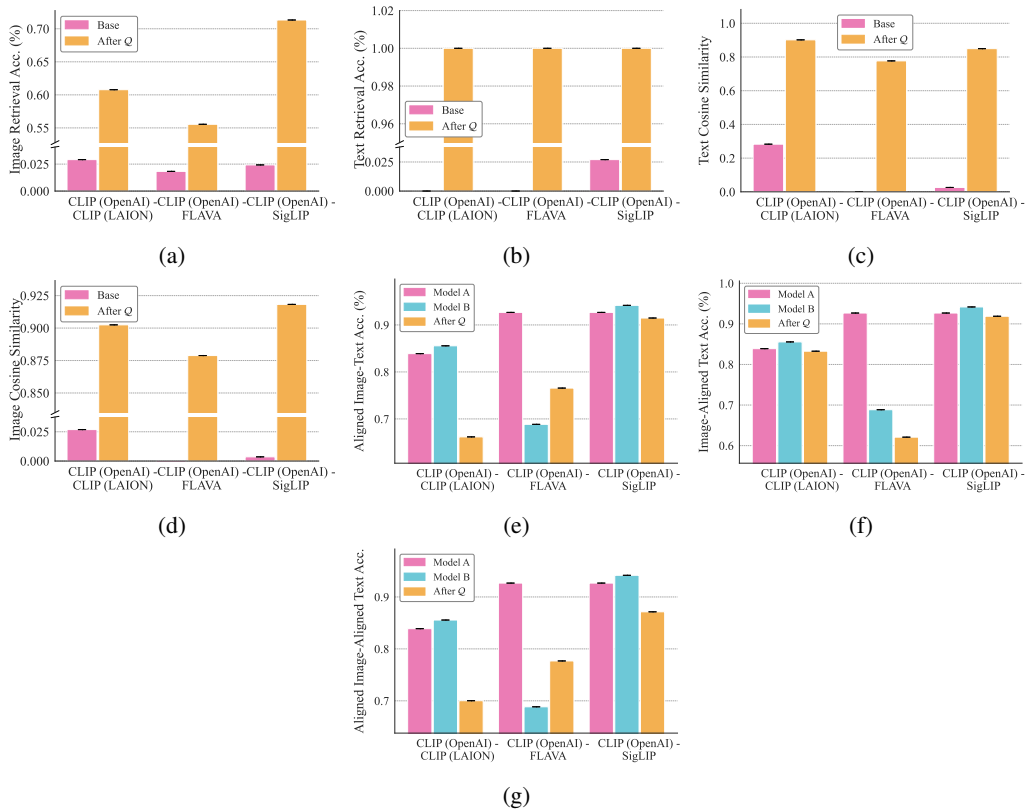


Figure 34: Cross-model alignment on Oxford Pets learned from text prototypes alone before and after fitting a single orthogonal map Q . (a) Image-image class retrieval and (b) text-text class retrieval (c) Mean text-text cosine similarity. (d) Mean image-image cosine similarity. (e) Image-text retrieval using aligned images from model A and text from model B. (f) Image-text retrieval using images from model B and aligned text from model A. (g) Image-text retrieval using aligned images and aligned text from model A.

E.7 CYCLE CONSISTENCY AND CONSISTENCY UNDER COMPOSITION

Our theory and experiment so far posit that different contrastive models represent the same underlying semantic geometry, differing only by a global orthogonal reparameterization. This would imply two forms of consistency: (1) aligning model (f, g) to (\tilde{f}, \tilde{g}) and then (\tilde{f}, \tilde{g}) to (\hat{f}, \hat{g}) should align (f, g) to (\hat{f}, \hat{g}) (composition), and aligning (f, g) to (\tilde{f}, \tilde{g}) should be undone by aligning (\tilde{f}, \tilde{g}) back to (f, g) (cycle). We therefore test transitivity and cycle consistency to ensure the learned maps reflect a shared geometry rather than pair-specific fitting.

In this section, we present complete results for cycle consistency and composition on Caltech-101 in [Figure 35](#) and [Figure 36](#) respectively, on CIFAR100 in [Figure 37](#) and [Figure 38](#) respectively, and on Oxford Pets in [Figure 39](#) and [Figure 40](#) respectively.

Across all composition figures i.e. [Figures 36, 38](#) and [40](#), we compare composing two learned isometries, $Q_{BC} \circ Q_{AB}$, with directly learning Q_{AC} between the endpoint models. Across all model triplets, the composed map almost matches the direct map. Strikingly, in some cases, the composed map outperforms the direct map, showing that the intermediate model’s stronger geometry propagates through the chain. Thus, isometries learned on disjoint model pairs compose reliably and can align a model pair that has never been calibrated together.

Across all cycle-consistency figures i.e. [Figures 35, 37](#) and [39](#), we compare the downstream performance of Model A to a round trip that applies the forward map followed by the reverse map (the composition labeled $Q_{BA} \circ Q_{AB}$). Since each map is estimated via closed-form Procrustes, Q_{BA} is the inverse of Q_{AB} by construction; we therefore fit them on independent 95% random subsets of the training pairs to test under finite-sample estimation. Across all pairs and datasets, image-text classification accuracies (across different aligned pairs) nearly matches the original model, indicating that the learned forward and reverse isometries behave as near-inverses and that the induced coordinate change is stable.

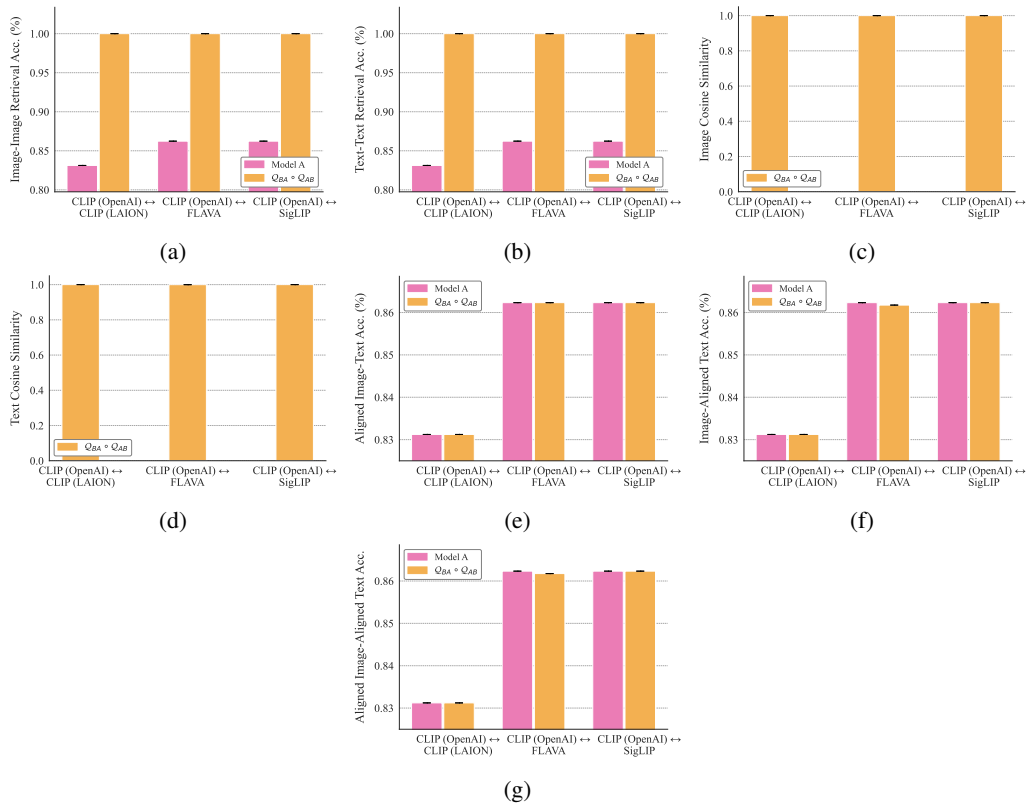


Figure 35: *Cycle consistency on Caltech-101.* (a) Image–image class retrieval. (b) text–text class retrieval (c) Mean image-image cosine similarity. (c) Mean text-text cosine similarity. (e) Image–text retrieval using aligned images from model A and text from model B. (f) Image–text retrieval using images from model B and aligned text from model A. (g) Image–text retrieval using aligned images and aligned text from model A.

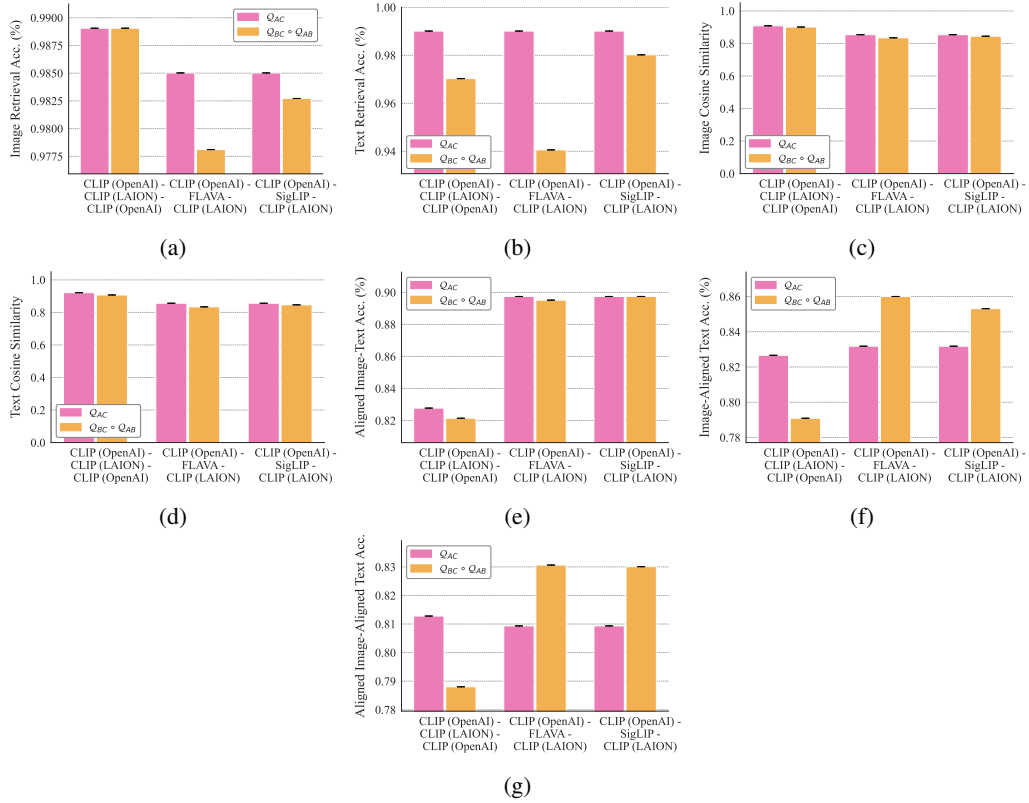


Figure 36: Compositionality on Caltech-101 by comparing $Q_{BC} \circ Q_{AB}$ with Q_{AC} . (a) Image-image class retrieval. (b) text-text class retrieval (c) Mean image-image cosine similarity. (c) Mean text-text cosine similarity. (e) Image-text retrieval using aligned images from model A and text from model B. (f) Image-text retrieval using images from model B and aligned text from model A. (g) Image-text retrieval using aligned images and aligned text from model A.

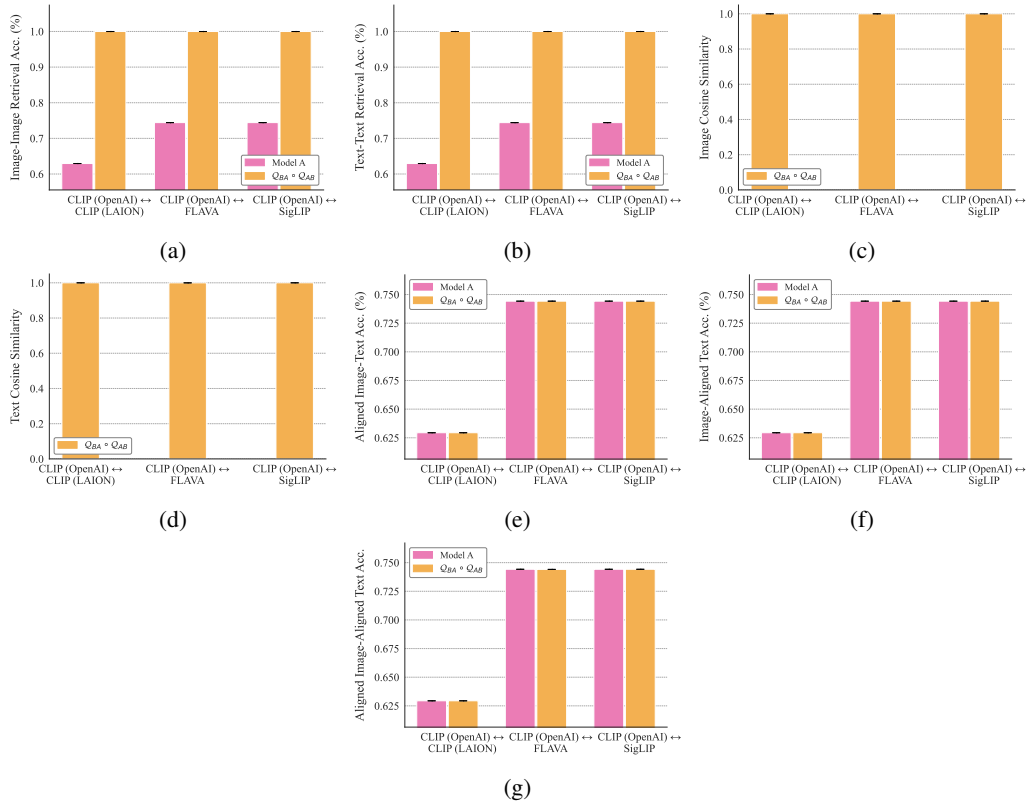


Figure 37: Cycle consistency on CIFAR-100. (a) Image–image class retrieval. (b) text–text class retrieval (c) Mean image-image cosine similarity. (c) Mean text-text cosine similarity. (e) Image–text retrieval using aligned images from model A and text from model B. (f) Image–text retrieval using images from model B and aligned text from model A. (g) Image–text retrieval using aligned images and aligned text from model A.

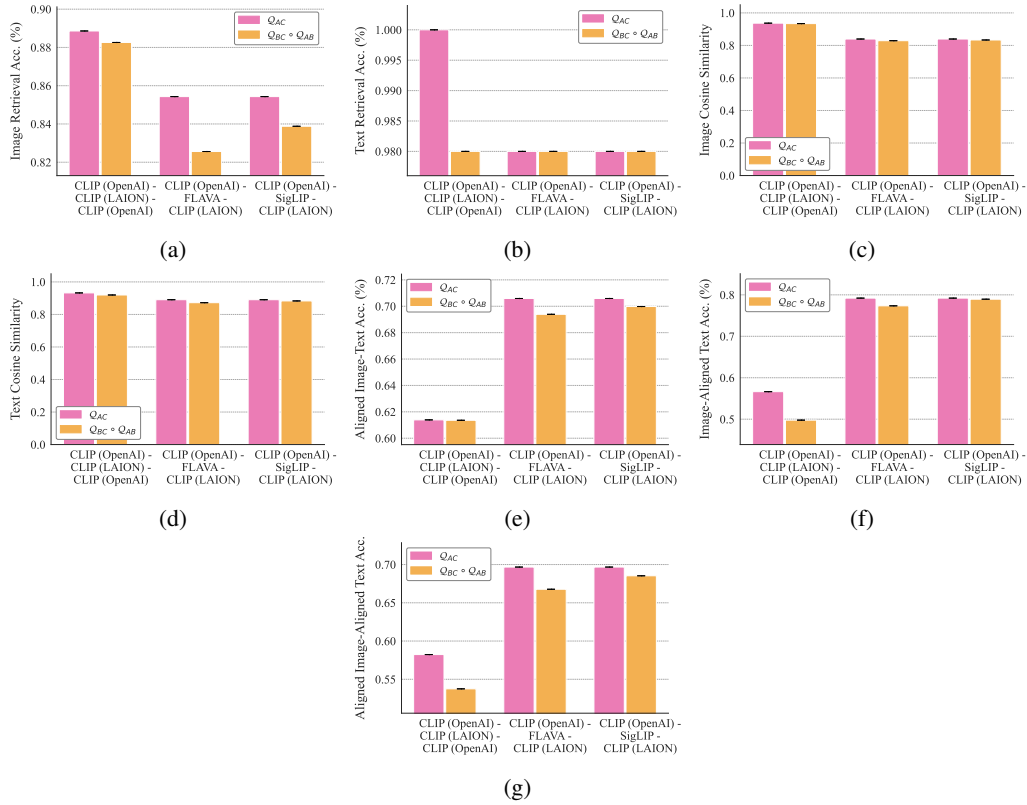


Figure 38: Compositionality on CIFAR-100 by comparing $Q_{BC} \circ Q_{AB}$ with Q_{AC} . (a) Image–image class retrieval. (b) text–text class retrieval (c) Mean image–image cosine similarity. (c) Mean text–text cosine similarity. (e) Image–text retrieval using aligned images from model A and text from model B. (f) Image–text retrieval using images from model B and aligned text from model A. (g) Image–text retrieval using aligned images and aligned text from model A.

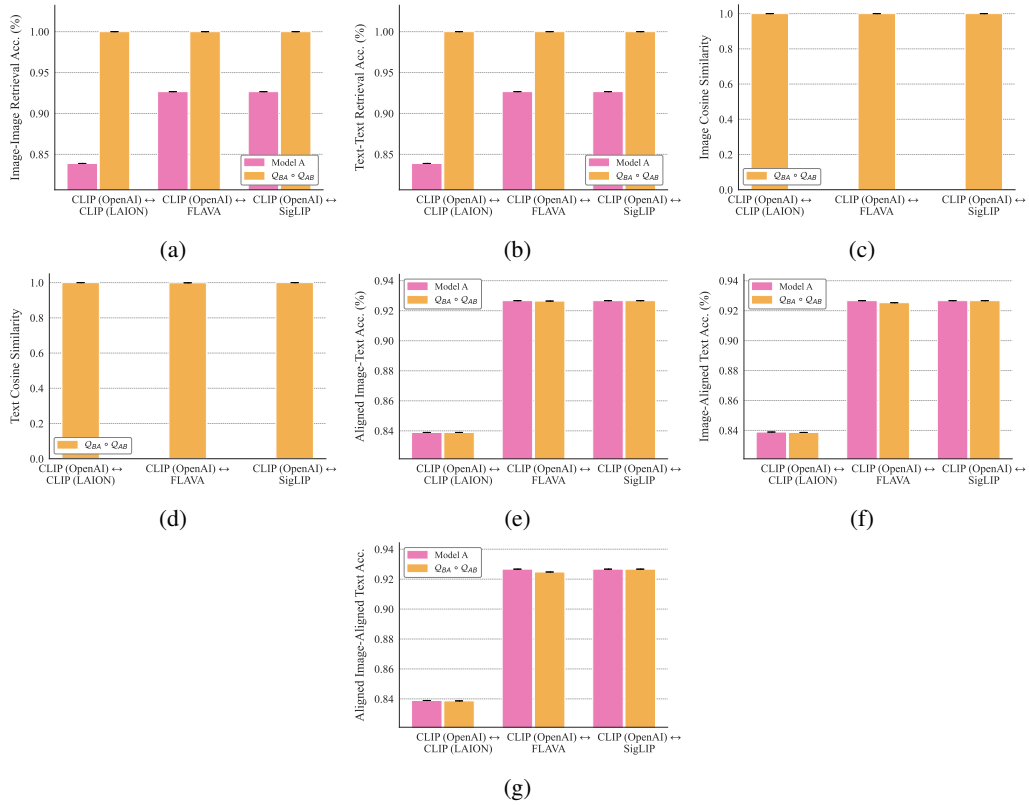


Figure 39: *Cycle consistency on Oxford Pets.* (a) Image–image class retrieval. (b) text–text class retrieval (c) Mean image-image cosine similarity. (c) Mean text-text cosine similarity. (e) Image–text retrieval using aligned images from model A and text from model B. (f) Image–text retrieval using images from model B and aligned text from model A. (g) Image–text retrieval using aligned images and aligned text from model A.

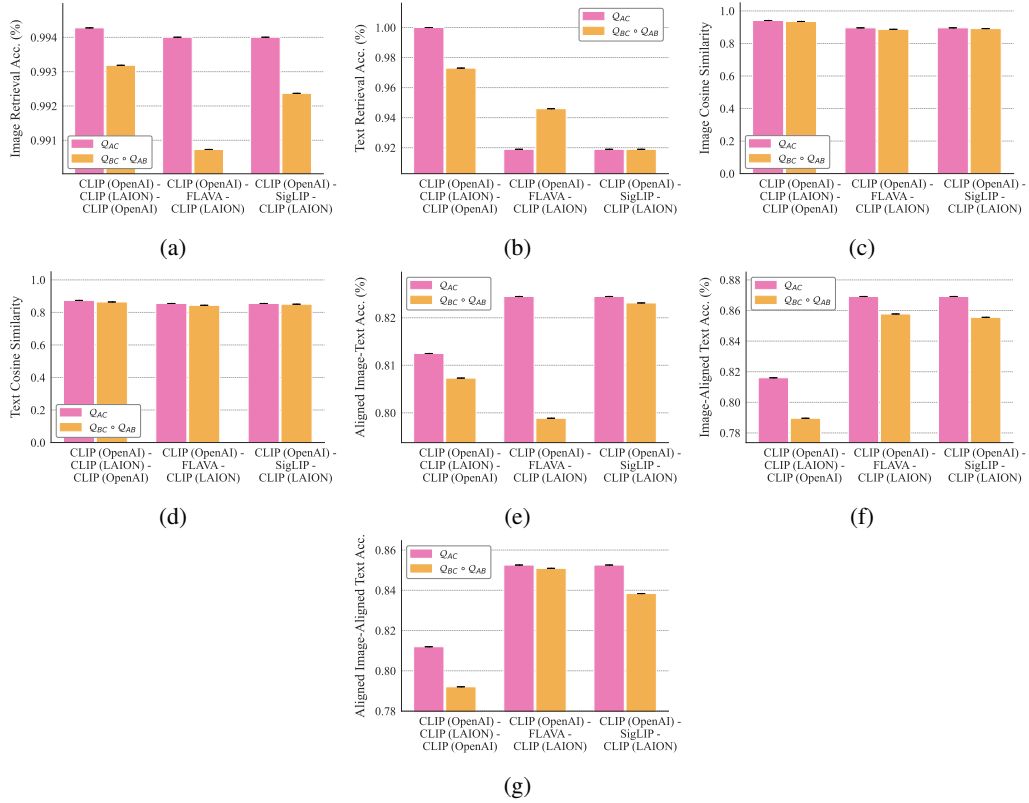


Figure 40: Compositionality on Oxford Pets by comparing $Q_{BC} \circ Q_{AB}$ with Q_{AC} . (a) Image-image class retrieval. (b) text-text class retrieval (c) Mean image-image cosine similarity. (c) Mean text-text cosine similarity. (e) Image-text retrieval using aligned images from model A and text from model B. (f) Image-text retrieval using images from model B and aligned text from model A. (g) Image-text retrieval using aligned images and aligned text from model A.

E.8 ALIGNMENT ACROSS EMBEDDING DIMENSIONS.

All results so far use matched embedding sizes, where an orthogonal map is well-defined. We next align CLIP ViT-B/32 ($d=512$) with CLIP ViT-L/14 ($d=768$) using a rectangular Q projected onto the Stiefel manifold, enforcing $Q^T Q \approx I$. In Figures 41 to 43, we report performance before and after learning Q on image embeddings, and evaluate how effectively the learned map transfers to texts. Across all figures, image-text accuracy remains high in the aligned space, whether using aligned images with native text, native images with aligned text, or aligning both. Text-text pointwise cosine and retrieval also remain strong, indicating that near-isometric maps preserve task-relevant geometry while enabling reliable transfer to the other modality (text in this case).

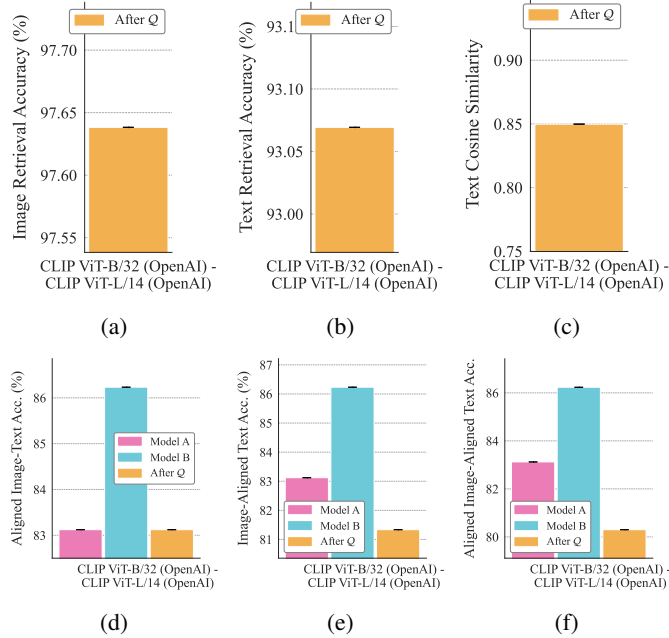


Figure 41: *Cross-model alignment across differing embedding dimensions (Caltech-101; CLIP ViT-B/32 to CLIP ViT-L/14, OpenAI) before and after fitting an orthogonal map Q .* (a) Image–image class retrieval and (b) text–text class retrieval. (c) Mean text–text cosine similarity. (d) Image–text retrieval using aligned images from model A and text from model B. (e) Image–text retrieval using images from model B and aligned text from model A. (f) Image–text retrieval using aligned images and aligned text from model A. Enforcing semi-orthogonality via Stiefel manifold projection allows the 512-dimensional source manifold to be isometrically embedded into the 768-dimensional target space, maintaining high cross-modal retrieval accuracy.

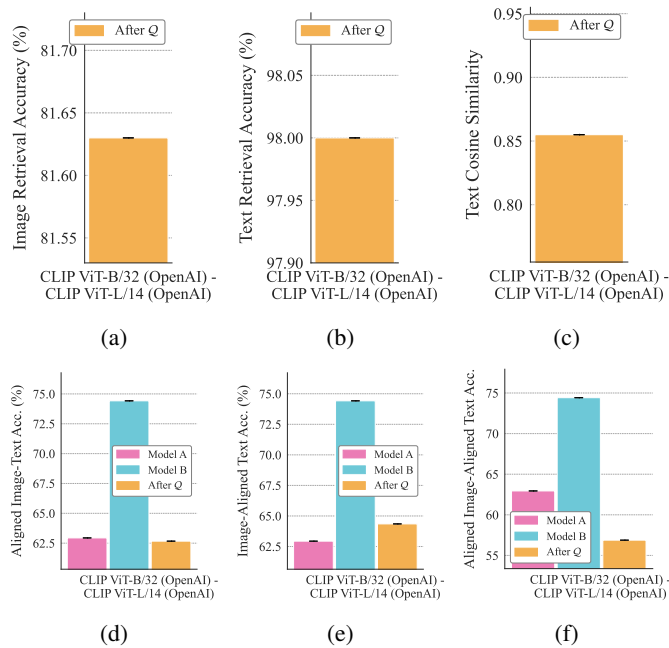


Figure 42: *Cross-model alignment across differing embedding dimensions (CIFAR-100; CLIP ViT-B/32 to CLIP ViT-L/14, OpenAI) before and after fitting an orthogonal map Q .* (a) Image–image class retrieval and (b) text–text class retrieval. (c) Mean text–text cosine similarity. (d) Image–text retrieval using aligned images from model A and text from model B. (e) Image–text retrieval using images from model B and aligned text from model A. (f) Image–text retrieval using aligned images and aligned text from model A. Enforcing semi-orthogonality via Stiefel manifold projection allows the 512-dimensional source manifold to be isometrically embedded into the 768-dimensional target space, maintaining high cross-modal retrieval accuracy.

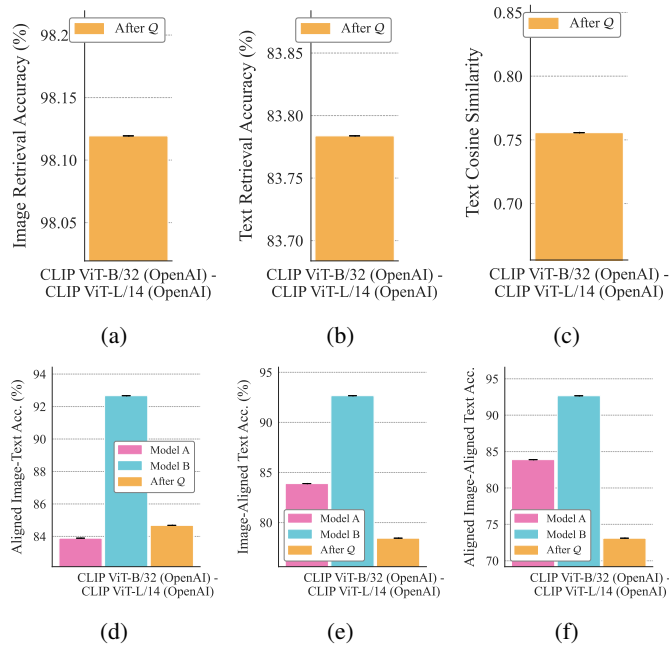


Figure 43: *Cross-model alignment across differing embedding dimensions (Oxford Pets; CLIP ViT-B/32 to CLIP ViT-L/14, OpenAI) before and after fitting an orthogonal map Q . (a) Image-image class retrieval and (b) text-text class retrieval. (c) Mean text-text cosine similarity. (d) Image-text retrieval using aligned images from model B and text from model A. (e) Image-text retrieval using aligned images from model B and aligned text from model A. (f) Image-text retrieval using aligned images and aligned text from model A. Enforcing semi-orthogonality via Stiefel manifold projection allows the 512-dimensional source manifold to be isometrically embedded into the 768-dimensional target space, maintaining high cross-modal retrieval accuracy.*

E.9 EVALUATING ALTERNATIVE ALIGNMENT MAPS THAN THE ORTHOGONAL MAPPING

In this section, we ablate the design choices for the alignment function and test three alignment strategies of increasing expressiveness: (i) an orthogonal Procrustes map Q , (ii) an unconstrained linear map, and (iii) a small MLP with a residual connection. Results across all model pairs are reported for Caltech-101 (Figure 44), CIFAR-100 (Figure 45), and Oxford Pets (Figure 46). We find that more expressive maps achieve better *pointwise* alignment: both the unconstrained linear map and the MLP yield higher text-text pointwise cosine similarity, reflecting their ability to approximate small distortions and scaling. However, this flexibility comes at the cost of geometric fidelity. The orthogonal map consistently outperforms the alternatives on geometry-sensitive metrics—most notably image to aligned text and aligned-image to aligned-text classification accuracy—which are the metrics that matter for downstream use. By contrast, text-text and image-image retrieval are comparable across all three approaches, with no clear winner. These results indicate that while flexible maps can improve local fit, preserving global geometry through orthogonality can be critical for reliable cross-modal transfer.

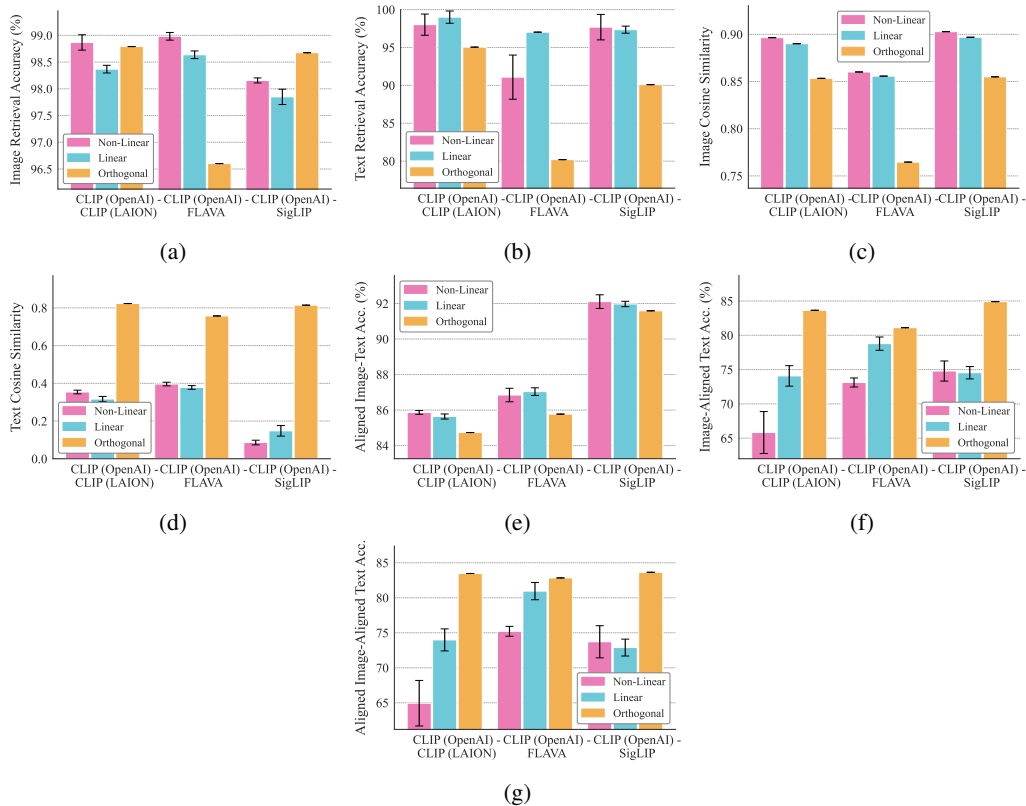


Figure 44: Comparison of alignment strategies (Non-Linear, Linear, and Orthogonal) on Caltech-101 across model pairs. (a) Image–image class retrieval and (b) text–text class retrieval. (c) Mean image–image cosine similarity. (d) Mean text–text cosine similarity. (e) Image–text retrieval using aligned images from model A and text from model B. (f) Image–text retrieval using images from model B and aligned text from model A. (g) Image–text retrieval using aligned images and aligned text from model A. Unlike linear and non-linear maps, orthogonal transformations preserve the underlying image-text geometry favorable for downstream tasks and transfers across modalities (text in this case).

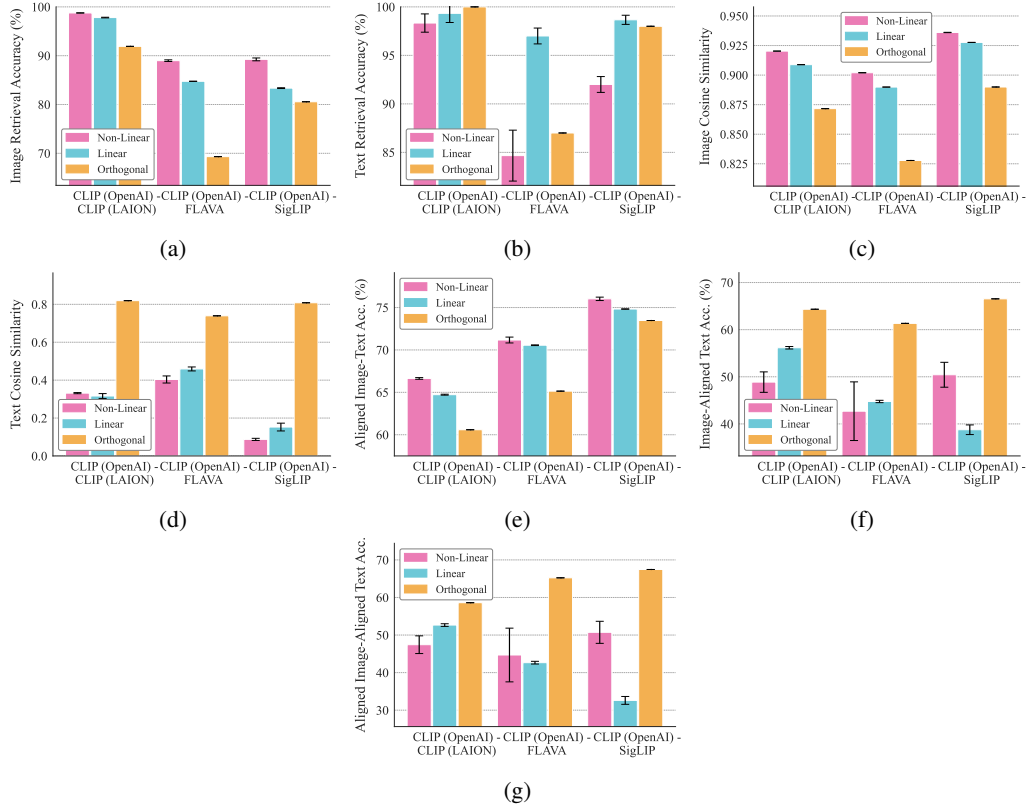


Figure 45: Comparison of alignment strategies (Non-Linear, Linear, and Orthogonal) on CIFAR-100 across model pairs. (a) Image–image class retrieval and (b) text–text class retrieval. (c) Mean image–image cosine similarity. (d) Mean text–text cosine similarity. (e) Image–text retrieval using aligned images from model A and text from model B. (f) Image–text retrieval using images from model B and aligned text from model A. (g) Image–text retrieval using aligned images and aligned text from model A. Unlike linear and non-linear maps, orthogonal transformations preserve the underlying image-text geometry favorable for downstream tasks and transfers across modalities (text in this case).

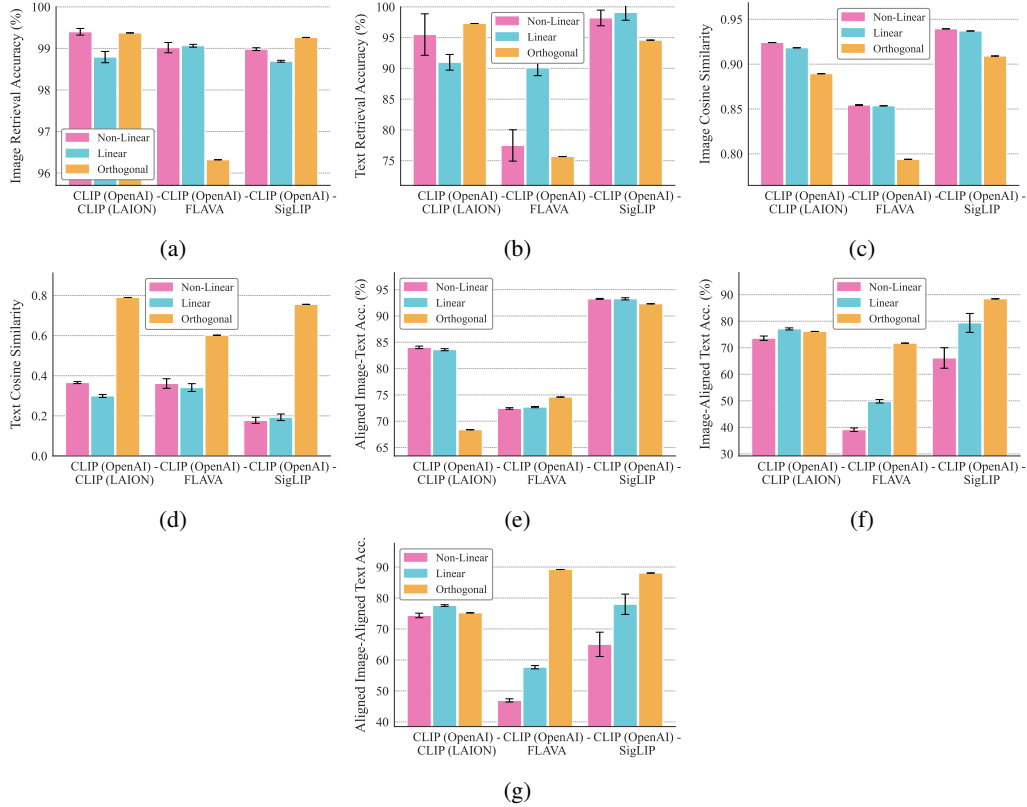


Figure 46: Comparison of alignment strategies (Non-Linear, Linear, and Orthogonal) on Oxford Pets across model pairs (CLIP OpenAI vs. LAION, FLAVA, and SigLIP). (a) Image–image class retrieval and (b) text–text class retrieval. (c) Mean image–image cosine similarity. (d) Mean text–text cosine similarity. (e) Image–text retrieval using aligned images from model A and text from model B. (f) Image–text retrieval using images from model B and aligned text from model A. (g) Image–text retrieval using aligned images and aligned text from model A. Unlike linear and non-linear maps, orthogonal transformations preserve the underlying image–text geometry favorable for downstream tasks and transfers across modalities (text in this case).

E.10 ORTHOGONAL ALIGNMENT WITH AND WITHOUT CENTERING

All results thus far use our *centered* orthogonal Procrustes map (i.e., we fit Q on mean-centered embeddings and re-center at deployment). Here, we ablate this centering step by comparing it to a pure orthogonal map applied directly to raw embeddings. Specifically, we evaluate: (i) *Orthogonal (no centering)*: $z \mapsto Qz$, a pure change of coordinates that preserves inner products exactly; and (ii) *Orthogonal (with centering)*: $z \mapsto Q(z - \mu_A^{(\cdot)}) + \mu_B^{(\cdot)}$, where $\mu_A^{(\cdot)}$ and $\mu_B^{(\cdot)}$ are modality-specific means of the source and target models (image or text, as appropriate). Throughout, we use “ Q ” to denote the centered variant, consistent with the rest of the paper.

Results on Caltech-101, CIFAR-100 and Oxford Pets are presented in Figures 47 to 49 respectively. Across all figures, the orthogonal map with mean centering (denoted by Q) improves downstream metrics such as aligned-image-to-text and image-to-aligned-text retrieval, suggesting that centering helps correct first-order inter-model offsets in these comparisons. In contrast, aligned-image-to-aligned-text accuracy is typically higher without centering, since orthogonal w/o centering exactly preserves model A’s geometry. That said, these gains are modest, and the overall retrieval performance remains largely similar with or without centering. Finally, pointwise image-image and text-text cosine similarities increase significantly under centering, as subtracting and re-adding the global centroid shift removes a global centroid shift that otherwise depresses cosine agreement.

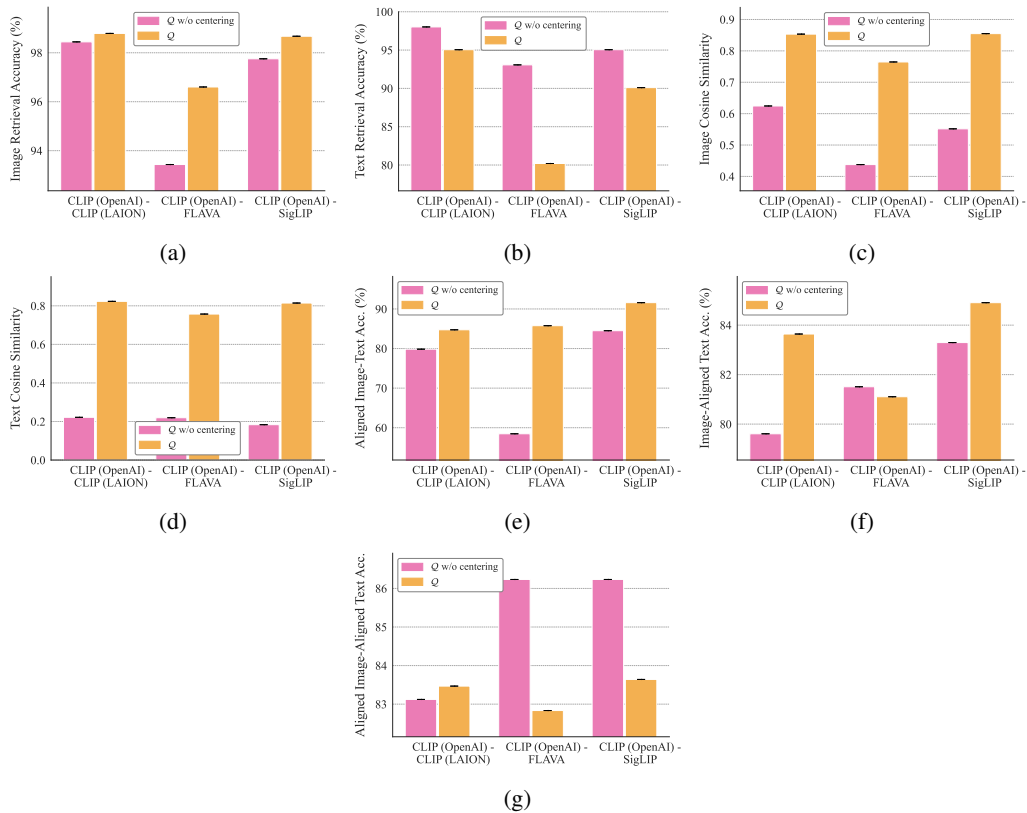


Figure 47: Comparison between an orthogonal map with and without mean adjustment across model pairs on Caltech-101. (a) Image–image class retrieval (b) text–text class retrieval (c) Mean image–image cosine similarity (d) Mean text–text cosine similarity. (e) Image–text retrieval using aligned images from model A and text from model B. (f) Image–text retrieval using images from model B and aligned text from model A. (g) Image–text retrieval using aligned images and aligned text from model A. Mean adjustment boosts pointwise image-image (c) and text-text (d) cosine and downstream performance (e,f), but can slightly underperform on aligned-image-to-aligned-text and text-text retrieval (b,g).

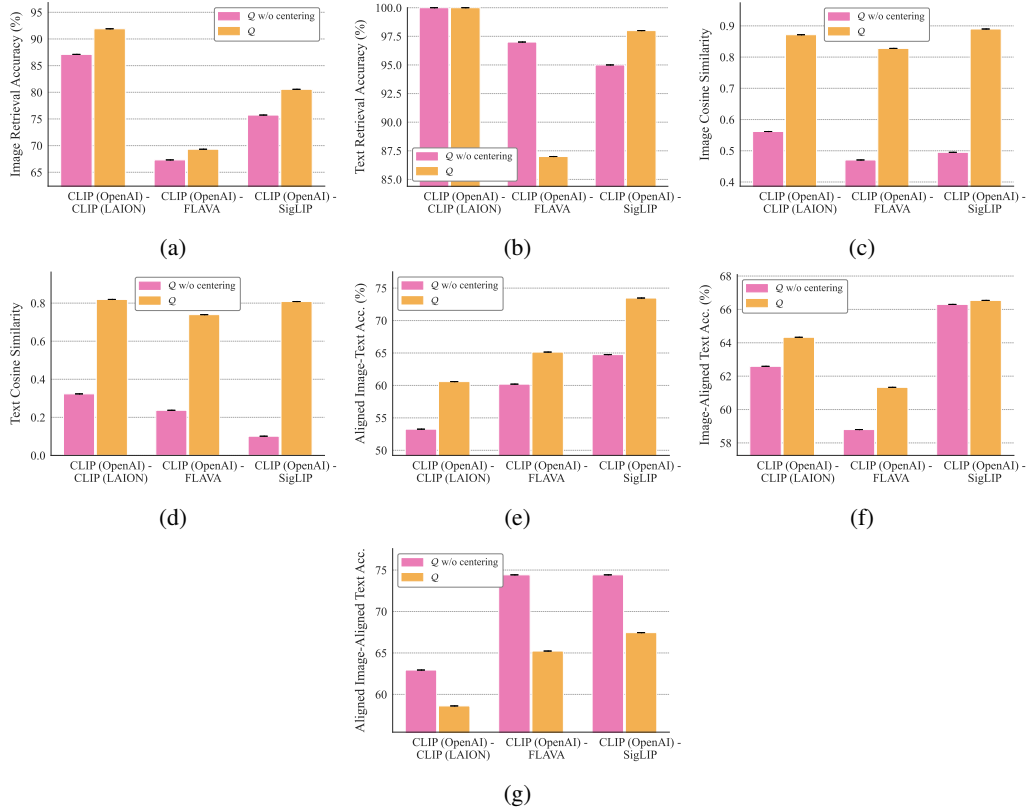


Figure 48: Comparison between an orthogonal map with and without mean adjustment across model pairs on CIFAR-100. (a) Image–image class retrieval (b) text–text class retrieval (c) Mean image–image cosine similarity (d) Mean text–text cosine similarity. (e) Image–text retrieval using aligned images from model A and text from model B. (f) Image–text retrieval using images from model B and aligned text from model A. (g) Image–text retrieval using aligned images and aligned text from model A. Mean adjustment boosts pointwise image-image (c) and text-text (d) cosine and downstream performance (e,f), but can slightly underperform on aligned-image-to-aligned-text and text-text retrieval (b,g).

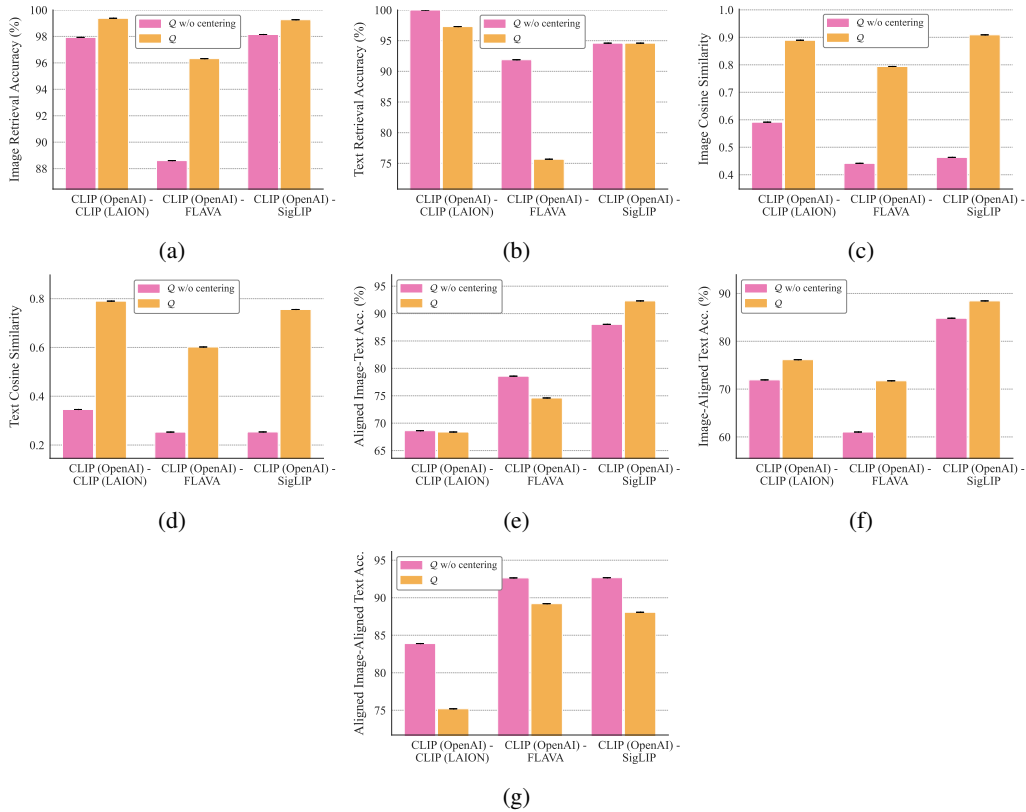


Figure 49: Comparison between an orthogonal map with and without mean adjustment across model pairs on Oxford Pets. (a) Image–image class retrieval (b) text–text class retrieval (c) Mean image–image cosine similarity (d) Mean text–text cosine similarity. (e) Image–text retrieval using aligned images from model A and text from model B. (f) Image–text retrieval using images from model B and aligned text from model A. (g) Image–text retrieval using aligned images and aligned text from model A. Mean adjustment boosts pointwise image–image (c) and text–text (d) cosine and downstream performance (e,f), but can slightly underperform on aligned-image-to-aligned-text and text–text retrieval (b,g).

E.11 ISOLATING ARCHITECTURE EFFECTS UNDER FIXED TRAINING DATA

All previous experiments compare model pairs that differ in both training distribution and design choices. Here we isolate the role of the data by aligning two CLIP models trained on (probably) the same distribution (dataset used by OpenAI (Radford et al., 2021)) but with different architectures (ViT-B/32 vs. ViT-B/16). Results on Caltech-101, CIFAR-100 and Oxford Pets are reported in Figures 50 to 52. As observed across these figures, in this controlled setting, transfer is consistently stronger, compared to when using models trained on different distributions. This aligns with our theoretical results showing that when two contrastive models optimize the same underlying P_{XY} , their induced multimodal kernels agree up to an additive constant at the population optimum Corollary C.6, tightening the conditions under which a single orthogonal map aligns both modalities.

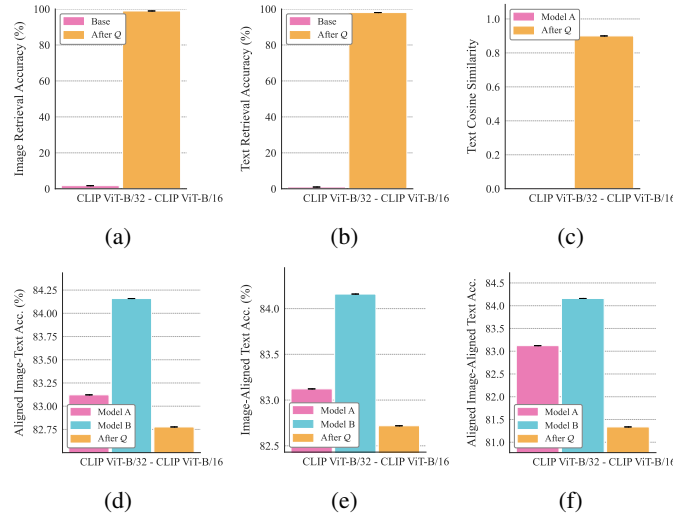


Figure 50: Cross-model alignment between identical training datasets with differing architectures (Caltech-101; CLIP ViT-B/32 to ViT-B/16, OpenAI) before and after fitting a single orthogonal map \mathcal{Q} . (a) Image–image class retrieval and (b) text–text class retrieval. (c) Mean text–text cosine similarity. (d) Image–text retrieval using aligned images from model A and text from model B. (e) Image–text retrieval using images from model B and aligned text from model A. (f) Image–text retrieval using aligned images and aligned text from model A. When both models are trained on the same data distribution and differ only in architecture, applying \mathcal{Q} yields substantially stronger modality transfer than in settings where the training distributions differ.




Figure 51: *Cross-model alignment between identical training datasets with differing architectures (CIFAR-100; CLIP ViT-B/32 to ViT-B/16, OpenAI) before and after fitting a single orthogonal map \mathcal{Q} . (a) Image–image class retrieval and (b) text–text class retrieval. (c) Mean text–text cosine similarity. (d) Image–text retrieval using aligned images from model A and text from model B. (e) Image–text retrieval using images from model B and aligned text from model A. (f) Image–text retrieval using aligned images and aligned text from model A. When both models are trained on the same data distribution and differ only in architecture, applying \mathcal{Q} yields substantially stronger modality transfer than in settings where the training distributions differ.*




Figure 52: *Cross-model alignment between identical training datasets with differing architectures (Oxford Pets; CLIP ViT-B/32 to ViT-B/16, OpenAI) before and after fitting a single orthogonal map \mathcal{Q} . (a) Image–image class retrieval and (b) text–text class retrieval. (c) Mean text–text cosine similarity. (d) Image–text retrieval using aligned images from model A and text from model B. (e) Image–text retrieval using images from model B and aligned text from model A. (f) Image–text retrieval using aligned images and aligned text from model A. When both models are trained on the same data distribution and differ only in architecture, applying \mathcal{Q} yields substantially stronger modality transfer than in settings where the training distributions differ.*

E.12 QUALITATIVE EVIDENCE OF COMMUTATIVITY OF IMAGE-TEXT ALIGNMENT PATHS

In this section, we qualitatively test whether the learned map Q induces a consistent, modality-invariant geometry by comparing two alignment paths from a source image x . In the *direct* path, we map the image embedding via Q and retrieve its top-5 nearest images in the target model’s space. In the *text-mediated* path, we first retrieve the nearest text to x in the source model, map this text embedding via Q , retrieve its top-1 nearest text in the target model, and then retrieve the top-5 images associated with that text.

As shown in Figure 53 and Figure 54, both paths yield highly consistent semantic neighborhoods: images retrieved via the text-mediated route closely match those obtained by direct image alignment. Equivalently, in terms of the commuting diagram induced by cross-modal nearest-neighbor operators, the two routes from x —direct image transport and transport via a retrieved caption—lead to nearly the same target-space neighborhood, suggesting that Q approximately commutes with these retrieval operators on this domain.



Figure 53: *Qualitative k-NN retrieval under two alignment paths for Oxford Pets.* We compare two routes from a source image x into the target model’s image space after alignment by Q : a direct image-alignment path (top row) and a text-mediated path (bottom row). Both routes yield highly consistent target-space neighborhoods, showing that the two alignment paths approximately close as a commuting diagram.




Figure 54: *Qualitative k-NN retrieval under two alignment paths for Caltech-101.* We compare two routes from a source image x into the target model’s image space after alignment by Q : a direct image-alignment path (top row) and a text-mediated path (bottom row). Both routes yield highly consistent target-space neighborhoods, showing that the two alignment paths approximately close as a commuting diagram.

E.13 FINE-GRAINED SEMANTIC PRESERVATION UNDER ORTHOGONAL MAPS.

We further probe whether the learned map Q preserves semantic structure beyond coarse class labels. Starting from a source image x , we retrieve its nearest caption in Model A’s text space and then construct a minimally perturbed variant by changing only the class token (e.g., *Birman* to *Russian Blue*) while keeping the remaining description fixed. We map this modified text via Q , retrieve its nearest text in Model B’s space, and finally retrieve the nearest image in Model B’s image space.

As shown in Figure 55, the retrieved images not only reflect the edited class (Russian Blue) but also somewhat preserve finer visual attributes implied by the original description. In particular, different captions lead to distinct Russian Blue images that emphasize different features (e.g., eye prominence

vs. body pose), despite sharing the same class label. Similarly, in Figure 56, the retrieved Basset Hound images emphasize body mass and posture consistent with the original Bulldog description, rather than collapsing to a generic class prototype. Together with the previous example, this indicates that alignment by \mathcal{Q} preserves fine-grained semantic structure encoded in text and image embeddings, enabling controlled semantic edits to transfer predictably across models.

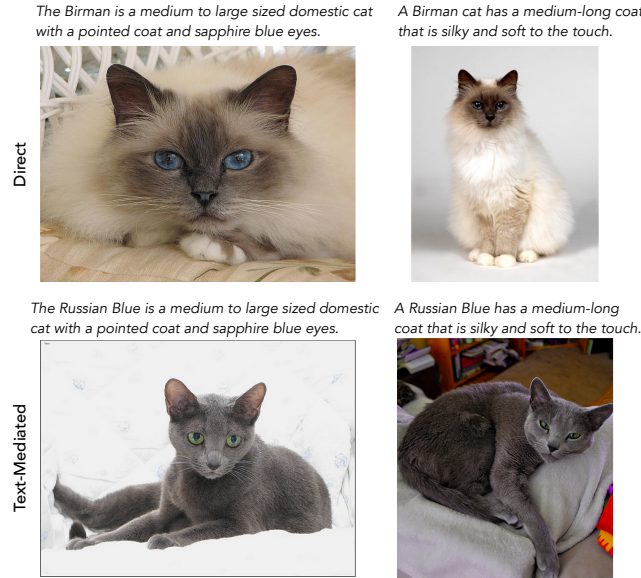


Figure 55: *Fine-grained semantic preservation under controlled text edits on Oxford Pets.* After changing only the class token in the caption (Birman to Russian Blue) while keeping descriptive attributes fixed, alignment via \mathcal{Q} retrieves target-model images that match the edited class while preserving caption-specific visual cues (e.g., eye prominence and facial structure), indicating preservation of fine-grained semantics beyond class identity.



Figure 56: *Fine-grained semantic preservation under controlled text edits on Oxford Pets.* After changing only the class token in the caption (American Bulldog to Basset Hound) while keeping descriptive attributes fixed, alignment via \mathcal{Q} retrieves target-model images that match both the edited class and the original attribute cues (e.g., muscular build and posture), indicating preservation of fine-grained semantics beyond class identity.