Do Vision and Language Encoders Represent the World Similarly?

Mayug Maniparambil* [1]   Raiymbek Akshulakov* [2]   Yasser Abdelaziz Dahou Djilali[3] [1]
Mohamed El Amine Seddik[3]   Sanath Narayan[3]   Karttikeya Mangalam[2]   Noel E. O'Connor[1]

Abstract

Aligned text-image encoders such as CLIP have become the de-facto model for vision-language tasks. Furthermore, modality-specific encoders achieve impressive performances in their respective domains. This raises a central question: does an alignment exist between uni-modal vision and language encoders since they fundamentally represent the same physical world? Analyzing the latent spaces structure of vision and language models on image-caption benchmarks using the Centered Kernel Alignment (CKA), we find that the representation spaces of unaligned and aligned encoders are semantically similar. In the absence of statistical similarity in aligned encoders like CLIP, we show that a possible matching of unaligned encoders exists without any training. We frame this as a seeded graph-matching problem exploiting the semantic similarity between graphs and propose two methods - a Fast Quadratic Assignment Problem optimization, and a novel localized CKA metric-based matching/retrieval. We demonstrate the effectiveness of this on several downstream tasks including cross-lingual, cross-domain caption matching and image classification. Code available at github.com/mayug/0-shot-llm-vision.

1. Introduction

The recent success of deep learning on vision-language tasks mainly relies on jointly trained language and image encoders following the success of CLIP and ALIGN [20, 40]. The standard procedure for training these models aims at aligning text and image representation using a contrastive loss that maximizes the similarity between image-text pairs while pushing negative captions away [10, 19, 36]. This achieves a statistical similarity across the two latent spaces, which is key to retrieving the closest cross-modal representations using cosine similarity. This property is not valid for unaligned encoders, hence, extra transformations are needed to bridge the gap. These transformations can be training a mapping network that captures the prior distribution over the text and image representations [31, 34, 35]. The work of [31] has shown that it is possible to train a linear mapping from the output embeddings of vision encoders to the input embeddings of language models and exhibit impressive performance on image captioning and VQA tasks. This indicates that the representations between the unaligned uni-modal vision and language encoders are sufficiently high level and differ only by a linear transformation. However, this linear layer is trained on CC-3M [9] consisting of three million image-caption pairs.

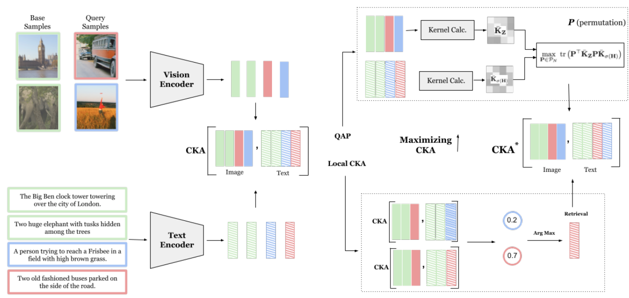

Figure 1. For matching, we calculate the kernels for image and text embeddings and employ QAP-based seeded matching to maximize CKA for obtaining the optimal permutation $\boldsymbol{P}$. For retrieval, we append query embeddings to base embeddings and retrieve the best caption that maximizes the local CKA for a query image.

Is this training step necessary? In an ideal scenario, we anticipate an alignment between vision and language encoders as they inherently capture representations of the same physical world. To this end, we employ Centered Kernel Alignment (CKA) [12, 22, 42], which is known for measuring representation similarity both within and between networks. As shown in Figure 2, we measure the CKA between a variety of unaligned vision and language encoders [8, 16, 28, 37, 47], on the image-caption pairs of the COCO [27] dataset and observe that some have comparable scores to that of aligned encoders like CLIP [40], affirmative of semantic similarities.

We then ask the question: If the unaligned image and text encoders are semantically similar, is there a way to connect them in a *zero-shot manner?* Do they build a similar representation graph over the same information coming from

*joint first authors
[1]ML Labs, Dublin City University
[2]University of California Berkeley
[3]Technological Innovation Institute

the two modalities? We study these questions, revealing key similarities between unaligned image and text encoders, and how these similarities can be exploited for downstream tasks. Furthermore, we devise a caption matching downstream task and show using two novel methods that latent space communication between unaligned encoders could be achieved by leveraging the semantic similarities between the cross-modal spaces. Our contributions are:

- We present a matching method that seeks to find the permutation of the captions that maximizes the CKA (see Fig. 1). Hence, We formulate maximizing CKA as a quadratic assignment problem and introduce transformations and normalizations that greatly improve the matching performance.
- We propose a local CKA metric and use it to perform retrieval between two unaligned embedding spaces, demonstrating superior performance with that of relative representations [34] on the COCO caption image retrieval.
- The method is benchmarked on COCO, NoCaps [2] cross-domain caption and image retrieval as well ImageNet-100 [15] classification tasks despite our method not being optimized to align the representation in any manner demonstrating zero-shot communication between the encoder's latent spaces.
- Finally, we show a practical application of our method on cross-lingual image retrieval by making use of sentence transformers trained in various languages and a CLIP vision encoder trained only in English.

2. Related Work

Recently, there has been an increasing consensus that good networks, when trained independently, learn general representations across different architectures and tasks. On the one hand, the works of [6, 22, 26, 33] show that these networks exhibit representation similarity by learning similar latent spaces when trained on similar tasks and data [3, 5, 11, 24, 32, 44, 46]. Specifically, [22] introduced centered kernel alignment (CKA) as a similarity metric for comparing the inner representations across networks. The CKA measure mitigates the limitation of canonical correlation analysis (CCA) [41] being invariant to an invertible linear transformation that often leads to difficulty in measuring meaningful similarities between representations. [48] uses CKA for comparing the representations from different layers of different language models and the effect of downstream task-finetuning on the representation similarities, while [6] utilizes CKA along with Procrustes similarity for understanding the ability of variational autoencoders (VAEs) [21] in learning disentangled representations. In general, these approaches study the representation similarity in unimodal models, either vision or language. Clearly, however, the use of CKA has been limited to visualization and analysis purposes, whereas we attempt at exploiting CKA as an optimization objective.

Recent works [34, 35] employ relative representations to match embeddings of unaligned encoders using the cosine similarity to a set of anchors. However, these relative representations are sensitive to the selection of anchors and noise in the original embeddings. Similarly, approaches [4, 14] analyze networks and empirically verify the "good networks learn similar representations" hypothesis by utilizing model stitching [24], which introduces trainable stitching layers to enable swapping parts of different networks. LiMBeR [31] can be seen as stitching the output of an image encoder to the input of a language model in the form of soft prompts [25]. However, these approaches involve training of stitching layers for evaluating the representation similarity between two models.

In this work, we argue that using an explicit similarity measure as done in [34, 35] is sensitive to the selection of anchors and noise in the original embeddings. One design choice is an implicit measure that captures the similarity of similarities, hence, inducing more robustness to the alignment process. Furthermore, we explore how this similarity can be leveraged for downstream cross-modal tasks in a *training-free* manner with the aid of CKA and a set of parallel anchors in the image and text latent embedding spaces.

3. Preliminaries

Centered Kernel Alignment (CKA) has shown its relevance in understanding and comparing the information encoded by different layers of a neural network [22]. Formally, CKA relies on two sets of data $\mathbf{X} \in \mathbb{R}^{p \times N}$ and $\mathbf{Y} \in \mathbb{R}^{q \times N}$ through their corresponding kernels $\mathbf{K} = k(\mathbf{X}^\top, \mathbf{X}) \in \mathbb{R}^{N \times N}$ and $\mathbf{L} = \ell(\mathbf{Y}^\top, \mathbf{Y}) \in \mathbb{R}^{N \times N}$ where k, ℓ are some kernel functions applied on the columns of $\mathbf{X}$ and $\mathbf{Y}$ respectively (e.g., linear or RBF kernels). Therefore, the CKA is computed in terms of $\mathbf{K}$ and $\mathbf{L}$ as:

$$\mathrm{CKA}(\mathbf{K}, \mathbf{L}) = \frac{\mathrm{HSIC}(\mathbf{K}, \mathbf{L})}{\sqrt{\mathrm{HSIC}(\mathbf{K}, \mathbf{K})\,\mathrm{HSIC}(\mathbf{L}, \mathbf{L})}}, \tag{1}$$

where $\mathrm{HSIC}(\cdot, \cdot)$ is the Hilbert-Schmidt Independence Criterion [18, 30] defined as:

$$\mathrm{HSIC}(\mathbf{K}, \mathbf{L}) = \frac{1}{(N-1)^2} \operatorname{tr}(\mathbf{KCLC}), \tag{2}$$

with $\mathbf{C} = \mathbf{I} - \frac{1}{N}\mathbf{1}\mathbf{1}^\top$ the centring matrix. We refer the reader to [22] for broader properties and studies of the CKA metric on neural network representations.

4. Proposed Method

Consider a set of N image-caption pairs, $\mathcal{S} = \{(\boldsymbol{x}_i, \boldsymbol{c}_i)\}_{i=1}^N$, where $\boldsymbol{x}_i \in \mathcal{X}$ and $\mathbf{c}_i \in \mathcal{C}$ represent the i-th image and its corresponding caption, respectively. In this

Table 1. **CKA reduces with shuffling.** We measure the CKA score between DINOv2 [37] and All-Roberta-large-v1 [28] on the 5k COCO [27] image-caption representations pairs of the valset. The exact ordering yields the best score, whereas randomly shuffling the representations reduces the CKA score.

Shuffling (%)	0	20	40	60	80	100
CKA Score	0.72	0.46	0.27	0.13	0.04	0.01

particular example, we are performing caption-to-image retrieval, but it is applicable for the reverse as well. Let $\boldsymbol{f} : \mathcal{X} \mapsto \mathbb{R}^{d_1}$ and $\boldsymbol{g} : \mathcal{C} \mapsto \mathbb{R}^{d_2}$ denote some vision and language encoders respectively. The image-caption pairs are mapped into their corresponding sets of representations $\mathbf{Z} = [\boldsymbol{z}_1, \dots, \boldsymbol{z}_N] \in \mathbb{R}^{d_1 \times N}$ and $\mathbf{H} = [\boldsymbol{h}_1, \dots, \boldsymbol{h}_N] \in \mathbb{R}^{d_2 \times N}$, where $\boldsymbol{z}_i = \boldsymbol{f}(\boldsymbol{x}_i)$ and $\boldsymbol{h}_i = \boldsymbol{g}(\boldsymbol{c}_i)$.

As shown in Table 1, the maximum CKA score is obtained on the ground-truth ordering of the representations $\mathrm{CKA}_{\text{max}} = \mathrm{CKA}(\mathbf{K_Z}, \mathbf{K_H})$, where $\mathbf{K_Z}$ and $\mathbf{K_H}$ are the kernels for the image and text representations, defined respectively as $\mathbf{K_Z} = k(\mathbf{Z}^\top, \mathbf{Z})$ and $\mathbf{K_H} = k(\mathbf{H}^\top, \mathbf{H})$. We find that the CKA is sensitive to the data ordering. Specifically, we shuffle x% of data to obtain wrong matches while keeping the remaining 100-x% aligned, measure the CKA on each new data set, and observe that it monotonically decreases with random shuffling. This motivates our methodology for finding an optimal permutation of the image data that maximizes the CKA.

Formally, let σ be some permutation of the set $\{1, \cdots, N\}$ and denote $\sigma(\mathbf{Z}) = [\boldsymbol{z}_{\sigma(1)}, \cdots, \boldsymbol{z}_{\sigma(N)}] \in \mathbb{R}^{d_1 \times N}$ the set of permuted image representations by σ. If σ is not identity, it disrupts the original ordering of the image representations leading to a lower CKA score as shown in Table 1. Therefore, our goal is to find a permutation σ^* that maximizes the CKA. Formally:

$$\sigma^* = \arg\max_{\sigma} \mathrm{CKA}(\mathbf{K}_{\sigma(\mathbf{Z})}, \mathbf{K_H}). \tag{3}$$

The solution to this problem seeks to realign the permuted set of images in a way that maximizes the CKA, potentially recovering the ground-truth pairing between images and their corresponding captions.

To solve the aforementioned optimization problem, we explore two main approaches (visualized in Fig. 1): the Quadratic Assignment Problem (QAP) algorithm and Local CKA-based retrieval and matching. The QAP algorithm provides a global matching solution, seeking the optimal permutation across the query set considered. On the other hand, Local CKA-based retrieval and matching focuses on aligning images and captions using a localized metric, facilitating retrieval on a more granular level. This approach is more suitable where a single query image is given for a set of captions or *vice versa*.

4.1. QAP Matching

For some random permutation σ, the optimization problem in Equation 3 can be reformulated as a quadratic optimization problem [45] which reads as:

$$\max_{\mathbf{P} \in \mathcal{P}_N} \operatorname{tr}\left(\mathbf{P}^\top \bar{\mathbf{K}}_{\sigma(\mathbf{Z})} \mathbf{P} \bar{\mathbf{K}}_{\mathbf{H}}\right), \tag{4}$$

where $\mathcal{P}_N$ is the set of all permutation matrices of size N and $\bar{\mathbf{K}} = \mathrm{HSIC}(\mathbf{K}, \mathbf{K})^{-\frac{1}{2}} \mathbf{KC}$ stands for the centered and re-scaled kernel. In principle, maximizing the above objective is a relaxation of a graph-matching problem. Moreover, finding a global maximum of Equation 4 is NP-hard due to the combinatorial nature of the problem and therefore optimizing it can lead to sub-optimal or approximate solutions.

To overcome the NP-hardness of QAP, in practice, we suppose that we have access to a base set $\mathcal{B} = \{(\boldsymbol{z}_i^b, \boldsymbol{h}_i^b)\}_{i=1}^M$ of image-caption representations pairs and solve an equivalent objective to Equation 4 only partially on some unmatched query set $\mathcal{Q} = \{\boldsymbol{z}_i^q\}_{i=1}^N \times \{\boldsymbol{h}_i^q\}_{i=1}^N$ using a seeded version of the fast QAP algorithm [17]. Formally, let $\mathbf{Z} = [\boldsymbol{z}_1^b, \cdots, \boldsymbol{z}_M^b, \boldsymbol{z}_1^q, \cdots, \boldsymbol{z}_N^q] \in \mathbb{R}^{d_1 \times (M+N)}$ and $\mathbf{H} = [\boldsymbol{h}_1^b, \cdots, \boldsymbol{h}_M^b, \boldsymbol{h}_1^q, \cdots, \boldsymbol{h}_N^q] \in \mathbb{R}^{d_2 \times (M+N)}$ be the matrix concatenating all base and query representations of images and captions respectively, and denote by $\bar{\mathbf{K}}_{\mathbf{Z}}, \bar{\mathbf{K}}_{\mathbf{H}} \in \mathbb{R}^{(M+N) \times (M+N)}$ the corresponding centered and re-scaled kernels. The partial matching for aligning the query samples is then performed by solving the following:

$$\max_{\mathbf{P} \in \mathcal{P}_N} \operatorname{tr}\left((\mathbf{I}_M \oplus \mathbf{P})^\top \bar{\mathbf{K}}_{\mathbf{Z}} (\mathbf{I}_M \oplus \mathbf{P}) \bar{\mathbf{K}}_{\mathbf{H}}\right), \tag{5}$$

where $\mathbf{I}_M \oplus \mathbf{P} \in \mathbb{R}^{(M+N) \times (M+N)}$ stands for the block-diagonal matrix having diagonal blocks $\mathbf{I}_M$ and $\mathbf{P}$.

4.2. Local CKA based Retrieval and Matching

The concept of a global CKA metric is extended to derive local similarity measures suitable for retrieval. This process begins with a base set $\mathcal{B} = \{(\boldsymbol{z}_i^b, \boldsymbol{h}_i^b)\}_{i=1}^M$ consisting of aligned pairs of images and captions representations. The objective is to facilitate caption-image retrieval/matching within an unaligned query set $\mathcal{Q} = \{\boldsymbol{z}_i^q\}_{i=1}^N \times \{\boldsymbol{h}_i^q\}_{i=1}^N$.

A local CKA score, denoted as $\mathrm{localCKA}(\boldsymbol{z}^q, \boldsymbol{h}^q)$ for a couple $(\boldsymbol{z}^q, \boldsymbol{h}^q) \in \mathcal{Q}$ is calculated by computing a global CKA score for the image-caption pairs in $\mathcal{B}$, augmented with the query pair $(\boldsymbol{z}^q, \boldsymbol{h}^q)$. The local CKA is computed as follows:

$$\mathrm{localCKA}(\boldsymbol{z}^q, \boldsymbol{h}^q) = \mathrm{CKA}(\mathbf{K}_{[\mathbf{Z}, \boldsymbol{z}^q]}, \mathbf{K}_{[\mathbf{H}, \boldsymbol{h}^q]}), \tag{6}$$

where $[\mathbf{M}, \boldsymbol{v}]$ denotes the concatenation of the matrix $\mathbf{M}$ and the vector $\boldsymbol{v}$ column-wise and $\mathbf{Z} = [\boldsymbol{z}_1^b, \cdots, \boldsymbol{z}_M^b] \in \mathbb{R}^{d_1 \times M}$ and $\mathbf{H} = [\boldsymbol{h}_1^b, \cdots, \boldsymbol{h}_M^b] \in \mathbb{R}^{d_2 \times M}$. In essence, a correctly matched image-caption pair in $\mathcal{Q}$ would exhibit a higher degree of alignment with the base set $\mathcal{B}$ in terms of the CKA score, resulting in an elevated localCKA score. This metric can be used to calculate a score between one source query and N target queries enabling effective retrieval. Furthermore, this framework allows for the use of linear sum assignment [23] for matching tasks.

4.3. Stretching and Clustering

The choice of base samples and the spread of the representations in each embedding space affect the performance of the QAP and Local CKA algorithms. To spread the representations out in each domain for matching, we introduce a stretching matrix that normalizes the features of each dimension by the variance calculated from the query and base sets. Given $\mathbf{X} = [\boldsymbol{x}_1, \cdots, \boldsymbol{x}_d]^\top \in \mathbb{R}^{d \times N}$, the stretched matrix $\mathbf{X}_s$ is computed as $\mathbf{X}_s = \mathbf{S}\mathbf{X}$, where the stretching matrix $\mathbf{S} \in \mathbb{R}^{d \times d}$ is a diagonal matrix with inverse empirical standard deviation of the feature dimension as entries, *i.e.*, $\mathbf{S} = \text{diag}\left(\frac{1}{\text{std}(\boldsymbol{x}_1)}, \cdots, \frac{1}{\text{std}(\boldsymbol{x}_d)}\right)$ and $\boldsymbol{x}_i \in \mathbb{R}^N$ is the i^{th} row of $\mathbf{X}$. This stretching operation is performed for both the image and text before calculating the kernels for both QAP and local CKA matching algorithms. For picking the most effective base samples, we find that the simple k-means clustering on the image embeddings works best. An ablation on how these affect the QAP and local CKA matching and retrieval accuracies is provided in Sec 7.

5. Experiments

We assess the performance of the proposed method using various vision and language encoders on a set of downstream tasks. We first detail the encoders, datasets, downstream tasks, and the baselines used.

5.1. Vision and Language Encoders

The experimental setup covers vision encoders of different architectures, such as ViTs [16] and ConvNeXt [29], trained in various ways: supervised, language-supervised, and self-supervised, across different training data regimes. For the language encoder, an encoder capable of producing a global embedding for a caption is essential. This includes encoders of multiple architectures varying in size, languages, and training data sizes. The Huggingface's sentence-transformers [43] library is utilized, where each sentence transformer is first pre-trained on the masked language modeling task using a large text corpus, followed by a finetuning stage on a sentence pairs dataset with a contrastive loss. It's not straightforward to acquire a global sentence embedding from decoder-only models like GPT models [7, 39], hence we did not study the semantic alignment of these class of models to vision encoders.

The CKA and Matching Score (MS) of the various combinations of vision and language encoders are reported in supplementary. The findings indicate that the All-Roberta-large-v1 [28] demonstrates the best CKA/MS across all vision models, establishing it as the primary language encoder for subsequent tasks, unless specified otherwise.

5.2. Baselines

Here, we briefly describe three baselines that we compare our methods against for caption matching/retrieval, image classification, and cross-lingual tasks.

Linear Regression: We propose a baseline that learns a linear transformation from the image embedding space to the text using M aligned base examples and apply the transformation to the query image embeddings. Concretely, given query image embeddings $\mathbf{Z}^q = [\boldsymbol{z}_1^q, \cdots, \boldsymbol{z}_N^q] \in \mathbb{R}^{d_1 \times N}$ and text embeddings $\mathbf{H}^q = [\boldsymbol{h}_1^q, \cdots, \boldsymbol{h}_N^q] \in \mathbb{R}^{d_2 \times N}$, and a set of aligned base samples $\mathbf{Z}^b = [\boldsymbol{z}_1^b, \cdots, \boldsymbol{z}_N^b] \in \mathbb{R}^{d_1 \times M}$ and $\mathbf{H}^b = [\boldsymbol{h}_1^b, \cdots, \boldsymbol{h}_N^b] \in \mathbb{R}^{d_2 \times M}$, we first construct a linear transformation between $\mathbf{Z}^b$ and $\mathbf{H}^b$ by minimizing the MSE loss as $\mathbf{W} = \arg\min_{\mathbf{W}} \|\mathbf{W}^\top \mathbf{Z}^b - \mathbf{H}^b\|_F^2$. Then we use $\mathbf{W}$ to transform the query image embeddings $\mathbf{Z}^q$ to the text domain as $\hat{\mathbf{H}}^q = \mathbf{W}^\top \mathbf{Z}^q$. Cosine similarity on $\hat{\mathbf{H}}^q$ and $\mathbf{H}^q$ can be used to perform caption retrieval.

Relative Representations [34]: enable latent space communication between unaligned encoders by representing each query point relative to an aligned base set. Concretely, let ℓ_2-normalized embeddings for image and text queries be $\mathbf{Z}^q = [\boldsymbol{z}_1^q, \cdots, \boldsymbol{z}_N^q] \in \mathbb{R}^{d_1 \times N}$ and $\mathbf{H} = [\boldsymbol{h}_1^q, \cdots, \boldsymbol{h}_N^q] \in \mathbb{R}^{d_2 \times N}$, respectively. Utilizing a set of aligned base sample ℓ_2-normalized embeddings $\mathbf{Z}^b = [\boldsymbol{z}_1^b, \cdots, \boldsymbol{z}_M^b] \in \mathbb{R}^{d_1 \times M}$ and $\mathbf{H}^b = [\boldsymbol{h}_1^b, \cdots, \boldsymbol{h}_M^b] \in \mathbb{R}^{d_2 \times M}$, we can construct relative image and text query representations as $\mathbf{Z}_{rel}^q = (\mathbf{Z}^b)^\top \mathbf{Z}^q$ and $\mathbf{H}_{rel}^q = (\mathbf{H}^b)^\top \mathbf{H}^q$. Relative representations are a single vector of dimension M for each query specifying the cosine similarity of a query sample with all the base samples. Now we can use the cosine similarity on the relative representations to perform retrieval. Sec D in appendix provides a further comparison with our method.

CLIP [40]: We also compare against CLIP which has been contrastively trained to obtain a joint embedding space- as an upper limit on performance for both retrieval and matching tasks. We perform retrieval using cosine similarity

For all 3 methods, caption matching can be achieved by constructing a cost matrix using cosine similarities and using linear sum assignment to find the permutation matrix.

5.3. Downstream Tasks

Caption Matching: Given N query images and their corresponding captions, a query set is constructed by shuffling the captions. The task involves finding the correct permutation over captions for perfect matching. In *Retrieval*, the objective is, given one caption, to retrieve the correct image from the overall set of N images. The alignment between unaligned vision and text encoders is investigated using our methods on the COCO and NoCaps validation sets.

Table 2. **Caption matching and retrieval task performance comparison in cross-domain and in-domain settings.** Base samples from COCO are utilized for matching/retrieval tasks on queries from NoCaps (cross-domain) and COCO (in-domain). CLIP-V denotes the vision encoder of CLIP [40]. We use the Large version of all vision encoders. Table A.5 shows the reverse setting.

Method	Vision Model	NoCaps [2]		COCO [27]	
		Matching accuracy	Top-5 retrieval	Matching accuracy	Top-5 retrieval
Cosine Similarity*	CLIP [40]	99.5	99.6	97.1	96.1
Linear regression	CLIP-V [40]	29.3	44.7	42.7	59.1
	ConvNeXt [47]	19.0	28.5	31.3	46.1
	DINOv2 [37]	38.1	50.3	45.1	65.4
Relative representations [34]	CLIP-V [40]	61.3	37.6	61.6	41.3
	ConvNeXt [47]	25.5	17.8	38.6	34.1
	DINOv2 [37]	46.0	46.4	47.7	52.3
Ours: QAP	CLIP-V [40]	**67.3**	-	**72.3**	-
	ConvNeXt [47]	46.7	-	66.1	-
	DINOv2 [37]	57.7	-	66.0	-
Ours: Local CKA	CLIP-V [40]	65.1	60.5	71.9	69.9
	ConvNeXt [47]	43.7	44.4	64.8	65.5
	DINOv2 [37]	58.7	**61.8**	64.3	**70.5**

The COCO dataset [27] comprises over 120,000 images with multiple captions per image. It is used for testing unimodal representation quality via a caption-matching task, utilizing a validation set of 5,000 image-caption pairs. The NoCaps dataset [2] is designed for testing image captioning models on unseen objects, with 166,100 captions for 15,100 images from OpenImages. Its validation set includes novel concepts absent from COCO.

Cross-lingual Caption Matching/Retrieval: The task mirrors prior matching and retrieval but uses multilingual captions, say German. Given N images and shuffled German captions, the objective is to match each image with the correct caption. In retrieval, the goal is to select the most fitting German caption for a given query image from the set.

The XTD-10 dataset [1] enhances COCO2014 with 1,000 human-annotated multi-lingual captions in ten languages for cross-lingual image retrieval and tagging, serving as a zero-shot model benchmark.

ImageNet-100 Classification. The task setup is similar to the conventional classification task with small differences to account for the methods used. Given N query images and their corresponding classes, image representations are obtained by processing them through a vision encoder. In parallel, textual representations are generated in a multi-step process. Initially, several text captions are derived from the class-associated Wordnet synsets' lemmas, definitions, and hypernyms. These captions are then passed through the language encoder and averaged to get the text representations. The classification task is performed by retrieving the closest text representations to each image representation using our local CKA metric. We employ the ImageNet-100 dataset. This dataset is a subset of the larger ImageNet dataset, featuring only 100 classes. It includes 130,000 training images, 50,000 validation images, and 100 classes.

5.4. Results

Importance of Good Initialization: For all tasks, we make use of a set of base samples of size S that is kept fixed at 320 samples. The size of the query set is analogously fixed at 500 samples (see Sec **??** for more details). These base samples are selected after clustering the image embeddings and choosing one closest sample to each of the S cluster centers. By aligning the initial samples with the diverse cluster centers, we ensure sufficient coverage of the sample space. This enhances the accuracy of the matching process, as the initial alignment closely mirrors the inherent structure and variability within the data. In the case of linear regression, uniform sampling is employed to select the base samples. For relative representations [34], the same clustering methodology is applied to select base samples, ensuring a fair and consistent comparison between all methods.

COCO and NoCaps Caption Matching: We present the results of cross-domain and in-domain caption matching/retrieval, as detailed in Table 2. We tested each baseline against three different vision models, while employing a consistent language model—specifically, the all-roberta-large-v1. The vision models utilized are OpenAI's CLIP ViT-L/14, the ConvNeXT-Base model (trained on the ImageNet-22k dataset at a resolution of 224x224), and the ViT-L/14 model trained using the DINOv2 method. It is important to note that the first row of the results table features vision and language models both being OpenAI's CLIP ViT-L/14. To effectively analyze cross-domain capabilities, our experiment design involved the use of the COCO validation set as the source of the base set and the NoCaps validation set for querying. Additionally, in-domain results are shown, when using COCO validation for both base and queries. We uniformly sample the query set and average the results over three different seeds. Although CLIP's cosine similarity metric emerges as the most robust due to the training paradigm inherent in CLIP models, our methods demonstrate commendable performance without necessitating any training. The DINOv2 model, trained solely through self-supervision, demonstrates the formation of semantic concepts independently of language supervision. This is evident in its remarkable top-5 retrieval scores of 70.5% and 61.8% on COCO and NoCaps datasets when coupled with an unaligned language encoder through our Local Kernel CKA method. However, the best-performing vision encoder is CLIP's vision encoder which has been trained using language supervision.

ImageNet-100 Classification: In Table 3, we detail the performance of our methods on the ImageNet-100 classification task. Mirroring our approach in cross-domain matching and retrieval, we evaluated three different vision models for each method. Notably, the first row of the table highlights the performance using CLIP's embedding cosine similarity. The results are averaged over three different seeds for sampling the query set. A significant observation from this table is the comparatively narrower performance gap between the CLIP's cosine similarity and our methods, as well as the baseline linear regression method, in contrast to the results observed in cross-domain caption matching/retrieval tasks.

It is interesting that ConvNeXt encoder trained on ImageNet has a classification top1 accuracy improvement of over 14% compared to CLIP and Dinov2 while on the caption matching task DinoV2 and CLIP perform much better.

Cross-lingual Caption Retrieval: The results of cross-lingual caption matching/retrieval are presented in Table 4 for the 10 languages in the XTD-dataset. OpenAI CLIP's ViT-L vision encoder, trained on English image-caption pairs, and a multilingual sentence transformer paraphrase-multilingual-mpnet-base-v2 were utilized for this task. The accuracy of CLIP's cosine retrieval method exhibits a significant drop when applied to languages other than English. *E.g.*, CLIP's retrieval at 5 experiences a drop of 30 points when switching from English to other Latin-alphabet languages (Spanish, French, German, and Italian). For non-Latin alphabet languages such as Korean, Chinese, Turkish, *etc.*, CLIP's performance decreases substantially, collapsing to zero, primarily due to most words resulting in unknown tokens. In contrast, the QAP and local CKA matching methods demonstrate consistent performance across all languages, including non-Latin languages, attributing to the robustness of a multilingual sentence transformer trained solely on text. On average, QAP surpasses CLIP by 12% in the caption matching task and also outperforms other baselines like relative representations and linear regression methods. For retrieval at 5, the local CKA-based method exceeds CLIP's performance by over 17%.

Table 3. **ImageNet-100 classification performance comparison.** We observe a narrow performance gap between the CLIP model and our methods. CLIP-V denotes the vision encoder of CLIP.

Method	Vision Model	Top 1	Top 5
Cosine Similarity*	CLIP	86.1	99.2
Linear Regression	CLIP-V	76.1	93.0
	ConvNeXt	84.5	95.4
	DINOv2	73.5	92.1
Relative representations [34]	CLIP-V	8.90	30.3
	ConvNeXt	7.20	15.7
	DINOv2	49.7	75.5
Local CKA	CLIP-V	68.7	91.2
	ConvNeXt	83.3	95.8
	DINOv2	67.7	88.3

It is possible to push the performance further by using language-specific sentence encoders and we report these results for a few languages in Sec **??** of supplementary. This is a practical application of our method as we can now turn a well-trained English CLIP model's vision encoder into a CLIP model for any low-resource language if a text-only Sentence Transformer trained on that language is available.

5.5. Matching complexity

In Table 5, we go over the time complexity and runtimes of QAP matching and local CKA based retrieval in comparison to the other baselines for matching when number of base samples and query samples are 320, 500 respectively. For all time complexities, we assume number of base samples m to be of the order of the number of query samples n. QAP uses the seeded version of the fast QAP algorithm from the SciPy library, which has a worst time complexity of $\mathcal{O}(n^3)$ [17], while local CKA retrieval requires constructing a graph over all the query image and text pairs, $\mathcal{O}(n^2)$, using local CKA, which is also $\mathcal{O}(n^2)$ resulting in $\mathcal{O}(n^4)$. Relative involves the calculation of the relative representations for every query image and text pair, resulting in a time complexity of $\mathcal{O}(n^2)$, but it's fast due to highly optimized algorithms for matrix multiplications in PyTorch [38]. Linear has a time complexity of $\mathcal{O}(nd)$, where n is the number of samples and d is the number of dimensions. It is to be noted that QAP runs on the CPU, and a CUDA-optimized version could bring the runtimes further down from 40 seconds. An efficient implementation of Local Kernel CKA is also possible, where the CKA of base samples is precalculated, and the graph is constructed in an additive manner, which would bring down the time complexity to $\mathcal{O}(n^3)$. For both relative and linear matching, we make use of SciPy's

Table 4. **Cross-Lingual caption matching and retrieval performance comparison.** Using QAP and local CKA-based methods we are able to do cross-lingual caption matching/retrieval using CLIP's ViT-L vision encoder and a multi-lingual sentence transformer paraphrase-multilingual-mpnet-base-v2. While CLIP performs well on the Latin languages, it degrades on non-Latin languages. In comparison, our QAP and Local-CKA-based methods perform comparably in Latin languages while outperforming non-Latin languages, highlighting the efficacy of our training-free transfer approach. See Table A.6 and Table A.7 in appendix for additional results.

Language		**Kernel CKA**		**Matching Accuracy**				**Retrieval @ 5**	
		CLIP	Ours	CLIP	Relative[34]	Linear	Ours (QAP)	CLIP	Ours (Local)
Latin	de	0.472	0.627	41.8	35.0	34.0	39.6	65.1	56.7
	en	0.567	0.646	81.5	52.5	40.9	51.6	92.5	69.0
	es	0.471	0.634	50.2	37.8	31.7	41.4	68.5	61.6
	fr	0.477	0.624	49.4	37.5	30.7	40.2	68.7	57.6
	it	0.472	0.638	41.0	37.2	34.9	38.5	61.3	59.7
Non-Latin	jp	0.337	0.598	13.2	28.3	23.5	30.5	30.0	49.4
	ko	0.154	0.620	0.50	30.4	23.5	30.9	3.30	53.4
	pl	0.261	0.642	5.40	36.6	30.2	40.2	18.8	59.5
	ru	0.077	0.632	0.80	31.9	30.7	35.1	4.10	53.2
	tr	0.301	0.624	4.30	35.8	29.6	38.9	15.2	59.3
	zh	0.133	0.641	2.70	36.5	31.1	40.3	8.90	57.8
	Avg.	–	–	26.4	36.3	30.9	**38.8**	39.6	**57.9**

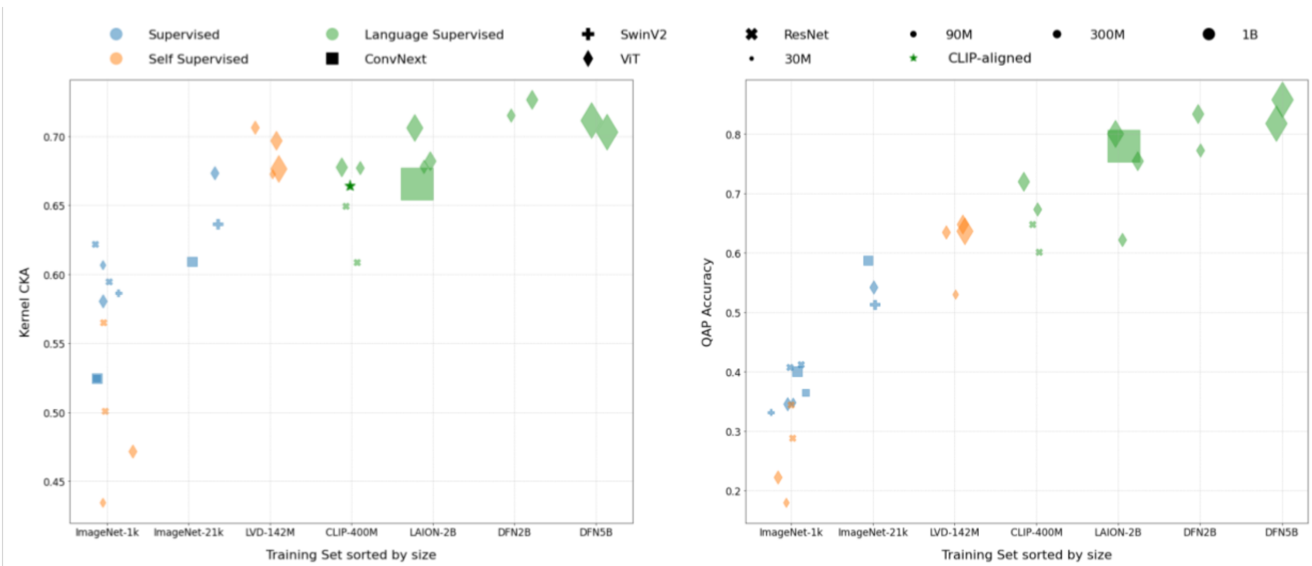

Figure 2. **Kernel CKA and QAP Matching accuracy are correlated with the training set size and quality of the training set.** Here the language encoder is kept constant to the best BERT-sentence encoder (i.e.All-Roberta-large-v1). There is a clear correlation between CKA and QAP Matching accuracy across all architectures, training paradigm and data regimes.

Table 5. Run times for different methods

Method	QAP	Local CKA	Relative	Linear
Run times	40 seconds	5 mins	1 second	1 second
Complexity	$\mathcal{O}(n^3)$	$\mathcal{O}(n^4)$	$\mathcal{O}(n^2)$	$\mathcal{O}(n \times d)$

modified Jonker-Volgenant algorithm [13] for linear sum assignment, which has the worst time complexity of $\mathcal{O}(n^3)$.

6. Analysis

This section focuses on how training paradigms, data regimes, and encoder size/architecture influence a vision encoder's ability to represent the world similarly to a language encoder. This is assessed by comparing the semantic alignment of their representation spaces using CKA as well as QAP matching accuracy. Figure 2 compares the kernel CKA and caption matching accuracy of different vision encoders with a fixed text-encoder (*i.e*., All-Roberta-large-v1), against the training datasets on which the vision encoder was trained for all pairs in the COCO captions validation set. The findings are summarized below:

Scale and quality of dataset results in encoders with high semantic alignment with the language space: It is observed that SSL methods like DINOv2 can learn semantic concepts in a relative manner even without language supervision during training. The CKA and QAP matching accuracy for DINOv2 embeddings are comparable to CLIP models, despite lacking language supervision and having significantly less data (LVD-142's 142M vs Open-AI-CLIP's 400M). A general trend emerges where more training data leads to semantically richer visual embeddings, evident when comparing CKA and QAP Accuracies from ImageNet1K to DFN-5B datasets. Notably, training on a curated dataset proves more effective than on an uncurated dataset of the same size, especially for smaller models. This is illustrated by the higher CKA and QAP accuracy of ViT-Large trained on the curated DFN-2B dataset compared to ViT-Large/Giant, and ConvNext-xxLarge trained on Laion 2B. Additionally, SSL methods show less semantic consistency when trained on ImageNet1K, as indicated by the clear difference in QAP accuracies between DINO trained on ImageNet1K and DINOv2 trained on LVD-142M.

Vision Encoders Trained with Language Supervision Exhibit Greater Semantic Alignment with Language Encoders: In line with the findings of Merullo et al.[31], it is observed in our experiments that vision encoders trained with more language supervision on datasets of comparable size exhibit a higher degree of semantic alignment with language encoders compared to self-supervised methods. For example, ViT-Large trained on CLIP-400M with language supervision demonstrates superior caption-matching capabilities compared to DINOv2's ViT-Large trained on LVD-142M. Similarly, we verify that class label supervision, like that from ImageNet, leads to more semantically aligned image encoders when compared to self-supervision when similarly sized models are compared on ImageNet-1k. For example, all supervised encoders trained on ImageNet-1k have higher CKA as well as QAP matching accuracy than all the self-supervised models.

7. Ablations

This section rationalizes our method choices through ablation studies on clustering, stretching, and the global CKA metric. We demonstrate the impact of these components on the performance of our methods, primarily through Table 6, which delineates the effectiveness of the QAP and the local CKA metric under various configurations. It shows the performance metrics in scenarios where each main component is either integrated or omitted. Notably, in instances where the CKA metric is not used, we opt for normalized correlation matrices for each graph. The empirical results presented are derived from the caption matching/retrieval task, utilizing both base and query sets extracted from the COCO validation set of size 320 and 500 respectively.

Table 6. **Impact of clustering and stretching.** The matching and retrieval performance is the best when both clustering and stretching are employed. Hence, justifying this choice.

Clustering	Stretching	CKA	QAP Matching	Local CKA Matching	Local CKA Retrieval @ 5
✗	✗	✗	10.1	16.2	1.0
✗	✗	✓	48.8	48.5	60.2
✗	✓	✓	57.3	56.7	73.0
✓	✗	✓	56.2	55.1	66.4
✓	✓	✓	**65.5**	**63.3**	**77.2**

Choice of the metric: CKA is more beneficial than using just the scaled correlation matrix to represent the semantic relationships in an embedding space as matching accuracy increases from 10.1% to 48.8%. The choice of a robust metric is core to aligning vision and language latent spaces.

Impact of Stretching: It is clear that stretching facilitates better alignment of embeddings in our methods as stretching spreads the representations out in each modality without sacrificing the relative positions of the different embeddings within each embedding space. This is reflected in the increase of QAP accuracy from 48.8% to 57.3%.

Clustering *vs.* Uniform Sampling: The choice of the base set is important in QAP matching and local CKA retrieval, as it measures any query pair alignment with the base set. A diverse base set is essential to capture a broad semantic range, and clustering within one of the embedding spaces aids in achieving this diversity. The third and fifth rows of the table demonstrate that clustering enhances the QAP performance from 57.3% to 65.5%. Consequently, these results highlight that all the components together significantly enhance the efficacy of our proposed approach.

8. Conclusion

In this work, we ask the question, *'Do vision encoders and language encoders represent the world similarly?'* and study this using CKA and a caption-matching task. We find that well-trained vision encoders on sufficiently large datasets exhibit surprisingly high semantic similarity with language encoders comparable to aligned encoders, irrespective of the training paradigm. Inspired by this, we draw parallels between CKA and the QAP matching objective and use seeded graph matching to align vision and language encoders by maximizing CKA. We also devise a local CKA-based metric to enable retrieval between unaligned vision and language encoders demonstrating a better performance than that of relative representations on cross-domain and cross-lingual caption matching/retrieval tasks, facilitating zero-shot latent space communication between unaligned encoders.

References

[1] Pranav Aggarwal and Ajinkya Kale. Towards zero-shot cross-lingual image retrieval. *arXiv preprint arXiv:2012.05107*, 2020. 5

[2] Harsh Agrawal, Karan Desai, Yufei Wang, Xinlei Chen, Rishabh Jain, Mark Johnson, Dhruv Batra, Devi Parikh, Stefan Lee, and Peter Anderson. Nocaps: Novel object captioning at scale. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8948–8957, 2019. 2, 5

[3] Richard Antonello, Javier S Turek, Vy Vo, and Alexander Huth. Low-dimensional structure in the space of language representations is reflected in brain responses. *Advances in neural information processing systems*, 2021. 2

[4] Yamini Bansal, Preetum Nakkiran, and Boaz Barak. Revisiting model stitching to compare neural representations. *Advances in neural information processing systems*, 2021. 2

[5] Serguei Barannikov, Ilya Trofimov, Nikita Balabin, and Evgeny Burnaev. Representation topology divergence: A method for comparing neural network representations. *arXiv preprint arXiv:2201.00058*, 2021. 2

[6] Lisa Bonheme and Marek Grzes. How do variational autoencoders learn? insights from representational similarity. *arXiv preprint arXiv:2205.08399*, 2022. 2

[7] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020. 4

[8] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9650–9660, 2021. 1

[9] Soravit Changpinyo, Piyush Sharma, Nan Ding, and Radu Soricut. Conceptual 12m: Pushing web-scale image-text pre-training to recognize long-tail visual concepts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3558–3568, 2021. 1

[10] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020. 1

[11] Alexis Conneau, Guillaume Lample, Marc'Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. Word translation without parallel data. *arXiv preprint arXiv:1710.04087*, 2017. 2

[12] Corinna Cortes, Mehryar Mohri, and Afshin Rostamizadeh. Algorithms for learning kernels based on centered alignment. *The Journal of Machine Learning Research*, 13(1):795–828, 2012. 1

[13] David F Crouse. On implementing 2d rectangular assignment algorithms. *IEEE Transactions on Aerospace and Electronic Systems*, 52(4):1679–1696, 2016. 7

[14] Adrián Csiszárik, Péter Korösi-Szabó, Akos K Matszangosz, Gergely Papp, and Dániel Varga. Similarity and matching of neural network representations. *arXiv preprint arXiv:2110.14633*, 2021. 2

[15] J. Deng, W. Dong, R. Socher, L. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2009. 2

[16] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 1, 4

[17] Donniell E Fishkind, Sancar Adali, Heather G Patsolic, Lingyao Meng, Digvijay Singh, Vince Lyzinski, and Carey E Priebe. Seeded graph matching. *Pattern recognition*, 87: 203–215, 2019. 3, 6

[18] Arthur Gretton, Olivier Bousquet, Alex Smola, and Bernhard Schölkopf. Measuring statistical dependence with hilbert-schmidt norms. In *International conference on algorithmic learning theory*, pages 63–77. Springer, 2005. 2

[19] Michael Gutmann and Aapo Hyvärinen. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 297–304. JMLR Workshop and Conference Proceedings, 2010. 1

[20] 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 *International conference on machine learning*, pages 4904–4916. PMLR, 2021. 1

[21] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *ICLR*, 2014. 2

[22] Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey Hinton. Similarity of neural network representations revisited. In *International conference on machine learning*, pages 3519–3529. PMLR, 2019. 1, 2

[23] Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955. 4

[24] 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*, pages 991–999, 2015. 2

[25] Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691*, 2021. 2

[26] Yixuan Li, Jason Yosinski, Jeff Clune, Hod Lipson, and John Hopcroft. Convergent learning: Do different neural networks learn the same representations? *arXiv preprint arXiv:1511.07543*, 2015. 2

[27] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. 1, 3, 5

[28] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019. 1, 3, 4

[29] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11976–11986, 2022. 4

[30] Wan-Duo Kurt Ma, JP Lewis, and W Bastiaan Kleijn. The hsic bottleneck: Deep learning without back-propagation. In *Proceedings of the AAAI conference on artificial intelligence*, pages 5085–5092, 2020. 2

[31] Jack Merullo, Louis Castricato, Carsten Eickhoff, and Ellie Pavlick. Linearly mapping from image to text space. *arXiv preprint arXiv:2209.15162*, 2022. 1, 2, 8

[32] Tomas Mikolov, Quoc V Le, and Ilya Sutskever. Exploiting similarities among languages for machine translation (2013). *arXiv preprint arXiv:1309.4168*, 2022. 2

[33] Ari Morcos, Maithra Raghu, and Samy Bengio. Insights on representational similarity in neural networks with canonical correlation. *Advances in neural information processing systems*, 31, 2018. 2

[34] Luca Moschella, Valentino Maiorca, Marco Fumero, Antonio Norelli, Francesco Locatello, and Emanuele Rodolà. Relative representations enable zero-shot latent space communication. In *The Eleventh International Conference on Learning Representations*, 2022. 1, 2, 4, 5, 6, 7

[35] Antonio Norelli, Marco Fumero, Valentino Maiorca, Luca Moschella, Emanuele Rodola, and Francesco Locatello. Asif: Coupled data turns unimodal models to multimodal without training. *arXiv preprint arXiv:2210.01738*, 2022. 1, 2

[36] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 1

[37] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 1, 3, 5

[38] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019. 6

[39] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019. 4

[40] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 1, 4, 5

[41] Maithra Raghu, Justin Gilmer, Jason Yosinski, and Jascha Sohl-Dickstein. Svcca: Singular vector canonical correlation analysis for deep learning dynamics and interpretability. *Advances in neural information processing systems*, 30, 2017. 2

[42] Maithra Raghu, Thomas Unterthiner, Simon Kornblith, Chiyuan Zhang, and Alexey Dosovitskiy. Do vision transformers see like convolutional neural networks? *Advances in Neural Information Processing Systems*, 34:12116–12128, 2021. 1

[43] Nils Reimers and Iryna Gurevych. Sentence-bert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 2019. 4

[44] Anton Tsitsulin, Marina Munkhoeva, Davide Mottin, Panagiotis Karras, Alex Bronstein, Ivan Oseledets, and Emmanuel Müller. The shape of data: Intrinsic distance for data distributions. *arXiv preprint arXiv:1905.11141*, 2019. 2

[45] Joshua T Vogelstein, John M Conroy, Vince Lyzinski, Louis J Podrazik, Steven G Kratzer, Eric T Harley, Donniell E Fishkind, R Jacob Vogelstein, and Carey E Priebe. Fast approximate quadratic programming for graph matching. *PLOS one*, 10(4):e0121002, 2015. 3

[46] Ivan Vulić, Sebastian Ruder, and Anders Søgaard. Are all good word vector spaces isomorphic? *arXiv preprint arXiv:2004.04070*, 2020. 2

[47] Sanghyun Woo, Shoubhik Debnath, Ronghang Hu, Xinlei Chen, Zhuang Liu, In So Kweon, and Saining Xie. Convnext v2: Co-designing and scaling convnets with masked autoencoders. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16133–16142, 2023. 1, 5

[48] John M Wu, Yonatan Belinkov, Hassan Sajjad, Nadir Durrani, Fahim Dalvi, and James Glass. Similarity analysis of contextual word representation models. *arXiv preprint arXiv:2005.01172*, 2020. 2