DYNAMIC REFLECTIONS: PROBING VIDEO REPRESENTATIONS WITH TEXT ALIGNMENT

Anonymous authorsPaper under double-blind review

000

001

002003004

010 011

012

013

014

015

016

017

018

019

021

024

025

026

027

028

029

031

032

034

035

037

040

041

042

043

044

046

047

051

052

ABSTRACT

The alignment of representations from different modalities has recently been shown to provide insights on the structural similarities and downstream capabilities of different encoders across diverse data types. While significant progress has been made in aligning images with text, the temporal nature of video data remains largely unexplored in this context. In this work, we conduct the first comprehensive study of video-text representation alignment, probing the capabilities of modern video and language encoders. Our findings reveal several key insights. First, we demonstrate that cross-modal alignment highly depends on the richness of both visual (static images vs. multi-frame videos) and text (single caption vs. a collection) data provided at test time, especially when using state-of-theart video encoders. We propose parametric test-time scaling laws that capture this behavior and show remarkable predictive power against empirical observations. Secondly, we investigate the correlation between semantic alignment and performance on both semantic and non-semantic downstream tasks, providing initial evidence that strong alignment against text encoders may be linked to generalpurpose video representation and understanding. Finally, we correlate temporal reasoning with cross-modal alignment providing a challenging test-bed for vision and language models. Overall, our work introduces video-text alignment as an informative zero-shot way to probe the representation power of different encoders for spatio-temporal data.

1 Introduction

A hallmark of general intelligence is the ability to reason about the world in its different manifestations and to integrate different sensory data into unified mental representations (Ernst & Bülthoff, 2004; Kudithipudi et al., 2022). The pursuit of such unified representations is also fundamental to the development of capable multimodal AI systems and for the advancement of embodied agents that must operate within a complex world (Xie et al., 2024; Fung et al., 2025).

In recent years, significant progress has been made towards this goal. The success of multimodal learning approaches has highlighted the potential of joint training paradigms (Radford et al., 2021; Yin et al., 2024; Liu et al., 2024b; Dubey et al., 2024; Team et al., 2023; 2025), while a parallel line of research has provided evidence that even *unimodal* models, when trained at scale, develop representations with strong structural similarities (Huh et al., 2024; Maniparambil et al., 2024a; Tjandrasuwita et al., 2025). This observation has given rise to the Platonic Representation Hypothesis (PRH), which posits that neural networks converge towards a shared statistical model of reality in their latent spaces (Huh et al., 2024). Building on this, several recent works have established a wide range of techniques to evaluate and exploit cross-modal alignment, in both zero-shot and few-shot scenarios (Maniparambil et al., 2024a; Jha et al., 2025; Schnaus et al., 2025; Hadgi et al., 2025).

However, prior work has focused almost exclusively on the alignment of *static* modalities: primarily images and text. In other words, although the Platonic Representation Hypothesis has been stated in full generality, so far, its validity has only been evaluated on static data, leaving its applicability to dynamic, temporal domains an open question. Consequently, the rich information contained in video: motion, causality and temporal dependencies, exhibited, e.g., in intricate human interactions (Gu et al., 2018; Wang et al., 2023b), has been largely overlooked in the context of representation alignment.

A notable limitation of previous studies focusing on static modalities was raised in Huh et al. (2024), where the authors pointed out: "the maximum theoretical value for the alignment metric is 1. Is a score of 0.16 indicative of strong alignment [...] or does it signify poor alignment with major differences left to explain? We leave this as an open question." In this work, we provide a partial answer to this open question by showing that the limited alignment observed previously is, in large part, due to the impoverished data given at test time. Specifically, we demonstrate that by considering multiple video frames instead of a single image, as well as a diverse set of captions instead of a single annotation, the alignment score can be improved significantly, achieving close to 0.4 in some cases, without modifying the underlying trained models. This result highlights, for the first time, that large improvements in alignment can be achieved through efforts at test time, which is complementary to the training-time resources (model size, amount of training data, etc.) considered in prior work. It also suggests that multi-frame approaches can capture different aspects of a scene better, and points to the possibility of a strong zero-shot evaluation metric capable of probing the representation power of both image and video encoders.

More broadly, in this work we extend previous cross-modal alignment studies *into the temporal domain* by conducting the first comprehensive investigation of video-text representation similarity. We introduce a robust evaluation framework to probe and compare the capabilities of modern video and language models. Our findings reveal that the rich, temporal nature of video provides a powerful signal for semantic grounding. Specifically, we demonstrate that (1) state-of-the-art self-supervised video encoders, e.g., VideoMAEv2 (Wang et al., 2023a), achieve competitive alignment with text compared to top-performing image encoders, e.g., DINOv2 (Oquab et al., 2023); (2) we observe that alignment quality is highly sensitive to the richness of the provided visual and textual data, and (3) there is a non-trivial relation between the video-text alignment and the performance of video models on downstream tasks. Finally, we highlight challenging situations with hard negative examples, in which current video and text models struggle, leaving room for future improvement.

2 RELATED WORK

Our work is situated at the intersection of three domains: the study of emergent representation alignment, self-supervised video understanding, and multimodal learning. We first review the recent work in representation convergence, then discuss the most prominent paradigms for learning from video data, and finally, contextualize our work by highlighting the specific gap in probing the alignment of unimodal trained video models.

The Platonic Representation Hypothesis and Emergent Alignment in Static Modalities A foundational concept for our investigation is the Platonic Representation Hypothesis (PRH), introduced by Huh et al. (2024). The PRH posits that as neural networks are scaled in terms of capacity, data diversity, and task variety, their learned internal representations converge toward a shared, universal statistical model of reality. This converged latent structure is termed the "Platonic representation," drawing an analogy to Plato's notion of an ideal reality that underlies our sensory observations. This hypothesis has been empirically tested by measuring the geometric similarity of representation spaces, obtained using a wide range of encoders. The primary method for this is comparing the representational kernels, which characterize how models measure dissimilarity between data points (Kornblith et al., 2019; Maniparambil et al., 2024b; Huh et al., 2024). The key finding supporting the PRH is that as unimodal vision and language models become more capable, the structure of their latent spaces becomes increasingly similar, even without explicit cross-modal supervision.

This convergence has been most robustly demonstrated in the alignment of static modalities, primarily between images and text (Maniparambil et al., 2024b; Tjandrasuwita et al., 2025). For example, the work of Maniparambil et al. (2024b) showed that powerful, independently trained vision and language encoders develop representations with a surprisingly high degree of semantic similarity. Furthermore, following prior work (Merullo et al., 2022), the authors of Maniparambil et al. (2024b) showed that these disparate latent spaces may differ only by a simple, learnable linear transformation. A more fundamental principle behind these and related results was offered in Huh et al. (2024), which posited that since all modalities are projections of the same physical world, models trained at sufficient scale will inevitably converge on similar latent structures that reflect this shared reality (Huh et al., 2024; Tjandrasuwita et al., 2025). The vision encoders used to validate these claims are typically state-of-the-art self-supervised models like DINOv2 (Oquab et al., 2023). Recent work has also examined the conditions and methods related to this emergent alignment (Tjandrasuwita et al.,

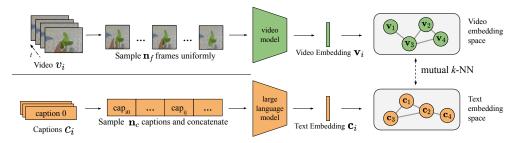


Figure 1: Scaling both the number of video frames and text captions at test time improves alignment. Given a paired video v_i and set of captions c_i , leveraging rich multi-frame and multi-caption information at test time leads to improved alignment, measured in terms of mutual k-NN.

2025). Cross-modal alignment has also been assessed via "alignment probing," which correlates emergent alignment potential with the representation's clustering quality (k-NN performance) over linear separability (Zhang et al., 2025). The underlying universal structure has even been shown to be sufficiently strong to enable both unsupervised and weakly supervised translation between disparate embedding spaces (Jha et al., 2025; Zhang et al., 2025; Schnaus et al., 2025). While these studies investigate and exploit emergent alignment, they focus almost exclusively on *static* modalities. Our work complements these findings by conducting the first systematic probe of emergent alignment in the *temporal* domain, and shedding light on the dependence of alignment on test-time data in both text and visual domains.

Learning Representations from Video Learning from video data presents unique challenges and opportunities due to the temporal dimension. Research in this area is divided into unimodal self-supervised learning, which focuses on spatiotemporal structure, and joint video-language pretraining, which focuses on explicit semantic alignment, e.g., Xu et al. (2021); Zhao et al. (2024). In the former category, masked autoencoding has proven highly effective (Tong et al., 2022; Bardes et al., 2024; Wang et al., 2023b). This paradigm, inspired by its success in language and image domains, involves masking portions of the input video and training a model to reconstruct the missing content in either pixel or feature spaces. Although initial approaches in masked video representation learning have focused on modest model sizes, recent efforts have also been made to scale such models to the billion parameter regime (Wang et al., 2023b; Carreira et al., 2024; Assran et al., 2025). Critically, such models are pre-trained on vast, unlabeled video datasets with self-supervised (e.g., reconstruction) objectives. For example, VideoMAEv2 has no exposure to textual supervision, yet it learns powerful spatiotemporal features that achieve top performance on downstream action understanding tasks (Wang et al., 2023b). This makes it the ideal subject for probing the existence of emergent semantic alignment. Furthermore, there is recent evidence that even unimodal video models, when trained at scale, tend to learn features that are useful in a broad range of tasks (Wu et al., 2025; Carreira et al., 2024). We also note that evaluating self-supervised video representations is challenging, and current approaches rely on expensive task-specific training (Wang et al., 2023b; Carreira et al., 2024; Hasson et al., 2025). There is thus a strong need for zero-shot metrics with strong predictive power for video models.

Bridging Unimodal Video and Text: Probing for Emergent Alignment Given the success of powerful unimodal encoders in both static (Oquab et al., 2023) and temporal (Wang et al., 2023b) domains, and the evidence for emergent alignment in static data (Maniparambil et al., 2024b), a natural question arises: does the Platonic Representation Hypothesis extend to the *temporal* domain? While our work is the first to conduct a systematic probe of this question, we acknowledge related efforts that bridge video and text representations, e.g., (Kim et al., 2023; Liu et al., 2024a; Li et al., 2025). However, existing works in this domain typically rely on an existing source of alignment rather than investigating the intrinsic properties of video encoders.

3 Our Approach

At a high-level, our approach follows the methodology introduced in Huh et al. (2024), which uses a mutual k-NN metric to measure the similarity across different modalities. We adapt this approach for our setting, extending it to measure the multi-frame (or multi-clip) and multi-caption similarity of video-text pairs. Our setting is described in Figure 1.

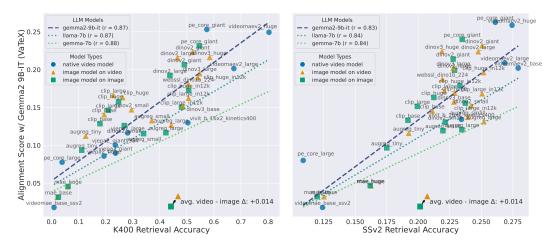


Figure 2: **Video representations are strongly aligned with text.** We measure alignment between modern vision encoders and recent text encoders (Gemma 2-9B-it shown, subset of all vision models for clarity) on the VaTeX video dataset using a single caption for each video. We differentiate between native video, multi-frame image and single frame vision models $(\bullet, \blacktriangle, \blacksquare)$. Vision representations are evaluated for video retrieval using weighted k-NN on the Kinetics-400 and SSv2 test sets. In addition to the points, we plot a linear regression of alignment scores to representation strength for three different LLMs. Our takeaways are threefold: (1) Image-only models are mostly limited to alignment scores of 20%, as noted in Huh et al. (2024). (2) Recent text models have improved alignment with vision models, shown by the regression lines. (3) Averaging over frames is a simple yet effective baseline for image models on video input.

At its core, the Platonic representation alignment is calculated in the following stages.

Dataset. We assume a **test set** of N video-caption pairs: $S = (V, C) = \{(v_1, c_1), \dots, (v_N, c_N)\}$. Here v_i is a video, and c_i is a corresponding set of text captions for this video. Each caption c_{ij} in the set c_i provides a textual description of the given video v_i .

Encoding. We embed each video using a video encoder into some video embedding space $\mathcal{E}_{\mathrm{vid}}(v_i) = \mathbf{v}_i \in \mathbb{R}^p$. Similarly, we encode each set of captions c_i into a text embedding space using a text encoder: $\mathcal{E}_{\mathrm{text}}(c_i) = \mathbf{c}_i \in \mathbb{R}^q$. Importantly, the dimensionality of the embedding spaces, p and q are typically different, making it difficult to directly compare distances. We define $\mathbf{X} \in \mathbb{R}^{N \times p}$ and $\mathbf{Y} \in \mathbb{R}^{N \times q}$ to be the matrices obtained by stacking all of the video and text embeddings respectively. For simplicity we denote $\mathbf{X} = \mathcal{E}_{\mathrm{vid}}(V)$, and $\mathbf{Y} = \mathcal{E}_{\mathrm{text}}(C)$.

Alignment metric. The Mutual k-NN (MkNN) metric introduced in Huh et al. (2024) measures the *agreement* between the nearest neighbor structure in two different embedding spaces. Given two feature sets $\mathbf{X} \in \mathbb{R}^{N \times p}$, $\mathbf{Y} \in \mathbb{R}^{N \times q}$, we first construct two binary indicator matrices $\mathbf{M}_{\mathbf{X}}$ and $\mathbf{M}_{\mathbf{Y}}$, both of size $N \times N$, s.t., $(\mathbf{M}_{\mathbf{X}})_{ij} = 1$ if element (row) j is among the k-nearest neighbors of element (row) i in space \mathbf{X} , and 0 otherwise. The Mutual k-NN alignment is then calculated as:

$$\mathcal{A}^{\text{MkNN}}(\mathbf{X}, \mathbf{Y}) = \frac{1}{kN} \sum_{i=1}^{N} \sum_{j=1}^{N} (\mathbf{M}_{\mathbf{X}} \odot \mathbf{M}_{\mathbf{Y}})_{ij}.$$
 (1)

Here, \odot denotes the Hadamard product (element-wise multiplication) of the two indicator matrices. Observe that the operation in Eq. (1) represents simply the *mean overlap* (alignment) in the sets of k nearest neighbors, normalized by a hyperparameter k (typically set to k=10 for a dataset of 1024 examples). Additionally, we also follow previous work and optimize over the choice of intermediate layers for both encoders, and pick the pair of layers that maximizes the alignment score.

Incorporating multi-instance data. As mentioned above, our main goal is to extend previous approaches, which focused on static instances to datasets of *videos* (treated as multi-frame sequences) paired with potentially *multiple* diverse captions; see Sec. 4 for a detailed description of our data sources. For a video v_i and a given video encoder \mathcal{E}_{vid} , which natively processes clips with n_o frames, we extract the indices of n_f frames through uniform linear interpolation to study the effect

of using an increasing number of frames n_f at test-time. Using one frame $(n_f=1)$ corresponds to the setup used in prior works that focus on image-text alignment. For videos longer than n_o , we use nearest neighbor interpolation to extract increasing numbers of frames that are multiple of n_o , e.g. for $n_o=16$, we consider $n_f\in\{16,32,64,80\}$. We pass the n_o -length sub-clips through the encoder and then average the representations over sub-clips.

Similarly, for a set of captions c_{ij} associated with a given video v_i , we can use anywhere from one caption to the full set of captions. We concatenate the set of selected captions into a single string and use text-based encoders (including LLM-based ones as discussed below) to extract their intermediate features. When an encoder produces per-token embedding, we average the features along the token dimension to obtain a [layer, hidden dimension]-shaped feature, associated with each video.

4 Data and Models

Data. We evaluate video-text alignment on several datasets that contain paired video/text data. Our main investigations use VATEX and the Perception Encoder Video Datasets (PVD) (Wang et al., 2019; Bolya et al., 2025). We construct two test sets using 1024 videos randomly sampled from VATEX and PVD. Each video in VATEX lasts around 10 seconds and is taken from a unique YouTube video. They are captioned by 10 different annotators in English and Chinese (we only use the English captions). In other words, each video contains 10 different captions (each caption is produced by a different annotator and is typically a sentence with around 15 words on average). This provides a rich source of diversity which we can use to vary the amount of textual information used for alignment. To achieve a similar effect with PVD, we use Gemini-2.5Pro to split each fine-grained caption into 10 separate captions in a similar style to VATEX; see Sec A.3 for more details.

Models. We experimented with a wide range of both video and image encoders, as well as aggregation strategies. For image models, we consider two variants: first, we simply use a single (middle) frame; second, to introduce naive temporal dynamics, we averaged the image features across the temporal dimension. We denote the resulting method as *image model on video*.

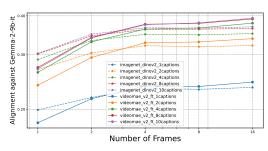
For video models, we consider a range of encoders, including self-supervised ones like VideoMAE and VideoMAEv2 (Tong et al., 2022; Wang et al., 2023b), as well as partially text-supervised ones like VideoPrism (Zhao et al., 2024). We also evaluate recent vision models, such as the Perception Encoder and DINOv3 (Bolya et al., 2025; Siméoni et al., 2025). We make a distinction between text-aligned and self-supervised models and use the former as an upper bound for our alignment evaluation. For image models which are trained on video data, i.e., Perception Encoder (Bolya et al., 2025), we denote them as video models when tested on video input. In total, we test 63 different models and variants; see full list in the Appendix, Section A.2.

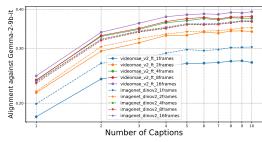
For language models, we experimented with a broad range of encoders, starting with the ones considered in Huh et al. (2024). We also consider recent *unimodal* language models from the Gemma 2 series (Gemma Team: Riviere et al., 2024) and evaluate their potential as text encoders. Given a caption, we encode it with a text encoder and then average the representation over the token sequence. As mentioned above, we store intermediate representations across all layers of the encoder and select the best pair of layers for each pair of vision and text models when measuring alignment.

5 VIDEO-TEXT ALIGNMENT RESULTS

Our first results on the VATEX dataset (Wang et al., 2019) are summarized in Figure 2. These results are obtained by using a randomly selected single (out of 10) caption for each video. We evaluate video understanding performance using nearest neighbour video retrieval on 10,000-video subsets of the Kinetics-400 and SSv2 test sets, respectively (Kay et al., 2017; Goyal et al., 2017). We use weighted k-NN accuracy as retrieval metric following NPID and DINO (Wu et al., 2018; Caron et al., 2021); note that this is different from our Mutual k-NN alignment metric. We sweep over the number of nearest neighbours and find 8 to give good results for majority of the models for the retrieval task. We also report the alignment of the same vision models against a wide range of language encoders on the VATEX dataset in the Appendix, Figure 6.

We make several observations: first, when considering pure image and text models that were studied in prior work Huh et al. (2024), the alignment scores that we obtain closely follow previously re-





- (a) Alignment Score vs. Number of Frames for Different Vision Models
- (b) Alignment Score vs. Number of Captions for Different Vision Models

Figure 3: Analysis of vision-text alignment sensitivity to the amount of visual and textual data available at test time. (a) Providing more *frames* increases vision-text alignment for both image and video models, with the latter being able to take advantage of more frames more effectively. (b) Providing more *captions* to the text model also significantly boosts alignment with vision models, across all frame counts.

ported values on other datasets (e.g., the alignment score between the best image model DINOv2 and the best text encoder outside of the Gemma family is 0.18). Secondly, we note that more recent text models in the Gemma-2 family lead to better alignment with video models, even if trained purely for text *generation*, highlighting the utility of these models as *text encoders*. Simply using a powerful text encoder already increases the best image-text alignment score to approximately 0.206. This is consistent with, e.g., Zhang et al. (2025) which has highlighted the importance of the language model for multimodal alignment. Third, we note that simple temporal averaging across multiple frames with powerful image models exhibits remarkably high video-text alignment reaching alignment of approximately 0.223. Finally, there is a very wide range of text-alignment scores across video models. Nevertheless, the highest alignment is achieved with a self-supervised VideoMAEv2 model. This last result is notable, since it shows, both that natively-trained video models can outperform extremely strong image-based models (DINOv2) in text alignment, and that temporal dynamics play an important role in semantic grounding.

Overall, these results above point to the potential, first, of using multi-frame video data for text alignment, and second, of using language models for computing informative embeddings. Since in all of our experiments, Gemma-based encoders achieved the best alignment, our subsequent experiments focus on that particular text encoder and investigate the role of vision encoders in text alignment.

6 VIDEO-TEXT ALIGNMENT AND DATA DEPENDENCE

A key observation of our work is that the quality of the vision-text alignment depends strongly on the amount of visual and caption data *given at test time*. This is in direct complement to the focus of previous works that analyzed the dependence of the alignment on the size of the *pre-training* datasets or number of trainable parameters in different models.

Figures 11a and 11b illustrate the dependence of vision-text alignment scores (against the Gemma 2-9b-it text encoder) on the number of frames from a video and the number of associated captions in the VATEX dataset (Wang et al., 2019). Figure 11a demonstrates how the average alignment score changes with an increasing number of frames for different caption counts and two vision models: VideoMAEv2 (Wang et al., 2023b) and DINOv2 (Oquab et al., 2023). Observe that across all settings, the alignment scores generally increases as more frames are incorporated, suggesting that a richer temporal context from the video leads to better alignment with textual descriptions.

Figure 11b complements the first by showing the alignment score's dependency on the number of captions used, for the same vision models and a selection of frame counts. Increasing the number of captions leads to a significant improvement in alignment scores, particularly when starting with a small number of captions. We note that since in the VATEX dataset each caption corresponds to a description of the same video provided by a different user, increasing the number of captions increases both the coverage of the captured visual concepts, as well as the *diversity* of perspectives

on the same concept and event. We include in Appendix, Figure 7 alignment scores against multiple Gemma models when using single vs all captions in VATEX, confirming this same observation.

These plots collectively underscore the importance of visual (frames) and textual (captions) data quantity in achieving high vision-text alignment, with distinct performance characteristics associated with different vision models. We emphasize that our results strongly complement the preliminary experiments in Sec. 6 of Huh et al. (2024), relating information density to alignment. This is because we investigate, for the first time, the use of multi-frame video data, but also in our study of multiple *complementary* captions instead of a single shorter description distilled from a longer one as done in that work. We observe that having *diverse* captions, possibly describing the same visual concept in different ways, can boost alignment significantly.

To illustrate this behavior further, we demonstrate that by using multiple text descriptions, artificially *synthesized* from a single long video caption, we can obtain alignment that is *higher* than that against the source caption. Namely, in Figure 8 (appendix), we find that our test time scaling approach improves alignment on the PVD dataset (Bolya et al., 2025).

Test-time Scaling Laws Building on the empirical observations in Figures 11a and 11b, we also quantify these dependencies using a predictive parametric model. We tested several formulations, but found that a saturation-based model (eq. 2) provided the best fit by a significant margin:

$$score(\mathbf{n}_f, \mathbf{n}_c) = S_{\infty} - (C_f \mathbf{n}_f^{-\alpha} + C_c \mathbf{n}_c^{-\beta}).$$
 (2)

Here ${\bf n}_f$ and ${\bf n}_c$ are the number of video frames and text captions given to a specific pair of vision and text encoders (as in Fig. 1), S_∞ represents the theoretical saturation score corresponding to ideal alignment of this vision/text model pair, whereas C_f, C_c, α and β are fitted scalar parameters. This formulation achieves remarkably high coefficients of determination for both VideoMAEv2 ($R^2=0.9838$) and DINOv2 ($R^2=0.9905$), with strong predictive power for test-time data scaling against the Gemma-2 text model.

The fitted parameters are also highly informative: VideoMAEv2 ($S_{\infty} \approx 0.40$, $C_f = 0.11$, $C_c = 0.14$, $\alpha = 1.22$, $\beta = 1.31$) versus DINOv2 ($S_{\infty} \approx 0.37$, $C_f = 0.06$, $C_c = 0.13$, $\alpha = 1.78$, $\beta = 1.36$). Notably, the frame coefficient (C_f) for VideoMAEv2 is nearly double that of DINOv2, while the caption coefficients (C_c) remain comparable, highlighting the video model's stronger ability to leverage temporal information from additional frames to improve alignment.

This parametric analysis resembles compute-optimal scaling laws, such as those identified for pretraining, e.g., in Hoffmann et al. (2022), which model the dependency of a loss function (e.g., Negative Log-Likelihood) on compute or data size. Whereas those laws are formulated in terms of a loss metric (where lower is better), our metric is an alignment score (where higher is better), making our additive saturation model conceptually equivalent. The saturation score S_{∞} (e.g., 0.40 for VideoMAEv2) can be interpreted as the base accuracy of an ideal alignment process, while the subtracted terms represent the respective error penalties incurred by the approximation of providing only a finite data at inference time.

We note that such predictive models for test-time scaling can be useful to design strategies for multimodal data acquisition (e.g., when collecting multiple high-quality annotations of video data, which can be costly), as well as to compare the ability of different encoders to incorporate diverse data modalities. We include test-time scaling laws for other vision models in the appendix, Sec. A.5.

7 Cross-modal alignment and downstream performance

Previous works, including Huh et al. (2024) and Maniparambil et al. (2024a) have highlighted that multimodal alignment between images and text is correlated with performance on *unimodal* tasks. I.e., image models that perform well on pure vision tasks (e.g., depth estimation) tend to also align well with language. Our goal is to evaluate whether a similar claim holds for video encoders.

We take advantage of recent work that trained large-scale unimodal video encoders and evaluated their performance on various semantic and non-semantic downstream tasks. We focus on video models trained without explicit text supervision, in order to better analyze the *emergent* alignment between representations, as opposed to models explicitly trained for video-language alignment (e.g. through video-text contrastive losses). In particular, we consider the best performing video models

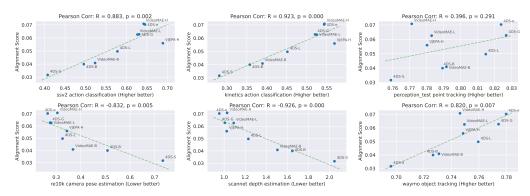


Figure 4: Correlation between video/text alignment and downstream video perception task performance, for SSL methods trained without text supervision.

reported in Carreira et al. (2024) and correlate their vision-text *feature alignment* against Gemma 2-9b-it on the VATEX dataset to their accuracy on video analysis tasks: point tracking on the Perception Test dataset (Patraucean et al., 2023), box tracking on the Waymo Open dataset (Sun et al., 2020), camera pose estimation on RealEstate10k (Zhou et al., 2018b), ScanNet depth estimation (Dai et al., 2017), and action classification on the SSv2 (Goyal et al., 2017) and Kinetics-700-2020 (Kay et al., 2017; Smaira et al., 2020) datasets. Downstream accuracy is computed by training a separate learnable attention-based decoder on top of the frozen video features for each task.

Figure 4 summarizes our key results. Observe that generally for self-supervised video models, there is a strong positive correlation between the cross-modal alignment scores and semantic tasks performance such as action classification on SSv2 and Kinetics. Interestingly, there is also a significant correlation between the alignment score and the accuracy on non-semantic perception tasks such as camera pose estimation, depth prediction, and object tracking. The point tracking task represents a notable exception, with weak correlation between text alignment and downstream performance. This can potentially be due to the highly *local* nature of the point tracking task, and also suggests room for improvement in terms of truly general purpose video encoders.

Overall, these results suggest that video-text alignment could potentially be used as a powerful zeroshot metric for probing the quality of video representations as an alternative or complementary to more expensive evaluation techniques that require repeatedly training multiple cross-modal decoders during self-supervised video model development.

8 TEMPORAL ANALYSIS & CROSS-MODEL ALIGNMENT

Given that both modalities we focus on follow the arrow of time, we set out to investigate their temporal abilities via cross-modal alignment. While LLMs have been trained natively to handle long sequences of text, video models traditionally are trained on short 16-32 frame videos lasting 2-4 seconds. Any understanding over longer horizons is therefore emergent. We analyze temporal alignment on two datasets: the simple, synthetic Test of Time (Bagad et al., 2023) and the challenging long-form VideoComp (Kim et al., 2025).

Test of Time. We first evaluate temporal alignment on the synthetic dataset from Test of Time (Bagad et al., 2023), specifically designed to probe the *temporal* abilities of video models. The dataset contains 180 synthetic (video, caption) pairs, with captions of the form "A c_1 circle appears {after, before} a c_2 circle," where c_1 , c_2 are selected colors, and the videos are 30 second 1fps clips enacting the caption with the shapes appearing in the same region of a black square (upper left, upper right, bottom left, bottom right). For each choice of shape and $\{c_1, c_2\}$, there are four pairs of related captions: two for each choice of preposition and color order. Within these, there are two sets of logically equivalent captions (e.g., c_1 before c_2 and c_2 after c_1), for which their videos may be identical or different depending on the location of the shapes.

We use video-text alignment on this dataset to understand how sensitive text and vision models are to this temporal reordering.

We find that when we take k = 3 neighbors, many models get near perfect alignment as expected, since each example has 3 distinctly closer neighbors (Figure 5). However, the order in which each modality ranks these three neighbors differs significantly, as shown by the differing k = 1, 2 alignments. We find that the language models tend to rank the neighbor which has the same words but ordered differently first, i.e. for "A c_1 square appears after a c_2 square", the closest neighbor will be "A c_2 square appears after a c_1 square." The other two related captions, " c_1 before c_2 " and " c_2 before c_1 ", the logically equivalent statement, are equally close to the original caption. Thus, in this case, LLMs measure closeness more akin to a bag-of-words than being temporally sensi-

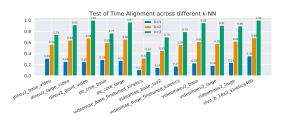


Figure 5: Video and text eventually are aligned, but encode temporal information differently. While the video text alignment is largely perfect for all of the models once k=3, they differ heavily before that when looking at k=1,2. Text models are often bag-of-words, while video models differ.

tive, at least at the shallower layers from where we are extracting the features.

VideoComp. We also test multimodal alignment on VideoComp (Kim et al., 2025). We use VideoComp's test set, sourced from YouTube and based on ActivityNet Captions and YouCook2 (Krishna et al., 2017; Zhou et al., 2018a). We consider videos at most 2 minutes long, and use the "segment mismatch" negatives which pair partial video segments with mismatched captions of alternative segments in the same video. This results in 285 video-caption pairs, so we use a lower k=3 for these experiments. To test the sensitivity to such negatives, we first compute the standard alignment using the text embeddings of positive captions. Then, we recompute neighbors in the text space, by considering the *negative* of each caption, say \tilde{c}_i , and computing its neighbors to the positive captions of other videos. Since negative captions are still from the same overall video, just temporally shuffled, this tests if the model is still aligned with the content in a temporally-agnostic manner. As shown in Figure 10 (appendix), the alignment with the negative captions is significantly lower than with the positive captions. Notably, models with larger alignment suffer more of a drop, suggesting that these models have well aligned structures which are adequately perturbed by the negative. These results again highlight that better alignment correlates with better downstream semantic understanding.

Cross-model alignment. We also investigated the alignment of features across different *video models*, and report the results in the Appendix A.6. We observe that text-supervised and self-supervised models, predictably, form clusters in pairwise alignment. However, interestingly, we note that some models (such as DINOv2) are able to *span* multiple clusters, and we hypothesize that such alignment against multiple models provides a strong indicator of the versatility of a given vision model. We leave a complete investigation of this phenomenon as interesting future work.

9 CONCLUSION, LIMITATIONS, AND FUTURE WORK

In this work, we conduct the first comprehensive study extending the Platonic Representation Hypothesis into the temporal domain. We demonstrate that alignment scores dramatically improve – doubling in some cases – simply by utilizing multiple video frames and diverse caption sets at inference, without any retraining. We quantify this phenomenon with a novel, highly accurate $(R^2>0.98)$ saturation-based scaling law, which quantitatively confirms that native video models (like VideoMAEv2) are fundamentally more able to leverage temporal information than static encoders (like DINOv2).

Through an extensive analysis involving 98 vision and language models, we show that alignment against text encoders strongly correlates with downstream performance on both semantic and non-semantic tasks, suggesting that vision-text alignment can be used as an informative zero-shot metric to guide video model development. At the same time, our analysis sheds light on some limitations of existing pure video foundation models, showing that many video models are outperformed by image models applied frame-by-frame.

REPRODUCIBILITY STATEMENT

All our results can be reproduced using models and datasets hosted on HuggingFace or open-sourced by their authors; see Section A.2 in the Appendix for full list of models.

REFERENCES

- Kelsey Allen, Carl Doersch, Guangyao Zhou, Mohammed Suhail, Danny Driess, Ignacio Rocco, Yulia Rubanova, Thomas Kipf, Mehdi SM Sajjadi, Kevin Murphy, et al. Direct motion models for assessing generated videos. *arXiv preprint arXiv:2505.00209*, 2025. 19
- Mido Assran, Adrien Bardes, David Fan, Quentin Garrido, Russell Howes, Matthew Muckley, Ammar Rizvi, Claire Roberts, Koustuv Sinha, Artem Zholus, et al. V-jepa 2: Self-supervised video models enable understanding, prediction and planning. *arXiv preprint arXiv:2506.09985*, 2025.
- Piyush Bagad, Makarand Tapaswi, and Cees G. M. Snoek. Test of Time: Instilling Video-Language Models with a Sense of Time. In CVPR, 2023. 8
- Adrien Bardes, Quentin Garrido, Jean Ponce, Xinlei Chen, Michael Rabbat, Yann LeCun, Mahmoud Assran, and Nicolas Ballas. Revisiting feature prediction for learning visual representations from video. *arXiv preprint arXiv:2404.08471*, 2024. 3
- Daniel Bolya, Po-Yao Huang, Peize Sun, Jang Hyun Cho, Andrea Madotto, Chen Wei, Tengyu Ma, Jiale Zhi, Jathushan Rajasegaran, Hanoona Rasheed, et al. Perception encoder: The best visual embeddings are not at the output of the network. *arXiv preprint arXiv:2504.13181*, 2025. 5, 7, 15
- 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*, pp. 9650–9660, 2021. 5
- João Carreira, Dilara Gokay, Michael King, Chuhan Zhang, Ignacio Rocco, Aravindh Mahendran, Thomas Albert Keck, Joseph Heyward, Skanda Koppula, Etienne Pot, et al. Scaling 4d representations. arXiv preprint arXiv:2412.15212, 2024. 3, 8, 19
- Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 5828–5839, 2017. 8
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. *arXiv e-prints*, pp. arXiv–2407, 2024. 1
- Marc O Ernst and Heinrich H Bülthoff. Merging the senses into a robust percept. *Trends in cognitive sciences*, 8(4):162–169, 2004.
- Pascale Fung, Yoram Bachrach, Asli Celikyilmaz, Kamalika Chaudhuri, Delong Chen, Willy Chung, Emmanuel Dupoux, Hongyu Gong, Hervé Jégou, Alessandro Lazaric, et al. Embodied ai agents: Modeling the world. *arXiv preprint arXiv:2506.22355*, 2025. 1
- Morgane Gemma Team: Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, et al. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*, 2024. 5
- Raghav Goyal, Samira Ebrahimi Kahou, Vincent Michalski, Joanna Materzynska, Susanne Westphal, Heuna Kim, Valentin Haenel, Ingo Fruend, Peter Yianilos, Moritz Mueller-Freitag, et al. The" something something" video database for learning and evaluating visual common sense. In *Proceedings of the IEEE international conference on computer vision*, pp. 5842–5850, 2017. 5, 8
- Chunhui Gu, Chen Sun, David A Ross, Carl Vondrick, Caroline Pantofaru, Yeqing Li, Sudheendra Vijayanarasimhan, George Toderici, Susanna Ricco, Rahul Sukthankar, et al. Ava: A video dataset of spatio-temporally localized atomic visual actions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 6047–6056, 2018.

- Souhail Hadgi, Luca Moschella, Andrea Santilli, Diego Gomez, Qixing Huang, Emanuele Rodolà,
 Simone Melzi, and Maks Ovsjanikov. Escaping plato's cave: Towards the alignment of 3d and
 text latent spaces. In *Proceedings of the Computer Vision and Pattern Recognition Conference*,
 pp. 19825–19835, 2025. 1
 - Yana Hasson, Pauline Luc, Liliane Momeni, Maks Ovsjanikov, Guillaume Le Moing, Alina Kuznetsova, Ira Ktena, Jennifer J Sun, Skanda Koppula, Dilara Gokay, et al. Scivid: Crossdomain evaluation of video models in scientific applications. *arXiv preprint arXiv:2507.03578*, 2025. 3
 - Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, Tom Hennigan, Eric Noland, Katie Millican, George van den Driessche, Bogdan Damoc, Aurelia Guy, Simon Osindero, Karen Simonyan, Erich Elsen, Jack W. Rae, Oriol Vinyals, and Laurent Sifre. Training compute-optimal large language models, 2022. 7
 - Minyoung Huh, Brian Cheung, Tongzhou Wang, and Phillip Isola. Position: The platonic representation hypothesis. In Hal Daumé III and Aarti Singh (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 20617–20642. PMLR, 21–27 Jul 2024. URL https://proceedings.mlr.press/v235/huh24a.html. 1, 2, 3, 4, 5, 7, 15
 - Rishi Jha, Collin Zhang, Vitaly Shmatikov, and John X Morris. Harnessing the universal geometry of embeddings. *arXiv preprint arXiv:2505.12540*, 2025. 1, 3
 - Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017. 5, 8
 - Dahun Kim, AJ Piergiovanni, Ganesh Mallya, and Anelia Angelova. Videocomp: Advancing fine-grained compositional and temporal alignment in video-text models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 29060–29070, 2025. 8, 9, 18
 - Dahye Kim, Jungin Park, Jiyoung Lee, Seongheon Park, and Kwanghoon Sohn. Language-free training for zero-shot video grounding. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pp. 1114–1123, 2023. 3
 - Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey Hinton. Similarity of neural network representations revisited. In *International conference on machine learning*, pp. 3519–3529. PMIR, 2019. 2
 - Ranjay Krishna, Kenji Hata, Frederic Ren, Li Fei-Fei, and Juan Carlos Niebles. Dense-captioning events in videos. In *International Conference on Computer Vision (ICCV)*, 2017. 9
 - Dhireesha Kudithipudi, Mario Aguilar-Simon, Jonathan Babb, Maxim Bazhenov, Douglas Blackiston, Josh Bongard, Andrew P Brna, Suraj Chakravarthi Raja, Nick Cheney, Jeff Clune, et al. Biological underpinnings for lifelong learning machines. *Nature Machine Intelligence*, 4(3):196–210, 2022. 1
 - Chau-Yi Li, Chia-Hao Li, Pin-Yu Chen, and Chen-Kuo Chiang. ELIOT: A zero-shot generative video-to-text retrieval framework, 2025. 3
 - Han Liu, Yinwei Wei, and Liqiang Nie. Dreaming user multimodal representation for micro-video recommendation, 2024a. 3
 - Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, et al. Mmbench: Is your multi-modal model an all-around player? In *European conference on computer vision*, pp. 216–233. Springer, 2024b. 1
 - Mayug Maniparambil, Raiymbek Akshulakov, Yasser Abdelaziz Dahou Djilali, Mohamed El Amine Seddik, Sanath Narayan, Karttikeya Mangalam, and Noel E O'Connor. Do vision and language encoders represent the world similarly? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14334–14343, 2024a. 1, 7

- Mayug Maniparambil, Raiymbek Akshulakov, Yasser Abdelaziz Dahou Djilali, Mohamed El Amine Seddik, Sanath Narayan, Karttikeya Mangalam, and Noel E O'Connor. Do vision and language encoders represent the world similarly? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14334–14343, 2024b. 2, 3
- Jack Merullo, Louis Castricato, Carsten Eickhoff, and Ellie Pavlick. Linearly mapping from image to text space. *arXiv preprint arXiv:2209.15162*, 2022. 2
- Maxime Oquab, Timothée Darcet, Theo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Russell Howes, Po-Yao Huang, Hu Xu, Vasu Sharma, Shang-Wen Li, Wojciech Galuba, Mike Rabbat, Mido Assran, Nicolas Ballas, Gabriel Synnaeve, Ishan Misra, Herve Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. DINOv2: Learning robust visual features without supervision, 2023. 2, 3, 6
- Viorica Patraucean, Lucas Smaira, Ankush Gupta, Adria Recasens, Larisa Markeeva, Dylan Banarse, Skanda Koppula, Mateusz Malinowski, Yi Yang, Carl Doersch, et al. Perception test: A diagnostic benchmark for multimodal video models. Advances in Neural Information Processing Systems, 36:42748–42761, 2023. 8
- 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*, pp. 8748–8763. PmLR, 2021. 1
- Dominik Schnaus, Nikita Araslanov, and Daniel Cremers. It's a (blind) match! towards vision-language correspondence without parallel data. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 24983–24992, 2025. 1, 3
- Oriane Siméoni, Huy V Vo, Maximilian Seitzer, Federico Baldassarre, and Maxime et al. Oquab. Dinov3. *arXiv preprint arXiv:2508.10104*, 3(4), 2025. 5
- Lucas Smaira, João Carreira, Eric Noland, Ellen Clancy, Amy Wu, and Andrew Zisserman. A short note on the kinetics-700-2020 human action dataset. *arXiv preprint arXiv:2010.10864*, 2020. 8
- Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for autonomous driving: Waymo open dataset. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 2446–2454, 2020. 8
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023. 1
- Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, et al. Gemma 3 technical report. *arXiv preprint arXiv:2503.19786*, 2025. 1
- Megan Tjandrasuwita, Chanakya Ekbote, Liu Ziyin, and Paul Pu Liang. Understanding the emergence of multimodal representation alignment. *arXiv preprint arXiv:2502.16282*, 2025. 1, 2
- Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. *Advances in neural information processing systems*, 35:10078–10093, 2022. 3, 5, 19
- Limin Wang, Bingkun Huang, Zhiyu Zhao, Zhan Tong, Yinan He, Yi Wang, Yali Wang, and Yu Qiao. Videomae v2: Scaling video masked autoencoders with dual masking. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 14549–14560, 2023a. 2
- Limin Wang, Bingkun Huang, Zhiyu Zhao, Zhan Tong, Yinan He, Yi Wang, Yali Wang, and Yu Qiao. VideoMAE v2: Scaling video masked autoencoders with dual masking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14549–14560, 2023b. 1, 3, 5, 6, 19

- Xin Wang, Jiawei Wu, Junkun Chen, Lei Li, Yuan-Fang Wang, and William Yang Wang. Vatex: A large-scale, high-quality multilingual dataset for video-and-language research. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 4581–4591, 2019. 5, 6
- Yao-Chih-Lee Wu, Yi-Ting Chen, Yi-Hsuan Tsai, and Wen-Hsiang Liao. VideoREPA: A video-based representation alignment framework for physics-enhanced video generation, 2025. 3
- Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3733–3742, 2018. 5
- Junlin Xie, Zhihong Chen, Ruifei Zhang, Xiang Wan, and Guanbin Li. Large multimodal agents: A survey. *arXiv preprint arXiv:2402.15116*, 2024. 1
- Hu Xu, Gargi Ghosh, Po-Yao Huang, Dmytro Okhonko, Armen Aghajanyan, Florian Metze, Luke Zettlemoyer, and Christoph Feichtenhofer. VideoCLIP: Contrastive pre-training for zero-shot video-text understanding. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 6787–6800. Association for Computational Linguistics, 2021. URL https://aclanthology.org/2021.emnlp-main.544.3
- Shukang Yin, Chaoyou Fu, Sirui Zhao, Ke Li, Xing Sun, Tong Xu, and Enhong Chen. A survey on multimodal large language models. *National Science Review*, 11(12):nwae403, 2024. 1
- Le Zhang, Qian Yang, and Aishwarya Agrawal. Assessing and learning alignment of unimodal vision and language models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 14604–14614, 2025. 3, 6
- Long Zhao, Nitesh B Gundavarapu, Liangzhe Yuan, Hao Zhou, Shen Yan, Jennifer J Sun, Luke Friedman, Rui Qian, Tobias Weyand, Yue Zhao, et al. Videoprism: A foundational visual encoder for video understanding. *arXiv preprint arXiv:2402.13217*, 2024. 3, 5, 18, 19
- Luowei Zhou, Chenliang Xu, and Jason Corso. Towards automatic learning of procedures from web instructional videos. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018a. 9
- Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. Stereo magnification: Learning view synthesis using multiplane images. *arXiv preprint arXiv:1805.09817*, 2018b. 8

A APPENDIX

A.1 ADDITIONAL RESULTS

We include in Figure 6 the alignment scores of all the vision and text encoders considered in our analysis on the VaTeX dataset – 98 models in total; see the full list in Section A.2. The rows correspond to vision encoders and the columns correspond to language models. The label for each model has the format: <model_name> (<max_alignment_score>) <num_layers>. Each ij entry in the matrix contains the alignment score (also colour-coded: lighter shades correspond to higher alignment scores) and a tuple (layer_i, layer_j) indicating from which layers in vision_model_i and language_model_j, respectively, the representations were extracted to get the best alignment for that pair; we recommend zooming in for better readability. We used a single caption for each video (out of 10) sampled randomly.

As mentioned in Section 6, the alignment scores improve significantly when considering more video frames and/or more text captions. In Figure 7, we show additional results that confirm this observation. We calculate alignment scores between visual representations produced by a subset of the vision encoders considered above and language representations given by different variants of Gemma text models, when processing a single caption (left) vs. multiple captions (right) from the VATEX dataset. As before, the rows correspond to vision encoders and the columns correspond to language models. Note that if we only have access to the final model output, num_layers in the model label is marked as 1. For vision models, we consider: (1) image models applied on the middle frame of each video, (2) image models applied on every frame of the video and the representations averaged

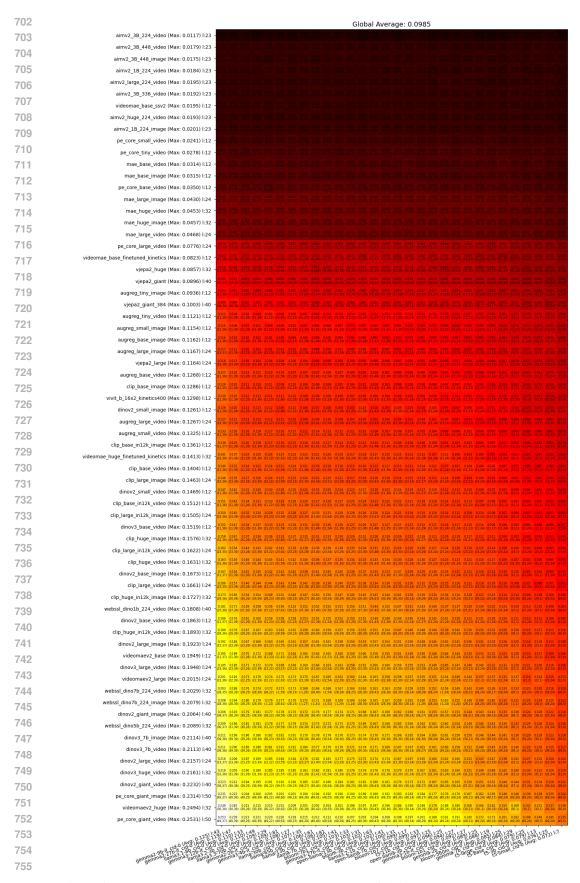


Figure 6: We measure alignment between a large number of vision and text encoders on the VATEX dataset.

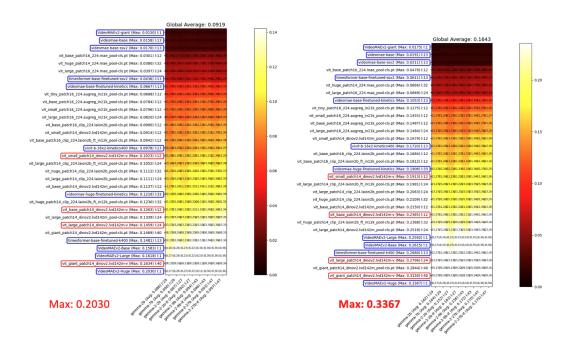


Figure 7: Video-text alignment results when using a single caption (left) vs. multiple captions (right) when measuring alignment. We observe a significant boost when integrating information across multiple captions. Notably, the best results on the right significantly exceed those reported in Huh et al. (2024), highlighting that the paucity of annotations in both visual space (images vs. videos) and text space (single caption vs. multiple descriptions) can help to explain the limited alignment observed in prior work.

over time (marked with red boxes), and (3) video models (marked with blue boxes). When using a single caption (left), we observe that the best alignment scores are generally obtained by video models, which are able to encode not only features about the scene appearance, but also details about the scene dynamics, leading to better alignment with caption representations. A notable exception is the DINOv2 image model applied on all the frames of the video, which obtains the second best alignment score. However, when we consider all the ten captions available for each video (right), the alignment scores improve significantly overall, and several image models applied on the full videos or even on single frames outperform many video models.

These findings hold on other datasets as well. We show the effect of adding more frames or using multiple captions at test time on the PVD dataset (Bolya et al., 2025). As the original dataset comes with a single detailed caption, we use an LLM to synthesize multiple captions; prompt and examples included in Section A.3. Figure 8 summarizes the results. On the left, we plot the alignment scores as we add more synthesized captions for different vision models. The dashed lines indicate the baseline alignment score of each vision model when considering the original detailed caption. We can observe that for all models, using more than 4 or 5 captions leads to better alignment compared to the baseline score. This saturates around 6 or 7 captions for most models, pointing to potential redundancy between the synthesized captions beyond this number. On the right, we show the effect of adding more frames.

A.2 FULL LIST OF MODELS

We use the following 68 open-sourced visual models and their variants for our analysis. All the models are available on HuggingFace or from the GitHub repositories of the authors.

```
aimv2_3B_224_video, aimv2_3B_448_video, aimv2_3B_448_image, aimv2_1B_224_video, aimv2_large_224_video, aimv2_3B_336_video, timesformer_base_k400, timesformer_base_ssv2, videomaev2_base, videomaev2_large, videomaev2_huge, videomae_base_ssv2, videomae_base_finetuned_kinetics, videomae_huge_finetuned_kinetics, vivit_b_16x2_kinetics400, dinov2_small_video, dinov2_small_image, dinov2_base_video, dinov2_base_image,
```

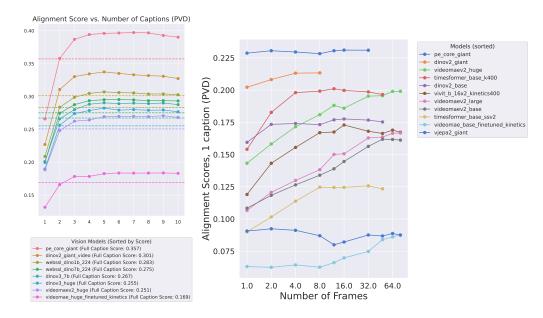


Figure 8: **Test Time Scaling for PVD works well.** Scaling the number of captions at test time outperforms using the original provided caption, and scaling with more clips is important to improve upon simple frame number scaling.

```
dinov2_large_video, dinov2_large_image, dinov2_giant_video, dinov2_giant_image, augreg_tiny_video, augreg_tiny_image, augreg_small_video, augreg_small_image, augreg_base_video, augreg_base_image, augreg_large_video, augreg_large_image, mae_base_video, mae_base_image, mae_large_video, mae_large_image, mae_huge_video, mae_buge_image, clip_base_video, clip_base_image, clip_large_video, clip_large_image, clip_base_image, clip_base_image, clip_base_image, clip_large_image, pe_core_tiny_video, pe_core_small_video, pe_core_base_video, pe_core_large_video, pe_core_large_video, dinov3_base_video, dinov3_large_video, dinov3_buge_video, dinov3_bvideo, dinov3_bimage, webssl_dino1b_224_video, webssl_dino1b_224_video, webssl_dino1b_224_video, vjepa2_giant_vijepa2_large, vjepa2_large, vjepa2_giant_vijepa2_giant_vijepa2_giant_vijeo.
```

For language models, we experiment with the following 30 open-sourced models and their variants:

```
bloom-560m, bloom-1b1, bloom-1b7, bloom-3b, bloom-7b1, open-llama-3b, open-llama-7b, open-llama-13b, llama-7b, llama-30b, llama-65b, gemma-2b, gemma-7b, gemma2-2b-it, gemma2-9b-it, gemma2-2b-it, gemma3-1b-it, gemma3-4b-it, gemma3-12b-it, gemma3-27b-it, t5-small, t5-base, t5-large, t5-3b, t5-11b, llama3.2-1b, llama3.2-3b, llama3.1-8b, llama3.3-70b.
```

A.3 SYNTHETIZING CAPTIONS FOR PVD DATASET USING LLMS

To study the effect of including an increasing number of captions on the alignment score between video representations and the language representation of the associated captions, we synthesized multiple captions for videos in the PVD dataset using LLMs, starting from the detailed caption provided in the original dataset. We relied on Gemini 2.5Pro. For reproducibility purposes, we provide below the prompt used to generate these captions, which was obtained from the prompt given to the VLM model when collecting the annotations for the original PVD dataset. We also include a qualitative example in Figure 9.

Create ten concise captions of a video using the provided detailed captions. Each caption should mention different details present. The goal is not to summarize all of the information 10 different ways, but to get 10 different descriptions all potentially containing different information.

TASK: Extract key information from the captions and combine it into a set of ten captions, each of which is a single phrase or set of phrases that includes some subset of the relevant details in alt text format.

Steps to Follow:

- 1. Review the caption for general context.
- 2. Extract the most relevant and concise information.
- 3. Don't include all of the key information in every caption. Pick some information to leave out between captions.
- 4. Combine extracted information into a alt text format using short phrase or set of phrases with approximately 120 tokens, considering special characters like comma as part of the token count.
- 5. Minimize the use of special characters and including everything in each caption.
- 6. Pick only a few things to include in each caption, it's okay to leave details out. What to Avoid:
- Avoid adding or inferring information not present in the original metadata and captions.
- Avoid using complex sentence structures or prioritizing sentence flow.

Return a list separated only by new lines, with only the captions, nothing else.

Figure 9 shows video frames, the original caption, and the synthesized captions for a video from the PVD dataset.



Human caption

The video shows a person making a burger. The person is adding mayonnaise to one bun and then placing different vegetables and a patty on top of it after that applies ketchup on another bun. The person is also using wooden boards to keep the ingredients and the burger. The person is wearing a red and white checkered shirt.

Synthesized captions

- 1. Person making a burger.
- 2. Adding mayonnaise and vegetables to a burger.
- 3. Burger assembled on a wooden board.
- 4. Tutorial showing how to make a burger.
- 5. Person in a red and white checkered shirt making a burger.
- 6. Well-lit video with vibrant colors showing burger preparation.
- 7. A person makes a burger using a wooden board.
- 8. Tutorial on adding mayonnaise and vegetables to a burger.
- 9. Person in a red and white checkered shirt preparing a burger.
- 10. Great tutorial for learning how to make a burger.

Figure 9: Example of video frames, the original human-provided caption, and LLM-synthesized captions for a video in the PVD dataset.

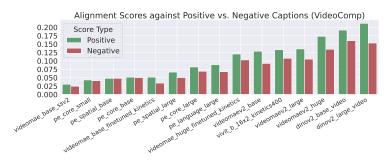
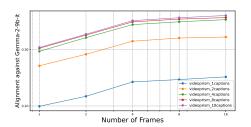
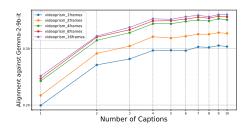


Figure 10: Video embeddings are sensitive to temporal caption reordering. Video-text alignment drops against the segment mismatch negatives in VideoComp, with the most aligned video representations showing the largest drop in alignment with the negatives (models towards the right).





(a) Alignment Score vs. Number of Frames for Video-Prism

(b) Alignment Score vs. Number of Captions for VideoPrism

Figure 11: Analysis of vision-text alignment sensitivity to the amount of visual and textual data available at test time for the VideoPrism model on the VATEX dataset

A.4 TEMPORAL ALIGNMENT IN VIDEOCOMP

Using the same mutual k-NN metric, we probe if the visual and text representations produced by different vision and language encoders are sensitive to temporal order and / or temporal alignment. We use VideoComp dataset (Kim et al., 2025), which contains tuplets of (video, positive_caption, negative_caption), and measure alignment for each video against its associated positive and negative captions; see Figure 10. We observe that the alignment with negatives is significantly lower. Importantly, as the visual-language alignment with the positive captions is stronger (towards the right of the figure), the gap in alignment with the negative captions increases, confirming that the mutual k-NN alignment metric correlates with temporal understanding.

A.5 TEST-TIME SCALING LAWS FOR OTHER MODELS AND DATASETS

Figure 11 shows the test-time scaling results for VideoPrism encoder Zhao et al. (2024) on the VATEX dataset.

R-squared: 0.9832

•
$$S_{\infty} = 0.40, C_f = 0.09, \alpha = 0.74, C_c = 0.14, \beta = 1.21$$

We also include some test time scaling laws for other models using PVD as our alignment dataset.

- dinov2_base_video ($R^2=0.9310$): $-S_{\infty}=0.242, C_f=0.031, \alpha=1.647, C_c=0.063, \beta=3.825$
- dinov2_giant_video ($R^2=0.8736$):
 - $S_{\infty}=0.310,$ $C_{f}=0.028,$ $\alpha=0.919,$ $C_{c}=0.087,$ $\beta=3.461$

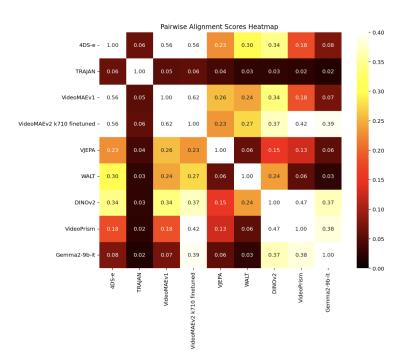


Figure 12: Video encoders cluster into semantic and geometric specialities. We find that video-video alignment exhibits two clusters of models: those which are closer aligned to language and semantic tasks (bottom right), and those which are closer to geometric abilities (top left). However, some models are able to span both axes, notably DINOv2 and VideoMAEv2 K710 finetuned.

```
 \begin{array}{l} \bullet \text{ timesformer\_base\_k400} \ (R^2=0.9079); \\ -S_\infty=0.285, C_f=0.066, \alpha=1.648, C_c=0.082, \beta=2.889 \\ \bullet \text{ timesformer\_base\_ssv2} \ (R^2=0.9589); \\ -S_\infty=0.172, C_f=0.056, \alpha=0.812, C_c=0.039, \beta=1.864 \\ \bullet \text{ videomaev2\_base} \ (R^2=0.9644); \\ -S_\infty=0.451, C_f=0.303, \alpha=0.065, C_c=0.052, \beta=2.277 \\ \bullet \text{ videomaev2\_huge} \ (R^2=0.9509); \\ -S_\infty=0.291, C_f=0.096, \alpha=0.484, C_c=0.074, \beta=2.573 \\ \bullet \text{ videomaev2\_large} \ (R^2=0.9722); \\ -S_\infty=0.292, C_f=0.146, \alpha=0.194, C_c=0.054, \beta=2.128 \\ \bullet \text{ vivit\_b\_16x2\_kinetics400} \ (R^2=0.9279); \\ -S_\infty=0.237, C_f=0.070, \alpha=1.124, C_c=0.064, \beta=2.971 \\ \end{array}
```

A.6 CROSS-MODEL (PAIRWISE VIDEO-VIDEO) ALIGNMENT)

Figure 12 shows the pairwise feature alignment between a selection of video models, as well as the Gemma2 text encoder. We include both MAE-based models Tong et al. (2022), a Kinetics-finetuned VideoMAEv2-G Wang et al. (2023b), self-supervised 4DS-e model Carreira et al. (2024) as well as language-supervised model Zhao et al. (2024), and a recent pure motion encoder TRAJAN Allen et al. (2025). We observe that there are several clusters appearing across video models, which seem to reflect language-alignment and spatio-temporal awareness (especially among MAE-based models). We hypothesize that although language alignment along is correlated with downstream performance, *cross-model* alignment might hold more predictive power. Specifically, strong, general-purpose models must align across multiple clusters to be useful in a broad range of downstream (including pixel or patch-level) applications.