

VL-JEPA: Joint Embedding Predictive Architecture for Vision-language

Delong Chen^{1,2,*} Mustafa Shukor^{1,3,*} Théo Moutakanni^{1,*} Willy Chung^{1,3,*}
 Jade Yu¹, Tejaswi Kasarla¹, Allen Bolourchi¹, Yann LeCun^{1,4}, Pascale Fung^{1,2}

¹ Meta FAIR ² HKUST ³ Sorbonne Université ⁴ NYU

* Equal contribution

delong.chen@connect.ust.hk

We introduce VL-JEPA, a vision-language model built on a Joint Embedding Predictive Architecture (JEPA). Instead of autoregressively generating tokens as in classical VLMs, VL-JEPA predicts *continuous embeddings* of the target texts. By learning in an abstract representation space, the model focuses on task-relevant semantics while abstracting away surface-level linguistic variability. In a strictly controlled comparison against standard token-space VLM training with the same vision encoder and training data, VL-JEPA achieves stronger performance while having 50% fewer trainable parameters. At inference time, a lightweight text decoder is invoked only when needed to translate VL-JEPA predicted embeddings into text. We show that VL-JEPA natively supports *selective decoding* that reduces the number of decoding operations by $\sim 2.85\times$ while maintaining similar performance compared to non-adaptive uniform decoding. Beyond generation, the VL-JEPA’s embedding space naturally supports open-vocabulary classification, text-to-video retrieval, and discriminative VQA *without any architecture modification*. On eight video classification and eight video retrieval datasets, the average performance VL-JEPA surpasses that of CLIP, SigLIP2, and Perception Encoder. At the same time, the model achieves comparable performance as classical VLMs (InstructBLIP, QwenVL) on four VQA datasets: GQA, TallyQA, POPE and POPEv2, despite only having 1.6B parameters.

1 Introduction

One of the most important aspects of advanced machine intelligence is the ability to understand the physical world that surrounds us. This ability enables AI systems to learn, reason, plan and act in the real world in order to assist humans [LeCun, 2022]. Intelligent systems that need to act in the real world includes wearable devices and robots [Fung et al., 2025]. Machine learning tasks that make up for this ability include captioning, retrieval, visual question answering, action tracking, reasoning and planning etc [Bordes et al., 2024, Chen et al., 2025b]. Systems for such real-world applications must have real-time response with low latency and inference cost.

Currently, the common approach to achieve these tasks is to use large token-generative Vision Language Models (VLMs) [Liu et al., 2023, Dai et al., 2023, Alayrac et al., 2022, Chen et al., 2024b, Cho et al., 2025, Chen et al., 2022], which takes visual input X_V , textual query X_Q to generate desired textual response Y autoregressively in token space, *i.e.*, $(X_V, X_Q) \mapsto Y$. This is straightforward but inadequate for two main reasons. First, VLMs are expensive to develop, because they are trained to generate responses Y to queries by capturing both task-relevant semantics with task-irrelevant surface linguistic features such as words choice, style or paraphrasing. During training, VLMs must model both aspects, which results in unnecessary

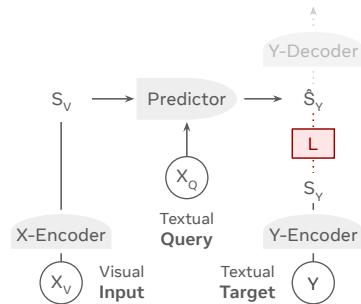


Figure 1. VL-JEPA model architecture

computing effort spent producing diverse token sequences that ultimately do not impact the correctness of the output. Second, real-time tasks involving live streaming video (*e.g.*, live action tracking) require sparse and selective decoding (*e.g.*, emitting a description only when a new event occurs) [Zhou et al., 2024]. However, VLMs rely on autoregressive token-by-token decoding, which must be completed before revealing the underlying semantics of Y . This process introduces unnecessary latency and hampers the ability to update semantics dynamically in real time.

This paper introduces the Joint Embedding Predictive Architecture for Vision-Language (VL-JEPA), turning expensive learning of data-space token generation into more

efficient latent-space semantic prediction. As illustrated in Fig. 1, the model employs **x-encoder** to map vision inputs X_V into embedding S_V , a **y-encoder** to map the textual target Y into an embedding S_Y , and a **predictor** that learns the mapping $(S_V, X_Q) \mapsto S_Y$ where X_Q is a textual query (*i.e.*, the prompt). The training objective is defined in the embedding space $\mathcal{L}_{\text{VL-JEPA}} = D(\hat{S}_Y, S_Y)$ instead of the data space $\mathcal{L}_{\text{VLM}} = D(\hat{Y}, Y)$. During inference, a **y-decoder** reads out the predicted embedding \hat{S}_Y to text space \hat{Y} when needed.

Thanks to its **non-generative** nature, VL-JEPA is not forced to reconstruct every surface detail of Y in the token space. Instead, it only needs to predict the abstract representation S_Y in the embedding space. In the raw one-hot token space, different plausible Y outputs for the same input often appear nearly orthogonal if they don't share overlapping tokens. However, in the embedding space, these diverse targets can be mapped to nearby points that share similar semantics. This simplifies the target distribution thus makes the learning process more efficient. In addition, unlike VLMs, this approach eliminates the need for learning language generation with a heavy decoder during training, resulting in significant efficiency gains.

Thanks to its **non-autoregressive** nature, VL-JEPA can produce continuous streams of target semantic embeddings within sliding windows with minimal latency as it only require a single forward pass without autoregressive decoding. This is particularly advantageous for real-time online applications such as live action tracking, scene recognition, or planning, where the embedding stream can be selectively decoded by a lightweight y-decoder, enabling efficient and prompt updates.

In this work, we empirically validate the advantages of VL-JEPA. We conduct a strictly controlled comparison against classical token-generative VLM [Liu et al., 2023, Cho et al., 2025]: both setups use the same vision encoder, spatial resolution, frame rate, training data, batch size, and number of iterations, etc., with the *only* difference being the objective in token space or embedding space. Under this matched training condition, VL-JEPA delivers consistently higher performance on zero-shot captioning and classification while using roughly half the trainable parameters, indicating that embedding-space supervision improves learning efficiency.

Beyond the training phase, VL-JEPA also delivers substantial inference-time efficiency improvement through *selective decoding*, where decoding happens only due to significant change in the predicted embedding stream. Empirically, this strategy reduces the number of decoding operations by $\sim 2.85\times$ while preserving overall output quality measured by average CIDEr scores.

Our final VL-JEPA models are trained in two stages: 1) a pretraining stage using caption data to establish robust

vision-language alignment, and 2) a supervised finetuning (SFT) stage that equips the model with VQA capabilities. The model resulting from the first stage, denoted as **VL-JEPA_{BASE}**, is evaluated on *zero-shot* classification and text-to-video retrieval. VL-JEPA_{BASE} outperforms CLIP [Radford et al., 2021], SigLIP2 [Tschanne et al., 2025], and Perception Encoder [Bolya et al., 2025] models in terms of average classification accuracy (across 8 datasets) and retrieval recall@1 (across 8 datasets). Following the second stage, the resulting **VL-JEPA_{SFT}** demonstrates significantly improved classification performance due to its exposure to in-domain training data. As a unified *generalist* model, VL-JEPA_{SFT} approaches the performance of *specialist* models optimized for individual benchmarks. Simultaneously, VL-JEPA_{SFT} exhibits effective VQA capabilities, achieving performance on par with established VLM families, such as InstructBLIP [Dai et al., 2023] and Qwen-VL [Bai et al., 2023], across four datasets covering compositional visual reasoning [Hudson and Manning, 2019], complex object counting [Acharya et al., 2019], and object hallucination [Li et al., 2023b, 2025b].

In summary, the contributions of this paper are as follows:

- We introduce VL-JEPA, the first non-generative model that can perform general-domain vision-language tasks in real-time, built on a joint embedding predictive architecture.
- We demonstrate in controlled experiments that VL-JEPA, trained with latent space embedding prediction, outperforms VLMs that rely on data space token prediction.
- We show that VL-JEPA delivers significant efficiency gains over VLMs for online video streaming applications, thanks to its non-autoregressive design and native support for selective decoding.
- We highlight that our VL-JEPA_{SFT} model, with an unified model architecture, can effectively handle a wide range of classification, retrieval, and VQA tasks at the same time.

2 Methodology

We propose **VL-JEPA** (Fig. 1), a model with the joint embedding predictive architecture (JEPA) for vision-language tasks. VL-JEPA is trained with triplets $\langle X_V, X_Q, Y \rangle$, where X_V denotes the **visual input** (a single image or a sequence of video frames), X_Q is a **textual query** (*i.e.*, a question) and Y is the **textual target** (*i.e.*, the answer) to be predicted. The VL-JEPA comprises of four components:

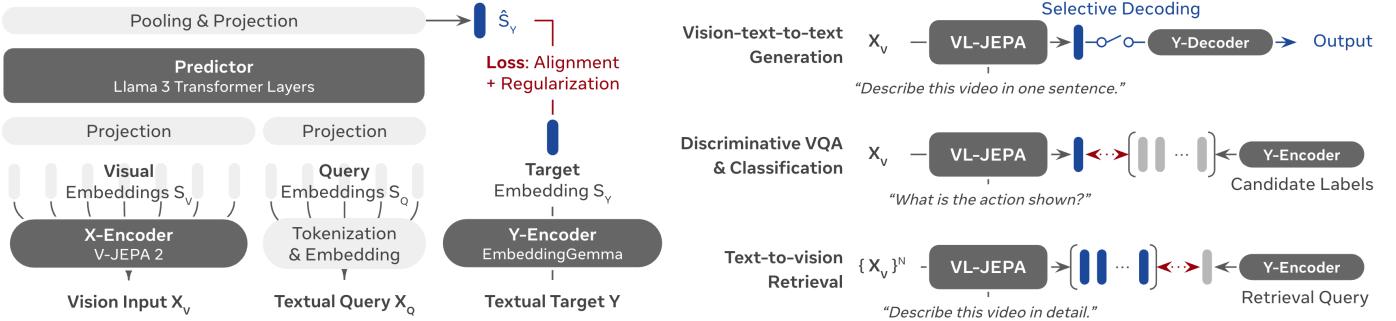


Figure 2. Left: VL-JEPA Architecture. It learns to predict the target embedding S_Y , instead of reconstructing the raw target Y in token space as in classical VLMs. Right: VL-JEPA Applications: It handles vision-text-to-text generation tasks (e.g., captioning) with selective decoding mechanism natively supported. Furthermore, VL-JEPA’s embedding space facilitates discriminative VQA, open-vocabulary classification and text-to-video retrieval tasks using a single unified model architecture.

1. **X-Encoder** ($X_V \mapsto S_V$) compresses high-volume visual inputs to compact visual embeddings—a sequence of continuous vectors analogous to “visual tokens” in classical VLMs.
2. **Predictor** ($(S_V, X_Q) \mapsto \hat{S}_Y$) is the core component of VL-JEPA. It maps visual embeddings to a prediction of target embedding, with a textual query as conditioning.
3. **Y-Encoder** ($Y \mapsto S_Y$) embeds the textual target into a continuous latent space as the prediction target. The target embedding is expected to abstract away task irrelevant information.
4. **Y-Decoder** ($\hat{S}_Y \mapsto \hat{Y}$) is not involved during the main training phrase of VL-JEPA. At inference time, it translates the predicted embedding as human-readable text when necessary.

Fig. 2 illustrates how we instantiate the VL-JEPA architecture in this paper. For the X-Encoder, we chose V-JEPA 2 [Assran et al., 2025], a Vision Transformer that outputs a sequence of visual tokens, which are then projected and fed into the Predictor initialized using Llama 3 Transformer layers. Query conditioning is achieved by tokenizing and embedding the textual query and feeding the resulting textual token embeddings into the Predictor along with the visual embeddings. The outputs of the Llama 3 Transformer layers are pooled and projected into the target embedding space produced by the Y-Encoder, which is initialized by EmbeddingGemma-300M [Vera et al., 2025]. We provide more technical details in §??.

Training Objective. JEPA models typically optimize two objectives jointly: 1) prediction error in the embedding space, and 2) additional regularization that avoids representation collapse [Bardes et al., 2021, Balestrieri and LeCun, 2025]. Any loss that implements these two properties can be applied to VL-JEPA. Alternatively, the

regularization term can be replaced by other anti-collapse strategies, such as using an exponential moving average (EMA) for the Y-Encoder [Assran et al., 2025] or freezing the Y-Encoder [Zhou et al., 2025].

In this work, we adopt the **InfoNCE loss** [Radford et al., 2021] due to its maturity in the vision-language domain. More advanced non-sample-contrastive regularization, such as VICReg [Bardes et al., 2021] and SIGReg [Balestrieri and LeCun, 2025] can also be applied but we leave the exploration to future works. InfoNCE loss can be mathematically divided [Wang and Isola, 2020] into: 1) a representation *alignment* term that minimizes the distance between normalized prediction and target embeddings, and 2) a *uniformity* regularization term that pushes embeddings in a batch apart from each other, thus avoiding representation collapse. We train the Predictor and the Y-Encoder jointly with bi-directional InfoNCE loss, enabling them to mutually learn from each other.

Compared to the token-space loss used by generative VLMs, calculating the training loss in the embedding space is beneficial due to the **simplified target distribution**. Specifically, many real-world prediction tasks are inherently ill-posed: for the same input X , there may exist multiple plausible targets Y that are all acceptable. For example, given the query “What will happen here if I flip this light switch down?”, both “the lamp is turned off” and “room will go dark” are valid answers. In the raw one-hot token space, however, the two sequences are orthogonal since they share no overlapping tokens. But when VL-JEPA’s Y-Encoder embeds them into nearby points (ideally yielding a compact unimodal distribution), the learning task becomes much easier: the model no longer needs to fit multiple disjoint high-density regions in sparse token space, but only a single coherent mode in a continuous embedding space.

Multi-tasking. VL-JEPA supports diverse tasks using a *single, unified* architecture (Fig. 2). For vision-text-to-text

generation tasks, such as captioning or open-ended VQA, the query X_Q is a captioning prompt or a question, and the predictor learns to predict the embedding of the target output, \hat{S}_Y , which is then decoded into text. VL-JEPA also supports CLIP-style open-vocabulary classification and discriminative VQA, where candidate label texts are encoded into embeddings and compared with prediction \hat{S}_Y to select the nearest match. For text-to-video retrieval, candidate videos are mapped to their predicted embeddings \hat{S}_Y using a retrieval a captioning prompt, and then ranked by similarity to the encoded textual retrieval query.

Selective Decoding. Real-world video applications often require online streaming inference, such as tracking user actions in smart glasses for procedural assistance [Chen et al., 2024c], monitoring world states for online planning, navigation and robotics [Shukor et al., 2025, Black et al., 2025, Song et al., 2025]. A central challenge is balancing two competing needs: the model must continuously update semantics as new frames arrive, but computational efficiency and latency are critical.

Existing VLMs typically rely on explicit memory mechanisms [Zhou et al., 2024, Qian et al., 2024] to decide when to decode or complex KV-cache optimizations [Di et al., 2025] for efficiency, since autoregressive language models are expensive to run continuously. VL-JEPA, in contrast, natively supports selective decoding. Since it predicts a semantic answer embedding non-autoregressively, the model provides a continuous semantic stream of \hat{S}_Y that can be monitored in real time. This stream can be stabilized with simple smoothing (e.g., average pooling) and decoded only when a significant semantic shift is detected, such as when the local window variance exceeds a threshold. In this way, VL-JEPA maintains always-on semantic monitoring while avoiding unnecessary decoding, achieving both responsiveness and efficiency.

3 Implementation of VL-JEPA

3.1 Model Architecture

X-Encoder. Unless otherwise specified, we use a frozen V-JEPA 2 ViT-L [Assran et al., 2025] with 304M parameters, a self-supervised vision model that excels at both image and video tasks. Each video input is uniformly sampled into frames at 256^2 resolution. For image inputs, the same image is duplicated to match the input shape.

Predictor. The predictor is initialized with the last 8 Transformer layers of Llama-3.2-1B, resulting in 490M trainable parameters. The text tokenizer and token embedding are also from Llama-3.2-1B. We allow maximum 512 query tokens, and put [PAD] tokens for short queries. We disable the causal attention mask so that both vision and query embeddings can be jointly attended. Linear projections connect the predictor with the vision and text

embeddings, and average pooling on non-[PAD] tokens is applied to obtain the predicted target embedding.

Y-Encoder. We use EmbeddingGemma-300M [Vera et al., 2025] as the initialization of the Y-Encoder. We set maximum context length of 512 to handle detailed captions. We found that setting a learning rate multiplier of $\times 0.05$ to all text encoder parameters improves performance, since the quality of embedding prediction would be suboptimal in the beginning of training. Linear projection head is applied to both Predictor and Y-Encoder, obtaining a shared embedding space with 1,536 dimensions, where the loss is calculated.

3.2 Two-stage Training

Large-scale Pretraining. VL-JEPA is trained with two stages. The first query-free pretraining stage aims to establish robust vision-language alignment using massive caption data. We use PLM-Image-Auto [Cho et al., 2025], Datacomp [Gadre et al., 2023] and YFCC-100M [Thomee et al., 2016] for image-text data. For video-text data, we include PLM-Video-Auto [Cho et al., 2025], Ego4D atomic action descriptions [Grauman et al., 2022], and an internal dataset Action100M consisting captions generated on HowTo100M videos [Chen et al., 2025b].

We first do image-only training on Datacomp and YFCC-100M with only 1 frame per visual input, which allows us to use a large batch size of 24k. After 100k iterations, the model has seen 2B samples and achieved 61.6% ImageNet zero-shot accuracy (without prompt ensembling). Then, we continue with joint image-video pretraining with 16 frames per input. The pretraining takes 2 weeks using 24 nodes with 8×NVIDIA H200 GPUs each. We adopt a constant learning rate of 5×10^{-5} to facilitate extended training. We call the resulting model **VL-JEPA_{BASE}** and measure zero-shot classification and retrieval performance with this model.

Supervised Finetuning. The second query-conditioned supervised finetuning (SFT) stage empowers VL-JEPA VQA capabilities while maintaining the pretrained vision-language alignment for classification and retrieval. The training data is selected from the PLM data mixture [Cho et al., 2025], including 25M VQA samples, 2.8M captioning samples, 1.8M classification samples, and downsampled pretraining stage data to avoid catastrophic forgetting.

We train the model for 35k steps with a batch size of 6k (~2 days with 24 nodes), with cosine learning rate annealing applied to improve convergence. Since excessive human labelled data is included in this SFT data mixture, we no longer emphasize zero-shot evaluation for the resulting **VL-JEPA_{SFT}** from this stage. Instead, we evaluate VQA capabilities and compare it with state-of-the-art *specialist* models.

Table 1. Video classification and text-to-video retrieval. Best *zero-shot* performance in each dataset are **highlighted**. Samples seen = training step \times effective batch size.

Model	# Parameters	# Samples Seen	Zero-shot	Generalist Model	Video Classification (Top-1 Accuracy)								Text-to-video Retrieval (Recall@1)									
					Average	SSv2	EK100	EgoExo4D	Kinetics-400	COIN (SR)	COIN (TR)	CrossTask (SR)	CrossTask (TR)	Average	MSR-VTT	ActivityNet	DiDeMo	MSVD	YouCook2	PVD-Bench	Dream-1k	VDC-1k
CLIP	RN50	75M	12.8B	✓ ✓	21.8	2.1	1.5	1.9	41.4	8.6	39.0	10.9	68.7	28.3	28.7	17.7	24.7	29.7	5.1	27.6	47.2	46.0
	ViT-B	124M	12.8B		25.3	3.1	1.3	2.4	49.5	11.2	47.3	16.2	71.5	29.3	31.0	19.5	25.7	34.0	6.1	27.0	48.5	42.9
	ViT-L	389M	12.8B		30.9	3.8	3.7	3.6	58.3	14.7	63.5	20.8	78.5	35.3	35.9	23.4	30.7	41.9	7.9	36.7	56.8	49.3
SigLIP2	ViT-B	375M	40B	✓ ✓	33.9	5.2	2.3	4.9	57.8	20.6	69.9	27.7	82.9	39.6	40.2	25.0	32.1	48.6	13.8	52.1	60.9	43.7
	ViT-L	882M	40B		38.7	5.9	4.5	7.0	63.6	24.2	78.5	35.1	90.8	45.4	41.6	32.7	35.1	53.5	19.0	59.2	71.6	50.9
	ViT-g	1.9B	40B		39.9	6.1	6.1	6.4	68.0	26.0	80.4	35.1	90.8	47.5	43.4	33.9	38.9	56.0	22.2	60.4	73.0	52.5
PE-Core	ViT-B	448M	58B	✓ ✓	37.3	5.8	3.3	6.3	65.4	21.5	77.1	26.9	91.8	44.9	46.5	35.4	35.3	49.1	15.2	59.8	68.7	49.2
	ViT-L	671M	58B		42.8	9.3	6.0	10.9	73.4	27.1	83.3	37.5	95.3	50.2	48.9	41.7	40.8	56.2	22.5	64.7	75.9	51.0
	ViT-G	2.3B	86B		44.6	9.0	6.4	13.0	76.4	29.0	86.0	40.3	97.2	58.1	51.6	49.1	44.5	58.7	26.0	77.0	89.2	68.5
VL-JEPA_{BASE}	ViT-L	1.6B	2.0B	✓ ✓	46.4	16.1	13.3	21.1	57.8	39.8	74.4	60.5	88.0	58.4	37.6	55.4	49.2	47.9	23.1	78.2	88.8	87.2
VL-JEPA_{SFT}	ViT-L	1.6B	2.5B	✗ ✓	70.7	68.2	38.8	59.5	81.4	60.3	86.8	77.1	93.0	59.5	43.7	53.8	46.2	49.1	28.8	81.1	86.4	86.7
<i>SoTA (including specialist models)</i>				✗ ✗	-	77.5	56.4	47.8	92.1	67.3	95.3	64.5	96.0	-	62.8	74.1	74.2	61.4	28.9	77.0	89.2	68.5

4 Experiments

4.1 Classification and Retrieval

We begin by evaluating VL-JEPA’s classification and retrieval performance in §4.1, and benchmark VL-JEPA on VQA datasets in §4.2. We demonstrate application of VL-JEPA for understanding the relationship between world state changes and action concepts (*i.e.*, inverse dynamics) in §4.3. In §4.4, we demonstrate the advantage of embedding prediction by comparing it with a token-predictive VLM baseline under a strictly controlled setting. In §4.5, we evaluate the effectiveness of VL-JEPA’s selective decoding, and show that it reduces decoding cost while maintaining the performance. Next, we analyze VL-JEPA’s Y-Encoder in §4.6.

Evaluation Setup. We evaluate VL-JEPA following the CLIP-style evaluation protocol (see Fig.2 and §2 “Multitasking”). We assess VL-JEPA on a broad suite of benchmarks, including 8 classification datasets and 8 retrieval datasets. For *zero-shot* evaluation, we compare against *generalist foundation models* CLIP [Radford et al., 2021], SigLIP2 [Tschannen et al., 2025], and Perception Encoder (PE-Core)[Bolya et al., 2025]. We additionally report reference numbers from *specialist models* that are individually optimized for each benchmark (summarized in Appendix??).

Results. Table 1 summarizes the results. In the strict zero-shot setting, VL-JEPA_{BASE} achieves higher average accuracy (46.4 vs 44.6) across the 8 classification datasets and higher average recall@1 (58.4 vs 58.1) across the 8 retrieval datasets than the best baseline PE-Core-G. Per-dataset scores show that VL-JEPA_{BASE} is particularly strong on *motion-centric* benchmarks (SSv2, EK-100, EgoExo4D, and step recognition on COIN and CrossTask), while relatively weaker on *appearance-centric* benchmarks (Kinetics-400 and

task recognition on COIN and CrossTask). This is due to VL-JEPA_{BASE} has seen substantially fewer vision-language pairs (only 2B in comparison with PE-Core-G’s 86B). After supervised finetuning, VL-JEPA_{SFT} improves significantly upon VL-JEPA_{BASE} since the model has seen in-domain training data. As a single *generalist* model, the performance of VL-JEPA_{SFT} is approaching *specialist* models optimized individually for each dataset.

4.2 Visual Question Answering

Evaluation Setup. We evaluate VL-JEPA_{SFT} on discriminative VQA tasks. The inference process involves encode candidate answers using the Y-Encoder and selecting the answer that minimizes the distance to the predicted embedding (see Fig. 2). We select four benchmarks that prioritize visual perception rather than knowledge and reasoning. We evaluate on GQA [Hudson and Manning, 2019], a dataset for real-world visual reasoning and compositional QA, reporting accuracy on the testdev-balanced split. For TallyQA [Acharya et al., 2019], which targets complex counting, we follow Chen et al. [2022] and report the weighted average accuracy across the “simple” and “complex” splits. Finally, to assess object hallucination, we utilize POPE [Li et al., 2023b] and POPEv2 [Li et al., 2025b]. For POPE, we report the average accuracy across the “random”, “popular”, and “adversarial” settings on MS-COCO.

Results. Table 4.2 compares VL-JEPA_{SFT} against established VLM families, including BLIP-2 [Li et al., 2023a], InstructBLIP [Dai et al., 2023], Qwen-VL [Bai et al., 2023], InternVL [Chen et al., 2024d], Llava-1.5 [Vallaeyns et al., 2024], SmolVLM [Marafioti et al., 2025], PaLI [Chen et al., 2022], PaliGemma [Beyer et al., 2024], and Video-LLaVA [Lin et al., 2024]. VL-JEPA_{SFT} outperforms many of these

Table 2. VQA benchmarks. We report accuracy on GQA [Hudson and Manning, 2019], TallyQA [Acharya et al., 2019], POPE [Li et al., 2023b], and POPEv2 [Li et al., 2025b]. Scores lower than our model are marked in red. Scores from SmolVLM are obtained by our evaluation, while other baselines are reported in the literature.

GQA: compositional visual reasoning		TallyQA: complex object counting		POPE: object hallucination		POPEv2: object hallucination	
Model	Accuracy	Model	Accuracy	Model	Accuracy	Model	Accuracy
BLIP-2 (OPT-2.7B)	33.9	SmolVLM-256M	32.3	SmolVLM2-256M	56.4	SmolVLM-256M	62.3
BLIP-2 (FlanT5XXL)	41.0	SmolVLM-500M	44.8	SmolVLM-256M	57.9	LLaVA-1.5-13B	72.7
InstructBLIP (FlanT5XL)	48.4	PaLI-700M	62.3	LLaVA-7B	72.9	InternVL2-8B	74.5
InstructBLIP (Vicuna-13B)	49.5	SmolVLM-2B	64.7	InstructBLIP (Vicuna-13B)	79.0	InternVL2-26B	76.1
Qwen-VL-Chat-7B	57.5	PaLI-3B	65.8	Video-LLaVA (7B)	83.4	Qwen2-VL-72B	79.4
Qwen-VL-7B	59.3	InstructBLIP (Vicuna-13B)	68.0	SmolVLM-500M	85.8	SmolVLM-500M	83.8
InternVL-Chat (Vicuna-7B)	59.5	PaLI-17B	71.9	LLaVA-1.5-7B	85.9	Qwen2-VL-7B	87.0
LLaVA-1.5 (Vicuna-7B)	62.0	LLaVA-1.5 (Vicuna-13B)	72.3	LLaVA-1.5-13B-HD	86.3	SmolVLM-2B	88.8
InternVL-Chat (Vicuna-13B)	66.6	PaliGemma (3B)	76.8	SmolVLM-2B	87.5	Qwen2-VL-2B	91.3
VL-JEPA_{SFT} (1.6B)	60.8	VL-JEPA_{SFT} (1.6B)	67.4	VL-JEPA_{SFT} (1.6B)	84.2	VL-JEPA_{SFT} (1.6B)	82.2

Table 3. WORLD PREDICTION-WM benchmark results. We compare the accuracy between large VLMs, socratic LLMs, and VL-JEPA. VL-JEPA_{SFT} achieves a new SoTA at 65.7%.

Vision Language Models										Socratic LLMs (w/ Qwen2.5-VL-72B captions)								VL-JEPA					
InternVL2.5					Qwen2.5-VL					Llama-3.1		Llama-4		Qwen2.5		GPT-4o		Claude-3.5		Gemini-2		BASE	SFT
2B	4B	26B	38B	3B	7B	32B	72B	8B	70B	109B	400B	3B	7B	72B	N/A	N/A	N/A	N/A	1.6B	1.6B			
20.0	29.8	30.2	50.3	21.6	45.5	49.0	57.0	48.7	49.8	52.7	53.6	44.0	49.1	48.5	52.0	53.3	55.6	63.9	<u>65.7</u>				

baselines despite requiring significantly less computational resources—classical VLMs rely on extensively pretrained CLIP backbones combined with multi-stage visual instruction tuning. In comparison, VL-JEPA_{SFT} employs a *unified architecture* and a *single embedding space* to seamlessly handle VQA, classification, and retrieval (Tab. 1).

4.3 WorldPrediction-WM

Evaluation Setup. We evaluate VL-JEPA on the “world modeling” task in the WORLD PREDICTION [Chen et al., 2025a] benchmark, where the model is provided with two images representing the initial and final world states and must identify, among four candidate video clips, the action that explains the observed transition. To adapt VL-JEPA, we duplicate and concatenate the initial and final state images to extract a *state embedding*, and encode each action candidate into *action embeddings*. The model then selects the candidate whose embedding is closest to the state embedding.

Results. Table 3 shows accuracy comparisons. VL-JEPA_{BASE} attains 63.9% and VL-JEPA_{SFT} attains 65.7% top-1 accuracy on WORLD PREDICTION-WM, establishing a new state of the art. Our VL-JEPA model not only substantially surpasses existing VLMs of comparable or larger scale but also exceeds the performance of frontier LLMs such as GPT-4o, Claude-3.5-sonnet, and Gemini-2.0.

4.4 Embedding Prediction vs. Token Prediction: A Controlled Comparison

Evaluation Setup. In this section, we compare VL-JEPA to a token-generative VLM baseline under a strictly aligned training conditions. Both models use the same Perception

Encoder [Bolya et al., 2025] (frozen ViT-L-14 with 336² resolution, no tiling, 16 frames per video) for vision inputs. We use the same training iterations with the same effective batch size of 128, same learning rate scheduler on the same pretraining data mixture described above (§3). The only difference is the prediction task: VL-JEPA predicts target embeddings [Duquenne et al., 2023] using a 0.5B predictor, whereas the VLM baseline performs next-token prediction with cross-entropy using a 1B LLM. For VLM, we use the standard training recipe and codebase of PerceptionLM [Cho et al., 2025], aligning frozen vision encoder and text-only LLM Llama-3.2-1B. For VL-JEPA, we initialize the predictor from the 8-16 layers of Llama-3.2-1B.

We evaluate both models at regular checkpoints throughout training spanning from 500K to 15M samples seen. At each checkpoint, we measure the performance on video captioning and video classification. For video captioning, we report CIDEr scores averaged across YouCook2 [Zhou et al., 2018], MSR-VTT [Xu et al., 2016] and PVD-Bench [Bolya et al., 2025]. VL-JEPA decodes the predicted embeddings while VLM generates the tokens directly. For video classification, we report top-5 accuracy averaged across CrossTask-Step, CrossTask-Task [Zhukov et al., 2019] and EgoExo4D [Grauman et al., 2024]. For VL-JEPA we choose the candidate with lowest cosine distance to the predicted embedding, while for VLM we pick the class with lowest perplexity.

Results. As shown in Fig. 3, both models yield comparable performance after 500K samples seen in both tasks, with respectively 1.23 and 1.35 CIDEr in video captioning and 14.9% and 14.0% top-5 accuracy for VL-JEPA and

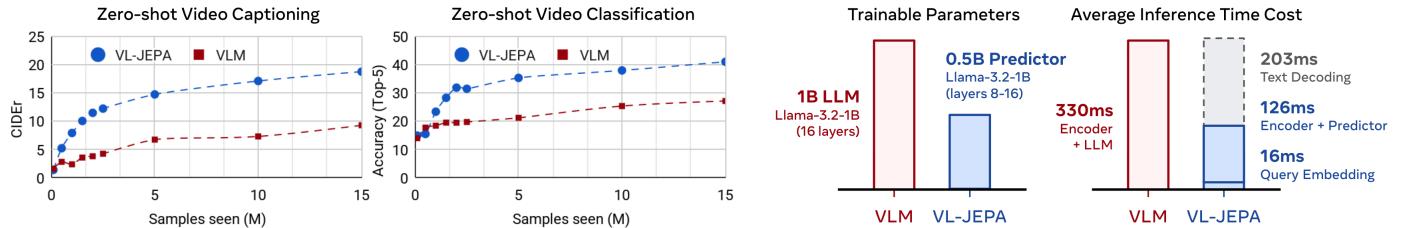


Figure 3. Comparison of embedding prediction (VL-JEPA) and token prediction (VLM). We conduct a fair comparison of under strictly aligned training settings (encoder, data, batchsize, etc.). **Left:** Zero-shot video captioning CIDEr score averaged over 3 datasets and zero-shot classification accuracy (top-5) averaged over 3 benchmarks. **Right:** Comparing the trainable parameters and average inference time cost.

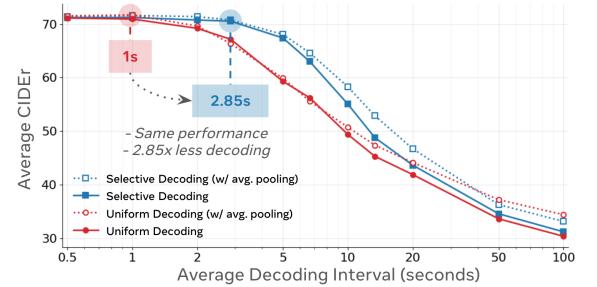
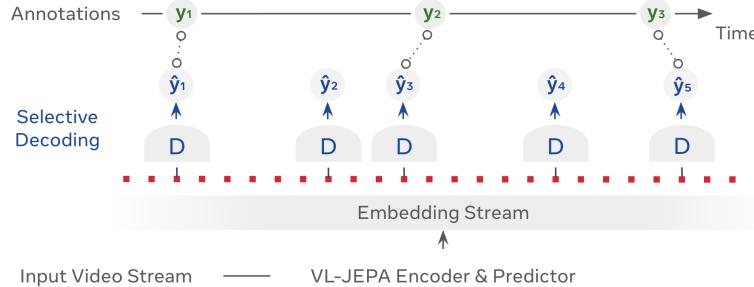


Figure 4. Evaluation of selective decoding. **Left:** We compare uniform sampling of decoding points at fixed intervals (red) and embedding-guided selective decoding (blue). Performance is measured by the average CIDEr score between each annotation y and its closest decoded output \hat{y} . **Right:** Results on EgoExo4D show that selective decoding achieves a Pareto improvement over uniform sampling: for the same performance level, it requires fewer decoding operations.

VLM. After a few iterations, we show that VL-JEPA’s performance increase is much sharper compared to VLM, reaching 14.7 CIDEr and 35.3% top-5 accuracy after 5M samples seen. This gap remains constant as training scales at 15M samples with 14.8 CIDEr and 41.0% top-5 accuracy for VL-JEPA, while the VLM baseline yield respectively 7.1 CIDEr and 27.2% top-5 accuracy. This controlled comparison highlights the benefit of predicting embeddings rather than tokens, showing both higher sample efficiency and stronger absolute performance.

We compare the inference cost of the above VL-JEPA and the VLM by pre-loading 64 video frames into memory and repeatedly decoding text 100 times with the same prompt, measuring the average time per sample. As shown in Fig. 3 (right most), both models exhibit comparable latency when generating text. What differentiates our model from classical VLM is the decoupling between the prompt processing (“Query Embedding”) and the video encoder (“Encoder + Predictor”) from the text generation module (“Decoder”). This allows us to only use the first part of the model to perform retrieval and decode text only when needed (see Section 4.5 below), making our model more scalable for online video inference.

4.5 Effectiveness of Selective Decoding

Evaluation Setup. We evaluate the effectiveness of VL-JEPA’s embedding-guided selective decoding on long-form video streams. To this end, we design a benchmark task

where the goal is to recover a temporal sequence of annotations while minimizing the number of text decoding operations, which dominate inference cost. As shown in Fig. 4 (left), decoding is performed only at selected points along the VL-JEPA embedding stream, yielding a sequence of N decoded outputs $[(\hat{t}_1, \hat{y}_1), (\hat{t}_2, \hat{y}_2), \dots, (\hat{t}_N, \hat{y}_N)]$. Each ground-truth annotation $[(t_1, y_1), (t_2, y_2), \dots, (t_T, y_T)]$ is then aligned to its nearest decoded output in time (illustrated as $\circ \dots \circ$ in Fig. 4), and CIDEr is computed between matched pairs. We use the EgoExo4D [Grauman et al., 2024] validation set in procedural activity domains, which consists of 218 videos with an average duration of 6 minutes and about $T = 143$ atomic action annotations per video.

As a baseline, we consider *uniform sampling*, where decoding points are placed at fixed intervals regardless of the underlying video content. Standard streaming VLMs are limited to this strategy, whereas VL-JEPA supports a more effective alternative: *adaptive selection* of decoding points guided by its predicted embeddings. We apply agglomerative clustering with temporal connectivity constraints [Murtagh and Contreras, 2012] to partition the embedding sequence into N segments of high intra-segment monosemantics [Chen et al., 2024a], measured by variance (i.e., Ward distance). The intuition is that within a semantically coherent segment, decoded outputs are highly similar, so decoding once per segment captures the essential information while greatly reducing overall decoding cost. The

Table 4. Comparison of text-encoders performance. We report triplet-based accuracy (%) on SugarCrepe++ and VISLA datasets.

Model	# Params. (total)	# Params. (text encoder)	Average	SugarCrepe++ [Dumpala et al., 2024a]					VISLA [Dumpala et al., 2024b]			
				Replace Attribute	Replace Object	Replace Relation	Swap Attribute	Swap Object	Average	Generic	Spatial	
CLIP	ViT-L	389M	85M	44.5	56.7	83.0	42.5	27.0	13.5	34.5	37.6	31.3
SigLIP2	ViT-g	1.9B	708M	56.5	66.9	74.4	52.1	58.4	30.6	40.4	48.7	32.1
PE-Core	ViT-G	2.3B	537M	58.6	73.6	90.6	48.9	53.2	26.5	38.3	45.2	31.4
VL-JEPA_{BASE}	ViT-L	1.6B	300M	63.9	72.2	90.1	52.2	62.9	42.0	42.9	49.8	35.9
VL-JEPA_{SFT}	ViT-L	1.6B	300M	58.4	68.5	90.9	47.4	55.4	29.8	39.5	44.8	34.2

Table 5. Ablation studies results. The default setting adopted by VL-JEPA is marked in blue. We calculate $\pm\delta$ within each group of ablations in comparison with the default setting.

	Classification (Accuracy)	Retrieval (Recall@1)	VQA (Accuracy)
VL-JEPA _{SFT}	59.1	70.6	53.2
<i>(a) Effectiveness of pretraining stage on caption data</i>			
w/ Pretraining	49.0	47.5	46.1
w/o Pretraining	27.3 <small>(-21.7)</small>	30.2 <small>(-17.3)</small>	42.5 <small>(-3.6)</small>
<i>(b) Learning rate multiplier for Y-Encoder</i>			
multiplier = 0.05	27.3	30.2	42.5
multiplier = 1.00	23.7 <small>(-3.6)</small>	28.8 <small>(-1.4)</small>	40.7 <small>(-1.8)</small>
multiplier = 0.10	26.9 <small>(-0.4)</small>	30.2 <small>(-0.0)</small>	42.9 <small>(+0.4)</small>
multiplier = 0.01	25.6 <small>(-1.7)</small>	27.7 <small>(-2.5)</small>	41.0 <small>(-1.5)</small>
multiplier = 0.00	20.0 <small>(-7.3)</small>	25.9 <small>(-4.3)</small>	41.4 <small>(-1.1)</small>
<i>(c) Loss function (with no projection head on top frozen text encoder)</i>			
InfoNCE	23.3	30.3	44.3
Cosine	16.5 <small>(-6.8)</small>	20.2 <small>(-10.1)</small>	46.6 <small>(+2.3)</small>
L1	14.8 <small>(-8.5)</small>	15.5 <small>(-14.8)</small>	41.9 <small>(-2.4)</small>
L2	13.5 <small>(-9.8)</small>	11.7 <small>(-18.6)</small>	43.7 <small>(-0.6)</small>

midpoint of each segment is then chosen as the decoding point, and decoding is performed either from the exact embedding or from the average-pooled embedding within the segment.

Results. As shown in Fig. 4 (right), we sweep the average decoding frequency from 2.0 Hz down to 0.01 Hz (*i.e.*, average intervals between consecutive decoding operations from 0.5s to 100s) by adjusting either the stride of uniform sampling or the number of clusters in adaptive selection. Across the entire range, adaptive selection consistently Pareto-dominates uniform sampling. In particular, selective decoding at 0.35 Hz (*i.e.*, \sim 2.85s interval) matches the performance of uniform decoding at 1 Hz, reducing decoding cost by \sim 2.85 \times . We further observe that average pooling provides consistent gains for both strategies, since it provides denoising and stabilization on embeddings prior feeding into the decoder.

4.6 Evaluation of Y-Encoder

Evaluation Setup. We evaluate whether the JEPA architecture improves the Y-Encoder by following the uni-modal text-only (TOT) evaluation setup. We use the hard-negative benchmarks SugarCrepe++ [Dumpala et al., 2024a] and VISLA [Dumpala et al., 2024b]. These datasets test sensitiv-

ity to semantic and lexical changes in image descriptions. Each dataset contains triplets: two semantically similar descriptions of the same image (p_1 and p_2), and one negative description (n) created by altering attributes, relations, or objects. We compare Y-Encoders from different models by computing the cosine similarity for all description pairs. We check that the similarity between positives $sim(p_1, p_2)$ is higher than both the similarity between each positive and the negative $sim(p_1, n)$ and $sim(p_2, n)$. We report accuracy (%) across all samples.

Results. Table 4 shows the performance of different models on text hard-negative benchmarks. VL-JEPA_{BASE} achieves a micro average accuracy of 63.9% on SugarCrepe++ and 42.9% on VISLA. This is higher than the best other models: PE-Core scores 58.6% on SugarCrepe++ and SigLIP2 scores 40.4% on VISLA. The finetuned VL-JEPA_{SFT} model also achieves competitive results, with 58.4% on SugarCrepe++ and 39.5% on VISLA. These results indicate that VL-JEPA_{BASE} has a Y-Encoder that is more resilient to text hard-negatives.

4.7 Ablation Study

Evaluation Setup. We study different design choices for VL-JEPA. Here we train all ablation models on the SFT stage

data for 10K steps with a batch size of 512 (5M samples seen) and constant learning rate. We report average classification top-1 accuracy of 8 datasets (Tab. 1), average text-to-video retrieval recall@1 of 8 datasets (Tab. 1), and average VQA accuracy of 4 datasets (CLEVR, GQA, TallyQA simple and complex). We report the results in Tab. 5.

Results. **(a) Pretraining.** Dropping the first query-free pretraining stage on image and video captions significantly hurt performance, especially on classification (-21.7) and retrieval (-17.3). **(b) LR Multiplier.** The sweet point of learning rate multiplier to the Y-Encoder is around 0.05 to 0.10. Either faster or slower learning degrades the performance. **(c) Loss Function.** InfoNCE generally give superior performance compared to cosine, L1, and L2 losses, with the only exception being cosine loss outperform InfoNCE on VQA. However, only InfoNCE has the anti-collapse regularization and can be applied with unfrozen Y-Encoder. **(d) Predictor.** In terms of predictor size, more layers yield better performance, especially on VQA performance. We also see that if using the original causal attention instead of updating to bi-direction attention hurt VQA performance (-1.9), since query tokens are appended after visual tokens, and visual tokens are no longer able to attend to query tokens. Finally, we also see that LLama-3 initialization is beneficial to VQA performance, although vision-language alignment (classification and retrieval) is a bit worse compared to randomly initialized Transformer layers. **(e) Y-Encoder.** We tried different text encoder as the Y-Encoder, and confirmed that VL-JEPA works well with other embedding models than EmbeddingGemma-300M. Generally, larger encoder leads to better performance, with visually aligned text encoders (PE models) has significant advantage in classification and retrieval.

5 Related Works

JEPA Models. JEPA model learns by predicting the representation of a target input Y from the representation of a context input X . Early instantiations include I-JEPA for image encoding [Assran et al., 2023] and V-JEPA for video encoding [Bardes et al., 2023], which demonstrated the effectiveness of this objective over pixel reconstruction approach in their respective modality. Recent JEPA work falls into two categories. One category of work emphasizes better unimodal representation learning [Assran et al., 2023, Bardes et al., 2023, Fei et al., 2023] or cross-modal alignment [Lei et al., 2025, Jose et al., 2025]. The other direction targets world modeling, where pretrained encoders are frozen and action-conditioned predictors are trained for conditional prediction of state representations [Zhou et al., 2025, Baldassarre et al., 2025, Assran et al., 2025]. This has shown good results but remains limited to narrow domains like mazes or robotic pick-and-place. Our proposed VL-JEPA is the first designed for general-purpose

vision-language tasks. It performs conditional latent prediction over vision and text, and preserves efficiency while enabling flexible, multitask architecture.

Vision Language Models. Existing vision-language models largely fall into two families: (1) CLIP-style models with a non-predictive joint-embedding architecture (JEA) [Radford et al., 2021, Zhai et al., 2023, Bolya et al., 2025, Liu et al., 2024, Chen et al., 2023] encode images and texts independently into a common latent space, $X_V \mapsto S_V$ and $Y \mapsto S_Y$. By minimizing $\mathcal{L}_{\text{CLIP}} = D(S_V, S_Y)$ with a contrastive loss (e.g., InfoNCE), CLIP learns aligned *representations* that support zero-shot classification and vision-language retrieval; (2) Generative VLMs [Liu et al., 2023, Chen et al., 2022, Dai et al., 2023, Alayrac et al., 2022, Chen et al., 2024b, Cho et al., 2025, Beyer et al., 2024] connect a vision encoder [Radford et al., 2021, Fini et al., 2025] with a language model (e.g., LLM). They are typically trained with $\mathcal{L}_{\text{VLM}} = D(\hat{Y}, Y)$, i.e., next token prediction with cross-entropy loss, and can learn to handle various vision-text-to-text generation tasks such as VQA.

Table 6. Task coverage comparison.

	CLIP	VLM	VL-JEPA
Generation	✗	✓	✓
Retrieval	✓	✗	✓

Our proposed VL-JEPA integrates the architectural advantages and task coverage of both CLIPs and VLMs (Table 6). Since VL-JEPA learns in embedding space, it can leverage web-scale noisy image–text pairs [Jia et al., 2021], yielding strong open-domain features. On the other hand, VL-JEPA supports conditional generation tasks with a readout text decoder. Meanwhile, compared to generative VLMs that optimize directly in data space, VL-JEPA is more efficient at learning in the latent space. In addition, it is also more efficient for online inference, as it allows naturally selective decoding.

Efficient Vision Language Models. The growing size and training cost of VLMs has motivated efforts to improve efficiency. On the training side, strong performance can be achieved by updating only a subset of parameters, such as the vision–language connector [Tsimpoukelli et al., 2021, Alayrac et al., 2022, Vallaey et al., 2024, Shukor et al., 2023, Koh et al., 2023, Merullo et al., 2022, Dai et al., 2023]. At inference, efficiency is pursued through pruning parameters or visual tokens [Cao et al., 2023, Shukor and Cord, 2024, Vasu et al., 2025]. For real-time use cases, recent work explores small VLMs [Yao et al., 2024, Marafioti et al., 2025] and heuristics to reduce query frequency in asynchronous inference [Shukor et al., 2025].

Latent-space Language Modeling. Current state-of-the-art LLMs are trained to decode and reason in text space

using autoregressive generation and chain-of-thought prompting [Wei et al., 2022]. Text-space LLMs have rapidly improved and now achieve strong results on a wide range of benchmarks. However, the discrete nature of their reasoning trace may limit both speed and performance in the long term. Several works have explored latent-space LLMs that process or reason in latent space, such as Large Concept Models [Barrault et al., 2024] and COCONUT [Hao et al., 2024]. These models focus on unimodal latent-space reasoning. With VL-JEPA, our goal is to align vision and text representations in a shared multi-modal latent space. This approach aims to enable better abstractions and improve both the performance and speed of vision-language models (VLMs). We hope VL-JEPA will serve as a foundation for future work on multi-modal latent space reasoning, including visual chain-of-thought methods [Li et al., 2025a].

6 Conclusion

We have presented VL-JEPA, a new vision-language model built upon the joint embedding predictive architecture. By shifting supervision from discrete token space to continuous semantic embedding space, VL-JEPA simplifies the learning target, avoids redundant modeling of surface linguistic variability, and enables non-autoregressive prediction. Through controlled experiments, we show that VL-JEPA outperforms generative VLMs trained with cross-entropy loss under matched training data budget, while achieving superior training efficiency and significantly lower inference latency. Beyond generation tasks, the embedding-based design further allows VL-JEPA to handle open-vocabulary classification and cross-modal retrieval within a single unified architecture. Its ability to emit continuous semantic embeddings also makes it particularly well suited for real-time video applications, where selective decoding can improve both responsiveness and efficiency. In this work, we demonstrated the advantages of VL-JEPA over standard VLMs, particularly in computational efficiency, streaming applications, and video-language tasks. Our goal at this stage, is not to propose a universal alternative to VLMs, as this would require broader evaluation on tasks such as reasoning, tool use, and agentic behaviors where current token generative VLMs excel. Finally, although our results show clear benefits from scaling parameters and dataset size, we did not fully explore this direction, leaving it for future work.

Acknowledgments

We would like to thank Yejin Bang, Adrien Bardes, Loïc Barrault, Lucas Beyer, Quentin Garrido, João Maria Janeiro, Yifu Qiu, Koustuv Sinha, Basile Terver, and François Yvon for providing valuable feedback and support to this work.

References

Manoj Acharya, Kushal Kafle, and Christopher Kanan. Tallyqa: Answering complex counting questions. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 8076–8084, 2019.

Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *Advances in neural information processing systems*, 35: 23716–23736, 2022.

Mahmoud Assran, Quentin Duval, Ishan Misra, Piotr Bojanowski, Pascal Vincent, Michael Rabbat, Yann LeCun, and Nicolas Ballas. Self-supervised learning from images with a joint-embedding predictive architecture. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15619–15629, 2023.

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.

Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. Qwen technical report. *arXiv preprint arXiv:2309.16609*, 2023.

Federico Baldassarre, Marc Szafraniec, Basile Terver, Vasil Khalidov, Francisco Massa, Yann LeCun, Patrick Labatut, Maximilian Seitzer, and Piotr Bojanowski. Back to the features: Dino as a foundation for video world models. *arXiv preprint arXiv:2507.19468*, 2025.

Randall Balestrieri and Yann LeCun. Lejepa: Provable and scalable self-supervised learning without the heuristics. *arXiv preprint arXiv:2511.08544*, 2025.

Adrien Bardes, Jean Ponce, and Yann LeCun. Vicreg: Variance-invariance-covariance regularization for self-supervised learning. *arXiv preprint arXiv:2105.04906*, 2021.

Adrien Bardes, Quentin Garrido, Jean Ponce, Xinlei Chen, Michael Rabbat, Yann LeCun, Mido Assran, and Nicolas Ballas. V-jepa: Latent video prediction for visual representation learning. 2023.

Loïc Barrault, Paul-Ambroise Duquenne, Maha Elbayad, Artyom Kozhevnikov, Belen Alastrauey, Pierre Andrews, Mariano Coria, Guillaume Couairon, Marta R Costa-jussà, David Dale, et al. Large concept models: Language

modeling in a sentence representation space. *arXiv preprint arXiv:2412.08821*, 2024.

Lucas Beyer, Andreas Steiner, André Susano Pinto, Alexander Kolesnikov, Xiao Wang, Daniel Salz, Maxim Neumann, Ibrahim Alabdulmohsin, Michael Tschannen, Emanuele Bugliarello, et al. Paligemma: A versatile 3b vlm for transfer. *arXiv preprint arXiv:2407.07726*, 2024.

Kevin Black, Manuel Y Galliker, and Sergey Levine. Real-time execution of action chunking flow policies. *arXiv preprint arXiv:2506.07339*, 2025.

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.

Florian Bordes, Richard Yuanzhe Pang, Anurag Ajay, Alexander C Li, Adrien Bardes, Suzanne Petryk, Oscar Mañas, Zhiqiu Lin, Anas Mahmoud, Bargav Jayaraman, et al. An introduction to vision-language modeling. *arXiv preprint arXiv:2405.17247*, 2024.

Qingqing Cao, Bhargavi Paranjape, and Hannaneh Hajishirzi. Pumer: Pruning and merging tokens for efficient vision language models. *arXiv preprint arXiv:2305.17530*, 2023.

Delong Chen, Zhao Wu, Fan Liu, Zaiquan Yang, Shaoqiu Zheng, Ying Tan, and Erjin Zhou. Protoclip: Prototypical contrastive language image pretraining. *IEEE Transactions on Neural Networks and Learning Systems*, 2023.

Delong Chen, Samuel Cahyawijaya, Jianfeng Liu, Baoyuan Wang, and Pascale Fung. Subobject-level image tokenization. *arXiv preprint arXiv:2402.14327*, 2024a.

Delong Chen, Jianfeng Liu, Wenliang Dai, and Baoyuan Wang. Visual instruction tuning with polite flamingo. In *Proceedings of the aaai conference on artificial intelligence*, volume 38, pages 17745–17753, 2024b.

Delong Chen, Willy Chung, Yebin Bang, Ziwei Ji, and Pascale Fung. Worldprediction: A benchmark for high-level world modeling and long-horizon procedural planning. *arXiv preprint arXiv:2506.04363*, 2025a.

Delong Chen, Theo Moutakanni, Willy Chung, Yebin Bang, Ziwei Ji, Allen Bolourchi, and Pascale Fung. Planning with reasoning using vision language world model. *arXiv preprint arXiv:2509.02722*, 2025b.

Joya Chen, Zhaoyang Lv, Shiwei Wu, Kevin Qinghong Lin, Chenan Song, Difei Gao, Jia-Wei Liu, Ziteng Gao, Dongxing Mao, and Mike Zheng Shou. Videollm-online: Online video large language model for streaming video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18407–18418, 2024c.

Xi Chen, Xiao Wang, Soravit Changpinyo, Anthony J Piergiovanni, Piotr Padlewski, Daniel Salz, Sebastian Goodman, Adam Grycner, Basil Mustafa, Lucas Beyer, et al. Pali: A jointly-scaled multilingual language-image model. *arXiv preprint arXiv:2209.06794*, 2022.

Zhe Chen, Jiannan Wu, Wenhui Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, et al. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 24185–24198, 2024d.

Jang Hyun Cho, Andrea Madotto, Effrosyni Mavroudi, Triantafyllos Afouras, Tushar Nagarajan, Muhammad Maaz, Yale Song, Tengyu Ma, Shuming Hu, Suyog Jain, et al. Perceptionlm: Open-access data and models for detailed visual understanding. *arXiv preprint arXiv:2504.13180*, 2025.

Wenliang Dai, Junnan Li, Dongxu Li, Anthony Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale N Fung, and Steven Hoi. Instructblip: Towards general-purpose vision-language models with instruction tuning. *Advances in neural information processing systems*, 36:49250–49267, 2023.

Shangzhe Di, Zhelun Yu, Guanghao Zhang, Haoyuan Li, Tao Zhong, Hao Cheng, Bolin Li, Wanggui He, Fangxun Shu, and Hao Jiang. Streaming video question-answering with in-context video kv-cache retrieval. *arXiv preprint arXiv:2503.00540*, 2025.

Sri Harsha Dumpala, Aman Jaiswal, Chandramouli Sastry, Evangelos Milios, Sageev Oore, and Hassan Sajjad. Sugarcrape++ dataset: vision-language model sensitivity to semantic and lexical alterations. In *Proceedings of the 38th International Conference on Neural Information Processing Systems, NIPS '24, Red Hook, NY, USA, 2024a. Curran Associates Inc. ISBN 9798331314385*.

Sri Harsha Dumpala, Aman Jaiswal, Chandramouli Sastry, Evangelos Milios, Sageev Oore, and Hassan Sajjad. Visla benchmark: Evaluating embedding sensitivity to semantic and lexical alterations. *arXiv preprint arXiv:2404.16365*, 2024b.

Paul-Ambroise Duquenne, Holger Schwenk, and Benoît Sagot. Sonar: sentence-level multimodal and language-

agnostic representations. *arXiv preprint arXiv:2308.11466*, 2023.

Zhengcong Fei, Mingyuan Fan, and Junshi Huang. A-jepa: Joint-embedding predictive architecture can listen. *arXiv preprint arXiv:2311.15830*, 2023.

Enrico Fini, Mustafa Shukor, Xiujun Li, Philipp Dufter, Michal Klein, David Haldimann, Sai Aitharaju, Victor G Turrisi da Costa, Louis Béthune, Zhe Gan, et al. Multimodal autoregressive pre-training of large vision encoders. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 9641–9654, 2025.

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.

Samir Yitzhak Gadre, Gabriel Ilharco, Alex Fang, Jonathan Hayase, Georgios Smyrnis, Thao Nguyen, Ryan Marten, Mitchell Wortsman, Dhruba Ghosh, Jieyu Zhang, et al. Datacomp: In search of the next generation of multimodal datasets. *Advances in Neural Information Processing Systems*, 36:27092–27112, 2023.

Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, et al. Ego4d: Around the world in 3,000 hours of egocentric video. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 18995–19012, 2022.

Kristen Grauman, Andrew Westbury, Lorenzo Torresani, Kris Kitani, Jitendra Malik, Triantafyllos Afouras, Kumar Ashutosh, Vijay Baiyya, Siddhant Bansal, Bikram Boote, et al. Ego-exo4d: Understanding skilled human activity from first-and third-person perspectives. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19383–19400, 2024.

Shibo Hao, Sainbayar Sukhbaatar, Dijia Su, Xian Li, Zhiting Hu, Jason Weston, and Yuandong Tian. Training large language models to reason in a continuous latent space. *arXiv preprint arXiv:2412.06769*, 2024.

Drew A Hudson and Christopher D Manning. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6700–6709, 2019.

Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and

Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International conference on machine learning*, pages 4904–4916. PMLR, 2021.

Cijo Jose, Théo Moutakanni, Dahyun Kang, Federico Baldassarre, Timothée Darcet, Hu Xu, Daniel Li, Marc Szafraniec, Michaël Ramamonjisoa, Maxime Oquab, et al. Dinov2 meets text: A unified framework for image-and pixel-level vision-language alignment. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 24905–24916, 2025.

Jing Yu Koh, Ruslan Salakhutdinov, and Daniel Fried. Grounding language models to images for multimodal inputs and outputs. In *International Conference on Machine Learning*, pages 17283–17300. PMLR, 2023.

Yann LeCun. A path towards autonomous machine intelligence. *Open Review*, 62(1):1–62, 2022.

Hongyang Lei, Xiaolong Cheng, Qi Qin, Dan Wang, Huazhen Huang, Qingqing Gu, Yetao Wu, and Luo Ji. M3-jepa: Multimodal alignment via multi-gate moe based on the joint-embedding predictive architecture. In *Forty-second International Conference on Machine Learning*, 2025.

Chengzu Li, Wenshan Wu, Huanyu Zhang, Yan Xia, Shaoguang Mao, Li Dong, Ivan Vulić, and Furu Wei. Imagine while reasoning in space: Multimodal visualization-of-thought. *arXiv preprint arXiv:2501.07542*, 2025a.

Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, pages 19730–19742. PMLR, 2023a.

Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Wayne Xin Zhao, and Ji-Rong Wen. Evaluating object hallucination in large vision-language models. *arXiv preprint arXiv:2305.10355*, 2023b.

Yifan Li, Kun Zhou, Wayne Xin Zhao, Lei Fang, and Ji-Rong Wen. Analyzing and mitigating object hallucination: A training bias perspective. *arXiv preprint arXiv:2508.04567*, 2025b.

Bin Lin, Yang Ye, Bin Zhu, Jiaxi Cui, Munan Ning, Peng Jin, and Li Yuan. Video-llava: Learning united visual representation by alignment before projection. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 5971–5984, 2024.

Fan Liu, Delong Chen, Zhangqinyun Guan, Xiaocong Zhou, Jiale Zhu, Qiaolin Ye, Liyong Fu, and Jun Zhou. Remoteclip: A vision language foundation model for remote sensing. *IEEE Transactions on Geoscience and Remote Sensing*, 62:1–16, 2024.

Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in neural information processing systems*, 36:34892–34916, 2023.

Andrés Marafioti, Orr Zohar, Miquel Farré, Merve Noyan, Elie Bakouch, Pedro Cuenca, Cyril Zakka, Loubna Ben Allal, Anton Lozhkov, Nouamane Tazi, et al. Smolvlm: Redefining small and efficient multimodal models. *arXiv preprint arXiv:2504.05299*, 2025.

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

Fionn Murtagh and Pedro Contreras. Algorithms for hierarchical clustering: an overview. *Wiley interdisciplinary reviews: data mining and knowledge discovery*, 2(1):86–97, 2012.

Rui Qian, Xiaoyi Dong, Pan Zhang, Yuhang Zang, Shuangrui Ding, Dahua Lin, and Jiaqi Wang. Streaming long video understanding with large language models. *Advances in Neural Information Processing Systems*, 37:119336–119360, 2024.

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.

Mustafa Shukor and Matthieu Cord. Skipping computations in multimodal llms. *arXiv preprint arXiv:2410.09454*, 2024.

Mustafa Shukor, Corentin Dancette, and Matthieu Cord. ep-alm: Efficient perceptual augmentation of language models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22056–22069, 2023.

Mustafa Shukor, Dana Aubakirova, Francesco Capuano, Pepijn Kooijmans, Steven Palma, Adil Zouitine, Michel Aractingi, Caroline Pascal, Martino Russi, Andres Marafioti, et al. Smolvla: A vision-language-action model for affordable and efficient robotics. *arXiv preprint arXiv:2506.01844*, 2025.

Wenxuan Song, Jiayi Chen, Pengxiang Ding, Han Zhao, Wei Zhao, Zhide Zhong, Zongyuan Ge, Jun Ma, and Haoang Li. Accelerating vision-language-action model

integrated with action chunking via parallel decoding. *arXiv preprint arXiv:2503.02310*, 2025.

Bart Thomee, David A Shamma, Gerald Friedland, Benjamin Elizalde, Karl Ni, Douglas Poland, Damian Borth, and Li-Jia Li. Yfcc100m: The new data in multimedia research. *Communications of the ACM*, 59(2):64–73, 2016.

Michael Tschannen, Alexey Gritsenko, Xiao Wang, Muhammad Ferjad Naeem, Ibrahim Alabdulmohsin, Nikhil Parthasarathy, Talfan Evans, Lucas Beyer, Ye Xia, Basil Mustafa, et al. Siglip 2: Multilingual vision-language encoders with improved semantic understanding, localization, and dense features. *arXiv preprint arXiv:2502.14786*, 2025.

Maria Tsimpoukelli, Jacob L Menick, Serkan Cabi, SM Es-lami, Oriol Vinyals, and Felix Hill. Multimodal few-shot learning with frozen language models. *Advances in Neural Information Processing Systems*, 34:200–212, 2021.

Théophane Vallaeyns, Mustafa Shukor, Matthieu Cord, and Jakob Verbeek. Improved baselines for data-efficient perceptual augmentation of llms. In *European Conference on Computer Vision*, pages 369–387. Springer, 2024.

Pavan Kumar Anasosalu Vasu, Fartash Faghri, Chun-Liang Li, Cem Koc, Nate True, Albert Antony, Gokula Santhanam, James Gabriel, Peter Grasch, Oncel Tuzel, et al. Fastvlm: Efficient vision encoding for vision language models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 19769–19780, 2025.

Henrique Schechter Vera, Sahil Dua, Biao Zhang, Daniel Salz, Ryan Mullins, Sindhu Raghuram Panyam, Sara Smoot, Iftekhar Naim, Joe Zou, Feiyang Chen, et al. Embeddinggemma: Powerful and lightweight text representations. *arXiv preprint arXiv:2509.20354*, 2025.

Tongzhou Wang and Phillip Isola. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *International conference on machine learning*, pages 9929–9939. PMLR, 2020.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837, 2022.

Jun Xu, Tao Mei, Ting Yao, and Yong Rui. Msr-vtt: A large video description dataset for bridging video and language. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5288–5296, 2016.

Yuan Yao, Tianyu Yu, Ao Zhang, Chongyi Wang, Junbo Cui, Hongji Zhu, Tianchi Cai, Haoyu Li, Weilin Zhao,

Zhihui He, et al. Minicpm-v: A gpt-4v level mllm on your phone. *arXiv preprint arXiv:2408.01800*, 2024.

Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language image pre-training. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 11975–11986, 2023.

Gaoyue Zhou, Hengkai Pan, Yann LeCun, and Lerrel Pinto. Dino-wm: World models on pre-trained visual features enable zero-shot planning. In *Forty-second International Conference on Machine Learning*, 2025.

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, 2018.

Xingyi Zhou, Anurag Arnab, Shyamal Buch, Shen Yan, Austin Myers, Xuehan Xiong, Arsha Nagrani, and Cordelia Schmid. Streaming dense video captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18243–18252, 2024.

Dimitri Zhukov, Jean-Baptiste Alayrac, Ramazan Gokberk Cinbis, David Fouhey, Ivan Laptev, and Josef Sivic. Cross-task weakly supervised learning from instructional videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3537–3545, 2019.