

InternVideo-Next: Towards World Understanding Video Models

Chenting Wang^{1,2} Yuhan Zhu^{2,5} Yicheng Xu²
Jiange Yang^{2,5} Ziang Yan² Yali Wang^{2,4} Yi Wang^{2,3} Limin Wang^{1,2,5,♠}

¹Shanghai Jiao Tong University ²Shanghai AI Laboratory ³Shanghai Innovation Institute

⁴Shenzhen Institutes of Advanced Technology, China ⁵Nanjing University

Abstract

Large-scale video-text pretraining achieves strong performance but depends on noisy, synthetic captions with limited semantic coverage, often overlooking implicit world knowledge such as object motion, 3D geometry, and physical cues. In contrast, masked video modeling (MVM) directly exploits spatiotemporal structures but trails text-supervised methods on general tasks. We find this gap arises from overlooked architectural issues: pixel-level reconstruction struggles with convergence and its low-level requirement often conflicts with semantics, while latent prediction often encourages shortcut learning. To address these, we disentangle the traditional encoder-decoder design into an Encoder-Predictor-Decoder (EPD) framework, where the predictor acts as a latent world model, and propose InternVideo-Next, a two-stage pretraining scheme that builds a semantically consistent yet detail-preserving latent space for this world model. First, conventional linear decoder in pixel MVM enforces the predictor’s output latent to be linearly projected to, thus separable in pixel space, causing the conflict with semantic abstraction. Our Stage 1 proposes a conditional diffusion decoder and injects reliable image-level semantic priors to enhance semantics and convergence, thus bridging pixel-level fidelity with high-level semantic abstraction. Stage 2 further learns world knowledge by predicting frozen Stage 1 targets within this space, mitigating shortcut learning. Trained on public, unlabeled videos, InternVideo-Next achieves state-of-the-art results across benchmarks and provides a scalable path toward general video representation learning. Code and models will be released at <https://github.com/OpenGVLab/InternVideo>.

1. Introduction

Video, as one of the most dynamic and information-rich

♠ Corresponding author (lmwang@nju.edu.cn).




Figure 1. **Comparison with previous SOTA methods of size ViT-Large.** With only public sources, our model excels at general video benchmarks within a probing setting where ViTs are frozen to directly show representation’s quality. The benchmarks involve scene-related, motion-related, complex video-language related and implicit world knowledge (3D geometric prior, causal relations and fine-grained object motion) related tasks.

modalities, offers a window into the physical world. It embeds not only spatiotemporal dynamics [19, 34, 53, 59] but also causal relationships, 3D geometrical priors and physical cues [2, 10]. These are all essential components for building genuine video understanding models. Such models are crucial for advancing embodied AI [21, 66], procedural reasoning [77], and the next generation of multimodal large language models [29, 30, 33, 36, 37]. However, learning such knowledge from raw data, especially in a scalable and unbiased manner, remains a challenge.

Recent progress in large-scale video representation learning often falls into two categories. 1) Text-supervised methods, leveraging large-scale video-text alignment [31, 38, 65, 70, 76], have achieved impressive performance on semantic, human-centric tasks like action recognition. However, due to expensive annotation costs, currently their reliance on noisy, synthetic annotations with limited semantic coverage introduces external biases, making it difficult to

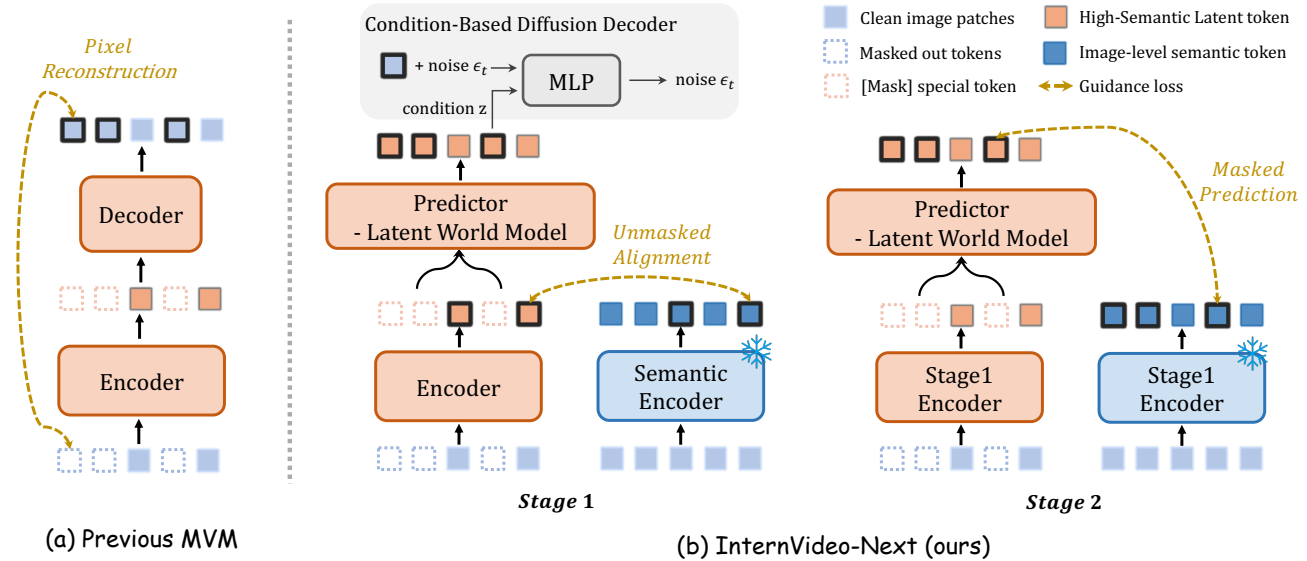


Figure 2. **Our two-stage self-supervised video pretraining framework.** For general video understanding and building real world-understanding video foundation models, we propose InternVideo-Next, which is simple, scalable, efficient and reproducible.

capture non-semantic vision features. They often fall short in tasks that require detailed object motion, causal structure or 3D geometrical priors. 2) Self-supervised video pre-training [2, 5, 52], in contrast, holds the promise of learning directly from the spatiotemporal structure of videos, offering a path towards truly unbiased and scalable world modeling. However, existing methods still lag behind text-supervised counterparts on benchmarks requiring strong semantics about main subject like Kinetics-400 [26].

We contend that this performance gap of self-supervised video pretraining is not an intrinsic limitation, but rather stems from fundamental architectural challenges unsolved in current methods: 1) **Bridge pixel-level fidelity with high-level semantic abstraction** 2) **Learn robust spatiotemporal dynamics, causal relations and 3D geometric priors from prediction without “shortcut solutions”**.

Popular MVM methods, like MAE [23], take a conventional Encoder-Decoder paradigm, where ViT-Decoder [18] directly takes Encoder output and generates reconstructed pixels. We explicitly decouple it into Predictor and Decoder, and consider the whole MVM paradigm as a new unified Encoder-Predictor-Decoder (EPD) overview. Such disentanglement enables a clearer inspection of the Predictor’s output latent space that is often neglected, which we find is essential in dealing with the unsolved challenges. It also reveals a key insight overlooked by previous works: the Encoder and Predictor should share a semantically rich yet detail-faithful output latent space, making the predictor a possible Latent World Model. Enforcing this compels the Predictor to complete missing content using genuine spatiotemporal relationships and implicit world knowledge rather than trivial correlations, further boosting seman-

tic abstraction with implicit world knowledge like geometry and motion in the Encoder’s representation.

In light of this, we propose InternVideo-Next, a novel two-stage video pretraining framework, as Figure 2:

In **Stage 1**, we construct such a semantically grounded, detail-preserving and structurally consistent latent space by: 1) leveraging a pixel reconstruction framework as foundation to learn fine-grained pixel details and spatiotemporal priors. 2) using a condition-based diffusion generative decoder in place of the popular Linear Decoder head in turning Predictor output latent to pixels. Such a diffusion decoder is helpful to maintain high semantics in the Predictor output space as common linear decoder may require the output to be linearly separable to pixels, which is harmful to balance semantic information with fine-grained details. 3) integrating semantic priors from a pretrained Image semantic model. Unlike video-text pretraining, which is limited by the sparsity and noisiness of video captions, image-text pretraining benefits from massive, high-quality corpora whose captions more comprehensively describe visual content. Such priors are cleaner and can accelerate model’s convergence by focusing on more temporal-centric information, meanwhile boosting multi-modal friendliness.

Building upon this, **Stage 2** then focuses on learning spatiotemporal dynamics and causal relations within this already coherent latent space. We employ a latent-space prediction objective, where a student model learns to predict the representations generated by a frozen teacher, where the two models are both initialized with weights learned in Stage 1. Crucially, the semantically rich and detail-preserving latent space established in Stage 1 prevents the “shortcut solutions” that plague traditional latent-prediction

methods, pushing the model to learn genuine predictive world knowledge rather than superficial temporal statistics.

As shown in **Figure 1**, **InternVideo-Next** achieves state-of-the-art results across a wide range of video understanding tasks. Remarkably, it is the **first video model trained without explicit video–text supervision** that surpasses video–text pretrained counterparts on both Kinetics-400 (action recognition) and SSv2 (fine-grained motion recognition). The learned representations also exhibit strong generalization to tasks requiring physical and 3D spatial intelligence, such as depth estimation and object-centric tracking, which are implicit world knowledge inherent in videos. In addition, with lightweight LiT-style [72] fine-tuning, **InternVideo-Next** delivers competitive zero-shot video-text retrieval performance, highlighting our model representation’s multimodal friendliness. Also, with preliminary exploration in chat-centric tasks in video, it holds the potential to be the foundation Encoder of next-generation multimodal video chat models. Our framework provides a unified and scalable pathway toward general-purpose world-understanding video models. All the code and model will be released upon publication.

2. Related Works

Text Supervised Pre-training. Clip-style pretraining has dominated current mainstream image pretraining frameworks [46, 54, 73], benefiting from large-scale web-crawled image-text pair data. Video–text pretraining methods [13, 15, 32, 41, 70] also achieve strong performance on benchmarks relying heavily on the main subject semantics like Kinetics-400. However, video captions are more difficult to be directly crawled from internet and are often generated by summarizing title and ASR information. Such synthetic annotations are typically noisier than images’ and lack sufficient motion or spatial diversity, considering that videos are harder to fully describe than images. As can be seen with **Figure 1**, although video-text models excel at semantic-heavy tasks, they often fall short in tasks related to non-semantic information like depth, fine-grained motion and causal relation. Our framework focuses more on unbiased learning from video itself, retaining high semantics, low-level details and world knowledge priors.

Masked Video Modeling. Masked video modeling (MVM) is the mainstream self-supervised video pretraining method, inspired by masked image modeling [4, 23, 44, 62]. VideoMAE [52] and MaskFeat [68] reconstruct masked video patches in the pixel domain, achieving strong performance but primarily capturing low-level appearance cues. Subsequent works such as VideoMAE-V2 [60] improve masking strategies and decoder designs but still struggle to preserve high-level semantics. Latent-space prediction approaches, such as V-JEPA [5], move toward semantic ab-

straction by predicting feature representations instead of pixels, yet symmetric teacher–student structures often lead to shortcut learning or semantic drift. As seen in **Figure 1**, it falls short in appearance-intensive tasks, depth estimation tasks that require low-level details, and high-semantic tasks, meaning that such shortcut learning leads to missing details and prevents higher-level information. Our approach unifies these paradigms under an Encoder–Predictor–Decoder (EPD) formulation and emphasizes maintaining semantic coherence between encoder and predictor latent spaces, combining the strengths of both.

The InternVideo Series. The InternVideo series [57, 63, 65] has explored multiple ways in integrating unsupervised video learning with text alignment to bridge pixel-level fidelity with high-level semantic abstraction. InternVideo merges weights from pretrained VideoMAE and CLIP models using model ensemble, while InternVideo2 introduces a video-only stage that aligns unmasked representations from VideoMAE and CLIP encoders. These methods can be viewed as weight- or embedding-level fusion of video priors and language knowledge. However, they still can not fully solve the conflict between fine-grained details and high-level semantics. InternVideo-Next revisits video-only pretraining from a task-level perspective by integrating CLIP-level semantic priors into an augmented video reconstruction framework, resolving such conflict.

3. Method

3.1. The EPD Disentanglement

We formulate our unified Encoder–Predictor–Decoder (EPD) overview as follows: E : A Vision Transformer that extracts spatiotemporal representations from an input video. P : A lightweight transformer that predicts latent representations for masked regions based on visible tokens. D : A reconstruction module that maps the predictor’s output latent to target space, either pixels or target latent.

3.2. S1: Semantic-Guided Pixel Reconstruction

Semantic Alignment Loss. Our approach leverages semantic knowledge from the image domain via a frozen SigLIP encoder. This allows us to inject a strong image-level semantic prior into our video-only pre-training, offsetting the need for noisy video-text annotations while focusing the model’s learning on spatiotemporal dynamics:

$$\mathcal{L}_{\text{sem}} = -\cos(E(X_{\text{vis}}), \text{vis}(\text{SigLIP}(X))), \quad (1)$$

We use a cosine similarity loss between the student’s embedding of a masked video and the corresponding visible region of the teacher’s embedding of the full video. Stage 1 jointly optimizes pixel reconstruction and semantic alignment with the same loss weight.

(a) Stage-1 module effect			(b) Predictor size & init			(c) Different Decoder settings		
Variant	K400	SSv2	Predictor	K400	SSv2	Decoder	K400	SSv2
Pixel Recon. Baseline	47.2	28.1	ModernBert-L last2	73.2	34.8	Linear Head	69.4	31.3
SigLIP Align only	70.7	32.1	ModernBert-L last5	75.8	36.9	3-Layer MLP Head	69.7	31.2
Pixel Rec. + Align	69.8	31.8	last5 w/o init	74.2	35.4	DiffMLP D3 W1024	73.4	33.2
+ Diffusion Decoder	74.2	35.4	ModernBert-L last8	75.6	37.0	DiffMLP D6 W1536	75.8	36.9
+ Text-Decoder Init	69.4	31.3	Depth-5 ViT	74.1	34.9	DiffMLP D9 W2048	75.5	36.4
+ Keep Both	75.8	36.9	Depth-12 ViT	73.2	34.4			

Table 1. **Ablations of Stage 1 design.** Linear probing results are shown. (a) Adding semantic alignment boosts overall performance, while directly merging it into pixel reconstruction framework causes performance drop. Combining our diffusion decoder and text decoder augmentation make them fuse perfectly and results in accuracy gain. (b) Moderate predictor depth with initialization performs best in our scheme, outperforming popular results with Depth-12 ViT. (c) D for number of MLP layers. and W for network width.

Semantic-Aware Masking. Semantic mask prioritizes temporally informative regions using attention scores derived from the semantic teacher with a top-k selection.

Diffusion Decoder. Unlike the popular linear decoder [23], we adopt a lightweight conditional diffusion [25] decoder to model the distribution of each patch independently. The diffusion process and loss function follow [42], where noise schedule has a cosine shape, with 1000 steps at training time. We use a small MLP consisting of a few residual blocks [22] for denoising. The denoising MLP is conditioned on a vector z produced by the Predictor and outputs corresponding pixels. Since it only models a patch’s latent distribution, a small MLP works perfectly and does not involve much overhead.

Text-Decoder Initialization. Conventional pixel reconstruction framework uses zero-init ViT. Our Predictor P is initialized with weights from a pretrained text decoder [38, 67], which provides better semantic priors and the smooth translation of high semantics between the two latent spaces, and proved to demand fewer layers than common methods.

3.3. S2: Semantically Coherent Latent Prediction

Stage 2 builds upon the semantically aligned latent spaces obtained in Stage 1 to further enhance temporal abstraction and representation generalization.

Initialization. Both the student and teacher networks are initialized with the weights learned in Stage 1. And the Stage 1 predictor is also kept in Stage 2.

Multi-Block Masking. To strengthen temporal reasoning, we apply a multi-block masking [5], where large contiguous spatiotemporal blocks in a video are masked. This strategy increases prediction difficulty and mitigates information leakage, which is better for Stage 2’s target of implicit world knowledge by prediction.

Latent Prediction Objective. The student predicts the teacher’s latent representations for masked regions. This objective enforces consistency in the latent space without direct pixel reconstruction, enabling the model to focus on abstract semantic and temporal patterns.

Frozen Teacher Target. Unlike V-JEPA, the teacher network is frozen with Stage 1 initialization to prevent shortcut learning, as the Stage 1 initialization is already detail-preserving and of high semantics.

4. Experiments

Implementation. For our pretraining, we utilize the last five layers of ModernBert-Large [67] as our Predictor. We use SigLIP2-1B [54] as our semantic teacher in our final version, and use SigLIP2-Large in the Ablation study. For Stage 1 pre-training, we use a mask ratio of 80% and a learning rate of $1e-3$. For ablation study, we use 32 A100 with a batch size of 1024 for 30 epochs in both Stage 1 and Stage 2. And for final training, we use 64 A100 with a batch size of 2048 for 50 epochs in Stage 1 and 100 epochs in Stage 2. For more, please refer to the Supplement.

Datasets. Unless otherwise stated, we use K-Mash [65] for training, including 1.1M video data from K400 [26], K600 [8], K700 [9], SSv2 [20], ANet [24], HACS [75], and MiT [39] with duplicates removed. For ablation, we use a subset K710 containing only the Kinetics datasets.

4.1. Ablation Study

We conduct comprehensive ablations to analyze performance of our training framework on appearance-based, motion-centric and depth-estimation benchmarks, including K400 [26], SSv2 [20], and ScanNet [16]. We leverage Linear Probing results on action recognition tasks for clear presentation of representation quality improvement. Results on ScanNet are based on probing results using learnable queries and a single-layer trainable cross-attention head.

Stage1-Variant	K400 Acc1↑	SSv2 Acc1↑	Scannet δ_1 ↑
DinoV2 Align only	68.4	29.5	43.3
Clip-ViT Align only	69.1	31.2	40.6
SigLIP Align only	70.7	32.1	42.1
<i>InternVideo2 Stage1 Strategy</i>			
SigLIP+VideoMAE align	70.4	32.5	50.1
Ours w/ <i>DinoV2 align</i>	69.3	31.5	58.3
Ours w/ <i>Clip-ViT align</i>	73.2	34.6	57.8
Ours w/ <i>SigLIP align</i>	75.8	36.9	59.4

Table 2. **Comparison with different supervision signals.** Also, we compare results with InternVideo2-S1, which directly aligns with two teachers as a first attempt to balance pixel-level detail and semantic abstraction, in fair training recipe.

Stage-2 Variant	K400	SSv2
Stage 1	75.8	36.9
Our Stage 2 setting	76.9	56.9
<i>replace w/ Zero-Init V-JEPA Predictor</i>	74.8	53.8
<i>replace w/ Momentum 0.9998 Target</i>	74.1	54.3
<i>replace w/ Frozen SigLIP2 Target</i>	75.4	45.7
<i>replace w/ Frozen InternVideo2 Target</i>	74.3	47.4
<i>if w/ Unmasked Token Alignment Loss</i>	75.7	51.1
<i>if w/ Pixel Reconstruction Loss</i>	76.8	57.0

Table 3. **Stage-2 ablation.** Our full configuration yields the strongest temporal abstraction and overall accuracy.

Effects of Stage 1 Settings. Table 1 presents detailed ablations on the Stage 1 design. As shown in (a), the pixel reconstruction baseline achieves 47.2% on K400, confirming its limited semantic abstraction ability. Using SigLIP alignment alone without reconstruction boosts recognition accuracy. However, naively combining pixel reconstruction and alignment hurts overall performance due to optimization conflicts. Introducing the *diffusion decoder* reverses such degradation into a +4.4% performance gain, demonstrating strong complementarity between the two supervision signals, if not hurt by the original Linear Decoder. It highlights the effect of our EPD disentanglement, which provides a novel reconsideration and inspires us to re-use the now seldom-utilized pixel reconstruction framework. (b) studies the impact of predictor architecture and initialization. The vanilla ViT predictor without initialization, as used in the original setting, underperforms our solution.

Panel (c) compares different decoder designs. A lightweight linear head underfits complex spatial relations, while our diffusion-based MLP decoders progressively improve as capacity grows. The medium-sized Depth 6,

Stage	Mask Type	#F	K400 Acc1↑	SSv2 Acc1↑	Scannet δ_1 ↑
<i>Effect of Masking Strategy (ViT-B)</i>					
Stage-1	Semantic	8	75.8	36.9	59.4
	Multi-block	8	74.4	36.3	58.1
Stage-2	Semantic	8	75.2	52.3	61.1
	Multi-block	8	76.9	56.9	66.1
<i>Effect of Temporal Frame Count (ViT-B)</i>					
Stage-1	Semantic	8	75.8	36.9	59.4
		16	76.4	38.4	59.8
		32	76.6	38.6	59.3
Stage-2	Multi-block	8	77.0	57.5	67.1
		16	77.4	58.4	68.7
		32	78.1	59.4	70.1

Table 4. **Impact of masking and frame settings.** Each Stage-2 model is based on the highlighted Stage-1 model in the same table.

Width 1536 variant achieves the best trade-off between representation richness and computational cost.

Different supervision for Stage 1’s semantic prior. Table 2 reveals that our framework shows generalizable performance gain with different semantic alignment targets, whether on common recognition tasks or 3D understanding tasks with low-level details. Supervision from semantic image models performs best. It also compares our unified Stage 1 strategy with the multi-teacher alignment scheme in InternVideo2, which aligns with both CLIP and VideoMAE to balance pixel-level detail and semantic abstraction.

Effects of Stage 2 Settings. Stage 2 further boosts temporal abstraction and causal relations via masked latent prediction. As shown in Table 3, our configuration achieves the strongest results. Replacing our predictor with a zero-initialized V-JEPA variant, unfreezing the target encoder or using other initialization of the target encoder all lead to degradation, confirming the importance of semantically coherent initialization and a stable target that encodes both high semantics and pixel-level details. Adding unmasked token alignment as in Stage 1 introduces noise to latent targets and harms motion modeling. Pixel-space reconstruction during this stage provides marginal benefit to SSv2, as the targets, produced by our Stage 1 encoder, already contain enough pixel-level details.

Masking Strategy and Temporal Window. Table 4 analyzes masking and temporal length configurations. Semantic masking in Stage 1 helps the model attend to spa-

Model Name	ViT Size	Data	GPU-hrs	K400 \uparrow	SSv2 \uparrow	Coin \uparrow
<i>Image Models</i>						
SigLIP2 [54]	Giant	10B	-	85.6	47.9	91.8
DinoV2 [7]	Giant	142M	-	83.1	50.0	87.8
<i>Models Trained with Video-Text Pairs</i>						
InternVideo2 _{s2} [65]	Base	25.5M	-	84.9	64.7	88.7
InternVideo2 _{s2} [65]	Large	25.5M	-	86.0	65.9	90.1
InternVideo2 _{s2} [65]	1B	25.5M	30K	87.9	67.3	91.7
InternVideo2 _{s2} [65]	6B	400M	200K	88.8	67.7	92.6
VideoPrism [76]	Base	618M	-	84.2	63.6	-
VideoPrism [76]	1B	618M	250K*	87.2	68.5	-
<i>Models Trained with Video Data Only</i>						
VideoMAEv2 [60]	Large	1.35M	-	80.9	54.9	83.2
VideoMAEv2 [60]	1B	1.35M	15K	82.8	56.1	84.6
V-JEPAv1 [5]	Large	2M	8K*	80.8	69.5	83.0
V-JEPAv2 [2]	Large	22M	10K*	83.3	72.0	85.9
InternVideo2 _{s1} [65]	1B	1.1M	15K	85.6	57.4	89.6
InternVideo2 _{s1} [65]	6B	2.1M	110K	86.0	59.0	90.3
InternVideo-Next_{s1}	Base	1.1M	1.5K	84.8	58.6	89.9
InternVideo-Next_{s2}	Base	1.1M	3.4K	85.9	70.1	91.4
InternVideo-Next_{s1}	Large	1.1M	3.6K	87.1	65.5	92.0
InternVideo-Next_{s2}	Large	1.1M	9.7K	88.4	73.0	93.6

Table 5. **Attentive Probing Results on Video Action Recognition Datasets K400, SSv2 and COIN.** We report top@1 results with a single probing head trainable. The training cost is estimated with equivalent A100 GPU Hours. * means estimated result, as not included in the paper. Models are tested with 16 frames, and with 1 clips \times 3 crops. Our models achieve the best performance across model sizes and with the smallest, public and annotation-free training data, highlighting our method’s efficiency.

tially meaningful regions and bridge high-level semantic abstraction with pixel-level fidelity. Multi-block masking in Stage 2 makes the prediction task harder and forces robust spatiotemporal dynamics and causal relations learned in the Predictor. Moreover, increasing the number of input frames consistently enhances performance, demonstrating that richer temporal diversity strengthens temporal abstraction. We finally chose F16 for Stage 1 and F32 for Stage 2 for balance between accuracy and efficiency.

4.2. Single-modality Tasks

We evaluate our method on several single-modality tasks: including the *Kinetics-400* [26], *Something-Something V2* [20] for action recognition, *COIN* [51] for long video classification, *ScanNet* [16] and *KITTI* [6] for monocular depth estimation, *WayMo*[50] for object tracking and *EK100*[17] for action prediction.

Video Classification. Table 5 reports probing results with different frozen backbones on general video classification tasks. We test the model in an ‘Attentive Probing’ setting where the encoders are frozen and a single-layer attention pooling head is trained. Such Frozen Encoder settings can test representation’s quality in an unbiased way. Our methods achieve the best results with only public data and less computation cost on these foundation tasks.

Depth Estimation. As discussed above, 3D geometric prior is also a part of the world knowledge that can be acquired from videos. Table 6 tests the Encoder’s probing performance on 3D-related Monocular Video Depth Estimation tasks *Scannet* [16] and *KITTI*[6] in two settings. One uses a single-layer cross-attention based Probing Head with learnable queries, as to directly show model’s spatiotemporal embedding quality regarding Video Depth. Our InternVideo-Next achieves far better results than image models as it learns more direct 3D spatial information in the backbone. And another uses a complex VDA head designed specifically for Image models with additional temporal fusion modules [14]. Our InternVideo-Next excels at both settings, even achieving nearly SoTA performance in such a ‘Probing Setting’ compared with the specifically designed and trained Video Depth Anything [14] model. Detailed test settings are left in the Supplement.

Object Tracking. We follow the setting in Scaling 4D representations [10] to test model embedding’s quality in capturing object-level motion with Object Tracking task on *Waymo Open* [50]. The task requires the model to predict Bounding Box given the coordinates of an object in the first frame. We use learnable queries and a single-layer cross-attention pooling head with the frozen ViT backbone for testing. We will leave more details in the Supp. Our model




Figure 3. **Video Depth estimation visualization.** All models are trained with a frozen setting with only a probing head (VDA head from Video Depth Anything [14]) module trainable. Graphs are from the first video of the ScanNet Dataset. Our model shows potential for building next-generation 3D-spatial intelligence models from video sources. Videos are resized with a spatial resolution of 224.

Model Name	Scannet		KITTI	
	AReI↓	δ_1 ↑	AReI↓	δ_1 ↑
<i>SOTA solution with VDA head and specific training</i>				
VideoDepthAnything [14]	8.7	92.6	8.3	94.6
<i>Image Models w/ a simple probing head [10]</i>				
SigLip2-L [54]	26.4	50.4	16.8	74.9
DinoV3-L [49]	24.0	59.1	13.5	81.8
<i>Video Models w/ a simple probing head [10]</i>				
VideoMAEv2-L [60]	19.3	66.9	11.9	84.6
InternVideo2-L [65]	23.4	57.0	12.9	83.0
V-JEPA2-L [2]	20.9	63.1	10.9	86.8
InternVideo-Next-L	13.7	81.5	9.9	88.6
<i>Image Models w/ VDA head with temporal fusion [14]</i>				
SigLip2-L [54]	19.9	65.6	12.2	84.2
DinoV3-L [49]	9.6	91.2	7.0	93.0
<i>Video Models w/ VDA head with temporal fusion [14]</i>				
VideoMAEv2-B [60]	21.8	60.6	9.6	89.7
InternVideo2-B [65]	15.7	76.2	8.8	91.3
InternVideo-Next-B	11.3	87.3	7.5	93.2
VideoMAEv2-L [60]	17.9	70.4	9.5	89.7
InternVideo2-L [65]	15.1	77.8	8.1	92.3
V-JEPA2-L [2]	14.4	79.6	8.7	91.2
InternVideo-Next-L	9.2	92.2	6.7	94.6

Table 6. **Probing Results on Video Depth-Estimation Datasets Scannet and KITTI.** Our InternVideo-Next with such a simple probing setting can achieve comparable results with the carefully designed and trained Video Depth Anything model. AReI for ‘absolute relative error’ and δ_1 for ‘relative accuracy’ [14].

achieves the best result in capturing object-level motion.

Video Prediction. Table 8 shows video prediction results on Epic-Kitchens-100 [17]. The model predicts the future

Model Name	Waymo Open mean IOU↑
<i>Image Models</i>	
SigLip2-L [54]	52.3
DinoV3-L [49]	59.7
<i>Video Models</i>	
InternVideo2-L [65]	63.0
V-JEPA2-L [2]	68.9
InternVideo-Next-L	72.4

Table 7. **Probing Results on Video Object Tracking Task Waymo Open.** We follow the task setting in Scaling4D [10] to test representation’s capability in capturing object-level motion.

action given a contextual video clip. The outputs of the predictor and encoder are concatenated along the token dimension and fed to an attentive probe layer and three classification heads. We follow most of the settings in V-JEPA2 [2] and report the mean-class recall-at-5 result. Our model performs better at predicting future actions.

4.3. Multi-modality Tasks

In this section, we evaluate the potential of our model in multi-modal learning scenarios through two representative experimental setups. First, following the LiT [72] methodology, we freeze the ViT backbone and train only a text encoder, initialized with MobileCLIP-B [55], on 25.5M [3, 11, 28, 35, 40, 43, 48] publicly available video-text pairs from UMT [31] for 5 epochs. This setup evaluates the model’s zero-shot video-text retrieval performance, serving as a direct probe into the completeness and correctness of the visual embeddings in the semantic space. We choose InternVideo2_{s2} and V-JEPA2 as our baseline model. Second, to assess the model’s applicability as a video foun-

Method	Verb	Noun	Action
<i>MLLM MCQ result</i>			
VideoLLaMa-7B [74]	52.9	52.0	26.0
<i>Frozen Backbone Evaluation</i>			
V-JEPA2-L [2]	57.8	53.8	32.7
InternVideo-Next-L	58.9	56.4	34.0

Table 8. **Results of Video Action Prediction task EK100.** We report mean-class recall-at-5 for verb, noun and action on the validation set of EK100, following the test pipeline in V-JEPA2.

Method	K400	K700	SSv2-MC	MiTv1
InternVideo2 _{s2} -B [65]	67.7	57.9	55.9	27.9
InternVideo-Next-B	68.9	59.0	61.2	29.0
V-JEPA2-L [2]	60.2	51.4	53.0	25.4
InternVideo2 _{s2} -L [65]	70.7	61.9	59.6	30.6
InternVideo-Next-L	72.1	63.0	64.2	32.0

Table 9. **Results of Zero-shot Action Recognition on K400, K700, SSv2-MC and MiT.** Both models use a MobileCLIP-B text encoder and train with ViT frozen to test their representations’ separability, alignment and completeness in text space.

ation model for high-level dialogue tasks, we connect the frozen ViT with a frozen large language model, Qwen2-7B [61], via a trainable MLP connector. Only the MLP is fine-tuned. This is consistent with the common Stage 1 training protocol for video-based chat models to validate the compatibility and transferability of our visual representations in a generative video-LLM framework.

Zero-shot Action Recognition. Table 9 shows results on zero-shot action recognition tasks including K400 [26], K700 [9], SSv2 [20] and MiTv1 [39]. Our model outperforms InternVideo2_{s2}, which is the previous SoTA encoder for multimodal retrieval tasks, especially on motion-intensive task SSv2-Multiple Choice.

Zero-shot Retrieval. Table 10 shows results on zero-shot text-to-video retrieval tasks, including MSRVT [71], DiDeMo [1], ActivityNet [27], LSMDC [47] and MSVD [12]. Our model shows comparable performance. This highlights our model’s embedding of the video is highly aligned with the semantic text space implicitly.

Chat-centric Tasks. We further evaluate our model’s potential as foundation Encoder for generative video-LLM frameworks. We follow the recipe of VideoChat-Flash [33] Stage 1 and utilize a large-scaled train-set including LLava-558K[37], S-MiT [40], 700k filtered subset of WebVid-10M [3], VidLN[56], and SSv2-open-ended [20] to fully train the connector. Such ‘Linear Probing’ setting is widely used in the Stage-1 training of multi-modal chat models.

Method	MSR	DDM	ANet	LSMDC	MSVD
<i>Popular methods with full video-text training</i>					
VINDLU-L [15]	32.0	36.9	30.9	-	-
InternVideo-L [63]	40.7	31.5	30.7	17.6	43.4
UMT-L [31]	40.7	48.6	41.9	24.9	49.0
ViClip-L [64]	42.4	18.4	15.1	20.1	49.1
LanguageBind-L [78]	42.8	39.7	38.4	-	54.1
<i>Frozen ViT with LiT Training using MobileCLIP-B</i>					
V-JEPA2-L [2]	34.4	36.3	35.7	19.2	40.1
InternVideo2 _{s2} -L [65]	42.1	42.8	43.6	21.4	44.5
InternVideo-Next-L	43.2	43.7	43.4	20.8	46.1

Table 10. **Results of zero-shot text-to-video retrieval on MSRVT, DiDeMo, ActivityNet, LSMDC and MSVD.** We only report the T2V R@1 accuracy here.

Encoder	MVBench	Percept Test	Dream1k
Clip-L [46]	45.6	45.1	28.4
SigLIP ₃₃₆ -L [73]	46.7	44.1	29.2
SigLIP ₂₃₃₆ -L [54]	46.9	45.8	29.6
UMT-L [31]	45.0	44.5	24.6
V-JEPA2-L [31]	44.3	44.2	24.3
InternVideo2-L [65]	47.0	46.7	28.7
InternVideo-Next-L	50.6	49.2	29.8

Table 11. **Results on Chat-Centric benchmarks MVBench [69], Perception Test [45] for General Perception and Dream1k [58] for Detailed Caption.** Models are trained in a multi-modal ‘linear prob’ setting where both the Encoder and the LLM are frozen, which is a common setting in the Stage 1 training for multi-modal chat models. For Dream1k we report the F1-score.

Our model achieves better results on general spatiotemporal perception Video QA tasks MVBench [69] and Perception Test [45], and caption task Dream1K [58].

5. Conclusion and Future Work

In this paper, we present InternVideo-Next, a two-stage video pretraining framework without video-text supervision that unifies pixel-level reconstruction and latent prediction under a coherent semantic space and learns by prediction with such a space. By enforcing semantic alignment between encoder and predictor, our method learns structured spatiotemporal priors from raw videos, surpassing previous video-only and video-text methods on general video tasks and proving InternVideo-Next’s capability as a strong general video foundation model. For future works, we plan to investigate the scalability of this method, and extend this framework toward multi-modal and interactive world modeling, exploring efficient scaling and open-world generalization for more versatile video foundation models.

Acknowledgement

This work is supported by the National Key R&D Program of China (No. 2022ZD0160900), and the Basic Research Program of Jiangsu (No. BK20250009).

References

- [1] Lisa Anne Hendricks, Oliver Wang, Eli Shechtman, Josef Sivic, Trevor Darrell, and Bryan Russell. Localizing moments in video with natural language. In *ICCV*, 2017. 8
- [2] Mido Assran, Adrien Bardes, David Fan, Quentin Garrido, Russell Howes, Mojtaba Komeili, Matthew Muckley, Ammar Rizvi, Claire Roberts, Koustuv Sinha, Artem Zhohus, Sergio Arnaud, Abha Gejji, Ada Martin, Francois Robert Hogan, Daniel Dugas, Piotr Bojanowski, Vasil Khalidov, Patrick Labatut, Francisco Massa, Marc Szafraniec, Kapil Krishnakumar, Yong Li, Xiaodong Ma, Sarath Chandar, Franziska Meier, Yann LeCun, Michael Rabbat, and Nicolas Ballas. V-jepa 2: Self-supervised video models enable understanding, prediction and planning. *ArXiv*, 2506.09985, 2025. 1, 2, 6, 7, 8
- [3] Max Bain, Arsha Nagraani, Gül Varol, and Andrew Zisserman. Frozen in time: A joint video and image encoder for end-to-end retrieval. In *ICCV*, 2021. 7, 8
- [4] Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei. Beit: Bert pre-training of image transformers. In *ICLR*, 2021. 3
- [5] 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*, 2404.08471, 2024. 2, 3, 4, 6
- [6] Jens Behley, Martin Garbade, Andres Milioto, Jan Quenzel, Sven Behnke, Cyrill Stachniss, and Juergen Gall. Semantickitti: A dataset for semantic scene understanding of lidar sequences. *ArXiv*, 2019. 6
- [7] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, 2021. 6
- [8] João Carreira, Eric Noland, Andras Banki-Horvath, Chloe Hillier, and Andrew Zisserman. A short note about kinetics-600. *ArXiv*, abs/1808.01340, 2018. 4
- [9] João Carreira, Eric Noland, Chloe Hillier, and Andrew Zisserman. A short note on the kinetics-700 human action dataset. *ArXiv*, abs/1907.06987, 2019. 4, 8
- [10] João Carreira, Dilara Gokay, Michael King, Chuhan Zhang, Ignacio Rocco, Aravindh Mahendran, Thomas Albert Keck, Joseph Heyward, Skanda Koppula, Etienne Pot, Goker Erdogan, Yana Hasson, Yi Yang, Klaus Greff, Guillaume Le Moing, Sjoerd van Steenkiste, Daniel Zoran, Drew A. Hudson, Pedro Vélez, Luisa Polanía, Luke Friedman, Chris Duvarney, Ross Goroshin, Kelsey Allen, Jacob Walker, Rishabh Kabra, Eric Aboussouan, Jennifer Sun, Thomas Kipf, Carl Doersch, Viorica Pătrăucean, Dima Damen, Pauline Luc, Mehdi S. M. Sajjadi, and Andrew Zisserman. Scaling 4d representations. *arXiv*, 2025. 1, 6, 7
- [11] Soravit Changpinyo, Piyush Sharma, Nan Ding, and Radu Soricut. Conceptual 12m: Pushing web-scale image-text pre-training to recognize long-tail visual concepts. In *CVPR*, 2021. 7
- [12] David L. Chen and William B. Dolan. Collecting highly parallel data for paraphrase evaluation. In *ACL*, 2011. 8
- [13] Sihan Chen, Handong Li, Qunbo Wang, Zijia Zhao, Mingzhen Sun, Xinxin Zhu, and Jing Liu. Vast: A vision-audio-subtitle-text omni-modality foundation model and dataset. In *NIPS*, 2023. 3
- [14] Sili Chen, Hengkai Guo, Shengnan Zhu, Feihu Zhang, Zilong Huang, Jiashi Feng, and Bingyi Kang. Video depth anything: Consistent depth estimation for super-long videos. In *CVPR*, 2025. 6, 7
- [15] Feng Cheng, Xizi Wang, Jie Lei, David J. Crandall, Mohit Bansal, and Gedas Bertasius. Vindlu: A recipe for effective video-and-language pretraining. *ArXiv*, abs/2212.05051, 2022. 3, 8
- [16] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. *ArXiv*, 2017. 4, 6
- [17] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, and Michael Wray. The epic-kitchens dataset: Collection, challenges and baselines. *TPAMI*, 43, 2020. 6, 7
- [18] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 2
- [19] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *ICCV*, 2019. 1
- [20] 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 *ICCV*, 2017. 4, 6, 8
- [21] Jing Gu, Eliana Stefani, Qi Wu, Jesse Thomason, and Xin Wang. Vision-and-language navigation: A survey of tasks, methods, and future directions. In *ACL*, 2022. 1
- [22] Kaiming He, X. Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *CVPR*, 2016. 4
- [23] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross B. Girshick. Masked autoencoders are scalable vision learners. In *CVPR*, 2022. 2, 3, 4
- [24] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *CVPR*, 2015. 4
- [25] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NIPS*, 2020. 4

- [26] Will Kay, João Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Apostol Natsev, Mustafa Suleyman, and Andrew Zisserman. The kinetics human action video dataset. *ArXiv*, abs/1705.06950, 2017. 2, 4, 6, 8
- [27] Ranjay Krishna, Kenji Hata, Frederic Ren, Li Fei-Fei, and Juan Carlos Niebles. Dense-captioning events in videos. In *ICCV*, 2017. 8
- [28] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *IJCV*, 2017. 7
- [29] Junnan Li, Dongxu Li, Silvio Savarese, and Steven C. H. Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *ICML*, 2022. 1
- [30] Kunchang Li, Yanan He, Yi Wang, Yizhuo Li, Wenhai Wang, Ping Luo, Yali Wang, Limin Wang, and Yu Qiao. Videochat: Chat-centric video understanding. *ArXiv*, abs/2305.06355, 2023. 1
- [31] Kunchang Li, Yali Wang, Yizhuo Li, Yi Wang, Yanan He, Limin Wang, and Yu Qiao. Unmasked teacher: Towards training-efficient video foundation models. In *ICCV*, 2023. 1, 7, 8
- [32] Tianhao Li and Limin Wang. Learning spatiotemporal features via video and text pair discrimination. *arXiv*, 2021. 3
- [33] Xinhao Li, Yi Wang, Jiashuo Yu, Xiangyu Zeng, Yuhan Zhu, Haian Huang, Jianfei Gao, Kunchang Li, Yanan He, Chenting Wang, Yu Qiao, Yali Wang, and Limin Wang. Videochat-flash: Hierarchical compression for long-context video modeling. *arXiv*, 2025. 1, 8
- [34] Ji Lin, Chuang Gan, and Song Han. Tsm: Temporal shift module for efficient video understanding. In *ICCV*, 2019. 1
- [35] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. 7
- [36] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning. *ArXiv*, abs/2310.03744, 2023. 1
- [37] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. In *NeurIPS*, 2023. 1, 8
- [38] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *NeurIPS*, 2019. 1, 4
- [39] Mathew Monfort, Bolei Zhou, Sarah Adel Bargal, Alex Andonian, Tom Yan, Kandan Ramakrishnan, Lisa M. Brown, Quanfu Fan, Dan Gutfreund, Carl Vondrick, and Aude Oliva. Moments in time dataset: One million videos for event understanding. *TPAMI*, 2020. 4, 8
- [40] Mathew Monfort, SouYoung Jin, Alexander Liu, David Harwath, Rogerio Feris, James Glass, and Aude Oliva. Spoken moments: Learning joint audio-visual representations from video descriptions. In *CVPR*, 2021. 7, 8
- [41] Bolin Ni, Houwen Peng, Minghao Chen, Songyang Zhang, Gaofeng Meng, Jianlong Fu, Shiming Xiang, and Haibin Ling. Expanding language-image pretrained models for general video recognition. *ArXiv*, abs/2208.02816, 2022. 3
- [42] Alex Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. *arXiv*, 2021. 4
- [43] Vicente Ordonez, Girish Kulkarni, and Tamara Berg. Im2text: Describing images using 1 million captioned photographs. In *NeurIPS*, 2011. 7
- [44] Zhiliang Peng, Li Dong, Hangbo Bao, Qixiang Ye, and Furu Wei. Beit v2: Masked image modeling with vector-quantized visual tokenizers. *ArXiv*, abs/2208.06366, 2022. 3
- [45] Viorica Pătrăucean, Lucas Smaira, Ankush Gupta, Adrià Recasens Contente, Larisa Markeeva, Dylan Banarse, Skanda Koppula, Joseph Heyward, Mateusz Malinowski, Yi Yang, Carl Doersch, Tatiana Matejovicova, Yury Sulsky, Antoine Miech, Alex Frechette, Hanna Klimczak, Raphael Koster, Junlin Zhang, Stephanie Winkler, Yusuf Aytar, Simon Osindero, Dima Damen, Andrew Zisserman, and João Carreira. Perception test : A diagnostic benchmark for multimodal models. In *NeurIPS*, 2023. 8
- [46] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 3, 8
- [47] Anna Rohrbach, Atousa Torabi, Marcus Rohrbach, Niket Tandon, Christopher Joseph Pal, H. Larochelle, Aaron C. Courville, and Bernt Schiele. Movie description. *International Journal of Computer Vision*, 2016. 8
- [48] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *ACL*, 2018. 7
- [49] Oriane Sim'eoni, Huy V. Vo, Maximilian Seitzer, Federico Baldassarre, Maxime Oquab, Cijo Jose, Vasil Khali-dov, Marc Szafraniec, Seungeun Yi, Michael Ramamonjisoa, Francisco Massa, Daniel Haziza, Luca Wehrstedt, Jianyuan Wang, Timothée Darcet, Théo Moutakanni, Leonel Sentana, Claire Roberts, Andrea Vedaldi, Jamie Tolan, John Brandt, Camille Couprie, Julien Mairal, Herv'e J'egou, Patrick Labatut, and Piotr Bojanowski. Dinov3. *arXiv*, 2025. 7
- [50] 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 *CVPR*, 2020. 6
- [51] Yansong Tang, Dajun Ding, Yongming Rao, Yu Zheng, Danyang Zhang, Lili Zhao, Jiwen Lu, and Jie Zhou. Coin: A large-scale dataset for comprehensive instructional video analysis. In *CVPR*, 2019. 6
- [52] Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. VideoMAE: Masked autoencoders are data-efficient learners for self-supervised video pre-training. In *NeurIPS*, 2022. 2, 3
- [53] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A closer look at spatiotemporal convolutions for action recognition. In *CVPR*, 2018. 1

- [54] Michael Tschanen, Alexey Gritsenko, Xiao Wang, Muhammad Ferjad Naeem, Ibrahim Alabdulmohsin, Nikhil Parthasarathy, Talfan Evans, Lucas Beyer, Ye Xia, Basil Mustafa, Olivier Hénaff, Jeremiah Harmsen, Andreas Steiner, and Xiaohua Zhai. Siglip 2: Multilingual vision-language encoders with improved semantic understanding, localization, and dense features. *ArXiv*, 2025. 3, 4, 6, 7, 8
- [55] Pavan Kumar Anasosalu Vasu, Hadi Pouransari, Fartash Faghri, Raviteja Vemulapalli, and Oncel Tuzel. Mobileclip: Fast image-text models through multi-modal reinforced training. In *CVPR*, 2024. 7
- [56] Paul Voigtlaender, Soravit Changpinyo, Jordi Pont-Tuset, Radu Soricut, and Vittorio Ferrari. Connecting vision and language with video localized narratives. *CVPR*, 2023. 8
- [57] Chenting Wang, Kunchang Li, Tianxiang Jiang, Xiangyu Zeng, Yi Wang, and Limin Wang. Make your training flexible: Towards deployment-efficient video models. *arXiv*, 2025. 3
- [58] Jiawei Wang, Liping Yuan, Yuchen Zhang, and Haomiao Sun. Tarsier: Recipes for training and evaluating large video description models. *arXiv*, abs/2407.00634, 2024. 8
- [59] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal segment networks: Towards good practices for deep action recognition. In *ECCV*, 2016. 1
- [60] 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 *CVPR*, 2023. 3, 6, 7
- [61] Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, et al. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *ArXiv*, abs/2409.12191, 2024. 8
- [62] Wenhui Wang, Hangbo Bao, Li Dong, Johan Bjorck, Zhiliang Peng, Qiang Liu, Kriti Aggarwal, Owais Khan Mohammed, Saksham Singhal, Subhojit Som, and Furu Wei. Image as a foreign language: Beit pretraining for all vision and vision-language tasks. *ArXiv*, abs/2208.10442, 2022. 3
- [63] Yi Wang, Kunchang Li, Yizhuo Li, Yinan He, Bingkun Huang, Zhiyu Zhao, Hongjie Zhang, Jilan Xu, Yi Liu, Zun Wang, Sen Xing, Guo Chen, Junting Pan, Jiashuo Yu, Yali Wang, Limin Wang, and Yu Qiao. Internvideo: General video foundation models via generative and discriminative learning. *ArXiv*, abs/2212.03191, 2022. 3, 8
- [64] Yi Wang, Yinan He, Yizhuo Li, Kunchang Li, Jiashuo Yu, Xin Jian Ma, Xinyuan Chen, Yaohui Wang, Ping Luo, Ziwei Liu, Yali Wang, Limin Wang, and Y. Qiao. Internvid: A large-scale video-text dataset for multimodal understanding and generation. *ArXiv*, 2023. 8
- [65] Yi Wang, Kunchang Li, Xinhao Li, Jiashuo Yu, Yinan He, Chenting Wang, Guo Chen, Baoqi Pei, Ziang Yan, Rongkun Zheng, Jilan Xu, Zun Wang, Yansong Shi, Tianxiang Jiang, Songze Li, Hongjie Zhang, Yifei Huang, Yu Qiao, Yali Wang, and Limin Wang. Internvideo2: Scaling foundation models for multimodal video understanding. In *ECCV*, 2024. 1, 3, 4, 6, 7, 8
- [66] Zun Wang, Jialu Li, Yicong Hong, Yi Wang, Qi Wu, Mohit Bansal, Stephen Gould, Hao Tan, and Yu Qiao. Scaling data generation in vision-and-language navigation. In *ICCV*, 2023. 1
- [67] Benjamin Warner, Antoine Chaffin, Benjamin Clavié, Orion Weller, Oskar Hallström, Said Taghadouini, Alexis Gallagher, Raja Biswas, Faisal Ladhak, Tom Aarsen, Nathan Cooper, Griffin Adams, Jeremy Howard, and Iacopo Poli. Smarter, better, faster, longer: A modern bidirectional encoder for fast, memory efficient, and long context finetuning and inference. *ArXiv*, 2412.13663, 2024. 4
- [68] Chen Wei, Haoqi Fan, Saining Xie, Chao-Yuan Wu, Alan Yuille, and Christoph Feichtenhofer. Masked feature prediction for self-supervised visual pre-training. In *CVPR*, 2022. 3
- [69] Longhui Wei, Lingxi Xie, Wen gang Zhou, Houqiang Li, and Qi Tian. Mvp: Multimodality-guided visual pre-training. In *ECCV*, 2022. 8
- [70] Hu Xu, Gargi Ghosh, Po-Yao (Bernie) Huang, Dmytro Okhonko, Armen Aghajanyan, and Florian Metze Luke Zettlemoyer Christoph Feichtenhofer. Videoclip: Contrastive pre-training for zero-shot video-text understanding. *ArXiv*, abs/2109.14084, 2021. 1, 3
- [71] Jun Xu, Tao Mei, Ting Yao, and Yong Rui. Msr-vtt: A large video description dataset for bridging video and language. In *CVPR*, 2016. 8
- [72] Xiaohua Zhai, Xiao Wang, Basil Mustafa, Andreas Steiner, Daniel Keysers, Alexander Kolesnikov, and Lucas Beyer. Lit: Zero-shot transfer with locked-image text tuning. In *CVPR*, 2022. 3, 7
- [73] Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language image pre-training. *CVPR*, 2023. 3, 8
- [74] Hang Zhang, Xin Li, and Lidong Bing. Video-llama: An instruction-tuned audio-visual language model for video understanding. *ArXiv*, abs/2306.02858, 2023. 8
- [75] Hang Zhao, Antonio Torralba, Lorenzo Torresani, and Zhicheng Yan. Hacs: Human action clips and segments dataset for recognition and temporal localization. In *ICCV*, 2019. 4
- [76] Long Zhao, Nitesh B. Gundavarapu, Liangzhe Yuan, Hao Zhou, Shen Yan, et al. Videoprism: A foundational visual encoder for video understanding. In *PMLR*, 2024. 1, 6
- [77] Gaoyue Zhou, Hengkai Pan, Yann LeCun, and Lerrel Pinto. Dino-wm: World models on pre-trained visual features enable zero-shot planning. *arXiv*, 2025. 1
- [78] Bin Zhu, Bin Lin, Munan Ning, Yang Yan, Jiayi Cui, Hongfa Wang, Yatian Pang, Wenhao Jiang, Junwu Zhang, Zong-Rui Li, Wancai Zhang, Zhifeng Li, Wei Liu, and Liejie Yuan. Languagebind: Extending video-language pretraining to n-modality by language-based semantic alignment. In *ICLR*, 2024. 8