

Generation Models Know Space: Unleashing Implicit 3D Priors for Scene Understanding

Xianjin Wu¹, Dingkan Liang^{1†}, Tianrui Feng¹, Kui Xia², Yumeng Zhang²,
Xiaofan Li², Xiao Tan², and Xiang Bai^{1✉}

¹ Huazhong University of Science and Technology, China

² Baidu Inc., China

[†] Project Lead ✉ Corresponding author

{wuxianjin,dkliang,xbai}@hust.edu.cn




Fig. 1: Comparison of existing paradigms. Unlike methods relying on (a) explicit 3D inputs or (b) complex geometric supervision, (c) our **VEGA-3D** extracts implicit priors from video generation models. By repurposing them as Latent World Simulators, we achieve (d) superior performance without external 3D dependencies.

Abstract. While Multimodal Large Language Models demonstrate impressive semantic capabilities, they often suffer from **spatial blindness**, struggling with fine-grained geometric reasoning and physical dynamics. Existing solutions typically rely on explicit 3D modalities or complex geometric scaffolding, which are limited by data scarcity and generalization challenges. In this work, we propose a paradigm shift by leveraging the implicit spatial prior within large-scale video generation models. We posit that to synthesize temporally coherent videos, these models inherently learn robust 3D structural priors and physical laws. We introduce **VEGA-3D** (Video **E**xtracted **G**enerative **A**wareness), a plug-and-play framework that repurposes a pre-trained video diffusion model as a Latent World Simulator. By extracting spatiotemporal features from intermediate noise levels and integrating them with semantic representations via a token-level adaptive gated fusion mechanism, we enrich MLLMs with dense geometric cues without explicit 3D supervision. Extensive

experiments across 3D scene understanding, spatial reasoning, and embodied manipulation benchmarks demonstrate that our method outperforms state-of-the-art baselines, validating that generative priors provide a scalable foundation for physical-world understanding. Code is publicly available at <https://github.com/H-EmbodVis/VEGA-3D>.

Keywords: Video Generation Models · 3D Scene Understanding · Spatial Reasoning · Embodied AI

1 Introduction

Recent advancements in video generation models [3, 40, 44, 51, 75] have reshaped our expectations of visual systems, moving beyond high-fidelity generation to acting as interactive world models [45, 74, 79, 80]. To generate a plausible video, the model inherently aligns appearance with 3D geometry: occlusion requires persistent object identity, camera motion reveals depth-dependent apparent motion, and interactions must follow consistent dynamics. These constraints encourage latent representations that encode geometry-consistent structure and motion, yielding a strong learned 3D prior without explicit 3D supervision [46, 69]. This raises a compelling research question: if video generators already possess an implicit understanding of space and physics, can these implicit physical priors be repurposed to improve downstream 3D visual understanding?

This perspective is particularly critical for domains that require granular 3D awareness [8, 26, 28, 30, 58, 82, 94, 100], such as scene understanding. To equip embodied agents with such capabilities, prevailing research has predominantly followed two explicit paradigms, as illustrated in Fig. 1. The first stream [81, 99] directly utilizes explicit 3D modalities (*e.g.*, point clouds or depth) to provide definitive geometric grounding. The second stream [17, 41, 76] focuses on geometric scaffolding, which lifts 2D features into 3D space via extra reconstruction or distillation. Alongside these methods, an underexplored yet increasingly promising paradigm (Fig. 1(c)) lies in modern video generation models trained on large-scale video datasets, whose training objective implicitly rewards representations consistent with 3D geometry and physical dynamics.

In this work, we explore a new paradigm: leveraging representations learned by video generation models as priors for geometric understanding. As illustrated in the Fig. 2 (a), video diffusion models demonstrate remarkable multi-view consistency. The model captures the structural integrity of objects across different frames, implying a robust internal representation of 3D geometry. While generative models lack the semantic alignment of contrastive pre-training [73, 88], their geometric priors offer unique spatial guidance. As further evidenced in Fig. 2(b), incorporating these priors sharpens the original scattered attention of the baseline, effectively serving as spatial anchors that enable precise localization for fine-grained 3D reasoning.

Motivated by these observations, we propose **VEGA-3D**, a plug-and-play framework that incorporates the strengths of semantic and generative representations. Specifically, we introduce a video generation model (*e.g.*, Wan2.1 [75],




Fig. 2: Visualization of implicit 3D priors. (a) The generation model demonstrates strong multi-view geometric consistency, evidenced by high correspondence scores and stable PCA feature representations across shifting camera views. (b) By leveraging these priors, our **VEGA-3D** overcomes the spatial ambiguity observed in the baseline, yielding precisely-located attention maps of the target object in the instruction.

Vmem [51]) as a *Latent World Simulator* to enrich the visual stream with spatiotemporal world-knowledge priors, complementary to the semantic encoder. To solve the distribution shift between generative and semantic space, we design a token-level adaptive gated fusion module that integrates the two features. This fusion enables the model to actively exploit the generative backbone’s 3D awareness to strengthen geometry-sensitive reasoning, while preserving discriminative semantic cues.

Extensive experiments on 3D scene understanding (e.g., visual grounding, dense captioning, and QA), spatial reasoning benchmarks (e.g., VSI-Bench [83]), and robotics manipulation tasks (LIBERO [55]) demonstrate that our method significantly outperforms larger spatially-enhanced models. Furthermore, in Fig. 3, we provide quantitative evidence for the strong correlation between multi-view correspondence and downstream understanding performance. Besides, as evidenced in Fig. 3(a), the gains stem from synergy rather than replacement: generative and semantic features are complementary, and their fusion yields substantial improvements. Our analysis further shows that the most informative spatial cues emerge from intermediate representations and mid-denoising time of the generative model, instead of the final pixel outputs, and that these priors are particularly beneficial for localization-centric tasks, effectively providing a spatial anchor for MLLMs.

In summary, our contributions are threefold.

- We investigate that modern video generators learn transferable spatiotemporal priors that encode geometry-consistent structure and motion, and we show that these priors are most informative in intermediate representations and mid-denoising stages.




Fig. 3: Feature Analysis. (a): Generative priors effectively complement semantic features with fusion, yielding consistent performance gains. (b): Multi-view correspondence strongly correlates with downstream 3D understanding performance. More details in Sec. 4.1

- We propose **VEGA-3D**, a plug-and-play framework that repurposes video generation models as *Latent World Simulator* for MLLMs, and introduce a token-level adaptive gated fusion module to align and integrate heterogeneous generative and semantic token spaces.
- Extensive experiments across 3D scene understanding, spatial reasoning, and embodied manipulation benchmarks demonstrate consistent gains, validating generative priors. Moreover, our framework is scalable: advances in video generation are readily transferable to stronger downstream 3D understanding.

2 Related Work

2.1 Scene Understanding with Large Language Models

Extending Large Language Models to 3D domains is a rapidly growing field [24]. Early approaches aligned point cloud encoders directly with LLMs [33, 37, 65, 81], as seen in PointLLM [81], Point-Bind [33], and GPT4Point [65]. While effective, they heavily depend on the availability of high-quality 3D data. To bypass the need for direct 3D input, multi-view approaches [41, 66, 76, 99, 101] like Video-3D LLM [99] and GPT4Scene [66] project 2D features into 3D space using positional embeddings or BEV rendering. More recent works attempt to lift 2D representations via auxiliary geometric supervision: Ross3D [76] utilizes reconstructive instruction tuning with bird’s-eye-view supervision, while 3DRS [41] and ThinkWith3D [17] distill knowledge from pre-trained 3D backbones. However, these methods typically require complex multi-stage training pipelines or task-specific geometric annotations (e.g., depth, camera pose). In contrast, our approach leverages the implicit physical priors already present in pre-trained video generation models, eliminating the need for explicit geometric supervision or complex rendering pipelines.

2.2 Spatial Reasoning

While MLLMs excel at semantic recognition, they often suffer from "spatial blindness" when tasked with geometric reasoning or determining spatial relationships, as highlighted by benchmarks [43, 53, 83, 84, 92] like Sphere [92] and VSI-Bench [83]. To mitigate this, one line of research focuses on scaling data: SpatialVLM [9] and VLM-3R [25] train on massive datasets of spatial reasoning instructions to embed geometric concepts. Another direction explores mental simulation or chain-of-thought prompting [32], where models like MindCube [86] and CVP [16] verify spatial logic through auxiliary cognitive maps or reconstruction.

Distinct from these approaches, which treat spatial reasoning as a linguistic or logical problem, we treat it as a representational problem. By fusing generative video priors, we ground the MLLM’s reasoning in a physically consistent world model, enabling intuitive spatial understanding akin to human perception.

2.3 Video Generation Models

Video generation has rapidly progressed to high-fidelity, long-horizon synthesis [5, 34, 49, 75, 85]. Recent large-scale models demonstrate strong temporal coherence, suggesting their latent spaces capture rich spatiotemporal regularities [5, 49, 75]. Beyond visual fidelity, an emerging line of research focuses on structuring and controlling these generators for downstream applications [3, 51, 69, 102].

Crucially, a recent wave of work connects video generation with 3D understanding, though typically by coupling generation with explicit geometric supervision. For instance, Ross3D [76] and Omni-View [36] treat generation or novel-view synthesis as trainable auxiliary tasks, relying on explicit geometric labels (*e.g.*, depth, camera poses, or bird’s-eye-view targets) to inject constraints into the model. Other works integrate world models for reasoning and control: MILO [7] utilizes implicit world modeling for spatial reasoning, GaussianDWM [23] couples a 3D Gaussian world model with unified scene understanding and generation, DyVA [90] pairs a video model with a VLM to infer dynamics, and JEPA-VLA [62] injects predictive embeddings into policies for robotic learning. In contrast, **VEGA-3D** keeps the video generator entirely *frozen* and performs no generation or geometry estimation during training. We instead extract the multi-view feature consistency inherent in a frozen diffusion model as a label-free 3D prior. Our approach is thus orthogonal and complementary: advances in video generation translate directly into stronger 3D priors at no annotation cost.

3 Preliminaries

Multimodal Large Language Models. Following standard protocols [57, 68], we consider a multimodal large language model with parameters Θ . Given a multimodal input consisting of text tokens $\mathbf{x}$ and visual inputs $\mathbf{V}$, the visual

content is mapped to a sequence of visual embeddings $\mathbf{v} = f_{\text{proj}}(f_{\text{enc}}(\mathbf{V}))$, where f_{enc} is a visual encoder (e.g., SigLIP [88]) and f_{proj} is a projector.

The MLLM is trained to maximize the likelihood of the response token sequence $\mathbf{y}$ given the context:

$$\mathcal{L}_{\text{CE}}(\Theta) = - \sum_{i=1}^L \log p_{\Theta}(y_i \mid y_{<i}, \mathbf{x}, \mathbf{v}), \quad (1)$$

where Θ denotes all trainable parameters (e.g., $\Theta = (\theta_{\text{lm}}, \theta_{\text{enc}}, \theta_{\text{proj}})$).

Crucially, this supervision is *sparse* and *discrete* [1, 11]. The loss is computed in the vocabulary space, where spatial errors (e.g., predicting “left” vs. “right”) are treated as generic token mismatches. Lacking geometric metric constraints, standard discriminative encoders f_{enc} often exhibit “spatial blindness,” focusing on semantic presence rather than a precise spatial structure.

Video Diffusion Models. Modern video generators (e.g., Wan2.1 [75]) are Diffusion Transformers trained with Flow Matching [54], which learns a continuous-time transport field in the latent space. Given a clean latent video $\mathbf{z}_0$, we sample Gaussian noise $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and a time $t \sim \mathcal{U}(0, 1)$, and train a flow network $v_{\psi}(\cdot)$ to regress the target velocity under MSE:

$$\mathcal{L}_{\text{FM}}(\psi) = \mathbb{E}_{\mathbf{z}_0, \epsilon, t} \left[\|\mathbf{u}_t - v_{\psi}(\mathbf{z}_t, t, \mathbf{c})\|_2^2 \right], \quad (2)$$

where $\mathbf{c}$ denotes conditioning signals. The corresponding target velocity is $\mathbf{u}_t = \frac{d\mathbf{z}_t}{dt}$. In implementation, we use a discrete timestep index $k \in \{0, \dots, K\}$ (with $K=1000$) and its normalized time $t_k = \frac{k}{K}$.

4 Method

Our goal is to endow Multimodal Large Language Models (MLLMs) with the implicit generative prior inherent in video generation models. As illustrated in Fig. 4, our framework introduces a dual-branch visual encoding mechanism that synergizes the high-level semantic capabilities of a discriminative encoder (e.g., SigLIP [88]) and dense 3D structure priors from a generative video diffusion model (e.g., Wan2.1 [75]). The methodology is organized into three logical stages: (1) **3D Awareness Analysis** (Sec. 4.1), where we identify multi-view feature consistency as the key indicator of geometric capability; (2) **Latent World Simulation** (Sec. 4.2), which operationalizes these insights by mining spatiotemporal geometry from the generator’s intermediate representations via noise injection; and (3) **Bridging the Generative and Semantic Gap** (Sec. 4.3), which adaptively integrates these heterogeneous features via a token-level adaptive gated fusion mechanism to align with the MLLM.

4.1 3D Awareness via Multi-view Feature Consistency

A pivotal factor in robust 3D scene understanding is the ability to maintain consistent representations of physical geometry across varying viewpoints. While




Fig. 4: Overview of the VEGA-3D framework. We repurpose a frozen video generation model as a **Latent World Simulator** to extract implicit 3D priors. These features are dynamically integrated with the semantic stream via **Adaptive Gated Fusion**, equipping the MLLM with dense 3D structural awareness.

traditional discriminative models excel at semantic invariance, we hypothesize that effective 3D reasoning often benefits from multi-view feature consistency, which maps the same physical 3D point to a unified latent representation across different views. To quantitatively verify this correlation and evaluate the geometric integrity of different feature backbones, we introduce **Multi-view Correspondence Score**.

Metric Definition. We utilize the ScanNet test dataset split [21], which provides posed RGB frames and dense depth maps (only for analysis). For a 3D scene observed from V views, we project the encoder features $\mathbf{F}_v$ from each view into a shared global voxel grid using the ground-truth camera extrinsics and depth. For a specific voxel m observed in two different views v_i and v_j , we extract the corresponding feature vectors $\mathbf{h}_{m,v_i}$ and $\mathbf{h}_{m,v_j}$. The consistency score for this voxel is defined as cosine similarity:

$$S_{\text{voxel}}^{(m)} = \frac{\mathbf{h}_{m,v_i}^\top \mathbf{h}_{m,v_j}}{\|\mathbf{h}_{m,v_i}\| \|\mathbf{h}_{m,v_j}\|}. \quad (3)$$

The final scene-level score is obtained by averaging $S_{\text{voxel}}^{(m)}$ over all valid voxel pairs across the scene. A higher score indicates that the model implicitly aligns distinct views of the same 3D structure.

Correlation and Architectural Analysis. To validate whether this consistency serves as a reliable indicator for downstream 3D capability, we define a Normalized Overall Score (NOS). It is calculated by normalizing the perfor-

mance metrics in Tab. 4 to $[0, 1]$ with Min-Max normalization across the discriminative and generative models, explicitly including the baseline results to establish a relative performance improvement, and then averaging them into a single scalar.

As illustrated in Fig. 3, plotting the Correspondence Score against NOS reveals a distinct positive correlation, confirming that multi-view consistency is a strong predictor of 3D performance. Furthermore, the results highlight a significant architectural divergence. Models based on UNet architectures (e.g., SVD [4], Stable Diffusion [70], Vmem [51]) exhibit notably lower consistency scores. We attribute this to the local inductive bias of convolutions and the insufficient scale of data, which limits the receptive field and hinders long-range geometric alignment. In contrast, DiT based models (e.g., Wan2.1 [75]) leverage global attention mechanisms to capture holistic context, achieving remarkably high consistency ($> 96\%$) and consequently superior downstream 3D understanding.

Guided by this evidence, **VEGA-3D** selects DiT-based architectures to provide robust spatial priors. Next, we detail how to actively extract these implicit geometric representations from a frozen generative model.

4.2 Video Generative Model as a Latent World Simulator

Building on the premise that generative models encapsulate physical laws, we adopt the pretrained, parameter-frozen Wan2.1-T2V 1.3B [75] as our default generative encoder for its simple text-conditioning interface and strong localization-centric performance. We additionally evaluate other video generative models in Tab. 4, demonstrating that our framework is compatible with different video generative backbones. While traditional visual encoders process raw pixel intensities, video generative models operate in a compressed latent space governed by diffusion dynamics.

Given an input video sequence $\mathbf{V} \in \mathbb{R}^{T \times H \times W \times 3}$ of T frames, we first map it to a low-dimensional latent space via the model’s Variational Autoencoder (VAE), yielding $\mathbf{z}_0 = E(\mathbf{V})$. However, a static latent representation $\mathbf{z}_0$ is insufficient to activate the generative model’s reasoning capabilities fully. Diffusion models are trained to enforce structural coherence primarily during active denoising of a corrupted signal; the process of restoration reveals the model’s understanding of structure. Therefore, we perturb the clean latent along the same Flow Matching noising path used by the pretrained backbone. Specifically, we choose a discrete timestep index




Fig. 5: Adaptive Gated Fusion. It dynamically integrates heterogeneous features using a token-level gating mechanism.

$k \in \{0, \dots, K\}$ and define the normalized time as $t_k = \frac{k}{K}$. We then sample $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and construct:

$$\mathbf{z}_k = (1 - t_k) \mathbf{z}_0 + t_k \epsilon. \quad (4)$$

We feed $\mathbf{z}_k$ into the backbone $\Phi(\cdot)$ using an empty text prompt ($\mathbf{c}_{\text{text}} = ""$). This ensures that the activated features rely solely on the visual signal and the model’s learned physics, minimizing semantic hallucination. We empirically select features from the specific intermediate DiT layer l , as they offer an optimal trade-off between spatial precision and abstract spatiotemporal context:

$$\mathbf{f}_{\text{raw}} = \Phi^{(l)}(\mathbf{z}_k, k; \mathbf{c}_{\text{text}} = ""). \quad (5)$$

After Adaptive Average Pooling to match the semantic tokenization, we obtain the generative representation $\mathbf{f}_{\text{gen}} \in \mathbb{R}^{T \times N \times D_{\text{gen}}}$. While this noise-driven process effectively extracts implicit 3D knowledge, these continuous physical features inherently misalign with the MLLM’s discrete semantic space. To bridge this semantic-geometric gap, **VEGA-3D** introduces a tailored fusion strategy.

4.3 Bridging the Generative and Semantic Gap

The generative features $\mathbf{f}_{\text{gen}}$ and semantic features $\mathbf{f}_{\text{sem}}$ reside in fundamentally different manifolds. To effectively synergize them, we introduce a mechanism to bridge this gap.

As shown in Fig. 5, we first project both streams into the LLM’s hidden dimension D_{llm} via independent MLP projectors P_{gen} and P_{sem} , aligning them to a shared embedding space:

$$\mathbf{F}_{\text{gen}} = P_{\text{gen}}(\mathbf{f}_{\text{gen}}), \mathbf{F}_{\text{sem}} = P_{\text{sem}}(\mathbf{f}_{\text{sem}}). \quad (6)$$

Here, $\mathbf{F}_{\text{gen}}, \mathbf{F}_{\text{sem}} \in \mathbb{R}^{T \times N \times D_{\text{llm}}}$, where T is the number of frames and N is the number of tokens per frame.

To avoid simply averaging conflicting signals, we employ an **Adaptive Gated Fusion** mechanism. This allows the model to adaptively weigh semantic versus structural cues for each specific token location. For the i -th spatial token $\mathbf{F}_i$, we compute a scalar gate $g_i \in [0, 1]$:

$$g_i = \sigma(\mathbf{W}_g^\top \text{Concat}(\text{LN}(\mathbf{F}_{\text{gen},i}), \text{LN}(\mathbf{F}_{\text{sem},i})) + b_g), \quad (7)$$

where $\sigma(\cdot)$ is the sigmoid function, LN denotes Layer Normalization, and $\mathbf{W}_g$ is a learnable weight vector. The final fused representation is a convex combination determined by this gate:

$$\mathbf{F}_i^{\text{fused}} = (1 - g_i) \cdot \mathbf{F}_{\text{gen},i} + g_i \cdot \mathbf{F}_{\text{sem},i}. \quad (8)$$

Crucially, this gate g_i acts as a semantic-geometric arbitrator: it enables the model to prioritize semantic priors for recognition tasks, while dynamically shifting attention to generative world knowledge for tasks requiring spatial reasoning. By seamlessly integrating continuous spatial priors with discrete semantic representations, **VEGA-3D** overcomes the spatial blindness of traditional encoders, achieving dense 3D understanding without explicit geometric supervision.

Table 1: Performance comparison on 3D scene understanding benchmarks. Specialists are single-task methods, while generalists target multiple tasks. † indicates the model is trained on extra datasets; ‡ indicates the model additionally uses BEV (bird’s-eye-view) reconstruction supervision rendered from depth and camera poses. **Bold** marks the best result in each column, and the Avg. Rank is recomputed over all listed methods. The baseline model is Video-3D LLM.

Method	Ref.	ScanRefer		Multi3DRefer		Scan2Cap		ScanQA		SQA3D	Avg.
		Acc. ₂₅	Acc. ₅	F1 _{.25}	F1 _{.5}	C _{.5}	B-4 _{.5}	C	EM	EM	Rank
Specialists											
ScanRefer [10]	ECCV 20	37.3	24.3	–	–	–	–	–	–	–	18.5
MVT [39]	CVPR 22	40.8	33.3	–	–	–	–	–	–	–	17.5
3DVG-Trans [95]	ICCV 21	45.9	34.5	–	–	–	–	–	–	–	16.0
ViL3DRel [12]	NeurIPS 21	47.9	37.7	–	–	–	–	–	–	–	14.5
M3DRef-CLIP [93]	ICCV 23	51.9	44.7	42.8	–	38.4	–	–	–	–	11.5
Scan2Cap [19]	CVPR 21	–	–	–	–	35.2	22.4	–	–	–	17.5
ScanQA [2]	CVPR 22	–	–	–	–	–	–	64.9	21.1	47.2	14.3
3D-VisTA [104]	ICCV 23	50.6	45.8	–	–	66.9	34.0	69.6	22.4	48.5	12.4
Generalists											
Chat-3D [78]	Arxiv 23	–	–	–	–	–	–	53.2	–	–	18.0
Chat-3D v2 [37]	Arxiv 23	42.5	38.4	45.1	41.6	63.9	31.8	87.6	–	54.7	12.5
LL3DA [13]	CVPR 24	–	–	–	–	62.9	36.0	76.8	–	–	14.0
LEO [38]	ICML 24	–	–	–	–	72.4	38.2	101.4	21.5	50.0	10.0
Grounded3D-LLM [14]	Arxiv 24	47.9	44.1	45.2	40.6	70.6	35.5	72.7	–	–	12.3
PQ3D [105]	ECCV 24	57.0	51.2	–	50.1	80.3	36.0	–	–	47.1	9.2
ChatScene [37]	NeurIPS 24	55.5	50.2	57.1	52.4	77.1	36.3	87.7	21.6	54.6	8.8
SceneLLM [29]	WACV 25	–	–	–	–	–	–	80.0	27.2	53.6	10.3
Inst3D-LLM [87]	CVPR 25	57.8	51.6	58.3	53.5	79.7	38.3	88.6	24.6	–	6.8
3D-LLaVA [22]	CVPR 25	51.2	40.6	–	–	78.8	36.9	92.6	–	54.5	9.8
3DRS [41]	NeurIPS 25	62.9	56.1	60.4	54.9	86.1	41.6	104.8	30.3	60.6	2.7
LLaVA-3D [103]	ICCV 25	50.1	42.7	49.8	43.6	84.1	42.6	–	30.6	60.1	6.6
Ross3D† [76]	ICCV 25	61.1	54.4	59.6	54.3	81.3	43.4	107.0	30.8	63.0	2.6
LLaVA-4D† [101]	ICLR 26	–	53.2	–	54.3	85.3	45.7	97.8	–	–	3.4
Omni-View† [36]	ICLR 26	50.8	45.0	–	–	–	–	103.0	29.5	59.2	7.0
Fase3D [61]	CVPR 26	–	–	–	–	78.1	41.3	91.7	–	54.3	8.8
Video-3D LLM [99]	CVPR 25	58.1	51.7	58.0	52.7	83.8	41.3	102.1	30.1	58.6	5.1
VEGA-3D (Ours)	-	63.2	56.2	60.8	55.1	83.2	42.2	106.3	30.4	61.3	2.2

5 Experiments

5.1 Implementation Details

We evaluate our framework on three representative axes: (i) **3D scene understanding** on ScanRefer [10], Multi3DRefer [93], Scan2Cap [19], ScanQA [2], and SQA3D [60]; (ii) **spatial reasoning** on VSI-Bench [83] with diverse capability categories and (iii) **robotic manipulation** on LIBERO [55] with four task suites and their average success rate. These benchmarks and reported metrics follow the standard protocols summarized in our main tables.

For 3D scene understanding, we build upon Video-3D LLM [99] as our baseline generalist and select Wan2.1-T2V 1.3B [75] as the latent world simulator plus an adaptive gated fusion module. For VSI-Bench [83], we adopt Qwen2.5VL-7B [67] as the baseline and attach the same plug-and-play generative branch, and the training datasets follow VG-LLM [97]. For LIBERO [55], we start from OpenVLA-OFT [47] and inject generative priors into the visual stream before policy learning. This design keeps the overall training and evaluation pipelines

Table 2: The comparison with state-of-the-art models on VSI-Bench. *Spatial-Enhanced Models* are models that are specialized for spatial reasoning. † indicates the baseline model’s performance is finetuned on the same training dataset configurations. The baseline model is Qwen2.5VL-7B.

Model	Avg.	Obj. Count	Abs. Dist.	Obj. Size	Room Size	Rel. Dist.	Rel. Dir.	Route Plan	Appr. Order
		Numerical Answer			Multiple-Choice Answer				
<i>Proprietary Models (API)</i>									
GPT-4o [42]	34.0	46.2	5.3	43.8	38.2	37.0	41.3	31.5	28.5
Gemini-1.5-Pro [31]	45.4	56.2	30.9	64.1	43.6	51.3	46.3	36.0	34.6
Gemini-1.5-Flash [31]	42.1	49.8	30.8	53.5	54.4	37.7	41.0	31.5	37.8
<i>Open-source Models</i>									
LongVA-7B [91]	29.2	38.0	16.6	38.9	22.2	33.1	43.3	25.4	15.7
LongVILA-8B [15]	21.6	29.1	9.1	16.7	0.0	29.6	30.7	32.5	25.5
InternVL2-8B [18]	34.6	23.1	28.7	48.2	39.8	36.7	30.7	29.9	39.6
InternVL2-40B [18]	36.0	34.9	26.9	46.5	31.8	42.1	32.2	34.0	39.6
VILA-1.5-40B [59]	31.2	22.4	24.8	48.7	22.7	40.5	25.7	31.5	32.9
LLaVA-OneVision-7B [50]	32.4	47.7	20.2	47.4	12.3	42.5	35.2	29.4	24.4
LLaVA-OneVision-72B [50]	40.2	43.5	23.9	57.6	37.5	42.5	39.9	32.5	44.6
LLaVA-NeXT-Video-7B [56]	35.6	48.5	14.0	47.8	24.2	43.5	42.4	34.0	30.6
LLaVA-NeXT-Video-72B [56]	40.9	48.9	22.8	57.4	35.3	42.4	36.7	35.0	48.6
<i>Spatial-Enhanced Models</i>									
Video-R1-7B [27]	37.1	-	-	-	-	-	-	-	-
vsGRPO-V-7B [52]	40.7	59.9	29.6	50.8	48.3	35.4	35.6	34.0	31.5
SPAR-8B [89]	41.1	-	-	-	-	-	-	-	-
SpaceR-7B [64]	45.6	-	-	-	-	-	-	-	-
3DRS-7B [41]	45.9	68.7	34.8	53.6	56.6	40.9	43.2	30.4	39.2
VG-LLM-4B [98]	45.9	65.6	37.4	54.8	60.2	42.3	46.3	33.0	25.9
VG-LLM-8B [98]	50.1	67.2	38.0	59.3	63.2	47.0	43.9	33.0	49.4
Qwen2.5VL-7B † [67]	48.9	68.3	37.0	57.4	58.7	39.7	43.0	29.4	57.8
VEGA-3D (Ours)	50.5	69.7	35.9	58.0	60.8	45.1	43.1	30.9	60.5

consistent with the corresponding baselines, while isolating the effect of generative priors. More details are provided in **Supplementary**.

For both training and inference, we uniformly sample 32 frames per scan to construct multi-view image sets. The Flow-Matching time interval $t \in [0, 1]$ is discretized into $K=1000$ steps in the pretrained Wan2.1 backbone. We denote the discrete timestep index as k and use $t_k = \frac{k}{K}$ as the normalized time. By default, we extract features at $k=300$ (i.e., $t_k=0.3$) from the 20th DiT layer. When calculating the correspondence score, we use a voxel size of 0.1 for voxelization. All models are optimized using Adam, with a batch size of 128 and a warm-up ratio of 0.03. The learning rates are set to a maximum of 1×10^{-5} for the language model and 2×10^{-6} for the visual backbone during the warm-up period. We use 8 H100 NVIDIA GPUs for all experiments.

5.2 Main Results on 3D Scene Understanding

Tab. 1 reports the main results on five 3D scene understanding benchmarks, covering spatial grounding, dense captioning, and question answering. Overall, our VEGA-3D consistently improves over the Video-3D LLM [99] baseline across most metrics, particularly excelling in localization-centric tasks (e.g., boosting ScanRefer Acc@0.5 from 51.7 to 56.2, SQA3D EM 58.6 to 61.3).

This performance divergence across different tasks suggest an informative pattern about when generative priors help most. We observe notable gains in

grounding and spatial QA, where the implicit 3D structural awareness extracted from the generative backbone acts as a robust spatial anchor, which helps reduce the spatial ambiguity of standard MLLMs (see Fig. 2(b)). The slight CIDEr drop in Scan2Cap may indicate a semantic-geometry trade-off: emphasizing structural cues may weaken fine-grained lexical details. Our Adaptive Gated Fusion aims to balance this by token-wise weighting between the two streams.

Notably, in contrast to prevailing state-of-the-art methods that rely on heavy geometric supervision [41] via external 3D teachers [77], BEV reconstruction targets rendered from depth and camera poses [76], or curated 3D-heavy datasets [101], our improvements are achieved entirely without explicit 3D annotations. In particular, while the reconstruction-based Ross3D [76] attains higher ScanQA and SQA3D scores through its BEV supervision, **VEGA-3D** still achieves the best overall average rank and leads on every grounding-oriented metric, suggesting that label-free generative priors are competitive with explicit geometric supervision for localization-centric understanding. This conveys a powerful insight: large-scale video generation models, through the pretext of temporal synthesis, have already internalized a robust 3D world model from the vast causality of the natural world. By repurposing these models as Latent World Simulators, we bypass the data-scarcity bottleneck posed by 3D-specific labels. This framework offers a highly scalable, data-efficient paradigm, demonstrating that the next frontier for 3D spatial awareness in MLLMs may not lie in more 3D data but in unleashing the physical priors already dormant within generative foundations.

5.3 Generalization to Spatial Reasoning and Manipulation

To validate the generalization ability of VEGA-3D, we extend our evaluation to 3D visual-spatial reasoning and embodied manipulation.

Spatial reasoning on VSI-Bench. Tab. 2 evaluates spatial reasoning on VSI-Bench [83], a comprehensive benchmark designed to diagnose diverse visual-spatial skills from videos, such as relative distance and route planning. By seamlessly augmenting the Qwen2.5VL-7B [67] baseline with our generative priors, we observe consistent gains across the overall average and multiple sub-categories. This trend aligns with the recent emphasis on geometry-aware mechanisms for spatial reasoning, yet our method remains lightweight and plug-and-play.

Table 3: The comparison of the simulation robotic manipulation benchmark LIBERO, the performance is evaluated with the average success rate SR (%).

Method	Reference	Spatial	Object	Goal	Long	Avg.
Diffusion Policy [20]	ICRR 23	78.3	92.5	68.3	50.5	72.4
Octo [72]	RSS 24	78.9	85.7	84.6	51.1	75.1
OpenVLA [48]	CoRL 24	84.7	88.4	79.2	53.7	76.5
DiT Policy [35]	ICCV 25	84.2	96.3	85.4	63.8	82.4
CoT-VLA [96]	CVPR 25	87.5	91.6	87.6	69.0	81.1
UniVLA [6]	RSS 25	96.5	96.8	95.6	92.0	95.2
OpenVLA-OFT [47]	RSS 25	97.5	98.3	97.8	94.4	97.0
VEGA-3D (Ours)	-	<u>97.4</u>	99.4	<u>97.0</u>	95.2	97.3

Robotic manipulation on LIBERO. Beyond passive reasoning, we assess whether our generative priors can ground active physical agents. Tab. 3 reports success rates on the LIBERO [55] suite, a challenging simulation benchmark for robotic manipulation. We treat **VEGA-3D** as a drop-in visual enhancement for Vision-Language-Action (VLA) models: starting from a pre-trained OpenVLA-OFT [47], we leave the language backbone and action head unchanged and only augment the visual stream, injecting the frozen Latent World Simulator features through the same adaptive gated fusion module used for scene understanding before policy learning. Despite being extracted without explicit action-conditioned training, the generative priors improve the already highly saturated baseline, with the clearest gains on the object-centric and long-horizon suites that demand precise contact localization and persistent spatial memory. This indicates that the spatial regularities embedded in the Latent World Simulator transfer to action grounding, supplying the policy with a geometry-aware visual representation rather than additional action supervision.

Towards World-Action Models. These results suggest a shift towards World-Action Models (WAMs): instead of relying on explicit 3D encoders, a frozen video generator acts as a foundational World Model to directly ground the action policy. This synergy offers three key advantages: **(i) Modularity:** Spatial priors are injected at the visual-token level, leaving the policy and language backbones intact. **(ii) Scalability:** Physical intelligence stems from large-scale video pre-training, meaning advances in world simulation directly translate to stronger action grounding without extra robotic annotations.

Real-robot manipulation. To test transfer beyond simulation, we deploy our model on real Agilex Piper arms following the RoboTwin [63] protocol. On the Click Bell task, VEGA-3D improves over the baseline policy from 40% (8/20) to 55% (11/20), indicating that the prior can transfer to real-world manipulation.

5.4 Ablation Studies

To validate our design choices, we conduct comprehensive ablation studies using Wan2.1-T2V as the default generative encoder. We note that VACE can achieve higher QA scores, while T2V provides stronger grounding-oriented performance; thus, we keep T2V as the default encoder.

Generative vs. Discriminative Priors.

Tab. 4 and Fig. 3(b) reveal a strong positive correlation between multi-view consistency and 3D scene understanding. While traditional discriminative models (e.g., DINO V3 [71], V-JEPA v2 [1], SigLIP [88]) offer rich semantics, they lack explicit 3D consistency. Conversely, DiT-based generative models and 3D foundation model like VGGT [77] excel at capturing robust spatial priors. DiTs significantly outperform UNet-based models, as their global attention mechanisms preserve long-range geometric dependencies better than local convolutions. This confirms that video generation models provide superior 3D prior for spatial reasoning compared to standard visual learners.

Table 4: Experiments on using different discriminative and generative foundation models. Bold denotes the best in each group.

Models	Params	ScanRefer		Multi3DRefer		Scan2Cap		ScanQA		SQA3D
		Acc. ₂₅	Acc. ₅	F1. ₂₅	F1. ₅	C. ₅	B-4. ₅	C	EM	EM
Baseline	-	58.1	51.7	58.0	52.7	83.8	41.3	102.1	30.1	58.6
<i>Discriminative Models</i>										
V-JEPA v2 [1]	1B	61.7	54.9	60.2	54.7	79.8	41.5	106.6	30.7	61.2
DinoV3-Large [71]	0.3B	61.1	54.2	59.6	54.1	80.6	41.1	105.9	30.5	61.9
VGGT [77]	1.1B	62.3	55.3	60.1	54.5	82.8	42.0	105.8	30.5	61.4
<i>Generative Models</i>										
Stable Video Diffusion [4]	1.5B	61.3	54.8	59.9	54.6	80.9	41.9	105.1	30.1	61.3
Stable Diffusion 2.1 [70]	0.9B	62.1	55.1	60.3	54.9	83.0	42.0	106.8	30.4	60.6
Vmem [51]	1.3B	62.5	55.7	60.2	54.7	82.0	41.9	106.0	30.0	61.4
SEVA [102]	1.3B	62.3	55.5	60.1	54.5	82.5	42.1	107.6	30.8	60.9
VAE [70]	0.08B	62.0	55.1	60.3	54.8	83.7	42.3	106.0	30.5	61.4
Wan2.1-VACE [44]	1.3B	62.2	55.3	60.3	55.0	82.8	42.7	107.8	31.0	61.8
Wan2.1-T2V [75]	1.3B	63.2	56.2	60.8	55.1	83.2	42.2	106.3	30.4	61.3

Dynamics of Internal Representations: Noise Levels and Layers. The generative prior varies significantly across the diffusion process and network depth (Fig. 6). **(i) Noise Ratio t_k :** Performance peaks at intermediate noise level. Clean latents underutilize the model’s denoising capabilities, while excessive noise destroys structural signals. Moderate noise optimally forces the model to engage its learned physics to restore underlying 3D structures. **(ii) Layer Selection:** Intermediate layers provide the optimal abstraction for spatial reasoning, effectively balancing the low-level textures of early layers and the pixel-level rendering of deeper layers.

Effectiveness of Adaptive Gated Fusion.

Tab. 5 demonstrates the necessity of our Adaptive Gated Fusion. Relying solely on generative features causes a substantial performance drop, confirming that generative priors complement rather than replace semantic representations. Among lightweight fusion variants, our method achieves the best overall trade-off: it outperforms other fusion baselines on most metrics, and matches the best results on ScanQA (C/EM). We note that a naive Add operation attains a slightly higher SQA3D EM (61.8 vs. 61.3), but it is consistently weaker on grounding and captioning metrics, suggesting that a fixed, non-adaptive fusion weight cannot reliably resolve the semantic–generative distribution gap. In contrast, our token-level gating dynamically balances semantic and geometric priors, yielding more consistent gains across diverse tasks.




Fig. 6: Ablation studies on noise injection and DiT depth. (a) Performance peaks at intermediate noise levels. (b) Specific intermediate layers capture the most robust geometric cues.

Table 5: Ablation study of the effects of different feature fusion modules. All models are finetuned with the same training data and built on Wan2.1-T2V 1.3B at sampling step $k = 300$ and the 20th layer feature.

Type	ScanRefer		Multi3DRefer		Scan2Cap		ScanQA		SQA3D
	Acc.25	Acc.5	F1.25	F1.5	C.5	B-4.5	C	EM	EM
Baseline	58.1	51.7	58.0	52.7	83.8	41.3	102.1	30.1	58.6
Only generative features	54.9	48.3	53.7	48.6	25.2	30.0	74.0	21.1	52.0
Add	61.5	54.6	59.6	54.1	81.4	41.6	106.3	30.4	61.8
Channel Concat+MLP	55.1	48.9	53.6	48.7	33.2	31.8	81.6	22.9	52.3
Sequence Concat	59.5	53.0	58.4	53.2	79.4	41.0	104.7	30.2	61.5
Cross-Attn (1 Layer)	58.5	51.9	57.9	52.6	48.8	34.7	104.9	29.6	61.0
Cross-Attn (3 Layers)	58.0	51.5	57.5	52.1	47.8	34.8	102.2	29.2	60.5
Channel-Level-Gated	61.8	54.9	60.0	54.4	82.2	41.9	105.7	30.3	61.2
Adaptive-Gated-Fusion(Ours)	63.2	56.2	60.8	55.1	83.2	42.2	106.3	30.4	61.3




Fig. 7: Inference overhead We cache the **Wan2.1-T2V** features once per scene and reuse them for all questions, substantially reducing inference overhead. The grey dashed line indicates the baseline result without the generative branch.

6 Conclusion

We introduce **VEGA-3D**, a plug-and-play framework that repurposes modern video generation models as *Latent World Simulators* to mitigate the spatial blindness of MLLMs. By activating these priors via noise injection and aligning them with semantic tokens through *Adaptive Gated Fusion*, VEGA-3D injects dense geometric anchors into MLLMs, consistently improving scene understanding, spatial reasoning, and manipulation *without* extra 3D supervision.

Limitations and Future Work. Incorporating a video diffusion backbone increases inference cost (Fig. 7). Nevertheless, it delivers substantial and consistent performance gains, making the trade-off acceptable in practice. Future work will distill these priors into lightweight encoders, extend the framework to more dynamic scene understanding, and scale the plug-and-play generative prior to larger Vision-Language-Action models for closed-loop embodied control.

Acknowledgements. This work was supported in part by the National Natural Science Foundation of China (NSFC) under Grants 62441615 and 623B2038, and in part by the Hubei Science and Technology Major Project under Grant 2024BAA007.

References

- [1] Assran, M., Bardes, A., Fan, D., Garrido, Q., Howes, R., Muckley, M., Rizvi, A., Roberts, C., Sinha, K., Zholus, A., et al.: V-jepa 2: Self-supervised video models enable understanding, prediction and planning. arXiv preprint arXiv:2506.09985 (2025)
- [2] Azuma, D., Miyanishi, T., Kurita, S., Kawanabe, M.: Scanqa: 3d question answering for spatial scene understanding. In: Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (2022), license: Creative Commons Attribution-NonCommercial-ShareAlike 3.0 Unported License
- [3] Ball, P.J., Bauer, J., Belletti, F., Brownfield, B., Ephrat, A., Fruchter, S., Gupta, A., Holsheimer, K., Holynski, A., Hron, J., Kaplanis, C., Limont, M., McGill, M., Oliveira, Y., Parker-Holder, J., Perbet, F., Scully, G., Shar, J., Spencer, S., Tov, O., Villegas, R., Wang, E., Yung, J., Baetu, C., Berbel, J., Bridson, D., Bruce, J., Buttimore, G., Chakera, S., Chandra, B., Collins, P., Cullum, A., Damoc, B., Dasagi, V., Gazeau, M., Gbadamosi, C., Han, W., Hirst, E., Kachra, A., Kerley, L., Kjems, K., Knoepfel, E., Koriakin, V., Lo, J., Lu, C., Mehring, Z., Moufarek, A., Nandwani, H., Oliveira, V., Pardo, F., Park, J., Pierson, A., Poole, B., Ran, H., Salimans, T., Sanchez, M., Saprykin, I., Shen, A., Sidhwani, S., Smith, D., Stanton, J., Tomlinson, H., Vijaykumar, D., Wang, L., Wingfield, P., Wong, N., Xu, K., Yew, C., Young, N., Zubov, V., Eck, D., Erhan, D., Kavukcuoglu, K., Hassabis, D., Gharamani, Z., Hadsell, R., van den Oord, A., Mosseri, I., Bolton, A., Singh, S., Rocktäschel, T.: Genie 3: A new frontier for world models (2025), <https://deepmind.google/blog/genie-3-a-new-frontier-for-world-models/>, accessed: 28 June 2026
- [4] Blattmann, A., Dockhorn, T., Kulal, S., Mendelevitch, D., Kilian, M., Lorenz, D., Levi, Y., English, Z., Voleti, V., Letts, A., et al.: Stable video diffusion: Scaling latent video diffusion models to large datasets. arXiv preprint arXiv:2311.15127 (2023)
- [5] Brooks, T., Peebles, B., Holmes, C., DePue, W., Guo, Y., Jing, L., Schnurr, D., Taylor, J., Luhman, T., Luhman, E., Ng, C., Wang, R., Ramesh, A.: Video generation models as world simulators (2024), <https://openai.com/research/video-generation-models-as-world-simulators>, accessed: 28 June 2026
- [6] Bu, Q., Yang, Y., Cai, J., Gao, S., Ren, G., Yao, M., Luo, P., Li, H.: Uni-vla: Learning to act anywhere with task-centric latent actions. Robotics: Science and Systems (RSS) (2025)
- [7] Cao, M., Lin, H., Li, H., Tang, H., Xu, R., An, D., Liu, X., Reid, I., Liang, X.: Seeing through imagination: Learning scene geometry via implicit spatial world modeling (2025), <https://arxiv.org/abs/2512.01821>, accessed: 28 June 2026
- [8] Cheang, C.L., Chen, G., Jing, Y., Kong, T., Li, H., Li, Y., Liu, Y., Wu, H., Xu, J., Yang, Y., et al.: Gr-2: A generative video-language-action model with web-scale knowledge for robot manipulation. arXiv preprint arXiv:2410.06158 (2024)

- [9] Chen, B., Xu, Z., Kirmani, S., Ichter, B., Sadigh, D., Guibas, L., Xia, F.: Spatialvlm: Endowing vision-language models with spatial reasoning capabilities. In: Proc. IEEE Conf. Comput. Vis. Pattern Recognit. pp. 14455–14465 (2024)
- [10] Chen, D.Z., Chang, A.X., Nießner, M.: Scanrefer: 3d object localization in rgb-d scans using natural language. In: Proc. Eur. Conf. Comput. Vis. (2020), license: Creative Commons Attribution-NonCommercial-ShareAlike 3.0 Unported License
- [11] Chen, D., Shukor, M., Moutakanni, T., Chung, W., Yu, J., Kasarla, T., Bolourchi, A., LeCun, Y., Fung, P.: Vl-jepa: Joint embedding predictive architecture for vision-language. In: Proc. Int. Conf. Learn. Representations (2026)
- [12] Chen, S., Guhur, P., Tapaswi, M., Schmid, C., Laptev, I.: Language conditioned spatial relation reasoning for 3d object grounding. In: Proc. Adv. Neural Inf. Process. Syst. (2022)
- [13] Chen, S., Chen, X., Zhang, C., Li, M., Yu, G., Fei, H., Zhu, H., Fan, J., Chen, T.: LL3DA: visual interactive instruction tuning for omni-3d understanding, reasoning, and planning. In: Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (2024)
- [14] Chen, Y., Yang, S., Huang, H., Wang, T., Xu, R., Lyu, R., Lin, D., Pang, J.: Grounded 3d-llm with referent tokens. arXiv preprint arXiv:2405.10370 (2024)
- [15] Chen, Y., Xue, F., Li, D., Hu, Q., Zhu, L., Li, X., Fang, Y., Tang, H., Yang, S., Liu, Z., et al.: Longvila: Scaling long-context visual language models for long videos. In: Proc. Int. Conf. Learn. Representations (2025)
- [16] Chen, Z., Zhang, X., Xu, H., Xie, J., Tu, Z.: Cvp: Central-peripheral vision-inspired multimodal model for spatial reasoning. arXiv preprint arXiv:2512.08135 (2025)
- [17] Chen, Z., Zhang, M., Yu, X., Luo, X., Sun, M., Pan, Z., Feng, Y., Pei, P., Cai, X., Huang, R.: Think with 3d: Geometric imagination grounded spatial reasoning from limited views. arXiv preprint arXiv:2510.18632 (2025)
- [18] Chen, Z., Wang, W., Tian, H., Ye, S., Gao, Z., Cui, E., Tong, W., Hu, K., Luo, J., Ma, Z., et al.: How far are we to gpt-4v? closing the gap to commercial multimodal models with open-source suites. *Science China Information Sciences* **67**(12), 220101 (2024)
- [19] Chen, Z., Gholami, A., Nießner, M., Chang, A.X.: Scan2cap: Context-aware dense captioning in rgb-d scans. In: Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (2021), license: Creative Commons Attribution-NonCommercial-ShareAlike 3.0 Unported License
- [20] Chi, C., Xu, Z., Feng, S., Cousineau, E., Du, Y., Burchfiel, B., Tedrake, R., Song, S.: Diffusion policy: Visuomotor policy learning via action diffusion. *The International Journal of Robotics Research* (2023)
- [21] Dai, A., Chang, A.X., Savva, M., Halber, M., Funkhouser, T., Nießner, M.: Scannet: Richly-annotated 3d reconstructions of indoor scenes. In: Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (2017), license: ScanNet Terms of Use

- [22] Deng, J., He, T., Jiang, L., Wang, T., Dayoub, F., Reid, I.: 3d-llava: Towards generalist 3d lmms with omni superpoint transformer. In: Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (2025)
- [23] Deng, T., Chen, X., Chen, Y., Chen, Q., Xu, Y., Yang, L., Xu, L., Zhang, Y., Zhang, B., Huang, W., Wang, H.: Gaussiandwm: 3d gaussian driving world model for unified scene understanding and multi-modal generation. In: Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (2026)
- [24] Deng, T., Pan, Y., Yuan, S., Li, D., Wang, C., Li, M., Chen, L., Xie, L., Wang, D., Wang, J., Civera, J., Wang, H., Chen, W.: What is the best 3d scene representation for robotics? from geometric to foundation models. arXiv preprint arXiv:2512.03422 (2025)
- [25] Fan, Z., Zhang, J., Li, R., Zhang, J., Chen, R., Hu, H., Wang, K., Qu, H., Wang, D., Yan, Z., et al.: Vlm-3r: Vision-language models augmented with instruction-aligned 3d reconstruction. In: Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (2026)
- [26] Fang, H., Li, S., Wang, S., Xi, X., Liang, D., Bai, X.: Towards generalizable robotic manipulation in dynamic environments. arXiv preprint arXiv:2603.15620 (2026)
- [27] Feng, K., Gong, K., Li, B., Guo, Z., Wang, Y., Peng, T., Wu, J., Zhang, X., Wang, B., Yue, X.: Video-r1: Reinforcing video reasoning in mllms. In: Proc. Adv. Neural Inf. Process. Syst. (2025)
- [28] Fu, H., Zhang, D., Zhao, Z., Cui, J., Xie, H., Wang, B., Chen, G., Liang, D., Bai, X.: Minddrive: A vision-language-action model for autonomous driving via online reinforcement learning. arXiv preprint arXiv:2512.13636 (2025)
- [29] Fu, R., Liu, J., Chen, X., Nie, Y., Xiong, W.: Scene-llm: Extending language model for 3d visual reasoning. In: Proceedings of the Winter Conference on Applications of Computer Vision (WACV). pp. 2195–2206 (February 2025)
- [30] Gao, J., Sarkar, B., Xia, F., Xiao, T., Wu, J., Ichter, B., Majumdar, A., Sadigh, D.: Physically grounded vision-language models for robotic manipulation. In: Proc. IEEE Int. Conf. Robotics Automation. pp. 12462–12469. IEEE (2024)
- [31] Gemini Team: Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. arXiv preprint arXiv:2403.05530 (2024)
- [32] Guan, Y., Yin, L., Liang, D., Ju, J., Luo, Z., Luan, J., Liu, Y., Bai, X.: Video streaming thinking: Videollms can watch and think simultaneously. arXiv preprint arXiv:2603.12262 (2026)
- [33] Guo, Z., Zhang, R., Zhu, X., Tang, Y., Ma, X., Han, J., Chen, K., Gao, P., Li, X., Li, H., et al.: Point-bind & point-llm: Aligning point cloud with multi-modality for 3d understanding, generation, and instruction following. arXiv preprint arXiv:2309.00615 (2023)
- [34] Hong, W., Ding, M., Zheng, W., Liu, X., Tang, J.: Cogvideo: Large-scale pretraining for text-to-video generation via transformers. In: Proc. Int. Conf. Learn. Representations (2023)

- [35] Hou, Z., Zhang, T., Xiong, Y., Duan, H., Pu, H., Tong, R., Zhao, C., Zhu, X., Qiao, Y., Dai, J., et al.: Dita: Scaling diffusion transformer for generalist vision-language-action policy. In: Proc. IEEE Int. Conf. Comput. Vis. pp. 7686–7697 (2025)
- [36] Hu, J., Zhao, S., Chen, Q.G., Qiu, X., Liu, J., Xu, Z., Luo, W., Zhang, K., Lu, Y.: Omni-view: Unlocking how generation facilitates understanding in unified 3d model based on multiview images. In: Proc. Int. Conf. Learn. Representations (2026)
- [37] Huang, H., Wang, Z., Huang, R., Liu, L., Cheng, X., Zhao, Y., Jin, T., Zhao, Z.: Chat-3d v2: Bridging 3d scene and large language models with object identifiers. In: Proc. Adv. Neural Inf. Process. Syst. (2024)
- [38] Huang, J., Yong, S., Ma, X., Linghu, X., Li, P., Wang, Y., Li, Q., Zhu, S., Jia, B., Huang, S.: An embodied generalist agent in 3d world. In: Proc. Int. Conf. Mach. Learn. (2024)
- [39] Huang, S., Chen, Y., Jia, J., Wang, L.: Multi-view transformer for 3d visual grounding. In: Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (2022)
- [40] Huang, S., Wu, J., Zhou, Q., Miao, S., Long, M.: Vid2world: Crafting video diffusion models to interactive world models. In: Proc. Int. Conf. Learn. Representations (2026)
- [41] Huang, X., Wu, J., Xie, Q., Han, K.: 3drs: Mllms need 3d-aware representation supervision for scene understanding. In: Proc. Adv. Neural Inf. Process. Syst. (2025)
- [42] Hurst, A., Lerer, A., Goucher, A.P., Perelman, A., Ramesh, A., Clark, A., Ostrow, A., Welihinda, A., Hayes, A., Radford, A., et al.: Gpt-4o system card. arXiv preprint arXiv:2410.21276 (2024)
- [43] Jia, M., Qi, Z., Zhang, S., Zhang, W., Yu, X., He, J., Wang, H., Yi, L.: Omnispatial: Towards comprehensive spatial reasoning benchmark for vision language models. In: Proc. Int. Conf. Learn. Representations (2026)
- [44] Jiang, Z., Han, Z., Mao, C., Zhang, J., Pan, Y., Liu, Y.: Vace: All-in-one video creation and editing. In: Proc. IEEE Int. Conf. Comput. Vis. pp. 17191–17202 (2025)
- [45] Kang, B., Yue, Y., Lu, R., Lin, Z., Zhao, Y., Wang, K., Huang, G., Feng, J.: How far is video generation from world model: A physical law perspective. In: Proc. Int. Conf. Mach. Learn. (2025)
- [46] Kim, G., Han, J., Cho, S.: Videofrom3d: 3d scene video generation via complementary image and video diffusion models. In: Proceedings of the SIGGRAPH Asia 2025 Conference Papers. pp. 1–11 (2025)
- [47] Kim, M.J., Finn, C., Liang, P.: Fine-tuning vision-language-action models: Optimizing speed and success. Robotics: Science and Systems (RSS) (2025)
- [48] Kim, M.J., Pertsch, K., Karamcheti, S., Xiao, T., Balakrishna, A., Nair, S., Rafailov, R., Foster, E., Lam, G., Sanketi, P., et al.: Openvla: An open-source vision-language-action model. In: Proceedings of the 8th Conference on Robot Learning. pp. 2679–2713. PMLR (2024)
- [49] Kondratyuk, D., Yu, L., Gu, X., Lezama, J., Huang, J., Schindler, G., Hornung, R., Birodkar, V., Yan, J., Chiu, M.C., et al.: Videopoet: A large language model for zero-shot video generation. In: Proc. Int. Conf. Mach. Learn. (2024)

- [50] Li, B., Zhang, Y., Guo, D., Zhang, R., Li, F., Zhang, H., Zhang, K., Zhang, P., Li, Y., Liu, Z., et al.: Llava-onevision: Easy visual task transfer. *Transactions on Machine Learning Research* (2025)
- [51] Li, R., Torr, P., Vedaldi, A., Jakab, T.: Vmem: Consistent interactive video scene generation with surfel-indexed view memory. In: *Proc. IEEE Int. Conf. Comput. Vis.* (2025)
- [52] Liao, Z., Xie, Q., Zhang, Y., Kong, Z., Lu, H., Yang, Z., Deng, Z.: Improved visual-spatial reasoning via r1-zero-like training. *arXiv preprint arXiv:2504.00883* (2025)
- [53] Lin, J., Zhu, C., Xu, R., Mao, X., Liu, X., Wang, T., Pang, J.: Ost-bench: Evaluating the capabilities of mllms in online spatio-temporal scene understanding. In: *Proc. Adv. Neural Inf. Process. Syst.* (2025)
- [54] Lipman, Y., Chen, R.T., Ben-Hamu, H., Nickel, M., Le, M.: Flow matching for generative modeling. In: *Proc. Int. Conf. Learn. Representations* (2023)
- [55] Liu, B., Zhu, Y., Gao, C., Feng, Y., Liu, Q., Zhu, Y., Stone, P.: Libero: Benchmarking knowledge transfer for lifelong robot learning. *Advances in Neural Information Processing Systems* **36** (2024)
- [56] Liu, H., Li, C., Li, Y., Li, B., Zhang, Y., Shen, S., Lee, Y.J.: Llava-next: Improved reasoning, ocr, and world knowledge (January 2024), <https://llava-vl.github.io/blog/2024-01-30-llava-next/>, accessed: 28 June 2026
- [57] Liu, H., Li, C., Wu, Q., Lee, Y.J.: Visual instruction tuning. *Advances in neural information processing systems* **36**, 34892–34916 (2023)
- [58] Liu, Y., Liu, L., Zheng, Y., Liu, Y., Dang, F., Li, N., Ma, K.: Embodied navigation. *Science China Information Sciences* **68**(4), 1–39 (2025)
- [59] Liu, Z., Zhu, L., Shi, B., Zhang, Z., Lou, Y., Yang, S., Xi, H., Cao, S., Gu, Y., Li, D., et al.: Nvila: Efficient frontier visual language models. In: *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.* pp. 4122–4134 (2025)
- [60] Ma, X., Yong, S., Zheng, Z., Li, Q., Liang, Y., Zhu, S., Huang, S.: SQA3D: situated question answering in 3d scenes. In: *Proc. Int. Conf. Learn. Representations* (2023), license: CC-BY-4.0
- [61] Mei, G., Lin, W., Riz, L., Wu, Y., Wang, Y., Poesi, F.: Efficient encoder-free fourier-based 3d large multimodal model. In: *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.* (2026)
- [62] Miao, S., Feng, N., Wu, J., Lin, Y., He, X., Li, D., Long, M.: JEPA-VLA: Video predictive embedding is needed for VLA models. *arXiv preprint arXiv:2602.11832* (2026)
- [63] Mu, Y., Chen, T., et al.: RoboTwin: Dual-arm robot benchmark with generative digital twins. In: *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.* (2025)
- [64] Ouyang, K., Liu, Y., Wu, H., Liu, Y., Zhou, H., Zhou, J., Meng, F., Sun, X.: Spacer: Reinforcing mllms in video spatial reasoning. *arXiv preprint arXiv:2504.01805* (2025)
- [65] Qi, Z., Fang, Y., Sun, Z., Wu, X., Wu, T., Wang, J., Lin, D., Zhao, H.: Gpt4point: A unified framework for point-language understanding and generation. In: *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.* (2024)

- [66] Qi, Z., Zhang, Z., Fang, Y., Wang, J., Zhao, H.: Gpt4scene: Understand 3d scenes from videos with vision-language models. In: Proc. Int. Conf. Learn. Representations (2026)
- [67] Qwen Team: Qwen2.5-vl technical report. arXiv preprint arXiv:2502.13923 (2025)
- [68] Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: Proc. Int. Conf. Mach. Learn. pp. 8748–8763. PMLR (2021)
- [69] Ren, X., Shen, T., Huang, J., Ling, H., Lu, Y., Nimier-David, M., Müller, T., Keller, A., Fidler, S., Gao, J.: Gen3c: 3d-informed world-consistent video generation with precise camera control. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 6121–6132 (2025)
- [70] Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proc. IEEE Conf. Comput. Vis. Pattern Recognit. pp. 10684–10695 (2022)
- [71] Siméoni, O., Vo, H.V., Seitzer, M., Baldassarre, F., Oquab, M., Jose, C., Khalidov, V., Szafraniec, M., Yi, S., Ramamonjisoa, M., et al.: Dinov3. arXiv preprint arXiv:2508.10104 (2025)
- [72] Team, O.M., Ghosh, D., Walke, H., Pertsch, K., Black, K., Mees, O., Dasari, S., Hejna, J., Kreiman, T., Xu, C., et al.: Octo: An open-source generalist robot policy. Robotics: Science and Systems (RSS) (2024)
- [73] Tschannen, M., Gritsenko, A., Wang, X., Naeem, M.F., Alabdulmohsin, I., Parthasarathy, N., Evans, T., Beyer, L., Xia, Y., Mustafa, B., et al.: Siglip 2: Multilingual vision-language encoders with improved semantic understanding, localization, and dense features. arXiv preprint arXiv:2502.14786 (2025)
- [74] Valevski, D., Leviathan, Y., Arar, M., Fruchter, S.: Diffusion models are real-time game engines. In: Proc. Int. Conf. Learn. Representations (2025)
- [75] Wan, T., Wang, A., Ai, B., Wen, B., Mao, C., Xie, C.W., Chen, D., Yu, F., Zhao, H., Yang, J., et al.: Wan: Open and advanced large-scale video generative models. arXiv preprint arXiv:2503.20314 (2025)
- [76] Wang, H., Zhao, Y., Wang, T., Fan, H., Zhang, X., Zhang, Z.: Ross3d: Reconstructive visual instruction tuning with 3d-awareness. In: Proc. IEEE Int. Conf. Comput. Vis. (2025)
- [77] Wang, J., Chen, M., Karaev, N., Vedaldi, A., Rupprecht, C., Novotny, D.: Vggt: Visual geometry grounded transformer. In: Proc. IEEE Conf. Comput. Vis. Pattern Recognit. pp. 5294–5306 (2025)
- [78] Wang, Z., Huang, H., Zhao, Y., Zhang, Z., Zhao, Z.: Chat-3d: Data-efficiently tuning large language model for universal dialogue of 3d scenes. arXiv preprint arXiv:2308.08769 (2023)
- [79] Xiao, Z., Lan, Y., Zhou, Y., Ouyang, W., Yang, S., Zeng, Y., Pan, X.: Worldmem: Long-term consistent world simulation with memory. In: Proc. Adv. Neural Inf. Process. Syst. (2025)
- [80] Xu, G., Zhang, Q., Zhou, J., Zhu, X., Shen, Y., Yang, X., Xu, Y.: Next forcing: Causal world modeling with multi-chunk prediction. arXiv preprint arXiv:2606.11187 (2026)

- [81] Xu, R., Wang, X., Wang, T., Chen, Y., Pang, J., Lin, D.: Pointllm: Empowering large language models to understand point clouds. In: Proc. Eur. Conf. Comput. Vis. (2024)
- [82] Xu, W., Shi, C., Tu, S., Zhou, X., Liang, D., Bai, X.: A unified framework for 3d scene understanding. *Advances in Neural Information Processing Systems* **37**, 59468–59490 (2024)
- [83] Yang, J., Yang, S., Gupta, A.W., Han, R., Fei-Fei, L., Xie, S.: Thinking in space: How multimodal large language models see, remember, and recall spaces. In: Proc. IEEE Conf. Comput. Vis. Pattern Recognit. pp. 10632–10643 (2025)
- [84] Yang, S., Xu, R., Xie, Y., Yang, S., Li, M., Lin, J., Zhu, C., Chen, X., Duan, H., Yue, X., et al.: Mmsi-bench: A benchmark for multi-image spatial intelligence. In: Proc. Int. Conf. Learn. Representations (2026)
- [85] Yang, Z., Teng, J., Zheng, W., Ding, M., Huang, S., Xu, J., Yang, Y., Hong, W., Zhang, X., Feng, G., et al.: Cogvideox: Text-to-video diffusion models with an expert transformer. In: Proc. Int. Conf. Learn. Representations (2025)
- [86] Yin, B., Wang, Q., Zhang, P., Zhang, J., Wang, K., Wang, Z., Zhang, J., Chandrasegaran, K., Liu, H., Krishna, R., et al.: Spatial mental modeling from limited views. In: Structural Priors for Vision Workshop at ICCV’25 (2025)
- [87] Yu, H., Li, W., Wang, S., Chen, J., Zhu, J.: Inst3d-lmm: Instance-aware 3d scene understanding with multi-modal instruction tuning. In: Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (2025)
- [88] Zhai, X., Mustafa, B., Kolesnikov, A., Beyer, L.: Sigmoid loss for language image pre-training. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 11975–11986 (2023)
- [89] Zhang, J., Chen, Y., Zhou, Y., Xu, Y., Huang, Z., Mei, J., Chen, J., Yuan, Y.J., Cai, X., Huang, G., et al.: From flatland to space: Teaching vision-language models to perceive and reason in 3d. In: Proc. IEEE Int. Conf. Comput. Vis. (2025)
- [90] Zhang, K., Ge, K., Chi, X., Zhang, R., Shi, S., Dong, Z., Han, S., Zhang, S.: Can world models benefit VLMs for world dynamics? arXiv preprint arXiv:2510.00855 (2025)
- [91] Zhang, P., Zhang, K., Li, B., Zeng, G., Yang, J., Zhang, Y., Wang, Z., Tan, H., Li, C., Liu, Z.: Long context transfer from language to vision. arXiv preprint arXiv:2406.16852 (2024)
- [92] Zhang, W., Ng, W.E., Ma, L., Wang, Y., Zhao, J., Koenecke, A., Li, B., Wanglu, W.: Sphere: Unveiling spatial blind spots in vision-language models through hierarchical evaluation. In: Proc. Annual Meeting of the Association for Computational Linguistics. pp. 11591–11609 (2025)
- [93] Zhang, Y., Gong, Z., Chang, A.X.: Multi3drefer: Grounding text description to multiple 3d objects. In: Proc. IEEE Int. Conf. Comput. Vis. (2023), license: MIT
- [94] Zhang, Z., Zheng, H., Wang, Y., Xu, L., Deng, T., Chen, X., Chen, Q., Zhang, B., Huang, W.: Omnidrive-r1: Reinforcement-driven interleaved

- multi-modal chain-of-thought for trustworthy vision-language autonomous driving. In: Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (2026)
- [95] Zhao, L., Cai, D., Sheng, L., Xu, D.: 3dvg-transformer: Relation modeling for visual grounding on point clouds. In: Proc. IEEE Int. Conf. Comput. Vis. (2021)
 - [96] Zhao, Q., Lu, Y., Kim, M.J., Fu, Z., Zhang, Z., Wu, Y., Li, Z., Ma, Q., Han, S., Finn, C., Handa, A., Liu, M.Y., Xiang, D., Wetzstein, G., Lin, T.Y.: Cot-vla: Visual chain-of-thought reasoning for vision-language-action models. In: Proc. IEEE Conf. Comput. Vis. Pattern Recognit. pp. 1702–1713 (2025)
 - [97] Zheng, D., Huang, S., Li, Y., Wang, L.: Learning from videos for 3d world: Enhancing mllms with 3d vision geometry priors. In: Proc. Adv. Neural Inf. Process. Syst. (2025)
 - [98] Zheng, D., Huang, S., Li, Y., Wang, L.: Learning from videos for 3d world: Enhancing mllms with 3d vision geometry priors. In: Proc. Adv. Neural Inf. Process. Syst. (2025)
 - [99] Zheng, D., Huang, S., Wang, L.: Video-3d llm: Learning position-aware video representation for 3d scene understanding. In: Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (2025)
 - [100] Zheng, D., Huang, S., Zhao, L., Zhong, Y., Wang, L.: Towards learning a generalist model for embodied navigation. In: Proc. IEEE Conf. Comput. Vis. Pattern Recognit. pp. 13624–13634 (2024)
 - [101] Zhou, H., Lee, G.H.: Llava-4d: Embedding spatiotemporal prompt into lmms for 4d scene understanding. In: Proc. Int. Conf. Learn. Representations (2026)
 - [102] Zhou, J., Gao, H., Voleti, V., Vasishta, A., Yao, C.H., Boss, M., Torr, P., Rupprecht, C., Jampani, V.: Stable virtual camera: Generative view synthesis with diffusion models. In: Proc. IEEE Int. Conf. Comput. Vis. (2025)
 - [103] Zhu, C., Wang, T., Zhang, W., Pang, J., Liu, X.: Llava-3d: A simple yet effective pathway to empowering lmms with 3d-awareness. In: Proc. IEEE Int. Conf. Comput. Vis. (2024)
 - [104] Zhu, Z., Ma, X., Chen, Y., Deng, Z., Huang, S., Li, Q.: 3d-vista: Pre-trained transformer for 3d vision and text alignment. In: Proc. IEEE Int. Conf. Comput. Vis. (2023)
 - [105] Zhu, Z., Zhang, Z., Ma, X., Niu, X., Chen, Y., Jia, B., Deng, Z., Huang, S., Li, Q.: Unifying 3d vision-language understanding via promptable queries. In: Proc. Eur. Conf. Comput. Vis. (2024)