

WorldArena: A Unified Benchmark for Evaluating Perception and Functional Utility of Embodied World Models

Yu Shang^{a,1}, Zhuohang Li^{a,1}, Yiding Ma^{a,1}, Weikang Su^{a,1}, Xin Jin², Ziyu Wang^{a,2}, Lei Jin^a, Xin Zhang^a, Yinzhou Tang^a, Haisheng Su^b, Chen Gao^a, Wei Wu^a, Xihui Liu^c, Dhruv Shah^d, Zhaoxiang Zhang^e, Zhibo Chen^f, Jun Zhu^a, Yonghong Tian^g, Tat-Seng Chua^h, Xin Wang^a, Wenwu Zhu^a and Yong Li^{a,*}

^a*Tsinghua University, Beijing, China*

^b*Shanghai Jiao Tong University, Shanghai, China*

^c*The University of Hong Kong, Hong Kong SAR, China*

^d*Princeton University, Princeton, NJ, USA*

^e*Chinese Academy of Sciences, Beijing, China*

^f*University of Science and Technology of China, Hefei, China*

^g*Peking University, Beijing, China*

^h*National University of Singapore, Singapore, Singapore*

ARTICLE INFO

Keywords:

Embodied world model

Benchmark

Embodied Intelligence

ABSTRACT

While world models have emerged as a cornerstone of embodied intelligence by enabling agents to reason about environmental dynamics through action-conditioned prediction, their evaluation remains fragmented. Current evaluation of embodied world models has largely focused on perceptual fidelity (e.g., video generation quality), overlooking the functional utility of these models in downstream decision-making tasks. In this work, we introduce WorldArena, a unified benchmark designed to systematically evaluate embodied world models across both perceptual and functional dimensions. WorldArena assesses models through three dimensions: video perception quality, measured with 16 metrics across six sub-dimensions; embodied task functionality, which evaluates world models as data engines, policy evaluators, and action planners integrating with subjective human evaluation. Furthermore, we propose EWMScore, a holistic metric integrating multi-dimensional performance into a single interpretable index. Through extensive experiments on 14 representative models, we reveal a significant perception–functionality gap, showing that high visual quality does not necessarily translate into strong embodied task capability. WorldArena benchmark with the public leaderboard is released at <https://world-arena.ai/>, providing a framework for tracking progress toward truly functional world models in embodied AI.

1. Introduction

In recent years, world models (Ding et al., 2025; Kong et al., 2025; Zhu et al., 2024) have emerged as a foundational component of embodied intelligence. A world model (WM) learns to predict future environment states conditioned on current observations and actions, enabling agents to reason about dynamics and interaction outcomes. Embodied World Model (EWM) forecasts future states based on robot actions and external instructions, effectively functioning as a mental simulator guiding robot action planning and decision-making, or an environment proxy to support scalable robotic training and evaluation (Shang et al., 2025a; Long et al., 2025). Unlike general-purpose video generation models, EWMs must capture not only perceptual fidelity but also physically grounded, action-consistent dynamics that are critical for downstream embodied tasks.

However, existing evaluation protocols suffer from significant limitations. First, they lack comprehensive evaluations oriented toward embodied tasks, including the role of world models as environment proxies and embodied agents. Current benchmarks (Yue et al., 2025; Li et al., 2025a; Lu et al., 2025) mainly focus on video-level quality

*Corresponding author.

✉ liyong07@tsinghua.edu.cn (Y. Li)

ORCID(s): 0000-0001-5617-1659 (Y. Li)

¹Equal contribution.

²Equal contribution.




Figure 1: EWMScore results (a) and performance comparisons across different evaluation dimensions (b) for representative embodied world models on RoboTwin 2.0 data.

metrics, which fail to reflect the real-world value of embodied world models for practical embodied applications, as shown in the comparison of Table 1. Although some recent studies (Qin et al., 2024; Zhang et al., 2025; Fan et al., 2026) evaluate embodied world models through closed-loop action execution, broader embodied capabilities such as their roles as synthetic data engines or tools for policy evaluation remain largely unassessed. Second, the coverage of evaluated models is insufficient. Most existing benchmarks (Yue et al., 2025; Fan et al., 2026) focus on general text-conditioned video generation models (Wan et al., 2025; Yang et al., 2024b), while many recent robot-specialized world models (Chi et al., 2025; Team et al., 2025; Liao et al., 2025; Zhen et al., 2025; Guo et al., 2025), have received little systematic attention and remain largely unevaluated.

To bridge this gap, we present **WorldArena**, the first embodied world model benchmark that integrates perceptual and three functional evaluations, combining both objective and subjective assessments. WorldArena provides a holistic evaluation framework across three complementary aspects: (1) multi-faceted video quality, comprising 16 numerical metrics across 6 key sub-dimensions, including visual quality, motion quality, content consistency, physics adherence, 3D accuracy, and controllability; (2) embodied task utility, which evaluates model performance in data synthesis, policy evaluation, and action planning; and (3) human evaluation, which complements automated metrics by capturing qualitative aspects of model behavior that are difficult to quantify, such as physical plausibility and instruction adherence. Additionally, we introduce EWMScore, a unified metric that combines multi-dimensional metrics into a single index, offering a comprehensive assessment of embodied world models' generative performance. An overview of the evaluation result is shown in Figure 1.

For the evaluation data, we use robotic manipulation as representative embodied scenarios and evaluate on two datasets: RobotTwin 2.0 (Chen et al., 2025) and LIBERO (Liu et al., 2023a). Across these settings, we conduct a unified evaluation of 14 representative world models, including both general video generation models and specialized embodied world models. The results reveal a substantial gap between visual fidelity and embodied task performance, indicating that visual quality alone remains insufficient for effectively supporting embodied tasks. Overall, our contributions can be summarized as follows:

- We introduce the first comprehensive benchmark tailored for embodied world models, enabling a unified evaluation of their perceptual and functional capabilities.
- We propose EWMScore, a unified objective metric for embodied world models, and conduct extensive human studies to validate its effectiveness. Results demonstrate that EWMScore highly aligns with subjective judgment, serving as a reliable and interpretable index.
- We conduct a systematic evaluation of 14 representative embodied world models and provide a multi-dimensional analysis of their strengths and limitations, offering insights and guidance for future research.

2. Related Works

2.1. Embodied World Models

Embodied world models are generative models that predict future observations of physical scenes involving robot locomotion and manipulation. These models can be broadly categorized into three types: video generation-based

Table 1

Comparison of existing world model benchmarks and WorldArena across three dimensions.

Benchmark	Video Quality						Embodied Tasks			Human
	Visual Quality	Motion Quality	Content Consist.	Physics Adher.	Control ability	3D Acc.	Data Engine	Policy Eval.	Action Planner	
WorldModelBench (Li et al., 2025a)	×	×	×	✓	✓	×	×	×	×	✓
WorldSimBench (Qin et al., 2024)	✓	✓	✓	×	✓	×	×	×	✓	✓
WorldScore (Duan et al., 2025)	✓	✓	✓	×	✓	✓	×	×	×	✓
4DWorldBench (Lu et al., 2025)	✓	✓	✓	✓	✓	✓	×	×	×	✓
EWMBench (Yue et al., 2025)	×	✓	✓	×	✓	×	×	×	×	✓
WorldEval (Li et al., 2025b)	×	×	×	×	×	×	×	✓	×	×
World-in-World (Zhang et al., 2025)	✓	✓	×	×	✓	×	×	×	✓	×
WoW-World-Eval (Fan et al., 2026)	✓	✓	✓	✓	✓	×	×	×	✓	✓
WorldArena (Ours)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

models (Liao et al., 2025; Shang et al., 2025b; Team et al., 2025), 3D reconstruction-based models (Huang et al., 2025; Qian et al., 2025), and latent-space world models (Assran et al., 2025; Liu and Chen, 2025; Hafner et al., 2025). In practice, embodied world models serve three key roles: (1) as data synthesis engines (Jang et al., 2025) for generating video-action sequences to augment robot policy training; (2) as policy evaluation environments (Shang et al., 2025b; Li et al., 2025b) for scalable virtual testing through policy-world model interaction; and (3) as action planners (Hu et al., 2024; Fan et al., 2026), where predicted states are decoded into executable actions for robot control. Given the diversity of model paradigms and functional roles, embodied world models are inherently challenging to evaluate comprehensively, underscoring the need for a holistic benchmark that systematically assesses them across both perceptual and functional dimensions, driving their future development.

2.2. World Model Benchmarks

Existing world model benchmarks can be broadly divided into general-purpose and embodied benchmarks. General-purpose benchmarks (Duan et al., 2025; Li et al., 2025a; Lu et al., 2025) mainly assess perceptual and generative quality, including visual fidelity, motion realism, content consistency, and sometimes physical or geometric plausibility. However, they largely treat world models as video generators and overlook their functional roles in decision-making and interaction. Recent embodied benchmarks (Yue et al., 2025; Li et al., 2025b; Zhang et al., 2025; Fan et al., 2026) extend evaluation to controllability, action conditioning, and limited closed-loop interaction. Yet they often focus on a single embodied role, primarily target text-conditioned video models, and provide limited coverage of action-conditioned and robot-centric world models. Most also evaluate fewer than ten models, limiting their comprehensiveness. In contrast, WorldArena offers a unified benchmark that systematically evaluates embodied world models across perceptual and functional dimensions, combining objective metrics with human assessment.

3. The WorldArena Benchmark

WorldArena consists of three key components. First, we assess video quality from 6 dimensions with 16 metrics, focusing on the world model’s open-loop prediction ability (Section 3.1). Second, we evaluate closed-loop performance across 3 typical embodied downstream tasks (Section 3.2). Third, to complement objective measurements with subjective judgment, we collect human annotations to assess qualitative aspects of model performance (Section 3.3). Finally, we integrate multi-dimensional metrics into an interpretable index EWMScore to reflect overall performance (Section 3.4).

3.1. Video Quality Evaluation

We begin by evaluating the quality of the videos generated by different embodied world models, considering 16 video metrics across six sub-dimensions, as shown in Figure 1 (b). The detailed metric explanations can be found in Appendix B and the case visualization is shown in Appendix J.




Figure 2: Illustrations of the video quality evaluations across six dimensions: visual quality, motion quality, content consistency, physics adherence, 3D accuracy, and controllability.

3.1.1. Visual Quality

Visual quality assesses whether generated videos are perceptually reliable for embodied scenarios, considering low-level fidelity, perceptual appeal, and similarity to real data. We evaluate it using three metrics:

Image Quality measures clarity and sharpness of frames using the MUSIQ (Ke et al., 2021) model, which detects distortions such as overexposure, noise, and compression artifacts (Huang et al., 2024). Higher scores indicate cleaner and more coherent images.

Aesthetic Quality evaluates the visual appeal of the video, considering lighting and color composition. Using the LAION aesthetic predictor (LAION-AI, 2022), we map frames to an aesthetic feature space and derive an average score (Huang et al., 2024), capturing both perceptual consistency and artistic quality.

JEPA Similarity quantifies similarity between feature distributions extracted by the pretrained V-JEPA encoder (Bardes et al., 2023), using maximum mean discrepancy (MMD) with a second-order polynomial kernel (Luo et al., 2024). Higher values indicate greater similarity to the ground-truth video.

3.1.2. Motion Quality

Motion quality reflects whether a model captures physically meaningful and temporally coherent dynamics. We assess both the strength of motion and its temporal continuity. To this end, we introduce the following three metrics:

Dynamic Degree measures motion intensity in the video. We use RAFT (Teed and Deng, 2020) to estimate optical flow between consecutive frames and compute the score over the top 5% most active pixels (Huang et al., 2024).

Flow Score measures the overall intensity of motion across the video by averaging optical flow magnitudes over time (Liu et al., 2023b). This score reflects the degree of dynamic interaction, where a higher value indicates greater motion intensity and more physically meaningful dynamics.

Motion Smoothness measures temporal coherence and physical consistency across frames. Following (Duan et al., 2025), we use a frame interpolation model (Zhang et al., 2024) to predict intermediate frames and compare them with real frames.

3.1.3. Content Consistency

Content consistency measures the stability of objects and scenes throughout the video, evaluated at both semantic and appearance levels using three metrics:

Subject Consistency assesses object consistency across frames by calculating cosine similarity between DINO (Carion et al., 2021) features from the first, current, and previous frames (Huang et al., 2024). Higher similarity scores indicate better consistency.

Background Consistency evaluates scene stability using CLIP (Radford et al., 2021) features, measuring cosine similarity between the current frame and the first and previous frames to assess scene stability (Huang et al., 2024).

Photometric Consistency measures pixel-level texture stability by calculating the average end-point error (AEPE) using optical flow (Duan et al., 2025). A higher AEPE indicates poorer alignment, while a higher score reflects better consistency.

3.1.4. Physics Adherence

Physics adherence evaluates whether generated behaviors conform to real-world physical constraints rather than merely appearing visually plausible. We therefore assess both local interaction realism and global motion correctness with the following two metrics:

Interaction Quality measures the physical plausibility of robot-object interactions, including contact behavior and force transmission. We use Qwen3-VL (Bai et al., 2025a) to assign a 1–5 score, normalized to $[0, 1]$, reflecting consistency with expected physical behavior. The reliability of VLM judgment is analyzed in Appendix F.

Trajectory Accuracy measures the accuracy of the robotic arm’s grasping trajectory. We use SAM 3 (Carion et al., 2025) to extract arm bounding boxes in each frame and compute the normalized dynamic time warping (NDTW) distance to the ground-truth trajectory (Yue et al., 2025). Higher scores indicate better spatiotemporal trajectory alignment.

3.1.5. 3D Accuracy

3D accuracy assesses whether generated videos preserve real-world spatial structure beyond image appearance. We evaluate geometric consistency and perspective plausibility with the following two metrics:

Depth Accuracy measures geometric consistency by comparing depth maps between generated and ground-truth videos. We use monocular depth estimation with median-based scaling to resolve scale ambiguity. Higher scores indicate better preservation of real-world spatial geometry.

Perspectivity measures 3D plausibility, including depth-dependent scale variation, lighting consistency, and occlusion relationships. We use Qwen3-VL to judge whether the video follows realistic 3D geometry, with higher scores indicating better perspective alignment with real-world scenes.

3.1.6. Controllability

Controllability measures the model’s ability to respond to external instructions. We evaluate whether generated videos align with intended actions and instructions using three metrics:

Instruction Following assesses the model’s accuracy in following instructions regarding action type, target object, and task state, measured by VLM and scores normalized to $[0, 1]$.

Semantic Alignment measures how well the generated video matches the semantic meaning of the instruction by computing cosine similarity between Qwen2.5-VL-generated (Bai et al., 2025b) descriptions of the generated and reference videos.

Action Response Sensitivity measures video diversity under different instructions. Given the same initial frame, we generate three distinct instructions and produce the corresponding videos. The score is computed as the average pairwise feature dissimilarity, with higher values indicating stronger instruction-sensitive diversity.

3.2. Embodied Task Evaluation

In this section, we evaluate the capabilities of world models through three embodied tasks, as illustrated in Figure 3.

Embodied Data Engine. World models can synthesize future observations from initial frames and instructions, providing additional training data for downstream embodied policies. We evaluate this capability by measuring the performance gain brought by world model-generated data. We adopt a two-stage pipeline. First, we fine-tune each world model and generate synthetic videos conditioned on the first frame and task instruction. Second, with the world model frozen, we use an inverse dynamics model (IDM) to recover actions from generated videos. Following VPP (Hu et al., 2024), the IDM combines a VideoFormer for fusing video observations and text instructions with a DiT-based diffusion policy for action denoising. It is trained on 5,000 video chunks per task and predicts actions in the Aloha-AgileX dual-arm 14D joint-position space for RoboTwin 2.0 and the 7D end-effector space for LIBERO. We then




Figure 3: Overview of embodied task evaluation, covering world models as data engines, policy evaluators, and action planners. We measure their performance by downstream policy success rate, correlation with real-world evaluation, and success rate of world model-based policies, respectively.

train a baseline $\pi_{0.5}$ (Intelligence et al., 2025) policy with different amounts of synthetic data. The resulting policy improvement reflects the world model’s effectiveness as an embodied data engine.

Embodied Policy Evaluator. In this section, we assess the capability of world models as environment proxies for evaluating policy performance. We train a series of policy models ($\pi_{0.5}$) with varying capabilities. These models are evaluated by interacting with an action-controllable world model, generating videos through a rollout process that continues until it exceeds 20% more frames than the corresponding ground truth video. Task success is evaluated using a VLM, which determines whether the embodied task was executed successfully. The used prompt for the VLM is shown in Appendix C. The success rate from the world model’s evaluation is compared to that from the result from simulator. A high correlation between the two suggests effective simulation of real-world dynamics, while a low correlation indicates a mismatch in environmental transition simulation.

Embodied Action Planner. By predicting future state transitions, world models can function as the action-planning “brain” of an embodied agent. In this part, we investigate the ability of world models to execute embodied tasks in a closed-loop manner. Similar to the data synthesis engine setup, we pair the world model with an inverse dynamics model, where the world model takes textual instructions and the initial frame as input and outputs the corresponding action sequence for future operations. This sequence is then executed in the RoboTwin simulator, and the task success rate is measured to evaluate the world model’s performance in closed-loop action execution.

3.3. Human Evaluation

Since video quality metrics alone cannot fully capture aspects like physical plausibility and instruction adherence, we incorporate two types of human evaluations. The first type involves scoring three key dimensions: overall video quality, instruction following, and physical adherence on a 1 to 5 scale, then normalizing to a 0-100 range. The second type is a head-to-head comparison, where annotators choose the superior video generated by two different models from the same prompt, yielding a win-rate metric. We recruited 70 annotators who evaluated a total of 3500 videos.

3.4. EWMScore Metric

We define three EWMScore variants to evaluate embodied world models from complementary perspectives: **EWMScore-P** for perceptual generation quality, **EWMScore-F** for downstream functional utility, and **EWMScore** for combined performance. This separation is needed because functional evaluation is not uniformly applicable

Table 2

Video quality evaluation results across visual quality, motion quality and content consistency dimensions.

Models	Visual Quality			Motion Quality			Content Consistency		
	Image Quality	Aesthetic Quality	JEPA Similarity	Dynamic Degree	Flow Score	Motion Smoothness	Subject Consist.	Background Consist.	Photometric Consist.
GigaWorld-0	0.5041	0.3991	0.4413	0.6709	0.3118	0.7811	0.7303	0.8563	0.1756
Genie Envioner	0.2305	0.3289	0.3340	0.6930	0.0855	0.6966	0.7760	0.9024	0.2006
TesserAct	0.3322	0.4590	0.4579	0.5150	0.2447	0.7579	0.8250	0.9238	0.2491
RoboMaster	0.3487	0.3842	0.2966	0.6124	0.1484	0.6940	0.8295	0.9123	0.3356
Vidar	0.4145	0.4068	0.5608	0.2767	0.1426	0.7973	0.7629	0.8300	0.2350
Cosmos-Predict 2.5 (text)	0.6668	0.4501	0.3126	0.5911	0.4302	0.7882	0.7488	0.8511	0.1383
Cosmos-Predict 2.5 (action)	0.4489	0.3576	0.9296	0.3994	0.0573	0.7100	0.8197	0.8894	0.3528
WoW	0.4587	0.3868	0.7440	0.4608	0.2706	0.7692	0.8161	0.9025	0.2170
CtrlWorld	0.3522	0.3893	0.9185	0.4257	0.3449	0.7377	0.8411	0.9057	0.1729
Wan 2.2	0.3884	0.3963	0.7575	0.4349	0.1269	0.7019	0.8388	0.9042	0.4776
CogvideoX	0.3582	0.3777	0.9384	0.3166	0.2189	0.7391	0.8083	0.8773	0.3580
IRASim	0.3489	0.3623	0.9330	0.4139	0.2083	0.7052	0.8312	0.9068	0.3522
Veo 3.1	0.6605	0.4632	0.5694	0.5450	0.1396	0.6989	0.7878	0.8710	0.3247
Wan 2.6	0.6824	0.4433	0.7229	0.7421	0.4532	0.8539	0.7517	0.8687	0.1904

Table 3

Video quality evaluation results across physics adherence, 3D accuracy and controllability dimensions.

Models	Physics Adherence		3D Accuracy		Controllability		
	Interaction Quality	Trajectory Acc.	Depth Acc.	Perspectivity	Instruction Following	Semantic Alignment	Action Response Sensitivity
GigaWorld-0	0.5368	0.1552	0.6316	0.7596	0.6156	0.8591	0.1134
Genie Envioner	0.2052	0.0679	0.8663	0.5284	0.2028	0.8544	0.0109
TesserAct	0.5800	0.1396	0.7159	0.7920	0.6152	0.8783	0.0311
RoboMaster	0.5364	0.1158	0.8335	0.7588	0.5772	0.8761	0.0352
Vidar	0.5348	0.1928	0.7872	0.7592	0.5912	0.8826	0.0819
Cosmos-Predict 2.5 (text)	0.3872	0.0816	0.7051	0.7964	0.2664	0.7733	0.1418
Cosmos-Predict 2.5(action)	0.5500	0.2945	0.8862	0.7644	0.5840	0.8879	0.0133
WoW	0.5564	0.2058	0.7283	0.7672	0.5692	0.8842	0.0434
CtrlWorld	0.6212	0.4766	0.9300	0.7960	0.7272	0.8912	0.0210
Wan 2.2	0.5184	0.1627	0.7768	0.7660	0.5376	0.8877	0.0512
CogvideoX	0.5940	0.3526	0.9097	0.7828	0.7268	0.8977	0.0076
IRASim	0.5656	0.3639	0.9312	0.7788	0.6604	0.8849	0.0526
Veo 3.1	0.7872	0.1231	0.7421	0.8276	0.9328	0.8607	0.0852
Wan 2.6	0.7280	0.1182	0.7144	0.8032	0.8536	0.8728	0.0992

across models: closed-source models may not support downstream tasks, while text-controlled and action-conditioned models correspond to different evaluation settings. EWMScore-P is computed from video-quality metrics spanning six perceptual dimensions. Each metric is linearly normalized using empirically defined boundaries, rescaled to [0, 100], and then averaged to obtain the final perceptual score. We use uniform weights for simplicity and comparability. EWMScore-F summarizes functional performance over applicable downstream tasks, including data-engine evaluation, policy evaluation, and action planning. For each model, we compute the score only over the tasks supported by its control interface, ensuring fair comparison within each evaluation setting. EWMScore aggregates EWMScore-P and EWMScore-F when both are available.

4. Experiments

4.1. Experimental Setup

Dataset. We evaluate robotic manipulation on RoboTwin 2.0 (Chen et al., 2025) and LIBERO (Liu et al., 2023a). Details of the dataset are shown in Appendix D.

Tested Models. We evaluate 14 representative world models, covering both general-purpose video world models and embodied-specific models. Detailed model information is shown in Appendix E.




Figure 4: Correlation of policy evaluation results from world models and the physical simulator.

Table 4

Task success rate of world models for downstream policy training and action planning in RoboTwin and LIBERO simulators.

(a) Data engine performance.				(b) Action planning performance.			
Model	RoboTwin		LIBERO	Model	RoboTwin		LIBERO
	Task 1	Task 2	Task 1		Task 1	Task 2	Task 1
GigaWorld-0	2%	13%	0%	GigaWorld-0	6%	19%	0%
Genie Envisioner	7%	21%	2%	Genie Envisioner	10%	20%	6%
Tesseract	1%	35%	34%	Tesseract	1%	35%	38%
Vidar	13%	53%	22%	Vidar	2%	19%	14%
WoW	45%	71%	16%	WoW	20%	21%	96%
Wan 2.2	15%	41%	10%	Wan 2.2	12%	20%	24%
CogVideoX	3%	28%	0%	CogVideoX	8%	16%	2%

4.2. Results

4.2.1. Video Quality Evaluation

Tables 2 and 3 summarize video-quality results across six dimensions on RoboTwin 2.0. Overall, embodied world models perform better on interaction-related metrics, while general-purpose video models excel mainly in perceptual quality. Among embodied models, CtrlWorld and Tesseract achieve strong subject consistency, background stability, and trajectory accuracy, indicating better alignment with manipulation dynamics. WoW shows strong action-following ability, and RoboMaster and Vidar maintain balanced motion smoothness and content consistency. CogVideoX performs well in visual quality and content consistency but lags in physics adherence and motion quality. Closed-source models, including Veo 3.1 and Wan 2.6, obtain the highest visual and aesthetic scores, yet bring limited gains on embodied-specific metrics. Qualitative results further show that visually strong models often suffer from semantic drift, whereas embodied world models generate more coherent and goal-consistent videos. Detailed results on LIBERO are shown in Appendix G.

4.2.2. Embodied Task Evaluation

In this section, we evaluate the capabilities of world models through three embodied tasks.

Embodied Data Engine. We evaluate 7 representative world models as data synthesis engines by measuring their effect on downstream policy learning. On RoboTwin 2.0, we use two manipulation tasks, *adjust bottle* (Task 1) and *click bell* (Task 2), each evaluated over 100 trials. For each task, we train a $\pi_{0.5}$ policy with 25 synthetic trajectories generated by each world model. As shown in Table 4, most synthetic data improves policy performance but still underperforms real data, except WoW on *click bell*. This suggests that current embodied world models provide useful but not yet reliable training data. To test whether the gap is mainly due to data scarcity, we increase the synthetic set to 100 trajectories per task for representative models. The larger scale brings moderate gains, while preserving the relative ranking across world models, indicating that limited data alone does not explain the performance gap. This scaling ablation is detailed in Appendix H.




Figure 5: Correlation between EWMScore-P with human evaluation and embodied task performance results on Robotwin Dataset.

Embodied Policy Evaluator. We evaluate whether world models can serve as proxy environments for policy evaluation. We train five $\pi_{0.5}$ policies with different performance levels and let each policy interact with an action-controllable world model, which generates rollouts conditioned on policy actions. As shown in Figure 4, CtrlWorld (Guo et al., 2025) and Roboscape (Shang et al., 2025b) correlate strongly with RoboTwin simulator results, suggesting effective transition modeling. In contrast, IRASim (Zhu et al., 2025) and Cosmos-Predict 2.5 (Gu, 2025) show weaker correlation, indicating less accurate dynamics modeling. All evaluated world models yield higher success rates than the simulator, suggesting systematic overestimation. To analyze this gap, we compare VLM-based and human success judgments and find strong agreement ($r = 0.91$), indicating that the VLM judge is not the main source of bias. Instead, the overestimation mainly stems from world models generating visually plausible rollouts even when the underlying actions are incorrect, thereby masking failure modes.

Embodied Action Planner. Following the data engine setting, we evaluate 7 representative world models as end-to-end action planners by executing their predicted action sequences in the RoboTwin 2.0 and LIBERO simulators. As shown in Table 4, several models achieve non-trivial success rates, but their overall performance remains far below VLA policies such as $\pi_{0.5}$. This suggests that current world models still struggle with reliable closed-loop control. To ensure that the IDM is not the main bottleneck, we conduct a video replay study on 100 real trajectories of *adjust bottle*. When real videos are encoded with Wan2.2-VAE and passed to the IDM, the predicted actions achieve an 84% simulator success rate, indicating that the planner evaluation primarily reflects world model quality.

4.2.3. Human Evaluation

As shown in Figure 1(b), commercial and large-scale general video models, such as Veo 3.1 and Wan 2.6, achieve the highest human scores in overall quality, instruction following, and physical adherence, reflecting strong perceptual realism and semantic alignment. Among embodied world models, action-conditioned methods such as CtrlWorld obtain better physical adherence and higher win rates than text-only models, highlighting the importance of explicit action modeling for plausible interactions. In contrast, earlier text-conditioned models, such as Genie Envioner, score lower across dimensions, revealing remaining gaps in long-horizon coherence and instruction compliance.

4.3. Inter-metric Analysis

Figure 5 analyzes the relationship between EWMScore, human evaluation, and embodied task performance. EWMScore strongly correlates with human judgments (Pearson $r = 0.825$), showing good alignment with subjective perceptual assessment. However, it correlates weakly with data synthesis and action planning performance. This suggests that perceptual realism improves human preference but does not necessarily yield proportional gains in downstream embodied tasks. In particular, the weak correlation with action planning indicates that visually plausible generations may still lack the predictive and decision-relevant signals needed for complex embodied reasoning.

5. Conclusion and Future Work

In this work, we present WorldArena, a unified benchmark for systematically evaluating embodied world models from both perceptual and functional perspectives, integrating multi-dimensional video quality metrics, embodied task

evaluations, and human assessments. Through an extensive evaluation of 14 representative models, we reveal consistent gaps between perceptual quality and embodied task performance, highlighting that strong visual generation alone is insufficient for reliable embodied decision-making. We further demonstrate that EWMScore effectively captures overall generative capability and correlates well with human judgments. In the future, we will continue to expand WorldArena, incorporating more models to support the advancement of perceptually strong and functionally reliable embodied world models.

A. Limitations

In the WorldArena benchmark, the current diversity of tasks and robot morphologies is restricted and requires further expansion. We are actively and continuously expanding the dataset and evaluation scope. To improve real-world grounding, our future work will introduce real-robot evaluation environments guided by both quantitative metrics and actual performance. Furthermore, as with any standardized evaluation, there is an inherent risk of leaderboard overfitting to chosen metrics.

B. Additional Details on Metrics

B.1. Image Quality

The per-frame sharpness of a generated video constitutes the foundation of its visual presentation. Unlike traditional reference-based metrics such as PSNR, which rely on ground-truth images, we adopt the **MUSIQ** (Multi-scale Image Quality Transformer) model (Ke et al., 2021) to evaluate technical distortions in a no-reference setting, including overexposure, sensor noise, and compression artifacts. MUSIQ leverages a multi-scale Transformer architecture to capture the relationships between local details and global composition.

For a video sequence $V = \{I_1, I_2, \dots, I_T\}$, where V represents a specific video and I_i represents the i th frame of the video V , the image quality score S_{img} is defined as:

$$S_{\text{img}} = \frac{1}{T} \sum_{t=1}^T \Phi_{\text{musiq}}(I_t) \quad (1)$$

where $\Phi_{\text{musiq}}(\cdot)$ denotes the pretrained quality prediction function. A higher value of S_{img} indicates greater visual purity and clarity at the level of digital image (Huang et al., 2024).

B.2. Aesthetic Quality

Beyond technical fidelity, generated videos are also required to conform to human aesthetic principles, such as harmonious lighting and visually pleasing color composition. We employ the **LAION Aesthetic Predictor** (LAION-AI, 2022) to perform aesthetic feature mapping for each frame. Similarly, for a video sequence $V = \{I_1, I_2, \dots, I_T\}$, the aesthetic quality score S_{aes} is defined as:

$$S_{\text{aes}} = \frac{1}{T} \sum_{t=1}^T \Psi_{\text{aes}}(I_t) \quad (2)$$

where $\Psi_{\text{aes}}(\cdot)$ maps each image into a high-dimensional feature space and predicts an aesthetic score. This formulation ensures that the evaluation extends beyond pixel-level sharpness to encompass perceptual coherence and artistic consistency (Huang et al., 2024).

B.3. JEPA Similarity

To evaluate video quality from a global feature-distribution perspective and detect high-level spatiotemporal collapse, we introduce the **JEPA Similarity**. Unlike traditional FVD metric, which relies on Gaussian assumptions, JEPA measures the maximum mean discrepancy (MMD) between feature distributions to provide evaluation results that better align with human perception:

$$S_{\text{JEPA}} = \exp \left(-\alpha \cdot \widehat{\text{MMD}}_{\text{poly}}^2(\mathcal{F}_{\text{gen}}, \mathcal{F}_{\text{ref}}) \right) \quad (3)$$

where $\alpha = 40$ is a scaling factor that enhances numerical distinguishability, $\mathcal{F}_{\text{gen}}$ and $\mathcal{F}_{\text{ref}}$ denote the feature space distributions of the generated video set and the reference expert demonstration (GT) set, respectively, extracted by a pretrained **V-JEPA** encoder (Bardes et al., 2023) which is pre-trained via masked prediction tasks and enables the model to capture high-level spatio-temporal causality and physical logic in videos, offering greater robustness to temporal warping and content variations. $\widehat{\text{MMD}}_{\text{poly}}^2$ represents the squared estimator of the Maximum Mean Discrepancy using a second-order polynomial kernel, defined as $k(\mathbf{x}, \mathbf{y}) = (\gamma \langle \mathbf{x}, \mathbf{y} \rangle + c_0)^2$, with $\gamma = 1, c_0 = 0$. It measures the distance between the two feature sets in the reproducing kernel Hilbert space (RKHS), computed as follows:

$$\widehat{\text{MMD}}_{\text{poly}}^2(\mathcal{F}_{\text{gen}}, \mathcal{F}_{\text{ref}}) = \frac{1}{m(m-1)} \sum_{i \neq j}^m k(\mathbf{f}_i^{\text{gen}}, \mathbf{f}_j^{\text{gen}}) + \frac{1}{n(n-1)} \sum_{i \neq j}^n k(\mathbf{f}_i^{\text{ref}}, \mathbf{f}_j^{\text{ref}}) - \frac{2}{mn} \sum_{i=1}^m \sum_{j=1}^n k(\mathbf{f}_i^{\text{gen}}, \mathbf{f}_j^{\text{ref}}) \quad (4)$$

where m and n denote the number of samples in the generated videos and the reference videos, $\mathbf{f}_i^{\text{gen}}$ and $\mathbf{f}_i^{\text{ref}}$ are the corresponding V-JEPA feature vectors. Higher values indicate closer alignment to reference demonstrations (Luo et al., 2024).

This metric is not only sensitive to breakdowns in high-level spatio-temporal logic but also avoids the Gaussian distribution assumption, offering significantly better sample efficiency than conventional metrics. Moreover, it achieves an enormous improvement in correlation with human subjective assessments, especially when evaluating complex embodied operation logic, thereby more reliably reflecting the physical plausibility and spatio-temporal consistency of the generated videos.

B.4. Dynamic Degree

We employ the **RAFT** (Recurrent All-Pairs Field Transforms) (Teed and Deng, 2020) optical flow model to extract motion vector fields between adjacent frames. To accurately capture the most representative and salient motions in a video, such as robotic arm grasping, we focus on pixels whose optical flow magnitudes fall within the top 5% (Huang et al., 2024).

Let $\mathbf{u}_{t,t+1}$ denote the two-dimensional optical flow field between consecutive frames. We define the average magnitude of the active pixel set as $\bar{v}_{\text{top5}}$. To obtain a smooth numerical mapping while introducing resolution adaptivity, we define the dynamic degree score as:

$$S_{\text{dyn}} = \frac{1}{1 + \exp\left(-\alpha \cdot \left(\frac{\bar{v}_{\text{top5}}}{\tau} - 1\right)\right)} \quad (5)$$

where $\tau = \frac{6}{256} \times \min(H, W)$ is a resolution-adaptive threshold constant, α controls the steepness of the mapping curve, and $S_{\text{dyn}} \in (0, 1)$. Values closer to 1 indicate more pronounced dynamic responses in the video.

B.5. Flow Score

To quantify overall physical motion intensity and dynamic activity, we compute an optical-flow-based motion score. Given a generated video $V = \{I_1, \dots, I_T\}$ with frame width W and height H , we use **RAFT** (Teed and Deng, 2020) to estimate dense optical flow fields $\mathbf{u}_t \in \mathbb{R}^{H \times W \times 2}$ between consecutive frames I_t and I_{t+1} . By averaging the magnitude of optical flow across all pixels, the average motion intensity is defined as:

$$S_{\text{flow_raw}} = \frac{1}{T-1} \sum_{t=1}^{T-1} \left(\frac{1}{H \cdot W} \sum_{i,j} \|\mathbf{u}_t(i, j)\|_2 \right) \quad (6)$$

where (i, j) indexes pixel locations, $\|\cdot\|_2$ denotes the Euclidean norm (L_2 norm), which quantifies the magnitude of pixel displacement per unit time (Liu et al., 2023b). In the context of embodied intelligence tasks, this metric serves a dual evaluative purpose: on the one hand, it effectively identifies whether a video degenerates into "static frames" or exhibits only "minimal drift" due to insufficient generative capability of the model; on the other hand, it captures whether unnatural, non-physical distortions are present in the overall scene.

Higher values of $S_{\text{flow_raw}}$ typically indicate more pronounced dynamic interaction and physically meaningful motion. Compared to **Dynamic Degree**, this metric focuses on assessing overall motion intensity and detecting

implausible global dynamics. To ensure consistent interpretation and comparability with other metrics, we will normalize $S_{\text{flow_raw}}$ to the range $[0, 1]$ in Section B.17, denoting the normalized value as S_{flow} , while preserving the property that higher values correspond to better performance.

B.6. Motion Smoothness

To evaluate whether motion is temporally coherent and consistent with physical inertia, we adopt a reconstruction-based strategy using a video frame interpolation model (**VFI-Mamba**) (Zhang et al., 2024). Given all frames of a video, we take the odd-indexed frames $\{I_1, I_3, \dots\}$ as inputs and predict the corresponding intermediate frames I_{mid} , which are then compared with the ground-truth (GT) frames. If the motion is physically plausible and smooth, the intermediate frame I_{mid} should be accurately reconstructed from its surrounding frames ($I_{\text{prev}}, I_{\text{next}}$) via nonlinear interpolation.

The key innovation lies in incorporating motion magnitude as a weighting factor to avoid overestimating static backgrounds. The final motion smoothness score is defined as:

$$S_{\text{smooth_raw}} = \frac{1}{N} \sum \text{SSIM}(\hat{I}_{\text{pred}}, I_{\text{mid}}) \cdot \ln(1 + \text{diff}(I_{\text{prev}}, I_{\text{next}})) \quad (7)$$

where N denotes the number of predicted intermediate frames (typically equal to or one less than the number of even-indexed frames), $\hat{I}_{\text{pred}}$ is the interpolated frame predicted by the model, I_{mid} is the real frame between I_{prev} and I_{next} , and $\text{diff}(\cdot)$ represents the mean raw pixel-wise difference between two frames. The logarithmic weighting $\ln(1 + x)$ compensates for the increased difficulty of interpolation under large motion, thereby assigning higher rewards to sequences that maintain high reconstruction fidelity even during rapid motion. For consistency with other evaluation metrics, $S_{\text{smooth_raw}}$ will be normalized to the range $[0, 1]$ in Section B.17, yielding the final smoothness score S_{smooth} , where higher values indicate superior temporal coherence and motion consistency.

B.7. Subject Consistency

For a video sequence $V = \{I_1, I_2, \dots, I_T\}$, we extract frame-level features using **DINO** (Caron et al., 2021), denoted as $f_i = \text{DINO}(I_i)$, which emphasize the spatial topological structure of objects. We compute the cosine similarity between the feature f_i of the current frame and both the first-frame feature f_1 and the previous-frame feature f_{i-1} , and average the similarities across all frames (Huang et al., 2024):

$$S_{\text{subj_raw}} = \sum_{i=2}^T \left(\frac{\cos(f_i, f_1) + \cos(f_i, f_{i-1})}{2} \right) \quad (8)$$

However, a common ‘‘shortcut’’ phenomenon in video generation evaluation is that models may produce nearly static videos to obtain artificially high consistency scores. To faithfully reflect dynamic generation capability in embodied scenarios, we introduce the dynamic degree S_{dyn} defined in Section B.4 as a weighting factor for subject consistency. Specifically, when the video’s dynamic degree falls below a predefined threshold γ , the raw score is penalized as:

$$S_{\text{subj}} = S_{\text{subj_raw}} \cdot \min\left(1, \frac{S_{\text{dyn}}}{\gamma}\right) \quad (9)$$

This mechanism ensures that static or near-static videos cannot achieve high scores even when frame-level similarity is extremely high, leading to more reasonable evaluation in embodied tasks.

B.8. Background Consistency

Analogous to subject consistency, for a video sequence $V = \{I_1, I_2, \dots, I_T\}$, we extract frame-level features using **CLIP** (Radford et al., 2021), denoted as $h_i = \text{CLIP}(I_i)$, which place greater emphasis on global scene semantics and prevent uncontrolled background variation during generation. We compute the cosine similarity between h_i and both h_1 and h_{i-1} , and average the results across all frames (Huang et al., 2024):

$$S_{\text{bg_raw}} = \sum_{i=2}^T \left(\frac{\cos(h_i, h_1) + \cos(h_i, h_{i-1})}{2} \right) \quad (10)$$

Similarly, the final background consistency score is adjusted using the dynamic degree:

$$S_{\text{bg}} = S_{\text{bg_raw}} \cdot \min\left(1, \frac{S_{\text{dyn}}}{\gamma}\right) \quad (11)$$

B.9. Photometric Consistency

Photometric consistency measures the physical stability of textures at the pixel level. For a video $V = \{I_1, I_2, \dots, I_T\}$, we use the forward optical flow field $\mathbf{u}_t$ between frames I_t and I_{t+1} , as well as the backward flow field $\mathbf{u}'_{t+1}$, to warp pixels from frame t to frame $t+1$ and then back to frame t . The average end-point error (AEPE) is defined as (Duan et al., 2025):

$$E_{\text{photo}} = \frac{1}{T} \sum_{t=1}^T \|\text{Warp}_{\text{back}}(\text{Warp}_{\text{fwd}}(I_t, \mathbf{u}_t), \mathbf{u}'_{t+1}) - I_t\|_2 \quad (12)$$

Since this metric quantifies pixel-level reconstruction error, lower values correspond to superior visual quality. To obtain a positively correlated measure that appropriately rewards sequences with meaningful motion while penalizing trivial solutions in static videos, we compute the pre-normalized photometric consistency score as:

$$S_{\text{photo_raw}} = \frac{1}{E_{\text{photo}}} \cdot \min\left(1, \frac{S_{\text{dyn}}}{\gamma}\right) \quad (13)$$

where S_{dyn} denotes the dynamic degree (formally defined in Section B.4), quantifying the overall motion intensity within a video sequence, and γ serves as a dynamic threshold that modulates the penalty for insufficient motion. The inclusion of S_{dyn} addresses a critical limitation of conventional photometric metrics: static or near-static sequences often achieve artificially high scores due to minimal frame-to-frame variations, even though they fail to demonstrate meaningful dynamic modeling.

By scaling the raw reciprocal score with the normalized dynamic degree, our formulation ensures that only videos with sufficient motion ($S_{\text{dyn}} \geq \gamma$) retain their full photometric consistency score, while static sequences are proportionally penalized. This encourages the model to maintain high reconstruction fidelity under actual motion rather than exploiting static scenarios. Subsequently, $S_{\text{photo_raw}}$ is normalized to the interval $[0, 1]$ in Section B.17, which produces the final photometric consistency metric S_{photo} , where higher values denote enhanced visual fidelity and temporal coherence.

B.10. Interaction Quality

This metric evaluates the physical plausibility of interactions between the robotic arm and environmental objects, including contact behavior, force transmission, friction, inertia, and boundary integrity. We employ the pretrained multimodal model **Qwen3-VL-8B** (Bai et al., 2025a) as a VLM-based judge. Given N_{sample} sampled frames and the task instruction, the model assigns a 1–5 Likert score, which is normalized to $[0, 1]$ to yield the final interaction quality score, the prompt used to evaluate interaction quality can be found in B.10.

VLM Evaluation Prompt

You are an expert evaluator for robot interaction videos. You are evaluating videos generated for ****embodied AI manipulation scenarios****, specifically focusing on robotic arms interacting with objects in tabletop environments.

****EVALUATION CONTEXT:****

- Target scenario: Robotic manipulation (e.g., pick-place, push, grasp)
- Expected agent: ****Robotic arm/end-effector****, NOT human hands
- Expected environment: Tabletop with objects, typical for robot manipulation tasks
- Expected physics: Realistic robot-object interactions following physical laws

****CRITICAL EVALUATION PRINCIPLES:****

1. Base ALL judgments ONLY on what is visually observable in the sampled frames
2. DO NOT infer information not shown (no assumptions about unseen parts)
3. Evaluate temporal coherence across the sampled frames
4. For instruction following: Compare STRICTLY against the provided text instruction

****EVALUATION DIMENSIONS & SCORING RUBRICS:****

1. Interaction_Quality (Quality of robot-object interactions)
 - Score 1: Objects pass through robot or other objects; no proper contact
 - Score 2: Contact exists but interaction is unrealistic (e.g., sliding without friction, incorrect force response)
 - Score 3: Mostly plausible interactions with minor issues (e.g., slight penetration, imperfect grasping)

- Score 4: Realistic contact physics (proper friction, force transfer, object deformation)
- Score 5: Perfect interaction physics; indistinguishable from real robot manipulation

2. PERSPECTIVITY (3D consistency and camera geometry)

- Score 1: Scene has no coherent 3D structure; objects float inconsistently
- Score 2: 3D structure is unstable (e.g., scale changes, incorrect occlusion)
- Score 3: Reasonable 3D consistency with minor issues (e.g., slight perspective drift)
- Score 4: Stable camera perspective with consistent depth relationships
- Score 5: Perfect camera geometry and 3D consistency

3. INSTRUCTION FOLLOWING (Adherence to given instruction:**VIDEO INSTRUCTION**)

- ****HALLUCINATION CHECK****: If the video shows human hands instead of robotic arms, score ≤ 2 immediately
- Score 1: Completely different from instruction (wrong action, wrong objects, wrong scene)
- Score 2: Partially related but major errors (e.g., wrong target object, incorrect manipulation type)
- Score 3: Follows general intent but with execution errors (e.g., correct action sequence but imprecise)
- Score 4: Mostly correct with minor deviations (e.g., slight position error, extra unnecessary motion)
- Score 5: Perfect execution of all specified elements (action, object, scene, outcome)

****SPECIFIC ROBOT-RELATED CHECKS:****

- Robotic arm should have mechanical appearance, NOT human limbs
- End-effector (gripper) should maintain consistent form throughout interaction
- Robot motion should show appropriate joint movement and kinematics
- Object manipulation should respect object mass and inertia
- Contact should be maintained appropriately during grasping/lifting

****OUTPUT FORMAT REQUIREMENTS:****

You MUST output a SINGLE, VALID JSON object with EXACTLY three keys:

- 'Interaction_Quality'
- 'Perspectivity'
- 'Instruction_Following'

Each value must be an object with exactly two keys:

- "score": integer 1-5
- "reason": concise explanation citing SPECIFIC visual evidence from frames

****EXAMPLE OUTPUT:****

```
{
  "Interaction_Quality": {
    "score": 2,
    "reason": "Object slides without friction during pushing;
              gripper penetrates object slightly"
  },
  "Perspectivity": {
    "score": 4,
    "reason": "Stable camera perspective with consistent depth ordering"
  },
  "Instruction_Following": {
    "score": 1,
    "reason": "Video shows human hand instead of robotic arm (hallucination)"
  }
}
```

****CRITICAL INSTRUCTIONS:****

1. Output ONLY the JSON object, no other text
2. Base scoring on observed visual evidence only
3. For instruction following: Strictly compare with the provided instruction
4. Consider temporal coherence across all sampled frames
5. Penalize hallucinations (e.g., human hands instead of robot) heavily

Now evaluate the provided video frames based on the above criteria.

B.11. Trajectory Accuracy

In embodied intelligence tasks, the accuracy of the robotic arm’s grasping trajectory is a core indicator of whether the model generates *effective actions*. Trajectories encode not only low-level physical consistency but also high-level task logic and interaction constraints. To quantify this property, we first apply **SAM3** (Segment Anything Model 3) (Carion et al., 2025) to extract bounding boxes of the robotic arm in each frame. After non-maximum suppression(nms) and confidence filtering, we construct the raw trajectory sequences using the centers of candidate boxes.

Let the ground-truth trajectory be $GT = (r_1, r_2, \dots, r_{|R|})$ and the generated trajectory be $P = (p_1, p_2, \dots, p_{|P|})$, where $|R|$ and $|P|$ denote the sequence lengths, respectively. To address missing detections caused by occlusion or tracking interruption, we apply linear interpolation to ensure temporal continuity. For a missing point p_i with $i \notin M$, its position is computed as:

$$p_i = (1 - \alpha)p_{\text{prev}} + \alpha p_{\text{next}}, \quad \alpha = \frac{i - \text{prev}}{\text{next} - \text{prev}} \quad (14)$$

where prev and next denote the nearest valid observation indices before and after i .

We then compute the normalized dynamic time warping distance (NDTW) (Müller, 2007) to evaluate global alignment between the generated trajectory and the ground-truth trajectory:

$$\text{NDTW}(GT, P) = \min_{\pi} \frac{1}{|R|} \sqrt{\sum_{(i,j) \in \pi} \|r_i - p_j\|^2} \quad (15)$$

where π denotes the optimal alignment path. This metric captures both temporal causality and task-stage ordering, enabling discrimination between correct and incorrect execution sequences such as "approach-grasp-move." Since lower NDTW values indicate better alignment, we first derive a pre-normalized trajectory alignment score (Yue et al., 2025):

$$S_{\text{traj_raw}} = \frac{1}{\text{NDTW}(GT, P)} \quad (16)$$

where higher values correspond to more accurate spatial-temporal alignment with the real trajectory and more accurate actions. To ensure consistency with our evaluation framework and facilitate direct comparison with other metrics, we normalize $S_{\text{traj_raw}}$ to the range $[0, 1]$ in Section B.17, yielding the final trajectory alignment score S_{traj} . This normalized metric preserves the property that higher values indicate superior trajectory fidelity and task-stage adherence.

B.12. Depth Accuracy

To evaluate whether the generated video preserves real-world spatial geometry, we compute depth discrepancies between the generated video and the ground-truth reference using the monocular depth estimation model **Depth-Anything** (Yang et al., 2024a). Since monocular depth prediction suffers from scale ambiguity, we adopt a **median-based scaling** strategy (Liang et al., 2025).

The procedure is as follows:

1. **Uniform Sampling:** We uniformly sample $T_{\text{target}} = 40$ frames from both the generated video and the ground-truth video to ensure temporal alignment.
2. **Scale Alignment:** Depth maps D_{gen} and D_{gt} are estimated for the generated and ground-truth frames, respectively. Their medians are computed as $m_{\text{gen}} = \text{median}(D_{\text{gen}})$ and $m_{\text{gt}} = \text{median}(D_{\text{gt}})$. The scaling factor $\alpha = \frac{m_{\text{gt}}}{m_{\text{gen}}}$ is applied to obtain the aligned depth $\hat{D}_{\text{gen}} = D_{\text{gen}} \cdot \alpha$.
3. **AbsRel Error:** Within the valid pixel mask $\mathcal{M}$ (which typically filters out noise and distant regions with ground-truth depth $D_{\text{gt}} < 1e - 3$), the absolute relative error is computed as follows:

$$E_{\text{Depth}} = \frac{1}{|\mathcal{M}|} \sum_{p \in \mathcal{M}} \frac{|\hat{D}_{\text{gen}}(p) - D_{\text{gt}}(p)|}{D_{\text{gt}}(p) + \epsilon} \quad (17)$$

where ϵ is a small constant to prevent division by zero. Lower values indicate stronger depth accuracy with the real-world scene. To align this metric with our evaluation framework where higher scores correspond to better performance, we will normalize E_{Depth} to the range $[0, 1]$ and invert its direction in Section B.17, such that higher values correspond to higher accuracy, resulting in the final normalized depth accuracy score S_{Depth} .

B.13. Perspectivity

This metric evaluates three-dimensional geometric plausibility. The VLM examines perspective cues such as scale variation with depth, lighting consistency, and occlusion relationships during camera motion. We use **Qwen3-VL-8B** as a judge and normalize the Likert-scale output to $[0, 1]$, which is normalized to $[0, 1]$ to yield the final perspectivity score, the prompt used to evaluate perspectivity can be found in B.10.

B.14. Instruction Following

This metric evaluates the semantic consistency between each generated video V_i and its corresponding instruction $Inst_i$, focusing on action type, target object, and final task state. We again use **Qwen3-VL-8B** as a VLM-based judge with a normalized 1–5 Likert scale, which is normalized to $[0, 1]$ to yield the final instruction following score, the prompt used to evaluate instruction following can be found in B.10.

B.15. Semantic Alignment

To assess whether the generated video truly understands and executes the given textual instruction, we evaluate semantic alignment as follows:

$$S_{\text{clip}} = w \cdot \max(\cos(f_{\text{gen}}, f_{\text{gt}}), 0) \quad (18)$$

where $f_{\text{gen}} \in \mathbb{R}^d$ denotes the semantic feature vector of the generated video. Specifically, we first employ a vision–language model (VLM), **Qwen2.5-VL** (Bai et al., 2025b), to produce a dense structured description L_{gen} of the generated video under task-oriented prompting, covering both task summary and action sequence. This text is then encoded by the CLIP (Radford et al., 2021) text encoder Φ_{txt} , yielding $f_{\text{gen}} = \Phi_{\text{txt}}(L_{\text{gen}})$.

Similarly, $f_{\text{gt}} \in \mathbb{R}^d$ denotes the semantic feature vector of the ground-truth (GT) video, obtained via the same pipeline as $f_{\text{gen}} = \Phi_{\text{txt}}(L_{\text{gt}})$, where L_{gt} is the structured description of the reference video. The scaling factor w ensures score normalization. A higher value of S_{clip} indicates stronger semantic alignment between the generated video and the reference execution.

B.16. Action Response Sensitivity

This metric evaluates the model’s ability to produce distinct and correct outcomes for different action instructions. In open-loop prediction tasks, a robust model should execute multiple instructions faithfully rather than collapsing into repetitive patterns. Given a single action instruction, we manually annotate or automatically generate multiple distinct action instructions and prompt the model to generate N corresponding videos.

For each generated video V_k , we extract a global CLIP feature vector f_k . The action-following diversity score is computed as the average pairwise feature dissimilarity (1-cosine similarity between two vectors f_i and f_j) (Yue et al., 2025):

$$S_{\text{div}} = \frac{1}{|\text{Pairs}(i, j)|} \sum_{i < j} \left(1 - \frac{f_i \cdot f_j}{\|f_i\| \|f_j\|} \right) \quad (19)$$

A higher value of this metric indicates stronger capability of the model in correctly executing action instructions, which is already normalized.

B.17. Score Normalization

Several metrics in our evaluation framework require normalization and direction alignment to ensure consistent interpretation and fair comparison across different models. Specifically, the **Flow Score**, **Trajectory Accuracy**, **Photometric Consistency**, and **Motion Smoothness** metrics are initially measured on different scales, while some metrics such as **JEPA Similarity** and **Depth Accuracy** represent error measures where lower values indicate better performance. To address these inconsistencies, we apply a two-step normalization procedure.

For the **Flow Score**, **Trajectory Accuracy**, **Photometric Consistency**, and **Motion Smoothness** metrics, we employ **empirical min-max** normalization based on the distribution of scores across all evaluated models. We compute the 99th and 1st percentiles of each metric across all videos generated by the 8 models, which serve as the empirical maximum and minimum bounds, respectively. And the specific numerical values for these empirical bounds are provided in Table 5. The final normalized score is calculated as:

Table 5

Empirical bounds for metric normalization. The values represent the 99th percentile (maximum) and 1st percentile (minimum) of each metric across all evaluated videos.

Metric	Empirical Maximum ($S_{\text{empirical}}^{\text{max}}$)	Empirical Minimum ($S_{\text{empirical}}^{\text{min}}$)
Photometric Consistency(Higher is Better)	6.7899	0.1257
Motion Smoothness (Higher is Better)	2.6413	0.0000
Trajectory Accuracy (Higher is Better)	40.8540	0.0000
Flow Score (Higher is Better)	8.9414	0.0531
Depth Accuracy (Lower is Better)	4.3711	0.2228

$$S_{\text{final}} = \max \left(0, \min \left(1, \frac{S_{\text{raw}} - S_{\text{empirical}}^{\text{min}}}{S_{\text{empirical}}^{\text{max}} - S_{\text{empirical}}^{\text{min}}} \right) \right) \quad (20)$$

where S_{raw} denotes the raw metric value, $S_{\text{empirical}}^{\text{max}}$ and $S_{\text{empirical}}^{\text{min}}$ represent the empirical bounds. This transformation ensures that all scores reside within the interval $[0, 1]$, with higher values indicating better performance.

For **Depth Accuracy**, which originally measures reconstruction error (lower values are better), we apply the same normalization but invert the direction:

$$S_{\text{final}} = 1 - \max \left(0, \min \left(1, \frac{S_{\text{raw}} - S_{\text{empirical}}^{\text{min}}}{S_{\text{empirical}}^{\text{max}} - S_{\text{empirical}}^{\text{min}}} \right) \right) \quad (21)$$

This comprehensive normalization strategy ensures that all metrics are scaled to the unit interval $[0, 1]$, aligned in direction (higher values always denote better performance), and comparable across different evaluation dimensions.

C. The Prompt of VLM-based Policy Success Judgement in Policy Evaluator Task

In the embodied policy evaluator task (Section 4.2.2), we assess whether world models can serve as proxy simulation environments for policy evaluation. To determine task success, we employ a VLM-based judge that compares the policy-generated video rollouts against ground-truth reference trajectories. The judge evaluates three critical aspects: (1) correct arm selection when specified in the instruction, (2) task completion by comparing final states between generated and ground-truth videos, and (3) overall action intent consistency. This evaluation approach accounts for visual artifacts inherent to world model rendering while focusing on functional correctness, enabling scalable and automated assessment of policy execution quality. The complete system prompt used for this VLM-based evaluation is provided below.

Policy Execution Evaluation System Prompt

You are a robot task execution judge. Please determine if the policy model correctly executed the instruction.

****Task Instruction**:** {instruction}

****INPUT DESCRIPTION**:**

- First 5 images: GT (Ground Truth) video frames (uniformly sampled: first frame, 3 middle frames, last frame), showing the correct task execution
- Last 5 images: Policy model generated video frames, showing the policy model's execution

****EVALUATION CRITERIA**** (by priority):

1. Arm Selection - If the instruction explicitly requires left/right arm, the correct arm must be used, otherwise fail
2. Task Completion - Compare GT's final state with Policy's final state:
 - GT's final frame shows the completed task state
 - Policy's final frame should show a similar completion state
 - If Policy's final frame differs significantly from GT's final frame, judge as failure
3. Action Intent - Is Policy's entire motion process consistent with the instruction's semantic meaning?

****TOLERABLE DIFFERENCES**:**

- Visual hallucinations from world model rendering (object deformation, color shifts)
- Minor differences in action trajectory
- Video length differences

****JUDGE AS SUCCESS (1)**:**

- Correct arm used
- Final state similar to GT (task basically completed)
- Correct action intent

****JUDGE AS FAILURE (0)**:**

- Wrong arm used
- Final state significantly different from GT (task not completed or completed incorrectly)
- Completely wrong action direction
- Grabbed/operated wrong object

Please carefully compare the ****final frames**** of GT and Policy to judge if the task is basically completed.

****OUTPUT FORMAT REQUIREMENTS**:**

Please respond in this format:

thinking: [Analysis: 1. Is arm correct? 2. Compare final frame task completion states 3. Is action intent consistent?]

answer: [0 or 1]

D. Dataset Details

RoboTwin 2.0 contains 50 task scenarios and 2,500 videos, of which 2,000 are used for world model training and 500 for video-quality testing, which runs on 8 A800 GPUs and takes approximately 8 hours on average to conduct the whole evaluation pipeline of 16 video quality metrics. For embodied data engine evaluation, we use two representative tasks: *adjust bottle*, where the robot picks up a randomly placed bottle and places it upright with the correct arm, and *click bell*, where the robot presses the top center of a randomly placed bell with the gripper end-effector. We train

Table 6

VLM vs. Human evaluation of task success.

Policy Checkpoint	VLM Success Rate	Human Success Rate
10% data	56%	60%
20% data	78%	74%
30% data	72%	74%
50% data	78%	86%
100% data	82%	90%

$\pi_{0.5}$ policies with 10%, 20%, 30%, 50%, and 100% of the data to obtain policies with varying performance. We also evaluate on LIBERO, a single-arm manipulation benchmark, using a LIBERO-Goal tasks: *Turn on the Stove* with 50 trajectories.

E. Details of Tested Models

The evaluated general video world models include CogvideoX (Yang et al., 2024b), Wan 2.2 (Wan et al., 2025), Wan 2.6 (Wan et al., 2025), and Veo 3.1³. The text-conditioned embodied world models consist of Genie Envisioner (Liao et al., 2025), GigaWorld (Team et al., 2025), Tesseract (Zhen et al., 2025), WOW (Chi et al., 2025), RoboMaster⁴, Cosmos-Predict 2.5 (text) (Gu, 2025), and Vidar (Feng et al., 2025). In addition, we include action-conditioned embodied world models, namely IRASim (Zhu et al., 2025), Cosmos-Predict 2.5 (action) (Gu, 2025), Roboscape (Shang et al., 2025b) and CtrlWorld (Guo et al., 2025). For fair comparison, all models with available training code are post-trained on the used dataset following their official implementations.

F. VLM-judge Reliability Analysis

We analyze the reliability of the VLM-based judge of 3 video quality metrics from both stability and human-alignment perspectives. For stability, we test four variants of the original evaluation setting: changing the random seed, reducing the sampled frames from 16 to 8, rewriting the prompt, and replacing Qwen3-VL with InternVL-3.5-8B. The resulting scores remain highly consistent with the original setting, achieving Pearson correlations of 0.982, 0.973, 0.979, and 0.913, respectively, while preserving the model ranking. In addition, repeating the Qwen3-VL evaluation with five different seeds shows that Interaction Quality varies by only 3.92% on average, with no ranking changes.

For human alignment, we collect 12 human annotations for the VLM-scored metrics. Qwen3-VL shows strong agreement with human judgments, with Pearson correlations of 0.818 for Interaction Quality and 0.853 for Perspectivity. These results indicate that our VLM-based evaluation is robust to implementation choices and broadly consistent with human assessment.

In addition to the VLM used in the policy evaluator, we further compare VLM and human success judgments on 50 episodes from 5 policy checkpoints trained with different data scales. The process is shown in Figure 6. The result in Table 6 shows that the two are highly consistent ($r = 0.91$), suggesting that the VLM is not the main source of overestimation.

G. Detailed Video Quality Results on LIBERO

We provide detailed visual metrics on LIBERO of tested models in Table 7 and Table 8 and EWMScore-P and EWMScore-F results are shown in Figure 7.

H. Synthetic-data Scaling Ablation

To assess the effect of synthetic data scale in the data-engine setting, we increase the number of synthetic trajectories from 25 to 100 per task and evaluate three representative world models on two tasks. As shown in Table 9, larger synthetic datasets consistently improve performance across both tasks. Notably, the relative ranking among models

³<https://aistudio.google.com/models/veo-3>

⁴<https://huggingface.co/datasets/robomaster2025/RoboMaster>



Figure 6: Qualitative evaluation of world model rollouts: We compare the keyframe sequences generated by policies trained on varying data scales (10% to 100%) against the GT. These sequences are used as visual prompts for both VLM-based and human-based success evaluation.

remains largely unchanged: WoW performs best overall, Vidar remains competitive, and RoboMaster lags behind despite benefiting from additional trajectories. These results indicate that increasing synthetic data scale provides moderate gains, but the observed performance differences are not solely attributable to extreme data scarcity.

I. Instructions Given to Participants

We provide text instructions for annotators in Figure 8 and Figure 9.

J. Case Comparison of Each Metric in EWMScore




Figure 7: EWMScore results for representative embodied world models on LIBERO data.

Table 7

Video quality evaluation results across visual quality, motion quality and content consistency dimensions on Libero dataset.

Models	Visual Quality			Motion Quality			Content Consistency		
	Image Quality	Aesthetic Quality	JEPA Similarity	Dynamic Degree	Flow Score	Motion Smoothness	Subject Consistency	Backgrd. Consistency	Photom. Consistency
GigaWorld-0	0.3056	0.4860	0.1693	0.2387	0.0814	0.5886	0.9082	0.9481	0.7859
Genie Envioner	0.3688	0.2518	0.0303	0.3330	0.2273	0.5776	0.6867	0.8670	0.2074
Tesseract	0.3779	0.5249	0.3524	0.5060	0.1433	0.6917	0.8659	0.9023	0.3294
RoboMaster	0.3530	0.3883	0.0468	0.3729	0.1263	0.5780	0.9080	0.9347	0.1964
Vidar	0.3071	0.4881	0.3725	0.0986	0.0342	0.5026	0.3607	0.3705	0.4354
Cosmos-Predict 2.5 (text)	0.7406	0.5014	0.0205	0.3174	0.6573	0.9590	0.6578	0.8210	0.3818
Cosmos-Predict 2.5 (action)	0.3642	0.5064	0.4205	0.1150	0.0374	0.4465	0.5550	0.5742	0.6496
WoW	0.3573	0.5061	0.4612	0.1972	0.0488	0.4773	0.7843	0.8184	0.8629
CtrlWorld	0.3530	0.4996	0.6087	0.2358	0.0698	0.5467	0.8396	0.8697	0.8227
Wan 2.2	0.3197	0.4753	0.2919	0.1341	0.0525	0.5597	0.3954	0.4253	0.4681
CogvideoX	0.3044	0.4614	0.4513	0.1334	0.0429	0.4793	0.5852	0.6036	0.6631
IRASim	0.3484	0.4843	0.4407	0.1364	0.0442	0.5012	0.5633	0.5941	0.7352

Table 8

Video quality evaluation results across physics adherence, 3D accuracy and controllability dimensions on Libero dataset.

Models	Physics Adherence		3D Accuracy		Controllability		
	Interaction Quality	Trajectory Accuracy	Depth Accuracy	Perspectivity	Instruction Following	Semantic Alignment	Action Response Sensitivity
GigaWorld-0	0.6225	0.3690	0.8272	0.8090	0.7281	0.9102	0.0100
Genie Envioner	0.2000	0.1495	0.7520	0.4920	0.2000	0.8758	0.0187
Tesseract	0.6000	0.3381	0.7039	0.7960	0.6380	0.9029	0.0185
RoboMaster	0.4060	0.1984	0.7918	0.6500	0.3580	0.9021	0.0363
Vidar	0.4100	0.2810	0.9133	0.7940	0.3820	0.9179	0.0389
Cosmos-Predict 2.5 (text)	0.2340	0.0015	0.5853	0.9400	0.2000	0.7637	0.1226
Cosmos-Predict 2.5 (action)	0.5600	0.4552	0.9236	0.8020	0.4200	0.9131	0.0017
WoW	0.6500	0.5859	0.9381	0.8060	0.7900	0.9147	0.0121
CtrlWorld	0.5160	0.4824	0.9431	0.7920	0.4040	0.9212	0.0087
Wan 2.2	0.5360	0.2467	0.8901	0.7960	0.5180	0.9102	0.0364
CogvideoX	0.5200	0.6676	0.9267	0.7800	0.4820	0.9182	0.0177
IRASim	0.6200	0.8427	0.9184	0.7800	0.5620	0.9240	0.0071

Table 9

Synthetic-data scaling ablation on data-engine tasks. Increasing the number of synthetic trajectories from 25 to 100 improves performance across models, while the relative ranking remains largely consistent.

Model	adjust_bottle		click_bell	
	25 traj.	100 traj.	25 traj.	100 traj.
WoW	25	34	26	38
Vidar	24	30	27	33
RoboMaster	2	10	18	25

[Evaluation Training](#) [Start Rating](#) [Start Comparing](#)

Please read carefully before the evaluation begins

1. Evaluate the interaction quality of the video

No Meaningful Interaction

No effective physical interaction occurs. Objects pass through the robot arm or other objects without any contact, or the robot interacts with the wrong object entirely. Actions such as grasping, pushing, opening, or moving do not actually happen, and no task-relevant object state change is observed. The robot arm may move, but the motion has no meaningful physical relation to the environment or object.

Poor Interaction

Contact exists but interaction is physically unrealistic or severely flawed. The robot makes contact, but the contact point, timing, or target is clearly wrong. Obvious physical errors occur, such as sliding without friction, incorrect force response, or unnatural inertia. Frequent failures (e.g., slipping, misaligned contact, interrupted manipulation) prevent effective task execution. The object may change state, but the interaction quality is clearly insufficient to achieve the intended goal.

Moderately Effective Interaction

Partially effective interaction with noticeable physical flaws. The video shows meaningful contact between the robot arm and the object, and part of the intended manipulation is achieved. However, the interaction exhibits noticeable issues, such as slight penetration, imperfect grasping that lacks proper friction, or unnatural force transmission. The manipulation is only moderately convincing, and task completion is incomplete due to a lack of precision or stability.

Mostly Effective Interaction

Realistic physical interaction with only minor issues. The robot arm makes reasonable contact with the target object, and the manipulation is clearly task-oriented. Physical behaviors such as friction, force transfer, and object response are mostly plausible. Actions like grasping, pushing, opening, and placing are generally natural and stable. Only minor physical artifacts may occur (e.g., subtle sliding or slight penetration), and the outcome is mostly consistent with the intended task.

Highly Effective Interaction

Perfect interaction physics, indistinguishable from real robot manipulation. Contact behavior, friction, force transmission, inertia, and boundary integrity are all accurate and coherent. Each manipulation action causes a physically reasonable and expected change in the object state. Grasping, manipulation, and release are seamlessly coordinated with no penetration, misalignment, interruption, or unnatural motion. The task is carried out naturally and completely, with physics that fully support the goal-directed manipulation.

Figure 8: The first part of text instructions for annotators.

2. Evaluate the perspective consistency of the video

Severely Broken Perspective

Severe perspective errors: The relationships of distance, scale, and depth in the scene are clearly broken, making the 3D space impossible to trust.
viewpoint changes: When the camera or viewpoint changes, object shapes, positions, or background structure collapse noticeably.
Incorrect occlusion relationships: Foreground and background ordering is confused, and occlusion violates basic spatial logic, severely harming spatial understanding.

Weak Perspective Consistency

Poor perspective consistency: Object scale, depth, and relative spatial positions often change unnaturally.
Insufficient viewpoint stability: As the camera or subject moves, spatial relationships drift, deform, or jump noticeably.
Local spatial inconsistencies: Occlusion, distance, and orientation are often incorrect, reducing the realism of the scenes 3D structure.

Moderately Consistent Perspective

Perspective is generally understandable: The scene conveys a basic sense of spatial structure and viewpoint, but noticeable instability remains.
Local errors are present: In some frames, object size, position, depth, or occlusion relationships are not fully correct.
Moderate 3D realism: The video shows some spatial plausibility, but not enough to support consistently reliable viewpoint perception.

Mostly Consistent Perspective

Perspective is mostly reasonable: Distance, scale, depth, and occlusion relationships are generally stable and support natural spatial understanding.
Viewpoint changes are mostly natural: During camera motion or viewpoint shifts, the geometric relationships between objects and background remain mostly consistent.
Only minor issues: Small-scale drift, slight occlusion errors, or limited spatial discontinuities may appear occasionally, but they do not significantly affect overall realism.

Highly Consistent Perspective

Fully consistent perspective: Distance relationships, scale changes, depth ordering, and occlusion remain accurate and natural throughout the video.
Viewpoint transitions are highly believable: No matter how the camera moves, the scene maintains stable and coherent 3D geometric structure.
Highly realistic spatial perception: The viewer can continuously perceive a realistic, stable, and coherent viewpoint and spatial layout, with almost no perspective-related errors.

Figure 9: The second part of text instructions for annotators.




Figure 10: Typical examples of Visual Quality. *Top:Image Quality.* The bad example on the right-hand-side exhibits significant motion blur and noise, while the good example preserves sharp structural details. *Middle:Aesthetic Quality.* The bad example suffers from severe geometric distortion and artifacts. Conversely, the good example demonstrates superior contrast and realistic lighting with clear reflections. *Bottom:JEPA Similarity.* In the good example, the style and morphology closely align with the GT, while in the bad example, the robotic gripper shows color discrepancies and introduces unintended grid artifacts not present in the GT.




Figure 11: Typical examples of Motion Quality. *Top:Dynamic Degree.* The good example shows the robotic arm exhibiting a complete and distinct motion sequence from picking up the bottle to placing it in the dustbin, while in the bad example, robotic arm remains static with only minor flickering of the bottle. *Middle:Flow Score.* The good example demonstrates a fluid manipulation of rotating the bottle with significant pixel-level movement and the bad example shows negligible motion, with only slight deformation at the top of the bottle. *Bottom:Motion Smoothness.* The good example features a stable and continuous translation of the hammer, but the bad example suffers from erratic shaking and disjointed, sharp movements immediately after grasping the object.




Figure 12: Typical examples of Content Consistency. *Top: Subjective Consistency.* In the good example, the bottle's shape, color, and packaging remain stable and coherent throughout the grasping process. In the bad example, the bottle suffers from severe deformation and structural chaos, losing its original identity. *Middle: Background Consistency.* The good example maintains a stable background and camera perspective during the cabinet interaction. Conversely, the bad example on the right exhibits a sudden camera shift to a top-down view, leading to an unstable and rapidly changing background. *Bottom: Photometric Consistency.* In the good example, the appearance and color of both the block and the robotic arm are consistently preserved. In the bad example, the grasped block undergoes an unnatural color transition from green to red, indicating poor photometric stability.




Figure 13: Typical examples of Physics Adherence. *Top: Interaction Quality.* In the good example, the robotic gripper interacts with the bread appropriately. In the bad example, the bread is lifted without any physical contact with the gripper, violating the fundamental physics laws. *Bottom: Trajectory Accuracy.* The good example demonstrates a movement trajectory that highly aligns with GT. Conversely, the bad example exhibits significant deviations from the GT trajectory, characterized by anomalous movements and jitter.

3D Accuracy




Figure 14: Typical examples of 3D Accuracy. *Top:Depth Accuracy.* In the good example, the generated depth map highly aligns with the GT, ensuring stable spatial and geometric structures, but the bad example suffers from severe geometric distortion, where the gripper unnaturally merges with the green block, leading to a collapse of spatial integrity. *Bottom:Perspectivity.* The good example maintains realistic perspective and lighting. Conversely, the bad example shows significant ghosting and blurring during movement, failing to preserve the object’s contour and exhibiting no shadow of robotic arm that deviate from physical reality.

Controllability




Figure 15: Typical examples of Controllability. *Top:Instruction Following.* In the good example, the model strictly adheres to the task instruction, but the bad example shows the movement of the incorrect object (knife), failing to execute instruction. *Middle:Semantic Alignment.* The good example demonstrates high semantic fidelity, but the bad example exhibits low alignment by transforming the QR code into a clothing tag and introducing irrational human hands not present in the target semantics. *Bottom:Action Following.* The good example successfully performs distinct actions based on varying prompts, placing the shoe at both the blue marker and to its left, but the bad example shows the model demonstrates limited discriminative ability, executing a singular action regardless of the instruction.

References

- Assran, M., Bardes, A., Fan, D., Garrido, Q., Howes, R., Muckley, M., Rizvi, A., Roberts, C., Sinha, K., Zholus, A., et al., 2025. V-jepa 2: Self-supervised video models enable understanding, prediction and planning. arXiv preprint arXiv:2506.09985 .
- Bai, S., Cai, Y., Chen, R., Chen, K., Chen, X., Cheng, Z., Deng, L., Ding, W., Gao, C., Ge, C., Ge, W., Guo, Z., Huang, Q., Huang, J., Huang, F., Hui, B., Jiang, S., Li, Z., Li, M., Li, M., Li, K., Lin, Z., Lin, J., Liu, X., Liu, J., Liu, C., Liu, Y., Liu, D., Liu, S., Lu, D., Luo, R., Lv, C., Men, R., Meng, L., Ren, X., Ren, X., Song, S., Sun, Y., Tang, J., Tu, J., Wan, J., Wang, P., Wang, P., Wang, Q., Wang, Y., Xie, T., Xu, Y., Xu, H., Xu, J., Yang, Z., Yang, M., Yang, J., Yang, A., Yu, B., Zhang, F., Zhang, H., Zhang, X., Zheng, B., Zhong, H., Zhou, J., Zhou, F., Zhou, J., Zhu, Y., Zhu, K., 2025a. Qwen3-vl technical report. arXiv preprint arXiv:2511.21631 .
- Bai, S., Chen, K., Liu, X., Wang, J., Ge, W., Song, S., Dang, K., Wang, P., Wang, S., Tang, J., Zhong, H., Zhu, Y., Yang, M., Li, Z., Wan, J., Wang, P., Ding, W., Fu, Z., Xu, Y., Ye, J., Zhang, X., Xie, T., Cheng, Z., Zhang, H., Yang, Z., Xu, H., Lin, J., 2025b. Qwen2.5-vl technical report. arXiv preprint arXiv:2502.13923 .
- Bardes, A., Garrido, Q., Ponce, J., Chen, X., Rabbat, M., LeCun, Y., Assran, M., Ballas, N., 2023. V-jepa: Latent video prediction for visual representation learning .
- Carion, N., Gustafson, L., Hu, Y.T., Debnath, S., Hu, R., Suris, D., Ryali, C., Alwala, K.V., Khedr, H., Huang, A., et al., 2025. Sam 3: Segment anything with concepts. arXiv preprint arXiv:2511.16719 .
- Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A., 2021. Emerging properties in self-supervised vision transformers, in: Proceedings of the IEEE/CVF international conference on computer vision, pp. 9650–9660.
- Chen, T., Chen, Z., Chen, B., Cai, Z., Liu, Y., Li, Z., Liang, Q., Lin, X., Ge, Y., Gu, Z., et al., 2025. Robotwin 2.0: A scalable data generator and benchmark with strong domain randomization for robust bimanual robotic manipulation. arXiv preprint arXiv:2506.18088 .
- Chi, X., Jia, P., Fan, C.K., Ju, X., Mi, W., Zhang, K., Qin, Z., Tian, W., Ge, K., Li, H., et al., 2025. Wow: Towards a world omniscient world model through embodied interaction. arXiv preprint arXiv:2509.22642 .
- Ding, J., Zhang, Y., Shang, Y., Zhang, Y., Zong, Z., Feng, J., Yuan, Y., Su, H., Li, N., Sukiennik, N., et al., 2025. Understanding world or predicting future? a comprehensive survey of world models. ACM Computing Surveys 58, 1–38.
- Duan, H., Yu, H.X., Chen, S., Fei-Fei, L., Wu, J., 2025. Worldscore: A unified evaluation benchmark for world generation. arXiv preprint arXiv:2504.00983 .
- Fan, C.K., Chi, X., Ju, X., Li, H., Bao, Y., Wang, Y.K., Chen, L., Jiang, Z., Ge, K., Li, Y., et al., 2026. Wow, wo, val! a comprehensive embodied world model evaluation turing test. arXiv preprint arXiv:2601.04137 .
- Feng, Y., Tan, H., Mao, X., Xiang, C., Liu, G., Huang, S., Su, H., Zhu, J., 2025. Vidar: Embodied video diffusion model for generalist manipulation. arXiv preprint arXiv:2507.12898 .
- Gu, J., 2025. Cosmos world foundation models for physical ai, in: Proceedings of the 3rd International Workshop on Rich Media With Generative AI, pp. 39–39.
- Guo, Y., Shi, L.X., Chen, J., Finn, C., 2025. Ctrl-world: A controllable generative world model for robot manipulation. arXiv preprint arXiv:2510.10125 .
- Hafner, D., Pasukonis, J., Ba, J., Lillicrap, T., 2025. Mastering diverse control tasks through world models. Nature , 1–7.
- Hu, Y., Guo, Y., Wang, P., Chen, X., Wang, Y.J., Zhang, J., Sreenath, K., Lu, C., Chen, J., 2024. Video prediction policy: A generalist robot policy with predictive visual representations. arXiv preprint arXiv:2412.14803 .
- Huang, S., Chen, L., Zhou, P., Chen, S., Jiang, Z., Hu, Y., Liao, Y., Gao, P., Li, H., Yao, M., et al., 2025. Enerverse: Envisioning embodied future space for robotics manipulation. arXiv preprint arXiv:2501.01895 .
- Huang, Z., He, Y., Yu, J., Zhang, F., Si, C., Jiang, Y., Zhang, Y., Wu, T., Jin, Q., Chanpaisit, N., et al., 2024. Vbench: Comprehensive benchmark suite for video generative models, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 21807–21818.
- Intelligence, P., Black, K., Brown, N., Darpinian, J., Dhabalia, K., Driess, D., Esmail, A., Equi, M., Finn, C., Fusai, N., Galliker, M.Y., Ghosh, D., Groom, L., Hausman, K., Ichter, B., Jakubczak, S., Jones, T., Ke, L., LeBlanc, D., Levine, S., Li-Bell, A., Mothukuri, M., Nair, S., Pertsch, K., Ren, A.Z., Shi, L.X., Smith, L., Springenberg, J.T., Stachowicz, K., Tanner, J., Vuong, Q., Walke, H., Walling, A., Wang, H., Yu, L., Zhilinsky, U., 2025. $\pi_{0.5}$: a vision-language-action model with open-world generalization URL: <https://arxiv.org/abs/2504.16054>, arXiv:2504.16054.
- Jang, J., Ye, S., Lin, Z., Xiang, J., Bjorck, J., Fang, Y., Hu, F., Huang, S., Kundalia, K., Lin, Y.C., et al., 2025. Dreamgen: Unlocking generalization in robot learning through video world models. arXiv preprint arXiv:2505.12705 .
- Ke, J., Wang, Q., Wang, Y., Milanfar, P., Yang, F., 2021. Musiq: Multi-scale image quality transformer, in: Proceedings of the IEEE/CVF international conference on computer vision, pp. 5148–5157.
- Kong, L., Yang, W., Mei, J., Liu, Y., Liang, A., Zhu, D., Lu, D., Yin, W., Hu, X., Jia, M., et al., 2025. 3d and 4d world modeling: A survey. arXiv preprint arXiv:2509.07996 .
- LAION-AI, 2022. Aesthetic predictor. <https://github.com/LAION-AI/aesthetic-predictor>. Accessed: 2024.
- Li, D., Fang, Y., Chen, Y., Yang, S., Cao, S., Wong, J., Luo, M., Wang, X., Yin, H., Gonzalez, J.E., et al., 2025a. Worldmodelbench: Judging video generation models as world models. arXiv preprint arXiv:2502.20694 .
- Li, Y., Zhu, Y., Wen, J., Shen, C., Xu, Y., 2025b. Worldval: World model as real-world robot policies evaluator. arXiv preprint arXiv:2505.19017 .
- Liang, A., Kong, L., Yan, T., Liu, H., Yang, W., Huang, Z., Yin, W., Zuo, J., Hu, Y., Zhu, D., Lu, D., Liu, Y., Jiang, G., Li, L., Li, X., Zhuo, L., Ng, L.X., Cottreau, B.R., Gao, C., Pan, L., Ooi, W.T., Liu, Z., 2025. WorldLens: Full-spectrum evaluations of driving world models in real world. arXiv preprint arXiv:2512.10958 .
- Liao, Y., Zhou, P., Huang, S., Yang, D., Chen, S., Jiang, Y., Hu, Y., Cai, J., Liu, S., Luo, J., et al., 2025. Genie envisioner: A unified world foundation platform for robotic manipulation. arXiv preprint arXiv:2508.05635 .
- Liu, B., Zhu, Y., Gao, C., Feng, Y., Liu, Q., Zhu, Y., Stone, P., 2023a. Libero: Benchmarking knowledge transfer for lifelong robot learning. Advances in Neural Information Processing Systems 36, 44776–44791.
- Liu, B.K., Chen, Z.P., 2025. Jepa-reasoner: Decoupling latent reasoning from token generation. arXiv preprint arXiv:2512.19171 .

- Liu, Y., Cun, X., Liu, X., Wang, X., Zhang, Y., Chen, H., Liu, Y., Zeng, T., Chan, R., Shan, Y., 2023b. Evalcrafter: Benchmarking and evaluating large video generation models. arXiv preprint arXiv:2310.11440 .
- Long, X., Zhao, Q., Zhang, K., Zhang, Z., Wang, D., Liu, Y., Shu, Z., Lu, Y., Wang, S., Wei, X., et al., 2025. A survey: Learning embodied intelligence from physical simulators and world models. arXiv preprint arXiv:2507.00917 .
- Lu, Y., Luo, W., Tu, P., Li, H., Zhu, H., Yu, Z., Wang, X., Chen, X., Peng, X., Li, X., et al., 2025. 4dworldbench: A comprehensive evaluation framework for 3d/4d world generation models. arXiv preprint arXiv:2511.19836 .
- Luo, G.Y., Favero, G., Luo, Z.H., Jolicoeur-Martineau, A., Pal, C., 2024. Beyond fvd: Enhanced evaluation metrics for video generation quality. URL: <https://arxiv.org/abs/2410.05203>, arXiv:2410.05203.
- Müller, M., 2007. Information Retrieval for Music and Motion. doi:10.1007/978-3-540-74048-3.
- Qian, Z., Chi, X., Li, Y., Wang, S., Qin, Z., Ju, X., Han, S., Zhang, S., 2025. Wristworld: Generating wrist-views via 4d world models for robotic manipulation. arXiv preprint arXiv:2510.07313 .
- Qin, Y., Shi, Z., Yu, J., Wang, X., Zhou, E., Li, L., Yin, Z., Liu, X., Sheng, L., Shao, J., et al., 2024. Worldsimbench: Towards video generation models as world simulators. arXiv preprint arXiv:2410.18072 .
- Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., Sutskever, I., 2021. Learning transferable visual models from natural language supervision. URL: <https://arxiv.org/abs/2103.00020>, arXiv:2103.00020.
- Shang, Y., Tang, Y., Zhang, X., Wang, S., Yan, Y., Zhang, H., Zheng, Z., Zhao, J., Feng, J., Gao, C., et al., 2025a. A survey of embodied world models .
- Shang, Y., Zhang, X., Tang, Y., Jin, L., Gao, C., Wu, W., Li, Y., 2025b. Roboscape: Physics-informed embodied world model. arXiv preprint arXiv:2506.23135 .
- Team, G., Ye, A., Wang, B., Ni, C., Huang, G., Zhao, G., Li, H., Zhu, J., Li, K., Xu, M., et al., 2025. Gigaworld-0: World models as data engine to empower embodied ai. arXiv preprint arXiv:2511.19861 .
- Teed, Z., Deng, J., 2020. Raft: Recurrent all-pairs field transforms for optical flow, in: European conference on computer vision, Springer. pp. 402–419.
- Wan, T., Wang, A., Ai, B., Wen, B., Mao, C., Xie, C.W., Chen, D., Yu, F., Zhao, H., Yang, J., et al., 2025. Wan: Open and advanced large-scale video generative models. arXiv preprint arXiv:2503.20314 .
- Yang, L., Kang, B., Huang, Z., Zhao, Z., Xu, X., Feng, J., Zhao, H., 2024a. Depth anything v2. arXiv:2406.09414 .
- Yang, Z., Teng, J., Zheng, W., Ding, M., Huang, S., Xu, J., Yang, Y., Hong, W., Zhang, X., Feng, G., et al., 2024b. Cogvideox: Text-to-video diffusion models with an expert transformer. arXiv preprint arXiv:2408.06072 .
- Yue, H., Huang, S., Liao, Y., Chen, S., Zhou, P., Chen, L., Yao, M., Ren, G., 2025. Ewmbench: Evaluating scene, motion, and semantic quality in embodied world models. arXiv preprint arXiv:2505.09694 .
- Zhang, G., Liu, C., Cui, Y., Zhao, X., Ma, K., Wang, L., 2024. Vfimbamba: Video frame interpolation with state space models. URL: <https://arxiv.org/abs/2407.02315>, arXiv:2407.02315.
- Zhang, J., Jiang, M., Dai, N., Lu, T., Uzunoglu, A., Zhang, S., Wei, Y., Wang, J., Patel, V.M., Liang, P.P., et al., 2025. World-in-world: World models in a closed-loop world. arXiv preprint arXiv:2510.18135 .
- Zhen, H., Sun, Q., Zhang, H., Li, J., Zhou, S., Du, Y., Gan, C., 2025. Tesseract: learning 4d embodied world models. arXiv preprint arXiv:2504.20995 .
- Zhu, F., Wu, H., Guo, S., Liu, Y., Cheang, C., Kong, T., 2025. Irasim: A fine-grained world model for robot manipulation, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 9834–9844.
- Zhu, Z., Wang, X., Zhao, W., Min, C., Li, B., Deng, N., Dou, M., Wang, Y., Shi, B., Wang, K., et al., 2024. Is sora a world simulator? a comprehensive survey on general world models and beyond. arXiv preprint arXiv:2405.03520 .