arXiv:2603.18811v1 [cs.RO] 19 Mar 2026

# V-Dreamer: Automating Robotic Simulation and Trajectory Synthesis via Video Generation Priors

Songjia He[1*], Zixuan Chen[1*], Hongyu Ding[1], Dian Shao[2], Jieqi Shi[1], Chenxu Li[1], Jing Huo[1], Yang Gao[1]

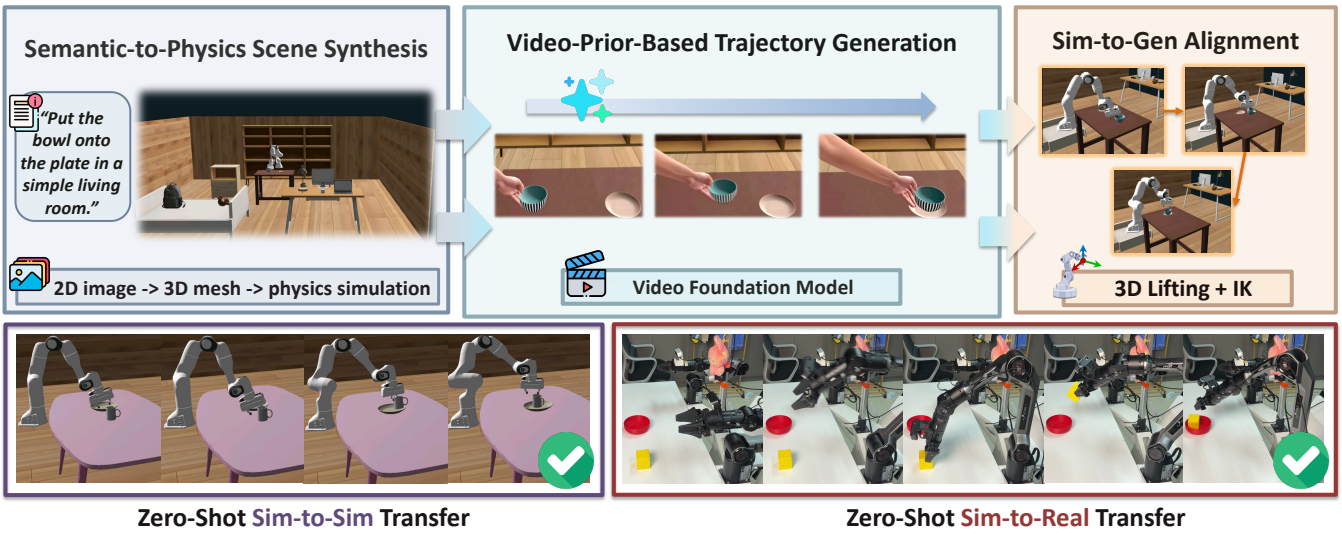


Fig. 1. We propose **V-Dreamer**, a fully automated full-cycle pipeline for instruction-driven, open-vocabulary robotic manipulation data synthesis. Given a natural language instruction (optionally with a real-scene photo), V-Dreamer generates a simulation-ready scene and an executable expert trajectory, which are then used to train a policy that can zero-shot generalize to unseen objects in simulation and transfer zero-shot to real robotic hardware.

***Abstract*— Training generalist robots demands large-scale, diverse manipulation data, yet real-world collection is prohibitively expensive, and existing simulators are often constrained by fixed asset libraries and manual heuristics. To bridge this gap, we present V-Dreamer, a fully automated framework that generates open-vocabulary, simulation-ready manipulation environments and executable expert trajectories directly from natural language instructions. V-Dreamer employs a novel generative pipeline that constructs physically grounded 3D scenes using large language models and 3D generative models, validated by geometric constraints to ensure stable, collision-free layouts. Crucially, for behavior synthesis, we leverage video generation models as rich motion priors. These visual predictions are then mapped into executable robot trajectories via a robust Sim-to-Gen visual-kinematic alignment module utilizing CoTracker3 and VGGT. This pipeline supports high visual diversity and physical fidelity without manual intervention. To evaluate the generated data, we train imitation learning policies on synthesized trajectories encompassing diverse object and environment variations. Extensive evaluations on tabletop manipulation tasks using the Piper robotic arm demonstrate that our policies robustly generalize to unseen objects in simulation and achieve effective sim-to-real transfer, successfully manipulating novel real-world objects. Project page: https://jia-handsome.github.io/v-Dreamer/.**

*These authors contributed equally to this work.
[1]Nanjing University.
[2]Northwestern Polytechnical University.

## I. Introduction

General-purpose robotic manipulation in unstructured environments requires policies capable of generalizing across diverse objects, scenes, and semantic goals. While recent advances in imitation learning have demonstrated impressive capabilities in acquiring complex behaviors from expert demonstrations [1], [2], their performance is fundamentally bottlenecked by the scale and diversity of the training data. Collecting large-scale real-world datasets via teleoperation is prohibitively expensive, time-consuming, and difficult to scale to long-tail scenarios [3], [4]. Although simulation offers a scalable alternative, existing simulators typically rely on fixed asset libraries and manually designed tasks, restricting the semantic and visual diversity necessary for robust sim-to-real transfer.

The emergence of generative foundation models offers a promising avenue to address this data scarcity. Large Language Models (LLMs) and diffusion models encapsulate extensive world knowledge regarding object semantics, spatial relationships, and visual affordances [5], [6]. While recent works have explored using foundation models to generate static 3D scenes [7] or high-level task plans [8], they rarely yield a fully executable robotic task. Video generation can produce visually plausible motion, yet often violates geometric consistency or kinematic constraints due to limited physical grounding [9]. The critical bottleneck is to automatically synthesize dynamic, physically feasible

manipulation trajectories without human supervision.

To bridge this gap, we introduce **V-Dreamer** 1, a fully automated framework that synthesizes open-vocabulary, simulation-ready manipulation environments and executable expert trajectories directly from natural language instructions. V-Dreamer addresses the dual challenges of scene diversity and behavioral realism through a unified generative pipeline. First, our *Semantic-to-Physics Scene Generation* module utilizes LLMs and 3D generative models to construct physically grounded 3D scenes. By integrating geometric reasoning with physics-based validation, the system ensures stable, collision-free object layouts that accurately reflect the user's semantic intent. Crucially, for behavior synthesis, we introduce a novel *Video-Prior-Based Trajectory Generation* approach. Building on the previously generated simulation scene, we leverage state-of-the-art video generation models as rich motion priors to predict the manipulation process, instead of relying on hand-crafted heuristics or traditional motion planners that often lack semantic nuance. To ground these generated visual motions into actionable data, we develop a robust *Sim-to-Gen* visual-kinematic alignment module utilizing CoTracker3 [10] and VGGT [11]. This module extracts physically valid end-effector trajectories, allowing us to generate diverse, policy-compliant demonstrations at scale across a wide range of object categories and interaction dynamics. Overall, V-Dreamer offers a fully automated pipeline that turns natural language instructions into simulation-ready, physically grounded scenes and executable expert trajectories. By combining physics-validated scene generation with video-model motion priors and robust visual-kinematic alignment, it achieves both open-vocabulary scene diversity and behavior realism, enabling scalable, policy-compliant demonstrations across diverse objects and interaction dynamics without human supervision.

In practice, V-Dreamer serves as a high-throughput data engine, generating approximately **600** LeRobot-formatted trajectories **per hour** on an **8×RTX 4090** workstation. This rapid generation bypasses the bottleneck of slow manual teleoperation, enabling highly efficient policy training. Furthermore, extensive evaluations on tabletop manipulation tasks using a Piper robotic arm demonstrate that policies trained on this synthetic data robustly generalize to unseen objects in simulation. To further bridge the reality gap, V-Dreamer also support a photo-conditioned setting that instantiates a scene-matched simulation directly from real-world images and a language instruction, making the real2sim2real pipeline highly practical. By training on just a single synthesized demonstration from this pipeline, we achieve effective **zero-shot sim-to-real transfer**. This highlights the immense potential of generative video priors for scalable and generalizable robot learning.

In summary, our main contributions are:

- We propose **V-Dreamer**, which, to our knowledge, is the first fully automated full-cycle pipeline for instruction-driven open-vocabulary data synthesis. By distilling natural language into simulation-ready environments and expert trajectories, V-Dreamer enables scalable policy learning without human-in-the-loop effort.
- We introduce a novel **video-prior-based behavior synthesis** method. By leveraging a robust visual-kinematic alignment module, we successfully ground generative visual motion into physically feasible end-effector trajectories, eliminating the reliance on hand-crafted heuristics.
- We empirically validate V-Dreamer as a high-throughput data engine capable of supporting robust zero-shot generalization and highly data-efficient one-shot imitation learning. Notably, via a zero-shot real2sim2real workflow, a policy trained on a single synthesized demonstration achieves successful sim-to-real transfer on physical hardware while remaining robust to local spatial perturbations.

## II. Related Work

### A. Generative Simulation for Robotics

Simulation is a cornerstone for scaling robot learning, enabling safe and rapid data collection [12], [13]. Early approaches relied on procedural generation with fixed asset libraries to create diverse environments [7], [14], but these methods are inherently limited by the diversity of hand-crafted asset databases and often lack photorealistic textures. Recent works have begun to leverage generative models to automate simulation design at a higher level. GenSim [15] and RoboGen [16] utilize LLMs to generate task codes and scene configurations, Holodeck [17] compositionally arranges assets into 3D environments from language descriptions, and MimicGen [18] automatically scales demonstrations to new object placements. However, all these methods are fundamentally constrained by fixed asset repositories, limiting visual variety and semantic open-endedness. V-Dreamer overcomes this by synthesizing novel 3D geometries and photorealistic textures entirely from scratch via diffusion models, enabling truly open-vocabulary scene diversity unconstrained by any asset library.

### B. Video Generation as Motion Priors

The rapid progress in video generation [19], [9] has sparked interest in using video models as world simulators or motion priors for robotics. UniSim [20] learns a universal interactive simulator from real-world action-annotated videos. Video language planning [21] and UniPi [22] use video diffusion models to imagine future frames as a planning signal. GR-2 [23] pre-trains on large-scale internet video and fine-tunes for manipulation, showing that video prediction priors enhance generalization. Genie [24] learns controllable world models from unlabeled play videos as interactive training grounds. Despite this progress, most approaches operate in 2D pixel space, making it difficult to extract precise, physically grounded 3D trajectories for robot control. V-Dreamer addresses this by employing a Sim-to-Gen alignment module based on CoTracker3 [10] and VGGT [11] to lift 2D visual motion into physically executable end-effector trajectories, directly bridging visual synthesis and kinematic control.

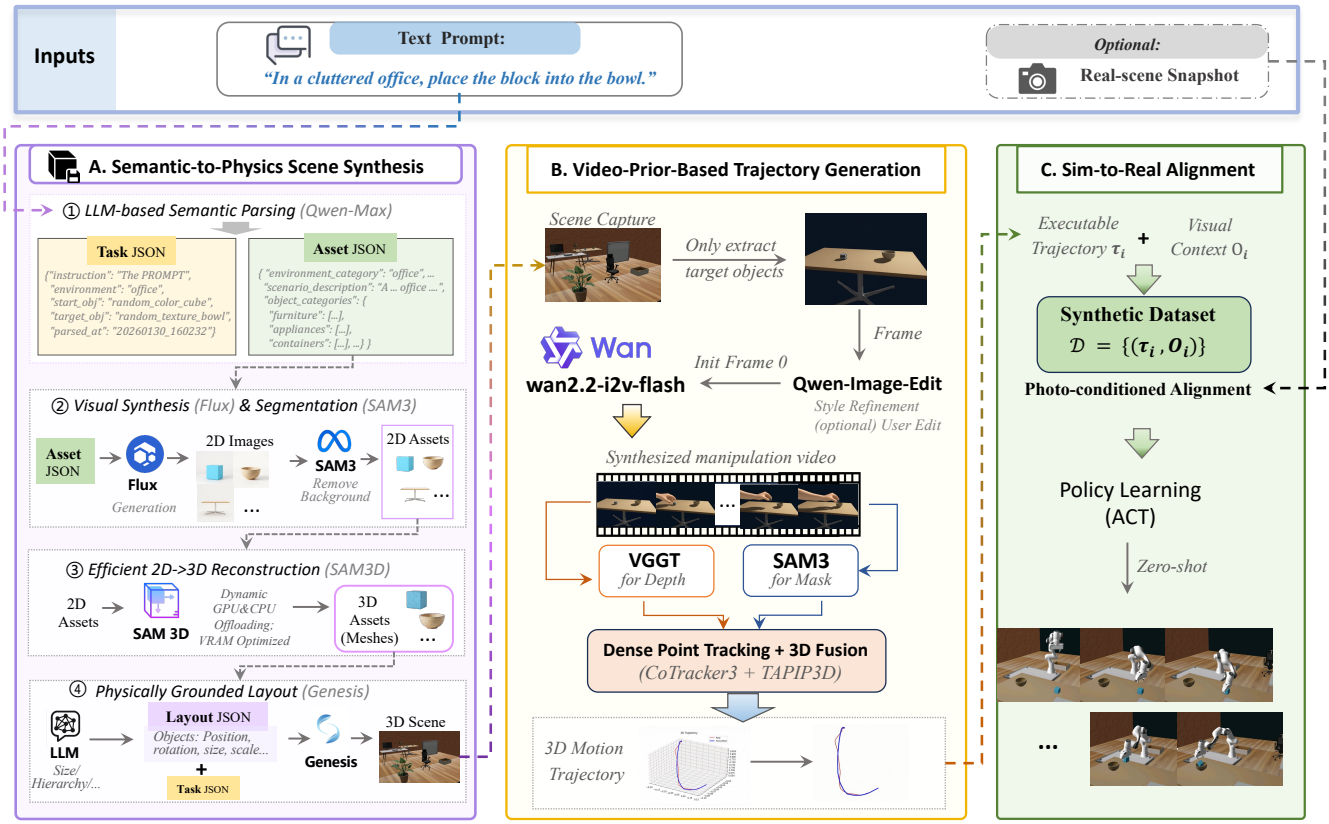


Fig. 2. **Overview of the V-Dreamer Pipeline.** The framework consists of three integrated stages: (1) Semantic-to-Physics Scene Synthesis. An LLM-based semantic planner decomposes prompts into asset manifests, which are then transformed into 3D meshes via generative visual synthesis and memory-efficient reconstruction before being assembled into a physics-validated layout. (2) Video-Prior-Based Trajectory Generation. Using the stabilized scene as a prior, we employ video foundation models with targeted negative prompting to dream up physically plausible manipulation motions. (3) Sim-to-Gen Alignment. Actionable 3D trajectories are extracted from the 2D pixel domain through mask-restricted tracking, metric depth estimation, and 3D motion lifting, ultimately mapping visual dreams into executable robot end-effector poses.

### C. Sim-to-Real Transfer and Generalization

Bridging the reality gap between simulation and the physical world remains a central challenge in robot learning. Domain Randomization (DR) [25], [26] randomizes physical and visual parameters to expose policies to diverse conditions, but requires manual tuning of randomization ranges and struggles to replicate real-world visual complexity. Real-to-Sim methods [27], [28] reconstruct specific real scenes in simulation, but are labor-intensive and tied to particular environments. TRANSIC [29] refines simulation-trained policies via online human corrections, while DrEureka [30] uses LLMs to automate randomization range design. V-Dreamer sidesteps these limitations by synthesizing a large variety of physically grounded scenes through generative diffusion models, naturally providing open-vocabulary visual diversity without manual specification or real-world interaction, and enabling zero-shot sim-to-real transfer to unseen objects.

## III. Methodology

**V-Dreamer** is a fully automated framework designed to synthesize large-scale, physically grounded manipulation datasets for training generalizable robot policies. The core challenge in automated data generation is bridging the substantial representational gap between abstract human intentions and executable physical simulations. As illustrated in Figure 2, V-Dreamer overcomes this challenge through a cohesive pipeline comprising three core modules: (1) *Semantic-to-Physics Scene Synthesis*, which translates natural language into interactive 3D environments; (2) *Video-Prior-Based Trajectory Generation*, which harnesses video foundation models to simulate expert manipulation behaviors; and (3) *Sim-to-Gen Alignment*, which extracts actionable 3D trajectories from 2D visual predictions.

### A. Semantic-to-Physics Scene Synthesis

To transform an open-ended semantic instruction into a physically interactive rigid-body simulation, we propose a multi-stage procedural generation pipeline including: ① Semantic Parsing and Asset Manifest, ② Generative Visual Synthesis, ③ Memory-Efficient 3D Reconstruction and ④ Physically Grounded Layout. Specifically:

① *Semantic Parsing and Asset Manifest.* Given a user prompt, we first employ a Large Language Model (Qwen-Max[31] in our experiment) acting as a high-level semantic planner. The LLM decomposes the abstract scene description into a structured *Asset Manifest* in JSON format. This manifest explicitly specifies object categories and stylistic elements. To ensure the downstream reconstruction quality, we inject specific negative constraints into the LLM prompt. For

example, we prohibit highly reflective or transparent surfaces like mirrors and glass in a cyberpunk workspace, thereby enforcing style consistency and avoiding rendering artifacts across all generated assets. Note that this part only performs object-level manifest generation without background and layouts.

② *Generative Visual Synthesis.* For each manifest entry, we perform visual synthesis to obtain high-fidelity 2D image assets using the Flux diffusion model [32]. To reduce environmental bias and avoid baked-in shadows which are critical for clean downstream 3D generation, we augment prompts with modifiers such as "studio lighting" and "solid white background". We then apply Segment Anything Model 3 (SAM3) [33] for precise background removal. For robustness, we adopt a hybrid segmentation scheme consisting of a semantic-first pathway that uses object names as prompts, and a vision-based fallback [34] when semantic prompting produces suboptimal masks. With these clean, segmentation-ready 2D assets in place, the next step is to lift them into detailed 3D meshes for simulation.

③ *Memory-Efficient 3D Reconstruction.* Lifting high-resolution 2D assets to 3D meshes is often the primary computational bottleneck in an automated pipeline. To address this, we introduce a memory-efficient reconstruction pipeline. This module dynamically offloads encoder and decoder sub-networks between the CPU and GPU during inference. By strategically trading computation time for VRAM optimization, this mechanism enables the reconstruction of complex, high-poly geometries on standard consumer-grade hardware, outputting raw 3D meshes ready for physical instantiation.

④ *Physically Grounded Layout.* After generating isolated 3D object meshes, the final step is to assemble them into a coherent and physically valid scene. We combine LLM spatial reasoning with rigorous physics-based validation. The LLM first infers approximate real-world metric dimensions for each object, after which we apply simple heuristics to correct distortions introduced during 3D lifting. Next, we generate the scene layout as a JSON specification using a foundation model, and then place each object at its corresponding 3D position accordingly. By this means, we construct a semantic scene graph that hierarchically categorizes objects into anchors and define the environment structure and accessories that serve as manipulable targets. Finally, we invoke a physics engine [35] to enforce constraint-based layout via AABB collision checking and gravity alignment. Anchors are snapped to the ground plane ($Z = 0$), while accessories are placed on supporting surfaces $Z = H_{\text{parent}} + \delta$ with procedural spatial jittering to systematically resolve bounding-box overlaps.

### B. Video-Prior-Based Trajectory Generation

Traditional motion planners often require task-specific engineering and may not directly benefit from the rich visual commonsense priors needed for open-vocabulary manipulation. In contrast, V-Dreamer exploits the rich, implicit motion priors encapsulated within video foundation models to synthesize diverse, context-aware manipulation trajectories.

Directly prompting a video model from a procedurally generated layout can lead to severe physical inconsistencies. Therefore, we first execute a brief dynamic settling phase within the Genesis [35] simulator to resolve remaining micro-penetrations between objects. We capture this stabilized scene as the initial frame ($I_0$). We then perform style refinement on $I_0$ using Qwen-Image-Edit [36] to obtain a *Style-Refined Initial Frame* $I'_0$. This step is used in two ways: (i) The system automatically applies a set of predefined style perturbations to increase visual diversity, such as adding stripes or modifying material appearance. (ii) It provides a human-in-the-loop interface for targeted, user-specified stylization when customization is desired.

Building upon the stabilized frame, we feed $I'_0$ into the Wan2.2-I2V-Flash [37] video generation model to synthesize a continuous temporal sequence of the specified manipulation task. To encourage physically plausible motion for robotic execution, we apply targeted negative prompts to suppress camera motion, object deformation, and background flicker. These linguistic constraints force the diffusion model to focus exclusively on the rigid-body kinematics of the manipulator and the target object, effectively suppressing the dream-like morphing artifacts that typically violate physical feasibility.

Since the synthesized video resides purely in the 2D pixel domain, the final step is to extract actionable 3D controls. We follow TraceGen [38] and design a tracking-based grounding module. We first re-run SAM3 on the visually processed initial frame to obtain a high-quality mask of the target object. We then estimate per-frame depth using Visual Geometry Grounded Transformer (VGGT) [11], which provides metric cues for 3D lifting. Next, we uniformly sample 2D points within the object mask and track them across frames using CoTracker3 [10], yielding dense but object-consistent 2D trajectories. These tracked points are subsequently fused into a coherent 3D motion trajectory using the TAPIP3D [39] module. Compared to TraceGen, which densely tracks pixels over the entire frame and follows moving regions, our mask-restricted tracking focuses on object-consistent motion and is empirically more precise.

Finally, we apply lightweight post-processing to remove outlier tracks and suppress temporal jitter. Subsequently, we utilize Graspgen [40] to generate a feasible robotic grasp pose for the target object, and then map the resulting 3D motion to the robot's end-effector trajectory via inverse kinematics (IK). This yields a physically executable trajectory $\tau = \{(p_t, q_t)\}_{t=0}^{T}$, where $p_t$ denotes the continuous end-effector pose and $q_t$ represents the discrete gripper state.

### C. Sim-to-Real Alignment

By continuously running the V-Dreamer pipeline, we automatically generate a large-scale, high-fidelity synthetic dataset $\mathscr{D} = \{(\tau_i, O_i)\}_{i=1}^{N}$, where each expert trajectory $\tau_i$ is paired with a diverse semantic and visual context $O_i$. For long instructions, the foundation model decomposes the instruction into several sub-tasks, and then iteratively

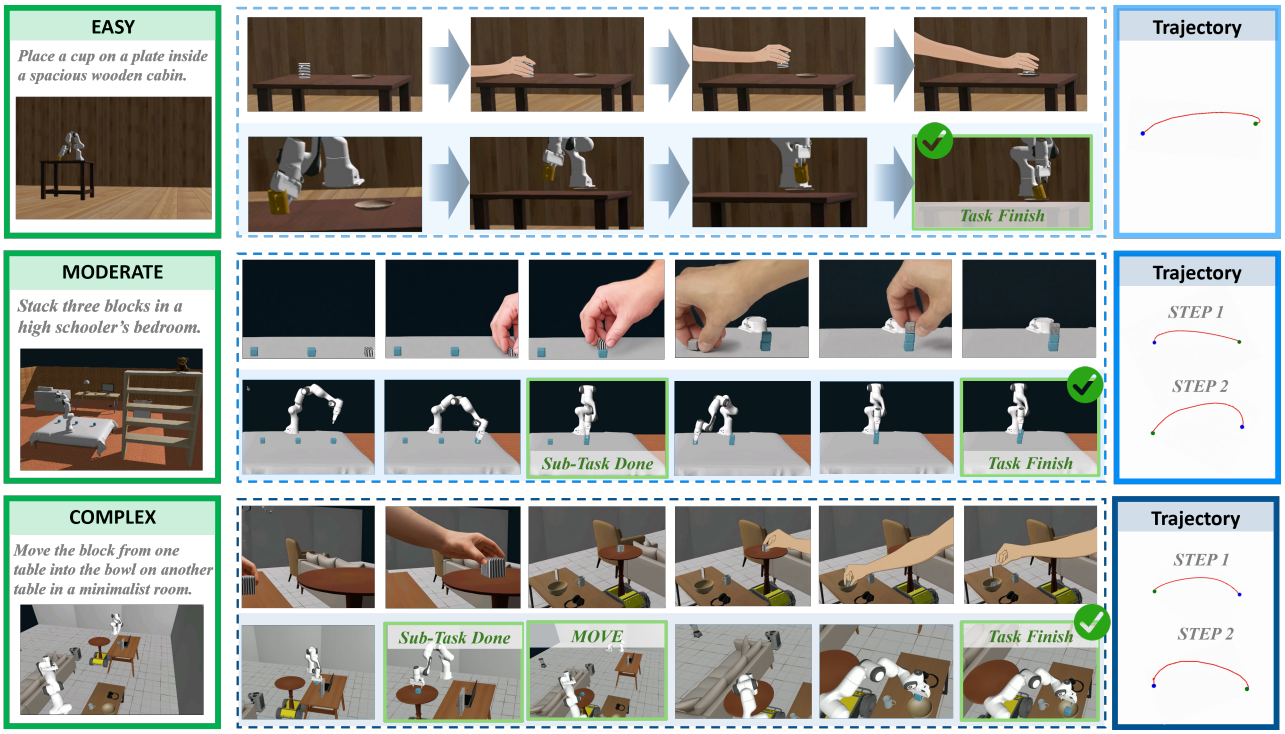


Fig. 3. Qualitative examples of V-Dreamer synthesized tasks of increasing difficulty. Left: Input Prompts and generated 3D scenes. Middle: Generated manipulation videos and the corresponding robot manipulation results. Right: Extracted end-effector trajectories for demonstrations.

performs video generation, trajectory generation and physical operation for each sub-task in simulation.

For physical deployment, we follow a strict sim-to-real transfer protocol. Specifically, we capture a snapshot of the real target scene and use it as an additional alignment input to instantiate a scene-matched, style-refined initial frame $I'_0$ in simulation. We also calibrate the simulator to match the real camera parameters and the robot workspace bounds. Conditioned on $I'_0$, all subsequent trajectory generation and policy training are performed entirely in simulation without any human demonstrations or target-domain fine-tuning. The policy trained solely on $\mathscr{D}$ is then deployed **zero-shot** on physical hardware for evaluation.

For safety, we optionally reconstruct coarse bounding geometries of novel objects in simulation for background collision checking, while the core manipulation behavior and closed-loop control are driven entirely by the generalized representations learned from V-Dreamer's synthesized data.

## IV. Experiments

Experiments are designed to assess **V-Dreamer** as a robust data-generation engine by answering three core questions.

**Q1. Scalability and Data Quality.** Can V-Dreamer synthesize large-scale, high-quality demonstrations for downstream policy learning?

**Q2. Semantic and Geometric Diversity.** Can V-Dreamer generate sufficiently diverse object instances and spatial layouts to support zero-shot generalization to novel geometries?

**Q3. Physical Grounding and Executability.** Are V-Dreamer's trajectories physically plausible and kinematically feasible, enabling zero-shot transfer on real hardware?

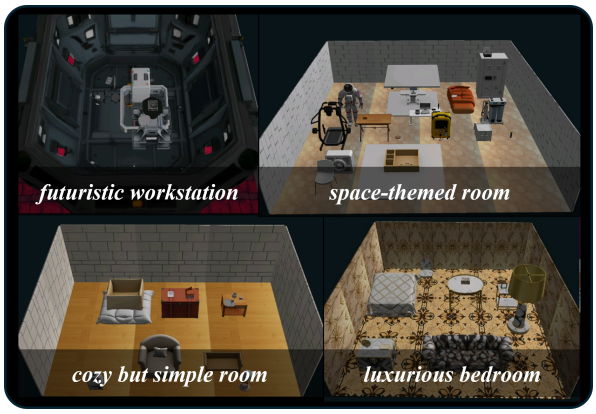


Fig. 4. Qualitative examples of V-Dreamer synthesized scenes with varying object instances, textures, and layouts.

To the best of our knowledge, V-Dreamer is the first automated system-level framework that unifies open-vocabulary scene synthesis, video-prior-based trajectory generation, and real-world deployment. As prior works usually address only a single stage of this pipeline, there is no directly comparable end-to-end baseline. We therefore evaluate V-Dreamer from a system perspective, focusing on scene diversity, trajectory executability, policy learning utility, and sim-to-real transfer.

### A. Experimental Setup

*a) Tasks and Dataset:* We systematically evaluate **V-Dreamer** on a fundamental yet challenging *tabletop manipulation* task, where the robot grasps a target object from a planar surface and places it onto a receptacle. To

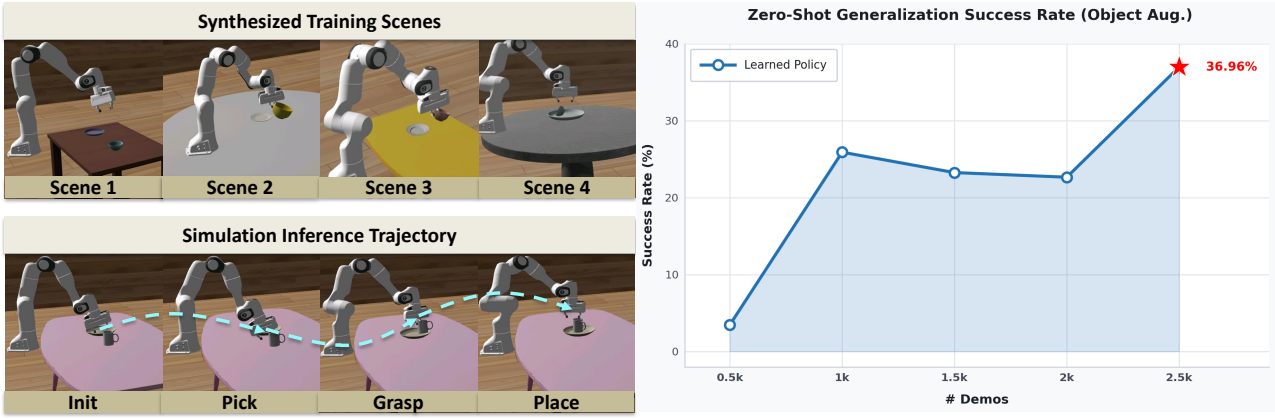


Fig. 5. **Simulation Evaluation.** *(Left top)* Representative synthesized training scenes with object augmentation, illustrating diverse spatial layouts used to generate demonstrations. *(Left bottom)* Zero-shot inference trajectory of the learned policy on an unseen mug, showing the manipulation sequence (Init $\rightarrow$ Pick $\rightarrow$ Grasp $\rightarrow$ Place). *(Right)* Success rate of the learned policy on 10 unseen mug instances under different training dataset sizes (500 trials each). Performance improves with increasing data scale and peaks at 2.5k synthesized demonstrations.

support Q2, our pipeline procedurally synthesizes a diverse training corpus spanning (i) **Objects**: 40 target objects and 40 receptacle objects with substantial variation in geometry, texture, and metric scale; (ii) **Environments**: 20 distinct table settings with randomized surface materials, such as wood, marble, and glass. To evaluate zero-shot generalization, we additionally curate a held-out test set of 10 novel mugs with entirely unseen shapes, handles, and textures, which are strictly excluded from the generative pipeline.

*b) Base Policy Model:* To evaluate whether data synthesized by V-Dreamer can effectively support robot policy learning, we adopt the classical imitation learning framework Action Chunking with Transformers (ACT) [2] as the downstream policy model. ACT is a widely used behavior cloning framework for robotic manipulation, providing a stable and reproducible training pipeline. Moreover, its sensitivity to demonstration quality and data distribution coverage makes it well suited for assessing the diversity and utility of synthesized data. Finally, since ACT does not involve exploration or reinforcement learning, performance differences can be more directly attributed to the training data itself.

*c) Evaluation Metrics:* We use **Success Rate (SR)** as the evaluation metric, defined as the percentage of trials in which the robot successfully grasps and places the object within the episode horizon. Results are averaged over randomized trials across object instances and initial conditions, using the same metric for both simulation and real-world experiments.

*d) Hardware Configuration:* In simulation, we use the 7-DoF Franka robotic arm. For real-world deployment, we evaluate the trained policies on the left arm of a Cobot-Magic platform, equipped with a 6-DoF Piper lightweight manipulator and a standard parallel-jaw gripper. For RGB perception, the system uses two Orbbec DaBai cameras: one mounted on the wrist and the other providing a top-down view. All synthetic training data are generated on a workstation equipped with 8$\times$RTX 4090 GPUs.

### B. Simulation Results

*a) Qualitative Visualization.:* To address **Q1**, we first visualize V-Dreamer's synthesis capability. Figure 3 further shows that V-Dreamer can generate coherent demonstrations across increasing levels of interaction difficulty, from *simple* motions to *hard* and *complex* cooperative sequential mobile manipulation task by three devices by presenting both the synthesized videos and the recovered end-effector trajectories. These examples indicate that V-Dreamer produces structured, executable scene-action pairs rather than merely randomizing assets. As a supplement, we show more generated scenes with substantial variation in texture, metric scale, and spatial layout in Figure 4.

*b) Data Scaling and Zero-Shot Generalization:* We further answer **Q2** and evaluate **V-Dreamer** as a data-generation engine along two axes, including (i) whether adding more synthesized demonstrations improves downstream learning, and (ii) whether the learned policy transfers to geometrically unseen objects. Concretely, we train the same diffusion-based imitation policy on subsets of the synthesized dataset ranging from 500 to 2,500 trajectories, and evaluate each variant on a held-out set of 10 novel mugs, conducting 50 trials per object with randomized initial configurations.

The left side of Figure 5 presents qualitative results of execution trajectories in simulation, where the policy is trained with synthesized data and evaluated on unseen mug instances. The quantitative results, summarized in the right side of Figure 5, reveal a non-linear correlation between generative data scale and policy success rate. The policy exhibits catastrophic failure when trained on only 500 trajectories (3.46%), indicating an insufficient coverage of the visual-spatial state distribution. Performance leaps significantly at 1,000 trajectories (25.90%) and peaks optimally at **2,500 trajectories** with a success rate of **36.96%**. Notably, this peak is achieved using purely synthesized demonstrations on entirely unseen mug geometries, highlighting V-Dreamer's ability to generate diverse, actionable training data. These results empirically confirm that V-Dreamer provides high-

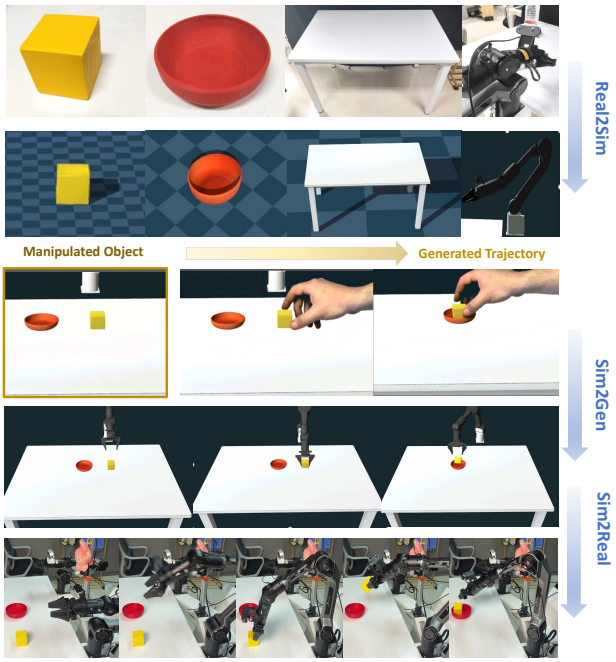


Fig. 6. **V-Dreamer enables a Real2Sim–Sim2Gen–Sim2Real pipeline for real-world robot manipulation.** Real-world observations are first converted into simulation environments (Real2Sim). V-Dreamer then synthesizes executable manipulation trajectories in simulation (Sim2Gen) for policy training. Finally, the learned policy is deployed on a physical robot arm (Sim2Real), successfully executing a pick-and-place task that transfers a novel yellow cube to the red target without any real-world fine-tuning.

quality data for robust manipulation primitives, and given that current imitation learning methods rarely exceed 30% success on unseen tasks [41], this further demonstrates its effectiveness for imitation learning.

### C. Real-World Experiments

To evaluate the extreme data efficiency and practical utility of **V-Dreamer**, we deploy the learned policy on a physical 6-DoF Piper robotic arm (left arm of an ALOHA setup). Unlike conventional methods requiring large-scale real demonstrations, we follow a strict sim-to-real protocol. The policy is trained on a single **one-shot** V-Dreamer trajectory in simulation, generated for a simple scene with a block, a bowl, and a table, and then transferred to the real world in a **zero-shot** manner.

We adopt a minimalist zero-shot sim-to-real setup for physical deployment. For low-cost domain alignment, we use the Piper URDF from Robot-Twin, refine joint limits, render the simulated arm links (except the base) in black, and apply black tape to the corresponding real links within the camera view to reduce robot-body appearance mismatch. The policy runs closed-loop using observations from an Orbbec DaBai RGB-D camera, while we intentionally feed only the raw RGB image as $O_t$ and discard depth and extra calibration to test robustness. The full pipeline is illustrated in Figure 6.

We evaluate four perturbation conditions on a one-shot policy trained using a single V-Dreamer-synthesized trajectory, focusing on Success Rate (SR). As shown in Table I, the conditions include: visual distractors on the table (50% success), out-of-distribution target objects such as apple, tape roll, mango, pear, and bottle (20% success), spatial perturbations of both target and receptacle poses (15% success), and camera occlusion during execution (0% success). The occlusion result confirms true closed-loop visual servoing rather than open-loop trajectory replay, while the other conditions highlight non-trivial robustness and limited generalization under extreme one-shot supervision.

TABLE I
ONE-SHOT SIM-TO-REAL ROBUSTNESS EVALUATION

| Perturbation Type | Real-World Setup | SR |
|---|---|---|
| Visual Distractors | Unmodeled background clutter in FOV | 50% |
| Object Generalization | Apple, tape, mango, pear, bottle | 20% |
| Spatial Perturbation | Shifted initial object and goal positions | 15% |
| Sensor Occlusion | Blocked Orbbec DaBai camera view | 0% |

All tests used a policy trained on a single synthetic trajectory. The 0% success rate under occlusion confirms reliance on closed-loop visual feedback rather than open-loop playback.

These real-world results primarily validate **V-Dreamer** as a data synthesis engine, rather than an algorithmic tweak aimed at maximizing task accuracy. Under a strict zero-shot sim-to-real protocol, a policy trained on a *single* V-Dreamer-synthesized demonstration can execute meaningful behaviors on physical hardware. This indicates that the generated trajectories are physically grounded, kinematically feasible, and directly usable for control, providing direct evidence for **Q3**. V-Dreamer distills generative video priors into executable demonstrations, bridging the reality gap even in an extreme low-data regime.

## V. CONCLUSION

In this work, we presented **V-Dreamer**, a fully automated framework for synthesizing large-scale, physically grounded simulation environments and executable manipulation trajectories for generalist robot learning. By integrating Large Language Models with advanced 2D/3D and video generative models, V-Dreamer translates open-vocabulary instructions into diverse interactive scenes and coherent manipulation behaviors. Unlike prior generative pipelines that primarily focus on scene synthesis, V-Dreamer tackles the full generation loop, producing training-ready data that is directly usable for policy learning.

Extensive experiments show that imitation policies trained exclusively on V-Dreamer-generated data achieve robust performance and zero-shot generalization to geometrically unseen objects in simulation. Furthermore, with photo-conditioned scene alignment, these policies transfer zero-shot to real hardware. These results suggest that the trajectories generated by V-Dreamer are not merely visually plausible, but physically executable and directly useful for robot control, highlighting the promise of full-cycle generative pipelines in alleviating the data bottleneck in robotic manipulation.

**Limitations and Future Work.** An important next step is to develop automated, physics-aware trajectory filtering to reduce the impact of low-quality generations. In addition, the

current framework is limited to rigid-body tabletop manipulation. Extending it to articulated and deformable objects is a key direction for future work. Finally, incorporating tighter physical feedback during generation and grounding may further improve physical realism and sim-to-real robustness.

## REFERENCES


[1] C. Chi, S. Feng, Y. Du, Z. Xu, E. Cousineau, B. Burchfiel, and S. Song, "Diffusion policy: Visuomotor policy learning via action diffusion," in *Robotics: Science and Systems (RSS)*, 2023.

[2] T. Z. Zhao, V. Kumar, S. Levine, and C. Finn, "Learning fine-grained bimanual manipulation with low-cost hardware," in *Robotics: Science and Systems (RSS)*, 2023.

[3] A. Brohan, *et al.*, "RT-1: Robotics transformer for real-world control at scale," *arXiv preprint arXiv:2212.06817*, 2022.

[4] A. Padalkar, *et al.*, "Open X-embodiment: Robotic learning datasets and RT-X models," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2024.

[5] S. Yang, O. Nachum, Y. Du, J. Wei, P. Abbeel, and D. Schuurmans, "Foundation models for decision making: Problems, methods, and opportunities," *arXiv preprint arXiv:2303.04129*, 2023.

[6] G. Wang, Y. Xie, Y. Jiang, A. Mandlekar, C. Xiao, Y. Zhu, L. Fan, and A. Anandkumar, "Voyager: An open-ended embodied agent with large language models," in *NeurIPS*, 2023.

[7] M. Deitke, E. VanderBilt, A. Herrasti, L. Weihs, K. Ehsani, J. Salvador, W. Han, E. Kolve, A. Kembhavi, and A. Farhadi, "ProcTHOR: Large-scale embodied AI using procedural generation," in *NeurIPS*, 2022.

[8] M. Ahn, *et al.*, "Do as I can, not as I say: Grounding language in robotic affordances," in *Conference on Robot Learning (CoRL)*, 2022.

[9] J. Z. Wu, Y. Ge, X. Wang, S. W. Lei, Y. Gu, Y. Shi, W. Hsu, Y. Shan, X. Qie, and M. Z. Shou, "Tune-A-video: One-shot tuning of image diffusion models for text-to-video generation," in *ICCV*, 2023.

[10] N. Karaev, I. Rocco, B. Graham, N. Neverova, A. Vedaldi, and C. Rupprecht, "CoTracker: It is better to track together," *arXiv preprint arXiv:2307.07635*, 2023.

[11] J. Wang, M. Chen, N. Karaev, A. Vedaldi, C. Rupprecht, and D. Novotny, "VGGT: Visual geometry grounded deep structured tokens," in *CVPR*, 2025.

[12] E. Todorov, T. Erez, and Y. Tassa, "MuJoCo: A physics engine for model-based control," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2012, pp. 5026–5033.

[13] V. Makoviychuk, *et al.*, "Isaac Gym: High performance GPU-based physics simulation for robot learning," in *NeurIPS Datasets and Benchmarks Track*, 2021.

[14] A. Szot, *et al.*, "Habitat 2.0: Training home assistants to rearrange their habitat," in *NeurIPS*, 2021.

[15] L. Wang, Y. Ling, Z. Yuan, M. Shridhar, C. Bao, Y. Qin, W. Wang, H. Xu, and X. Wang, "GenSim: Generating robotic simulation tasks via large language models," in *ICLR*, 2024.

[16] Y. Wang, Z. Xian, F. Chen, T.-H. Wang, Y. Wang, K. Fragkiadaki, Z. Erickson, D. Held, and C. Gan, "RoboGen: Towards unleashing infinite data for automated robot learning via generative simulation," in *ICML*, 2024.

[17] Y. Yang, *et al.*, "Holodeck: Language guided generation of 3D embodied AI environments," in *CVPR*, 2024.

[18] A. Mandlekar, S. Nasiriany, B. Wen, I. Akinola, Y. Narang, L. Fan, Y. Zhu, and D. Fox, "MimicGen: A data generation system for scalable robot learning using human demonstrations," in *Conference on Robot Learning (CoRL)*, 2023.

[19] A. Blattmann, *et al.*, "Stable video diffusion: Scaling latent video diffusion models to large datasets," *arXiv preprint arXiv:2311.15127*, 2023.

[20] M. Yang, Y. Du, K. Ghasemipour, J. Tompson, L. Kaelbling, D. Schuurmans, and P. Abbeel, "Learning interactive real-world simulators," in *ICLR*, 2024.

[21] Y. Ding, J. Liu, Z. Hu, and C. Gan, "Video language planning," in *ICLR*, 2024.

[22] Y. Du, M. Yang, B. Dai, H. Dai, O. Nachum, J. B. Tenenbaum, P. Abbeel, and D. Schuurmans, "Learning universal policies via text-guided video generation," in *NeurIPS*, 2023.

[23] S. Gu, C. Yuan, Z. Xue, Y. Wang, S. Huang, Y. Ma, and P. Luo, "GR-2: Generalizing robot policies via video prediction," *arXiv preprint arXiv:2412.09726*, 2024.

[24] J. Bruce, *et al.*, "Genie: Generative interactive environments," in *ICML*, 2024.

[25] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel, "Domain randomization for transferring deep neural networks from simulation to the real world," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2017.

[26] X. B. Peng, M. Andrychowicz, W. Zaremba, and P. Abbeel, "Sim-to-real transfer of robotic control with dynamics randomization," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2018.

[27] H. Zhang, J. Li, Y. Liu, G. Chen, X. Zhu, and F. Xia, "Real-to-sim: Learning real-world robot policies by dreaming in simulation," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2023.

[28] M. Torne, A. Simeonov, Z. Li, A. Chan, T. Chen, A. Gupta, and P. Agrawal, "Reconciling reality through simulation: A real-to-sim-to-real approach for robust manipulation," in *Conference on Robot Learning (CoRL)*, 2024.

[29] Y. Jiang, C. Wang, R. Zhang, J. Wu, and L. Fei-Fei, "TRANSIC: Sim-to-real policy transfer by learning from online correction," in *Conference on Robot Learning (CoRL)*, 2023.

[30] Y. J. Ma, W. Liang, H.-J. Wang, S. Wang, Y. Zhu, L. Fan, O. Bastani, and D. Jayaraman, "DrEureka: Language model guided sim-to-real transfer," in *Robotics: Science and Systems (RSS)*, 2024.

[31] A. Yang, *et al.*, "Qwen2.5 technical report," *arXiv preprint arXiv:2412.15115*, 2024.

[32] Black Forest Labs, "FLUX," https://github.com/black-forest-labs/flux, 2024.

[33] N. Carion, *et al.*, "SAM 3: Segment anything with concepts," 2025. [Online]. Available: https://arxiv.org/abs/2511.16719

[34] D. Gatis, "rembg: Rembg is a tool to remove images background," https://github.com/danielgatis/rembg, 2023, accessed: 2025.

[35] Genesis Authors, "Genesis: A generative and universal physics engine for robotics and beyond," https://github.com/Genesis-Embodied-AI/Genesis, December 2024.

[36] C. Wu, *et al.*, "Qwen-image technical report," 2025. [Online]. Available: https://arxiv.org/abs/2508.02324

[37] Team Wan, *et al.*, "Wan: Open and advanced large-scale video generative models," *arXiv preprint arXiv:2503.20314*, 2025.

[38] S. Lee, *et al.*, "Tracegen: World modeling in 3d trace space enables learning from cross-embodiment videos," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2026.

[39] B. Zhang, L. Ke, A. W. Harley, and K. Fragkiadaki, "Tapip3d: Tracking any point in persistent 3d geometry," *arXiv preprint arXiv:2504.14717*, 2025.

[40] A. Murali, B. Sundaralingam, Y.-W. Chao, W. Yuan, J. Yamada, M. Carlson, F. Ramos, S. Birchfield, D. Fox, and C. Eppner, "Graspgen: A diffusion-based framework for 6-dof grasping with on-generator training," *arXiv preprint arXiv:2507.13097*, 2025.

[41] Y. Chen, Z. Chen, N. T. Chan, J. Chen, J. Yin, J. Shi, Y. Gao, Y.-L. Li, and J. Huo, "Robohiman: A hierarchical evaluation paradigm for compositional generalization in long-horizon manipulation," *arXiv preprint arXiv:2510.13149*, 2025.