

GWM: Towards Scalable Gaussian World Models for Robotic Manipulation

Guanxing Lu^{1,2*}, Baoxiong Jia^{2*}, Puhao Li^{2*}, Yixin Chen²

Ziwei Wang³, Yansong Tang^{1†}, Siyuan Huang^{2†}

*Equal contributor †Corresponding author

¹ Tsinghua University, ² State Key Lab of General AI, Beijing Institute for General AI

³ School of Electrical and Electronic Engineering, Nanyang Technological University

[gaussian-world-model.github.io](https://github.com/gaussian-world-model)

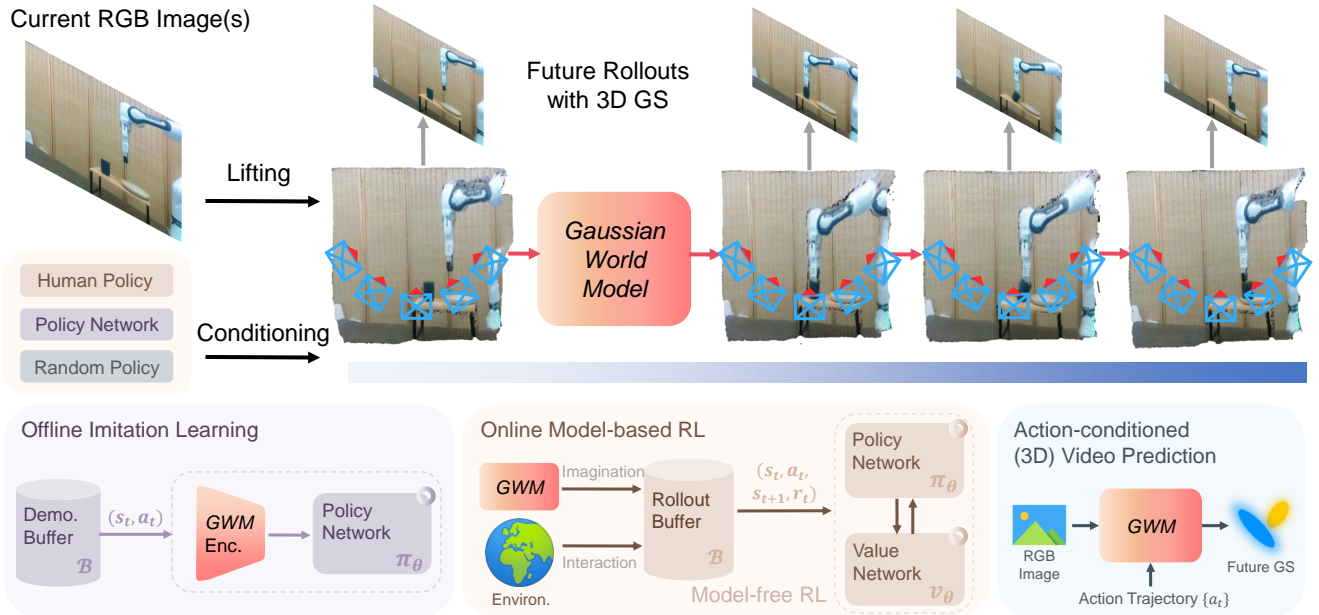


Figure 1. **Gaussian World Model (GWM)** is a novel branch of world model that predicts dynamic future states and enables robotic manipulation based on the 3D Gaussian Splatting representation. It facilitates action-conditioned 3D video prediction, enhances visual representation learning for imitation learning, and serves as a robust neural simulator for model-based reinforcement learning.

Abstract

Training robot policies within a learned world model is trending due to the inefficiency of real-world interactions. The established image-based world models and policies have shown prior success, but lack robust geometric information that requires consistent spatial and physical understanding of the three-dimensional world, even pre-trained on internet-scale video sources. To this end, we propose a novel branch of world model named Gaussian World Model (GWM) for robotic manipulation, which reconstructs the future state by inferring the propagation of

Gaussian primitives under the effect of robot actions. At its core is a latent Diffusion Transformer (DiT) combined with a 3D variational autoencoder, enabling fine-grained scene-level future state reconstruction with Gaussian Splatting. GWM can not only enhance the visual representation for imitation learning agent by self-supervised future prediction training, but can serve as a neural simulator that supports model-based reinforcement learning. Both simulated and real-world experiments depict that GWM can precisely predict future scenes conditioned on diverse robot actions, and can be further utilized to train policies that outperform the state-of-the-art by impressive margins, showcasing the

1. Introduction

Humans construct predictive *world models* from limited sensory input, allowing them to anticipate future outcomes and adapt to new situations [11, 17]. Inspired by this capability, world model learning has driven major advances in intelligent agents, enabling strong performance in autonomous driving [13, 25, 26, 76, 94] and gaming [1, 17–21, 62, 85]. As intelligent agents increasingly engage with the physical world, advancing world model learning for robotic manipulation becomes an essential research direction, as it could ideally empower robots to reason about interactions, predict physical dynamics, and adapt to diverse unseen environments.

This naturally raises the following question: *How to effectively represent, construct, and leverage the world model to enhance robotic manipulation?* Such demand poses significant challenges to existing representations and models.

- **Necessity of 3D Representation** High-capacity architectures [24, 73] and internet-scale pre-training have established video-based generative models as powerful tools for capturing world dynamics information, which significantly enhances policy learning [78, 83]. However, their reliance on image inputs makes them susceptible to unseen visual variations (*e.g.*, lighting, camera pose, textures, *etc.*) [39], as they lack 3D geometric and spatial understanding. While RGB-D and multi-view [15, 16] setups attempt to mitigate this gap, implicitly aligning image patch features within a coherent 3D space remains challenging [58, 96], leaving the robustness concern unresolved. This *underscores the need for representations that integrate fine visual details with 3D spatial information* to enhance world modeling for robotic manipulation.
- **Efficiency and Scalability** To identify a 3D representation that preserves both 3D geometric structure and fine visual details from 2D images, multi-view 3D reconstruction methods such as Neural Radiance Field (NeRF) [53] and 3D Gaussian Splatting (3D-GS) [34] offer natural solutions. Among them, 3D-GS is particularly appealing due to its explicit per-Gaussian modeling of 3D scenes, marrying efficient 3D representations like point clouds with high-fidelity rendering. However, since these methods primarily rely on offline per-scene reconstruction, their computational demands pose significant challenges [47, 87] on applying them in robotic manipulation, especially for Model-based Reinforcement Learning (MBRL), limiting their scalability.

To this end, we propose **Gaussian World Model (GWM)**, a novel 3D world model that integrates 3D-GS with high-capacity generative models for robotic manipulation. Specifically, our approach combines recent advance-

ments in feed-forward 3D-GS reconstruction with Diffusion Transformers (DiTs), enabling fine-grained future scene reconstruction through Gaussian rendering conditioned on current observations and robot actions. To achieve real-time training and inference, we design a 3D Gaussian Variational Autoencoder (VAE) to extract latent representations from 3D Gaussians, allowing the diffusion-based world model to operate efficiently in a compact latent space. With this novel design, we demonstrate that GWM enhances visual representation learning, improving its role as a visual encoder for imitation learning while also serving as a robust neural simulator for model-based Reinforcement Learning (RL).

To comprehensively evaluate GWM, we conduct extensive experiments in action-conditioned video prediction, imitation learning, and model-based RL settings, covering 31 diverse robotic tasks across 3 domains. For real-world evaluation, we introduce a Franka PnP task suite with 20 variations, encompassing both in-domain and out-of-domain settings. For the ablation study, we evaluate both perceptual metrics and success rates to verify the effectiveness of each building blocks. GWM consistently outperforms previous baselines, including state-of-the-art image-based world models, offering notable advantages and highlighting its data-scaling potential.

In summary, our main contributions are threefold.

- We introduce GWM, a novel 3D world model that is instantiated with a Gaussian diffusion transformer and a Gaussian VAE for efficient dynamic modeling. GWM learns to predict accurate future states and dynamics in a scalable end-to-end manner without human intervention.
- GWM can be easily integrated into offline imitation learning and online reinforcement learning with superior efficiency, depicting impressive scaling potential in learning-based robotic manipulation.
- We demonstrate the efficacy of GWM through extensive experiments in two challenging simulation environments, which improves the previous state-of-the-art baselines by a large margin of 16.25%. Furthermore, we validate its practicality in real-world scenarios, where GWM improves a typical diffusion policy by 30% with 20 trials.

2. Related Work

World Models World models capture scene dynamics and enable efficient learning by predicting future states based on current observations and actions. They have been widely explored in autonomous driving [13, 25, 26, 76, 94, 98], game agents [1, 17–21, 62, 85], and robotic manipulation [22, 63, 79]. Early works [17–22, 52, 61–63, 85, 92] learn a latent space for future prediction, achieving strong results in both simulated and real-world settings [79]. However, while simplifying modeling, latent representations struggle to capture the world’s fine details. Recent advances in diffusion models [23, 67, 68] and transformers [60, 73]

have shifted world modeling toward direct pixel-space modeling [1, 83], enabling fine-grained detail capture and large scale learning from internet videos. Yet, image-based models often lack physical commonsense [4], thus limiting their applicability in robotic manipulation.

Gaussian Splatting 3D-GS [34] represents scenes using 3D Gaussians, which are efficiently projected onto 2D planes via differentiable splitting. Compared to implicit representations like NeRF [53], it offers greater efficiency, benefiting applications such as invasive surgery [45], SLAM [33], and autonomous driving [95]. This advantage extends to 4D dynamic modeling [27, 48, 80] as the 3D Gaussians, similar to point clouds, are spatially meaningful. However, the offline per-scene reconstruction required by these methods imposes computational challenges for real-time applications like robotic manipulation. Recent works [6, 12, 70, 81, 82, 90, 93, 97] address this issue by learning generative mappings from pixels to Gaussians using large-scale datasets, but still rely on known camera poses, restricting scalability. A parallel effort [36, 66, 75] explores feed-forward novel-view synthesis from unposed images, leveraging the predicted point map as a proxy for explicit multi-view alignment. Building on these advances, this work develops a scalable Gaussian world model from unposed images, ensuring spatial awareness and scalability for policy training.

Visual Manipulation Building vision-driven robots with human-like capabilities is a long-standing challenge. Visual imitation learning methods [5, 35, 38, 44, 71] mimic expert demonstrations using various visual representations, such as point clouds [7, 14], voxels [43, 65], NeRFs [10, 31, 40, 42, 64, 87], and 3D-GS [47]. While effective for learned tasks, these models struggle in unseen real-world scenarios [49, 50]. RL fills in this gap by refining policies through trial and error but requires costly real-world rollouts. Therefore, many methods adopt sim-to-real transfer, *i.e.*, deploying RL policies learned in digital twins of the world for task execution. Nonetheless, scalability remains a challenge due to their reliance on predefined assets [3, 54, 74] or labor-intensive conversion of real-world objects into simulation [9, 37, 41, 46, 59]. To address limitations, GWM focuses on providing both a stronger visual representation for imitation learning and an efficient neural simulator for visual RL to enable more effective and scalable robotic manipulation.

3. Gaussian World Model

The overall pipeline of our GWM method is shown in Figure 2, in which we construct a Gaussian world model to infer the future scene reconstruction represented by 3D Gaussian primitives. Specifically, we encode the real-world vision inputs into latent 3D Gaussian representations

(Sec. 3.1) and leverage a diffusion-based conditional generative model to learn the dynamics over representations given robot states and actions (Sec. 3.2). We demonstrate that GWM can be flexibly integrated into both offline imitation learning and online model-based reinforcement learning for diverse robotic manipulation tasks (Sec. 3.3).

3.1. World State Encoding

Feed-forward 3D Gaussian Splatting Given single or two-view image inputs $\mathcal{I} = \{\mathcal{I}_{i=\{1,2\}}\}$ of a world state, our goal is to first encode the scene into 3D Gaussian representations for dynamics learning and prediction. 3D-GS represents a 3D scene with multiple unstructured 3D Gaussian kernels $\mathbf{G} = \{\mathbf{x}_p, \sigma_p, \Sigma_p, \mathcal{C}_p\}_{p \in \mathcal{P}}$, where \mathbf{x}_p , σ_p , Σ_p , and \mathcal{C}_p represent the centers, opacities, covariance matrices, and spherical harmonic coefficients of the Gaussians, respectively. To obtain the color of each pixel from a given viewpoint, 3D-GS projects the 3D Gaussians onto the image plane and computes the pixel color as:

$$\mathbf{C}(\mathbf{G}) = \sum_{p \in \mathcal{P}} \alpha_p \text{SH}(\mathbf{d}_p; \mathcal{C}_p) \prod_{j=1}^{p-1} (1 - \alpha_j), \quad (1)$$

where α_p represents the z -depth ordered effective opacities, *i.e.*, products of the 2D Gaussian weights derived from Σ_p and their overall opacities σ_p ; \mathbf{d}_p stands for the view direction from the camera to \mathbf{x}_p ; $\text{SH}(\cdot)$ is the spherical harmonics function. Since vanilla 3D-GS relies on time-consuming per-scene offline optimization, we employ generalizable 3D-GS to learn feed-forward mappings from images to 3D Gaussians to accelerate the process. Specifically, we obtain the 3D Gaussian world state \mathbf{G} using Splatt3R [66], which first employs the stereo reconstruction model Mast3R [36] to generate 3D point maps from input images and then predicts the parameters of each 3D Gaussian given these point maps using an additional prediction head.

3D Gaussian VAE Since the number of learned 3D Gaussians for each world state can vary significantly across different scenes and tasks, we adopt a 3D Gaussian VAE (E_θ, D_θ) to encode the reconstructed 3D Gaussians \mathbf{G} into a fixed length of N latent embeddings $\mathbf{x} \in \mathbb{R}^{N \times D}$. Specifically, we first downsample the reconstructed 3D Gaussians \mathbf{G} to a fixed number of N Gaussians \mathbf{G}_N using Farthest Point Sampling (FPS): $\mathbf{G}_N = \text{FPS}(\mathbf{G})$. Next, we use these sampled Gaussians \mathbf{G}_N as queries to attend and aggregate information from all Gaussians \mathbf{G} to latent embedding \mathbf{x} using a L layer cross-attention-based encoder E_θ like [89]:

$$\begin{aligned} \mathbf{X} &= E_\theta(\mathbf{G}_N, \mathbf{G}) = E_\theta^{(L)} \circ \dots \circ E_\theta^{(1)}(\mathbf{G}_N, \mathbf{G}), \\ E_\theta^{(l)}(\mathbf{Q}, \mathbf{G}) &= \text{LayerNorm}(\text{CrossAttn}(\mathbf{Q}, \text{PosEmbed}(\mathbf{G}))). \end{aligned} \quad (2)$$

With latent encoding \mathbf{x} , we employ a mirrored transformer-based decoder D_θ to propagate and aggregate information

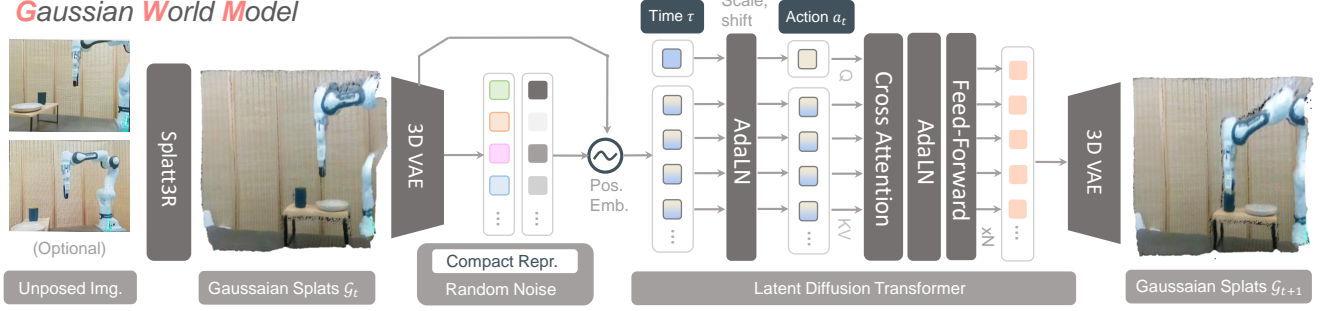


Figure 2. **The overall pipeline of GWM**, which primarily consists of a 3D variational encoder and a latent diffusion transformer. The 3D variational encoder embeds the Gaussian Splats estimated by a foundational reconstruction model to a compact latent space, and the diffusion transformer operates on the latent patches to interactively imagine the future Gaussian Splats conditioned on the robot action and denoising time step.

within the latent code set and leverage to obtain the reconstructed Gaussians \hat{G} :

$$\hat{G} = D_{\theta}(\mathbf{x}) = \text{LayerNorm}(\text{SelfAttn}(\mathbf{x}, \mathbf{x})) \quad (3)$$

For learning the 3D Gaussian VAE (E_{θ}, D_{θ}), we use the Chamfer loss between the centers of our reconstructed Gaussians \hat{G} and the original ones G for supervision. We also add a rendering loss of our reconstructed Gaussians \hat{G} to achieve high-fidelity rendering for image-based policy:

$$\mathcal{L}_{\text{VAE}} = \text{Chamfer}(\hat{G}, G) + \|C(\hat{G}) - C(G)\|_1 \quad (4)$$

3.2. Diffusion-based Dynamics Modeling

With the encoded world state embeddings \mathbf{x}_t at time t and its future state \mathbf{x}_{t+1} , we aim to learn the world dynamics $p(\mathbf{x}_{t+1}|\mathbf{x}_{\leq t}, a_{\leq t})$, where $\mathbf{x}_{\leq t}$ and $a_{\leq t}$ denote history states and actions, respectively. Specifically, we leverage a diffusion-based dynamics model where we convert dynamics learning into a conditional generation problem, generating future state \mathbf{x}_{t+1} from noise with history states and actions $\mathbf{y}_t = (\mathbf{x}_{\leq t}, a_{\leq t})$ as conditions.

Diffusion Formulation To generate the future state, we begin with the formulation of the diffusion process. Specifically, we first add noise to the ground truth future state $\mathbf{x}_{t+1}^0 = \mathbf{x}_{t+1}$ to obtain noised future state samples \mathbf{x}_{t+1}^{τ} via a Gaussian perturbation kernel:

$$p^{0 \rightarrow \tau}(\mathbf{x}_{t+1}^{\tau}|\mathbf{x}_{t+1}^0) = \mathcal{N}(\mathbf{x}_{t+1}^{\tau}; \mathbf{x}_{t+1}^0, \sigma^2(\tau)\mathbf{I}), \quad (5)$$

where τ is the noise step index and $\sigma(\tau)$ is the noise schedule. This diffusion process can be described as the solution to a stochastic differential equation (SDE) [68]:

$$d\mathbf{x} = \mathbf{f}(\mathbf{x}, \tau)d\tau + g(\tau)d\mathbf{w}, \quad (6)$$

where \mathbf{w} represents the standard Wiener process, \mathbf{f} is the drift coefficient, and g is the diffusion coefficient. Under this formulation, the effect of the Gaussian perturbation kernel is equivalent to setting $\mathbf{f}(\mathbf{x}, \tau) = \mathbf{0}$ and $g(\tau) =$

$\sqrt{2\sigma(\tau)\sigma'(\tau)}$. To generate samples from noises, we can reverse Eq. (6) using the reverse-time SDE [2] for sampling:

$$d\mathbf{x} = [\mathbf{f}(\mathbf{x}, \tau) - g(\tau)^2 \nabla_{\mathbf{x}} \log p^{\tau}(\mathbf{x})]d\tau + g(\tau)d\bar{\mathbf{w}}, \quad (7)$$

where $\bar{\mathbf{w}}$ denotes the reverse-time Wiener process and $\nabla_{\mathbf{x}} \log p^{\tau}(\mathbf{x})$ is the score function, *i.e.*, the gradient of log-marginal probability with respect to \mathbf{x} [29]. As the score function could be estimated by a network, we learn the conditional denoising model \mathcal{D}_{θ} by minimizing the difference between the sampled future state $\hat{\mathbf{x}}_{t+1}^0 = \mathcal{D}_{\theta}(\mathbf{x}_{t+1}^{\tau}, \mathbf{y}_t)$ and the ground truth future state \mathbf{x}_{t+1}^0 :

$$\mathcal{L}(\theta) = \mathbb{E} \left[\left\| \mathcal{D}_{\theta}(\mathbf{x}_{t+1}^{\tau}, \mathbf{y}_t^{\tau}) - \mathbf{x}_{t+1}^0 \right\|_2^2 \right]. \quad (8)$$

Learning with EDM As pointed out in [32], directly learning the denoiser $\mathcal{D}_{\theta}(\mathbf{x}_{t+1}^{\tau}, \mathbf{y}_t)$ can be affected by problems like varying noise magnitudes. Therefore, we follow [1] and adopt the practice in EDM [32] to learn a network \mathcal{F}_{θ} with preconditioning instead. Specifically, we parameterize the denoiser $\mathcal{D}_{\theta}(\mathbf{x}_{t+1}^{\tau}, \mathbf{y}_t^{\tau})$ with:

$$\mathcal{D}_{\theta}(\mathbf{x}_{t+1}^{\tau}, \mathbf{y}_t^{\tau}) = c_{\text{skip}}^{\tau} \mathbf{x}_{t+1}^{\tau} + c_{\text{out}}^{\tau} \mathcal{F}_{\theta}(c_{\text{in}}^{\tau} \mathbf{x}_{t+1}^{\tau}, \mathbf{y}_t^{\tau}; c_{\text{noise}}^{\tau}), \quad (9)$$

where the preconditioner c_{in}^{τ} and c_{out}^{τ} scale the input and output magnitudes, c_{skip}^{τ} modulates the skip connection, and c_{noise}^{τ} maps noise levels as an additional conditioning input into \mathcal{F}_{θ} . We provide details for these preconditioners in ???. With this conversion, we can rewrite the objective in Eq. (8) with:

$$\mathcal{L}(\theta) = \mathbb{E} \left[\left\| \mathbf{F}_{\theta}(c_{\text{in}}^{\tau} \mathbf{x}_{t+1}^{\tau}, \mathbf{y}_t^{\tau}) - \frac{1}{c_{\text{out}}^{\tau}} (\mathbf{x}_{t+1}^0 - c_{\text{skip}}^{\tau} \mathbf{x}_{t+1}^{\tau}) \right\|_2^2 \right]. \quad (10)$$

A crucial insight of this conversion is creating a new training target for better learning the network \mathcal{F}_{θ} by adaptively mixing signal and noise depending on the noise schedule $\sigma(\tau)$. Intuitively, at high noise levels ($\sigma(\tau) \gg \sigma_{\text{data}}$),

Algorithm 1: Monotonic Model-Based Policy Optimization (MBPO) with Gaussian World Model

Initialize policy $\pi(\mathbf{a}_t|\mathbf{s}_t)$, Gaussian world model $p_\theta(\mathbf{s}_{t+1}, \mathbf{r}_t|\mathbf{s}_t, \mathbf{a}_t)$, empty replay buffer \mathcal{B} ;
for N epochs **do**
 Collect data with π in real environment:
 $\mathcal{B} = \mathcal{B} \cup \{(\mathbf{s}_t, \mathbf{a}_t, \mathbf{s}_{t+1}, \mathbf{r}_t)\}_t$;
 Train Gaussian world model p_θ on dataset \mathcal{B} via maximum likelihood:
 $\theta \leftarrow \arg \max_\theta \mathbb{E}_{\mathcal{B}}[\log p_\theta(\mathbf{s}_{t+1}, \mathbf{r}_t|\mathbf{s}_t, \mathbf{a}_t)]$;
 Optimize policy under predictive model:
 $\pi \leftarrow \arg \max_\pi \mathbb{E}_\pi[\sum_{t \geq 0} \gamma^t \mathbf{r}_t]$;
end

$c_{\text{skip}}^\tau \rightarrow 0$ and the network primarily learns to predict the clean signal. Conversely, at low noise levels ($\sigma(\tau) \rightarrow 0$), $c_{\text{skip}}^\tau \rightarrow 1$, the target becomes the noise component, preventing the objective from becoming trivial.

Implementation Technically, we implement the network \mathcal{F}_θ with a DiT [56]. Given a sequence of actual world state latent embeddings $\{\mathbf{x}_t^0 = \mathbf{x}_t\}_{t=1}^T$, we first create latents with noise $\{\mathbf{x}_t^\tau\}_{t=1}^T$ following the Gaussian perturbation described in Eq. (5). Next, we concat the noise latent embeddings with rotary position embedding (RoPE [69]) and pass it to the DiT as inputs. In terms of conditions $\mathbf{y}_t = (\mathbf{x}_{\leq t}^0, a_{\leq t}, c_{\text{noise}}^\tau)$, the time embeddings are modulated by adaptive layer normalization (AdaLN [57]), and the current robot actions are used as the keys and values for the cross-attention layers within the DiT for conditional generation. For stability and efficiency across all attention mechanisms, we employ Root Mean Square Normalization (RMSNorm [88]) with learnable scales to stabilize training that processes spatial representations while incorporating temporal action sequences as conditions.

3.3. GWM for Policy Learning

GWM for Reinforcement Learning We demonstrate that GWM can be seamlessly integrated into existing model-based RL methods. Formally, a Markov Decision Process (MDP) is defined by the tuple $(\mathcal{S}, \mathcal{A}, p, r, \gamma, \rho_0)$. \mathcal{S} and \mathcal{A} are the state and action spaces, γ is the discount factor, and $r(\mathbf{s}, \mathbf{a})$ is the reward function. The goal of model-based RL [30] is to learn a policy π that maximizes the expected sum of discounted rewards $\pi^* = \arg \max_\pi \mathbb{E}_\pi[\sum_{t=0}^{\infty} \gamma^t \mathbf{r}_t]$ while constructing a model of the dynamics $p_\theta(\mathbf{s}_{t+1}, \mathbf{r}_t|\mathbf{s}_t, \mathbf{a}_t)$ using the policy roll-outs. We provide the pseudo-code for the model-based RL policy learning in Algorithm 1. Under this formulation, we add an additional reward prediction head over GWM to parameterize the dynamics model $p_\theta(\mathbf{s}_{t+1}, \mathbf{r}_t|\mathbf{s}_t, \mathbf{a}_t)$. To improve performance in visual manipulation tasks, we build our base

RL policy following design choices discussed in [78].

GWM for Imitation Learning In imitation learning, we use GWM as a more effective encoder to provide better features for policy learning. Specifically, we use the feature vector after the first denoising step in the diffusion process as the input for downstream policy models like BC-transformer [55] and diffusion policy [8]. The first denoising step carries out the representative spatial information to deal with the severe noise level. In our implementations, we predict actions in sequential chunks to promote consistency in robotic control.

4. Experiments

In our experiments, we focus on the following questions:

1. How is the quality of the action-conditioned video prediction results across different domains?
2. Does Gaussian world model benefits downstream imitation and reinforcement learning? Does it show greater robustness compared with image-based world model?
3. How does the Gaussian world model help typical policies (e.g., diffusion policy [8]) in real-world robotic manipulation tasks?

In the following sections, we describe in detail the model performance regarding these key topics. Specifically, we leverage the following three testing environments and four tasks in our experiments:

Environments To provide a comprehensive analysis of GWM’s capability, we evaluate our method on two synthetic and one real-world environment: (1) META-WORLD [86], a synthetic environment for learning RL policies for robotic manipulation; (2) ROBOCASA [55], a large-scale multi-scale synthetic imitation learning benchmark featuring diverse robotic manipulation tasks in the kitchen environment; and (3) FRANKA-PNP, a real-world pick-and-place environment using a Franka Emika FR3 robot arm.

Tasks We meticulously design four tasks to evaluate GWM across various testing environments systematically: (1) Action-conditioned scene prediction assesses GWM’s effectiveness in world modeling and future prediction; (2) GWM-based imitation learning examines the representation quality and its benefits for imitation-learning-based robotic manipulation; (3) GWM-based RL explores its potential for model-based reinforcement learning; and (4) real-world task deployment evaluates GWM’s robustness in real-world robot manipulation.

4.1. Action-conditioned Scene Prediction

Experiment Setup The capability of a world model to generate high-fidelity and action-aligned rollouts is critical for effective policy optimization. To evaluate this capability, we train GWM on human demonstrations available on all considered real and synthetic environments, and evaluate

Table 1. **Quantitative results for future state prediction** on Meta-World and FRANKA PNP. LPIPS and SSIM scores are scaled by 100. Best results are highlighted in bold.

Dataset	Method	FVD↓	PSNR↑	SSIM↑	LPIPS↓
META-WORLD	iVideoGPT	75.0	20.4	82.3	9.5
	GWM	73.0	20.6	82.8	9.0
FRANKA-PNP	iVideoGPT	63.2	27.8	90.6	4.9
	GWM	61.5	28.0	91.0	4.5




Figure 3. **Qualitative comparison between models on META-WORLD.** GWM successfully predicts better details on the gripper movement (highlighted in blue).

the future prediction quality by conditioning the model on unseen action trajectories sampled from the validation set. For quantitative evaluation, we employ common metrics for generation quality, including FVD [72] to measure temporal consistency, image-based metrics including PSNR [28] for pixel-level accuracy, alongside SSIM [77] and LPIPS [91] for perceptual quality.

Results and Analyses We provide quantitative comparison between our method and iVideoGPT in Tab. 1. As shown in Tab. 1, our method consistently outperforms the current state-of-the-art image-based world modeling method iVideoGPT on both synthetic and real-world environments, demonstrating the effectiveness of our diffusion-based Gaussian world model learning pipeline. Notably, as shown in Fig. 3, image-based models like iVideoGPT are prone to failures in capturing dynamics details (*e.g.*, the gripper). Though these details might not cause large differences in visual metrics, they will significantly affect policy learning as we later discuss in Sec. 4.3. We provide more qualitative visualizations of GWM’s prediction result on ROBOCASA and FRANKA-PNP in Fig. 4.

4.2. GWM-based Imitation Learning

Experiment Setup As discussed in Sec. 3.3, GWM can be used to extract informative representation from image observation, which is expected to benefit imitation learning. We verify this property by testing GWM’s effectiveness for imitation learning on ROBOCASA. The task suite in ROBOCASA comprises 24 atomic tasks with related language instructions for kitchen environments, including actions such as pick-and-place, open, and close. Each task is provided




Figure 4. **Qualitative visualization on future state prediction of GWM on FRANKA-PNP and ROBOCASA.** All predictions are rolled out by applying the unseen action trajectory from the valid dataset. Zoom in for details.

with a limited set of 50 human demonstrations and a set of 3000 generated demonstrations from MimicGen [51]. We train our GWM on these demonstrations and pass it as the state encoding for the state-of-the-art BC-transformer [55] for quantitative comparison on the success rate metrics.

Results and Analyses Our experimental results on the ROBOCASA benchmark are presented in Table 2, which demonstrate the effectiveness of our method in multi-task imitation learning scenarios. Across 24 kitchen manipulation tasks, our approach consistently outperforms the BC-Transformer baseline. With limited human demonstrations (H-50), our method shows an average improvement of 10.5% in success rates. When trained on generated demonstrations (G-3000), our method maintains scalable performance with an average gain of 7.6%. Notably, our approach exhibits particular strengths in complex manipulation tasks such as pick-and-place operations and interactive tasks like turning on/off appliances, where the performance gains are most significant. These results confirm that our method’s ability to extract informative representations from visual observations effectively enhances imitation learning capabilities in practical robotic manipulation scenarios.

4.3. GWM-based Reinforcement Learning

Experiment Setup We evaluate GWM’s capabilities for RL policies on six Meta-World [86] robotic manipulation tasks with increasing complexity. We implement a model-based RL approach inspired by MBPO [30], using GWM to generate synthetic rollouts that augment the replay buffer of a DrQ-v2 [84] actor-critic algorithm. We include the state-of-the-art image-based world model iVideoGPT [78] as a strong baseline. For fair comparisons, we do not utilize pre-

Table 2. **Multi-Task Imitation Learning Results in Robocasa.** Average success rates (%) of multi-task agents trained with 50 human demonstrations or 3000 generated demonstrations per task. Results are evaluated over 50 episodes with different floor plans and styles.

Method	PnP CabToCounter		PnP CounterToCab		PnP CounterToMicrowave		PnP CounterToSink		PnP CounterToStove		PnP MicrowaveToCounter	
	H-50	G-3000	H-50	G-3000	H-50	G-3000	H-50	G-3000	H-50	G-3000	H-50	G-3000
BC-transformer	2	18	6	28	2	18	2	44	2	6	2	8
GWM	18	32	4	22	14	44	20	38	2	18	20	26
Δ	+16	+14	-2	-6	+12	+26	+18	-6	0	+12	+18	+18
	PnP SinkToCounter		PnP StoveToCounter		Open SingleDoor		Open DoubleDoor		Close DoubleDoor		Close SingleDoor	
	H-50	G-3000	H-50	G-3000	H-50	G-3000	H-50	G-3000	H-50	G-3000	H-50	G-3000
BC-transformer	8	42	6	28	46	50	28	48	28	46	56	94
GWM	22	38	18	44	58	62	28	42	50	58	54	90
Δ	+14	-4	+12	+16	+12	+12	0	-6	+22	+12	-2	-4
	Open Drawer		Close Drawer		TurnOn Stove		TurnOff Stove		TurnOn SinkFaucet		TurnOff SinkFaucet	
	H-50	G-3000	H-50	G-3000	H-50	G-3000	H-50	G-3000	H-50	G-3000	H-50	G-3000
BC-transformer	42	74	80	96	32	46	4	24	38	34	50	72
GWM	56	90	80	90	46	80	22	40	52	48	44	66
Δ	+14	+16	0	-6	+14	+24	+18	+16	+14	+14	-6	-6
	Turn SinkSpout		CoffeePress Button		TurnOn Microwave		TurnOff Microwave		CoffeeServe Mug		CoffeeSetup Mug	
	H-50	G-3000	H-50	G-3000	H-50	G-3000	H-50	G-3000	H-50	G-3000	H-50	G-3000
BC-transformer	54	96	48	74	62	90	70	60	22	34	0	12
GWM	72	90	76	90	64	84	70	54	36	50	16	28
Δ	+18	-6	+28	+16	+2	-6	0	-6	+14	+16	+16	+16

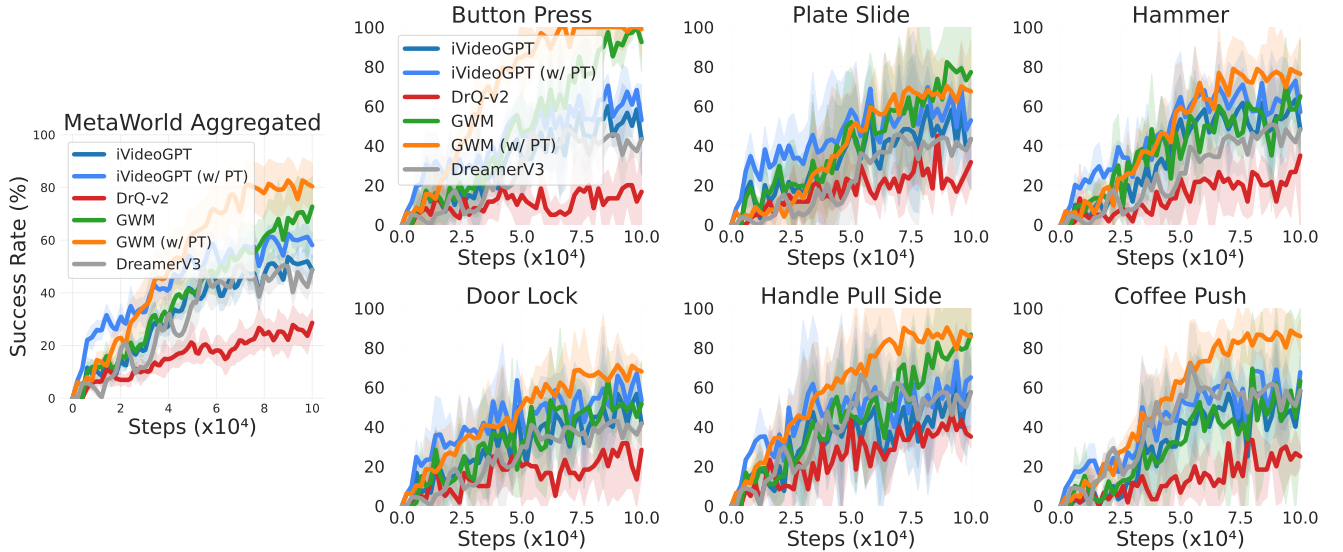


Figure 5. **Model-based RL Results of GWM and iVideoGPT [78] on METAWORLD.** The shadow area represents 95% confidence interval (CI) across three random seeds. Each data point is evaluated over 20 episodes.

trained initialization of both methods. For fair comparisons, all compared methods use the same context length, horizon, and are trained to a maximum of 1×10^5 steps.

Results and Analyses Figure 5 demonstrates that GWM consistently outperforms iVideoGPT across all six MetaWorld tasks. On average, GWM converges approximately 2

faster than iVideoGPT and reaches higher asymptotic performance on complex manipulation tasks. The superior performance stems from GWM’s 3D Gaussian representation, which allows more accurate prediction of contact dynamics and object movement under manipulation, compared to purely image-based approaches. The results confirm that




Figure 6. **Real-World Experiment Setup.** Left: using a Franka Emika Panda robotic arm equipped with an RGB camera, we evaluate the performance of the diffusion policy [8] both with and without our proposed method. Right: the robot’s visual inputs in the task completion.

Table 3. **Real-world Experiment Results.** We report the number of successful trials out of all 20 trials in FRANKA PNP.

FRANKA-PNP	Diffusion Policy	GWM (Ours)
Cup distractor	6/10	7/10
Plate distractor	1/5	3/5
Table distractor	0/5	3/5
Total	7/20	13/20

explicit 3D representation offers substantial advantages for robotic control tasks requiring precise spatial reasoning.

4.4. Real-world Deployment

Experiment Setup We deploy a Franka Emika FR3 robotic arm and a Panda gripper for our real-robot experiments. We focus on the real-world task of picking a colored cup, and placing it onto a plate on the table. We collect a small set of 30 demonstrations using the Mujoco AR teleoperation interfaces. We also setup one third-view Realsense D435i camera to provide unposed RGB-only images for observation. We provide an overview of the real-world task setting in Fig. 6. Similar to the experiment setting in Sec. 4.2, we compare the performance of the state-of-the-art RGB-based policy Diffusion Policy [8] with or without our GWM representation on task success rate for quantitative analysis.

Results and Analysis As shown in Table 3, GWM outperforms Diffusion Policy (**65% vs. 35%** success rate) on 20 trials with different initial start positions and object locations (*i.e.* distractors). The performance gap widens for novel distractors, demonstrating GWM’s superior generalization capabilities. Our approach maintains consistent performance across task variations due to its effective world model that captures task-relevant dynamics while being robust to visual differences. The real-world rollouts are shown

Table 4. **Ablation Study** on PnP CabToCounter in ROBO-CASA task suite. We report the reconstruction metrics and the success rates (SR) of imitation learning on the Human-50 dataset.

GS	3D VAE	FVD↓	PSNR↑	SSIM↑	LPIPS↓	SR ↑
✗	✗	67.8	27.2	88.2	5.1	4
✓	✗	65.3	26.9	89.5	4.9	18
✓	✓	62.4	28.1	90.8	4.6	24

in the supplementary file, where the advantage stems primarily from more precise object localization and accurate placement operations. The results demonstrate GWM’s robust spatial-temporal understanding in real-world robotic manipulation tasks.

4.5. Ablation Analysis

We conduct additional experiments on ROBOCASA to further verify our design choices.

Choice of Gaussian Splatting As shown in Table 4, compared to directly building image-based world model with diffusion transformer on par with [1], introducing Gaussian Splatting significantly improves the success rate (SR) from 4% to 18%. While PSNR shows a slight decrease, both SSIM and LPIPS metrics improve, suggesting that Gaussian Splatting provides better 3D consistency across different time steps. This validates our hypothesis that explicit 3D representation enhances spatial understanding for robot learning compared to pure 2D approaches.

Choice of 3D Gaussian VAE Further incorporating the 3D VAE component yields consistent improvements across all metrics, including PSNR. The success rate further improves from 18% to 24%. The results confirm that our 3D Gaussian VAE efficiently captures the latent structure of the scene, enabling more compact scene representation while maintaining spatial understanding.

5. Conclusion

In this paper, we introduce a novel Gaussian World Model (GWM) for robotic manipulation that addresses limitations of image-based world models by incorporating robust geometric information. Our approach reconstructs future states by modeling the propagation of Gaussian primitives under robot actions. The method combines a DiT with a 3D-aware variational autoencoder for precise scene-level future state reconstruction via Gaussian Splatting. We develop a scalable data processing pipeline to facilitate test-time updates within a model-based reinforcement learning framework, extracting aligned Gaussian splats from unposed images. Experiments in both simulated and real-world settings demonstrate the effectiveness of GWM in predicting future scenes and training superior policies.

Acknowledgement

This work was supported by Shenzhen Science and Technology Program (JCYJ20240813111903006) and Guangdong Natural Science Funds for Distinguished Young Scholar (No. 2025B1515020012).

References

- [1] Eloi Alonso, Adam Jelley, Vincent Micheli, Anssi Kanervisto, Amos Storkey, Tim Pearce, and François Fleuret. Diffusion for world modeling: Visual details matter in atari. *arXiv preprint arXiv:2405.12399*, 2024. 2, 3, 4, 8
- [2] Brian DO Anderson. Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12(3):313–326, 1982. 4
- [3] Genesis Authors. Genesis: A universal and generative physics engine for robotics and beyond, 2024. 3
- [4] Hritik Bansal, Zongyu Lin, Tianyi Xie, Zeshun Zong, Michal Yarom, Yonatan Bitton, Chenfanfu Jiang, Yizhou Sun, Kai-Wei Chang, and Aditya Grover. Videophy: Evaluating physical commonsense for video generation. *arXiv preprint arXiv:2406.03520*, 2024. 3
- [5] Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, et al. A vision-language-action flow model for general robot control. *arXiv preprint arXiv:2410.24164*, 2024. 3
- [6] David Charatan, Sizhe Li, Andrea Tagliasacchi, and Vincent Sitzmann. pixelsplat: 3d gaussian splats from image pairs for scalable generalizable 3d reconstruction. *arXiv preprint arXiv:2312.12337*, 2023. 3
- [7] Shizhe Chen, Ricardo Garcia, Cordelia Schmid, and Ivan Laptev. Polarnet: 3d point clouds for language-guided robotic manipulation. *arXiv preprint arXiv:2309.15596*, 2023. 3
- [8] Cheng Chi, Zhenjia Xu, Siyuan Feng, Eric Cousineau, Yilun Du, Benjamin Burchfiel, Russ Tedrake, and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. *International Journal of Robotics Research (IJRR)*, page 02783649241273668, 2023. 5, 8
- [9] Tianyuan Dai, Josiah Wong, Yunfan Jiang, Chen Wang, Cem Gokmen, Ruohan Zhang, Jiajun Wu, and Li Fei-Fei. Acdc: Automated creation of digital cousins for robust policy learning. In *8th Annual Conference on Robot Learning*, 2024. 3
- [10] Danny Driess, Ingmar Schubert, Pete Florence, Yunzhu Li, and Marc Toussaint. Reinforcement learning with neural radiance fields. *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*, 2022. 3
- [11] Jay W Forrester. Counterintuitive behavior of social systems. *Theory and decision*, 2(2):109–140, 1971. 2
- [12] Yang Fu, Sifei Liu, Amey Kulkarni, Jan Kautz, Alexei A Efros, and Xiao-long Wang. Colmap-free 3d gaussian splatting. *arXiv preprint arXiv:2312.07504*, 2023. 3
- [13] Zeyu Gao, Yao Mu, Ruoyan Shen, Chen Chen, Yangang Ren, Jianyu Chen, Shengbo Eben Li, Ping Luo, and Yanfeng Lu. Enhance sample efficiency and robustness of end-to-end urban autonomous driving via semantic masked world model. *arXiv preprint arXiv:2210.04017*, 2022. 2
- [14] Theophile Gervet, Zhou Xian, Nikolaos Gkanatsios, and Katerina Fragkiadaki. Act3d: 3d feature field transformers for multi-task robotic manipulation. In *Conference on Robot Learning (CoRL)*, pages 3949–3965, 2023. 3
- [15] Ankit Goyal, Jie Xu, Yijie Guo, Valts Blukis, Yu-Wei Chao, and Dieter Fox. Rvt: Robotic view transformer for 3d object manipulation. *arXiv preprint arXiv:2306.14896*, 2023. 2
- [16] Ankit Goyal, Valts Blukis, Jie Xu, Yijie Guo, Yu-Wei Chao, and Dieter Fox. Rvt-2: Learning precise manipulation from few demonstrations. *arXiv preprint arXiv:2406.08545*, 2024. 2
- [17] David Ha and Jürgen Schmidhuber. Recurrent world models facilitate policy evolution. *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*, 31, 2018. 2
- [18] Danijar Hafner, Timothy Lillicrap, Jimmy Ba, and Mohammad Norouzi. Dream to control: Learning behaviors by latent imagination. *arXiv preprint arXiv:1912.01603*, 2019.
- [19] Danijar Hafner, Timothy Lillicrap, Mohammad Norouzi, and Jimmy Ba. Mastering atari with discrete world models. *arXiv preprint arXiv:2010.02193*, 2020.
- [20] Danijar Hafner, Kuang-Huei Lee, Ian Fischer, and Pieter Abbeel. Deep hierarchical planning from pixels. *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*, 35:26091–26104, 2022.
- [21] Danijar Hafner, Jurgis Pasukonis, Jimmy Ba, and Timothy Lillicrap. Mastering diverse domains through world models. *arXiv preprint arXiv:2301.04104*, 2023. 2
- [22] Nicklas Hansen, Hao Su, and Xiaolong Wang. Td-mpc2: Scalable, robust world models for continuous control. *arXiv preprint arXiv:2310.16828*, 2023. 2
- [23] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020. 2
- [24] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 2
- [25] Anthony Hu, Gianluca Corrado, Nicolas Griffiths, Zachary Murez, Corina Gurau, Hudson Yeo, Alex Kendall, Roberto Cipolla, and Jamie Shotton. Model-based imitation learning for urban driving. *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*, 35:20703–20716, 2022. 2
- [26] Anthony Hu, Lloyd Russell, Hudson Yeo, Zak Murez, George Fedoseev, Alex Kendall, Jamie Shotton, and Gianluca Corrado. Gaia-1: A generative world model for autonomous driving. *arXiv preprint arXiv:2309.17080*, 2023. 2
- [27] Yi-Hua Huang, Yang-Tian Sun, Ziyi Yang, Xiaoyang Lyu, Yan-Pei Cao, and Xiaojuan Qi. Sc-gs: Sparse-controlled gaussian splatting for editable dynamic scenes. In *Proceedings of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 3
- [28] Quan Huynh-Thu and Mohammed Ghanbari. Scope of validity of psnr in image/video quality assessment. *Electronics letters*, 44(13):800–801, 2008. 6

- [29] Aapo Hyvärinen and Peter Dayan. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research (JMLR)*, 6(4), 2005. 4
- [30] Michael Janner, Justin Fu, Marvin Zhang, and Sergey Levine. When to trust your model: Model-based policy optimization. In *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*, 2019. 5, 6
- [31] Zhenyu Jiang, Yifeng Zhu, Maxwell Svetlik, Kuan Fang, and Yuke Zhu. Synergies between affordance and geometry: 6-dof grasp detection via implicit representations. *arXiv preprint arXiv:2104.01542*, 2021. 3
- [32] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. *Advances in Neural Information Processing Systems*, 35:26565–26577, 2022. 4
- [33] Nikhil Keetha, Jay Karhade, Krishna Murthy Jatavallabhula, Gengshan Yang, Sebastian Scherer, Deva Ramanan, and Jonathon Luiten. Splatam: Splat track & map 3d gaussians for dense rgb-d slam. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21357–21366, 2024. 3
- [34] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *TOG*, 42(4), 2023. 2, 3
- [35] Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair, Rafael Rafailov, Ethan Foster, Grace Lam, Pannag Sanketi, et al. Openvla: An open-source vision-language-action model. *arXiv preprint arXiv:2406.09246*, 2024. 3
- [36] Vincent Leroy, Yohann Cabon, and Jérôme Revaud. Grounding image matching in 3d with mast3r. *arXiv preprint arXiv:2406.09756*, 2024. 3
- [37] Puhao Li, Tengyu Liu, Yuyang Li, Muzhi Han, Haoran Geng, Shu Wang, Yixin Zhu, Song-Chun Zhu, and Siyuan Huang. Ag2manip: Learning novel manipulation skills with agent-agnostic visual and action representations. In *Proceedings of International Conference on Intelligent Robots and Systems (IROS)*, pages 573–580. IEEE, 2024. 3
- [38] Puhao Li, Yingying Wu, Ziheng Xi, Wanlin Li, Yuzhe Huang, Zhiyuan Zhang, Yinghan Chen, Jianan Wang, Song-Chun Zhu, Tengyu Liu, et al. Controlvla: Few-shot object-centric adaptation for pre-trained vision-language-action models. *arXiv preprint arXiv:2506.16211*, 2025. 3
- [39] Xuanlin Li, Kyle Hsu, Jiayuan Gu, Karl Pertsch, Oier Mees, Homer Rich Walke, Chuyuan Fu, Ishikaa Lunawat, Isabel Sieh, Sean Kirmani, et al. Evaluating real-world robot manipulation policies in simulation. *arXiv preprint arXiv:2405.05941*, 2024. 2
- [40] Yunzhu Li, Shuang Li, Vincent Sitzmann, Pulkit Agrawal, and Antonio Torralba. 3d neural scene representations for visuomotor control. In *Conference on Robot Learning (CoRL)*, pages 112–123, 2022. 3
- [41] Vincent Lim, Huang Huang, Lawrence Yunliang Chen, Jonathan Wang, Jeffrey Ichnowski, Daniel Seita, Michael Laskey, and Ken Goldberg. Real2sim2real: Self-supervised learning of physical single-step dynamic actions for planar robot casting. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 8282–8289. IEEE, 2022. 3
- [42] Yen-Chen Lin, Pete Florence, Andy Zeng, Jonathan T Barron, Yilun Du, Wei-Chiu Ma, Anthony Simeonov, Alberto Rodriguez Garcia, and Phillip Isola. Mira: Mental imagery for robotic affordances. In *Conference on Robot Learning (CoRL)*, pages 1916–1927, 2023. 3
- [43] I Liu, Chun Arthur, Sicheng He, Daniel Seita, and Gaurav Sukhatme. Voxact-b: Voxel-based acting and stabilizing policy for bimanual manipulation. *arXiv preprint arXiv:2407.04152*, 2024. 3
- [44] Songming Liu, Lingxuan Wu, Bangguo Li, Hengkai Tan, Huayu Chen, Zhengyi Wang, Ke Xu, Hang Su, and Jun Zhu. Rdt-1b: a diffusion foundation model for bimanual manipulation. *arXiv preprint arXiv:2410.07864*, 2024. 3
- [45] Yifan Liu, Chenxin Li, Chen Yang, and Yixuan Yuan. Endogaussian: Gaussian splatting for deformable surgical scene reconstruction. *arXiv preprint arXiv:2401.12561*, 2024. 3
- [46] Haozhe Lou, Yurong Liu, Yike Pan, Yiran Geng, Jianteng Chen, Wenlong Ma, Chenglong Li, Lin Wang, Hengzhen Feng, Lu Shi, et al. Robo-gs: A physics consistent spatial-temporal model for robotic arm with hybrid representation. *arXiv preprint arXiv:2408.14873*, 2024. 3
- [47] Guanxing Lu, Shiyi Zhang, Ziwei Wang, Changliu Liu, Jiwen Lu, and Yansong Tang. Manigaussian: Dynamic gaussian splatting for multi-task robotic manipulation. In *European Conference on Computer Vision*, pages 349–366. Springer, 2024. 2, 3
- [48] Jonathon Luiten, Georgios Kopanas, Bastian Leibe, and Deva Ramanan. Dynamic 3d gaussians: Tracking by persistent dynamic view synthesis. In *Proceedings of International Conference on 3D Vision (3DV)*, 2024. 3
- [49] Jianlan Luo, Zheyuan Hu, Charles Xu, You Liang Tan, Jacob Berg, Archit Sharma, Stefan Schaal, Chelsea Finn, Abhishek Gupta, and Sergey Levine. Serl: A software suite for sample-efficient robotic reinforcement learning. In *Proceedings of International Conference on Robotics and Automation (ICRA)*, pages 16961–16969. IEEE, 2024. 3
- [50] Jianlan Luo, Charles Xu, Jeffrey Wu, and Sergey Levine. Precise and dexterous robotic manipulation via human-in-the-loop reinforcement learning. *arXiv preprint arXiv:2410.21845*, 2024. 3
- [51] Ajay Mandlekar, Soroush Nasiriany, Bowen Wen, Iretiayo Akinola, Yashraj Narang, Linxi Fan, Yuke Zhu, and Dieter Fox. Mimicgen: A data generation system for scalable robot learning using human demonstrations. *arXiv preprint arXiv:2310.17596*, 2023. 6
- [52] Vincent Micheli, Eloi Alonso, and François Fleuret. Transformers are sample-efficient world models. *International Conference on Learning Representations*, 2023. 2
- [53] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *CACM*, 65(1):99–106, 2021. 2, 3
- [54] Yao Mu, Tianxing Chen, Zanxin Chen, Shijia Peng, Zhiqian Lan, Zeyu Gao, Zhixuan Liang, Qiaojun Yu, Yude Zou, Mingkun Xu, et al. Robotwin: Dual-arm robot benchmark with generative digital twins. In *Proceedings of Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 27649–27660, 2025. 3

- [55] Soroush Nasiriany, Abhiram Maddukuri, Lance Zhang, Adeet Parikh, Aaron Lo, Abhishek Joshi, Ajay Mandlekar, and Yuke Zhu. Robocasa: Large-scale simulation of everyday tasks for generalist robots. *arXiv preprint arXiv:2406.02523*, 2024. 5, 6
- [56] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of International Conference on Computer Vision (ICCV)*, pages 4195–4205, 2023. 5
- [57] Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. Film: Visual reasoning with a general conditioning layer. In *Proceedings of AAAI Conference on Artificial Intelligence (AAAI)*, 2018. 5
- [58] Skand Peri, Iain Lee, Chanh Kim, Li Fuxin, Tucker Hermans, and Stefan Lee. Point cloud models improve visual robustness in robotic learners. In *Proceedings of International Conference on Robotics and Automation (ICRA)*, 2024. 2
- [59] Mohammad Nomaan Qureshi, Sparsh Garg, Francisco Yandun, David Held, George Kantor, and Abhishek Silwal. SplatSim: Zero-shot sim2real transfer of rgb manipulation policies using gaussian splatting. *arXiv preprint arXiv:2409.10161*, 2024. 3
- [60] Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. Improving language understanding by generative pre-training. 2018. 2
- [61] Jan Robine, Marc Höftmann, Tobias Uelwer, and Stefan Harmeling. Transformer-based world models are happy with 100k interactions. *International Conference on Learning Representations*, 2023. 2
- [62] Julian Schrittwieser, Ioannis Antonoglou, Thomas Hubert, Karen Simonyan, Laurent Sifre, Simon Schmitt, Arthur Guez, Edward Lockhart, Demis Hassabis, Thore Graepel, et al. Mastering atari, go, chess and shogi by planning with a learned model. *Nature*, 588(7839):604–609, 2020. 2
- [63] Younggyo Seo, Danijar Hafner, Hao Liu, Fangchen Liu, Stephen James, Kimin Lee, and Pieter Abbeel. Masked world models for visual control. In *Conference on Robot Learning (CoRL)*, pages 1332–1344, 2023. 2
- [64] Dongseok Shim, Seungjae Lee, and H Jin Kim. Snerl: Semantic-aware neural radiance fields for reinforcement learning. In *Proceedings of International Conference on Machine Learning (ICML)*, 2023. 3
- [65] Mohit Shridhar, Lucas Manuelli, and Dieter Fox. Perceiver-actor: A multi-task transformer for robotic manipulation. In *Conference on Robot Learning (CoRL)*, 2023. 3
- [66] Brandon Smart, Chuanxia Zheng, Iro Laina, and Victor Adrian Prisacariu. Splatt3r: Zero-shot gaussian splatting from uncalibrated image pairs. *arXiv preprint arXiv:2408.13912*, 2024. 3
- [67] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. *International Conference on Machine Learning*, 2015. 2
- [68] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *International Conference on Learning Representations*, 2020. 2, 4
- [69] Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024. 5
- [70] Stanislaw Szymanowicz, Christian Rupprecht, and Andrea Vedaldi. Splatter image: Ultra-fast single-view 3d reconstruction. *arXiv preprint arXiv:2312.13150*, 2023. 3
- [71] Octo Model Team, Dibya Ghosh, Homer Walke, Karl Pertsch, Kevin Black, Oier Mees, Sudeep Dasari, Joey Hejna, Tobias Kreiman, Charles Xu, et al. Octo: An open-source generalist robot policy. *arXiv preprint arXiv:2405.12213*, 2024. 3
- [72] Thomas Unterthiner, Sjoerd Van Steenkiste, Karol Kurach, Raphael Marinier, Marcin Michalski, and Sylvain Gelly. Towards accurate generative models of video: A new metric & challenges. *arXiv preprint arXiv:1812.01717*, 2018. 6
- [73] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*, 2017. 2
- [74] Lirui Wang, Yiyang Ling, Zhecheng Yuan, Mohit Shridhar, Chen Bao, Yuzhe Qin, Bailin Wang, Huazhe Xu, and Xiaolong Wang. Gensim: Generating robotic simulation tasks via large language models. *arXiv preprint arXiv:2310.01361*, 2023. 3
- [75] Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. Dust3r: Geometric 3d vision made easy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20697–20709, 2024. 3
- [76] Xiaofeng Wang, Zheng Zhu, Guan Huang, Xinze Chen, and Jiwen Lu. Drivedreamer: Towards real-world-driven world models for autonomous driving. *arXiv preprint arXiv:2309.09777*, 2023. 2
- [77] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *Proceedings of Transactions on Image Processing (TIP)*, 13(4):600–612, 2004. 6
- [78] Jialong Wu, Shaofeng Yin, Ningya Feng, Xu He, Dong Li, Jianye Hao, and Mingsheng Long. ivideopt: Interactive videopts are scalable world models. *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*, 37:68082–68119, 2025. 2, 5, 6, 7
- [79] Philipp Wu, Alejandro Escontrela, Danijar Hafner, Pieter Abbeel, and Ken Goldberg. Daydreamer: World models for physical robot learning. In *Conference on Robot Learning (CoRL)*, pages 2226–2240, 2023. 2
- [80] Tianyi Xie, Zeshun Zong, Yuxing Qiu, Xuan Li, Yutao Feng, Yin Yang, and Chenfanfu Jiang. Physgaussian: Physics-integrated 3d gaussians for generative dynamics. In *Proceedings of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 3
- [81] Dejia Xu, Ye Yuan, Morteza Mardani, Sifei Liu, Jiaming Song, Zhangyang Wang, and Arash Vahdat. Agg: Amortized generative 3d gaussians for single image to 3d. *arXiv preprint arXiv:2401.04099*, 2024. 3

- [82] Yinghao Xu, Zifan Shi, Wang Yifan, Hansheng Chen, Ceyuan Yang, Sida Peng, Yujun Shen, and Gordon Wetstein. Grm: Large gaussian reconstruction model for efficient 3d reconstruction and generation. *arXiv preprint arXiv:2403.14621*, 2024. 3
- [83] Mengjiao Yang, Yilun Du, Kamyar Ghasemipour, Jonathan Tompson, Dale Schuurmans, and Pieter Abbeel. Learning interactive real-world simulators. *arXiv preprint arXiv:2310.06114*, 2023. 2, 3
- [84] Denis Yarats, Rob Fergus, Alessandro Lazaric, and Lerrel Pinto. Mastering visual continuous control: Improved data-augmented reinforcement learning. In *Proceedings of International Conference on Learning Representations (ICLR)*, 2022. 6
- [85] Weirui Ye, Shaohuai Liu, Thanard Kurutach, Pieter Abbeel, and Yang Gao. Mastering atari games with limited data. *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*, 34:25476–25488, 2021. 2
- [86] Tianhe Yu, Deirdre Quillen, Zhanpeng He, Ryan Julian, Karol Hausman, Chelsea Finn, and Sergey Levine. Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning. In *Conference on Robot Learning (CoRL)*, pages 1094–1100, 2020. 5, 6
- [87] Yanjie Ze, Ge Yan, Yueh-Hua Wu, Annabella Macaluso, Yuying Ge, Jianglong Ye, Nicklas Hansen, Li Erran Li, and Xiaolong Wang. Gnfactor: Multi-task real robot learning with generalizable neural feature fields. In *Conference on Robot Learning (CoRL)*, pages 284–301. PMLR, 2023. 2, 3
- [88] Biao Zhang and Rico Sennrich. Root mean square layer normalization. *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*, 32, 2019. 5
- [89] Biao Zhang, Jiapeng Tang, Matthias Niessner, and Peter Wonka. 3dshape2vecset: A 3d shape representation for neural fields and generative diffusion models. *ACM Transactions On Graphics (TOG)*, 42(4):1–16, 2023. 3
- [90] Bowen Zhang, Yiji Cheng, Jiaolong Yang, Chunyu Wang, Feng Zhao, Yansong Tang, Dong Chen, and Baining Guo. Gaussiancube: Structuring gaussian splatting using optimal transport for 3d generative modeling. *arXiv preprint arXiv:2403.19655*, 2024. 3
- [91] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 6
- [92] Weipu Zhang, Gang Wang, Jian Sun, Yetian Yuan, and Gao Huang. Storm: Efficient stochastic transformer based world models for reinforcement learning. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. 2
- [93] Shunyuan Zheng, Boyao Zhou, Ruizhi Shao, Boning Liu, Shengping Zhang, Liqiang Nie, and Yebin Liu. Gps-gaussian: Generalizable pixel-wise 3d gaussian splatting for real-time human novel view synthesis. *arXiv preprint arXiv:2312.02155*, 2023. 3
- [94] Wenzhao Zheng, Weiliang Chen, Yuanhui Huang, Borui Zhang, Yueqi Duan, and Jiwen Lu. Occworld: Learning a 3d occupancy world model for autonomous driving. *arXiv preprint arXiv:2311.16038*, 2023. 2
- [95] Xiaoyu Zhou, Zhiwei Lin, Xiaojun Shan, Yongtao Wang, Deqing Sun, and Ming-Hsuan Yang. Drivinggaussian: Composite gaussian splatting for surrounding dynamic autonomous driving scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21634–21643, 2024. 3
- [96] Haoyi Zhu, Yating Wang, Di Huang, Weicai Ye, Wanli Ouyang, and Tong He. Point cloud matters: Rethinking the impact of different observation spaces on robot learning. *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*, 2025. 2
- [97] Zi-Xin Zou, Zhipeng Yu, Yuan-Chen Guo, Yangguang Li, Ding Liang, Yan-Pei Cao, and Song-Hai Zhang. Triplane meets gaussian splatting: Fast and generalizable single-view 3d reconstruction with transformers. *arXiv preprint arXiv:2312.09147*, 2023. 3
- [98] Sicheng Zuo, Wenzhao Zheng, Yuanhui Huang, Jie Zhou, and Jiwen Lu. Gaussianworld: Gaussian world model for streaming 3d occupancy prediction. In *Proceedings of Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6772–6781, 2025. 2