

WonderPlay: Dynamic 3D Scene Generation from a Single Image and Actions

Zizhang Li^{1*} Hong-Xing Yu^{1*} Wei Liu¹
 Yin Yang² Charles Herrmann¹ Gordon Wetzstein¹ Jiajun Wu¹
¹Stanford University ²University of Utah

<https://kyleleey.github.io/WonderPlay/>

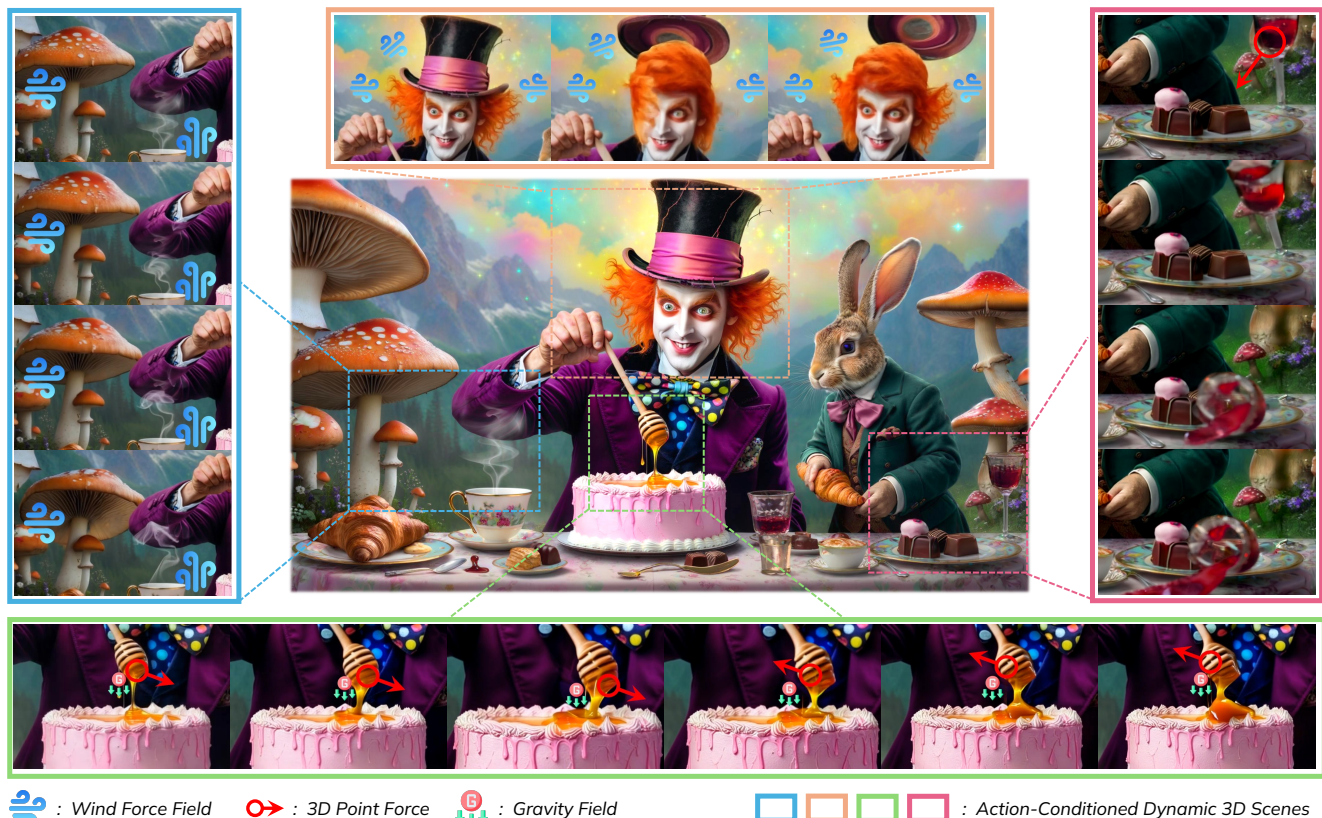


Figure 1. We propose **WonderPlay**, a framework that takes a single image and **actions** as inputs, and then generates dynamic 3D scenes that depict the consequence of the actions. WonderPlay allows users to interact with various scenes of diverse physical materials, e.g., the hat and wine glass (rigid body), the hair (thin strands), the steam (gas), the mushroom (elastic), honey (liquid), and more. See <https://kyleleey.github.io/WonderPlay/> for interactive video results.

Abstract

WonderPlay is a novel framework integrating physics simulation with video generation for generating action-conditioned dynamic 3D scenes from a single image. While prior works are restricted to rigid body or simple elastic

dynamics, WonderPlay features a hybrid generative simulator to synthesize a wide range of 3D dynamics. The hybrid generative simulator first uses a physics solver to simulate coarse 3D dynamics, which subsequently conditions a video generator to produce a video with finer, more realistic motion. The generated video is then used to update the simulated dynamic 3D scene, closing the loop between the physics solver and the video generator. This approach

*Equal contribution

enables intuitive user control to be combined with the accurate dynamics of physics-based simulators and the expressivity of diffusion-based video generators. Experimental results demonstrate that WonderPlay enables users to interact with various scenes of diverse content, including cloth, sand, snow, liquid, smoke, elastic, and rigid bodies – all using a single image input. Code will be made public.

1. Introduction

Recent years have seen rapid progress in image, video, and 3D and 4D scene generation, culminating in models that achieve great visual quality [6, 46] and dynamic realism [11]. This has directly motivated recent interest in generative world models, which, beyond their relevance to AR/VR and embodied AI [51], can also be created and explored as standalone experiences. However, while significant efforts have been devoted to enhancing the generation quality [12, 14], relatively little attention has been paid to enabling action-based interaction. In this work, we study **action-conditioned dynamic 3D scene generation from a single image**: given an input image and a 3D action, such as wind or a point force, we aim to generate the resulting dynamic 3D scene in the near future. In particular, we focus on three types of actions: gravity, force fields like wind, and point forces like pushes or pulls.

Existing methods often rely exclusively on physics simulation for computing dynamic 3D scenes given user action input [59, 67, 73]. These methods face two critical limitations. First, they require *accurate physics solvers for all types of dynamics* involved in the scene. Nevertheless, accurate physics solvers such as solid-fluid two-way coupling [16] still remain an open problem. Second, they require *full reconstruction of physical states from limited observations*. However, reconstructing complete physical states for materials like snow, sand, cloth, and fluids from a single image is often infeasible. Consequently, existing methods are constrained to a narrow range of dynamics types, primarily rigid body dynamics [43] and simple elasticity [59, 73].

This motivates us to incorporate data priors from video generation models [8, 11, 70], which are trained on extensive real-world videos of diverse physical phenomena. However, video generation models cannot accept precise 3D actions as inputs and simulate the resulting dynamics. In this work, we reconsider the relationship between physics simulators and video generation models for action-conditioned dynamic 3D scene generation. Our novel framework, *WonderPlay*, enables users to interact with 3D scenes encompassing diverse materials—including rigid bodies, cloth, liquids, gases, and granular substances—from a single input image, as shown in Figure 1.

Our core technical idea is a **hybrid generative simulator**. First, we let the physics simulator provide a coarse sim-

ulation of action-induced dynamic consequences to a video generator. Conditioned on the coarse simulation, the video generator synthesizes a video with realistic motion. Finally, the synthesized video is used to update the coarse simulation.

In the conditional video generation stage, we explore a novel strategy to optimally use the simulator conditioning signal: a motion–appearance *bimodal control* scheme, designed to improve the quality and realism of the dynamics in the generated video. Additionally, to reduce the video generator hallucination in simulator-trustable spatial regions such as static backgrounds, we introduce a *spatially varying* masking scheme for the bimodal control.

In summary, our contributions are three-fold:

- We tackle the challenging problem of single-image, action-conditioned dynamic 3D scene generation with diverse physical materials.
- We propose WonderPlay, featuring a hybrid generative simulator that integrates a physics solver and video diffusion to acquire both high simulation fidelity in response to actions and high visual quality.
- We demonstrate that WonderPlay significantly outperforms both pure physics-based methods and adapted video generation models in terms of visual quality and physical plausibility under various interactions.

2. Related Work

Action-conditioned dynamic scene generation. Early work on generating action-conditioned dynamic scenes approached the problem by extracting modal bases of vibrating objects in 2D image space [19, 20], essentially representing motion as a series of vibrations with different frequencies and intensities. Following the advent of generative diffusion modeling [25, 56, 57], this approach was later extended by retaining the same motion representation but generating the modal basis with a diffusion model [38]. While this representation can be effective for motions similar to vibrations, modal basis functions struggle to represent more general motions, prompting the emergence of an alternative line of research that explicitly uses physics solvers. For example, PhysGen [43] focused on the 2D domain using a rigid-body physics solver to handle colliding objects.

Recently, several physics-based approaches have been developed to synthesize dynamic 3D scenes [2, 15, 17, 28, 34, 39, 41, 67, 73, 78]. However, due to the requirements of physics solvers, all of these techniques require complete 3D geometric reconstructions of the scene, requiring complex, multi-view captures. For example, Virtual Elastic Objects [17] reconstructs the geometry, appearances, and physical parameters of elastic objects from a multi-view capture setup. Later work, such as PAC-NeRF [36], PhysGaussian [67], and PhysDreamer [73], integrates physics-based

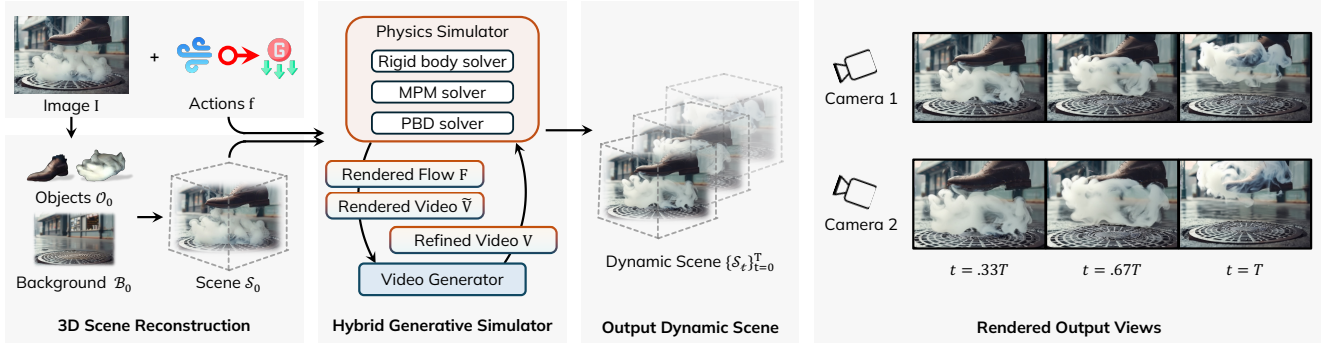


Figure 2. **Overview of WonderPlay.** Given a single image, we first reconstruct the 3D scene and estimate material properties. Then our hybrid generative simulator uses physics solver and input actions to infer coarse 3D dynamics. The simulated appearance and motion signals are used to condition the video generator through spatially varying bimodal control to synthesize the realistic motion. The dynamic 3D scene is refined using the synthesized video, finishing the hybrid generative simulation.

simulations with NeRF or 3D Gaussians from multi-view reconstruction. A concurrent approach, PhysMotion [59], is closest to our work. Both approaches take a single image as input and use a combination of a 3D physics solver and a video generation model. Unlike ours, however, PhysMotion [59] relies on a physics solver to compute the dynamics for the entire scene, only using the video generator to refine the appearance. Due to the restrictive assumption that the physics solver will specify all the dynamics, PhysMotion is limited to rigid and elastic dynamics. In contrast, WonderPlay uses both a physics solver and the video generator to compute dynamics, enabling realistic action-conditioned dynamics for various types of physical phenomena, and generates a dynamic 3D scene as opposed to a video.

Controllable video generation. In recent years, video generation has rapidly improved, making significant strides in both visual quality and realistic dynamics [8, 10, 11, 22, 26, 54, 70]. Recent video generation methods, such as Sora [11], have demonstrated great promise in generating diverse real-world physical phenomena. However, despite this promise, these models are conditioned with text and/or images and lack controllability regarding general actions and other physical inputs. While there has been considerable work on adding controls to video models, most of this work has focused on camera control [24, 27, 58, 66, 68, 77] and various types of motion control [18, 21, 37, 49, 50, 53, 62, 65, 71, 74], including drag-based, trajectory-based, and optical flow-based approaches. However, most of these motion control models require the resulting dynamics of an action as input to generate a trajectory-following video. Concurrently, Motion Prompting [21] uses temporally sparse trajectories as conditioning to generate videos that adhere to an initial trajectory and then continue using a generative video prior. Nevertheless, many actions, such as those involving fluids or wind, are difficult or impossible to represent as trajectory signals. In contrast, we aim to take physics-based 3D actions as input and model a dynamic 3D

scene, rather than just a video.

World models. Along with the rise in video models, there has been a growing interest in interactive world models [23], which recurrently generate world states from prior state and actions. While this area has seen considerable research, the focus has been on video game domains [12, 14, 60] due to the data availability of action-video pairs. Consequently, both the generated worlds and the actions considered are centered around those found in video games. Only a few concurrent works have explored this for the real world (e.g., [1, 7]) and none provides interaction beyond camera control or text. In contrast, we focus on physics-based actions in realistic worlds.

Dynamic 3D scene generation. Text-conditioned generation of 3D scenes with motion has primarily been tackled by distilling video generation into dynamic 3D representations [55]. Most of the existing works focus on objects or simple scenes composed of a few objects [4, 5, 40, 52, 75], while recent methods also attempt to deal with more complex scenes including nature [35, 42]. Recent work [63, 64, 76] has focused on training generators which model 4D itself, conditioning on time and camera pose. Yet, they do not have the ability to simulate dynamics in response to actions.

3. WonderPlay

Formulation. Our goal is action-conditioned 3D scene dynamics synthesis. The input is a single image I and actions. We model actions as three types of forces: gravity \mathbf{f}_g , 3D force fields $\mathbf{f}_w(x, y, z, t)$ such as wind, and 3D point forces $\mathbf{f}_p(t)$ which is defined on a point of an object. The output is a dynamic 3D scene $\{\mathcal{S}_t\}_{t=0}^T$ that is the consequence of applying the actions to the input scene, where \mathcal{S}_0 denotes our initial 3D scene representation recovered from the input image I , and T denotes the total simulation time steps.

Overview. We aim to simulate the dynamics of diverse ma-

materials including rigid, elastic, cloth, smoke, liquid, granular, and their interactions. To this end, we propose WonderPlay. As illustrated in Figure 2, we first reconstruct the 3D scene \mathcal{S}_0 from the input image \mathbf{I} (top left of Figure 2). Our main technical innovation is the hybrid generative simulator (middle left of Figure 2). It takes the 3D scene \mathcal{S}_0 and the actions as input and predicts the 3D dynamics $\{\mathcal{S}_t\}_{t=1}^T$ (middle right of Figure 2).

3.1. 3D Scene Reconstruction

Our 3D scene representation $\mathcal{S}_t = \mathcal{B}_t \cup \mathcal{O}_t$ consists of a background \mathcal{B}_t and objects \mathcal{O}_t at a timestep t . Our first step is to reconstruct/generate \mathcal{S}_0 from the input image \mathbf{I} . We reconstruct the 3D background \mathcal{B}_0 and the 3D objects \mathcal{O}_0 separately to jointly form $\mathcal{S}_0 = \mathcal{B}_0 \cup \mathcal{O}_0$.

Background. We represent the background with Fast Layered Gaussian Surfels (FLAGS) [72]. Formally, the background $\mathcal{B}_t = \{\mathbf{p}^B, \mathbf{q}^B, \mathbf{s}^B, \mathbf{o}^B, \mathbf{c}_t^B\}$ consists of N_B Gaussian surfels, parameterized by 3D spatial positions $\mathbf{p}^B \in \mathbb{R}^{3N_B}$, orientation quaternions \mathbf{q}^B , scales \mathbf{s}^B , opacities \mathbf{o}^B , and view-independent RGB colors \mathbf{c}_t^B . We detail the reconstruction of background in Appendix B.2 and treat it as a static boundary in simulation.

Objects. An “object” in WonderPlay refers to a dynamic entity we simulate in the physics solver, including rigid object, cloth, granular material, and fluids. To represent a simulatable object that is compatible with our physics solvers, we build a simulation-ready representation on top of the Gaussian surfels by adding connectivity to them, turning them into “topological Gaussian surfels”. Formally, the topological Gaussian surfels consist of N_O Gaussian surfels with edges and velocities, $\mathcal{O}_t = \{\mathbf{E}, \mathbf{v}_t, \mathbf{p}_t^O, \mathbf{q}_t^O, \mathbf{s}_t^O, \mathbf{o}_t^O, \mathbf{c}_t^O\}$, where the edge matrix $\mathbf{E} \in \{0, 1\}^{N_O \times N_O}$ indicates the topological connectivity of the surfels, and $\mathbf{v}_t \in \mathbb{R}^{3N_O}$ denotes the velocity.

We create the initial topological Gaussian surfels \mathcal{O}_0 by first generating an object mesh from an image segment of the object using an image-to-mesh model InstantMesh [69]. Then, we bind a Gaussian surfel to each of the mesh vertices. We detail this process in Appendix B.3.

Materials. Besides the geometry and appearance representation \mathcal{O} , an object also has material properties \mathbf{m} . The definition of object material depends on the object type, which follows a 6-way classification: rigid, elastic, cloth, smoke, liquid, and granular. We detail the material properties in Appendix B.4.

3.2. Hybrid Generative Simulator

Main idea. The reconstructed scene geometry \mathcal{S}_0 and estimated material properties \mathbf{m} are inherently inaccurate and incomplete, and accurate physics solvers for all materials and their complex interactions are still an open prob-

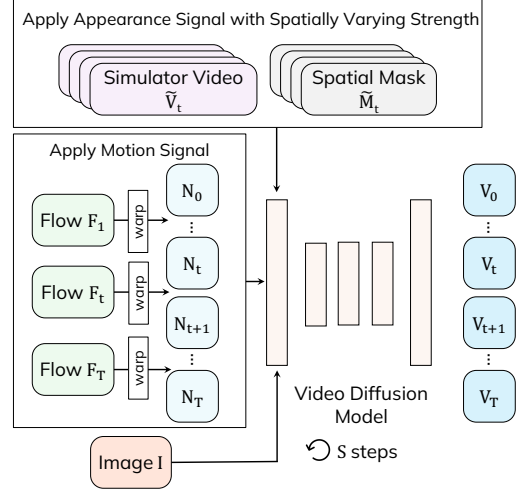


Figure 3. Illustration on our spatially varying bimodal control, which drives the video generator with input image \mathbf{I} , pixel-space flow \mathbf{F} and simulation rendered \tilde{V} .

lem. Therefore, existing methods are limited to simple rigid/elastic simulations [43, 59, 73]. Our main idea to address this challenge is extracting the dynamics knowledge from a video generator which has been trained on numerous videos of real-world physics.

In particular, we use physics solvers to estimate a coarse and incomplete dynamic scene $\{\tilde{\mathcal{S}}_t\}_{t=1}^T$ given initial scene \mathcal{S}_0 and actions $\mathbf{f}_g, \mathbf{f}_w, \mathbf{f}_p$. The coarse dynamic scene is used to drive the video generator to synthesize a video V that has realistic dynamics. We obtain the output dynamic scene $\{\mathcal{S}_t\}_{t=1}^T$ by updating $\{\tilde{\mathcal{S}}_t\}_{t=1}^T$ to match the video V through differentiable rendering.

Physics solvers. At each simulation time step, a physics solver takes the current scene $\tilde{\mathcal{S}}_t$ and forces $\mathbf{f}_g, \mathbf{f}_w(t), \mathbf{f}_p(t)$ as input, and solves for the object dynamics attributes including the velocity \mathbf{v}_{t+1} , position \mathbf{p}_{t+1}^O , and orientation \mathbf{q}_{t+1}^O at the next time step:

$$\mathbf{v}_{t+1}, \mathbf{p}_{t+1}^O, \mathbf{q}_{t+1}^O = \text{solver}(\tilde{\mathcal{S}}_t, \mathbf{f}_g, \mathbf{f}_w(t), \mathbf{f}_p(t)), \quad (1)$$

where $\tilde{\mathcal{S}}_0 = \mathcal{S}_0$. Then, we construct the coarse scene at the next time step $\tilde{\mathcal{S}}_{t+1}$ by

$$\tilde{\mathcal{S}}_{t+1} = \mathcal{B}_0 \cup \{\mathbf{E}, \mathbf{v}_{t+1}, \mathbf{p}_{t+1}^O, \mathbf{q}_{t+1}^O, \mathbf{s}_0^O, \mathbf{o}_0^O, \mathbf{c}_0^O\}, \quad (2)$$

where we keep all non-dynamics attributes the same as \mathcal{S}_0 . To compute the dynamics attributes of various materials, we employ multiple types of physics solvers. These solvers are coupled to tackle multi-physics scenes, e.g., fluid and rigid as shown in Figure 2. Please see details for different solvers in Appendix B.4.

Conditioning the video generator. Given the coarse dynamic scene $\{\tilde{\mathcal{S}}_t\}_{t=0}^T$, we condition a video generator to synthesize a video $\mathbf{V} \in \mathbb{R}^{(T+1) \times H \times W \times 3}$ that has more detailed motion while adhering to the action consequence




Figure 4. **Qualitative comparisons** between WonderPlay (ours) and the baseline methods. The top row shows the input images, actions, and the texts describing the actions for CogVideoX [70].

depicted by the coarse dynamic scene. To this end, we introduce a *bimodal control* scheme that uses two modalities for control: motion (represented by flow) and appearance (RGB). In particular,

$$\mathbf{V} = g(\mathbf{F}, \tilde{\mathbf{V}}, \mathbf{I}), \quad (3)$$

where g denotes the video generator, $\mathbf{F} \in \mathbb{R}^{T \times H \times W \times 2}$ denotes the pixel-space flow rendered using the velocity $\{\mathbf{v}_t\}_{t=1}^T$, $\tilde{\mathbf{V}} \in \mathbb{R}^{(T+1) \times H \times W \times 3}$ denotes the video rendered from the coarse scene $\{\tilde{\mathcal{S}}_t\}_{t=0}^T$, and \mathbf{I} denotes the input image. We show an illustration in Figure 3.

Motion control. We leverage a pre-trained motion-controlled image-to-video diffusion model, Go-with-the-Flow [13], as our g . The motion control is based on noise warping. In short, instead of using an unstructured random Gaussian noise distribution, it uses a warping-based structured noise $\mathbf{N}(\mathbf{F}) \in \mathbb{R}^{(T+1) \times H \times W \times 3}$. $\mathbf{N}(\mathbf{F})$ is created by first sampling a random Gaussian $\mathbf{N}_0 \in \mathbb{R}^{H \times W \times 3}$ and then iteratively doing warping such that $\mathbf{N}_{t+1} = \text{warp}(\mathbf{N}_t, \mathbf{F}_{t+1})$ where $\mathbf{F}_{t+1} \in \mathbb{R}^{H \times W \times 2}$ denotes the flow at $t+1$. The structured noise $\mathbf{N}(\mathbf{F})$ is then fused with some random noise to improve visual quality,

controlled by a degradation factor γ [13].¹

RGB control. To incorporate additional control with RGB frames $\tilde{\mathbf{V}}$, we use SDEdit [45]. Specifically, the diffusion-based generation process is gradually denoising $\mathbf{N}(\mathbf{F})$, such that $\mathbf{V}_{s-1} = \text{Denoise}(\mathbf{V}_s, s)$ where $s = S, S-1, \dots, 1$ denotes the diffusion timestep, $\mathbf{V}_S = \mathbf{N}(\mathbf{F})$ is the initial noise, and the generated video is given by $\mathbf{V} = \mathbf{V}_0$. We control this process by skipping first several steps and directly starting the denoising from step $s_1 < S$ with

$$\mathbf{V}_{s_1} = \alpha_{s_1} \tilde{\mathbf{V}} + \sqrt{1 - \alpha_{s_1}^2} \mathbf{N}(\mathbf{F}), \quad (4)$$

where α_i denotes the diffusion coefficient at timestep i . This has been shown to control the main content of the generation, while allowing details to be synthesized [45].

Discussion. With our bimodal control, we pass the coarse motion and appearance information from the physics simulator to the video generator. This allows not only gener-

¹The video model is a latent diffusion model [70] which downsamples the spatiotemporal dimensions H, W, T and upsamples feature dimension from 3 to C . But for notational simplicity, we do not distinguish the latent space from the pixel space.

ating more realistic motion, but also fixes appearance artifacts caused by imperfect 3D scene reconstruction and the lack of lighting information to resolve appearance changes. The question is how much we trust the generator to overwrite the coarse information. Intuitively, this is dictated by s_1 : If s_1 is close to 0, then the generator only modifies the coarse video $\tilde{\mathbf{V}}$ a bit; if s_1 is close to S , then it can overwrite $\tilde{\mathbf{V}}$ more and hallucinate new contents. Thus, s_1 positively corresponds to the responsibility of the video generator.

Spatially varying responsibility. The responsibility of the video generator is inherently uneven across spatial regions in every frame. For example, most of our background remains static in the dynamic process, which we want to trust the simulator output $\tilde{\mathbf{V}}$ more, rather than the video generator, because the video generator may hallucinate incorrect details such as ghost objects. To this end, we introduce the spatially varying bimodal control.

In this work, we consider two responsibility levels in our spatially varying bimodal control for the background and the dynamic objects, respectively, such that we set a lower responsibility $s_2 < s_1$ of the video generator on modifying the background. Specifically, at the step s_2 , we compute

$$\hat{\mathbf{V}}_{s_2} = \mathbf{M} \odot \mathbf{V}_{s_2} + (\mathbf{1} - \mathbf{M}) \odot (\alpha_{s_2} \tilde{\mathbf{V}} + \sqrt{1 - \alpha_{s_2}^2} \mathbf{N}(\mathbf{F})), \quad (5)$$

where \mathbf{V}_{s_2} is computed from gradually denoising \mathbf{V}_{s_1} . $\mathbf{M} \in \{0, 1\}^{(T+1) \times H \times W \times 3}$ denotes the binary mask that uses 1 to mark a pixel of dynamic objects and 0 to mark a pixel of the background, which is rendered from the coarse scene. After this computation, we set $\mathbf{V}_{s_2} \leftarrow \hat{\mathbf{V}}_{s_2}$ and continue the denoising to generate \mathbf{V} .

Updating scene dynamics. Finally, we use the generated video \mathbf{V} as a supervision to update the coarse dynamic scene $\{\tilde{\mathcal{S}}_t\}_{t=0}^T$. This is done by minimizing a photometric L1 loss: $\min_{\{c_t^p, \mathcal{O}_t\}_{t=0}^T} \|\mathbf{V} - \tilde{\mathbf{V}}\|_1$ over the foreground object’s motion trajectory and appearance $\{\mathcal{O}_t\}_{t=0}^T$. We also update the background color c_t^b for shading effects. This optimization yields the final dynamic scene $\{\mathcal{S}_t\}_{t=0}^T$.

4. Experiments

Implementation details. For physics simulation, we adopt the Genesis [3] framework, which unifies several different physics solvers. For all scenes, we run physical simulation for 960 steps and render one frame for each 20 steps. We include additional implementation details in Appendix B.1.

Baselines. We compare against two types of baselines for action-conditioned 3d dynamic scene generation: physics-based and conditional video generation methods. For physics-based methods, we compare with PhysGen [43] and PhysGaussian [67]. PhysGen decomposes an image into 2D rigid bodies and run rigid simulation given certain action. PhysGaussian models the 3D scene as elastic objects

	Physics Plausibility	Motion Fidelity	Visual Quality
Over PhysGen [43]	78.0%	78.0%	80.1%
Over PhysGaussian [67]	80.2%	81.2%	85.2%
Over Tora [74]	77.0%	72.0%	71.0%
Over CogVideoX-I2V [70]	80.2%	73.0%	74.6%

Table 1. Human study 2AFC results of favor rate of WonderPlay (Ours) over baseline methods.

Methods	Imaging (↑)	Aesthetic (↑)	Motion (↑)	Consistency (↑)	PhysReal (↑)
PhysGen	<u>0.692</u>	0.593	0.992	0.212	0.545
PhysGaussian	0.492	0.564	<u>0.994</u>	0.206	0.350
CogVideoX	0.686	0.574	0.993	0.219	0.670
Tora	0.644	0.620	0.992	0.210	0.530
Ours	0.695	<u>0.610</u>	0.995	<u>0.217</u>	0.700

Table 2. Quantitative comparison to baselines on 15 scenes.

with the MPM [30] framework. Since PhysGen only requires a single image as the input, we directly follow their preprocessing for image decomposition. PhysGaussian requires multiview images to reconstruct the underlying scene first, so we provide as input our reconstructed 3D scene and then run simulation with given actions. For conditional video generation methods, we compare against two methods: CogVideoX-I2V [70] with text prompts and Tora [74] with drag-based conditioning. For Tora, we use the trajectories from our simulation as the drag input.

Metrics. We render videos from the input viewpoint to compute quantitative metrics. We adopt the imaging, aesthetic, motion quality, and consistency metrics from VBench [29]. We also adopt the GPT-4o-based physical realism metric [15]. We curate 15 examples, including 7 real photos and 8 realistic synthetic images, covering diverse types of scenes contents including cloth, rigid body, elastic objects, liquid, gas, granular substance, etc.

4.1. Results

Comparison to baselines. We show side-by-side comparisons on two scenes in Figure 4. The top row shows input images, actions (gravity for duck dropping, force to pull the red boat towards the right) and the text prompt for CogVideoX-I2V [70], followed by the action-conditioned generated dynamics from our method and the baselines.

Despite their ability to produce plausible visual quality, video generation methods struggle to adhere to the actions. In the duck-dropping-into-water scene, Tora [74] submerges the duck under the water and then changes its shape after it re-emerges. CogVideoX-I2V struggles to generate realistic dynamics for the duck’s drop and adds undesirable dynamics by moving the duck to the left. Both models also struggle with the boats scene. Tora completely alters the scene mid-video, while CogVideoX-I2V fails to generate meaningful dynamics.

As for the physics-based method, PhysGen [43] is limited to rigid body simulation in 2D space, making it hard to




Figure 5. **Qualitative results** of the proposed WonderPlay. In the left column we show the input scene image and actions, where G , W , P indicate gravity action, wind field action and 3D point force action, respectively.



Figure 6. **Different actions** on the same scene. WonderPlay supports to use different 3D actions on the same scene. Here we show four different scenes and the corresponding dynamics from two different actions within each scene.

handle scene with complex materials such as water. Phys-Gaussian [43] typically requires multiview images and as a result, struggles to produce a reasonable 3D representation with only one input image. Also due to the lack of

a complete 3D physical state, both physics-based methods fail to properly handle the shading effect in the boats scene and make the reflections move with the boats. Our method, in contrast, offers the advantages of both physical simula-



Figure 7. **Ablation** on hybrid generative simulator. Top row: Coarse simulation (i.e., only physics solver is used without video generator for refinement). Bottom row: Refined dynamic scene.



Figure 8. **Ablation** on the motion signal and the appearance signal to condition the video generator.

tion and video generation: the physical simulation handles a wide range of materials and ensures the desired dynamics, and the video generation model provides visual realism by successfully synthesizing water waves and bubbles surrounding the duck, as well as the following reflections. Also shown in Table 2, WonderPlay (ours) achieve the best or second-best performance across all metrics, showing strong motion quality, visual quality, and physical plausibility.

User study. To evaluate the generation results with human preference, we recruit 200 participants and conduct a user study. We employ a Two-alternative Forced Choice (2AFC) protocol. Each participant evaluates 10 scenes. The participants view an action description alongside a randomly ordered side-by-side comparison video: one from our method and one from a baseline. Participants then select which video demonstrates superior performance in one of three criteria: *physics plausibility* which measures the correctness of the predicted motion in response to the action, *motion fidelity* that reflects the quality and naturalness of the generated motion, and *visual quality*.

We show the averaged results on all scenes in Table 1. In comparison to all baselines, about 70% to 80% of the participants prefer WonderPlay (ours) across all three aspects, proving the superior performance of combining the physical simulator and video generator for dynamics with fidelity in response to actions and realistic visual appearance.

Diverse scenes and materials. In Figure 5, we present the generated dynamic 3D scenes on a variety of input images with diverse actions. It is important to note that achieving

realistic visual quality in simulations of complex materials from a single image input with limited physical state information is extremely challenging. However, with the aid of the video generator, the sticky jam appears vivid as it pours onto the cake, and the river waves look natural in response to the boat’s movement. Notably, the underlying physical simulator ensures that all dynamics follow the input actions. For example, the roses are initially blown to the right by the wind and then move back due to their elasticity.

Condition on different actions. A significant advantage of our method is that it enables generating different interactions with different actions in the same scene. In Figure 6, we present four scenes, each with two different actions and their corresponding output dynamic 3D scenes.

4.2. Ablation on Hybrid Generative Simulator

In the following we discuss an ablation study on the hybrid generative simulator. We leave quantitative numbers and further ablation in Appendix A.

Video generator refines both dynamics and appearance.

In Figure 7, we compare a dynamic scene created solely with the physics simulator, i.e., the coarse simulation (top row), and the refined dynamic scene created by our full model (bottom row). In the coarse simulation, we observe unrealistic motion: the motion of the smoke looks too sticky due to numerical viscosity that exists for almost all fluid solvers. The video model refines it so that the fluid motion looks smooth with swirls. There are also appearance artifacts in the coarse simulation such as the grainy smoke, where the video model can also refine them.

Both signals to condition the video generator are necessary.

To demonstrate the benefit of conditioning the video generator on both motion and appearance signals, we show ablation results in Figure 8. The top row shows the synthesized video from our full model with both signals; “w/o RGB” uses motion but no appearance; and “w/o flow” uses appearance but no motion. Using only RGB conditioning (“w/o flow”), the video model fails to retain or improve detailed dynamics in the sand grains. Using only a motion signal (“w/o RGB”) leads to unexpected hallucinations beyond user action input, e.g., it hallucinates a pile of sand standing in the back and the background texture unexpectedly changes. In contrast, using both signals produces the best results.

5. Conclusion

In this work, we propose WonderPlay, a novel framework for action-conditioned dynamic 3D scene generation from a single image. WonderPlay features a hybrid generative simulator for simulation fidelity and visual quality. We showcase superior performance of WonderPlay on diverse scenes with various interactions.

Acknowledgments. We thank Guandao Yang, Yunzhi Zhang, and Zhehao Li for the comments and fruitful discussions, and Hadi Alzayer for help reviewing the draft. This work is in part supported by the Stanford Institute for Human-Centered AI (HAI), the Okawa Foundation Research Grant, NSF RI #2211258 and #2338203, ONR YIP N00014-24-1-2117, and ONR MURI N00014-22-1-2740.

References

- [1] Niket Agarwal, Arslan Ali, Maciej Bala, Yogesh Balaji, Erik Barker, Tiffany Cai, Prithvijit Chattopadhyay, Yongxin Chen, Yin Cui, Yifan Ding, et al. Cosmos world foundation model platform for physical ai. *arXiv preprint arXiv:2501.03575*, 2025. 3
- [2] Luca Savant Aira, Antonio Montanaro, Emanuele Aiello, Diego Valsesia, and Enrico Magli. Motioncraft: Physics-based zero-shot video generation. *arXiv preprint arXiv:2405.13557*, 2024. 2
- [3] Genesis Authors. Genesis: A universal and generative physics engine for robotics and beyond, 2024. 6
- [4] Sherwin Bahmani, Xian Liu, Wang Yifan, Ivan Skorokhodov, Victor Rong, Ziwei Liu, Xihui Liu, Jeong Joon Park, Sergey Tulyakov, Gordon Wetzstein, Andrea Tagliasacchi, and David B. Lindell. Tc4d: Trajectory-conditioned text-to-4d generation. *arXiv*, 2024. 3
- [5] Sherwin Bahmani, Ivan Skorokhodov, Victor Rong, Gordon Wetzstein, Leonidas Guibas, Peter Wonka, Sergey Tulyakov, Jeong Joon Park, Andrea Tagliasacchi, and David B Lindell. 4d-fy: Text-to-4d generation using hybrid score distillation sampling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7996–8006, 2024. 3
- [6] Jason Baldrige, Jakob Bauer, Mukul Bhutani, Nicole Brich-tova, Andrew Bunner, Kelvin Chan, Yichang Chen, Sander Dieleman, Yuqing Du, Zach Eaton-Rosen, et al. Imagen 3. *arXiv preprint arXiv:2408.07009*, 2024. 2
- [7] Amir Bar, Gaoyue Zhou, Danny Tran, Trevor Darrell, and Yann LeCun. Navigation world models. *arXiv preprint arXiv:2412.03572*, 2024. 3
- [8] Bar-Tal et al. Lumiere: A space-time diffusion model for video generation. *arXiv:2401.12945*, 2024. 2, 3
- [9] Jan Bender, Matthias Müller, and Miles Macklin. Position-based simulation methods in computer graphics. In *Eurographics (tutorials)*, 2015. S2
- [10] Andreas Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendelevitch, Maciej Kilian, Dominik Lorenz, Yam Levi, Zion English, Vikram Voleti, Adam Letts, et al. Stable video diffusion: Scaling latent video diffusion models to large datasets. *arXiv preprint arXiv:2311.15127*, 2023. 3
- [11] Tim Brooks, Bill Peebles, Connor Homes, Will DePue, Yufei Guo, Li Jing, David Schnurr, Joe Taylor, Troy Luhman, Eric Luhman, Clarence Ng, Ricky Wang, and Aditya Ramesh. Video generation models as world simulators, 2024. 2, 3
- [12] Jake Bruce, Michael D Dennis, Ashley Edwards, Jack Parker-Holder, Yuge Shi, Edward Hughes, Matthew Lai, Aditi Mavalankar, Richie Steigerwald, Chris Apps, et al. Genie: Generative interactive environments. In *Forty-first International Conference on Machine Learning*, 2024. 2, 3
- [13] Ryan Burgert, Yuancheng Xu, Wenqi Xian, Oliver Pilarski, Pascal Clausen, Mingming He, Li Ma, Yitong Deng, Lingxiao Li, Mohsen Mousavi, et al. Go-with-the-flow: Motion-controllable video diffusion models using real-time warped noise. *arXiv preprint arXiv:2501.08331*, 2025. 5
- [14] Haoxuan Che, Xuanhua He, Quande Liu, Cheng Jin, and Hao Chen. Gamegen-x: Interactive open-world game video generation. *arXiv preprint arXiv:2411.00769*, 2024. 2, 3
- [15] Boyuan Chen, Hanxiao Jiang, Shaowei Liu, Saurabh Gupta, Yunzhu Li, Hao Zhao, and Shenlong Wang. Physgen3d: Crafting a miniature interactive world from a single image. *CVPR*, 2025. 2, 6
- [16] Duowen Chen, Zhiqi Li, Junwei Zhou, Fan Feng, Tao Du, and Bo Zhu. Solid-fluid interaction on particle flow maps. *ACM Transactions on Graphics (TOG)*, 2024. 2
- [17] Hsiao-yu Chen, Edith Tretschk, Tuur Stuyck, Petr Kadlecek, Ladislav Kavan, Etienne Vouga, and Christoph Lassner. Virtual elastic objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 2
- [18] Tsai-Shien Chen, Chieh Hubert Lin, Hung-Yu Tseng, Tsung-Yi Lin, and Ming-Hsuan Yang. Motion-conditioned diffusion model for controllable video synthesis. *arXiv preprint arXiv:2304.14404*, 2023. 3
- [19] Abe Davis, Katherine L Bouman, Justin G Chen, Michael Rubinstein, Fredo Durand, and William T Freeman. Visual vibrometry: Estimating material properties from small motion in video. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015. 2
- [20] Abe Davis, Justin G Chen, and Frédo Durand. Image-space modal bases for plausible manipulation of objects in video. *ACM Transactions on Graphics (TOG)*, 2015. 2
- [21] Daniel Geng, Charles Herrmann, Junhwa Hur, Forrester Cole, Serena Zhang, Tobias Pfaff, Tatiana Lopez-Guevara, Carl Doersch, Yusuf Aytar, Michael Rubinstein, et al. Motion prompting: Controlling video generation with motion trajectories. *arXiv preprint arXiv:2412.02700*, 2024. 3
- [22] Rohit Girdhar, Mannat Singh, Andrew Brown, Quentin Duval, Samaneh Azadi, Sai Saketh Rambhatla, Akbar Shah, Xi Yin, Devi Parikh, and Ishan Misra. Emu video: Factorizing text-to-video generation by explicit image conditioning. *arXiv preprint arXiv:2311.10709*, 2023. 3
- [23] David Ha and Jürgen Schmidhuber. World models. *arXiv preprint arXiv:1803.10122*, 2018. 3
- [24] Hao He, Yinghao Xu, Yuwei Guo, Gordon Wetzstein, Bo Dai, Hongsheng Li, and Ceyuan Yang. Cameractrl: Enabling camera control for text-to-video generation, 2024. 3
- [25] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *NeurIPS*, 33:6840–6851, 2020. 2
- [26] Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, et al. Imagen video: High definition video generation with diffusion models. *arXiv preprint arXiv:2210.02303*, 2022. 3

- [27] Chen Hou, Guoqiang Wei, Yan Zeng, and Zhibo Chen. Training-free camera control for video generation. *arXiv preprint arXiv:2406.10126*, 2024. 3
- [28] Tianyu Huang, Haoze Zhang, Yihan Zeng, Zhilu Zhang, Hui Li, Wangmeng Zuo, and Rynson WH Lau. Dreamphysics: Learning physical properties of dynamic 3d gaussians with video diffusion priors. *arXiv preprint arXiv:2406.01476*, 2024. 2
- [29] Ziqi Huang, Yinan He, Jiashuo Yu, Fan Zhang, Chenyang Si, Yuming Jiang, Yuanhan Zhang, Tianxing Wu, Qingyang Jin, Nattapol Chanpaisit, et al. Vbench: Comprehensive benchmark suite for video generative models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21807–21818, 2024. 6
- [30] Chenfanfu Jiang, Craig Schroeder, Joseph Teran, Alexey Stomakhin, and Andrew Selle. The material point method for simulating continuum materials. *Acm siggraph 2016 courses*, 2016. 6, S2
- [31] Bingxin Ke, Anton Obukhov, Shengyu Huang, Nando Metzger, Rodrigo Caye Daudt, and Konrad Schindler. Repurposing diffusion-based image generators for monocular depth estimation. In *CVPR*, 2024. S1
- [32] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM TOG*, 2023. S1
- [33] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *ICCV*, 2023. S1
- [34] Simon Le Cleac’h, Hong-Xing Yu, Michelle Guo, Taylor Howell, Ruohan Gao, Jiajun Wu, Zachary Manchester, and Mac Schwager. Differentiable physics simulation of dynamics-augmented neural objects. *IEEE Robotics and Automation Letters*, 2023. 2
- [35] Yao-Chih Lee, Yi-Ting Chen, Andrew Wang, Ting-Hsuan Liao, Brandon Y Feng, and Jia-Bin Huang. Vividdream: Generating 3d scene with ambient dynamics. *arXiv preprint arXiv:2405.20334*, 2024. 3
- [36] Xuan Li, Yi-Ling Qiao, Peter Yichen Chen, Krishna Murthy Jatavallabhula, Ming Lin, Chenfanfu Jiang, and Chuang Gan. Pac-nerf: Physics augmented continuum neural radiance fields for geometry-agnostic system identification. *arXiv preprint arXiv:2303.05512*, 2023. 2
- [37] Yaowei Li, Xintao Wang, Zhaoyang Zhang, Zhouxia Wang, Ziyang Yuan, Liangbin Xie, Yuexian Zou, and Ying Shan. Image conductor: Precision control for interactive video synthesis. *arXiv preprint arXiv:2406.15339*, 2024. 3
- [38] Zhengqi Li, Richard Tucker, Noah Snavely, and Aleksander Holynski. Generative image dynamics. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024. 2
- [39] Jiajing Lin, Zhenzhong Wang, Shu Jiang, Yongjie Hou, and Min Jiang. Phys4dgen: A physics-driven framework for controllable and efficient 4d content generation from a single image. *arXiv preprint arXiv:2411.16800*, 2024. 2
- [40] Huan Ling, Seung Wook Kim, Antonio Torralba, Sanja Fidler, and Karsten Kreis. Align your gaussians: Text-to-4d with dynamic 3d gaussians and composed diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8576–8588, 2024. 3
- [41] Fangfu Liu, Hanyang Wang, Shunyu Yao, Shengjun Zhang, Jie Zhou, and Yueqi Duan. Physics3d: Learning physical properties of 3d gaussians via video diffusion. *arXiv preprint arXiv:2406.04338*, 2024. 2
- [42] Jinxiu Liu, Shaoheng Lin, Yinxiao Li, and Ming-Hsuan Yang. Dynamicscaler: Seamless and scalable video generation for panoramic scenes. *arXiv preprint arXiv:2412.11100*, 2024. 3
- [43] Shaowei Liu, Zhongzheng Ren, Saurabh Gupta, and Shenglong Wang. Physgen: Rigid-body physics-grounded image-to-video generation. In *ECCV*, 2024. 2, 4, 6, 7, S1
- [44] Miles Macklin and Matthias Müller. Position based fluids. *ACM Transactions on Graphics (TOG)*, 2013. S2
- [45] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. Sdedit: Guided image synthesis and editing with stochastic differential equations. *ICLR*, 2022. 5
- [46] Midjourney. <https://www.midjourney.com/>, 2023. 2
- [47] Matthias Müller, Bruno Heidelberger, Matthias Teschner, and Markus Gross. Meshless deformations based on shape matching. *ACM transactions on graphics (TOG)*, 2005. S2
- [48] Matthias Müller, Bruno Heidelberger, Marcus Hennix, and John Ratcliff. Position based dynamics. *Journal of Visual Communication and Image Representation*, 2007. S2
- [49] Koichi Namekata, Sherwin Bahmani, Ziyi Wu, Yash Kant, Igor Gilitschenski, and David B Lindell. Sg-i2v: Self-guided trajectory control in image-to-video generation. *arXiv preprint arXiv:2411.04989*, 2024. 3
- [50] Muyao Niu, Xiaodong Cun, Xintao Wang, Yong Zhang, Ying Shan, and Yinqiang Zheng. Mofa-video: Controllable image animation via generative motion field adaptations in frozen image-to-video diffusion model. *arXiv preprint arXiv:2405.20222*, 2024. 3
- [51] Mohammad Nomaan Qureshi, Sparsh Garg, Francisco Yandun, David Held, George Kantor, and Abhishesh Silwal. Splatsim: Zero-shot sim2real transfer of rgb manipulation policies using gaussian splatting. *arXiv preprint arXiv:2409.10161*, 2024. 2
- [52] Jiawei Ren, Liang Pan, Jiayang Tang, Chi Zhang, Ang Cao, Gang Zeng, and Ziwei Liu. Dreamgaussian4d: Generative 4d gaussian splatting. *arXiv preprint arXiv:2312.17142*, 2023. 3
- [53] Xiaoyu Shi, Zhaoyang Huang, Fu-Yun Wang, Weikang Bian, Dasong Li, Yi Zhang, Manyuan Zhang, Ka Chun Cheung, Simon See, Hongwei Qin, et al. Motion-i2v: Consistent and controllable image-to-video generation with explicit motion modeling. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–11, 2024. 3
- [54] Uriel Singer, Adam Polyak, Thomas Hayes, Xi Yin, Jie An, Songyang Zhang, Qiyuan Hu, Harry Yang, Oron Ashual, Oran Gafni, et al. Make-a-video: Text-to-video generation without text-video data. *arXiv preprint arXiv:2209.14792*, 2022. 3

- [55] Uriel Singer, Shelly Sheynin, Adam Polyak, Oron Ashual, Iurii Makarov, Filippos Kokkinos, Naman Goyal, Andrea Vedaldi, Devi Parikh, Justin Johnson, et al. Text-to-4d dynamic scene generation. *arXiv preprint arXiv:2301.11280*, 2023. 3
- [56] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *ICML*, 2015. 2
- [57] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *ICLR*, 2021. 2, S1
- [58] Wenqiang Sun, Shuo Chen, Fangfu Liu, Zilong Chen, Yueqi Duan, Jun Zhang, and Yikai Wang. Dimensionx: Create any 3d and 4d scenes from a single image with controllable video diffusion. *arXiv preprint arXiv:2411.04928*, 2024. 3
- [59] Xiyang Tan, Ying Jiang, Xuan Li, Zeshun Zong, Tianyi Xie, Yin Yang, and Chenfanfu Jiang. Physmotion: Physics-grounded dynamics from a single image. *arXiv preprint arXiv:2411.17189*, 2024. 2, 3, 4
- [60] Dani Valevski, Yaniv Leviathan, Moab Arar, and Shlomi Fruchter. Diffusion models are real-time game engines. *arXiv preprint arXiv:2408.14837*, 2024. 3
- [61] 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 (CVPR)*, 2024. S1
- [62] Zhouxia Wang, Ziyang Yuan, Xintao Wang, Yaowei Li, Tianshui Chen, Menghan Xia, Ping Luo, and Ying Shan. Motionctrl: A unified and flexible motion controller for video generation. In *ACM SIGGRAPH 2024 Conference Papers*, 2024. 3
- [63] Daniel Watson, Saurabh Saxena, Lala Li, Andrea Tagliasacchi, and David J Fleet. Controlling space and time with diffusion models. *arXiv preprint arXiv:2407.07860*, 2024. 3
- [64] Rundi Wu, Ruiqi Gao, Ben Poole, Alex Trevithick, Changxi Zheng, Jonathan T Barron, and Aleksander Holynski. Cat4d: Create anything in 4d with multi-view video diffusion models. *arXiv preprint arXiv:2411.18613*, 2024. 3
- [65] Weijia Wu, Zhuang Li, Yuchao Gu, Rui Zhao, Yefei He, David Junhao Zhang, Mike Zheng Shou, Yan Li, Tingting Gao, and Di Zhang. Draganything: Motion control for anything using entity representation. In *European Conference on Computer Vision*, 2024. 3
- [66] Zeqi Xiao, Yifan Zhou, Shuai Yang, and Xingang Pan. Video diffusion models are training-free motion interpreter and controller. *arXiv preprint arXiv:2405.14864*, 2024. 3
- [67] Tianyi Xie, Zeshun Zong, Yuxing Qiu, Xuan Li, Yutao Feng, Yin Yang, and Chenfanfu Jiang. Physgaussian: Physics-integrated 3d gaussians for generative dynamics. In *CVPR*, 2024. 2, 6, S1, S2
- [68] Dejjia Xu, Weili Nie, Chao Liu, Sifei Liu, Jan Kautz, Zhangyang Wang, and Arash Vahdat. Camco: Camera-controllable 3d-consistent image-to-video generation. *arXiv preprint arXiv:2406.02509*, 2024. 3
- [69] Jiale Xu, Weihao Cheng, Yiming Gao, Xintao Wang, Shenghua Gao, and Ying Shan. Instantmesh: Efficient 3d mesh generation from a single image with sparse-view large reconstruction models. *arXiv preprint arXiv:2404.07191*, 2024. 4, S1
- [70] Zhuoyi Yang, Jiayan Teng, Wendi Zheng, Ming Ding, Shiyu Huang, Jiazheng Xu, Yuanming Yang, Wenyi Hong, Xiaohan Zhang, Guanyu Feng, et al. Cogvideox: Text-to-video diffusion models with an expert transformer. *arXiv preprint arXiv:2408.06072*, 2024. 2, 3, 5, 6
- [71] Shengming Yin, Chenfei Wu, Jian Liang, Jie Shi, Houqiang Li, Gong Ming, and Nan Duan. Dragnuwa: Fine-grained control in video generation by integrating text, image, and trajectory. *arXiv preprint arXiv:2308.08089*, 2023. 3
- [72] Hong-Xing Yu, Haoyi Duan, Charles Herrmann, William T Freeman, and Jiajun Wu. Wonderworld: Interactive 3d scene generation from a single image. *arXiv preprint arXiv:2406.09394*, 2024. 4, S1
- [73] Tianyuan Zhang, Hong-Xing Yu, Rundi Wu, Brandon Y Feng, Changxi Zheng, Noah Snaveley, Jiajun Wu, and William T Freeman. Physdreamer: Physics-based interaction with 3d objects via video generation. In *ECCV*, 2024. 2, 4
- [74] Zhenghao Zhang, Junchao Liao, Menghao Li, Zuozhuo Dai, Bingxue Qiu, Siyu Zhu, Long Qin, and Weizhi Wang. Tora: Trajectory-oriented diffusion transformer for video generation. *arXiv preprint arXiv:2407.21705*, 2024. 3, 6
- [75] Yuyang Zhao, Zhiwen Yan, Enze Xie, Lanqing Hong, Zhen-guo Li, and Gim Hee Lee. Animate124: Animating one image to 4d dynamic scene. *arXiv preprint arXiv:2311.14603*, 2023. 3
- [76] Yuyang Zhao, Chung-Ching Lin, Kevin Lin, Zhiwen Yan, Linjie Li, Zhengyuan Yang, Jianfeng Wang, Gim Hee Lee, and Lijuan Wang. Genxd: Generating any 3d and 4d scenes. *arXiv preprint arXiv:2411.02319*, 2024. 3
- [77] Guangcong Zheng, Teng Li, Rui Jiang, Yehao Lu, Tao Wu, and Xi Li. Cami2v: Camera-controlled image-to-video diffusion model. *arXiv preprint arXiv:2410.15957*, 2024. 3
- [78] Licheng Zhong, Hong-Xing Yu, Jiajun Wu, and Yunzhu Li. Reconstruction and simulation of elastic objects with spring-mass 3d gaussians. In *European Conference on Computer Vision*, 2024. 2