

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 PHYSWORLD: FROM REAL VIDEOS TO WORLD MODELS OF DEFORMABLE OBJECTS VIA PHYSICS-AWARE DEMONSTRATION SYNTHESIS

Anonymous authors

Paper under double-blind review

ABSTRACT

Interactive world models that simulate object dynamics are crucial for robotics, VR, and AR. However, it remains a significant challenge to learn physics-consistent dynamics models from limited real-world video data, especially for deformable objects with spatially-varying physical properties. To overcome the challenge of data scarcity, we propose PhysWorld, a novel framework that utilizes a simulator to synthesize physically plausible and diverse demonstrations to learn efficient world models. Specifically, we first construct a physics-consistent digital twin within MPM simulator via constitutive model selection and global-to-local optimization of physical properties. Subsequently, we apply part-aware perturbations to the physical properties and generate various motion patterns for the digital twin, synthesizing extensive and diverse demonstrations. Finally, using these demonstrations, we train a lightweight GNN-based world model that is embedded with physical properties. The real video can be used to further refine the physical properties. PhysWorld achieves accurate and fast future predictions for various deformable objects, and also generalizes well to novel interactions. Experiments show that PhysWorld has competitive performance while enabling inference speeds 47 times faster than the recent state-of-the-art method, *i.e.*, PhysTwin. The code and pre-trained models will be publicly available.

1 INTRODUCTION

Humans possess the ability to predict object movements to achieve specific goals during interaction, which is developed through accumulated experience from observations and experiments. This skill is not only fundamental to human intelligence but also a crucial requirement for various technological applications, including robotics, virtual reality (VR), and augmented reality (AR). Thus, it becomes increasingly popular to construct world models for dynamics modeling of objects.

Existing works explore this mainly from learning-based and physics-based simulation perspectives. **First**, learning-based methods generally take neural networks such as Graph Neural Networks (GNN) (Sanchez-Gonzalez et al., 2020) and Multilayer Perceptron (MLP) (Zhu et al., 2024) for dynamics modeling. Such models can achieve real-time inference and be applicable to a variety of objects, including plasticine, cloth, and fluids. However, they rely on extensive training data represented by particle (Wang et al., 2023), mesh (Pfaff et al., 2020), and 3D point (Zhang et al., 2025b). Although the data can be derived from either simulators or real videos, the simulated data can be physically inconsistent with real-world ones, and sufficient real-world data are labor-intensive to acquire. Recently, AdaptiGraph (Zhang et al., 2024a) further introduces the physical property-conditioned GNN to adapt unseen objects during interaction, but its physics-inconsistent data synthesis and global physical parameters restrict applicability to objects with spatially varying properties. **Second**, physics-based simulation methods (Zhang et al., 2024c; Huang et al., 2024a; Zhong et al., 2024; Lin et al.) model object dynamics using established simulators such as the Material Point Method (MPM) (Jiang et al., 2016; Stomakhin et al., 2013; Bardenhagen et al., 2000) or the Mass-Spring System (MSS), often in conjunction with the optimization of physical properties of objects. For instance, recent PhysTwin (Jiang et al., 2025) leverages the way to construct digital twins of objects from sparse videos, enabling effective resimulation. Benefiting from the physics-

054 based prior knowledge embedded in the simulator, these methods can achieve relatively realistic
 055 simulations. Nevertheless, their real-time inference capability is still limited.
 056

057 It can be seen that large amounts of data or powerful simulators are required for the accuracy of
 058 world models, while lightweight modeling manners (*e.g.*, GNN) are necessary for the efficiency.
 059 When real-world observation sequences are short, it remains challenging to construct both accurate
 060 and fast world models for deformable objects. To address this issue, we propose to build a bridge
 061 between powerful simulators and lightweight modeling manner, utilizing simulators to synthesize a
 062 large amount of data to learn a lightweight world model. Therein, we argue that the *physical plausibility*
 063 and *diversity* of synthesized data are crucial to obtain satisfactory performance, thus propose a
 a physics-consistent digital twin construction and diverse demonstration generation methods.
 064

065 Specifically, we propose a framework named PhysWorld. It consists of three main stages. **Firstly**, to
 066 establish a physics-consistent digital twin, we first utilize a Vision-Language Model (VLM) to auto-
 067 matically select the optimal constitutive models for deformable objects within the MPM simulation.
 068 We then introduce a global-to-local optimization strategy to refine physical properties (*e.g.*, friction,
 069 density, and Young's modulus), ensuring the MPM simulations align with the observed video. **Sec-
 070 ondly**, since the motion trajectory of the real video is single and the learned physical parameters of
 071 the digital twin inevitably exist errors, we propose Various Motion Pattern Generation (VMP-Gen)
 072 and Part-aware Physical Property Perturbation (P^3 -Pert) methods. It enables synthesizing exten-
 073 sive and diverse 4D demonstrations. **Thirdly**, the demonstrations are used to train a GNN-based
 074 world model embedded with spatially varying physical properties. The original real videos can also
 075 be used to fine-tune the physical property values, thereby enhancing the alignment between GNN
 076 models and real-world object dynamics.

077 The learned world model can perform accurate and fast future prediction, and also generalizes well
 078 to novel interactions. Experiments on 22 scenarios show our model has competitive performance
 079 while enabling more efficient inference. In particular, it is 47 times faster than the recent state-of-
 080 the-art method, *i.e.*, PhysTwin (Jiang et al., 2025).

081 Our contributions can be summarized as follows:
 082

- (1) We propose a framework that constructs not only accurate but also fast world models for de-
 083 formable objects from short real-world videos via physics-consistent digital twin construction and
 084 diverse demonstration generation, named PhysWorld.
- (2) We propose the Various Motion Pattern Generation (VMP-Gen) and Part-aware Physical Prop-
 085 erty Perturbation (P^3 -Pert) methods for the diversity of synthesized demonstrations.
- (3) Experiments on 22 scenarios show our model has competitive performance while enabling infer-
 086 ence speeds 47 times faster than the recent state-of-the-art method, *i.e.*, PhysTwin.

091 2 RELATED WORK

092 2.1 PHYSICS-BASED SIMULATION OF DEFORMABLE OBJECTS

093 Recent advances in dynamic scene reconstruction (Li et al., 2022; Wu et al., 2024a; Luiten et al.,
 094 2024; Wu et al., 2024b) have integrated physics-based simulators with modern 3D representations
 095 such as NeRF (Mildenhall et al., 2021) and 3DGS (Kerbl et al., 2023), utilizing their physical
 096 mechanics modeling to learn dynamics grounded in physical principles. For example, PIE-NeRF (Feng
 097 et al., 2024) combines nonlinear elastodynamics with NeRF to generate plausible animations. Phys-
 098 Gaussian (Xie et al., 2024) integrates material point methods (MPM) (Jiang et al., 2016; Stom-
 099 akhin et al., 2013; Bardenhagen et al., 2000) with 3DGS for interactive object representations.
 100 VR-GS (Jiang et al., 2024) employs extended position-based dynamics alongside 3DGS to develop
 101 physics-aware virtual reality systems. However, these methods require manually specified simula-
 102 tion parameters, risking a mismatch with real-world observations. Subsequent research (Chen et al.,
 103 2022a; Li et al., 2023; Qiao et al., 2022; Zhang et al., 2024c; Zhong et al., 2024; Huang et al.,
 104 2024a) addresses this via system identification, estimating parameters directly from video during
 105 reconstruction. PhysTwin (Jiang et al., 2025) exemplifies this, using a mass-spring model optimized
 106 from interaction videos to enable action-conditioned prediction of elastic objects.
 107

108 While physics-grounded simulations offer greater realism through inherent physical priors, they face
 109 challenges. High-fidelity simulators like MPM incur prohibitive computational costs, hindering real-
 110 time applications. A persistent domain gap also exists between the dynamics of simplified simulators
 111 and complex real phenomena, frequently causing motion prediction inaccuracies.
 112

113 2.2 LEARNING-BASED SIMULATION OF DEFORMABLE OBJECTS

114
 115 Learning dynamics models directly from data has emerged as a prominent research direc-
 116 tion (Sanchez-Gonzalez et al., 2020; Chen et al., 2022b; Evans et al., 2022; Ma et al., 2023; Wu
 117 et al., 2019; Xu et al., 2019; Zhang et al., 2024a; 2025a; 2024b). Compared to traditional physics-
 118 based approaches, these methods can utilize lightweight neural networks to achieve significantly
 119 greater efficiency. In particular, Graph-based neural networks (GNNs) have proven particularly
 120 effective at learning dynamics across diverse deformable objects, including plasticine (Shi et al.,
 121 2023), cloth (Lin et al., 2022; Pfaff et al., 2020), and fluids (Li et al., 2018; Sanchez-Gonzalez
 122 et al., 2020). For instance, GS-Dynamics (Zhang et al., 2024b) employs tracking and appearance
 123 priors derived from Dynamic Gaussians (Luiten et al., 2024) to train GNN-based world models on
 124 real-world interaction videos with deformable objects.
 125

126 However, learning-based methods require vast data and lack generality. In this work, our PhysWorld
 127 synergizes physics and learning-based methods. PhysWorld uses the MPM simulator as a data
 128 factory, generating physically plausible and diverse demonstrations from real videos to train a GNN-
 129 based world model. It can take a better trade-off between accuracy and efficiency.
 130

131 3 PHYSWORLD

132 PhysWorld builds a physics-consistent MPM-based digital twin from short real-world videos by
 133 leveraging a VLM to determine the appropriate constitutive models and a global-to-local physical
 134 property optimization strategy to align the simulation with real observations. The calibrated twin
 135 is then used to generate a variety of interaction data, which trains a GNN-based world model en-
 136 abling real-time simulation. The physical properties embedded in GNNs are finally fine-tuned on
 137 the original real data to mitigate the sim-to-real gap.
 138

139 3.1 PHYSICS-CONSISTENT DIGITAL TWIN

140 Empirical research (Zhang et al., 2024b; Cao et al., 2024; Cai et al., 2024) reveals that short real inter-
 141 action videos alone are insufficient for learning comprehensive world models of deformable objects,
 142 whereas estimating physical properties proves comparatively tractable. This motivates our proposi-
 143 tion, *i.e.*, leveraging a high-fidelity physics engine to impose fundamental physical constraints (*e.g.*,
 144 mass and momentum conservation laws) as strong priors. Rather than learning physics from first
 145 principles, we parameterize the object within the physics engine and optimize only its digital twin’s
 146 constitutive model and physical properties using the real data.
 147

148 Construction of Initial Digital Twin Following object point cloud extraction from real interaction

149 videos (Jiang et al., 2025), we construct a physics-consistent MPM-based digital twin. To ensure

150 simulation fidelity under continuum mechanics principles, we utilize a tetrahedral mesh refinement

151 method (Wang et al., 1996) to fill interior voids in the extracted point cloud, thereby creating com-
 152 plete volumetric representations. We implement two robotic manipulation primitives within the
 153 MPM framework: (1) **Grasping**, which is achieved by directly assigning velocities to the Eule-
 154 rian grids near the contact points; and (2) **Pushing**, which is implemented via geometric controllers
 155 defined by a Signed Distance Function (SDF) that transfer motion to the proximal grid nodes. Mean-
 156 while, the videos are analyzed using Qwen3 (Yang et al., 2025) to automatically identify the most
 157 suitable material constitutive models from our physics library. This VLM-based method dynami-
 158 cally adapts material representations according to the observed deformation behaviors. The consti-
 159 tutive models implemented and the associated prompts are provided in the Appendix B.1.
 160

161 Physical Property Optimization To ensure physics-consistent alignment between our digital twin
 162 and real-world objects, we optimize the twin’s physical properties (*e.g.*, friction coefficients, den-
 163 sity, and Young’s modulus) using real-world observational data. More specifically, the densified
 164 point cloud is registered as MPM particles. Using the implemented manipulation primitives, we

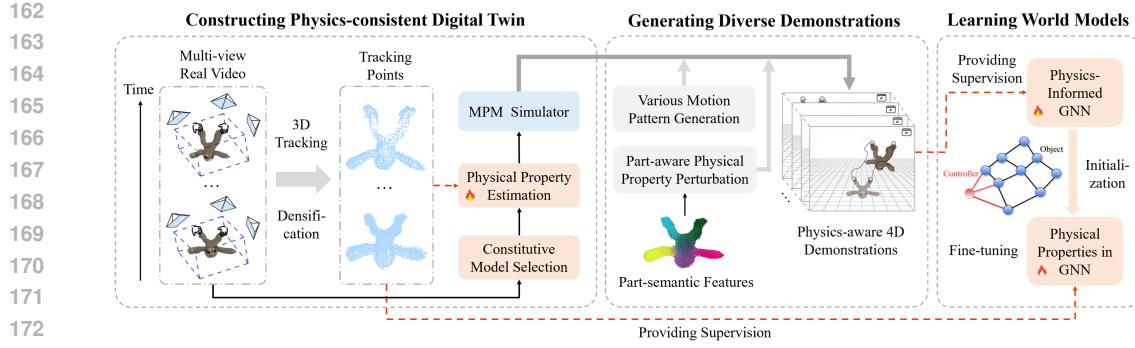


Figure 1: Overview of PhysWorld. The framework first constructs a physics-consistent digital twin from videos, then uses it to generate diverse 4D demonstrations, and finally trains a GNN-based world model for real-time future state prediction.

drive the system to re-simulate the hand motion trajectory extracted from real-world data. The entire MPM simulation framework is differentiable, enabling the computation of gradients of the Chamfer Distance and \mathcal{L}_1 loss between the simulated outcomes and the ground truth. This differentiability facilitates the optimization of the physical properties of MPM particles to achieve enhanced alignment. The loss function used for optimization is provided in Appendix B.2. Experimental results show that parameter initialization critically affects optimization performance: well-chosen initial values lead to faster convergence and higher performance. Thus, we employ a global-to-local optimization strategy. First, on the global stage, we optimize homogeneous physical properties across all particles. Then, on the local stage, we refine per-particle heterogeneous properties.

3.2 AUGMENTED INTERACTION DEMONSTRATION SYNTHESIS

To train a robust GNN-based world model, we generate various interaction demonstrations by simultaneously varying motion patterns and applying part-aware perturbations to physical properties.

Various Motion Pattern Generation Addressing the need for complex motion patterns during inference, control point trajectories are generated via curvature-constrained Bézier (Prautzsch et al., 2002) curves with integrated velocity regimes. One spatial trajectory $\mathbf{x}(t)$ can be defined by a cubic Bézier curve, *i.e.*,

$$\mathbf{B}(u) = (1-u)^3 \mathbf{p}_0 + 3(1-u)^2 u \mathbf{p}_1 + 3(1-u)u^2 \mathbf{p}_2 + u^3 \mathbf{p}_3. \quad (1)$$

\mathbf{p}_0 and \mathbf{p}_3 are start and end points of the trajectory. \mathbf{p}_1 and \mathbf{p}_2 are curvature control points generated through randomized curvature parameters. $u \in [0, 1]$ is the Bézier parameter. The temporal evolution of the trajectory is governed by the time-to-parameter mapping $u(t)$, which is constructed using normalized arc-length parameterization (Hartlen & Cronin, 2022). The normalized arc length $s(t)$ can be formulated as:

$$s(t) = \frac{\int_0^t v(\tau) d\tau}{\int_0^T v(\tau) d\tau}, \quad (3)$$

where $v(t)$ represents the velocity profile along the trajectory. This ensures smooth and physically consistent motion along the curve. To enable diverse and natural motion generation, we implement a three-phase velocity profile consisting of acceleration, uniform motion, and deceleration phases. Crucially, these phases are connected with smooth C^1 transitions, ensuring continuous velocity and avoiding abrupt changes in acceleration.

Part-aware Physical Property Perturbation To address potential inaccuracies in physical properties optimized in MPM solely from short real videos, we impose semantic-partition-guided stochastic perturbations on the optimized physical properties during demonstration synthesis. This strategy incorporates controlled diversity in physical property distribution while preserving continuity to match real-world distributions and maintain simulation stability. We extract the part-semantic feature vector $\{\mathbf{F}_i\}_{i=1}^N$ for each of the N MPM particles via PartField (Liu et al., 2025). The feature similarity between particles i and j is computed as,

$$S_{ij} = \exp \left(-\frac{\|\mathbf{F}_i - \mathbf{F}_j\|_2^2}{2\ell^2} \right), \quad (4)$$

216 where ℓ controls the feature similarity decay rate. The covariance matrix Σ is constructed using
 217 pairwise similarities, which can be written as,
 218

$$\Sigma_{ij} = \sigma^2 S_{ij}, \quad (5)$$

220 where σ governs the perturbation intensity. Stochastic perturbations sampled from $\mathcal{N}(\mathbf{0}, \Sigma)$ are then
 221 applied to physical properties across particles. However, exact computation of the covariance matrix
 222 becomes intractable for large-scale particle systems. We therefore employ the Nyström approxima-
 223 tion (Fowlkes et al., 2004), which constructs a low-rank approximation through subset sampling,
 224 achieving significant computational efficiency. We thus generate various augmented interaction
 225 demonstrations $\{\mathbf{X}_t, \Phi, \mathbf{a}_t\}_{t=0}^T$ for each scenario, where $\mathbf{X}_t = \{\mathbf{x}_t^{(i)}\}_{i=1}^N$ represents the object's
 226 N -particle point cloud at timestep t ; $\Phi = \{\phi^{(i)}\}_{i=1}^N$ encapsulates per-particle physical properties
 227 (e.g., Young's modulus E_i and density ρ_i); and $\mathbf{a}_t = \{\mathbf{a}_t^{(k)}\}_{k=1}^K$ denotes the velocities of K control
 228 points interacting with the object at timestep t . These data are then used to train a world model.
 229

230 3.3 GNN-BASED WORLD MODELS

232 Following the construction of a physics-consistent MPM digital twin, we recognize that computa-
 233 tional latency is a main limitation to prevent its direct use as a world model. MPM simulations
 234 exhibit prohibitive inference times, making them impractical for real-time applications like model-
 235 based planning, which require fast responses. To address the issue, we propose to train lightweight
 236 and fast GNN-based world models using the augmented synthetic interaction demonstrations gen-
 237 erated by the constructed MPM digital twin. The model achieves heterogeneous material-aware
 238 dynamics prediction with real-time inference capabilities.

239 **Model Architecture** The GNN-based world model f can be formulated as:

$$\mathbf{X}_{t+1} = f(\mathbf{X}_{t-h:t}, \mathbf{a}_t, \Phi), \quad (6)$$

242 where t denotes the current timestep; h specifies the history window size; \mathbf{a}_t represents the ap-
 243 plied action(the velocities of control points) at time t ; and Φ represents the time-invariant phys-
 244 ical properties embedded in GNNs. The GNN construction proceeds through three key steps:
 245 First, farthest point sampling (FPS) is applied to the raw point clouds to extract n control parti-
 246 cles $\hat{\mathbf{X}}_t = \{\hat{\mathbf{x}}_t^{(i)}\}_{i=1}^n$ with minimum inter-particle separation d_v , while their corresponding physical
 247 properties $\Phi = \{\phi^{(i)}\}_{i=1}^N$ are subsampled accordingly to yield $\hat{\Phi} = \{\hat{\phi}^{(i)}\}_{i=1}^n$. Second, control
 248 points are incorporated as additional graph vertices. Finally, bidirectional edges $\hat{\mathbf{E}}_t$ are established
 249 between vertices within connection radius d_e .

250 The GNN architecture integrates three core components: vertex/edge encoders for feature extraction,
 251 a multi-step message propagator, and a motion prediction decoder. Designed to ensure translation
 252 equivariance, the vertex encoder only processes vertex type indicators (object/controller) encoded
 253 as one-hot vectors, control point velocities \mathbf{a}_t and time-invariant physical properties $\hat{\phi}^{(i)}$ without
 254 vertex positions. Simultaneously, the edge encoder handles history distance between two edge nodes
 255 $\{\mathbf{x}_{t-\tau}^{(i)} - \hat{\mathbf{x}}_{t-\tau}^{(j)}\}_{\tau=0}^h$ and a edge type identifier(object-object or object-controller). The output features
 256 then undergo p ($p = 7$ in our experiments) iterations of message passing through the propagator's
 257 independent MLP blocks with unshared parameters. And the decoder finally generates per-particle
 258 3D motion predictions $\Delta\hat{\mathbf{x}}_t^{(i)}$. The entire process can be represented as:

$$\hat{\mathbf{X}}_{t+1,pred} = \hat{\mathbf{X}}_t + f_\theta(\hat{\mathbf{X}}_{t-h:t}, \mathbf{a}_t, \hat{\Phi}), \quad (7)$$

262 where f_θ denotes the GNN model parameterized by θ .

263 **Model Training** The generated interaction demonstrations serve as training data for the GNN, with
 264 the mean squared error between predicted and ground-truth particle positions constituting the pri-
 265 mary loss, *i.e.*,

$$\mathcal{L}_{pred} = \sum_{i=1}^{\tau} \|\hat{\mathbf{X}}_{t+i,pred} - \hat{\mathbf{X}}_{t+i}\|^2. \quad (8)$$

266 The look-forward horizon τ determines the multi-step prediction window size, balancing accuracy
 267 and computational efficiency. Training sequences are initialized at randomly sampled timesteps t .
 268

270 Since our model is trained on ground-truth short-term data pairs, it is susceptible to error accumulation
 271 during long-horizon rollouts. To alleviate this issue, we inject noise into historical vertex
 272 positions, thereby aligning the training distribution more closely with that encountered during roll-
 273 out generation. Noise is also introduced to the GNN’s physical properties $\hat{\Phi}$ during training, with
 274 the objective of enhancing the stability of subsequent physical property fine-tuning.

275 **Physical Property Fine-Tuning** Due to the proneness to gradient explosion and training instability
 276 when optimizing the physical properties $\hat{\Phi}$ of MPM particles—a result of employing a large number
 277 of simulation substeps per frame—the estimated values of these properties are not yet sufficiently
 278 accurate. In contrast, GNNs do not suffer from gradient explosion during training and demonstrate
 279 greater stability. Therefore, based on a pre-trained physical property-conditioned GNN, we fine-
 280 tune the physical properties $\hat{\Phi}$ at the GNN vertices using real-world data while freezing the network
 281 parameters θ . This approach enhances the alignment between the GNN and real objects in terms of
 282 physical motion characteristics, while also further narrowing the sim-to-real domain gap resulting
 283 from training the model solely on synthetic data.

284 **Action-Conditioned Video Prediction** To achieve action-conditioned video prediction, we addi-
 285 tionally optimize object appearance representations using 3DGS (Kerbl et al., 2023). Each object is
 286 modeled as a set $\mathcal{G} = \{\mathcal{G}^{(k)}\}_{k=1}^K$ of 3D Gaussian kernels, where each $\mathcal{G}^{(k)} = (\mu_k, \mathbf{q}_k, \mathbf{s}_k, \alpha_k, \mathbf{c}_k)$
 287 consists of: center position $\mu_k \in \mathbb{R}^3$, rotation quaternion $\mathbf{q}_k \in \mathbb{SO}(3)$, scaling vector $\mathbf{s}_k \in \mathbb{R}^3$,
 288 opacity $\alpha_k \in [0, 1]$, and RGB color coefficients $\mathbf{c}_k \in \mathbb{R}^3$. We optimize the 3D Gaussian Splatting
 289 representation solely at $t = 0$. For subsequent timesteps ($t > 0$), the representation is derived via
 290 Linear Blend Skinning (LBS) (Sumner et al., 2007), *i.e.*,

$$\mathcal{G}_t = \text{LBS} \left(\mathcal{G}_0, \{\Delta \hat{\mathbf{X}}_\tau\}_{\tau=1}^t \right), \quad (9)$$

294 where LBS interpolates Gaussian motions using the vertex motion field predicted by our GNN. A
 295 more detailed description of LBS is provided in Appendix B.4.

297 4 EXPERIMENT

300 4.1 EXPERIMENTAL SETUP

301 **Implementation Details** Our MPM simulator is built upon PhysGaussian (Xie et al., 2024). We
 302 implement additional constitutive models for diverse materials (*e.g.*, anisotropic cloth) and incorpo-
 303 rates grasping and pushing capabilities. During MPM simulations, for the majority of simulation
 304 scenarios, we set the total simulation duration to 0.04 seconds and the per-frame duration to 0.0001
 305 seconds. This indicates that the state between consecutive frames is advanced through 400 simula-
 306 tion substeps, implicitly satisfying the CFL condition for numerical stability. To mitigate the risk of
 307 gradient explosion and ensure training stability during the optimization of physical properties, we
 308 compute gradients using only the final 10 simulation substeps, supplemented by gradient clipping.
 309 For each scene, we generate 500 interaction demonstration episodes to train the GNN. The frame
 310 count for each episode closely matches that of the original real-world data. During GNN training,
 311 the number of GNN nodes was maintained at approximately 100 to 150 through FPS downsampling.
 312 Meanwhile, the number of message passing steps was set to 7 in our experiments.

313 **Dataset** We utilize an open-access dataset (Jiang et al., 2025) featuring human interactions with a
 314 variety of deformable objects. The dataset consists of 1-10 second videos capturing diverse interac-
 315 tions such as rapid lifting, stretching, pushing, and bimanual squeezing. Each one is with distinct
 316 physical properties, including ropes, stuffed animals, cloth, and packages. For each scenario, we par-
 317 tition the whole video into training and test frames with a 7:3 ratio. We only use the training frames
 318 exclusively for developing MPM digital twins and fine-tuning the GNN-based world models.

319 **Evaluation Configurations** Following PhysTwin (Jiang et al., 2025), we employ metrics in both 3D
 320 and 2D spaces for evaluation. In 3D space, we utilize the single-direction Chamfer Distance (*i.e.*,
 321 CD) and a tracking error (*i.e.*, Track) derived from manually annotated ground-truth points. In 2D
 322 space, we evaluate image reconstruction quality using the PSNR, SSIM, and LPIPS (Zhang et al.,
 323 2018) metrics, while silhouette alignment is measured by the Intersection over Union (*i.e.*, IoU). To
 324 quantify the prediction quality of the model on unseen interactions in the absence of ground truth, we

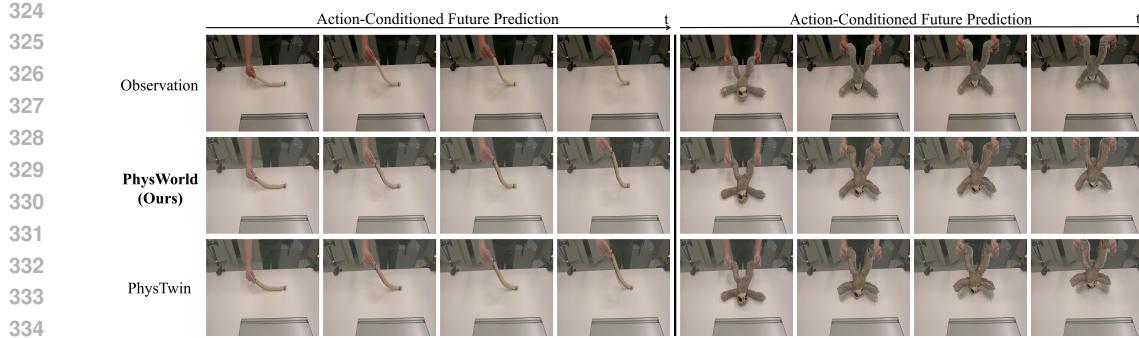


Figure 2: Visual results of action-conditioned future prediction. Our method’s predicted positions show closer alignment with ground truth compared to PhysTwin.

Table 1: Quantitative results on action-conditioned future prediction and inference speed (FPS). Best metric values are **bolded**, while second-best ones are underlined.

Methods	CD \downarrow	Track \downarrow	IoU \uparrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	FPS \uparrow
Spring-Gaus (ECCV 2024)	0.062	0.094	46.4	22.488	0.924	0.113	2
GS-Dynamics (CoRL 2024)	0.041	0.070	49.8	22.540	0.924	0.097	<u>236</u>
PhysTwin (ICCV 2025)	0.012	0.022	72.5	25.617	0.941	<u>0.055</u>	17
Our MPM	<u>0.011</u>	0.021	74.7	26.231	0.942	0.052	3
Our GNN w/o Finetuning	0.012	0.025	70.0	25.345	0.939	0.059	799
PhysWorld(Our GNN w Finetuning)	0.010	0.021	<u>73.3</u>	<u>25.940</u>	<u>0.941</u>	<u>0.055</u>	799

follow DreamPhysics (Huang et al., 2024a) and employ the following metrics from VBench (Huang et al., 2024c): aesthetic quality score, motion smoothness score, and subject consistency score.

Comparison Configurations To the best of our knowledge, there is a scarcity of methods for learning object-centric world models from short video clips (under 10 seconds). PhysTwin (Jiang et al., 2025) currently stands as the state-of-the-art approach by employing a spring-mass model as its foundational dynamics model, where the physical properties are optimized through real-world video observations. To provide a more comprehensive evaluation, we include several other approaches, *i.e.*, Spring-Gauss (Zhong et al., 2024) and GS-Dynamics (Zhang et al., 2025b), where Spring-Gauss is a physics-based simulation method also based on spring-mass models and GS-Dynamics is a learning-based approach employing a GNN-based model to learn dynamics from real videos.

4.2 EXPERIMENTAL RESULTS

Action-Conditioned Future Prediction The qualitative results in Fig. 2 and quantitative results in Table 1 for the action-conditioned future prediction task demonstrate the superior performance of our proposed methods. Our MPM method, enhanced with optimized physical properties, surpasses other approaches across most evaluated metrics, which underscores MPM’s ability to simulate a wide variety of deformable objects with high fidelity. Furthermore, our finetuned GNN achieves competitive results, outperforming the baseline model PhysTwin in the majority of benchmarks and achieving the lowest CD and Track loss with the fastest inference speed. This highlights the accuracy and effectiveness of PhysWorld in both predicting motion and generating realistic images.

Inference Time Comparison We perform the inference speed comparisons with an NVIDIA GeForce RTX 4060 Ti (16GB) GPU in the *double_lift_cloth_3* scene, which consists of 118 frames (a relatively long sequence) and contains 171,602 3DGS kernels. The results are shown in Table 1. The results demonstrate our PhysWorld maintains competitive accuracy while achieving fast inference speed. Spring-Gaus exhibits limitations in both prediction accuracy and inference speed. Although GS-Dynamics is observed to have a fast inference speed, it suffers from low predictive accuracy. The recent state-of-the-art PhysTwin lags behind in both speed ($47\times$ slower) and prediction accuracy. MPM achieves slightly better prediction accuracy but at a prohibitive computational cost, running about $400\times$ slower than PhysWorld. The real-time inference capability of PhysWorld

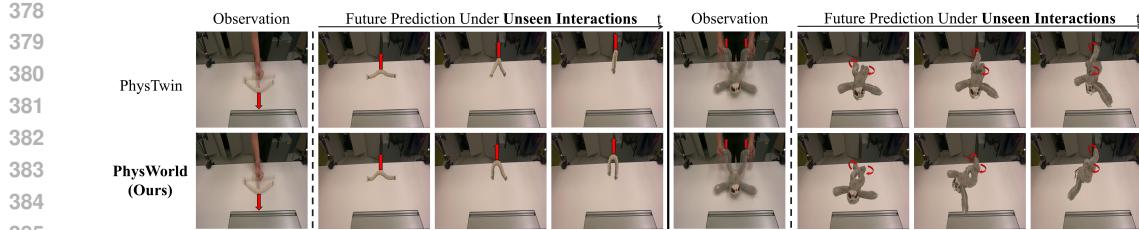


Figure 3: Generalization to unseen interactions. As representative examples, we consider two unseen interaction scenarios: lifting a pushed rope and rotating a lifted sloth. The results show that PhysWorld generates physically plausible predictions, while PhysTwin suffers from artifacts such as fracture-like rope distortions and unnatural foot folding.

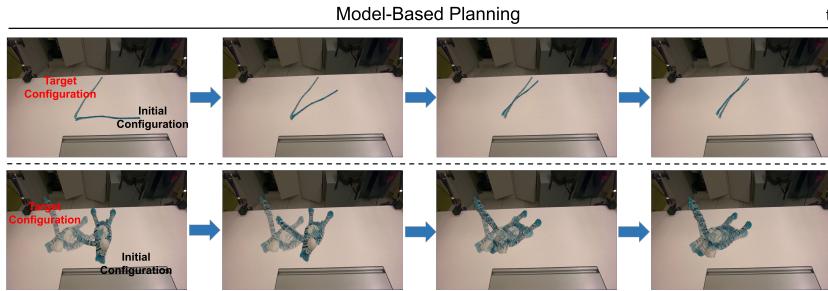


Figure 4: Examples of model-based planning. With MPPI control, the rope and the zebra doll are transferred from the initial configurations to the target ones.

Table 2: Quantitative results on unseen interactions.

Methods	Aesthetic Quality↑	Motion Smoothness↑	Subject Consistency↑
Phystwin	0.4315	0.9971	0.9155
PhysWorld (Ours)	0.4440	0.9973	0.9312

enables its deployment in computationally demanding scenarios, such as model-based robotic planning, interactive simulation systems, and real-time optimization tasks.

Generalization to Unseen Interactions We further compare PhysWorld’s generalization performance on unseen interactions. Using *single.push.rope* and *double.lift.sloth* as test cases, we apply randomized action sequences to objects. Visual results in Fig. 3 show that PhysWorld consistently generates high-fidelity physical predictions, significantly outperforming PhysTwin in dynamic accuracy and deformation realism. The quantitative results in Table 2 also indicate that PhysWorld generates more plausible predictions when confronted with unseen interactions.

Model-Based Planning PhysWorld enables real-time trajectory optimization for model-based robotic planning, as shown in the case studies in Fig. 4. Given initial and target configurations of the deformable objects, we utilize a Model-Predictive Path Integral (MPPI) (Williams et al., 2017) planning framework, which successfully generates controlling trajectories to guide the objects towards their target configurations.

4.3 ABLATION STUDY

Effect of Global-to-Local Physical Property Optimization Here we conduct an ablation study on optimization strategies of physical properties within MPM. As summarized in Table 3, experimental results demonstrate that global-only optimization fails to capture localized material variations, while local optimization exhibits convergence challenges. Our global-to-local approach significantly enhances optimization stability in the second stage by initializing local properties with global property priors, yielding superior overall performance.

Effect of VMP-Gen The ablation studies (Table 4) demonstrate that the proposed VMP-Gen modules significantly enhance the trained GNN’s prediction accuracy and robustness. This improvement

432

433

Table 3: Effect of different physical property optimization strategies for MPM.

434

435

Strategy	CD \downarrow	Track \downarrow	IoU \uparrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
Global	0.012	0.024	72.4	25.854	0.940	0.055
Local	0.016	0.032	66.5	24.901	0.935	0.066
Global-to-Local (Ours)	0.010	0.021	74.7	26.231	0.942	0.052

436

437

438

439

440

Table 4: Effect of different motion patterns.

441

442

443

444

Motion Pattern	CD \downarrow	Track \downarrow	IoU \uparrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
Uniform and linear	0.0114	0.0175	76.95	24.475	0.920	0.067
VMP-Gen (Ours)	0.0100	0.0154	78.66	24.666	0.921	0.067

445

446

447

Table 5: Effect of different physical property perturbation strategies.

448

449

450

451

452

453

Perturbation	CD \downarrow	Track \downarrow	IoU \uparrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
\times	0.0111	0.0179	75.84	23.984	0.919	0.070
Random	0.0153	0.0216	70.19	23.073	0.915	0.082
Uniform	0.0147	0.0258	72.00	23.101	0.914	0.079
P^3 -Pert (Ours)	0.0100	0.0154	78.66	24.666	0.921	0.067

454

455

Table 6: Performance comparison with a GNN directly trained on real data.

456

457

458

459

Methods	CD \downarrow	Track \downarrow	IoU \uparrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
GNN (Directly trained on real data)	0.0530	0.0802	41.32	19.925	0.881	0.132
PhysWorld (Ours)	0.0100	0.0154	78.66	24.666	0.921	0.067

460

461

462

is achieved by increasing motion diversity, enabling the model to capture a wider range of patterns beyond simple uniform linear motion.

463

Effect of P^3 -Pert Ablation studies demonstrate the effectiveness of the proposed P^3 -Pert modules, with quantitative results detailed in Table 5. The key advantage of P^3 -Pert lies in its generation of physically realistic part-aware property variations. As evidenced by the results, our method surpasses the random and uniform perturbation baselines (e.g., AdaptiGraph (Zhang et al., 2024a)) by producing more plausible and diverse physical property distributions.

464

465

466

467

468

Training GNN Directly on Real Data To evaluate the role of synthetic data in GNN training, we compared PhysWorld against traditional real-data-only training. Results on the *double_lift_sloth* task in Table 6 indicate that real-data scarcity caused severe overfitting in GNNs, which, as a result, led to compromised performance in future state prediction due to poor generalization. In contrast, incorporating various synthetic data used in PhysWorld significantly improved prediction accuracy and model robustness.

469

470

471

472

473

474

5 CONCLUSION

475

476

477

478

479

480

481

482

483

484

485

In this work, we introduce PhysWorld, a novel framework that constructs accurate and efficient world models for deformable objects through physics-aware demonstration synthesis from real-world videos. It first establishes a MPM-based digital twin and then generates extensive augmented interaction demonstrations via strategic perturbations of physical parameters, control trajectories, and velocity profiles. We employ the generated data to train a real-time GNN world model that learns spatially heterogeneous physical properties, with real video data subsequently refining the model’s parameters. The world model supports diverse downstream applications such as model-based planning and robotic manipulation. Extensive experiments show that PhysWorld outperforms state-of-the-art methods.

486 REFERENCES
487

488 SG Bardenhagen, JU Brackbill, and Deborah Sulsky. The material-point method for granular mate-
489 rials. *Computer methods in applied mechanics and engineering*, 187(3-4):529–541, 2000.

490 Junhao Cai, Yuji Yang, Weihao Yuan, Yisheng He, Zilong Dong, Liefeng Bo, Hui Cheng, and Qifeng
491 Chen. Gic: Gaussian-informed continuum for physical property identification and simulation.
492 *Advances in Neural Information Processing Systems*, 37:75035–75063, 2024.

493 Junyi Cao, Shanyan Guan, Yanhao Ge, Wei Li, Xiaokang Yang, and Chao Ma. Neuma: Neural
494 material adaptor for visual grounding of intrinsic dynamics. *Advances in Neural Information
495 Processing Systems*, 37:65643–65669, 2024.

496 Hsiao-yu Chen, Edith Tretschk, Tuur Stuyck, Petr Kadlecak, Ladislav Kavan, Etienne Vouga, and
497 Christoph Lassner. Virtual elastic objects. In *Proceedings of the IEEE/CVF Conference on Com-
498 puter Vision and Pattern Recognition*, pp. 15827–15837, 2022a.

499 Zhenfang Chen, Kexin Yi, Yunzhu Li, Mingyu Ding, Antonio Torralba, Joshua B Tenenbaum, and
500 Chuang Gan. Comphy: Compositional physical reasoning of objects and events from videos.
501 *arXiv preprint arXiv:2205.01089*, 2022b.

502 Ben Evans, Abitha Thankaraj, and Lerrel Pinto. Context is everything: Implicit identification for
503 dynamics adaptation. In *2022 International Conference on Robotics and Automation (ICRA)*, pp.
504 2642–2648. IEEE, 2022.

505 Yutao Feng, Yintong Shang, Xuan Li, Tianjia Shao, Chenfanfu Jiang, and Yin Yang. Pie-nerf:
506 Physics-based interactive elastodynamics with nerf. In *Proceedings of the IEEE/CVF Conference
507 on Computer Vision and Pattern Recognition*, pp. 4450–4461, 2024.

508 Charless Fowlkes, Serge Belongie, Fan Chung, and Jitendra Malik. Spectral grouping using the
509 nyström method. *IEEE transactions on pattern analysis and machine intelligence*, 26(2):214–
510 225, 2004.

511 Devon C Hartlen and Duane S Cronin. Arc-length re-parametrization and signal registration to de-
512 termine a characteristic average and statistical response corridors of biomechanical data. *Frontiers
513 in Bioengineering and Biotechnology*, 10:843148, 2022.

514 Tianyu Huang, Yihan Zeng, Hui Li, Wangmeng Zuo, and Rynson WH Lau. Dreamphysics: Learning
515 physical properties of dynamic 3d gaussians with video diffusion priors. *arXiv e-prints*, pp. arXiv–
516 2406, 2024a.

517 Yi-Hua Huang, Yang-Tian Sun, Ziyi Yang, Xiaoyang Lyu, Yan-Pei Cao, and Xiaojuan Qi. Sc-
518 gs: Sparse-controlled gaussian splatting for editable dynamic scenes. In *Proceedings of the
519 IEEE/CVF conference on computer vision and pattern recognition*, pp. 4220–4230, 2024b.

520 Ziqi Huang, Yinan He, Jiashuo Yu, Fan Zhang, Chenyang Si, Yuming Jiang, Yuanhan Zhang, Tianx-
521 ing Wu, Qingyang Jin, Nattapol Chanpaisit, et al. Vbench: Comprehensive benchmark suite for
522 video generative models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and
523 Pattern Recognition*, pp. 21807–21818, 2024c.

524 Chenfanfu Jiang, Craig Schroeder, Joseph Teran, Alexey Stomakhin, and Andrew Selle. The mate-
525 rial point method for simulating continuum materials. In *Acm siggraph 2016 courses*, pp. 1–52.
526 2016.

527 Hanxiao Jiang, Hao-Yu Hsu, Kaifeng Zhang, Hsin-Ni Yu, Shenlong Wang, and Yunzhu Li. Phys-
528 twin: Physics-informed reconstruction and simulation of deformable objects from videos. *arXiv
529 preprint arXiv:2503.17973*, 2025.

530 Ying Jiang, Chang Yu, Tianyi Xie, Xuan Li, Yutao Feng, Huamin Wang, Minchen Li, Henry Lau,
531 Feng Gao, Yin Yang, et al. Vr-gs: A physical dynamics-aware interactive gaussian splatting
532 system in virtual reality. In *ACM SIGGRAPH 2024 Conference Papers*, pp. 1–1, 2024.

533 Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splat-
534 ting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139–1, 2023.

540 Tianye Li, Mira Slavcheva, Michael Zollhoefer, Simon Green, Christoph Lassner, Changil Kim,
 541 Tanner Schmidt, Steven Lovegrove, Michael Goesele, Richard Newcombe, et al. Neural 3d video
 542 synthesis from multi-view video. In *Proceedings of the IEEE/CVF conference on computer vision*
 543 and pattern recognition, pp. 5521–5531, 2022.

544

545 Xuan Li, Yi-Ling Qiao, Peter Yichen Chen, Krishna Murthy Jatavallabhula, Ming Lin, Chenfanfu
 546 Jiang, and Chuang Gan. Pac-nerf: Physics augmented continuum neural radiance fields for
 547 geometry-agnostic system identification. *arXiv preprint arXiv:2303.05512*, 2023.

548 Yunzhu Li, Jiajun Wu, Russ Tedrake, Joshua B Tenenbaum, and Antonio Torralba. Learning par-
 549 ticle dynamics for manipulating rigid bodies, deformable objects, and fluids. *arXiv preprint*
 550 *arXiv:1810.01566*, 2018.

551

552 Xingyu Lin, Yufei Wang, Zixuan Huang, and David Held. Learning visible connectivity dynamics
 553 for cloth smoothing. In *Conference on Robot Learning*, pp. 256–266. PMLR, 2022.

554

555 Yuchen Lin, Chenguo Lin, Jianjin Xu, and Yadong MU. Omniphysgs: 3d constitutive gaussians
 556 for general physics-based dynamics generation. In *The Thirteenth International Conference on*
 557 *Learning Representations*.

558

559 Minghua Liu, Mikaela Angelina Uy, Donglai Xiang, Hao Su, Sanja Fidler, Nicholas Sharp, and
 560 Jun Gao. Partfield: Learning 3d feature fields for part segmentation and beyond. *arXiv preprint*
 561 *arXiv:2504.11451*, 2025.

562

563 Jonathon Luiten, Georgios Kopanas, Bastian Leibe, and Deva Ramanan. Dynamic 3d gaussians:
 564 Tracking by persistent dynamic view synthesis. In *2024 International Conference on 3D Vision*
 (3DV), pp. 800–809. IEEE, 2024.

565

566 Pingchuan Ma, Peter Yichen Chen, Bolei Deng, Joshua B Tenenbaum, Tao Du, Chuang Gan, and
 567 Wojciech Matusik. Learning neural constitutive laws from motion observations for generalizable
 568 pde dynamics. In *International Conference on Machine Learning*, pp. 23279–23300. PMLR,
 2023.

569

570 Xu Ma, Can Qin, Haoxuan You, Haoxi Ran, and Yun Fu. Rethinking network design and local
 571 geometry in point cloud: A simple residual mlp framework. *arXiv preprint arXiv:2202.07123*,
 572 2022.

573

574 Fanqing Meng, Jiaqi Liao, Xinyu Tan, Wenqi Shao, Quanfeng Lu, Kaipeng Zhang, Yu Cheng,
 575 Dianqi Li, Yu Qiao, and Ping Luo. Towards world simulator: Crafting physical commonsense-
 576 based benchmark for video generation. *arXiv preprint arXiv:2410.05363*, 2024.

577

578 Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and
 579 Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications*
 of the ACM, 65(1):99–106, 2021.

580

581 Tobias Pfaff, Meire Fortunato, Alvaro Sanchez-Gonzalez, and Peter Battaglia. Learning mesh-based
 582 simulation with graph networks. In *International conference on learning representations*, 2020.

583

584 Hartmut Prautzsch, Wolfgang Boehm, and Marco Paluszny. *Bézier and B-spline techniques*.
 Springer Science & Business Media, 2002.

585

586 Yi-Ling Qiao, Alexander Gao, and Ming Lin. Neuphysics: Editable neural geometry and physics
 587 from monocular videos. *Advances in Neural Information Processing Systems*, 35:12841–12854,
 588 2022.

589

590 A Sanchez-Gonzalez, J Godwin, T Pfaff, R Ying, J Leskovec, and P Battaglia. Learning to simulate
 591 complex physics with graph networks. In *International Conference on Machine Learning (ICML)*,
 592 2020.

593

Haochen Shi, Huazhe Xu, Samuel Clarke, Yunzhu Li, and Jiajun Wu. Robocook: Long-horizon
 elasto-plastic object manipulation with diverse tools. *arXiv preprint arXiv:2306.14447*, 2023.

594 Haochen Shi, Huazhe Xu, Zhiao Huang, Yunzhu Li, and Jiajun Wu. Robocraft: Learning to see,
 595 simulate, and shape elasto-plastic objects in 3d with graph networks. *The International Journal
 596 of Robotics Research*, 43(4):533–549, 2024.

597 Alexey Stomakhin, Craig Schroeder, Lawrence Chai, Joseph Teran, and Andrew Selle. A material
 598 point method for snow simulation. *ACM Transactions on Graphics (TOG)*, 32(4):1–10, 2013.

600 Robert W Sumner, Johannes Schmid, and Mark Pauly. Embedded deformation for shape manipula-
 601 tion. In *ACM siggraph 2007 papers*, pp. 80–es. 2007.

602 Everett X Wang, Martin D Giles, Sia Yu, FA Leon, A Hiroki, and S Odanaka. Recursive m-tree
 603 method for 3-d adaptive tetrahedral mesh refinement and its application to brillouin zone dis-
 604 cretization. In *1996 International Conference on Simulation of Semiconductor Processes and
 605 Devices. SISPAD’96 (IEEE Cat. No. 96TH8095)*, pp. 67–68. IEEE, 1996.

606 Yixuan Wang, Yunzhu Li, Katherine Driggs-Campbell, Li Fei-Fei, and Jiajun Wu. Dynamic-
 607 resolution model learning for object pile manipulation. In *Robotics: Science and Systems*, 2023.

609 Grady Williams, Andrew Aldrich, and Evangelos A Theodorou. Model predictive path integral
 610 control: From theory to parallel computation. *Journal of Guidance, Control, and Dynamics*, 40
 611 (2):344–357, 2017.

612 Guanjun Wu, Taoran Yi, Jiemin Fang, Lingxi Xie, Xiaopeng Zhang, Wei Wei, Wenyu Liu, Qi Tian,
 613 and Xinggang Wang. 4d gaussian splatting for real-time dynamic scene rendering. In *Proceed-
 614 ings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 20310–20320,
 615 2024a.

617 Renlong Wu, Zhilu Zhang, Mingyang Chen, Xiaopeng Fan, Zifei Yan, and Wangmeng Zuo. De-
 618 blur4dgs: 4d gaussian splatting from blurry monocular video. *arXiv preprint arXiv:2412.06424*,
 619 2024b.

620 Xiaoyang Wu, Li Jiang, Peng-Shuai Wang, Zhijian Liu, Xihui Liu, Yu Qiao, Wanli Ouyang, Tong
 621 He, and Hengshuang Zhao. Point transformer v3: Simpler faster stronger. In *Proceedings of the
 622 IEEE/CVF conference on computer vision and pattern recognition*, pp. 4840–4851, 2024c.

623 Yilin Wu, Wilson Yan, Thanard Kurutach, Lerrel Pinto, and Pieter Abbeel. Learning to manipulate
 624 deformable objects without demonstrations. *arXiv preprint arXiv:1910.13439*, 2019.

626 Tianyi Xie, Zeshun Zong, Yuxing Qiu, Xuan Li, Yutao Feng, Yin Yang, and Chenfanfu Jiang.
 627 Physgaussian: Physics-integrated 3d gaussians for generative dynamics. In *Proceedings of the
 628 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4389–4398, 2024.

630 Zhenjia Xu, Jiajun Wu, Andy Zeng, Joshua B Tenenbaum, and Shuran Song. Densephysnet: Learn-
 631 ing dense physical object representations via multi-step dynamic interactions. *arXiv preprint
 632 arXiv:1906.03853*, 2019.

633 An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang
 634 Gao, Chengan Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu,
 635 Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin
 636 Yang, Jiaxi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang,
 637 Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui
 638 Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang
 639 Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger
 640 Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan
 641 Qiu. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025.

642 Kaifeng Zhang, Baoyu Li, Kris Hauser, and Yunzhu Li. Adaptigraph: Material-adaptive graph-based
 643 neural dynamics for robotic manipulation. *arXiv preprint arXiv:2407.07889*, 2024a.

644 Kaifeng Zhang, Baoyu Li, Kris Hauser, and Yunzhu Li. Particle-grid neural dynamics for learning
 645 deformable object models from rgb-d videos. *arXiv preprint arXiv:2506.15680*, 2025a.

647 Mingtong Zhang, Kaifeng Zhang, and Yunzhu Li. Dynamic 3d gaussian tracking for graph-based
 648 neural dynamics modeling. *arXiv preprint arXiv:2410.18912*, 2024b.

648 Mingtong Zhang, Kaifeng Zhang, and Yunzhu Li. Dynamic 3d gaussian tracking for graph-based
649 neural dynamics modeling. In *Conference on Robot Learning*, pp. 1851–1862. PMLR, 2025b.
650

651 Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable
652 effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on*
653 *computer vision and pattern recognition*, pp. 586–595, 2018.

654 Tianyuan Zhang, Hong-Xing Yu, Rundi Wu, Brandon Y Feng, Changxi Zheng, Noah Snavely, Jiajun
655 Wu, and William T Freeman. Physdreamer: Physics-based interaction with 3d objects via video
656 generation. In *European Conference on Computer Vision*, pp. 388–406. Springer, 2024c.
657

658 Licheng Zhong, Hong-Xing Yu, Jiajun Wu, and Yunzhu Li. Reconstruction and simulation of elastic
659 objects with spring-mass 3d gaussians. In *European Conference on Computer Vision*, pp. 407–
660 423. Springer, 2024.

661 Xiangming Zhu, Huayu Deng, Haochen Yuan, Yunbo Wang, and Xiaokang Yang. Latent intuitive
662 physics: Learning to transfer hidden physics from a 3d video. In *ICLR*, 2024.

663

664

665

666

667

668

669

670

671

672

673

674

675

676

677

678

679

680

681

682

683

684

685

686

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

APPENDIX

702
703
704
705
706 The content of this supplementary material includes:
707

- 708 • Details about Material Point Method (MPM) in Sec. A;
- 709 • Details of our method in Sec. B;
- 710 • [Additional ablation studies in Sec. C.](#)
- 711 • [Additional analysis and results in Sec. D.](#)

714 A MATERIAL POINT METHOD (MPM)
715

716 The Material Point Method (MPM) (Jiang et al., 2016; Stomakhin et al., 2013; Bardenhagen et al.,
717 2000) is particularly powerful for simulating complex nonlinear mechanics in highly deformable
718 objects, including soft materials (e.g. rubber, foam, biological tissues), granular media (e.g., soil,
719 snow), and fluids undergoing extreme plastic flows. Its core advantage lies in its hybrid Lagrangian-
720 Eulerian formulation: Lagrangian material points inherently track mass, momentum, complete de-
721 formation history (captured via the deformation gradient tensor), and evolving material state (stress,
722 damage) through arbitrary motion and topological changes, while a background computational grid
723 solves momentum equations on-the-fly automatically avoiding mesh entanglement issues inherent
724 to traditional mesh-based methods like FEM. This unique synergy enables MPM to natively handle
725 extreme deformations, automatic material separation/fracture, and history-dependent constitutive
726 behavior without remeshing or special failure criteria.

726 MPM operates through three key phases: Particle-to-Grid (P2G) Transfer, Grid Operations(GridOp),
727 and Grid-to-Particle (G2P) Transfer.
728

729 A.1 PARTICLE-TO-GRID (P2G) TRANSFER.
730

731 During the Particle-to-Grid (P2G) phase, material points transfer their mass and momentum to grid
732 nodes via interpolation functions. Using a fixed timestep Δt such that $t^n = n\Delta t$, the mass m_i^n at
733 grid node i at timestep n is computed as:

$$734 \quad 735 \quad 736 \quad m_i^n = \sum_p w_{ip}^n m_p \quad (10)$$

737 where m_p is the mass of material point p , and w_{ip}^n denotes the B-spline interpolation weight between
738 material point p and grid node i . Similarly, grid momentum is calculated through:

$$739 \quad 740 \quad m_i^n \mathbf{v}_i^n = \sum_p w_{ip}^n m_p (\mathbf{v}_p^n + \mathbf{C}_p^n (\mathbf{x}_i - \mathbf{x}_p^n)) \quad (11)$$

741 where \mathbf{v}_p^n is the material point's velocity, \mathbf{C}_p^n represents its affine momentum matrix, and \mathbf{x}_i , \mathbf{x}_p^n
742 denote grid node and material point positions respectively.
743

744 A.2 GRID OPERATIONS (GRIDOP).
745

746 The Grid Operations (GridOp) phase constitutes the computational core of MPM, where the gov-
747 erning equations of continuum mechanics are solved on the background grid. Following Particle-to-
748 Grid (P2G) transfer, this phase performs two critical operations:

749 First, grid velocities are updated by resolving both internal stresses and external forces:
750

$$751 \quad 752 \quad 753 \quad \hat{\mathbf{v}}_i^{n+1} = \mathbf{v}_i^n + \Delta t \left(\frac{1}{m_i^n} \sum_p \boldsymbol{\sigma}_p^n \nabla w_{ip}^n V_p^0 + \mathbf{g} \right) \quad (12)$$

754 where $\boldsymbol{\sigma}_p^n$ denotes the Cauchy stress tensor at material point p (derived from its deformation state),
755 V_p^0 is the material point's initial volume, and \mathbf{g} represents gravitational acceleration. This explicit
integration accounts for internal forces and external body forces.

756 Subsequently, boundary conditions are enforced through velocity modifications:
 757

$$\mathbf{v}_i^{n+1} = \mathcal{P}_{\partial\Omega}(\hat{\mathbf{v}}_i^{n+1}) \quad (13)$$

759 where $\mathcal{P}_{\partial\Omega}$ is a projection operator imposing constraints at domain boundaries $\partial\Omega$. For Dirichlet
 760 boundaries (e.g., collider surfaces), signed distance fields (SDF) modify velocities to satisfy pre-
 761 scribed kinematic conditions.
 762

763 A.3 GRID-TO-PARTICLE (G2P) TRANSFER. 764

765 The Grid-to-Particle (G2P) phase transfers the updated grid kinematics back to material points,
 766 completing the time step. Material point velocities are first reconstructed through grid velocity
 767 interpolation:

$$\mathbf{v}_p^{n+1} = \sum_i \mathbf{v}_i^{n+1} w_{ip}^n \quad (14)$$

770 where w_{ip}^n is the consistent B-spline interpolation weight. Material point positions are then updated
 771 via explicit advection:

$$\mathbf{x}_p^{n+1} = \mathbf{x}_p^n + \Delta t \mathbf{v}_p^{n+1} \quad (15)$$

773 To preserve local deformation characteristics and reduce numerical diffusion, the affine velocity
 774 field matrix is reconstructed:

$$\mathbf{C}_p^{n+1} = \frac{4}{(\Delta x)^2} \sum_i w_{ip}^n \mathbf{v}_i^{n+1} (\mathbf{x}_i - \mathbf{x}_p^n)^T \quad (16)$$

778 where Δx is the grid spacing. Finally, the deformation gradient tensor undergoes incremental up-
 779 dating to track finite-strain kinematics:

$$\mathbf{F}_p^{n+1} = (\mathbf{I} + \Delta t \mathbf{C}_p^{n+1}) \mathbf{F}_p^n \quad (17)$$

782 This deformation gradient update enables history-dependent constitutive evaluation in the subse-
 783 quent time step, closing the MPM computational loop.
 784

785 B METHOD DETAILS

787 B.1 VLM-ASSISTED CONSTITUTIVE MODEL SELECTION

789 Our framework leverages Qwen3 Yang et al. (2025)’s multimodal capabilities to automate constitutive
 790 model selection for MPM simulations. By processing object manipulation videos, we extract
 791 deformation sequences through optical flow variance detection. Qwen3 then analyzes the material
 792 responses observed in these videos, matching the underlying physics descriptors against our curated
 793 library of constitutive laws following Jiang et al. (2016); Xie et al. (2024). This library includes
 794 plasticity models for describing irreversible deformation and elasticity models for computing Kirch-
 795 hoff stress:

- 796 1. **Plasticity Models:** Encompass Drucker-Prager for pressure-sensitive yielding in geomaterials, Von Mises for J^2 plasticity in metals, and St. Venant-Kirchhoff for finite-strain
 797 elasticity preceding yield or purely elastic behavior.
- 800 2. **Elasticity Models:** Feature Fixed Co-rotated (FCR) for robustness under large rotations,
 801 Neo-Hookean with volumetric coupling, St. Venant-Kirchhoff for direct mapping to Green
 802 strain, Drucker-Prager Elastic with pressure-dependent stiffness and Anisotropic Hypere-
 803 lasticity for cloth.

804 Qwen3 performs model selection through structured prompting. The system receives the following
 805 instruction:

807 Analyze the video of object-hand interaction and recommend
 808 optimal MPM constitutive models from the following library:

- 809 • Elastic Models: FCR, Neo-Hookean, St. Venant-Kirchhoff,
 Anisotropic Hyperelastic

810 • Plastic Models: Von Mises, Drucker-Prager,
 811 St. Venant-Kirchhoff, Purely Elastic
 812 Provide recommendations in the format:
 813 Elastic: [model] | Plastic: [model] | Reason:
 814 [justification]

815
 816 When video evidence demonstrates matching physical behaviors, Qwen3 outputs the best-matched
 817 constitutive models from our library. We evaluated the accuracy of Qwen3 in predicting plasticity
 818 models and elasticity models across 22 scenarios. The experimental results demonstrate that Qwen3
 819 achieved perfect prediction of 100% accuracy for both plasticity models and elasticity models in all
 820 22 scenarios. Specifically, elasticity models for 8 cloth-type objects were consistently predicted as
 821 anisotropic hyperelastic, while Neo-Hookean was assigned to the other 14 objects. Additionally, the
 822 plasticity models for all objects were correctly identified as purely elastic.

823 B.2 PHYSICAL PROPERTY OPTIMIZATION

824
 825 We implement a global-to-local optimization strategy to calibrate the physical properties of an
 826 MPM-based digital twin through two sequential computational stages. The initial global homogenization
 827 stage optimizes fundamental domain-wide parameters—including Young’s modulus E ,
 828 friction coefficient μ , Poisson’s ratio ν , yield stress σ_y , and mass density ρ —to establish stable
 829 dynamic behavior. This provides robust initialization for the subsequent heterogeneous refinement
 830 stage, which selectively optimizes only E , μ , and ρ with spatial variation to ensure convergence
 831 while avoiding over-parameterization. Given the predicted point positions $\hat{X}_t = \{\hat{x}_t^{(i)}\}_{i=1}^N$ and
 832 the ground truth $X_t = \{x_t^{(i)}\}_{i=1}^N$ at time t , both stages employ identical composite loss metrics
 833 combining geometric matching and kinematic consistency:

$$\begin{aligned} \mathcal{L}_{\text{CD}}(X_t, \hat{X}_t) &= \frac{1}{|X_t|} \sum_{x \in X_t} \min_{y \in \hat{X}_t} \|x - y\| \\ &+ \frac{1}{|\hat{X}_t|} \sum_{y \in \hat{X}_t} \min_{x \in X_t} \|y - x\| \end{aligned} \quad (18)$$

$$\mathcal{L}_{\text{track}}(X_t, \hat{X}_t) = \frac{1}{3|\mathcal{V}|} \sum_{i \in \mathcal{V}} \sum_{c \in \{x, y, z\}} \ell_{\text{smooth_l1}}(\|\hat{x}_{t,c}^{(i)} - x_{t,c}^{(i)}\|) \quad (19)$$

840
 841 where \mathcal{V} represents point indices, and the smooth ℓ_1 component is defined as:

$$\ell_{\text{smooth_l1}}(d) = \begin{cases} \frac{1}{2}d^2 & \text{if } d < 1 \\ d - 0.5 & \text{otherwise} \end{cases} \quad (20)$$

842 The total optimization objective combines these metrics through weighted summation:

$$\mathcal{L}_{\text{total}} = \lambda_1 \mathcal{L}_{\text{CD}} + \lambda_2 \mathcal{L}_{\text{track}} \quad (21)$$

843 with experimentally configured weights $\lambda_1 = 1.0$ and $\lambda_2 = 0.1$.

844 B.3 FINE-TUNING PHYSICAL PROPERTY FOR GNN-BASED WORLD MODELS

845
 846 When training GNN-based world models using generated demonstrations, we concurrently input
 847 the physical properties of objects from each demonstration into the Graph Neural Network (GNN).
 848 This process yields a physics-informed GNN capable of dynamically adapting its physical dynamics
 849 predictions according to varying input physical features. To align the simulated physical properties
 850 with real-world objects for consistent kinematic behavior, we subsequently fine-tune these properties
 851 using tracking points extracted from real videos. During fine-tuning, we freeze the GNN model
 852 parameters θ while optimizing only the physical properties $\hat{\Phi}$ associated with each graph vertex.
 853 Due to the discrepancy between the number of GNN vertices and ground truth tracking points, we
 854 employ Linear Blend Skinning (LBS) to interpolate positions for all target track points from the
 855 GNN vertex outputs. The optimization utilizes a composite loss function combining unidirectional
 856 Chamfer distance and Mean Squared Error (MSE) Loss.

864 B.4 3D GAUSSIAN UPDATE THROUGH LBS
865

866 To predict the appearance of deformable objects at time t represented by a set of 3D Gaussian kernels
867 \mathcal{G}_{t+1} , we first compute the 6-DoF transformation for each vertex $\hat{\mu}_i^t \in \hat{X}_t$, consisting of a translation
868 \mathbf{T}_i^t and rotation \mathbf{R}_i^t . The translation is directly obtained from the vertex displacement:

$$869 \quad \mathbf{T}_i^t = \hat{\mu}_i^{t+1} - \hat{\mu}_i^t = \hat{\mathbf{x}}_{t+1}^i - \hat{\mathbf{x}}_t^i. \quad (22)$$

870 For 3D rotation, we estimate a rigid local rotation \mathbf{R}_i^t for each vertex by minimizing the motion
871 discrepancy of its neighborhood $\mathcal{N}(i)$ between t and $t+1$:

$$873 \quad \mathbf{R}_i^t = \arg \min_{\mathbf{R} \in \text{SO}(3)} \sum_{j \in \mathcal{N}(i)} \left\| \mathbf{R}(\hat{\mu}_j^t - \hat{\mu}_i^t) - (\hat{\mu}_j^{t+1} - \hat{\mu}_i^{t+1}) \right\|^2, \quad (23)$$

875 We then transform Gaussian kernels via Linear Blend Skinning (LBS) Huang et al. (2024b); Sumner
876 et al. (2007); Zhang et al. (2024b); Jiang et al. (2025) by interpolating transformations from their
877 neighboring GNN vertices. The 3D position and rotation of each Gaussian can be computed by:

$$878 \quad \mu_j^{t+1} = \sum_{k \in \mathcal{N}(j)} w_{jk}^t (\mathbf{R}_k^t (\mu_j^t - \hat{\mu}_k^t) + \hat{\mu}_k^t + \mathbf{T}_k^t) \quad (24)$$

$$881 \quad \mathbf{q}_j^{t+1} = \left(\sum_{k \in \mathcal{N}(j)} w_{jk}^t \mathbf{r}_k^t \right) \otimes \mathbf{q}_j^t, \quad (25)$$

883 where $\mathbf{R}_k^t \in \mathbb{R}^{3 \times 3}$ and $\mathbf{r}_k^t \in \mathbb{R}^4$ respectively denote the rotation matrix and quaternion of vertex
884 k ; \otimes represents the quaternion multiplication; $\mathcal{N}(j)$ indicates K -nearest GNN vertices to Gaussian
885 kernel j ; and w_{jk}^t defines the interpolation weights between a Gaussian μ_j^t and a corresponding
886 GNN vertex $\hat{\mu}_k^t$, which is inversely proportional to their Euclidean distance:

$$888 \quad w_{jk}^t = \frac{\|\mu_j^t - \hat{\mu}_k^t\|^{-1}}{\sum_{k \in \mathcal{N}(j)} \|\mu_j^t - \hat{\mu}_k^t\|^{-1}} \quad (26)$$

890 to ensure that spatially closer Gaussian-vertex pairs receive higher weighting influence. With the
891 updated 3D Gaussian parameters, we render the deformable objects to obtain their appearance at
892 $t+1$.

893
894 C ADDITIONAL ABLATION STUDIES
895896 C.1 BACKBONE ARCHITECTURE ABLATION
897

901 To examine the effect of backbone choice on deformable object dynamics learning, we consider two
902 alternative architectures in place of our GNN:

- 903 • **MLP backbone.** We implement a point-based MLP model following PointMLP (Ma et al.,
904 2022).
- 905 • **Transformer backbone.** We implement a Transformer model following PointTrans-
906 formerV3 (Wu et al., 2024c).

908 All variants are trained on the *double_lift_sloth* scenario under the same data, losses, and optimization
909 protocol as PhysWorld. We then evaluate their prediction accuracy and inference speed using
910 the metrics in the main paper. The results are demonstrated in Table. 7. The results show that the
911 GNN backbone achieves the best prediction accuracy across all metrics, and is substantially faster
912 at inference. We attribute this advantage to the inductive bias of GNNs for deformable dynamics:
913 the radius-graph construction and multi-step message passing explicitly model sparse, local
914 particle-particle and particle-tool interactions, which aligns well with how forces propagate in de-
915 formable media. In contrast, MLPs must learn such interactions implicitly from point features, while
916 Transformers incur dense attention overhead without explicit locality priors. These properties make
917 GNNs a practical and effective choice for real-time, material-aware deformable world modeling.

918
 919 Table 7: Comparison of different backbone architectures for action-conditioned future prediction
 920 and inference speed (FPS).

Methods	CD \downarrow	Track \downarrow	IoU \uparrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	FPS \uparrow
Transformer	0.0497	0.0656	42.43	20.291	0.891	0.133	56
MLP	0.0213	0.0343	70.13	23.206	0.914	0.084	107
Our GNN	0.0100	0.0154	78.66	24.666	0.921	0.067	784

921
 922
 923
 924
 925
 926
 927 Table 8: GNN distillation results using PhysTwin as the teacher simulator. Incorporating VMP-Gen
 928 and P^3 -Pert improves performance, while MPM-based PhysWorld achieves the best overall results.
 929
 930
 931
 932
 933
 934
 935
 936

Method	VMP-Gen & P^3 -Pert	CD \downarrow	Track \downarrow	IoU \uparrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
PhysTwin + GNN	\times	0.0296	0.0679	55.66	21.220	0.895	0.121
PhysTwin + GNN	\checkmark	0.0187	0.0320	67.87	22.686	0.911	0.082
MPM + GNN (Ours)	\checkmark	0.0100	0.0154	78.66	24.666	0.921	0.067

C.2 PHYS TWIN + GNN DISTILLATION EXPERIMENTS

To further examine the role of our data-generation modules (VMP-Gen and P^3 -Pert) and to contextualize the benefits of using an MPM-based digital twin as the teacher model, we conducted an additional set of distillation experiments in which the GNN is trained using PhysTwin as the supervisory simulator. For the *double_lift_sloth* scene, we replace the MPM simulator with PhysTwin’s spring–mass model and generate synthesized demonstrations using the same data-generation protocol as PhysWorld.

We consider two settings: (1) training on synthesized demonstrations without VMP-Gen or P^3 -Pert, and (2) training with both modules enabled. In all cases, the GNN architecture, losses, and optimization schedule are identical to those used in the main paper. Table 8 summarizes the performance of the distilled GNN. Incorporating VMP-Gen and P^3 -Pert improves the quality of PhysTwin-generated training data, leading to a stronger GNN even when PhysTwin serves as the teacher. When compared to PhysWorld, the GNN distilled from MPM still achieves the highest accuracy across all metrics. These results show that VMP-Gen and P^3 -Pert substantially enhance the effectiveness of GNN distillation even when the supervisory simulator is PhysTwin and the MPM-based digital twin provides higher-fidelity training signals, leading to the best-performing distilled GNN in our framework.

D ADDITIONAL ANALYSIS AND RESULTS

D.1 ADDITIONAL QUANTITATIVE RESULTS FOR MPPI-BASED PLANNING

We use an MPPI-based planner, where at each planning step we sample 2000 candidate trajectories, roll them out with our learned dynamics model, compute the distance between the terminal state of each trajectory and the target configuration, assign higher weights to trajectories with smaller terminal errors, and obtain the control by taking a weighted average over all trajectories; only the first action of this averaged control sequence is executed, after which we re-sample and repeat the procedure in a receding-horizon fashion. The results in Fig. 5 compare our method(GNN) with PhysTwin and MPM in two deformable object relocating tasks by plotting the distance error between the objects’s current and target positions as a function of the planning steps. Fig. 6 further compares these three methods in terms of success-rate curves under different error thresholds. A comparison of the wall-clock time required for 10-step MPPI planning is also illustrated in Table. 9, where each step samples 2000 trajectories. Thanks to the GNN’s ability to perform parallel inference over multiple trajectories, our method requires significantly less planning time while still achieving strong control performance.

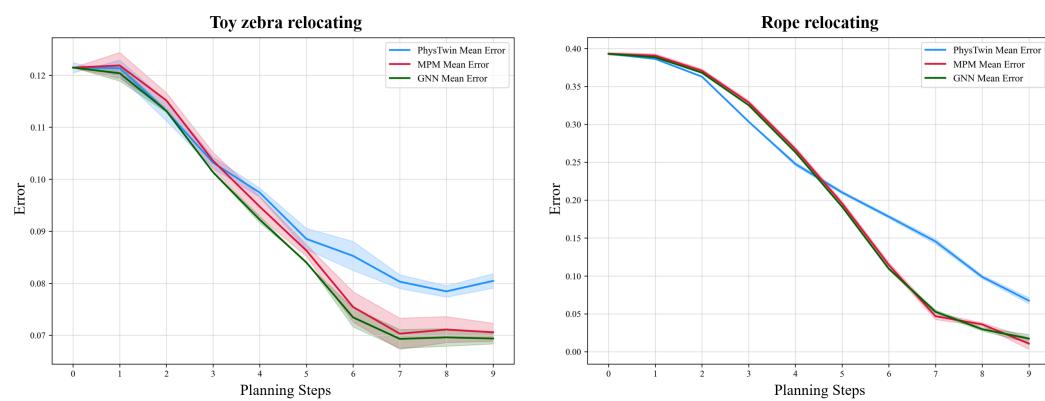


Figure 5: Mean error curves over planning steps for different methods (PhysTwin, our MPM, and our GNN) on the toy-zebra relocation (left) and rope relocation (right) tasks.

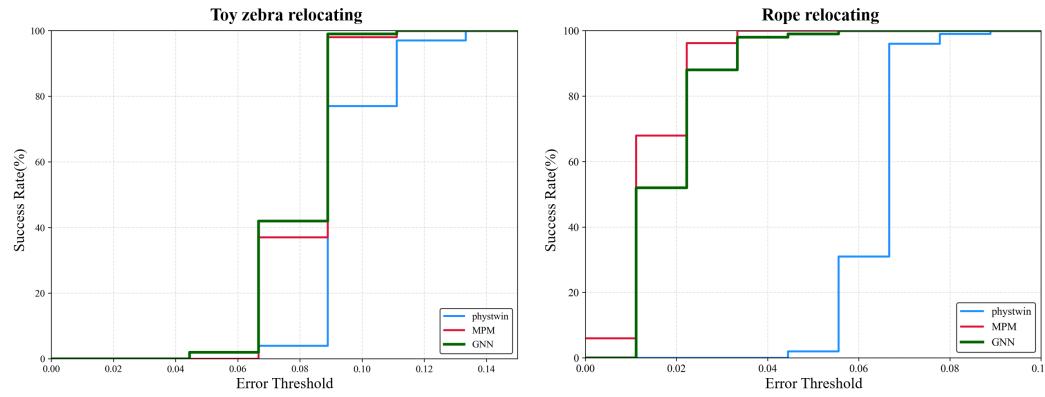


Figure 6: Success rate curves with respect to different error thresholds for different methods (PhysTwin, our MPM, and our GNN) on the toy-zebra relocation (left) and rope relocation (right) tasks.

D.2 PHYSICAL REALISM OF SYNTHESIZED DEMONSTRATIONS

To assess the physical realism of our synthesized demonstrations, we adopt the evaluation protocol of PhyGenBench (Meng et al., 2024). This benchmark measures the consistency of object motion with 27 physical laws, including those related to gravity, friction, elastic forces, and other basic dynamics. We apply these metrics to both our generated demonstrations and the real ground-truth videos. Our synthesized demonstrations achieve an average Realism score of **1.6010**, which is very close to the real videos' score of **1.6087**, indicating that the generated object motions are highly physically realistic.

D.3 ROBUSTNESS OF VLM-ASSISTED CONSTITUTIVE MODEL SELECTION

To further evaluate the robustness of our VLM-based constitutive model selection, we collected additional videos from the Internet showing hand/robot-arm interactions with deformable objects (as shown in Fig. 7) and asked Qwen3 to choose the most suitable constitutive model for the manipulated object from our material library, containing 4 elastic models and 4 plastic models. We then compared Qwen3's predictions with expert annotations and observed consistently high agreement (100%) accuracy on this extended set. This suggests that, with a reasonably designed mate-

1026
1027 **Table 9: Per-step planning time comparison for different methods.**
1028
1029
1030

Methods	PhysTwin	MPM	GNN(Ours)
Planning time per step (s)	998	4687	22

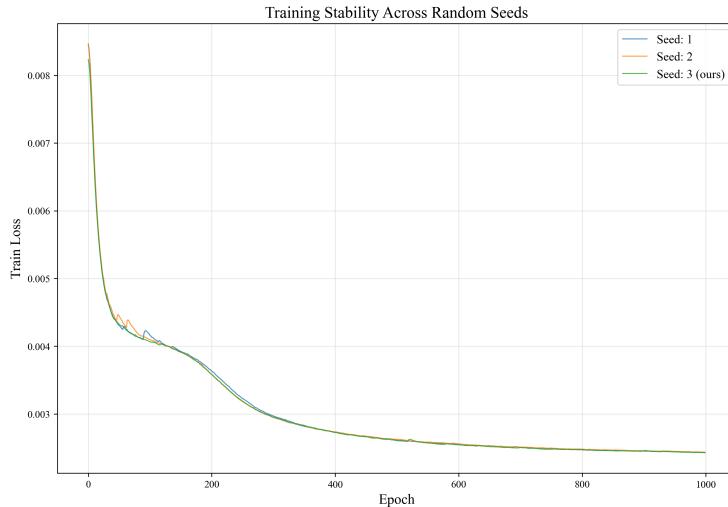
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1050
1051
1052
1053 **Figure 7: Examples of additional collected video clips showing hand/robot-arm interactions with**
1054 **deformable objects.**1055
1056
1057
1058 **rial library, constitutive-model selection using a modern VLM is a manageable subproblem in our**
1059 **pipeline.**1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079 We also note that picking a “second-best” model does not necessarily break the system. In practice, some objects can be described reasonably well by more than one constitutive law. For instance, many cloth-like objects are best matched by an anisotropic hyperelastic model, but a simpler Neo-Hookean elastic model can still yield plausible behavior for thicker fabrics. To test robustness, we conducted this study on the *single_lift_cloth_1* scene, where the theoretically optimal constitutive model is Anisotropic Hyperelastic. We intentionally replaced it with a suboptimal model-Neo-Hookean and reran the pipeline. The results in Table. 10 show that our framework can re-optimize the physical parameters under this suboptimal choice and still recover accurate dynamics, indicating tolerance to occasional non-optimal model selection.

D.4 GNN TRAINING SENSITIVITY

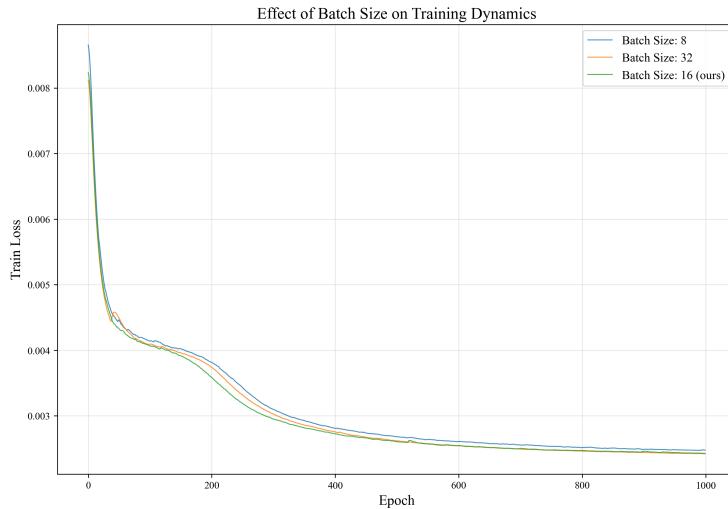
To assess the robustness of our GNN training procedure, we ran additional experiments with three different random seeds in Fig. 8 (leading to different parameter initializations) and varied the batch size and learning rate in Fig. 9 and Fig. 10 respectively. We found that changing random seeds or batch size has negligible impact on final performance, and the training loss curves consistently converge to similar values. The learning rate is also stable within a reasonable range; however, when it is set too small, convergence becomes noticeably slower (though it still reaches a comparable minimum). This indicates that PhysWorld is not overly sensitive to initialization or most hyperparameters.

1080
 1081 **Table 10: Robustness to suboptimal VLM constitutive-model selection: comparison of PhysTwin**
 1082 **and our pipeline using the optimal vs. an intentionally suboptimal constitutive model.**

1083 Methods	CD \downarrow	Track \downarrow	IoU \uparrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
1084 PhysTwin	0.0221	0.0286	55.75	23.919	0.958	0.062
1085 Ours (Suboptimal material)	0.0193	0.0339	62.15	26.512	0.968	0.046
1086 Ours (Optimal material)	0.0186	0.0241	75.54	27.700	0.974	0.039



1105
 1106 **Figure 8: Training loss curves over epochs with different random seeds.**



1125
 1126 **Figure 9: Training loss curves over epochs with different batch sizes.**

1128 D.5 OVERALL COMPUTATIONAL COST

1130 For completeness, we report the end-to-end computational cost of PhysWorld for a representative
 1131 scene, *double_lift_slope*, evaluated on a single RTX 4060 Ti (16 GB) GPU. The pipeline consists
 1132 of four stages: (1) MPM-based digital twin optimization, (2) demonstration synthesis, (3) GNN
 1133 training, and (4) GNN fine-tuning on real videos. On average, the end-to-end process requires a
 total of **8.89 h** per object, with the stage-wise breakdown summarized in Table 11. While this one-



Figure 10: Training loss curves over epochs with different learning rates.

Table 11: Overall computational cost.

Stage	Duration (h)	Share of total
MPM optimization	0.38	4.3%
Demonstration synthesis	3.00	33.7%
GNN training	5.48	61.7%
GNN fine-tuning	0.03	0.3%
Total	8.89	100%

time preparation cost is non-trivial, it yields a fast learned world model that runs orders of magnitude faster than the underlying MPM simulator and can be directly used for downstream tasks such as model-based planning and control, where high inference speed is crucial.

D.6 EVALUATION ON ELASTO-PLASTIC OBJECTS FROM ROBOCRAFT

To assess the advantage of using an MPM-based digital twin on materials beyond the elastic regime, we conducted additional experiments on elasto-plastic objects from RoboCraft (Shi et al., 2024). These objects undergo large and partially irreversible deformations, which are challenging to capture with spring-mass dynamics. For each RoboCraft sequence, we calibrate physical parameters independently for (1) an MPM simulator and (2) PhysTwin’s spring-mass simulator, using the same optimization procedure as in our main pipeline. We then roll out both calibrated simulators under the recorded actions and compare their predicted object trajectories against ground-truth point clouds. Table 12 reports quantitative errors. The calibrated MPM digital twin consistently achieves lower Chamfer, EMD, and Hausdorff losses than PhysTwin, indicating more faithful reproduction of elasto-plastic dynamics, especially under large and irreversible deformations.

Furthermore, as illustrated in Fig. 11, when the two control points push inward on the elasto-plastic object and then move back, the behaviors of the two simulators differ markedly. PhysTwin, built on a spring-mass formulation, tends to enforce elastic recovery, causing the object to automatically return toward its original shape after the push. In contrast, the calibrated MPM digital twin accurately predicts the persistent, plastic deformation induced by this interaction, producing trajectories that remain consistent with the irreversible shape changes observed in the real sequence.

1188

1189 Table 12: Comparison of calibrated MPM and PhysTwin on elasto-plastic objects from RoboCraft.

1190

1191

1192

1193

1194

1195

1196

1197

1198

1199

1200

1201

1202

1203

1204

1205

1206

1207

1208

1209

1210

1211 Figure 11: Qualitative results of PhysTwin and calibrated MPM on an elasto-plastic RoboCraft
1212 object. When two control points push inward and then retract, PhysTwin’s spring–mass simulator
1213 exhibits elastic recovery toward the original shape, whereas MPM captures the persistent plastic
1214 deformation, matching the irreversible shape changes observed in the real sequence.

1215

1216

1217

D.7 ADDITIONAL QUALITATIVE RESULTS

1218

1219

1220

1221

1222

1223

1224

1225

1226

1227

1228

1229

1230

1231

1232

1233

1234

1235

1236

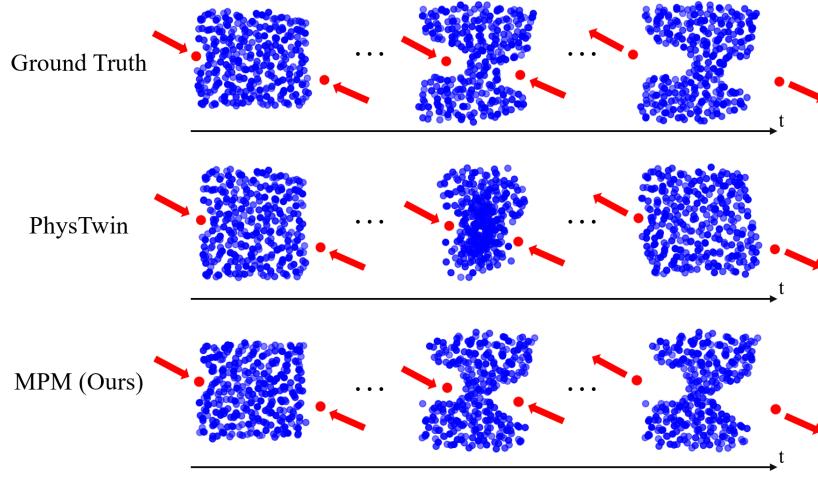
1237

1238

1239

1240

1241



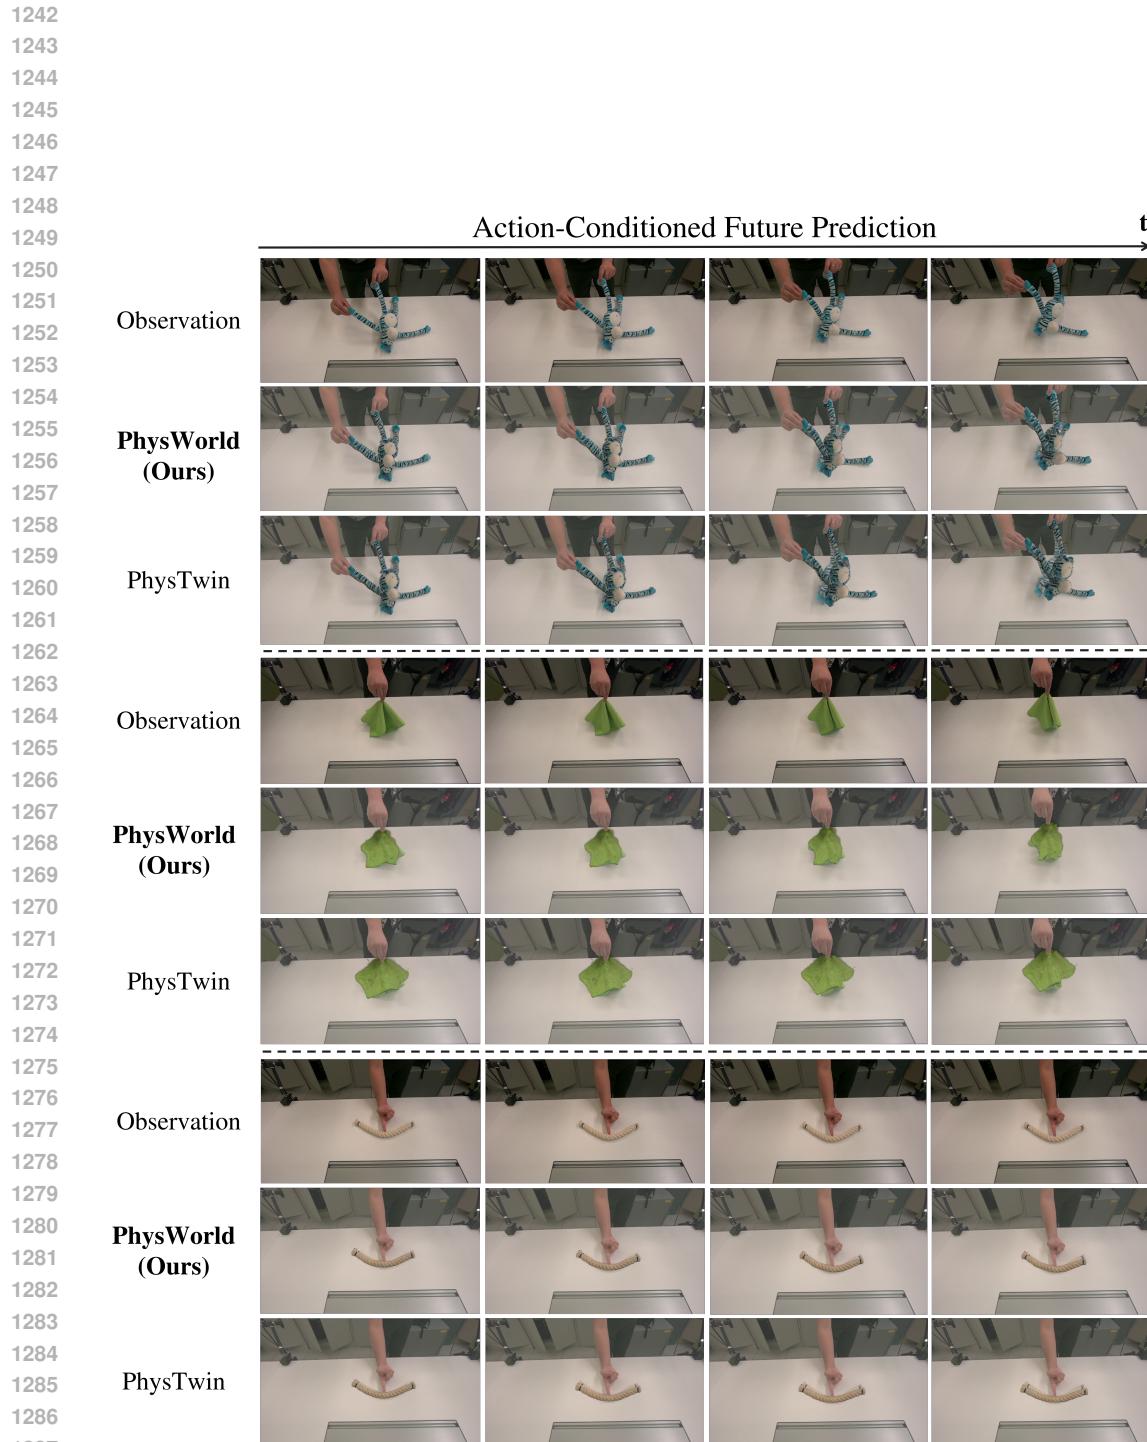


Figure 12: Additional visual results of action-conditioned future prediction. Our method's predicted positions show closer alignment with ground truth compared to PhysTwin.