

# NeuSpring: Neural Spring Fields for Reconstruction and Simulation of Deformable Objects from Videos

Qingshan Xu<sup>1\*</sup>, Jiao Liu<sup>1\*</sup>, Shangshu Yu<sup>2†</sup>, Yuxuan Wang<sup>1</sup>, Yuan Zhou<sup>1†</sup>, Junbao Zhou<sup>1</sup>, Jiequan Cui<sup>3</sup>, Yew-Soon Ong<sup>1,4</sup>, Hanwang Zhang<sup>1</sup>

<sup>1</sup>Nanyang Technological University, Singapore

<sup>2</sup>Northeastern University, China

<sup>3</sup>Hefei University of Technology, China

<sup>4</sup>A\*STAR, Singapore

## Abstract

In this paper, we aim to create physical digital twins of deformable objects under interaction. Existing methods focus more on the physical learning of current state modeling, but generalize worse to future prediction. This is because existing methods ignore the intrinsic physical properties of deformable objects, resulting in the limited physical learning in the current state modeling. To address this, we present NeuSpring, a neural spring field for the reconstruction and simulation of deformable objects from videos. Built upon spring-mass models for realistic physical simulation, our method consists of two major innovations: 1) a piecewise topology solution that efficiently models multi-region spring connection topologies using zero-order optimization, which considers the material heterogeneity of real-world objects. 2) a neural spring field that represents spring physical properties across different frames using a canonical coordinate-based neural network, which effectively leverages the spatial associativity of springs for physical learning. Experiments on real-world datasets demonstrate that our NeuSpring achieves superior reconstruction and simulation performance for current state modeling and future prediction, with Chamfer distance improved by **20%** and **25%**, respectively.

**Code** — <https://github.com/GhiXu/NeuSpring>

## Introduction

Reconstruction and simulation of deformable objects is a fundamental problem in physical AI, with wide applications in robotics, animation, virtual reality, etc. The objective is to recover the geometry, appearance and physical properties of an object under interactions, allowing not only the reconstruction to closely match the input observations but also the future simulation to conform the realistic physical laws. This is critical for future prediction and planning. However, because of the complex physics involved in real-world deformable objects under interactions, it is difficult to recover the accurate physical properties of deformable objects, posing great challenges to the reconstruction and simulation of deformable objects.

\*These authors contributed equally.

†Corresponding authors

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Many works (Park et al. 2021; Kratimenos, Lei, and Daniilidis 2024; Yang et al. 2024; Wu et al. 2024) have put effort into dynamic 3D reconstruction from videos. These methods focus on reconstructing the geometry, appearance and motion from input videos but ignore the underlying physics of deformable objects. This greatly limits their ability to predict future states. To model the intricate physics of deformable objects, there exist mainly two lines of research. Some neural dynamics methods (Zhang, Zhang, and Li 2024; Ma et al. 2023) adopt neural networks to model object dynamics, which take advantage of the powerful fitting capabilities of neural networks to learn complex physics. These pure learning-driven methods omit the basic physics modeling and require extensive training samples. Moreover, these methods are often limited to specific objects or motions.

As another line of research, physics-driven methods (Xie et al. 2024; Zhong et al. 2024; Jiang et al. 2025) build their physics models upon some powerful simulators, such as Material Point Method (MPM) and spring-mass models. These methods typically assume that objects exhibit certain regular physical properties, enabling the use of physical models—such as continuum mechanics or Newtonian mechanics—to describe their motion. By reconstructing the underlying properties and parameters of these models from video data, it becomes possible to reconstruct and predict the motion of deformable bodies. However, accurately identifying the appropriate parameters from sparse observations remains a significant challenge. A common issue is that, while such methods often align well with the observed motion in the input data, they tend to **generalize poorly** to future states or novel scenes, as illustrated in the top of Figure 1. This phenomenon can be attributed to:

- *Independent learning*: existing physics-driven methods often adopt independent physics properties learning for each physical element, ignoring their spatial associativity and lacking structural regularization;
- *Homogeneous assumption*: these methods often assume that the material of an object is homogeneous, resulting in a single material parameter or topology.

This raises the key question we ask in this work: *how to accurately learn the physics properties of deformable objects from the input observations?*

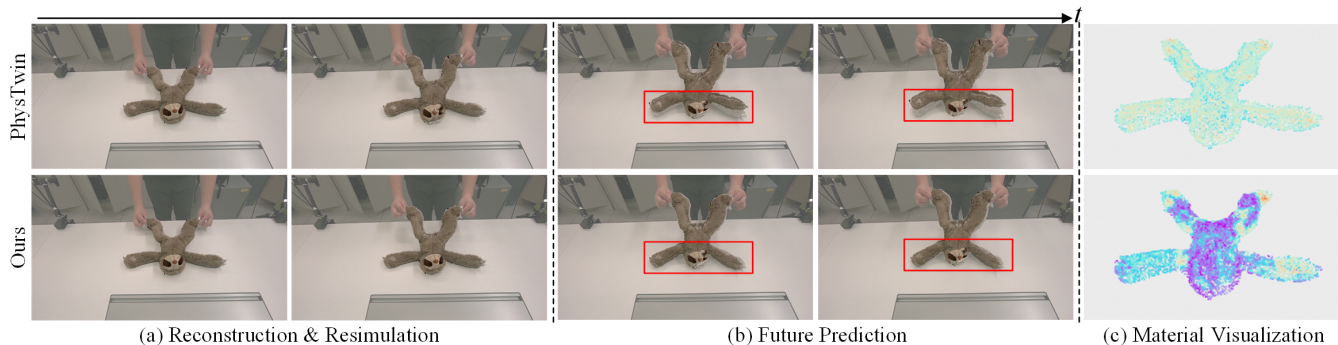


Figure 1: Motivation of our NeuSpring. We show the reconstruction & resimulation and future prediction performance for PhysTwin (Jiang et al. 2025) and our proposed method. We also visualize the approximated material based on the estimated physical property of springs–stiffness.

To address this problem, based on spring-mass models, we introduce a novel implicit neural representation—**neural spring field** for reconstruction and simulation of deformable objects from videos, dubbed NeuSping. However, combining neural fields with spring-mass models to learn physics properties from videos is challenging due to: 1) spring-mass models contain non-differentiable parameters—spring connection topologies; 2) the spatial position of the spring varies across video frames. To this end, we introduce a two-stage learning strategy:

- In the first stage, we present a piecewise topology solution. It divides deformable objects into different parts and models multi-region spring connection topologies under the assumption of homogeneous spring physics properties, with efficient zero-order optimizers.
- In the second stage, we propose a canonical coordinate-based neural network to represent the physics properties of springs across different frames. We take as input the midpoint between spring endpoints in the canonical space, and estimate the spring physical properties through our carefully designed neural network.

With the multi-region spring topologies and canonical spring physical properties, physical motions across frames are simulated and supervised by input trajectories to optimize spring properties. Integrated with the geometry and appearance modeled by 3D Gaussian splatting (Kerbl et al. 2023), our proposed NeuSpring enables accurate reconstruction and simulation of deformable objects from videos.

Our approach inherits the benefits of neural methods and physics-driven methods: the expressive power of neural networks and the reduced requirement for training samples enabled by physical priors. Crucially, our method models the **material heterogeneity** of real-world objects and the **spatial associativity** of physics properties. As a result, our method can not only reconstruct dynamic 3D objects from input videos, but also generalize well to future prediction, as illustrated in the bottom of Figure 1. In particular, NeuSpring achieves **20%** and **25%** improvements in Chamfer distance for current state modeling and future prediction, respectively, demonstrating superior physical learning, see

Table 1. In summary, our main contributions are as follows:

- We present a neural spring field for reconstruction and simulation of deformable objects from videos.
- We develop a canonical coordinate-based neural network to learn the physical properties of springs across different frames. It leverages the spatial continuity and coherence of springs to facilitate physical learning.
- We propose a piecewise topology solution to model multi-region spring connection topologies, which enables us to consider material heterogeneity.

## Related Work

**Dynamic Scene Reconstruction.** Reconstructing dynamic scenes from videos has achieved significant advances in recent years (Kratimenos, Lei, and Daniilidis 2024; Li et al. 2023; Park et al. 2021; Tretschk et al. 2021). Different from static scene reconstruction, dynamic scene reconstruction needs to consider the integration of motion modeling. Inspired by Neural Radiance Fields (NeRF) (Mildenhall et al. 2021), early approaches (Park et al. 2021; Pumarola et al. 2021) mainly focus on training a canonical NeRF and optimizing a deformable field to model dynamic scenes. Recently, 3D Gaussian Splatting (GS) (Kerbl et al. 2023) has emerged as a powerful technique due to its real-time efficiency and high-quality rendering, driving significant progress in dynamic scene reconstruction. Similarly, GS-based dynamic scene reconstruction methods (Yang et al. 2024; Wu et al. 2024) usually deform canonical 3D Gaussian primitives over time by optimizing deformation networks. Since these methods model motion information by conditioning the deformable field on timestep, they cannot capture real-world physics dynamics, preventing them from predicting future states and interactive modeling.

**Neural Dynamics Modeling of Deformable Objects.** Several studies have explored the use of neural networks to approximate the dynamics of deformable objects (Chen et al. 2022; Evans, Thankaraj, and Pinto 2022; Ma et al. 2023; Wu et al. 2019; Xu et al. 2019). A key advantage of these approaches lies in their ability to model complex state spaces and capture diverse physical properties. In particular, graph

neural networks have been applied to simulate the dynamics of various materials, including plasticine (Shi et al. 2023, 2024), cloth (Lin et al. 2022; Pfaff et al. 2020), fluids (Li et al. 2018), and stuffed animals (Zhang, Zhang, and Li 2024). More recently, physics-informed neural networks (Raissi, Perdikaris, and Karniadakis 2019; Wong et al. 2022; Jin et al. 2024; Wong, Gupta, and Ong 2021) have also shown promise in modeling complex physical behaviors (Xu et al. 2024b). However, as noted in (Jiang et al. 2025), these methods often require large training data and are typically limited to a small number of predefined scenarios, which significantly constrains their generalizability.

**Physics-Driven Simulation of Deformable Objects.** Physical simulation is fundamental to physical AI (Karniadakis et al. 2021; Chiu et al. 2022), enabling physically plausible and expressive object deformations (de Avila Belbute-Peres et al. 2018). Classical models such as the mass-spring system (Liu et al. 2013), SPH (Monaghan 2005), FEM (Chen and Zeltzer 1992), and MPM (Jiang et al. 2016) rely on prescanned or clean geometric data for realistic results. Recent advances in 3D representations (Mildenhall et al. 2021; Kerbl et al. 2023; Xu et al. 2024c,a) have integrated physics into geometric learning frameworks, including SDFs (Xu et al. 2024b, 2025), NeRF (Feng et al. 2024), and Gaussians (Xie et al. 2024). However, these approaches typically handle only small deformations under fixed forces, limiting real-world applicability. The most related works, Spring-Gaussian (Zhong et al. 2024) and PhysTwin (Jiang et al. 2025), reconstruct deformable objects using mass-spring systems. Yet, Spring-Gaussian suffers from over-regularization, while PhysTwin’s separately learned parameters and homogeneity assumptions restrict physical accuracy. These limitations motivate our NeuSpring framework.

## Preliminary: Spring-Mass Model

The spring-mass model is a classical framework for modeling the physical behavior of elastic and deformable objects. In this model, a deformable object is represented using a system of springs and masses, which can be formulated as a graph structure  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ , where  $\mathcal{V}$  denotes the set of mass points defining the object’s geometry, and  $\mathcal{E}$  represents the set of springs connecting mass points. Both the mass points and springs are associated with physical parameters that characterize their behavior. Each mass point  $i \in \mathcal{V}$  is associated with a position  $\mathbf{x}_i \in \mathbb{R}^3$  and a velocity  $\mathbf{v}_i \in \mathbb{R}^3$ , which describe its state of motion at a given time step  $t$ . Each spring  $e \in \mathcal{E}$  connects two mass points and is characterized by physical properties such as Young’s modulus and damping ratio, which govern its elastic and dissipative behavior. This model provides a simple yet powerful basis for simulating the dynamics of deformable bodies.

In general, the spring-mass model simulates object deformation based on Newton’s laws of motion. By computing the force acting on each mass point at every time step, the position and velocity of the mass point can be updated for the next time step. The force acting on the  $i$ -th mass point in  $\mathcal{V}$  is determined by the combined effects of the neighboring

mass points connected via springs:

$$\mathbf{F}_i = \sum_{i,j \in \mathcal{E}} \left( \mathbf{F}_{i,j}^{\text{spring}} + \mathbf{F}_{i,j}^{\text{dashpot}} \right) + \mathbf{F}_i^{\text{ext}}, \quad (1)$$

where the spring force  $\mathbf{F}_{i,j}^{\text{spring}}$  is calculated as:

$$\mathbf{F}_{i,j}^{\text{spring}} = k_{ij} \left( \|\mathbf{x}_j - \mathbf{x}_i\| - l_{ij} \right) \frac{\mathbf{x}_j - \mathbf{x}_i}{\|\mathbf{x}_j - \mathbf{x}_i\|}, \quad (2)$$

the dashpot damping force  $\mathbf{F}_{i,j}^{\text{dashpot}}$  is calculated as:

$$\mathbf{F}_{i,j}^{\text{dashpot}} = -\gamma(\mathbf{v}_i - \mathbf{v}_j), \quad (3)$$

and the external force  $\mathbf{F}_i^{\text{ext}}$  is incorporated to account for factors such as gravity, collisions, and user interactions. In Eqs. (2) and (3),  $k_{ij}$  denotes the stiffness of the spring connecting mass points  $(i, j)$ ,  $l_{ij}$  represents its rest length, and  $\gamma$  is the damping coefficient associated with the dashpot. For collision handling—whether between mass points or between the object and external colliders—an impulse-based approach is employed.

The state of the spring system is updated using explicit Euler integration. Specifically, for all  $i \in \mathcal{V}$ , the velocities and positions at time step  $t$  are updated as follows:

$$\begin{aligned} \mathbf{v}_i^{t+1} &= \delta \left( \mathbf{v}_i^t + \Delta t \cdot \frac{\mathbf{F}_i}{m_i} \right), \\ \mathbf{x}_i^{t+1} &= \mathbf{x}_i^t + \Delta t \cdot \mathbf{v}_i^{t+1}, \end{aligned} \quad (4)$$

where  $\delta$  is the drag damping. The simulation results of the spring-mass system are highly sensitive to physical parameters such as spring stiffness, damping, collision responses, and spring topology. This sensitivity becomes particularly pronounced with heterogeneous springs, where each spring may have distinct physical properties. As a result, the number of parameters grows substantially, posing significant challenges for parameter optimization in spring-mass-model-based dynamic 3D reconstruction.

## Problem Formulation

Our NeuSpring takes as input three RGBD videos of a deformable object under interaction to learn the appearance, geometry and physical properties for the deformable object. Inspired by (Jiang et al. 2025), we formulate the dynamic reconstruction task as a two-stage optimization problem.

In the first stage, we aim to enable the spring-mass model to accurately capture both the geometry and the deformation behavior of the target object. To this end, we jointly optimize the **physical property parameters**  $\alpha$  (e.g., spring stiffness, damping coefficients, collision parameters, etc.) and the **topology** of the spring-mass model  $\mathcal{G}$ , allowing the physical system to support a faithful and expressive reconstruction of dynamic deformable objects. At each time frame  $t$ , the depth observation  $\mathbf{D}_t$  obtained from the camera is converted into a partial 3D point cloud  $\mathbf{X}_t$ , which serves as a reference for adjusting the spring-mass system so that its simulation aligns with the observed data. Formally, the optimization problem for the first stage is defined as:

$$\begin{aligned} \min_{\alpha, \mathcal{G}} \sum_t \left( \mathcal{L}_{\text{geometry}}(\mathbf{X}_t, \hat{\mathbf{X}}_t) + \mathcal{L}_{\text{motion}}(\mathbf{X}_t, \hat{\mathbf{X}}_t) \right), \\ \text{s.t. } \hat{\mathbf{X}}_{t+1} = f_{\alpha, \mathcal{G}}(\hat{\mathbf{X}}_t, a_t), \end{aligned} \quad (5)$$

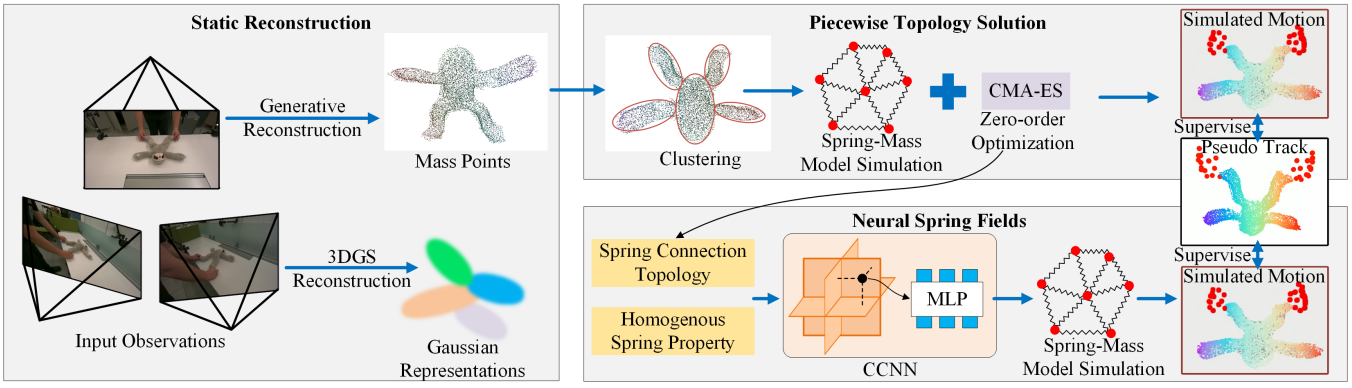


Figure 2: Overview of Our NeuSpring. CCNN: Canonical Coordinate-based Neural Network.

where  $\mathcal{L}_{\text{geometry}}$  quantifies the single-direction Chamfer distance between the inferred state  $\hat{\mathbf{X}}_t$  and the partially observed point cloud  $\mathbf{X}_t$ , encouraging the simulated geometry to align with the observed shape. Meanwhile,  $\mathcal{L}_{\text{motion}}$  measures the motion tracking error between the predicted point  $\hat{\mathbf{x}}_i^t$  and its corresponding observed trajectory point  $\mathbf{x}_i^t$ , enforcing temporal consistency of the simulated motion. The state of the spring-mass system, i.e., the velocity and position of the mass points, is updated via  $f_{\alpha, \mathcal{G}}$  by using the explicit Euler integration. Following (Jiang et al. 2025), the observed trajectories are extracted using the vision foundation model CoTracker3 (Karaev et al. 2025), and then lifted into 3D space via depth map unprojection.

In the second stage, since the spring-mass model captures only physical states without rendering, we establish a mapping from mass points to image space to bridge physics and observation. This ensures alignment between the simulated object’s **appearance** and the observed frames for accurate dynamic reconstruction. Let  $g_\phi$  denote the mapping function from the mass points to the rendered image, where  $\phi$  represents the learnable parameters. The goal of the optimization in this stage is to minimize the discrepancy between the rendered predictions and the corresponding ground-truth observations. Accordingly, the overall optimization problem can be formulated as:

$$\min_{\phi} \sum_{t,c} \mathcal{L}_{\text{render}}(\mathbf{I}_{t,c}, \hat{\mathbf{I}}_{t,c}), \text{ s.t. } \hat{\mathbf{I}}_{t,c} = g_\phi(\hat{\mathbf{X}}_t, c), \quad (6)$$

where  $c$  denotes the camera index,  $\mathbf{I}_{t,c}$  and  $\hat{\mathbf{I}}_{t,c}$  are the ground truth and rendered images from view  $c$  at time  $t$ , respectively. The loss term  $\mathcal{L}_{\text{render}}$  combines the pixel-wise  $L_1$  loss with a D-SSIM component to encourage both photometric accuracy and structural consistency between the rendered and ground truth images.

In physics-driven dynamic reconstruction, physical simulation serves as the backbone for modeling the dynamic deformation of an object. The objective of this work is to improve the simulation fidelity by introducing NeuSpring, a more accurate and adaptable physical model for complex deformations. Consequently, our focus lies on the first-stage optimization, while the second stage follows the method in

(Jiang et al. 2025) without modification. The following sections detail the proposed NeuSpring framework.

## Method

Figure 2 illustrates the pipeline of NeuSpring. Given input sparse RGB inputs, we first utilize the 3D generative model TRELIS (Xiang et al. 2025) to generate mass points. Then, we adopt our piecewise topology solution to build spring topologies and initialize the physical properties of springs. Next, these outcomes are input to our neural spring fields proposed to further optimize physical properties. In this way, the dynamics of the mass points is obtained to guide the dynamic rendering of Gaussian representations.

### Neural Spring Fields

As analyzed above, the core of reconstruction and simulation of deformable objects is to estimate accurate physical properties for deformable objects, which capture real-world physics dynamics. Previous spring-mass-based dynamic reconstruction methods (Jiang et al. 2025; Zhong et al. 2024) typically adopt independent learning to optimize the physical property parameters of each spring. This ignores the spatial associativity of spring properties, leading to the limited physics learning in the current state modeling. To address this problem, based on spring-mass models, we introduce a neural spring field to learn the physical properties of springs from input observations.

Our neural spring field is conditioned on the spatial information of springs so as to consider the spatial associativity of physical properties. Specifically, our neural spring field  $\mathcal{S}$  is defined in the canonical space to represent the physical properties of springs across different frames. For each spring  $e \in \mathcal{E}$  that connects two mass points  $(i, j)$ , we compute the midpoint between these points to represent the spatial information of each spring:

$$\mathbf{x}_{\text{mid}}(e) = \frac{(\mathbf{x}_i + \mathbf{x}_j)}{2}. \quad (7)$$

Then, we use a neural network  $F_\theta$  to learn the physical properties of each spring by querying its midpoint coordinate:

$$\mathcal{S}(e) = F_\theta(\mathbf{x}_{\text{mid}}(e)), \quad (8)$$

where  $\theta$  represents learnable parameters. More concretely, the  $F_\theta$  adopts a tri-plane representation (Chan et al. 2022) which is a hybrid explicit-implicit network architecture and is efficient for evaluation. It contains three 2D feature planes ( $\mathbf{f}_{xy}, \mathbf{f}_{yz}, \mathbf{f}_{xz}$ ) with a spatial resolution of  $N \times N$  and  $C$  feature channels each and a shallow Multi-Layer Perceptron (MLP). To learn  $\mathcal{S}(e)$  at position  $\mathbf{x}_{\text{mid}}(e)$ , we project  $\mathbf{x}_{\text{mid}}(e)$  onto each of the three planes and query the corresponding features on each plane by bilinear interpolation. By aggregating these features through the sum operation and passing them into the MLP with the Fourier feature transform, we obtain  $\mathcal{S}(e)$ .

### Piecewise Topology Solution

In the spring field, the input is defined as the midpoints between pairs of connected points, which requires prior knowledge of the spring topology. Since this topological information is not encoded inherently in the neural spring field itself, it must be determined before training the neural spring field. Moreover, the spring topology significantly impacts simulation behaviors and thus affects dynamic reconstruction accuracy. In previous work (Jiang et al. 2025; Zhong et al. 2024), K-nearest neighbor (KNN) algorithms are commonly employed to construct topological structures by connecting nearby mass points with springs. The choice of KNN hyperparameters determines the final topology.

While KNN provides an effective mechanism for defining connectivity, using a uniform set of KNN hyperparameters across the entire system results in a homogeneous spring connection density. However, real-world objects often exhibit heterogeneous material properties, and such uniform connection schemes may fail to capture this variability.

To address this limitation, we propose a piecewise topology construction method for the spring-mass system. Specifically, we first cluster the mass points in the spring system. After clustering, we assign independent KNN hyperparameters—namely, the maximum number of neighbors and the search radius—to each cluster. This allows different regions of the object to have varying spring connection densities, thereby better reflecting material heterogeneity. To determine the optimal KNN hyperparameters for each cluster, we treat these hyperparameters as decision variables and solve the optimization problem defined in Eq. (5). Since the KNN hyperparameters are non-differentiable, we adopt a gradient-free optimization approach. In particular, we utilize the Covariance Matrix Adaptation Evolution Strategy (CMA-ES), a widely used evolutionary algorithm, to solve this problem.

It is important to note that evaluating the objective function in Eq. (5) requires not only the spring system topology but also the physical parameters of the springs. Therefore, during the optimization process, we assume homogeneous spring properties and include these physical parameters as part of the decision variables. The resulting optimized topology and homogeneous physical parameters are then used as the initialization for training the neural spring field. In this way, our neural spring field is modified as:

$$\mathcal{S}(e) = \mathcal{S}_0 + F_\theta(\mathbf{x}_{\text{mid}}(e)), \quad (9)$$

where  $\mathcal{S}_0$  is the obtained physical parameters in this phase,

which provides a good initialization for our neural spring field learning.

### Reconstruction and Simulation with NeuSpring

With our proposed NeuSpring, the physical properties parameters are learned to support the dynamic rendering of deformable objects. Specifically, following the practices in (Jiang et al. 2025; Zhong et al. 2024), we adopt 3D Gaussian Splatting (Kerbl et al. 2023) to compute parameterized Gaussian points in the canonical frame, which models the static appearance and geometry of deformable objects. For the appearance and geometry in other frames, we leverage Linear Blend Skinning (LBS) to interpolate the motions of Gaussian points based on the motions of neighboring mass points (Zhang, Zhang, and Li 2024; Jiang et al. 2025). By splatting these moved Gaussian points and implementing  $\alpha$ -blending, we efficiently render deformable objects across frames for realistic reconstruction and simulation.

## Experiments

### Experimental Settings

**Dataset.** We evaluate NeuSpring on the dataset collected by PhysTwin (Jiang et al. 2025), which contains 22 scenarios. Each scenario provides three RGBD videos capturing human interactions with various deformable objects with different physical properties. These scenarios can be categorized into six types of deformable objects, including cloth, ropes, packages, stuffed zebra, stuffed sloth and stuffed dinosaur. For each scenario, the RGBD videos are split into a training set and a testing set following a 7:3 ratio. The training set is used to train our NeuSpring for reconstruction and resimulation, while the test set is used to evaluate its generalization ability, i.e., future prediction.

**Baselines.** We compare our NeuSpring with state-of-the-art methods, including GS-Dynamics (Zhang, Zhang, and Li 2024), Spring-Gaus (Zhong et al. 2024) and PhysTwin (Jiang et al. 2025). The first one is a neural dynamics method while the last two are physics-driven methods.

**Evaluation Metrics.** To evaluate how our predictions match the observations, we adopt Chamfer Distance (CD), tracking error and image assessment metrics. The image assessment metrics include Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index Measure (SSIM) (Wang et al. 2004), Learned Perceptual Image Patch Similarity (LPIPS) (Zhang et al. 2018) and silhouette alignment using IoU.

**Implementation Details.** Our implementation is based on PyTorch framework (Paszke et al. 2019) and the simulation is built on the warp implementation of spring-mass models (Macklin 2022). In the piecewise topology solution stage, we empirically select five clusters to represent flexible topologies of deformable objects with agglomerative clustering. In terms of our canonical coordinate-based neural network architecture, the spatial resolution of tri-plane representation is adaptively set to  $N = 0.85 \cdot \sqrt{|\mathcal{E}|}$ , and the feature channel is set to  $C = 32$ . We adopt a 3-layer MLP with 128 hidden units. All experiments are performed on one NVIDIA RTX 3090 GPU.

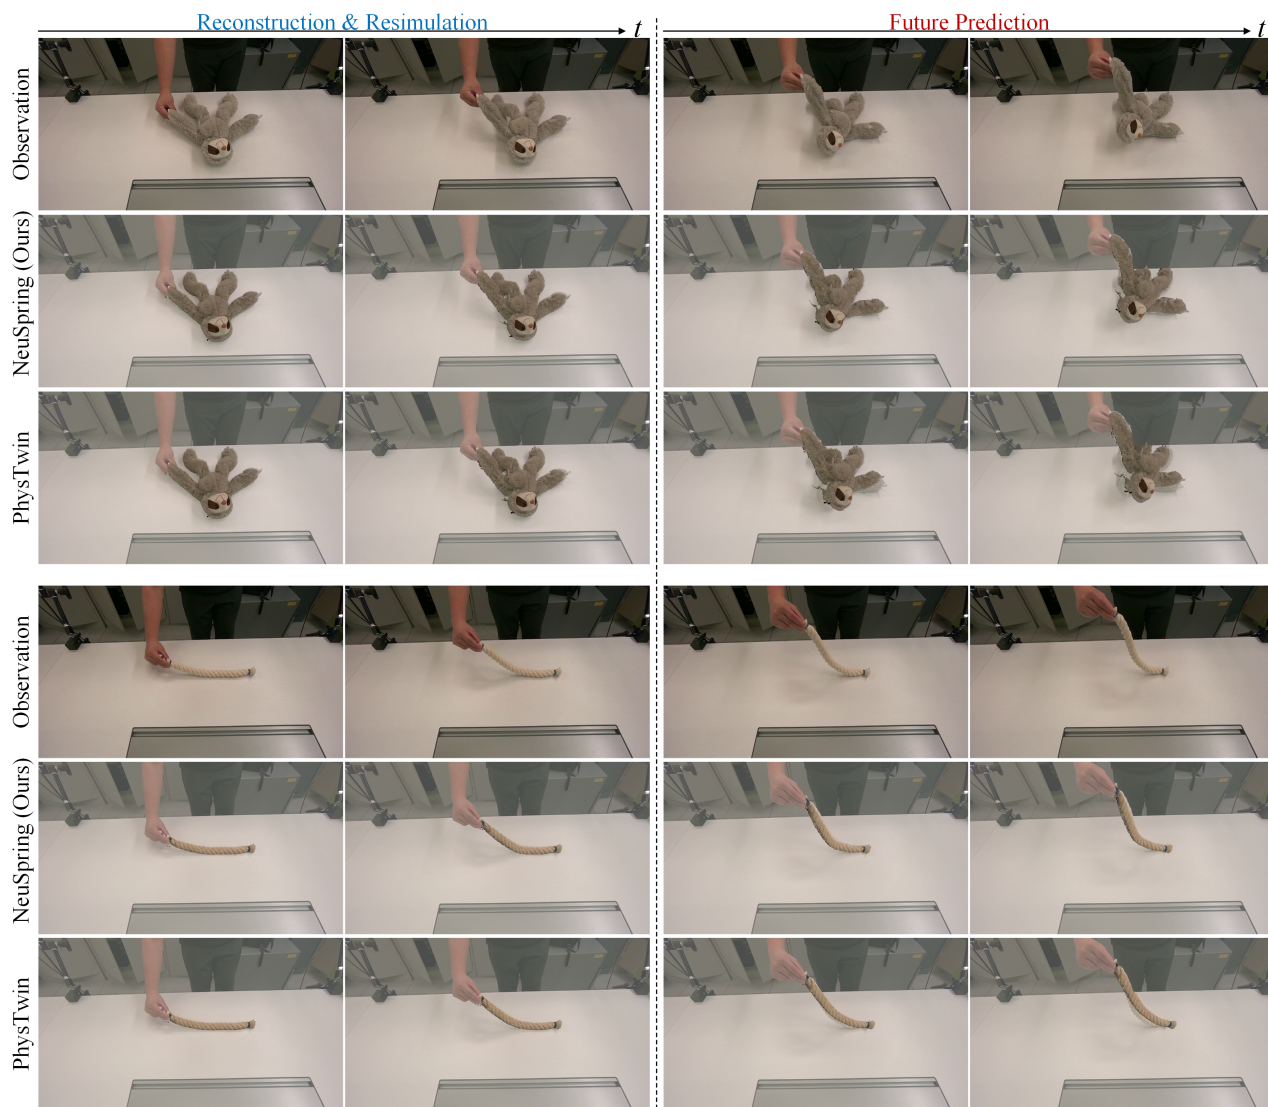


Figure 3: Qualitative Results on Reconstruction & Resimulation and Future Prediction.

## Comparison

**Reconstruction & Resimulation.** The left column of Table 1 reports the reconstruction & resimulation results of different methods. We observe that our method achieves the best performance across all metrics. In particular, NeuSpring outperforms PhysTwin by **20.0%**, **11.1%**, **2.1%** and **8.8%** in CD, tracking error, IoU and LPIPS, respectively. As shown in Eq. (5), the CD and tracking error metrics reflect the accuracy of physical learning, our noticeable improvements in these two metrics demonstrate that NeuSpring can better learn the physical properties of deformable objects. Moreover, our better physical learning further boosts the performance of dynamic rendering. The qualitative results in Figure 3 further validate the effectiveness of our NeuSpring.

**Future Prediction.** The right column of Table 1 further shows the future prediction capability of different methods, which is an important component for digital twins. As can

be seen, our method outperforms baseline methods across all metrics. In particular, NeuSpring outperforms PhysTwin by **25.0%**, **22.7%**, **4.5%** and **10.9%** in CD, tracking error, IoU and LPIPS, respectively. We show qualitative visualizations of the future prediction of different methods in Figure 3. We note that although PhysTwin can achieve competitive reconstruction and resimulation, it cannot align its future predictions well with the observations in unseen frames. This demonstrates that the learning of PhysTwin is limited to the seen frames. In contrast, NeuSpring can not only reconstruct seen frames well, but also predict unseen frames precisely. This provides strong evidence that our learned physics dynamics is more accurate than the baseline methods.

## Analysis

**Ablation Study.** To evaluate the effect of our proposed contributions, we adopt PhysTwin as our baseline model to carry

Task	Reconstruction & Resimulation						Future Prediction					
	CD ↓	Track Error ↓	IoU % ↑	PSNR ↑	SSIM ↑	LPIPS ↓	CD ↓	Track Error ↓	IoU % ↑	PSNR ↑	SSIM ↑	LPIPS ↓
Spring-Gaus	0.041	0.050	57.6	23.445	0.928	0.102	0.062	0.094	46.4	22.488	0.924	0.113
GS-Dynamics	0.014	0.022	72.1	26.260	0.940	0.052	0.041	0.070	49.8	22.540	0.924	0.097
PhysTwin	<u>0.005</u>	<u>0.009</u>	<u>84.4</u>	<u>28.214</u>	<u>0.945</u>	<u>0.034</u>	<u>0.012</u>	<u>0.022</u>	<u>72.5</u>	<u>25.617</u>	<u>0.941</u>	<u>0.055</u>
NeuSpring (Ours)	<b>0.004</b>	<b>0.008</b>	<b>86.2</b>	<b>28.664</b>	<b>0.947</b>	<b>0.031</b>	<b>0.009</b>	<b>0.017</b>	<b>75.9</b>	<b>26.361</b>	<b>0.944</b>	<b>0.049</b>

Table 1: Quantitative Results on Reconstruction & Resimulation and Future Prediction. Our NeuSpring framework consistently outperforms the baselines across all metrics.

Task	Method	PT NSF	Reconstruction & Resimulation						Future Prediction					
			CD ↓	Track Error ↓	IoU % ↑	PSNR ↑	SSIM ↑	LPIPS ↓	CD ↓	Track Error ↓	IoU % ↑	PSNR ↑	SSIM ↑	LPIPS ↓
Baseline			0.0054	0.0095	83.5	28.033	0.944	0.035	0.0118	0.0230	70.2	25.281	0.939	0.058
Model-A	✓		0.0050	0.0091	84.8	28.278	0.945	0.033	0.0114	0.0223	72.3	25.720	0.941	0.054
Model-B		✓	0.0048	0.0085	83.9	28.114	0.945	0.034	0.0101	0.0205	71.6	25.638	0.940	0.056
NeuSpring (Ours)	✓	✓	<b>0.0043</b>	<b>0.0077</b>	<b>86.2</b>	<b>28.664</b>	<b>0.947</b>	<b>0.031</b>	<b>0.0087</b>	<b>0.0175</b>	<b>75.9</b>	<b>26.361</b>	<b>0.944</b>	<b>0.049</b>

Table 2: Ablations of Our NeuSpring. PT: Piecewise Topology solution. NSF: Neural Spring Fields.




Figure 4: Effect of the Cluster Number.

out ablation studies. Note that, we rerun the baseline model in our machine. Then, we progressively replace the corresponding modules in PhysTwin by our designs to investigate their efficacy. The results are reported in Table 2.

By adding our designed piecewise topology solution to the baseline, all metrics are improved. This shows that this design better considers the material heterogeneity of deformable objects, providing more flexibility for topology optimization. Therefore, it facilitates physical learning. On the other hand, our proposed neural spring fields enhance both the baseline and Model-A. This strategy introduces spatial associativity for springs, serving as a structural regularizer for physical learning. Therefore, the improvements in the CD and tracking error metrics are more pronounced.

**Effect of the Cluster Number.** Figure 4 shows the effect of the cluster number on two scenarios. As the number of clusters increases, the CD metrics of reconstruction & resimulation and future prediction almost decrease. This further demonstrates that equipping the spring-mass model with piecewise topologies is beneficial to physical learning. Generally, more clusters will allow the spring-mass model to consider more topology structures, providing more flexibility. However, this increases the number of parameters to be optimized, thus exacerbating the complexity of zeroth-order optimization. As a result, it typically requires more function evaluations and exhibits slower convergence. Thus, we em-

Task	Reconstruction & Resimulation				Future Prediction			
	CD ↓	TE ↓	IoU % ↑	LPIPS ↓	CD ↓	TE ↓	IoU % ↑	LPIPS ↓
w/o TP	0.0052	0.0091	84.1	0.035	0.0112	0.0216	72.1	0.056
w/o AR	<u>0.0045</u>	<u>0.0083</u>	<u>85.7</u>	<u>0.032</u>	<u>0.0095</u>	<u>0.0186</u>	<u>74.9</u>	<u>0.052</u>
Ours	<b>0.0043</b>	<b>0.0077</b>	<b>86.2</b>	<b>0.031</b>	<b>0.0087</b>	<b>0.0175</b>	<b>75.9</b>	<b>0.049</b>

Table 3: Effect of Network Design. TP: Tri-Plane. AR: Adaptive Resolution of TP. TE: Track Error.

pirically set the cluster number to five for all scenarios.

**Effect of Network Design.** To investigate how our designed network learns the spatial associativity of springs, we conduct two experiments: removing the tri-plane representation (TP) and replacing adaptive resolution (AR) with fixed resolution  $N = 256$ . The results are reported in Table 3. We observe that, without explicit spatial modeling, the pure MLP cannot effectively learn the spatial associativity for springs, obviously degrading the performance. With AR, our method can better consider the shape of deformable objects to model the spatial associativity, further improving the performance.

## Conclusion

In this paper, we propose Neural Spring fields, NeuSpring, for reconstruction and simulation of deformable objects from videos. We design the piecewise topology solution and neural spring field learning to consider the material heterogeneity and spatial associativity for spring-mass models, respectively. This facilitates the physical learning of deformable objects, enabling more accurate reconstruction and simulation. Extensive experiments demonstrate that NeuSpring achieves state-of-the-art performance in current state modeling and future prediction, offering valuable insights to construct effective physical digital twin for deformable objects from sparse videos.

## Acknowledgments

This research is supported by National Research Foundation, Singapore, NRF-NRFI10-2024-0004, and the National Research Foundation (NRF), Singapore through the AI Singapore Programme under the project titled “AI-based Urban Cooling Technology Development” (Award No. AISG3-TC-2024-014-SGKR).

## References

- Chan, E. R.; Lin, C. Z.; Chan, M. A.; Nagano, K.; Pan, B.; De Mello, S.; Gallo, O.; Guibas, L. J.; Tremblay, J.; Khamis, S.; et al. 2022. Efficient geometry-aware 3d generative adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 16123–16133.
- Chen, D. T.; and Zeltzer, D. 1992. Pump it up: computer animation of a biomechanically based model of muscle using the finite element method. *SIGGRAPH Comput. Graph.*, 26(2): 89–98.
- Chen, Z.; Yi, K.; Li, Y.; Ding, M.; Torralba, A.; Tenenbaum, J. B.; and Gan, C. 2022. Comphy: Compositional physical reasoning of objects and events from videos. *arXiv preprint arXiv:2205.01089*.
- Chiu, P.-H.; Wong, J. C.; Ooi, C.; Dao, M. H.; and Ong, Y.-S. 2022. CAN-PINN: A fast physics-informed neural network based on coupled-automatic-numerical differentiation method. *Computer Methods in Applied Mechanics and Engineering*, 395: 114909.
- de Avila Belbute-Peres, F.; Smith, K.; Allen, K.; Tenenbaum, J.; and Kolter, J. Z. 2018. End-to-End Differentiable Physics for Learning and Control. In Bengio, S.; Wallach, H.; Larochelle, H.; Grauman, K.; Cesa-Bianchi, N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc.
- Evans, B.; Thankaraj, A.; and Pinto, L. 2022. Context is everything: Implicit identification for dynamics adaptation. In *2022 International Conference on Robotics and Automation (ICRA)*, 2642–2648. IEEE.
- Feng, Y.; Shang, Y.; Li, X.; Shao, T.; Jiang, C.; and Yang, Y. 2024. PIE-NeRF: Physics-based Interactive Elastodynamics with NeRF. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 4450–4461.
- Jiang, C.; Schroeder, C.; Teran, J.; Stomakhin, A.; and Selle, A. 2016. The material point method for simulating continuum materials. In *ACM SIGGRAPH 2016 Courses*, SIGGRAPH '16. New York, NY, USA: Association for Computing Machinery. ISBN 9781450342896.
- Jiang, H.; Hsu, H.-Y.; Zhang, K.; Yu, H.-N.; Wang, S.; and Li, Y. 2025. Phystwin: Physics-informed reconstruction and simulation of deformable objects from videos. *ICCV*.
- Jin, G.; Wong, J. C.; Gupta, A.; Li, S.; and Ong, Y.-S. 2024. Fourier warm start for physics-informed neural networks. *Engineering Applications of Artificial Intelligence*, 132: 107887.
- Karaev, N.; Makarov, Y.; Wang, J.; Neverova, N.; Vedaldi, A.; and Ruppert, C. 2025. Cotracker3: Simpler and better point tracking by pseudo-labelling real videos. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 6013–6022.
- Karniadakis, G. E.; Kevrekidis, I. G.; Lu, L.; Perdikaris, P.; Wang, S.; and Yang, L. 2021. Physics-informed machine learning. *Nature Reviews Physics*, 3(6): 422–440.
- Kerbl, B.; Kopanas, G.; Leimkühler, T.; and Drettakis, G. 2023. 3D Gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4): 139–1.
- Kratimenos, A.; Lei, J.; and Daniilidis, K. 2024. Dynmf: Neural motion factorization for real-time dynamic view synthesis with 3d gaussian splatting. In *European Conference on Computer Vision*, 252–269. Springer.
- Li, Y.; Wu, J.; Tedrake, R.; Tenenbaum, J. B.; and Torralba, A. 2018. Learning particle dynamics for manipulating rigid bodies, deformable objects, and fluids. *arXiv preprint arXiv:1810.01566*.
- Li, Z.; Wang, Q.; Cole, F.; Tucker, R.; and Snavely, N. 2023. Dynibar: Neural dynamic image-based rendering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4273–4284.
- Lin, X.; Wang, Y.; Huang, Z.; and Held, D. 2022. Learning visible connectivity dynamics for cloth smoothing. In *Conference on Robot Learning*, 256–266. PMLR.
- Liu, T.; Bargteil, A. W.; O’Brien, J. F.; and Kavan, L. 2013. Fast simulation of mass-spring systems. *ACM Trans. Graph.*, 32(6).
- Ma, P.; Chen, P. Y.; Deng, B.; Tenenbaum, J. B.; Du, T.; Gan, C.; and Matusik, W. 2023. Learning neural constitutive laws from motion observations for generalizable pde dynamics. In *International Conference on Machine Learning*, 23279–23300. PMLR.
- Macklin, M. 2022. Warp: A high-performance python framework for gpu simulation and graphics. In *NVIDIA GPU Technology Conference (GTC)*, volume 3.
- Mildenhall, B.; Srinivasan, P. P.; Tancik, M.; Barron, J. T.; Ramamoorthi, R.; and Ng, R. 2021. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1): 99–106.
- Monaghan, J. J. 2005. Smoothed particle hydrodynamics. *Reports on Progress in Physics*, 68(8): 1703.
- Park, K.; Sinha, U.; Barron, J. T.; Bouaziz, S.; Goldman, D. B.; Seitz, S. M.; and Martin-Brualla, R. 2021. Nerfies: Deformable neural radiance fields. In *Proceedings of the IEEE international conference on computer vision*, 5865–5874.
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32.
- Pfaff, T.; Fortunato, M.; Sanchez-Gonzalez, A.; and Battaglia, P. 2020. Learning mesh-based simulation with graph networks. In *International conference on learning representations*.

- Pumarola, A.; Corona, E.; Pons-Moll, G.; and Moreno-Noguer, F. 2021. D-nerf: Neural radiance fields for dynamic scenes. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10318–10327.
- Raissi, M.; Perdikaris, P.; and Karniadakis, G. E. 2019. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational physics*, 378: 686–707.
- Shi, H.; Xu, H.; Clarke, S.; Li, Y.; and Wu, J. 2023. Robocook: Long-horizon elasto-plastic object manipulation with diverse tools. *arXiv preprint arXiv:2306.14447*.
- Shi, H.; Xu, H.; Huang, Z.; Li, Y.; and Wu, J. 2024. Robocraft: Learning to see, simulate, and shape elasto-plastic objects in 3d with graph networks. *The International Journal of Robotics Research*, 43(4): 533–549.
- Tretschk, E.; Tewari, A.; Golyanik, V.; Zollhöfer, M.; Lassner, C.; and Theobalt, C. 2021. Non-rigid neural radiance fields: Reconstruction and novel view synthesis of a dynamic scene from monocular video. In *Proceedings of the IEEE international conference on computer vision*, 12959–12970.
- Wang, Z.; Bovik, A. C.; Sheikh, H. R.; and Simoncelli, E. P. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4): 600–612.
- Wong, J. C.; Gupta, A.; and Ong, Y.-S. 2021. Can transfer neuroevolution tractably solve your differential equations? *IEEE Computational Intelligence Magazine*, 16(2): 14–30.
- Wong, J. C.; Ooi, C. C.; Gupta, A.; and Ong, Y.-S. 2022. Learning in sinusoidal spaces with physics-informed neural networks. *IEEE Transactions on Artificial Intelligence*, 5(3): 985–1000.
- Wu, G.; Yi, T.; Fang, J.; Xie, L.; Zhang, X.; Wei, W.; Liu, W.; Tian, Q.; and Wang, X. 2024. 4d gaussian splatting for real-time dynamic scene rendering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 20310–20320.
- Wu, Y.; Yan, W.; Kurutach, T.; Pinto, L.; and Abbeel, P. 2019. Learning to manipulate deformable objects without demonstrations. *arXiv preprint arXiv:1910.13439*.
- Xiang, J.; Lv, Z.; Xu, S.; Deng, Y.; Wang, R.; Zhang, B.; Chen, D.; Tong, X.; and Yang, J. 2025. Structured 3d latents for scalable and versatile 3d generation. In *Proceedings of the Computer Vision and Pattern Recognition*, 21469–21480.
- Xie, T.; Zong, Z.; Qiu, Y.; Li, X.; Feng, Y.; Yang, Y.; and Jiang, C. 2024. Physgaussian: Physics-integrated 3d gaussians for generative dynamics. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4389–4398.
- Xu, Q.; Cui, J.; Yi, X.; Wang, Y.; Zhou, Y.; Ong, Y.-S.; and Zhang, H. 2024a. Pushing Rendering Boundaries: Hard Gaussian Splatting. *arXiv preprint arXiv:2412.04826*.
- Xu, Q.; Liu, J.; Wong, M.; Chen, C.; and Ong, Y.-S. 2024b. Precise-physics driven text-to-3d generation. *arXiv preprint arXiv:2403.12438*.
- Xu, Q.; Liu, J.; Wong, M.; Chen, C.; and Ong, Y.-S. 2025. Looks great, functions better: Physics compliance text-to-3d shape generation.
- Xu, Q.; Yi, X.; Xu, J.; Tao, W.; Ong, Y.-S.; and Zhang, H. 2024c. Few-shot nerf by adaptive rendering loss regularization. In *European Conference on Computer Vision*, 125–142. Springer.
- Xu, Z.; Wu, J.; Zeng, A.; Tenenbaum, J. B.; and Song, S. 2019. Densephysnet: Learning dense physical object representations via multi-step dynamic interactions. *arXiv preprint arXiv:1906.03853*.
- Yang, Z.; Gao, X.; Zhou, W.; Jiao, S.; Zhang, Y.; and Jin, X. 2024. Deformable 3d gaussians for high-fidelity monocular dynamic scene reconstruction. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 20331–20341.
- Zhang, M.; Zhang, K.; and Li, Y. 2024. Dynamic 3D Gaussian Tracking for Graph-Based Neural Dynamics Modeling. *arXiv preprint arXiv:2410.18912*.
- Zhang, R.; Isola, P.; Efros, A. A.; Shechtman, E.; and Wang, O. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 586–595.
- Zhong, L.; Yu, H.-X.; Wu, J.; and Li, Y. 2024. Reconstruction and simulation of elastic objects with spring-mass 3d gaussians. In *Proceedings of the European Conference on Computer Vision*, 407–423. Springer.