

# Learning Physics-Guided Residual Dynamics for Deformable Object Simulation

Shivansh Patel<sup>1,2</sup> Kaifeng Zhang<sup>2,3</sup> Sanjay Pokkali<sup>1</sup> Svetlana Lazebnik<sup>1</sup> Yunzhu Li<sup>2,3</sup>

<sup>1</sup>UIUC

<sup>2</sup>SceniX Inc.

<sup>3</sup>Columbia University

**Abstract**—Simulating deformable objects is essential for a wide range of robotic manipulation applications. Yet, accurately predicting their dynamics remains challenging. Physics-based simulations are interpretable, but they require a known model class and precise identification of physical parameters, which are often difficult to obtain. Learning-based approaches can be more expressive in capturing dynamics, but they often generalize poorly outside the training distribution, require substantial training data, and lack physical consistency. We propose Physics-Guided Residual Dynamics (PReD), a hybrid simulation framework that combines an optimizable spring–mass simulator as a backbone with a learned neural network that predicts residual corrections to compensate for discrepancies between physics-based predictions and reality. Our velocity-based formulation ensures stable simulation, while a sliding-window transformer captures temporal dependencies. We validate our approach on diverse real-world deformable objects, demonstrating that PReD outperforms both purely physics-based simulators and learning-based methods. We further demonstrate the practical utility of our framework in action-conditioned 3D video prediction using 3D Gaussian Splatting and in Model Predictive Control for manipulation planning on challenging tasks such as cable rerouting, where purely physics-based simulation fails.

## I. INTRODUCTION

Accurately predicting the dynamics of deformable objects remains a persistent challenge in robotics and computer vision. These can deform significantly under external forces, undergoing changes in shape such as stretching, bending, crumpling, and twisting, often complicated by self-collisions and friction. The difficulty stems not only from material behaviors such as heterogeneous elasticity and damping, but also from the complexity of contact-rich interactions.

To address this challenge of modeling these complex behaviors, two dominant paradigms have emerged: physics-based methods and learning-based approaches. Physics-based methods use mathematical equations that describe how materials deform and respond to forces, yielding physically plausible, interpretable simulations. However, these methods require precise material parameters that are difficult to obtain in practice [42, 77, 53]. The complex mathematics of these simulators typically restricts practitioners to gradient-free optimization methods [42], which can effectively tune only a handful of parameters before becoming intractable [77, 86]. Even with optimal parameters, physics-based models are inherently limited by their discretization and constitutive modeling assumptions. Coarse meshes miss fine-scale deformations [63, 48], and linear elasticity fails under large strains [69, 49].

On the other hand, learning-based approaches have shown promise by directly learning from data without requiring

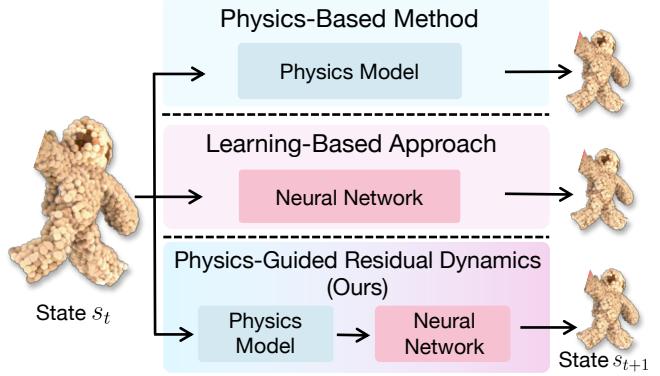


Fig. 1: **Comparison of simulation paradigms for deformable objects.** Physics-based methods (top) are interpretable but limited by discretization and constitutive modeling assumptions. Learning-based methods (middle) can capture complex dynamics but often generalize poorly beyond the training distribution. Physics-Guided Residual Dynamics (bottom) combines a physics model with a learned neural network that predicts residual velocity corrections, enabling accurate and robust simulation.

physics simulations. These methods can capture complex, arbitrary behaviors that are challenging to model analytically. However, purely data-driven models suffer from poor generalization to unseen scenarios [12, 65], require substantial amounts of training data [79], and often lack the physical consistency [12] required for reliable deployment in real-world applications. Moreover, without the inductive bias provided by physical priors, these models may learn spurious correlations that do not reflect the real-world [68].

We propose Physics-Guided Residual Dynamics (PReD), a hybrid simulation framework that combines interpretability and generalization of physics-based models with flexibility and expressiveness of learning-based approaches. PReD uses a spring–mass simulator as its physics backbone and follows a two-stage training procedure: first, we optimize the simulator’s parameters to match real-world observations; second, we train a neural network to predict residual corrections that compensate for discrepancies between the physics model and the real world. Crucially, training residual dynamics models is nontrivial. Directly injecting learned corrections into simulated states can destabilize the dynamics and cause error accumulation over time. Instead, we predict residual velocities and integrate them forward in time, yielding residual position corrections that respect the underlying dynamical structure. This velocity-based formulation ensures smooth, stable simulation and enables effective training of the residual model without destabilizing the underlying physics simulation. To model temporal

dependencies, we incorporate a sliding-window transformer that refines velocity corrections by modeling residuals across multiple timesteps. This temporal history enables the network to learn dynamic phenomena, such as momentum, while a gating mechanism ensures stable training.

PReD offers several advantages. First, the physics backbone provides priors that improve generalization to new scenarios. Second, by leveraging a physics backbone, our model requires significantly less training data than purely learning-based approaches, as the network only needs to learn corrections rather than learn to predict the entire dynamics from scratch. Third, our approach maintains computational efficiency suitable for real-time applications while achieving higher accuracy than physics-based methods alone.

We validate our approach on a diverse set of real-world deformable objects, including rope, paper, plush toys, flag, and duster. Our experiments demonstrate that PReD significantly outperforms both purely physics-based methods and purely learning-based approaches in tracking accuracy across prehensile and non-prehensile manipulation. Notably, only PReD performs well on the duster, which is a heterogeneous object with rigid stem and soft feathers. Here, uniform material assumptions in most physics-based deformable simulators fail and learning-based methods struggle with the complexity.

By attaching 3D Gaussians to the simulated particles, we can render photorealistic images and evaluate performance using both vision-based and geometric metrics. Beyond state tracking, we demonstrate two practical applications: action-conditioned 3D video prediction for photorealistic rendering, and model predictive control for manipulation planning on challenging tasks such as cable rerouting, where purely physics-based simulation fails.

In summary, our contributions are as follows: (1) We propose Physics-Guided Residual Dynamics, a hybrid simulation framework that combines an optimizable spring-mass model with learned residual corrections to accurately model deformable object dynamics. (2) We introduce a two-stage training procedure that first optimizes physics parameters using black-box optimization and then trains a neural network to predict residual corrections. (3) We conduct extensive experiments on diverse real-world deformable objects, demonstrating that our approach achieves superior performance compared to existing physics-based and learning-based methods.

## II. RELATED WORKS

**Physics-Based Simulation for Deformable Objects.** Traditional physics-based simulation methods rely on analytical models to discretize deformable objects and numerical solvers to solve equations of motion of the models. For deformable objects, spring-mass models [4, 41, 37] represent one of the most intuitive and efficient approaches, where objects are modeled as networks of point masses interconnected by springs. Consequently, this representation has been extensively employed to effectively model the dynamics of diverse deformable objects in both computer graphics and robotics [4, 41, 37, 84, 42, 83, 30]. Our work leverages the

spring-mass model as a backbone precisely because of its qualities of efficiency and interpretability. While more complex approaches like Finite Element Methods (FEM) [61, 51], Position-Based Dynamics (PBD) [52, 47], and Material Point Methods (MPM) [62, 29] offer higher physical fidelity, they often incur high computational costs and complexity.

Despite their physical foundations, these methods require precise material parameters that are difficult to obtain in practice. A recent line of work addresses this challenge through differentiable physics approaches [37, 64, 24, 9, 13, 18, 23, 28, 55, 56, 44, 40, 16, 30], which reformulate physics simulators to be compatible with automatic differentiation frameworks. These methods enable gradient-based optimization to identify optimal physics parameters by backpropagating through the simulation process, solving inverse problems to find material properties that best match observed data. While this parameter identification can improve simulation accuracy, it imposes a strong requirement: the simulator must be fully differentiable. This constraint is often impractical in contact-rich robotics tasks where discontinuities from collisions [8, 38, 84], friction mode transitions [38, 32, 21], and non-smooth contact geometry [70, 78, 38] lead to exploding or vanishing gradients. Unlike these approaches, our method does not require the simulator to be differentiable, enabling broader applicability to realistic manipulation scenarios.

**Learning-Based Simulation for Deformable Objects.** This has emerged as a powerful alternative to traditional physics-based methods, particularly excelling in scenarios requiring real-time performance with complex nonlinear dynamics [76, 10, 34, 85, 75, 6, 1, 25, 11, 15, 45, 73, 74, 43, 26]. These approaches leverage data-driven models, typically neural networks, to learn deformation patterns directly from observation data, bypassing the need for explicit material parameters or complex numerical solvers [35, 50].

Graph Neural Networks (GNNs) have become particularly prominent in this domain due to their natural ability to represent mesh structures as graphs, where nodes correspond to vertices and edges represent connectivity. GNNs excel at capturing long-range interactions and complex dependencies through message passing [80, 82, 57, 54, 35, 66, 36], allowing them to effectively simulate internal and external forces causing deformation [5, 33]. Notably, GBND [82] exploits this graph structure by establishing a topology over sparse particles, enabling the GNN to efficiently learn and propagate the dynamics across the object. However, these methods face significant challenges: the effectiveness of message passing is sensitive to graph structure and is vulnerable to partial observations, while insufficient steps fail to capture global information and excessive steps cause oversmoothing [67, 3].

Recent advances have explored architectural improvements, including transformer-based models for handling dense particle systems [58, 71], recurrent neural networks for temporal consistency [46], and attention mechanisms for improved long-range dependency modeling [57]. These methods have demonstrated success across diverse materials, including cloth [37,

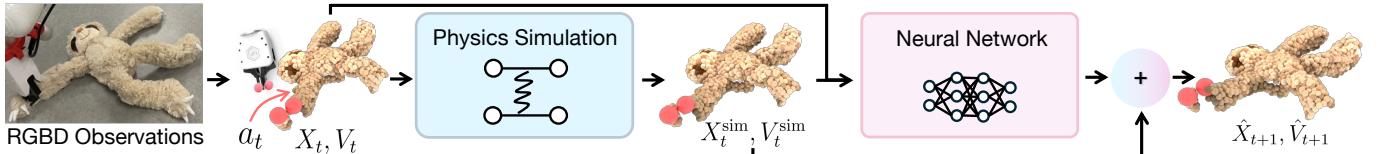


Fig. 2: **Physics-Guided Residual Dynamics Overview.** Given RGBD observations, we extract surface points to instantiate the simulation. Red dots indicate the point where gripper held the object. The framework rolls out an optimized spring-mass physics backbone from the current object state  $X_t, V_t$  and robot actions  $a_t$  to produce simulation prediction  $X_t^{\text{sim}}, V_t^{\text{sim}}$ . Subsequently, our residual dynamics network predicts per particle residual velocities, which are added to the simulator velocities and time integrated to obtain the final positions  $\hat{X}_{t+1}, \hat{V}_{t+1}$ .

39, 54], fluids [57, 35], plasticine [27, 60, 59], and granular materials [35]. The current best method is Particle-Grid Neural Dynamics [81] (PGND), which combines particle representations with spatial grids to learn dynamics while maintaining spatial continuity. These learning-based approaches typically require substantial training data and may struggle with generalization to unseen scenarios or material properties that are significantly different from the training distributions. Our work addresses these limitations by grounding the neural network in physical priors, allowing it to focus on learning only residual corrections rather than the full dynamics from scratch.

**Residual Dynamics.** Residual learning has been effectively employed to model the dynamics of robots [19, 20, 7]. In these works, the robot’s model is known, and the residuals primarily reflect simple modeling errors, such as PID gains or backlash. These discrepancies are easier to capture because they often manifest as consistent deviations from the nominal model. In contrast, we model the dynamics of deformable objects, where the residual must compensate for complex, high-dimensional phenomena, including nonlinear stiffness, spatially varying material properties, and contact-rich interactions. The most closely related work is [2], which models the environment’s residual dynamics rather than the robot’s. However, it is limited to low-dimensional toy problems, such as planar pushing, where the state space is small and the dynamics are relatively simple. PReD extends this residual paradigm to the high-dimensional, complex states of deformable objects.

### III. METHOD

We propose Physics-Guided Residual Dynamics, which augments a physics-based dynamics backbone with a learned neural network. Given RGBD observations, we extract surface points to instantiate the simulation. Our method overview is shown in Fig. 2. We first explain our physics backbone, followed by the neural network.

#### A. Physics Backbone

We represent deformable objects using a spring-mass model, wherein the object is discretized into a graph structure consisting of point masses (nodes) interconnected by springs (edges). Each node  $i$  has a position  $\mathbf{x}_i \in \mathbb{R}^3$  and velocity  $\mathbf{v}_i \in \mathbb{R}^3$  that evolve over time according to Newtonian mechanics with a semi-implicit Euler solver. This representation provides computational efficiency and straightforward implementation while capturing essential elastic deformation behaviors. To model manipulation, we apply constraints to particles near the gripper: for prehensile manipulation, they move rigidly

with the gripper, whereas for non-prehensile manipulation, the gripper acts as a spherical collision shape that pushes penetrating particles to its surface.

The fidelity of the spring-mass simulator depends critically on five physical parameters collectively denoted as  $\theta$ : 1) *Stiffness*, which defines the elastic resistance of springs; 2) *Damping*, which controls energy dissipation during motion; 3) *Threshold*, which determines the maximum distance for creating springs between nodes; 4) *Max springs per node*, which limits the connectivity of each mass point for simulation stability; and 5) *Ground friction*, which governs contact interactions with the environment. These parameters collectively encode the material properties and environmental conditions that govern the dynamics of the simulated object.

Given a batch of trajectories sampled from the dataset, each with initial state  $(\mathbf{X}_0, \mathbf{V}_0)$  and action sequences  $\mathcal{A} = \{a_t\}_{t=1}^T$ , we roll out the simulator for  $T$  time steps with actions  $\{a\}_{t=1}^T$  under candidate parameters  $\theta$  to obtain predicted point cloud positions  $\{\hat{\mathbf{X}}_t\}_{t=1}^T$ . At each time step  $t$ , we compute a pointwise mean squared error between predicted and ground truth configurations. We average this error over the trajectory horizon to obtain the optimization objective. We minimize this objective using Covariance Matrix Adaptation Evolution Strategy (CMA-ES) [22], which is particularly effective for our problem as it optimizes parameters without requiring gradients. We use CMA-ES because contacts and collisions cause sudden changes in forces, preventing reliable gradient computation in the simulator.

For volumetric objects such as plush toys, modeling only the surface shell often leads to unrealistic collapse, as the lack of internal support points causes the object to flatten under gravity or compression. To address this, we augment the spring-mass model with internal points that provide structural volume. Starting from our multi-view observations, we first reconstruct the complete surface point cloud using RaySt3R [14] and extract a watertight mesh via marching cubes. We then uniformly sample particles within this mesh to populate the interior, ensuring the spring-mass model preserves the object’s volume and structural integrity during simulation.

#### B. Residual Dynamics Framework

While our optimized spring-mass backbone provides a physically grounded baseline, it inherently struggles to capture complex behaviors such as nonlinear stiffness, non-uniform contact friction, and heterogeneous material properties. Instead of attempting to model these intricacies analytically, we introduce a learned residual module that models the discrepancy

between the physics model and reality.

Formally, at each time step  $t$  and step size  $dt$ , we use the physics backbone to generate an immediate prediction  $(X_t^{\text{sim}}, V_t^{\text{sim}})$ . We then employ a neural network to predict a per-particle residual velocity  $\Delta V_t^{\text{residual}}$  based on the current simulator state and history. We choose a velocity-based residual formulation because it ensures smooth integration and avoids the instability often associated with direct position corrections. The final velocities are obtained by adding the learned residuals to the simulator’s predictions:

$$\hat{V}_{t+1} = V_t^{\text{sim}} + \Delta V_t^{\text{residual}}.$$

These corrected velocities are then time-integrated to obtain the final position updates:

$$\hat{X}_{t+1} = X_t^{\text{sim}} + \Delta V_t^{\text{residual}} \cdot dt.$$

For particles that are rigidly held by grippers, we enforce a boundary condition by zeroing out their residuals. This formulation effectively bridges the gap between the optimized spring-mass model and the observed real-world dynamics.

We train the residual network using a supervised learning objective. To ensure the model is robust to error accumulation over long horizons, we employ a multi-step rollout training scheme. During training, we do not reset the simulator to the ground truth state at every step. Instead, we feed the hybrid simulator’s *predicted* state from time  $t$  back as the input for time  $t+1$ . This exposes the network to its own past predictions, allowing it to learn how to recover from drift.

### C. Network Architecture

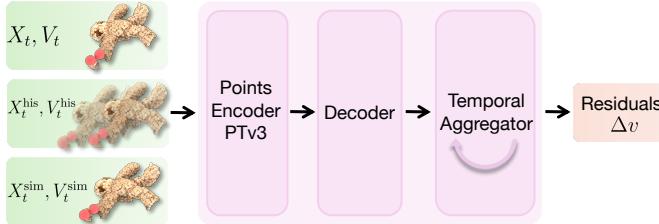


Fig. 3: **Network Architecture.** The points encoder extracts spatiotemporal features, which the decoder projects to estimate initial residual velocity. Finally, the temporal aggregator refines these estimates using a recurrent history to output the final residuals  $\Delta v$ .

Our network is shown in Fig. 3. To parameterize the residual velocity, we design a network consisting of three primary components: encoder, decoder, and temporal aggregator.

**Encoder.** We employ a Point Transformer V3 (PTv3) [72] architecture to extract per-particle features. At time step  $t$ , we construct input features by concatenating the current state  $(X_t, V_t)$ , history from previous time steps  $(X_t^{\text{his}}, V_t^{\text{his}})$ , and the simulator’s immediate predictions  $(X_t^{\text{sim}}, V_t^{\text{sim}})$ . The encoder processes these inputs through multiple layers of patch-based self-attention, where each particle attends to others within its local neighborhood. We demonstrate experimentally that the attention mechanism computes position-dependent features for each particle based on its local geometric configuration.

**Decoder.** To decode per-particle residuals, we utilize a Neural Radiance Field (NeRF) style architecture. We apply Fourier positional encoding to the particle positions and concatenate the resulting embeddings with the features extracted by the encoder. These concatenated features are passed through an MLP decoder to produce an initial velocity correction estimate  $\Delta V_t^{\text{initial}}$ . Empirically, we found that including the decoder improves performance compared with directly using PTv3 outputs as velocity estimates.

**Temporal Aggregator.** To incorporate dynamics from the previous predictions, we refine the base predictions using a sliding-window transformer. At each step  $t$ , we update a temporal buffer with the current initial correction  $\Delta V_t^{\text{initial}}$ . If the buffer contains fewer than  $W$  frames, we pad it by replicating the most recent frame. The sequence is projected to a latent embedding, augmented with sinusoidal positional encodings, and processed by the transformer. We extract the output feature vector  $\mathbf{h}_t$  corresponding to the final time step and apply two parallel learned projections:

$$\delta_t = \tanh(\text{Proj}_\delta(\mathbf{h}_t)), \quad g_t = \sigma(\text{Proj}_g(\mathbf{h}_t)),$$

where  $\text{Proj}$  denotes a linear layer,  $\tanh$  produces a bounded temporal offset  $\delta_t$ , and the sigmoid function  $\sigma$  generates gating weights  $g_t$ . The final residual velocity is computed as:

$$\Delta V_t^{\text{residual}} = 0.1 \cdot [(1 - g_t) \odot \Delta V_t^{\text{initial}} + g_t \odot \delta_t],$$

where  $\odot$  denotes element-wise multiplication. This gating mechanism allows the network to adaptively blend the local base prediction with the temporally refined correction based on the transformer’s context.

### D. Action-Conditioned Video Prediction

Given an initial observation and a sequence of actions, we generate realistic future frames showing how the object will deform. We leverage 3D Gaussian Splatting as the rendering module. At the initial timestep, we fit Gaussian primitives to the observed point cloud, where each Gaussian is parameterized by its 3D position, covariance, opacity, and color. As PReD predicts future point cloud configurations given the action sequence, we update only the Gaussian positions and covariances, keeping their appearance attributes fixed. This allows the rendering to maintain visual consistency with the initial observation while reflecting the predicted deformations. To render a frame at timestep  $t$ , we project the updated Gaussians onto the camera view and alpha-blend them in depth order to produce the output image.

### E. Planning with PReD

PReD can serve as a forward model for model-based control. Given a state  $s_t$  and action  $a_t$ , the model predicts the next state  $s_{t+1}$ . Actions specify gripper positions that interact with the object. We integrate PReD with Model Predictive Path Integral (MPPI) control [17]. At each timestep, MPPI samples a batch of candidate action sequences  $\{\tau_i\}$ , where  $\tau_i = \{a_t^i, a_{t+1}^i, \dots, a_{t+H}^i\}$  spans a planning horizon  $H$ . For each sequence, PReD rolls out the dynamics to predict future

states. MPPI then scores trajectories using a task-specific cost function and updates the action distribution toward lower-cost regions. Following MPC, we execute only the first action from the final action and replan at the next time step with the newly observed states.

#### IV. EXPERIMENTS

##### A. Experimental Setup

We experiment with a UFACTORY xArm 7. We evaluate PReD on six objects shown in Fig. 4:



Fig. 4: **Experiment Objects.** We evaluate PReD on a diverse objects.

- a) **Rope.** A mostly one-dimensional object grasped at one end and manipulated by dragging across the table or lifting and lowering motions.
- b) **Paper.** A thin two-dimensional sheet grasped at the top and manipulated by suspending it in the air while performing waving motions, exhibiting bending and twisting.
- c) **Plush Toy.** A three-dimensional volumetric object with long, highly flexible limbs that are challenging to model.
- d) **Duster.** A heterogeneous object with a rigid stem and highly deformable feathers. This is especially challenging because it requires different dynamics predictions for the rigid stem compared to the flexible feathers.
- e) **Flag.** A cloth flag, demonstrating large area deformation dynamics that are challenging to model due to the waving motion. We chose it because its top stick prevents the cloth's surface from collapsing. We model the stick as a rigid constraint, assuming that all points along the stick follow the gripper trajectory.
- f) **Teddy Toy.** A volumetric object with higher rigidity than the plush toy. We deform it by poking rather than holding it, representing non-prehensile manipulation.

##### B. Data Collection

We follow the data collection procedure from [81]. Four Intel RealSense D455 cameras capture synchronized multi-view observations while manipulating the object. We apply Grounded SAM 2 for object segmentation and CoTracker [31] for 2D trajectory prediction. These 2D trajectories are lifted to 3D using depth information from all camera views, and persistent point tracks are extracted through an iterative rollout approach. Additional details are provided in the appendix.

##### C. Metrics

We evaluate our method using two categories of metrics: tracking metrics that measure dynamics prediction accuracy in 3D space, and visual metrics that assess rendering quality

when combined with 3D Gaussian Splatting for action conditioned video prediction.

**Tracking Metrics.** These measure dynamics prediction accuracy in 3D space. We employ three metrics computed between predicted and ground truth point clouds:

- **Mean Distance Error (MDE)** measures the average distance (in cm) between corresponding points in predicted and ground truth configurations, providing a direct assessment of positional accuracy.
- **Chamfer Distance (CD)** computes the bidirectional nearest neighbor distance between point clouds, capturing coverage and precision without requiring point correspondences.
- **Earth Mover's Distance (EMD)** quantifies the minimum cost of transforming one point cloud into another, providing a holistic measure of distributional similarity.

**Visual Metrics.** To evaluate rendering quality for action-conditioned video prediction Sec. III-D, we assess the rendered images obtained by updating Gaussian positions with our predicted point clouds. We report three metrics:

- **J-Score (IoU)** measures intersection over union of predicted and ground truth masks, quantifying spatial overlap.
- **F Score** evaluates contour accuracy, determining how well the boundary of the predicted object mask matches the boundary of the ground truth.
- **Learned Perceptual Image Patch Similarity (LPIPS)** assesses perceptual similarity using features extracted from predicted and ground truth images, capturing visual differences closer to human perception.

##### D. Baselines

We compare PReD against four baselines: two analytical physics-based simulations and two learning-based approaches.

**Spring-Mass Model [30]:** This is the same as the physics backbone we use for our model. Improvements over this baseline demonstrate that the learned residual model effectively captures the dynamics missed by the physics simulator alone.

**Material Point Method (MPM) [62] :** MPM combines Eulerian and Lagrangian representations to handle material behaviors and topological changes. This hybrid approach discretizes materials into particles while using a background grid for computing spatial derivatives and enforcing conservation laws. MPM is particularly effective for materials exhibiting both solid and fluid-like behaviors but requires careful parameter tuning and can be computationally expensive.

**Graph Based Neural Dynamics (GBND) [82] :** This is a learned approach that employs Graph Neural Networks to learn object dynamics directly from data without explicit material parameters. The object is represented as a graph where message passing mechanisms propagate information between nodes to capture long range interactions.

**Particle-Grid Neural Dynamics (PGND) [81] :** This represents the state-of-the-art in learning-based methods, employing a hybrid representation of particles and spatial grids inspired by MPM to model deformable object dynamics. In this, particles capture the object geometry while the spatial grid

Method	Metric	Rope	Paper	Plush Toy	Duster	Flag	Teddy Toy
Spring-Mass [30]	MDE $\downarrow$	4.4 $\pm$ 2.0	2.3 $\pm$ 2.0	3.2 $\pm$ 0.8	3.8 $\pm$ 2.0	5.7 $\pm$ 2.5	5.8 $\pm$ 3.4
MPM [62]		7.3 $\pm$ 2.5	15.5 $\pm$ 4.2	7.4 $\pm$ 1.8	6.0 $\pm$ 2.6	23.1 $\pm$ 7.0	—
GBND [82]		5.5 $\pm$ 1.7	3.0 $\pm$ 1.4	7.7 $\pm$ 2.6	5.1 $\pm$ 2.2	30.9 $\pm$ 6.2	1.5 $\pm$ 0.4
PGND [81]		3.3 $\pm$ 1.8	2.1 $\pm$ 0.5	4.0 $\pm$ 1.3	3.8 $\pm$ 0.1	3.2 $\pm$ 1.4	1.6 $\pm$ 1.3
Ours		<b>2.6<math>\pm</math>1.31</b>	<b>1.7<math>\pm</math>0.7</b>	<b>2.7<math>\pm</math>0.4</b>	<b>2.9<math>\pm</math>1.8</b>	<b>2.8<math>\pm</math>1.2</b>	<b>1.3<math>\pm</math>0.3</b>
Spring-Mass [30]	CD $\downarrow$	0.045 $\pm$ 0.024	0.017 $\pm$ 0.012	0.029 $\pm$ 0.007	0.047 $\pm$ 0.025	0.091 $\pm$ 0.047	0.051 $\pm$ 0.030
MPM [62]		0.061 $\pm$ 0.024	0.148 $\pm$ 0.065	0.071 $\pm$ 0.017	0.052 $\pm$ 0.017	0.236 $\pm$ 0.104	—
GBND [82]		0.066 $\pm$ 0.024	0.050 $\pm$ 0.017	0.065 $\pm$ 0.018	0.063 $\pm$ 0.028	0.491 $\pm$ 0.130	0.027 $\pm$ 0.003
PGND [81]		0.027 $\pm$ 0.014	0.022 $\pm$ 0.009	0.032 $\pm$ 0.008	0.040 $\pm$ 0.001	0.043 $\pm$ 0.024	0.016 $\pm$ 0.010
Ours		<b>0.023<math>\pm</math>0.013</b>	<b>0.013<math>\pm</math>0.006</b>	<b>0.026<math>\pm</math>0.003</b>	<b>0.038<math>\pm</math>0.025</b>	<b>0.043<math>\pm</math>0.022</b>	<b>0.013<math>\pm</math>0.002</b>
Spring-Mass [30]	EMD $\downarrow$	0.024 $\pm$ 0.013	0.018 $\pm$ 0.012	0.015 $\pm$ 0.004	0.019 $\pm$ 0.012	0.049 $\pm$ 0.026	0.029 $\pm$ 0.018
MPM [62]		0.039 $\pm$ 0.016	0.114 $\pm$ 0.055	0.039 $\pm$ 0.009	0.028 $\pm$ 0.012	0.194 $\pm$ 0.078	—
GBND [82]		0.030 $\pm$ 0.013	0.021 $\pm$ 0.009	0.031 $\pm$ 0.011	0.028 $\pm$ 0.014	0.245 $\pm$ 0.063	0.008 $\pm$ 0.002
PGND [81]		0.014 $\pm$ 0.006	0.015 $\pm$ 0.005	0.016 $\pm$ 0.004	0.016 $\pm$ 0.004	0.023 $\pm$ 0.013	0.008 $\pm$ 0.006
Ours		<b>0.012<math>\pm</math>0.006</b>	<b>0.014<math>\pm</math>0.007</b>	<b>0.014<math>\pm</math>0.002</b>	<b>0.014<math>\pm</math>0.011</b>	<b>0.022<math>\pm</math>0.012</b>	<b>0.006<math>\pm</math>0.001</b>

TABLE I. **Tracking accuracy across diverse objects.** Physics-Guided Residual Dynamics achieves the lowest error across all objects and metrics, outperforming both physics-based methods and learning-based approaches.

Method	Metric	Rope	Paper	Plush Toy	Duster	Flag	Teddy Toy
Spring-Mass [30]	$\mathcal{J}$ -Score / IoU $\uparrow$	0.317 $\pm$ 0.151	0.682 $\pm$ 0.147	0.530 $\pm$ 0.087	0.666 $\pm$ 0.082	0.460 $\pm$ 0.195	0.543 $\pm$ 0.199
MPM [62]		0.244 $\pm$ 0.145	0.451 $\pm$ 0.094	0.522 $\pm$ 0.088	0.592 $\pm$ 0.087	0.378 $\pm$ 0.112	—
GBND [82]		0.053 $\pm$ 0.055	0.557 $\pm$ 0.140	0.370 $\pm$ 0.149	0.513 $\pm$ 0.104	0.508 $\pm$ 0.142	0.752 $\pm$ 0.066
PGND [81]		0.318 $\pm$ 0.213	0.708 $\pm$ 0.090	0.586 $\pm$ 0.093	0.669 $\pm$ 0.046	0.547 $\pm$ 0.166	0.713 $\pm$ 0.102
Ours		<b>0.395<math>\pm</math>0.169</b>	<b>0.753<math>\pm</math>0.082</b>	<b>0.621<math>\pm</math>0.078</b>	<b>0.677<math>\pm</math>0.083</b>	<b>0.554<math>\pm</math>0.189</b>	<b>0.782<math>\pm</math>0.114</b>
Spring-Mass [30]	$\mathcal{F}$ -Score $\uparrow$	0.659 $\pm$ 0.160	0.471 $\pm$ 0.261	0.570 $\pm$ 0.112	0.532 $\pm$ 0.097	0.208 $\pm$ 0.181	0.473 $\pm$ 0.272
MPM [62]		0.492 $\pm$ 0.245	0.473 $\pm$ 0.082	0.513 $\pm$ 0.086	0.500 $\pm$ 0.108	0.249 $\pm$ 0.097	—
GBND [82]		0.195 $\pm$ 0.191	0.495 $\pm$ 0.135	0.359 $\pm$ 0.115	0.411 $\pm$ 0.109	0.254 $\pm$ 0.204	0.771 $\pm$ 0.096
PGND [81]		0.587 $\pm$ 0.296	0.539 $\pm$ 0.165	0.553 $\pm$ 0.111	0.529 $\pm$ 0.097	0.265 $\pm$ 0.198	0.727 $\pm$ 0.154
Ours		<b>0.738<math>\pm</math>0.185</b>	<b>0.591<math>\pm</math>0.176</b>	<b>0.630<math>\pm</math>0.107</b>	<b>0.551<math>\pm</math>0.102</b>	<b>0.268<math>\pm</math>0.214</b>	<b>0.803<math>\pm</math>0.175</b>
Spring-Mass [30]	LPIPS $\downarrow$	0.026 $\pm$ 0.009	0.056 $\pm$ 0.022	0.059 $\pm$ 0.020	0.070 $\pm$ 0.019	0.099 $\pm$ 0.030	0.036 $\pm$ 0.017
MPM [62]		0.038 $\pm$ 0.014	0.065 $\pm$ 0.014	0.077 $\pm$ 0.017	0.068 $\pm$ 0.017	0.091 $\pm$ 0.025	—
GBND [82]		0.052 $\pm$ 0.017	0.057 $\pm$ 0.013	0.105 $\pm$ 0.020	0.064 $\pm$ 0.013	0.088 $\pm$ 0.013	0.022 $\pm$ 0.005
PGND [81]		0.026 $\pm$ 0.013	0.052 $\pm$ 0.011	0.065 $\pm$ 0.018	0.060 $\pm$ 0.009	0.089 $\pm$ 0.019	0.020 $\pm$ 0.008
Ours		<b>0.023<math>\pm</math>0.011</b>	<b>0.049<math>\pm</math>0.014</b>	<b>0.054<math>\pm</math>0.021</b>	<b>0.059<math>\pm</math>0.019</b>	<b>0.085<math>\pm</math>0.029</b>	<b>0.017<math>\pm</math>0.009</b>

TABLE II. **Video prediction with 3DGS.** PReD produces the most accurate visual predictions, demonstrating superior rendering quality.

discretizes the 3D domain to ensure spatial continuity and improve learning efficiency.

### E. Results

We evaluate PReD across two primary dimensions: dynamics tracking accuracy and action-conditioned video prediction.

**Dynamics and Tracking Accuracy.** We first analyze our model’s ability to track and simulate object states over time, with results summarized in Tab. I. PReD consistently achieves the lowest error across all objects and metrics, demonstrating that learned residual corrections on physics backbones improve accuracy beyond either pure physics or pure learning approaches. Among physics methods, the optimized spring-mass model generally outperforms MPM, particularly for objects with coherent structure like the rope and plush toy. Among learning methods, PGND substantially outperforms GBND across all objects, indicating that particle representations with physics priors are more effective than graph neural approaches for deformable object simulation.

Qualitatively, Fig. 5 reveals consistent patterns in how different approaches handle deformation. The optimized spring-mass model captures overall motion but lacks the flexibility to represent local variations in material properties, often ap-

pearing too stiff. MPM often struggles to maintain structural cohesion, with object particles separating rather than behaving as connected solids. Learning methods show complementary strengths: PGND generates smooth predictions but accumulates drift over time, while GBND struggles with graph topology, producing unphysical artifacts. PReD leverages its physics backbone for stability while using learned residuals to correct local dynamics, maintaining both plausibility and accuracy throughout the rollout.

The qualitative results reveal several interesting observations that provide further insights into the limitations of different approaches. MPM fails catastrophically for thin objects like paper and flag, where object tear down and lose structure. This happens because it relies on grid-based discretization, which can cause material points to lose cohesion when particle distribution is non-uniform. The duster object, with its rigid handle and soft feathers, reveals distinct failure modes across all baselines: the optimized spring-mass model cannot assign spatially varying stiffness and causes the rigid stem to bend unnaturally, MPM fails to maintain volume with material points collapsing toward the floor, GBND predicts erroneous upward motion for all feather particles, and PGND incorrectly compresses the rigid stem. Only PReD simulates



Fig. 5: **Qualitative comparison on diverse objects.** The first column shows the states with gripper positions (red spheres). The second column shows the GT object visualization after taking the action. PReD accurately captures deformations across all objects. For the duster, only PReD maintains stem rigidity while allowing deformation for the feathers.

it correctly. For the teddy, under non-prehensile manipulation, the object slips out of the optimized spring-mass model. MPM cannot realistically simulate he constraints for non-prehensile manipulation, thus we exclude it from this comparison. GBND predicts minimal particle displacement, producing an overly stiff response. PGND shows the gripper’s entire surface in our visualization, but the gripper fingers themselves are not visible because they have pushed into its interior. Only PReD successfully captures the deformation, properly maintaining contact while allowing realistic compression.

**Action Conditioned Video Prediction.** PReD translates superior dynamics accuracy to improved photorealistic rendering. Across all objects, PReD consistently achieves the best J-Score and F-Score metrics, demonstrating better spatial overlap and foreground detection compared to both physics based and learning based baselines. PReD also maintains the lowest LPIPS scores across objects, indicating that the physics backbone enables the residual model to preserve objects visually during rendering.

**Interactive Photo-Realistic Simulation.** PReD enables photo-realistic interactive simulation by integrating the predicted dynamics with 3D Gaussian Splatting for rendering. Users provide manipulation commands via keyboard input, which drives the physics backbone, applies learned residual corrections at each step, and renders the updated Gaussian representations. This combination allows users to manipulate objects with visual realism from the rendering. The results are shown in the appendix.

#### F. Planning using PReD

We use PReD for planning using MPPI. As the objective, we minimize the Chamfer Distance (CD) between the predicted configuration and a pre-collected target configuration, both represented as point clouds. Other implementation details are provided in the appendix.

**Standard Manipulation Tasks.** We evaluate planning on the rope with four start and goal configurations shown in Fig. 6 (top): (1) lifting the rope up, (2) lowering it, (3) rotating it such

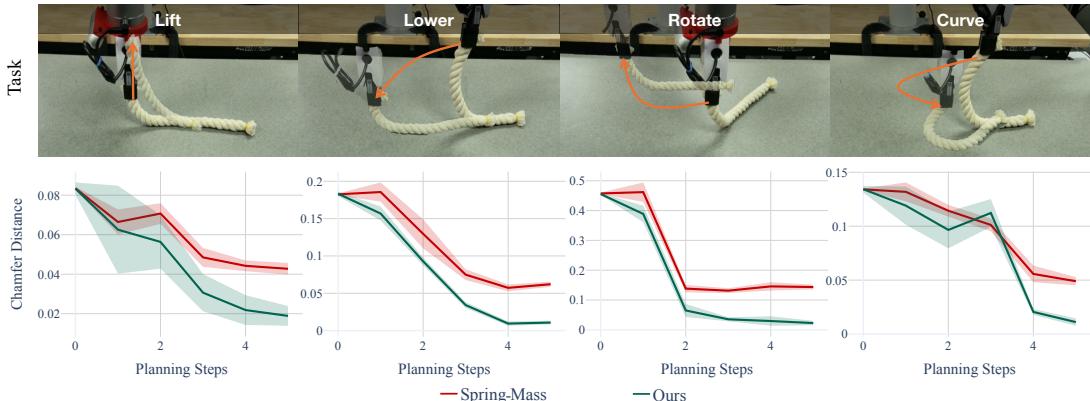


Fig. 6: **Planning results on rope manipulation tasks.** Top row shows start and end configurations for four tasks. The bottom row shows the Chamfer Distance to the target during MPPI planning. PReD (green) converges faster and achieves lower final error than the optimized spring-mass baseline (red) across all tasks.

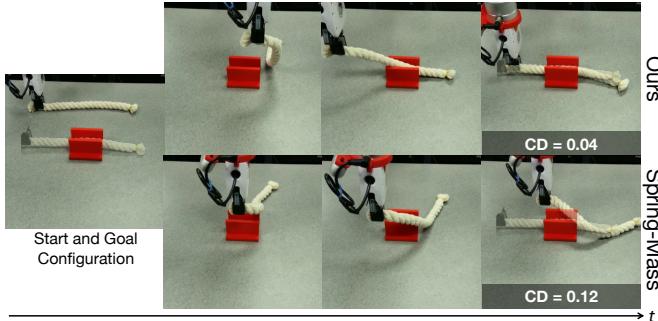


Fig. 7: **Cable Rerouting Execution.** Comparison of PReD (top) and the Spring-Mass backbone (bottom) for rope rerouting through a slot. PReD navigates the rope through the opening while the Spring-Mass backbone exhibits repeated collisions with the slot edges.

that one end remains roughly stationary while the other moves, and (4) deforming it into a curved shape. We compare PReD against the spring-mass backbone, providing both methods with the same MPPI planning budget. We run 10 trials for each task using the same start and end configurations. As shown in Fig. 6, PReD’s CD decreases smoothly and saturates at a value close to the target, whereas the spring-mass model converges more slowly and sometimes fails to reach a comparable final error. Additional planning results on the plush toy will be provided in the appendix.

**Cable Rerouting.** We further test PReD on a more challenging task: cable rerouting through a narrow slot. This requires the robot to move the rope and thread it through the slot. We run 10 trials with different start and end configurations. The target configuration is provided to the planner.

As demonstrated in Fig. 7, PReD successfully navigates the rope through the slot. The learned residuals compensate for inaccuracies in the physics backbone, enabling precise trajectory execution that avoids collisions with the edges. In contrast, the spring-mass backbone struggles with these fine dynamics. Small errors cause the rope to collide with the edges, forcing the planner to retry the motion multiple times before completing the task.

Quantitatively, Fig. 8 shows the distributions of start and end configurations overlaid (left) and comparisons of chamfer distances (right). PReD achieves significantly lower CD than

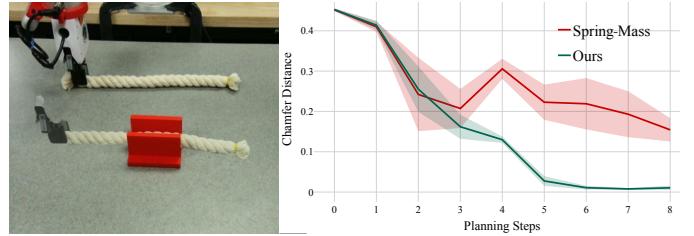


Fig. 8: **Cable rerouting through narrow slot.** Left: distribution of start and end configurations. Right: Chamfer Distance during MPPI planning steps. PReD achieves lower distance than the spring-mass baseline, successfully threading the rope through the slot in 8 out of 10 trials compared to 2 out of 10 for the baseline.

the optimized spring-mass. We also evaluate success rates for this task. A trial is considered successful if the rope passes entirely through the slot and the gripper remains below the slot height in the final configuration. Under this criterion, PReD succeeds in 8/10 trials, whereas the spring-mass model succeeds in only 2/10 trials. These results demonstrate that the improved model enables reliable execution of tasks that would otherwise be infeasible.

## V. CONCLUSIONS

We presented PReD, a hybrid simulation framework that bridges physics-based and learning-based approaches for deformable object simulation. By combining an optimized spring-mass backbone with learned velocity residuals, PReD achieved superior accuracy while maintaining physical plausibility. Our velocity-based residuals, combined with a sliding-window transformer for temporal aggregation, yielded stable predictions, avoiding the instability issues plaguing naive residual learning. Extensive experiments across real-world objects demonstrated that PReD consistently outperformed both physics-based methods and SOTA learning-based approaches in tracking. PReD handled challenging scenarios involving heterogeneous materials, volumetric deformations, and contact-rich interactions where baselines failed. Beyond tracking, we demonstrated applications in action-conditioned video prediction and MPC. Rerouting experiments highlighted that improved dynamics translated to higher success, with physics-based simulation alone proving insufficient.

## REFERENCES

[1] Bo Ai, Stephen Tian, Haochen Shi, Yixuan Wang, Tobias Pfaff, Cheston Tan, Henrik I Christensen, Hao Su, Jiajun Wu, and Yunzhu Li. A review of learning-based dynamics models for robotic manipulation. *Science Robotics*, 10(106):eadt1497, 2025.

[2] Anurag Ajay, Jiajun Wu, Nima Fazeli, Maria Bauza, Leslie P Kaelbling, Joshua B Tenenbaum, and Alberto Rodriguez. Augmenting physical simulators with stochastic neural networks: Case study of planar pushing and bouncing. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3066–3073. IEEE, 2018.

[3] Hugo Attali, Davide Buscaldi, and Nathalie Pernelle. Rewiring techniques to mitigate oversquashing and oversmoothing in gnn: A survey. *arXiv preprint arXiv:2411.17429*, 2024.

[4] David Baraff and Andrew Witkin. Large steps in cloth simulation. In *Proceedings of 25th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '98)*, pages 43 – 54. SIGGRAPH, July 1998.

[5] Peter Battaglia, Razvan Pascanu, Matthew Lai, Danilo Jimenez Rezende, et al. Interaction networks for learning about objects, relations and physics. *Advances in neural information processing systems*, 29, 2016.

[6] Dominik Bauer, Zhenjia Xu, and Shuran Song. Doughnet: A visual predictive model for topological manipulation of deformable objects. In *European Conference on Computer Vision*, pages 92–108. Springer, 2024.

[7] Leonard Bauersfeld, Elia Kaufmann, Philipp Foehn, Si-hao Sun, and Davide Scaramuzza. Neurobem: Hybrid aerodynamic quadrotor model. *arXiv preprint arXiv:2106.08015*, 2021.

[8] Onur Beker, Nico Gürler, Ji Shi, A René Geist, Amirreza Razmjoo, Georg Martius, and Sylvain Calinon. A smooth analytical formulation of collision detection and rigid body dynamics with contact. *arXiv preprint arXiv:2503.11736*, 2025.

[9] Filipe De Avila Belbute-Peres, Thomas Economou, and Zico Kolter. Combining differentiable pde solvers and graph neural networks for fluid flow prediction. In *international conference on machine learning*, pages 2402–2411. PMLR, 2020.

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

[11] Kurtland Chua, Roberto Calandra, Rowan McAllister, and Sergey Levine. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. *Advances in neural information processing systems*, 31, 2018.

[12] Loris Di Natale, Bratislav Svetozarevic, Philipp Heer, and Colin N Jones. Physically consistent neural networks for building thermal modeling: theory and analysis. *Applied Energy*, 325:119806, 2022.

[13] Tao Du, Kui Wu, Pingchuan Ma, Sebastien Wah, Andrew Spielberg, Daniela Rus, and Wojciech Matusik. Diffpd: Differentiable projective dynamics. *ACM Transactions on Graphics (ToG)*, 41(2):1–21, 2021.

[14] Bardienus P. Duisterhof, Jan Oberst, Bowen Wen, Stan Birchfield, Deva Ramanan, and Jeffrey Ichniowski. Rayst3r: Predicting novel depth maps for zero-shot object completion, 2025. URL <https://arxiv.org/abs/2506.05285>.

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

[16] Junpeng Gao, Mike Y Michelis, Andrew Spielberg, and Robert K Katzschmann. Sim-to-real of soft robots with learned residual physics. *IEEE Robotics and Automation Letters*, 2024.

[17] Carlos E Garcia, David M Prett, and Manfred Morari. Model predictive control: Theory and practice—a survey. *Automatica*, 25(3):335–348, 1989.

[18] Moritz Geilinger, David Hahn, Jonas Zehnder, Moritz Bächer, Bernhard Thomaszewski, and Stelian Coros. Add: Analytically differentiable dynamics for multi-body systems with frictional contact. *ACM Transactions on Graphics (TOG)*, 39(6):1–15, 2020.

[19] Florian Golemo, Adrien Ali Taiga, Aaron Courville, and Pierre-Yves Oudeyer. Sim-to-real transfer with neural-augmented robot simulation. In *Conference on Robot Learning*, pages 817–828. PMLR, 2018.

[20] Joshua Gruenstein, Tao Chen, Neel Doshi, and Pulkit Agrawal. Residual model learning for microrobot control. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 7219–7226. IEEE, 2021.

[21] Mathew Halm and Michael Posa. Set-valued rigid body dynamics for simultaneous frictional impact. *arXiv preprint arXiv:2103.15714*, 2021.

[22] Nikolaus Hansen and Andreas Ostermeier. Completely derandomized self-adaptation in evolution strategies. *Evolutionary computation*, 9(2):159–195, 2001.

[23] Eric Heiden, Miles Macklin, Yashraj Narang, Dieter Fox, Animesh Garg, and Fabio Ramos. Disect: A differentiable simulation engine for autonomous robotic cutting. *arXiv preprint arXiv:2105.12244*, 2021.

[24] Philipp Holl, Vladlen Koltun, and Nils Thuerey. Learning to control pdes with differentiable physics. *arXiv preprint arXiv:2001.07457*, 2020.

[25] Suning Huang, Qianzhong Chen, Xiaohan Zhang, Jiankai Sun, and Mac Schwager. Particleformer: A 3d point cloud world model for multi-object, multi-material robotic manipulation. *arXiv preprint arXiv:2506.23126*, 2025.

[26] Wenlong Huang, Yu-Wei Chao, Arsalan Mousavian, Ming-Yu Liu, Dieter Fox, Kaichun Mo, and Li Fei-Fei. Pointworld: Scaling 3d world models for in-the-wild robotic manipulation. *arXiv preprint arXiv:2601.03782*, 2026.

2026.

- [27] Zhaio Huang, Yuanming Hu, Tao Du, Siyuan Zhou, Hao Su, Joshua B Tenenbaum, and Chuang Gan. Plasticinelab: A soft-body manipulation benchmark with differentiable physics. *arXiv preprint arXiv:2104.03311*, 2021.
- [28] Krishna Murthy Jatavallabhula, Miles Macklin, Florian Golemo, Vikram Voleti, Linda Petrini, Martin Weiss, Breandan Considine, Jérôme Parent-Lévesque, Kevin Xie, Kenny Erleben, et al. gradsim: Differentiable simulation for system identification and visuomotor control. *arXiv preprint arXiv:2104.02646*, 2021.
- [29] Chenfanfu Jiang, Craig Schroeder, Andrew Selle, Joseph Teran, and Alexey Stomakhin. The affine particle-in-cell method. *ACM Transactions on Graphics (TOG)*, 34(4):1–10, 2015.
- [30] Hanxiao Jiang, Hao-Yu Hsu, Kaifeng Zhang, Hsin-Ni Yu, Shenlong Wang, and Yunzhu Li. Phystwin: Physics-informed reconstruction and simulation of deformable objects from videos. *arXiv preprint arXiv:2503.17973*, 2025.
- [31] Nikita Karaev, Yuri Makarov, Jianyuan Wang, Natalia Neverova, Andrea Vedaldi, and Christian Rupprecht. Cotracker3: Simpler and better point tracking by pseudo-labelling real videos. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6013–6022, 2025.
- [32] Hajun Kim, Dongyun Kang, Min-Gyu Kim, Gijeong Kim, and Hae-Won Park. Online friction coefficient identification for legged robots on slippery terrain using smoothed contact gradients. *IEEE Robotics and Automation Letters*, 2025.
- [33] Thomas Kipf, Ethan Fetaya, Kuan-Chieh Wang, Max Welling, and Richard Zemel. Neural relational inference for interacting systems. In *International conference on machine learning*, pages 2688–2697. Pmlr, 2018.
- [34] Thomas Kipf, Elise Van der Pol, and Max Welling. Contrastive learning of structured world models. *arXiv preprint arXiv:1911.12247*, 2019.
- [35] Yunzhu Li, Jiajun Wu, Russ Tedrake, Joshua B Tenenbaum, and Antonio Torralba. Learning particle dynamics for manipulating rigid bodies, deformable objects, and fluids. *arXiv preprint arXiv:1810.01566*, 2018.
- [36] Yunzhu Li, Jiajun Wu, Jun-Yan Zhu, Joshua B Tenenbaum, Antonio Torralba, and Russ Tedrake. Propagation networks for model-based control under partial observation. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 1205–1211. IEEE, 2019.
- [37] Junbang Liang, Ming Lin, and Vladlen Koltun. Differentiable cloth simulation for inverse problems. *Advances in neural information processing systems*, 32, 2019.
- [38] Quentin Le Lidec, Louis Montaut, Yann de Mont-Marin, Fabian Schramm, and Justin Carpentier. End-to-end and highly-efficient differentiable simulation for robotics. *arXiv preprint arXiv:2409.07107*, 2024.
- [39] Xingyu Lin, Yufei Wang, Zixuan Huang, and David Held. Learning visible connectivity dynamics for cloth smoothing. In *Conference on Robot Learning*, pages 256–266. PMLR, 2022.
- [40] Min Liu, Gang Yang, Siyuan Luo, and Lin Shao. Softmac: Differentiable soft body simulation with forecast-based contact model and two-way coupling with articulated rigid bodies and clothes. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 12008–12015. IEEE, 2024.
- [41] Tiantian Liu, Adam W Bargteil, James F O’Brien, and Ladislav Kavan. Fast simulation of mass-spring systems. *ACM Transactions on Graphics (TOG)*, 32(6):1–7, 2013.
- [42] Bryn Lloyd, Gábor Székely, and Matthias Harders. Identification of spring parameters for deformable object simulation. *IEEE transactions on visualization and computer graphics*, 13(5):1081–1094, 2007.
- [43] Jonathon Luiten, Georgios Kopanas, Bastian Leibe, and Deva Ramanan. Dynamic 3d gaussians: Tracking by persistent dynamic view synthesis. In *2024 International Conference on 3D Vision (3DV)*, pages 800–809. IEEE, 2024.
- [44] Pingchuan Ma, Tao Du, Joshua B Tenenbaum, Wojciech Matusik, and Chuang Gan. Risp: Rendering-invariant state predictor with differentiable simulation and rendering for cross-domain parameter estimation. *arXiv preprint arXiv:2205.05678*, 2022.
- [45] Pingchuan Ma, Peter Yichen Chen, Bolei Deng, Joshua B Tenenbaum, Tao Du, Chuang Gan, and Wojciech Matusik. Learning neural constitutive laws from motion observations for generalizable pde dynamics. In *International Conference on Machine Learning*, pages 23279–23300. PMLR, 2023.
- [46] Xiao Ma, David Hsu, and Wee Sun Lee. Learning latent graph dynamics for visual manipulation of deformable objects. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 8266–8273. IEEE, 2022.
- [47] Miles Macklin, Matthias Müller, Nuttapong Chentanez, and Tae-Yong Kim. Unified particle physics for real-time applications. *ACM Transactions on Graphics (TOG)*, 33(4):1–12, 2014.
- [48] Dominique Madier. An introduction to the fundamentals of mesh generation in finite element analysis. Technical report, FEA Academy, Montreal, Canada, December 2023. URL <https://wwwfea-academy.com/pdf/FEA%20Academy%20-%20The%20Fundamentals%20of%20Mesh%20Generation%20in%20Finite%20Element%20Analysis.pdf>. Published by FEA Academy.
- [49] L Angela Mihai and Alain Goriely. How to characterize a nonlinear elastic material? a review on nonlinear constitutive parameters in isotropic finite elasticity. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 473(2207):20170607, 2017.
- [50] Damian Mrowca, Chengxu Zhuang, Elias Wang, Nick Haber, Li Fei-Fei, Josh Tenenbaum, and Daniel L Yamins. Flexible neural representation for physics prediction. *Advances in neural information processing*

systems, 31, 2018.

- [51] Matthias Müller and Markus H Gross. Interactive virtual materials. In *Graphics interface*, volume 2004, pages 239–246, 2004.
- [52] Matthias Müller, Bruno Heidelberger, Marcus Hennix, and John Ratcliff. Position based dynamics. *Journal of Visual Communication and Image Representation*, 18(2):109–118, 2007.
- [53] J Krishna Murthy, Miles Macklin, Florian Golemo, Vikram Voleti, Linda Petrini, Martin Weiss, Breandan Considine, Jérôme Parent-Lévesque, Kevin Xie, Kenny Erleben, et al. gradsim: Differentiable simulation for system identification and visuomotor control. In *International conference on learning representations*, 2020.
- [54] Tobias Pfaff, Meire Fortunato, Alvaro Sanchez-Gonzalez, and Peter Battaglia. Learning mesh-based simulation with graph networks. In *International conference on learning representations*, 2020.
- [55] Yiling Qiao, Junbang Liang, Vladlen Koltun, and Ming Lin. Differentiable simulation of soft multi-body systems. *Advances in Neural Information Processing Systems*, 34:17123–17135, 2021.
- [56] Junior Rojas, Eftychios Sifakis, and Ladislav Kavan. Differentiable implicit soft-body physics. *arXiv preprint arXiv:2102.05791*, 2021.
- [57] Alvaro Sanchez-Gonzalez, Jonathan Godwin, Tobias Pfaff, Rex Ying, Jure Leskovec, and Peter Battaglia. Learning to simulate complex physics with graph networks. In *International conference on machine learning*, pages 8459–8468. PMLR, 2020.
- [58] Yidi Shao, Chen Change Loy, and Bo Dai. Transformer with implicit edges for particle-based physics simulation. In *European conference on computer vision*, pages 549–564. Springer, 2022.
- [59] 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.
- [60] Haochen Shi, Huazhe Xu, Zhiao Huang, Yunzhu Li, and Jiajun Wu. 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, 2024.
- [61] Eftychios Sifakis and Jernej Barbic. Fem simulation of 3d deformable solids: a practitioner’s guide to theory, discretization and model reduction. In *Acm siggraph 2012 courses*, pages 1–50. 2012.
- [62] Alexey Stomakhin, Craig Schroeder, Lawrence Chai, Joseph Teran, and Andrew Selle. A material point method for snow simulation. *ACM Transactions on Graphics (TOG)*, 32(4):1–10, 2013.
- [63] Shijie Tan, Hongjun Zhou, and Jinjin Zheng. A hybrid model method for accurate surface deformation and incision based on fem and pbd. *Scientific Programming*, 2021(1):8343312, 2021.
- [64] Kiwon Um, Robert Brand, Yun Raymond Fei, Philipp Holl, and Nils Thuerey. Solver-in-the-loop: Learning from differentiable physics to interact with iterative pde-solvers. *Advances in neural information processing systems*, 33:6111–6122, 2020.
- [65] Vaiva Vasiliauskaitė and Nino Antulov-Fantulin. Generalization of neural network models for complex network dynamics. *Communications Physics*, 7(1):348, 2024.
- [66] Changhao Wang, Yuyou Zhang, Xiang Zhang, Zheng Wu, Xinghao Zhu, Shiyu Jin, Te Tang, and Masayoshi Tomizuka. Offline-online learning of deformation model for cable manipulation with graph neural networks. *IEEE Robotics and Automation Letters*, 7(2):5544–5551, 2022.
- [67] Xinjue Wang, Esa Ollila, and Sergiy A Vorobyov. Graph neural network sensitivity under probabilistic error model. In *2022 30th European Signal Processing Conference (EUSIPCO)*, pages 2146–2150. IEEE, 2022.
- [68] Zizhao Wang, Xuesu Xiao, Zifan Xu, Yuke Zhu, and Peter Stone. Causal dynamics learning for task-independent state abstraction. *arXiv preprint arXiv:2206.13452*, 2022.
- [69] Robert J Webster III and Bryan A Jones. Design and kinematic modeling of constant curvature continuum robots: A review. *The International Journal of Robotics Research*, 29(13):1661–1683, 2010.
- [70] Keenon Werling, Dalton Omens, Jeongseok Lee, Ioannis Exarchos, and C Karen Liu. Fast and feature-complete differentiable physics for articulated rigid bodies with contact. *arXiv preprint arXiv:2103.16021*, 2021.
- [71] William F Whitney, Jacob Varley, Deepali Jain, Krzysztof Choromanski, Sumeet Singh, and Vikas Sindhwani. Modeling the real world with high-density visual particle dynamics. *arXiv preprint arXiv:2406.19800*, 2024.
- [72] Xiaoyang Wu, Li Jiang, Peng-Shuai Wang, Zhijian Liu, Xihui Liu, Yu Qiao, Wanli Ouyang, Tong He, and Hengshuang Zhao. Point transformer v3: Simpler faster stronger. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4840–4851, 2024.
- [73] Yilin Wu, Wilson Yan, Thanard Kurutach, Lerrel Pinto, and Pieter Abbeel. Learning to manipulate deformable objects without demonstrations. *arXiv preprint arXiv:1910.13439*, 2019.
- [74] Zhenjia Xu, Jiajun Wu, Andy Zeng, Joshua B Tenenbaum, and Shuran Song. Densephysnet: Learning dense physical object representations via multi-step dynamic interactions. *arXiv preprint arXiv:1906.03853*, 2019.
- [75] Shangjie Xue, Shuo Cheng, Pujith Kachana, and Danfei Xu. Neural field dynamics model for granular object piles manipulation. In *Conference on Robot Learning*, pages 2821–2837. PMLR, 2023.
- [76] Mengyuan Yan, Yilin Zhu, Ning Jin, and Jeannette Bohg. Self-supervised learning of state estimation for manipulating deformable linear objects. *IEEE robotics and automation letters*, 5(2):2372–2379, 2020.
- [77] Xintong Yang, Ze Ji, and Yu-Kun Lai. Differentiable physics-based system identification for robotic manipulation of elastoplastic materials. *The International Journal*

*of Robotics Research*, page 02783649251334661, 2025.

- [78] Xiaohan Ye, Kui Wu, Zherong Pan, and Taku Komura. Efficient differentiable contact model with long-range influence. *arXiv preprint arXiv:2509.20917*, 2025.
- [79] Rose Yu and Rui Wang. Learning dynamical systems from data: An introduction to physics-guided deep learning. *Proceedings of the National Academy of Sciences*, 121(27):e2311808121, 2024.
- [80] Kaifeng Zhang, Baoyu Li, Kris Hauser, and Yunzhu Li. Adaptigraph: Material-adaptive graph-based neural dynamics for robotic manipulation. *arXiv preprint arXiv:2407.07889*, 2024.
- [81] Kaifeng Zhang, Baoyu Li, Kris Hauser, and Yunzhu Li. Particle-grid neural dynamics for learning deformable object models from rgb-d videos. *arXiv preprint arXiv:2506.15680*, 2025.
- [82] Mingtong Zhang, Kaifeng Zhang, and Yunzhu Li. Dynamic 3d gaussian tracking for graph-based neural dynamics modeling. *arXiv preprint arXiv:2410.18912*, 2024.
- [83] Licheng Zhong, Hong-Xing Yu, Jiajun Wu, and Yunzhu Li. Reconstruction and simulation of elastic objects with spring-mass 3d gaussians. In *European Conference on Computer Vision*, pages 407–423. Springer, 2024.
- [84] Yaofeng Desmond Zhong, Jiequn Han, Biswadip Dey, and Georgia Olympia Brikis. Improving gradient computation for differentiable physics simulation with contacts. In *Learning for dynamics and control conference*, pages 128–141. PMLR, 2023.
- [85] Gaoyue Zhou, Hengkai Pan, Yann LeCun, and Lerrel Pinto. Dino-wm: World models on pre-trained visual features enable zero-shot planning. *arXiv preprint arXiv:2411.04983*, 2024.
- [86] Yicheng Zhu, David Yang, and Yangming Lee. Deformable and fragile object manipulation: A review and prospects. *Sensors*, 25(17):5430, 2025.

# Appendix

We structure the supplement into the following sections:

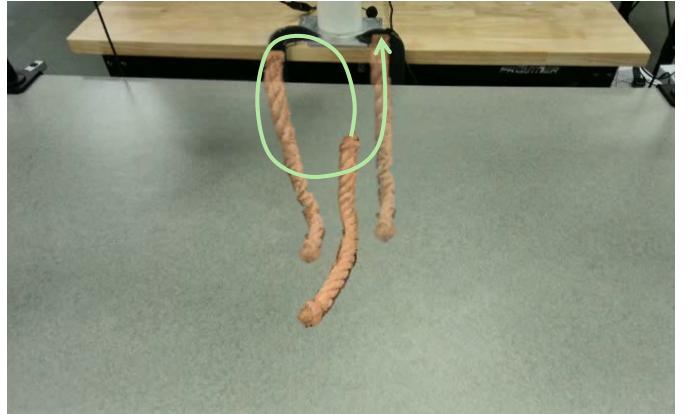
- A** Data collection and processing pipeline for obtaining consistent particle trajectories from multi-view RGBD observations.
- B** Interactive photorealistic simulation results combining PReD with 3D Gaussian Splatting for real-time rendering.
- C** Planning experiments on plush toy manipulation demonstrating PReD’s effectiveness on volumetric objects with flexible limbs.
- D** Comparison to PhysTwin on heterogeneous objects, highlighting limitations of parameter optimization versus learned residual corrections.
- E** Network architecture details and hyperparameter specifications for the encoder, decoder, and temporal aggregator.
- F** Ablation study on decoder architecture, isolating the contribution of the NeRF-style decoder to tracking accuracy.

## A. DATA COLLECTION AND PROCESSING

We mount four Intel RealSense D455 cameras around the workspace to capture synchronized RGBD observations at 30 Hz. The cameras are calibrated to a shared world coordinate frame using a checkerboard calibration procedure, enabling the fusion of observations across views. For each camera view, we apply Grounded SAM 2 to extract object masks. Using text prompts containing object descriptions, it detects and segments the object in the first frame of a sequence and propagates the mask across subsequent frames. We employ CoTracker [31] as the tracking model due to its stable tracking performance. It predicts a set of 2D trajectories, initialized from uniformly sampled grid locations within the first-frame object mask. For each pixel in the segmented region at time  $t$ , we compute its 2D displacement to time  $t + 1$  from the predicted trajectories, and convert it to a 2D velocity by dividing by the time step.

For each camera view, we lift the 2D pixel velocities to 3D using the corresponding depth measurements. Specifically, for a pixel at position  $(u, v)$  with depth  $d$  and 2D velocity  $(v_u, v_v)$ , we compute the 3D velocity by back-projecting both the current and next pixel positions to 3D coordinates using camera intrinsics, then computing the difference. This yields per-pixel 3D velocities in the camera frame, which we transform to the world frame using the calibration parameters. We aggregate velocities across all camera views by averaging the 3D velocity estimates for overlapping spatial regions, thereby improving robustness to noise and partial observations from individual views.

Given the 3D point cloud with per-point velocities at each timestep, we extract temporally consistent particle trajectories using an iterative rollout procedure. Starting from frame 0, we initialize particles at the point cloud positions  $\{x_i^0\}$ . For each subsequent timestep  $t$ , we propagate particles forward using their current velocities:  $x_i^{t+1} = x_i^t + v_i^t \cdot \Delta t$ . To update velocities at the new positions, we perform k-nearest



**Fig. 9: Interactive rope manipulation with PReD rendered using 3DGS.** The visualization overlays three timesteps, with later configurations in lighter opacity. As PReD predicts the dynamics, the 3D Gaussians attached to the particles are updated to render photorealistic images that maintain visual consistency with the observed appearance while accurately reflecting the predicted deformations.

neighbor search (with  $k = 5$ ) between the propagated particle positions and the observed point cloud at time  $t + 1$ . Each particle’s velocity is updated to the mean velocity of its  $k$  nearest neighbors in the observed cloud. This iterative process maintains correspondence across frames while being robust to tracking noise, yielding persistent point tracks suitable for training our dynamics model. The entire processing pipeline runs at approximately 2 Hz for sequences with 1000 particles.

## B. INTERACTIVE PHOTO-REALISTIC SIMULATION

Fig. 9 demonstrates the real-time interactive manipulation capability of PReD. By integrating accurate dynamics prediction with photorealistic rendering through 3D Gaussian Splatting, users can interactively manipulate objects while observing visually realistic deformations in real time. This differs from traditional simulation interfaces that rely on simplified geometric representations or offline rendering pipelines. The photorealistic feedback allows users to immediately assess both the geometric accuracy and visual plausibility of predicted deformations, making it valuable for scenarios where appearance matters alongside physical correctness, such as training data generation for vision-based manipulation.

## C. PLANNING ON PLUSH TOY

We extend our planning experiments to the plush toy, a volumetric object with long, flexible limbs. The plush toy requires accurate prediction of both volumetric deformation and limb dynamics, testing the framework’s ability to handle three-dimensional objects. Similar to rope, we design four manipulation objectives with varying complexity, as visualized in Fig. 10 (top): (1) lifting the object, (2) lowering it, (3) rotating it, and (4) bending it into a curved configuration.

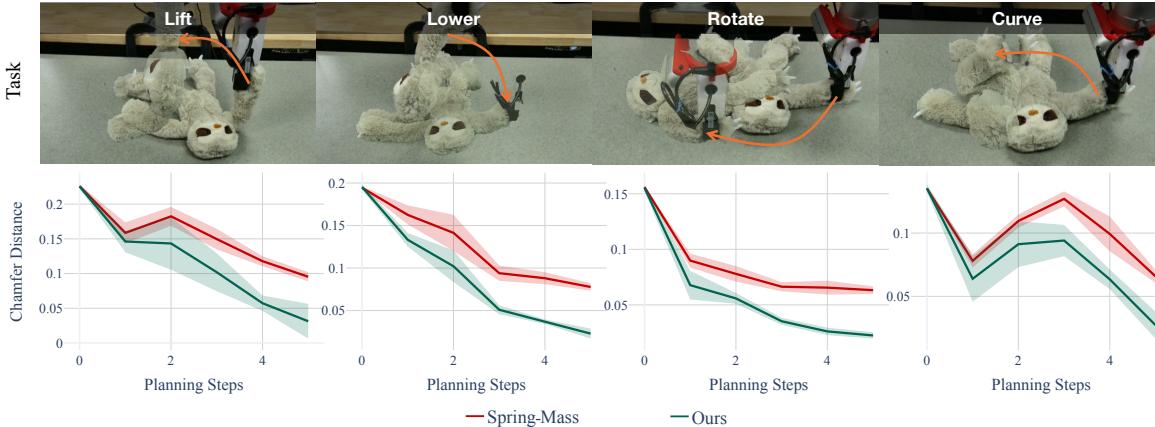


Fig. 10: **Planning results on plush toy manipulation tasks.** Top row displays the initial and target configurations for four manipulation objectives. Bottom row plots the Chamfer Distance evolution throughout MPPI planning. PReD (green) consistently achieves superior convergence compared to the tuned spring-mass baseline (red).

Following the same experimental setup, we conduct 10 trials per objective with identical initial and target states, allocating equal computational budget to both PReD and the spring-mass baseline. The results in Fig. 10 demonstrate that PReD exhibits consistent convergence behavior, reaching target configurations with reasonable residual error. The spring-mass baseline shows slower error reduction and occasionally plateaus at suboptimal solutions, particularly for tasks involving complex limb deformation, where the learned residuals provide crucial corrections to the physics backbone predictions.

#### D. COMPARISON TO PHYSTWIN

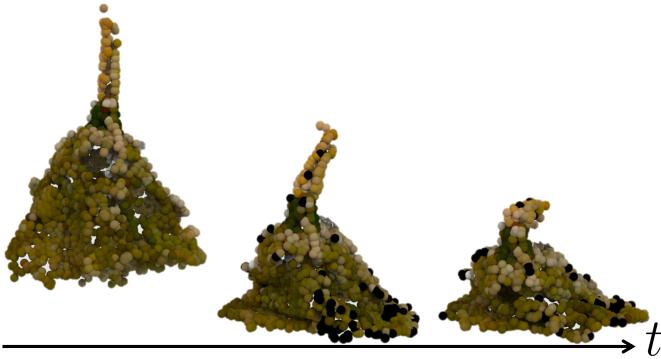


Fig. 11: **PhysTwin execution on duster.** Time progresses from left to right. The duster initially maintains its shape, but the rigid stem begins to bend unnaturally at intermediate time steps and completely collapses by the final frame. This demonstrates that spatially varying stiffness parameters recovered through first-order optimization fail to maintain the structural integrity of heterogeneous objects.

PhysTwin [30] represents another approach to modeling deformable objects using spring-mass representations. Similar to PReD, it employs a spring-mass simulator as its physical backbone. However, it attempts to recover spatially varying material parameters through first-order gradient-based optimization, with a differentiable simulator. While this strategy can identify heterogeneous material properties in principle, we find that it struggles with objects exhibiting extreme variations in properties, such as our duster which has a rigid stem and highly deformable feathers.

We evaluate PhysTwin on the duster object to assess its ability to handle heterogeneous materials. The execution results are shown in Fig. 11. As time progresses from left to right in the figure, the simulation exhibits catastrophic failure. In the leftmost frame, the duster appears correctly initialized. However, by the middle frame, the rigid stem begins to bend unnaturally at the junction between the stem and feathers. By the final frame, the rigid stem completely collapses, with the entire structure folding in on itself. This failure mode reveals limitation of relying solely on first-order optimization to infer spatially varying parameters.

In contrast, our approach as discussed in Sec. IV-E successfully models the duster by combining a tuned spring-mass simulator with learned residual corrections. Rather than attempting to capture all material heterogeneity through parameter optimization alone, our residual network learns to compensate for discrepancies between the physics model and reality.

#### E. NETWORK ARCHITECTURE DETAILS

Our network network consists of three primary components: an encoder, a decoder, and a temporal aggregator. The hyperparameter values for each component are provided in Table III. The encoder uses a Point Transformer V3 (PTv3) [72] architecture to extract per-particle features through patch-based self-attention, where the feature dimension determines the size of the latent representation. The input channels correspond to the concatenated position and velocity information from the current state, history timesteps, and the physics simulator predictions. The decoder employs a conditional NeRF-style MLP to map spatial locations to residual velocities, where hidden layers refers to the number of fully connected layers, hidden dimension specifies the width of each layer, and output channels indicates the dimensionality of the predicted residual velocity (3 for xyz components). The decoder takes as input the per-particle features from the encoder, concatenated with Fourier positional encodings of particle positions. The temporal aggregator employs a transformer encoder to refine predictions using temporal context. Here, the embedding dimension defines the size of the latent space for temporal

features, attention heads specifies the number of parallel attention mechanisms, transformer layers indicates the depth of the transformer stack, and feedforward dimension determines the width of the intermediate feedforward network within each transformer layer. Finally, the gating scale controls the magnitude of temporal corrections applied to the base velocity predictions, preventing large adjustments that could destabilize the simulation.

Parameter	Value
<i>Encoder (PTv3)</i>	
Feature dimension	64
Input channels	18
<i>Decoder</i>	
Hidden layers	2
Hidden dimension	64
Output channels	3
<i>Temporal Aggregator</i>	
Embedding dimension ( $d$ )	64
Attention heads ( $H$ )	4
Transformer layers ( $L$ )	2
Feedforward dimension ( $d_{ff}$ )	128
Gating scale	0.1

TABLE III. Network architecture hyperparameters.

#### F. ABLATION STUDY: DECODER ARCHITECTURE

In Sec. III-C, we describe the use of a NeRF-style decoder to produce per-particle residual velocities from the PTv3 features. We empirically found that including this decoder improves performance compared with directly using PTv3 outputs as velocity estimates. To isolate the decoder’s contribution, we compare two architectural variants on the rope object: (1) **PTv3 direct**, where the encoder features are directly projected to velocity corrections without the NeRF-style decoder, and (2) **PReD (full)**, which is our model. Both variants use the same encoder and temporal aggregator.

The results demonstrate that the decoder improves tracking accuracy. Without the decoder, the model achieves 2.8 cm MDE, 0.026 CD, and 0.13 EMD on rope. In contrast, the full architecture with the decoder achieves 2.6 cm MDE, 0.023 CD, and 0.012 EMD. These results confirm that the positional encoding and MLP decoder effectively refine the spatially varying residual corrections, particularly improving the distributional similarity measured by EMD.