

EMPM: Embodied MPM for Modeling and Simulation of Deformable Objects

Yunuo Chen^{1,2*}, Yafei Hu^{1*}, Lingfeng Sun¹, Tushar Kusnur¹, Laura Herlant¹, Chenfanfu Jiang²

*equal contribution

¹Robotics and AI Institute ²UCLA

Abstract

Modeling deformable objects – especially continuum materials – in a way that is physically plausible, generalizable, and data-efficient remains challenging across 3D vision, graphics, and robotic manipulation. Many existing methods oversimplify the rich dynamics of deformable objects or require large training sets, which often limits generalization. We introduce embodied MPM (EMPM), a deformable object modeling and simulation framework built on a differentiable Material Point Method (MPM) simulator that captures the dynamics of challenging materials. From multi-view RGB-D videos, our approach reconstructs geometry and appearance, then uses an MPM physics engine to simulate object behavior by minimizing the mismatch between predicted and observed visual data. We further optimize MPM parameters online using sensory feedback, enabling adaptive, robust, and physics-aware object representations that open new possibilities for robotic manipulation of complex deformables. Experiments show that EMPM outperforms spring-mass baseline models. Project website: <https://embodied-mpm.github.io>.




Figure 1: Overview. Our differentiable MPM simulator estimates material parameters from real observations such as 3D object shape. It handles diverse, complex deformable objects and can update parameters online from streaming video.

1. Introduction

Accurately modeling deformable objects, including their geometry, appearance, physical properties, and dynamics from visual information, and using these models for differentiable physics simulation, remains a long-standing and fundamental challenge in 3D vision, graphics, and robotics. By using a differentiable real-to-sim-to-real pipeline, we can employ gradient-based optimization to perform accurate interactive simulation of deformable objects directly from visual inputs, paving the way for constructing photo- and physics-realistic digital twins, as well as enabling

physics-aware robotic manipulation tasks.

Joint modeling and simulation of deformable objects often involve integrating geometry reconstruction, appearance rendering, and physics-based simulation. With recent advances in radiance field methods such as NeRF and Gaussian Splatting (GS) [1, 2], appearance modeling for both static and dynamic scenes and objects has seen significant progress through differentiable rendering techniques [3, 4]. However, physics-based simulation remains a core challenge in this domain. Previous approaches have tackled this problem using particle-based

physics models [5–7] and spring-mass systems [8]. While effective for certain scenarios, these methods often oversimplify the complex dynamics of deformable objects and struggle to accurately simulate various real-world objects. Although learning-based intuitive physics models [9–12] offer the potential to capture more complex behaviors through statistical approximation, they typically require large amounts of training data and often generalize poorly to novel objects—limiting their real-world applicability. Recently, more advanced methods based on the Material Point Method (MPM) [13] have shown promise in simulating deformable objects composed of continuum materials directly from visual information [14–16].

In this paper, we introduce EMPM, a real-to-sim-to-real pipeline that jointly performs MPM-based differentiable physics modeling and simulation, alongside geometry reconstruction and appearance rendering of deformable objects, including both elastic and elastoplastic materials from multi-view RGB-D video inputs. Using a pre-calibrated multi-view RGB-D camera rig, we reconstruct a 3D point cloud to represent the object’s geometry. This point cloud is then used to build a 3D Gaussian Splatting (3DGS) model for high-fidelity appearance rendering. The resulting geometry-appearance model is integrated with an MPM simulator to capture the object’s dynamic behavior. MPM physical parameters—such as elasticity and stiffness—are optimized by minimizing the discrepancy between real-world observations and simulated rollouts, measured in both 3D geometry and visual appearance throughout the object’s deformation process. Our approach includes both offline training of the MPM physics model parameters using pre-collected RGB-D data and online correction/adaptation based on live sensory observations. While concurrent work [17] also applies MPM simulation to robotic manipulation, it focuses solely on elastoplastic objects and lacks visual rendering capabilities. In contrast, our method handles a broader range of deformable materials and integrates photorealistic rendering. Experimental results on various deformable objects demonstrate the effectiveness of our MPM-based differentiable physics simulation, outperforming baseline spring-mass models [8].

In summary, the major contributions of this paper are as summarized as follows:

- We present a real-to-sim-to-real system based on MPM physics simulation, featuring differentiable parameter optimization grounded in real sensory observations.
- We present an online parameter optimization scheme that adapts to incoming observations, unlocking key power for potential downstream robotic control tasks.
- We demonstrate our method on a wide range of deformable objects, including both elastic and elastoplastic materials, which are challenging for existing real-to-sim-to-real simulation approaches.

2. Related Work

2.1. 4D visual reconstruction

4D reconstruction is the problem of modeling 3D geometry of dynamic objects/scene over time to capture their motion. Many methods track dynamic objects in 3D and differentiate the static background from the moving objects. Recent 3D point tracking feedforward models are able to generalize to unseen scenarios and handle certain degrees of occlusion [18, 19]. Recent radiance field approaches also adapt this traditional 4D reconstruction framework, extending the original NeRF to dynamics alternatives [20]. Radiance field methods based on explicit scene representation, such as dynamic Gaussian Splatting-based methods, adapt more naturally with the conventional 4D reconstruction more naturally by regularizing the rigid-body motion of underlying Gaussian kernels [3]. Our reconstruction pipeline follows these aforementioned methods and consists of point cloud lifting, 3D tracking and GS-based appearance modeling.

2.2. Physics simulation for deformable objects

Although physics simulation has been studied for decades [13, 21–23], methods that leverage sensory input and enable differentiability through observations have only gained popularity in recent years [8, 14, 16, 24]. Early methods follow particle-based models to model rigid objects and a few types of deformable objects [6]. Subsequent approaches have adopted the spring-mass model for simulating deformable objects [8, 25], and learning-based particle/grid-based methods [9, 10, 12, 26, 27]. As the dynamics of deformable objects and interactive motions such as tearing become more complex, spring-mass and learning-based models may fail to accurately capture the underlying physics. Subsequently, MPM-based approaches emerged as one of the leading approaches in this domain, which are well-suited for handling objects with continuum materials, and interactions like large deformations, fracture and tearing. Differentiable MPM-based methods been successfully demonstrated in prior work [16], which likewise underpins our work.

2.3. Differentiable MPM simulation

To accurately model and simulate real-world deformable objects, it is often necessary to optimize physical parameters such as material stiffness based on real-world observations, a process commonly referred to as system identification. The sources of discrepancy used to guide this optimization can vary, including force measurements [28], 3D geometry alignment [8], and photometric loss from image-based rendering [15, 16]. Our method follows this line of research by leveraging 3D geometry to optimize MPM parameters in a differentiable manner. While deriving analytical gradients for multi-stage physical simulators is prohibitively complex, recent advances have shown that auto-differentiation can effectively make MPM simulation a differentiable layer [29]. Most recent efforts in applying




Figure 2: Method Overview of EMPM. Our MPM simulation engine takes a reconstructed pointcloud and tracked 3D points as inputs. The model parameters are optimized using the discrepancy between the model’s prediction and 3D shape reconstruction and tracking. The control inputs applied to the dynamics model of MPM is computed as the velocity extracted from the reconstructions. In the inference stage, we test the model’s performance in terms of predicting 3D point positions and RGB image rendering of the 3D Gaussian splats.

differentiable MPM systems have leveraged frameworks such as Taichi [30, 31] and Warp [32], demonstrating reliable capabilities in physical parameter estimation and the advantages of GPU acceleration [15, 16, 33, 34]. Similarly, we also leverage Warp [32] as the tool for our differentiable MPM engine to differentiate through the model outputs.

3. Method

3.1. Overview

We model deformable object dynamics with an action-conditioned MPM simulator that is *embodied* through robot (or hand) interactions. Let X_t denote particle states at time t (with positions $\mathbf{x}_p$, velocities $\mathbf{v}_p$, and deformation gradients F_p), and let $\mathbf{u}_t$ denote controller (hands or grippers) actions obtained from hand tracking or robot telemetry. The MPM model advances state prediction at the next step as, $\hat{X}_{t+1} = f_\theta(X_t, \mathbf{u}_t)$, where θ denotes material parameters (e.g., Young’s modulus E , Poisson’s ratio ν , density ρ , plastic yield stress y).

We use multi-view RGB-D images to provide observations $O_t^i = \{I_t^i, D_t^i\}$, where I and D denote color and depth images, and i denotes the camera index. We reconstruct the point cloud $\tilde{X}_t \in \mathbb{R}^{3 \times N}$ by fusing the points and uplifting them to 3D world frame. Due to potential occlusions, we employ a point tracking model using the initial frame as a reference to obtain the 3D point trajectories $\tilde{X}_t^{\text{trk}}$.

Our goal is to optimize the MPM model parameters θ , based on the loss function of observed points $\tilde{X}$, $\tilde{X}^{\text{trk}}$ and model predictions $\hat{X}$, hence solving the general optimization problem in the following form:

$$\underset{\theta}{\operatorname{argmin}} L(\tilde{X}, \tilde{X}^{\text{trk}}, f_\theta(X, \mathbf{u})) \quad (1)$$

3.2. MPM and Continuum Mechanics

We briefly review the Material Point Method (MPM) [13] used in our work. MPM couples continuum balance laws with a hybrid particle–grid discretization, making it well-suited for large deformations, contact, and elastoplastic flow in embodied interactions.

MPM represents objects as particles while performing force computations on a background grid. At each step, particle mass and momentum are transferred to grid nodes (P2G), in eq. (2). Grid velocities are then updated with action conditions (eq. (3)) and transferred back to the particles for advection (G2P), in eq. (4). We summarize the explicit MPM update at time step t^n as follows, with p denoting particles and i denoting grid nodes:

P2G (particle to grid). We transfer mass and momentum using the affine particle-in-cell method (APIC) [35] scheme:

$$m_i^n = \sum_p w_{ip}^n m_p, \quad m_i^n \mathbf{v}_i^n = \sum_p w_{ip}^n m_p (\mathbf{v}_p^n + C_p^n (\mathbf{x}_i - \mathbf{x}_p^n)), \quad (2)$$

where w_{ip} are the B-spline weights and C_p^n is the affine velocity term.

Grid update (embodiment and contact). Grid velocities advance under stress, gravity, and constraints:

$$\mathbf{v}_i^{n+1} = \mathbf{v}_i^n + \frac{\Delta t}{m_i} \mathbf{f}_i(\mathbf{x}_i^n; \theta). \quad (3)$$

Here, $\mathbf{f}_i$ denotes the forces at grid node i , which are internally computed within f_θ , the full simulation step in eq. (1). We enforce *Dirichlet boundary velocities* on nodes in contact with the controller (from reconstructed or robot-measured end-effector motion $\mathbf{u}_t$), and apply Coulomb friction for table/gripper contact.

G2P (grid to particle). Grid velocities are transferred back to particles:

$$\mathbf{v}_p^{n+1} = \sum_i \mathbf{w}_{ip}^n \mathbf{v}_i^{n+1}, \quad \mathbf{x}_p^{n+1} = \mathbf{x}_p^n + \Delta t \mathbf{v}_p^{n+1}. \quad (4)$$

Deformation update. Deformation gradient is updated by:

$$\mathbf{F}_p^{n+1} = \left(\mathbf{I} + \Delta t \sum_i \mathbf{v}_i^{n+1} (\nabla \mathbf{w}_{ip}^n)^\top \right) \mathbf{F}_p^n. \quad (5)$$

Elastoplasticity. Most real-world deformable objects are not purely elastic and they undergo permanent shape changes under deformation. In continuum mechanics, this can be modeled by factoring the deformation gradient as $\mathbf{F} = \mathbf{F}^E \mathbf{F}^P$. Here the plastic part $\mathbf{F}^P$ serves as an intermediate state that stores the new rest shape, and only the elastic part $\mathbf{F}^E$ is used to compute stress. During MPM updates, we use a trial deformation gradient to track this process:

$$\mathbf{F}_p^{\text{trial}, n+1} = \text{updateF}(\mathbf{F}_p^{E, n}), \quad \mathbf{F}_p^{E, n+1} = \text{returnMap}(\mathbf{F}_p^{\text{trial}, n+1}) \quad (6)$$

In this paper, we use Fixed Corotated stress as elasticity model and use von Mises return mapping to model plasticity.

Connection to embodiments. Robot or hand motion $\mathbf{u}_t$ can enter as boundary velocities at the grid step. This makes the simulator *action-conditioned*: given candidate end-effector trajectories $\mathbf{u}_{t:t+H}$, we can predict object responses $\hat{\mathbf{X}}_{t:t+H}$ and score task costs (target geometry, collision). The same interface supports system identification (align to observations) and decision-making (rollout for planning). We next instantiate this in the offline identification setting.

3.3. Offline Simulation and Optimization

We instantiate the action-conditioned simulator for parameter identification from multi-view RGB-D videos. At time step t , RGB-D image observations $\{\mathbf{I}_t^i, \mathbf{D}_t^i\}_{i=0}^T$ are fused into point sets $\tilde{\mathbf{X}}_t$, and end-effector motion is reconstructed as $\mathbf{u}_{0:T-1}$, which is enforced as boundary velocities in the grid update.

Problem. Let $\mathbf{X}_0$ denote particles sampled from $\tilde{\mathbf{X}}_0$ (with point densification applied when required for enhanced dynamics). For material parameters $\theta = \{\mathbf{E}, \mathbf{v}, \rho, \mathbf{y}\}$, the MPM rollout recursively produces $\hat{\mathbf{X}}_{1:T} = f_\theta(\mathbf{X}_0, \mathbf{u}_{0:T-1})$. The identification problem is

$$\theta^* = \arg \min_{\theta} L_{\text{offline}}(\tilde{\mathbf{X}}_{1:T}, \tilde{\mathbf{X}}_{1:T}^{\text{trk}}, \hat{\mathbf{X}}_{1:T}), \quad (7)$$

which conditions learning on the observed interaction through $\mathbf{u}_{0:T-1}$.

Real-to-sim data processing. We use grounded-SAM2 [36] to compute the object and controllers (e.g. hands) masks for each image frame based on language prompts. Next we lift the 3D object point cloud measurement $\tilde{\mathbf{X}}_t$ with depth and camera extrinsics. 3D temporal correspondences $\tilde{\mathbf{X}}_t^{\text{trk}}$ are computed on offline sequences with an off-the-shelf point tracking model. The controller actions $\mathbf{u}_t$ are bound to grid nodes as Dirichlet boundary

velocities, and contacts with the table and grippers are modeled using Coulomb friction.

Objective. To bridge the discrepancy between MPM model predictions and real observations, we minimize the predicted position $\hat{\mathbf{X}}_t$ with lifted 3D points and their trackings:

$$L_{\text{offline}} = \lambda_{\text{dist}} \sum_{t=1}^T \text{Chamfer}(\hat{\mathbf{X}}_t, \tilde{\mathbf{X}}_t) + \lambda_{\text{trk}} \sum_{t=1}^T \sum_{j \in \mathcal{T}_t} \|\hat{\mathbf{X}}_{t,j} - \tilde{\mathbf{X}}_{t,j}^{\text{trk}}\|_2^2 \quad (8)$$

In this loss function, the first term measures the 3D Chamfer distance between the point clouds, while the second term computes the $\mathcal{L}^2$ squared distance between each particle, serving to align the motions of individual particles. We use occlusion masks from the pre-processing stage to filter out the invalid particles and exclude them from loss computation.

Optimization. We differentiate through the model rollout with Warp [32] to obtain the loss gradient $\nabla_{\theta} L_{\text{offline}}$ and update θ by gradient-based optimization methods. While gradient-based optimization enables smooth recovery of MPM physics parameters, we also provide an option of zero-order optimization using CMA-ES on the forward simulator. This strategy may be preferable for reducing memory usage or enabling a more aggressive parameter search.

3.4. Online Adaptive Identification

We also optimize the model parameters online during real interactions using streaming 3D observations.

Streaming supervision from cameras. At the beginning of streaming, we use multi-view RGB-D images to fuse an initial point cloud $\tilde{\mathbf{X}}_0$, which is then sampled as the MPM particles $\mathbf{X}$. Similar to offline optimization, we use Grounded SAM2 model to segment target objects. In subsequent streaming steps, we continuously perform segmentation and back-projection to lift the video frames to get 3D point clouds, serving as the 3D observation $\tilde{\mathbf{X}}_t$. Additionally, we record the current positions of the robot grippers as the control inputs $\mathbf{u}_t$ to define the boundary conditions. During online streaming, we observe that 3D point tracking becomes less reliable as tracked points tend to become invalid over time due to occlusion in indefinite incoming image sequences. Therefore, we do not perform online key point tracking in this stage.

Simulation and Online Correction. With the object geometry reconstructed in real time, we perform simultaneous simulation on the initial point cloud $\tilde{\mathbf{X}}_0$ and the current controller motion $\mathbf{u}_t$. Through continuous dynamic updates, the simulated object $\hat{\mathbf{X}}_t$ serves as a synchronized digital twin, closely aligning with the real object's motion.

To narrow the gap in real-to-sim object deformation and extract estimation of material parameters of the modeled objects, we perform differentiable-simulation steps along with the synchronized simulation, generating gradient flows to direct the object to 3D observation. To avoid unstable gradients, we restrict the optimizing steps to only *quasi-static* states – assuming objects are close to stable-equilibrium. More specifically, when an object is under

balanced forces, such as being pinned by two end-effectors, we perform a fixed-step H forward simulation, using the current particle states as input: $\tilde{\mathbf{X}}_{t+H} = f_{\theta}(\tilde{\mathbf{X}}_t, \mathbf{u}_t)$. Under the equilibrium assumption, the object should not deviate from its current rest shape (i.e., $\tilde{\mathbf{X}}_{t+H} \approx \hat{\mathbf{X}}_t$) if the material parameters are estimated correctly. To penalize the discrepancy caused by inaccurate parameters, we propose a loss function similar to the one used in offline optimization. We use the same 3D Chamfer distance L_{dist} to align the point clouds. Since we do not have direct correspondence between simulated and observed particles during online optimization, we additionally employ a 2D segmentation mask loss L_{mask} as a regularizer to align the simulated particles $\tilde{\mathbf{X}}_{t+H}$ with the segmented mask $\mathbf{M}_{t+H}$. Together, we have:

$$L_{\text{online}} = \lambda_{\text{dist}} L_{\text{dist}} + \lambda_{\text{mask}} L_{\text{mask}}. \quad (9)$$

Similarly, we back-propagate the gradient to obtain a new estimation of parameter θ' . After each step, the new parameter θ' will replace the current parameter in the simulation.

3.5. Applications

The EMPM simulator, potentially combined with offline and online optimization, offers a direct interface for incorporating deformable dynamics into embodied workflows. Because end-effector motions are bound to grid boundary conditions, EMPM produces action-conditioned rollouts that can potential be queried for both interactive simulation and robotic decision making.

With the digital twin synchronized to observed motions, operators can visualize the responses of objects under teleoperation or hypothetical controller inputs. This enables rapid what-if analysis of manipulation strategies and immediate feedback on feasibility, contact regimes, and deformation patterns. Candidate end-effector trajectories $\{\mathbf{u}_{t:t+H}\}$ are applied as boundary velocities to generate predictions $\tilde{\mathbf{X}}_{t:t+H}$. The same interface could support both gradient-free sampling and trajectory optimization. End-effector control points and simulator boundary conditions are aligned in EMPM, which transforms the simulator into a predictive engine that seamlessly connects perception to planning and execution.

4. Experiments

4.1. Experiment Setup

To demonstrate the coverage and capability of the MPM simulation engine in modeling deformable objects, we select objects from two major categories:

- **Elastic** objects: a rope segment; a piece of cloth; a soft elastic toy that bounces back when pushed or poked
- **Elastoplastic** objects: a piece of plasticine; a lump of bread dough; a pita bread

Figure 3 shows the types of objects we simulate in our experiments. We perform various manipulations on these

objects—including extending, stretching, squeezing, and tearing—to demonstrate the range of deformations achievable in our framework.




Figure 3: Objects used in our experiments. first row: elastic objects to simulate bending, poking and pulling; second row: elastoplastic objects to simulate fracture, stretch, and squeezing.

For the sensor setup used in data collection and online optimization, we employ a multi-view stereo RGB-D camera rig as the perception system. In our experiments, we set up three RealSense D455 cameras. We first collect data using human hands to facilitate the complicated interactions with the deformable objects that robot grippers struggle. This data is used in the offline simulation and optimization pipeline. For the robotic manipulation experiments within the online optimization framework, we use a bimanual setup with two Franka arms equipped with their original grippers.

Implementation details: We perform training/online optimization of the MPM model parameters on a single Nvidia A6000 GPU. We optimize the Young’s modulus E and Poisson’s ratio ν of the objects and assume the parameters remain constant across the material field. E is normalized to $[0, 1]$ by a factor of 10^6 . We leverage Warp’s automatic differentiation capability via its PyTorch integration, using the AdamW optimizer with a default learning rate of 10^{-4} .

4.2. Offline Simulation and Optimization

Using the collected videos, we follow the data processing procedure outlined above to extract the 3D information and perform offline simulation and optimization.

We visualize the simulated particles overlaid on the input video in Figure 4. Our simulated results exhibit physically plausible deformations consistent with the properties of each object type, closely aligning with the behavior observed in the real world. Additionally, we present Gaussian-rendered visualizations of the dynamic simulations in Figure 5, offering a photorealistic reproduction of the visual observations. The Gaussian splatting is based on the original 3DGS and the gsplat [37] library’s implementations and the motion of the Gaussian centers is interpolated with the motions of neighboring MPM particle nodes using Linear Blend Skinning (LBS) [8].




Figure 4: Qualitative results showing the simulation predictions and real observations for different objects, comparing our method and PhysTwin [8]. Our method especially shines at modeling fracture, illustrated in the pita bread example, where PhysTwin model prediction completely fails. In the last frame of the plasticine example, the PhysTwin model prediction fails due to the over-bend of the internal spring model, where our model can still successfully model the squeezing of the object.

To demonstrate the advantages of our method, we compare our results with PhysTwin [8], which models deformable objects using a spring-mass system. We visualize the qualitative results and photorealistic rendering in Figures 4 and 5. To quantitatively evaluate the performance, we report the reconstruction qualities in terms of geometric differences of point clouds, as well as visual rendering qualities. We also compare our results with a learning-based dynamic model PGND [10], reported in Table 1.

Leveraging our MPM-based simulator, our method achieves higher simulation accuracy and narrows the gap between real-world captured data and the simulated prediction. Compared to the spring-mass system baseline [8] and neural dynamic baseline [10], our approach demonstrates better alignment in both geometric and visual metrics.

In addition to offering higher overall simulation accuracy, another key advantage of EMPM is its ability to handle a broad spectrum of materials and physical phenomena within our framework. Our method can model both hyper-elastic and elastoplastic materials, encompassing a wide range of everyday objects—from purely elastic jelly to plasticine, which exhibits largely different behavior under external forces. In contrast, spring-mass systems are limited to modeling elastic springs and cannot capture intrinsic non-elastic behaviors, such as permanent shape change beyond the elastic threshold in elastoplastic materials, as demonstrated in our experiments. Moreover, spring-mass systems are unable to simulate more complex phenomena like material damage (e.g., in the pita experiment), since elastic springs do not naturally accommodate fracture. On the other hand, our method accurately captures fracture

dynamics and produces realistic and desirable outcomes.

4.3. Online Optimization

To demonstrate the efficacy of our online optimization, we conduct experiments on one elastic object (rope) and one elastoplastic object (bread dough). We teleoperate a pair of bimanual Franka arms to control the objects. Specifically, we control the arms to move along various trajectories, while also allowing the objects to remain steady between motions. With an initial point cloud generated at the beginning of the loop, we simulate the object point clouds under the same manipulations. As shown in Figure 6, our simulated twins closely follow the motion of the real objects.

We perform online differentiable simulation to optimize the material stiffness of the objects during the simulation, with the loss function shown in Eq. 9. In our experiments, we perform one optimization step every 5 online streaming step and provided the object is relatively steady at that step. In each optimization step, we simulate for a total of 10 forward simulation steps. With the optimization, the alignments between simulated objects and real observations improved over time, while we gradually adjust our estimation of the physical parameters. To quantitatively evaluate the improvement with simulation correction, we report the 2D mask and 3D distance errors in the online experiments for both objects, with and without optimization, in Table 2. We observe improved simulation accuracy with the presence of our online correction.

We also include running time analysis results comparing EMPM with the baseline methods PhysTwin [8] and PGND [10] in Table 3. For PhysTwin, we conducted

Method	Elastic Objects						Elastoplastic Objects					
	Dist ↓	Track↓	IoU % ↑	PSNR ↑	SSIM ↑	LPIPS ↓	Dist ↓	Track↓	IoU % ↑	PSNR ↑	SSIM ↑	LPIPS ↓
PGND [10]	0.0618	-	0.6898	21.06	0.8938	0.0977	0.0245	-	0.5069	21.26	0.9472	0.0738
PhysTwin [8]	0.0227	0.1467	0.6981	24.12	0.9312	0.0756	0.0177	0.1108	0.6918	27.01	0.9695	0.0353
Ours	0.0222	0.1377	0.7095	24.19	0.9319	0.0711	0.0082	0.1014	0.7768	27.82	0.9725	0.0291

Table 1: Quantitative Comparison on Offline Simulation. We compare our prediction results with PhysTwin [8] and PGND [10] in terms of both reconstruction accuracy (measured by distance and tracking of the 3D point cloud; tracking errors from [10] are omitted because we do not have point correspondence between their results and ours) and visual quality (evaluated using IoU, PSNR, SSIM, and LPIPS) through photorealistic rendering with Gaussian Splatting. The results are categorized by object type: elastic and elastoplastic. Our method outperforms baselines across all metrics, with a particularly notable advantage in handling elastoplastic objects.



Figure 5: Qualitative results demonstrate photorealistic reconstruction using Gaussian Splatting. Our reconstructed dynamics closely match real observations, whereas the spring-mass-based PhysTwin [8] fails to accurately model these materials.

Experiment	Without Optimization		With Optimization	
	$L_{\text{mask}} \downarrow$	$L_{\text{dist}} \downarrow$	$L_{\text{mask}} \downarrow$	$L_{\text{dist}} \downarrow$
Rope	0.0456	0.0057	0.0428	0.0054
Bread Dough	0.0031	0.0060	0.0024	0.0059

Table 2: Quantitative results demonstrate improvements in geometric accuracy achieved through progressive correction during online simulation.

both CMA-ES and gradient-based optimization and report only the gradient training time. For PGND, we train 10 episodes following the default setup. The training of EMPM is the most efficient due to its optimization over entire training frames as a batch. PGND achieves the lowest inference time due to the feedforward pass of a relatively small network, but with the cost of very long training time.

4.4. Applications

In our experiment, we demonstrate a proof-of-concept application with real robotics hardware to show how our method would be used for deformable object manipulation

Running time (s)	Elastic Objects		Elastoplastic Objects	
	Training ↓	Testing ↓	Training ↓	Testing ↓
PGND [10]	50675.50	5.06	48464.40	5.34
PhysTwin [8]	496.84	11.49	589.51	15.94
Ours	161.80	14.71	171.60	22.21

Table 3: Running time (in seconds) analysis of EMPM and baseline approaches using the same training data.

tasks, shown in Figure 7. With accurate dynamics model for deformable objects proposed in this work, we can potentially deploy model-based robotic motion planning and control and enable autonomous manipulations.

5. Conclusions & Discussions

We presented EMPM, an embodied, differentiable MPM framework for modeling geometry, appearance, and dynamics of complex deformable objects – including elastic and soft/hard elastoplastic materials that are difficult for spring-mass baselines. EMPM updates material parameters from real observations in both offline and online settings and improves over spring-mass baselines qualitatively.




Figure 6: Qualitative results of the online simulation demonstrate the tracking performance of our MPM model against real observations. With online optimization guided by real-time reconstruction, the alignment between simulation and observation improves progressively over time. The losses are normalized from 0 to 1 for better illustration. Updates in the bread dough case are not obvious due to a relatively small motion and the occlusion of the grippers, it is recommended to refer to the video for a better dynamic visualization.




Figure 7: Application of EMPM in potential model-based planning and control for manipulation. Here we simulate model rollout and let robot track these changes.

tively and quantitatively. Future work includes enhanced appearance modeling for Gaussian Splatting and integrating model-based control for autonomous deformable manipulation.

Limitation We observe that a key obstacle during optimization is obtaining reliable point tracking, whose quality degrades under occlusion and large deformations, leading to unreliable 3D trajectories. Tracking conditioned on the first frame often yields zero reliably tracked points after just a few seconds, rendering it unsuitable for online optimization. We note that emerging tracking models hold promise for mitigating this issue, potentially enabling more faithful 3D guidance for reconstruction and facilitating the integration of tracking-based losses into online optimization.

References

- [1] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020.
- [2] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4), July 2023.
- [3] Jonathon Luiten, Georgios Kopanas, Bastian Leibe, and Deva Ramanan. Dynamic 3d gaussians: Tracking by persistent dynamic view synthesis. In *3DV*, 2024.
- [4] Hidenobu Matsuki, Riku Murai, Paul H. J. Kelly, and Andrew J. Davison. Gaussian Splatting SLAM. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024.
- [5] Jad Abou-Chakra, Feras Dayoub, and Niko Sünderhauf. Particlenet: A particle-based encoding for online neural radiance fields. In *WACV*, pages 5975–5984, January 2024.
- [6] Jad Abou-Chakra, Krishan Rana, Feras Dayoub, and Niko Sünderhauf. Physically embodied gaussian splatting: A real-time correctable world model for robotics. In *8th Annual Conference on Robot Learning*, 2024.
- [7] Jad Abou-Chakra, Lingfeng Sun, Krishan Rana, Brandon May, Karl Schmeckpeper, Niko Sünderhauf, Maria Vittoria Minniti, and Laura Herlant. Real-is-sim: Bridging the sim-to-real gap with a dynamic digital twin, 2025.
- [8] 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. *ICCV*, 2025.
- [9] Mingtong Zhang, Kaifeng Zhang, and Yunzhu Li. Dynamic 3d gaussian tracking for graph-based neural dynamics modeling. In *8th Annual Conference on Robot Learning*, 2024.
- [10] Kaifeng Zhang, Baoyu Li, Kris Hauser, and Yunzhu Li. Particle-grid neural dynamics for learning deformable object models from rgb-d videos. In *Proceedings of Robotics: Science and Systems (RSS)*, 2025.
- [11] Jan Obrist, Miguel Zamora, Hehui Zheng, Ronan Hinchet, Firat Ozdemir, Juan Zarate, Robert K. Katzschmann, and Stelian Coros. Pokeflex: A real-world dataset of volumetric deformable objects for robotics, 2025.
- [12] 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. In *9th Conference on Robot Learning (CoRL 2025)*, 2025.
- [13] Chenfanfu Jiang, Craig Schroeder, Joseph Teran, Alexey Stomakhin, and Andrew Selle. The material point method

-
- for simulating continuum materials. In *Acm siggraph 2016 courses*, pages 1–52. 2016.
- [14] Tianyi Xie, Zeshun Zong, Yuxing Qiu, Xuan Li, Yutao Feng, Yin Yang, and Chenfanfu Jiang. Physgaussian: Physics-integrated 3d gaussians for generative dynamics. *CVPR*, 2024.
 - [15] Xuan Li, Yi-Ling Qiao, Peter Yichen Chen, Krishna Murthy Jatavallabhula, Ming Lin, Chenfanfu Jiang, and Chuang Gan. Pac-nerf: Physics augmented continuum neural radiance fields for geometry-agnostic system identification. 2023.
 - [16] Tianyuan Zhang, Hong-Xing Yu, Rundui Wu, Brandon Y. Feng, Changxi Zheng, Noah Snively, Jiajun Wu, and William T. Freeman. PhysDreamer: Physics-based interaction with 3d objects via video generation. In *ECCV*. Springer, 2024.
 - [17] Xintong Yang, Ze Ji, and Yu-Kun Lai. Differentiable physics-based system identification for robotic manipulation of elastoplastic materials. In *IJRR*, 2025.
 - [18] Zhengqi Li, Richard Tucker, Forrester Cole, Qianqian Wang, Linyi Jin, Vickie Ye, Angjoo Kanazawa, Aleksander Holynski, and Noah Snively. MegaSaM: Accurate, fast and robust structure and motion from casual dynamic videos. *CVPR*, 2025.
 - [19] Nikita Karaev, Iurii Makarov, Jianyuan Wang, Natalia Neverova, Andrea Vedaldi, and Christian Rupprecht. CoTracker3: Simpler and better point tracking by pseudo-labelling real videos. 2024.
 - [20] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-NeRF: Neural Radiance Fields for Dynamic Scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021.
 - [21] Demetri Terzopoulos, John Platt, Alan Barr, and Kurt Fleischer. Elastically deformable models. 1987.
 - [22] James F. O’Brien, Jessica K. Hodgins. Graphical modeling and animation of brittle fracture. 1999.
 - [23] Eftychios Sifakis and Jernej Barbic. Fem simulation of 3d deformable solids: A practitioner’s guide to theory, discretization and model reduction. 2012.
 - [24] Yutao Feng, Yintong Shang, Xuan Li, Tianjia Shao, Chenfanfu Jiang, and Yin Yang. Pie-nerf: Physics-based interactive elastodynamics with nerf. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4450–4461, June 2024.
 - [25] Licheng Zhong, Hong-Xing Yu, Jiajun Wu, and Yunzhu Li. Reconstruction and simulation of elastic objects with spring-mass 3d gaussians. *ECCV*, 2024.
 - [26] Yunzhu Li, Jiajun Wu, Russ Tedrake, Joshua B Tenenbaum, and Antonio Torralba. Learning particle dynamics for manipulating rigid bodies, deformable objects, and fluids. In *ICLR*, 2019.
 - [27] Kaifeng Zhang, Baoyu Li, Kris Hauser, and Yunzhu Li. Adaptigraph: Material-adaptive graph-based neural dynamics for robotic manipulation. In *Proceedings of Robotics: Science and Systems (RSS)*, 2024.
 - [28] Wenqiang Xu, Dongzhe Zheng, Yutong Li, Jieji Ren, and Cewu Lu. Differentiable fluid physics parameter identification by stirring and for stirring. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2024.
 - [29] Yuanming Hu, Jiancheng Liu, Andrew Spielberg, Joshua B Tenenbaum, William T Freeman, Jiajun Wu, Daniela Rus, and Wojciech Matusik. Chainqueen: A real-time differentiable physical simulator for soft robotics. In *2019 International conference on robotics and automation (ICRA)*, pages 6265–6271. IEEE, 2019.
 - [30] Yuanming Hu, Tzu-Mao Li, Luke Anderson, Jonathan Ragan-Kelley, and Frédo Durand. Taichi: a language for high-performance computation on spatially sparse data structures. *ACM Transactions on Graphics (TOG)*, 38(6):1–16, 2019.
 - [31] Yuanming Hu, Luke Anderson, Tzu-Mao Li, Qi Sun, Nathan Carr, Jonathan Ragan-Kelley, and Frédo Durand. DiffTaichi: Differentiable programming for physical simulation. *arXiv preprint arXiv:1910.00935*, 2019.
 - [32] Miles Macklin. Warp: A high-performance python framework for gpu simulation and graphics. <https://github.com/nvidia/warp>, March 2022. NVIDIA GPU Technology Conference (GTC).
 - [33] Haozhe Su, Xuan Li, Tao Xue, Chenfanfu Jiang, and Mridul Aanjaneya. A generalized constitutive model for versatile mpm simulation and inverse learning with differentiable physics. *Proceedings of the ACM on Computer Graphics and Interactive Techniques*, 6(3):1–20, 2023.
 - [34] Zeshun Zong, Xuan Li, Minchen Li, Maurizio M Chiaramonte, Wojciech Matusik, Eitan Grinspun, Kevin Carlberg, Chenfanfu Jiang, and Peter Yichen Chen. Neural stress fields for reduced-order elastoplasticity and fracture. In *SIGGRAPH Asia 2023 Conference Papers*, pages 1–11, 2023.
 - [35] 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.
 - [36] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, Eric Mintun, Junting Pan, Kalyan Vasudev Alwala, Nicolas Carion, Chaoyuan Wu, Ross Girshick, Piotr Dollár, and Christoph Feichtenhofer. Sam 2: Segment anything in images and videos, 2024.
 - [37] Vickie Ye, Ruilong Li, Justin Kerr, Matias Turkulainen, Brent Yi, Zhuoyang Pan, Otto Seiskari, Jianbo Ye, Jeffrey Hu, Matthew Tancik, and Angjoo Kanazawa. gsplat: An open-source library for gaussian splatting, 2024.