# MoSA: Motion-constrained Stress Adaptation for Mitigating Real-to-Sim Gap in Continuum Dynamics via Learning Residual Anisotropy

# Introduction

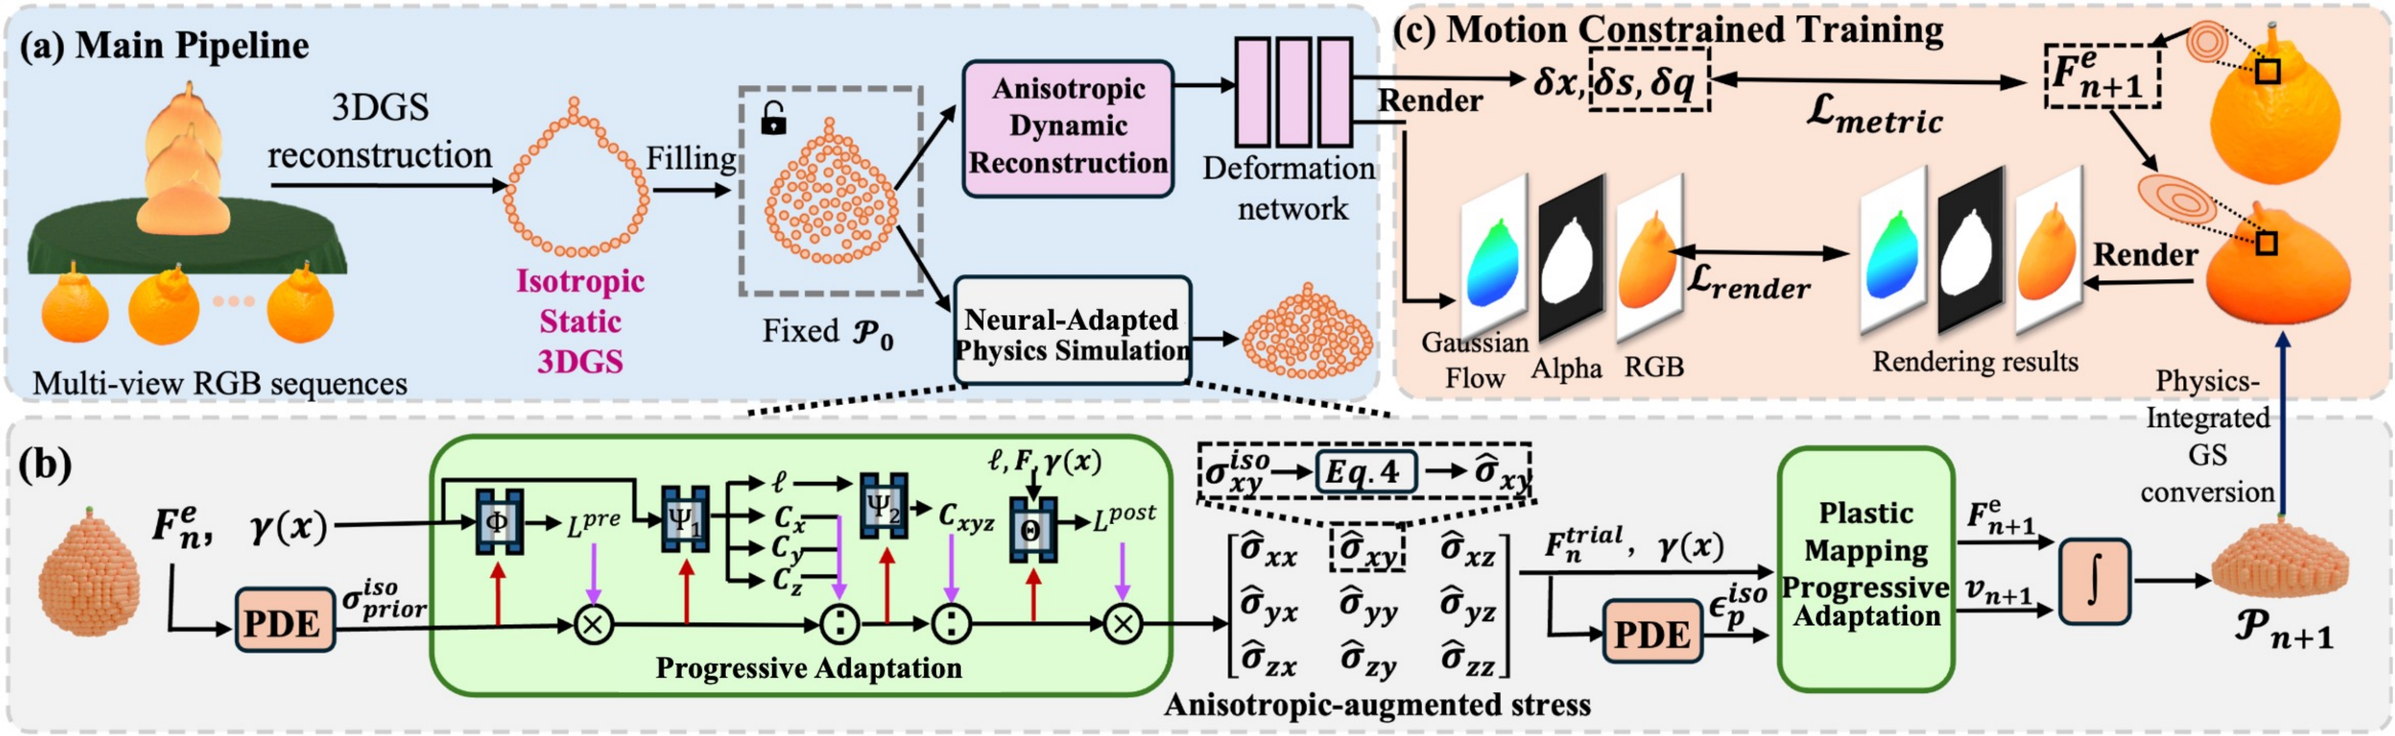


Overview of the pipeline. (a) Two-stage dynamic reconstruction. (b)Simulation with progressive anisotropic stress adaptation (c) Motion-constrained optimization strategy

# Methodology

## Structured, Progressive Stress Adaptation for modeling anisotropic residual effect

## Motivation:

- Classic constitutive models often assume materials are homogeneous and isotropic. However, this approximation is rarely exact in the real world
- Linear anisotropic models use a fourth-order stiffness tensor. Accounting for symmetry properties, this tensor still has 36 independent parameters, making it far more expressive

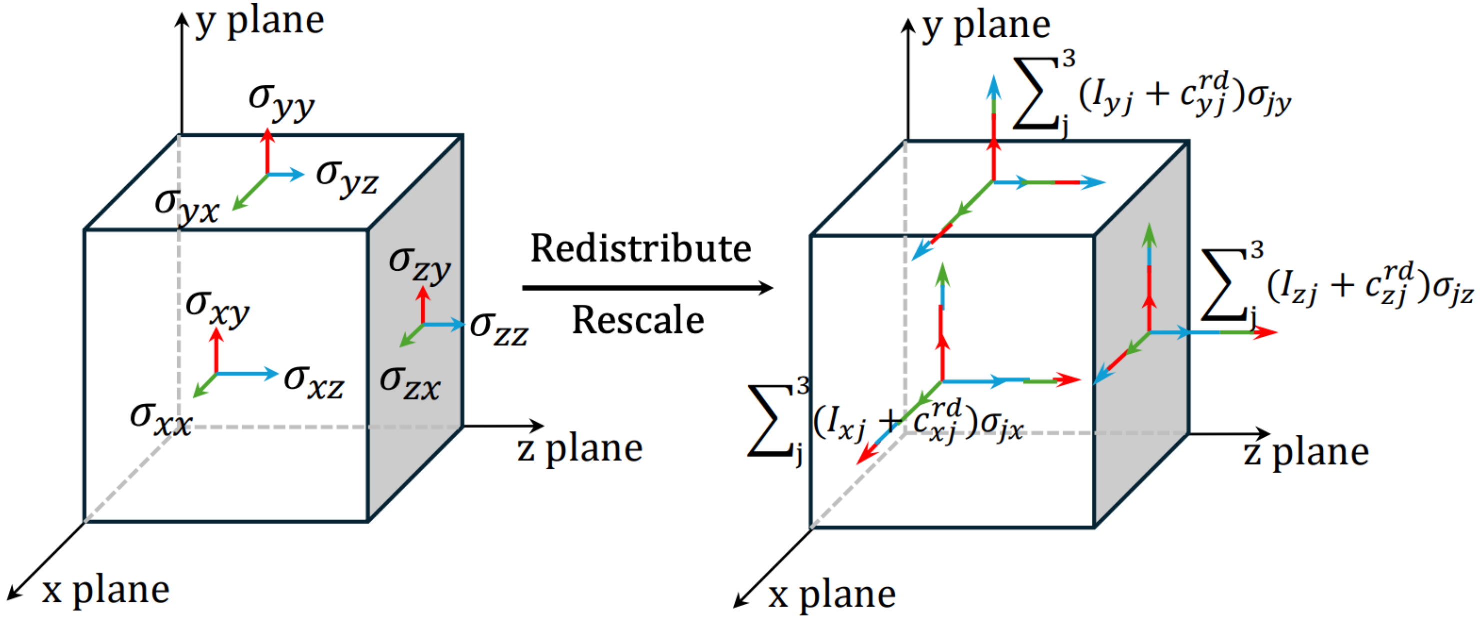


Rescaling and redistributing of stress tensor

## Solution:

We correct a prior isotropic constitutive model with a physics-informed network that projects isotropic stress onto anisotropic configurations given the material state. Corrections are applied progressively to adjust the prior stress.

# Methodology

## Motion-constrained Optimization Strategy

**Motivation:**

- Rendering losses alone give weak, indirect supervision for deformation.

**Solution:**

- **Two-stage reconstruction**
  - First, build a static 3DGS at time 0 to get a clean particle scaffold.
  - Then, train a dynamic 3DGS on top of it, letting Gaussians deform and rotate over time.
- **Velocity supervision**
  - Extract “Gaussian flow” maps from each camera view.
  - Use them to supervise the simulated particle velocities – sharper motion inside the silhouette.
- **Deformation gradient supervision**
  - Relate the covariance change of Gaussians to the deformation gradient F.
  - Break F into scaling and rotation parts (via SVD), then supervise rotation and scale separately

# What does MoSA learn?

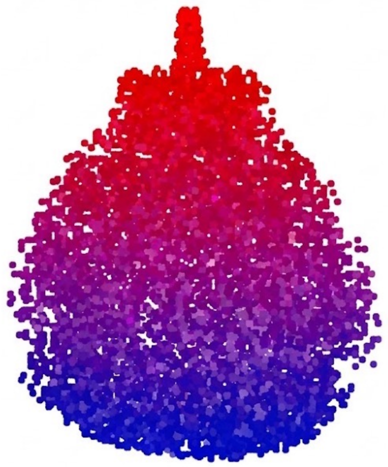

To check that MoSA captures *physically meaningful* residual effects rather than overfitting, we probe two complementary aspects of the learned model: **(i) directional anisotropy** of the stress-adaptation operator, and **(ii) spatial heterogeneity** of the local material field.

**Directional anisotropy.** On the synthetic *PAC-NeRF* benchmark — generated by a purely isotropic simulator — $w(\theta)$ stays near zero with no directional pattern, so MoSA does *not* introduce spurious anisotropy. On the real *Mandarin* sequence, $w(\theta)$ exhibits pronounced directionality aligned with the object's principal axis, meaning the operator adaptively recovers anisotropic effects *only when they truly exist in the data*.

**Spatial heterogeneity.** The learned field $\eta(\mathbf{x})$ locally modulates the global material parameter and produces smooth, object-dependent stiffness patterns — confirming that the continuous field captures structured material variation rather than fitting unstructured residual noise.

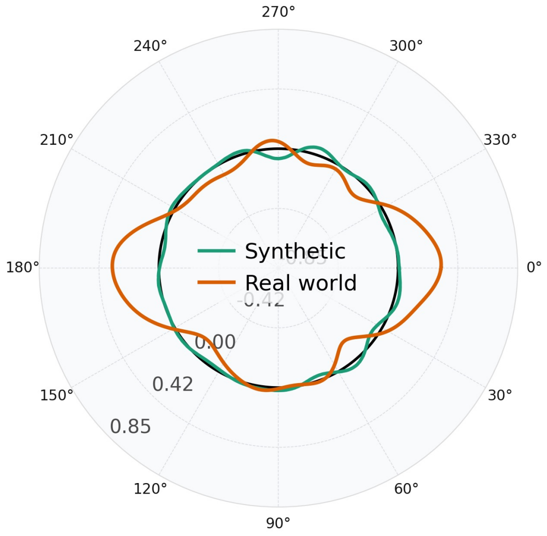

# Introduction

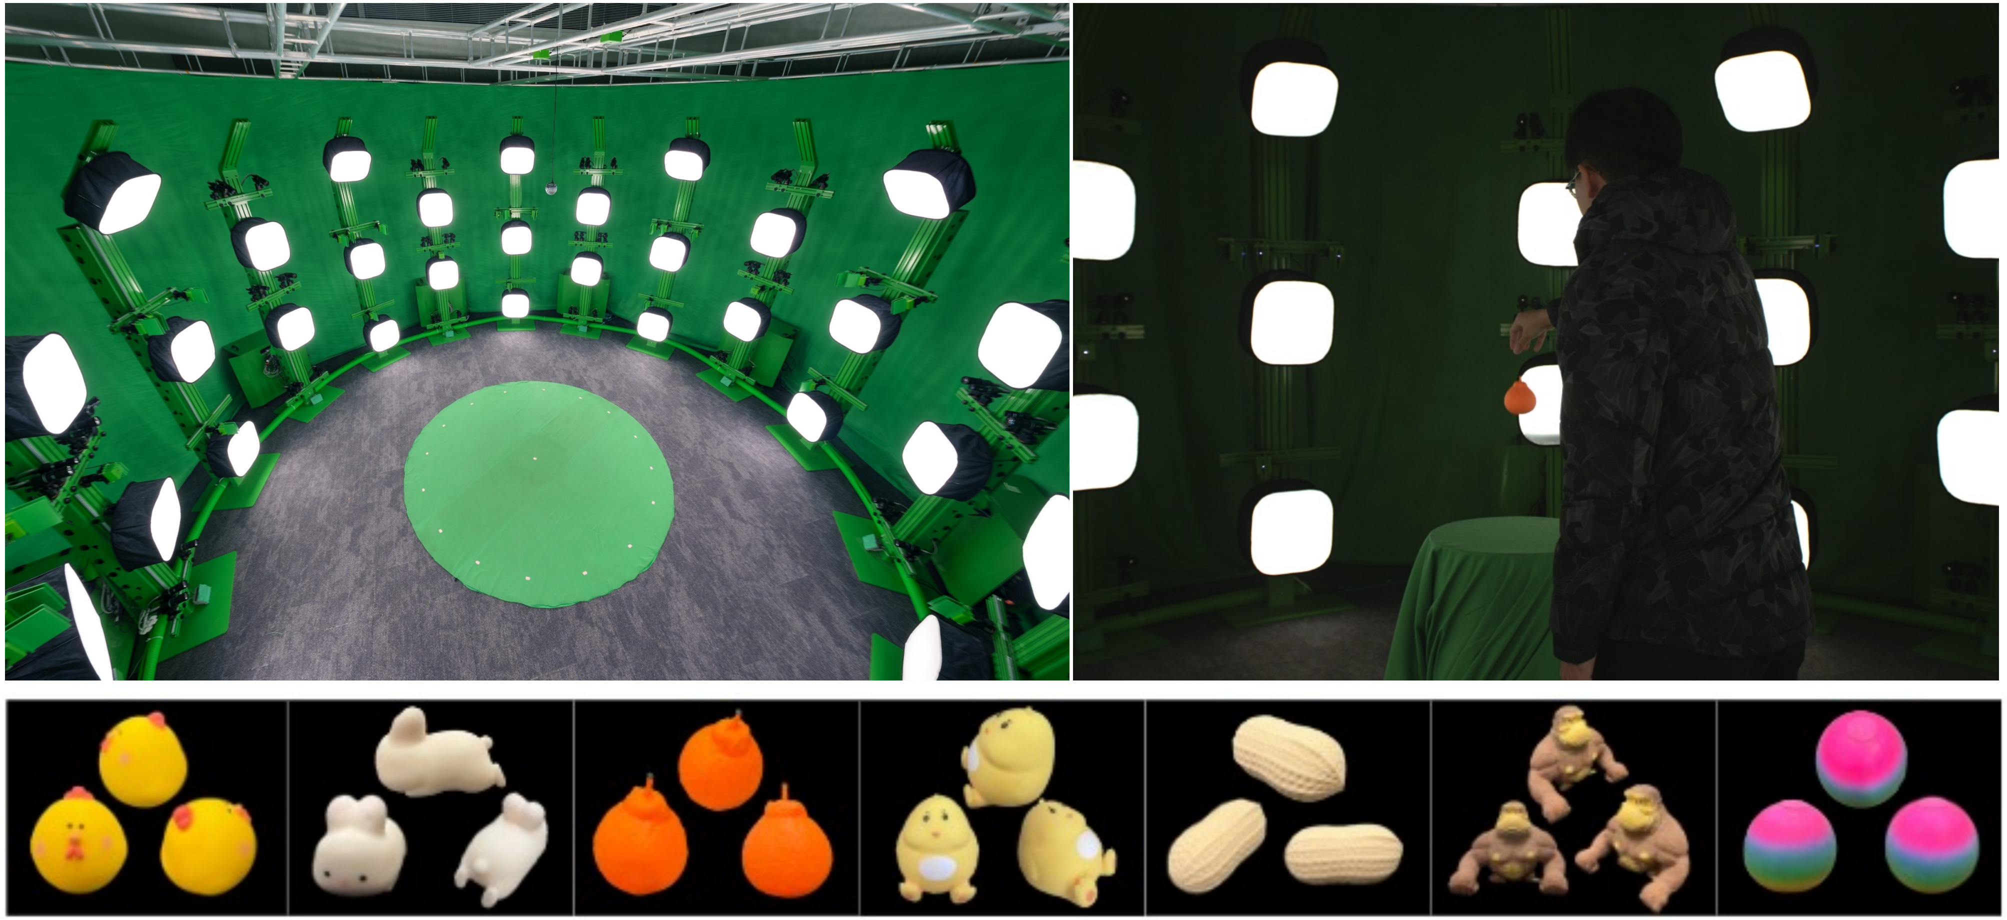

Our advanced light field reconstruction system and visualization of the objects in our realistic dataset

# Experiment Results

| Methods | torus CD↓ | torus EMD↓ | cat CD↓ | cat EMD↓ | playdoh CD↓ | playdoh EMD↓ | droplet CD↓ | droplet EMD↓ | Cream CD↓ | Cream EMD↓ | Bird CD↓ | Bird EMD↓ | Letter CD↓ | Letter EMD↓ | Mean CD↓ | Mean EMD↓ |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| PAC | 21.8 | 11.6 | 9.8 | 14.4 | 18.6 | 5.6 | 10.4 | 3.2 | 20.5 | 12.7 | 19.3 | 21.1 | 12.8 | 8.5 | 16.2 | 11.0 |
| DEL | 21.7 | 10.7 | 7.9 | 12.8 | 12.2 | 2.5 | 9.8 | 1.7 | 19.8 | 9.8 | 17.8 | 20.2 | 12.6 | 7.2 | 14.5 | 9.3 |
| GIC | 20.2 | 9.9 | 7.6 | 12.6 | 12.3 | 2.5 | 10.2 | 1.9 | 19.5 | 10.1 | 16.5 | 19.5 | 10.3 | 7.5 | 13.8 | 9.1 |
| Ours | **20.1** | **9.8** | **7.3** | **12.4** | **11.4** | **2.3** | **9.6** | **1.5** | **19.4** | **9.5** | **16.3** | **19.2** | **10.1** | **6.6** | **13.5** | **8.8** |

Table 1: Dynamic Grounding on PAC-NeRF dataset. All metrics are scaled by 100 for clarification.

| Methods | **Chick1** PSNR↑ | **Chick1** SSIM↑ | **Gorilla** PSNR↑ | **Gorilla** SSIM↑ | **Mandarin** PSNR↑ | **Mandarin** SSIM↑ | **Chick2** PSNR↑ | **Chick2** SSIM↑ | **Peanut** PSNR↑ | **Peanut** SSIM↑ | **Rabbit** PSNR↑ | **Rabbit** SSIM↑ | **RBball** PSNR↑ | **RBball** SSIM↑ | **Mean** PSNR↑ | **Mean** SSIM↑ |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| DEL | 28.32 | 91.6 | 29.11 | 90.4 | 31.92 | 92.4 | 28.59 | 90.9 | 28.38 | 91.3 | 28.57 | 91.4 | 30.95 | 91.7 | 29.41 | 91.4 |
| NeuMA | 30.73 | 92.4 | 29.78 | 91.1 | 31.85 | 92.4 | 28.92 | 91.0 | 30.70 | 91.8 | 28.30 | 92.3 | 30.74 | 91.1 | 30.00 | 91.7 |
| GIC | 30.93 | 92.5 | 29.75 | 91.0 | 29.54 | 91.6 | 28.17 | 90.9 | 32.88 | 91.3 | 28.08 | 91.3 | 30.78 | 91.7 | 30.02 | 91.5 |
| Vid2Sim | 26.71 | 90.9 | 29.02 | 90.0 | 25.85 | 89.8 | 25.70 | 89.5 | 30.69 | 94.6 | 26.85 | 90.3 | 31.71 | 91.7 | 28.08 | 91.0 |
| Ours | **32.05** | **92.7** | **30.19** | **91.9** | **32.83** | **92.7** | **30.17** | **91.6** | **33.01** | **92.0** | **30.35** | **92.1** | **32.06** | **92.7** | **31.35** | **92.3** |

Table 2: Quantitative comparisons of the initial state generalizations on the real-world dataset

| | PAC-NeRF | GIC | Ours | Ground Truth |
|---|---|---|---|---|
| Droplet | $\mu, \kappa = 2.09, 1.085$ | $\mu, \kappa = 2.01, 0.18,$ | $\mu, \kappa = \mathbf{2.01}, \mathbf{1.07}$ | $\mu, \kappa = 2, 1$ |
| Letter | $\mu, \kappa = 83.85, 1.35$ | $\mu, \kappa = 95.05, 1.00$ | $\mu, \kappa = \mathbf{97.00}, \mathbf{1.00}$ | $\mu, \kappa = 100, 10^5$ |
| Cream | $\mu, \kappa = 1.21, 1.57$ | $\mu, \kappa = 1.03, 1.48$ | $\mu, \kappa = \mathbf{1.01}, \mathbf{1.39}$ | $\mu, \kappa = 1, 1$ |
| | $\tau_Y, \eta = 3.16, 5.6$ | $\tau_Y, \eta = \mathbf{2.98}, 6.6$ | $\tau_Y, \eta = 2.97, \mathbf{7.75}$ | $\tau_Y, \eta = 3, 10$ |
| Toothpaste | $\mu, \kappa = 6.51, 2.22$ | $\mu, \kappa = 4.19, 9.24$ | $\mu, \kappa = \mathbf{4.52}, \mathbf{9.37}$ | $\mu, \kappa = 5, 10$ , |
| | $\tau_Y, \eta = 228, \mathbf{9.77}$ | $\tau_Y, \eta = 226, 9.1$ | $\tau_Y, \eta = \mathbf{212}, 9.63$ | $\tau_Y, \eta = 200, 10$ |
| Torus | $E, \nu = 1.04, 0.322$ | $E, \nu = \mathbf{0.99}, 0.295$ | $E, \nu = \mathbf{0.99}, \mathbf{0.298}$ | $E, \nu = 1, 0.3$ |
| Bird | $E, \nu = 2.78, 0.273$ | $E, \nu = 3.08, 0.284$ | $E, \nu = \mathbf{3.02}, \mathbf{0.29}$ | $E, \nu = 3, 0.3$ |
| Playdoh | $E, \nu, \tau_Y = 3.84, 0.272, 1.69$ | $E, \nu, \tau_Y = 1.58, 0.322, 1.56$ | $E, \nu, \tau_Y = \mathbf{1.72}, \mathbf{0.289}, \mathbf{1.55}$ | $E, \nu, \tau_Y = 2, 0.3, 1.54$ |
| Cat | $E, \nu, \tau_Y = 1.61, 0.293, 3.57$ | $E, \nu, \tau_Y = \mathbf{0.98}, 0.296, 3.76$ | $E, \nu, \tau_Y = \mathbf{0.98}, \mathbf{0.297}, \mathbf{3.77}$ | $E, \nu, \tau_Y = 1, 0.3, 3.85$ |
| Trophy | $\theta^0_{fric} = 36.1°$ | $\theta^0_{fric} = 38.0°$ | $\theta^0_{fric} = \mathbf{38.5}°$ | $\theta^0_{fric} = 40°$ |

Table 3: System Identification Performance on PAC-NeRF Dataset. All the values and quantities in the table are scaled based on the magnitude of the ground truth values.

# Experiment Results

## Chick1

Groundtruth

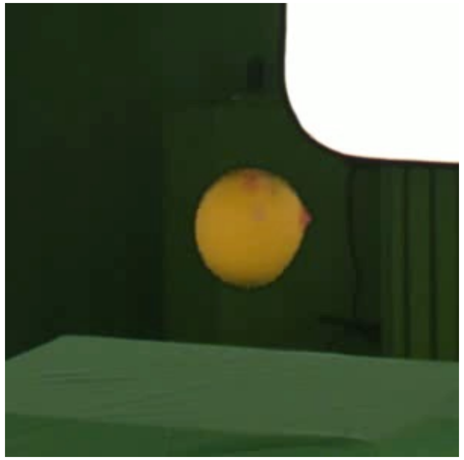

Ours

GIC

DEL

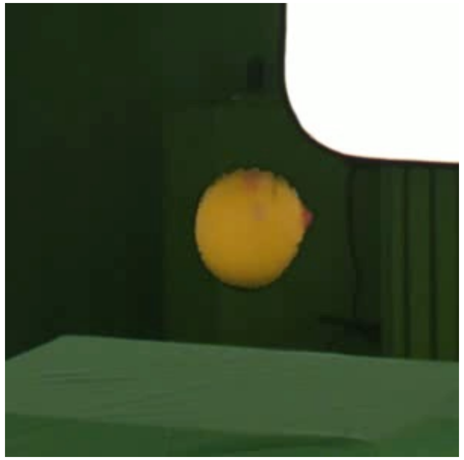

NeuMA

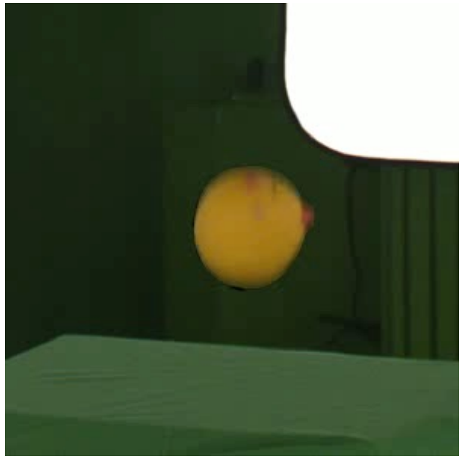

Vid2Sim

# Experiment Results

## Chick2

Groundtruth

Ours

GIC

DEL

NeuMA

Vid2Sim

# Experiment Results

## Gorilla

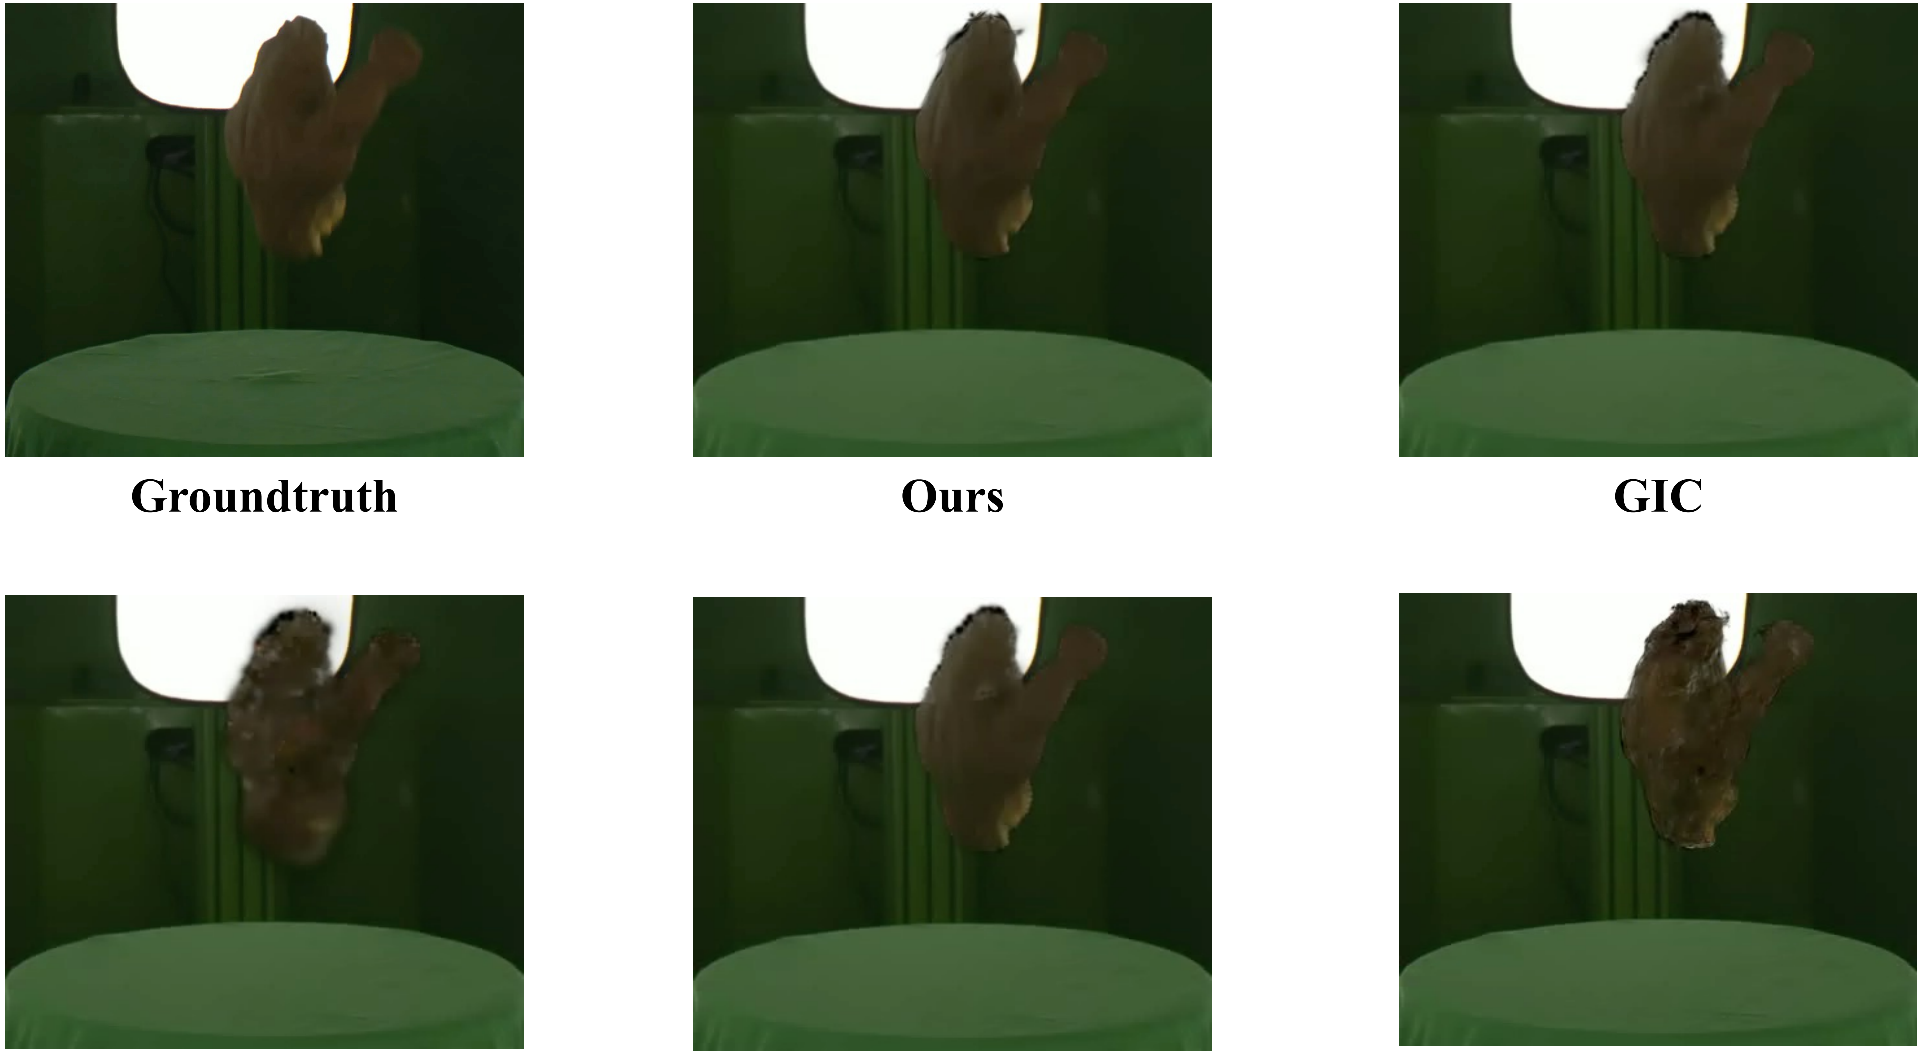

Groundtruth

Ours

GIC

DEL

NeuMA

Vid2Sim

# Experiment Results

## Mandarin

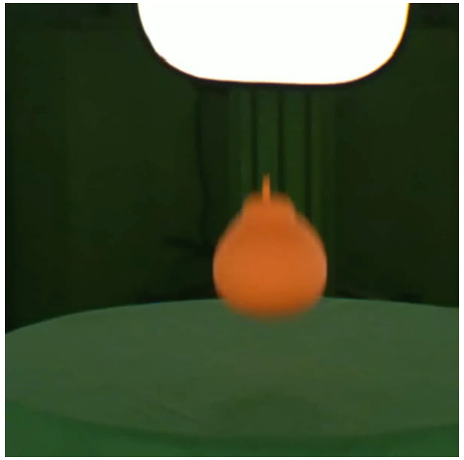

Groundtruth

Ours

GIC

DEL

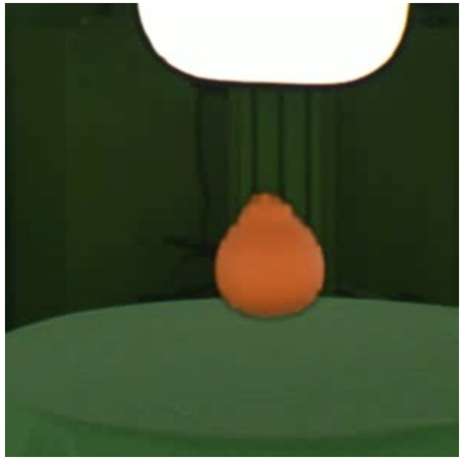

NeuMA

Vid2Sim

# Experiment Results

## Peanut

Groundtruth

Ours

GIC

DEL

NeuMA

Vid2Sim

# Experiment Results

## Rabbit

Groundtruth

Ours

GIC

DEL

NeuMA

Vid2Sim

# Experiment Results

## Rainbowball

Groundtruth

Ours

GIC

DEL

NeuMA

Vid2Sim

# Experiment Results

Prior

Ours