GaussianFlesh: A Total Lagrangian APIC Substrate for Unified Geometry, Material State, and Continuum Simulation

Chukwudalu Dumebi-Kachikwu
Star Labs University of Maryland
daludk@terpmail.umd.edu

Abstract

We present a unified particle substrate in which 3D Gaussian primitives simultaneously encode geometric shape, material anisotropy, and simulation state—removing the representational boundary between scene geometry and physical dynamics. Building on the insight shared with concurrent work that Gaussian covariance can be evolved as $\mathbf{C}' = \mathbf{F}\mathbf{C}\mathbf{F}^\top$ *to reflect local strain, we identify two architectural commitments that distinguish our approach: a Total Lagrangian formulation with a fixed reference-configuration grid, and the Affine Particle-in-Cell (APIC) transfer scheme with additive deformation gradient updates* $\mathbf{F}_{new} = \mathbf{F}_{old} + \Delta t\, \mathbf{C}$*, correct for Total Lagrangian MPM where* $\mathbf{C}$ *approximates* $\partial v/\partial X$ *in the fixed reference frame. Together these choices, combined with per-frame WLS* $\mathbf{F}$ *recomputation, eliminate the numerical drift that we demonstrate empirically accumulates in a PhysGaussian-style UL-MLS-MPM baseline under sustained dynamics. Constitutive behaviour is fully swappable at the material level—corotated elasticity and Neo-Hookean solids share an identical particle structure and neighbourhood topology, differing only in the stress function* $\Psi(\mathbf{F})$*. Corotated elasticity provides implicit shape restoration through the* $(\mathbf{F} - \mathbf{R})$ *stress term, substantially reducing reliance on explicit Procrustes corrections, which are eliminated entirely for stiff materials. A passive triangle mesh, bound to the particle cloud via full-cloud inverse-distance weighting with locally* $\mathbf{U}$*-transformed rest offsets (*$\mathbf{U}$ *the symmetric stretch factor from polar decomposition of* $\mathbf{F}$*), acts as a rendering skin decoupled from the physics. The system runs stably at real-time rates across three material regimes spanning three orders of stiffness magnitude, demonstrated on both sphere and deformed sphere geometries from a single shared substrate.*

1. Introduction

Contemporary physically-grounded scene representations maintain a division of labour: geometry systems encode shape and appearance while physics systems simulate dynamics on separate proxy representations. The coupling between them—skinning, cage embedding, mesh-to-particle transfer—is an engineering afterthought. For applications demanding real-time, physically plausible dynamics—robotics, interactive simulation, embodied world models—this separation is architecturally limiting.

PhysGaussian [8] demonstrated that 3D Gaussian primitives [3] can be enriched with mechanical attributes and evolved under MPM dynamics, establishing $\mathbf{C}' = \mathbf{F}\mathbf{C}\mathbf{F}^\top$ to keep Gaussian shape consistent with local strain. However, these approaches are reconstruction-first: they couple physics to an existing 3DGS scene. The simulation substrate and rendering representation remain conceptually distinct layers.

We propose a substrate-first architecture in which the Gaussian particle is the primary design primitive, simultaneously encoding geometric shape (covariance matrix), material anisotropy (covariance orientation), and simulation state ($\mathbf{F}$, $\mathbf{C}$). Two commitments distinguish us from prior work. First, a Total Lagrangian MPM formulation with a fixed reference-configuration grid, where per-frame WLS recomputation of $\mathbf{F}$ prevents additive drift accumulation. Second, corotated linear elasticity as the default constitutive model, providing implicit shape restoration through the $(\mathbf{F} - \mathbf{R})$ stress term that substantially reduces or eliminates reliance on explicit Procrustes correction. A passive triangle mesh bound via full-cloud IDW with $\mathbf{U}$-transformed rest offsets provides rendering compatibility.

We propose this substrate as a foundational primitive for native 4D world models: a representation in which any object can be instantiated, assigned a material identity, and simulated forward in the same Gaussian space that 4D vision systems use for reconstruction and tracking—closing the loop between scene understanding and physical simulation without representational conversion.

GaussianFlesh is complementary to reconstruction-first approaches like PhysGaussian [8]: those couple physics to captured 3DGS scenes, while we provide a unified substrate from which physically-grounded scenes can be instantiated

and simulated in a single representation.

Our contributions are:

- A Total Lagrangian APIC substrate in which 3D Gaussian covariance simultaneously encodes material anisotropy and simulation state, with $\mathbf{C}' = \mathbf{F}\mathbf{C}\mathbf{F}^\top$ grounding Gaussian shape in physical strain at every frame.
- An empirical demonstration that corotated elasticity substantially reduces Procrustes correction reliance, eliminating it for stiff materials and reducing it to a small residual for high-restitution rubber.
- A passive mesh skin with $\mathbf{U}$-transformed rest offsets and full-cloud IDW binding with mean drift below $0.05\,\mathrm{m}$ across all tested regimes, on both sphere and deformed sphere geometries.
- An empirical comparison against a PhysGaussian-style UL-MLS-MPM baseline implemented under identical conditions, demonstrating that TL-APIC bounds surface drift (12% lower final-frame mean over 3000 frames) while preserving per-particle high-strain state under sustained dynamics.

2. Related Work

Material Point Methods. MPM was introduced to graphics by Stomakhin et al. [6], who coupled a hybrid Eulerian/Lagrangian discretisation with the corotated constitutive model for snow, and demonstrated that constitutive law is a modular parameter [7]. Jiang et al. [2] introduced APIC, conserving linear and angular momentum across transfers; we adopt APIC and use its affine matrix $\mathbf{C}$ as the material-frame deformation rate. Hu et al. [1] introduced MLS-MPM with a multiplicative $\mathbf{F}$ update $\mathbf{F}_{\text{new}} = (\mathbf{I} + \Delta t\, \nabla \mathbf{v})\, \mathbf{F}_{\text{old}}$, used by PhysGaussian [8]. Our Total Lagrangian formulation uses an additive $\mathbf{F}$ update against a fixed reference grid; §4.4 demonstrates empirically that this design bounds drift under sustained dynamics where the multiplicative scheme accumulates measurable drift.

Gaussian-physics coupling and mesh skinning. 3DGS [3] represents scenes as anisotropic Gaussian primitives rendered via tile-based rasterisation, extended to dynamic settings by Luiten et al. [4] through per-frame optimisation. PhysGaussian [8] established the covariance update $\mathbf{C}' = \mathbf{F}\mathbf{C}\mathbf{F}^\top$ for Gaussian-physics coupling from reconstructed scenes. GaussianFlesh is substrate-first rather than reconstruction-first. For mesh binding we follow McAdams et al. [5] but apply only the symmetric stretch $\mathbf{U}$ to rest offsets, avoiding double-counting the rotational component already captured by particle positions.

3. Method

GaussianFlesh represents a deformable body as a cloud of N Gaussian particles. Each particle p carries: position $\mathbf{x}_p$, rest position $\mathbf{X}_p$, velocity $\mathbf{v}_p$, affine velocity matrix $\mathbf{C}_p \in \mathbb{R}^{3\times3}$, deformation gradient $\mathbf{F}_p \in \mathbb{R}^{3\times3}$, and covariance $\mathbf{\Sigma}_p \in \mathbb{R}^{3\times3}$.

3.1. Gaussian Particle Representation

Rest positions $\mathbf{X}_p$ are sampled uniformly inside the target volume via rejection sampling. Each particle has reference volume $V_0 = \text{total volume}/N$. The rest covariance is a flat disc aligned to the local surface normal:

$$\mathbf{\Sigma}_{0,p} = \mathbf{R}_p\, \mathrm{diag}(\sigma_r,\, \sigma_t,\, \sigma_t)\mathbf{R}_p^\top, \quad \sigma_r \ll \sigma_t,$$

encoding surface geometry as material anisotropy. Each frame, covariance is updated as the push-forward $\mathbf{\Sigma}_p = \mathbf{F}_p\, \mathbf{\Sigma}_{0,p}\, \mathbf{F}_p^\top$ [8], making each Gaussian a visual strain gauge.

3.2. Total Lagrangian APIC Solver

We use MPM in Total Lagrangian (TL) form: shape functions and gradients are computed once from the rest configuration Ω_0 and never rebuilt. Internal forces use the first Piola-Kirchhoff stress $\mathbf{P}$.

Grid construction. A background grid (spacing $h = r/4$) is precomputed from rest positions. For each particle p and each of the 27 stencil nodes i we store: B-spline weight w_{ip}, reference-frame gradient $\nabla_{\mathbf{X}} w_{ip}$, and rest-frame displacement $\delta\mathbf{X}_{ip} = \mathbf{X}_i - \mathbf{X}_p$. All quantities are constant for the simulation duration.

APIC transfer and F update. Each substep proceeds via APIC [2]. P2G scatter:

$$\mathbf{p}_i = \sum_p w_{ip}\, m_p \big(\mathbf{v}_p + \mathbf{C}_p\, \delta\mathbf{X}_{ip}\big), \tag{1}$$

$$\mathbf{f}_i = -V_0 \sum_p \mathbf{P}_p\, \nabla_{\mathbf{X}} w_{ip}. \tag{2}$$

G2P gather after grid update:

$$\mathbf{v}_p^- = \textstyle\sum_i w_{ip}\, \mathbf{v}_i^g, \quad \mathbf{C}_p^- = \frac{4}{h^2} \sum_i w_{ip}\, \mathbf{v}_i^g \otimes \delta\mathbf{X}_{ip}. \tag{3}$$

The deformation gradient updates additively:

$$\mathbf{F}_{\text{new}} = \mathbf{F}_{\text{old}} + \Delta t\, \mathbf{C}. \tag{4}$$

This is correct for TL-MPM where $\mathbf{C} \approx \partial v/\partial X$ in the fixed reference frame. Multiplicative updates appropriate for Updated Lagrangian MPM cause $\mathbf{F}$ to blow up in the TL setting. To prevent long-term drift, $\mathbf{F}$ is recomputed once per display frame via WLS over actual particle positions. SVD clamping clips singular values of $\mathbf{F}$ to $[0.2, 5.0]$; stress magnitude is capped at $\tau = 100$ before P2G scatter.

3.3. Constitutive Laws

$\Psi(\mathbf{F})$ is the sole material-specific component; the particle struct, grid, and mesh binding are unchanged across materials.

Corotated linear elasticity (all solid materials). $\Psi = \mu\|\tilde{\mathbf{F}} - \mathbf{R}\|^2 + (\lambda/2)(J-1)^2$, with $\mathbf{R} = \mathbf{U}\mathbf{V}^\top$ from SVD, where $\tilde{\mathbf{F}}$ denotes the elastic part of the deformation gradient ($\tilde{\mathbf{F}} = \mathbf{F}$ for the purely elastic materials studied here). PK1 stress:

$$\mathbf{P} = 2\mu(\tilde{\mathbf{F}} - \mathbf{R}) + \lambda(J-1)\,J\,\tilde{\mathbf{F}}^{-\top}. \tag{5}$$

The $(\tilde{\mathbf{F}} - \mathbf{R})$ term generates restoring stress proportional to μ, providing implicit shape restoration. High-μ materials (metal, jelly) require no explicit Procrustes correction; rubber uses a small residual (shape_alpha $= 0.02$).

3.4. Passive Mesh Skin

A triangle mesh with V vertices is bound to all N particles via inverse-distance weighting. The rest offset $\delta\mathbf{v} = \mathbf{x}_v - \sum_k w_{vk}\,\mathbf{X}_{p_k}$ is stored at initialisation. Each frame:

$$\mathbf{x}_v = \sum_k w_{vk}\,\mathbf{x}_{p_k} + \mathbf{U}_v^{\text{local}}\,\delta\mathbf{v}, \tag{6}$$

where $\mathbf{U}_v^{\text{local}} = \sum_k w_{vk}\,\mathbf{U}_{p_k}$ is the weighted-average symmetric stretch. Applying $\mathbf{U}$ rather than $\mathbf{F}$ avoids double-counting the rotational component already captured by particle positions. The mesh skin is a compatibility layer with surface-based rendering pipelines; the particle substrate is architecturally independent of it, and native Gaussian splatting integration can omit the skinning layer entirely. Ground collision uses a penalty spring with damping coefficient $c = -2\sqrt{k_{\text{floor}}\,m}\ln r/\sqrt{\pi^2 + \ln^2 r}$ derived from desired restitution r.

4. Results

We evaluate on a unit sphere ($N = 120$ particles) and a deformed sphere ($N = 120$ particles), each dropped from 4.0 m. Three material regimes are tested: rubber ($\mu = 40$, $\lambda = 20$), metal ($\mu = 300$, $\lambda = 1500$), and jelly ($\mu = 1$, $\lambda = 5$). All parameters other than the material preset are identical across materials and geometries. The physics hot path runs via Taichi CUDA kernels on an RTX 4070 at real-time rates (42–72 ms/frame across all materials).

4.1. Qualitative Material Behaviour

Figure 1 shows the three materials at peak squash. The same 120-particle substrate produces qualitatively distinct behaviours through constitutive law swap alone.

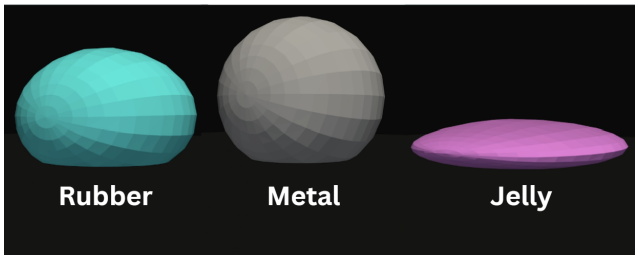

Figure 1. Three material regimes from a single particle substrate. Constitutive law is the only difference between panels.

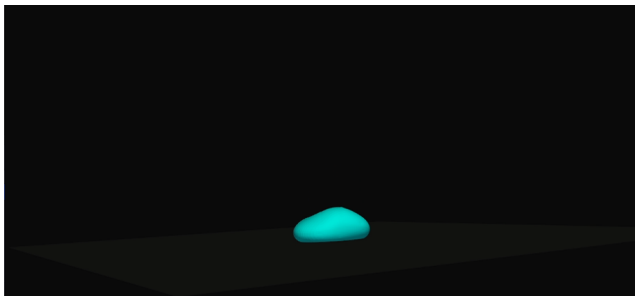

Figure 2. Deformed sphere at peak squash. The substrate and binding protocol operate unchanged across geometry types.

Rubber ($\mu = 40$). Compresses to $U_{yy,\min} = 0.561$ of rest height and rebounds elastically (one bounce, 260 frames). Contact-face ellipsoids flatten ($\sigma_{\min}$ mean $= 0.854$) while top-hemisphere ellipsoids remain near-spherical, as shown in Figure 3.

Metal ($\mu = 300$). Barely deforms visibly (U_{yy} spread $= 0.349$); the $(\mathbf{F} - \mathbf{R})$ term at $\mu = 300$ provides sufficient implicit rigidity without explicit constraint ($\sigma_{\min}$ mean $= 0.948$).

Jelly ($\mu = 1$). Reaches U_{yy} spread $= 1.482$—the largest of the three—and spreads laterally before slowly recovering with residual oscillation. SVD clamping activates for 6 particles at peak.

Deformed sphere geometry. Figure 2 shows rubber on the deformed sphere geometry. The same substrate and binding protocol operate unchanged on this topology (peak drift mean 0.046 m).

4.2. Quantitative Results

Table 1 reports mesh-to-isosurface drift and deformation diagnostics at peak squash. All materials share rest-state drift (mean 0.079 m) by construction. Jelly shows the highest peak drift (0.108 m), consistent with its extreme lateral spread. $\sigma_{\min}$ and U_{yy} spread are both zero at rest for all materials, confirming no spurious pre-drop deformation.

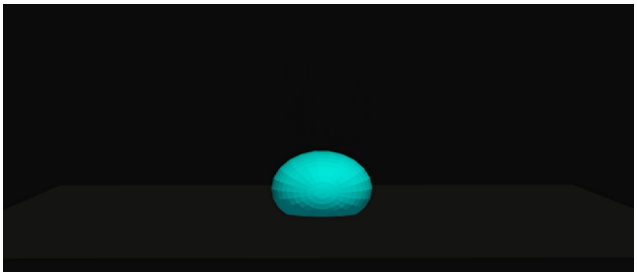

Figure 3. Covariance ellipsoids at peak squash (U_{yy} spread $=$ 0.611). Contact-face ellipsoids flatten ($\sigma_{\min}$ min $=$ 0.237); top ellipsoids retain rest shape.

Table 1. Peak-squash diagnostics (sphere geometry). Drift: mesh-to-isosurface mean/max (m). Spread: U_{yy} spread. $\bar{\sigma}$: $\sigma_{\min}$ mean.

Material	Drift mean	Drift max	Spread	$\bar{\sigma}$	Clamped
Rubber	0.083	0.202	0.611	0.854	0
Metal	0.080	0.191	0.349	0.948	0
Jelly	0.108	0.274	1.482	0.568	6

Table 2. Procrustes ablation for rubber, 800 frames.

α	Drift mean (m)	Drift max (m)	U_{yy} spread
0.00	0.070	0.215	1.978
0.02	0.083	0.202	0.611
0.04	0.082	0.194	0.566

4.3. Corotated Elasticity and Procrustes Correction

Table 2 ablates Procrustes correction strength for rubber over 800 frames. With shape_alpha $= 0$, the $(\mathbf{F} - \mathbf{R})$ stress term alone is insufficient at high restitution ($r = 0.85$), yielding U_{yy} spread of 1.978. A small correction ($\alpha = 0.02$) reduces this to 0.611 with negligible drift cost. Metal and jelly require no correction ($\alpha = 0$) due to higher μ.

4.4. TL vs UL Solver Comparison

To isolate the contribution of the Total Lagrangian formulation we implement a PhysGaussian-style Updated Lagrangian MLS-MPM baseline within the same Taichi pipeline. The baseline follows Xie et al. [8] and Hu et al. [1]: a multiplicative deformation gradient update $\mathbf{F}_{\text{new}} = (\mathbf{I} + \Delta t\, \nabla \mathbf{v})\, \mathbf{F}_{\text{old}}$, a world-frame background grid rebuilt every substep against current particle positions, and Kirchhoff stress $\boldsymbol{\tau} = \mathbf{P}\mathbf{F}^{\top}$ scattered with world-frame gradients $\nabla_{\mathbf{x}} w$. The particle struct, covariance update, mesh skin, boundary conditions, gravity, restitution, timestep, and material parameters are identical to the TL configuration; only the solver internals differ. Two rubber spheres ($N = 120$, $\mu = 40$, $\lambda = 20$) are dropped side-by-side from 4.0 m and simulated for 3000 frames (~ 50 s of sim time).

Through the first $\sim$2500 frames the two solvers track closely (within 1–2% on drift mean), but by the final frame

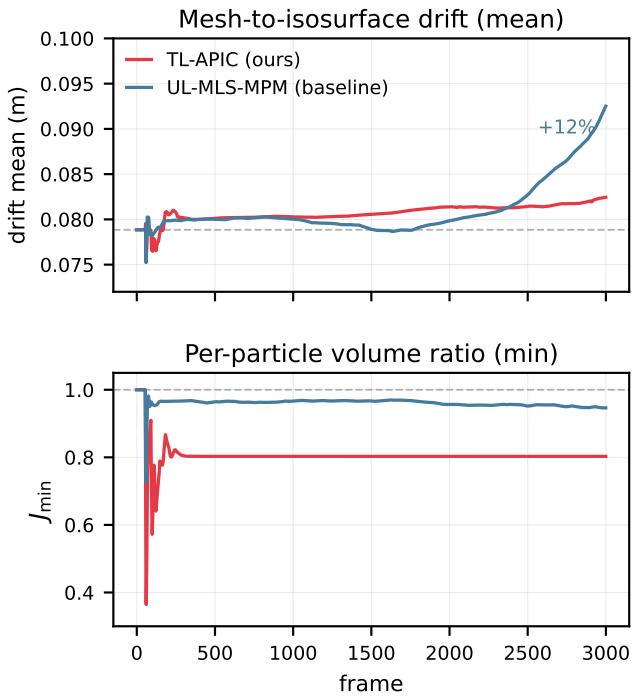

Figure 4. Drift trajectory over 3000 frames. Top: mesh-to-isosurface drift mean. TL (ours, red) holds near the rest-state baseline (dashed line at 0.079 m); the UL-MLS-MPM baseline (blue) tracks closely through $\sim$frame 2500 then accumulates drift, ending 12% higher. Bottom: per-particle $J_{\min}$ trajectory (dashed reference at $J = 1$). TL settles to a stable fixed point at $J_{\min} = 0.803$; UL's $J_{\min}$ hovers closer to unity but drifts steadily downward over the run, mirroring the upward drift in surface drift mean.

Table 3. TL vs UL-MLS-MPM solver comparison over 3000 frames, rubber sphere. Δ shows TL improvement (positive $=$ TL better).

Metric	TL (ours)	UL baseline	Δ
Drift mean, final frame	0.0824	0.0925	$+12\%$
Drift max, worst frame	0.206	0.229	$+11\%$
Drift mean, average	0.0807	0.0810	tie
$J_{\min}$ mean	0.80	0.96	—
U_{yy} spread mean	0.17	0.07	—

UL's drift mean has risen to 0.0925 m versus TL's 0.0824 m, a 12% gap consistent with multiplicative-$\mathbf{F}$ error accumulating under sustained dynamics. TL's per-particle volume diagnostic exhibits a striking fixed-point behaviour: post-impact, the most-compressed contact particle settles to $J_{\min} = 0.803$ and holds that value to five-decimal precision across the remaining 2700 frames, while UL's $J_{\min}$ drifts steadily from 0.97 to 0.94 over the same interval. This is the multiplicative-$\mathbf{F}$ volume accumulation predicted by Total Lagrangian theory: per-frame WLS recomputation anchors the deformation gradient to actual particle geometry,

while the multiplicative update loses volume through compounded small errors. The two solvers thus differ not only in surface-drift magnitude but in mechanism: TL captures and preserves high-strain contact state; UL smooths it during impact and then accumulates volume drift across the bulk.

Figure 5 extends the qualitative demonstration of §4.1 to a non-spherical, feature-bearing geometry: a rubber duck under the same drop conditions and the same constitutive law as the sphere experiments.

5. Conclusion

We have presented and empirically validated GaussianFlesh, a unified substrate in which a single Gaussian particle simultaneously serves as geometry proxy, material state encoding, and continuum mechanics integration point. Constitutive law is a hot-swap parameter; rubber, metal, and jelly emerge from identical particle structure. The empirical comparison in §4.4 demonstrates that this unified substrate admits bounded drift under sustained dynamics, validated against a faithful PhysGaussian-style UL-MLS-MPM baseline.

The Total Lagrangian formulation is strongest for elastic and small-to-moderate plastic deformation around a well-defined rest configuration. It is less suited to problems involving topology change, fracture, or large plastic flow, where the fixed reference grid loses physical meaning. PhysGaussian's UL-MLS-MPM formulation [8] accommodates these regimes naturally through plasticity return maps and the deforming grid; fracture, granular flow, and viscoplastic materials remain outside GaussianFlesh's current scope.

A natural extension is a per-particle solver dispatch driven by material identity: particles with elastic constitutive laws (corotated, Neo-Hookean) integrate via TL-APIC, while particles undergoing plasticity (Drucker-Prager, Herschel-Bulkley) fall back to UL-MLS-MPM. The two solvers share the same particle struct, covariance update $\mathbf{C}' = \mathbf{F}\mathbf{C}\mathbf{F}^\top$, and mesh binding; only the $\mathbf{F}$ update rule and grid frame differ. Systematic exploration of this hybrid dispatch — and of mesh-free rendering directly off the simulated Gaussians — is left to future work.

Future work will also couple a VLM material inference pipeline to existing 3DGS scenes, enabling physics to be assigned from visual observation.

References

[1] Yuanming Hu, Yu Fang, Ziheng Ge, Ziyin Qu, Yixin Zhu, Andre Pradhana, and Chenfanfu Jiang. A moving least squares material point method with displacement discontinuity and two-way rigid body coupling. *ACM TOG*, 37(4), 2018. 2, 4

[2] Chenfanfu Jiang, Craig Schroeder, Andrew Selle, Joseph Teran, and Alexey Stomakhin. The affine particle-in-cell method. *ACM TOG*, 34(4), 2015. 2

[3] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3D Gaussian splatting for real-time radiance field rendering. *ACM TOG*, 42(4), 2023. 1, 2

[4] Jonathon Luiten, Georgios Kopanas, Bastian Leibe, and Deva Ramanan. Dynamic 3D Gaussians: Tracking by persistent dynamic view synthesis. In *International Conference on 3D Vision (3DV)*, 2024. 2

[5] Aleka McAdams, Yongning Zhu, Andrew Selle, Mark Empey, Rasmus Tamstorf, Joseph Teran, and Eftychios Sifakis. Efficient elasticity for character skinning with contact and collisions. *ACM TOG*, 30(4), 2011. 2

[6] Alexey Stomakhin, Craig Schroeder, Lawrence Chai, Joseph Teran, and Andrew Selle. A material point method for snow simulation. *ACM TOG*, 32(4), 2013. 2

[7] Alexey Stomakhin, Craig Schroeder, Chenfanfu Jiang, Lawrence Chai, Joseph Teran, and Andrew Selle. Augmented MPM for phase-change and varied materials. *ACM TOG*, 33 (4), 2014. 2

[8] Tianyi Xie, Zeshun Zong, Yuxing Qiu, Xuan Li, Yutao Feng, Yin Yang, and Chenfanfu Jiang. PhysGaussian: Physics-integrated 3D Gaussians for generative dynamics. In *CVPR*, 2024. 1, 2, 4, 5

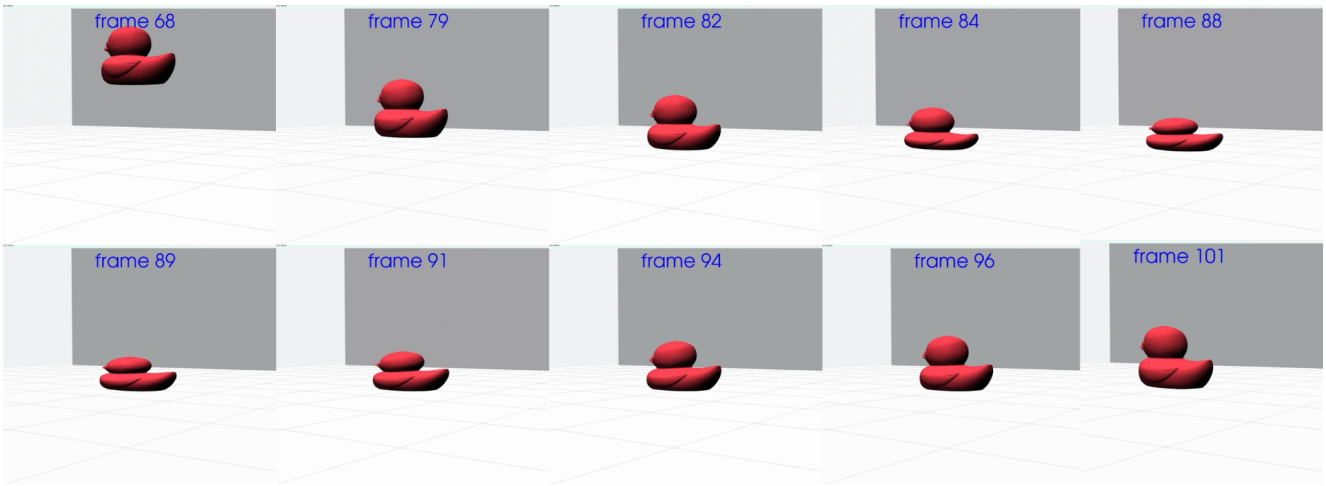

Figure 5. Rubber duck under gravity-driven impact ($\mu = 40$, $\lambda = 20$, identical solver and parameters to the sphere experiments in §4.1). The substrate handles non-spherical, feature-bearing geometry without modification: head and body deform under a single shared constitutive law with no geometry-specific tuning.