

# Real-Time Physics Simulation with Dynamic Mesh-Gaussian Reconstructions

Adrian Ramlal  
University of Waterloo  
adrian.ramlal@uwaterloo.ca

John S. Zelek  
University of Waterloo  
jzelek@uwaterloo.ca

## Abstract

*Integrating dynamic 3D reconstructions into physics simulation requires fixed mesh topology for efficient collision detection, but state-of-the-art methods like DG-Mesh produce varying topology optimized for geometric quality. We investigate whether topology conversion can enable physics integration while preserving reconstruction fidelity. We propose a dual-representation framework combining fixed-topology meshes for physics with Gaussian splatting for rendering, achieving  $4.65\times$  speedup over varying-topology baselines through runtime vertex buffer updates. We evaluate two conversion strategies, temporal correspondence tracking and template-based projection, against native fixed-topology methods (MaGS) on the DG-Mesh dataset. Our evaluation reveals that both conversion approaches incur 65–80% geometric degradation, producing results inferior to MaGS despite DG-Mesh’s superior initial quality. This demonstrates that high-quality reconstruction and physics-compatible topology represent fundamentally distinct objectives that cannot be reconciled through post-processing. Our findings inform future development of physics-aware reconstruction methods and our framework enables real-time simulation with any fixed-topology approach.*

## 1. Introduction

Physics simulation with photorealistic rendering is essential for robotics [21], virtual reality [16], and reinforcement learning [1]. Training robots in simulation requires realistic collision dynamics and visual feedback to enable effective sim-to-real transfer. Similarly, embodied AI agents need physically accurate environments to develop transferable behaviors. Recent advances in neural rendering, particularly 3D Gaussian Splatting (3DGS) [8] and its dynamic extensions [9, 12], enable photorealistic reconstruction of dynamic scenes from multi-view video with explicit mesh geometry. However, integrating them into physics engines faces a critical barrier: *mesh topology consistency*.

Modern physics engines require *fixed topology* (consistent vertex count and face connectivity) to maintain effi-

cient collision detection structures. Varying topology necessitates expensive per-frame rebuilds of spatial acceleration structures, preventing real-time simulation.

This exposes a fundamental trade-off in current methods. DG-Mesh [9] achieves state-of-the-art quality via iso-surface extraction but produces *varying topology*. MaGS [12] maintains *fixed topology* by jointly optimizing a template mesh with Gaussians, enabling physics simulation at the cost of fidelity. Can high-quality varying-topology reconstructions be converted to fixed topology while preserving geometric quality? Bridging this gap would enable leveraging the highest-quality reconstructions for physics simulation in robotics and interactive applications.

We investigate this question through systematic evaluation in the Genesis simulator [1]. Our contributions are:

- Dual-Representation Framework.** We develop an architecture combining fixed-topology meshes for physics with Gaussians for rendering. Runtime vertex buffer updates achieve 147.8 FPS versus 31.8 FPS for varying topology ( $4.65\times$  speedup), enabling efficient integration of any fixed-topology method.
- Topology Conversion Evaluation.** We propose two conversion strategies (temporal tracking and template projection) and provide the first systematic comparison against native fixed-topology methods on the DG-Mesh dataset [9], measuring geometric accuracy across six dynamic scenes.

Our experiments reveal that topology conversion incurs 65–80% geometric degradation, making both strategies inferior to native fixed-topology methods. This demonstrates that high-quality reconstruction and physics-compatible topology are fundamentally distinct objectives. Our framework provides valuable infrastructure for fixed-topology methods and our evaluation informs future physics-aware reconstruction approaches.

## 2. Related Works

### 2.1. Dynamic Neural Rendering

Neural Radiance Fields (NeRF) [13] and extensions [4, 15, 17] model dynamic scenes with volumetric represen-

tations but lack explicit geometry. 3D Gaussian Splatting (3DGS) [8] provides explicit point-based rendering with real-time performance. Dynamic 3DGS methods [11, 23, 25] learn deformation networks to model temporal motion, with controllable variants like SC-GS [7] using sparse control points for editing. While achieving high rendering quality, these point-based methods produce meshes with inconsistent topology across frames, preventing direct physics integration.

## 2.2. Mesh-Gaussian Hybrid Representations

Recent work combines meshes with 3DGS for structured geometry. SuGaR [5] regularizes Gaussians to surfaces and extracts meshes via Poisson reconstruction. GaMeS [22] introduces mesh-Gaussian binding for topology-preserving edits through triangle soups. Domain-specific methods leverage parametric models: GaussianAvatars [18] uses FLAME for animatable heads, while PGC [6] and Gaussian Garments [19] bind Gaussians to cloth meshes for physics-based simulation.

For general dynamic scenes, DG-Mesh [9] achieves state-of-the-art geometric quality through differentiable Poisson reconstruction and Marching Cubes, with Gaussian-Mesh Anchoring enforcing one-to-one Gaussian-face correspondence. However, the iso-surface extraction produces varying topology across frames. In contrast, MaGS [12] maintains fixed topology by jointly optimizing a template mesh with mesh-adsorbed Gaussians, incorporating deformation priors like ARAP [20] and SMPL [10]. This prioritizes topological stability and physics compatibility over geometric fidelity. Our work provides the first systematic comparison of these approaches for physics simulation.

## 2.3. Physics Integration with Neural Representations

Physics engines require fixed topology for efficient collision detection through spatial acceleration structures (BVH trees, SDFs). Traditional approaches use explicit meshes with simulators like MuJoCo [21], Bullet [3], or Genesis [1]. Recent neural-physics integration includes Phys-Gaussian [24] with mass-spring dynamics on Gaussian particles, GASP [2] for soft-body simulation on point clouds, and Splatting Physical Scenes [14] optimizing coarse meshes through differentiable simulation. However, these works either avoid mesh extraction or use simplified static geometry, leaving a gap in integrating high-fidelity time-varying meshes.

Our framework enables systematic evaluation of topology strategies for physics simulation in Genesis [1]. We investigate whether topology conversion can preserve geometric quality while enabling physics integration, and provide infrastructure for real-time simulation with fixed-

topology methods through a dual-representation architecture that decouples physics meshes from photorealistic rendering.

## 3. Method

Our method enables integration of state-of-the-art dynamic reconstruction techniques into real-time physics simulation through a dual-representation framework that decouples physics-compatible mesh topology from high-fidelity rendering.

### 3.1. Background: Dynamic Reconstruction

We build upon two complementary approaches representing different points in the quality-topology trade-off space.

**DG-Mesh** [9] achieves state-of-the-art reconstruction quality (CD = 0.697) by representing scenes using canonical 3D Gaussians deformed via learned temporal networks. Geometry extraction via Differentiable Poisson Surface Reconstruction and Marching Cubes produces high-quality meshes with *varying topology*: both vertex count  $|\mathbf{V}_t|$  and connectivity  $\mathbf{F}_t$  change across frames, preventing direct physics integration.

**MaGS** [12] maintains *fixed topology* by jointly optimizing a template mesh and Gaussian primitives. Gaussians are adsorbed to mesh faces and displaced via learned networks, while mesh vertices deform through a deformation network. The fixed template topology enables direct physics integration but achieves lower geometric fidelity (CD = 1.108) than DG-Mesh.

DG-Mesh’s superior geometric quality (36% lower CD) motivates our investigation of topology conversion strategies to enable physics integration while preserving reconstruction fidelity.

### 3.2. Physics-Rendering Dual Representation

Our framework maintains two synchronized representations: a **physics mesh** ( $\mathcal{M}_{\text{phys}}$ ) with fixed topology for collision detection, and a **rendering representation** ( $\mathcal{R}_{\text{render}}$ ) using high-fidelity Gaussian splatting for photorealistic visualization. This decoupling enables optimal performance for both physics and rendering while supporting comparative analysis of different reconstruction methods.

### 3.3. Fixed-Topology Mesh Integration

**MaGS: Native Fixed-Topology.** MaGS meshes directly provide fixed topology. At simulation time  $t$ , we load corresponding deformed vertices:  $k = \lfloor t \cdot f_{\text{recon}} \rfloor$  where  $f_{\text{recon}}$  is the reconstruction framerate.

**DG-Mesh: Topology Conversion Strategies.** To enable physics integration of DG-Mesh’s varying-topology sequences  $\{\mathcal{M}_t = (\mathbf{V}_t, \mathbf{F}_t)\}$ , we explore two conversion strategies to establish fixed correspondences.

**Strategy 1: Temporal Correspondence Tracking.** Given an initial template  $\mathcal{M}_0 = (\mathbf{V}_0, \mathbf{F}_0)$ , we track vertices forward through time:

$$\tilde{\mathbf{V}}_{t+1} = \{\text{NearestSurface}(\mathbf{v}_i^t, \mathcal{M}_{t+1})\}_{i=1}^{N_v} \quad (1)$$

where  $\text{NearestSurface}(\mathbf{v}, \mathcal{M})$  projects point  $\mathbf{v}$  onto mesh  $\mathcal{M}$  via k-NN weighted surface projection. To mitigate drift accumulation, we periodically re-anchor to the template every 50 frames by blending 80% tracked motion with 20% direct template projection.

**Strategy 2: Template-Based Projection.** We independently project a static template onto each frame:

$$\tilde{\mathbf{V}}_t = \{\text{NearestSurface}(\mathbf{v}_i^{(0)}, \mathcal{M}_t)\}_{i=1}^{N_v} \quad (2)$$

using barycentric interpolation for sub-face accuracy. This eliminates drift but may lose detail in regions with varying geometric complexity.

### 3.4. Efficient Runtime Updates

Our key technical contribution enables time-varying mesh geometry without rebuilding collision structures.

**Initialization.** We construct spatial acceleration structures (BVH, SDF) from template topology  $\mathbf{F}_0$ . These structures store only vertex and face *indices*, enabling reuse across all timesteps.

**Per-Frame Update.** At each simulation step:

1. Compute frame index  $k = \lfloor t \cdot f_{\text{recon}} \rfloor$  and load vertices  $\mathbf{V}_k$
2. Update vertex buffer via  $O(N_v)$  memory copy
3. Recompute per-face normals on-demand during collision queries

This approach scales as  $O(N_v)$  per frame versus  $O(N_v \log N_v)$  for BVH reconstruction, enabling real-time rates for typical meshes ( $N_v \approx 10^3\text{--}10^4$ ).

## 4. Experiments

We evaluate our framework on the DG-Mesh dataset [9], which provides ground truth mesh sequences for six dynamic scenes. Each sequence contains 200 frames with varying-topology meshes extracted from multi-view captures.

### 4.1. Baseline: Varying-Topology Meshes

To establish the computational necessity of fixed topology, we implement a baseline that uses DG-Mesh’s original varying-topology meshes directly in physics simulation. At each frame, this baseline must destroy and rebuild all collision structures (BVH trees, contact caches) to accommodate topology changes.

We measure performance on a test scene containing the *horse* sequence with 5 rigid spheres falling onto the deforming mesh. The varying-topology baseline achieves 31.8

Table 1. Geometric reconstruction quality on DG-Mesh dataset averaged over 6 scenes. Lower is better. Our topology conversion strategies result in significant degradation compared to both original DG-Mesh and MaGS native output.

Method	CD ↓	EMD ↓
DG-Mesh (Original)	<b>0.697</b>	0.130
Temporal Tracking (Ours)	1.152	0.185
Template Projection (Ours)	1.254	0.200
MaGS (Native)	1.108	<b>0.113</b>

FPS, with collision rebuild overhead consuming 23.8 ms per frame (75.6% of total frame time). While this provides the highest geometric quality (CD = 0.697), the performance is insufficient for real-time applications. In contrast, fixed-topology approaches achieve 147.8 FPS by eliminating rebuild overhead through vertex buffer updates, demonstrating a 4.65× speedup.

### 4.2. Geometric Accuracy Analysis

We compare four approaches: DG-Mesh original (varying topology), our two topology conversion strategies (temporal tracking and template projection), and MaGS native (fixed topology). Like DG-Mesh [9] and MaGS [12], we measure quality using Chamfer Distance (CD, in  $10^{-3}$  units) and Earth Mover’s Distance (EMD, in  $10^{-1}$  units) by sampling 10K points from reconstructed and ground truth meshes.

Table 1 shows that DG-Mesh’s original varying-topology meshes achieve the best accuracy (CD = 0.697, EMD = 0.130). MaGS achieves CD = 1.108 and EMD = 0.113. While MaGS has higher CD, it achieves lower EMD, indicating better global shape preservation consistent with its fixed-topology design.

Our topology conversion strategies result in substantial degradation. Temporal tracking achieves CD = 1.152 and EMD = 0.185 (65% and 42% worse than original DG-Mesh). Template projection performs worse with CD = 1.254 and EMD = 0.200 (80% and 54% degradation).

Critically, both conversion strategies produce geometrically inferior results compared to MaGS native output. Temporal tracking achieves 4% worse CD and 64% worse EMD than MaGS, while template projection is 13% worse in CD and 77% worse in EMD. This demonstrates a fundamental limitation: converting high-quality varying-topology reconstruction to fixed-topology incurs geometric losses that negate the original quality advantage.

### 4.3. Qualitative Analysis

Figure 1 visualizes geometric quality at frame 20 of the *horse* sequence. The original DG-Mesh mesh exhibits crisp surface details and accurate thin structures. Our temporal tracking conversion introduces noticeable smoothing and

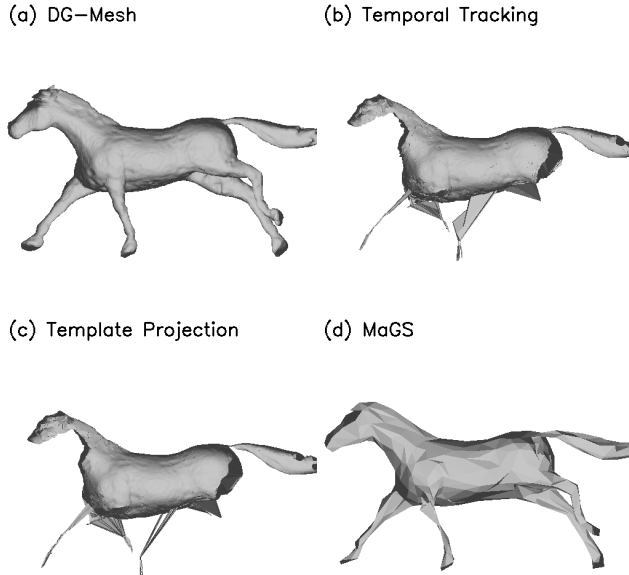


Figure 1. Qualitative comparison at frame 20 on the *horse* sequence. From left to right: (a) DG-Mesh original varying-topology mesh, (b) our temporal tracking conversion, (c) our template projection conversion, (d) MaGS native fixed-topology. Both conversion strategies introduce visible artifacts compared to the original DG-Mesh quality. MaGS maintains better geometric fidelity despite having native fixed topology.

vertex drift, particularly visible in the neck. Template projection shows more severe degradation with over-smoothed surfaces and loss of fine features. In comparison, MaGS native output maintains substantially better geometric fidelity than either conversion strategy.

The temporal tracking method suffers from accumulated drift over the 200-frame sequence despite periodic re-anchoring, while template projection loses detail by forcing a fixed vertex distribution onto varying geometric complexity. These results reveal that the topology conversion process fundamentally compromises the geometric quality that made DG-Mesh attractive for this application.

## 5. Discussion

Our experiments reveal a critical finding: converting varying-topology reconstructions to fixed topology for physics simulation fundamentally degrades geometric quality (65–80% increase in CD, 42–54% increase in EMD). Both our conversion strategies incur substantial geometric losses that eliminate DG-Mesh’s quality advantage and produce results inferior to MaGS’s native fixed-topology method.

This finding demonstrates that high-quality reconstruction and physics-compatible topology represent *distinct de-*

*sign objectives* that cannot be easily bridged through post-processing. DG-Mesh optimizes for geometric fidelity without topological constraints, while MaGS balances quality with topological stability from the outset. The underlying cause is that topology conversion requires establishing correspondences between meshes with different geometric complexity distributions. Temporal tracking accumulates drift as errors compound over frames, while template projection forces a fixed vertex distribution onto varying geometric features.

Despite these conversion limitations, our work provides valuable contributions. The dual-representation architecture and runtime vertex buffer update mechanism achieve a 4.65× performance improvement, enabling real-time physics simulation at 147.8 FPS. This infrastructure supports any fixed-topology dynamic reconstruction method. Our systematic evaluation provides the first direct comparison between varying-topology conversion strategies and native fixed-topology approaches in physics simulation contexts, offering concrete evidence that informs future research directions.

**Limitations.** Our topology conversion strategies assume smooth deformations without topological events (self-contact, tears). Temporal tracking accumulates drift beyond 200 frames, while template projection cannot adapt vertex density to local complexity. Our evaluation focuses on geometric accuracy (CD, EMD) rather than physics-specific metrics like collision fidelity or task performance. The framework is tested only with rigid body physics; soft-body dynamics and deformable materials remain unexplored. Fundamentally, fixed-topology representations cannot capture phenomena requiring topology changes.

## 6. Conclusion

We identified mesh topology consistency as the critical barrier preventing integration of state-of-the-art dynamic 3D reconstructions into real-time physics simulation. We demonstrated that varying-topology meshes impose a 4.65× performance penalty due to per-frame collision structure rebuilds. We explored two topology conversion strategies to enable physics integration of high-quality reconstructions, revealing that both approaches incur 65–80% geometric degradation and produce results inferior to native fixed-topology methods like MaGS. This demonstrates that high-quality reconstruction and physics-compatible topology are fundamentally distinct objectives that cannot be easily reconciled through post-processing. Our dual-representation framework provides valuable infrastructure for integrating fixed-topology methods into physics simulation, achieving 147.8 FPS with runtime vertex buffer updates. Future work should focus on developing reconstruction methods that directly optimize for high-fidelity, fixed-topology constraints during training rather than attempting post-hoc conversion.

## References

- [1] Genesis Authors. Genesis: A generative and universal physics engine for robotics and beyond, 2024. 1, 2
- [2] Piotr Borycki, Weronika Smolak, Joanna Waczyńska, Marcin Mazur, Sławomir Tadeja, and Przemysław Spurek. Gasp: Gaussian splatting for physic-based simulations. 2024. 2
- [3] Erwin Coumans and Yunfei Bai. Pybullet, a python module for physics simulation for games, robotics and machine learning. <http://pybullet.org>, 2016–2021. 2
- [4] Chen Gao, Ayush Saraf, Johannes Kopf, and Jia-Bin Huang. Dynamic view synthesis from dynamic monocular video. In *Proceedings of the IEEE International Conference on Computer Vision*, 2021. 1
- [5] Antoine Guédon and Vincent Lepetit. Sugar: Surface-aligned gaussian splatting for efficient 3d mesh reconstruction and high-quality mesh rendering. *CVPR*, 2024. 2
- [6] Michelle Guo, Matt Jen-Yuan Chiang, Igor Santesteban, Nikolaos Sarafianos, Hsiao-Yu Chen, Oshri Halimi, Aljaž Božič, Shunsuke Saito, Jiajun Wu, C. Karen Liu, Tuur Stuyck, and Egor Larionov. Pgc: Physics-based gaussian cloth from a single pose. In *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 21215–21225, 2025. 2
- [7] Yi-Hua Huang, Yang-Tian Sun, Ziyi Yang, Xiaoyang Lyu, Yan-Pei Cao, and Xiaojuan Qi. Sc-gs: Sparse-controlled gaussian splatting for editable dynamic scenes. *arXiv preprint arXiv:2312.14937*, 2023. 2
- [8] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4), 2023. 1, 2
- [9] Isabella Liu, Hao Su, and Xiaolong Wang. Dynamic gaussians mesh: Consistent mesh reconstruction from monocular videos, 2024. 1, 2, 3
- [10] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *ACM Trans. Graphics (Proc. SIGGRAPH Asia)*, 34(6):248:1–248:16, 2015. 2
- [11] Jonathon Luiten, Georgios Kopanas, Bastian Leibe, and Deva Ramanan. Dynamic 3d gaussians: Tracking by persistent dynamic view synthesis. In *3DV*, 2024. 2
- [12] Shaojie Ma, Yawei Luo, Wei Yang, and Yi Yang. Mags: Reconstructing and simulating dynamic 3d objects with mesh-adsorbed gaussian splatting. 2024. 1, 2, 3
- [13] 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. 1
- [14] Ben Moran, Mauro Comi, Arunkumar Byravan, Steven Bohez, Tom Erez, Zhibin Li, and Leonard Hasenclever. Splatting physical scenes: End-to-end real-to-sim from imperfect robot data, 2025. 2
- [15] Keunhong Park, Utkarsh Sinha, Jonathan T. Barron, Sofien Bouaziz, Dan B Goldman, Steven M. Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. *ICCV*, 2021. 1
- [16] Anastasiya Pechko, Piotr Borycki, Joanna Waczyńska, Daniel Barczyk, Agata Szymańska, Sławomir Tadeja, and Przemysław Spurek. Gs-verse: Mesh-based gaussian splatting for physics-aware interaction in virtual reality, 2025. 1
- [17] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. *arXiv preprint arXiv:2011.13961*, 2020. 1
- [18] Shenhan Qian, Tobias Kirschstein, Liam Schoneveld, Davide Davoli, Simon Giebenhain, and Matthias Nießner. Gaussianavatars: Photorealistic head avatars with rigged 3d gaussians. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20299–20309, 2024. 2
- [19] Boxiang Rong, Artur Grigorev, Wenbo Wang, Michael J. Black, Bernhard Thomaszewski, Christina Tsalicoglou, and Otmar Hilliges. Gaussian Garments: Reconstructing simulation-ready clothing with photorealistic appearance from multi-view video. In *International Conference on 3D Vision 2025*, 2025. 2
- [20] Olga Sorkine and Marc Alexa. As-rigid-as-possible surface modeling. In *Proceedings of EUROGRAPHICS/ACM SIGGRAPH Symposium on Geometry Processing*, pages 109–116, 2007. 2
- [21] Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033. IEEE, 2012. 1, 2
- [22] Joanna Waczyńska, Piotr Borycki, Sławomir Tadeja, Jacek Tabor, and Przemysław Spurek. Games: Mesh-based adapting and modification of gaussian splatting. 2024. 2
- [23] Guanjun Wu, Taoran Yi, Jiemin Fang, Lingxi Xie, Xiaopeng Zhang, Wei Wei, Wenyu Liu, Qi Tian, and Xinggang Wang. 4d gaussian splatting for real-time dynamic scene rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 20310–20320, 2024. 2
- [24] Tianyi Xie, Zeshun Zong, Yuxing Qiu, Xuan Li, Yutao Feng, Yin Yang, and Chenfanfu Jiang. Physgaussian: Physics-integrated 3d gaussians for generative dynamics. *arXiv preprint arXiv:2311.12198*, 2023. 2
- [25] Ziyi Yang, Xinyu Gao, Wen Zhou, Shaohui Jiao, Yuqing Zhang, and Xiaogang Jin. Deformable 3d gaussians for high-fidelity monocular dynamic scene reconstruction. *arXiv preprint arXiv:2309.13101*, 2023. 2