DecoupledGaussian: Object-Scene Decoupling for Physics-Based Interaction

Miaowei Wang[1] ✠ Yibo Zhang[3] Weiwei Xu[4]
Rui Ma[3] Changqing Zou[2,4] ✉ Daniel Morris[5]
[1] University of Edinburgh [2] Zhejiang Lab [3] Jilin University
[4] State Key Lab of CAD&CG, Zhejiang University [5] Michigan State University

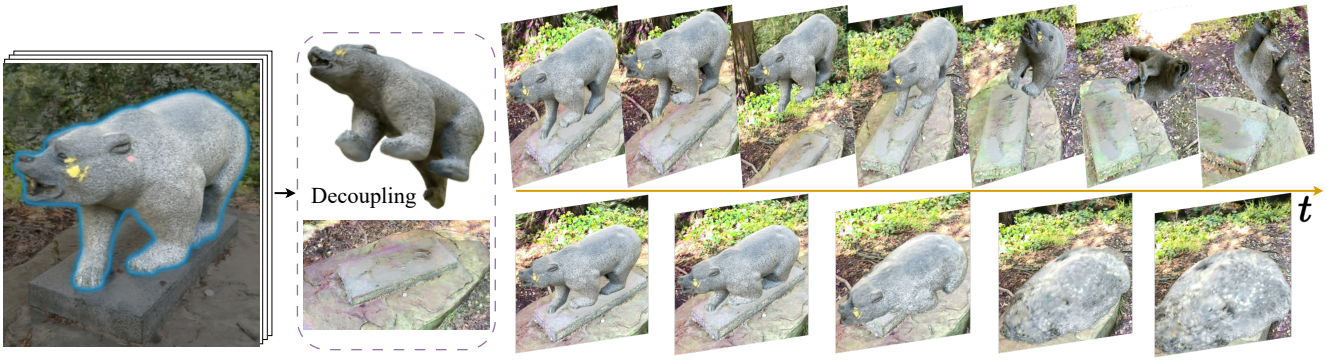

Figure 1. **DecoupledGaussian** decomposes static objects and contacted scenes from videos or multi-view images, enabling simulations like scene collisions (Top) and object melting with material adjustments (Bottom). See the supplementary video for the full sequences.

Abstract

We present DecoupledGaussian, a novel system that decouples static objects from their contacted surfaces captured in-the-wild videos, a key prerequisite for realistic Newtonian-based physical simulations. Unlike prior methods focused on synthetic data or elastic jittering along the contact surface, which prevent objects from fully detaching or moving independently, DecoupledGaussian allows for significant positional changes without being constrained by the initial contacted surface. Recognizing the limitations of current 2D inpainting tools for restoring 3D locations, our approach proposes joint Poisson fields to repair and expand the Gaussians of both objects and contacted scenes after separation. This is complemented by a multi-carve strategy to refine the object's geometry. Our system enables realistic simulations of decoupling motions, collisions, and fractures driven by user-specified impulses, supporting complex interactions within and across multiple scenes. We validate DecoupledGaussian through a comprehensive user study and quantitative benchmarks. This system enhances digital interaction with objects and scenes in real-world environments, benefiting industries such as VR, robotics, and autonomous driving. Our project page is at: https://wangmiaowei.github.io/DecoupledGaussian.github.io/.

✠: Partially conducted at Zhejiang Lab. ✉: Corresponding author.

1. Introduction

Interactive reconstruction and simulation of target objects and their surrounding scenes have become increasingly sophisticated recently. These can provide 4D assets for autonomous driving [74] and robotics [64], and also enable immersive applications in Virtual Reality (VR) [66] and the entertainment industry [95].

Advances in realism have been made by moving beyond traditional representations, such as point clouds [47], meshes [69], grids [12], and signed distance fields [77]. Neural Radiance Fields (NeRF) [62] use neural rendering techniques for novel view synthesis from videos, enabling interactive games [60], animation [96], and simulations [50], where *what is simulated directly stems from what is captured*. And Gaussian Splatting (GS) [43], known for its rapid rendering and reconstruction speeds, leverages discrete Gaussian kernels, making it easier to directly manipulate and process [11, 21] objects reconstructed from videos.

However, current physics-based simulation methods that use NeRF [17, 50] or Gaussian splatting [6, 38, 89] either focus on synthetic objects, allowing for full-view observations during reconstruction, or simulate elastic deformations and jittering, in which objects remain constrained to the contacted surface. This prevents objects from truly detaching under user-specified impulses.

To allow objects to move without being constrained by

the initial contacted surface, we need to decouple objects from the contacted surface before simulation. In real-life settings, objects are influenced by gravity and typically rest on other surfaces, such as the sculpture on the pedestal in Fig. 1. During imaging, an object and its contacted surfaces will be joined, resulting in hidden parts and occlusions and a fragmented representation of the object's surface. The primary challenge in decoupling, therefore, is accurately restoring and completing the 3D structures of both the object and its surrounding scene before simulation.

To tackle this challenge, we introduce **DecoupledGaussian**, a system that restores 3D geometry and textures of objects and contacted surfaces from in-the-wild videos using GS, laying the groundwork for realistic object-scene interactive simulations (see Fig. 1). Notably, 2D inpainting (Fig. 2) often struggles with 3D restoration, especially in accurately capturing geometric positions. Our approach overcomes this by leveraging geometric priors assuming closed surfaces and multi-view observations from training viewpoints to restore realistic object and scene geometry.

Our method employs joint Poisson fields to reconstruct shape indicators for objects and scenes, resolving intersecting regions. Using Gaussian centers directly can introduce surface deviations due to blended rendering, causing artifacts in object reconstruction. To avoid this, we use unbiased depth maps from planar-based GS to create proxy points for realistic object reconstruction and reduce the scene's floaters through geometry regularization with flattened 3D Gaussians. To alleviate geometry expansion in Poisson reconstruction, we introduce a unilateral negative cross-entropy (UNCE) method for multi-view carving, refining the geometry to align with the observed views.

DecoupledGaussian is the first to restore both object and contacted surface geometry independently of 2D inpainting, which we use only for texture properties refinement. Extensive experiments on real-world videos, a new decoupling benchmark, user studies, quantitative comparisons, and ablations demonstrate our approach's effectiveness in restoring accurate 3D properties and enabling precise interactive simulations. In summary, our contributions include:

- Development of an object-scene interactive simulation system that allows objects to detach from their contacted surfaces when constructed from in-the-wild videos and represented using GS.
- Introduction of geometric priors via joint Poisson fields and multi-view observations with UNCE for more realistic restoration (see Fig. 6) of geometric properties in GS.

2. Related Work

GS Editing A variety of methods have been proposed to modify or edit scenes built by Gaussian Splatting. Wu et al. [88] enhance GS textures with learnable lighting adjustments, while Fiebelman et al. [18] refine 4D video playback using human language prompts. Texture-GS [91] supports texture modifications via UV mapping decoupled from the original GS, and Ma et al. [57] introduce deformations by aligning Gaussians to a proxy mesh with as-rigid-as-possible regularization. Additionally, GaussianEditor [11] allows object addition and removal in GS scenes through 2D segmentation [45] and inpainting techniques [79]. SC-GS [34] enables object deformation using sparse control points learned from dynamic video data, while Modi et al. [65] learns skinning weights for elastic deformations. Huang and Yu [29] apply a bounding cage as a control proxy to deform GS representations.

GS Simulation Gaussian Splatting can be incorporated into traditional simulation frameworks. For instance, PhysGaussian [89] uses the Material Point Method (MPM) [37, 46, 78] to simulate Gaussian kernels motion directly, while VR-GS [38] applies eXtended Position-based Dynamics (XPBD) [59] to control GS via a bounding mesh. Feng et al. [16] combine XPBD to model interactions between liquids and solids in GS, and Borycki et al. [5] utilize MPM with triangle soups derived from GS. Additionally, Abou-Chakra et al. [2] apply GS in robotic decision-making through XPBD. For accurate physical property estimation in these simulations, GIC [6] derives physical properties from multiview video captures, building on techniques from Guan et al. [20], Li et al. [50] for system identification. Besides, Liu et al. [52] and Huang et al. [33] estimate material properties from synthetic video generated from static images using generative models. Whitney et al. [85] developed simulators trained on dynamic multi-view RGB-D video, and Qiu et al. [72] use visual language models to classify objects as elastic or rigid for text-driven physics simulations. However, these methods do not address the challenge of simulating an object detached from the contact surface when a user-provided impulse is applied.

Figure 2. Inpainting tools (LaMa [79]; PhotoRoom [1]) introduce artifacts and inconsistent textures across frames.

GS Restoration Restoration techniques have addressed occluded mesh completion [28], single-view depth point cloud completion [31, 32, 39, 81, 87], and SDF-based reconstructions [19, 30, 56, 68]. In the context of Gaussian Splatting, recent work has focused on restoring scene sur-

faces after object removal [54, 83] and reconstructing GS objects from sparse views [92]. However, a key challenge remains unaddressed: no GS or NeRF-based methods currently restore objects occluded by the scene or missed due to limited training viewpoints—a challenge mesh restoration techniques have begun to tackle [24]. Similarly, video generation methods like PhysGen [53], which rely on static cameras, also overlook this issue. Current GS scene restoration techniques [54, 63, 83] depend heavily on 2D inpainting tools [1, 71, 79] to fill gaps in geometry and texture post-object removal. However, these methods face two major issues (Fig. 2): (1) inpainted regions often fail to blend seamlessly with surrounding geometry, creating artifacts, and (2) texture inconsistencies across frames due to the lack of robust video inpainting tools. Our approach addresses these limitations by prioritizing geometry restoration, leveraging intrinsic GS geometry priors to ensure a coherent surface even when texture inpainting is imperfect.

3. Preliminaries

3.1. 3D Gaussian Splatting

Gaussian Splatting [43] represents a 3D scene with possible features by constructing 3D Gaussian kernels $\{\boldsymbol{k}_g, \sigma_g, \boldsymbol{\Sigma}_g, \mathcal{C}_g\}_{g\in\mathcal{G}}$, where $\boldsymbol{k}_g$, σ_g, $\boldsymbol{\Sigma}_g$, and $\mathcal{C}g$ denote Gaussian centers, opacities (encoding density), covariance matrices, and spherical harmonic (SH) representing color coefficients, respectively. The covariance matrix $\boldsymbol{\Sigma}_g$ at a Gaussian g is factorized as $\boldsymbol{\Sigma}_g = \boldsymbol{R}_g\boldsymbol{S}_g\boldsymbol{S}_g^T\boldsymbol{R}_g^T$, where $\boldsymbol{R}_g$ is a rotation matrix, and corresponding scaling $\boldsymbol{S}_g = diag(s_1, s_2, s_3)$ is a diagonal matrix. Like NeRF [62], GS is optimized for novel-view synthesis, where for a given 2D image plane, an integrated quantity $\boldsymbol{q}$ at a pixel p is obtained by the following front-to-back α-blending:

$$\boldsymbol{q}(p) = \sum_{i\in\mathcal{G}} \boldsymbol{q}_i\alpha_i \left[\prod_{j=1}^{i-1}(1-\alpha_j)\right] \tag{1}$$

where $\boldsymbol{q}_i$ is the quantity (for instance, SH-evaluated color c_i), and α_i is the termination probability derived from opacity σ_i and affine-projected 2D Gaussian weights from $\boldsymbol{\Sigma}_i$.

3.2. Continuum Simulation

We use the MLS-MPM [27] framework to solve Gaussian kernel governing equations (mass and momentum conservation) [89]. The continuum is discretized into Lagrangian particles p, with time steps of Δt for deformation. At each step, particle mass and momentum are transferred to an Eulerian grid (P2G), where momentum is updated using the first Piola-Kirchhoff stress (PK1), and velocities v are advanced via forward Euler integration (Grid Operation). These grid velocities are then interpolated back to particles (G2P) for position updates during advection. MLS-MPM employs affine C_p as a first-order approximation of ∇v, optimizing computation time. The elastic deformation gradient F^E is updated as $F_\mathrm{p}^{n+1} = (I + \Delta t\, C_\mathrm{p}^n)F_\mathrm{p}^n$. Material parameters such as Young's modulus E and shear modulus μ [37] influence PK1 in grid momentum updates.

4. DecoupledGaussian

The DecoupledGaussian system starts with a reconstructed GS scene and allows an object resting on a planar surface to be moved off its surface in a realistic manner as shown in Fig. 3. First an object is segmented and a planar-based GS aligns Gaussians $\mathcal{G}$ to the underlying surface geometry. Joint Poisson fields, informed by geometric priors, then repair fragmented surfaces of both the scene and object after separation. For the object, proxy points serve as input to the Poisson fields, and the output is carved using our UNCE method to ensure geometry aligns with training observations. The Gaussians' texture properties ($\{\sigma_g, \mathcal{C}_g\}$) are refined with 2D inpainting, and this is followed by a real-time interactive simulation of the decoupled object and scene via MLS-MPM. Each stage is detailed in this section.

4.1. 3D Gaussians Preparation

The scene is freely recorded with a consumer-level camera. The frame sequence is then processed in COLMAP [75, 76] to obtain intrinsic and extrinsic calibrations and to generate initial Gaussian centers for the next section.

4.1.1 Planar-based Gaussian Splatting

Optimizing vanilla 3D Gaussian models [43] with only image reconstruction loss often results in local optima, complicating accurate geometry extraction, which is vital for the subsequent restoration stage. To avoid this, we adopt PGSR [8] for unbiased depth D estimation. Given the inherent disorder of vanilla Gaussian distributions, we initially compress the Gaussians into an approximate local plane that aligns with the scene surface. This is achieved by penalizing the minimum scaling term $||\min(s_1, s_2, s_3)||_1$ during training, allowing for a tolerable loss in rendering quality to enhance geometric accuracy. After compression into plane-like Gaussians (see the right inset), we assign normals $\boldsymbol{n}_g$ along the shortest axis, with orientation disambiguated by viewing directions [8]. The distance to the image plane is calculated as $d_g = ||\boldsymbol{n}_g^T\boldsymbol{d}'_g||$ where $\boldsymbol{d}'_g$ is the vector from the camera center to the Gaussian center $\boldsymbol{k}_g$. The final unbiased depth at pixel p after α-blending (see Eq. (1)) is then given by

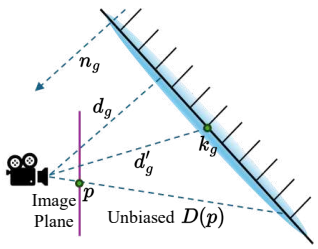

$$D(p) = \frac{d(p)}{\boldsymbol{n}(p)\boldsymbol{K}^{-1}p'} \tag{2}$$

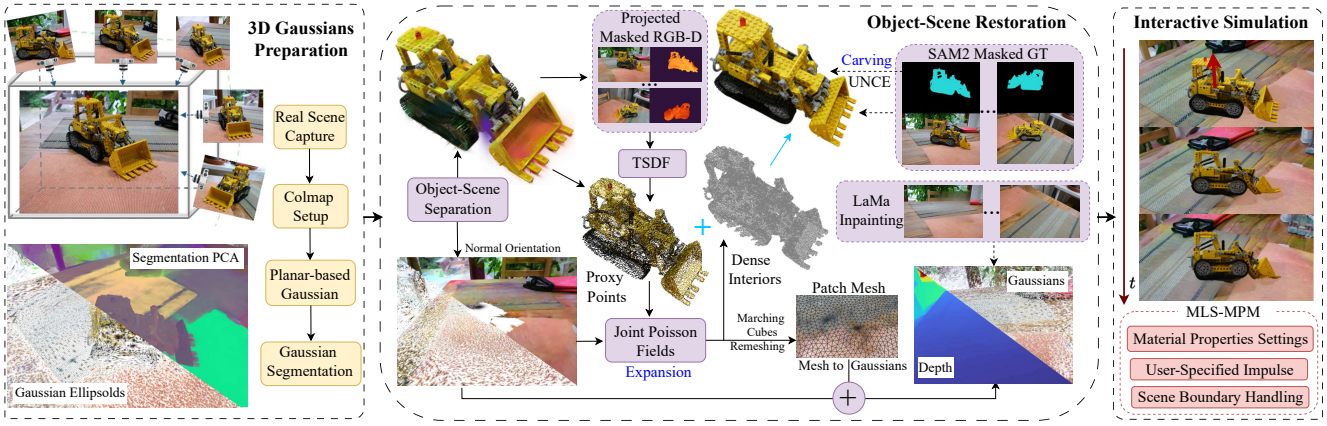

Figure 3. **System Overview**. DecoupledGaussian is an interactive simulation system that enables objects to detach from their initial contact surfaces after applying our proposed restoration pipeline, driven by user-specified impulses (red arrow on the right).

where $\boldsymbol{K}$ is the camera's intrinsic matrix and p' is the homogeneous coordinate of p. Flattened Gaussians provide single- and multi-view geometry regularization for consistent geometry, with exposure compensation applied to address illumination variations (see Chen et al. [8] for details).

4.1.2 Gaussian Segmentation

To implement GS segmentation [7, 94], each kernel g is assigned semantic affinity features $\boldsymbol{\xi}_g \in \mathbb{R}^{32}$. A gating network, a single-layer MLP $\zeta : \mathbb{R}^{32} \rightarrow \mathbb{R}^C$ [94], maps α-blended features $\boldsymbol{\xi}(p)$ to C segmentation class probabilities via softmax [36]. The network is trained with cross-entropy loss using multi-view 2D segmentation labels from SAM2 [73]. To reduce artifacts among nearby Gaussians, we apply local feature smoothing [7] and initialize segmentation by manually selecting classes in the first frame [38].

4.2. Object-Scene Restoration

To simulate an object $\mathcal{O}$ interacting with its surrounding scene surface $\mathcal{S}$, we first separate $\mathcal{O}$ from $\mathcal{S}$ by identifying its Gaussians through comparing affinity features with α-blended $\boldsymbol{\xi}(p)$ at a user-specified click position p. We then apply KNN to remove nearby Gaussians representing residual artifacts [94]. For realistic simulation, we should repair and complete both $\mathcal{O}$ and $\mathcal{S}$, as detailed next.

4.2.1 Joint Poisson Fields

The main contribution is the novel restoration of the geometric properties $\{\boldsymbol{k}_g, \boldsymbol{\Sigma}_g\}$ of GS, assuming that both the object $\mathcal{O}$ and scene $\mathcal{S}$ are smooth, closed shapes. Inspired by the equivalence between Poisson surface reconstruction and winding number field construction [15, 90], we introduce **joint Poisson fields** $\mathcal{W}$, which incorporate heterogeneous constraints to enable the simultaneous restoration of both $\mathcal{O}$ and $\mathcal{S}$ (see Fig. 4). The process is as follows:

(1) Solve the indicator functions $\mathcal{X}_\mathcal{S}$ and $\mathcal{X}_\mathcal{O}$ for the scene and object surfaces, respectively, via screened Poisson reconstruction [41] implicitly enforcing a minimum curvature surface. We implement this with Adaptive Multigrid Solvers [40] in corresponding canonical grid spaces, $\mathcal{W}_\mathcal{S}$ and $\mathcal{W}_\mathcal{O}$, where $\mathcal{X} > 0.5$ indicates interiors while $\mathcal{X} < 0.5$ for exteriors. Each Poisson field with 128^3 grid size is processed in under 20 seconds in experiments.

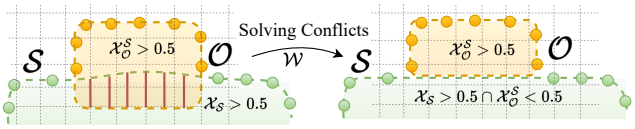

Figure 4. Joint Poisson Fields $\mathcal{W}$ first reconstruct $\mathcal{O}$ and $\mathcal{S}$ independently, then resolve conflicts (red area) by defining a boundary that separates them into distinct, non-intersecting entities.

(2) Transform $\mathcal{X}_\mathcal{O}$ to $\mathcal{X}_\mathcal{O}^\mathcal{S}$ by mapping it from $\mathcal{W}_\mathcal{O}$ to $\mathcal{W}_\mathcal{O}^\mathcal{S}$ via world-coordinate transformation to the canonical coordinates of $\mathcal{S}$. To resolve conflicts $\{x \mid \mathcal{X}_\mathcal{S}(x) > 0.5 \cap \mathcal{X}_\mathcal{O}^\mathcal{S}(x) > 0.5\}$ (intersection regions), we prioritize $\mathcal{S}$ (details in Suppl.) due to its simpler, more reliable geometry. Conflicting regions in $\mathcal{W}_\mathcal{O}^\mathcal{S}$ are then discarded.

(3) Dense interior points $P_\mathcal{O}$ (for continuum simulation) are extracted from $\mathcal{W}_\mathcal{O}^\mathcal{S}$. We apply marching cubes [55] to $\mathcal{W}_\mathcal{S}$ and then re-meshing [70] and further cropped by $P_\mathcal{O}$-scaled bounding box to get a mesh patch $\mathcal{M}_\mathcal{S}$. Both $\mathcal{M}_\mathcal{S}$ and $P_\mathcal{O}$ are subsequently converted to world coordinates.

To solve $\mathcal{X}_\mathcal{S}$, we use Gaussian centers $\{\boldsymbol{k}_g\}_{g\in\mathcal{S}}$ as input (see suppl. for normals). For $\mathcal{X}_\mathcal{O}$, due to the geometric complexity of $\mathcal{O}$, we introduce proxy points $\mathcal{P}_\mathcal{O}$ as input.

4.2.2 Proxy Points

Due to α-blending, Gaussian centers $\{\boldsymbol{k}_g\}_{g\in\mathcal{O}}$ fail to accurately represent the complex surface of $\mathcal{O}$. Our proposed proxy points $\mathcal{P}_\mathcal{O}$ can enhance geometry estimations of $\mathcal{X}_\mathcal{O}$ ablated as shown in Fig. 5.

We first render RGB images and unbiased depth maps D in Eq. (2) for the entire layout under all training views. Next, we obtain the projected mask $M_\mathcal{O}^{\text{proj}}$ by setting zero-opacity for all other Gaussians $(\mathcal{G} \setminus \mathcal{O})$, where zero in $M_\mathcal{O}^{\text{proj}}$ indicates no accumulated opacity, while one signifies existing opacity at a pixel location. Using the masked projected

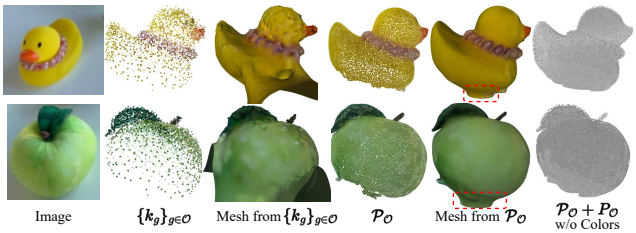

Figure 5. **Ablation for $\mathcal{P}_\mathcal{O}$.** Independent Poisson reconstruction of object $\mathcal{O}$ using Gaussian centers $\{k_g\}_{g\in\mathcal{O}}$ yields poor mesh quality compared to using proxy points $\mathcal{P}_\mathcal{O}$. Our joint Poisson field $\mathcal{W}$, which integrates the scene surface $\mathcal{S}$, effectively removes the overextended regions (highlighted in red). The final dense points $P_\mathcal{O}$ are then combined with proxy points $\mathcal{P}_\mathcal{O}$ for Gaussian restoration and continuum simulation.

depth map $D \circ M_\mathcal{O}^{\text{proj}}$, the TSDF fusion algorithm [67] concentrates on the object area, rapidly integrating information from training views within 100 seconds, followed by standard post-processing [8]. However, the integrated result still includes points from $(\mathcal{G} \setminus \mathcal{O})$ due to boundaries smearing [82] from $M_\mathcal{O}^{\text{proj}}$. To address this, we segment the final proxy points $\mathcal{P}_\mathcal{O}$ by inheriting features $\boldsymbol{\xi}_g$ from raw Gaussian kernels of $\mathcal{O}$ with nearest neighbor search.

4.2.3 Unilateral Negative Cross Entropy

Despite closing broken surfaces and filling internal dense points $\mathcal{P}_\mathcal{O}$, the over-smoothness of Poisson fields leads to *geometry expansion*, introducing particles beyond observable viewpoints. To address this, we apply *multi-view carving*. Specifically, we propose a Unilateral Negative Cross Entropy (UNCE) loss at each rendered pixel p for the **isometric** dense object Gaussians $G_\mathcal{O}$. This loss measures the discrepancy between the α-blended opacity $\mathbb{1}_\mathcal{O}$ (see Eq. (1)) during fine-tuning and the 2D ground truth object mask $M_\mathcal{O}^{\text{GT}}$ from SAM2, defined as:

$$\text{UNCE}(p) = -(1 - M_\mathcal{O}^{\text{GT}}(p)) \log(1 - \mathbb{1}_\mathcal{O}(p)). \quad (3)$$

Every 100 iterations, we clean Gaussians $\{\sigma_g \leq 0.05\}$. These isometric Gaussians $G_\mathcal{O}$ are defined by centers $\{k_g \in \mathcal{P}_\mathcal{O} \cup P_\mathcal{O}\}$, each associated with an opacity of $\sigma_g = 0.1$ and an isometric covariance matrix $\boldsymbol{\Sigma}_g = \text{diag}(s_g^2, s_g^2, s_g^2)$ [6]. Here, $s_g = c\left(\frac{3}{4\pi}\right)^{\frac{1}{3}}$ [89], where c is the Poisson grid cell length in world coordinates. For centers $\{k_g \in \mathcal{P}_\mathcal{O}\}$, view-independent SHs are derived from the colors of points already integrated by the TSDF. In contrast, view-independent SHs for $\{k_g \in P_\mathcal{O}\}$ are Gaussian-weighted interpolations based on the 15 nearest neighbors in $\mathcal{P}_\mathcal{O}$. We zero all other coefficients in each $\mathcal{C}_g$.

4.2.4 Gaussian Restoration

During multi-view carving, we also fine-tune $\{\sigma_g, \mathcal{C}_g\}_{g\in G_\mathcal{O}}$ as described by Li et al. [51], but using $M_\mathcal{O}^{\text{GT}}$-masked training images to mitigate influence from other areas in the scene. In each iteration, a random background is applied for image rendering and is consistently used for the other regions of the masked ground truth.

To restore the Gaussians (i.e., the holes) in the scene surface $\mathcal{S}$, we first bind new 3D flattened Gaussians $G_\mathcal{S}$ [5, 21] to the patch mesh $\mathcal{M}_S$ (see our Mesh to Gaussian algorithm in Supplementary) with minimal scaling ϵ along the mesh face normals. At this stage, we finalize and fix the geometric properties $\{k_g, \boldsymbol{\Sigma}_g\}_{g\in G_\mathcal{S}}$. During fine-tuning, we adjust only the texture properties $\{\sigma_g, \mathcal{C}_g\}_{g\in G_\mathcal{S}}$, initialized from the nearest neighbors of the raw broken $\mathcal{S}$, guided by 2D inpainted images in the masked areas $M_\mathcal{O}^{\text{GT}}$ using LaMa [79]. Finally, we fill holes in $\mathcal{S}$ by adding the patch $G_\mathcal{S}$.

4.3. Interactive Simulation

We simulate and render all restored Gaussians $G_\mathcal{O}$ to enable a range of interactive simulations, including user-specified impulses as external forces for elastic deformation, scene collisions with $\mathcal{S}$, and effects like shape fracturing and material changes, all based on MLS-MPM. To enforce a *Dirichlet boundary condition* [4], we set the velocities of grid nodes containing Gaussians from the restored scene $\mathcal{S}$ to zero during Grid Operation stage in MLS-MPM, creating a sticky boundary effect. To simulate gravity, we automatically align the z-axis by segmenting Gaussians for planar objects (e.g., ground or desk surfaces) and estimating the plane normals with RANSAC [49]. We then apply the rotation matrix derived from plane normals to all $\{k_g, \boldsymbol{\Sigma}_g\}_{g\in\mathcal{G}}$ directly, while view-dependent SHs are rotated through Wigner D-matrices [86] (see details in Suppl.)

5. Experiments

5.1. Implementation Details

Input resolutions range from 720p to 1K. Gaussian restoration of $\mathcal{S}$ and $\mathcal{O}$ uses L_1 and L_{SSIM} losses with UNCE regularization at 10^{-4}, fine-tuned for 1000 iterations for $\mathcal{S}$ and 3000 for $\mathcal{O}$, skipping iterations without valid masks, totaling under 4 minutes. For LaMa inpainting, masked areas are dilated with a 21×21 kernel to reduce boundary artifacts in $M_\mathcal{O}^{\text{GT}}$[54]. The simulation area and physical parameters (e.g., E, μ) are manually set following [38, 89] (see Supplementary). Based on Warp [58], the simulation runs on an 18-core Intel Xeon Gold 5220 CPU and NVIDIA A40 GPU, achieving ~10 FPS for 50-frame videos.

5.2. Evaluating Object-Scene Interaction

Dataset We evaluate our system for generating diverse object-scene interactive simulations using several real-world scenario sources. Our evaluation includes the following datasets: BICYCLE, GARDEN, BONSAI, ROOM, and KITCHEN from the Mip-NeRF360 [3] dataset; TRUCK from the Tanks&Temples dataset [48]; PLAYROOM from

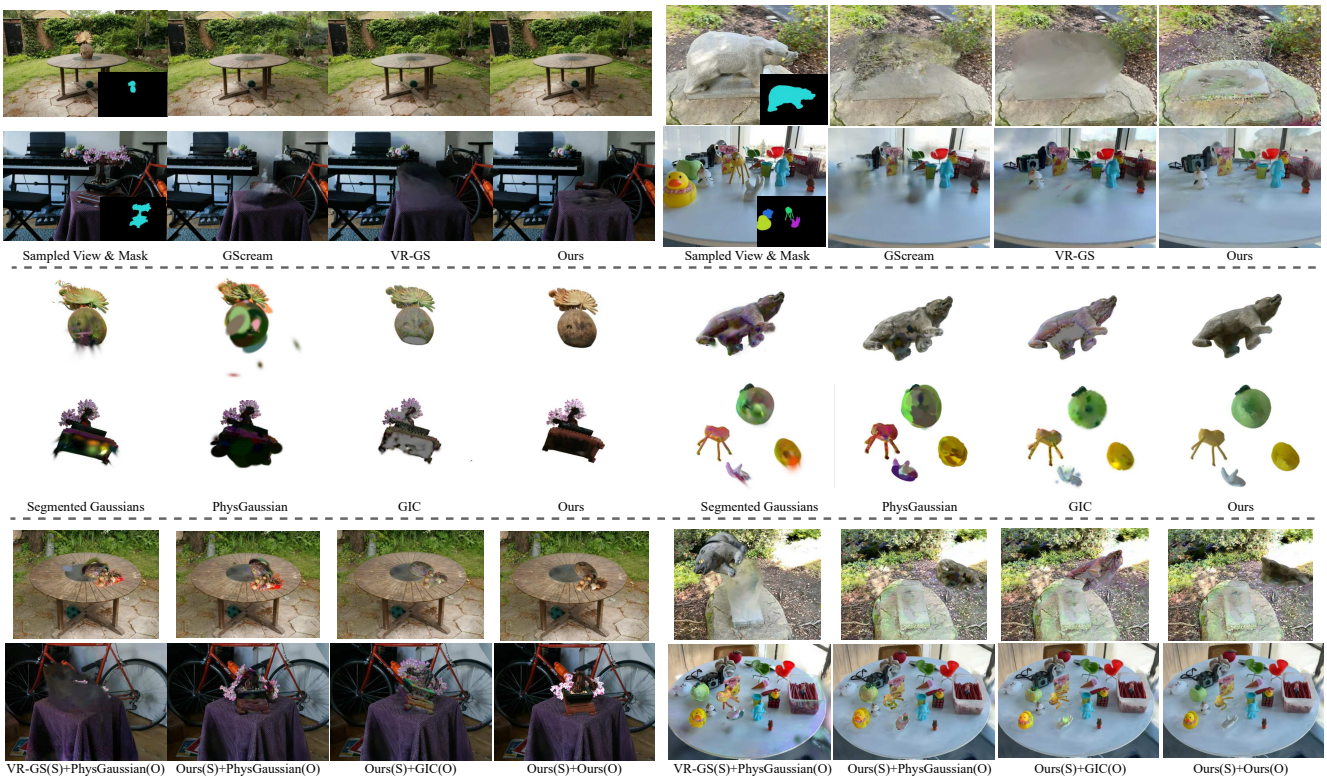

Figure 6. **Qualitative Comparisons.** We demonstrate **Scene Restoration** (Top), **Object Restoration** (Middle), and **Object-Scene Interactive Simulation** (Bottom) on real-world scenes, including GARDEN, BEAR, BONSAI, and FIGURINES. For each comparison, a single test viewpoint is selected, with the object initially suspended before being impacted by an external force during simulation.

the Deep Blending dataset [23]; FIGURINES from Lerf [44]; and BEAR from Instruct-NeRF2NeRF [22].

Baselines We compare our method with SOTA simulation frameworks based on Gaussian splatting, incorporating necessary adaptations: 1) **PhysGaussian** [89] uses anisotropy regularization to prevent narrow kernels and applies a user-defined opacity field (based on $\mathcal{O}$'s bounding box) for interior filling. 2) **GIC** [6] employs isotropic Gaussians with a coarse-to-fine density field to fill interior points, assigning scales and opacities with zero color. 3) **VR-GS** [38] is closely aligned with our approach; however, due to unavailable simulation code, we adapt their methods to restore scene Gaussians, $\mathcal{S}$, with geometry and texture properties guided by LaMa. 4) **GScream** [83], a SOTA technique for $\mathcal{S}$ restoration, integrates monocular depth estimation [42] from the inpainted reference view for training guidance.

User Study We conducted a human evaluation to assess both visual realism and simulation fidelity, following methods from prior work [9, 53, 84]. Ten participants with varying experience in simulation and 3D vision rated three aspects: 1) **Scene Restoration Quality (SRQ)**, which evaluates the accuracy of scene restoration, $\mathcal{S}$, after object removal; 2) **Object Restoration Quality (ORQ)**, assessing the realism of restored objects, $\mathcal{O}$; and 3) **Interactive Simulation Fidelity (ISF)**, checking if the object scenes response to a user-specified impulse is both realistic and as expected. Rendered videos of $\mathcal{S}$, $\mathcal{O}$, and interactive simulations were presented in random order, with participants rating each on a five-point scale (1 = poor, 5 = excellent). Mean scores are reported, with supplementary material containing additional statistics and video examples.

Table 1. **User Study**. Participants rated the fidelity of restoration and interactive simulation in a moving-camera video.

Scene Restoration			**Object Restoration**	
Methods	SRQ ↑	Time ↓	Methods	ORQ ↑
GScream [83]	1.94	∼70m	PhysGaussian [89]	1.40
VR-GS [38]	2.12	∼7m	GIC [6]	1.60
Ours	**3.48**	**∼1m**	**Ours**	**4.03**

Object-Scene Interactive Simulation	
Methods	ISF ↑
VR-GS($\mathcal{S}$) + PhysGaussian($\mathcal{O}$)	1.50
Ours($\mathcal{S}$) + PhysGaussian($\mathcal{O}$)	2.60
Ours($\mathcal{S}$) + GIC($\mathcal{O}$)	2.73
Ours($\mathcal{S}$) + Ours($\mathcal{O}$)	**4.35**

Results Our method (Tab. 1) achieves the highest ratings and shortest training time (1 minute) for scene restoration. VR-GS and GScream rely on 2D inpainting for $\mathcal{S}$ restoration, leading to geometry inaccuracies (e.g., BEAR, BONSAI in Fig. 6 Top) when inpainting quality is poor. GScream's use of a single reference image limits view con-

Figure 7. **Interactive Versatility.** Our object-scene decoupling method enables a variety of user-specified interactions both within a single scene (e.g., ROOM Top) and across different scenes (e.g., TRUCK in BICYCLE, Bottom).

sistency, causing issues in non-forward-facing views. In contrast, our approach uses planar-based GS geometry priors, ensuring precise structural restoration while limiting 2D inpainting to texture properties. For $\mathcal{O}$ restoration, we are the first to restore $\mathcal{O}$ in both single and complex multi-object scenes (e.g., FIGURINES in Fig. 6 Middle), maintaining input quality. Unlike PhysGaussian, which suffers from artifacts due to incomplete opacity assumptions, and GIC, which shows artifacts from non-zero internal opacities (white dots in Fig. 6 Middle), our method produces stable, high-quality results. For interactive simulations, GIC (BONSAI in Fig. 6 Bottom) exhibits unintended motion due to particle imbalance. VR-GS, relying on 2D inpainting, shows flawed geometry of $\mathcal{S}$, limiting object-scene interactions and causing artifacts or pass-through issues (e.g., BEAR, BONSAI in Fig. 6 Bottom). Our video demonstrates dynamic effects, and Fig. 7 showcases simulations with user-specified impulses, including cross-scene interactions (e.g., TRUCK in BICYCLE scene), highlighting our method's high controllability and motion realism.

5.3. Decoupling Benchmark Evaluation

Dataset To address the lack of ground truth for object-scene decoupling interactions, we utilize real reconstructed scenes and objects from the PEGASET dataset [61] and the PLAYROOM and SOFA SUITE environments from BlenderNeRF [10]. Test cases with realistic elements are created by placing objects within scenes using PyBullet [13] and rendering object-scene setups as input from the raw scene's training viewpoints. Ground truth for object restoration is provided by well-reconstructed objects, internally filled [89], and without internal textures, while ground truth for scene restoration is based on the raw scenes with no objects. For object-scene interaction, we render multi-view MLS-MPM simulations by dropping these ground truth objects from a height to the ground, following [6].

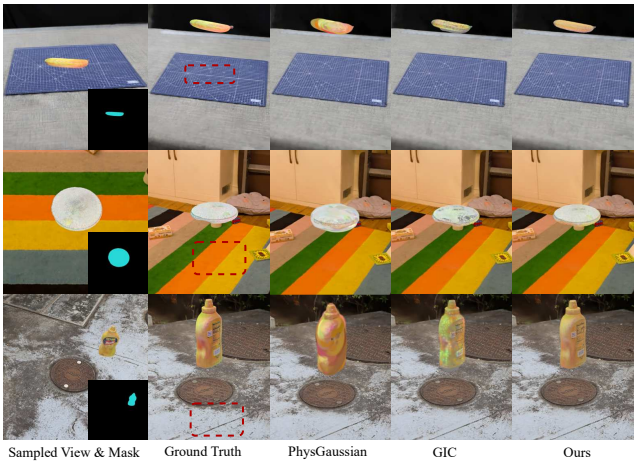

Figure 8. **Benchmark Comparisons.** A test viewpoint visualizes comparisons using the restored scene from our method, with inpainting regions marked by a red rectangle in the Ground Truth.

Metrics We use **PSNR** [26], **LPIPS** [99], and **FID** [25] as primary metrics to evaluate reconstruction quality. For restoring scene $\mathcal{S}$ and object $\mathcal{O}$, we additionally apply Chamfer Distance (**CD**) [14] to measure the geometric accuracy of Gaussian centers within inpainted regions, critical for accurate physics-based simulation. Viewpoints outside training views are used for $\mathcal{O}$ captures while training viewpoints are retained for $\mathcal{S}$. To evaluate motion accuracy in interactive simulations, we compute **Motion-FID** [53] by extracting and colorizing optical flow using RAFT [80] and calculating FID on the resulting flow images.

Results For object-scene interaction, we use our restored $\mathcal{S}$ across all methods to ensure fair comparison. Quantitative results are shown in Tab. 2 and qualitative examples in Fig. 8, with sample views of objects attached to scene surfaces from the input. Our restoration of $\mathcal{S}$ closely matches ground truth, and $\mathcal{O}$ significantly outperforms other methods in GS simulation. Competing methods often produce artifacts due to inadequate handling of broken surfaces and

Table 2. **Quantitative Comparisons & Ablations.** We create a decoupling benchmark with comprehensive metrics comparing baselines and ablations to validate design choices.

Scene Restoration				
Methods	PSNR ↑	LPIPS ↓	FID ↓	CD ($\times 10^{-3}$) ↓
GScream [83]	17.82	0.56	42.28	44.00
VR-GS [38]	25.13	0.32	58.50	6.41
Ours	**27.32**	**0.30**	**32.07**	**4.40**
Object Restoration				
Methods	PSNR ↑	LPIPS ↓	FID ↓	CD ($\times 10^{-3}$) ↓
PhysGaussian [89]	24.46	0.07	227.60	0.53
GIC [6]	26.62	0.06	201.91	0.73
Ours	**30.32**	**0.04**	**138.75**	**0.17**
Object-Scene Interaction Simulation				
Methods	PSNR ↑	LPIPS ↓	FID ↓	Motion-FID ↓
PhysGaussian [89]	19.48	0.37	112.55	54.79
GIC [6]	20.90	0.31	134.56	47.47
w/o dense $P_{\mathcal{O}}$	21.19	0.29	98.19	48.39
w/o Proxy $\mathcal{P}_{\mathcal{O}}$	21.08	0.30	90.26	36.01
w/o $\mathcal{W}$	20.97	0.30	96.16	42.27
Ours	**21.33**	**0.29**	**86.98**	**31.69**

hidden areas, which degrades interactive simulation quality. Although our approach excels in interactive simulation, the object's FID is high due to reliance on training view interpolation for texture restoration. Future work will explore 3D AI-based texture generative inpainting to improve this.

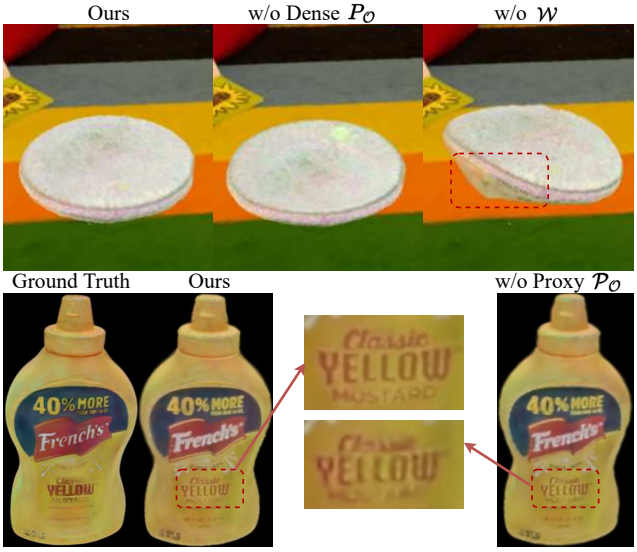

Figure 9. **Dense** $P_{\mathcal{O}}$ prevents collapse under gravity, **Joint Point Fields** $\mathcal{W}$ remove red-highlighted intersection regions, and **Proxy Points** $\mathcal{P}_{\mathcal{O}}$ enhance texture details (see zoom-ins).

Ablations We quantitatively evaluate several design choices (see Tab. 2): 1) **Dense Interior Points** ($P_{\mathcal{O}}$), our internal filling strategy, prevent collapse under gravity or external forces, unlike objects without internal particles (see Fig. 9, Top-Middle). 2) **Proxy Points** ($\mathcal{P}_{\mathcal{O}}$) enhance geometry recovery in Poisson reconstruction (see Fig. 5), and their combination with $P_{\mathcal{O}}$ improves texture details over $P_{\mathcal{O}}$ alone (see Fig. 9, Bottom). 3) **Joint Poisson Fields** ($\mathcal{W}$) reduce artifacts and resolve intersection regions better than independent Poisson reconstructions (see Fig. 9, Top-Right).

5.4. Additional Qualitative Ablations

UNCE Poisson reconstruction prioritizing smoothness in $\mathcal{W}$ can introduce artifacts, even with opacity filtering. Our Unilateral Negative Cross Entropy (UNCE) method (shown in Fig. 10, Top) leverages negative labels from SAM2 to carve and remove these artifacts, aligning $G_{\mathcal{O}}$ with the underlying geometry for accurate simulation.

Planar-based GS Unlike standard Gaussian splatting [43] with low-opacity filtering ($\sigma_g \leq 0.02$) [89], planar-based GS with compressed kernels enhances geometry regularization, reducing floaters (Fig. 10, Bottom) without the need for opacity filtering. This method enables unrestricted object motion in the simulation area, free from scene artifacts.

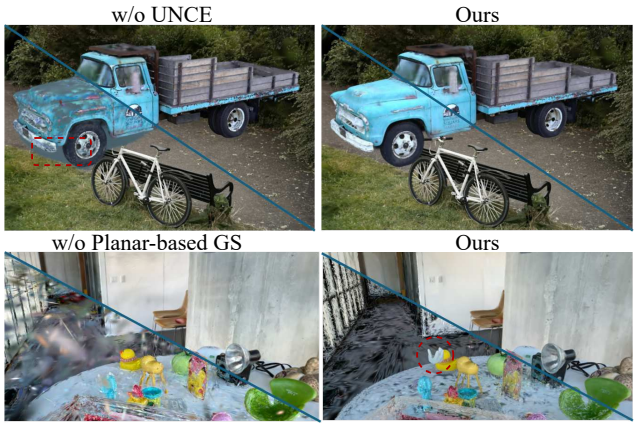

Figure 10. **UNCE** (Top) removes artifacts from Poisson expansion via multi-view carving. Opacity is set to one for TRUCK to highlight artifacts. **Planar-based GS** (Bottom) avoids floaters and artifacts compared to Vanilla GS [43], which limits motion (e.g., red-circled figurine). Opacity is set to 1, with $\times 0.4$ scaling for better Gaussian kernel visualization.

6. Discussion

Conclusion This paper presents DecoupledGaussian, a fast and robust approach for decoupling static objects from contact surfaces and restoring geometry and texture for object-scene interaction using the MLS-MPM simulator.

Limitations Our evaluation does not address complex scenes with multiple objects in varying contact configurations. High-frequency texture completion for object restoration is challenging, and GS-based texture generative approaches [35, 97, 98] may offer potential solutions. Additionally, decoupling the fine-grained components [93] of individual objects presents further difficulties.

Acknowledgements This research was also supported by Zhejiang Provincial Natural Science Foundation of China under Grant No. LD24F020007.

References

[1] Remove object from photo tool - photoroom. https://www.photoroom.com/tools/remove-object-from-photo. Accessed: 2024-11-03. 2, 3

[2] Jad Abou-Chakra, Krishan Rana, Feras Dayoub, and Niko Suenderhauf. Physically embodied gaussian splatting: A visually learnt and physically grounded 3d representation for robotics. In *8th Annual Conference on Robot Learning*, 2024. 2

[3] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5470–5479, 2022. 5

[4] Yuri Bazilevs and Thomas JR Hughes. Weak imposition of dirichlet boundary conditions in fluid mechanics. *Computers & fluids*, 36(1):12–26, 2007. 5

[5] Piotr Borycki, Weronika Smolak, Joanna Waczyńska, Marcin Mazur, Sławomir Tadeja, and Przemysław Spurek. Gasp: Gaussian splatting for physic-based simulations. *arXiv preprint arXiv:2409.05819*, 2024. 2, 5

[6] Junhao Cai, Yuji Yang, Weihao Yuan, Yisheng He, Zilong Dong, Liefeng Bo, Hui Cheng, and Qifeng Chen. Gaussian-informed continuum for physical property identification and simulation. *Advances in Neural Information Processing Systems*, 2024. 1, 2, 5, 6, 7, 8

[7] Jiazhong Cen, Jiemin Fang, Chen Yang, Lingxi Xie, Xiaopeng Zhang, Wei Shen, and Qi Tian. Segment any 3d gaussians, 2024. 4

[8] Danpeng Chen, Hai Li, Weicai Ye, Yifan Wang, Weijian Xie, Shangjin Zhai, Nan Wang, Haomin Liu, Hujun Bao, and Guofeng Zhang. Pgsr: Planar-based gaussian splatting for efficient and high-fidelity surface reconstruction. *arXiv preprint arXiv:2406.06521*, 2024. 3, 4, 5

[9] Haoxin Chen, Menghan Xia, Yingqing He, Yong Zhang, Xiaodong Cun, Shaoshu Yang, Jinbo Xing, Yaofang Liu, Qifeng Chen, Xintao Wang, et al. Videocrafter1: Open diffusion models for high-quality video generation. *arXiv preprint arXiv:2310.19512*, 2023. 6

[10] Shuo Chen, Wu Liu, Binbin Yan, Xinzhu Sang, Alicia Li, and Xiangcheng Yi. Blender-nerf: A monocular dynamic human body explicit reconstruction and rendering method. In *2024 IEEE International Conference on Multimedia and Expo Workshops (ICMEW)*, pages 1–6. IEEE, 2024. 7

[11] Yiwen Chen, Zilong Chen, Chi Zhang, Feng Wang, Xiaofeng Yang, Yikai Wang, Zhongang Cai, Lei Yang, Huaping Liu, and Guosheng Lin. Gaussianeditor: Swift and controllable 3d editing with gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21476–21485, 2024. 1, 2

[12] Jonathan M Cohen, Sarah Tariq, and Simon Green. Interactive fluid-particle simulation using translating eulerian grids. In *Proceedings of the 2010 ACM SIGGRAPH symposium on Interactive 3D Graphics and Games*, pages 15–22, 2010. 1

[13] Erwin Coumans and Yunfei Bai. Pybullet, a python module for physics simulation for games, robotics and machine learning, 2016. 7

[14] Haoqiang Fan, Hao Su, and Leonidas J Guibas. A point set generation network for 3d object reconstruction from a single image. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 605–613, 2017. 7

[15] Nicole Feng, Mark Gillespie, and Keenan Crane. Winding numbers on discrete surfaces. *ACM Trans. Graph.*, 42(4): 36–1, 2023. 4

[16] Yutao Feng, Xiang Feng, Yintong Shang, Ying Jiang, Chang Yu, Zeshun Zong, Tianjia Shao, Hongzhi Wu, Kun Zhou, Chenfanfu Jiang, and Yin Yang. Gaussian splashing: Unified particles for versatile motion synthesis and rendering. *arXiv preprint arXiv:2401.15318*, 2024. 2

[17] 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*, pages 4450–4461, 2024. 1

[18] Gal Fiebelman, Tamir Cohen, Ayellet Morgenstern, Peter Hedman, and Hadar Averbuch-Elor. 4-legs: 4d language embedded gaussian splatting, 2024. 2

[19] Craig Gotsman and Kai Hormann. A linear method to consistently orient normals of a 3d point cloud. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–10, 2024. 2

[20] Shanyan Guan, Huayu Deng, Yunbo Wang, and Xiaokang Yang. Neurofluid: Fluid dynamics grounding with particle-driven neural radiance fields. In *International Conference on Machine Learning*, pages 7919–7929. PMLR, 2022. 2

[21] Antoine Guédon and Vincent Lepetit. Sugar: Surface-aligned gaussian splatting for efficient 3d mesh reconstruction and high-quality mesh rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5354–5363, 2024. 1, 5

[22] Ayaan Haque, Matthew Tancik, Alexei A Efros, Aleksander Holynski, and Angjoo Kanazawa. Instruct-nerf2nerf: Editing 3d scenes with instructions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 19740–19750, 2023. 6

[23] Peter Hedman, Julien Philip, True Price, Jan-Michael Frahm, George Drettakis, and Gabriel Brostow. Deep blending for free-viewpoint image-based rendering. *ACM Transactions on Graphics (ToG)*, 37(6):1–15, 2018. 6

[24] Marina Hernández-Bautista and FJ Melero. 3d hole filling using deep learning inpainting. *arXiv preprint arXiv:2407.17896*, 2024. 3

[25] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017. 7

[26] Alain Hore and Djemel Ziou. Image quality metrics: Psnr vs. ssim. In *2010 20th international conference on pattern recognition*, pages 2366–2369. IEEE, 2010. 7

[27] 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 Transactions on Graphics (TOG)*, 37(4):1–14, 2018. 3

[28] Yubin Hu, Sheng Ye, Wang Zhao, Matthieu Lin, Yuze He, Yu-Hui Wen, Ying He, and Yong-Jin Liu. Oˆ 2-recon: Completing 3d reconstruction of occluded objects in the scene with a pre-trained 2d diffusion model. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 2285–2293, 2024. 2

[29] Jiajun Huang and Hongchuan Yu. Gsdeformer: Direct cage-based deformation for 3d gaussian splatting. *arXiv preprint arXiv:2405.15491*, 2024. 2

[30] Shi-Sheng Huang, Guo Chen, CHEN LI HENG, and Hua Huang. Neuralindicator: Implicit surface reconstruction from neural indicator priors. In *Forty-first International Conference on Machine Learning*, 2024. 2

[31] Tianxin Huang, Hao Zou, Jinhao Cui, Xuemeng Yang, Mengmeng Wang, Xiangrui Zhao, Jiangning Zhang, Yi Yuan, Yifan Xu, and Yong Liu. Rfnet: Recurrent forward network for dense point cloud completion. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 12508–12517, 2021. 2

[32] Tianxin Huang, Zhiwen Yan, Yuyang Zhao, and Gim Hee Lee. Zero-shot point cloud completion via 2d priors. *arXiv preprint arXiv:2404.06814*, 2024. 2

[33] Tianyu Huang, Yihan Zeng, Hui Li, Wangmeng Zuo, and Rynson WH Lau. Dreamphysics: Learning physical properties of dynamic 3d gaussians with video diffusion priors. *arXiv preprint arXiv:2406.01476*, 2024. 2

[34] 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. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4220–4230, 2024. 2

[35] Sahil Jain, Avik Kuthiala, Prabhdeep Singh Sethi, and Prakanshul Saxena. Stylesplat: 3d object style transfer with gaussian splatting. *arXiv preprint arXiv:2407.09473*, 2024. 8

[36] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. In *International Conference on Learning Representations*, 2022. 4

[37] 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. 2, 3

[38] Ying Jiang, Chang Yu, Tianyi Xie, Xuan Li, Yutao Feng, Huamin Wang, Minchen Li, Henry Lau, Feng Gao, Yin Yang, et al. Vr-gs: A physical dynamics-aware interactive gaussian splatting system in virtual reality. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–1, 2024. 1, 2, 4, 5, 6, 8

[39] Yoni Kasten, Ohad Rahamim, and Gal Chechik. Point cloud completion with pretrained text-to-image diffusion models. *Advances in Neural Information Processing Systems*, 36, 2024. 2

[40] Michael Kazhdan. Adaptive multigrid solvers. `https://github.com/mkazhdan/PoissonRecon.git`, 2024. Version 18.41, accessed: 2024-11-02. 4

[41] Michael Kazhdan and Hugues Hoppe. Screened poisson surface reconstruction. *ACM Trans. Graph.*, 32(3), 2013. 4

[42] Bingxin Ke, Anton Obukhov, Shengyu Huang, Nando Metzger, Rodrigo Caye Daudt, and Konrad Schindler. Repurposing diffusion-based image generators for monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 6

[43] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics (ToG)*, 42(4):1–14, 2023. 1, 3, 8

[44] Justin Kerr, Chung Min Kim, Ken Goldberg, Angjoo Kanazawa, and Matthew Tancik. Lerf: Language embedded radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 19729–19739, 2023. 6

[45] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4015–4026, 2023. 2

[46] Gergely Klár, Theodore Gast, Andre Pradhana, Chuyuan Fu, Craig Schroeder, Chenfanfu Jiang, and Joseph Teran. Drucker-prager elastoplasticity for sand animation. *ACM Transactions on Graphics (TOG)*, 35(4):1–12, 2016. 2

[47] Jan Klein and Gabriel Zachmann. Point cloud collision detection. In *Computer Graphics Forum*, pages 567–576. Wiley Online Library, 2004. 1

[48] Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics (ToG)*, 36 (4):1–13, 2017. 5

[49] Lin Li, Fan Yang, Haihong Zhu, Dalin Li, You Li, and Lei Tang. An improved ransac for 3d point cloud plane segmentation based on normal distribution transformation cells. *Remote Sensing*, 9(5):433, 2017. 5

[50] 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. In *The Eleventh International Conference on Learning Representations*, 2023. 1, 2

[51] 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. In *The Eleventh International Conference on Learning Representations*, 2023. 5

[52] Fangfu Liu, Hanyang Wang, Shunyu Yao, Shengjun Zhang, Jie Zhou, and Yueqi Duan. Physics3d: Learning physical properties of 3d gaussians via video diffusion. *arXiv preprint arXiv:2406.04338*, 2024. 2

[53] Shaowei Liu, Zhongzheng Ren, Saurabh Gupta, and Shenlong Wang. Physgen: Rigid-body physics-grounded image-to-video generation. In *European Conference on Computer Vision*, pages 360–378. Springer, 2025. 3, 6, 7

[54] Zhiheng Liu, Hao Ouyang, Qiuyu Wang, Ka Leong Cheng, Jie Xiao, Kai Zhu, Nan Xue, Yu Liu, Yujun Shen, and Yang Cao. Infusion: Inpainting 3d gaussians via learn-

ing depth completion from diffusion prior. *arXiv preprint arXiv:2404.11613*, 2024. 3, 5

[55] William E. Lorensen and Harvey E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. *SIGGRAPH Comput. Graph.*, 21(4):163–169, 1987. 4

[56] Wenjia Lu, Zuoqiang Shi, Jian Sun, and Bin Wang. Surface reconstruction based on the modified gauss formula. *ACM Transactions on Graphics (TOG)*, 38(1):1–18, 2018. 2

[57] Shaojie Ma, Yawei Luo, and Yi Yang. Reconstructing and simulating dynamic 3d objects with mesh-adsorbed gaussian splatting. *arXiv preprint arXiv:2406.01593*, 2024. 2

[58] Miles Macklin. Warp: A high-performance python framework for gpu simulation and graphics. https://github.com/nvidia/warp, 2022. NVIDIA GPU Technology Conference (GTC). 5

[59] Miles Macklin, Matthias Müller, and Nuttapong Chentanez. Xpbd: position-based simulation of compliant constrained dynamics. In *Proceedings of the 9th International Conference on Motion in Games*, page 49–54, New York, NY, USA, 2016. Association for Computing Machinery. 2

[60] Willi Menapace, Aliaksandr Siarohin, Stéphane Lathuilière, Panos Achlioptas, Vladislav Golyanik, Sergey Tulyakov, and Elisa Ricci. Promptable game models: Text-guided game simulation via masked diffusion models. *ACM Trans. Graph.*, 43(2), 2024. 1

[61] Lukas Meyer, Floris Erich, Yusuke Yoshiyasu, Marc Stamminger, Noriaki Ando, and Yukiyasu Domae. Pegasus: Physical enhanced gaussian splatting simulation system for 6dof object pose dataset generation. *IROS*, 2024. 7

[62] 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. *Communications of the ACM*, 65(1):99–106, 2021. 1, 3

[63] Ashkan Mirzaei, Riccardo De Lutio, Seung Wook Kim, David Acuna, Jonathan Kelly, Sanja Fidler, Igor Gilitschenski, and Zan Gojcic. Reffusion: Reference adapted diffusion models for 3d scene inpainting. *arXiv preprint arXiv:2404.10765*, 2024. 3

[64] Mayank Mittal, Calvin Yu, Qinxi Yu, Jingzhou Liu, Nikita Rudin, David Hoeller, Jia Lin Yuan, Ritvik Singh, Yunrong Guo, Hammad Mazhar, et al. Orbit: A unified simulation framework for interactive robot learning environments. *IEEE Robotics and Automation Letters*, 8(6):3740–3747, 2023. 1

[65] Vismay Modi, Nicholas Sharp, Or Perel, Shinjiro Sueda, and David IW Levin. Simplicits: Mesh-free, geometry-agnostic elastic simulation. *ACM Transactions on Graphics (TOG)*, 43(4):1–11, 2024. 2

[66] Rania Moussa, Amira Alghazaly, Nebras Althagafi, Rawah Eshky, and Sary Borzangy. Effectiveness of virtual reality and interactive simulators on dental education outcomes: systematic review. *European journal of dentistry*, 16(01): 14–31, 2022. 1

[67] Richard A Newcombe, Shahram Izadi, Otmar Hilliges, David Molyneaux, David Kim, Andrew J Davison, Pushmeet Kohi, Jamie Shotton, Steve Hodges, and Andrew Fitzgibbon. Kinectfusion: Real-time dense surface mapping and tracking. In *2011 10th IEEE international symposium on mixed and augmented reality*, pages 127–136. Ieee, 2011. 5

[68] Songyou Peng, Chiyu Jiang, Yiyi Liao, Michael Niemeyer, Marc Pollefeys, and Andreas Geiger. Shape as points: A differentiable poisson solver. *Advances in Neural Information Processing Systems*, 34:13032–13044, 2021. 2

[69] Tobias Pfaff, Meire Fortunato, Alvaro Sanchez-Gonzalez, and Peter W. Battaglia. Learning mesh-based simulation with graph networks. In *International Conference on Learning Representations*, 2021. 1

[70] Nico Pietroni, Marco Tarini, and Paolo Cignoni. Almost isometric mesh parameterization through abstract domains. *IEEE Transactions on Visualization and Computer Graphics*, 16(4):621–635, 2009. 4

[71] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image synthesis. In *The Twelfth International Conference on Learning Representations*, 2023. 3

[72] Ri-Zhao Qiu, Ge Yang, Weijia Zeng, and Xiaolong Wang. Language-driven physics-based scene synthesis and editing via feature splatting. In *European Conference on Computer Vision (ECCV)*, 2024. 2

[73] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, et al. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*, 2024. 4

[74] Guodong Rong, Byung Hyun Shin, Hadi Tabatabaee, Qiang Lu, Steve Lemke, Mārtiņš Možeiko, Eric Boise, Geehoon Uhm, Mark Gerow, Shalin Mehta, et al. Lgsvl simulator: A high fidelity simulator for autonomous driving. In *2020 IEEE 23rd International conference on intelligent transportation systems (ITSC)*, pages 1–6. IEEE, 2020. 1

[75] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 3

[76] Johannes Lutz Schönberger, Enliang Zheng, Marc Pollefeys, and Jan-Michael Frahm. Pixelwise view selection for unstructured multi-view stereo. In *European Conference on Computer Vision (ECCV)*, 2016. 3

[77] Dario Seyb, Alec Jacobson, Derek Nowrouzezahrai, and Wojciech Jarosz. Non-linear sphere tracing for rendering deformed signed distance fields. *ACM Transactions on Graphics*, 38(6), 2019. 1

[78] Alexey Stomakhin, Craig Schroeder, Lawrence Chai, Joseph Teran, and Andrew Selle. A material point method for snow simulation. *ACM Transactions on Graphics (TOG)*, 32(4): 1–10, 2013. 2

[79] Roman Suvorov, Elizaveta Logacheva, Anton Mashikhin, Anastasia Remizova, Arsenii Ashukha, Aleksei Silvestrov, Naejin Kong, Harshith Goka, Kiwoong Park, and Victor Lempitsky. Resolution-robust large mask inpainting with fourier convolutions. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 2149–2159, 2022. 2, 3, 5

[80] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *Computer Vision–ECCV*

2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16, pages 402–419. Springer, 2020. 7

[81] Hanchen Wang, Qi Liu, Xiangyu Yue, Joan Lasenby, and Matt J Kusner. Unsupervised point cloud pre-training via occlusion completion. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9782–9792, 2021. 2

[82] Miaowei Wang and Daniel Morris. Self-annotated 3d geometric learning for smeared points removal. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3494–3503, 2024. 5

[83] Yuxin Wang, Qianyi Wu, Guofeng Zhang, and Dan Xu. Gscream: Learning 3d geometry and feature consistent gaussian splatting for object removal. In *ECCV*, 2024. 3, 6, 8

[84] Yujie Wei, Shiwei Zhang, Zhiwu Qing, Hangjie Yuan, Zhiheng Liu, Yu Liu, Yingya Zhang, Jingren Zhou, and Hongming Shan. Dreamvideo: Composing your dream videos with customized subject and motion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6537–6549, 2024. 6

[85] William F. Whitney, Tatiana Lopez-Guevara, Tobias Pfaff, Yulia Rubanova, Thomas Kipf, Kim Stachenfeld, and Kelsey R. Allen. Learning 3d particle-based simulators from rgb-d videos. In *International Conference on Learning Representations (ICLR)*, 2024. Available at: https://openreview.net/forum?id=4rBEgZCubP. 2

[86] Eugene Wigner. *Group theory: and its application to the quantum mechanics of atomic spectra*. Elsevier, 2012. 5

[87] Lintai Wu, Xianjing Cheng, Junhui Hou, Yong Xu, and Huanqiang Zeng. Self-supervised 3d point cloud completion via multi-view adversarial learning. *arXiv preprint arXiv:2407.09786*, 2024. 2

[88] Tong Wu, Jia-Mu Sun, Yu-Kun Lai, Yuewen Ma, Leif Kobbelt, and Lin Gao. Deferredgs: Decoupled and editable gaussian splatting with deferred shading. *arXiv preprint arXiv:2404.09412*, 2024. 2

[89] Tianyi Xie, Zeshun Zong, Yuxing Qiu, Xuan Li, Yutao Feng, Yin Yang, and Chenfanfu Jiang. Physgaussian: Physics-integrated 3d gaussians for generative dynamics. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4389–4398, 2024. 1, 2, 3, 5, 6, 7, 8

[90] Rui Xu, Zhiyang Dou, Ningna Wang, Shiqing Xin, Shuangmin Chen, Mingyan Jiang, Xiaohu Guo, Wenping Wang, and Changhe Tu. Globally consistent normal orientation for point clouds by regularizing the winding-number field. *ACM Trans. Graph.*, 42(4), 2023. 4

[91] Tian-Xing Xu, Wenbo Hu, Yu-Kun Lai, Ying Shan, and Song-Hai Zhang. Texture-gs: Disentangling the geometry and texture for 3d gaussian splatting editing. In *European Conference on Computer Vision*, pages 37–53. Springer, 2025. 2

[92] Chen Yang, Sikuang Li, Jiemin Fang, Ruofan Liang, Lingxi Xie, Xiaopeng Zhang, Wei Shen, and Qi Tian. Gaussianobject: High-quality 3d object reconstruction from four views with gaussian splatting. *ACM Transactions on Graphics*, 43 (6), 2024. 3

[93] Yunhan Yang, Yukun Huang, Yuan-Chen Guo, Liangjun Lu, Xiaoyang Wu, Lam Edmund Y., Yan-Pei Cao, and Xihui Liu. Sampart3d: Segment any part in 3d objects. *arXiv preprint arXiv:2411.07184*, 2024. 8

[94] Mingqiao Ye, Martin Danelljan, Fisher Yu, and Lei Ke. Gaussian grouping: Segment and edit anything in 3d scenes. In *ECCV*, 2024. 4

[95] Thomas Y Yeh, Petros Faloutsos, and Glenn Reinman. Enabling real-time physics simulation in future interactive entertainment. In *Proceedings of the 2006 ACM SIGGRAPH Symposium on Videogames*, pages 71–81, 2006. 1

[96] Yu Yin, Kamran Ghasedi, HsiangTao Wu, Jiaolong Yang, Xin Tong, and Yun Fu. Nerfinvertor: High fidelity nerf-gan inversion for single-shot real image animation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8539–8548, 2023. 1

[97] Xin-Yi Yu, Jun-Xin Yu, Li-Bo Zhou, Yan Wei, and Lin-Lin Ou. Instantstylegaussian: Efficient art style transfer with 3d gaussian splatting. *arXiv preprint arXiv:2408.04249*, 2024. 8

[98] Dingxi Zhang, Yu-Jie Yuan, Zhuoxun Chen, Fang-Lue Zhang, Zhenliang He, Shiguang Shan, and Lin Gao. Stylizedgs: Controllable stylization for 3d gaussian splatting. *arXiv preprint arXiv:2404.05220*, 2024. 8

[99] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. 7