

MeGAS: Thermomechanical Dynamic Gaussian Splatting for Thermophysical Scene Editing

Zesong Yang^{1*}, Yuanhang Lei^{1*}, Liyuan Cui¹, Yihang Chen¹, Jiaer Huang¹, Boming Zhao¹, Peter Yichen Chen², Hujun Bao¹, and Zhaopeng Cui^{1†}

¹State Key Laboratory of CAD&CG, Zhejiang University, China

²University of British Columbia, Canada

Project Page: zju3dv.github.io/MeGAS




Fig. 1: Thermophysical editing and photorealistic rendering of real scenes. MeGAS integrates thermomechanical dynamics into 3D Gaussian Splatting, enabling physically grounded editing with temporally evolving temperature fields, while maintaining smooth and continuous geometry under large deformations.

Abstract. Recent advances integrate physically grounded Newtonian dynamics with neural rendering frameworks, narrowing the gap between photorealistic scene reconstruction and physics-based animation. However, existing approaches focus on mechanically driven dynamics while neglecting temperature, a fundamental yet invisible physical factor underlying phenomena such as melting, solidification, and other thermomechanical processes. In this paper, we propose **MeGAS**, a novel framework that incorporates thermomechanical phase-change dynamics into 3D Gaussian Splatting (3DGS). Specifically, we propose a new thermomechanical dynamic Gaussian Splatting representation that augments 3DGS with temperature attributes and employs a heat advection-diffusion solver with MPM dynamics incorporating phase transitions, enabling physically plausible and visually realistic synthesis of thermophysical phenomena. Furthermore, a new topology-adaptive Gaussian rendering strategy is proposed to mitigate cracking and floaters under extreme

* Equal contribution. † Corresponding Author.

deformation. Extensive experiments demonstrate that MeGAS produces physically consistent thermomechanical behavior while maintaining high-fidelity photorealistic rendering, advancing toward physics-integrated world models.

Keywords: 3D Gaussian Splatting · 3D Scene Editing · Physics-based Animation

1 Introduction

Recent advances in neural scene representations, such as Neural Radiance Fields (NeRF) [31] and 3D Gaussian Splatting (3DGS) [19], have significantly advanced photorealistic scene reconstruction and rendering from real-world imagery, substantially promoting the development toward visually faithful digital twins and world models. While these approaches achieve impressive visual fidelity, a practical open-world simulator requires not only visual realism but also the controllable synthesis of physically plausible phenomena that faithfully adhere to the physical laws. To this end, recent research has integrated physical simulators into neural rendering, enabling the simulation of diverse phenomena such as elastic deformation [12, 24, 38, 42, 44, 48], fluid animation [11], and weather effects [10, 25, 37].

However, most existing approaches primarily focus on mechanically driven dynamics, while thermomechanical dynamics remain largely unexplored. In the physical world, temperature, an invisible yet fundamental quantity, governs a broad spectrum of visually perceivable phenomena, ranging from the melting of ice cream to the solidification of lava. Although several works have incorporated thermal information, they typically treat temperature as an auxiliary sensing modality for multispectral reconstruction [27, 29, 33, 46] or estimate spatio-temporal temperature fields restricted in static scenes [9, 45]. Nevertheless, since thermomechanical dynamics are among some of the most intricate natural phenomena, integrating thermomechanical effects into neural rendering frameworks remains highly challenging. **First**, thermomechanical dynamics necessitate not only the enforcement of momentum conservation but also the coupled modeling of heat transfer and phase transitions to produce physically plausible outcomes. **Second**, thermomechanical coupling (e.g., in melting) introduces extreme deformations and topological changes, posing a significant challenge for neural rendering representations that typically assume continuous or smoothly deforming geometry. Such extreme deformations cause sharp and spiky artifacts in splatting-based rendering, and the object surface tends to crack and expose interior regions, further breaking the geometric continuity for stable rendering. Consequently, a straightforward combination of existing Gaussian pipelines with thermomechanical dynamics fails to ensure visually stable results.

To address these challenges, we propose **MeGAS**, a physics-integrated neural rendering framework that integrates thermomechanical dynamics with 3D Gaussian Splatting, enabling controllable and thermophysically grounded editing of real scenes as shown in Fig. 1. To this end, we first propose a new thermomechanical dynamic Gaussian Splatting representation that augments each

Gaussian with a per-splat temperature attribute and supports direct rendering of a consistent thermal field within the standard 3DGS formulation. Temperatures can be flexibly initialized and manipulated through user-specified heat sources, providing intuitive and physically grounded editing control. To evolve this thermal field, we adopt a heat advection-diffusion solver defined on a simulation grid. The solver propagates heat diffusion through the reconstructed scene, yielding a smooth, physical temperature distribution. We further couple this thermal field to a phase-change-aware constitutive model switching, enabling particle-level melting behavior: once a particle’s temperature exceeds the melting threshold, its constitutive model transitions from an elastic solid to a viscoplastic formulation, producing physically grounded and spatially localized melting dynamics.

Since melting induces large-scale deformation and surface topology changes, directly applying such extreme motion to vanilla or PhysGaussian-style [44] 3DGS leads to severe artifacts such as cracking, floaters, and interior exposure. We therefore design a novel Topology-Adaptive Gaussian Rendering strategy to ensure stable rendering throughout the melting process. Specifically, during reconstruction, we regularize Gaussian anisotropy to suppress excessively elongated kernels, enforce volumetric consistency, and prune internal Gaussians that would otherwise become visible under deformation. During animation, we perform an online surface-aware refinement that combines adaptive densification with local Gaussian kernel refitting guided by the evolving surface, which seals emerging cracks and maintains a coherent splat distribution as the object undergoes drastic thermomechanical change.

Our main contributions are summarized as follows:

- We introduce **MeGAS**, a novel framework that integrates thermomechanical phase-change dynamics into 3D Gaussian Splatting (3DGS), enabling physically plausible and photorealistic synthesis of thermophysical phenomena from real-world imagery.
- We propose a new thermomechanical dynamic Gaussian Splatting representation that jointly models appearance and a controllable thermal field, while integrating heat advection-diffusion and temperature-controlled constitutive model switching, enabling melting dynamics controlled by phase change.
- We develop a novel topology-adaptive Gaussian Rendering strategy with anisotropy regularization, internal-free filtering, and online surface-aware refinement to effectively mitigate cracking, floaters, and interior exposure under extreme deformation.
- Extensive experiments on real and synthetic scenes demonstrate that our method achieves thermophysically consistent scene editing while preserving photorealistic rendering quality.

2 Related Work

Physics-Integrated Neural Rendering. Recent studies have integrated physical simulators into neural rendering frameworks, bridging photorealistic recon-




Fig. 2: Thermo-geometry evolution with MeGAS. Under user-specified heating source, MeGAS performs thermomechanically consistent phase-change simulation, steering the evolution of temperature fields and shape and delivering controllable, targeted melting without compromising geometric smoothness.

struction with physically grounded realism. Early works [12, 24] embed continuum and elastodynamic models into NeRF, enabling geometry-consistent simulation of elastic and interactive deformations. Subsequent researches [18, 23, 26, 35, 40, 42, 44, 48] adapt 3D Gaussian Splatting (3DGS) to model mechanical dynamics while preserving rendering realism. Beyond elasticity, [11] unifies position-based dynamics with 3DGS for fluid animation, and [10, 25, 37] synthesize environmental effects such as rainfall, snowfall and fire within neural rendering frameworks. Despite these advances, existing approaches primarily focus on *mechanically driven* or *environmentally induced* phenomena. In contrast, many natural phenomena are *thermally driven*, where invisible temperature variations govern visually perceivable behaviors such as softening and melting. Our work makes the first exploration of thermomechanical phase-change dynamics in neural rendering, addressing this underexplored yet fundamental dimension of physics integration.

Thermal-Integrated 3DGS. Recent works extend neural rendering to thermal modalities by incorporating infrared imagery for multimodal scene reconstruction. [27, 29, 33, 46] leverage paired RGB-thermal inputs to improve novel view synthesis via multispectral supervision, yet they rely purely on image-level supervision without explicit thermodynamic modeling. [9, 43, 45] couple 3DGS with simplified heat diffusion to estimate spatiotemporal temperature variations, achieving dynamic thermal field reconstruction but still limited to static scenes without considering more complex thermodynamic phenomena. In contrast, we integrate a physically grounded thermomechanical model into 3DGS, capturing the coupled heat transfer and deformation dynamics for realistic melting synthesis.

3 Preliminaries

3D Gaussian Splatting. Given posed RGB image sequences as input, we first reconstruct the scene with 3D Gaussian Splatting (3DGS) [19]. 3DGS represents the scene with a set of explicit 3D Gaussians, each parameterized by a center μ and a full covariance matrix Σ :

$$G(x) = e^{-\frac{1}{2}(x-\mu)^\top \Sigma^{-1}(x-\mu)}. \quad (1)$$

For differentiable rendering optimization, Σ is decomposed into a scaling matrix S and rotation matrix R , i.e., $\Sigma = RSS^T R^T$, where S and R are stored by a 3D scale vector s and a quaternion q respectively. Given a viewing transformation W , the 3D Gaussians are projected to the image plane, and we obtain the 2D covariance matrix Σ' and 2D center location μ' as:

$$\Sigma' = JW\Sigma W^T J^T, \quad \mu' = JW\mu, \quad (2)$$

where J is the Jacobian of the affine approximation of the projective transformation. Then we can use the point-based volume rendering to render the color C of each pixel with N ordered 3D Gaussians:

$$C = \sum_{i \in N} T_i c_i \alpha_i, \quad T_i = \prod_{j=1}^{i-1} (1 - \alpha_j), \quad (3)$$

with alpha α_i defined as:

$$\alpha_i = o_i e^{-\frac{1}{2}(x - \mu'_i)^\top \Sigma'^{-1}_i (x - \mu'_i)}. \quad (4)$$

Here, o_i and c_i denote the learned opacity and color parameterized by spherical harmonics for each Gaussian. Besides, we define the normal n_i as the shortest axis direction.

Constitutive Governing Equations. We use the hybrid Eulerian-Lagrangian framework, Material Point Method (MPM) [15, 17], to simulate continuum dynamics. The governing equations follow the conservation of mass and momentum:

$$\frac{D\rho}{Dt} + \rho \nabla \cdot \mathbf{v} = 0, \quad \rho \frac{D\mathbf{v}}{Dt} = \nabla \cdot \boldsymbol{\sigma} + \mathbf{f}^{ext}, \quad (5)$$

where ρ , $\mathbf{v}$, and $\boldsymbol{\sigma}$ denote the density, velocity, and Cauchy stress tensor respectively, and $\mathbf{f}^{ext}$ is the external force. The stress tensor $\boldsymbol{\sigma}(\mathbf{x}, t)$ is derived from the hyperelastic constitutive model as:

$$\boldsymbol{\sigma}(\mathbf{x}, t) = \frac{1}{\det(\mathbf{F})} \frac{\partial \Psi}{\partial \mathbf{F}^E} (\mathbf{F}^E)^T, \quad (6)$$

where $\Psi : \mathbb{R}^{3 \times 3} \rightarrow \mathbb{R}$ is the hyperelastic energy density function, and $\mathbf{F}^E \in \mathbb{R}^{3 \times 3}$ represents the elastic part of the deformation gradient $\mathbf{F}$. By specifying different formulations of Ψ , various constitutive models can be incorporated to characterize distinct material behaviors.

4 Methods

As shown in Fig. 3, given posed RGB image sequences as input, we first reconstruct the scene with 3D Gaussian Splatting (3DGS) [19]. Since melting serves as a representative and visually salient stress test of thermomechanical coupling, we adopt it as the primary instantiation to present our framework. To simulate temperature-driven editing, we propose a thermomechanical dynamic Gaussian Splatting representation. Each Gaussian is augmented with a temperature attribute and we couple 3DGS to an MPM simulator with a heat advection-diffusion solver and a phase-change-aware constitutive model switching, enabling melting dynamics at the particle level. To maintain stable rendering




Fig. 3: System overview. Starting from posed RGB sequences, we reconstruct the scene with 3DGS. To mitigate floaters and interior artifacts under extreme deformation, we introduce an interior-free uniform-Gaussian regularization that suppresses anisotropy and prunes unsupported interior splats. We then provide volumetric support via ray-tracing-based internal filling. Next, we augment 3DGS with per-Gaussian temperature and, driven by user-specified heat sources, MeGAS updates temperatures using a grid-based heat advection-diffusion solver within an MPM simulator, and performs temperature-dependent phase-aware constitutive switching. Deformations are mapped back to 3DGS, and we apply adaptive surface densification to handle topology change and seal cracks. Finally, we synthesize physically plausible thermomechanical dynamics with photorealistic appearance and temporally smooth geometry.

under extreme deformation, we introduce a topology-adaptive Gaussian rendering strategy to obtain a uniform-distributed, volume-consistent, floater-free 3DGS, and incorporate a Ray-Tracing-based internal filling strategy to provide volumetric support for MPM. Finally, to prevent surface cracking during large deformations, we employ implicit-surface-guided densification that adaptively fills hollow regions with smoothly aligned Gaussian splats.

4.1 Thermomechanical Dynamic Gaussian Splats

We augment each Gaussian with a temperature attribute $T_i \in \mathbb{R}$, and the corresponding temperature rendering of each pixel is obtained following Eq. 3:

$$\hat{T} = \sum_{i \in N} w_i T_i, \quad (7)$$

where w_i are the Gaussian weights along the viewing ray.

Thermal Advection-Diffusion Solver. To drive melting, we require a physically grounded evolution of the temperature field. Instead of heuristically editing, we solve heat diffusion on a simulation grid. Specifically, we model heat transport with the advection-diffusion equation (ADE) [3]:

$$\frac{\partial T}{\partial t} + \mathbf{u} \cdot \nabla T = \alpha \nabla^2 T + S, \quad (8)$$

where T is the temperature field, $\mathbf{u}$ is the thermal velocity field, α is the thermal diffusivity, and S represents external heat sources. The velocity $\mathbf{u}$ is modeled by the Navier-Stokes equations (NSE) [5]:

$$\rho \left(\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} \right) = -\nabla p + \mu \nabla^2 \mathbf{u} + \mathbf{f}, \quad (9)$$

with density ρ , pressure p , dynamic viscosity μ , and external forces $\mathbf{f}$.

Algorithm 1 Thermomechanical MPM for Thermophysical Dynamics

```

1: Notation:  $p, i$ : particle & grid indices;  $m$ : mass;  $x$ : position;  $w_{ip}$ : P2G2P
   weight function;  $T$ : temperatures;  $\mathbf{F}^E$ : elastic deformation gradient;  $\tau$ : Kirch-
   hoff stress;
2: while time_step < total_steps do
3:   for all grid  $i$  do                                     ▷ Particle-to-Grid Transfer
4:     Transfer Mass & Momentum to Grid via APIC
5:      $m_i T_i \leftarrow \sum_p w_{ip} m_p (T_p + (\mathbf{x}_i - \mathbf{x}_p) \cdot \nabla T_p)$   ▷ Transfer Thermal Field
6:   end for
7:   for all grid  $i$  do                                     ▷ Grid Update
8:     Update Grid Velocities
9:     Update Thermal Boundary
10:     $T_i = \text{LBM\_Update}(T_i)$                                ▷ Thermal Evolution
11:  end for
12:  for all particle  $p$  do                                   ▷ Grid-to-Particle Transfer
13:    Interpolate Velocity from Grid
14:    Update Particle Mechanical Properties
15:     $T_p \leftarrow \sum_i w_{ip} T_i$ 
16:     $\nabla T_p \leftarrow \sum_i T_i (\nabla w_{ip})^T$ 
17:    if  $T_p > T_{\text{melt}}$  then                                   ▷ Phase Transition Update
18:       $\{\tau_p, \mathbf{F}_p^E\} \leftarrow \tau_{\text{liquid}}(\mathbf{F}_p^E)$            ▷ Switch Material
19:    else
20:       $\{\tau_p, \mathbf{F}_p^E\} \leftarrow \tau_{\text{solid}}(\mathbf{F}_p^E)$ 
21:    end if
22:  end for
23: end while

```

To advance the coupled system (8)–(9), we adopt the Lattice Boltzmann Method (LBM) that solves the flow and temperature fields and naturally captures both diffusion and advection [39]. Additionally, thermal buoyancy is incorporated via the Boussinesq approximation as a body force [13]:

$$\mathbf{F}_b = \beta \rho_0 (T - T_0) \mathbf{g}, \quad (10)$$

where β is the thermal expansion coefficient, ρ_0 the reference density at temperature T_0 , and $\mathbf{g}$ gravity. We add $\mathbf{F}_b$ to $\mathbf{f}$ in (9) for correction, avoiding explicit density updates while enabling accurate temperature-driven convection. See **supplementary material** for more details.

Thermal-Driven Material Point Method. While we have defined how temperature is represented and evolved, building upon the standard MPM framework [15, 17, 49], we extend the particle and grid states with thermal attributes, enabling temperature-dependent material behavior. During each time step, in addition to the standard MPM updates of mechanical properties, particle temperatures are transferred to the grid (P2G), followed by thermal boundary updates and temperature diffusion on the grid solved via a Lattice Boltzmann Method (LBM). The updated grid temperature is then interpolated back to the particles




Fig. 4: Naïve combination fails under extreme topology changes. Our thermo-mechanical augmentation can directly apply to Internet-pretrained splats, nevertheless, naïve 3DGS cracks under large deformations. With low-cost self fine-tuning (1000 steps $\approx$ 10s), **MeGAS** makes pretrained splats robust and yields smoother geometry.

(G2P). If the particle temperature exceeds a predefined melting threshold T_{melt} , we switch its material phase and update the constitutive model accordingly. For the solid state, we adopt a StVK elasticity model [20], while the melted state can be treated as a viscoplastic material and we model it using a Herschel–Bulkley plasticity formulation [47]. The thermal boundaries are user-defined heat source regions that directly control the spatial extent and intensity of melting, as shown in Fig. 2. The overall thermomechanical pipeline is summarized in Algorithm 1 and an expanded version with details is provided in **supplementary material**. **Mapping Deformation to 3D Gaussians.** After each simulation step t , following [44], the elastic deformation gradient $\mathbf{F}_p^{E,t}$ for splat G_p updates the covariance via an affine push-forward:

$$\Sigma_p^t = \mathbf{F}_p^{E,t} \Sigma_p^0 (\mathbf{F}_p^{E,t})^\top. \quad (11)$$

Then we decompose $\Sigma_p^t = \mathbf{R}_p^t (\mathbf{S}_p^t \mathbf{S}_p^t)^T \mathbf{R}_p^t$ to obtain the updated rotation matrix $\mathbf{R}_p^t$ and scaling matrix $\mathbf{S}_p^t$, and the orientation-dependent appearance is updated by the relative rotation $\Delta \mathbf{R}_p^t = \mathbf{R}_p^t \mathbf{R}_p^0{}^T$: we rotate the view direction d for SH color as: $c_p^t(d) = c_p^0((\Delta \mathbf{R}_p^t)^\top d)$.

4.2 Topology-Adaptive Gaussian Rendering

After obtaining particle-level melting animation through our thermomechanical MPM-LBM simulation, we achieve a unified framework for elastoplastic simulation and photorealistic rendering. However, since melting entails extreme non-rigid deformation and topology change (e.g., surface cracking, internal exposure), naïvely applying PhysGaussian under such conditions leads to severe visual artifacts, e.g., hollow surfaces and needle-like internal floaters, as shown in Fig. 4. To address these issues, we propose a topology-adaptive 3D Gaussian Splatting (3DGS) framework specifically designed for stable rendering under large-scale melting deformation with the following innovations.

Internal-Free Uniform Gaussian Regularization. The original 3D Gaussian Splatting (3DGS) reconstructs object details through outside-in multi-view photometric supervision, which leads to undesirable internal distribution of redundant Gaussians, and consequently degrades rendering quality when hidden Gaussians are exposed under large deformation. Additionally, the reconstructed surface Gaussians typically exhibit highly irregular anisotropy and non-uniform spatial distribution, resulting in needle-like visual artifacts during melting.

We build upon the planar-based Gaussian Splatting [16] and further enhance the training stage. Previous work [4] observes that enforcing a delta-like per-pixel alpha weight distribution along each ray effectively suppresses floating artifacts in volumetric rendering. Inspired by this observation, since 3DGS employs an explicit point-based representation, we propose a direct yet effective strategy—pruning low-contribution Gaussians during training. Specifically, we traverse all training views and record the maximum weight $w_i^{\max}$ of each Gaussian G_i across all pixels. Then we prune Gaussians whose $w_i^{\max}$ falls below a predefined threshold τ_{prune} which contribute negligibly to the rendered appearance: $\mathcal{G}' = \{G_i \in \mathcal{G} \mid w_i^{\max} > \tau_{\text{prune}}\}$. The remaining subset $\mathcal{G}'$ is subsequently fine-tuned until convergence, eliminating redundant internal Gaussians and continuing the original 3DGS training pipeline.

Irregular anisotropic Gaussians often lead to needle-like kernels or large volumetric discrepancies, which deteriorate rendering stability under large deformation. Following [11, 44], we adopt an anisotropy loss that penalizes excessive scale ratios and encourages the kernel to approach a disk-like shape:

$$\mathcal{L}_{\text{aniso}} = \frac{1}{|\mathcal{P}|} \sum_{p \in \mathcal{P}} \max \left(\frac{S_p^1}{S_p^2} - a, 0 \right), \quad (12)$$

where a is a ratio threshold, and $\mathbf{S}_p = \{S_p^1, S_p^2, S_p^3\}$ denotes the principal scalings of Gaussian G_p , with S_p^1 being the largest and S_p^3 the smallest scaling factor.

We further apply scale clamping with a threshold τ_{scale} to prevent large Gaussian volumetric discrepancies: $S_p \leftarrow \min(S_p, \tau_{\text{scale}})$. Densityfication in 3DGS naturally compensates for potential sparse regions caused by the clamping, yielding a uniform, volume-consistent, and surface-aligned Gaussian distribution. As illustrated in Fig. 8, our regularization produces a uniform Gaussian distribution and reduces internal floaters compared to the regularization used in PhysGaussian [44] and VR-DoH [30] without compromising rendering fidelity.

Ray-Tracing-Based Internal Filling.

To enable a physically consistent melting simulation, the object interior needs to be properly filled for volumetric support. Since our reconstruction is internal-floater-free, we adopt a simple yet robust Gaussian-ray-tracing-based strategy [32] to detect interior regions. For each sampled ray, if the accumulated alpha mask indicates full opacity and the blended normal aligns with the ray direction, the ray is classified as interior-originated. As shown in Fig. 5, we divide the object’s bounding box into uniform grids, each grid casts rays along principal directions. Grids with a majority of interior-originated rays are supplemented into internal particles.




Fig. 5: Our ray-tracing internal filling. Starting from an interior-free, uniformly distributed GS model, we cast rays from uniform grids and detect interior samples by enforcing directional consistency with Gaussian ray-traced normals. This produces uniform, well-conditioned volumetric filling for physically plausible melting.

Grids with a majority of interior-originated rays are supplemented into internal particles.




Fig. 6: Implicit-surface-guided adaptive densification. Internal particles migrating to cracks are detected and projected onto an IMLS-fitted implicit surface to reorient Gaussian kernels with surface-aligned normals. Locality-aware densification seals emerging cracks and maintains coherent splat distributions under severe thermomechanical topology change. Integrating Gaussian Splatting with IMLS exports temporally more coherent and smoother surfaces while preserving fine geometric details.

Adaptive Surface Cracking Densification. Melting induces extreme non-rigid deformation and frequent topology changes, which often cause surface cracking and hollow artifacts, as shown in Fig. 6. We observe that the internal particles filled naturally migrate toward surface voids during the MPM evolution, partially compensating for the missing geometry. However, these internal particles lack valid Gaussian kernel parameters (covariance/appearance), thus cannot be rendered coherently. Let the initial surface Gaussians be denoted as $\mathcal{G}_s = \{G_i^s = (\mathbf{x}_i^s, \mathbf{R}_i^s, \mathbf{S}_i^s, \mathbf{o}_i^s)\}$. After deformation step t , the internal filled particles $\mathcal{P}_t = \{\mathbf{p}_j^t\}$ geometrically occupy the cracked regions of $\mathcal{G}_s$ but without valid kernel parameters. To reconstruct a coherent surface representation, we leverage the deformed surface Gaussians $\mathcal{G}_s^t = \Phi_t(\mathcal{G}_s)$ to guide the local Gaussian kernel fitting of nearby internal particles.

Surface Detection. We first identify which internal particles are suitable for use in cracking filling with surface detection. Specifically, we place virtual cameras uniformly distributed over 360° . Internal particles $\mathbf{p}_j^t$ are initialized as isotropic Gaussian splats G_j^p with opacity $\mathbf{o}_j = 1$ and radius r_j :

$$r_j = \kappa \min \left\{ \min_i \|\mathbf{p}_j^t - \mathbf{x}_i^s\|, \min_{k \neq j} \|\mathbf{p}_j^t - \mathbf{p}_k^t\| \right\}, \quad (13)$$

where $\kappa \in (0, 1)$ is a shrinkage ratio. For each pixel, we record the Gaussian splat G_i with maximal alpha contribution along the ray, and yields the surface subset $\{\hat{\mathcal{G}}_s^t, \hat{\mathcal{P}}^t\}$.

Implicit-Surface Guidance. As shown in Fig. 6, to restore plausible Gaussian kernel for $\hat{\mathcal{P}}^t$, we fit an implicit surface $\mathcal{S}_{\text{IMLS}}$ from $\hat{\mathcal{G}}_s^t$ via Implicit Moving Least Squares (IMLS) [2, 22, 34]:

$$f(\mathbf{x}) = \frac{\sum_i \phi(\|\mathbf{x} - \tilde{\mathbf{x}}_i^s\|) (\tilde{\mathbf{n}}_i^s)^\top (\mathbf{x} - \tilde{\mathbf{x}}_i^s)}{\sum_i \phi(\|\mathbf{x} - \tilde{\mathbf{x}}_i^s\|)}, \quad (14)$$




Fig. 7: Melting-style scene editing comparisons. Qualitative results with multi-view, temporally evolving melting edits across DGE, AutoVFX, Runway Gen-4.5, and our MeGAS show that prior methods lack physical realism, photorealism, or multi-view consistency, whereas MeGAS preserves photorealistic appearance and delivers physically plausible, temporally consistent melting. **Please refer to our supp. video for more vivid dynamic comparisons.**

where $\phi(\cdot)$ is a Gaussian weighting kernel and $\tilde{\mathbf{n}}_i^s$ is the normal derived from $\tilde{\mathcal{G}}_i^s \subset \tilde{\mathcal{G}}_s^t$. Each internal particle is then projected onto $\mathcal{S}_{\text{IMLS}}$ to obtain its surface normal $\mathbf{n}_j$ and rotation matrix $\mathbf{R}_j$, and we further assign color $\mathbf{c}_j$ as the average of K nearest Gaussians in $\tilde{\mathcal{G}}_s^t$, yielding a surface-aligned kernel:

$$G_j^p = (\mathbf{p}_j^t, \mathbf{R}_j, \mathbf{S}_j, \mathbf{o}_j, \mathbf{c}_j), \text{ where } \mathbf{S}_j = r_j \mathbf{I}. \quad (15)$$

This adaptive densification seamlessly re-covers cracked regions with correctly oriented internal Gaussians and yields smooth, stable rendering under extreme melting deformation. Meanwhile, the fitted implicit surface $\mathcal{S}_{\text{IMLS}}$ also enables direct extraction of high-quality meshes. See more details in **supplementary**.

5 Experiments

We evaluate our MeGAS on diverse synthetic and real-world scenes to demonstrate our capability to produce physically plausible, photorealistic thermophysical editing and robustness under extreme deformations.

5.1 Evaluation on Thermophysical Scene Editing

Datasets and Comparison. To evaluate MeGAS as a practical editing tool, we perform melting-style edits on in-the-wild scenes from MipNeRF360 [4] and Tanks and Temples [21] datasets. For each scene, we first obtain a high-quality static 3DGS reconstruction, and then compare our melting edits against three representative state-of-the-art approaches: 1) DGE [8], a text-driven diffusion-based editing method for 3DGS; 2) AutoVFX [14], which leverages an LLM [1] to generate scripts for external physics engines to perform mesh-based melting; 3) Runway Gen-4.5 [36], a leading commercial image-to-video generation model that augments a single input image with text-guided dynamic effects. **Notably**, existing physics-integrated neural rendering methods do not support thermophysical animation or phase-change-aware dynamics, therefore, we compare against these general editing baselines (including a commercial open-domain video generator) to contextualize our thermophysical editing. All methods are




Fig. 8: Extreme melting deformation comparisons. Internal cross-sections and dynamic frames show that Original 3DGS, PhysGaussian, and VR-DoH accumulate internal floaters, needle-like artifacts, and surface cracking under large deformations, whereas MeGAS yields volume-consistent, crack-free geometry and stable rendering throughout melting. **Please refer to supp. video for vivid dynamic comparison.**

provided with comparable editing instructions, and detailed prompts and parameter settings are reported in **supplementary material**.

Results. Qualitative comparisons are shown in Fig. 7. DGE fails to induce coherent dynamic melting behavior and exhibits limited physical realism. AutoVFX generates plausible mesh-based melting, but decoupled geometry/shading from radiance causes non-photorealistic compositions and inconsistencies with the original scene. Runway Gen-4.5 produces visually vivid effects, yet the motion is often physically implausible, object identities drift over time, and it lacks multi-view consistency. These limitations stem from the absence of explicit physical modeling, which makes it difficult to capture complex physical phenomena. In contrast, by integrating thermomechanical simulation with neural scene representations, MeGAS preserves the original photorealistic appearance and produces multi-view consistent, physically plausible thermophysical dynamics. **Please refer to supp. for more edited results on real-world data.**

User Study. We further conduct a user study as shown in Tab. 1 across physical plausibility, photorealism, and multi-view/temporal consistency, supporting our thermomechanical dynamic Gaussian Splatting modeling yields more convincing and reliable melting edits (see **supp.** for full config.).

Table 1: User Study. The rankings assigned by participants to each method are converted to integer scores (best=4, worst=1) and averaged over all videos and raters.

Methods	DGE [8]	AutoVFX [14]	Gen-4.5 [36]	Ours
Physical-Realism $\uparrow$	1.538	2.692	2.038	3.731
Photo-Realism $\uparrow$	2.038	2.154	2.154	3.654
Multi-View Consistency $\uparrow$	2.154	2.385	1.885	3.577

5.2 Evaluation on Extreme Melting Deformation

Datasets and Comparison. We evaluate rendering quality under extreme thermomechanical deformations with a set of Blender-synthetic objects and real-world scenes captured with an iPhone. We compare the following 3DGS-based deformation pipelines: 1) Original 3DGS [19]; 2) PhysGaussian [44], a physics-integrated 3DGS with isotropic constrained kernels; 3) VR-DoH [30], which en-




Fig. 9: Ablations on Topology-Adaptive strategy. Our topology-adaptive strategy yields an interior-free, uniformly distributed Gaussian representation with consistent volume, mitigating needle-like artifacts and internal floaters. The adaptive densification fills surface cracking during animation, and our ray-tracing-based internal filling uniformly populates the interior, providing robust volumetric support.

forces volume-consistency constraints. All baselines are coupled with the same MPM-LBM simulator, and differ only in how Gaussians are regularized, internally filled, or adapted to deformation.

Implementation Details. For each scene, we first reconstruct 3DGS with the respective methods, adopting the internal filling strategies proposed for each baseline (e.g., density-field-based schemes [44]) and employing our ray-tracing-based filling strategy for MeGAS. We then manually place identical heat sources across methods, and evolve the thermal field using the same LBM-based advection-diffusion solver with consistent thermal diffusivity, boundary conditions, and phase-transition parameters. The resulting thermal field drives MPM particles according to phase-aware constitutive models, and Gaussians are updated at each step via deformation-gradient-based transformations. All methods are simulated for the same duration with identical material and solver settings (see **supp. for full configurations and computational cost**).

Results. Fig. 8 visualizes internal cross-sections and dynamic frames under large-scale melting. Baseline methods accumulate internal floaters, produce severe needle-like artifacts, and suffer from surface cracking that exposes under-specified interiors. In contrast, our MeGAS yields an internal-floater-free and volume-consistent Gaussian distribution, with stable, crack-free surfaces refined by our adaptive surface densification. The resulting geometry remains smooth and coherent throughout the melting process, allowing direct extraction of high-quality meshes from the evolving Gaussian representation, as illustrated in Fig. 2.

5.3 Ablation Studies

Topology-Adaptive Strategy. As shown in Fig. 9, we progressively incorporate our topology-adaptive strategy into the backbone model [7]. Internal-Prune effectively suppresses internal floaters, while anisotropy loss and scale clamp-




Fig. 10: Heat diffusion with constitutive model switching. Phase-change-driven constitutive switching couples melting to temperature evolution, while different heat-source configurations yield distinct melting patterns.




Fig. 11: Ablations on phase-change switching. Without phase change, melting is triggered globally and collapses prematurely.




Fig. 12: Controllable melting via parameter adjustment. By modifying thermophysical parameters, users controllably achieve different melting effects.

ing yield a uniform, volume-consistent Gaussian distribution that avoids needle-like artifacts under large deformation. During animation, adaptive densification compensates for surface cracking and maintains smooth rendering. Since our representation is interior-free, directly applying density-field-based filling [44] leaves the interior hollow and prone to collapse, whereas our ray-tracing filling uniformly populates the volume and provides robust volumetric support.

Phase-Change Switching. As shown in Fig. 11, we demonstrate the necessity of phase-change-driven constitutive switching by comparing an init-melted strategy, where all Gaussians are initialized in the melted state and follow the viscoplastic constitutive model, while we incrementally change each Gaussian’s phase state based on the evolving temperature field. Init-melted strategy exhibits premature global slumping, while our temperature-controlled transition yields a more physically plausible melting progression. Moreover, we present visualizations of pure heat diffusion in Fig. 10, where the temperature field progressively propagates from high- to low-temperature regions over time. Building on this diffusion process, our phase-change-aware constitutive switching produces melting dynamics explicitly coupled to temperature evolution. We additionally show that interactive heat-source specification (Fig. 10) and thermophysical parameter tuning (Fig. 12) enable controllable melting with distinct patterns and effects.

5.4 Scalability to Complex Thermomechanical Interactions.

Built upon an MPM-based simulator, our framework supports broader thermomechanical interactions beyond single-object melting. In addition, the cooling-driven solidification shown in Fig. 13 demonstrates that our framework naturally supports bidirectional phase transitions, extending thermomechanical editing beyond melting. As illustrated in Fig. 14, we further showcase collisions with




Fig. 13: Bidirectional phase-change editing. Controlled heating induces melting, while subsequent cooling drives solidification, demonstrating consistent thermomechanical evolution across phase transitions and enabling general thermophysical editing.




Fig. 14: Scalability to complex scenes. Our method naturally supports multi-object interactions with multi-material and collisions (e.g., an ice cream ball). **Please view the dynamic videos in Adobe Acrobat Reader!**

real-scene environments, interactions with complex geometric colliders and elastic bodies, and heat-transfer interactions among multiple melting objects. Our method preserves the inherent extensibility of MPM while maintaining stable rendering under severe deformations, thereby broadening dynamic neural rendering toward richer thermomechanical dynamics modeling.

6 Conclusion

We have proposed MeGAS, a novel framework that integrates thermomechanical dynamics into 3D Gaussian Splatting. MeGAS jointly models appearance and a controllable thermal field, coupling a heat advection-diffusion solver with temperature-conditioned constitutive switching to realize phase-change-driven thermophysical editing. To handle extreme deformations that induce topology changes and surface cracking, we combine training-time anisotropy regularization and interior-free filtering with an online deformation-stage refinement that we fit local implicit surfaces to seal cracks and yield smooth renderings. To our knowledge, MeGAS provides the first exploration of thermomechanical phase-change dynamics within a neural rendering pipeline, addressing an underexplored yet fundamental dimension of physics integration.

Limitation. We currently do not model material-dependent appearance changes induced by phase transitions (e.g., PBR property variations during melting), as jointly coupling thermomechanics with dynamic material appearance remains highly challenging. A promising direction is to integrate our framework with video diffusion models [6, 28, 41] to enhance the final appearance rendering.

Acknowledgments

This work was partially supported by the National Natural Science Foundation of China (NSFC) under Grants 62572425 and 624B2132. We thank the anonymous reviewers for their valuable comments and suggestions.

References

1. Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F.L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., et al.: Gpt-4 technical report. arXiv preprint arXiv:2303.08774 (2023)
2. Alexa, M., Behr, J., Cohen-Or, D., Fleishman, S., Levin, D., Silva, C.T.: Point set surfaces. In: *Proceedings Visualization, 2001. VIS'01.* pp. 21–29. IEEE (2001)
3. Ashgriz, N., Mostaghimi, J.: An introduction to computational fluid dynamics. *Fluid Flow Handbook* **1**, 1–49 (2002)
4. Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P., Hedman, P.: Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022)
5. Batchelor, G.K.: *An Introduction to Fluid Dynamics*. Cambridge University Press (2000)
6. Blattmann, A., Dockhorn, T., Kulal, S., Mendelevitch, D., Kilian, M., Lorenz, D., Levi, Y., English, Z., Voleti, V., Letts, A., et al.: Stable video diffusion: Scaling latent video diffusion models to large datasets. arXiv preprint arXiv:2311.15127 (2023)
7. Chen, D., Li, H., Ye, W., Wang, Y., Xie, W., Zhai, S., Wang, N., Liu, H., Bao, H., Zhang, G.: Pgsr: Planar-based gaussian splatting for efficient and high-fidelity surface reconstruction. *IEEE Transactions on Visualization and Computer Graphics* (2024)
8. Chen, M., Laina, I., Vedaldi, A.: Dge: Direct gaussian 3d editing by consistent multi-view editing. In: *European Conference on Computer Vision*. pp. 74–92. Springer (2024)
9. Chen, Q., Shu, S., Bai, X.: Thermal3d-gs: Physics-induced 3d gaussians for thermal infrared novel-view synthesis. In: *European Conference on Computer Vision*. pp. 253–269. Springer (2024)
10. Dai, Q., Ni, X., Shen, Q., Chen, W., Chen, B., Chu, M.: Rainygs: Efficient rain synthesis with physically-based gaussian splatting. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 16153–16162 (2025)
11. Feng, Y., Feng, X., Shang, Y., Jiang, Y., Yu, C., Zong, Z., Shao, T., Wu, H., Zhou, K., Jiang, C., et al.: Gaussian splashing: Unified particles for versatile motion synthesis and rendering. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 518–529 (2025)
12. Feng, Y., Shang, Y., Li, X., Shao, T., Jiang, C., Yang, Y.: Pie-nerf: Physics-based interactive elastodynamics with nerf. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 4450–4461 (2024)
13. Guo, Z., Zheng, C., Shi, B.: Discrete lattice effects on the forcing term in the lattice boltzmann method. *Physical Review E* **65**(4), 046308 (2002)
14. Hsu, H.Y., Lin, C.H., Zhai, A.J., Xia, H., Wang, S.: Autovfx: Physically realistic video editing from natural language instructions. In: *2025 International Conference on 3D Vision (3DV)*. pp. 769–780. IEEE (2025)
15. Hu, Y., Fang, Y., Ge, Z., Qu, Z., Zhu, Y., Pradhana, A., Jiang, C.: A moving least squares material point method with displacement discontinuity and two-way rigid body coupling. *ACM Transactions on Graphics* **37**(4), 1–14 (2018)
16. Huang, B., Yu, Z., Chen, A., Geiger, A., Gao, S.: 2d gaussian splatting for geometrically accurate radiance fields. In: *ACM SIGGRAPH 2024 Conference Papers*. pp. 1–11 (2024)

17. Jiang, C., Schroeder, C., Selle, A., Teran, J., Stomakhin, A.: The affine particle-in-cell method. *ACM Transactions on Graphics* **34**(4), 1–10 (2015)
18. Jiang, Y., Yu, C., Xie, T., Li, X., Feng, Y., Wang, H., Li, M., Lau, H., Gao, F., Yang, Y., et al.: Vr-gs: A physical dynamics-aware interactive gaussian splatting system in virtual reality. In: *ACM SIGGRAPH 2024 Conference Papers*. pp. 1–1 (2024)
19. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics* **42**(4) (July 2023), <https://repo-sam.inria.fr/fungraph/3d-gaussian-splatting/>
20. Klár, G., Gast, T., Pradhana, A., Fu, C., Schroeder, C., Jiang, C., Teran, J.: Drucker-prager elastoplasticity for sand animation. *ACM Transactions on Graphics* **35**(4), 1–12 (2016)
21. Knapitsch, A., Park, J., Zhou, Q.Y., Koltun, V.: Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics* **36**(4), 1–13 (2017)
22. Kolluri, R.: Provably good moving least squares. *ACM Transactions on Algorithms (TALG)* **4**(2), 1–25 (2008)
23. Lei, Y., Zhao, B., Yang, Z., Li, X., Cheng, T., Peng, H., Zhang, R., Yang, Y., Huang, S., Shen, Y., et al.: Diffwind: Physics-informed differentiable modeling of wind-driven object dynamics. In: *International Conference on Learning Representations*
24. Li, X., Qiao, Y.L., Chen, P.Y., Jatavallabhula, K.M., Lin, M., Jiang, C., Gan, C.: Pac-nerf: Physics augmented continuum neural radiance fields for geometry-agnostic system identification. *arXiv preprint arXiv:2303.05512* (2023)
25. Li, Y., Lin, Z.H., Forsyth, D., Huang, J.B., Wang, S.: Climatenerf: Extreme weather synthesis in neural radiance field. In: *IEEE/CVF International Conference on Computer Vision*. pp. 3227–3238 (2023)
26. Lin, Y., Lin, C., Xu, J., MU, Y.: Omniphysgs: 3d constitutive gaussians for general physics-based dynamics generation. In: *International Conference on Learning Representations*
27. Lin, Y.Y., Pan, X.Y., Fridovich-Keil, S., Wetzstein, G.: Thermalnerf: Thermal radiance fields. In: *2024 IEEE International Conference on Computational Photography (ICCP)*. pp. 1–12. IEEE (2024)
28. Liu, S., Ren, Z., Gupta, S., Wang, S.: Physgen: Rigid-body physics-grounded image-to-video generation. In: *European Conference on Computer Vision*. pp. 360–378. Springer (2024)
29. Lu, R., Chen, H., Zhu, Z., Qin, Y., Lu, M., Zhang, L., Yan, C., Xue, A.: Thermal-gaussian: Thermal 3d gaussian splatting. *arXiv preprint arXiv:2409.07200* (2024)
30. Luo, Z., Cui, Z., Luo, S., Chu, M., Li, M.: Vr-doh: Hands-on 3d modeling in virtual reality. *ACM Transactions on Graphics* **44**(4) (Jul 2025). <https://doi.org/10.1145/3731154>
31. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM* **65**(1), 99–106 (2021)
32. Moenne-Loccoz, N., Mirzaei, A., Perel, O., de Lutio, R., Martinez Esturo, J., State, G., Fidler, S., Sharp, N., Gojcic, Z.: 3d gaussian ray tracing: Fast tracing of particle scenes. *ACM Transactions on Graphics* (2024)
33. Özer, M., Weiherer, M., Hundhausen, M., Egger, B.: Exploring multi-modal neural scene representations with applications on thermal imaging. In: *European Conference on Computer Vision*. pp. 82–98. Springer (2024)
34. Oztireli, C., Guennebaud, G., Gross, M.: Feature preserving point set surfaces based on non-linear kernel regression. In: *Computer Graphics Forum*. vol. 28, pp. 493–501. Wiley (2009)

35. Qiu, R.Z., Yang, G., Zeng, W., Wang, X.: Language-driven physics-based scene synthesis and editing via feature splatting. In: European Conference on Computer Vision (2024)
36. Runway AI, Inc.: Introducing runway gen-4. <https://runwayml.com/research/introducing-runway-gen-4.5> (2025), accessed: 2025-12-1
37. Shen, Q., Tao, N., Dai, Q., Chen, T., Qin, M., Zhang, Y., Chu, M., Chen, W., Chen, B.: Fierygs: In-the-wild fire synthesis with physics-integrated gaussian splatting. In: International Conference on Learning Representations (ICLR) (2026)
38. Shen, S., Shao, T., Zhou, K., Jiang, C., Yang, Y.: Enlivenings: Active locomotion of 3dgs. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 896–905 (2025)
39. Succi, S.: The Lattice Boltzmann Equation: For Fluid Dynamics and Beyond. Oxford University Press (2001)
40. Vasile, F., Qiu, R.Z., Natale, L., Wang, X.: Gaussian-augmented physics simulation and identification with complex colliders. In: Advances in Neural Information Processing Systems (2025)
41. Wan, Team, Wang, A., Ai, B., Wen, B., Mao, C., Xie, C.W., Chen, D., Yu, F., Zhao, H., Yang, J., et al.: Wan: Open and advanced large-scale video generative models. arXiv preprint arXiv:2503.20314 (2025)
42. Wang, M., Zhang, Y., Xu, W., Ma, R., Zou, C., Morris, D.: Decoupledgaussian: Object-scene decoupling for physics-based interaction. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11361–11372 (2025)
43. Wang, Z., Ling, H., Zhang, W., Sun, Y., Sun, Q.: ETGS: Explicit thermodynamics gaussian splatting for dynamic thermal reconstruction. In: International Conference on Learning Representations (2026), <https://openreview.net/forum?id=P2Nw2LMkjH>
44. Xie, T., Zong, Z., Qiu, Y., Li, X., Feng, Y., Yang, Y., Jiang, C.: Physgaussian: Physics-integrated 3d gaussians for generative dynamics. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4389–4398 (2024)
45. Yang, K., Liu, Y., Cui, Z., Liu, Y., Zhang, M., Yan, S., Wang, Q.: Ntr-gaussian: Nighttime dynamic thermal reconstruction with 4d gaussian splatting based on thermodynamics. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 691–700 (2025)
46. Ye, T., Wu, Q., Deng, J., Liu, G., Liu, L., Xia, S., Pang, L., Yu, W., Pei, L.: Thermal-nerf: Neural radiance fields from an infrared camera. In: 2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 1046–1053. IEEE (2024)
47. Yue, Y., Smith, B., Batty, C., Zheng, C., Grinspun, E.: Continuum foam: A material point method for shear-dependent flows. ACM Transactions on Graphics **34**(5), 1–20 (2015)
48. Zhong, L., Yu, H.X., Wu, J., Li, Y.: Reconstruction and simulation of elastic objects with spring-mass 3d gaussians. In: European Conference on Computer Vision. pp. 407–423. Springer (2024)
49. Zong, Z., Li, X., Li, M., Chiaramonte, M.M., Matusik, W., Grinspun, E., Carlberg, K., Jiang, C., Chen, P.Y.: Neural stress fields for reduced-order elastoplasticity and fracture. In: SIGGRAPH Asia 2023 Conference Papers. pp. 1–11 (2023)