


# 000 001 002 003 004 005 FIERYGS: IN-THE-WILD FIRE SYNTHESIS WITH 006 PHYSICS-INTEGRATED GAUSSIAN SPLATTING 007 008 009

010 **Anonymous authors**  
011 Paper under double-blind review  
012  
013  
014  
015  
016  
017  
018  
019  
020  
021  
022  
023  
024  
025  
026  
027  
028  
029

## 030 ABSTRACT 031

032 We consider the problem of synthesizing photorealistic, physically plausible  
033 combustion effects in in-the-wild 3D scenes. Traditional CFD and graphics pipelines  
034 can produce realistic fire effects but rely on handcrafted geometry, expert-tuned  
035 parameters, and labor-intensive workflows, limiting their scalability to the real  
036 world. Recent scene modeling advances like 3D Gaussian Splatting (3DGS) en-  
037 able high-fidelity real-world scene reconstruction, yet lack physical grounding  
038 for combustion. To bridge this gap, we propose FieryGS, a physically-based  
039 framework that integrates physically-accurate and user-controllable combustion  
040 simulation and rendering within the 3DGS pipeline, enabling realistic fire synthesis  
041 for real scenes. Our approach tightly couples three key modules: (1) multimodal  
042 large-language-model-based physical material reasoning, (2) efficient volumetric  
043 combustion simulation, and (3) a unified renderer for fire and 3DGS. By unifying  
044 reconstruction, physical reasoning, simulation, and rendering, FieryGS removes  
045 manual tuning and automatically generates realistic, controllable fire dynamics  
046 consistent with scene geometry and materials. Our framework supports complex  
047 combustion phenomena—including flame propagation, smoke dispersion, and sur-  
048 face carbonization—with precise user control over fire intensity, airflow, ignition  
049 location and other combustion parameters. Evaluated on diverse indoor and outdoor  
050 scenes, FieryGS outperforms all comparative baselines in visual realism, physical  
051 fidelity, and controllability. We will release codes to facilitate future research.  
052  
053



054  
055 Figure 1: FieryGS synthesizes physically-grounded fire effects from multi-view image, enabling controllable  
056 and realistic fire for in-the-wild scenes.  
057

## 058 1 INTRODUCTION 059

060 Synthesizing realistic and controllable combustion effects grounded in in-the-wild 3D scenes is  
061 critical for applications ranging from AR/VR, gaming, and film production to virtual fire drills,  
062 heritage preservation, and robotics perception under adverse conditions, where fire must be visually  
063 convincing, physically plausible, interactively controllable, and well-aligned with the real world.  
064 Existing approaches, however, fall short of meeting these requirements (Table 1).  
065

066 The most authentic option—full-scale fire experiments, such as burning life-sized structures  
067 (Fig. 2)—is prohibitively expensive, risky, and irreproducible, making systematic exploration under  
068 varying conditions infeasible. Alternatively, digital approaches like computational fluid dynamics  
069



Figure 2: Left: Real-world combustion in a live-fire drill (5280 Fire Science); Right: Full-scale combustion test measuring flame spread time (Zhang et al., 2021)

Table 1: Applicability comparison of combustion approaches. FieryGS offers accessible fire simulation for real-world scenes by combining scene-aligned physics, visual fidelity, efficiency, and user control.

Method	Real-world Applicability	Visual Fidelity	Physical Fidelity	Parameter Control	User Friendliness	Scalability
Full-scale Experiments	✓	✓	✓	✗	✗	✗
CFD Methods	✗	sim-to-real gap	✓	✓	expert-only	✗
VFX Tools	✗	sim-to-real gap	✓	✓	expert-only	✗
Commercial Software	✓	✗	pre-stored	✗	✓	✗
Large Video Models	✓	✓	✗	✗	✓	✓
FieryGS (Ours)	✓	✓	✓	✓	✓	✓

**Notes:** **Real-world Applicability** indicates ease of use in real scenes, where CFD/VFX requires manual modeling. **Visual Fidelity** measures perceptual realism, where CFD/VFX suffer sim-to-real gaps and commercial software overlays pre-computed results. **Physical Fidelity** checks consistency with physics, where large video models are data-driven and commercial software uses pre-stored effects. **Parameter Control** reflects the ability to vary conditions, where full-scale experiments are costly to repeat, large video models offer little precise control, and commercial software is limited to pre-stored effects. **User Friendliness** considers usability, where full-scale experiments are dangerous and CFD/VFX requires experts. **Scalability** is automatic adaptation to new scenes at low cost. Full-scale experiments are expensive, CFD/VFX needs manual modeling, and commercial software is limited to pre-stored effects.

(CFD) or visual effect (VFX) software (e.g., Houdini, Blender) incorporate physics-based simulation but depend on asset construction, detailed material annotation, carefully discretized geometry, and brittle simulation–rendering pipelines (Lakkonen, 2024; Mahadika & Utami, 2025). Thus, targeting real-world scenes demands impractical manual specification, and each step remains an incomplete approximation, inevitably producing a pronounced sim-to-real gap that limits practical deployment. With the rise of Large Video Models (LVM), it has become possible to add fire effects directly to footage, but the results lack physical consistency and precise controllability. Due to these limitations, current commercial software (SimsUshare, 2025; Digital Combustion, 2025) instead relies on overlaying pre-stored fire effects onto scenes, without ensuring physical fidelity.

Recent advances in scene modeling present new opportunities. Methods such as Neural Radiance Fields (NeRF)(Mildenhall et al., 2020) and 3DGS(Kerbl et al., 2023) enable high-fidelity 3D reconstruction from multi-view images, providing highly detailed surface information with strong real-world alignment. Although primarily designed for static appearance capture, their visual fidelity and rendering efficiency suggest potential for further material inference and physics-informed modeling. Some prior works leverage such reconstructions to incorporate physical properties (Li et al.; Cai et al., 2024; Li et al., 2023; Feng et al., 2024; Dai et al., 2025; Hsu et al., 2024) to model related phenomena such as fluid dynamics or deformable objects. However, realistic combustion remains out of reach, as it requires accurate scene-level material inference, complex simulation tightly coupled with scene representation, and fine-grained controllability over fire behavior.

To bridge this gap, we introduce FieryGS, a physically based framework that integrates accurate and controllable combustion simulation into the 3DGS pipeline. Our method automatically generates photorealistic, dynamic fire in reconstructed scenes while allowing precise control over fire intensity, airflow, ignition location, and other parameters. The framework tightly couples three components:

- Multimodal-large-language-model(MLLM)-based material reasoning, zero-shot inferring combustion-relevant reliable properties from 3DGS reconstructions;
- Controllable volumetric combustion simulation with wood charring via a principled balance of computational cost and visual realism;
- A novel unified renderer, combining fire, smoke, and 3DGS for seamless photorealistic emission and illumination.

Tightly coupling these modules enables realistic fire effects to emerge directly from real-world data without expert design or handcrafted inputs. FieryGS is, to our knowledge, the first framework that

generates visually and physically realistic combustion in in-the-wild scenes, while being efficient and supporting precise user controls over ignition location, fire intensity, airflow, and other parameters. Experiments across tabletop, indoor, and outdoor scenarios show that FieryGS outperforms state-of-the-art baselines in visual realism, physical fidelity, and user controllability, advancing fire synthesis from labor-intensive, expert-heavy workflows to automatic, real-world aligned process.

## 2 RELATED WORK

**Challenges in Combustion Simulation** Combustion simulation has long been studied in both CFD and computer graphics, with physically based models developed to replicate fire behavior (Husain & Srivastava, 2018; Nguyen et al., 2002; Nielsen et al., 2022; Feldman et al., 2003; Kwatra et al., 2010), material changes such as pyrolysis and charring of wood (Liu et al., 2024a), and volumetric rendering of flames and smoke (Huang et al., 2014; Nguyen et al., 2002; Pegoraro & Parker, 2006). While these methods excel in specific aspects, they rely heavily on manual inputs, such as detailed geometry and material properties, and often require expert knowledge to combine multiple tools, resulting in limited flexibility and sim-to-real gaps in diversity and fidelity. Existing commercial software (SimsUshare, 2025; Digital Combustion, 2025) supports real-world case studies but relies on pre-stored fire effects, lacking both physical consistency and control over fire parameters. These limitations motivate a combustion framework that can automatically align with real-world scenes while maintaining efficiency, controllability, and physical **plausibility**.

**Neural Scene Representations for Physically-Grounded Editing** Recent NeRF and 3DGS representations have enabled high-fidelity 3D reconstruction and inspired extensions to physical property inference. Some estimate parameters like Young’s modulus, fluid viscosity, friction or stiffness from videos (Li et al.; Cai et al., 2024; Cao et al., 2024; Zhong et al., 2024), while others (Zhang et al., 2024; Huang et al., 2024a; Liu et al., 2024b; Lin et al., 2025; Liu et al., 2025) exploit dynamics in video models to infer material properties. LLMs provide a complementary direction to physical property reasoning, as in NeRF2Physics (Zhai et al., 2024), GaussianProperty (Xu et al., 2024), and PUGS (Shuai et al., 2025). However, they remain object-centric and do not address combustion-related attributes. Parallel efforts integrate explicit simulation with neural representations, including deformable bodies via Material Point Method (MPM) (Xie et al., 2024; Zhang et al., 2024; Huang et al., 2024a; Liu et al., 2024b), weather phenomena (Li et al., 2023), fluid–solid interactions (Feng et al., 2024), and rainfall (Dai et al., 2025). AutoVFX (Hsu et al., 2024) supports flame effects using Blender’s built-in physics, but its dynamics are driven by LLM-generated scripts rather than spatiotemporal physical interactions, lacking physical consistency and control. We address this gap by introducing the first framework that integrates combustion simulation with 3DGS, enabling controllable and physically faithful fire synthesis.

## 3 METHODS

Given multi-view images, we reconstruct 3DGS scenes and infer combustion properties through zero-shot MLLM reasoning (Sec. 3.1). The properties guide a physics-based combustion simulation (Sec. 3.2), which is rendered together with the scene using unified volumetric rendering (Sec. 3.3). Fig. 3 illustrates the pipeline.

### 3.1 SCENE MODELING WITH COMBUSTION PROPERTY REASONING

High-fidelity 3D modeling of appearance, geometry, and physical properties in in-the-wild scenes is essential for realistic combustion simulation. We adopt PGSR (Chen et al., 2024), a recent 3DGS-based method that jointly reconstructs photorealistic appearance and accurate geometry, for scene reconstruction. To enable physically plausible fire simulation, we estimate combustion-relevant material properties for each Gaussian in reconstructed 3DGS, including material type, burnability, thermal diffusivity, and smoke color. Recent MLLMs have shown strong capabilities in inferring material from 2D images. However, extending the capabilities to in-the-wild 3DGS scenes remains challenging. Intuitively, nearby Gaussians with visual similarities are likely to share same material properties. Inspired by recent work in 3DGS segmentation (Ye et al., 2024; Cen et al., 2025), we first partition Gaussians into coherent 3D regions, each with a shared material. Then, each region is rendered to 2D and passed to an MLLM for material inference. To ensure reliable MLLM prediction, inference is performed from the viewpoint where the target 3D region has the highest visibility.




Figure 3: **Overall Pipeline of FieryGS.** Given multi-view images as input, we first apply PGSR (Chen et al., 2024) to reconstruct scenes with high-quality normal and depth. Next, we leverage MLLM to infer combustion-related properties, such as material type and burnability. Based on these, we conduct combustion simulations, enabling fire and charring effects with user control. A unified volumetric renderer seamlessly integrates 3DGS and fire, accounting for smoke scattering, fire illumination, and charring, producing realistic fire results.

**Preliminary of 3D Gaussian Splatting** 3DGS (Kerbl et al., 2023) models a scene as a set of anisotropic Gaussians, each parameterized by its center, covariance ( $\Sigma$ ), opacity, and view-dependent color encoded with spherical harmonics. During rendering, a Gaussian is projected into screen space with covariance  $\Sigma' = JV\Sigma V^\top J^\top$ , where  $V$  is the camera extrinsic matrix and  $J$  the Jacobian of the projection. Pixel colors are obtained by alpha blending over depth-sorted Gaussians.

**3D Gaussian Segmentation** Given a reconstructed 3DGS model, we first assign each Gaussian a learnable feature vector  $f_g \in \mathbb{R}^D$ , where  $D$  is the feature dimension. These features are rendered into 2D feature maps via 3DGS alpha blending. We then apply SAM (Kirillov et al., 2023), a foundation model for 2D segmentation, to obtain segmentation maps across multiple views. Following SAGA (Cen et al., 2025), we adopt contrastive learning to train the feature vectors  $f_g$ , encouraging pixels within the same mask to share similar embeddings. After training, Gaussians associated with the same 3D region exhibit similar features. We then apply HDBSCAN algorithm (McInnes et al., 2017) to cluster these feature vectors into instance-level 3D segments, each assumed to correspond to a distinct material region (See Appendix A.1.1 for hyperparameter details).

**MLLM-based Combustion Property Reasoning** For each segmented region in 3D Gaussians, we rasterize it into 2D and perform material inference using an MLLM. In real-world scenes, complex occlusions cause large visibility differences across viewpoints, and limited exposure to the target region can degrade MLLM prediction accuracy. To address this, we select the viewpoint where the target 3D region has the highest visibility, determined by counting the number of unoccluded Gaussians based on rendered depth maps. [We then feed GPT-4o \(Hurst et al., 2024\)](#) a three-panel image composite rendered from the selected viewpoint, including (1) the full scene rendering; (2) the same rendering with the target region highlighted by a bounding box and mask overlay; and (3) an isolated, zoomed-in view of the segmented region, along with a tailored prompt, to infer the material type and combustion-relevant physical properties (See Appendix A.1.2 and Fig. 8 for prompt details). The predicted attributes are projected back to 3D by directly assigning to all Gaussians in the corresponding region. On average, GPT-4o API calls cost about \$0.55 per scene, making our pipeline highly economical (see Appendix B.2). We further validate the robustness and accuracy of the material reasoning results (see Appendix B.3).

The result is a 3DGS augmented with physical and combustion-aware attributes (Fig. 4). An occupancy grid is then constructed, where a voxel is labeled as occupied if it overlaps with one or more 3D Gaussians whose opacity exceeds a given threshold, and further labeled as combustible if any of these Gaussians are burnable. This grid defines the domain for combustion simulation, with unoccupied voxels representing air regions and occupied voxels representing solid regions.

### 3.2 COMBUSTION SIMULATION

Given the occupancy grid obtained in Section 3.1, we run combustion simulation in two parts. [Fire simulation is performed only in the air regions, with solid regions treated as boundary conditions to](#)




Figure 4: **Combustion Property Reasoning.** Given an RGB input (a), our method reliably predicts material types (b) and burnability (d). In a complex region with metal spoons inside a mug surrounded by various materials, the method distinguishes the spoons and correctly infers their non-flammable metallic nature. These results drive the combustion simulation and rendering, where material-specific behaviors are applied—for instance, combustion produces white smoke for the wooden box and black smoke for the plastic Lego (c).

ensure that the velocity field does not penetrate into solid voxels. The implementation details of these boundary conditions are provided in Appendix A.2. Charring simulation updates combustible regions with the degree of charring, supporting the rendering of charred surfaces. Focusing on efficiency, our method employs simplified physical models, primarily an incompressible formulation for fire simulation and a basic charring model, along with several additional minor simplifications, while maintaining visually plausible results. Compared to CFD and VFX methods, which require manual geometry modeling and explicit specification of combustible regions, our pipeline leverages scene modeling and material reasoning to automatically initialize geometry, infer material properties, and identify combustible areas. Meanwhile, users retain flexible control over key parameters, making the simulation workflow largely automated and easy to customize. In the following, we present fire simulation, charring simulation, and user control, while further implementation details are provided in Appendix A.2.

**Fire Simulation** We model flame dynamics using the following equations:

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} = -\frac{1}{\rho} \nabla p + \mathbf{f}, \quad \text{s.t. } \nabla \cdot \mathbf{u} = 0; \quad \frac{\partial Y}{\partial t} + \mathbf{u} \cdot \nabla Y = -k. \quad (1)$$

where  $\mathbf{u}$  is the divergence-free velocity field,  $\rho$  and  $p$  denote density and pressure, and  $Y$  is the reaction coordinate variable ( $Y = 1$  for burning material,  $Y = 0$  for unburnt material). At the beginning of the simulation, all voxels are initialized with  $Y = 0$ . Only the voxels corresponding to user-specified ignition points, which are also predicted as combustible in the occupancy grid, are set to  $Y = 1$ , indicating the onset of combustion.

In this formulation, we choose an incompressible flow model (Nguyen et al., 2002) to balance physical plausibility with computational simplicity, in contrast to compressible formulations (Liu et al., 2024a) that provide higher physical fidelity but at the cost of greater complexity. Among the external forces  $\mathbf{f}$  in Eq. 1, we consider buoyancy force  $\mathbf{f}_{buo} = \alpha(T - T_{air})\mathbf{z}$  and vorticity confinement force  $\mathbf{f}_{vor}$  (Nguyen et al., 2002). To further improve efficiency, the temperature  $T$  is approximated as a quadratic function of reaction coordinate variable  $Y$ , rather than solved through PDE-based thermal models (Nguyen et al., 2002; Nielsen et al., 2022). This simplification makes the simulation pipeline more concise while still capturing the correlation between combustion progress and temperature.

**Charring Simulation** For combustible solids, we simulate temperature evolution by solving a simplified heat transfer equation:

$$\frac{\partial T_m}{\partial t} = \beta \nabla^2 T_m + \gamma_m (T_{amb}^4 - T_m^4) + S_{T_m}, \quad (2)$$

where  $T_m$  denotes the material temperature,  $\beta$  is the thermal diffusivity, and  $\gamma_m$  is the radiative cooling coefficient. To avoid the high cost of explicitly modeling internal heat generation,  $S_{T_m}$  is

approximated by clamping  $T_m$  to  $T_{burn}$  once the ignition threshold  $T_{ign}$  is exceeded. Based on the simulated temperature, the relative char mass is computed as  $\frac{\partial M_c}{\partial t} = \varepsilon_c \xi(T_m)$ , where  $M_c$  denotes the relative char mass, with  $M_c = 1$  representing a fully charred state and  $M_c = 0$  indicating the opposite. The parameter  $\varepsilon_c$  represents the charring rate, while  $\xi(T_m)$  equals 1 if  $T_m \geq T_{ign}$  and 0 otherwise. Unlike prior work that incorporates more detailed mechanisms such as insulation-layer formation or volatile release (Liu et al., 2024a), our formulation deliberately omits these processes. This simplification makes the simulation more efficient while still capturing the visually dominant aspects of charring. Subsequently, each 3D Gaussians directly inherits the  $M_c$  value from its containing grid voxel, providing a simple mapping to guide charring visualization.

**User Control** Our combustion simulation framework provides users with a high degree of control over key aspects of the simulation, including ignition location, fire intensity, and airflow, as demonstrated in Fig. 7. Specifically, users can accurately set the ignition point by assigning reaction coordinate variable  $Y = 1$  to the target ignition voxel. The perceived fire intensity can be adjusted by increasing the buoyancy force coefficient  $\alpha$ , which lifts the flames higher, and decreasing the reaction rate  $k$ , which extends flame visibility—both contributing to a visually stronger fire effect. Airflow can be flexibly controlled by adding an external wind force, enabling users to steer the fire as desired. In addition to these core controls, all other combustion parameters such as thermal diffusivity  $\beta$ , charring rate  $\varepsilon_c$  are also accessible, allowing users to fine-tune the simulation for customized effects.


### 3.3 COMBUSTION RENDERING

We introduce the first rendering framework that jointly integrates simulated fire, smoke, and reconstructed 3DGS into a unified volumetric pipeline. It builds upon the reconstructed 3DGS and the grids obtained from Section 3.2, including the reaction coordinate variable  $Y$  for fire and smoke and the relative char mass  $M_c$  for charring. Using this information, the framework generates the final rendered image that seamlessly combines combustion effects with scene geometry.

Our framework builds upon volumetric rendering (Fong et al., 2017) with targeted simplifications tailored to combustion. Since fire is modeled as a blackbody radiator with negligible scattering, and smoke is treated as a low-albedo medium, we omit scattering terms (Nguyen et al., 2002; Pegoraro & Parker, 2006). The 3DGS is rendered as an opaque background where charring effects are incorporated through  $M_c$ . Under these assumptions, the radiance  $L$  at each pixel is computed as:

$$L = L_{\text{fire}} + L_{\text{smoke}} + \hat{T}(L_{\text{GS}} + L_{\text{phong}}). \quad (3)$$

Here,  $L_{\text{fire}}$  and  $L_{\text{smoke}}$  are accumulated along the ray before reaching the 3DGS,  $\hat{T}$  is the transmittance describing remaining energy,  $L_{\text{GS}}$  is the 3DGS radiance with charring, and  $L_{\text{phong}}$  models fire illumination on the geometry. The contribution of each term is visualized in Fig. 5. Their computation is given in subsequent rendering passes, with details in Appendix A.3.



**Figure 5: Rendering Components Breakdown.** Starting with the original view (a), we first add the charring effect (b). Next, we incorporate the simulated smoke (c), followed by the simulated fire (d). Finally, Phong illumination enhances the ground lighting effect caused by the fire, allowing the originally dark shadow to be brightened (e). An optional generative refinement can further enhance the ground reflection (f).

**Fire Rendering** In fire rendering, the dominant visual effect arises from self-emission, which we model based on Planck’s blackbody radiation law (Nguyen et al., 2002). The absorption coefficient  $\sigma_a$  is set to a fixed positive value when the reaction coordinate  $Y > 0$ , indicating active combustion, and zero otherwise. Spectral volumetric rendering is then performed by integrating the emission term along the ray, and the resulting spectral distribution is converted to RGB color space with chromatic adaptation, following the approach in (Nguyen et al., 2002), to obtain perceptually plausible colors.

**Smoke Rendering** Smoke becomes visible as the flame cools down during combustion. We render the smoke when the reaction coordinate variable satisfies  $Y \leq Y_{\text{smoke}}$ . The smoke color is determined by the type of burning material from material reasoning in Section 3.1. For example, smoke from wood combustion is white, while smoke of burning plastic is black (Fig. 4c). By incorporating this model into the volume rendering pipeline, smoke can be presented along with the fire.

**3DGS Rendering** To implement the charring effect in 3DGS, we apply a scaling factor to the color of 3DGS points where the relative char mass satisfies  $M_c \geq M_c^{\text{dark}}$ . Specifically, the color is dimmed by  $r^{\text{dark}} \frac{M_c - M_c^{\text{dark}}}{1 - M_c^{\text{dark}}}$ , where  $r^{\text{dark}}$  is a user-defined factor (typically less than 1) that controls the degree of color dimming when the char mass reaches its maximum ( $M_c = 1$ ). This approach allows the charred regions to progressively darken as the char mass increases, visually simulating the accumulation of charring on the material surface.

**Phong Illumination** We adopt the traditional Phong illumination model (Phong, 1998) to simulate the lighting effect of fire on 3DGS. Specifically, we treat voxels with temperatures exceeding a given threshold as volumetric light sources. For each 3D Gaussian, we consider only the diffuse and specular components. The accumulated spectral radiance at each wavelength  $\lambda$  is:

$$L_\lambda = \sum_i L_{e,\lambda}^{(i)} \cdot [k_d (\mathbf{n} \cdot \mathbf{l}_i) + k_s (\mathbf{r}_i \cdot \mathbf{v})^s], \quad (4)$$

where  $L_{e,\lambda}^{(i)}$  is the spectral radiance emitted by voxel  $i$ ,  $\mathbf{n}$  is the surface normal obtained from the normal map rendered by the 3DGS in Section 3.1,  $\mathbf{l}_i$  is the light direction,  $\mathbf{v}$  is the view direction, and  $\mathbf{r}_i$  is the reflection direction.  $k_d$  and  $k_s$  are the diffuse and specular reflection coefficients, and  $s$  controls the sharpness of the specular highlight. Finally, the accumulated spectral radiance is converted into RGB color space using the same way in fire rendering, resulting in the perceived illumination effect on the 3DGS. **To further enhance realism, we introduce a Perlin noise-based fluctuation in the emitted light intensity, producing a natural flickering effect. This effect is clearly observable in an additional experiment on the *Firewood* scene with reduced background brightness, as shown in the supplementary video.**

**Optional Generative Refinement** While our method captures key physical aspects of fire, real-world combustion involves additional complexities such as indirect illumination, flickering, and subtle light–material interactions, which remain difficult for physics-based pipelines. To enhance realism, we introduce an optional generative refinement module based on Wan2.1 (Wang et al., 2025), a diffusion video model supporting image and text conditioning. Inspired by SDEdit (Meng et al., 2022) and PhysGen (Liu et al., 2024c), we encode the simulated video into the model’s latent space, perturb it with noise, and then denoise it with the first frame as image condition, guided by classifier-free guidance (Dhariwal & Nichol, 2021; Ho & Salimans, 2022). This process adds high-frequency details and more realistic illumination, as shown in Fig. 5f. However, it may also alter background content and lacks strong 3D consistency, so we treat it as an optional refinement step and provide further discussion in Appendix B.6.

## 4 EXPERIMENTS

In this section, we evaluate FieryGS across diverse scenes and compare it with baselines. We further demonstrate the flexible user control of FieryGS. Results highlight FieryGS’s strengths in high-fidelity rendering, physical plausibility, and controllable fire synthesis. Please refer to our supplementary video for high-quality dynamic visualizations.

**Experimental Details** We evaluate FieryGS on 6 real-world scenes, including 4 custom-captured scenes (*Firewood*, *Kitchen*, *Chair*, *Stool*) recorded with an iPhone, the *Garden* scene from the MipNeRF360 dataset, and the *Playground* scene from the Tanks and Temples dataset. These scenes cover both indoor and outdoor environments and feature diverse object geometries, materials, and spatial arrangements, validating our method in complex, in-the-wild settings.

378 Table 2: Quantitative comparisons.  
379

Method	Aesthetic Quality↑	Imaging Quality↑	DINO Structure↓
AutoVFX	0.488	0.603	1.04
Runway-V2V	0.605	0.701	0.68
Instruct-GS2GS	0.451	0.394	0.66
<b>Ours</b>	<b>0.624</b>	<b>0.702</b>	<b>0.38</b>

380 Table 3: User Studies results.  
381

Baseline	Perceptual Realism		Physical Plausibility	
	Image	Video	Image	Video
vs AutoVFX	88.9	77.8	86.6	85.5
vs Runway-V2V	79.4	66.5	85.3	79.0
vs Instruct-GS2GS	85.5	63.0	83.2	84.5

382 *Note:* Values = % of cases where FieryGS is preferred.  
383

384 We compare FieryGS against 3 representative baselines: an automatic VFX pipeline (AutoVFX (Hsu  
385 et al., 2024)), a video-to-video generation model (Runway-V2V (Runway, 2024a;b)), and a text-  
386 driven 3DGS editing method (Instruct-GS2GS (Vachha & Haque, 2024)). AutoVFX enables dynamic  
387 editing in 3DGS scenes via language instructions using Blender’s physics engine. Runway-V2V  
388 refers to the leading commercial model of Runway for video-to-video synthesis. Instruct-GS2GS  
389 performs text-driven editing on 3DGS models via a 2D diffusion model. All support fire synthesis,  
390 enabling a comprehensive comparison with our method. All prompts are in Appendix B.1.  
391

392 **Qualitative Evaluation** Fig. 6 presents a comparison of FieryGS against baselines on *Kitchen*  
393 scene, demonstrating dynamic fire synthesis over time. Runway-V2V produces visually appealing  
394 fire videos, but significantly alters the original scene’s appearance and structure—for instance, a plate  
395 originally placed on the table is transformed into a circular groove on the tabletop, and Lego bricks are  
396 turned into a pile of wooden blocks. Furthermore, its fire lacks physical plausibility, failing to capture  
397 core combustion dynamics such as flame propagation, and it cannot generate smoke colors that vary  
398 with different burning materials. AutoVFX incorporates dynamic fire through Blender’s physics  
399 engine. However, in complex indoor environments, the resulting flames fail to achieve a convincing  
400 level of realism. Instruct-GS2GS cannot localize fire edits and supports only static modifications of  
401 3DGS models. In contrast, FieryGS generates temporally coherent fire effects that are both visually  
402 authentic and physically grounded, faithfully reproducing the evolution of ignition, flame spread, and  
403 scene illumination. More qualitative comparisons are presented in Appendix B.1.  
404

405 **Quantitative Evaluation** We report **Aesthetic Quality** and **Imaging Quality** scores from  
406 VBench (Huang et al., 2024b) to assess visual fidelity, and **DINO Structure Score** (Parmar et al.,  
407 2024) to evaluate structure preservation. As shown in Table 2, our method achieves the highest scores  
408 in both visual quality metrics and the lowest DINO Structure Score among all baselines, indicating  
409 that it produces visually compelling results while faithfully preserving the input scene structure.  
410

411 **User Studies** We conducted two user studies to evaluate both the perceptual realism and physical  
412 plausibility. In the first study (86 participants), users compared 31 randomly sampled image or video  
413 pairs and selected the one with more realistic fire that better preserved the background scene. The  
414 second study (88 participants) followed the same setup but asked users to judge which result appeared  
415 more physically plausible. Results in Table 3 demonstrate a consistent preference for our method.  
416 Additional setup details are provided in Appendix B.5.  
417

418 **Runtime** The average runtime of FieryGS during the simulation and rendering stage is 2.37 seconds  
419 per frame on an NVIDIA RTX 4090D GPU. A detailed timing breakdown and comparisons with  
420 baselines are provided in Appendix B.4.  
421

422 **User Control Analysis** A key advantage of FieryGS is its fine-grained user controllability over  
423 combustion behavior. Users can adjust the full combustion-related physical parameters—ignition  
424 location, airflow, fire intensity, thermal diffusivity, charring rate, and more. Fig. 7 illustrates how  
425 varying these parameters produces semantically meaningful and physically consistent changes in  
426 fire behavior. For example, altering the ignition location results in different flame propagation paths,  
427 while adjusting airflow direction directs the spread of flames accordingly. These controls enable  
428 precise authoring of dynamic fire effects without manual 3D modeling or complex simulation setup.  
429 Compared to baselines, which either lack explicit control (Runway-V2V), or support only limited,  
430 coarse-grained edits (AutoVFX, Instruct-GS2GS), FieryGS offers a significantly more flexible and  
431 intuitive editing workflow for physically plausible fire synthesis.  
432

## 433 5 LIMITATIONS AND CONCLUSIONS

434 While FieryGS demonstrates strong performance in multi-object scenes, it incorporates several  
435 simplifications for efficiency. Specifically, the framework does not explicitly model mass loss or  
436 thermal degradation, simplifies certain fire dynamics, and focuses more on multi-object scenes rather  
437 than complex indoor environments.  
438




Figure 6: Fire synthesis results over time on *Kitchen* scene. AutoVFX shows limited fire realism in complex indoor environments. Runway-V2V generates visually plausible flames but significantly alters the scene and omits ignition dynamics. Instruct-GS2GS produces static, low-fidelity edits without temporal evolution. In contrast, FieryGS synthesizes physically grounded, time-evolving fire with realistic ignition, spread, and scene illumination.




Figure 7: Controllability of FieryGS. Rows vary ignition location: under (Bottom), behind (Behind), and in front of the table (Front). Columns show simulation settings: baseline (Original), increased intensity via stronger buoyancy ( $\uparrow \alpha$ ) and lower reaction rate ( $\downarrow k$ ) (Intensified), and added rightward wind (Airflow). FieryGS enables intuitive control over ignition, intensity, and airflow.

than modeling large-scale conflagrations. In addition, the uneven distribution of reconstructed 3DGS points can introduce artifacts, and misclassifications in material reasoning may lead to incorrect combustion behavior. Despite these limitations, FieryGS provides an automated pipeline for in-the-wild fire synthesis, with broad potential for simulation, safety training, and immersive content. Code and data will be released upon acceptance. For a more detailed discussion of limitations and potential directions for future work, we refer readers to Appendix C.

## ETHICS STATEMENT

This work adheres to the ICLR Code of Ethics. We conducted two user studies on Amazon Mechanical Turk to evaluate perceptual realism and physical plausibility. The studies followed platform guidelines, and no personally identifiable information was collected. Beyond these studies, no human subjects or animal experiments were involved. All datasets used in this work were either publicly available or captured in controlled environments, ensuring no violation of privacy or copyright.

One potential societal risk of this research is the misuse of fire synthesis for misinformation or malicious visual manipulation. We explicitly acknowledge this risk and strongly encourage responsible and ethical use. At the same time, we believe that high-quality fire synthesis has significant positive applications. It can benefit a wide range of domains, from AR/VR, gaming, and film production to virtual fire drills, heritage preservation, and robotics perception under adverse conditions, by providing controllable, safe, and realistic fire effects without requiring real-world flame generation, thereby reducing potential risks. We are committed to transparency, integrity, and the responsible dissemination of research outcomes.

## REFERENCES

5280 Fire Science. Adams 12 school district washington square fire science live burn. URL <https://5280fire.com/2024-incidents/adams-12-school-district-washington-square-fire-science-live-burn/>. Accessed: 2024-06-21.

515 Blender Online Community. Blender - a 3d modelling and rendering package. URL <https://www.blender.org>. Blender Foundation, Stichting Blender Foundation, Amsterdam.

518 Junhao Cai, Yuji Yang, Weihao Yuan, Yisheng He, Zilong Dong, Liefeng Bo, Hui Cheng, and Qifeng Chen. Gic: Gaussian-informed continuum for physical property identification and simulation. *arXiv preprint arXiv:2406.14927*, 2024.

521 Junyi Cao, Shanyan Guan, Yanhao Ge, Wei Li, Xiaokang Yang, and Chao Ma. NeuMA: Neural material adaptor for visual grounding of intrinsic dynamics. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL <https://openreview.net/forum?id=AvWB40qXZh>.

525 Jiazhong Cen, Jiemin Fang, Chen Yang, Lingxi Xie, Xiaopeng Zhang, Wei Shen, and Qi Tian. Segment any 3d gaussians. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 1971–1979, 2025.

529 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. *IEEE Transactions on Visualization and Computer Graphics*, 2024.

533 Qiyu Dai, Xingyu Ni, qianfan Shen, Wenzheng Chen, Baoquan Chen, and Mengyu Chu. Rainygs: Efficient rain synthesis with physically-based gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2025.

537 Prafulla Dhariwal and Alexander Quinn Nichol. Diffusion models beat gans on image synthesis. In Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pp.

540 8780–8794, 2021. URL <https://proceedings.neurips.cc/paper/2021/hash/49ad23d1ec9fa4bd8d77d02681df5cfa-Abstract.html>.

541

542

543 Digital Combustion. Fire studio 7: Fire simulator software, 2025. URL <https://digitalcombustion.com/>. Accessed: 2025-09-22.

544

545

546 Bryan E Feldman, James F O’brien, and Okan Arikan. Animating suspended particle explosions. In *ACM SIGGRAPH 2003 Papers*, pp. 708–715. 2003.

547

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

551

552 Randima Fernando et al. *GPU gems: programming techniques, tips, and tricks for real-time graphics*,  
553 volume 590. Addison-Wesley Reading, 2004.

554

555 Julian Fong, Magnus Wrenninge, Christopher Kulla, and Ralf Habel. Production volume rendering:  
556 Siggraph 2017 course. In *ACM SIGGRAPH 2017 Courses*, pp. 1–79. 2017.

557

558 Torsten Hädrich, Daniel T Banuti, Wojtek Pałubicki, Sören Pirk, and Dominik L Michels. Fire in  
559 paradise: Mesoscale simulation of wildfires. *ACM Transactions on Graphics (TOG)*, 40(4):1–15,  
2021.

560

561 Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance, 2022. URL <https://arxiv.org/abs/2207.12598>.

562

563 Hao-Yu Hsu, Zhi-Hao Lin, Albert Zhai, Hongchi Xia, and Shenlong Wang. Autovfx: Physically  
564 realistic video editing from natural language instructions. *arXiv preprint arXiv:2411.02394*, 2024.

565

566 Yuanming Hu, Tzu-Mao Li, Luke Anderson, Jonathan Ragan-Kelley, and Frédo Durand. Taichi: a  
567 language for high-performance computation on spatially sparse data structures. *ACM Transactions  
568 on Graphics (TOG)*, 38(6):1–16, 2019.

569

570 Tianyu Huang, Haoze Zhang, Yihan Zeng, Zhilu Zhang, Hui Li, Wangmeng Zuo, and Rynson WH  
571 Lau. Dreampysics: Learning physical properties of dynamic 3d gaussians with video diffusion  
572 priors. *arXiv preprint arXiv:2406.01476*, 2024a.

573

574 Zhanpeng Huang, Guanghong Gong, and Liang Han. Physically-based modeling, simulation and  
575 rendering of fire for computer animation. *Multimedia tools and applications*, 71(3):1283–1309,  
2014.

576

577 Ziqi Huang, Yinan He, Jiashuo Yu, Fan Zhang, Chenyang Si, Yuming Jiang, Yuanhan Zhang, Tianxing  
578 Wu, Qingyang Jin, Nattapol Chanpaisit, Yaohui Wang, Xinyuan Chen, Limin Wang, Dahua Lin,  
579 Yu Qiao, and Ziwei Liu. Vbench: Comprehensive benchmark suite for video generative models.  
580 In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA,  
USA, June 16-22, 2024*, pp. 21807–21818. IEEE, 2024b. doi: 10.1109/CVPR52733.2024.02060.  
URL <https://doi.org/10.1109/CVPR52733.2024.02060>.

581

582 Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Os-  
583 trow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. *arXiv preprint  
584 arXiv:2410.21276*, 2024.

585

586 Shad Husain and Bhuvan Bhasker Srivastava. A review analysis on combustion modelling and  
587 implementation of cfd. *IOSR Journal of Applied Physics (IOSR-JAP)*, 10(6):65–70, 2018.  
588 doi: 10.9790/4861-1006016570. URL <https://www.iosrjournals.org/iosr-jap/papers/Vol10-issue6/Version-1/M1006016570.pdf>. Accessed: 2025-09-22.

589

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

592

593 Chung Min Kim, Mingxuan Wu, Justin Kerr, Ken Goldberg, Matthew Tancik, and Angjoo Kanazawa.  
594 Garfield: Group anything with radiance fields. In *Proceedings of the IEEE/CVF Conference on  
Computer Vision and Pattern Recognition*, pp. 21530–21539, 2024.

594 Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete  
 595 Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *Proceedings*  
 596 *of the IEEE/CVF international conference on computer vision*, pp. 4015–4026, 2023.

597

598 Nipun Kwatra, Jon Gretarsson, and Ronald Fedkiw. Practical animation of compressible flow for  
 599 shockwaves and related phenomena. In *Symposium on Computer Animation*, pp. 207–215, 2010.

600

601 Max Lakkonen. IWMA Position Paper on CFD (Computational Fluid Dynamics) and Water Mist  
 602 Fire Fighting. Presented at the International Water Mist Conference (IWMC), 2024. Available  
 603 at: [iwma.net/fileadmin/user\\_upload/general/IWMA\\_position\\_paper\\_on\\_CFD\\_v0.2f.pdf](http://iwma.net/fileadmin/user_upload/general/IWMA_position_paper_on_CFD_v0.2f.pdf).

604

605 Caroline Larboulette, Pablo Quesada, and Olivier Dumas. Burning paper: Simulation at the fiber's  
 606 level. In *Proceedings of Motion on Games*, pp. 47–52. 2013.

607

608 Xuan Li, Yi-Ling Qiao, Peter Yichen Chen, Krishna Murthy Jatavallabhula, Ming Lin, Chenfanfu  
 609 Jiang, and Chuang Gan. Pac-nerf: Physics augmented continuum neural radiance fields for  
 610 geometry-agnostic system identification. In *The Eleventh International Conference on Learning  
 611 Representations*.

612

613 Yuan Li, Zhi-Hao Lin, David Forsyth, Jia-Bin Huang, and Shenlong Wang. Climatenerf: Ex-  
 614 treme weather synthesis in neural radiance field. In *Proceedings of the IEEE/CVF International  
 615 Conference on Computer Vision*, pp. 3227–3238, 2023.

616

617 Yuchen Lin, Chenguo Lin, Jianjin Xu, and Yadong MU. OmniphysGS: 3d constitutive gaussians for  
 618 general physics-based dynamics generation. In *The Thirteenth International Conference on Learning  
 619 Representations*, 2025. URL <https://openreview.net/forum?id=9HZtP6I51v>.

620

621 Daoming Liu, Jonathan Klein, Florian Rist, ŠA Pirk, Dominik L Michels, et al. Flameforge:  
 622 Combustion of generalized wooden structures. *arXiv preprint arXiv:2412.16735*, 2024a.

623

624 Fangfu Liu, Hanyang Wang, Shunyu Yao, Shengjun Zhang, Jie Zhou, and Yueqi Duan. Physics3d:  
 625 Learning physical properties of 3d gaussians via video diffusion. *arXiv preprint arXiv:2406.04338*,  
 626 2024b.

627

628 Shaowei Liu, Zhongzheng Ren, Saurabh Gupta, and Shenlong Wang. Physgen: Rigid-body physics-  
 629 grounded image-to-video generation. In *European Conference on Computer Vision (ECCV)*,  
 630 2024c.

631

632 Zhuoman Liu, Weicai Ye, Yan Luximon, Pengfei Wan, and Di Zhang. Unleashing the potential of  
 633 multi-modal foundation models and video diffusion for 4d dynamic physical scene simulation.  
 634 *CVPR*, 2025.

635

636 Arya Luthfi Mahadika and Ema Utami. A comparative study of embergen and blender in fire  
 637 explosion simulations. *Jurnal Sisfokom (Sistem Informasi dan Komputer)*, 14(2):216–221, 2025.

638

639 Leland McInnes, John Healy, and Steve Astels. hdbscan: Hierarchical density based clustering. *The  
 640 Journal of Open Source Software*, 2, 03 2017. doi: 10.21105/joss.00205.

641

642 Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon.  
 643 Sdedit: Guided image synthesis and editing with stochastic differential equations. In *The Tenth  
 644 International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29,  
 645 2022*. OpenReview.net, 2022. URL [https://openreview.net/forum?id=aBsCjcPu\\_tE](https://openreview.net/forum?id=aBsCjcPu_tE).

646

647 Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and  
 648 Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020.

649

650 Krzysztof Narkowicz. Aces filmic tone mapping curve. <https://knarkowicz.wordpress.com/2016/01/06/aces-filmic-tone-mapping-curve/>, 2016. Accessed: 2025-05-21.

648 Duc Quang Nguyen, Ronald Fedkiw, and Henrik Wann Jensen. Physically based modeling and  
 649 animation of fire. In *Proceedings of the 29th annual conference on Computer graphics and*  
 650 *interactive techniques*, pp. 721–728, 2002.

651

652 Michael B Nielsen, Morten Bojsen-Hansen, Konstantinos Stamatelos, and Robert Bridson. Physics-  
 653 based combustion simulation. *ACM Transactions on Graphics (TOG)*, 41(5):1–21, 2022.

654

655 Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz,  
 656 and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. In *Proceedings of the*  
 657 *IEEE/CVF international conference on computer vision*, pp. 5865–5874, 2021.

658

659 Gaurav Parmar, Taesung Park, Srinivasa Narasimhan, and Jun-Yan Zhu. One-step image translation  
 660 with text-to-image models. *arXiv preprint arXiv:2403.12036*, 2024.

661

662 Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor  
 663 Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style,  
 664 high-performance deep learning library. *Advances in neural information processing systems*, 32,  
 2019.

665

666 Vincent Pegoraro and Steven G Parker. Physically-based realistic fire rendering. In *NPH*, pp. 51–59,  
 2006.

667

668 Bui Tuong Phong. Illumination for computer generated pictures. In *Seminal graphics: pioneering*  
 669 *efforts that shaped the field*, pp. 95–101. 1998.

670

671 Runway. Introducing gen-3 alpha: A new frontier for video generation.  
 672 <https://runwayml.com/research/introducing-gen-3-alpha>, 2024a.

673

674 Runway. Gen-3 alpha video to video. <https://academy.runwayml.com/gen3-alpha/gen3-alpha-video-to-video>, 2024b.

675

676 Yinghao Shuai, Ran Yu, Yuntao Chen, Zijian Jiang, Xiaowei Song, Nan Wang, Jv Zheng, Jianzhu  
 677 Ma, Meng Yang, Zhicheng Wang, et al. Pugs: Zero-shot physical understanding with gaussian  
 678 splatting. *arXiv preprint arXiv:2502.12231*, 2025.

679

680 SimsUshare. Simsushare: Emergency fire simulation & training software, 2025. URL <https://simsushare.com/>. Accessed: 2025-09-22.

681

682 Jos Stam. Stable fluids. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, pp. 779–786.  
 683 2023.

684

685 Andrew Staniforth and Jean Côté. Semi-lagrangian integration schemes for atmospheric models—a  
 686 review. *Monthly weather review*, 119(9):2206–2223, 1991.

687

688 Cyrus Vachha and Ayaan Haque. Instruct-gs2gs: Editing 3d gaussian splats with instructions, 2024.  
 689 URL <https://instruct-gs2gs.github.io/>.

690

691 Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, Di Chen, Feiwu Yu, Haiming Zhao,  
 692 Jianxiao Yang, Jianyuan Zeng, Jiayu Wang, Jingfeng Zhang, Jingren Zhou, Jinkai Wang, Jixuan  
 693 Chen, Kai Zhu, Kang Zhao, Keyu Yan, Lianghua Huang, Mengyang Feng, Ningyi Zhang, Pandeng  
 694 Li, Pingyu Wu, Ruihang Chu, Ruili Feng, Shiwei Zhang, Siyang Sun, Tao Fang, Tianxing Wang,  
 695 Tianyi Gui, Tingyu Weng, Tong Shen, Wei Lin, Wei Wang, Wei Wang, Wenmeng Zhou, Wente  
 696 Wang, Wenting Shen, Wenyuan Yu, Xianzhong Shi, Xiaoming Huang, Xin Xu, Yan Kou, Yangyu  
 697 Lv, Yifei Li, Yijing Liu, Yiming Wang, Yingya Zhang, Yitong Huang, Yong Li, You Wu, Yu Liu,  
 698 Yulin Pan, Yun Zheng, Yuntao Hong, Yupeng Shi, Yutong Feng, Zeyinzi Jiang, Zhen Han, Zhi-Fan  
 699 Wu, and Ziyu Liu. Wan: Open and advanced large-scale video generative models. *arXiv preprint*  
 700 *arXiv:2503.20314*, 2025.

701 Tianyi Xie, Zeshun Zong, Yuxing Qiu, Xuan Li, Yutao Feng, Yin Yang, and Chenfanfu Jiang.  
 702 Physgaussian: Physics-integrated 3d gaussians for generative dynamics. In *Proceedings of the*  
 703 *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4389–4398, 2024.

702 Qingfeng Xu and Zhuojun Nan. Pyrolysis model to simulate the thermomechanical behaviour of  
 703 cross-laminated timber structures in fire. In *International Conference on Engineering Structures*,  
 704 pp. 1671–1681. Springer, 2024.

705

706 Xinli Xu, Wenhong Ge, Dicong Qiu, ZhiFei Chen, Dongyu Yan, Zhuoyun Liu, Haoyu Zhao, Hanfeng  
 707 Zhao, Shunsi Zhang, Junwei Liang, et al. Gaussianproperty: Integrating physical properties to 3d  
 708 gaussians with Imms. *arXiv preprint arXiv:2412.11258*, 2024.

709

710 Mingqiao Ye, Martin Danelljan, Fisher Yu, and Lei Ke. Gaussian grouping: Segment and edit  
 711 anything in 3d scenes. In *European Conference on Computer Vision*, pp. 162–179. Springer, 2024.

712

713 Albert J Zhai, Yuan Shen, Emily Y Chen, Gloria X Wang, Xinlei Wang, Sheng Wang, Kaiyu Guan,  
 714 and Shenlong Wang. Physical property understanding from language-embedded feature fields.  
 715 In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp.  
 28296–28305, 2024.

716

717 Tianyuan Zhang, Hong-Xing Yu, Rundi Wu, Brandon Y Feng, Changxi Zheng, Noah Snavely, Jiajun  
 718 Wu, and William T Freeman. Physdreamer: Physics-based interaction with 3d objects via video  
 719 generation. In *European Conference on Computer Vision*, pp. 388–406. Springer, 2024.

720

721 Yuchun Zhang, Xiaolong Yang, Yueyang Luo, Yunji Gao, Haiyan Liu, and Tao Li. Experi-  
 722 mental study of compartment fire development and ejected flame thermal behavior for a large-  
 723 scale light timber frame construction. *Case Studies in Thermal Engineering*, 27:101133, 2021.  
 724 ISSN 2214-157X. URL <https://www.sciencedirect.com/science/article/pii/S2214157X21002963>.

725

726 Licheng Zhong, Hong-Xing Yu, Jiajun Wu, and Yunzhu Li. Reconstruction and simulation of elastic  
 727 objects with spring-mass 3d gaussians. *European Conference on Computer Vision (ECCV)*, 2024.

728

729

730

731

732

733

734

735

736

737

738

739

740

741

742

743

744

745

746

747

748

749

750

751

752

753

754

755

756 APPENDIX OVERVIEW  
757

758 This appendix provides supplementary materials to support and extend the main content of FieryGS.  
759 Section A elaborates on implementation specifics for each core component of FieryGS, including  
760 scene modeling with combustion property reasoning, combustion simulation, and rendering. Section B  
761 presents extended experimental results and analyses, covering additional qualitative comparisons,  
762 cost and accuracy analyses of combustion property reasoning, runtime and resource usage, user study  
763 setup, and the discussion of optional generative refinement. Section C discusses the current limitations  
764 and future directions of our method. We further include a supplementary video, showcasing the  
765 dynamic fire synthesis results generated by FieryGS.

766  
767 A METHOD DETAILS  
768769 A.1 SCENE MODELING WITH COMBUSTION PROPERTY REASONING  
770

771 As outlined in Section 3.1 of the main paper, we first reconstruct a high-quality 3DGS model from  
772 multi-view images, accurately capturing both the appearance and geometry of the scene. We then  
773 segment the 3D Gaussians and infer combustion-relevant physical properties for each segmented  
774 region using a multimodal large language model (MLLM). Below, we provide further implementation  
775 details on segmentation and prompt design.

776  
777 A.1.1 HDBSCAN HYPERPARAMETER SETUP  
778

779 To obtain instance-level 3D segments, we employ HDBSCAN (McInnes et al., 2017) to cluster the fea-  
780 ture vectors of 3D Gaussians. We adopt the HDBSCAN parameter settings used in SAGA (Cen et al.,  
781 2025), including a minimum cluster size of 10 and an epsilon of 0.01. Inspired by GARField (Kim  
782 et al., 2024), we further construct a hierarchy of 3D clusters by recursively applying HDBSCAN  
783 at multiple affinity feature scales—specifically 0.9, 0.5, and 0.1. These parameters were selected  
784 through empirical validation and remain fixed across all experiments. We found this configuration to  
785 generalize well across the diverse scenes in our dataset.

786 A.1.2 PROMPTS FOR COMBUSTION PROPERTY REASONING  
787

788 A carefully crafted combination of visual and textual prompts is critical to enable accurate material  
789 reasoning by the MLLM.


790 Inspired by previous work (Xu et al., 2024), we design a specialized prompt for GPT-4o tailored to  
791 combustion property inference (see Fig. 8). The visual prompt includes a [three-panel image composite](#),  
792 ranging from global to local perspectives: (1) a full-scene rendering, (2) the same rendering with the  
793 target region highlighted using a bounding box and mask overlay, and (3) an isolated and zoomed-in  
794 view of the segmented region. This visual hierarchy encourages the MLLM to reason about each part  
795 in relation to its global spatial context.

796 The textual prompt guides the model through a step-by-step reasoning process: it first generates a  
797 brief caption describing the segmented region, then selects the most appropriate material type from a  
798 predefined material library, and finally infers physical combustion attributes such as burnability and  
799 thermal diffusivity. This prompt design enables the MLLM to connect local and global visual cues,  
800 and incrementally construct semantic understanding of the scene, facilitating more accurate physical  
801 property inference.

802  
803 A.2 COMBUSTION SIMULATION  
804

805 We implement our simulation framework from scratch using the Taichi programming language (Hu  
806 et al., 2019), where all variables—including the velocity field  $\mathbf{u}$ , reaction coordinate  $Y$ , material  
807 temperature  $T_m$ , and relative char mass  $M_c$ —are stored at the center of the grid with a resolution of  
808  $256 \times 256 \times 256$ , following the convention in (Fernando et al., 2004). Based on the operator splitting  
809 method (Stam, 2023) for time discretization, the combustion simulation within a single time step  $\Delta t$   
can be summarized as follows:

810  
811  
812  
813  
814  
815  
816  
817  
818  
819  
820  
821  
822  
823  
824  
825  
826  
827  
828  
829  
830  
831  
832  
833  
834  
835  
836  
837  
838  
839  
840  
841  
842  
843  
844  
845  
846  
847  
848  
849  
850  
851  
852  
853  
854  
855  
856  
857  
858  
859  
860  
861  
862  
863



yellow toy dump truck bucket, plastic, burnable, 0.2

Figure 8: Visual and textual prompts used in the MLLM-based combustion property reasoning. The visual input is a three-panel composite rendered from the reconstructed 3DGS: from left to right, a full-scene view, the same view with the target region highlighted using a bounding box and semi-transparent mask, and an isolated, zoomed-in view of the segmented region. The accompanying text prompt guides the MLLM through a step-by-step reasoning process: it first generates a brief caption describing the segmented region, then selects the most likely material from a predefined material library, and finally infers combustion-relevant physical properties such as burnability and thermal diffusivity.



Figure 9: Effect of solid-voxel geometry constraints on fire behavior. Left (a) shows the original simulation without a brick. Right (b) shows the simulation with a virtual brick placed above the campfire, causing the fire to split into two streams and demonstrating that obstacle boundaries are properly enforced.

1. **Advection.** The velocity field  $\mathbf{u}$  and the reaction coordinate variable  $Y$  are advected using the semi-Lagrangian method (Staniforth & Côté, 1991):

$$\mathbf{u}^* := \text{SemiLagrangian}(\mathbf{u}^n, \Delta t, \mathbf{u}^n), \quad (5)$$

$$Y^* := \text{SemiLagrangian}(Y^n, \Delta t, \mathbf{u}^n). \quad (6)$$

2. **External Forces and Reaction.** We then account for external forces  $\mathbf{f}$  acting on the velocity field  $\mathbf{u}$ , and for the reaction consumption on  $Y$ :

$$\mathbf{u}^* := \mathbf{u}^* + \mathbf{f} \Delta t, \quad (7)$$

$$Y^{n+1} := Y^* - k \Delta t. \quad (8)$$

3. **Pressure Projection.** To enforce the incompressibility condition ( $\nabla \cdot \mathbf{u}^{n+1} = 0$ ), we solve the Poisson equation  $\nabla^2 p = \nabla \cdot \mathbf{u}^*$  using Gauss-Seidel iteration to obtain the pressure field  $p$ . The velocity field is then updated as:

$$p := \text{GaussSeidel}(\mathbf{u}^*), \quad (9)$$


$$\mathbf{u}^{n+1} := \mathbf{u}^* - \frac{\Delta t}{\rho} \nabla p. \quad (10)$$

For boundary conditions, we apply open boundary condition on the simulation bounding box. For obstacles, open boundary conditions are used when velocity points outward, while no-through (Neumann) boundary conditions are enforced when velocity points inward. This encourages fluid to flow out of obstacles freely but prevents it from entering them. As shown in Fig. 9, placing a virtual brick above the campfire in the *Firewood* scene splits the fire into two streams, producing behavior markedly different from the original simulation and confirming that solid-voxel geometry constraints are correctly enforced.

4. **Charring Effect.** The material temperature  $T_m$  and relative char mass  $M_c$  are updated explicitly. Since the thermal diffusion term in the update of  $T_m$  corresponds to solving a Poisson equation, we subdivide the time step into smaller sub-steps to ensure stability. In this way, our formulation allows us to capture temperature exchange between objects and enables fire propagation between adjacent combustible solids, as illustrated in Fig. 10

To evaluate the validity of the physical simplifications introduced in our simulator, we conducted a comparison with the expert-level VFX tool Blender Blender Online Community. In this experiment, Blender was provided with the same occupancy grid obtained from our 3DGS reconstruction combined with MLLM-based material reasoning, ensuring that both systems start from identical scene geometry and material properties. Both simulations were run at the same spatial resolution ( $256 \times 256 \times 256$ ) and under identical ignition conditions. The output of Blender’s simulation was rendered using our renderer with consistent lighting and camera settings, allowing for a direct visual comparison.

As shown in Fig. 11, our simplified simulation produces flame behavior comparable to Blender’s, with only minor deviations in highly turbulent regions. These deviations remain within an acceptable

918  
919  
920  
921  
922  
923  
924  
925  
926  
927928  
929 Figure 10: Fire propagation between contacting combustible objects. The three images (left to right) show the  
930 gradual spread of fire across different objects. They demonstrate that our model accurately captures thermal  
931 diffusion, which enables realistic flame transmission between neighboring flammable materials.  
932933 accuracy range, indicating that the simplifications do not substantially compromise realism. In  
934 addition, our method is significantly more efficient, requiring 1.27 seconds per frame on average,  
935 compared with 5.25 seconds per frame for Blender at the same resolution. Furthermore, Blender  
936 does not model heat conduction inside the wood or the resulting flame spread, leading to a stationary  
937 flame. In contrast, our framework naturally captures internal heat transfer and flame propagation.  
938939  
940  
941  
942  
943  
944  
945  
946  
947

(a) Ours

(b) Blender

948  
949 Figure 11: Comparison between our fire simulation and Blender’s fire simulation for the Firewood scene. Both  
950 simulations use the same reconstructed geometry and inferred material properties from our 3DGS and MLLM  
951 pipeline, and are rendered with consistent lighting and camera settings. Our method (a) produces dynamic  
952 flame behavior with internal heat conduction and flame propagation. Blender’s simulation (b) does not model  
953 heat transfer within the wood, resulting in stationary flames. The comparison demonstrates that our simplified  
954 simulation achieves comparable visual realism while being more efficient (1.27 s/frame vs. 5.25 s/frame at  
955 256 × 256 × 256 resolution).956 Overall, this comparison demonstrates that the physical simplifications in our simulator have a limited  
957 effect on visual realism, while enabling higher computational efficiency and supporting physical  
958 processes not available in existing open-source tools.  
959960 A.3 COMBUSTION RENDERING  
961962 To render fire in a physically accurate manner, we first integrate its self-emission spectrum and  
963 convert the result into the RGB color space following the approach in (Nguyen et al., 2002; Pegoraro  
964 & Parker, 2006). Specifically, the emitted spectral radiance at a given wavelength  $\lambda$  is modeled using  
965 Planck’s blackbody radiation law:

966  
967 
$$L_{e,\lambda}(T) = \frac{2hc^2}{\lambda^5} \frac{1}{e^{\frac{hc}{\lambda kT}} - 1}, \quad (11)$$

968 where  $T$  denotes the local temperature, and  $h$ ,  $c$ , and  $k$  are the Planck constant, the speed of light,  
969 and the Boltzmann constant, respectively.  
970971 To reduce computational cost, the spectral radiance is first converted to the CIE XYZ color space using  
972 the standard tristimulus curves defined by the Commission Internationale de l’Éclairage (CIE), prior

972 to volume rendering integration (Nguyen et al., 2002; Pegoraro & Parker, 2006). The integrated XYZ  
 973 values are then transformed into the LMS cone response space using the M\_CAT02 transformation  
 974 matrix. Chromatic adaptation is applied in this space based on the maximum temperature present in  
 975 the fire (Nguyen et al., 2002). Finally, the result is converted back to the RGB color space, followed  
 976 by gamma correction for display.

977 To further enhance the quality of volume rendering, we adopt a coarse-to-fine sampling strategy (Park  
 978 et al., 2021). We first sample 128 points along each ray uniformly, followed by 1024 points via  
 979 importance sampling based on the reaction coordinate variable  $Y$ . Both sets of samples are used for  
 980 the joint rendering of fire and smoke. To address potential exposure issues when compositing their  
 981 RGB outputs, we apply ACES tone mapping curve (Narkowicz, 2016) to remap the colors into the  
 982  $[0, 1]$  range. All these rendering procedures are implemented from scratch in PyTorch (Paszke et al.,  
 983 2019).



994  
 995 Figure 12: Larger specular coefficient  $k_s$  produces clear reflection effects on the 3DGS geometry. Left to right:  
 996 temporal evolution showing pronounced specular highlights under our Phong illumination model.  
 997

998 In addition to the physically based volume emission, we account for the lighting of the surrounding  
 999 3DGS geometry using the Phong illumination model. To better demonstrate that this formulation  
 1000 supports both diffuse and specular components, we increase the specular reflection coefficient  $k_s$  in  
 1001 the *Chair* scene. As shown in Fig. 12, this produces pronounced specular highlights, confirming that  
 1002 our lighting model can generate reflection effects rather than only diffuse shading.

## 1003 B EXTENDED EXPERIMENTAL DETAILS AND RESULTS

### 1004 B.1 MORE QUALITATIVE COMPARISONS ACROSS DIVERSE SCENES

1005 **Datasets and Baselines** We provide additional qualitative comparisons across 6 real-world scenes:  
 1006 4 custom-captured scenes (*Firewood*, *Kitchen*, *Chair*, *Stool*), the *Garden* scene from MipNeRF360,  
 1007 and the *Playground* scene from Tanks and Temples. For comparison, we consider 3 baselines:  
 1008 AutoVFX (Hsu et al., 2024), a language-driven automatic VFX pipeline; Runway-V2V (Runway,  
 1009 2024a;b), a commercial video-to-video generation model; and Instruct-GS2GS (Vachha & Haque,  
 1010 2024), an instruction-based 3DGS editing method.

1011 **Prompt Design** For Runway-V2V, the prompt is “Add fire to the  $\{Target\}$ , showing a full burning  
 1012 process — from ignition to full blaze to smoldering ashes. Flames gradually grow, engulf the object,  
 1013 then slowly fade as smoke rises and embers glow.”; for AutoVFX and Instruct-GS2GS, we use a  
 1014 shared prompt: “The  $\{Target\}$  in the scene is engulfed in roaring flames. The firelight illuminates  
 1015 the surroundings. The smoke billows into the air.” In both cases,  $\{Target\}$  refers to the manually  
 1016 specified object to be ignited.

1017 **More Qualitative Evaluation** Beyond the *kitchen* scene comparison shown in the main paper  
 1018 (Fig. 6), we present qualitative results for the remaining 5 scenes in Figs. 17– 21. Runway-V2V  
 1019 generates visually appealing fire effects but significantly alters the rest of the scene—including geom-  
 1020 etry and appearance of both the background and the burning object—and fails to depict physically  
 1021 plausible combustion dynamics such as ignition, spread, and dissipation. Although AutoVFX is  
 1022 based on Blender’s built-in physics engine, it is not specifically designed for fire synthesis and lacks

1026 Table 4: GPT-4o API call counts per scene for combustion property reasoning  
1027

Scene	Firewood	Stool	Chair	Kitchen	Garden	Playground	Avg.
Times	26	9	46	46	209	169	<b>84</b>

1031  
1032 fine-grained control over combustion behavior, resulting in limited visual realism. Instruct-GS2GS  
1033 performs only coarse, static global edits and is not capable of producing realistic dynamic flames.  
1034

1035 In contrast, FieryGS produces photorealistic and physically grounded fire effects that faithfully  
1036 capture the full progression of combustion, including ignition, flame spread, surface carbonization,  
1037 and eventual burnout.

## 1038 B.2 COST ANALYSIS OF COMBUSTION PROPERTY REASONING

1039 As described in Section 3.1 of the main paper, we employ GPT-4o (Hurst et al., 2024) to perform  
1040 zero-shot material property reasoning. In our pipeline, the number of API calls corresponds to the  
1041 number of segmented regions. As summarized in Table 4, FieryGS requires between 9 and 209 calls  
1042 per scene, depending on scene complexity.

1043 We adopt the ChatGPT-4o-Latest API, which is officially priced at \$5 per million input tokens and  
1044 \$15 per million output tokens. On average, each query uses 1,282 input tokens and generates 18  
1045 output tokens, resulting in a cost of approximately \$0.0066 per call. For a typical scene (mean = 84  
1046 calls), the total cost amounts to approximately **\$0.55**.

1047 Overall, our GPT-4o-based reasoning pipeline is highly cost-efficient and substantially more economical  
1048 than manual annotation.

## 1049 B.3 ACCURACY ANALYSIS OF COMBUSTION PROPERTY REASONING


1050 Accurate and robust combustion property reasoning is essential for physically plausible fire simulation.  
1051 Here, we quantitatively and qualitatively evaluate the accuracy and robustness of our approach.

1052 Since no public scene-level benchmark currently offers reliable ground-truth labels for combustion-  
1053 relevant materials, we perform a manual evaluation on our 6 test scenes. Specifically, we annotated  
1054 the material type for each segmented region and compared these annotations against predictions  
1055 from the MLLM-based material reasoning module. A prediction is deemed correct if it matches the  
1056 ground-truth label. As summarized in Table 5, our method achieves an average accuracy of **89.31%**  
1057 across six diverse scenes, demonstrating strong material reasoning capability.

1058 Most material reasoning errors occur in (i) distant background regions or very small objects that  
1059 are difficult to discern, (ii) heavily occluded areas where the initial segmentation is unreliable, and  
1060 (iii) occasional reconstruction artifacts in 3DGS that distort texture under the GPT-4o inference  
1061 view. These limitations are consistent with those of current 3DGS segmentation and vision–language  
1062 models—limitations shared by current 3DGS segmentation methods and vision–language models.  
1063 Nevertheless, the overall accuracy is sufficient to support downstream combustion simulation with  
1064 minimal human intervention.

1065 **Choice of MLLM** GPT-4o yields the best performance and is therefore our default choice, while  
1066 other frontier MLLMs can also follow our text–visual instructions and work well within our pipeline.  
1067 To evaluate model dependence, we also evaluate the material reasoning accuracy of **Qwen3-VL-Plus**  
1068 under exactly the same text–visual instructions. As shown in Table 5, Qwen3-VL-Plus achieves a  
1069 respectable average accuracy of 82.73%, confirming that our text–visual prompting is robust and  
1070 generalizes well to other frontier MLLMs. However, GPT-4o still provides a clear advantage (89.31%)  
1071 and shows more reliable material reasoning. Given that GPT-4o is already highly cost-effective  
1072 ( $\sim \$0.55$  per scene; see Sec. B.2) and offers the highest accuracy for physics simulation, we use it as  
1073 the optimal choice, while remaining compatible with future MLLMs.

1074 **MLLM Robustness on Challenging Materials** Beyond standard household and outdoor scenes, we  
1075 further evaluate the MLLM-based reasoning on additional real-world examples containing uncommon



1088 (a) Uncommon and composite materials. GPT-4o prediction: {car exhaust diffuser, composite, unburnable}.



1101 (b) Objects with unusual textures. GPT-4o prediction: {barbed wire, metal, unburnable}.

1102  
1103  
1104  
1105  
1106  
1107  
1108  
1109  
1110  
1111  
1112  
1113  
1114  
1115  
1116  
1117  
1118  
1119  
1120  
1121  
1122  
1123  
1124  
1125  
1126  
1127  
1128  
1129  
1130  
1131  
1132  
1133

Figure 13: **MLLM-based material reasoning on challenging materials.** For each example, we show the original image, the segmentation mask overlay, and the cropped part image fed to GPT-4o, together with its prediction. The model correctly infers the car diffuser as composite and unburnable and the barbed wire as metal and unburnable.

Table 5: Accuracy(%) of MLLM-based material reasoning across test scenes.

Scene	Firewood	Stool	Chair	Kitchen	Garden	Playground	Avg.
GPT-4o (Default)	88.46	88.89	82.61	91.30	89.95	89.94	<b>89.31</b>
<b>Qwen3-VL-Plus</b>	<b>88.46</b>	<b>77.78</b>	<b>76.09</b>	<b>84.78</b>	<b>79.90</b>	<b>89.35</b>	<b>82.73</b>

materials and unusual textures (Fig. 13). Representative cases include (1) a car exhaust diffuser, whose surface is a mixture of carbon fiber and hard plastic, and (2) barbed wire, which exhibits dense, repetitive geometric patterns. In these examples, GPT-4o correctly maps the diffuser region to the “composite” class and infers it as unburnable, and identifies the barbed wire as metal and unburnable. These results indicate that the MLLM reasoning generalizes well to challenging materials and textures, providing a robust basis for downstream combustion simulation.

#### B.4 RUNTIME AND COMPUTATIONAL RESOURCES

In our pipeline, 3DGS reconstruction and combustion property reasoning are performed offline, after which pre-frame combustion simulation and rendering are executed. As summarized in Table 6, we report a detailed runtime breakdown of key components, including combustion simulation, Gaussian splatting rendering, and fire and smoke rendering, on a single NVIDIA RTX 4090D GPU. On average, the simulation and rendering stage runs at 2.37 seconds per frame, with peak GPU memory usage below 10.0 GB, demonstrating that our method is both computationally efficient and hardware-friendly.

**Comparison with Baselines** Compared to existing baselines, our method offers a favorable balance of speed and visual quality: AutoVFX Hsu et al. (2024) relies on Blender Blender Online Community for simulation and rendering and requires approximately 4–10 minutes per frame, making it significantly slower than our method. Instruct-GS2GS Vachha & Haque (2024), which directly edits

1134 Table 6: Runtime breakdown (s/frame) of key components in FieryGS across different scenes.  
1135

Scene	Firewood	Stool	Chair	Kitchen	Garden	Playground	Avg.
Simulation	1.27	1.33	1.31	2.56	1.30	1.34	<b>1.52</b>
GS Render	0.010	0.0045	0.0077	0.0043	0.034	0.013	<b>0.012</b>
Fire & Smoke Render	0.75	0.90	0.86	0.69	0.45	1.37	<b>0.84</b>

1141  
1142 3DGS, runs at a fast speed comparable to vanilla 3DGS. However, it produces only coarse, static edits,  
1143 making it unsuitable for synthesizing realistic dynamic flames. Runway-V2V Runway (2024a;b)  
1144 is a closed-source model, preventing direct runtime comparisons; according to its official website,  
1145 generating a 10-second video takes about 30 seconds, but while it produces vivid flame effects, it  
1146 often alters the background content and lacks both physical plausibility and parameter controllability.  
1147

## 1148 B.5 USER STUDY SETUP DETAILS

1149 We conduct two user studies on Amazon Mechanical Turk to assess the key aspects of our method:  
1150 **perceptual realism** and **physical plausibility**. Both studies use an A/B comparison setup, where  
1151 participants were shown 31 randomly sampled image or video pairs. Each pair included one result  
1152 from FieryGS and one from a baseline, with randomized left-right placement to avoid positional bias.  
1153 An example of the evaluation interface is shown in Fig. 14.  
1154

1155 **Study 1: Perceptual Realism** This study involved 86 participants. In each trial, users were asked  
1156 to select the result that exhibited more visually realistic fire effects while maintaining the integrity of  
1157 the original scene. Results are summarized in Table 3 under “Perceptual Realism” (Image/Video).  
1158

1159 **Study 2: Physical Plausibility** We recruited 88 participants using the same evaluation protocol.  
1160 This time, participants were instructed to choose the version that appeared more physically plausible,  
1161 based on how consistent the fire behavior was with real-world expectations, while also preserving  
1162 scene structure. Results are reported in Table 3 under “Physical Plausibility” (Image/Video).  
1163

1164 Across both studies, FieryGS consistently outperforms all baselines in user preference for both images  
1165 and videos. These results indicate that our method produces fire effects that are not only visually  
1166 compelling but also more aligned with human perception of physical realism.  
1167

## 1168 B.6 DISCUSSION OF OPTIONAL GENERATIVE REFINEMENT

1169 To balance efficiency and realism, we simplify our combustion simulation and rendering pipeline by  
1170 omitting certain computationally intensive modules. As a result, our method struggles to capture some  
1171 high-frequency visual effects, such as complex lighting interactions (e.g., multi-bounce reflections),  
1172 fine-scale flame textures, and realistic charring patterns.  
1173

1174 To address these limitations, we introduce a video refinement module based on a pre-trained diffusion-  
1175 based generative model (Wan2.1 (Wang et al., 2025)), as described in Section 3.4. Rather than  
1176 replacing physics-based simulation, this model is used to complement it—enhancing visual fidelity  
1177 while preserving physically grounded motion. In practice, we find that this refinement leads to more  
1178 natural lighting, sharper flame boundaries, and a more compelling overall appearance.  
1179

1180 While the refinement model improves visual quality in many aspects, it can introduce two notable side  
1181 effects that warrant further investigation. First, selectively enhancing fire effects without affecting  
1182 the background is inherently difficult. Generative models tend to alter surrounding areas along with  
1183 the target region, and due to the complex and diffuse nature of flame boundaries, masking proves  
1184 unreliable. Second, as illustrated in Fig. 15, maintaining temporal and 3D consistency remains a  
1185 challenge, especially for long videos—a limitation rooted in the current capabilities of generative  
1186 video models themselves.  
1187

1188 In summary, while generative refinement opens up new possibilities for achieving photorealistic  
1189 fire videos, it is still a complementary step that must be carefully integrated with physically-based  
1190 simulation. We view this as a promising direction for future research, particularly as generative video  
1191 models continue to evolve in quality and controllability.  
1192

1188  
1189  
1190  
1191  
1192  
1193  
1194  
1195



1196  
1197  
1198  
1199  
1200

## Instruction - **MUST READ**

An artificial intelligence agent is trying to add realistic fire to real-world scenes, showing what it might look like if something in the scene caught fire. The goal is to make the flames, smoke, and lighting look natural and believable, while keeping the rest of the scene unchanged.

You will see two images/video clips (Trial 1 and Trial 2). Please zoom in if needed, and consider the following when making your choice:

1. **Check how realistic the fire looks.** Does the fire spread in a natural way? Do the flames and smoke look real? Does the fire light up nearby objects?
2. **Check whether the background still looks like the input scene.** Are the objects, textures, and layout outside the fire area mostly preserved?

Pick the image/video that shows realistic fire, while NOT obviously changing the scene background. Important: Repeatedly selecting the same option will not be considered valid.

- Trial 1
- Trial 2

20 / 31

**Submit**

1211

Figure 14: Visualization of interface for user study.

1212  
1213  
1214  
1215  
1216  
1217  
1218  
1219  
1220



1221  
1222  
1223

Figure 15: Temporal inconsistency in generative refinement across a fire sequence. While the fire visually improves realism, the underlying table texture—occluded during peak fire—changes after the flame dissipates, revealing the diffusion model’s limitations in preserving scene consistency over longer time spans.

1224

## C DISCUSSION OF LIMITATIONS

1227  
1228  
1229

Although FieryGS performs effectively on object-level scenes, it also exhibits several limitations that affect both the physical realism of simulated fire and the generality of the framework.

1230  
1231  
1232  
1233  
1234  
1235  
1236  
1237  
1238



1239  
1240  
1241

Figure 16: Fire consumption results on the *Chair* scene. The three images (left to right) illustrate the material being gradually consumed. This demonstrates that implementing voxel consumption is straightforward within our simulation framework.

1242     **Material Degradation and Mass Loss** To maintain efficiency, we do not simulate thermal degra-  
 1243     dation such as shrinkage, crumpling, or disintegration. Prior works (Liu et al., 2024a; Larboulette et al.,  
 1244     2013) attempt to capture these effects, but doing so requires high computational cost and complex,  
 1245     manually intensive frameworks. Accurately modeling structural changes with unknown internal  
 1246     material properties (e.g., weakening or collapse) also remains extremely challenging (Lakkonen,  
 1247     2024; Xu & Nan, 2024).

1248     We have also experimented with enabling mass loss, i.e., allowing fire to consume material in the voxel  
 1249     grid. As shown in Fig. 16, consuming the lower crossbeam of the wooden chair is straightforward  
 1250     to implement within our simulation framework. However, this introduces several limitations. First,  
 1251     removing surface voxels exposes interior regions reconstructed using 3DGS. While 3DGS provides  
 1252     high-quality surface geometry, it does not model meaningful internal volume. As a result, the newly  
 1253     revealed interior often contains artifacts and produces visually implausible results. Second, mass loss  
 1254     alters scene geometry in ways that significantly affect shadows and illumination, which our current  
 1255     renderer cannot model reliably.

1256     For these reasons, neither thermal degradation nor mass loss is incorporated into our main pipeline.  
 1257     Developing robust treatments for them is an important direction for future work.

1258

1259     **Simplified Flame and Charring Behavior** While our model captures turbulent flames and smoke,  
 1260     it simplifies detailed physical processes for efficiency. For example, we do not model how flames  
 1261     ignite surrounding materials, and more physically grounded approaches such as the thin flame  
 1262     model (Nguyen et al., 2002) could better capture dynamic fire behavior. These are important  
 1263     directions for improving the physical fidelity of the simulation.

1264

1265     **Limitations in Scene Scale** FieryGS is currently tailored to object-level scenes and cannot be  
 1266     directly applied to large-scale scenarios, such as forest or building fires. Extending the framework  
 1267     would require redesigning the fire modeling pipeline and solving new governing equations (Hädrich  
 1268     et al., 2021).

1269

1270     **Non-Uniform 3DGS Distribution** The reconstructed 3DGS points are unevenly distributed, pri-  
 1271     marily concentrated on obstacle surfaces, which can introduce artifacts in volumetric simulation and  
 1272     rendering. Achieving a more uniform distribution throughout obstacle volumes is therefore another  
 1273     important direction for improvement.

1274

1275     **Misclassifications in Material Reasoning** As discussed in Section B.3, while FieryGS achieves  
 1276     high accuracy in material property reasoning, misclassifications occur in (i) tiny or distant background  
 1277     objects with limited visual cues, (ii) heavily occluded regions where segmentation quality degrades,  
 1278     and (iii) occasional 3DGS reconstruction artifacts that distort appearance in the GPT-4o inference  
 1279     view; these issues are inherent limitations of current 3DGS segmentation and vision–language models,  
 1280     and future work will focus on improving robustness and reliability under low visibility and occlusion,  
 1281     as well as resilience to reconstruction imperfections.

1282

1283     Despite these limitations, FieryGS provides an automated approach for fire synthesis in complex  
 1284     scenes, enabling applications in simulation, safety, and immersive content creation. Future work will  
 1285     aim to address these constraints to enhance both physical fidelity and scene generalization.

1286

1287

1288

1289

1290

1291

1292

1293

1294

1295

1296  
 1297  
 1298  
 1299  
 1300  
 1301  
 1302  
 1303  
 1304  
 1305  
 1306  
 1307  
 1308  
 1309  
 1310  
 1311  
 1312  
 1313  
 1314  
 1315  
 1316  
 1317  
 1318  
 1319  
 1320  
 1321  
 1322  
 1323  
 1324  
 1325  
 1326  
 1327  
 1328  
 1329  
 1330  
 1331  
 1332  
 1333  
 1334



1335  
 1336  
 1337  
 1338  
 Figure 17: Fire synthesis results over time on *Firewood* scene. AutoVFX produces unrealistic fire and smoke.  
 Runway-V2V generates visually realistic fire, but it completely alters the scene and lacks a gradual ignition  
 process, showing only fully developed flames. Instruct-GS2GS produces static and unrealistic results. In contrast,  
 FieryGS generates realistic, time-evolving fire with a natural ignition and growth process.

1339  
 1340  
 1341  
 1342  
 1343  
 1344  
 1345  
 1346  
 1347  
 1348  
 1349

1350

1351

1352

1353

1354

1355

1356

1357

1358

1359

1360

1361

1362

1363

1364

1365

1366

1367

1368

1369

1370

1371

1372

1373

1374

1375

1376

1377

1378

1379

1380

1381

1382

1383

1384

1385

1386

1387

1388

1389

1390

1391

1392

1393

1394

1395

1396

1397

1398

1399

1400

1401

1402

1403




Figure 18: Fire synthesis results over time on *Stool* scene. AutoVFX yields visually implausible results, with exaggerated flames and smoke. Runway-V2V produces realistic-looking fire, but heavily distorts the scene geometry and skips the ignition phase, showing only fully developed flames. Instruct-GS2GS outputs blurry, static edits without dynamic behavior. In contrast, FieryGS produces physically plausible, temporally coherent fire that evolves naturally from ignition to flame spread and decay.

1404

1405

1406

1407

1408

1409

1410

1411

1412

1413

1414

1415

1416

1417

1418

1419

1420

1421

1422

1423

1424

1425

1426

1427

1428

1429

1430

1431

1432

1433

1434

1435

1436

1437

1438

1439

1440

1441

1442

1443




Figure 19: Fire synthesis results over time on *Chair* scene. AutoVFX exhibits exaggerated and implausible fire behavior, with little integration into the scene. Runway-V2V produces visually plausible flames but significantly modifies the scene’s appearance and omits the ignition phase. Instruct-GS2GS yields static, glowing effects lacking realistic dynamics. In contrast, FieryGS produces physically grounded fire that evolves naturally—capturing ignition, spread, and burnout—while preserving the underlying scene.

1444

1445

1446

1447

1448

1449

1450

1451

1452

1453

1454

1455

1456

1457

1458

1459

1460

1461

1462

1463

1464

1465

1466

1467

1468

1469

1470

1471

1472

1473

1474

1475

1476

1477

1478

1479

1480

1481

1482

1483

1484

1485

1486

1487

1488

1489

1490

1491

1492

1493

1494

1495

1496

1497

Figure 20: Fire synthesis results over time on *Garden* scene. AutoVFX produces unrealistic, oversized flames and dense smoke that fail to integrate with the environment. Runway-V2V generates visually compelling fire but alters scene details and skips the ignition phase, displaying only intense, fully developed flames. Instruct-GS2GS results in static, overly saturated outputs with no temporal dynamics. In contrast, FieryGS produces physically plausible fire that evolves naturally over time—capturing ignition, spread, and gradual decay—while preserving the original scene context.

1502

1503

1504

1505

1506


1507

1508

1509

1510

1511



1512

1513

1514

1515

1516

1517

1518

1519

1520

1521

1522

1523

1524

1525

1526

1527

1528

1529

1530

1531

1532

1533

1534

1535

1536

1537

1538

1539

1540

1541

1542

1543

1544

1545

1546

1547

1548

1549

1550

Figure 21: Fire synthesis results over time on *Playground* scene. AutoVFX generates exaggerated fire and dense smoke that appear detached from the physical structure. Runway-V2V produces high-quality flames but drastically alters the geometry and texture of the playground, lacking any notion of progressive ignition. Instruct-GS2GS results in temporally static and visually distorted outputs. In contrast, FieryGS synthesizes physically realistic fire that evolves smoothly over time, preserving scene structure while capturing natural ignition, flame spread, carbonization, and decay.

1555

1556

1557

1558

1559

1560

1561

1562

1563

1564

1565

