

Where, What, Why: Toward Explainable 3D-GS Watermarking

Mingshu Cai¹ Jiajun Li² Osamu Yoshie¹ Yuya Ieiri¹ Yixuan Li^{3*}
¹Waseda University ²Southeast University ³Nanyang Technological University
 mignshucaif@fuji.waseda.jp, yoshie@waseda.jp, ieyuharu@ruri.waseda.jp
 jiajun.li@seu.edu.cn, yixuan.li@ntu.edu.sg

Abstract

As 3D Gaussian Splatting becomes the de facto representation for interactive 3D assets, robust yet imperceptible watermarking is critical. We present a representation-native framework that separates where to write from how to preserve quality. A Trio-Experts module operates directly on Gaussian primitives to derive priors for carrier selection, while a Safety and Budget Aware Gate (SBAG) allocates Gaussians to watermark carriers—optimized for bit resilience under perturbation and bitrate budgets—and to visual compensators that are insulated from watermark loss. To maintain fidelity, we introduce a channel-wise group mask that controls gradient propagation for carriers and compensators, thereby limiting Gaussian parameter updates, repairing local artifacts, and preserving high-frequency details without increasing runtime. Our design yields view-consistent watermark persistence and strong robustness against common image distortions such as compression and noise, while achieving a favorable robustness–quality trade-off compared with prior methods. In addition, the decoupled finetuning provides per-Gaussian attributions that reveal where the message is carried and why those carriers are selected, enabling auditable explainability. Compared with state-of-the-art methods, our approach achieves a PSNR improvement of +0.83 dB and a bit-accuracy gain of +1.24%.

1. Introduction

3D Gaussian Splatting (3D-GS) [10], with its explicit parameterization, real-time performance, and high fidelity, is emerging as the mainstream paradigm for 3D content creation and deployment—succeeding NeRF [22]—and is being widely adopted across film, gaming, autonomous driving, digital humans, and world models [5, 12, 14, 28, 39, 48].

However, with the large-scale generation and widespread

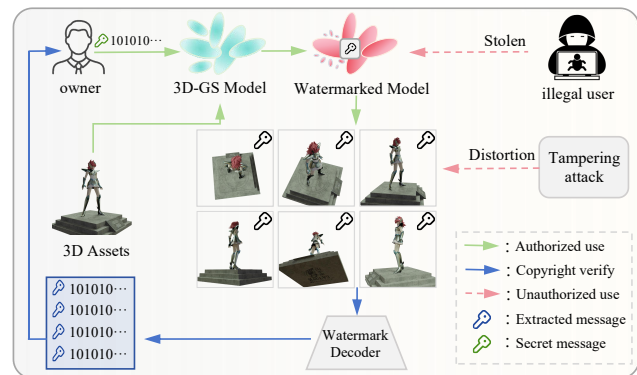


Figure 1. **Usage of our model.** The owner embeds a secret message into a 3D-GS model using our method; even if the model is stolen and undergoes distortion attacks, the rendered views can still be decoded to verify copyright.

distribution of 3D assets, their susceptibility to copying and tampering has made copyright protection increasingly critical—once a model is illicitly obtained or redistributed, the original creator’s rights become extremely difficult to enforce. This challenge is particularly acute in 3D Gaussian Splatting (3D-GS): its core strength—explicit, directly editable Gaussian parameters—while enhancing representational expressiveness and rendering efficiency, also introduces significant security risks. Attackers can easily copy the model, tamper with its content, strip away authorship information, and illegally redistribute it, thereby severely undermining copyright attribution and provenance tracking.

Effective watermarking for 3D-GS is therefore urgent. Prior work on radiance-field watermarking, such as WaterRF and 3DGSW [3, 7, 8], leaves two core gaps for explicit, discretized Gaussians:

(i) Carrier selection: From a large, heterogeneous set of Gaussian primitives, select watermark carriers by jointly considering multi-view visibility, frequency-domain cues, and the stability of geometry and appearance, with an emphasis on visual stability and information security.

(ii) Robust & imperceptible embedding: Embed a robust watermark without degrading visual or rendering quality,

*Corresponding author.

and ensure it remains extractable after common distortions such as cropping, compression, and format conversion.

Motivated by these considerations, we aim to build a framework unifying where to write, what to write, and why it matters: selecting stable 3D Gaussian carriers (where), embedding a distortion-resilient yet invisible signal (what), and achieving interpretable, auditable watermarking (why). A brief usage example is shown in Fig. 1

To address carrier selection, we introduce Trio-Experts, which analyzes intrinsic 3D parameters rather than rendered images. It extracts representation aligned priors: the Geometry expert scores structural stability, the Appearance expert gauges frequency characteristics for imperceptible edits, and the Redundancy expert estimates spatial replaceability for robustness. An uncertainty-aware fusion yields robustness-optimized Gaussian priors. These priors feed the Safety and Budget Aware Gate (SBAG), which finalizes **where to write** by routing Gaussians to the watermark set (WM) only when multi-view stability and safety budgets are met, and by expanding/densifying WM within the permitted visual-quality envelope. Non-carriers are assigned as Visual Compensators (VIS) for watermark embedding.

To minimize visual degradation from both WM and VIS Gaussians and improve embedding efficiency, we introduce a channel-wise Group Mask that specifies **what to write** by selecting watermark-eligible parameter channels while constraining gradient flow and per-channel update magnitudes.

Crucially, we completely decouple WM and VIS during training. VIS points are excluded from the watermark loss to avoid conflicts between carrier optimization and rendering fidelity, which explains **why** the system remains stable under adversarial procedures such as EOT; if VIS were coupled they would counter WM updates, harming visual quality and destabilizing extraction. Decoupling preserves image quality while maintaining watermark accuracy. In essence, this dual-role architecture enables robust & imperceptible watermark embedding.

Extensive experiments show that our method embeds watermarks consistently across all rendered views of a 3D-GS model while remaining robust to attacks on both images and the underlying representation. Compared with state-of-the-art approaches [3, 7, 8], it achieves superior results across all major metrics. Our core contributions are summarized as follows:

- We propose a highly interpretable, attack oriented **Decoupled Finetuning** framework that fully separates carriers and compensators updating, with a **channel wise Group Mask** to route WM and VIS gradients, suppress harmful changes, preserve high frequency details, and achieve secure, robust embedding.
- We introduce **Trio-Experts**, which operate directly on large 3D-GS point sets to extract high-quality geometry/appearance/redundancy priors for densification,

channel-mask construction and watermark-carrier selection.

- We present the **Safety and Budget-Aware Gate (SBAG)** that, under an adaptive budget, uses representation priors and lightweight rendering to select and densify watermark carriers, while cleanly separating them from visual compensators to enable decoupled optimization.
- Our method achieves **state-of-the-art performance** and remains robust under diverse image distortion attacks.

2. Related Work

2.1. 3D Gaussian Splatting

3D Gaussian Splatting (3D-GS) [10] pairs an explicit point representation with differentiable splatting for real-time neural radiance-field rendering. Recent progress spans static reconstruction—shading/reflectance, large-scale scenes, structured/view-adaptive models, multi-scale anti-aliasing [9, 16, 34, 44]—and dynamic/editable settings, including sparsely controllable editing, geometry-prior deformation, real-time 4D, and pose-conditioned human avatars [6, 15, 32, 41]. On the systems side, volumetrically consistent rasterization, 4K-scale efficiency, and training-time point dropping enhance physical fidelity, throughput, and regularization [27, 33, 37]; under weak constraints, depth-GS coupling and unposed joint estimation enable feed-forward reconstruction and cross-scene generalization [4, 42, 46], while 3D-GS also probes 2D foundation features and unifies 2D/3D tone mapping for HDR [2, 17]. These trends sharpen the need for copyright protection and trustworthy use of 3D-GS assets.

2.2. Discrete Wavelet Transform

Discrete Wavelet Transform (DWT) captures spatial-frequency locality via multiscale decomposition and has long supported image denoising/detail recovery [24]. In NeRFs and related 3D pipelines, wavelet-space priors improve sparsity, generalization, and training stability—covering coefficient sparsification and generalizable synthesis [31, 43], direction-aware dynamics [18], and sparse-view wavelet losses plus 3D-GS frequency regularization [8, 25, 26]. We therefore enforce multiscale and local high-frequency consistency in the wavelet domain to boost quality and robustness while preserving global structure.

2.3. Digital Watermarking

Digital watermarking protects digital assets by identifying copyrights. The main difference lies in the priority of data embedding: watermarking prioritizes robustness—ensuring detection after distortions—whereas steganography prioritizes invisibility. To achieve robustness, traditional methods embed data in DWT subbands [1, 30, 35, 38]. HiDDen

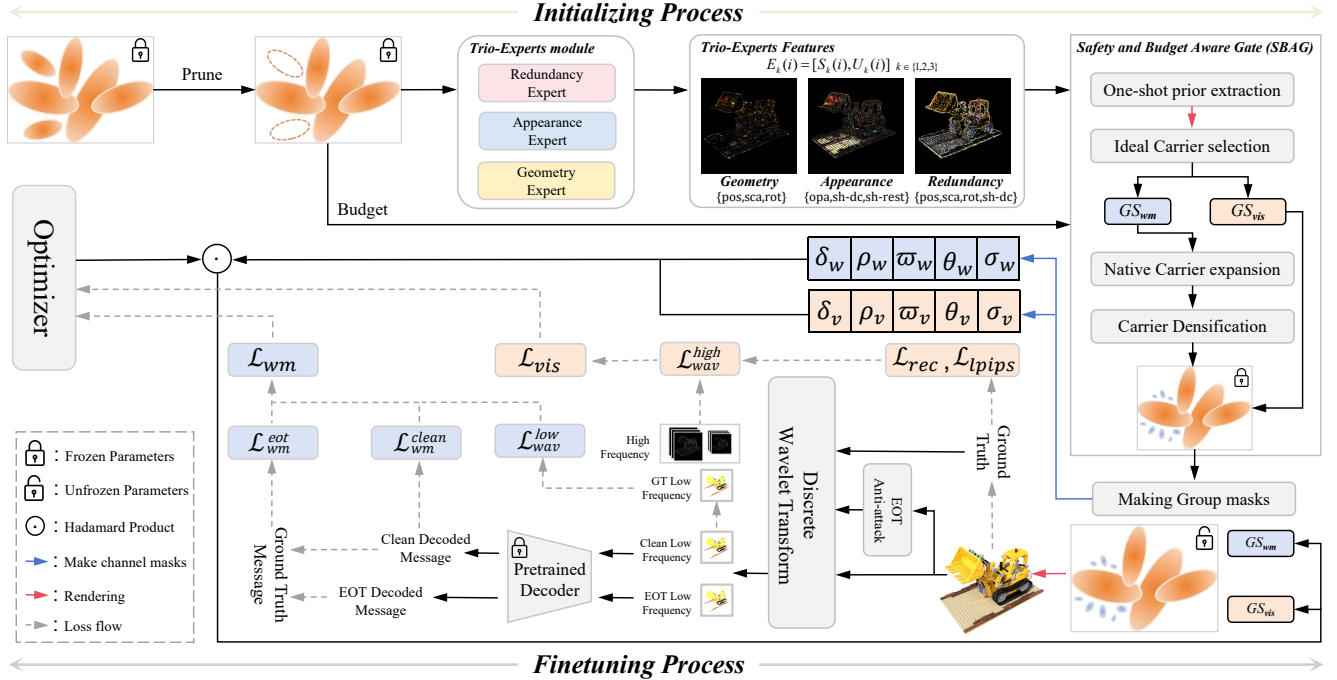


Figure 2. **Pipeline of Our Method.** During initialization, we prune redundant Gaussians based on their rendering contribution. The **Trio-Experts** module extracts geometry/appearance/redundancy priors and aggregates them into an evidence package $E_k(i)$. The **SBAG** decouples ranking and budgeting and uses this evidence and a one-shot render to select, expand, and densify watermark carriers. In finetuning, a channel-wise **group mask** enforces disjoint gradient routing for watermark carriers (GS_{wm}) and visual compensators (GS_{vis}). EOT attacks render both clean and attacked views: \mathcal{L}_{vis} preserves appearance, while low-frequency subbands form the watermark loss \mathcal{L}_{wm} , with \mathcal{L}_{wav}^{low} penalizing over-editing. The separate optimization of \mathcal{L}_{vis} and \mathcal{L}_{wm} improves both fidelity and robustness.

[47] introduces an end-to-end deep watermarking framework with a noise layer. For radiance fields, CopyRNeRF [19] embeds messages into images rendered from implicit NeRFs. StegaNeRF [13] focuses on hiding information within rendered views with minimal visual impact, emphasizing imperceptibility. NeRFProtector [36] approaches the problem from a protection and authentication perspective, adding verifiable information to NeRFs without altering scene structure. WaterF [7], on the other hand, leverages DWT-based frequency embedding to enhance both fidelity and robustness. In the 3D-GS representation, 3DGSW [8] improves robustness through joint regularization based on rendering contribution and wavelet-domain constraints, while GuardSplat [3] focuses on real-world attack scenarios, achieving robust watermark embedding and verification with a strong pretrained CLIP [29] decoder. Unlike prior frequency-cue methods, we watermark directly in 3D-GS parameter space with gated, decoupled, attack-aware training for secure, robust, visually lossless embedding.

3. Method

3.1. Preliminary

3D Gaussian Splatting. We work under the standard 3D Gaussian Splatting (3D-GS) formulation [10], where a

scene is represented by a set of 3D Gaussian primitives parameterized by mean μ and covariance Σ :

$$G(\mathbf{x}; \mu, \Sigma) = \exp\left(-\frac{1}{2}(\mathbf{x} - \mu)^\top \Sigma^{-1}(\mathbf{x} - \mu)\right). \quad (1)$$

For rendering, each 3D Gaussian is projected to the image plane of a viewpoint π . Given depth ordering, pixels are composited by front-to-back alpha blending:

$$I_\pi[x, y] = \sum_{i \in \mathcal{N}_G} c_i \alpha_i T_i, \quad T_i = \prod_{j=1}^{i-1} (1 - \alpha_j), \quad (2)$$

where \mathcal{N}_G is the set of visible 2D Gaussians at π sorted by depth, c_i denotes color, α_i is the per-pixel contribution, and T_i is accumulated transmittance.

Prune by Contribution of Rendering Quality. Pre-trained 3D-GS models often contain redundant 3D Gaussians. We adopt the contribution-based pruning method from 3D-GSW [8], which quantifies contribution using an auxiliary loss with temporary color parameters C' . The gradient $V_\pi = \partial L_\pi^{aux} / \partial C'$ serves as the contribution score to identify and prune low-impact Gaussians.

3.2. 3D Feature-Aware Trio-Experts

Prior methods often select carrier locations in the image or frequency domain based on gradient or high-frequency

heuristics, but this evidence tends to drift with viewpoint changes. To address this issue, we propose a representation-native Trio-Experts system, whose decision evidence is fully anchored in the 3D-GS parameter space.

For the i -th Gaussian, its parameters are denoted as $\mathbf{x}_i \in \mathbb{R}^3$ (position), $\mathbf{s}_i \in \mathbb{R}_+^3$ (scale), $\mathbf{q}_i \in \mathbb{H}$ (rotation), $\alpha_i \in (0, 1)$ (opacity), and $\mathbf{h}_i = [\mathbf{h}_i^{(0)}, \mathbf{h}_i^{(\geq 1)}]$ (SH coefficients). We only extract features from these native parameters, without relying on pixel-domain gradients, thereby ensuring view-consistent evaluation and reducing computational overhead. We group the native 3D-GS parameters by semantics: $\mathcal{C}_{\text{geo}} = \{\mathbf{x}, \mathbf{s}, \mathbf{q}\}$, $\mathcal{C}_{\text{app}} = \{\alpha, \mathbf{h}^{(0)}, \mathbf{h}^{(\geq 1)}\}$, and $\mathcal{C}_{\text{red}} = \{\mathbf{x}, \mathbf{s}, \mathbf{q}, \mathbf{h}^{(0)}\}$. After quantile-based min-max normalization, we construct a k -NN neighborhood $\mathcal{N}_k(i)$ in the 3D position space of \mathbf{x} to model local density, and compute:

Geometric Features z_1 . Based on \mathcal{C}_{geo} , we capture structural decomposition and boundary cues in 3D space via scale isotropy, rotational consistency, and a compact footprint, to measure geometric stability:

$$z_1(i) = \text{Norm}(\text{Iso}_i, \text{RotCons}_i, 1 - \overline{\text{fp}}_i), \quad (3)$$

where $\text{Iso}_i = \min(\mathbf{s}_i) / \max(\mathbf{s}_i)$ measures scale isotropy, $\text{RotCons}_i = (1/k) \sum_{j \in \mathcal{N}_k(i)} |\langle \mathbf{q}_i, \mathbf{q}_j \rangle|$ is neighborhood quaternion consistency, and $\overline{\text{fp}}_i = \exp((1/3) \sum_d \log s_{i,d})$ is the geometric-mean footprint.

Appearance Features z_2 . Based on \mathcal{C}_{app} , we measure cross-view appearance consistency from color and opacity cues via DC band-pass, opacity gating, and high-frequency suppression:

$$z_2(i) = \text{Norm}(1 - \rho_i^{\text{hf}}, g(\alpha_i), c_i), \quad (4)$$

where ρ_i^{hf} is the AC high-frequency energy ratio, $g(\alpha)$ is a double-sided opacity gate, and c_i is a Gaussian band-pass on DC strength.

Redundancy Features z_3 . Based on \mathcal{C}_{red} , we characterize distributional density among 3D Gaussians and estimate substitutability via overlap-weighted neighborhood similarity in DC color and shape:

$$z_3(i) = \text{Norm}\left(\frac{1}{k} \sum_{j \in \mathcal{N}_k(i)} w_{ij} r_{ij}\right), \quad (5)$$

where r_{ij} combines color and shape similarity, and $w_{ij} = \exp(-d_{ij}^2 / (\sigma_o^2(\bar{s}_i^2 + \bar{s}_j^2)))$ approximates projected overlap using spatial distance and scale.

Each expert k maps its features $z_k(i)$ into an evidence package $E_k(i) = [U_k(i), S_k(i)]$, separating quality from

certainty:

$$\begin{aligned} U_k(i) &= \text{Norm}(\text{Disp}_{\mathcal{N}(i)}(z_k) + \text{Penalty}_k(i)), \\ S_k(i) &= \text{Norm}(z_k(i)), \quad k \in \{1, 2, 3\}. \end{aligned} \quad (6)$$

where $U_k \in [0, 1]$ measures uncertainty from neighborhood dispersion (Disp) and expert-specific penalties, and $S_k \in [0, 1]$ is the quality score; this decoupling enables confidence-aware expert gating.

3.3. Safety and Budget-Aware Gate (SBAG)

To select robust watermark carriers from a large 3D-GS set, we decouple the pipeline into *ranking* and *budgeting*. Ranking relies solely on Trio-Experts evidence packages $E_k(i) = [U_k(i), S_k(i)]$, mapped to proxy scores aligned with $\{\mathcal{C}_{\text{geo}}, \mathcal{C}_{\text{app}}, \mathcal{C}_{\text{red}}\}$:

$$R_k(i) = \text{clip}(S_k(i) - \beta U_k(i), 0, 1), \quad k \in \{1, 2, 3\}, \quad (7)$$

where β is a fixed constant shared across experts. We interpret R_1 as geometric stability, R_2 as appearance safety, and R_3 as redundancy certainty; in all cases, higher scores are preferred. To avoid assuming expert dominance, we define a symmetric point-wise utility:

$$u_i = \left(R_1(i) \cdot R_2(i) \cdot R_3(i)\right)^{\frac{1}{3}}. \quad (8)$$

We render all training views once using DC+opacity to obtain view-corrected visibility and distribution priors. The rasterizer provides the non-negative compositing weight $w_{i,p}^{(t)} \geq 0$ of Gaussian i at pixel p in view t . We estimate a scene-level crowding factor $\eta \in (0, 1]$ as

$$\eta = \frac{1}{V} \sum_{t=1}^V \frac{\sum_p \min(1, \sum_i w_{i,p}^{(t)})}{\sum_p \sum_i w_{i,p}^{(t)} + \epsilon}, \quad w_{i,p}^{(t)} \geq 0, \quad (9)$$

and obtain per-Gaussian visibility $v_i \in [0, 1]$ by accumulating its screen-space contribution and normalizing. We further compute the scene-average visibility $\bar{v} = \frac{1}{N} \sum_{i=1}^N v_i$.

Given message length M (bits), we model the effective bits contributed by one carrier as a scene-adaptive coefficient

$$\kappa_{\text{eff}} = \kappa_0 \cdot \bar{v} \cdot \eta, \quad B = \left\lceil \frac{M}{\kappa_{\text{eff}}} \right\rceil, \quad (10)$$

where κ_0 is a constant determined by the embedding design.

We define a generic feasible set using quantile-based constraints:

$$\mathcal{F} = \left\{ i \mid \begin{aligned} R_1(i) &\geq Q_q(R_1), R_2(i) \geq Q_q(R_2), \\ R_3(i) &\geq Q_q(R_3), v_i \geq Q_q(v) \end{aligned} \right\}, \quad (11)$$

and perform deterministic water-level selection within \mathcal{F} by u_i :

$$\mathcal{WM}_0 = \text{top-}B \{ u_i \mid i \in \mathcal{F} \}. \quad (12)$$

To bridge viewpoint-induced coverage gaps, we retain a **Prototype-based Proximity Extension**. We build a compact evidence vector where $c_i = \|\mathbf{h}_i^{(0)}\|_2$ is DC strength and $h_i = \rho_i^{\text{hf}}$ is the AC high-frequency ratio:

$$\mathbf{e}_i = \text{Norm}(R_1(i), R_2(i), R_3(i), v_i, h_i, c_i), \quad (13)$$

compute the prototype $\boldsymbol{\mu} = \frac{1}{|\mathcal{WM}_0|} \sum_{i \in \mathcal{WM}_0} \mathbf{e}_i$, recruit proximal neighbors $\mathcal{WM}_{\text{prox}}$ by cosine similarity, and form

$$\mathcal{WM}_{\text{parent}} = \mathcal{WM}_0 \cup \mathcal{WM}_{\text{prox}}. \quad (14)$$

Finally, each parent is split into N_s visually equivalent children, routing one child to the watermark branch:

$$\mathcal{WM}_* = \bigcup_{i \in \mathcal{WM}_{\text{parent}}} \mathcal{C}_{\text{wm}}(i), \quad (15)$$

while the remaining children act as visual compensators during finetuning to neutralize embedding artifacts. The resulting \mathcal{WM}_* and its complement \mathcal{VIS} are used for subsequent channel-wise mask routing.

3.4. Channel-wise Group Mask

To avoid visible degradation, we assign channel-wise masks to watermark carriers and visual compensators, and optimize a *visual loss* and a *watermark loss* under separate gradient routes. For Gaussian i , we precompute per-group masks $m_g^{\text{wm}}(i), m_g^{\text{vis}}(i) \in [0, 1]$ from Trio evidence and one-shot priors. We derive the two masks from point-wise channel weights $w_g^{(\text{vis})}(j)$ and $w_g^{(\text{wm})}(j)$ as

$$\begin{aligned} m_g^{\text{vis}} &= \max(\text{clip}(\text{mean}(w_g^{\text{vis}}[\mathcal{VIS}]), 0, \text{cap}_g), \text{floor}_g), \\ m_g^{\text{wm}} &= \text{clip}(\text{med}(w_g^{\text{wm}}[\mathcal{WM}_*]), 0, \text{cap}_g). \end{aligned} \quad (16)$$

where $Q_{0.5}$ is the median, $\mathcal{V} = \{1, \dots, N\} \setminus \mathcal{WM}_*$, and $\text{cap}_g, \text{floor}_g$ are per-channel bounds.

and gate gradients as

$$\nabla_{\theta_i^g} \mathcal{L} = \begin{cases} m_g^{\text{wm}}(i) \nabla_{\theta_i^g} \mathcal{L}_{\text{wm}}, & i \in \mathcal{WM}_*, \\ m_g^{\text{vis}}(i) \nabla_{\theta_i^g} \mathcal{L}_{\text{vis}}, & i \in \mathcal{VIS}, \end{cases} \quad (17)$$

where $g \in \{\delta_{\text{dc}}, \boldsymbol{\rho}_{\text{rest}}, \boldsymbol{\omega}_{\text{opa}}, \boldsymbol{\theta}_{\text{rot}}, \boldsymbol{\sigma}_{\text{sca}}\}$ denotes five parameter channels.

To ensure complete gradient separation, we perform two passes of \mathcal{L}_{vis} and \mathcal{L}_{wm} , so that \mathcal{WM}_* and \mathcal{VIS} receive their gradients in an orthogonal manner.

3.5. Decoupled Watermark Finetuning

To resolve the gradient conflict between rendering fidelity and watermark robustness, we implement a **decoupled finetuning** strategy. It isolates gradient propagation by ensuring that watermark carriers \mathcal{WM}_* are exempt from visual reconstruction constraints, while visual compensators \mathcal{VIS}

remain unaffected by watermark-related objectives. Specifically, the visual objective is

$$\mathcal{L}_{\text{vis}} = \lambda_{\text{rec}} \mathcal{L}_{\text{rec}} + \lambda_{\text{lpips}} \mathcal{L}_{\text{lpips}} + \lambda_{\text{wav}}^{\text{high}} \mathcal{L}_{\text{wav}}^{\text{high}}, \quad (18)$$

where \mathcal{L}_{rec} is the L1 reconstruction loss, $\mathcal{L}_{\text{lpips}}$ is LPIPS [45], and $\mathcal{L}_{\text{wav}}^{\text{high}}$ penalizes multi-level DWT high-frequency subbands ($S \in \{LH, HL, HH\}$) with an L1 distance.

Simultaneously, we optimize the watermark branch using Expectation Over Transformation (EOT). Let $\mathbf{b}_i \in \{0, 1\}^B$ denote the target bits of carrier i , \mathcal{V}_i its visible views, and $D_\psi(\cdot)$ a decoder producing logits. For clean and transformed renderings, we average logits over visible views:

$$\bar{\mathbf{z}}_i^{\text{cl}} = \frac{1}{|\mathcal{V}_i|} \sum_{v \in \mathcal{V}_i} D_\psi(\hat{I}_v), \quad \bar{\mathbf{z}}_i(\mathbf{t}) = \frac{1}{|\mathcal{V}_i|} \sum_{v \in \mathcal{V}_i} D_\psi(\mathbf{t}(\hat{I}_v)), \quad (19)$$

where $\mathbf{t} \sim p(\mathbf{t})$ is sampled from a standard degradation family (e.g., blur, rotation, scaling, crop, noise, JPEG). We decode bits by $\hat{\mathbf{y}}_i^{\text{cl}} = \sigma(\bar{\mathbf{z}}_i^{\text{cl}})$ and $\hat{\mathbf{y}}_i^{\text{eot}}(\mathbf{t}) = \sigma(\bar{\mathbf{z}}_i(\mathbf{t}))$. We embed watermarks *only* in the DWT low-frequency (LL) subband and regularize low-frequency distortion with $\mathcal{L}_{\text{wav}}^{\text{low}}$ (multi-level DWT LL subband with an L1 distance; see supp.). The clean and EOT watermark losses are

$$\begin{aligned} \mathcal{L}_{\text{wm}}^{\text{clean}} &= \frac{1}{|\mathcal{W}_*|} \sum_{i \in \mathcal{W}_*} \text{BCE}(\hat{\mathbf{y}}_i^{\text{cl}}, \mathbf{b}_i), \\ \mathcal{L}_{\text{wm}}^{\text{eot}} &= \mathbb{E}_{\mathbf{t} \sim p} \left[\frac{1}{|\mathcal{W}_*|} \sum_{i \in \mathcal{W}_*} \text{BCE}(\hat{\mathbf{y}}_i^{\text{eot}}(\mathbf{t}), \mathbf{b}_i) \right], \end{aligned} \quad (20)$$

and the total watermark objective is

$$\mathcal{L}_{\text{wm}} = \lambda_{\text{wm}}^{\text{clean}} \mathcal{L}_{\text{wm}}^{\text{clean}} + \lambda_{\text{wm}}^{\text{eot}} \mathcal{L}_{\text{wm}}^{\text{eot}} + \lambda_{\text{wav}}^{\text{low}} \mathcal{L}_{\text{wav}}^{\text{low}}. \quad (21)$$

Finally, although we report the joint scalar objective $\mathcal{L}_{\text{total}} = \lambda_{\text{vis}} \mathcal{L}_{\text{vis}} + \lambda_{\text{wm}} \mathcal{L}_{\text{wm}}$, the actual parameter updates follow the masked routing rule in Eq. (17): watermark gradients act only on \mathcal{W}_* , and visual gradients act only on \mathcal{VIS} . This disjoint gradient routing effectively eliminates the optimization interference in joint finetuning.

4. Experiments

4.1. Experimental Setting

Dataset & Pre-trained 3D-GS. We evaluate on the standard benchmarks in the NeRF [22] and 3D-GS [10] literature—Blender [21], LLFF [20], and Mip-NeRF 360 [23]. Following common practice, we report results on 25 scenes drawn from the full versions of these datasets.

Baseline. We compare our method against three strategies for fairness: (i) WaterRF [7], an innovative watermarking method deployable in both the NeRF representations; for

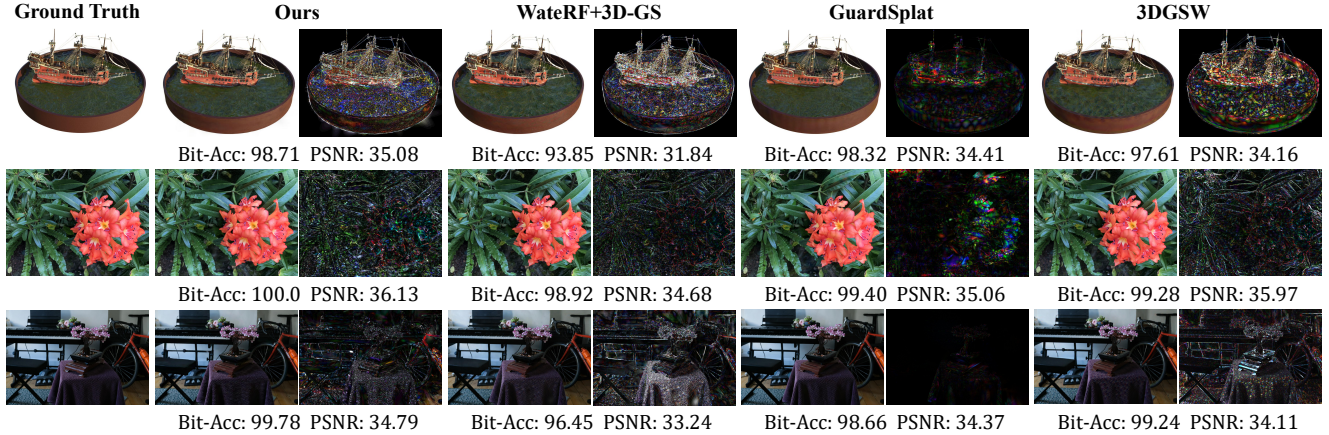


Figure 3. **Rendering-quality comparison.** We compare our method with all baselines using 32-bit messages. The difference maps are shown at 10× scale. Our approach achieves higher bit accuracy and better visual fidelity than competing methods.

Methods	32 bits				48 bits				64 bits			
	Bit Acc↑	PSNR ↑	SSIM ↑	LPIPS ↓	Bit Acc↑	PSNR ↑	SSIM ↑	LPIPS ↓	Bit Acc↑	PSNR ↑	SSIM ↑	LPIPS ↓
WateRF [7]+3D-GS [10]	93.28	30.57	0.954	0.052	84.39	30.06	0.949	0.056	74.92	25.73	0.887	0.105
GuardSplat [3]	95.58	35.32	0.978	0.043	93.29	33.36	0.969	0.045	90.14	32.25	0.963	0.048
3D-GSW [8]	97.22	35.15	0.977	0.044	93.59	33.26	0.972	0.047	91.31	32.52	0.966	0.050
Ours	98.46	35.98	0.982	0.041	94.29	33.45	0.973	0.044	91.65	32.71	0.969	0.047

Table 1. **Bit accuracy and rendering quality compared with baselines.** We report results for 32, 48, and 64-bit messages, averaged over the Blender, LLFF, and Mip-NeRF 360 datasets. The best scores are shown in **bold**.

fairness, we instantiate and evaluate it in the 3D-GS setting. (ii) GuardSplat [3], a CLIP-guided and SH-aware robust watermarking model for 3DGS, enabling invisible, secure, and distortion-resistant embedding with high visual fidelity. (iii) 3D-GSW [8], a state-of-the-art watermark embedding model operating in the frequency domain.

Implementation Details. We train all models on a single NVIDIA A800 GPU for 2–10 epochs using Adam [11]. The decoder is a frozen HiDDeN [47] model (32/48/64 bits). We prune Gaussians with $V_\pi < 10^{-8}$. All experiments are repeated with three random seeds.

Evaluation. We evaluate our watermarking framework along three key dimensions. **(1) Invisibility.** Visual fidelity is measured using Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index (SSIM) [40], and Learned Perceptual Image Patch Similarity (LPIPS) [45]. **(2) Robustness.** We assess robustness by reporting bit accuracy under a variety of image-space degradations: Gaussian noise ($\sigma = 0.1$), rotation (randomly sampled from $\pm\pi/6$), scaling to 75% of the original size, Gaussian blur ($\sigma = 0.1$), cropping to 40%, JPEG compression (quality 50%), and a composite attack that combines Gaussian noise, crop, and JPEG compression. We also evaluate model-space distortions by removing 20% of Gaussians, cloning 20% of Gaussians, and adding Gaussian noise ($\sigma = 0.1$) to 3D-GS parameters. **(3) Capacity.** We further investigate bit accuracy under different payload sizes, with message lengths

$$M_b \in \{32, 48, 64\}.$$

4.2. Results

Rendering Quality and Bit Accuracy. We compare our approach with existing methods in terms of visual fidelity and watermark extraction. As shown in Fig. 3, our results are closest to the original views while still achieving high bit accuracy. This is especially notable on real-world scenes, where complex geometry makes faithful rendering harder. Unlike prior methods that struggle to trade off image quality and watermark reliability, our approach maintains both. Quantitative results in Tab. 1 further confirm that our method consistently preserves rendering quality and bit accuracy across all datasets.

Robustness to Image Distortions. In this section, we evaluate how well our method withstands post-processing on rendered images, where such operations may alter or weaken the embedded watermark. We measure bit accuracy under various distortion types. As shown in Table. 2, most existing methods fail to maintain high robustness. Moreover, WateRF+3D-GS and 3D-GSW without EOT-based adversarial training still struggle to reliably preserve embedded messages in rendered views. GuardSplat does introduce EOT to improve robustness, but its heavy reliance on the CLIP decoder and SH-space embedding makes it less effective under more complex or compound distortions. In contrast, our method performs strict selection and de-



Figure 4. **Rendering-quality comparison under different message capacities.** With 32-bit, 48-bit, and 64-bit embedded messages, (differences $\times 10$), our method maintains high bit-acc across different capacities while preserving perceptual quality.

Methods	Bit Accuracy(%) \uparrow							
	No Distortion	Gaussian Noise ($\sigma = 0.1$)	Rotation ($\pm\pi/6$)	Scaling (75%)	Gaussian Blur ($\sigma = 0.1$)	Crop (40%)	JPEG Compression (50% quality)	Combined
WaterRF [7]+3D-GS [10]	93.28	78.12	81.47	84.63	87.09	84.58	82.03	64.73
GuardSplat [3]	95.58	90.11	95.87	94.93	97.16	95.05	89.92	88.64
3D-GSW [8]	97.22	83.71	88.05	94.58	95.94	92.73	92.54	90.96
Ours	98.46	91.22	96.18	95.06	97.75	95.88	92.95	91.30

Table 2. **Quantitative robustness comparison under different attacks against baseline methods.** Results are averaged over the Blender, LLFF, and Mip-NeRF 360 datasets using 32-bit messages. The best scores are marked in **bold**.

Methods	Bit Accuracy(%) \uparrow			
	No Distortion	Adding Gaussian Noise ($\sigma = 0.1$)	Removing 3D Gaussians (20%)	Cloning 3D Gaussians (20%)
WaterRF [7]+3D-GS [10]	93.28	62.35	60.91	76.16
GuardSplat [3]	95.58	79.96	88.49	92.30
3D-GSW [8]	97.22	89.93	97.23	96.85
Ours	98.46	90.52	98.16	97.48

Table 3. Robustness under model-level distortions with a 32-bit payload. Best results are shown in **bold**.

coupled embedding, writing watermarks only into safe 3D Gaussians, thereby achieving consistently robust extraction across all distortion settings.

Robustness to Model-level Distortions. To evaluate robustness under malicious model-level tampering, we perturb the 3D-GS representation by adding parameter noise and randomly removing or cloning 3D Gaussians. As reported in Tab. 3, our method consistently outperforms prior works across these distortions, indicating that the embedded message is not tied to a fragile subset of primitives. Notably, it maintains reliable decoding even when the underlying model is partially corrupted, demonstrating robust protection for both the 3D-GS model and its rendered outputs.

Message Capacity. Since there is an inherent trade-off among bit accuracy, rendering quality, and payload size, we investigate different message lengths 32, 48, 64. As shown in Table. 1, increasing the message length consistently leads to a slight drop in both bit accuracy and rendering quality. Nevertheless, our method maintains a favorable balance between invisibility and capacity, and the performance gap over other methods becomes more evident as the message length grows. Fig. 4 further illustrates that our approach

achieves a better compromise between bit accuracy and visual fidelity.

4.3. Ablation Study

In this section, we ablate the **SBAG**, **Group Mask**, and **Decoupled Finetuning** to clarify their necessity. Fig. 5 and Tab. 4 show that removing any one of them leads to either lower bit accuracy or degraded rendering, and removing all of them breaks the quality-robustness balance. This indicates that all three are required to stabilize watermark embedding.

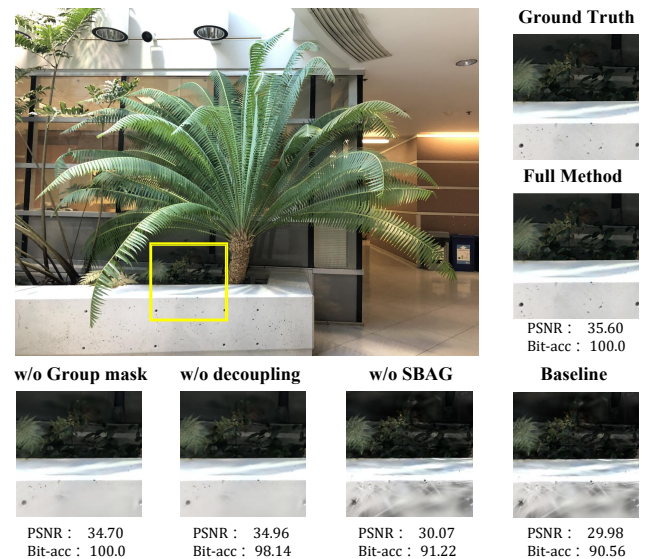


Figure 5. Comparison of rendering quality among the full method (ours), without SBAG, without group mask, without decoupled finetuning, and the baseline model. All images are embedded with 32-bit messages.

Methods			Ours			
SBAG	Group Mask	Decoupling	Bit Acc(%) \uparrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
-	-	-	94.70	30.00	0.952	0.073
\checkmark	\checkmark	-	96.80	34.10	0.974	0.048
\checkmark	-	\checkmark	97.10	33.70	0.972	0.051
-	\checkmark	\checkmark	95.10	33.10	0.968	0.056
\checkmark	\checkmark	\checkmark	97.80	35.20	0.979	0.042

Table 4. **Quantitative ablation results.** The best performance is obtained when all components are activated, with evaluation conducted on 32-bit message embedding.

Budget	Bit Acc \uparrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	Storage \downarrow
1%	97.19	34.98	0.980	0.046	164MB
Adaptive	98.46	35.98	0.982	0.041	178MB
10%	98.12	34.24	0.975	0.048	195MB

Table 5. Ablation study on Budget for Watermark Carrier Gaussians. Results represent the average score across Blender, LLFF, and Mip-NeRF 360 datasets using 32-bit messages.

SBAG. To jointly improve bit accuracy and rendering quality during watermark embedding, we introduce the SBAG to select, expand and create high-quality carrier Gaussians. The Gaussians selected by SBAG lie in visually insensitive regions and have multiple channels that can be safely modified. Combined with decoupled training, this effectively alleviates the conflict between visual fidelity and watermark embedding. As shown in Tab. 4 and Fig. 5, SBAG-based selection consistently boosts both rendering quality and bit accuracy. Moreover, Tab. 5 shows that controlled densification within an adaptive budget can further improve watermark performance with lower storage cost.




Figure 6. **Training dynamics of coupled vs. decoupled finetuning.** Decoupled optimization converges faster and consistently achieves higher PSNR and BitAcc across finetuning steps.

Group mask. Before watermark embedding, the pretrained 3D-GS already achieves high rendering quality, so we want to avoid large updates on the majority of non-carrier (i.e., visual) Gaussians. In contrast, carrier Gaussians should have more freedom to be perturbed across channels to encode the watermark. Our gradient group mask enforces this asymmetry during backpropagation: gradients on visual Gaussians are suppressed below those of prior methods, while gradients on carrier Gaussians are adaptively scaled per channel. This preserves rendering quality while



Figure 7. **Qualitative comparison of our group mask.** For background-free objects, it confines 3D Gaussian updates to object boundaries, preserving rendering quality.(difference $\times 8$)

maintaining watermark accuracy. As shown in Fig. 7, our group-aware mask achieves a better balance between visual fidelity and watermark performance than alternative masking strategies.

Decoupled Finetuning. To verify the necessity of decoupled finetuning, we analyze the optimization conflict between the visual and watermark objectives by running coupled and decoupled finetuning under identical settings and periodically evaluating PSNR and BitAcc on the same validation set. As shown in Fig. 6, coupled finetuning exhibits a transient degradation in BitAcc while PSNR keeps improving, revealing a clear trade-off, whereas decoupled optimization converges faster and maintains consistently higher PSNR and BitAcc throughout finetuning.

We provide additional ablation results in the supplementary material.

5. Conclusion

We conclude with a representation-native framework that unifies where to write, what to write, and why it matters. **Where:** Trio-Experts derive geometry, appearance, and redundancy priors from Gaussian parameters, and together with a safety and budget aware SBAG plus single-pass rendering cues, select carriers that are robust yet visually safe, training them separately from visual Gaussians. **What:** A channel wise group mask constrains gradients to designated parameter channels, preserving rendering quality while encoding the signal. **Why:** This separation makes watermarking controllable and interpretable, sustains high bit accuracy under image-space and model-space attacks, and yields per Gaussian attribution for auditable verification; the design naturally extends to dynamic scenes and multimodal payloads.

Limitations. Our frequency decoupled training requires careful tuning of loss weights to balance quality and robustness; we provide empirically validated defaults, but extreme configurations can degrade performance. In addition, as with other neural watermarking methods, the approach depends on a pretrained decoder, and its robustness bounds the overall system performance.

Acknowledgements

This work was supported by JST BOOST (Japan), Grant Number JPMJBS2429.

References

- [1] Mauro Barni, Franco Bartolini, and Alessandro Piva. Improved wavelet-based watermarking through pixel-wise masking. *IEEE transactions on image processing*, 10(5): 783–791, 2001. 2
- [2] Yue Chen, Xingyu Chen, Anpei Chen, Gerard Pons-Moll, and Yuliang Xiu. Feat2gs: Probing visual foundation models with gaussian splatting. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 6348–6361, 2025. 2
- [3] Zixuan Chen, Guangcong Wang, Jiahao Zhu, Jianhuang Lai, and Xiaohua Xie. Guardsplat: Efficient and robust watermarking for 3d gaussian splatting. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 16325–16335, 2025. 1, 2, 3, 6, 7
- [4] Yuanyuan Gao, Hao Li, Jiaqi Chen, Zhengyu Zou, Zhihang Zhong, Dingwen Zhang, Xiao Sun, and Junwei Han. Citygs-x: A scalable architecture for efficient and geometrically accurate large-scale scene reconstruction, 2025. 2
- [5] Georg Hess, Carl Lindström, Maryam Fatemi, Christoffer Petersson, and Lennart Svensson. Splatad: Real-time lidar and camera rendering with 3d gaussian splatting for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11982–11992, 2025. 1
- [6] Yi-Hua Huang, Yang-Tian Sun, Ziyi Yang, Xiaoyang Lyu, Yan-Pei Cao, and Xiaojuan Qi. Sc-gs: Sparse-controlled gaussian splatting for editable dynamic scenes. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4220–4230, 2024. 2
- [7] Youngdong Jang, Dong In Lee, MinHyuk Jang, Jong Wook Kim, Feng Yang, and Sangpil Kim. Waterf: Robust watermarks in radiance fields for protection of copyrights. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12087–12097, 2024. 1, 2, 3, 5, 6, 7
- [8] Youngdong Jang, Hyunje Park, Feng Yang, Heeju Ko, Euijin Choo, and Sangpil Kim. 3d-gsw: 3d gaussian splatting for robust watermarking. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 5938–5948, 2025. 1, 2, 3, 6, 7
- [9] Yingwenqi Jiang, Jiadong Tu, Yuan Liu, Xifeng Gao, Xiaoxiao Long, Wenping Wang, and Yuexin Ma. Gaussian-shader: 3d gaussian splatting with shading functions for reflective surfaces. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5322–5332, 2024. 2
- [10] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139–1, 2023. 1, 2, 3, 5, 6, 7
- [11] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017. 6
- [12] Muhammed Kocabas, Jen-Hao Rick Chang, James Gabriel, Oncel Tuzel, and Anurag Ranjan. Hugs: Human gaussian splats. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 505–515, 2024. 1
- [13] Chenxin Li, Brandon Y. Feng, Zhiwen Fan, Panwang Pan, and Zhangyang Wang. Steganerf: Embedding invisible information within neural radiance fields. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 441–453, 2023. 3
- [14] Wanhua Li, Yujie Zhao, Minghan Qin, Yang Liu, Yuanhao Cai, Chuang Gan, and Hanspeter Pfister. Langsplatv2: High-dimensional 3d language gaussian splatting with 450+ fps. *arXiv preprint arXiv:2507.07136*, 2025. 1
- [15] Zhe Li, Zerong Zheng, Lizhen Wang, and Yebin Liu. Animatable gaussians: Learning pose-dependent gaussian maps for high-fidelity human avatar modeling. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 19711–19722, 2024. 2
- [16] Jiaqi Lin, Zhihao Li, Xiao Tang, Jianzhuang Liu, Shiyong Liu, Jiayue Liu, Yangdi Lu, Xiaofei Wu, Songcen Xu, Youliang Yan, et al. Vastgaussian: Vast 3d gaussians for large scene reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5166–5175, 2024. 2
- [17] Jinfeng Liu, Lingtong Kong, Bo Li, and Dan Xu. Gausshdr: High dynamic range gaussian splatting via learning unified 3d and 2d local tone mapping. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 5991–6000, 2025. 2
- [18] Ange Lou, Benjamin Planche, Zhongpai Gao, Yamin Li, Tianyu Luan, Hao Ding, Terrence Chen, Jack Noble, and Ziyang Wu. Darenerf: Direction-aware representation for dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5031–5042, 2024. 2
- [19] Ziyuan Luo, Qing Guo, Ka Chun Cheung, Simon See, and Renjie Wan. Copyrnerf: Protecting the copyright of neural radiance fields. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 22401–22411, 2023. 3
- [20] Ben Mildenhall, Pratul P. Srinivasan, Rodrigo Ortiz-Cayon, Nima Khademi Kalantari, Ravi Ramamoorthi, Ren Ng, and Abhishek Kar. Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. *ACM Transactions on Graphics (TOG)*, 2019. 5
- [21] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *European Conference on Computer Vision (ECCV)*, pages 405–421. Springer, 2020. 5
- [22] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. 1, 5
- [23] Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, Peter Hedman, Ricardo Martin-Brualla, and Jonathan T. Barron.

- MultiNeRF: A Code Release for Mip-NeRF 360, Ref-NeRF, and RawNeRF, 2022. 5
- [24] S Kother Mohideen, S Arumuga Perumal, and M Mohamed Sathik. Image de-noising using discrete wavelet transform. *International Journal of Computer Science and Network Security*, 8(1):213–216, 2008. 2
- [25] Hung Nguyen, Blark Runfa Li, and Truong Nguyen. Dwtnerf: Boosting few-shot neural radiance fields via discrete wavelet transform. *arXiv preprint arXiv:2501.12637*, 2025. 2
- [26] Hung Nguyen, Runfa Li, An Le, and Truong Nguyen. Dwtgs: Rethinking frequency regularization for sparse-view 3d gaussian splatting. *arXiv preprint arXiv:2507.15690*, 2025. 2
- [27] Hyunwoo Park, Gun Ryu, and Wonjun Kim. Dropgaussian: Structural regularization for sparse-view gaussian splatting. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 21600–21609, 2025. 2
- [28] Minghan Qin, Wanhua Li, Jiawei Zhou, Haoqian Wang, and Hanspeter Pfister. Langsplat: 3d language gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20051–20060, 2024. 1
- [29] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision, 2021. 3
- [30] Mehul S Raval and Priti P Rege. Discrete wavelet transform based multiple watermarking scheme. In *TENCON 2003. Conference on Convergent Technologies for Asia-Pacific Region*, pages 935–938. IEEE, 2003. 2
- [31] Daniel Rho, Byeonghyeon Lee, Seungtae Nam, Joo Chan Lee, Jong Hwan Ko, and Eunbyung Park. Masked wavelet representation for compact neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20680–20690, 2023. 2
- [32] Fei Shen and Jinhui Tang. Imagpose: A unified conditional framework for pose-guided person generation. *Advances in neural information processing systems*, 37:6246–6266, 2024. 2
- [33] Fei Shen, Xin Jiang, Xin He, Hu Ye, Cong Wang, Xiaoyu Du, Zechao Li, and Jinhui Tang. Imagdressing-v1: Customizable virtual dressing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 6795–6804, 2025. 2
- [34] Fei Shen, Jian Yu, Cong Wang, Xin Jiang, Xiaoyu Du, and Jinhui Tang. Imaggarment-1: Fine-grained garment generation for controllable fashion design. *arXiv preprint arXiv:2504.13176*, 2025. 2
- [35] Mark J Shensa. The discrete wavelet transform: wedding the a trous and mallat algorithms. *IEEE Transactions on signal processing*, 40(10):2464–2482, 2002. 2
- [36] Qi Song, Ziyuan Luo, Ka Chun Cheung, Simon See, and Renjie Wan. Protecting nerfs’ copyright via plug-and-play watermarking base model. In *ECCV*, 2024. 3
- [37] Chinmay Talegaonkar, Yash Belhe, Ravi Ramamoorthi, and Nicholas Antipa. Volumetrically consistent 3d gaussian rasterization. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 10953–10963, 2025. 2
- [38] Peining Tao and Ahmet M Eskicioglu. A robust multiple watermarking scheme in the discrete wavelet transform domain. In *Internet Multimedia Management Systems V*, pages 133–144. SPIE, 2004. 2
- [39] Hanzhang Tu, Zhanfeng Liao, Boyao Zhou, Shunyuan Zheng, Xilong Zhou, Liuxin Zhang, QianYing Wang, and Yebin Liu. Gbc-splat: Generalizable gaussian-based clothed human digitalization under sparse rgb cameras. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 26377–26387, 2025. 1
- [40] Zhou Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4): 600–612, 2004. 6
- [41] Guanjun Wu, Taoran Yi, Jiemin Fang, Lingxi Xie, Xiaopeng Zhang, Wei Wei, Wenyu Liu, Qi Tian, and Xinggang Wang. 4d gaussian splatting for real-time dynamic scene rendering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 20310–20320, 2024. 2
- [42] Haofei Xu, Songyou Peng, Fangjinhua Wang, Hermann Blum, Daniel Barath, Andreas Geiger, and Marc Pollefeys. Depthspat: Connecting gaussian splatting and depth. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 16453–16463, 2025. 2
- [43] Muyu Xu, Fangneng Zhan, Jiahui Zhang, Yingchen Yu, Xiaojin Zhang, Christian Theobalt, Ling Shao, and Shijian Lu. Wavenerf: Wavelet-based generalizable neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 18195–18204, 2023. 2
- [44] Zhiwen Yan, Weng Fei Low, Yu Chen, and Gim Hee Lee. Multi-scale 3d gaussian splatting for anti-aliased rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20923–20931, 2024. 2
- [45] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. 5, 6
- [46] Shangzhan Zhang, Jianyuan Wang, Yinghao Xu, Nan Xue, Christian Rupprecht, Xiaowei Zhou, Yujun Shen, and Gordon Wetzstein. Flare: Feed-forward geometry, appearance and camera estimation from uncalibrated sparse views. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 21936–21947, 2025. 2
- [47] Jiren Zhu, Russell Kaplan, Justin Johnson, and Li Fei-Fei. Hidden: Hiding data with deep networks. In *Proceedings of the European conference on computer vision (ECCV)*, pages 657–672, 2018. 3, 6
- [48] Sicheng Zuo, Wenzhao Zheng, Yuanhui Huang, Jie Zhou, and Jiwen Lu. Gaussianworld: Gaussian world model for streaming 3d occupancy prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6772–6781, 2025. 1