

PrinGS: Privacy Preservation in Gaussian Splatting

Gaurav Behera*
IIIT Hyderabad
Telangana, India
gaurav.behera@research.iiit.ac.in

Abhinav Raundhal*
IIIT Hyderabad
Telangana, India
abhinav.raundhal@students.iiit.ac.in

Archisha Panda*
IIIT Hyderabad
Telangana, India
archisha.panda@research.iiit.ac.in

Saumya Balina*
IIIT Hyderabad
Telangana, India
saumya.balina@students.iiit.ac.in

Gargi Shroff*
IIIT Hyderabad
Telangana, India
gargi.shroff@research.iiit.ac.in

Ponnuram Kumaraguru
IIIT Hyderabad
Telangana, India
pk.guru@iiit.ac.in




Figure 1: Given only a trained splat and user-selected reference views, our method localizes the Gaussians corresponding to sensitive content and directly anonymizes the splat without requiring the original training images. The resulting privacy-preserving splat consistently hides faces across arbitrary novel viewpoints while preserving the fidelity of the remaining scene.

ABSTRACT

3D Gaussian Splatting (3DGS) has rapidly become the preferred representation for high-fidelity novel view synthesis and 3D scene reconstruction. However, the learned representation encodes sensitive visual information including human faces, license plates, documents, and computer screens, raising significant privacy concerns when trained splats are shared or deployed. In this work, we address the underexplored problem of privacy preservation in Gaussian Splatting by proposing a representation-level anonymization framework that directly anonymizes a trained splat. Given a trained splat and a user-provided reference view identifying sensitive content, our method localizes the corresponding regions in the underlying 3D representation and directly anonymizes them, ensuring that the sensitive information is absent from all subsequently rendered novel views while preserving the overall scene geometry and visual fidelity. We achieve this post-training anonymization in less than 1 second. Furthermore, we extend our framework with a privacy-by-design training strategy that prevents sensitive content from being embedded into the representation during optimization, enabling the generation of privacy-preserving splats directly from training images. Our method achieves a remarkably low Face Similarity Measure (FSM) of 9.79% at a retained Face Detection Rate (FDR) of 39.47% and makes sensitive text completely unreadable with an OCR Character Error Rate of 1.00. Our work aims to serve as the primary stepping stone towards privacy-aware 3D representations and motivates further research along these directions.

*Equal Contribution. Project Page: <https://privacypreserving3dgs.github.io>

ICVGIP2026, December 21-25, 2026, Kolkata, India
2026. ACM ISBN 978-x-xxxx-xxxx-x/YY/MM...\$15.00
<https://doi.org/10.1145/nnnnnnnn.nnnnnnn>

CCS CONCEPTS

• **Computing methodologies** → **Computer vision**; *Reconstruction*; • **Security and privacy** → Pseudonymity, anonymity and untraceability.

KEYWORDS

Gaussian Splatting, Privacy Preservation, Anonymization

ACM Reference Format:

Gaurav Behera, Abhinav Raundhal, Archisha Panda, Saumya Balina, Gargi Shroff, and Ponnuram Kumaraguru. 2026. PrinGS: Privacy Preservation in Gaussian Splatting. In *Proceedings of 17th Indian Conference on Computer Vision, Graphics and Image Processing (ICVGIP2026)*. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/nnnnnnnn.nnnnnnn>

1 INTRODUCTION

The widespread adoption of 3D Gaussian Splatting (3DGS) [15] has enabled creating and sharing highly realistic 3D scene representations across applications such as robotics, autonomous navigation, digital twins, cultural heritage preservation, and virtual reality [14, 32, 39]. While these representations excel at preserving visual fidelity, they also make it possible to recover information that users may expect to remain private, creating new challenges for the safe sharing and deployment of learned scene representations. Since this information is embedded within the learned 3D representation it remains accessible from arbitrary novel viewpoints, introducing significant privacy risks and raising challenges for compliance with regulations such as the General Data Protection Regulation (GDPR) [7] and the Digital Personal Data Protection (DPDP) Act [23].

An increasing number of splats are being publicly shared and deployed across real-world applications. A challenge arises when an individual notices their identity captured in a scene without consent and requests its removal to comply with privacy regulations. Retraining from scratch is computationally infeasible. Additionally, in many scenarios, the original training images used for reconstruction are unavailable, inaccessible, or prohibitively expensive to reacquire. Consequently, there is a pressing need for techniques that can anonymize existing splats without requiring access to the original training data.

A straightforward solution is to anonymize rendered images after novel-view synthesis. However, such image-space anonymization is inherently view-dependent and leaves the underlying 3D representation unchanged, allowing sensitive information to remain visible when the scene is rendered from different viewpoints. This limitation motivates a representation-level approach that directly anonymizes the learned splat, ensuring that sensitive information is removed across all novel views. We present PrinGS, a framework for post-training privacy preservation in Gaussian Splatting. Given a trained splat and user-selected reference views containing sensitive content, our method localizes the corresponding Gaussians in the underlying 3D representation and directly anonymizes them. By operating on the representation itself rather than rendered images, sensitive information is removed from the splat, preventing its recovery from arbitrary novel viewpoints while preserving scene geometry and fidelity as shown in Figure 1.

While post-training anonymization is essential for existing splats, the rapid adoption of 3DGS is expected to result in an ever-growing number of new 3D scene representations. This presents an opportunity to incorporate privacy preservation directly into the reconstruction pipeline rather than addressing it retrospectively. Motivated by this, we further introduce a privacy-by-design framework that integrates our representation-level anonymization mechanism into the optimization process, preventing sensitive information from being encoded into the splat during training.

Existing 3DGS benchmarks predominantly lack privacy-sensitive content like faces, creating a critical gap that hinders the evaluation of privacy aware 3D scene reconstructions. To address this, we construct, PrinGS-Humans, a dataset of synthetic indoor scenes containing 9 human subjects across varying demographics like skin tones and genders, establishing the first benchmark for systematic evaluation of privacy preservation of faces in Gaussian Splatting.

Furthermore, conventional rendering metrics primarily assess visual fidelity and do not capture privacy preservation. We therefore quantitatively evaluate our privacy-by-design approach using novel task-specific metrics based on Face Similarity Measure (FSM) and Face Detection Rate (FDR) for faces and OCR Character Error Rate to quantify the removal of textual sensitive information. In terms of facial privacy, our method achieves a Face Similarity Metric (FSM) of only 9.79%, representing a substantial reduction (74.08%) compared to standard 3DGS (83.87%). Furthermore, for textual content, our approach attains a Character Error Rate (CER) of 1.00, indicating complete unreadability of private text, significantly outperforming baseline 3DGS, which yields a CER of 0.208.

Our main contributions are summarized as follows:

- **PrinGS.** We present the first framework for post-training privacy preservation in Gaussian Splatting that directly anonymizes the learned representation without requiring the original training images, preventing sensitive information from being recovered from novel viewpoints in just 1 second.
- **Privacy-by-Design.** We introduce a privacy-by-design extension that incorporates anonymization into the Gaussian Splatting optimization process, preventing sensitive information from being encoded during training.
- **PrinGS-Humans.** To address the lack of static scenes with humans due to privacy reasons, we construct a benchmark of synthetic indoor scenes containing 9 human subjects across varying demographics, enabling systematic evaluation of privacy-preserving Gaussian Splatting methods.
- **Privacy Metrics.** We introduce novel evaluation protocols tailored to the 3D Gaussian Splatting domain, specifically utilizing the Face Similarity Measure (FSM), Face Detection Rate (FDR) and OCR Character Error Rate (CER) to rigorously quantify privacy preservation, demonstrating a 74.08% reduction in FSM and a perfect CER of 1.00.

2 RELATED WORK

2.1 3D Gaussian Splatting

Recent advancements in novel view synthesis have shifted from continuous volumetric fields (NeRF) [22] to explicit and efficient representations. 3D Gaussian Splatting (3DGS) [15] models scenes as collections of anisotropic 3D Gaussians parameterized by spatial positions, covariance matrices, opacity, and spherical harmonic coefficients. This formulation bypasses costly volumetric ray marching, enabling real-time rendering and interactive scene exploration.

Building upon the foundational 3DGS framework, subsequent research has branched into multiple directions to enhance rendering quality and efficiency. Notable quality improvements include anti-aliased formulations such as Mip-Splatting [36] and specialized regularizations like AbsGS [35] that handle unbounded scenes and reduce artifacts. Another direction modifies the densification strategy such as 3DGS-MCMC [16]. To bridge the gap toward resource-constrained settings, methods like Taming 3DGS [21] introduce controlled densification to curb primitive redundancy. Building upon these efficiency measures, compact variants like Mini-Splatting [8, 9] introduce memory-efficient pruning and compression mechanisms. Despite these rapid developments in performance and visual fidelity, existing literature predominantly focuses on optimization efficiency and rendering quality while largely overlooking the security and privacy implications of capturing real-world environments.

A parallel line of work adds semantics to 3DGS by learning auxiliary per-Gaussian attributes during training, such as the SAM-distilled affinity features of SAGA [3], the identity encodings of Gaussian Grouping [34], and the feature fields or semantic tracing of Feature-3DGS [38] and GaussianEditor [4]. Since these attributes must be learned during optimization, such methods cannot be applied to an already-trained splat that does not carry them. PrinGS instead localizes sensitive Gaussians directly from the rasterizer’s transmittance-weighted contributions, and anonymizes them.

2.2 Privacy Preservation in 2D Visual Data

The rapid growth of visual data as images and videos on the internet has made privacy preservation an increasingly important research problem. To mitigate the exposure of personally identifiable information (PII), regulatory frameworks have accelerated research into automated anonymization; however, existing efforts remain predominantly confined to the 2D domain.

For images, anonymization is achieved as a two-step process: detection or segmentation followed by blurring. For detection and segmentation, frameworks such as YOLO [13, 33], Faster R-CNN [25], Mask R-CNN [12], and SAM [17] are utilized to accurately localize sensitive regions. Following detection, privacy is enforced by applying conventional filtering and obfuscation mechanisms such as localized spatial blurring, pixelation, masking, or identity scrambling to protect sensitive regions [28]. In videos, privacy preservation builds upon these image-based techniques by incorporating temporal consistency through multi-object tracking algorithms [5, 31] to prevent identity leakage across successive frames.

2.3 Privacy Preservation in 3D Scenes

While privacy preservation has been extensively studied in 2D images and videos, extending these mechanisms to 3D scene representations remains a largely underexplored non-trivial task. In 3D environments, privacy preservation requires ensuring that sensitive content is completely anonymized across all arbitrary novel viewpoints rather than just a fixed set of frames.

Completely removing sensitive objects from a 3D scene is often impractical, as it degrades the semantic richness required by downstream applications. For example, autonomous driving systems require pedestrians to learn safe navigation, but their facial identities are unnecessary. Similarly, many robotics and digital twin applications benefit from preserving human presence while suppressing personally identifiable information. This motivates representation-level anonymization that retains scene semantics while removing identity-bearing visual attributes.

A naive baseline is to apply image-space anonymization directly to the training images prior to 3D reconstruction. However, this approach introduces fundamental limitations. First, masking or blurring training images degrades local feature quality, severely hindering feature matching and camera pose estimation in Structure-from-Motion (SfM) pipelines like COLMAP [29, 30], which serve as a critical prerequisite for 3DGS. Second, per-view image-space modifications often yield inconsistent appearances across viewpoints, destabilizing the optimization landscape for 3DGS (refer supplement for detailed analysis).

To address these limitations, our work addresses privacy directly at the representation level. We are the first to formalize a privacy-preserving framework explicitly designed for 3DGS enabling view-consistent anonymization across arbitrary viewpoints while preserving scene geometry and rendering quality.

3 METHODOLOGY

This section details our framework for privacy-preserving 3D Gaussian Splatting. We first introduce our approach for accurately segmenting sensitive Gaussians in 3D space (Section 3.1). Next, we define the localized anonymization operators for blurring sensitive

Algorithm 1 3D Segmentation of Sensitive Gaussians

Input: Set of views $\mathcal{V}$, corresponding binary masks $\mathcal{M} = \{M_v\}_{v \in \mathcal{V}}$, 3D Gaussians $\mathcal{G} = \{g_1, \dots, g_N\}$, transmittance threshold τ_T , global visibility threshold τ_v

Output: Sensitive Gaussian subset $\mathcal{S}$

```

1: Initialize global visibility counters  $C_i \leftarrow 0$  for all  $g_i \in \mathcal{G}$ 
2: for each view  $v \in \mathcal{V}$  do
3:   A Single-view identification
4:   Initialize  $c_{\text{true},i} \leftarrow 0$  and  $c_{\text{false},i} \leftarrow 0$  for all  $g_i \in \mathcal{G}$ 
5:   for each pixel  $p$  in view  $v$  do
6:     for each overlapping Gaussian  $g_i$  at pixel  $p$  do
7:        $T_i \leftarrow \text{EvaluateTransmittance}()$ 
8:       if  $T_i > \tau_T$  then
9:         if  $M_v(p) == 1$  then
10:           $c_{\text{true},i} \leftarrow c_{\text{true},i} + 1$ 
11:        else
12:           $c_{\text{false},i} \leftarrow c_{\text{false},i} + 1$ 
13:        end if
14:      end if
15:    end for
16:  end for
17:   $\mathcal{S}_v \leftarrow \{g_i \in \mathcal{G} \mid c_{\text{true},i} > c_{\text{false},i}\}$ 
18:   $\mathcal{S}'_v \leftarrow \text{IsolationForest}(\mathcal{S}_v)$  on  $\mu_i$ 
19:  B Multi-view aggregation
20:  for each  $g_i \in \mathcal{S}'_v$  do
21:     $C_i \leftarrow C_i + 1$ 
22:  end for
23: end for
24: return  $\mathcal{S} \leftarrow \{g_i \in \mathcal{G} \mid C_i \geq \tau_v\}$ 

```

content (Section 3.2). Finally, we describe how these techniques are applied in two distinct scenarios: modifying an existing Gaussian Splat post-training (Section 3.3), and an end-to-end privacy-by-design pipeline that prevents sensitive features from being learned during optimization (Section 3.4). An overview of the proposed framework is illustrated in Figures 2 and 3. Figure 2 presents the segmentation of sensitive Gaussians in 3D along with the localized anonymization operators applied to them. Figure 3 provides an overview of our two techniques: post-training anonymization of an existing splat and privacy-by-design training that prevents sensitive content from being encoded during optimization.

3.1 3D Segmentation of Sensitive Gaussians

To remove sensitive information from a 3D scene, we must first accurately segment in 3D the specific Gaussians responsible for rendering those regions. While heuristic COLMAP based 2D-to-3D projection methods [24] exist for 3D text segmentation, they often struggle with complex volumetric occlusions and background artifacts. To overcome these limitations, we integrate the segmentation directly into the differentiable rasterization pipeline, constructing a precise representation-level mapping by tracking splat contributions during alpha compositing stage of the rasterizer.

Our downstream tasks demand either single-view segmentation for post-training anonymization, or a robust accumulation throughout the entire scene to handle multiple views for privacy-by-design.




Figure 2: Overview of core components for our privacy preservation pipeline. **A** Single-view segmentation identifies sensitive Gaussians via rasterization against 2D masks combined with Isolation Forest outlier rejection. **B** Multi-view aggregation uses a visibility counter threshold to isolate the final set of sensitive Gaussians. Once identified, localized volumetric blur is applied using k -NN neighborhood. **C** Face anonymization modifies the spherical harmonic coefficients with scale-dependent weights. **D** Text anonymization, on the other hand, modifies both SH coefficients and spatial means with uniform weight. The final output is a splat with sensitive information anonymized.

To support these two scenarios, our segmentation process offers two corresponding methods: single-view identification **A** and multi-view aggregation **B**, respectively. The overall logic is shown in Algorithm 1.

A For a given view v , let $M_v \in \{0, 1\}^{H \times W}$ be the binary mask indicating the sensitive 2D region. During the tile-based forward pass, for each pixel p and each overlapping Gaussian g_i , we evaluate its accumulated transmittance T_i . If $T_i > \tau_T$, we atomically update one of two contribution counters for g_i : c_{true} if $p \in M_v$ (i.e., $M_v(p) = 1$), and c_{false} if $p \notin M_v$. A Gaussian g_i is classified as sensitive for view v if $c_{\text{true}} > c_{\text{false}}$. Additionally, as rendering is projective, the background geometries may coincidentally project into M_v . To robustly filter these spurious floaters, we apply an Isolation Forest algorithm [20] to the 3D spatial coordinates μ_i of all flagged Gaussians in the current frame, discarding spatial outliers.

B To aggregate these filtered identifications across multiple views, we maintain a global visibility counter C_i for each Gaussian g_i , incrementing it whenever the Gaussian is marked sensitive in a view. The final sensitive subset S is defined as $S = \{g_i \in \mathcal{G} \mid C_i \geq \tau_v\}$, where τ_v is a global visibility threshold.

3.2 Gaussian Anonymization Formulation

Let the 3D Gaussian representation be denoted as $\mathcal{G} = \{g_1, g_2, \dots, g_N\}$, where each Gaussian g_i is defined by its explicit parameters, including its spatial mean $\mu_i \in \mathbb{R}^3$, spherical harmonics (SH) coefficients c_i , and scaling vector s_i . Once the segmentation process identifies

the sensitive subset $S \subset \mathcal{G}$, we anonymize these targeted Gaussians by applying a localized volumetric blur to their parameters.

For each sensitive Gaussian $g_i \in S$, we find its k -nearest spatial neighbors $\mathcal{N}_k(i)$. We then iteratively update a target parameter θ_i over t steps using a weighted average of this neighborhood:

$$\theta_i^{(t+1)} = \frac{1}{\sum_{j \in \mathcal{N}_k(i)} w_j} \sum_{j \in \mathcal{N}_k(i)} w_j \theta_j^{(t)}$$

The specific parameter θ_i updated and the weight w_j applied depend on the type of sensitive content being anonymized.

C For face anonymization, preserving the fundamental 3D structure of the head is necessary to maintain scene coherence. Facial identity characteristics are predominantly tied to high-frequency visual details such as eyes, nose, and lips, which are naturally represented by Gaussians with smaller scales. Therefore, we restrict the anonymization operator strictly to the appearance parameters (the SH coefficients, $\theta_i = c_i$) and use a scale-dependent weight $w_j = f(s_j)$. This approach ensures that larger, structural background splats appropriately guide the color blending by stripping away identifying details while leaving the spatial geometry intact.

D For text anonymization, readable characters are mostly planar, and preserving their local geometric structure can inadvertently leave readable physical cues. We therefore expand the anonymization to target both appearance and geometry ($\theta_i = \{c_i, \mu_i\}$) using a uniform, unweighted filter ($w_j = 1$). This effectively dissolves both the text's coloring and its physical shape, collapsing distinct character geometries into unreadable text.




Figure 3: Overview of our two operational settings. Top: Post-Training Privacy Preservation takes an already-trained 3D Gaussian Splat along with a reference camera view and binary mask, to produce an anonymized splat without accessing original training images. Bottom: Privacy-by-Design integrates our anonymization pipeline directly into the training iterations along with loss masking to yield a anonymized splat.

3.3 Post-Training Privacy Preservation

Our framework can seamlessly modify an already optimized 3D Gaussian Splat without requiring the original training data. The inputs for this post-training setting are simply the existing 3D Gaussian Splat, a single reference camera view capturing the sensitive region and its corresponding 2D binary privacy mask.

Because this setting relies on a single user-provided view, we utilize the single-view segmentation process (A) to identify the sensitive Gaussians. Once isolated, the corresponding anonymization operator is applied directly to the static representation: facial regions undergo scale-weighted appearance blurring (C), while text regions receive uniform appearance and structural blurring (D). By altering the core representation, the sensitive information is permanently erased across all novel views.

3.4 Privacy-by-Design

While post-training edits are effective, retrospectively modifying a highly optimized scene may sometimes leave back traces of private information, if sensitive Gaussians are deeply entangled with the scene structure. Our Privacy-by-Design approach mitigates this by integrating anonymization directly into the standard 3DGS training pipeline, ensuring sensitive features are never encoded from the input images.

Face Anonymization: For facial regions, interrupting the optimization gradients mid-training degrades the underlying geometric

fidelity of the head, which is needed for view-dependent coherence. Therefore, we apply anonymization only after the standard training iterations (e.g., 30K) when the geometry of the faces has converged. This involves multi-view segmentation (B) followed by applying the scale-weighted facial anonymization operator (C) to the final splat. This preserves the properly learned geometry while completely removing high-frequency facial details.

Text Anonymization: Text geometries are highly sensitive to even weak gradients, meaning sharp, readable artifacts can easily form during training. To prevent this, we explicitly exclude text regions from the photometric loss computation using masked non-sensitive background regions:

$$\mathcal{L} = (1 - \lambda) \mathcal{L}_1(I \odot \bar{M}_v, \hat{I} \odot \bar{M}_v) + \lambda \mathcal{L}_{D-SSIM}(I \odot \bar{M}_v, \hat{I} \odot \bar{M}_v)$$

where I and $\hat{I}$ denote the ground-truth and rendered images, respectively, and $\bar{M}_v = 1 - M_v$ represents the inverted mask for view v . Furthermore, to suppress any accidental readable structures formed by unsupervised floating geometry, we periodically interrupt the training loop to apply the multi-view text segmentation (B) followed by the uniform structural anonymization operator (D). This continuous intervention ensures readable text is never optimized.

4 RESULTS AND EXPERIMENTS

We evaluate our approach in both the post-training privacy preservation and privacy-by-design settings. The former demonstrates the ability to anonymize existing Gaussian Splats without requiring access to the original training images, while the latter enables quantitative evaluation through the availability of ground-truth images. We compare against standard 3D Gaussian Splatting and analyze both qualitative and quantitative aspects of privacy preservation.

4.1 Experimental Setup

All experiments are conducted on a single NVIDIA RTX 3090 Ti GPU with 24 GB VRAM. Following the 3DGS [15] literature, every eighth frame is reserved for evaluation while the remaining frames are used for optimization. For all experiments, we use the optimization strategies for faces and text as described in Section 3. For detailed implementation details, refer to the supplement.

4.2 Datasets

4.2.1 PrinGS-Humans Dataset. Existing Gaussian Splatting benchmarks contain no human subjects suitable for evaluating facial privacy preservation. We therefore construct a synthetic benchmark in Blender comprising two indoor scenes containing 9 unique human subjects [26, 27]. The dataset includes subjects with diverse appearances, including varying skin tones, ages, genders, hairstyles, facial hair making the benchmark representative of real-world humans. We will publicly release this dataset upon acceptance.

4.2.2 Text Privacy Dataset. For evaluating textual privacy, we select five scenes containing readable text from existing public benchmarks. Specifically, we use the *Truck*, *Barn*, and *M60* scenes from Tanks and Temples [18] together with two scenes from the DL3DV-10K benchmark [19] which we call *Car* and *Piano*. These scenes contain diverse private text instances like number plates, house




Figure 4: Qualitative results for post-training anonymization on scenes from our PrinGS-Humans dataset (rows 1-4) and Tanks and Temples [18] (row 5). Starting from a single user-selected *reference view* (left), our method directly anonymizes the underlying Gaussian representation without requiring access to the original training images. The remaining columns show novel views rendered after representation-level anonymization. Sensitive facial and textual information is consistently removed across arbitrary viewpoints, demonstrating that privacy preservation naturally propagates through the edited 3D representation rather than being limited to the selected view. Notice the geometry of the faces being preserved across all novel views, with the silhouette sharply defining face contours. (Best seen on screen).

numbers, and brand names with varying font sizes, colours, orientations, and viewing distances, providing a challenging benchmark for text anonymization.

4.3 Evaluation Protocol

4.3.1 Post-training Anonymization. The post-training setting assumes that only a trained Gaussian Splat is available, without access to the original calibrated training images. Given an arbitrary reference view rendered using the SIBR viewer [1], the user specifies the sensitive region to be anonymized, which is segmented in 2D using SAM3 [2]. Our method then localizes the corresponding Gaussians in the learned representation and performs representation-level anonymization as described in Section 3.3. Since the modification is

performed directly in the Gaussian representation, the anonymization naturally propagates to all rendered viewpoints.

As the original training images are unavailable in this setting, quantitative evaluation is infeasible. We therefore demonstrate the effectiveness of our post-training framework through extensive qualitative results across multiple scenes and viewpoints in Figure 4.

4.3.2 Privacy-by-Design. To enable quantitative evaluation, we consider the privacy-by-design setting where sensitive regions are anonymized during optimization. Given input images, sensitive regions including faces, license plates, documents, and other privacy-sensitive content, are segmented using SAM3 [2] with user-provided text prompts. These masks are used by the optimization pipeline to output a privacy-preserving Gaussian representation,




Figure 5: Qualitative comparison of privacy-by-design training on scenes from DL3DV-10K benchmark [19] (rows 1-3) and our PrinGS-Humans Dataset (rows 4-6). From left to right: ground truth image, novel rendered images from 3DGS, and our method PrinGS. While 3DGS reconstructs sensitive content, our approach prevents such information from being encoded during training, effectively anonymizing the scene while preserving the appearance of non-sensitive regions.

as described in Section 3.4. Since ground-truth images are available in this setting, both qualitative and quantitative evaluation become possible. We first describe the evaluation metrics below.

4.4 Evaluation Metrics

Conventional rendering metrics such as PSNR, SSIM, and LPIPS [37] primarily measure rendering fidelity and are not designed to quantify privacy preservation. To address this, we employ three task-specific metrics: Face Similarity Measure (FSM), Face Detection Rate (FDR), and OCR Character Error Rate (CER).

4.4.1 Face Similarity Measure (FSM). To evaluate the anonymization of facial recognition features, we use the industry-standard InsightFace framework on the novel-view rendered images.

Table 1: Face Similarity Measure and Face Detection Rate on our PrinGS-Humans benchmark. Lower FSM indicates stronger anonymization, while higher FDR indicates that facial geometry is preserved. 3DGS reconstructs facial identity, whereas our method substantially suppresses identity-related features while keeping the faces detectable.

Scene	3DGS		Ours	
	FDR $\uparrow$	FSM $\downarrow$	FDR $\uparrow$	FSM $\downarrow$
House	91.67	78.68	36.11	6.55
Hospital	97.50	88.55	42.50	12.70
Overall	94.74	83.87	39.47	9.79

The evaluation pipeline utilizes an SCRFD-10GF [11] detector to scan the entire rendered image and detect faces. Following detection, we extract 512-dimensional embeddings via ArcFace [6] (ResNet-50) CNN backbone. Face Similarity Measure (FSM) is then quantified using Cosine Similarity between the embeddings of the ground-truth and anonymized faces. A lower cosine similarity indicates successful and robust anonymization.

4.4.2 Face Detection Rate (FDR). Since a ground-truth face with no corresponding detection in the render is assigned a similarity of 0, FSM alone rewards a method that destroys the face entirely rather than anonymizing it. We therefore also report the Face Detection Rate (FDR), the fraction of ground-truth faces across the evaluation views that remain detectable by the SCRFD-10GF [11] detector in the anonymized render.

4.4.3 OCR Character Error Rate (CER). To evaluate text readability, we utilize the Google Cloud Vision API [10] as a robust Optical Character Recognition (OCR) tool. Following STRINGS [24], we quantify text readability using recall-based Character Error Rate (CER), computed via the Levenshtein edit distance between text extracted from the ground-truth images and the anonymized renders. In the context of privacy preservation, a higher CER demonstrates that the underlying textual geometry has been successfully destroyed.

4.5 Privacy-by-Design Results

Figure 5 qualitatively compares the rendered images produced by standard 3D Gaussian Splatting and our privacy-preserving framework against the ground truth. Our approach effectively removes sensitive facial and textual information while maintaining high rendering quality for the remaining scene content.

Table 1 reports the Face Similarity Measure (FSM) and Face Detection Rate (FDR) on PrinGS-Humans. Because subjects appear across a varying number of evaluation frames, per-subject FSM and FDR values are computed and aggregated via weighted averages to derive scene-level metrics. Standard 3DGS yields high overall similarity at an FDR of 94.74%, indicating severe privacy leakage, as the faces are both present and identifiable. In contrast, our method drastically reduces embedding similarity to an overall average of 9.79% while still retaining an FDR of 39.47%, meaning the anonymized faces are detected as faces, yet the recognizer can no longer match them to the correct identity. FDR is based on a

Table 2: OCR Character Error Rate (OCR-CER) comparison. Higher values indicate stronger anonymization, with a value of 1.0 corresponding to complete failure of text recognition. Our method consistently suppresses readable text across all scenes, achieving the maximum OCR-CER, whereas standard 3DGS reconstructs text with substantially lower error.

OCR-CER $\uparrow$	Truck	Barn	M60	Piano	Car
3DGS	0.500	0.028	0.207	0.149	0.158
Ours	1.00	1.00	1.00	1.00	1.00

face detector which takes both facial features and geometry into account for inference, justifying the 39.47% FDR of retained geometry and lost facial features. Quantitative metrics comparing PrinGS with baselines mentioned in Subsection 2.3 are presented in the supplement. Our method is the only one that drops FSM while maintaining a reasonable FDR.

Table 2 reports the CER across all text scenes. Our approach substantially increases OCR error rates compared to standard Gaussian Splatting, indicating successful suppression of readable text. Specifically, while standard 3DGS yields an average CER of 0.208 across five evaluation scenes, our method achieves a perfect recall-based CER of 1.00, completely disrupting text recognition.

In summary, our quantitative and qualitative evaluations confirm that the proposed framework successfully masks sensitive facial and textual information across arbitrary novel viewpoints.

4.6 Ablation Studies

4.6.1 Runtime Analysis. Finally, we report the runtime of our post-training anonymization framework in Table 3 along with Number of Sensitive Gaussians, Neighbourhood Size (k) and Anonymization Iterations (t). Since our method directly operates on the learned Gaussian representation without requiring retraining, anonymization can be performed efficiently on existing splats, in less than 1 second, making the approach practical for real-world deployment.

4.6.2 Neighborhood size (k) and Number of Anonymization Iterations (t). Our method updates each sensitive Gaussian using information aggregated from its neighboring Gaussians. We evaluate

Table 3: Runtime analysis of post-training anonymization. We report the total number of Gaussians in each scene, the number of Gaussians identified for anonymization, Neighbourhood Size (k) and Anonymization Iterations (t), and the runtime of our post-training anonymization stage. Since, we directly anonymize the reconstructed splat without retraining, we achieve this in less than 1 second.

Scene	#G	#SensitiveG	k	t	Time (s)
House (P1)	592K	1531	25	30	0.70
House (P2)	592K	1213	40	30	0.57
Hospital (P4)	538K	1334	40	30	0.62
Hospital (P2)	538K	1127	40	30	0.52
Truck	2052K	2551	100	1	0.24




Figure 6: Impact of neighborhood size (k) and anonymization iterations (t) on FSM. Evaluation on the House dataset shows that higher values of k and t yield lower FSM scores (stronger privacy preservation), with gains eventually saturating.

different neighborhood sizes k and report the resulting Face Similarity Measure (FSM) while simultaneously studying the effect of the number of anonymization iterations t .

As illustrated in the heatmap in Figure 6 for the House dataset, increasing both the neighborhood size k and the optimization duration t consistently yields better quantitative privacy results (lower FSM scores). However, after a certain threshold, these quantitative improvements gradually saturate. Beyond this convergence point, further increases in k and t transition from being strictly necessary for privacy gains to serving as a control for visual aesthetics, allowing users to strike the right balance between local color retention and a more uniform blending.

5 CONCLUSION

We presented a unified framework for privacy preservation in 3D Gaussian Splatting through representation-level anonymization. Unlike existing approaches that rely on modifying input images prior to reconstruction, our method directly edits the learned Gaussian representation, enabling privacy preservation even when the original training images are unavailable. This post-training capability allows existing Gaussian Splats to be anonymized without requiring retraining, while the same formulation naturally extends to a privacy-by-design setting during Gaussian Splat optimization. To facilitate systematic evaluation, we introduced benchmarks for both facial and textual privacy and evaluated our approach using task-specific metrics based on Face Re-Identification similarity and OCR Character Error Rate, complemented by qualitative evaluation of post-training anonymization. Experimental results demonstrate that our method consistently removes sensitive information across novel viewpoints while preserving the visual fidelity of the remaining scene. As 3D scene representations become increasingly prevalent in applications such as robotics, digital twins, autonomous systems, and immersive visualization, incorporating privacy into their design will become increasingly important. We hope this work serves as a step toward practical privacy-aware 3D representations and motivates further research on privacy-preserving methods, benchmarks, and evaluation protocols for emerging 3D scene understanding and reconstruction systems.

REFERENCES

- [1] Sebastien Bonopera, Jerome Esnault, Siddhant Prakash, Simon Rodriguez, Theo Thonat, Mehdi Benadel, Gaurav Chaurasia, Julien Philip, and George Drettakis. 2020. sibr: A System for Image Based Rendering. https://gitlab.inria.fr/sibr/sibr_core
- [2] Nicolas Carion, Laura Gustafson, Yuan-Ting Hu, Shoubhik Debnath, Ronghang Hu, Didac Suris, Chaitanya Ryali, Kalyan Vasudev Alwala, Haitham Khedr, Andrew Huang, et al. 2025. Sam 3: Segment anything with concepts. *arXiv preprint arXiv:2511.16719* (2025).
- [3] Jiazhong Cen, Jiemin Fang, Chen Yang, Lingxi Xie, Xiaopeng Zhang, Wei Shen, and Qi Tian. 2025. Segment any 3d gaussians. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 39. 1971–1979.
- [4] Yiwen Chen, Zilong Chen, Chi Zhang, Feng Wang, Xiaofeng Yang, Yikai Wang, Zhongang Cai, Lei Yang, Huaping Liu, and Guosheng Lin. 2024. Gaussianeditor: Swift and controllable 3d editing with gaussian splatting. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 21476–21485.
- [5] Gioele Ciaparrone, Francisco Luque Sánchez, Siham Tabik, Luigi Troiano, Roberto Tagliaferri, and Francisco Herrera. 2020. Deep learning in video multi-object tracking: A survey. *Neurocomputing* 381 (2020), 61–88.
- [6] Jiankang Deng, Jia Guo, Xue Niannan, and Stefanos Zafeiriou. 2019. ArcFace: Additive Angular Margin Loss for Deep Face Recognition. In *CVPR*.
- [7] European Parliament and Council of the European Union. 2016. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation). , 88 pages. <http://europa.eu>
- [8] Guangchi Fang and Bing Wang. 2024. Mini-splatting: Representing scenes with a constrained number of gaussians. In *European conference on computer vision*. Springer, 165–181.
- [9] Guangchi Fang and Bing Wang. 2024. Mini-splatting2: Building 360 scenes within minutes via aggressive gaussian densification. *arXiv preprint arXiv:2411.12788* (2024).
- [10] Google LLC. 2026. Cloud Vision API (Optical Character Recognition). <https://google.com>. Accessed: July 27, 2026.
- [11] Jia Guo, Jiankang Deng, Alexandros Lattas, and Stefanos Zafeiriou. 2021. Sample and computation redistribution for efficient face detection. *arXiv preprint arXiv:2105.04714* (2021).
- [12] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. 2017. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*. 2961–2969.
- [13] Petr Hurtik, Vojtech Molek, Jan Hula, Marek Vajgl, Pavel Vlasanek, and Tomas Nejezchleba. 2022. Poly-YOLO: higher speed, more precise detection and instance segmentation for YOLOv3. *Neural Computing and Applications* 34, 10 (2022), 8275–8290.
- [14] Osama Jamil and AnnMarie Brennan. 2025. Immersive heritage through gaussian splatting: A new visual aesthetic for reality capture. *Frontiers in Computer Science* 7 (2025), 1515609.
- [15] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 2023. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Transactions on Graphics* 42, 4 (July 2023). <https://repo-sam.inria.fr/fungraph/3d-gaussian-splatting/>
- [16] Shakiba Kheradmand, Daniel Rebain, Gopal Sharma, Weiwei Sun, Yang-Che Tseng, Hossam Isack, Abhishek Kar, Andrea Tagliasacchi, and Kwang Moo Yi. 2024. 3d gaussian splatting as markov chain monte carlo. *Advances in Neural Information Processing Systems* 37 (2024), 80965–80986.
- [17] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. 2023. Segment anything. In *Proceedings of the IEEE/CVF international conference on computer vision*. 4015–4026.
- [18] Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. 2017. Tanks and Temples: Benchmarking Large-Scale Scene Reconstruction. *ACM Transactions on Graphics* 36, 4 (2017).
- [19] Lu Ling, Yichen Sheng, Zhi Tu, Wentian Zhao, Cheng Xin, Kun Wan, Lantao Yu, Qianyu Guo, Zixun Yu, Yawen Lu, et al. 2024. DL3dv-10k: A large-scale scene dataset for deep learning-based 3d vision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 22160–22169.
- [20] Fei Tony Liu, Kai Ming Ting, and Zhi-Hua Zhou. 2008. Isolation forest. In *2008 eighth IEEE international conference on data mining*. Ieee, 413–422.
- [21] Saswat Subhajyoti Mallick, Rahul Goel, Bernhard Kerbl, Markus Steinberger, Francisco Vicente Carrasco, and Fernando De La Torre. 2024. Taming 3dgs: High-quality radiance fields with limited resources. In *SIGGRAPH Asia 2024 Conference Papers*. 1–11.
- [22] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. 2020. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. In *ECCV*.
- [23] Parliament of India. 2023. The Digital Personal Data Protection Act, 2023. <https://www.indiacode.nic.in/bitstream/123456789/22037/1/a2023-22.pdf>
- [24] Abhinav Raundhal, Gaurav Behera, P. J. Narayanan, Ravi Kiran Sarvadevabhatla, and Makarand Tapaswi. 2026. STRinGS: Selective Text Refinement in Gaussian Splatting. In *Proceedings of the Winter Conference on Applications of Computer Vision (WACV)*.
- [25] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2016. Faster R-CNN: Towards real-time object detection with region proposal networks. *IEEE transactions on pattern analysis and machine intelligence* 39, 6 (2016), 1137–1149.
- [26] Renderpeople GmbH. 2025. Renderpeople: Scanned 3D People Models. <https://renderpeople.com/>. Accessed: 2026-07-28.
- [27] restore50. 2026. 3D Models Profile on Sketchfab. <https://sketchfab.com/restore50/models>. Accessed: 2026-07-28.
- [28] Slobodan Ribaric, Aladdin Ariyaeeinia, and Nikola Pavesic. 2016. De-identification for privacy protection in multimedia content: A survey. *Signal Processing: Image Communication* 47 (2016), 131–151.
- [29] Johannes Lutz Schönberger and Jan-Michael Frahm. 2016. Structure-from-Motion Revisited. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [30] Johannes Lutz Schönberger, Enliang Zheng, Marc Pollefeys, and Jan-Michael Frahm. 2016. Pixelwise View Selection for Unstructured Multi-View Stereo. In *European Conference on Computer Vision (ECCV)*.
- [31] Andrew Senior, Sharath Pankanti, Arun Hampapur, Lisa Brown, Ying-Li Tian, Ahmet Ekin, Jonathan Connell, Chiao Fe Shu, and Max Lu. 2005. Enabling video privacy through computer vision. *IEEE Security & Privacy* 3, 3 (2005), 50–57.
- [32] Suguru Shimomura, Kazuki Yamanouchi, and Jun Tanida. 2025. Digital-twin imaging based on descattering Gaussian splatting. *Optics Express* 33, 14 (2025), 29351–29359.
- [33] Chien-Yao Wang and Hong-Yuan Mark Liao. 2024. YOLOv1 to YOLOv10: The fastest and most accurate real-time object detection systems. *arXiv preprint arXiv:2408.09332* (2024).
- [34] Mingqiao Ye, Martin Danelljan, Fisher Yu, and Lei Ke. 2024. Gaussian grouping: Segment and edit anything in 3d scenes. In *European conference on computer vision*. Springer, 162–179.
- [35] Zongxin Ye, Wenyu Li, Sidun Liu, Peng Qiao, and Yong Dou. 2024. Absgs: Recovering fine details in 3d gaussian splatting. In *Proceedings of the 32nd ACM international conference on multimedia*. 1053–1061.
- [36] Zehao Yu, Anpei Chen, Binbin Huang, Torsten Sattler, and Andreas Geiger. 2024. Mip-splatting: Alias-free 3d gaussian splatting. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 19447–19456.
- [37] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 586–595.
- [38] Shijie Zhou, Haoran Chang, Sicheng Jiang, Zhiwen Fan, Zehao Zhu, Dejia Xu, Pradyumna Chari, Suyu You, Zhangyang Wang, and Achuta Kadambi. 2024. Feature 3dgs: Supercharging 3d gaussian splatting to enable distilled feature fields. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 21676–21685.
- [39] Siting Zhu, Guangming Wang, Xin Kong, Dezhi Kong, and Hesheng Wang. 2024. 3d gaussian splatting in robotics: A survey. *arXiv preprint arXiv:2410.12262* (2024).