






GTAvatar: Bridging Gaussian Splatting and Texture Mapping for Relightable and Editable Gaussian Avatars

Kelian Baert  Mae Younes  Francois Bourel  Marc Christie  Adnane Boukhayma 

Univ Rennes, Inria, CNRS, IRISA, France



Figure 1: *GTAvatar* broadens applications of monocular Gaussian Splatting head avatars beyond reenactment and relighting, enabling interactive editing of textures for precise control of intrinsic appearance, while preserving training efficiency, rendering speed and visual fidelity.

Abstract

Recent advancements in Gaussian Splatting have enabled increasingly accurate reconstruction of photorealistic head avatars, opening the door to numerous applications in visual effects, videoconferencing, and virtual reality. This, however, comes with the lack of intuitive editability offered by traditional triangle mesh-based methods. In contrast, we propose a method that combines the accuracy and fidelity of 2D Gaussian Splatting with the intuitiveness of UV texture mapping. By embedding each canonical Gaussian primitive's local frame into a patch in the UV space of a template mesh in a computationally efficient manner, we reconstruct continuous editable material head textures from a single monocular video on a conventional UV domain. Furthermore, we leverage an efficient physically based reflectance model to enable relighting and editing of these intrinsic material maps. Through extensive comparisons with state-of-the-art methods, we demonstrate the accuracy of our reconstructions, the quality of our relighting results, and the ability to provide intuitive controls for modifying an avatar's appearance and geometry via texture mapping without additional optimization.

1. Introduction

Gaussian Splatting based avatars have revolutionized the capture and rendering of digital humans by delivering an unprecedented

level of photorealism with real-time rendering capability, enabling new possibilities for reanimating content from video inputs. However, this realism comes with a trade-off: unlike traditional texture-map-based modeling, Gaussian avatars offer little flexibility for intuitive appearance editing. This gap becomes critical in practice. In film and visual effects, artists routinely adjust the smallest details of a face – smoothing skin to create a flawless appearance, removing distracting high-frequency features, or sculpting age and character through wrinkles, scars and bruises. In gaming and virtual production, creators seek the same level of control to personalize avatars with tattoos, makeup, or stylized patterns using artist-friendly tools that operate directly on material texture maps. Without such editing capabilities, even the most realistic avatar remains a fixed reproduction, rather than a medium for creative expression.

This limitation stems from the fact that each Gaussian splat carries its own color or material properties in isolation, without the shared structure that texture maps naturally provide. As a result, the surface lacks the local coherence needed to treat it as a continuous canvas, making it difficult – if not impossible – to apply edits consistently, whether to a small region of the face or to its overall appearance. Conversely, while mesh-based 3D inverse rendering methods [MHS*22; BZH*23; BBC*24] inherit a coherent default UV-domain fit for artist-friendly edits, they can suffer from topological rigidity when representing high-frequency details and are hampered in their reconstruction capacity due to the limitations of differentiable rasterization when handling complex or translucent geometries (cf. Table 1). Furthermore, naively embedding Gaussians in UV domain leads to discontinuous texture maps that are hard to edit successfully, as witnessed in our results (e.g. Figure 10) and also in [ZWL*25]. The authors of FATE [ZWL*25] proposed a second U-Net neural baking stage to alleviate this issue in the context of non-relightable avatars.

Our goal is to reconstruct a head avatar from a single monocular RGB video, that can be rendered in real time, relit under arbitrary environment maps, animated with new poses and expressions, and directly edited through texture mapping, yet without the added postprocessing complexity of training a two-stage model or baking lighting information in texture directly.

Following seminal work [QKS*24; LKB*25a], we adopt FLAME-anchored [LBB*17] Gaussian splats. Our key observation is that, unlike vanilla 3DGS [KKLD23b], which operates directly in screen space, the 2DGS variant [HYC*24] defines a local splat coordinate frame and computes ray-splat intersections to query kernels. We leverage this property to map ray-splat intersections into the UV domain via an approximate orthographic projection from the splat tangent space to its corresponding FLAME mesh triangle, effectively mapping each splat plane to a continuous UV patch. We devise an efficient method to compute this mapping using a single matrix multiplication per intersection, which is highly important to maintain the method’s efficiency given the very high number of intersections involved in rendering. PBR material attributes are then sampled smoothly from learnable texture maps (albedo, roughness, and specular reflectance) using bilinear filtering. The splat orientations are combined with a residual normal map to produce the final shading normals. We jointly optimize the FLAME parameters, texture maps, 2D Gaussian positional parameters and environ-

ment lighting. This enables differentiable expression of outgoing radiance via the Cook-Torrance BRDF [CT82] under the split-sum approximation [Kar13], using splatted geometry and material G-buffers in a deferred shading framework. Finally, we introduce a novel UV regularization that is crucial for maintaining representation integrity by enforcing alignment of UV coordinates of ray-splat intersections for a given ray. Figure 10 underlines the benefits of our UV mapping and regularization compared to a naive UV embedding strategy based on the projection of splat origins.

The advantages of our novel representation are twofold:

- **Expressive Gaussian primitives.** Spatially varying attributes and normals, in a Phong shading [Pho75] fashion, as opposed to single values, yield higher reconstruction quality with a smaller total model size than state-of-the-art relightable avatar competition [ZLL*25] (Figure 7b). While similar expressiveness can also be achieved with small per-primitive textures (e.g. [RCB*25; SMAD24; CTP*25; YB25b; XCW*24; WB24; SLL*24]), our approach uses a compact global map, offering both efficiency and editability.
- **UV-based semantic control.** Defining primitive attributes continuously on a standard 3DMM template UV map enables a wide range of semantically grounded manipulations and enhancements, including:
 - **Intrinsic material supervision**, such as our albedo regularization with the FLAME albedo model that guides physical decomposition mid-training and reduces artifacts (Figure 10). We note that we also incorporate a diffusion-based albedo prior in screen space, similar to [ZLL*25], to further disambiguate the decomposition.
 - **Model compression** at test time by downsampling the converged textures with controllable loss in rendering quality (Figure 8).
 - **Intuitive editing** of material and normals directly in the FLAME UV domain (e.g. Figures 1, 3, 6 and supplementary video), without requiring any further optimization, while seamlessly integrating with the avatar and maintaining compatibility with the physical lighting representation.

We demonstrate the benefits of our method compared to traditional splat-based reconstruction and rendering techniques, and illustrate its versatility through various texture editing examples. Our method produces highly realistic head reconstructions that not only match or surpass the quality of state-of-the-art monocular relightable avatar HRAvatar [ZLL*25] – despite its lack of explicit texture representations – but also unlocks new possibilities for artistic control.

2. Related work

Avatar Representations. The prior success of NeRFs [MST*20; MESK22] brought impressive levels of photorealism in neural avatars based on differentiable volume rendering [HPX*22; GTZN21; GCL*21]. Rasterization approaches can alleviate the computational rendering burden of implicit neural ray casting. However, mesh-based rasterization [MHS*22] can yield subpar photometric reconstruction performance (e.g. FLARE [BZH*23])

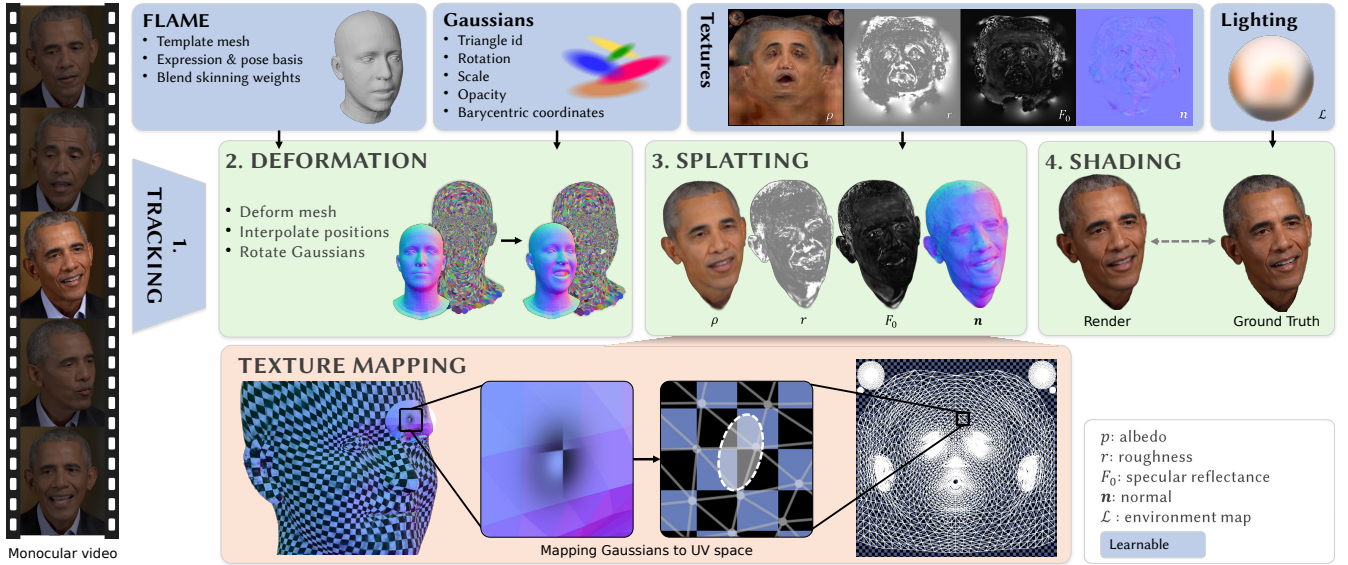


Figure 2: Method overview. GTAavatar takes as input a monocular video sequence and reconstructs it by optimizing parameters of 2D Gaussians, alongside physically-based rendering material and environment lighting. The resulting head avatar can be animated, relit under arbitrary lighting and edited in texture space. The key contribution lies in our efficient texture mapping technique that relates a splat’s tangent space to a patch in canonical UV domain (see Section 3.4).

in Table 1) due to the limitations of meshes in modeling transparent and complex geometry. The recent advent of Gaussian Splatting [KKLD23b] has enabled state-of-the-art reconstruction quality with real-time rendering [XCL*24; YLH*24; QKS*24]. Gaussian Splatting extends the EWA volume resampling framework [ZPVG02; ZPVG01] to learnable [KB15] inverse rendering, modeling 3D scenes with explicit anisotropic Gaussian kernel primitives that can be sorted and rasterized efficiently in tiles. In this respect, the 2DGS variant [HYC*24] leverages planar 2D primitives instead of volumetric ones, and performs precise 2D kernel evaluation in object space as opposed to approximate ones in screen space (3DGS), thus leading to superior geometry and multi-view consistency. The superior normals provided by 2DGS [LKB*25a; SCD*25; JOYB25; DXX*24] have shown to be beneficial in physics-based inverse rendering [YZG*25; YB25b; SGX*25; TLM*25; GWZ*25] where inferring precise reflection directions [VHM*22] is pivotal. 2DGS also defines a local splat coordinate frame that allows us to map ray-splat intersections to our canonical UV space continuously, which is key in our novel avatar representation. Generalizable models for 2DGS [CZC*25; JVB25] have been recently proposed similarly to 3DGS [CLTS24; CXZ*24] before.

Monocular & Relightable Avatars. Learning avatars from casual consumer grade videos paves the road to democratizing volumetric capture, in contrast to the costly requirements of light stages or multi-view setups [Deb12; GLD*19; YZF*23]. Recent advances enable avatar learning from monocular video input (e.g. [ZYW*23; XGGZ24; CWL*24; SWL*24; MWSZ24]), thanks to built-in inductive biases and robust monocular facial

tracking (e.g. [FFBB21; DBB22; RFD*24; ZBT22]) and reconstruction (e.g. [GKR*25; KBM*24; GKG*24; FWS*18; FBTB20]) methods. Rigging avatars with underlying 3D morphable models (3DMMs) [LBB*17; PKA*09; ABWB19; BRZ*18; RBSB18; TL18; OBD*23] enables efficient and robust learning and parametric animation control. They define a canonical space for implicit and explicit representations and factor out expression and pose deformations. Several works leverage data captured under controlled lighting conditions to produce relightable head avatars [SSS*24; YZM*24] with dynamic radiance fields and physical reflectance models. Relightable Gaussian Codec Avatar [SSS*24] has shown good relighting capabilities for full multi-view OLAT captures through its material parameterization, and has also been extended to the phone scan setting [LCS*24]. Despite the increased ill-posedness of the problem, physics-based inverse rendering and relighting from mere monocular videos or sparse images is possible through further regularizations or appearance priors, such as modeling reflectance via simplified bidirectional reflectance distribution functions (BRDF) [BZH*23; BBC*24], or relying on large data corpora [LCS*24]. Closest to our context, HRAAvatar [ZLL*25] employs deferred shading of FLAME-rigged 3D Gaussians augmented with material attributes. In contrast to the latter, we build on 2DGS deferred shading with the aim of extending controllability beyond relighting and reenactment, to additionally include intuitive local editing on conformal material maps, in the interest of more expressive and user-controllable avatar manipulation without sacrificing performance.

UV Mapping & Editing. Recent research has sought to combine explicit surface parameterization with neural representa-




Figure 3: Examples of simple albedo texture editing using our method; adding a star decal, changing hair, teeth or eye colors, removing skin imperfections or adding make-up. Changes remain consistent across poses and expressions.

tions. Within the NeRF [MST*20] framework, works such as NeuTex [XXH*21], Neural Gauge Field [ZLKT23] and Nuvo [SGV*24] learn neural UV mappings that establish bijective correspondences between surfaces and textures, enabling surface-aware rendering across general or category-specific domains. Simultaneously, Gaussian Splatting can support 3D domain editing through semantic grouping and manipulation [YDYK24], dynamic 4D content [YJM*24], text-driven scene modifications [CCZ*24], or brush painting [PHF*25]. Editing via screen-space gradients [SCL*24; WKS*25; GXZ*24], which relies mostly on extensive optimization with pretrained diffusion models [ZYC*24; ZYL*25] such as Instruct-Pix2Pix [BHE23], suffers from high computational costs and lacks precise controllability. UV-domain editing requires continuous texture maps. While some methods structure Gaussians in the UV domain [KGT*24; AYS*24; SWL*24; XGGZ24; WKS*25], most still operate with discontinuous sparse maps. Texture-GS [XHL*24] learns a neural UV mapping for

3DGS primitives for static objects, whilst we aim to use a predefined UV map for semantically grounded edits of a dynamic avatar. Contemporary work FATE [ZWL*25] attempts to alleviate the texture map continuity issue for non-relightable dynamic Gaussian head avatars. However, it requires a separate learning stage involving neural baking via additional networks to obtain smooth editable textures and fails to preserve sharp details when edited, as shown in Figure 12. Our method avoids these multiple stages, bypassing the need for any neural UV mapping or neural baking. Moreover, to our knowledge, our approach is the first Gaussian head avatar to conjointly enable relighting and UV editing of appearance and normals.

2D Image-based relighting. A significant area of research addresses portrait relighting from monocular inputs directly, bypassing the need for full 3D reconstruction. Many approaches decompose the input image into various intrinsic components such as albedo, normals and specular reflectance [WZL*09; SYH*17;

[SKCJ18; LK19; PEL*21], leveraging priors for accurate lighting decomposition. These priors often include parametric face models [BV99; LBB*17; ABWB19; BRZ*18; RBSB18; TL18; OBD*23], light-stage datasets [DHT*00] or simplified physical models [Pho75; CT82]. In recent years, many works have used encoder-decoder architectures [MHP*19; SBT*19; ZHSJ19; WYL*20; NLML20; PEL*21; RFM*24] or diffusion models [PTS23; ZDP*24; CRH*25] that can directly synthesize relit images. While these methods can capture complex illumination effects (e.g. hard shadows, inter-reflections and subsurface scattering), disentanglement of lighting from other factors remains difficult in the under-constrained monocular setting. Additionally, the lack of explicit 3D representations severely limits their ability to render novel viewpoints.

3. Method

Our method reconstructs a relightable and animatable 3D head avatar from a single monocular video by embedding 2D Gaussians onto the surface of a template mesh, where each Gaussian inherits physical material properties from texture maps dynamically. An overview of the pipeline is shown in Figure 2. The approach builds on the FLAME morphable model, which provides a parametric representation of head geometry and facial deformations. We first review the FLAME model and the principles of Gaussian splatting. We then describe how Gaussians are bound to mesh triangles, introduce our novel UV-space mapping that enables differentiable texturing, and present our real-time image-based lighting formulation for physically plausible shading.

3.1. FLAME

The FLAME [LBB*17] 3D morphable model provides a template head mesh with joints for eyes, jaw and neck, and a statistical model for identity and expressions. Identity and expression deformations are expressed as blendshapes while the joints transform vertices through Linear Blend Skinning [MLT88]. For simplicity, we refer to FLAME as template vertex positions $V_t \in \mathbb{R}^{3V}$, triangle topology $[0..V - 1]^{3F}$ and a deformation function \mathcal{F} :

$$V_d = \mathcal{F}(V_t, \Psi)$$

where the identity parameters are fixed and Ψ is the concatenation of pose and expression parameters. We refer the reader to [LBB*17] for exhaustive details on this 3DMM. Additionally, FLAME provides per-corner UV coordinates for texturing (i.e. UV parameterization for the unwrapped template mesh).

3.2. Gaussian Splatting

3D Gaussian Splatting [KKLD23a] (3DGS) uses 3D Gaussians primitives to reconstruct a 3D scene from a set of images. Each Gaussian is defined with a position μ , a rotation matrix R (parameterized as a quaternion), scales $s \in \mathbb{R}^3$, opacity o and view-dependent color c as spherical harmonics coefficients. For rendering, Gaussians are projected to screen space and alpha-blended front-to-back, yielding the final color:

$$C = \sum_i c_i o_i G_i(x) \prod_{j=1}^{i-1} (1 - o_j G_j(x))$$

where $G(x) = \exp(-\frac{1}{2}(x - \mu)^T \Sigma^{-1}(x - \mu))$ is the Gaussian function, $\Sigma = RSS^T R^T$ the covariance matrix and S the diagonal scaling matrix derived from s .

2D Gaussian Splatting [HYC*24] (2DGS) extends this by using flat 2D surfels better suited for reconstructing high-fidelity surfaces. Each gaussian defines a local tangent plane in world space, which is parameterized as:

$$P(s, t) = p + s s + t t$$

where $s \in \mathbb{R}^3$ and $t \in \mathbb{R}^3$ are the scaled tangential vectors of the splat and $p \in \mathbb{R}^3$ its center position. Contrary to 3DGS, 2DGS enables computation of closed-form normal vectors $\mathbf{n} = \frac{s \times t}{\|s \times t\|}$ and precise depth at ray-splat intersections. Please note that we purposefully use (s, t) coordinates to describe the Gaussian tangent planes in place of the more common (u, v) to avoid confusion with (u, v) coordinates used for texture sampling.

3.3. Mesh Binding

Representation. Similarly to existing Gaussian head avatars [QKS*24; LKB*25b; ZLL*25; XGGZ24], we bind Gaussians to triangles of the FLAME mesh. Within each triangle, n splats are initialized with random barycentric coordinates within $[0, 1]$. Each Gaussian has a learnable rotation r relative to the orientation of its triangle (represented as a quaternion), a displacement $d \in \mathbb{R}$ along the triangle normal, scales $s \in \mathbb{R}^2$, and opacity $o \in [0, 1]$.

Deformation. Given pose and expression parameters for FLAME, we first compute the deformed vertex positions $V_d \in \mathbb{R}^{3V}$ and face normals $N_d \in \mathbb{R}^{3F}$. The position of each Gaussian is obtained by interpolating the deformed vertex positions of its triangle with barycentric coordinates. The rotation is the Gaussian's relative rotation r multiplied by the triangle's orientation matrix, whose columns are a tangent, bi-tangent and normal vector of the triangle plane.

3.4. UV-Mapping for Textured Gaussians




Figure 4: From ray-splat intersection to (u, v) texture coordinates. The local 2D splat coordinate (s, t) is orthogonally projected onto the splat's reference triangle, expressed as triangle barycentric coordinates and subsequently mapped to a (u, v) position.

Our method extends Gaussian Splatting by replacing the color

of Gaussians with texture patches. Contrary to existing works that enhance the representation power of Gaussians with small per-primitive textures [RCB*25; SMAD24; CTP*25], we seek to embed Gaussians into a single texture space with a semantic mapping predefined on the template mesh. To achieve this, we map the tangent plane of each splat to a continuous patch in the mesh's UV space. We modify the Gaussian Splatting rasterizer to retrieve colors by sampling a texture at the mapped UV coordinates efficiently with bilinear filtering. This texture is learned alongside the Gaussian parameters.

Given a ray-splat intersection position (s, t) in the tangent space of the splat, an intuitive and practical solution for the mapping to (u, v) coordinates is to use the nearest mesh point. For ease of computation, we approximate this point as the nearest point on the primitive's bound triangle plane, *i.e.* the orthogonal projection of the ray-splat intersection point on its triangle plane as illustrated in Figure 4. The final UV coordinates are thus obtained by linearly interpolating the UV coordinates of the triangle vertices at this projection location. Irregularities resulting from the projection falling outside the triangle are mitigated by the splat scales remaining generally smaller than triangles, and by our UV regularization loss introduced in Section 4.3.

Minimizing the number of operations per ray-splat intersection is key to maintaining high rendering speed. To this end, we devise a method that only requires a single matrix multiplication per intersection to compute the final UV coordinates, as naive alternatives do not scale well with the large number of ray-splat intersections involved in rendering. For a splat attached to a mesh triangle with UV coordinates $UV_X \in [0, 1]^2$ and FLAME-deformed vertices $V_X \in \mathbb{R}^3$ ($X \in \{A, B, C\}$), we adopt triangle barycentric coordinates with origin at V_A and define two functions that map from barycentric coordinates (b, c) to world space and UV respectively:

$$v(b, c) = V_A + J_v \begin{bmatrix} b \\ c \end{bmatrix} \quad uv(b, c) = UV_A + J_{uv} \begin{bmatrix} b \\ c \end{bmatrix}$$

where J_v and J_{uv} denote the Jacobians of the two transformations:

$$J_v = [V_B - V_A \mid V_C - V_A] \in \mathbb{R}^{3 \times 2}$$

$$J_{uv} = [UV_B - UV_A \mid UV_C - UV_A] \in \mathbb{R}^{2 \times 2}$$

Similarly, let f denote the transformation from the splat's tangential space to world space:

$$f(s, t) = p + J_{st} \begin{bmatrix} s \\ t \end{bmatrix}, \quad J_{st} = [s \mid t] \in \mathbb{R}^{3 \times 2}$$

Assuming a non-degenerate triangle, v is invertible and its inverse is defined on the triangle plane. Thus, if the Gaussian is aligned with the triangle plane, we can define the mapping $(u, v) = \mathcal{M}(s, t) = (uv \circ v^{-1} \circ f)(s, t)$. We exploit the linearity of uv, v, f and \mathcal{M} to rewrite it as follows:

$$uv_0 = \mathcal{M}(0, 0) \quad J_{st \rightarrow uv} = \frac{\partial \mathcal{M}}{\partial st} = J_{uv} J_v^\dagger J_{st} \quad (1)$$

$$\mathcal{M}(s, t) = uv_0 + J_{st \rightarrow uv} \begin{bmatrix} s \\ t \end{bmatrix} \quad (2)$$

where J_v^\dagger is the pseudo-inverse of J_v and uv_0 the UV coordinates

at the center of the splat, computed using the Gaussian's barycentric coordinates directly. We extend this to Gaussians that are not aligned with their triangle by noting that $J_v J_v^\dagger$ is the orthogonal projector onto the column space of J_v (the triangle plane) and $x \mapsto J_v^\dagger(x - V_A)$ yields the barycentric coordinates of the projection. Thus, \mathcal{M} maps (s, t) to the UV coordinates of this projection (the $J_v^\dagger V_A$ term is accounted for in uv_0). The pseudo-inversion is implemented by solving a least-squares problem, which does not back-propagate gradients to the template vertices.

This lightweight mapping only requires a single matrix multiplication at each ray-splat intersection during rendering. The largest computational costs are the pseudo-inversion of J_v^\dagger for each triangle and the bilinear texture filtering; the latter is a well-understood cost of graphics pipelines that can be mitigated with hardware acceleration. Notice that J_v – and its pseudo-inverse J_v^\dagger – are mesh and expression-dependent but do not need to be updated for view-point changes. See Section 5.5 for more details on rendering speed and a comparison with a naive projection baseline.

3.5. Appearance Modeling

The final per-pixel color is derived using deferred shading on the G-buffers obtained through Gaussian Splatting. For a given pixel, let $\mathbf{n} \in \mathbb{R}^3$ denote the splatted normal vector (normalized) and $\omega_0 \in \mathbb{R}^3$ the view direction. We replace typical per-Gaussian spherical harmonics with a 5-channel material texture: albedo $\rho \in [0, 1]^3$, roughness $r \in [0, 1]$ and specular reflectance $f_0 \in [0, 1]$. After rasterization, we retrieve those values per-pixel in the G-buffers.

3.5.1. Real-time Physically-Based Rendering

We adopt the real-time shading model of [Kar13], which we briefly summarize in this section. The rendering equation expresses the outgoing light L_0 leaving surface point x in the camera direction ω_0 by integrating the light reaching it from all directions:

$$L_0(x, \omega_0) = \int_{\Omega} f_r(x, \omega_i, \omega_0) L_i(x, \omega_i) (\omega_i \cdot \mathbf{n}) d\omega_i$$

where f_r is the bidirectional reflectance distribution function of the material (BRDF). We decompose the BRDF as a diffuse Lambertian term independent of incoming light direction:

$$L_{0, \text{diffuse}} = \rho \int_{\Omega} L_i(x, \omega_i) (\omega_i \cdot \mathbf{n}) d\omega_i \quad (3)$$

and a view-dependent specular term, calculated using the split-sum approximation [Kar13]:

$$L_{0, \text{spec}} = \left(\int_{\Omega} f_{r, \text{spec}}(x, \omega_i, \omega_0) d\omega_i \right) \left(\int_{\Omega} L_i(x, \omega_i) (\omega_i \cdot \mathbf{n}) d\omega_i \right) \quad (4)$$

Similarly to previous relightable head avatars [MHS*22; BZH*23; ZLL*25], we adopt the Cook-Torrance microfacet reflectance model [CT82], which parameterizes the specular BRDF with surface roughness r and specular reflectance at normal incidence f_0 . In practice, the second integral incorporates parts of $f_{r, \text{spec}}$ that are not dependent on ω_0 , and is stored as a pre-filtered environment map $\mathcal{L}(\omega_r, r)$ with mip levels sampled by spatially-varying roughness r , where ω_r is the reflection of ω_0 onto the normal. We learn this function as a RGB cubemap that represents the lighting of the scene. The first integral is only dependent on material properties

and viewing angle and independent of lighting. It is pre-computed in a look-up table; for conciseness, we refer the reader to [Kar13] for details and denote this reflectance term as $R(f_0, r, \omega_0 \cdot \mathbf{n})$. Note that due to properties of the BRDF omitted here, the integral in Eq. 3 can be computed as $\mathcal{L}(\mathbf{n}, 0)$. Hence the final color is:

$$\begin{aligned} L_0(x, \omega_0) &= L_{0,\text{diffuse}} + L_{0,\text{spec}} \\ &= \rho \cdot \mathcal{L}(\mathbf{n}, 0) + R(f_0, r, \omega_0 \cdot \mathbf{n}) \cdot \mathcal{L}(\omega_{\mathbf{r}}, r) \end{aligned} \quad (5)$$

3.5.2. Normal Mapping

The use of Texture Mapping for appearance properties increases the representation power of each Gaussian. However, the shading model described in Section 3.5.1 requires precise surface normals to achieve photo-realistic results. Relying solely on increasing the number of Gaussians to model this high-frequency geometry is inefficient and can lead to a prohibitively large memory footprint. Instead, inspired by classical normal mapping in real-time graphics, we augment our material texture with two additional channels that define a local perturbation relative to the splat's own geometric normal. This allows a single, relatively large Gaussian to represent a surface patch with intricate geometric detail.

From 2D Texture to 3D Tangent-Space Normal. The two normal map channels, (n_x, n_y) , parameterize a 3D unit normal vector, \mathbf{n}_t , within the local tangent space of the splat. The splat's local coordinate system is defined by its orthonormal basis: the two principal tangential vectors (\mathbf{s}, \mathbf{t}) and its geometric normal $\mathbf{n}_g = \frac{\mathbf{s} \times \mathbf{t}}{\|\mathbf{s} \times \mathbf{t}\|}$. To reconstruct the full 3D vector from the 2-channel texture, the third component n_z is derived assuming the normal points away from the surface ($z > 0$):

$$\mathbf{n}_t = \begin{bmatrix} n_x & n_y & \sqrt{1 - n_x^2 - n_y^2} \end{bmatrix}^T \quad (6)$$

The resulting unit vector \mathbf{n}_t represents the high-frequency surface normal in the splat's local space. A value of $(0, 0)$ in the texture corresponds to the tangent-space normal $[0, 0, 1]^T$, representing no perturbation from the splat's geometric normal \mathbf{n}_g .

Transformation to World Space. During rendering, we sample the tangent-space normal \mathbf{n}_t from our texture. This local normal is then transformed into world space by multiplying it with the splat's rotation matrix R :

$$\mathbf{n}_w = R\mathbf{n}_t, \quad \text{where } R = \begin{bmatrix} \mathbf{s} & \mathbf{t} & \mathbf{n}_g \end{bmatrix} \quad (7)$$

This final perturbed world-space normal, \mathbf{n}_w , is then used for all shading calculations as described in Section 3.5.1. This strategy enables higher-quality shading and more accurate reconstruction of fine details with a significantly reduced number of primitives. Figure 9 shows examples of the detailed normal maps recovered by our method and demonstrates the resulting improvement in rendering quality compared to using only the base splat normals.

4. Training

Our model is trained end-to-end from a monocular RGB video. Following previous work [ZYW*23; BZH*23; ZLL*25], the input video is preprocessed prior to training to extract alpha masks using a matting method [LYSS21] and initial FLAME parameters using an off-the-shelf facial tracker [RFD*24]. During training,

we simultaneously optimize the parameters of our 2D Gaussians (barycentric coordinates, relative rotation, scale, displacement and opacity), the textures (albedo, roughness, specular reflectance and normals) and the scene's lighting environment map (a $6 \times 32^2 \times 3$ RGB cubemap). Gaussians undergo densification and pruning as in other Gaussian-based methods [KKLD23b; HYC*24; QKS*24; ZLL*25]. We also optimize the FLAME model (template mesh and deformation basis) for higher fidelity and fine-tune the expression encoder from the FLAME tracker [RFD*24] while continuously predicting updated expression parameters during training, following recent work [ZLL*25]. We train for 15 epochs using Adam optimizer [KB15], requiring one to two hours on an NVIDIA RTX A5000 GPU. Learning rates are given in Table 3a.

Our total loss function is a weighted sum of terms designed to supervise the image reconstruction, guide the physical property disentanglement, regularize our novel UV mapping, and ensure geometric stability. The overall objective is:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{photo}} + \mathcal{L}_{\text{PBR}} + \mathcal{L}_{\text{uv}} + \mathcal{L}_{\text{geom}} \quad (8)$$

Values used for the weights referenced in this section are provided in Table 3b.

4.1. Image Reconstruction Losses

The primary supervisory signal comes from comparing the rendered image with the ground truth. Let I and M be the ground-truth image and foreground mask for a given frame. Let \hat{I} and \hat{M} be our corresponding rendered image and alpha mask. The photometric loss $\mathcal{L}_{\text{photo}}$ is a combination of an L1 loss, a structural similarity (SSIM) loss and a mask loss.

$$\begin{aligned} \mathcal{L}_{\text{photo}} &= \lambda_{\text{L1}} \|I - \hat{I}\|_1 + \lambda_{\text{SSIM}} (1 - \text{SSIM}(I, \hat{I})) \\ &\quad + \lambda_{\text{mask}} \|M - \hat{M}\|_1 \end{aligned} \quad (9)$$

4.2. Priors for Physical Disentanglement

Decomposing appearance into intrinsic physical properties from a monocular video is a highly ill-posed problem. To guide this decomposition, we introduce a set of priors, \mathcal{L}_{PBR} .

First, we supervise the rendered albedo \hat{I}_ρ with a diffusion-based prior I_ρ [XSD*24] precomputed for every third frame of the video, as recent work has proven this effective for disentangling albedo and lighting [ZLL*25].

$$\mathcal{L}_{\text{diff_albedo}} = M \|I_\rho - \hat{I}_\rho\|_1$$

Second, we enforce smoothness on the learned material maps, which is a natural prior for skin appearance. We apply a total variation (TV) loss to the roughness (\mathcal{T}_r) and specular reflectance (\mathcal{T}_{f_0}) textures:

$$\mathcal{L}_{\text{smooth}} = \text{TV}(\mathcal{T}_r) + \text{TV}(\mathcal{T}_{f_0})$$

Finally, we regularize the detailed normal map (\mathcal{T}_n) toward the default up-vector $[0, 0, 1]^T$, encouraging it to capture only necessary deviations from the base splat normal:

$$\mathcal{L}_{\text{normal_reg}} = \|\mathcal{T}_n - [0, 0, 1]^T\|_1$$




Figure 5: Reconstruction and relighting examples obtained with our method. From left to right: original frame, reconstruction, rendered normals and albedo, relighting under various environment maps.

The total PBR loss writes:

$$\mathcal{L}_{\text{PBR}} = \lambda_{\text{diff_albedo}} \mathcal{L}_{\text{diff_albedo}} + \lambda_{\text{smooth}} \mathcal{L}_{\text{smooth}} + \lambda_{\text{normal_reg}} \mathcal{L}_{\text{normal_reg}} \quad (10)$$

4.3. UV-Mapping Regularization

A key goal of our method is to produce continuous, artifact-free texture maps suitable for editing. Since multiple semi-transparent Gaussian splats can contribute to a single pixel, it is crucial that they all map to a consistent location in the UV texture for a given ray. To enforce this, we introduce \mathcal{L}_{uv} , a set of novel regularization terms.

UV Distortion Loss. Inspired by the depth distortion loss in

2DGS [HYC*24], we propose a UV distortion loss to concentrate the texture coordinates of all ray-splat intersections contributing to a single ray. For a given pixel ray, we minimize the weighted pairwise distance between the UV coordinates of all ray-splat intersections. Let uv_i be the (u, v) coordinate derived from the i -th splat intersection along a ray. The loss then writes:

$$\mathcal{L}_{\text{uv_dist}} = \sum_{\text{rays}} \sum_{i,j} \omega_i \omega_j \|uv_i - uv_j\|_2$$

where ω_i is the volumetric blending weight of the i -th splat intersection: $\omega_i = \alpha_i \mathcal{G}_i \prod_{j=1}^{i-1} (1 - \alpha_j \mathcal{G}_j)$, where α_i and \mathcal{G}_i are the opacity and Gaussian value at the intersection respectively. This loss penalizes rays where multiple highly-weighted splats contribute to the same pixel but map to distant UV coordinates. Minimizing this

Method	Relighting	Texture editing	INSTA dataset			HDTF dataset		
			PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
[ZBT23] INSTA	\times	\times	27.85	0.9110	0.1047	25.03	0.8475	0.1614
[ZYW*23] Point-avatar	\times	\times	26.84	0.8970	0.0926	25.14	0.8385	0.1278
[SWL*24] Splatting-avatar	\times	\times	28.71	0.9271	0.0862	26.66	0.8611	0.1351
[XGGZ24] Flash-avatar	\times	\times	29.13	0.9255	0.0719	27.58	0.8664	0.1095
[MWSZ24] GBS	\times	\times	29.64	0.9394	0.0823	27.81	0.8915	0.1297
[BZH*23] FLARE	\checkmark	\times	26.80	0.9063	0.0816	25.55	0.8479	0.1183
[ZLL*25] HRAvatar	\checkmark	\times	30.36	0.9482	0.0569	28.55	0.9089	0.0825
GTAavatar (ours)	\checkmark	\checkmark	30.52	0.9537	0.0552	28.83	0.9130	0.0794

Table 1: Results of various methods for the self-reenactment task on the INSTA and HDTF datasets. Our method outperforms all others in PSNR, SSIM and LPIPS.




Figure 6: Texture editing with off-the-shelf PBR material maps. Our method enables consistent rendering of edited avatars under varying illumination with conventional material definitions. The last column underlines the importance of normal mapping for creating edits that interact with lighting convincingly.

term forces the UV projections of overlapping Gaussians to converge to a single, sharp point in the texture map, preventing blurry or "ghosting" artifacts and creating an editable texture.

UV Mask Loss. To prevent Gaussians from sampling invalid areas of the texture atlas (i.e. the gaps between UV islands), we introduce a boundary loss. We pre-compute a binary mask M_{uv} that is 0 inside valid triangle regions of the UV map and 1 elsewhere. The loss regularizes the intensity of texels outside of valid areas:

$$\mathcal{L}_{uv_boundary} = M_{uv} \|\mathcal{T} - \mathcal{T}_{init}\|_1$$

where \mathcal{T}_{init} is the initial value for that texture – $[0, 0, 0]$ for albedo, 0.5 for roughness, 0.05 for specular reflectance and $[0, 0, 1]$ for normals.

Statistical Albedo The albedo texture (\mathcal{T}_ρ) is further regularized using the texture-space PCA albedo prior provided by FLAME. PCA coefficients ω_ρ are optimized during training, yielding pseudo ground-truth texture $\tilde{\mathcal{T}}_\rho = \tilde{\mathcal{T}}_{\rho, \text{mean}} + \omega_\rho \tilde{\mathcal{T}}_{\rho, \text{basis}}$, used in:

$$\mathcal{L}_{stat_albedo} = \|\mathcal{T}_\rho - \tilde{\mathcal{T}}_\rho\|_1$$

Together, $\mathcal{L}_{uv} = \lambda_{uv_dist} \mathcal{L}_{uv_dist} + \lambda_{boundary} \mathcal{L}_{uv_boundary} +$

$\lambda_{stat_albedo} \mathcal{L}_{stat_albedo}$. These terms are critical for producing high-quality editable textures as shown in Figure 10.

4.4. Geometric and Gaussian Regularization


Finally, to maintain a stable and plausible underlying structure, we use a set of geometric regularizers, \mathcal{L}_{geom} .

Surface Normal Consistency. To ensure that our 2D splats align locally with the overall reconstructed surface, we adopt the normal consistency loss from 2DGS [HYC*24]. This loss aligns the normal of each individual splat \mathbf{n}_i with the macro-scale surface normal \mathbf{N} , derived from the screen-space gradients of the rendered depth map. The loss for each ray is expressed as:


$$\mathcal{L}_{normal_consist} = \sum_i \omega_i (1 - \mathbf{n}_i^\top \mathbf{N})$$

This encourages the formation of a smooth, coherent geometric surface, which is an important foundation for a stable UV mapping.

FLAME mesh regularization. Following prior work on drivable avatars [BZH*23], we maintain the structure of the underlying mesh through a Laplacian smoothing loss that regularizes the offsets between the original and the fine-tuned FLAME mesh:



(a) LPIPS for varying number of Gaussians and texture resolutions.



(b) LPIPS for varying number of learned parameters, accounting for Gaussians and texture resolution. For our method, we vary both simultaneously to find a trade-off between quality and model size. As indicated by the dotted line, our method achieves the same quality with significantly lower size.

Figure 7: Reconstruction quality on unseen frames averaged over 10 videos of the INSTA dataset, compared with HRAvatar [ZLL*25] (LPIPS, lower is better). Densification and pruning of Gaussians are disabled in these experiments for manual control of splat count.




Figure 8: Normals and render for different levels of texture scaling after training with a 1024×1024 texture. As the texture is down-scaled, the reconstruction quality decreases gradually.

$\mathcal{L}_{\text{lap}} = \|L(\mathcal{F}(V_i, \Psi) - \tilde{\mathcal{F}}(\tilde{V}_i, \tilde{\Psi}))\|_2^2$ where $\tilde{\mathcal{F}}$ denotes FLAME with optimized basis, \tilde{V}_i the tuned template vertices, $\tilde{\Psi}$ the FLAME parameters with expression computed using the tuned encoder and L the graph Laplacian of the mesh [Sor05]. We additionally regularize the difference in all attributes of FLAME with a L2 loss $\mathcal{L}_{\text{FLAME}}$. Finally, we bias the predicted expression parameters towards the initial values to prevent excessive divergence from the initial tracker: $\mathcal{L}_{\text{expr}} = \|\tilde{\Psi}_{\text{expression}} - \Psi_{\text{expression}}\|_2^2$. Together, these losses ensure our underlying mesh maintains geometric integrity and the surface UV mapping is preserved.

Lastly, we encourage Gaussian primitives to remain near the center of their parent triangle by regularizing their barycentric coordinates \mathbf{b} : $\mathcal{L}_{\text{bary}} = \|\mathbf{b} - [\frac{1}{3}, \frac{1}{3}, \frac{1}{3}]^T\|_2^2$.

In summary, the total geometric loss is $\mathcal{L}_{\text{geom}} = \lambda_{\text{normal}}\mathcal{L}_{\text{normal_consist}} + \lambda_{\text{lap}}\mathcal{L}_{\text{lap}} + \lambda_{\text{FLAME}}\mathcal{L}_{\text{FLAME}} + \lambda_{\text{expr}}\mathcal{L}_{\text{expr}} + \lambda_{\text{bary}}\mathcal{L}_{\text{bary}}$.



Figure 9: Ablation test on the normal map. 3D reconstruction is performed without relying on the normal map, i.e. normals are only computed from the Gaussian splats and fine surface details are baked in the other channels (top row), compared to reconstruction with a normal map (bottom row).

5. Results

Dataset. We evaluate various aspects of our method on two common datasets for monocular avatar reconstruction: INSTA [ZBT23] and HDTF [ZLDF21]. Both provide 2-3 minute 512×512 resolution talking head videos. Our experimental setting aligns with that of our main baseline, HRAvatar [ZLL*25]: we use 10 subjects from the INSTA dataset, and 8 from HDTF. For self-reenactment evaluation, the last 350 frames are left out of training for the former, and the last 500 frames for the latter.

5.1. Self-reenactment

In this section, we validate the quality of our reconstructions in the self-reenactment setting. Unseen frames from the video are tracked to recover FLAME parameters using the tuned encoder, then reconstructed. In Table 1 we report average PSNR, SSIM and VGG-based LPIPS metrics for the INSTA and HDTF datasets, compared to state-of-the-art monocular avatar reconstruction methods as reported by the authors of HRAvatar [ZLL*25]. Our method displays state of the art performance, comparably with HRAvatar, outperforming all baselines. Qualitative reconstruction examples are shown in Figures 5 and 13. Figure 11 showcases our method's ability to render realistic avatars from novel viewpoints.

Moreover, we assess the effect of primitive count and texture resolution in Figure 7. Average LPIPS on the INSTA dataset is reported for various primitive counts and texture resolutions. For fair comparison with HRAvatar which does not have textures, we also compare with equal parameter counts. These results underline the efficiency of our method, enabling on-par reconstruction




Figure 10: Ablation of key aspects of our texturing approach. While the reconstruction quality remains similar, the UV distortion regularization yields a sharper UV mapping and helps maintain high-frequency details from the texture. The statistical albedo smooths the texture and reduces artifacts. In the last row, we set $J_{st \rightarrow uv}$ to zero (see Section 3.4), resulting in a discontinuous texture that only uses a sparse set of texels. Please zoom in for details.

quality with a significantly smaller model size. The added representation power of textures reduces the number of primitives required through a compact representation of normal and appearance attributes enabled by the UV-space locality of adjacent splats.

In Figure 8, we render an avatar with textures down-scaled to lower resolutions at test time without any additional tuning. As shown, our texture-based modeling enables intuitive control of model size, an important capability for bandwidth-adaptive applications. Note that training with lower resolution textures directly would yield better results than down-scaling the textures post-facto, as the optimization process can adapt to the reduced representation power.

5.2. Texture Editing

In this section, we show that our texture mapping approach enables several editing use cases that would be very difficult if not impossible to implement with standard Gaussian Splatting rasterization.

Figure 3 demonstrates overlay of sharp decals, precise editing of local features and color-shifting face regions. Figure 6 shows cases more substantial edits using off-the-shelf PBR materials, ren-




Figure 11: Examples of reconstructions and novel viewpoints.




Figure 12: Visual comparison of our results with FATE [ZWL*25]. While both approaches provide texture editing features, our UV mapping technique yields sharper renders with high-frequency details. Numbers indicate the respective number of splats.

dered under varying illumination to validate the accuracy of our physically-based rendering with conventional material definitions. Our results show that the edited avatars maintain sharpness and consistency across variations in pose, expression and illumination.

In Figure 10, we perform an ablation study of several aspects of our method aimed at improving the texture mapping and the quality of reconstructed textures. First, we validate that the UV distortion loss enhances the sharpness of our render. While the reconstruction converges well without the regularization, mapping a texture with high-frequency details reveals some blurriness. Next, we show that the statistical albedo regularization is able to fill-in the holes in the textures, remove distracting artifacts and better preserve the structure of the texture space. Finally, we set our Jacobians (see Section 3.4) to zero, which equates to using a single UV coordinate per splat. This per-primitive mapping yields a discontinuous texture with only a sparse set of texels being used (note that we disable the statistical albedo regularization to better visualize those texels), and cannot be edited at high resolution unless the number of primitives is impractically high.

FATE [ZWL*25] allows texture-based editing of Gaussian Avatars. However, their method relies on per-primitive color, which severely limits the quality of rendering for high-frequency textures. Figure 12 compares renders with the same texture for FATE and



Figure 13: Visual comparison of our method with HRAvatar [ZLL*25] and FLARE [BZH*23] for reconstruction and environment map relighting.

GTAAvatar. Our method outperforms FATE even while increasing the number of primitives FATE uses well beyond its default.

5.3. Normal mapping

In Figure 9, we demonstrate the impact of our normal mapping on the sharpness of rasterized normals and the quality of relit renders. Furthermore, Figure 6 shows how handcrafted normal maps substantially improve the realism for material editing under dynamic lighting.

5.4. Relighting

Figure 13 presents relighting examples compared with HRAAvatar and FLARE. Our method achieves comparable or superior quality relative to both baselines, which employ the same physically-based rendering formulation. We provide further examples in Figure 5 to illustrate the relighting capabilities of our method.

5.5. Rendering time

Sampling a texture at every ray-splat intersection comes with a computational cost. To minimize this cost, our implementation uses hardware acceleration at test time for more efficient access via CUDA texture objects. In Table 2, we compare rendering speed with other methods, and with variants of our method without hardware acceleration, with more Gaussians and with naive projection instead of the fast UV mapping described in Section 3.4. In the latter, we explicitly perform orthogonal projection to the triangle plane and calculate barycentric coordinates at ray-splat intersection. We report FPS for both static and dynamic geometry, as we found FLAME deformations and, in our case, updating the Jacobians, to be a significant contributor to rendering time. Note that the static geometry case still enables camera movement and relighting. While our method performs slower than the HRAAvatar baseline for a given number of primitives, ours requires less (as shown in Figure 7a), yielding on par rendering time for comparable or better reconstruction quality. The final trained avatar can be rendered at more than **170 FPS** with static geometry, or **80 FPS** with dynamic geometry on a RTX A5000.

6. Ethics

This research aims at pushing forward the precision and authenticity of 3D facial reconstruction for legitimate applications such as visual effects or virtual interactions. We do not condone the use of our work for producing unconsented deepfakes or deceptive content of any kind. Our focus remains on contributing to scientific progress and industry innovation in alignment with ethical standards. We also encourage ongoing dialogue and the development of regulations to safeguard individual rights as this technology evolves.

7. Conclusion

We presented a novel UV-domain Gaussian splatting framework that combines the fidelity and efficiency of EWA volume resampling enabled physics-based inverse rendering with the intuitive-

Method	FPS (static)	FPS (dynamic)	LPIPS ↓
FLARE [BZH*23]	31	27	0.082
HRAAvatar 80k [ZLL*25]	136	88	0.062
HRAAvatar 10k [ZLL*25]	165	105	0.071
Ours 10k (no hw. acceleration)	121	64	0.060
Ours 10k (naive projection)	152	67	0.060
Ours 80k (more Gaussians)	109	64	0.058
Ours 10k	175	83	0.060

Table 2: Comparison of rendering speed at inference for static and dynamic geometry on a NVIDIA RTX A5000 GPU, and average LPIPS for unseen frames of the INSTA dataset. Numbers next to method names indicate the number of Gaussians used, when applicable (*densification and pruning are disabled for manual control*). Despite the additional overhead of texture sampling, our method achieves competitive speed by reducing the number of Gaussians required to achieve a photorealistic reconstruction.

Parameter	LR	Weight	Value
Barycentric coords.	1e-3	λ_{L1}	0.80
Rotation	1e-3	λ_{SSIM}	0.20
Scale	5e-3	λ_{mask}	0.10
Displacement	2e-5	λ_{diff_albedo}	0.25
Opacity	5e-2	λ_{stat_albedo}	0.0001
Expression encoder	5e-5	λ_{expr}	0.01
Material texture	5e-3	λ_{smooth}	0.01
Normal texture	1e-3	λ_{normal_reg}	0.01
Environment map	2e-2	$\lambda_{normal_consist}$	0.05
FLAME template vertices	1e-5	λ_{uv_dist}	50
FLAME LBS weights	1e-4	$\lambda_{boundary}$	1
FLAME expr. and pose shapes	1e-6	λ_{bary}	0.1
FLAME statistical albedo	5e-2	λ_{lap}	200
		λ_{FLAME}	0.001

(a) Learning rates.

(b) Objective function weights.

Table 3: Hyperparameters used for training our method.

ness of texture-based editing, enabling photorealistic, easily editable, and relightable head avatars from monocular videos. Our method achieves state-of-the-art reconstruction and relighting quality while introducing efficient, semantically meaningful solutions for material and geometry control.

Several limitations remain. The use of image-based lighting to represent scene illumination enables intuitive manipulation of the environment, but comes with shortcomings. Most importantly, this approach can only model light at infinity and is unable to represent point lights. Furthermore, effects such as indirect illumination and shadows are not represented, and will thus be baked into other properties, yielding inaccurate physical decomposition.

As part of **future work**, several avenues remain to further improve our representation. The texture-mapping contribu-

tion could be integrated in a Gaussian rasterizer that enables more advanced lighting representations, e.g. through ray-tracing [MMP*24; XCX*24; GWZ*25; SGX*25]. Reducing aliasing artifacts is also a key challenge: incorporating anti-aliasing strategies, either through trilinear texture filtering in the UV domain or directly within the 2DGS splatting stage [YB25a] could enhance visual smoothness and stability of avatars and edits. In addition, extending this representation to the better-constrained multi-view setting could enable further uses. Such extensions would make our method even more suitable for production-grade applications.

References

- [ABWB19] ABREVAYA, VICTORIA FERNÁNDEZ, BOUKHAYMA, ADNANE, WUHRER, STEFANIE, and BOYER, EDMOND. “A Decoupled 3D Facial Shape Model by Adversarial Training”. *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. Oct. 2019, 9419–9428 **3**, **5**.
- [AYS*24] ABDAL, RAMEEN, YIFAN, WANG, SHI, ZIFAN, et al. “Gaussian Shell Maps for Efficient 3D Human Generation”. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2024, 1–11. DOI: [10.1109/CVPR52688.2024.00001](https://doi.org/10.1109/CVPR52688.2024.00001). URL: <https://arxiv.org/abs/2311.17857> **4**.
- [BBC*24] BAERT, KELIAN, BHARADWAJ, SHRISHA, CASTAN, FABIEN, et al. “SPARK: Self-supervised Personalized Real-time Monocular Face Capture”. *ACM SIGGRAPH Asia 2024 Conference Papers*. ACM, 2024. DOI: [10.1145/3450626.3459936](https://doi.org/10.1145/3450626.3459936) **2**, **3**.
- [BHE23] BROOKS, TIM, HOLYNSKI, ALEKSANDER, and EFROS, ALEXEI A. “InstructPix2Pix: Learning To Follow Image Editing Instructions”. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2023, 18392–18402. DOI: [10.1109/CVPR46556.2023.01307](https://doi.org/10.1109/CVPR46556.2023.01307) **4**.
- [BRZ*18] BOOTH, JAMES, ROUSSOS, ANASTASIOS, ZAFEIRIOU, STEFANOS, et al. “Large Scale 3D Morphable Models”. *International Journal of Computer Vision* 126 (2018), 233–254 **3**, **5**.
- [BV99] BLANZ, VOLKER and VETTER, THOMAS. “A Morphable Model for the Synthesis of 3D Faces”. *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)*. 1999 **5**.
- [BZH*23] BHARADWAJ, SHRISHA, ZHENG, YUFENG, HILLIGES, OTMAR, et al. “FLARE: Fast learning of Animatable and Relightable Mesh Avatars”. *ACM Transactions on Graphics* 42 (Dec. 2023), 15. DOI: <https://doi.org/10.1145/3618401> **2**, **3**, **6**, **7**, **9**, **12**, **13**.
- [CCZ*24] CHEN, YIWEN, CHEN, ZILONG, ZHANG, CHI, et al. “GaussianEditor: Swift and controllable 3d editing with gaussian splatting”. *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 2024, 21476–21485 **4**.
- [CLTS24] CHARATAN, D., LI, S., TAGLIASACCHI, A., and SITZMANN, V. “PixelSplat: 3D Gaussian Splatting from Image Pairs for Scalable Generalizable 3D Reconstruction”. *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2024 **3**.
- [CRH*25] CHATURVEDI, SUMIT, REN, MENGWEI, HOLD-GEOFFROY, YANNICK, et al. “SynthLight: Portrait Relighting with Diffusion Model by Learning to Re-render Synthetic Faces”. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2025)* **5**.
- [CT82] COOK, ROBERT L and TORRANCE, KENNETH E. “A Reflectance Model for Computer Graphics”. *ACM Transactions on Graphics (TOG)* 1.1 (1982), 7–24 **2**, **5**, **6**.
- [CTP*25] CHAO, BRIAN, TSENG, HUNG-YU, PORZI, LORENZO, et al. *Textured Gaussians for Enhanced 3D Scene Appearance Modeling*. 2025. arXiv: [2411.18625](https://arxiv.org/abs/2411.18625) [cs.CV]. URL: <https://arxiv.org/abs/2411.18625> **2**, **6**.
- [CWL*24] CHEN, YUFAN, WANG, LIZHEN, LI, QIJING, et al. “Mono-GaussianAvatar: Monocular Gaussian Point-based Head Avatar”. *ACM SIGGRAPH 2024 Conference Papers*. 2024, 1–9 **3**.
- [CXZ*24] CHEN, YUEDONG, XU, HAOFEI, ZHENG, CHUANXIA, et al. “MVSplat: Efficient 3D Gaussian Splatting from Sparse Multi-View Images”. *Lecture Notes in Computer Science (ECCV 2024)*. 2024, 370–386 **3**.
- [CZC*25] CHANG, HANZHI, ZHU, RUIJIE, CHANG, WENJIE, et al. “MeshSplat: Generalizable Sparse-View Surface Reconstruction via Gaussian Splatting”. *arXiv preprint arXiv:2508.17811* (2025) **3**.
- [DDB22] DANECEK, RADEK, BLACK, MICHAEL, and BOLKART, TIMO. “EMOCA: Emotion Driven Monocular Face Capture and Animation”. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. New Orleans, LA, USA: IEEE, June 2022, 20279–20290. ISBN: 978-1-66546-946-3. DOI: [10.1109/CVPR52688.2022.01967](https://doi.org/10.1109/CVPR52688.2022.01967). (Visited on 12/22/2023) **3**.
- [Deb12] DEBEVEC, PAUL. “The Light Stages and Their Applications to Photoreal Digital Actors”. *SIGGRAPH Asia*. Singapore, Nov. 2012 **3**.
- [DHT*00] DEBEVEC, PAUL, HAWKINS, TIM, TCHOU, CHRIS, et al. “Acquiring the Reflectance Field of a Human Face”. *ACM Transactions on Graphics (Proc. SIGGRAPH)*. 2000, 145–156 **5**.
- [DXX*24] DAI, PINXUAN, XU, JIAMIN, XIE, WENXIANG, et al. “High-quality Surface Reconstruction using Gaussian Surfels”. *SIGGRAPH 2024 Conference Papers*. 2024, 22. DOI: [10.1145/3641519.3657441](https://doi.org/10.1145/3641519.3657441) **3**.
- [FBTB20] FERNÁNDEZ ABREVAYA, VICTORIA, BOUKHAYMA, ADNANE, TORR, PHILIP H.S., and BOYER, EDMOND. “Cross-Modal Deep Face Normals with Deactivable Skip Connections”. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2020, 4979–4989 **3**.
- [FFBB21] FENG, YAO, FENG, HAIWEN, BLACK, MICHAEL J., and BOLKART, TIMO. “Learning an Animatable Detailed 3D Face Model from In-The-Wild Images”. *ACM SIGGRAPH 2021 Conference Proceedings*. ACM, 2021. DOI: [10.1145/3450626.3459936](https://doi.org/10.1145/3450626.3459936) **3**.
- [FWS*18] FENG, YAO, WU, FAN, SHAO, XIAOHU, et al. “Joint 3D Face Reconstruction and Dense Alignment with Position Map Regression Network”. *Proceedings of the European Conference on Computer Vision (ECCV)*. 2018, 547–563 **3**.
- [GCL*21] GUO, YUDONG, CHEN, KEYU, LIANG, SEN, et al. “AD-NeRF: Audio Driven Neural Radiance Fields for Talking Head Synthesis”. *IEEE/CVF International Conference on Computer Vision (ICCV)*. arXiv preprint arXiv:2103.11078. 2021 **2**.
- [GKG*24] GIEBENHAIN, SIMON, KIRSCHSTEIN, TOBIAS, GEORGIOPOULOS, MARKOS, et al. “MonoNPHM: Dynamic Head Reconstruction from Monocular Videos”. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2024, 10747–10758 **3**.
- [GKR*25] GIEBENHAIN, SIMON, KIRSCHSTEIN, TOBIAS, RÜNZ, MARTIN, et al. *Pixel3DMM: Versatile Screen-Space Priors for Single-Image 3D Face Reconstruction*. 2025. arXiv: [2505.00615](https://arxiv.org/abs/2505.00615) [cs.CV]. URL: <https://arxiv.org/abs/2505.00615> **3**.
- [GLD*19] GUO, KAIWEN, LINCOLN, PETER, DAVIDSON, PHILIP, et al. “The Relightables: Volumetric Performance Capture of Humans with Realistic Relighting”. *ACM Transactions on Graphics (TOG)* 38.6 (2019), 217:1–217:19. DOI: [10.1145/3355089.3356571](https://doi.org/10.1145/3355089.3356571) **3**.
- [GTZN21] GAFNI, GUY, THIES, JUSTUS, ZOLLHÖFER, MICHAEL, and NIESSNER, MATTHIAS. “Dynamic Neural Radiance Fields for Monocular 4D Facial Avatar Reconstruction”. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2021, 8649–8658 **2**.
- [GWZ*25] GU, CHUN, WEI, XIAOFEI, ZENG, ZIXUAN, et al. “IRGS: Inter-Reflective Gaussian Splatting with 2D Gaussian Ray Tracing”. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2025, 10943–10952 **3**, **14**.

- [GXZ*24] GAO, XUAN, XIAO, HAIYAO, ZHONG, CHENGLAI, et al. "Portrait Video Editing Empowered by Multimodal Generative Priors". *ACM SIGGRAPH Asia 2024 Conference Papers*. 2024. DOI: [10.1145/3680528.3687601](https://doi.org/10.1145/3680528.3687601). URL: <https://doi.org/10.1145/3680528.3687601>.
- [HPX*22] HONG, YANG, PENG, BO, XIAO, HAIYAO, et al. "HeadNeRF: A Real-time NeRF-based Parametric Head Model". *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2022. 2.
- [HYC*24] HUANG, BINBIN, YU, ZEHAO, CHEN, ANPEI, et al. "2D Gaussian Splatting for Geometrically Accurate Radiance Fields". *SIGGRAPH 2024 Conference Papers*. Association for Computing Machinery, 2024. DOI: [10.1145/3641519.3657428](https://doi.org/10.1145/3641519.3657428), 3, 5, 7–9.
- [JOYB25] JENA, SHUBHENDU, OUSAFI, AMINE, YOUNES, MAE, and BOUKHAYMA, ADNANE. "Sparfels: Fast Reconstruction from Sparse Unposed Imagery". *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. Oct. 2025, 27476–27487. 3.
- [JV B25] JENA, SHUBHENDU, VUTUKUR, SHISHIR REDDY, and BOUKHAYMA, ADNANE. "SparSplat: Fast Multi-View Reconstruction with Generalizable 2D Gaussian Splatting". *arXiv preprint arXiv:2505.02175* (2025) 3.
- [Kar13] KARIS, BRIAN. *Real Shading in Unreal Engine 4*. Tech. rep. Epic Games, 2013 2, 6, 7.
- [KB15] KINGMA, D.P. and BA, J.B. "Adam: A Method for Stochastic Optimization". *International Conference on Learning Representations (ICLR)* (2015). URL: <https://arxiv.org/abs/1412.6980>, 3, 7.
- [KBM*24] KHIRODKAR, RAWAL, BAGAUTDINOV, TIMUR, MARTINEZ, JULIETA, et al. "Sapiens: Foundation for Human Vision Models". *Proceedings of the European Conference on Computer Vision (ECCV)*. 2024. 3.
- [KGT*24] KIRSCHSTEIN, TOBIAS, GIEBENHAIN, SIMON, TANG, JIA-PENG, et al. "GGHead: Fast and Generalizable 3D Gaussian Heads". *Proceedings of the ACM SIGGRAPH Asia 2024 Conference Papers*. 2024, 1–11. DOI: [10.1145/3586781.3587976](https://doi.org/10.1145/3586781.3587976). URL: <https://arxiv.org/abs/2406.09377>.
- [KKLD23a] KERBL, BERNHARD, KOPANAS, GEORGIOS, LEIMKÜHLER, THOMAS, and DRETTAKIS, GEORGE. "3D Gaussian Splatting for Real-Time Radiance Field Rendering". *ACM Transactions on Graphics* 42.4 (July 2023). URL: <https://repo-sam.inria.fr/fungraph/3d-gaussian-splatting/> 5.
- [KKLD23b] KERBL, BERNHARD, KOPANAS, GEORGIOS, LEIMKÜHLER, THOMAS, and DRETTAKIS, GEORGE. "3d gaussian splatting for real-time radiance field rendering". *ACM Transactions on Graphics* 42.4 (2023), 1–14 2, 3, 7.
- [LBB*17] LI, TIANYE, BOLKART, TIMO, BLACK, MICHAEL. J., et al. "Learning a model of facial shape and expression from 4D scans". *ACM Transactions on Graphics, (Proc. SIGGRAPH Asia)* 36.6 (2017), 194:1–194:17. URL: <https://doi.org/10.1145/3130800.3130813>, 2, 3, 5.
- [LCS*24] LI, JUNXUAN, CAO, CHEN, SCHWARTZ, GABRIEL, et al. "URAvatar: Universal Relightable Gaussian Codec Avatars". *Proceedings of the ACM SIGGRAPH Asia 2024 Conference*. ACM, 2024. DOI: [10.1145/3680528.3687653](https://doi.org/10.1145/3680528.3687653). URL: <https://arxiv.org/abs/2410.24223>.
- [LK19] LE, HA A. and KAKADIARIS, IOANNIS A. "Illumination-Invariant Face Recognition With Deep Relit Face Images". *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*. 2019, 2146–2155. DOI: [10.1109/WACV.2019.00232](https://doi.org/10.1109/WACV.2019.00232) 5.
- [LKB*25a] LEE, JAESEONG, KANG, TAEWOONG, BUEHLER, MARCEL, et al. "SurFHead: Affine Rig Blending for Geometrically Accurate 2D Gaussian Surfel Head Avatars". *The Thirteenth International Conference on Learning Representations*. 2025. URL: <https://openreview.net/forum?id=1x1gGg49jr.2>, 3.
- [LKB*25b] LEE, JAESEONG, KANG, TAEWOONG, BUEHLER, MARCEL, et al. "SurFHead: Affine Rig Blending for Geometrically Accurate 2D Gaussian Surfel Head Avatars". *The Thirteenth International Conference on Learning Representations*. 2025. URL: <https://openreview.net/forum?id=1x1gGg49jr.5>.
- [LYSS21] LIN, SHANCHUAN, YANG, LINJIE, SALEEMI, IMRAN, and SENGUPTA, SOUMYADIP. *Robust High-Resolution Video Matting with Temporal Guidance*. 2021. arXiv: [2108.11515](https://arxiv.org/abs/2108.11515) [cs.CV] 7.
- [MESK22] MÜLLER, THOMAS, EVANS, ALEX, SCHIED, CHRISTOPH, and KELLER, ALEXANDER. "Instant Neural Graphics Primitives with a Multiresolution Hash Encoding". *ACM Transactions on Graphics (Proc. SIGGRAPH)* 41.4 (July 2022), 102:1–102:15. 2.
- [MHP*19] MEKA, ABHIMITRA, HAENE, CHRISTIAN, PANDEY, ROHIT, et al. "Deep Reflectance Fields - High-Quality Facial Reflectance Field Inference From Color Gradient Illumination". Vol. 38. 4. July 2019. DOI: [10.1145/3306346.3323027](https://doi.org/10.1145/3306346.3323027). URL: <http://gvv.mpi-inf.mpg.de/projects/DeepReflectanceFields/> 5.
- [MHS*22] MUNKBERG, JACOB, HASSELGREN, JON, SHEN, TIAN-CHANG, et al. "Extracting Triangular 3D Models, Materials, and Lighting From Images". *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2022. 2, 6.
- [MLT88] MAGNENAT-THALMANN, NADIA, LAPERRIÈRE, RICHARD, and THALMANN, DANIEL. "Joint-dependent local deformations for hand animation and object grasping". *Proceedings on Graphics Interface*. 1988, 26–33. 5.
- [MMP*24] MOENNE-LOCCOZ, NICOLAS, MIRZAEI, ASHKAN, PEREL, OR, et al. "3D Gaussian Ray Tracing: Fast Tracing of Particle Scenes". *ACM Transactions on Graphics and SIGGRAPH Asia* (2024) 14.
- [MST*20] MILDENHALL, BEN, SRINIVASAN, PRATUL P., TANCIK, MATTHEW, et al. "NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis". *Proceedings of the European Conference on Computer Vision (ECCV)*. 2020. 2, 4.
- [MWSZ24] MA, SHENGJIE, WENG, YANLIN, SHAO, TIANJIA, and ZHOU, KUN. "3D Gaussian Blendshapes for Head Avatar Animation". *ACM SIGGRAPH 2024 Conference Proceedings*. ACM, 2024. DOI: [10.1145/3641519.3657462](https://doi.org/10.1145/3641519.3657462), 3, 9.
- [NLML20] NESTMEYER, THOMAS, LALONDE, JEAN-FRANCOIS, MATTHEWS, IAIN, and LEHRMANN, ANDREAS. "Learning Physics-Guided Face Relighting Under Directional Light". *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2020. 5.
- [OBD*23] OLIVIER, NICOLAS, BAERT, KELIAN, DANIEAU, FABIEN, et al. "FaceTuneGAN: Face autoencoder for convolutional expression transfer using neural generative adversarial networks". *Comput. Graph.* 110.C (Feb. 2023), 69–85. ISSN: 0097-8493. DOI: [10.1016/j.cag.2022.12.004](https://doi.org/10.1016/j.cag.2022.12.004). URL: <https://doi.org/10.1016/j.cag.2022.12.004>, 3, 5.
- [PEL*21] PANDEY, ROHIT, ESCOLANO, SERGIO ORTS, LEGENDRE, CHLOE, et al. "Total relighting: learning to relight portraits for background replacement". *ACM Transactions on Graphics (TOG)* 40.4 (2021), 1–21. 5.
- [PHF*25] PANDEY, KARRAN, HU, ANITA, FUJI-TSANG, CLEMENT, et al. "Painting with 3D Gaussian Splat Brushes". *ACM SIGGRAPH 2025 Conference Papers*. 2025. 4.
- [Pho75] PHONG, BUI TUONG. "Illumination for Computer Generated Pictures". *Communications of the ACM* 18.6 (1975), 311–317. DOI: [10.1145/360825.360839](https://doi.org/10.1145/360825.360839), 2, 5.
- [PKA*09] PAYSAN, PASCAL, KNOTHE, REINHARD, AMBERG, BRIAN, et al. "A 3D face model for pose and illumination invariant face recognition". *2009 sixth IEEE international conference on advanced video and signal based surveillance*. Ieee. 2009, 296–301. 3.
- [PTS23] PONGLERTNAPAKORN, PUNTAWAT, TRITRONG, NONTAWAT, and SUWAJANAKORN, SUPASORN. "DiFaReli: Diffusion Face Relighting". 2023. 5.

- [QKS*24] QIAN, SHENHAN, KIRSCHSTEIN, TOBIAS, SCHONEVELD, LIAM, et al. "GaussianAvatars: Photorealistic Head Avatars with Rigged 3D Gaussians". *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2024, 20299–20309 **2**, **3**, **5**, **7**.
- [RBSB18] RANJAN, ANURAG, BOLKART, TIMO, SANYAL, SOUBHIK, and BLACK, MICHAEL J. "Generating 3D Faces using Convolutional Mesh Autoencoders". *European Conference on Computer Vision (ECCV)*. 2018, 725–741. DOI: [10.1007/978-3-030-01219-9_43](https://doi.org/10.1007/978-3-030-01219-9_43) **3**, **5**.
- [RCB*25] RONG, VICTOR, CHEN, JINGXIANG, BAHMANI, SHERWIN, et al. "Gstex: Per-primitive texturing of 2d gaussian splatting for decoupled appearance and geometry modeling". *2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*. IEEE. 2025, 3508–3518 **2**, **6**.
- [RFD*24] RETSINAS, GEORGE, FILNTISIS, PANAGIOTIS P., DANECSEK, RADEK, et al. "3D Facial Expressions through Analysis-by-Neural-Synthesis". *Conference on Computer Vision and Pattern Recognition (CVPR)*. 2024 **3**, **7**.
- [RFM*24] RAO, PRAMOD, FOX, GEREON, MEKA, ABHIMITRA, et al. "Lite2Relight: 3D-aware Single Image Portrait Relighting". *ACM SIGGRAPH 2024 Conference Papers*. SIGGRAPH '24. Denver, CO, USA: Association for Computing Machinery, 2024. ISBN: 9798400705250. DOI: [10.1145/3641519.3657470](https://doi.org/10.1145/3641519.3657470). URL: <https://doi.org/10.1145/3641519.3657470>.
- [SBT*19] SUN, TIANCHENG, BARRON, JONATHAN T., TSAI, YUN-TA, et al. "Single image portrait relighting". *ACM Trans. Graph.* 38.4 (July 2019). ISSN: 0730-0301. DOI: [10.1145/3306346.3323008](https://doi.org/10.1145/3306346.3323008). URL: <https://doi.org/10.1145/3306346.3323008>.
- [SCD*25] SCHONEVELD, LIAM, CHEN, ZHE, DAVOLI, DAVIDE, et al. "SHeaP: Self-Supervised Head Geometry Predictor Learned via 2D Gaussians". *arXiv preprint arXiv:2504.12292* (2025) **3**.
- [SCL*24] SONG, LUCHUAN, CHEN, LELE, LIU, CELONG, et al. "Text-Toon: Real-Time Text Toonify Head Avatar from Single Video". *Proceedings of the ACM SIGGRAPH Asia 2024 Conference Papers*. ACM, 2024, 1–11. DOI: [10.1145/3680528.3687632](https://doi.org/10.1145/3680528.3687632). URL: <https://dl.acm.org/doi/10.1145/3680528.3687632>.
- [SGV*24] SRINIVASAN, PRATUL P, GARBIN, STEPHAN J, VERBIN, DOR, et al. "Nuvo: Neural uv mapping for unruly 3d representations". *European Conference on Computer Vision*. Springer. 2024, 18–34 **4**.
- [SGX*25] SUN, HANXIAO, GAO, YUPENG, XIE, JIN, et al. "SVG-IR: Spatially-Varying Gaussian Splatting for Inverse Rendering". *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2025, 16143–16152 **3**, **14**.
- [SKC]18] SENGUPTA, SOUMYADIP, KANAZAWA, ANGIOO, CASTILLO, CARLOS D., and JACOBS, DAVID W. "SfSNet: Learning Shape, Reflectance and Illuminance of Faces in the Wild". *Computer Vision and Pattern Recognition (CVPR)*. 2018 **5**.
- [SLL*24] SONG, YUNZHOU, LIN, HEGUANG, LEI, JIAHUI, et al. "HDGS: Textured 2D Gaussian Splatting for Enhanced Scene Rendering". *arXiv preprint arXiv:2412.01823* (2024). URL: <https://arxiv.org/abs/2412.01823>.
- [SMAD24] SVITOV, DAVID, MORERIO, PIETRO, AGAPITO, LOURDES, and DEL BUE, ALESSIO. "BillBoard Splatting (BBSplat): Learnable Textured Primitives for Novel View Synthesis". *arXiv preprint arXiv:2411.08508* (2024) **2**, **6**.
- [Sor05] SORKINE, OLGA. "Laplacian Mesh Processing". *Eurographics (State of the Art Reports)* 4.4 (2005) **10**.
- [SSS*24] SAITO, SHUNSUKE, SCHWARTZ, GABRIEL, SIMON, TOMAS, et al. "Relightable Gaussian Codec Avatars". *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2024, 130–141. DOI: [10.1109/CVPR52733.2024.000213](https://doi.org/10.1109/CVPR52733.2024.000213).
- [SWL*24] SHAO, ZHIJING, WANG, ZHAOLONG, LI, ZHUANG, et al. "SplattingAvatar: Realistic Real-Time Human Avatars with Mesh-Embedded Gaussian Splatting". *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2024, 159–169. DOI: [10.1109/CVPR52733.2024.001593](https://doi.org/10.1109/CVPR52733.2024.001593) **3**, **4**, **9**.
- [SYH*17] SHU, Z., YUMER, E., HADAP, S., et al. "Neural Face Editing with Intrinsic Image Disentangling". *Computer Vision and Pattern Recognition, 2017. CVPR 2017. IEEE Conference on*. IEEE. 2017, - **4**.
- [TL18] TRAN, LUAN and LIU, XIAOMING. "On Learning 3D Face Morphable Model from In-the-Wild Images". *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2018 **3**, **5**.
- [TLM*25] TONG, JINGUANG, LI, XUESONG, MAKEN, FAHIRA AFZAL, et al. "GS-2DGS: Geometrically Supervised 2DGS for Reflective Object Reconstruction". *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2025, 21547–21557 **3**.
- [VHM*22] VERBIN, DOR, HEDMAN, PETER, MILDENHALL, BEN, et al. "Ref-NeRF: Structured View-Dependent Appearance for Neural Radiance Fields". *CVPR* (2022) **3**.
- [WB24] WEISS, SEBASTIAN and BRADLEY, DEREK. "Gaussian Billboards: Expressive 2D Gaussian Splatting with Textures". *arXiv preprint arXiv:2412.12734* (2024). URL: <https://arxiv.org/abs/2412.12734>.
- [WKS*25] WANG, CONG, KANG, DI, SUN, HEYI, et al. "Mega: Hybrid mesh-gaussian head avatar for high-fidelity rendering and head editing". *Proceedings of the Computer Vision and Pattern Recognition Conference*. 2025, 26274–26284 **4**.
- [WYL*20] WANG, ZHIBO, YU, XIN, LU, MING, et al. "Single image portrait relighting via explicit multiple reflectance channel modeling". *ACM Trans. Graph.* 39.6 (Nov. 2020). ISSN: 0730-0301. DOI: [10.1145/3414685.3417824](https://doi.org/10.1145/3414685.3417824). URL: <https://doi.org/10.1145/3414685.3417824>.
- [WZL*09] WANG, YANG, ZHANG, LEI, LIU, ZICHENG, et al. "Face Relighting from a Single Image under Arbitrary Unknown Lighting Conditions". *IEEE Transactions on Pattern Analysis and Machine Intelligence* 31.11 (2009), 1968–1984. DOI: [10.1109/TPAMI.2008.2444](https://doi.org/10.1109/TPAMI.2008.2444).
- [XCL*24] XU, YUELANG, CHEN, BENWANG, LI, ZHE, et al. "Gaussian Head Avatar: Ultra High-fidelity Head Avatar via Dynamic Gaussians". *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2024 **3**.
- [XCW*24] XU, RUI, CHEN, WENYUE, WANG, JIEPENG, et al. "Super-Gaussians: Enhancing Gaussian Splatting Using Primitives with Spatially Varying Colors". *arXiv preprint arXiv:2411.18966* (2024). URL: <https://arxiv.org/abs/2411.18966>.
- [XCX*24] XIE, TAO, CHEN, XI, XU, ZHEN, et al. "EnvGS: Modeling View-Dependent Appearance with Environment Gaussian". *arXiv preprint arXiv:2412.15215* (2024) **14**.
- [XGGZ24] XIANG, JUN, GAO, XUAN, GUO, YUDONG, and ZHANG, JUYONG. "FlashAvatar: High-fidelity Head Avatar with Efficient Gaussian Embedding". *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2024 **3–5**, **9**.
- [XHL*24] XU, TIAN-XING, HU, WENBO, LAI, YU-KUN, et al. "TextureGS: Disentangling the Geometry and Texture for 3D Gaussian Splatting Editing". *Computer Vision – ECCV 2024*. Springer, 2024, 1–16. DOI: [10.1007/978-3-031-72698-9_3](https://doi.org/10.1007/978-3-031-72698-9_3). URL: <https://arxiv.org/abs/2403.10050>.
- [XSD*24] XI, CHEN, SIDA, PENG, DONGCHEN, YANG, et al. "IntrinsicAnything: Learning Diffusion Priors for Inverse Rendering Under Unknown Illumination". *arxiv:2404.11593* (2024) **7**.
- [XXH*21] XIANG, FANBO, XU, ZEXIANG, HASAN, MILOS, et al. "Neutex: Neural texture mapping for volumetric neural rendering". *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2021, 7119–7128 **4**.

- [YB25a] YOUNES, MAE and BOUKHAYMA, ADNANE. “Anti-Aliased 2D Gaussian Splatting”. *Advances in Neural Information Processing Systems (NeurIPS)*. Poster Presentation. 2025. URL: <https://neurips.cc/virtual/2025/poster/119938> 14.
- [YB25b] YOUNES, MAE and BOUKHAYMA, ADNANE. “TextureSplat: Per-Primitive Texture Mapping for Reflective Gaussian Splatting”. *arXiv preprint arXiv:2506.13348* (2025) 2, 3.
- [YDYK24] YE, MINGQIAO, DANELLJAN, MARTIN, YU, FISHER, and KE, LEI. “Gaussian Grouping: Segment and Edit Anything in 3D Scenes”. *Computer Vision – ECCV 2024*. Springer, 2024, 419–438. DOI: [10.1007/978-3-031-19778-9_15](https://doi.org/10.1007/978-3-031-19778-9_15) 4.
- [YJM*24] YU, HENG, JULIN, JOEL, MILACSKI, ZOLTÁN Á, et al. “Cogs: Controllable gaussian splatting”. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2024, 21624–21633 4.
- [YLH*24] YUAN, YE, LI, XUETING, HUANG, YANGYI, et al. “GAvatar: Animatable 3D Gaussian Avatars with Implicit Mesh Learning”. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2024 3.
- [YZF*23] YANG, HAOTIAN, ZHENG, MINGWU, FENG, WANQUAN, et al. “Towards Practical Capture of High-Fidelity Relightable Avatars”. *ACM SIGGRAPH Asia 2023 Conference Papers*. ACM, 2023, 1–11. DOI: [10.1145/3610548.3618138](https://doi.org/10.1145/3610548.3618138) 3.
- [YZG*25] YAO, YUXUAN, ZENG, ZIXUAN, GU, CHUN, et al. “Reflective Gaussian Splatting”. *ICLR*. 2025 3.
- [YZM*24] YANG, HAOTIAN, ZHENG, MINGWU, MA, CHONGYANG, et al. “VRMM: A Volumetric Relightable Morphable Head Model”. *ACM SIGGRAPH 2024 Conference Proceedings*. 2024, 1–11. DOI: [10.1145/3641519.3657406](https://doi.org/10.1145/3641519.3657406) 3.
- [ZBT22] ZIELONKA, WOJCIECH, BOLKART, TIMO, and THIES, JUSTUS. “Towards Metrical Reconstruction of Human Faces”. *Computer Vision – ECCV 2022*. Springer, 2022, 249–266. DOI: [10.1007/978-3-031-19778-9_15](https://doi.org/10.1007/978-3-031-19778-9_15) 3.
- [ZBT23] ZIELONKA, WOJCIECH, BOLKART, TIMO, and THIES, JUSTUS. “Instant Volumetric Head Avatars”. *CVPR*. 2023, 4574–4584 9, 10.
- [ZDP*24] ZENG, CHONG, DONG, YUE, PEERS, PIETER, et al. “DiLightNet: Fine-grained Lighting Control for Diffusion-based Image Generation”. *ACM SIGGRAPH 2024 Conference Papers*. 2024 5.
- [ZHS]19] ZHOU, HAO, HADAP, SUNIL, SUNKAVALLI, KALYAN, and JACOBS, DAVID W. “Deep single-image portrait relighting”. *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2019, 7194–7202 5.
- [ZLDF21] ZHANG, ZHIMENG, LI, LINCHENG, DING, YU, and FAN, CHANGJIE. “Flow-Guided One-Shot Talking Face Generation With a High-Resolution Audio-Visual Dataset”. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2021, 3661–3670 10.
- [ZLKT23] ZHAN, FANGNENG, LIU, LINGJIE, KORTYLEWSKI, ADAM, and THEOBALT, CHRISTIAN. “General neural gauge fields”. *arXiv preprint arXiv:2305.03462* (2023) 4.
- [ZLL*25] ZHANG, DONGBIN, LIU, YUNFEI, LIN, LIJIAN, et al. “HRAAvatar: High-Quality and Relightable Gaussian Head Avatar”. *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR)*. June 2025, 26285–26296 2, 3, 5–7, 9, 10, 12, 13.
- [ZPVG01] ZWICKER, MATTHIAS, PEISTER, HANSPETER, VAN BAAR, JEROEN, and GROSS, MARKUS. “EWA Volume Splatting”. *Proceedings of the IEEE Conference on Visualization (VIS)*. IEEE, 2001, 29–36 3.
- [ZPVG02] ZWICKER, MATTHIAS, PEISTER, HANSPETER, VAN BAAR, JEROEN, and GROSS, MARKUS. “EWA splatting”. *IEEE Transactions on Visualization and Computer Graphics* 8.3 (2002), 223–238 3.
- [ZWL*25] ZHANG, JIAWEI, WU, ZIJIAN, LIANG, ZHIYANG, et al. “Fate: Full-head gaussian avatar with textural editing from monocular video”. *Proceedings of the Computer Vision and Pattern Recognition Conference*. 2025, 5535–5545 2, 4, 11.
- [ZYC*24] ZHANG, XINCHEN, YANG, LING, CAI, YAQI, et al. “Real-Compo: Balancing Realism and Compositionality Improves Text-to-Image Diffusion Models”. *Advances in Neural Information Processing Systems (NeurIPS)*, 37. arXiv:2402.12908. Curran Associates, Inc., 2024, -. URL: <https://arxiv.org/abs/2402.12908> 4.
- [ZYL*25] ZHANG, XINCHEN, YANG, LING, LI, GUOHAO, et al. “Iter-Comp: Iterative Composition-Aware Feedback Learning from Model Gallery for Text-to-Image Generation”. *Proceedings of the International Conference on Learning Representations (ICLR)*. Poster Presentation. 2025. URL: <https://openreview.net/forum?id=4w99NAikOE> 4.
- [ZYW*23] ZHENG, YUFENG, YIFAN, WANG, WETZSTEIN, GORDON, et al. “PointAvatar: Deformable Point-based Head Avatars from Videos”. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2023 3, 7, 9.