

Relightable Gaussian Codec Avatars

Shunsuke Saito, Gabriel Schwartz, Tomas Simon, Junxuan Li, Giljoo Nam

Codec Avatars Lab, Meta

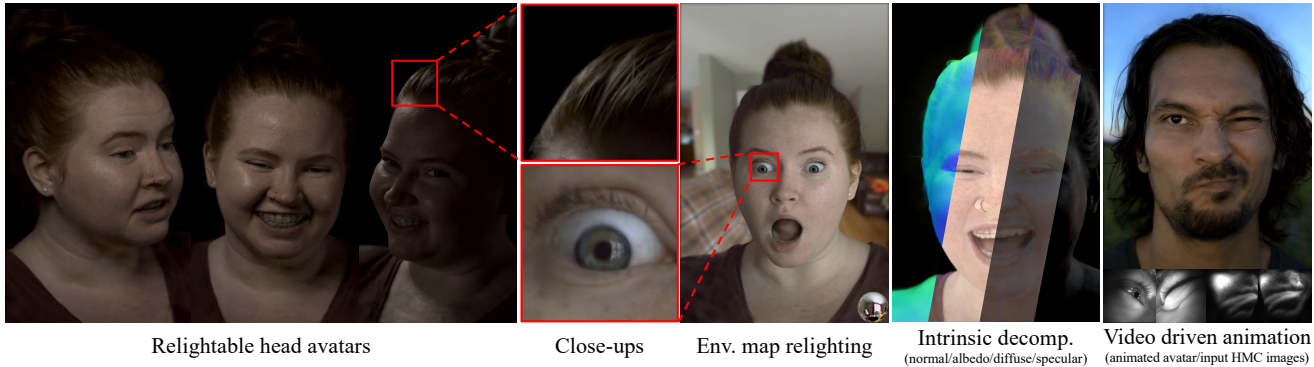


Figure 1. **Relightable Gaussian Codec Avatars.** Our approach enables real-time relighting of human head avatars with all-frequency reflections and detailed hair reconstruction using 3D Gaussians and learnable radiance transfer. Our dynamic avatars can be driven live in real-time from images captured with head mounted cameras (HMC). <https://shunsukesaito.github.io/rgca/>

Abstract

The fidelity of relighting is bounded by both geometry and appearance representations. For geometry, both mesh and volumetric approaches have difficulty modeling intricate structures like 3D hair geometry. For appearance, existing relighting models are limited in fidelity and often too slow to render in real-time with high-resolution continuous environments. In this work, we present Relightable Gaussian Codec Avatars, a method to build high-fidelity relightable head avatars that can be animated to generate novel expressions. Our geometry model based on 3D Gaussians can capture 3D-consistent sub-millimeter details such as hair strands and pores on dynamic face sequences. To support diverse materials of human heads such as the eyes, skin, and hair in a unified manner, we present a novel relightable appearance model based on learnable radiance transfer. Together with global illumination-aware spherical harmonics for the diffuse components, we achieve real-time relighting with all-frequency reflections using spherical Gaussians. This appearance model can be efficiently relit under both point light and continuous illumination. We further improve the fidelity of eye reflections and enable explicit gaze control by introducing relightable explicit eye models. Our method outperforms existing approaches without compromising real-time performance. We also demonstrate real-time relighting of avatars on a tethered consumer VR headset, showcasing the efficiency and fidelity of our avatars.

1. Introduction

What makes avatar relighting so challenging? Our visual perception is highly sensitive to facial appearance. Convincing the visual system requires modeling each part of the head in sufficient detail that is coherent with an environment, and this synthesis typically needs to be performed in real-time for primary applications of photorealistic avatars including games and telecommunication [39, 57]. Real-time relighting of animatable human heads with convincing details remains an open challenge for three reasons.

The first challenge is that human heads are composed of highly complex and diverse materials that exhibit different properties of scattering and reflectance. For example, skin produces intricate reflections due to micro-geometry as well as significant subsurface scattering [53, 78], hair exhibits out-of-plane scattering with multiple reflections due to its translucent fiber structure [46], and the eyes have multiple layers with highly reflective membranes [33, 67]. By and large, there is no single material representation that accurately represents them all, especially in real-time. Moreover, precise tracking and modeling of the underlying geometry in motion is extremely challenging because deformations do not always contain sufficient visual markers to track. Finally, the real-time requirement severely limits the algorithmic design. Increase in photorealism traditionally results in an exponential increase in the cost of transporting light and tracking motion. Our goal is to design a learn-

ing framework that builds real-time renderable head avatars with accurate scattering and reflections under illuminations of any frequency.

Given exhaustive measurements obtained using a light-stage [10, 16, 45], physically-based rendering methods [68, 78] can generalize to novel illuminations. However, it remains non-trivial to extend these methods to dynamic performance capture and non-skin parts such as hair and eyeballs. Additionally, acquiring sufficiently accurate geometry and material parameters is a laborious process with a significant amount of manual cleanup required [68].

More recently, neural relighting approaches sidestep the need for accurate geometry and material modeling by only modeling the direct relationship between the input (i.e., illumination) and output (i.e., outgoing radiance) with neural networks and approximated geometry using meshes [7], volumetric primitives [20, 35, 87], and neural fields [33, 66]. Typically these models are learned from one-light-at-a-time (OLAT) [66] or grouped lights [7, 20, 87] controlled by a light-stage, and supporting real-time rendering with continuous illumination requires expensive teacher-student distillation [7, 20] or physically-inspired appearance models that explicitly maintain key properties of light transport such as linearity [33, 87]. Despite promising results, we observe that the existing approaches lead to suboptimal performance due to insufficient expressiveness of both the geometric and appearance representations. In particular, none of the methods achieve all-frequency reflections on hair and eyes, and submillimeter thin structures such as hair strands are often blurred out or glued into larger blobs, making hair rendering less than photorealistic.

To address the aforementioned issues, we present three contributions: (1) drivable avatars based on 3D Gaussians that can be efficiently rendered with intricate geometric details, (2) a relightable appearance model based on learned radiance transfer that supports global light transport and all-frequency reflections in real-time, and (3) a relightable explicit eye model that enables disentangled control of gaze from other facial movements as well as all-frequency eye-reflections for the first time in a fully data-driven manner.

3D Gaussian Avatars. Our geometric representation is based on 3D Gaussians [24] that can be rendered in real-time using splatting. To achieve a drivable avatar, we decode 3D Gaussians on a shared UV space for a template head using 2D convolutional neural networks. We encode the driving signals such as facial expressions in a self-supervised manner akin to traditional *codecs*. This allows us to track the moving heads in a temporally coherent manner with intricate geometric details such as hair strands.

Learned Radiance Transfer. For appearance, inspired by precomputed radiance transfer [69], we introduce a relightable appearance model based on learnable radiance transfer that consists of diffuse spherical harmonics and

specular spherical Gaussians. We learn diffuse radiance transfer parameterized by dynamic spherical harmonics coefficients for each 3D Gaussian. This transfer precomputes visibility and global light transport, including multi-bounce and subsurface scattering. For specular reflection, we introduce a novel parameterization of spherical Gaussians [74, 89] with view-dependent visibility that effectively approximates the combined effects of occlusion, Fresnel, and geometric attenuation without explicitly estimating the individual contributions. Our specular Gaussian lobe is aligned with the reflection vector and computed using the view direction and per-Gaussian view-dependent normals. Most importantly, spherical Gaussians support all-frequency reflection under high-resolution illuminations in real-time. Both diffuse and specular representations satisfy the linearity of light transport, hence supporting real-time rendering under both point lights and environment illumination without additional training. In addition, the proposed learnable radiance transfer supports global light transport and all-frequency reflection of the eyes, skin, and hair with the unified representation, significantly simplifying the learning process while achieving extremely high-fidelity relighting.

Relightable Explicit Eye Model. To reproduce reflections on the cornea, our relightable Gaussian avatar incorporates an explicit eye model [67] that also enables explicit control of the eyeballs with better disentanglement. In addition, our appearance model naturally supports relighting of the eyes with all-frequency reflections, which is crucial for photorealism under natural environments.

We run an evaluation of various pairs of geometry and relightable appearance models in this work and other real-time renderable baseline methods. Our experiments show that the combination of 3D Gaussians with our relighting appearance model outperforms any other combination.

2. Related Work

Face Modeling. Facial avatar modeling has been an active research topic for over half a century [61]. We refer to [61] for a comprehensive overview on research and tools for artist-friendly authoring of 3D facial models. Advancements in image-based 3D reconstruction [12] enabled the precise and more automated acquisition of 3D faces using multi-view capture systems, especially in high-end film production [2, 4, 9, 16, 90]. These approaches primarily focus on the facial skin region, and more tailored solutions are required for the reconstruction and modeling of different components such as teeth [79], lips [15], facial hair [3], eyes [5, 51], and hair [19, 43, 55, 60], which are difficult to scale for dynamic and complete head avatars.

More recently, learning-based approaches emerged to holistically represent human heads without requiring precise input geometry [39, 44, 94, 96]. In particular, volu-

metric representations [40, 50] show the promise of representing both skin and more complex geometric structures like hair with a single representation [27, 94, 96]. To enable real-time rendering with volumes, several hybrid approaches are proposed to partition the space for efficient raymarching using mixture of volumetric primitives [41] or tetrahedra [14]. Point clouds are also utilized to model head avatars [95]. However, we observe that the existing shape representations for real-time renderable avatars struggle with modeling extremely thin structures such as hair strands. To address this limitation, we extend a state-of-the-art efficient scene representation based on 3D Gaussian splatting [24] to animatable facial avatar modeling. While several works already show dynamic modeling of 3D Gaussians [42, 80], we are the first to enable the modeling of animatable and (most importantly) relightable 3D Gaussians.

Facial Reflectance Capture. In the early 2000s, visual production was a great driver for facial reflectance capture and relighting research for composing actors into virtual environments. A seminal work by Debevec *et al.* [10] demonstrated that one-light-at-a-time (OLAT) captures can be used to obtain reflectance properties and relight faces in novel illuminations by leveraging the linearity of light transport. Follow-up work further extended the method to dynamic relighting [62], and fast acquisition of reflectance maps using spherical gradient illuminations [16, 18, 45]. Subsequently, the collection of larger datasets allowed estimating reflectance from a single image using neural networks [31, 32, 34, 36, 59, 86]. However, accurate reflectance estimation is typically limited to skin regions because the intricate hair and eye structure make the inverse rendering intractable. While inverse rendering with various scene representations has also been proposed to estimate spatially-varying BRDFs (SVBRDFs) [7, 52, 54, 89, 93], it remains a challenge to photorealistically model complete human heads due to the highly complex and diverse material and geometric composition. The lack of photorealism is also evident in recent relightable head modeling in the wild using simple BRDF and shading models [11, 65, 95].

Neural Relighting. Instead of modeling BRDF parameters, learning-based relighting approaches attempt to directly learn relightable appearance from a light-stage capture [13, 48, 49, 66, 83, 84, 92]. While these approaches show promising relighting for static [66, 83, 84, 92] and dynamic scenes [48, 49], they do not support generating novel animations, which is an essential requirement for avatars. Portrait relighting methods [58, 70, 71, 76, 88] also enable relighting under novel illuminations given a single image. However, they cannot produce novel view synthesis or temporally coherent dynamic relighting. Bi *et al.* [6] propose a neural-rendering method that supports global illumination for animatable facial avatars. To enable real-time rendering with natural environments, they distillate a slow

teacher model conditioned with individual point lights into a light-weight student model that can be conditioned with environment maps. This work is later extended to articulated hand modeling [20], compositional modeling of heads and eyeglasses [35], and scalable training by eliminating the need of teacher-student distillation [87]. These relightable avatars take as **input** the lighting information, which we discover is the main limiting factor for expressiveness in all-frequency relighting. In contrast, inspired by Precomputed Radiance Transfer (PRT) [69, 74], we propose to integrate a target illumination at the **output** of our neural decoder, improving quality and simplifying the learning pipeline.

Precomputed Radiance Transfer. In computer graphics, rendering a scene with global illumination is an expensive process due to iterative path tracing or multiple bounce computations. To enable real-time rendering with global light transport, Sloan *et al.* [69] propose to precompute a part of light transport that only depends on intrinsic scene properties, such as geometry and reflectance, and then integrate the precomputed intrinsic factor with an extrinsic illumination at runtime. For fast integration, they utilize spherical harmonics as an angular basis. To overcome the limited frequency band in spherical harmonics, follow-up works introduce wavelets [56], spherical radial basis functions [73], spherical Gaussians [17, 74], anisotropic spherical Gaussians [81], and neural network-based decompositions [85]. Similarly, Neural PRT [63] applies the same principle to screen-space relighting based on neural deferred rendering [72]. Despite many desirable properties, these methods primarily focus on static scenes due to the dependency on knowing geometry and material properties. Unfortunately, we neither know the geometry and material properties for human heads *a priori*, nor are they static. Thus, we propose to **learn** the intrinsic radiance transfer from dynamic real-data observations without explicitly assuming any material types or underlying geometry. The closest to our work in terms of the appearance representation is EyeNeRF [33], where they learn view-independent spherical harmonics for diffuse and view-conditioned spherical harmonics for specular components from image observations to build a relightable eye model. However, this appearance model suffers from the limited expressiveness of spherical harmonics for specular reflections. Empirically, we find that their proposed model does not generalize well to novel view and light directions. Please refer to Sec. 4 for our analysis.

3. Method

In this section, we provide details of the data acquisition process (Sec. 3.1), geometry and appearance representations (Sec. 3.2–3.3), the relightable explicit eye model (Sec. 3.4), as well as the training method (Sec. 3.5).

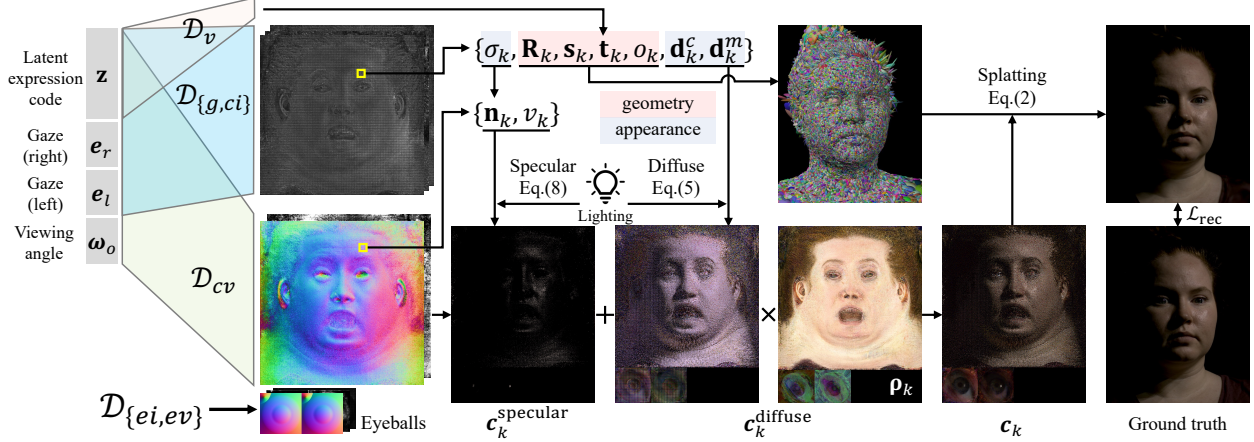


Figure 2. **Overview.** Given an expression latent code \mathbf{z} , gaze $\mathbf{e}_{\{l,r\}}$, and view direction $\boldsymbol{\omega}_o$, our model decodes the parameters of 3D Gaussians (rotation \mathbf{R}_k , translation \mathbf{t}_k , scale \mathbf{s}_k , and opacity o_k) and learned radiance transfer functions (colored and monochrome diffuse SH coefficients $\mathbf{d}_k^c, \mathbf{d}_k^m$, roughness σ_k , normal \mathbf{n}_k , and visibility v_k). We integrate the decoded radiance transfer functions with the input light to compute the final color \mathbf{c}_k , which we then render via splatting and supervise in image space. The coarse vertex decoder \mathcal{D}_v and geometry decoder \mathcal{D}_g are described in Sec. 3.2, the appearance decoders $\mathcal{D}_{\{ci,cv\}}$ in Sec. 3.3, and eyeball decoders $\mathcal{D}_{\{ei,ev\}}$ in Sec. 3.4.

3.1. Data Acquisition

We use a setup similar to [6, 35], where we capture calibrated and synchronized multiview images at 4096×2668 resolution with 110 cameras and 460 white LED lights at 90 Hz. Each participant is asked to perform a predefined set of various facial expressions, sentences, and gaze motions for about 144,000 frames. To collect diverse illumination patterns while enabling stable facial tracking, we employ time-multiplexed illumination. In particular, the full-on illumination is interleaved at every third frame to allow tracking, and the rest is lit with grouped or random sets of 5 lights.

As in [6, 35], we perform a topologically consistent coarse mesh tracking using multi-view full-on images. We further stabilize the head pose using the mode pursuit method of Lamarre *et al.* [30]. We also estimate the eye gaze of both eyes using the method described in [67]. The tracked mesh, head pose, unwrapped averaged texture, and gaze are interpolated to adjacent partially lit frames for the following avatar training.

3.2. Geometry: 3D Gaussian Avatars

The core of our geometric representation is the mixture of 3D anisotropic Gaussians [24], which supports representing varying topology and can represent thin volumetric structures. We first review the underlying parameterization and key idea in 3D Gaussian Splatting [24] and then highlight major changes to enable animatable avatar modeling.

We render avatars as collections of 3D Gaussians, where each Gaussian $\mathbf{g}_k = \{\mathbf{t}_k, \mathbf{R}_k, \mathbf{s}_k, o_k, \mathbf{c}_k\}$ is defined by a translation $\mathbf{t}_k \in \mathbb{R}^3$, a rotation matrix $\mathbf{R}_k \in SO(3)$ parameterized as a quaternion, per-axis scale factors $\mathbf{s}_k \in \mathbb{R}^3$, an opacity value $o_k \in \mathbb{R}$, and a color $\mathbf{c}_k \in \mathbb{R}^3$. The spatial extent of each Gaussian is defined in 3D by its covariance

$$\boldsymbol{\Sigma}_k = \mathbf{R}_k \text{diag}(\mathbf{s}_k) \text{diag}(\mathbf{s}_k)^T \mathbf{R}_k^T.$$

This representation allows efficient rendering using the Elliptical Weighted Average (EWA) splatting technique proposed by Zwicker *et al.* [97] by computing the 2D projection of each 3D Gaussian,

$$\boldsymbol{\Sigma}'_k = \mathbf{J} \mathbf{V} \boldsymbol{\Sigma}_k \mathbf{V}^T \mathbf{J}^T, \quad (1)$$

where $\mathbf{J} \in \mathbb{R}^{2 \times 3}$ is the Jacobian of the projective transformation, $\mathbf{V} \in \mathbb{R}^{3 \times 3}$ the viewing transformation, and $\boldsymbol{\Sigma}'_k \in \mathbb{R}^{2 \times 2}$ is the covariance matrix of the projected 2D Gaussian (a ‘‘splat’’). The final color \mathbf{C}_p at pixel p given \mathcal{N} ordered splats is computed with point-based cumulative volumetric rendering [24, 28, 29] as follows:

$$\mathbf{C}_p = \sum_{k \in \mathcal{N}} \mathbf{c}_k \alpha_k \prod_{j=1}^{k-1} (1 - \alpha_j), \quad (2)$$

where the transparency α_k is evaluated using the 2D covariance $\boldsymbol{\Sigma}'_k$ and multiplied by the per-Gaussian opacity o_k . Please refer to [24] for more details.

In contrast to 3D Gaussian Splatting [24], which focuses on static scene reconstruction, our objective is to build an animatable 3D avatar representation that can span the dynamic facial expressions of a person and also be relit under novel illuminations. This necessitates a relightable appearance model to recolor \mathbf{c}_k of all Gaussians as a function of the environmental illumination, allowing for a realistic adaptation of the avatar’s appearance under varying lighting conditions. Additionally, it is essential to register the geometry $\{\mathbf{g}_k\}$ for all Gaussians in response to the state of any facial expressions, ensuring that the avatar’s expressions remain consistent with the user’s actual facial movements. Enabling the encoding and decoding of any facial movements is crucial for animating and driving avatars.

To this end, we parameterize the 3D Gaussians on a shared UV texture map of a coarse template mesh, and decode their transformation and opacity using 2D convolutional neural networks. As facial expressions are highly non-linear and non-trivial to precisely define, in the spirit of Lombardi *et al.* [39] and Xu *et al.* [82], we employ a conditional variational auto-encoder (CVAE) [26] to learn the latent distribution of facial expressions from the data. Given the eye gaze directions of both eyes $\mathbf{e}_{\{l,r\}} \in \mathbb{R}^3$ in head-centric coordinates, coarse tracked mesh vertices \mathbf{V} , an unwrapped averaged texture \mathbf{T} , our encoder \mathcal{E} and geometry decoder \mathcal{D}_g are defined as:

$$\begin{aligned} \mu_e, \sigma_e &= \mathcal{E}(\mathbf{V}, \mathbf{T}; \Theta_e), \\ \{\delta \mathbf{t}_k, \mathbf{R}_k, \mathbf{s}_k, o_k\}_{k=1}^M &= \mathcal{D}_g(\mathbf{z}, \mathbf{e}_{\{l,r\}}; \Theta_g), \end{aligned} \quad (3)$$

where Θ_e and Θ_g are the learnable parameters for the encoder and decoder respectively, M is the total number of Gaussians, and σ_e and μ_e are the mean and standard deviation of a normal distribution $\mathbf{z} \sim \mathcal{N}(\mu_e, \sigma_e)$. The sampled latent vector $\mathbf{z} \in \mathbb{R}^{256}$ is computed using the reparameterization trick proposed by Kingma and Welling [26]. We also decode mesh vertices from \mathbf{z} such that we can animate the avatars from a headset [77] or latent-space manipulation [1]:

$$\mathbf{V}' = \mathcal{D}_v(\mathbf{z}; \Theta_v). \quad (4)$$

Note that while we directly infer the rotation \mathbf{R} and scale \mathbf{s} , we use the coarse geometry g as guidance to avoid poor local minima under large motions. The final Gaussian position \mathbf{t}_k is computed as $\mathbf{t}_k = \hat{\mathbf{t}}_k + \delta \mathbf{t}_k$, where $\hat{\mathbf{t}}_k$ is the interpolated coarse mesh position of the corresponding UV-coordinates using barycentric interpolation of the vertices in \mathbf{V}' . To assign UV-coordinates to Gaussians, we map one Gaussian to each texel in the UV map.

While the aforementioned parameterization is similar to the Mixture of Volumetric Primitives (MVP) [41], where a collection of voxel grids is anchored on a template mesh and used as the renderable primitive, there are two important differences: (1) Unlike MVP, which requires raymarching, the 3D Gaussians can be more efficiently rendered using splatting [24, 97]. (2) Additionally, the Gaussians have a greater ability to recreate thin structures, yielding sharper appearance for hair, which we show in Sec. 4.

3.3. Appearance: Learned Radiance Transfer

An appearance model for faces must accurately model a wide range of light transport effects, including subsurface scattering in skin, and specular reflections on the skin, eyes, as well as multi-bounced scattering on the hair. As discussed in early works [69, 74], while a diffuse transfer operator is a low-pass filtering of incident illumination, a specular transfer operator requires the ability to represent all-frequency information for mirror-like reflections. To effec-

tively allocate the capacity of the network to each component, we decompose the final color \mathbf{c}_k of each 3D Gaussian into a view-independent diffuse term $\mathbf{c}_k^{\text{diffuse}}$, and a view-dependent specular term $\mathbf{c}_k^{\text{specular}}(\omega_o)$, such that $\mathbf{c}_k = \mathbf{c}_k^{\text{diffuse}} + \mathbf{c}_k^{\text{specular}}(\omega_o)$, where ω_o is the viewing direction.

Diffuse Color. Our diffuse term is based on spherical harmonics (SH), and incorporates global light transport effects including occlusion, subsurface scattering, and multi-bounce scattering [69]. The diffuse color contribution $\mathbf{c}_k^{\text{diffuse}}$ is computed by the spherical integration of an (extrinsic) incident illumination, $\mathbf{L}(\cdot)$, with an intrinsic function that models the radiance transfer, $\mathbf{d}(\cdot)$. By representing both functions in the SH basis, this can be efficiently computed as a dot product of the coefficient vectors due to orthonormality of the basis:

$$\mathbf{c}_k^{\text{diffuse}} = \rho_k \odot \int_{\mathbb{S}^2} \mathbf{L}(\omega_i) \odot \mathbf{d}_k(\omega_i) d\omega_i = \rho_k \odot \sum_{i=1}^{(n+1)^2} \mathbf{L}_i \odot \mathbf{d}_k^i, \quad (5)$$

where $\mathbf{L} = \{\mathbf{L}_i\}$ and $\mathbf{d}_k = \{\mathbf{d}_k^i\}$ are the n -th order SH coefficients of the incident light and the intrinsic radiance transfer function, with $\mathbf{d}_k^i \in \mathbb{R}^3$, and $\rho_k \in \mathbb{R}^3$ a learnable albedo which we statically define on each Gaussian to encourage temporally consistent reconstructions. While diffuse light transport is a low-pass filter that requires only 2nd or 3rd order SH [64], this is not sufficient to represent shadows. To enable higher-frequency shadows while decoding a manageable number of coefficients that can fit on consumer GPU memory, we propose to decode RGB intrinsic SH coefficients \mathbf{d}_k^c up to the 3rd order, and only monochrome intrinsic SH coefficients \mathbf{d}_k^m from 4-th to 8-th order.

Specular Reflection. To achieve sharp, mirror-like reflections, for the view-dependent specular term we use a spherical Gaussian (SG) as an angular basis. In particular, we propose a normalized, angle-based spherical Gaussian:

$$G_s(\mathbf{p}; \mathbf{q}, \sigma) = C e^{-\frac{1}{2} \left(\frac{\arccos(\frac{\mathbf{p} \cdot \mathbf{q}}{\sigma})}{\sigma} \right)^2}, \quad (6)$$

where $\sigma \in \mathbb{R}^+$ is the standard deviation of angular decay, $\mathbf{q} \in \mathbb{S}^2$ is the lobe axis, $\mathbf{p} \in \mathbb{S}^2$ is the direction of evaluation, and $C = 1/(\sqrt{2}\pi^{2/3}\sigma)$ is a normalization factor to preserve the integral of the Gaussian.

Note that this parameterization is different from the more widely used parameterization $G(\mathbf{p}; \mathbf{q}, \lambda, \mu) = \mu e^{\lambda(\mathbf{p} \cdot \mathbf{q} - 1)}$ [74, 89], but we observed that this choice often fails to model highly reflective surfaces, such as the cornea.

As the majority of specular BRDFs have a lobe that is axis-aligned with a particular reflection direction, we compute a reflection vector as the lobe axis:

$$\mathbf{q}_k = 2(\omega_k^o \cdot \mathbf{n}_k) \mathbf{n}_k - \omega_k^o, \quad (7)$$

where ω_k^o is the viewing direction evaluated at the Gaussian center, and \mathbf{n}_k is a normal direction computed for each

Gaussian, and \mathbf{q}_k is the lobe axis. Our final color for the specular term of each 3D Gaussian is defined as follows:

$$\mathbf{c}_k^{\text{specular}}(\boldsymbol{\omega}_o) = v_k(\boldsymbol{\omega}_o) \int_{\mathbb{S}^2} \mathbf{L}(\boldsymbol{\omega}_i) G_s(\boldsymbol{\omega}_i; \mathbf{q}_k, \sigma_k) d\boldsymbol{\omega}_i, \quad (8)$$

where $v_k(\boldsymbol{\omega}_o) \in (0, 1)$ is a learnable view-dependent visibility term that accounts for Fresnel effects and geometric occlusion integrated within the BRDF lobe. Please refer to Supp. Mat. for the connection to the rendering equation [21]. For point lights, we use a Dirac delta function multiplied by a light intensity as the incident light $\mathbf{L}(\boldsymbol{\omega}_i)$ for fast evaluation, but an SG parameterization with closed form integrals [74] would also be possible for directional area lights. Importantly, Eq. 8 can be efficiently evaluated with negligible overhead for all-frequency continuous illumination by prefiltering the environment maps [23, 47] as demonstrated in Fig. 1. This requires only a single mipmap texture look up per 3D Gaussian, which is a critical property for real-time relighting with all-frequency reflections.

Although the aforementioned formulation works well for surfaces, we use 3D Gaussians to also represent thin fiber-like structures such as hair strands, where Eq. 8 is incorrect if a viewer rotates along a tangent vector of each fiber [22]. To support specular reflection of both surfaces and fibers in a unified manner, we propose to learn a view-conditioned surface normal. While the learned normal can remain constant under view changes for surface regions, the normal can rotate along the tangent axis based on the view direction for fibers. This way, each 3D Gaussian can flexibly choose its underlying reflection behavior without requiring predefining it a priori. Our learnable view-dependent normal also supports the case where BRDF lobes are not exactly aligned with the surface normal [46] by adjusting the normal orientation.

Decoder. Similar to the geometric decoder, we decode radiance transfer parameters using a view-independent decoder \mathcal{D}_{ci} and a view-dependent decoder \mathcal{D}_{cv} as follows:

$$\begin{aligned} \{\mathbf{d}_k^c, \mathbf{d}_k^m, \sigma_k\}_{k=1}^M &= \mathcal{D}_{ci}(\mathbf{z}, \mathbf{e}_{\{l,r\}}; \Theta_{ci}), \\ \{\delta \mathbf{n}_k, v_k\}_{k=1}^M &= \mathcal{D}_{cv}(\mathbf{z}, \mathbf{e}_{\{l,r\}}, \boldsymbol{\omega}_o; \Theta_{cv}), \end{aligned} \quad (9)$$

where Θ_{ci} and Θ_{cv} are the learnable parameters of each decoder, and the normal residual $\delta \mathbf{n}_k$ is added to the barycentric interpolated coarse mesh normal $\hat{\mathbf{n}}_k$ to obtain the final normal \mathbf{n}_k as follows: $\mathbf{n}_k = (\hat{\mathbf{n}}_k + \delta \mathbf{n}_k) / \|\hat{\mathbf{n}}_k + \delta \mathbf{n}_k\|$. In practice, since \mathcal{D}_g in Eq. 3 and \mathcal{D}_{ci} in Eq. 9 take the same input and produce per-Gaussian values, we model them using a single decoder.

3.4. Relightable Explicit Eye Model

To enable better disentanglement and high-fidelity eye relighting, we use an explicit eye model proposed by Schwartz *et al.* [67] as the underlying geometric representation of eyes. In particular, we parameterize eyeballs as

the smooth blending of two spheres; one accounts for the eyeball and the other for the cornea. They are explicitly rotated based on a gaze direction. Each eyeball is parameterized by $\mathbf{E} = \{r_e, r_c, d, c_e\}$ with the radii of the eyeball r_e and cornea r_c , the offset d along the optical axis from the center of the eyeball to the center of the cornea sphere, and the center of the eyeball c_e relative to the head in a canonical space. We first optimize E following the optimization presented in Schwartz *et al.* [67], and then jointly refine it end-to-end with the other parameters of an avatar.

While we use the same geometric and appearance representations for eyeballs as the rest of the head, we observe that additional modification is required to enable high-fidelity eye relighting. Since the cornea exhibits mirror-like reflections, the discrete point lights of our capture system create reflections that span only a few pixels (often a single Bayer cell) and saturate the sensor. The remaining region has nearly zero contribution. Due to this highly discrete signal, 3D Gaussians quickly fall into poor local minima and fail to correctly model eye glints if we independently optimize the position and surface normal of each Gaussian. Therefore, we freeze the position of Gaussians on the surface of the eyeballs and fix their normals to be the surface normals of the eyeball mesh. In addition, the iris is observed only through the transparent cornea, with significant refraction. To support refraction with the same underlying appearance representation, we use a view-conditioned albedo for the eyes. This effectively allows the eye diffuse color to account for refraction based on the input viewpoint.

To incorporate these modifications, our geometry and appearance eye decoders for each eye are defined as follows:

$$\begin{aligned} \{\mathbf{R}_k, \mathbf{s}_k, o_k, \mathbf{d}_k^c, \mathbf{d}_k^m, \sigma_k\}_{k=1}^{M_e} &= \mathcal{D}_{ei}(\mathbf{e}, \mathbf{h}_p, \mathbf{h}_r; \Theta_{ei}), \\ \{\boldsymbol{\rho}_k, v_k\}_{k=1}^{M_e} &= \mathcal{D}_{ev}(\mathbf{e}, \mathbf{h}_p, \mathbf{h}_r, \boldsymbol{\omega}_o; \Theta_{ev}), \end{aligned} \quad (10)$$

where \mathcal{D}_{ei} and \mathcal{D}_{ev} are view-independent and view-conditioned eye decoders with parameters Θ_{ei} and Θ_{ev} respectively, and the relative head position $\mathbf{h}_p \in \mathbb{R}^3$ and rotation $\mathbf{h}_r \in SO(3)$ and the gaze \mathbf{e} are used to absorb tracking errors.

3.5. Training

Given multiview video data of a person illuminated with known point light patterns, we jointly optimize all trainable network parameters Θ , the static albedo $\boldsymbol{\rho}$, and the eyeball parameters $\mathbf{E}_{\{l,r\}}$ with the following loss function \mathcal{L} :

$$\mathcal{L} = \mathcal{L}_{\text{rec}} + \mathcal{L}_{\text{reg}} + \lambda_{\text{kl}} \mathcal{L}_{\text{kl}}, \quad (11)$$

where \mathcal{L}_{kl} is the KL-divergence loss on our encoder outputs. The reconstruction loss \mathcal{L}_{rec} consists of L1 and D-SSIM on the rendered image as in the original 3DGS paper [24, 75] as well as L2 loss on the coarse geometry V' as follows:

$$\mathcal{L}_{\text{rec}} = \lambda_{\text{l1}} \mathcal{L}_{\text{l1}} + \lambda_{\text{ssim}} \mathcal{L}_{\text{ssim}} + \lambda_{\text{geo}} \mathcal{L}_{\text{geo}}. \quad (12)$$

Table 1. **Comparison on held-out segments.** The top three techniques are highlighted in red, orange, and yellow, respectively.

		Geometry	Appearance	PSNR \uparrow	Metrics SSIM \uparrow	LPIPS \downarrow
A	Ours w/ EEM	EyeNeRF [33]		34.550	0.939	0.115
B		Ours		36.501	0.943	0.110
C	Ours	EyeNeRF [33]		35.110	0.938	0.113
D		Linear [87]		33.831	0.936	0.184
E		Ours		36.529	0.943	0.111
F	MVP [41]	EyeNeRF [33]		27.594	0.922	0.151
G		Linear [87]		36.294	0.942	0.140
H		Ours		35.789	0.943	0.134

The regularization loss is defined as:

$$\mathcal{L}_{\text{reg}} = \lambda_s \mathcal{L}_s + \lambda_{c-} \mathcal{L}_{c-} + \lambda_{es} \mathcal{L}_{es} + \lambda_{ev} \mathcal{L}_{ev} + \lambda_{eo} \mathcal{L}_{eo}. \quad (13)$$

The scale regularization term \mathcal{L}_s encourages the Gaussian scale $\{s_k\}$ to stay within a reasonable range as follows:

$$\mathcal{L}_s = \text{mean}(l_s), l_s = \begin{cases} 1/\max(s, 10^{-7}) & \text{if } s < 0.1 \\ (s - 10.0)^2 & \text{if } s > 10.0 \\ 0 & \text{otherwise,} \end{cases} \quad (14)$$

where s denotes the scale value of each axis in each Gaussian, and $\text{mean}(\cdot)$ is the average operation across all dimensions. The negative color loss \mathcal{L}_{c-} penalizes negative color in the diffuse term as SH can yield negative values:

$$\mathcal{L}_{c-} = \text{mean}(l_{c-}), l_{c-} = \min(\mathbf{c}_k^{\text{diffuse}}, 0)^2. \quad (15)$$

Also, three regularization terms are used to prevent eye Gaussians from becoming transparent as follows:

$$\begin{aligned} \mathcal{L}_{es} &= \text{mean}(l_{es}), l_{es} = \max(s - 0.1, 0)^2, \\ \mathcal{L}_{eo} &= \text{mean}(l_{eo}), l_{eo} = (1 - o_k)^2, \\ \mathcal{L}_{ev} &= \text{mean}(l_{ev}), l_{ev} = (1 - v_k)^2. \end{aligned} \quad (16)$$

The relative weights are $\lambda_{\text{geo}} = \lambda_{11} = 10$, $\lambda_{\text{ssim}} = 0.2$, $\lambda_s = \lambda_{c-} = \lambda_{es} = 1.0 \times 10^{-2}$, $\lambda_{eo} = \lambda_{ev} = 1.0 \times 10^{-4}$, and $\lambda_{\text{kl}} = 2.0 \times 10^{-3}$. We use the Adam optimizer [25] with a learning rate of 0.0005. We train our model on 4 NVIDIA A100 GPUs with a batch size of 16 for 200k iterations. Please refer to Supp. Mat. for network architecture.

4. Experiments

Evaluation Protocol. We selected three subjects for quantitative evaluations and three more subjects for qualitative results with diverse races, genders, and hairstyles. Our evaluation included around 9,000 conversational expression frames and about 100 *disgust* expression frames not seen during training. We also exclude 10 unique frontally-biased light patterns entirely from the training. This corresponds to approximately 1800 out of 144,000 frames. We report PSNR, SSIM, and LPIPS [91] on images masked with the face region to avoid influence from the background.

Table 2. **Comparison on held-out lights.** The top three techniques are highlighted in red, orange, and yellow, respectively.

		Geometry	Appearance	PSNR \uparrow	Metrics SSIM \uparrow	LPIPS \downarrow
A	Ours w/ EEM	EyeNeRF [33]		30.7976	0.828	0.162
B		Ours		34.042	0.858	0.148
C	Ours	EyeNeRF [33]		30.836	0.815	0.163
D		Linear [87]		32.829	0.870	0.202
E		Ours		33.845	0.831	0.148
F	MVP [41]	EyeNeRF [33]		28.030	0.812	0.210
G		Linear [87]		33.444	0.726	0.192
H		Ours		33.778	0.877	0.168

4.1. Qualitative Results

Fig. 1 shows that our reconstructed avatars generalize to novel views, expressions, and illuminations including point lights and high-resolution environment maps. Notice the mirror-like reflections in the eyes that faithfully represent the environment without losing high-frequency details. As our model is drivable and supports real-time relighting, real-time driving from a headset is also possible [77].

While this is not our primary goal, as a bi-product, our approach estimates intrinsic properties of reflectance including albedo, geometry, surface normal, multi-bounce scattering, and specular components in a self-supervised manner. As shown in Fig. 3, our method enables 3D consistent and high-fidelity intrinsics decomposition.

4.2. Discussion

Geometric Representation. We evaluate the geometry component by comparing three variations: our proposed method, our method excluding the explicit eye model (EEM) [67], and voxel-based primitives [41]. For fair comparison, we use the same appearance model and only change the geometric representation (Tab. 1 and Tab. 2 B, D, H). Fig. 4 clearly demonstrates that our geometry based on 3D Gaussians can better model skin details and hair strands than MVP. Further, our full model, when combined with an EEM, achieves convincing eye glints.

Appearance Representation. For appearance representation, we compare our appearance model with existing relightable appearance representations that support rendering with environment maps in real-time. The model presented by Yang *et al.* [87] is a linear neural network that explicitly retains the linearity of light transport (denoted as Linear), demonstrating superior performance than a previous state-of-the-art method [6]. For this reason, we omit the comparison with [6]. To evaluate the effectiveness of our specular reflection model, we also replace our specular component with view-dependent spherical harmonics proposed by EyeNeRF [33]. Tab. 1 and Tab. 2 C, D, E show that our appearance representation outperforms existing appearance models in most of the metrics. As shown in Fig. 5, while




Figure 3. **Intrinsic decomposition.** The full render (a) is produced by addition of a diffuse (b) and a specular component (c) (intensity multiplied by 2 for clarity). The diffuse component is obtained by multiplying a learned albedo (d) with shading computed by SH-based radiance transfer (e). The specular lobes direction is computed using a per-Gaussian normal (f).

the linear model produces correct overall color, the relighting result is blurry and lacks high-frequency details. This is primarily limited by the bottleneck lighting representation. The view-dependent spherical harmonics in EyeNeRF shows more detailed reflections, but its expressiveness is limited due to the use of spherical harmonics for specularity. Additionally, we observe that view-dependent spherical harmonics are more prone to overfitting, resulting in flickering artifacts in animation. Please refer to our supplemental video for more details. In contrast, our approach based on spherical Gaussians is not band-limited and thus achieves high-frequency reflections.

Impact of Data Quality. We discover that our approach works even for more relaxed setups such as using a generic template mesh as the base geometry regardless of expressions, and ablating up to 90% of cameras or 95% of light patterns from the training data. While we recommend to use our setup to achieve the best quality, this indicates that the proposed method can be applied to much more modest setups. Please refer to Supp. Mat. for the experiments.

5. Conclusion

We presented Relightable Gaussian Codec Avatars, a novel appearance and geometric representation for relightable 3D head avatars that supports real-time rendering. Our experiments show that high-fidelity relighting of hair, skin, and eyes in all-frequency illuminations is now possible in real-time with the proposed radiance transfer basis composed of spherical harmonics and spherical Gaussians. We have also shown that our choice of the geometric representa-



Figure 4. **Geometric representation comparison.** Compared to a held out frame, (a), our Gaussian splatting decoded geometry (b,c) shows improved resolution over MVP [41] (d), especially in fine details like eyelashes and pores. The explicit eyeball model (b) additionally improves realism in eye glints. All methods use the appearance model described in Sec. 3.3.

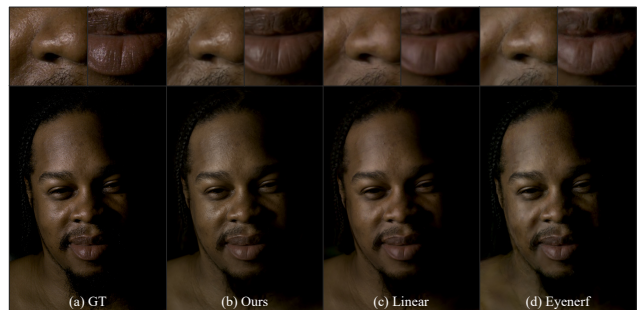


Figure 5. **Appearance representation comparison.** Compared to a held out frame (a), our appearance model (Sec. 3.3) shows sharper pore-level specularities than methods using only a linear neural network [87] or the spherical harmonics-only method “Eyenerf” [33]. All methods use the geometric representation described in Sec. 3.2 (without explicit eyeballs.)

tion based on 3D Gaussian Splatting is critical for strand-accurate hair reconstruction and relighting. Our approach achieves a significant quality improvement in comparison to existing real-time-renderable geometry and appearance models, both qualitatively and quantitatively.

Limitations and Future Work. The current approach requires a coarse mesh and gaze tracking as a preprocessing step, which may be sensitive to tracking failures. Similar to [87], end-to-end learning together with topology consistent tracking [8, 37, 38] is an interesting future work direction to enable scalable training. Extending our approach to in-the-wild inputs also remains a challenge due to the lack of precisely known illumination information. Lastly, rendering a large number of our Gaussian avatars would be difficult, as the relighting operation is performed per individual 3D Gaussian, and scales linearly with the number of avatars. Offloading computation in a per-pixel fragment shader, similar to [44], is also an exciting future research.

References

- [1] Rameen Abdal, Yipeng Qin, and Peter Wonka. Image2stylegan: How to embed images into the stylegan latent space? In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4432–4441, 2019. 5
- [2] Thabo Beeler, Bernd Bickel, Paul Beardsley, Bob Sumner, and Markus Gross. High-quality single-shot capture of facial geometry. In *ACM SIGGRAPH 2010 papers*, pages 1–9, 2010. 2
- [3] Thabo Beeler, Bernd Bickel, Gioacchino Noris, Paul Beardsley, Steve Marschner, Robert W Sumner, and Markus Gross. Coupled 3d reconstruction of sparse facial hair and skin. *ACM Transactions on Graphics (ToG)*, 31(4):1–10, 2012. 2
- [4] Thabo Beeler, Fabian Hahn, Derek Bradley, Bernd Bickel, Paul A Beardsley, Craig Gotsman, Robert W Sumner, and Markus H Gross. High-quality passive facial performance capture using anchor frames. *ACM Trans. Graph.*, 30(4):75, 2011. 2
- [5] Pascal Bérard, Derek Bradley, Markus Gross, and Thabo Beeler. Lightweight eye capture using a parametric model. *ACM Transactions on Graphics (TOG)*, 35(4):1–12, 2016. 2
- [6] Sai Bi, Stephen Lombardi, Shunsuke Saito, Tomas Simon, Shih-En Wei, Kevyn Mcphail, Ravi Ramamoorthi, Yaser Sheikh, and Jason Saragih. Deep relightable appearance models for animatable faces. *ACM Transactions on Graphics (TOG)*, 40(4):1–15, 2021. 3, 4, 7
- [7] Sai Bi, Zexiang Xu, Kalyan Sunkavalli, Miloš Hašan, Yannick Hold-Geoffroy, David Kriegman, and Ravi Ramamoorthi. Deep reflectance volumes: Relightable reconstructions from multi-view photometric images. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*, pages 294–311. Springer, 2020. 2, 3
- [8] Timo Bolkart, Tianye Li, and Michael J Black. Instant multi-view head capture through learnable registration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 768–779, 2023. 8
- [9] Derek Bradley, Wolfgang Heidrich, Tiberiu Popa, and Alla Sheffer. High resolution passive facial performance capture. In *ACM SIGGRAPH 2010 papers*, pages 1–10, 2010. 2
- [10] Paul Debevec, Tim Hawkins, Chris Tchou, Haarm-Pieter Duiker, Westley Sarokin, and Mark Sagar. Acquiring the reflectance field of a human face. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 145–156, 2000. 2, 3
- [11] Boyang Deng, Yifan Wang, and Gordon Wetzstein. Luminan: Unconditional generation of relightable 3d human faces. *arXiv preprint arXiv:2304.13153*, 2023. 3
- [12] Yasutaka Furukawa and Jean Ponce. Accurate, dense, and robust multiview stereopsis. *IEEE transactions on pattern analysis and machine intelligence*, 32(8):1362–1376, 2009. 2
- [13] Duan Gao, Guojun Chen, Yue Dong, Pieter Peers, Kun Xu, and Xin Tong. Deferred neural lighting: free-viewpoint relighting from unstructured photographs. *ACM Transactions on Graphics (TOG)*, 39(6):1–15, 2020. 3
- [14] Stephan J Garbin, Marek Kowalski, Virginia Estellers, Stanislaw Szymonowicz, Shideh Rezaeifar, Jingjing Shen, Matthew Johnson, and Julien Valentin. Voltmorph: Real-time, controllable and generalisable animation of volumetric representations. *arXiv preprint arXiv:2208.00949*, 2022. 3
- [15] Pablo Garrido, Michael Zollhöfer, Chenglei Wu, Derek Bradley, Patrick Pérez, Thabo Beeler, and Christian Theobalt. Corrective 3d reconstruction of lips from monocular video. *ACM Trans. Graph.*, 35(6):219–1, 2016. 2
- [16] Abhijeet Ghosh, Graham Fyffe, Borom Tunwattanapong, Jay Busch, Xueming Yu, and Paul Debevec. Multiview face capture using polarized spherical gradient illumination. *ACM Transactions on Graphics (TOG)*, 30(6):1–10, 2011. 2, 3
- [17] Paul Green, Jan Kautz, Wojciech Matusik, and Frédo Durand. View-dependent precomputed light transport using nonlinear gaussian function approximations. In *Proceedings of the 2006 symposium on Interactive 3D graphics and games*, pages 7–14, 2006. 3
- [18] Kaiwen Guo, Peter Lincoln, Philip Davidson, Jay Busch, Xueming Yu, Matt Whalen, Geoff Harvey, Sergio Orts-Escolano, Rohit Pandey, Jason Dourgarian, et al. The relightables: Volumetric performance capture of humans with realistic relighting. *ACM Transactions on Graphics (ToG)*, 38(6):1–19, 2019. 3
- [19] Liwen Hu, Chongyang Ma, Linjie Luo, and Hao Li. Robust hair capture using simulated examples. *ACM Transactions on Graphics (TOG)*, 33(4):1–10, 2014. 2
- [20] Shun Iwase, Shunsuke Saito, Tomas Simon, Stephen Lombardi, Timur Bagautdinov, Rohan Joshi, Fabian Prada, Takaaki Shiratori, Yaser Sheikh, and Jason Saragih. Relightablehands: Efficient neural relighting of articulated hand models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16663–16673, 2023. 2, 3
- [21] James T Kajiya. The rendering equation. In *Proceedings of the 13th annual conference on Computer graphics and interactive techniques*, pages 143–150, 1986. 6
- [22] James T Kajiya and Timothy L Kay. Rendering fur with three dimensional textures. *ACM Siggraph Computer Graphics*, 23(3):271–280, 1989. 6
- [23] Jan Kautz, Pere-Pau Vázquez, Wolfgang Heidrich, and Hans-Peter Seidel. A unified approach to prefiltered environment maps. In *Rendering Techniques 2000: Proceedings of the Eurographics Workshop in Brno, Czech Republic, June 26–28, 2000 11*, pages 185–196. Springer, 2000. 6
- [24] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics (ToG)*, 42(4):1–14, 2023. 2, 3, 4, 5, 6
- [25] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 7
- [26] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 5
- [27] Tobias Kirschstein, Shenhan Qian, Simon Giebenhain, Tim Walter, and Matthias Nießner. Nersemble: Multi-view radiance field reconstruction of human heads. *arXiv preprint arXiv:2305.03027*, 2023. 3
- [28] Georgios Kopanas, Thomas Leimkühler, Gilles Rainer, Clément Jambon, and George Drettakis. Neural point catacaustics for novel-view synthesis of reflections. *ACM Trans-*

- actions on Graphics (TOG)*, 41(6):1–15, 2022. 4
- [29] Georgios Kopanas, Julien Philip, Thomas Leimkühler, and George Drettakis. Point-based neural rendering with per-view optimization. In *Computer Graphics Forum*, volume 40, pages 29–43. Wiley Online Library, 2021. 4
- [30] Mathieu Lamarre, John P Lewis, and Etienne Danvoye. Face stabilization by mode pursuit for avatar construction. In *2018 International Conference on Image and Vision Computing New Zealand (IVCNZ)*, pages 1–6. IEEE, 2018. 4
- [31] Alexandros Lattas, Stylianos Moschoglou, Baris Gecer, Stylianos Ploumpis, Vasileios Triantafyllou, Abhijeet Ghosh, and Stefanos Zafeiriou. Avatarme: Realistically renderable 3d facial reconstruction” in-the-wild”. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 760–769, 2020. 3
- [32] Alexandros Lattas, Stylianos Moschoglou, Stylianos Ploumpis, Baris Gecer, Abhijeet Ghosh, and Stefanos Zafeiriou. Avatarme++: Facial shape and brdf inference with photorealistic rendering-aware gans. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(12):9269–9284, 2021. 3
- [33] Gengyan Li, Abhimitra Meka, Franziska Mueller, Marcel C Buehler, Otmar Hilliges, and Thabo Beeler. Eyenerf: a hybrid representation for photorealistic synthesis, animation and relighting of human eyes. *ACM Transactions on Graphics (TOG)*, 41(4):1–16, 2022. 1, 2, 3, 7, 8
- [34] Jiaman Li, Zhengfei Kuang, Yajie Zhao, Mingming He, Kalle Bladin, and Hao Li. Dynamic facial asset and rig generation from a single scan. *ACM Transactions on Graphics (TOG)*, 39:1 – 18, 2020. 3
- [35] Junxuan Li, Shunsuke Saito, Tomas Simon, Stephen Lombardi, Hongdong Li, and Jason Saragih. Megane: Morphable eyeglass and avatar network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12769–12779, 2023. 2, 3, 4
- [36] Ruilong Li, Kalle Bladin, Yajie Zhao, Chinmay Chinara, Owen Ingraham, Pengda Xiang, Xinglei Ren, Pratusha Bhuvana Prasad, Bipin Kishore, Jun Xing, and Hao Li. Learning formation of physically-based face attributes. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3407–3416, 2020. 3
- [37] Tianye Li, Shichen Liu, Timo Bolkart, Jiayi Liu, Hao Li, and Yajie Zhao. Topologically consistent multi-view face inference using volumetric sampling. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3824–3834, 2021. 8
- [38] Shichen Liu, Yunxuan Cai, Haiwei Chen, Yichao Zhou, and Yajie Zhao. Rapid face asset acquisition with recurrent feature alignment. *ACM Transactions on Graphics (TOG)*, 41(6):1–17, 2022. 8
- [39] Stephen Lombardi, Jason Saragih, Tomas Simon, and Yaser Sheikh. Deep appearance models for face rendering. *ACM Transactions on Graphics (ToG)*, 37(4):1–13, 2018. 1, 2, 5
- [40] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: Learning dynamic renderable volumes from images. *arXiv preprint arXiv:1906.07751*, 2019. 3
- [41] Stephen Lombardi, Tomas Simon, Gabriel Schwartz, Michael Zollhoefer, Yaser Sheikh, and Jason Saragih. Mix-
ture of volumetric primitives for efficient neural rendering. *ACM Transactions on Graphics (ToG)*, 40(4):1–13, 2021. 3, 5, 7, 8
- [42] Jonathon Luiten, Georgios Kopanas, Bastian Leibe, and Deva Ramanan. Dynamic 3d gaussians: Tracking by persistent dynamic view synthesis. *arXiv preprint arXiv:2308.09713*, 2023. 3
- [43] Linjie Luo, Hao Li, and Szymon Rusinkiewicz. Structure-aware hair capture. *ACM Transactions on Graphics (TOG)*, 32(4):1–12, 2013. 2
- [44] Shugao Ma, Tomas Simon, Jason Saragih, Dawei Wang, Yuecheng Li, Fernando De La Torre, and Yaser Sheikh. Pixel codec avatars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 64–73, 2021. 2, 8
- [45] Wan-Chun Ma, Tim Hawkins, Pieter Peers, Charles-Felix Chabert, Malte Weiss, Paul E Debevec, et al. Rapid acquisition of specular and diffuse normal maps from polarized spherical gradient illumination. *Rendering Techniques*, 2007(9):10, 2007. 2, 3
- [46] Stephen R Marschner, Henrik Wann Jensen, Mike Cammarano, Steve Worley, and Pat Hanrahan. Light scattering from human hair fibers. *ACM Transactions on Graphics (TOG)*, 22(3):780–791, 2003. 1, 6
- [47] David K McAllister, Anselmo Lastra, and Wolfgang Heidrich. Efficient rendering of spatial bi-directional reflectance distribution functions. In *Proceedings of the ACM SIGGRAPH/EUROGRAPHICS conference on Graphics hardware*, pages 79–88, 2002. 6
- [48] Abhimitra Meka, Christian Haene, Rohit Pandey, Michael Zollhöfer, Sean Fanello, Graham Fyffe, Adarsh Kowdle, Xueming Yu, Jay Busch, Jason Dourgarian, et al. Deep reflectance fields: high-quality facial reflectance field inference from color gradient illumination. *ACM Transactions on Graphics (TOG)*, 38(4):1–12, 2019. 3
- [49] Abhimitra Meka, Rohit Pandey, Christian Haene, Sergio Orts-Escolano, Peter Barnum, Philip David-Son, Daniel Erickson, Yinda Zhang, Jonathan Taylor, Sofien Bouaziz, et al. Deep relightable textures: volumetric performance capture with neural rendering. *ACM Transactions on Graphics (TOG)*, 39(6):1–21, 2020. 3
- [50] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. 3
- [51] Erick Miller and Dmitry Pinskiy. Realistic eye motion using procedural geometric methods. In *SIGGRAPH 2009: Talks*, pages 1–1. 2009. 2
- [52] Jacob Munkberg, Jon Hasselgren, Tianchang Shen, Jun Gao, Wenzheng Chen, Alex Evans, Thomas Müller, and Sanja Fidler. Extracting triangular 3d models, materials, and lighting from images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8280–8290, 2022. 3
- [53] Koki Nagano, Graham Fyffe, Oleg Alexander, Jernej Barbic, Hao Li, Abhijeet Ghosh, and Paul E Debevec. Skin microstructure deformation with displacement map convolution. *ACM Trans. Graph.*, 34(4):109–1, 2015. 1

- [54] Giljoo Nam, Joo Ho Lee, Diego Gutierrez, and Min H Kim. Practical svbrdf acquisition of 3d objects with unstructured flash photography. *ACM Transactions on Graphics (TOG)*, 37(6):1–12, 2018. [3](#)
- [55] Giljoo Nam, Chenglei Wu, Min H Kim, and Yaser Sheikh. Strand-accurate multi-view hair capture. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 155–164, 2019. [2](#)
- [56] Ren Ng, Ravi Ramamoorthi, and Pat Hanrahan. All-frequency shadows using non-linear wavelet lighting approximation. In *ACM SIGGRAPH 2003 Papers*, pages 376–381. 2003. [3](#)
- [57] Sergio Orts-Escolano, Christoph Rhemann, Sean Fanello, Wayne Chang, Adarsh Kowdle, Yury Degtyarev, David Kim, Philip L Davidson, Sameh Khamis, Mingsong Dou, et al. Holoportation: Virtual 3d teleportation in real-time. In *Proceedings of the 29th annual symposium on user interface software and technology*, pages 741–754, 2016. [1](#)
- [58] Rohit Pandey, Sergio Orts Escolano, Chloe Legendre, Christian Haene, Sofien Bouaziz, Christoph Rhemann, Paul Debevec, and Sean Fanello. Total relighting: learning to relight portraits for background replacement. *ACM Transactions on Graphics (TOG)*, 40(4):1–21, 2021. [3](#)
- [59] Foivos Paraperas Papantoniou, Alexandros Lattas, Stylianos Moschoglou, and Stefanos Zafeiriou. Relightify: Relightable 3d faces from a single image via diffusion models. *arXiv preprint arXiv:2305.06077*, 2023. [3](#)
- [60] Sylvain Paris, Will Chang, Oleg I Kozhushnyan, Wojciech Jarosz, Wojciech Matusik, Matthias Zwicker, and Frédo Durand. Hair photobooth: geometric and photometric acquisition of real hairstyles. *ACM Trans. Graph.*, 27(3):30, 2008. [2](#)
- [61] Frederic I Parke and Keith Waters. *Computer facial animation*. CRC press, 2008. [2](#)
- [62] Pieter Peers, Naoki Tamura, Wojciech Matusik, and Paul Debevec. Post-production facial performance relighting using reflectance transfer. *ACM Transactions on Graphics (TOG)*, 26(3):52–es, 2007. [3](#)
- [63] Gilles Rainer, Adrien Bousseau, Tobias Ritschel, and George Drettakis. Neural precomputed radiance transfer. In *Computer Graphics Forum*, volume 41, pages 365–378. Wiley Online Library, 2022. [3](#)
- [64] Ravi Ramamoorthi and Pat Hanrahan. An efficient representation for irradiance environment maps. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 497–500, 2001. [5](#)
- [65] Anurag Ranjan, Kwang Moo Yi, Jen-Hao Rick Chang, and Oncel Tuzel. Facelit: Neural 3d relightable faces. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8619–8628, 2023. [3](#)
- [66] Kripasindhu Sarkar, Marcel C. Buehler, Gengyan Li, Daoye Wang, Delio Vicini, Jérémy Riviere, Yinda Zhang, Sergio Orts-Escolano, Paulo Gotardo, Thabo Beeler, and Abhimitra Meka. Litnerf: Intrinsic radiance decomposition for high-quality view synthesis and relighting of faces. In *ACM SIGGRAPH Asia 2023*. 2023. [2](#), [3](#)
- [67] Gabriel Schwartz, Shih-En Wei, Te-Li Wang, Stephen Lombardi, Tomas Simon, Jason Saragih, and Yaser Sheikh. The eyes have it: An integrated eye and face model for photo-realistic facial animation. *ACM Transactions on Graphics (TOG)*, 39(4):91–1, 2020. [1](#), [2](#), [4](#), [6](#), [7](#)
- [68] Mike Seymour, Chris Evans, and Kim Libreri. Meet mike: epic avatars. In *ACM SIGGRAPH 2017 VR Village*, pages 1–2. 2017. [2](#)
- [69] Peter-Pike Sloan, Jan Kautz, and John Snyder. Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. *ACM Trans. Graph.*, 21(3):527–536, jul 2002. [2](#), [3](#), [5](#)
- [70] Tiancheng Sun, Jonathan T Barron, Yun-Ta Tsai, Zexiang Xu, Xueming Yu, Graham Fyffe, Christoph Rhemann, Jay Busch, Paul Debevec, and Ravi Ramamoorthi. Single image portrait relighting. *ACM Transactions on Graphics (TOG)*, 38(4):1–12, 2019. [3](#)
- [71] Ayush Tewari, Mohamed Elgharib, Gaurav Bharaj, Florian Bernard, Hans-Peter Seidel, Patrick Pérez, Michael Zollhofer, and Christian Theobalt. Stylerig: Rigging stylegan for 3d control over portrait images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6142–6151, 2020. [3](#)
- [72] Justus Thies, Michael Zollhöfer, and Matthias Nießner. Deferred neural rendering: Image synthesis using neural textures. *ACM Transactions on Graphics (TOG)*, 38(4):1–12, 2019. [3](#)
- [73] Yu-Ting Tsai and Zen-Chung Shih. All-frequency precomputed radiance transfer using spherical radial basis functions and clustered tensor approximation. *ACM Transactions on graphics (TOG)*, 25(3):967–976, 2006. [3](#)
- [74] Jiaping Wang, Peiran Ren, Minmin Gong, John Snyder, and Baining Guo. All-frequency rendering of dynamic, spatially-varying reflectance. In *ACM SIGGRAPH Asia 2009 papers*, pages 1–10. 2009. [2](#), [3](#), [5](#), [6](#)
- [75] Zhou Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004. [6](#)
- [76] Zhibo Wang, Xin Yu, Ming Lu, Quan Wang, Chen Qian, and Feng Xu. Single image portrait relighting via explicit multiple reflectance channel modeling. *ACM Transactions on Graphics (TOG)*, 39(6):1–13, 2020. [3](#)
- [77] Shih-En Wei, Jason Saragih, Tomas Simon, Adam W Harley, Stephen Lombardi, Michal Perdoch, Alexander Hypes, Dawei Wang, Hernan Badino, and Yaser Sheikh. Vr facial animation via multiview image translation. *ACM Transactions on Graphics (TOG)*, 38(4):1–16, 2019. [5](#), [7](#)
- [78] Tim Weyrich, Wojciech Matusik, Hanspeter Pfister, Bernd Bickel, Craig Donner, Chien Tu, Janet McAndless, Jinho Lee, Addy Ngan, Henrik Wann Jensen, et al. Analysis of human faces using a measurement-based skin reflectance model. *ACM Transactions on Graphics (ToG)*, 25(3):1013–1024, 2006. [1](#), [2](#)
- [79] Chenglei Wu, Derek Bradley, Pablo Garrido, Michael Zollhöfer, Christian Theobalt, Markus H Gross, and Thabo Beeler. Model-based teeth reconstruction. *ACM Trans. Graph.*, 35(6):220–1, 2016. [2](#)
- [80] Guanjun Wu, Taoran Yi, Jiemin Fang, Lingxi Xie, Xiaopeng Zhang, Wei Wei, Wenyu Liu, Qi Tian, and Xinggang Wang. 4d gaussian splatting for real-time dynamic scene rendering. *arXiv preprint arXiv:2310.08528*, 2023. [3](#)

- [81] Kun Xu, Wei-Lun Sun, Zhao Dong, Dan-Yong Zhao, Run-Dong Wu, and Shi-Min Hu. Anisotropic spherical gaussians. *ACM Transactions on Graphics (TOG)*, 32(6):1–11, 2013. 3
- [82] Yuelang Xu, Hongwen Zhang, Lizhen Wang, Xiaochen Zhao, Han Huang, Guojun Qi, and Yebin Liu. Latentavatar: Learning latent expression code for expressive neural head avatar. *arXiv preprint arXiv:2305.01190*, 2023. 5
- [83] Yingyan Xu, Gaspard Zoss, Prashanth Chandran, Markus Gross, Derek Bradley, and Paulo Gotardo. Renerf: Relightable neural radiance fields with nearfield lighting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22581–22591, 2023. 3
- [84] Zexiang Xu, Kalyan Sunkavalli, Sunil Hadap, and Ravi Ramamoorthi. Deep image-based relighting from optimal sparse samples. *ACM Transactions on Graphics (TOG)*, 37(4):1–13, 2018. 3
- [85] Zilin Xu, Zheng Zeng, Lifan Wu, Lu Wang, and Ling-Qi Yan. Lightweight neural basis functions for all-frequency shading. In *SIGGRAPH Asia 2022 Conference Papers*, pages 1–9, 2022. 3
- [86] Shugo Yamaguchi, Shunsuke Saito, Koki Nagano, Yajie Zhao, Weikai Chen, Kyle Olszewski, Shigeo Morishima, and Hao Li. High-fidelity facial reflectance and geometry inference from an unconstrained image. *ACM Transactions on Graphics (TOG)*, 37(4):1–14, 2018. 3
- [87] Haotian Yang, Mingwu Zheng, Wanquan Feng, Haibin Huang, Yu-Kun Lai, Pengfei Wan, Zhongyuan Wang, and Chongyang Ma. Towards practical capture of high-fidelity relightable avatars. In *SIGGRAPH Asia 2023 Conference Proceedings*, 2023. 2, 3, 7, 8
- [88] Yu-Ying Yeh, Koki Nagano, Sameh Khamis, Jan Kautz, Ming-Yu Liu, and Ting-Chun Wang. Learning to relight portrait images via a virtual light stage and synthetic-to-real adaptation. *ACM Transactions on Graphics (TOG)*, 41(6):1–21, 2022. 3
- [89] Kai Zhang, Fujun Luan, Qianqian Wang, Kavita Bala, and Noah Snavely. Physg: Inverse rendering with spherical gaussians for physics-based material editing and relighting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5453–5462, 2021. 2, 3, 5
- [90] Li Zhang, Noah Snavely, Brian Curless, and Steven M Seitz. Spacetime faces: high resolution capture for modeling and animation. In *ACM SIGGRAPH 2004 Papers*, pages 548–558. 2004. 2
- [91] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. 7
- [92] Xiuming Zhang, Sean Fanello, Yun-Ta Tsai, Tiancheng Sun, Tianfan Xue, Rohit Pandey, Sergio Orts-Escolano, Philip Davidson, Christoph Rhemann, Paul Debevec, et al. Neural light transport for relighting and view synthesis. *ACM Transactions on Graphics (TOG)*, 40(1):1–17, 2021. 3
- [93] Xiuming Zhang, Pratul P Srinivasan, Boyang Deng, Paul Debevec, William T Freeman, and Jonathan T Barron. Nerfactor: Neural factorization of shape and reflectance under an unknown illumination. *ACM Transactions on Graphics (TOG)*, 40(6):1–18, 2021. 3
- [94] Xiaochen Zhao, Lizhen Wang, Jingxiang Sun, Hongwen Zhang, Jinli Suo, and Yebin Liu. Havatar: High-fidelity head avatar via facial model conditioned neural radiance field. *ACM Transactions on Graphics*, 2023. 2, 3
- [95] Yufeng Zheng, Wang Yifan, Gordon Wetzstein, Michael J Black, and Otmar Hilliges. Pointavatar: Deformable point-based head avatars from videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21057–21067, 2023. 3
- [96] Wojciech Zielonka, Timo Bolkart, and Justus Thies. Instant volumetric head avatars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4574–4584, 2023. 2, 3
- [97] Matthias Zwicker, Hanspeter Pfister, Jeroen Van Baar, and Markus Gross. Ewa splatting. *IEEE Transactions on Visualization and Computer Graphics*, 8(3):223–238, 2002. 4, 5