
BecomingLit: Relightable Gaussian Avatars with Hybrid Neural Shading

Jonathan Schmidt Simon Giebenhain Matthias Nießner

Technical University of Munich

jonathansch.github.io/becominglit




Figure 1: **BecomingLit**: Our approach effectively reconstructs detailed human head avatars that can be animated from videos and relighted in real-time using our hybrid neural shading approach. Besides our method, we introduce a new high-quality, multi-view OLAT dataset of faces.

Abstract

We introduce *BecomingLit*, a novel method for reconstructing relightable, high-resolution head avatars that can be rendered from novel viewpoints at interactive rates. Therefore, we propose a new low-cost light stage capture setup, tailored specifically towards capturing faces. Using this setup, we collect a novel dataset consisting of diverse multi-view sequences of numerous subjects under varying illumination conditions and facial expressions. By leveraging our new dataset, we introduce a new relightable avatar representation based on 3D Gaussian primitives that we animate with a parametric head model and an expression-dependent dynamics module. We propose a new hybrid neural shading approach, combining a neural diffuse BRDF with an analytical specular term. Our method reconstructs disentangled materials from our dynamic light stage recordings and enables all-frequency relighting of our avatars with both point lights and environment maps. In addition, our avatars can easily be animated and controlled from monocular videos. We validate our approach in extensive experiments on our dataset, where we consistently outperform existing state-of-the-art methods in relighting and reenactment by a significant margin.

1 Introduction

The creation of photorealistic, relightable 3D head avatars from real-world data is a core problem of computer vision with applications across a wide range of graphics tasks, such as cinematography, virtual reality, or the metaverse in general. Traditionally, this requires professional, room-scale capture setups that only a handful of institutions can afford [34, 4, 23, 24], as the joint estimation of geometry, intrinsic material parameters, and lighting is an extremely under-constrained problem.

At the same time, with the progressing growth of virtual reality applications at the consumer level, creating photorealistic avatars is becoming more important than ever. While there has been immense progress over the recent years in terms of geometric representations [15, 30, 10, 18, 25], visual quality and rendering speed thanks to the availability of custom datasets [16], most 3D avatars do not have a disentangled representation of the material properties and bake the radiance properties of the training environment into the avatar, which makes relighting impossible. As a result, placing the avatar in a novel virtual environment dramatically lowers the visual quality of the renderings. In comparison, research on relightable avatars is scarce. One of the major reasons for this is the lack of publicly available and free-to-use datasets that come with controlled light captures in order to broadly study the reconstruction of facial appearance.

To this end, we introduce an OLAT dataset and propose *BecomingLit*, a novel approach to reconstruct photorealistic, relightable head avatars from short multi-view light stage sequences. We represent the head with expression-dependent Gaussian primitives and model the complex reflection behavior of faces by learning a hybrid neural BRDF. Thanks to our efficient parameterization and regularization, our method requires a capture setup that is an order of magnitude more economical compared to previous work, and outperforms state-of-the-art methods in self-reenactment under novel illuminations. To address the lack of data, we introduce a new multi-view video dataset of different participants in a light stage setting, which we will make publicly available for research purposes. Overall, our contributions are two-fold:

- We introduce a novel, publicly available dataset, combining high-resolution, high-framerate, multi-view recordings of different subjects in a calibrated light stage setting.
- We propose a relightable, photorealistic avatar representation based on 3D Gaussian primitives and hybrid neural shading, which can be relighted and rendered from novel viewpoints in real-time and animated from monocular videos.

2 Related Work

Human Head Modeling addresses the problem of representing and modeling the geometry and appearance of human heads. Traditional methods learn morphable models from head scans via PCA [5, 28, 17]. While being strong in generalization, PCA-based 3DMMs have a limited expressiveness and can fail to represent fine geometric details such as skin wrinkles or hair. As an alternative, [4, 20] propose to learn the geometry and appearance space with autoencoders. More recently, volumetric approaches based on NeRF [25], represent heads with more detailed appearance, despite not requiring explicit input geometry [16, 54, 18, 57, 8, 32]. Another line of work uses 3D Gaussian primitives [15] to model human heads [10, 19, 47], some of them in combination with a 3D morphable model [30, 34].

Facial Appearance Capture. Capturing the appearance of human faces is a long-standing problem in computer vision. Debevec et al. [7] introduced the light stage and demonstrated how the reflectance field of a human face can be reconstructed from one-light-at-a-time captures, and relighted using image-based rendering [7, 43]. Subsequent work leveraged polarized light to decompose specular and diffuse reflectance [21, 9, 11, 33, 2]. [35, 46] approach the intrinsic decomposition problem with radiance fields [25]. [4, 49, 49] propose a learnable, data-driven appearance model that learns avatar relighting in an end-to-end manner with a neural lighting model. In contrast, [34, 40] propose to learn radiance transfer properties of 3D Gaussian primitives [15].

Neural Shading is concerned with learning light reflectance functions instead of using analytical models developed in computer graphics. This has been successfully applied to static scenes [35, 46, 48, 53, 31] and dynamic objects [23, 24]. Image-based methods enable relighting of a single portrait [12, 26, 38, 41, 51], but fail to synthesize novel views and struggle with temporal consistency, which are key requirements for head avatars. While [34] learns the coefficients of an explicit precomputed radiance transfer (PRT) function, [31] proposes to learn the PRT function with a neural network. In contrast, we propose a hybrid neural shading approach, combining implicitly learned diffuse radiance transfer with a well-established analytical specular term.

3 Multi-View OLAT Dataset of Faces

Capturing human faces under known, calibrated illumination enables efficient estimation of skin properties such as reflectance [7, 34, 35] and pore-level normals [21]. We therefore introduce a novel dataset, which consists of multi-view recordings of different subjects in a light stage setting.

Table 1: Existing light stage datasets of human heads. ICT-3DRFE [37] contains only processed data and no raw footage or calibration data.

Dataset	# IDs	# Views	# Lights	FPS	Resolution	Setup Cost
3DRFE [37]	23		only processed data			\$\$
Goliath [22]	4	144	460	9 (90) ¹	1334x2048	\$\$\$
Ours	10	16	40	72	2200x3208	\$




Figure 2: **OLAT Dataset:** (a) Our custom light-stage rig we used to capture (b) our dataset consisting of high-resolution, high frame rate, multi-view recordings of faces under both OLAT and fully-lit conditions.

The dataset offers an unprecedented combination of high-resolution, high-frame-rate multi-view recordings of many sequences under numerous calibrated lighting conditions.

3.1 Capture Setup

As our primary capture target are human faces, we build a light stage setup that covers the frontal hemisphere of the subject’s head. The setup consists of 16 machine vision cameras and 40 custom-built LED modules that are driven by microcontrollers. The cameras cover a field of view of 93° horizontally and 32° vertically. The lights are placed uniformly around the subject, covering a range of 180° horizontally and 60° vertically. All cameras and lights face towards the subject’s face. See Figure 2 for a visualization of the capture rig. Each LED emits enough luminance to run both one-light-at-a-time (OLAT) and more complex light patterns, while maintaining a low shutter speed of 3ms, thus reducing motion blur to a minimum. We use high-quality LEDs with a Color Rendering Index (CRI) of over 98, which closely approximates natural white light.

We control the LEDs using microcontrollers that we synchronize with the cameras using a vendor-specific logic. Our capture rig is equipped with 16 machine vision cameras, which we internally synchronize using the Precision Time Protocol (PTP), leading to multi-view frames captured with a deviation of less than one microsecond. Each of the cameras records images with a resolution of 2,200x3,208 pixels at 72 frames per second, sufficient to capture specular reflections on the skin at pore-level detail as depicted in Figure 2b.

3.2 Data Acquisition

Using our light stage setup, we capture several sequences of different participants for a few minutes in total. During the capture sessions, each participant performs a predefined set of facial expressions, emotions, and reads out several sentences. Please refer to the supplementary for more details about our capture script. In total, we record around 150 seconds for each subject, which is divided into 6 blocks. In addition, we capture another sequence where every participant is free to perform arbitrary expressions for 20 seconds. For each frame, we activate a new light from the set of available OLAT configurations. To enable tracking, we follow previous work [43, 34] and interleave our cycle of light patterns with fully-lit tracking frames. More specifically, every third frame is a tracking frame, which

¹90 FPS is only available for a short test segment.

results in tracking sequences captured at effectively 24 frames per second. See the rightmost image of Figure 2b for an example of a tracking frame.

3.3 Data Processing

The camera poses and intrinsic parameters are obtained using a checkerboard and bundle adjustment. Both the position and the intensity of the LEDs are calibrated using a mirror sphere whose shape and reflection properties are known. We follow the procedure of [45] and find the 3D position of each light source using ray-tracing in a multi-view capture of the mirror sphere. To account for differences in colors among the camera sensors, we use a color checker board and compute a color correction matrix for each camera. We use BiRefNet [55] for obtaining high-resolution foreground masks and obtain semantic segmentation with Facer [56].

3.4 Data Privacy

Our dataset contains highly personal information, which requires distributing it with extreme caution. We will only share the data with approved academic institutions and exclusively for non-commercial research purposes. All participants signed an agreement for publication, yet retain the right to have their data deleted at any time in the future, which we will enforce when distributing the dataset.

4 Method

Our method reconstructs relightable avatars from multi-view light stage sequences. Figure 3 provides an overview of our method. After preliminary information (Sec. 4.1), we describe our geometry (Sec. 4.2) and appearance (Sec. 4.3) model. In Sec. 4.4 and 4.5, we provide details about the optimization strategy and implementation details, respectively.

4.1 Preliminaries

3D Gaussian Splatting [15] introduces a point-based radiance field representation, that defines a 3D scene with a set of anisotropic 3D Gaussians parameterized by mean μ , covariance Σ and opacity σ . In addition, each Gaussian can hold an arbitrary number of features. Unlike continuous representations such as NeRF [25] that require ray marching for rendering, 3D Gaussians can be efficiently projected onto the image plane and rasterized in real-time on consumer-grade GPUs. We refer to the original paper of [15] for a more thorough overview.

Physically-based Rendering aims at synthesizing images by simulating the physical transport of light from the emitter to the camera sensor. The core is the rendering equation [13] that is defined as follows:

$$L_o(\mathbf{x}, \omega_o) = \int_{\Omega} f_r(\mathbf{x}, \omega_i, \omega_o) L_i(\mathbf{x}, \omega_i) (\omega_i \cdot \mathbf{n}) d\omega_i \quad (1)$$

where L_o is the outgoing radiance observed by the camera, L_i is the incident radiance at point $\mathbf{x}$ from direction ω_i , and f_r is the BRDF. Our goal is to recover the BRDF f_r from data observations, such that the resulting avatar can be integrated with novel illuminations.

4.2 Geometry

We model the geometry of our avatar with a fixed set of anisotropic Gaussians [15] that we define on the UV map of a tracked template mesh. Inspired by [40, 34, 10], we employ an expression-dependent dynamics module $\mathcal{F}_g$, and a view and expression-dependent module $\mathcal{F}_v$ to model fine-grained geometric expression details beyond the scope of the template mesh.

As our base geometry, we use the parametric head model FLAME [17], which models coarse deformations over time. Given the fully-lit tracking frames of our dataset, we obtain shape, expression, and pose parameters using the photometric tracker VHAP [30, 29]. For the remaining OLAT frames, we linearly interpolate the FLAME parameters of the nearest tracking frames. To get the proxy geometry for the Gaussian primitives, we obtain the posed mesh $\mathcal{M} = (\mathcal{V}, \mathcal{F})$ from the FLAME parameters and compute tangents $\mathbf{t}_k$, bitangents $\mathbf{b}_k$ and normals $\mathbf{n}_k$ for every texel k on the UV map. In addition, we obtain the interpolated 3D position of texel k , denoted as $\hat{\mu}_k$.




Figure 3: **Method Overview:** Given estimated FLAME coefficients, we obtain posed 3D Gaussian primitives with our expression-dependent dynamics module $\mathcal{F}_g$. To render photorealistic appearance, we combine the neural diffuse BRDF $\mathcal{F}_d$ with an analytical specular shading term. The parameters for the specular shading are predicted by the view-dependent $\mathcal{F}_v$ network. The avatar is optimized from light stage sequences using a photometric loss term.

Given the tracked FLAME expression parameters θ_{FLAME} , we define $\mathcal{F}_g$ as a convolutional neural network which predicts per-gaussian attributes in UV-space:

$$\{\delta\mu, q, s, \sigma, f^{expr}\}_{k=1}^M = \mathcal{F}_g(\theta_{FLAME}) \quad (2)$$

The final Gaussian center μ_k is defined as $\mu_k = \hat{\mu}_k + R_k^{TBN} \mu_k^{local} + \delta\mu$, where $R_k^{TBN} = [b_k, t_k, n_k]$ is the orientation of the shading frame of texel k , and μ_k^{local} is a parameter learned statically for each gaussian. The purpose of μ_k^{local} is to define most of the offsets expression-independent, such that we can regularize the expression-dependent offsets $\delta\mu_k$ to be small, which avoids artifacts when synthesizing novel expressions. The remaining Gaussian parameters q_k, s_k, σ_k are directly predicted by $\mathcal{F}_g$. f_k^{expr} is an expression-dependent feature vector for shading, which we describe in Sec. 4.3.

4.3 Material

Modeling the reflectance properties of faces with common analytical models from computer graphics inevitably leads to insufficient quality due to their lack of modeling global illumination effects, such as subsurface scattering, which is omnipresent on human skin. We observe that such global illumination effects primarily affect the low-frequency, view-independent diffuse part. Thus, we propose a hybrid shading scheme, which learns diffuse light transport implicitly with a small neural network, while modeling specular reflectance with a well-established analytical model. To this end, we decompose the reflectance function $f_r(\omega_o, \omega_i)$ from Eq. (1) into a view-independent diffuse term $f_d(\omega_i)$ and a view-dependent specular term $f_s(\omega_o, \omega_i)$.

Diffuse. The view-independent diffuse term models subsurface scattering and self-shadowing effects. At its core is a tiny neural network $\mathcal{F}_d$, shared among all primitives, and jointly trained with the avatar. The final diffuse color c_k^d is computed by multiplying statically learned albedo a_k with the predicted reflectance of $\mathcal{F}_d$:

$$c_k^d = a_k \mathcal{F}_d(SH_m(L_i), f_k^{expr}) \quad (3)$$

where $SH_m(L_i)$ are the coefficients from the spherical harmonics parameterization of the incident light of degree m , and f_k^{expr} are the expression-dependent feature vectors. We empirically set SH degree m to 6 in all experiments.

$\mathcal{F}_d$ is parameterized as a monochrome BRDF function, mapping single-channel incident light to a scalar reflectance value. This parameterization is necessary as the model only sees white light during

training, yet must also handle colored illumination at inference time. Here, we evaluate $\mathcal{F}_d$ separately for each color channel and concatenate the results into a single reflectance vector, which we multiply element-wise with the albedo. The architecture of $\mathcal{F}_d$ is detailed in the supplementary.

Specular. Our specular term is based on the Cook-Torrance model [6] which is generally defined as

$$f_s(\omega_i, \omega_o, r) = k_s \frac{D(\omega_o, \omega_i, r) G(\omega_o, \omega_i) F(\omega_o, \omega_i)}{4(\mathbf{n} \cdot \omega_o)(\mathbf{n} \cdot \omega_i)} \quad (4)$$

$$D(\cdot) = \alpha D_{12}(\cdot) + (1 - \alpha) D_{48}(\cdot) \quad (5)$$

where k_s is the specular intensity, D is the Normal Distribution Function (NDF), and G is the masking and shadowing term, which is derived from the NDF [39]. F models the Fresnel effect, for which we use Schlick's approximation [36]. As the NDF, we use the 2-Blinn-Phong-lobe mix introduced by Riviere et al. [33]. The advantage of this NDF representation is that roughness r is a linear parameter, which is beneficial during the optimization.

Due to the ellipsoidal shape of 3D Gaussian primitives, it is non-trivial to associate a single normal vector to them. As suggested by Saito et al. [34], we observe that the normal varies with the viewing direction. We, therefore, use a second, smaller CNN $\mathcal{F}_v$, which takes expression features $\mathbf{f}_{expr}$ and the viewing direction ω_o , and predicts specular intensity s_k and normal offsets $\delta \mathbf{n}$. While the general idea behind $\mathcal{F}_v$ is similar to [34], our U-Net architecture requires fewer network parameters, improving performance on consumer-level hardware. The final shading normals are obtained by adding the normal offsets $\delta \mathbf{n}_k$ to the mesh normals, followed by normalization. During training, we evaluate the specular term with the point light pattern of the current frame. For environment map relighting, we use the split-sum approximation [14]. We provide details in the supplementary.

4.4 Optimization

Given calibrated multi-view sequences from our dataset and corresponding estimated FLAME parameters, we jointly optimize $\mathcal{F}_g$, $\mathcal{F}_d$, $\mathcal{F}_v$, and static parameters μ_k^{local} , $\mathbf{a}_k$, r_k with the following loss term:

$$\mathcal{L} = \mathcal{L}_{rgb} + \mathcal{L}_{reg} \quad (6)$$

$$\mathcal{L}_{reg} = \lambda_{normal} \mathcal{L}_{normal} + \lambda_{alpha} \mathcal{L}_{alpha} + \lambda_{scale} \mathcal{L}_{scale} + \lambda_{pos} \mathcal{L}_{pos} \quad (7)$$

where $\mathcal{L}_{rgb} = \lambda_{l1} \mathcal{L}_{l1} + \lambda_{SSIM} \mathcal{L}_{SSIM}$ is the photometric loss term consisting of an L1 and SSIM term as proposed by [15]. We set $\{\lambda_{l1}, \lambda_{SSIM}\}$ to $\{1.0, 0.2\}$ in all experiments. Our regularization loss $\mathcal{L}_{reg}$ consists of the normal loss $\mathcal{L}_{normal} = \|\delta \mathbf{n}\|$, which encourages the predicted normal offsets to be small, and thus, be close to the normals of the FLAME mesh. Our capture rig only contains lights on the frontal hemisphere, which would lead to artifacts when we render the avatars with lights from the rear or environment maps. We find that a simple L2 loss $\mathcal{L}_{alpha}$ between the rendered alpha maps and the alpha masks from background matting prevents the avatar from becoming too transparent. The scale loss is adapted from [34] and promotes the primitive scales to remain in a reasonable range. $\mathcal{L}_{pos}$ is another L2 term which drives $\mathcal{F}_g$ to predict small delta means. We set $\{\lambda_{alpha}, \lambda_{scale}, \lambda_{pos}\}$ to $\{2e-2, 2e-2, 1e-5\}$ in all experiments.

4.5 Implementation Details

We implement all networks and optimization logic in PyTorch [27], and write custom GPU kernels for the specular shading using the SLANG.D shading language [3]. For rendering the Gaussian primitives, we use *gsplat* [50]. We use a texture resolution of 512^2 in all experiments, which results in 202k primitives after masking out texels from the FLAME UV map that are not assigned to any surface point. We use the 2023 version of FLAME [17] with the manually added teeth from Qian et al. [30]. We train our avatars at 1100x1604 resolution for 250k iterations with a batch size of 4, which takes approximately 30 hours on a single NVIDIA RTX A6000 GPU.

4.6 Differences to RGCA

RGCA [34] uses a variational autoencoder, which learns a personalized expression space and predicts the parameters of the precomputed radiance transfer function. In contrast, our method builds directly on top of FLAME [17], which has a shared expression space across identities, enabling applications

Table 2: Quantitative results on held-out lights on both the training and held-out segments.

Method	Relighting			Relighting + Self-Reenactment		
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
RGCA	29.21	0.8462	0.1659	26.31	0.8206	0.1917
RGCA _{FLAME}	29.78	0.8464	0.1444	26.91	0.8282	0.1667
Ours	31.38	0.8956	0.1040	28.08	0.8730	0.1317

such as cross-reenactment. Generally speaking, our avatars can be animated with FLAME parameters from any source without the need for a personalized encoder. Further, we model diffuse light transport with a small MLP and use a Cook-Torrance [6, 33] variant for specular reflection.

5 Experiments

We evaluate our method on 4 subjects from our dataset, where our focus lies on relighting and self-reenactment. From the 16 available camera views, we use 15 for training, and hold out the center camera for testing. We further hold out 4 light patterns from training altogether. From the available sequences, we use all scripted sequences for training and use the *free* sequence for testing. As the test metrics, we use the Peak-Signal-to-Noise Ratio (PSNR), Structural-Similarity-Index-Measure (SSIM) [42] and the Learned Perceptual Image Patch Similarity (LPIPS) [52]

Baselines. Our main baseline is *Relightable Gaussian Codec Avatars* (RGCA) [34], a recent method that builds head avatars by decoding learned expression codes to 3D Gaussian attributes for geometry and intrinsic radiance transfer. The resulting avatars can be relighted by integrating the predicted intrinsic radiance properties with novel light sources. The input to RGCA are the vertices of a coarse template mesh together with unwrapped average textures, for which we use the FLAME meshes and textures from the VHAP [29, 30] tracking. Since the expression space of RGCA is learned per identity, it requires comprehensive training sequences, while our method can leverage the existing FLAME expression space. Therefore, we introduce a second baseline denoted *RGCA_{FLAME}*, where we replace the learned expression latent space with FLAME expression coefficients.

5.1 Relighting and Self-Reenactment

Our primary target application is reenactment under novel illuminations. Therefore, we animate our trained avatars with the FLAME parameters of the held-out sequence and select those frames with a light pattern not seen during training. We then render the avatars from the held-out camera view. In Table 2 we report the quantitative results of relighting for both a training and test sequence. Figure 4 presents the qualitative results of our avatars rendered from the test camera, with an unseen lighting condition and expression. A comparison under environment map relighting is shown in Figure 7. Our rendered avatars match the target appearance more closely in terms of the color and fine geometric details, which enables more realistic specular reflections. Notably, we observe that RGCA conditioned on FLAME parameters performs strictly better than the original version with the personalized expression space. We hypothesize that learning an expression space per subject is suboptimal for reenactment tasks.

In Figure 5, we qualitatively compare the intrinsic decomposition performed by our approach to the baselines. Our method recovers cleaner albedo and sharper specular highlights from the data observations and faithfully decomposes the diffuse and specular parts of the material. In addition, we recommend to watch our supplementary video for more results, which allow for a more complete comparison, including the temporal axis.

5.2 Ablation Study

We verify the key components of our method with ablation experiments, which we conduct with the same subjects. A qualitative and quantitative comparison is presented in Figure 6 and Table 3, respectively.

PBR. We compare our full model to a version where we replace the hybrid neural shading with a classic PBR shading model using a Lambertian term for diffuse, and a Cook-Torrance [6] for specular reflection. This simple appearance model cannot reproduce the complex appearance of skin since subsurface scattering is not modeled, which results in synthetically looking renderings.




Figure 4: **Relighting and Self-Reenactment:** Qualitative comparison on held-out segments and held-out illuminations.

PRT Diffuse. We compare our neural diffuse component to the learned precomputed radiance transfer (PRT) model introduced in [34]. The respective SH coefficients are directly predicts by $\mathcal{F}_g$. Although, we observe good results on the training frames, PRT struggles to generalize to novel illuminations, which aligns with our findings from Section 5.1.

SG. To justify the choice of the specular term, we ablate our specular BRDF to a simple Spherical Gaussian term. Compared to our Cook-Torrance term, we observe slightly better PSNR and LPIPS scores as well as marginal sharper details under point light illumination. However, as depicted in Fig. 6, we observe significantly more natural and detailed pore-level reflection with our Cook-Torrance based specular model.

Alpha Loss. The key component to prevent the avatars from becoming too transparent is the alpha loss using estimated foreground segmentation masks. Since our capture setup only has lights and cameras on the frontal hemisphere, we observe artifacts with environment map relighting when using no regularization. We want to highlight that this simple regularization scheme effectively reduces the complexity of our capture setup, the resulting dataset, and computational cost during training by half.

Expression Features. We compare our expression-dependent features against static feature vectors as proposed by Giebenhain et al. [10]. Here, we use static features for the diffuse BRDF network $\mathcal{F}_d$, and use the FLAME parameters as a condition for $\mathcal{F}_v$. We observe that without the expression-dependent features, the model fails to accurately reproduce pore-level details and specular highlights. We can further notice worse color and reflections compared to the full model.

5.3 Application

Once trained, the only inference parameters are FLAME expression and pose parameters, which can be obtained from monocular videos [29, 30]. We demonstrate this by animating our avatars with short




Figure 5: **Comparison of Intrinsic Decomposition:** We compare the recovered albedo (a) and normals (b), as well as the diffuse (c) and specular (d) contributions on a training frame that sum up to the final rendering (e). Note that the reference image (f) is identical in all rows.

Table 3: **Ablations:** We conduct ablations with the same 4 subjects and report relighting and reenactment results on the training and test expressions.

Method	Relighting			Relighting + Self-Reenactment		
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
w/ PBR shading	29.42	0.8719	0.1344	26.31	0.8448	0.1665
w/ PRT diffuse	29.23	0.8374	0.1577	25.47	0.8074	0.1918
w/ SG	31.55	0.8953	0.1031	28.09	0.8729	0.1310
w/o alpha loss	31.34	0.8955	0.1043	28.07	0.8729	0.1328
w/o expr. features	31.23	0.8928	0.1071	28.13	0.8717	0.1332
Ours (full)	31.38	0.8956	0.1040	28.08	0.8730	0.1317



Figure 6: **Ablation Study:** With only PBR shading, the avatar has a synthetic, plastic-like appearance. Without the alpha loss, we observe artifacts when rendering with environment maps. The expression-dependent features further improve both appearance and fine geometric details.



Figure 7: *Left*: Qualitative comparison on environment map relighting. *Right*: Animation using monocular videos.

Table 4: **Runtime comparison**: We report the component-wise inference time in milliseconds.

Method	CNNs	Diffuse Shading	Specular Shading	Splatting	Total
RGCA	9ms	1ms	1ms	9ms	20ms
Ours	4ms	3ms	1ms	9ms	17ms

video sequences from the VFHQ dataset [44]. We obtain the FLAME parameter with the monocular version of the VHAP tracker [29], and relight our avatars with environment maps collected from PolyHaven [1]. We present the results in Figure 7 and highly encourage the reader to watch the accompanying video for temporal results.

5.4 Runtime

In Table 5, we summarize the runtime of the components of our method and compare it with RGCA [34]. We conducted all measurements on a single NVIDIA RTX A6000 GPU using a UV resolution of 512^2 , which results in 202k primitives (due to masking of texels not assigned to any surface location). We render images at a resolution of 1100×1604 , corresponding to the training resolution of our avatars.

5.5 Discussion

Limitations. While our method delivers state-of-the-art results and enables more practical animation than previous methods, our approach is still not without limitations. While our capture setup is an order of magnitude more economical than existing setups [34, 4, 12], avatar training still requires several thousand frames and a diverse set of training expressions. Obtaining photorealistic avatars from causal phone captures with uncalibrated lighting remains an open challenge for future work. The FLAME base geometry is limited in its expressiveness, is sensitive to tracking failures (particularly with respect to gaze direction), and does not model the mouth interior. These limitations are consequently inherited by our avatars. As of now, the neural diffuse shading model is trained from scratch jointly with the avatar. Using our dataset to learn an appearance prior of human faces and heads is an interesting direction for future work.

Ethical Considerations. Creating photorealistic, relightable avatars entails the potential for various malicious use cases, such as identity theft, deepfakes, and privacy violations. This is a particular concern when avatars can be driven from simple video sequences, as in our case. However, to create an avatar with our method, the respective subject must first be scanned in our capture setup, which is only applied to a limited number of consenting individuals. Further, we will be restrictive with access to our dataset as outlined in Section 3.4.

6 Conclusion

We have presented *BecomingLit*, a novel framework for reconstructing photorealistic, relightable avatars from a capture setup, orders of magnitude more economical than previous state-of-the-art methods. We have proposed a new hybrid shading approach for 3D Gaussian primitives, which enables better generalization to novel illuminations and expressions. Our relightable avatars can be animated from simple videos and relighted with both point lights and environment maps. Along with our method, we will publish a new dataset of faces under OLAT conditions, which is unprecedented in terms of resolution and frame rate. We believe that this will democratize research on facial appearance modeling and serve as a valuable contribution to the community.

Acknowledgments and Disclosure of Funding

This work was supported by the ERC Consolidator Grant Gen3D (101171131) and the German Research Foundation (DFG) Research Unit “Learning and Simulation in Visual Computing”. Additionally, we would like to thank Angela Dai for the video voice-over.

References

- [1] Poly haven - hdris, 2023. URL <https://polyhaven.com/hdris>.
- [2] Dejan Azinović, Olivier Maury, Christophe Hery, Matthias Nießner, and Justus Thies. High-res facial appearance capture from polarized smartphone images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16836–16846, 2023.
- [3] Sai Praveen Bangaru, Lifan Wu, Tzu-Mao Li, Jacob Munkberg, Gilbert Bernstein, Jonathan Ragan-Kelley, Frédo Durand, Aaron Lefohn, and Yong He. Slang. d: Fast, modular and differentiable shader programming. *ACM Transactions on Graphics (TOG)*, 42(6):1–28, 2023.
- [4] 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.
- [5] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3d faces. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, pages 157–164, 2023.
- [6] Robert L Cook and Kenneth E. Torrance. A reflectance model for computer graphics. *ACM Transactions on Graphics (ToG)*, 1(1):7–24, 1982.
- [7] 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.
- [8] Stephan J Garbin, Marek Kowalski, Virginia Estellers, Stanislaw Szymanowicz, Shideh Rezaeifar, Jingjing Shen, Matthew A Johnson, and Julien Valentin. Voltemorph: Real-time, controllable and generalizable animation of volumetric representations. In *Computer Graphics Forum*, volume 43, page e15117. Wiley Online Library, 2024.
- [9] Abhijeet Ghosh, Graham Fyffe, Borom Tunwattanapong, Jay Busch, Xueming Yu, and Paul Debevec. Multiview face capture using polarized spherical gradient illumination. In *Proceedings of the 2011 SIGGRAPH Asia Conference*, pages 1–10, 2011.
- [10] Simon Giebenhain, Tobias Kirschstein, Martin Rünz, Lourdes Agapito, and Matthias Nießner. Npga: Neural parametric gaussian avatars. In *SIGGRAPH Asia 2024 Conference Papers*, pages 1–11, 2024.
- [11] 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.
- [12] Mingming He, Pascal Clausen, Ahmet Levent Taşel, Li Ma, Oliver Pilarski, Wenqi Xian, Laszlo Rikker, Xueming Yu, Ryan Burgert, Ning Yu, et al. Diffrelight: Diffusion-based facial performance relighting. *arXiv preprint arXiv:2410.08188*, 2024.
- [13] James T Kajiya. The rendering equation. In *Proceedings of the 13th annual conference on Computer graphics and interactive techniques*, pages 143–150, 1986.
- [14] Brian Karis and Epic Games. Real shading in unreal engine 4. *Proc. Physically Based Shading Theory Practice*, 4(3):1, 2013.
- [15] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139–1, 2023.
- [16] Tobias Kirschstein, Shenhan Qian, Simon Giebenhain, Tim Walter, and Matthias Nießner. Nersemble: Multi-view radiance field reconstruction of human heads. *ACM Transactions on Graphics (TOG)*, 42(4):1–14, 2023.
- [17] Tianye Li, Timo Bolkart, Michael J Black, Hao Li, and Javier Romero. Learning a model of facial shape and expression from 4d scans. *ACM Trans. Graph.*, 36(6):194–1, 2017.

- [18] Stephen Lombardi, Tomas Simon, Gabriel Schwartz, Michael Zollhoefer, Yaser Sheikh, and Jason Saragih. Mixture of volumetric primitives for efficient neural rendering. *ACM Transactions on Graphics (ToG)*, 40(4):1–13, 2021.
- [19] Shengjie Ma, Yanlin Weng, Tianjia Shao, and Kun Zhou. 3d gaussian blendshapes for head avatar animation. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–10, 2024.
- [20] 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.
- [21] 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*, 9(10):2, 2007.
- [22] Julieta Martinez, Emily Kim, Javier Romero, Timur Bagautdinov, Shunsuke Saito, Shou-I Yu, Stuart Anderson, Michael Zollhöfer, Te-Li Wang, Shaojie Bai, Chenghui Li, Shih-En Wei, Rohan Joshi, Wyatt Borsos, Tomas Simon, Jason Saragih, Paul Theodosis, Alexander Greene, Anjani Josyula, Silvio Mano Maeta, Andrew I. Jewett, Simon Venshtain, Christopher Heilman, Yueh-Tung Chen, Sidi Fu, Mohamed Ezzeldin A. Elshaer, Tingfang Du, Longhua Wu, Shen-Chi Chen, Kai Kang, Michael Wu, Youssef Emad, Steven Longay, Ashley Brewer, Hitesh Shah, James Booth, Taylor Koska, Kayla Haidle, Matt Andromalos, Joanna Hsu, Thomas Dauer, Peter Selednik, Tim Godisart, Scott Ardisson, Matthew Cipperly, Ben Humberston, Lon Farr, Bob Hansen, Peihong Guo, Dave Braun, Steven Krenn, He Wen, Lucas Evans, Natalia Fadeeva, Matthew Stewart, Gabriel Schwartz, Divam Gupta, Gyeongsik Moon, Kaiwen Guo, Yuan Dong, Yichen Xu, Takaaki Shiratori, Fabian Prada, Bernardo R. Pires, Bo Peng, Julia Buffalini, Autumn Trimble, Kevyn McPhail, Melissa Schoeller, and Yaser Sheikh. Codec Avatar Studio: Paired Human Captures for Complete, Driveable, and Generalizable Avatars. *NeurIPS Track on Datasets and Benchmarks*, 2024.
- [23] 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.
- [24] 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.
- [25] 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.
- [26] Rohit Pandey, Sergio Orts-Escolano, Chloe Legendre, Christian Haene, Sofien Bouaziz, Christoph Riemann, Paul E Debevec, and Sean Ryan Fanello. Total relighting: learning to relight portraits for background replacement. *ACM Trans. Graph.*, 40(4):43–1, 2021.
- [27] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, volume 32, 2019.
- [28] Pascal Paysan, Reinhard Knothe, Brian Amberg, Sami Romdhani, and Thomas Vetter. A 3d face model for pose and illumination invariant face recognition. In *2009 sixth IEEE international conference on advanced video and signal based surveillance*, pages 296–301. Ieee, 2009.
- [29] Shenhan Qian. Vhap: Versatile head alignment with adaptive appearance priors, sep 2024. URL <https://github.com/ShenhanQian/VHAP>.
- [30] Shenhan Qian, Tobias Kirschstein, Liam Schoneveld, Davide Davoli, Simon Giebenhain, and Matthias Nießner. Gaussianavatars: Photorealistic head avatars with rigged 3d gaussians. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20299–20309, 2024.
- [31] 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.
- [32] Pramod Rao, Mallikarjun BR, Gereon Fox, Tim Weyrich, Bernd Bickel, Hans-Peter Seidel, Hanspeter Pfister, Wojciech Matusik, Ayush Tewari, Christian Theobalt, et al. Vorf: Volumetric relightable faces. In *33rd British Machine Vision Conference*, 2022.

- [33] Jérémy Riviere, Paulo FU Gotardo, Derek Bradley, Abhijeet Ghosh, and Thabo Beeler. Single-shot high-quality facial geometry and skin appearance capture. *ACM Trans. Graph.*, 39(4):81, 2020.
- [34] Shunsuke Saito, Gabriel Schwartz, Tomas Simon, Junxuan Li, and Giljoo Nam. Relightable gaussian codec avatars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 130–141, 2024.
- [35] Kripasindhu Sarkar, Marcel C Bühler, Gengyan Li, Daoye Wang, Delio Vicini, Jérémy Riviere, Yinda Zhang, Sergio Orts-Escolano, Paulo Gotardo, Thabo Beeler, et al. Litnerf: Intrinsic radiance decomposition for high-quality view synthesis and relighting of faces. In *SIGGRAPH Asia 2023 Conference Papers*, pages 1–11, 2023.
- [36] Christophe Schlick. An inexpensive brdf model for physically-based rendering. In *Computer graphics forum*, volume 13, pages 233–246. Wiley Online Library, 1994.
- [37] Giota Stratou, Abhijeet Ghosh, Paul Debevec, and Louis-Philippe Morency. Effect of illumination on automatic expression recognition: a novel 3d relightable facial database. In *2011 IEEE International Conference on Automatic Face & Gesture Recognition (FG)*, pages 611–618. IEEE, 2011.
- [38] Tiancheng Sun, Jonathan T Barron, Yun-Ta Tsai, Zexiang Xu, Xueming Yu, Graham Fyffe, Christoph Rhemann, Jay Busch, Paul E Debevec, and Ravi Ramamoorthi. Single image portrait relighting. *ACM Trans. Graph.*, 38(4):79–1, 2019.
- [39] Bruce Walter, Stephen R Marschner, Hongsong Li, and Kenneth E Torrance. Microfacet models for refraction through rough surfaces. *Rendering techniques*, 2007:18th, 2007.
- [40] Shaofei Wang, Tomas Simon, Igor Santesteban, Timur Bagautdinov, Junxuan Li, Vasu Agrawal, Fabian Prada, Shou-I Yu, Pace Nalbone, Matt Gramlich, Roman Lubachersky, Chenglei Wu, Javier Romero, Jason Saragih, Michael Zollhoefer, Andreas Geiger, Siyu Tang, and Shunsuke Saito. Relightable full-body gaussian codec avatars. *arXiv.org*, 2501.14726, 2025.
- [41] 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.
- [42] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.
- [43] Andreas Wenger, Andrew Gardner, Chris Tchou, Jonas Unger, Tim Hawkins, and Paul Debevec. Performance relighting and reflectance transformation with time-multiplexed illumination. *ACM Transactions on Graphics (TOG)*, 24(3):756–764, 2005.
- [44] Liangbin Xie, Xintao Wang, Honglun Zhang, Chao Dong, and Ying Shan. Vfhq: A high-quality dataset and benchmark for video face super-resolution. In *The IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2022.
- [45] Yingyan Xu, Jérémy Riviere, Gaspard Zoss, Prashanth Chandran, Derek Bradley, and Paulo Gotardo. Improved lighting models for facial appearance capture. *EG 2022-Short Papers*, pages 5–8, 2022.
- [46] Yingyan Xu, Gaspard Zoss, Prashanth Chandran, Markus Gross, Derek Bradley, and Paulo Gotardo. Renef: Relightable neural radiance fields with nearfield lighting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22581–22591, 2023.
- [47] Yuelang Xu, Benwang Chen, Zhe Li, Hongwen Zhang, Lizhen Wang, Zerong Zheng, and Yebin Liu. Gaussian head avatar: Ultra high-fidelity head avatar via dynamic gaussians. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1931–1941, 2024.
- [48] 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.
- [49] 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 Papers*, pages 1–11, 2023.
- [50] Vickie Ye, Ruilong Li, Justin Kerr, Matias Turkulainen, Brent Yi, Zhuoyang Pan, Otto Seiskari, Jianbo Ye, Jeffrey Hu, Matthew Tancik, and Angjoo Kanazawa. gsplat: An open-source library for Gaussian splatting. *arXiv preprint arXiv:2409.06765*, 2024. URL <https://arxiv.org/abs/2409.06765>.
- [51] 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.

- [52] 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.
- [53] 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.
- [54] 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*, 43(1):1–16, 2023.
- [55] Peng Zheng, Dehong Gao, Deng-Ping Fan, Li Liu, Jorma Laaksonen, Wanli Ouyang, and Nicu Sebe. Bilateral reference for high-resolution dichotomous image segmentation. *CAAI Artificial Intelligence Research*, 3:9150038, 2024.
- [56] Yinglin Zheng, Hao Yang, Ting Zhang, Jianmin Bao, Dongdong Chen, Yangyu Huang, Lu Yuan, Dong Chen, Ming Zeng, and Fang Wen. General facial representation learning in a visual-linguistic manner. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18697–18709, 2022.
- [57] 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.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: Our main contributions from the abstract are a novel dataset of faces and an avatar representation with a new hybrid shading scheme. The dataset is introduced in Sec. 3, the avatar in Sec. 4. We validate our contributions with experiments in Sec. 5.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: A discussion of limitations can be found in the respective paragraph in Sec. 5.5.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: The paper does not include theoretical results.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.

- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We list implementation details including hyperparameters in Sec. 4.5 and Sec. 4.4. Detailed network architectures are included in the supplementary. In addition, code and data will be publicly released upon acceptance.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: Code and data will be released upon acceptance. Following GDPR, we will set up a download form to manage the distribution. Unfortunately, this is not possible in an anonymous form.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).

- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: See Sec. 4.5 of the main paper. Detailed network architectures and hyperparameters are provided in the supplementary. Data splits will be released with the dataset upon acceptance.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: The quantitative results do not contain statistical significance tests due to limited computational resources.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: Details are listed in Section 4.5 of the main paper. The resources are the same for all experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [No]

Justification: We deviate from the Code of ethics with respect to human research subjects to whom we cannot provide financial compensation. All captured subjects participated on a voluntary basis.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: A discussion of ethical concerns can be found in Sec. 5.5 in the respective paragraph.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [Yes]

Justification: We discuss the sensitivity of our dataset in Sec. 3.4 and outline how we minimize the risk for misuse when distributing it.

Guidelines:

- The answer NA means that the paper poses no such risks.

- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [\[Yes\]](#)

Justification: We cited all used code in the main paper (Sec. 4.5). Detailed licence information is included in the supplementary.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [\[Yes\]](#)

Justification: Our dataset is thoroughly documented in Sec. 3. Additional documentation will be made available with the release of the dataset upon acc

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [\[Yes\]](#)

Justification: The script given to the participants of our dataset is provided in the supplementary material. Data collection and curation was handled exclusively by the authors. All captured subjects participated on a voluntary basis.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.

- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [Yes]

Justification: We confirmed that our data collection is compliant with local privacy regulations.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.

BecomingLit: Relightable Gaussian Avatars with Hybrid Neural Shading

Supplementary Material

Jonathan Schmidt Simon Giebenhain Matthias Nießner
Technical University of Munich

jonathsch.github.io/becominglit

A Network Architecture

Our geometry module $\mathcal{F}_g$ maps FLAME [17] expression, jaw and eyes pose coefficients, and predicts per-voxel attributes $\{\delta\mu, \mathbf{q}, \mathbf{s}, \sigma, \mathbf{f}^{expr}\}_{k=1}^M$, where $\mathbf{f}_k^{expr}$ has a dimension of 32 in all experiments. We use the first 100 principal components for the expression parameter, and a rodrigues parameterization for jaw and both eye rotations. Hence, our input of shape $\mathbb{R}^{109}$ is transformed by a linear layer and then reshaped to $256 \times 8 \times 8$. A set of transposed convolutional layers then gradually upsamples the feature maps to the final output of shape $43 \times 512 \times 512$. We use leaky-ReLU as activation function for all layers except for the final output. For all convolutional layers, we adopt untied bias [4].

$\mathcal{F}_v$ takes as input the per-Gaussian feature map $\mathbf{f}^{expr}$, and the view direction, which is encoded using a single linear layer (8-dim output shape) and then expanded to the height and with dimension of the feature map. We concatenate the feature map and encoded view direction and feed it through a single convolutional layer which downsamples the input by half. Finally, a transposed convolutional layer maps the latent feature map back to its original resolution with 4 output channels.

Our diffuse BRDF network $\mathcal{F}_d$ is a 3-layer MLP with hidden dimension 64 and leaky-ReLU activation in every layer, except the last one. The input is the concatenation of $\mathbf{f}_k^{expr}$ and the spherical harmonics coefficients of the incident light.

B Environment Map Rendering

In this section, we provide further details on how we render our avatars with all-frequency continuous illumination in the form of environment maps. While our diffuse BRDF trivially adopts to continuous illumination due to the spherical harmonics parameterization, we need to adopt the specular shading of the primitives.

For the specular pre-integration, we follow the split-sum approximation [14]. Karis et al. [14] propose to assume that the view direction ω_o and the surface normal $\mathbf{n}$ are identical. With that assumption, the specular reflection is no longer view-dependent, and we can pre-integrate the environment map for different roughness values using a mipmap. In each mipmap level, we numerically integrate L_i with importance sampling using the Blinn-Phong distribution:

$$L_o^{specular}(\mathbf{x}, \omega_o) = \int_{\Omega} L_i(\omega_i) D(\mathbf{h}, \mathbf{n}, r^2) (\omega_i \cdot \mathbf{n}) d\omega_i * \int_{\Omega} k_s \frac{DGF}{4(\omega_o \mathbf{n})(\mathbf{n} \omega_i)} d\omega_i \quad (8)$$

The incoming illumination $L_i(\omega_i)$ is now stored in the pre-integrated environment map $\hat{L}_{specular}(\omega, r)$. During rendering, we linearly interpolate the mip levels to obtain the final radiance value for the roughness parameter. Hence, the new specular term becomes:

$$L_o^{specular}(\mathbf{x}, \omega_o) \approx \hat{L}_{specular}(\omega, r) \int_{\Omega} k_s \frac{DGF}{4(\omega_o \mathbf{n})(\mathbf{n} \omega_i)} d\omega_i \quad (9)$$

The remaining integral is essentially the integration of the BRDF with a completely white environment light. We can substitute the Fresnel term $F(\omega_o, \mathbf{h})$ with the Schlick approximation [36] and factor out F_0 :




Figure 8: **Goliath-4 Evaluation:** Qualitative comparison on held-out frames and held-out light patterns from an unseen viewpoint.

$$\int_{\Omega} f_s(\omega_i, \omega_o)(\mathbf{n} \cdot \omega_i) d\omega_i = F_0 \int_{\Omega} \frac{f_s(\omega_i, \omega_o)}{F(\omega_o, \mathbf{h})} (1 - (1 - \omega_o \cdot \mathbf{h})^5) (\omega_i \cdot \mathbf{n}) d\omega_i + \int_{\Omega} \frac{f_s(\omega_i, \omega_o)}{F(\omega_o, \mathbf{h})} (1 - \omega_o \cdot \mathbf{h})^5 (\omega_i \cdot \mathbf{n}) d\omega_i \quad (10)$$

These integrals depend on the two inputs $(\omega_i \cdot \mathbf{n})$ and the roughness parameter r and act as a scale and bias to F_0 . We pre-integrate both terms for all possible input combinations in $[0, 1]^2$ and store the two outputs in the 2D texture map $\hat{f}_s(\omega, r)$

During rendering, we can now compute the shaded color by evaluating the following terms:

$$\omega_r = -\omega_o - 2(-\omega_o \cdot \mathbf{n})\mathbf{n} \quad (11)$$

$$a, b = \hat{f}_s((\omega_r \cdot \mathbf{n}), r) \quad (12)$$

$$c_k^s(\omega_o) = (a k_s + b) \hat{L}_{\text{specular}}(\omega_r, r) \quad (13)$$

C Goliath-4 Dataset

To demonstrate that our method also generalize well to data domains beyond our new light stage dataset, we perform additional experiments on the Goliath-4 dataset [22]. We train both our method and RGCA [34] in two different configurations. (1) With the full available set of cameras, holding out 10 random views for evaluation, and (2) with a random subset of 16 train cameras and 4 cameras for evaluation. We limit ourselves to one of the subjects (QZX685) and use the provided train/test split for evaluating unseen expressions and hold out a random subset of 10% of the available light patterns to evaluate relighting capabilities. Our method consistently outperforms RGCA on the SSIM and LPIPS metrics. We report quantitative results in Table 5, and a qualitative comparison in Figure 8. We want to point out, that the publicly released dataset [22] is subsampled to every 10th frame, and the provided images are heavily compressed, which inevitably leads to a drop in quality compared to the results shown in [34].

D Capture Script

For each participant of our dataset, we record 7 sequences in total. The first 6 consist of a predefined set of facial expressions, emotions and sentences that we ask the subjects to perform and read out. In the 7th sequence the participant is free to perform any facial expression for 20s. The instructions are given via a screen that is placed

Table 5: **Goliath-4 Dataset:** Quantitative results on Self-Reenactment and Relighting. We evaluate self-reenactment on the validation camera views.

Method	Full Cam Set			Random Cam Subset (10%)		
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
RGCA	29.89	0.8869	0.1392	28.74	0.8753	0.1468
Ours	29.70	0.9080	0.1165	28.68	0.8966	0.1298

in front of the subject. In the following we provide a list of the single components, which during the capture sessions are accompanied with images.

- **Expressions-1:**
 - Head rotation with mouth open and closed
 - Eyes blink
 - Eyes squint
 - Eyebrows up / down
 - Puffed Cheeks
 - Mouth Vacuum
 - Nose Wrinkle
 - Lip bite
- **Expressions-2**
 - Grin (multiple variations)
 - Jaw movement
 - Lip licking
 - Tongue
- **Emotions**
 - Shout
 - Laugh
 - Surprise
 - Fear
 - Angry
 - Sad
 - Disgust
 - Happy
 - Confusion
 - Amazement
 - Embarrassment
- **Sentences-1**
 - A cramp is no small danger on a swim.
 - He said the same phrase thirty times.
 - Pluck the bright rose without leaves.
 - Two plus seven is less than ten.
 - The glow deepened in the eyes of the sweet girl.
 - By eating yogurt you may live longer.
- **Sentences-2**
 - Bring your problems to the wise chief.
 - Write a fond note to the friend you cherish.
 - Clothes and lodging are free to new men.
 - We frown when events take a bad turn.
 - Port is a strong wine with a smoky taste.
 - They had slapped their thighs.
- **Sentences-3**

- She always jokes about too much garlic in his food.
- Why put such a high value on being top dog.
- All your wishful thinking won't change that.
- Take charge of choosing her bridesmaids gowns.
- Why buy oil when you always use mine.