

NeuFace: Realistic 3D Neural Face Rendering from Multi-view Images

Mingwu Zheng^{1,2}, Haiyu Zhang^{1,2}, Hongyu Yang³, Di Huang^{1,2,4*}

¹State Key Laboratory of Software Development Environment, Beihang University, Beijing, China

²School of Computer Science and Engineering, Beihang University, Beijing, China

³Institute of Artificial Intelligence, Beihang University, Beijing, China

⁴Hangzhou Innovation Institute, Beihang University, Hangzhou, China

{zhengmingwu, zhyzhy, hongyuyang, dhuang}@buaa.edu.cn

Abstract

Realistic face rendering from multi-view images is beneficial to various computer vision and graphics applications. Due to complex spatially-varying reflectance properties and geometry characteristics of faces, however, it remains challenging to recover 3D facial representations both faithfully and efficiently in the current studies. This paper presents a novel 3D face rendering model, namely **NeuFace**, to learn accurate and physically-meaningful underlying 3D representations by neural rendering techniques. It naturally incorporates the neural BRDFs into physically based rendering, capturing sophisticated facial geometry and appearance clues in a collaborative manner. Specifically, we introduce an approximated BRDF integration and a simple yet new low-rank prior, which effectively lower the ambiguities and boost the performance of the facial BRDFs. Extensive experiments are performed to demonstrate the superiority of NeuFace in human face rendering, along with a decent generalization ability to common objects. Code is released at [NeuFace](#).

1. Introduction

Rendering realistic human faces with controllable viewpoints and lighting is now becoming ever increasingly important with its applications ranging from game production, movie industry, to immersive experiences in the Metaverse. Various factors, including the sophisticated geometrical differences among individuals, the person-specific appearance idiosyncrasies, along with the spatially-varying reflectance properties of skins, collectively make faithful face rendering a rather difficult problem.

According to photogrammetry, the pioneering studies on this issue generally leverage complex active lighting setups, e.g., LightStage [8], to build 3D face models from multiple




Figure 1. Demonstration of the face rendering results and recovered underlying 3D representations.

photos of an individual, where accurate shape attributes and high-quality diffuse and specular reflectance properties are commonly acknowledged as the premises of its success. An elaborately designed workflow is required, typically involving a series of stages such as camera calibration, dynamic data acquisition, multi-view stereo, material estimation, and texture parameterization [42]. While a compelling and convincing 3D face model can be finally obtained, this output highly depends on the expertise of the engineers and artists with significant manual efforts, as the multi-step process inevitably brings diverse optimization goals.

Recently, 3D neural rendering, which offers an end-to-end alternative, has demonstrated promising performance in recovering scene properties from real-world imageries, such as view-dependent radiance [26, 28, 36, 38, 47] and geometry [33, 48, 49, 54, 55]. It's mainly credited to the disentanglement of the learnable 3D representations and the differentiable image formation process, free of the tedious photogrammetry pipeline. However, like classical function fitting, inverse rendering is fundamentally under-constrained, which may incur badly-conditioned fits of the underlying 3D representations, especially for intricate cases, e.g., non-Lambertian surfaces with view-dependent highlights. With the trend in the combination of computer graphics and learning techniques, several attempts take advantage of physically motivated inductive biases and present Phys-

*Corresponding author.

ically Based Rendering (PBR) [14, 31, 47, 56], where Bidirectional Reflectance Distribution Functions (BRDFs) are widely adopted. By explicitly mimicking the interaction of the environment light with the scene, they facilitate network optimization and deliver substantial gains. Unfortunately, the exploited physical priors are either heuristic or analytic [7, 20, 44], limited to a small set of real-world materials, *e.g.*, metal, incapable of describing human faces.

For realistic face rendering, the most fundamental issue lies in accurately modeling the optical properties of multi-layered facial skin [21]. In particular, the unevenly distributed fine-scale oily layers and epidermis reflect the incident lights irregularly, leading to complex view-dependent and spatially-varying highlights. This characteristic and the low-textured nature of facial surfaces strongly amplify the shape-appearance ambiguity. Moreover, subsurface scattering between the underlying dermis and other skin layers further complicates this problem.

In this paper, we follow the PBR paradigm for its potential in learning 3D representations and make the first step towards realistic 3D neural face rendering, mainly targeting complex skin reflection modeling. Our method, namely *NeuFace*, is able to recover faithful facial reflectance and geometry from only multi-view images. Concretely, we establish a PBR framework to learn neural BRDFs to describe facial skin, which simulates physically-correct light transport with a much higher representation capability. By using a differentiable Signed Distance Function (SDF) based representation, *i.e.*, *ImFace* [61], as the shape prior, the facial appearance and geometry field can be synchronously optimized in inverse rendering.

Compared to the analytic BRDFs, the neural ones allow richer representations for sophisticated material like facial skin. In spite of this superiority, such representations pose challenges to computational cost and data demand during training. To tackle these difficulties, the techniques in real-time rendering [1] are adapted to separate the hemisphere integral of neural BRDFs, where the material and light integrals are individually learned instead, bypassing the massive Monte-Carlo sampling phase [34] required by numerical solutions. Furthermore, a low-rank prior is introduced into the spatially-varying facial BRDFs, which greatly restricts the solution space thereby diminishing the need for large-scale training observations. These model designs indeed enable *NeuFace* to accurately and stably describe how the light interacts with the facial surface as in the real 3D space. Fig. 1 displays an example.

The main contributions of this study include: 1) A novel framework with naturally-bonded PBR as well as neural BRDF representations, which collaboratively captures facial geometry and appearance properties in complicated facial skin. 2) A new and simple low-rank prior, which significantly facilitates the learning of neural BRDFs and im-

proves the appearance recovering performance. 3) Impressive face rendering results from only multi-view images, applicable to various applications such as relighting, along with a decent generalization ability to common objects.

2. Related Work

We restrict the discussion specifically to static facial geometry and appearance capturing and 3D neural rendering. Please refer to [9, 21, 43] for more in-depth discussion.

Face Capturing. Its goal is to render a realistic 3D face under arbitrary lighting condition. Existing methods generally take advantage of photogrammetry techniques to estimate facial geometries and appearances, requiring massive manual efforts. In this case, they typically decompose the problem where facial geometry is pre-captured by an intricate multi-view stereo process [3, 4]. However, facial appearances with reflectances are still hard to acquire due to complex interactions between light and skin. The initial attempts [8, 16, 51] tackle the challenge by densely capturing the per-pixel facial reflectance at the cost of extensive data acquisition and specialized equipment. Subsequently, gradient [11, 19, 25] or polarized illumination [11, 12, 25, 39] is explored to reduce the memory cost, where most efforts are paid to the well-conditioned fitting of BRDFs. In contrast to the studies above, our solution is truly end-to-end and only observes facial skin under a single, unknown illumination condition, without cumbersome capturing settings.

3D Neural Rendering. The recent advance in this field, like NeRF [29], have already revolutionized the paradigm of multi-view reconstruction. With learnable volumetric representations (*e.g.* neural field [29, 55], grid [10], and hybrid [30]) and analytic differentiable forward mapping, scene properties can be directly inferred from only 2D imageries. While reporting impressive results on novel view synthesis [26, 28, 36, 38, 47] or geometry reconstruction [33, 48, 49, 54, 55], those studies still suffer from 2D-3D ambiguities, leading to that realistic appearances and accurate geometries can hardly be established simultaneously [57]. Such an issue is more prominent for facial skin due to the rather complicated reflection properties. Ref-NeRF [47] makes a step forward to the goal of both accurate surface normals and glossy appearances, which is achieved by reparameterizing the radiance in classical NeRF with reflection direction. It validates the significance of physical laws in disambiguity. [14, 22, 31, 41, 56] go further with PBR pipelines, delivering improved quality and supporting relighting simultaneously. Nevertheless, simplified material models are utilized or assumed, which are incapable of handling complex materials like skin. [24, 59, 60] train a neural material model from measurement data. But acquiring such data is impractical for live biotissue. By contrast, our method directly builds neural skin BRDFs and reconstructs accurate geometrical clues without any external data.




Figure 2. **Method overview.** *NeuFace* is composed of an appearance module implemented by *Spatial MLP* and *Integrated Basis MLP*, a geometry module achieved by *SDF MLP*, a physically-based forward rendering, and a neural photometric calibration module. In forward rendering, the complex appearance reflectance properties are explicitly decomposed to (a) diffuse albedo, (b) light integral, and (c) BRDFs integral. Closed-form solutions are provided for the first two terms under learnable Spherical Harmonics (SH) lighting. For the BRDFs integral term, a low-rank prior is introduced to learn a set of neural bases which describe the complex material characteristics of facial skin.

3. Preliminaries

We resort to PBR [18] to explicitly predict complex skin reflectances, where outgoing radiance L_o at surface position x along direction ω_o can be expressed as:

$$L_o(x, \omega_o) = \int_{S^2} L_i(\omega_i) f(x, \omega_i, \omega_o) (\omega_i \cdot \mathbf{n})^+ d\omega_i, \quad (1)$$

where L_o is calculated by an integral of the product of incident radiance L_i from direction ω_i , skin BRDF f , and half-cosine function $(\omega_i \cdot \mathbf{n})^+$ between surface normal \mathbf{n} and ω_i over sphere S^2 . Single-bounce direct illumination without shadows is assumed, and L_i is thus independent to x . Besides, as a pixel-level subsurface scattering is assumed, L_o is only considered within a single, small region.

Spherical Harmonics (SH) [6] and Spherical Gaussians (SG) [56] are commonly exploited as efficient representations of incident illumination. SH is selected in this study for its completeness and compactness. More importantly, it can implicitly facilitate specular separation with lower orders [45]. Accordingly, unknown L_i can be approximated by spherical basis functions $Y_{\ell m}$ of the first ten orders multiplied by corresponding learnable coefficients $c_{\ell m}$:

$$L_i(\omega_i) \approx \sum_{\ell=0}^{10} \sum_{m=-\ell}^{\ell} c_{\ell m} Y_{\ell m}(\omega_i). \quad (2)$$

f is modeled as two components, akin to prior art [13, 39]:

$$f(x, \omega_i, \omega_o) = \frac{\mathbf{a}(x)}{\pi} + \rho f_s(x, \omega_i, \omega_o), \quad (3)$$

where the left term is a diffuse component (Lambertian) that is only determined by albedo $\mathbf{a}(x)$, with the residual term mainly accounting for the glossy reflection, which is represented by $f_s(x, \omega_i, \omega_o)$ with a scale factor ρ indicating

the specular intensity. Consequently, the rendering equation (Eq. (1)) can be split into two terms:

$$L_o(x, \omega_o) = \underbrace{\frac{\mathbf{a}(x)}{\pi} \int_{S^2} L_i(\omega_i) (\omega_i \cdot \mathbf{n})^+ d\omega_i}_{\text{diffuse term } L_d} + \rho \underbrace{\int_{S^2} L_i(\omega_i) f_s(x, \omega_i, \omega_o) (\omega_i \cdot \mathbf{n})^+ d\omega_i}_{\text{residual (specular) term } L_s}. \quad (4)$$

4. Method

We incorporate neural BRDFs into PBR to collaboratively learn complex appearance properties (*i.e.*, diffuse and specular albedo) and sophisticated geometry attributes. The proposed *NeuFace* performs 3D face capturing from multi-view images, which only requires additional camera poses. As Fig. 2 illustrates, it exploits a *Spatial MLP* and an *Integrated Basis MLP* to learn the skin BRDFs and simulates the physically-correct light transport by explicitly decomposing it into diffuse albedo, light integral, and BRDFs integral. Facial appearance modeling can be formulated by Eq. (4), where we make great efforts to tackle the challenges caused by high computational cost and data demand during the learning procedure. For more accurate estimation, *NeuFace* further makes use of an SDF-based representation, *i.e.*, *ImFace* [61], as the facial geometry prior in forward mapping. The facial appearance and geometry field are jointly recovered by an end-to-end inverse rendering.

4.1. Specular Modeling

The specular term *i.e.* L_s with BRDF $f_s(x, \omega_i, \omega_o)$ is intractable to solve due to the view-dependent characteristic

of ω_o . In prior studies dealing with non-face neural rendering, the analytic BRDFs, *e.g.* microfacet models [31, 40, 56], have been explored for an approximated integral. Nevertheless, these explicit models can only recover a small set of materials while excluding facial skin which exhibits quite unique properties [32]. For a stronger representation ability, NeRFactor [59] employs the data-driven BRDFs learned from the real-world captured MERL database [27], but the live tissues like facial skin with spatially-varying BRDFs are extremely difficult to measure in vivo [15]. Moreover, heavy Monte Carlo sampling has to be performed in NeRFactor to solve the rendering equation, which is not practical in our case as the facial geometry is originally unknown.

In this study, we propose an alternative to render complicated facial highlights without external data and numerical integration. In particular, inspired by the split-integral approximation in real-time rendering [1], the specular is decomposed, where L_s can be approximated by:

$$L_s \approx \varrho \int_{s^2} f_s(x, \omega_i, \omega_o) (\omega_i \cdot \mathbf{n})^+ d\omega_i \int_{s^2} D(h) L_i(\omega_i) d\omega_i, \quad (5)$$

where $D(h)$ is a distribution function suggesting how light bounces at the half vector $h = \frac{\omega_i + \omega_o}{\|\omega_i + \omega_o\|_2}$. The entire equation is flanked by the integral terms of material and light, respectively, which are solved individually in the subsequent.

Material Integral. The first split integral term in Eq. (5) is only related to the material property, which is parameterized by a learnable network in our method for a higher representation capacity to facial skin. Considering that the 2D observations can only constrain the integrated values, rather than modeling $f_s(x, \omega_i, \omega_o)$ as in [59, 60], we directly formulate the entire integral as a smoother function F :

$$\int_{s^2} f_s(x, \omega_i, \omega_o) (\omega_i \cdot \mathbf{n})^+ d\omega_i = F(x, \omega_o, \mathbf{n}). \quad (6)$$

It should be noted that an enormous number of ground-truth samples are needed to achieve robust fitting with MLP for such a spatially-varying 9-variable function, which is not available in practice. Based on the previous studies on face appearance measurements [12, 51], we put forward a key assumption that all the facial surface positions of an individual share a similar specular structure, and the spatially-varying properties can thus be represented by diverse linear combinations of only a few (low-rank) learnable BRDFs:

$$f_s(x, \omega_i, \omega_o) \approx \sum_{j=1}^k c_j(x) b_j(\omega_i, \omega_o), \quad (7)$$

where $\{b_j(\omega_i, \omega_o)\}_{j=1}^k$ denotes k global space-independent BRDF bases, and $\mathbf{c}(x) = [c_1(x), c_2(x), \dots, c_k(x)]^T$ means the corresponding linear coefficients for each surface position x . As such, the material integral can be formulated as:

$$\int_{s^2} f_s(x, \omega_i, \omega_o) (\omega_i \cdot \mathbf{n})^+ d\omega_i = \mathbf{c}(x) \cdot \mathbf{B}(\omega_o, \mathbf{n}), \quad (8)$$

where $\mathbf{B}(\omega_o, \mathbf{n}) = [B_1, B_2, \dots, B_k]^T$ indicates k integrated BRDF basis b_j multiplied by half-cosine function $(\omega_i \cdot \mathbf{n})^+$:

$$B_j(\omega_o, \mathbf{n}) = \int_{s^2} b_j(\omega_i, \omega_o) (\omega_i \cdot \mathbf{n})^+ d\omega_i, j = 1, 2, \dots, k. \quad (9)$$

We leverage an MLP, namely *Integrated Basis MLP*, to fit $\mathbf{B}(\omega_o, \mathbf{n})$. $\omega_o \cdot \mathbf{n}$ is also fed into it to account for the Fresnel effects. Meanwhile, coefficient vector $\mathbf{c}(x)$ is predicted by a *Spatial MLP* with x as input. Note that facial skin is dielectric and cannot color the highlight; we thus use a single channel B_j to represent the monochromatic facial BRDFs.

As depicted in Fig. 2 (c), the low-rank prior globally restricts the solution space for all facial specularities without enforcing any spatial smoothness [53] or clustering of sampling positions [31]. With such a prior, the material integral term is much easier to fit and interpolate, producing impressive results with moderate training data.

Light Integral. For glossy surfaces, the reflected energy is mostly concentrated near the reflection direction, denoted by ω_r , as illustrated in Fig. 2 (b). Accordingly, we model $D(h)$ with a 3-dimensional normalized Spherical Gaussian distribution (vMF), centered at ω_r with a concentration parameter κ indicating the shininess:

$$D(h) \approx \text{vMF}(\omega_i; \omega_r, \kappa). \quad (10)$$

κ is predicted by *Spatial MLP*, and a larger value suggests a sharper underlying BRDF lobe pointing to ω_r , making facial skin look shinier. As the proof in our *Supplementary Material* shows, the second split integral in Eq. (5) which describes specular light transport can be approximated by:

$$\int_{s^2} D(h) L_i(\omega_i) d\omega_i \approx \sum_{\ell=0}^{\ell} \sum_{m=-\ell}^{\ell} e^{-\frac{\ell(\ell+1)}{2\kappa}} c_{\ell m} Y_{\ell m}(\omega_r). \quad (11)$$

The entire specular term in the neural BRDFs (Eq. (4)) can then be efficiently and differentially computed as:

$$L_s \approx \varrho (\mathbf{c}(x) \cdot \mathbf{B}(\omega_o, \mathbf{n})) \sum_{\ell=0}^{\ell} \sum_{m=-\ell}^{\ell} e^{-\frac{\ell(\ell+1)}{2\kappa}} c_{\ell m} Y_{\ell m}(\omega_r). \quad (12)$$

4.2. Diffuse Modeling

$\text{qrefsh}_l, \text{thedi f fuseradiance} L_d$ at a surface position x can be rewritten as: Based on Eq. (2), diffuse radiance L_d at surface position x can be rewritten as:

$$L_d(x) = \frac{\mathbf{a}(x)}{\pi} \sum_{\ell=0}^{\ell} \sum_{m=-\ell}^{\ell} c_{\ell m} \int_{s^2} Y_{\ell m}(\omega_i) (\omega_i \cdot \mathbf{n})^+ d\omega_i. \quad (13)$$

Following the Funk-Hecke theorem [2], the convolution of $(\omega_i \cdot \mathbf{n})^+$ with spherical harmonics $Y_{\ell m}$ can be analytically calculated by:

$$\int_{s^2} Y_{\ell m}(\omega_i) (\omega_i \cdot \mathbf{n})^+ d\omega_i \approx \Lambda_{\ell m} Y_{\ell m}(\mathbf{n}), \quad (14)$$

where

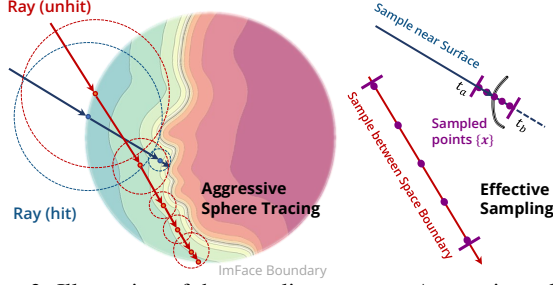


Figure 3. Illustration of the sampling strategy. Aggressive sphere tracing is performed to accelerate the sampling process for volume rendering. For the rays unhit the surface, the sphere defined by the geometry prior of *ImFace* is used to constrain the sampling space.

$$\Lambda_{\ell m} = \begin{cases} \frac{2\pi}{3}, & \text{if } \ell = 1, \\ \frac{(-1)^{\ell/2+1}\pi}{2^{\ell-1}(\ell-1)(\ell+2)} \binom{\ell}{\ell/2}, & \text{if } \ell \text{ is even,} \\ 0, & \text{if } \ell \text{ is odd.} \end{cases} \quad (15)$$

Refer to [2] for details. In *NeuFace*, diffuse albedo $\mathbf{a}(x)$ is modeled by *Spatial MLP*. With learnable coefficients $c_{\ell m}$ of SH lighting, the diffuse term can be directly computed as:

$$L_d(x) \approx \frac{\mathbf{a}(x)}{\pi} \sum_{\ell=0}^{\ell} \sum_{m=-\ell}^{\ell} \Lambda_{\ell m} c_{\ell m} Y_{\ell m}(\mathbf{n}). \quad (16)$$

In summary, the appearance component of *NeuFace* is composed of a *Spatial MLP*: $x \mapsto (\varrho, \mathbf{c}, \kappa, \mathbf{a})$, an *Integrated Basis MLP*: $(\omega_o, \mathbf{n}, \omega_o \cdot \mathbf{n}) \mapsto \mathbf{B}$, and the learnable environment light coefficients *i.e.* $c_{\ell m}$. Radiance L_o can finally be estimated through Eq. (12) and Eq. (16).

4.3. Geometry Modeling

To achieve end-to-end training, a differentiable geometry representation is essential. Similar to most neural rendering practices [48, 54–56], a neural SDF can be used to implicitly define the facial geometry. Here, we leverage *ImFace* [61] as a direct facial SDF prior to facilitate sampling and training. To capture the geometries outside the prior distribution, we fine-tune *ImFace* $\mathcal{I} : x \mapsto \text{SDF}$ and introduce a neural displacement field $\mathcal{D}(x)$ to correct the final results:

$$\text{SDF}(x) = \mathcal{I}(x) + \mathcal{D}(x). \quad (17)$$

Based on the property of SDF, surface normal \mathbf{n} can be extracted by auto-gradient: $\mathbf{n} = \nabla \text{SDF}(x)$.

4.4. Sampling and Rendering

In order to cope with multi-layered facial skin, we volumetrically render the radiance value as in *VoISDF* [54]. For a ray $x(t)$ emanated from a camera position $\mathbf{o} \in \mathbb{R}^3$ in direction ω_o , defined by $x(t) = \mathbf{o} + t\omega_o, t > 0$, the density function is defined as $\sigma(t) = \beta^{-1}\Psi_{\beta}(-\text{SDF}(x(t)))$. Ψ_{β} is the cumulative distribution function of the Laplace distribution with zero mean and learned scale parameters β . As such, the integrated radiance for each ray is evaluated by:

$$\mathbf{I}(\mathbf{o}, \omega_o) = \int_{t_a}^{t_b} L_o(x(t), \omega_o) \sigma(t) T(t) dt, \quad (18)$$

where $T(t) = \exp(-\int_0^t \sigma(s) ds)$ is the transparency.

For rendering acceleration, instead of densely sampling points along the ray as in [54], we first perform aggressive sphere tracing [23] to quickly find position t_0 near the surface (with a 0.05mm threshold), and 32 points are then uniformly sampled from $t \in [t_a, t_b]$, where $t_a = t_0 - 0.5\text{mm}$, $t_b = t_0 + 0.5\text{mm}$. As Fig. 3 shows, for the rays that do not hit the surface, we sparsely sample points within sphere Ω defined by the geometry prior from *ImFace* [61] and calculate the accumulation, bypassing the mask loss [55] which requires heavy sampling. By combining volume and surface rendering, it makes a good trade-off balancing the rendering quality, geometry precision, and sampling efficiency.

Neural Photometric Calibration. To auto-calibrate the inconsistent color response and white balance among cameras, we apply per-image linear mapping matrices $\mathbf{A}_n \in \mathbb{R}^{3 \times 3}$ to the rendered radiance values:

$$\mathbf{I}_n = \mathbf{A}_n \mathbf{I}, \quad (19)$$

where \mathbf{A}_n is predicted by a lightweight MLP conditioned on learnable per-image embedding.

4.5. Loss Functions

NeuFace is trained with a compound critic to learn accurate appearance and geometry properties:

$$\mathcal{L} = \mathcal{L}_{RGB} + \mathcal{L}_{white} + \mathcal{L}_{spec} + \mathcal{L}_{geo}. \quad (20)$$

Let P_n be the pixels of the n -th multi-view image, $\bar{\mathbf{I}}_n$. Each term of \mathcal{L} is described as follows, where λ denotes the trade-off hyper-parameter.

Reconstruction Loss. It is used to supervise the rendered 2D faces being close to the real observations:

$$\mathcal{L}_{RGB} = \lambda_1 \sum_n \sum_{\omega_o \in P_n} |\mathbf{I}_n - \bar{\mathbf{I}}_n|. \quad (21)$$

Light Regularization. We assume a nearly white environment light in capturing, achieved by:

$$\mathcal{L}_{white} = \lambda_2 \sum_n \sum_{\omega_o \in P_n} |L_i - \bar{L}_i|, \quad (22)$$

where \bar{L}_i is calculated by averaging RGB channels.

Facial Specular Regularization. Only a small ratio of the incident light on facial surfaces (roughly 6% [46]) reflects directly, and we thus penalize the specular energy by:

$$\mathcal{L}_{spec} = \lambda_3 \sum_n \sum_{\omega_o \in P_n} L_s. \quad (23)$$

Geometry Losses. Following [61], the embedding regularization, Eikonal regularization, and a new residual constraint are exploited for accurate geometry modeling:

$$\mathcal{L}_{geo} = \lambda_4 \|\mathbf{z}\|^2 + \lambda_5 \sum_{x \in \Omega} \|\|\nabla \text{SDF}(x)\| - 1\| + \lambda_6 \sum_{x \in \Omega} |\mathcal{D}(x)|. \quad (24)$$

where \mathbf{z} represents the embeddings of *ImFace* [61].




Figure 4. Comparison of novel view synthesis with Ref-NeRF [47], PhySG [56], VoISDF [54] and DIFFREC [31]. Zoom in for a better view. *NeuFace* captures faces with much richer skin details and high-fidelity highlights.

w/ our calibration	F.	S.N. ↓	Subject 1 (Rich Reflection)				Subject 2 (Moderate Reflection)				Subject 3 (Low Reflection)			
			PSNR ↑	SSIM ↑	LPIPS ↓	Chamfer ↓	PSNR ↑	SSIM ↑	LPIPS ↓	Chamfer ↓	PSNR ↑	SSIM ↑	LPIPS ↓	Chamfer ↓
NeRF [29]	✗	256	13.98	0.712	.1378	3.666	13.67	0.724	.1493	1.910	-	-	-	-
Ref-NeRF [47]	✗	256	19.83	0.845	.1134	2.042	20.01	0.846	.1410	1.782	17.33	0.820	.1381	3.682
VoISDF [54]	✗	610	29.03	0.941	.0502	0.516	28.40	0.937	.0730	0.589	26.74	0.926	.0580	0.626
PhySG [56]	✓	40.2	27.77	0.928	.0611	0.864	26.36	0.913	.0864	0.977	23.07	0.902	.0735	1.776
DIFFREC* [31]	✓	39.5	27.72	0.942	.0244	0.546	27.61	0.934	.0353	0.606	27.02	0.930	.0372	0.727
<i>NeuFace</i>	✓	35.8	31.25	0.958	.0237	0.447	31.87	0.961	.0314	0.481	30.92	0.953	.0240	0.533

'F.' refers to factorization. 'S.N.' refers to the number of sampled points per ray. DIFFREC* is implemented with neural surface. NeRF fails to converge on Subject 3.

Table 1. Quantitative comparison with NeRF [29], Ref-NeRF [47], VoISDF [54], PhySG [56], and DIFFREC [31]. *NeuFace* outperforms the counterparts by a large margin in both appearance and geometry metrics. Red, orange, and yellow: the best, second-best, and third-best.

5. Experiments

We extensively conduct experiments for both subjective and objective evaluation and compare our *NeuFace* with the state-of-the-art neural rendering methods. Ablation studies are performed to validate the specifically designed modules. For more implementation details on the network architecture and training procedure, please refer to the *Supplement*. FaceScape [52] is adopted, and it is a large-scale 3D face dataset including high-quality multi-view images of 4K resolution, which is collected from 359 real subjects with 20 expressions. During capturing, around 58 cameras with different types and settings are equally distributed on a spherical structure with uniform illumination. We use the authorized data from 3 individuals for model evaluation, with 43

Shading Model	PSNR ↑	SSIM ↑	LPIPS ↓	Chamfer ↓
SH+Phong Model	29.31	0.951	.0300	0.556
Cubemap+Disney BRDF	30.45	0.954	.0242	0.524
Cubemap+Neural Bases	30.82	0.958	.0256	0.426
SH+Neural Bases (Ours)	31.25	0.958	.0237	0.447

Table 2. Quantitative ablation study on the shading model.

images for training and 11 images for testing. The images are downsampled to 1K resolution.

5.1. Comparison

To validate the ability of *NeuFace* for simultaneous appearance and geometry capturing, we compare it with the state-of-the-art approaches to view synthesis [29,47], geometry reconstruction [54] and appearance factorization [31,




Figure 5. Comparison of geometry reconstruction to PhySG [56], VolSDF [54], and DIFFREC [31]. *NeuFace* achieves lower errors.




Figure 6. Comparison with the Phong [35] and Disney BRDFs [5].

Designs	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	Chamfer \downarrow
w/o low rank (directly fitting)	29.45	0.950	.0304	0.553
w/o low rank (256-D feature)	29.43	0.951	.0302	0.544
w/o calibration	24.04	0.943	.0284	0.451
w/o geometry prior	28.11	0.942	.0452	0.647
w/o specular loss	31.03	0.957	.0248	0.453
<i>NeuFace</i>	31.25	0.958	.0237	0.447

Table 3. Quantitative ablation study on the low-rank prior, neural calibration, geometry prior, and specular loss.

[56]. Considering that neural rendering is extremely sensitive to camera inconsistency, we apply the proposed neural photometric calibration to all the counterparts for fair comparison. Because DIFFREC [31] fails to reconstruct all the faces with explicit tetrahedral grids, we implement a more stable neural surface version instead. PSNR, SSIM [50] and LPIPS [58] are used as the metrics to evaluate the view synthesis quality, and the Chamfer distance to the ground-truth 3D model is used to measure the geometry accuracy.

Novel View Synthesis. Fig. 4 displays the results. It can be seen that as a recent advanced view-synthesis method, Ref-NeRF [47] fails to recover facial appearances, mainly due to the strong shape-appearance ambiguity on faces. Designed for geometry reconstruction, VolSDF [54] sacrifices the appearance quality, resulting in blurred face renderings. PhySG [56] and DIFFREC [31] are also physically-based neural rendering methods as ours, which aim at end-to-end appearance and geometry recovery. However, limited by the inadequate capacity of the analytic BRDFs, they can only recover rough facial appearances with obvious artifacts. By contrast, *NeuFace* achieves realistic face appearances with much richer skin details and high-fidelity highlights, *e.g.*,

the areas around lips, glabellar, and nose. The quantitative results in Tab. 1 confirm its effectiveness, which are evaluated on smile subjects 1, neutral subjects 2 and 3.

Geometry Reconstruction. Fig. 5 makes comparison of geometry reconstruction with PhySG [56], VolSDF [54], DIFFREC [31], and ground-truth, measured by color-coded distance. *NeuFace* reports the lowest errors and even outperforms the geometry-oriented method, *i.e.*, VolSDF, suggesting that our PBR framework benefits not only facial appearances but also facial shapes. The quantitative results in Tab. 1 further validate this claim.

Tab. 1 demonstrates the efficiency of our rendering strategy (Section 4.4) in terms of average sampling number per ray. Our method requires fewer sampling points than general volume rendering and is on par with surface rendering, which needs additional sampling to minimize mask loss.

5.2. Ablation Study

The following experiments are performed on Subject 1 as the face images exhibit more complex highlights.

On Shading Model. Our neural BRDFs module is designed for a higher representation capability to facial skin, and we carry out an ablation study by replacing it with the well-reputed Phong [35] and Disney BRDFs [5]. In particular, the specular term in the Phong BRDFs is approximated as in [37], controlled by a single shininess parameter. The Disney BRDFs are implemented by using the code provided by DIFFREC [31]. We show novel view rendering results in Fig. 6, which demonstrate that the proposed neural BRDFs module can capture high-fidelity reflections of facial skin, such as the sheen around grazing angles, while such light transport cannot be well modeled by the Phong and Disney BRDFs. Tab. 2 also reveals the higher representation ability of neural BRDFs in learning sophisticated skin reflections.

Moreover, we show factorization and relighting results in Fig. 7, where SH lighting is compared with Cubemap. As shown, Cubemap indeed achieves comparable results of novel view synthesis and even slightly better geometries by combining our other designs. However, in our unconstrained settings, it fails to accurately relight faces under new environment maps. In contrast, *NeuFace* can well adapt to arbitrary lighting conditions.

On Neural Integrated Basis. The low-rank prior is introduced to constrain the solution space and learn accurate neural BRDFs with moderate observations. We compare it with two variants: 1) directly fitting the 9-variable material integral function formulated in Eq. (6) (“directly fitting”); and 2) following IDR [55], we use *Spatial MLP* to predict a 256-dimensional feature and then feed it into *Integrated Basis MLP* to learn the material integral (“256-D feature”). Tab. 3 shows that without low-rank regularization, both variants perform significantly worse. Besides, as shown in our *Supplement*, the basis number $k = 3$ yields the

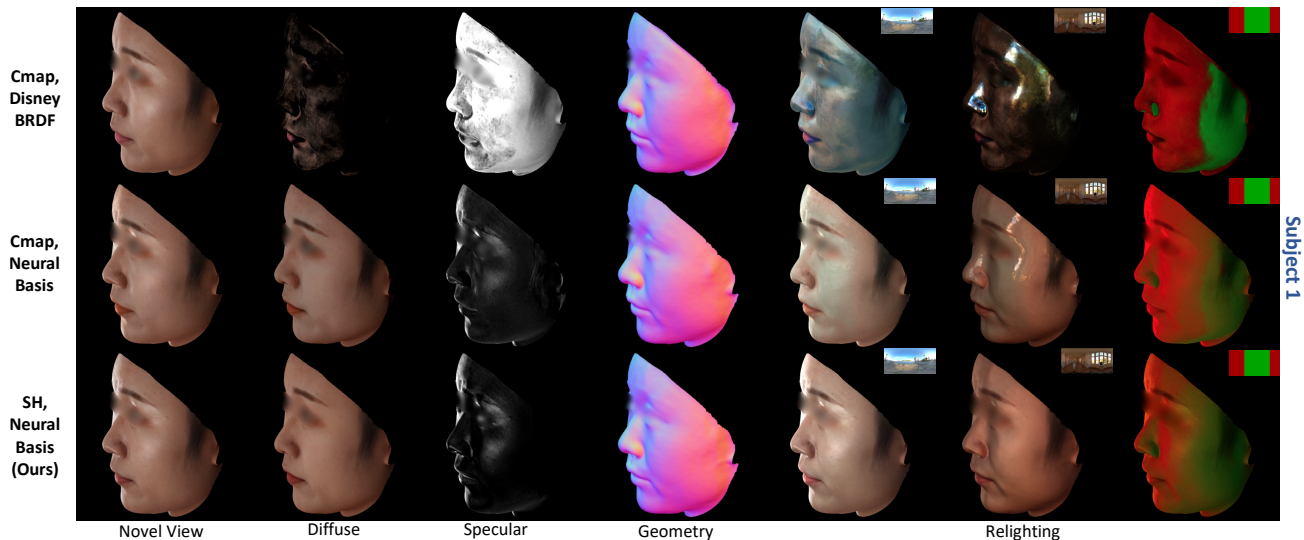


Figure 7. Ablation on shading models. Ours (bottom) compares with improved DIFFREC [31] (upper) and Cubemap lighting (middle).




Figure 8. Comparison with DIFFREC on the DTU database [17]. Upper row: DIFFREC; bottom row: ours.

best scores. Increasing the basis number makes the underlying BRDFs more difficult to fit with limited observations, leading to performance declines for both metrics.

Miscellaneous Ablation. We analyze the neural photometric calibration, the *ImFace* prior, and the specular energy loss. Table 3 demonstrates the importance of each component in achieving realistic face rendering and accurate geometry recovery. Neural calibration remarkably benefits the inverse rendering procedure by accounting for camera inconsistency. The *ImFace* prior significantly improves the recovered geometry and appearance quality by constraining sampling within a valid 3D space during forward rendering. Additionally, the specular energy loss includes real-world empirical evidence on face skin and stabilizes the factorization by preventing energy influx into the specular term.

5.3. Extension to Common Objects

To validate the generalization ability of *NeuFace*, we extend it to common objects in the DTU database [17], which contains large-scale real-world multi-view captures. In this evaluation, the classical IDR [55] is modified with our rendering pipeline. As Fig. 8 shows, our model achieves more impressive reconstruction and factorization qualities than DIFFREC [31]. More results are in the *Supplement*.

6. Limitation

The fidelity of the appealing rendering results achieved in this study can be further improved since *NeuFace* mainly focuses on reflection modeling of complex facial skin without explicitly tackling the more challenging issue of subsurface scattering, and only applies a simplified shading model instead. Moreover, *NeuFace* currently delivers a static 3D face rather than a drivable one, and although the geometry model, *i.e.* *ImFace*, is a non-linear morphable model and theoretically supports controllable expression editing, more work is needed for decent performance.

7. Conclusion

This paper presents a novel 3D neural rendering model, *i.e.*, *NeuFace*, to simultaneously capture realistic facial appearances and geometries from only multi-view images. To cope with complex facial skin reflectance, we bond physically based rendering with neural BRDFs for more accurate and physically-meaningful 3D representations. Moreover, to facilitate the optimization of the underlying BRDFs, we introduce a split-integral technique as well as a simple yet new low-rank prior, which significantly improve the recovering performance. Extensive experiments demonstrate the superiority of *NeuFace* in face rendering and its decent generalization ability to common objects.

Acknowledgements

This work is partly supported by the National Natural Science Foundation of China (No. 62176012, 62202031, 62022011), the Beijing Natural Science Foundation (No. 4222049), the Research Program of State Key Laboratory of Software Development Environment (SKLSDE-2021ZX-04), and the Fundamental Research Funds for the Central Universities.

References

- [1] Tomas Akenine-Moller, Eric Haines, and Naty Hoffman. *Real-time rendering*. AK Peters/crc Press, 2019. 2, 4
- [2] Ronen Basri and David W Jacobs. Lambertian reflectance and linear subspaces. *IEEE TPAMI*, 25(2):218–233, 2003. 4, 5
- [3] Thabo Beeler, Bernd Bickel, Paul Beardsley, Bob Sumner, and Markus Gross. High-quality single-shot capture of facial geometry. In *ACM SIGGRAPH*. 2010. 2
- [4] Derek Bradley, Wolfgang Heidrich, Tiberiu Popa, and Alla Sheffer. High resolution passive facial performance capture. In *ACM SIGGRAPH*, 2010. 2
- [5] Brent Burley and Walt Disney Animation Studios. Physically-based shading at Disney. In *ACM SIGGRAPH*, 2012. 7
- [6] Wenzheng Chen, Huan Ling, Jun Gao, Edward Smith, Jaakko Lehtinen, Alec Jacobson, and Sanja Fidler. Learning to predict 3D objects with an interpolation-based differentiable renderer. *NeurIPS*, 2019. 3
- [7] Robert L Cook and Kenneth E. Torrance. A reflectance model for computer graphics. *ACM TOG*, 1(1):7–24, 1982. 2
- [8] Paul Debevec, Tim Hawkins, Chris Tchou, Haarm-Pieter Duiker, Westley Sarokin, and Mark Sagar. Acquiring the reflectance field of a human face. In *ACM SIGGRAPH*, 2000. 1, 2
- [9] Bernhard Egger, William AP Smith, Ayush Tewari, Stefanie Wuhler, Michael Zollhoefer, Thabo Beeler, Florian Bernard, Timo Bolkart, Adam Kortylewski, Sami Romdhani, Christian Theobalt, Volker Blanz, and Thomas Vetter. 3D morphable face models—past, present, and future. *ACM TOG*, 39(5):1–38, 2020. 2
- [10] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *CVPR*, 2022. 2
- [11] Abhijeet Ghosh, Graham Fyffe, Borom Tunwattapanong, Jay Busch, Xueming Yu, and Paul Debevec. Multiview face capture using polarized spherical gradient illumination. In *ACM SIGGRAPH Asia*, 2011. 2
- [12] Abhijeet Ghosh, Tim Hawkins, Pieter Peers, Sune Frederiksen, and Paul Debevec. Practical modeling and acquisition of layered facial reflectance. In *ACM SIGGRAPH Asia*. 2008. 2, 4
- [13] Paulo Gotardo, Jérémy Riviere, Derek Bradley, Abhijeet Ghosh, and Thabo Beeler. Practical dynamic facial appearance modeling and acquisition. *ACM TOG*, 37(6):1–13, 2018. 3
- [14] Jon Hasselgren, Nikolai Hofmann, and Jacob Munkberg. Shape, light & material decomposition from images using Monte Carlo rendering and denoising. In *NeurIPS*, 2022. 2
- [15] Takanori Igarashi, Ko Nishino, and Shree K Nayar. The appearance of human skin. 2005. 4
- [16] Henrik Wann Jensen, Stephen R Marschner, Marc Levoy, and Pat Hanrahan. A practical model for subsurface light transport. In *ACM SIGGRAPH*, 2001. 2
- [17] Rasmus Jensen, Anders Dahl, George Vogiatzis, Engin Tola, and Henrik Aanæs. Large scale multi-view stereopsis evaluation. In *CVPR*, 2014. 8
- [18] James T Kajiya. The rendering equation. In *ACM SIGGRAPH*, 1986. 3
- [19] Christos Kampouris, Stefanos Zafeiriou, and Abhijeet Ghosh. Diffuse-specular separation using binary spherical gradient illumination. In *EGSR*, 2018. 2
- [20] Brian Karis and Epic Games. Real shading in Unreal Engine 4. *Proc. Physically Based Shading Theory Practice*, 4(3):1, 2013. 2
- [21] Oliver Klehm, Fabrice Rousselle, Marios Papas, Derek Bradley, Christophe Hery, Bernd Bickel, Wojciech Jarosz, and Thabo Beeler. Recent advances in facial appearance capture. In *CGF*, 2015. 2
- [22] 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 TOG*, 41(4):1–16, 2022. 2
- [23] Shaohui Liu, Yinda Zhang, Songyou Peng, Boxin Shi, Marc Pollefeys, and Zhaopeng Cui. DIST: Rendering deep implicit signed distance function with differentiable sphere tracing. In *CVPR*, 2020. 5
- [24] Linjie Lyu, Ayush Tewari, Thomas Leimkühler, Marc Habermann, and Christian Theobalt. Neural radiance transfer fields for relightable novel-view synthesis with global illumination. In *ECCV*, 2022. 2
- [25] Wan-Chun Ma, Tim Hawkins, Pieter Peers, Charles-Felix Chabert, Malte Weiss, and Paul E Debevec. Rapid acquisition of specular and diffuse normal maps from polarized spherical gradient illumination. *ECRT*, 2007(9):183–194, 2007. 2
- [26] Ricardo Martin-Brualla, Noha Radwan, Mehdi SM Sajjadi, Jonathan T Barron, Alexey Dosovitskiy, and Daniel Duckworth. NeRF in the Wild: Neural radiance fields for unconstrained photo collections. In *CVPR*, 2021. 1, 2
- [27] Wojciech Matusik, Hanspeter Pfister, Matt Brand, and Leonard McMillan. A data-driven reflectance model. *ACM TOG*, 22(3):759–769, 2003. 4
- [28] Ben Mildenhall, Peter Hedman, Ricardo Martin-Brualla, Pratul P Srinivasan, and Jonathan T Barron. NeRF in the Dark: High dynamic range view synthesis from noisy raw images. In *CVPR*, 2022. 1, 2
- [29] B Mildenhall, PP Srinivasan, M Tancik, JT Barron, R Ramamoorthi, and R Ng. NeRF: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020. 2, 6
- [30] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM TOG*, 41(4):102:1–102:15, 2022. 2
- [31] 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 *CVPR*, 2022. 2, 4, 6, 7, 8
- [32] H Murakami, T Horii, N Tsumura, and Y Miyake. Measurement and simulation of 3D gonio spectral reflectance of skin surface. *DBJ*, 93:21–26, 2002. 4

- [33] Michael Oechsle, Songyou Peng, and Andreas Geiger. UNISURF: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. In *CVPR*, 2021. 1, 2
- [34] Matt Pharr, Wenzel Jakob, and Greg Humphreys. *Physically based rendering: From theory to implementation*. Morgan Kaufmann, 2016. 2
- [35] Bui Tuong Phong. Illumination for computer generated pictures. *Communications of the ACM*, 18(6):311–317, 1975. 7
- [36] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-NeRF: Neural radiance fields for dynamic scenes. In *CVPR*, 2021. 1, 2
- [37] Ravi Ramamoorthi and Pat Hanrahan. Frequency space environment map rendering. In *ACM SIGGRAPH*, 2002. 7
- [38] Christian Reiser, Songyou Peng, Yiyi Liao, and Andreas Geiger. KiloNeRF: Speeding up neural radiance fields with thousands of tiny MLPs. In *CVPR*, 2021. 1, 2
- [39] Jérémy Riviere, Paulo Gotardo, Derek Bradley, Abhijeet Ghosh, and Thabo Beeler. Single-shot high-quality facial geometry and skin appearance capture. *ACM TOG*, 39(4):81:1–81:12, 2020. 2, 3
- [40] Pratul P Srinivasan, Boyang Deng, Xiuming Zhang, Matthew Tancik, Ben Mildenhall, and Jonathan T Barron. NeRV: Neural reflectance and visibility fields for relighting and view synthesis. In *CVPR*, 2021. 4
- [41] Tiancheng Sun, Kai-En Lin, Sai Bi, Zexiang Xu, and Ravi Ramamoorthi. Nelf: Neural light-transport field for portrait view synthesis and relighting. *EGSR*, 2021. 2
- [42] Blago Taskov, Eugene d’Eon, Morteza Ramezani, and Simon Yuen. Realistic digital human rendering with Omniverse RTX Renderer. <https://www.nvidia.com/en-us/on-demand/session/siggraph2021-sigg21-s-09/>. 1
- [43] Ayush Tewari, Justus Thies, Ben Mildenhall, Pratul Srinivasan, Edgar Tretschk, Wang Yifan, Christoph Lassner, Vincent Sitzmann, Ricardo Martin-Brualla, Stephen Lombardi, Tomas Simon, Christian Theobalt, Matthias Niessner, Jonathan T Barron, Gordon Wetzstein, Michael Zollhoefer, and Vladislav Golyanik. Advances in neural rendering. In *CGF*, 2022. 2
- [44] Kenneth E Torrance and Ephraim M Sparrow. Theory for off-specular reflection from roughened surfaces. *JOSA*, 57(9):1105–1114, 1967. 2
- [45] Borom Tunwattanapong, Graham Fyffe, Paul Graham, Jay Busch, Xueming Yu, Abhijeet Ghosh, and Paul Debevec. Acquiring reflectance and shape from continuous spherical harmonic illumination. *ACM TOG*, 32(4):1–12, 2013. 3
- [46] Tuchin Valery. Tissue optics: Light scattering methods and instruments for medical diagnosis. *SPIE Tutorial Texts in Optical Engineering*, 2000. 5
- [47] Dor Verbin, Peter Hedman, Ben Mildenhall, Todd Zickler, Jonathan T Barron, and Pratul P Srinivasan. Ref-NeRF: Structured view-dependent appearance for neural radiance fields. In *CVPR*, 2022. 1, 2, 6, 7
- [48] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. NeuS: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *NeurIPS*, 2021. 1, 2, 5
- [49] Yiqun Wang, Ivan Skorokhodov, and Peter Wonka. HF-NeuS: Improved surface reconstruction using high-frequency details. In *NeurIPS*, 2022. 1, 2
- [50] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE TIP*, 13(4):600–612, 2004. 7
- [51] Tim Weyrich, Wojciech Matusik, Hanspeter Pfister, Bernd Bickel, Craig Donner, Chien Tu, Janet McAndless, Jinho Lee, Addy Ngan, Henrik Wann Jensen, and Markus Gross. Analysis of human faces using a measurement-based skin reflectance model. *ACM TOG*, 25(3):1013–1024, 2006. 2, 4
- [52] Haotian Yang, Hao Zhu, Yanru Wang, Mingkai Huang, Qiu Shen, Ruigang Yang, and Xun Cao. Facescape: A large-scale high quality 3D face dataset and detailed riggable 3D face prediction. In *CVPR*, 2020. 6
- [53] Yao Yao, Jingyang Zhang, Jingbo Liu, Yihang Qu, Tian Fang, David McKinnon, Yanghai Tsin, and Long Quan. NeLF: Neural incident light field for physically-based material estimation. In *ECCV*, 2022. 4
- [54] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. *NeurIPS*, 2021. 1, 2, 5, 6, 7
- [55] Lior Yariv, Yoni Kasten, Dror Moran, Meirav Galun, Matan Atzmon, Basri Ronen, and Yaron Lipman. Multiview neural surface reconstruction by disentangling geometry and appearance. *NeurIPS*, 2020. 1, 2, 5, 7, 8
- [56] 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 *CVPR*, 2021. 2, 3, 4, 5, 6, 7
- [57] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. NeRF++: Analyzing and improving neural radiance fields. *arXiv preprint arXiv:2010.07492*, 2020. 2
- [58] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. 7
- [59] 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 TOG*, 40(6):1–18, 2021. 2, 4
- [60] Chuankun Zheng, Ruzhang Zheng, Rui Wang, Shuang Zhao, and Hujun Bao. A compact representation of measured BRDFs using neural processes. *ACM TOG*, 41(2):1–15, 2021. 2, 4
- [61] Mingwu Zheng, Hongyu Yang, Di Huang, and Liming Chen. ImFace: A nonlinear 3D morphable face model with implicit neural representations. In *CVPR*, 2022. 2, 3, 5