

GroomLight: Hybrid Inverse Rendering for Relightable Human Hair Appearance Modeling

Yang Zheng^{1,2} Menglei Chai² Delio Vicini² Yuxiao Zhou^{2,3} Yinghao Xu¹

Leonidas Guibas¹ Gordon Wetzstein¹ Thabo Beeler²

¹Stanford University ²Google ³ETH Zurich



Figure 1. *GroomLight* achieves high-fidelity reconstruction of human hair appearance from real-world images, enabling realistic rendering under diverse lighting conditions. Here, in each column, we present relighting results of a same subject under two different environments, using the appearance model reconstructed by *GroomLight*. One input view with the similar head pose is shown at the top right corner.

Abstract

We present *GroomLight*, a novel method for relightable hair appearance modeling from multi-view images. Existing hair capture methods struggle to balance photorealistic rendering with relighting capabilities. Analytical material models, while physically grounded, often fail to fully capture appearance details. Conversely, neural rendering approaches excel at view synthesis but generalize poorly to novel lighting conditions. *GroomLight* addresses this challenge by combining the strengths of both paradigms. It employs an extended hair BSDF model to capture primary light transport and a light-aware residual model to reconstruct the remaining details. We further propose a hybrid inverse rendering pipeline to optimize both components, enabling high-fidelity relighting, view synthesis,

and material editing. Extensive evaluations on real-world hair data demonstrate state-of-the-art performance of our method. Our project website is at: <https://syntec-research.github.io/GroomLight>.

1. Introduction

Hair capture is a long-standing problem in computer graphics and vision, and plays an essential role in human digitization. While substantial progress has been made in strand-accurate geometry reconstruction, hair appearance modeling and photorealistic relighting remain open challenges.

As part of efforts towards physics-based photorealistic rendering, analytical material models for hair have been extensively studied. Following this path, inverse rendering approaches [43] seek to match renderings to input images by

optimizing hair material properties, specifically the bidirectional scattering distribution function (BSDF). A key advantage of such approaches lies in their exceptional ability to generalize to novel viewpoints, unseen lighting, and altered hair geometry. However, even with sophisticated differentiable renderers, the reconstruction methods suffer from system errors stemming from factors such as imperfect geometry reconstruction, inaccurate lighting calibration, simplified light scattering models without diffraction and interference effects, etc. Consequently, these methods struggle to reproduce critical appearance details, hindering their successful application to real images.

Meanwhile, recent advances in image-based neural rendering have achieved impressive success in hair view synthesis. By employing primitives specifically tailored for strand structures, such as hybrid neural volumes [48], neural strand features [37], and linked cylindrical Gaussians [27], these methods exhibit a superior capability in capturing fine-grained appearance details from multi-view inputs. However, lacking a physics-grounded appearance parameterization, these primitives tend to entangle geometry, appearance, and associated illumination conditions, leading to overfitting of input images at the expense of generalization capability. Consequently, significant quality degradation is constantly observed in their relighting results.

Towards high-quality hair relighting, our motivation is to leverage the complementary strengths of both physically based material models and neural rendering representations. This involves employing analytical BSDF models to reconstruct the primary light transport while relying on neural rendering to capture subtle appearance details.

In this work, we introduce *GroomLight*, a hair appearance modeling method that reconstructs relightable hair appearance from OLAT (one-light-at-a-time) images. For the very first time, *GroomLight* enables high-fidelity and fully functional relighting for photorealistic human hair. Our appearance representation consists of two key components: 1) an *extended hair BSDF model* robust to natural hair color variations and imperfections in the reference hair geometry; and 2) a *light-aware residual model* that captures remaining details utilizing light- and view-conditioned dual-level spherical harmonics (SH) associated with a Gaussian-based hair representation. On top of that, we propose a hybrid inverse rendering pipeline that combines a physics-based differentiable renderer with neural image-based rendering to optimize the appearance.

We extensively evaluate our method on real-world hair data, demonstrating state-of-the-art relighting and view synthesis results under diverse lighting conditions. Moreover, our pipeline naturally supports versatile applications, including material editing and dynamic animation rendering, and seamlessly integrates with conventional CG workflows.

2. Related Work

2.1. Hair Appearance Modeling

Besides geometry capture from single- [5, 6, 8, 16] or multi-view [18, 28, 33, 34, 56], hair appearance modeling has also been a long-standing challenge. We provide a brief overview of physics- and neural-based methods that model the complex appearance of hair fibers.

Physically Based Models. Physically based hair appearance models describe light transport resulting from scattering on and within hair fibers. Building upon the foundational work on fiber reflectance [23], Marschner et al. [30] and subsequent research [11–13, 17, 25, 38, 62–64] have derived hair BSDFs that encapsulate hair scattering effects into analytical reflection and transmission lobes. These BSDF models, when used in conjunction with a path tracer [22], can generate highly realistic images of hair. Variations of hair BSDF models have also been employed to render textile fibers [1] and animal fur [53, 54]. While most models rely on geometric optics, incorporating wave-optical effects can further enhance realism [2, 50, 51]. Nevertheless, analytical physically based BSDFs may lack the expressive capacity and flexibility to perfectly match the appearance of real hair, potentially leading to inaccuracies in re-rendering crucial appearance features such as specular highlights and color variations.

Neural Representations. Neural networks offer a powerful means of representing high-fidelity human hair appearance. Generative models [9, 21, 36, 44, 49] have been trained for realistic hair synthesis. Neural rendering techniques [45] leverage scene representations such as neural radiance field (NeRF) [29, 31, 37], 3D Gaussian Splatting (3DGS) [24, 27, 35, 52], and implicit geometry and texture networks [10, 15, 41, 58] to learn portrait appearance from images. While neural models can produce high-quality, photorealistic results, they often lack robust relighting capabilities and can be challenging to animate or edit manually.

2.2. Inverse Rendering of Hair

Inverse rendering aims to estimate appearance parameters from visual observations using *analysis-by-synthesis* techniques. Sun et al. [43] estimate hair BSDF [11] parameters by measuring differences between sparse hair samples and synthetic ground truth observations. Their approach relies on an approximation of the observed color based on idealized lighting conditions and demonstrates results on synthetic data with accurate hair geometry reconstructed from ground truth segmentation, camera and lighting information. Chai et al. [7] propose an image-based rendering technique for hair relighting. Neural-based relighting [3, 32, 42, 55, 57] and Gaussian-based inverse rendering methods [4, 14, 20, 39, 40, 59] often combine neural representations with simplified physically based rendering

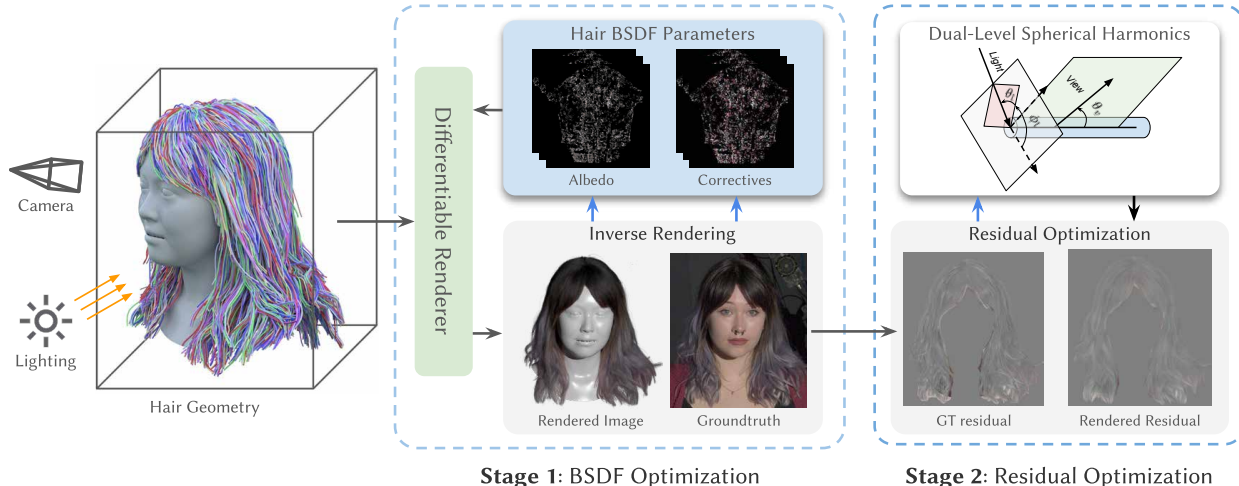


Figure 2. **Our hybrid inverse rendering pipeline** employs a two-stage optimization scheme. Given input OLAT images and the reconstructed hair geometry (here we visualize the 1/10 downsampled version), we first estimate the physical parameters of the extended hair BxDF model (Stage 1, §3.2), and then leverage a light-aware residual model to capture fine-grained appearance details (Stage 2, §3.3).

functions, which may not accurately model complex phenomena such as multiple scattering and shadowing. Luo et al. [27] leverage 3D Gaussians to capture hair appearance and design a scattering approximation method for relighting. However, their method remains fundamentally based on Gaussian splatting with baked-in appearance, and the approximation function may not fully capture the intricacies of real-world light scattering in human hair, potentially leading to suboptimal rendering results. In contrast, we propose a hybrid inverse rendering pipeline that fully leverages the strengths of both physically based and neural rendering representations.

3. Method

3.1. Overview

Our approach employs a two-level representation of hair appearance. This comprises: 1) an extended BxDF model (§3.2), which utilizes a physics-based parameterization to capture the primary appearance characteristics of real-world hair strands; and 2) an implicit light-aware residual model (§3.3) that captures light- and view-dependent details not fully encompassed by the analytical BxDF formulation.

Accordingly, our hybrid inverse rendering pipeline (§3.4) constructs these appearance models in two stages. First, inverse rendering optimization is performed within the analytical BxDF parameter space, which is extended to accommodate spatial variations in hair material and improve robustness against potential misalignment in the input hair geometry. Subsequently, leveraging a 3D Gaussian-based representation of individual strands, we optimize light-aware dual-level spherical harmonics to model the remaining appearance residuals.




Figure 3. **Impact of light scattering.** We compare the effects of varying the maximum number of ray bounces during path tracing (∞ means the integrator will not enforce any hard cutoff on the number of bounces). Our method employs 8 bounces to achieve high quality rendering at practical efficiency.

Input Data. Our method operates on multi-camera OLAT captures. We capture our input data in a light stage, featuring an array of 50 cameras and 331 light sources, uniformly distributed over the dome. From the complete set of frames \mathcal{I} , we exclude views where cameras do not prominently feature hair or where the subject is strictly back-lit (to mitigate artifacts arising from low exposure). The reference hair geometry \mathcal{S} and parametric head mesh \mathcal{H} are pre-reconstructed from the same OLAT input using the state-of-the-art strand-level hair capture technique of *Groom-Cap* [61], which typically yields 150K scalp-rooted strands. Additionally, we perform 2D segmentation using *SAM* [26] to obtain a set of hair masks \mathcal{M} for input frame images.

3.2. Extended Hair BxDF Model

Our strand-based hair BxDF is based on the analytical fiber reflectance model by Chiang et al. [11]. Each strand’s geometry is approximated as a cylindrical B-spline curve with a uniform radius, which we set to the average radius of human hair ($40\mu m$). For each point on the fiber, the perceived

reflectance is parameterized as a function of the incident θ_i and exitant θ_o lighting inclinations, their relative azimuth $\phi = \phi_o - \phi_i$, and the physical parameters Ψ of the material:

$$S(\theta_i, \theta_o, \phi; \Psi) = \sum_{p=0}^{\infty} M_p(\theta_i, \theta_o; \Psi) N_p(\phi; \Psi), \quad (1)$$

where the reflectance S is decomposed into separate lobes according to the path of light propagation p (i.e., $p \in \{R, TT, TRT, \dots\}$), which counts the reflections and transmissions in the fiber cylinder. Each lobe consists of M_p and N_p , i.e., the longitudinal and azimuthal scattering functions, respectively. Fig. 3 shows the light scattering effects in hair rendering, where single-bounce illumination (only $p = 0$ is evaluated) leads to degenerated results, while more scattering events contribute to more realistic rendering.

The material parameters Ψ of the BSDF model, to be optimized during inverse rendering, include absorption coefficients σ_a , longitudinal β_M and azimuthal β_N roughnesses, and interior i_{ior} and exterior e_{ior} indices of refraction [11].

Scalp-Aligned Albedo Parameterization. By assuming homogeneous material parameters Ψ (Eq. 1), a uniform hair albedo is enforced across all strand points, which cannot model the nuanced color variations commonly observed in real world, naturally or artificially, ranging from root-to-tip gradients within a single strand to subtle differences across wisps. To address this limitation, following previous work [37, 47, 60, 61], we propose a scalp-aligned parameterization that introduces sufficient degrees of freedom to capture spatial color variation, while maintaining natural regularization and preventing over-parameterization.

Specifically, recall that all input strands \mathcal{S} are rooted on the parametric head \mathcal{H} . We pre-define a scalp region \mathcal{H}_S on the head surface and unwrap it to a 2D parameterization, namely the *scalp space*. For each strand $s \in \mathcal{S}$, we project its root to the corresponding attachment point on \mathcal{H}_S and embed it into the scalp space at coordinates u_s via barycentric interpolation. We also convert each strand s to a B-spline curve s^* by uniformly sampling n_c control points. The hair is finally formulated as a regular 3D volume, conceptually equivalent to n_c layers of textures with resolutions of (w_u, h_u) . An arbitrary point p on strand s is mapped to this volume based on the strand’s root coordinate u_s and the point’s relative on-strand coordinate w_p^s . We utilize this scalp-aligned parameterization to encode absorption coefficients $\sigma_a \in \mathbb{R}^{n_c \times w_u \times h_u}$, which determine the albedo.

Compared to world-space voxelization, this parameterization ensures that points on the same strand share the same texture coordinate, and nearby texels map to points on adjacent strands. Decoupled from specific geometric styles, this inherent structure leads to more controllable degree of freedom and effective regularization across texture layers.

Rotation Correctives. In contrast to previous works [43] that benefit from clean images and relatively precise hair geometry obtained from synthetic renderings, our pipeline must contend with imperfections in hair reconstruction that introduces discrepancies between the reconstructed and ground-truth strand shapes. This poses challenges for inverse rendering, particularly problematic for specular highlights, which are highly sensitive to the accuracy of strand tangents. To mitigate sub-optimal roughness optimization due to misplaced highlights, we propose a lightweight corrective to the BSDF model that, rather than directly refining the curve geometry, optimizes a constrained rotation transformation for each strand segment to compensate for inaccuracies in the tangent directions.

Specifically, we formulate this rotation corrective R within the same scalp-aligned parameterization used for the absorption coefficients σ_a . This corrective encodes a rotation matrix R for each strand control point, represented using the Rodrigues formula. When evaluating a BSDF sample, its local shading frame $\{n, t, b\}$, defined by the normal n , tangent t , and bitangent b , is rotated by R to yield a corrected frame $\{Rn, Rt, Rb\}$ for shading.

3.3. Light-Aware Residual Model

While our BSDF model effectively reconstruct the large-scale appearance of hair and generalizes well to novel views and lighting conditions, relying solely on this analytical model leads to over-smoothed results due to a broad range of imperfect assumptions such as simulated lighting, missing diffraction and interference in the BSDF model, etc., which limits its capacity to capture subtle, yet visually critical, local details.

We introduce a residual term to model the remaining discrepancies between the physics-based rendering obtained the optimized BSDF and the real-world observations. To ensure desired consistency and generalization, this residual model is designed to be both light-aware and strand-aligned, to enable robust relighting capabilities and facilitate a compact, geometry-aware parameterization.

Specifically, we leverage 3D Gaussian Splatting (3DGS) [24] for scene representation. However, instead of optimizing free-form Gaussians as in the vanilla 3DGS formulation, we enforce identical anisotropic Gaussians for all kernels, and place them along the input strand geometry to form a chained sequence. During optimization, we fix the geometric parameters of all Gaussians (position, covariance, and opacity) and focus solely on optimizing their appearance parameters, as detailed below. This strategy is motivated by our primary focus on appearance modeling, recognizing that the initial Gaussian configuration already inherits good strand geometry from the input.

Dual-Level Spherical Harmonics. We introduce a dual-level spherical harmonics (SH) function as the appearance representation for the residual model, which is both view- and light-dependent, achieved by nesting two levels of SH. Given a dual-level SH, denoted as C_{vl} , the final Gaussian color c is calculated as a function of the camera view d_v and incident lighting d_l directions (in global coordinates):

$$\begin{aligned} c &= C_v(d_v), \\ C_v &= C_{vl}(d_l). \end{aligned} \quad (2)$$

With both SH levels using the same degree l , the parameter size of each dual-level SH is $C_{vl} \in \mathbb{R}^{(l+1)^2 \times (l+1)^2 \times 3}$ and $C_v \in \mathbb{R}^{(l+1)^2 \times 3}$. To maintain compact appearance parameters, we attach k dual-level SHs to each strand as anchor SHs. During rasterization, for each Gaussian primitive at location p belonging to strand s , its dual-level SH parameters are linearly interpolated from its nearest two anchor SHs based on its relative on-strand coordinate w_p^s .

3.4. Hybrid Inverse Rendering

Given the input hair geometry \mathcal{S} and multi-camera OLAT images \mathcal{I} , we optimize the hair appearance, including both the analytical BSDF model and the residual model, using our hybrid inverse rendering in two stages.

Stage 1: BSDF Optimization. The first stage optimizes the physical parameters Ψ of the BSDF model using differentiable path tracing. Due to the large number of scattering events, we use path replay backpropagation [46] to efficiently compute gradients of the used photometric loss:

$$\mathcal{L}_{rgb} = \sum_{I \in \mathcal{I}} \|I - \mathbf{R}(\Psi, \mathcal{S}, \mathcal{C}(I), \mathcal{L}(I))\|_1, \quad (3)$$

where \mathbf{R} represents the rendering function, $I \in \mathcal{I}$ denotes one ground truth view image, and $\mathcal{C}(I)$ and $\mathcal{L}(I)$ are the calibrated camera and lighting information associated with the image I , respectively.

With our albedo parameterization and rotation correctives, the complete set of parameters for the hair BSDF model is represented as $\Psi = \{\sigma, \mathbf{R}, \beta_N, \beta_M, e_{ior}, i_{ior}\}$. To ensure smooth variation along each strand, we add regularization terms for the spatially varying variables in the scalp-aligned volume:

$$\mathcal{L}_\sigma = \sum_{i=1}^{n_c-1} \|\sigma_i - \sigma_{i-1}\|_1, \quad (4)$$

$$\mathcal{L}_R = \sum_{i=1}^{n_c-1} \|\mathbf{R}_i - \mathbf{R}_{i-1}\|_1. \quad (5)$$

Finally, the complete loss function for the first inverse rendering stage is:

$$\mathcal{L}_{BSDF} = \mathcal{L}_{rgb} + \lambda_\sigma \mathcal{L}_\sigma + \lambda_R \mathcal{L}_R, \quad (6)$$

where $\lambda_\sigma = 0.5$ and $\lambda_R = 10$ control the strength of both regularization terms.

Stage 2: Residual Optimization. Once the BSDF optimization in Stage 1 is complete, we proceed to optimize the residual model. This involves fitting the remaining differences between the rendered images and the input views. Specifically, we optimize the dual-level spherical harmonics parameters of the residual appearance by minimizing the following residual loss:

$$\mathcal{L}_{res} = \sum_{I \in \mathcal{I}} \|I - (I_{render} + I_{res})\|_2^2, \quad (7)$$

where I_{render} represents the path-traced rendering using the optimized BSDF parameters from Stage 1, and I_{res} is the residual image rendered via 3D Gaussian rasterization. Note that in stage 2 we employ L_2 loss which we find leads to more robust performance.

Implementation Details. The inverse rendering pipeline in Stage 1 is built upon the differentiable renderer of Mitsuba 3 [19]. We approximate the lighting using measured IES profiles, with lighting positions initialized via OLAT calibration. We jointly optimize the BSDF parameters Ψ and the lighting parameters, including positions, orientations, and intensities since the lighting calibration might not be perfectly accurate. In each optimization iteration, we randomly select 8 images from all camera views and lighting conditions, and aggregate gradients across these images. For a given capture, we run 1000 iterations, which takes about 20 hours on a single NVIDIA A100 GPU (40 GiB).

In Stage 2, we downsample the input hair to approximately 52K strands, each represented by 100 segments. Cylindrical Gaussians are attached to these strand segments, with their positions and covariance matrices derived from the location and orientation of the the corresponding hair segments. Each strand is associated with 8 anchor SHs with SH degree $l = 3$, and per-segment SH parameters are obtained through interpolation between its two nearest anchor SHs. During optimization, we optimize the SH appearance parameters and keep all geometric properties fixed. We train the residual Gaussians for 160K iterations on 8 NVIDIA A100 GPUs (40 GiB), which takes about 12 hours.

4. Experiments

In this section, we first explain our experimental setup, including the dataset, baselines, and evaluation metrics. Then we present and discuss the experiment results, including comparisons and ablation study, in detail.

4.1. Setup

Dataset. We capture our test data from 10 subjects (§3.1). These subjects feature diverse hair characteristics, including

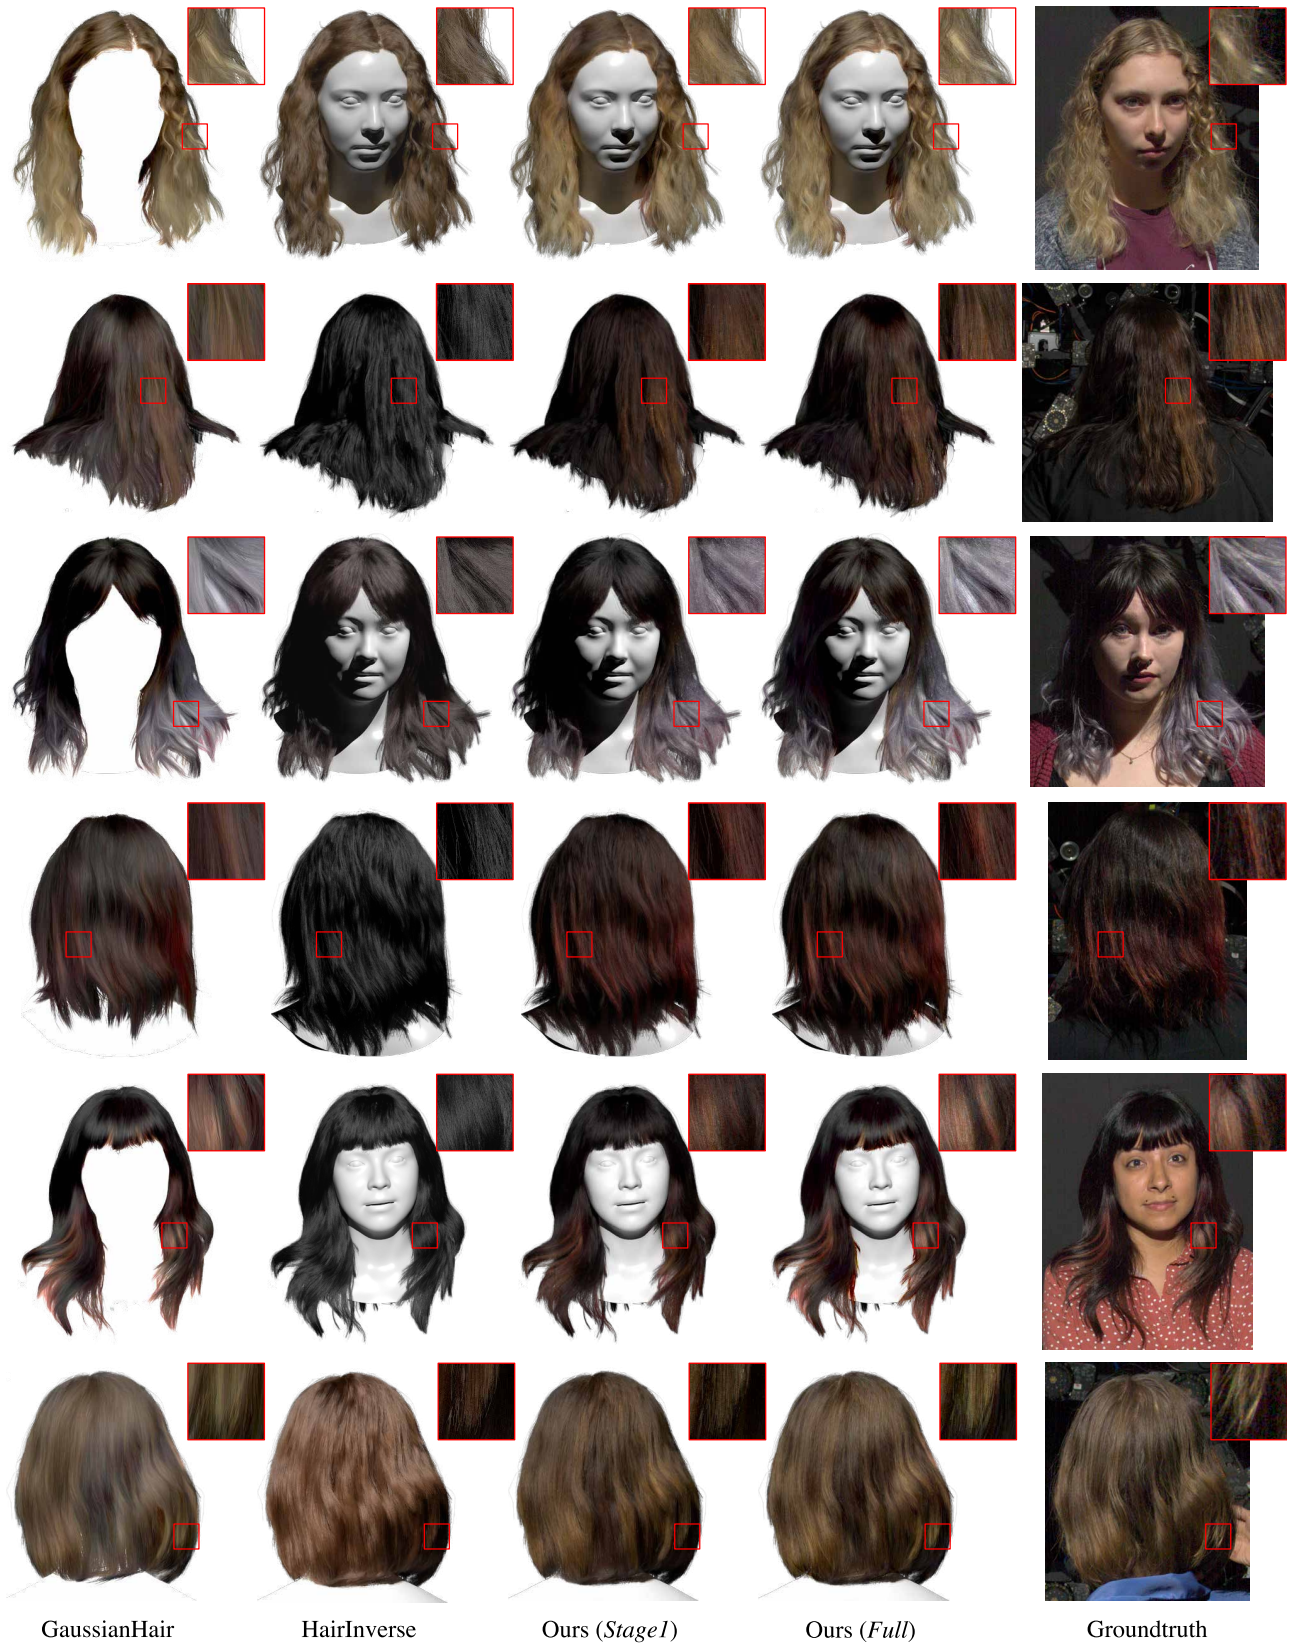


Figure 4. **Qualitative evaluation** of rendering results at novel testing views under novel lighting conditions of our method and baselines.



Figure 5. **Ablation study** of components in our method. Starting from Stage 1 only with original BSDF model (*Stage1 origBSDF*), we sequentially add spatially-varying albedo (*Stage1 NoCorr*), rotation correctives (*Stage1 Full*), and the residual model (*Full*).

Method	PSNR (\uparrow)	SSIM (\uparrow)	LPIPS (\downarrow)
HairInverse [43]	26.50	0.6767	0.3171
GaussianHair [27]	27.22	0.7340	0.2920
Ours (<i>Stage1 OrigBSDF</i>)	27.75	0.7350	0.2772
Ours (<i>Stage1 NoCorr</i>)	28.58	0.7410	0.2690
Ours (<i>Stage1 Full</i>)	28.67	0.7423	0.2653
Ours (<i>Full</i>)	29.71	0.7727	0.2461

Table 1. **Quantitative results on view reconstruction** of baseline comparisons and ablation study.

variations in length (short, medium, and long), style (curly, wavy, and straight), and color (black, blond, and gray). The hair geometry is reconstructed by *GroomCap* [61] from views under uniform lighting. For each subject, we use 160 lighting conditions and approximately 48 cameras views for training. We designate 2 novel views (one front and one back view) and 160 novel lightings for evaluation.

Baselines. To assess the quality of view synthesis and re-lighting, we compare *GroomLight* with two adapted state-of-the-art baselines quantitatively and qualitatively:

- *HairInverse* [43]: An inverse rendering method based on an analytical BSDF material model. The original method performs image-space optimization on selected pixels without involving differentiable renderers and presents results only on synthetic data. We re-implement this baseline, adapting to our algorithm, which iterative optimizes BSDF parameters based on photometric losses.
- *GaussianHair* [27]: A 3DGS-based [24] neural rendering method for hair. We re-implement this baseline by reusing our chained cylindrical Gaussians and applying our residual appearance model, i.e., dual-level SH, to the full image. This differs from their hair scattering model, which employs optimized SH base colors and manually tuned material parameters.



Figure 6. **Appearance editing.** Top row: decreasing roughness of hair BSDF to add more highlights. Bottom row: editing the albedo parameters to change the color for specific hair wisps.

Evaluation Metrics. We evaluate all methods using the following photometric error metrics: Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index Measure (SSIM), and Learned Perceptual Image Patch Similarity (LPIPS).

4.2. Comparisons

We compare our method with both *HairInverse* and *GaussianHair* baselines on testing view reconstruction under diverse lightings. As demonstrated in Tab. 1, our method outperforms both baselines by a large margin across all metrics. We also show qualitative comparisons in Fig. 4.

HairInverse [43], based on homogeneous hair BSDF material, estimates its parameters from the average color of the hair region. This model lacks the flexibility to represent the diverse textures found in real-world hair.

While 3DGS-based *GaussianHair* [27] achieves reasonable view reconstruction results, it often overfits to training



Figure 7. **Dynamic rendering.** We show rendering results of animated hair sequences using our optimized appearance model.

views and lightings, leading to poor generalization.

As shown in Fig. 4, *GroomLight* effectively reconstructs coarse hair appearances in stage 1 (*Ours Stage1*). By incorporating the residual model in Stage 2, we achieve photorealistic final renderings (*Ours Full*) with enhanced details.

4.3. Ablation Study

We conduct an ablation study to evaluate the contribution of each proposed component in our method, including the spatially varying albedo parameterization and rotation correctives in Stage 1, and the residual module of Stage 2.

Extended BSDF. We augment the base hair BSDF with a scalp-aligned albedo parameterization to enable spatially-varying hair color and introduce rotation correctives to compensate for inaccuracies in the reference geometry. We present the performance of different variants in Tab. 1: including *Stage1 OrigBSDF* (Stage 1 only using the original version of hair BSDF [11]), *Stage1 NoCorr* (adding the albedo parameterization), *Stage1 Full* (adding both the albedo parameterization and rotation correctives). The results demonstrate that each component contributes to improved reconstruction quality. As shown in Fig. 5, the original hair BSDF, assuming homogeneous materials, produces a uniform color across all strands. The spatially-varying albedo facilitates the reconstruction of more detailed hair texture, while the rotation correctives further enhance the accuracy of light reflectance, particularly for highlights.

Residual Model. We incorporate the residual model to address the inherent limitations of analytical BSDF models, which are based on approximations of physical behaviors. This residual model helps capture fine-grained texture details and highlights (*Full*), leading to more accurate results (Tab. 1) and enhanced photorealism as shown in Fig. 5.

4.4. Applications

The generalization capabilities of our hair appearance model naturally lends itself to various downstream applications that are previously challenging.

Relighting. Fig. 1 shows our relighting results under diverse environment lighting conditions. To provide the lighting direction input for our residual module, we approximate the light source in HDR environment maps as the position with maximum radiance. Our method achieves high-fidelity rendering quality across a wide range of lighting scenarios.

Appearance Editing. Our spatially-varying albedo parameterization enables convenient appearance editing (Fig. 6). By directly modifying the absorption coefficients for hair strands on the scalp space embedding, we can achieve various color effects. In addition to albedo editing, our representation supports rendering diverse highlight patterns by changing the roughness parameters.

Dynamic Rendering. Our method also generalizes to novel hair geometry. By animating the reconstructed hair using off-the-shelf simulation tools, we can render the sequence using the optimized appearance model, including BSDF and residual parameters, as shown in Fig. 7. Please refer to the supplementary video for the entire sequence.

5. Conclusion

We introduce *GroomLight*, a novel method for capturing and relighting photorealistic hair. Our approach combines the strengths of both physically based rendering and neural representation learning within a hybrid inverse rendering framework. We propose an extended analytical BSDF model that effectively captures the primary appearance of hair while remaining robust to variations in hair color and inaccuracies in the underlying geometry. To further enhance visual fidelity, we introduce a light-aware residual model based on dual-level spherical harmonics and a 3D Gaussian-based scene representation. *GroomLight* achieves state-of-the-art performance on inverse rendering for relightable human hair appearance modeling from real-world input data, which could benefit a wide spectrum of applications in virtual reality, digital fashion, movie industry, and more.

Limitations and Future Work. *GroomLight* relies on light stage capture setup to obtain lighting and camera calibrations, which limits its applicability to in-the-wild data where lighting and camera information are unknown. In addition, our pipeline is based on offline optimization and renders images at about 1 FPS during inference, which cannot support real-time applications. Furthermore, our method requires relatively accurate geometry and segmentation masks to prevent artifacts arising from the inclusion of unwanted textures from clothing, the face, or the background. We aim as future work to improve the efficiency of our pipeline, leverage large-scale datasets and neural networks to enable generalization of our method for in-the-wild image inputs.

Acknowledgement

We thank Tiancheng Sun, Stefanos Zafeiriou, and Erroll Wood for their helpful discussions; Jay Busch, Cynthia Herrera, Christoph Rhemann, Di Qiu, and Xu Chen for their assistance with data capture and processing; and all of our capture models. We are also grateful for the insightful feedback from the anonymous reviewers.

References

- [1] Carlos Aliaga, Carlos Castillo, Diego Gutierrez, Miguel A. Otaduy, Jorge Lopez-Moreno, and Adrián Jarabo. An appearance model for textile fibers. *Comput. Graph. Forum*, 36(4):35–45, 2017. [2](#)
- [2] Alexis Benamira and Sumanta N. Pattanaik. A combined scattering and diffraction model for elliptical hair rendering. *Comput. Graph. Forum*, 40(4):163–175, 2021. [2](#)
- [3] Sai Bi, Zexiang Xu, Pratul P. Srinivasan, Ben Mildenhall, Kalyan Sunkavalli, Milos Hasan, Yannick Hold-Geoffroy, David J. Kriegman, and Ravi Ramamoorthi. Neural reflectance fields for appearance acquisition. *CoRR*, abs/2008.03824, 2020. [2](#)
- [4] Zoubin Bi, Yixin Zeng, Chong Zeng, Fan Pei, Xiang Feng, Kun Zhou, and Hongzhi Wu. Gs^3 : Efficient relighting with triple gaussian splatting. In *SIGGRAPH Asia 2024*, pages 12:1–12:12, 2024. [2](#)
- [5] Menglei Chai, Lvdi Wang, Yanlin Weng, Yizhou Yu, Baining Guo, and Kun Zhou. Single-view hair modeling for portrait manipulation. *ACM Trans. Graph.*, 31(4):116:1–116:8, 2012. [2](#)
- [6] Menglei Chai, Lvdi Wang, Yanlin Weng, Xiaogang Jin, and Kun Zhou. Dynamic hair manipulation in images and videos. *ACM Trans. Graph.*, 32(4):75:1–75:8, 2013. [2](#)
- [7] Menglei Chai, Linjie Luo, Kalyan Sunkavalli, Nathan Carr, Sunil Hadap, and Kun Zhou. High-quality hair modeling from a single portrait photo. *ACM Trans. Graph.*, 34(6):204:1–204:10, 2015. [2](#)
- [8] Menglei Chai, Tianjia Shao, Hongzhi Wu, Yanlin Weng, and Kun Zhou. Autohair: fully automatic hair modeling from a single image. *ACM Trans. Graph.*, 35(4):116:1–116:12, 2016. [2](#)
- [9] Menglei Chai, Jian Ren, and Sergey Tulyakov. Neural hair rendering. In *ECCV 2020*, pages 371–388, 2020. [2](#)
- [10] Eric R. Chan, Connor Z. Lin, Matthew A. Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas J. Guibas, Jonathan Tremblay, Sameh Khamis, Tero Karras, and Gordon Wetzstein. Efficient geometry-aware 3d generative adversarial networks. In *CVPR 2022*, pages 16102–16112, 2022. [2](#)
- [11] Matt Jen-Yuan Chiang, Benedikt Bitterli, Chuck Tappan, and Brent Burley. A practical and controllable hair and fur model for production path tracing. *Comput. Graph. Forum*, 35(2):275–283, 2016. [2](#), [3](#), [4](#), [8](#)
- [12] Eugene d’Eon, Guillaume François, Martin Hill, Joe Letteri, and Jean-Marie Aubry. An energy-conserving hair reflectance model. *Comput. Graph. Forum*, 30(4):1181–1187, 2011.
- [13] Eugene d’Eon, Steve Marschner, and Johannes Hanika. A fiber scattering model with non-separable lobes. In *SIGGRAPH ’14 Talks*, page 46:1, 2014. [2](#)
- [14] Jian Gao, Chun Gu, Youtian Lin, Zhihao Li, Hao Zhu, Xun Cao, Li Zhang, and Yao Yao. Relightable 3d gaussians: Realistic point cloud relighting with BRDF decomposition and ray tracing. In *ECCV 2024*, pages 73–89, 2024. [2](#)
- [15] Philip-William Grassal, Malte Prinzler, Titus Leistner, Carsten Rother, Matthias Nießner, and Justus Thies. Neural head avatars from monocular RGB videos. In *CVPR 2022*, pages 18632–18643, 2022. [2](#)
- [16] Liwen Hu, Chongyang Ma, Linjie Luo, and Hao Li. Single-view hair modeling using a hairstyle database. *ACM Trans. Graph.*, 34(4):125:1–125:9, 2015. [2](#)
- [17] Weizhen Huang, Matthias B. Hullin, and Johannes Hanika. A microfacet-based hair scattering model. *Comput. Graph. Forum*, 41(4):79–91, 2022. [2](#)
- [18] Wenzel Jakob, Jonathan T. Moon, and Steve Marschner. Capturing hair assemblies fiber by fiber. *ACM Trans. Graph.*, 28(5):164, 2009. [2](#)
- [19] Wenzel Jakob, Sébastien Speierer, Nicolas Roussel, Merlin Nimier-David, Delio Vicini, Tizian Zeltner, Baptiste Nicolet, Miguel Crespo, Vincent Leroy, and Ziyi Zhang. Mitsuba 3 renderer, 2022. <https://mitsuba-renderer.org>. [5](#)
- [20] Yingwenqi Jiang, Jiadong Tu, Yuan Liu, Xifeng Gao, Xiaoxiao Long, Wenping Wang, and Yuexin Ma. Gaussianshader: 3d gaussian splatting with shading functions for reflective surfaces. In *CVPR 2024*, pages 5322–5332, 2024. [2](#)
- [21] Youngjoo Jo and Jongyoul Park. SC-FEGAN: face editing generative adversarial network with user’s sketch and color. In *ICCV 2019*, pages 1745–1753, 2019. [2](#)
- [22] James T. Kajiya. The rendering equation. In *SIGGRAPH 1986*, pages 143–150, 1986. [2](#)
- [23] James T. Kajiya and Timothy L. Kay. Rendering fur with three dimensional textures. In *SIGGRAPH 1989*, pages 271–280, 1989. [2](#)
- [24] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139:1–139:14, 2023. [2](#), [4](#), [7](#)
- [25] Pramook Khungurn and Steve Marschner. Azimuthal scattering from elliptical hair fibers. *ACM Trans. Graph.*, 36(2):13:1–13:23, 2017. [2](#)
- [26] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloé Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross B. Girshick. Segment anything. In *ICCV 2023*, pages 3992–4003, 2023. [3](#)
- [27] Haimin Luo, Min Ouyang, Zijun Zhao, Suyi Jiang, Longwen Zhang, Qixuan Zhang, Wei Yang, Lan Xu, and Jingyi Yu. Gaussianhair: Hair modeling and rendering with light-aware gaussians. *CoRR*, abs/2402.10483, 2024. [2](#), [3](#), [7](#)
- [28] Linjie Luo, Hao Li, and Szymon Rusinkiewicz. Structure-aware hair capture. *ACM Trans. Graph.*, 32(4):76:1–76:12, 2013. [2](#)
- [29] Shugao Ma, Tomas Simon, Jason M. Saragih, Dawei Wang, Yuecheng Li, Fernando De la Torre, and Yaser Sheikh. Pixel codec avatars. In *CVPR 2021*, pages 64–73, 2021. [2](#)

- [30] Stephen R. Marschner, Henrik Wann Jensen, Mike Cammarano, Steve Worley, and Pat Hanrahan. Light scattering from human hair fibers. *ACM Trans. Graph.*, 22(3):780–791, 2003. 2
- [31] 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. In *ECCV 2020*, pages 405–421, 2020. 2
- [32] Jacob Munkberg, Wenzheng Chen, Jon Hasselgren, Alex Evans, Tianchang Shen, Thomas Müller, Jun Gao, and Sanja Fidler. Extracting triangular 3d models, materials, and lighting from images. In *CVPR 2022*, pages 8270–8280, 2022. 2
- [33] Giljoo Nam, Chenglei Wu, Min H. Kim, and Yaser Sheikh. Strand-accurate multi-view hair capture. In *CVPR 2019*, pages 155–164, 2019. 2
- [34] Sylvain Paris, Will Chang, Oleg I. Kozhushnyan, Wojciech Jarosz, Wojciech Matusik, Matthias Zwicker, and Frédo Durand. Hair photobooth: geometric and photometric acquisition of real hairstyles. *ACM Trans. Graph.*, 27(3):30, 2008. 2
- [35] Shenhan Qian, Tobias Kirschstein, Liam Schoneveld, Davide Davoli, Simon Giebenhain, and Matthias Nießner. Gaussianavatars: Photorealistic head avatars with rigged 3d gaussians. In *CVPR 2024*, pages 20299–20309, 2024. 2
- [36] Haonan Qiu, Chuan Wang, Hang Zhu, Xiangyu Zhu, Jinjin Gu, and Xiaoguang Han. Two-phase hair image synthesis by self-enhancing generative model. *Comput. Graph. Forum*, 38(7):403–412, 2019. 2
- [37] Radu Alexandru Rosu, Shunsuke Saito, Ziyang Wang, Chenglei Wu, Sven Behnke, and Giljoo Nam. Neural strands: Learning hair geometry and appearance from multi-view images. In *ECCV 2022*, pages 73–89, 2022. 2, 4
- [38] Iman Sadeghi, Heather Pritchett, Henrik Wann Jensen, and Rasmus Tamstorf. An artist friendly hair shading system. *ACM Trans. Graph.*, 29(4):56:1–56:10, 2010. 2
- [39] Shunsuke Saito, Gabriel Schwartz, Tomas Simon, Junxuan Li, and Giljoo Nam. Relightable gaussian codec avatars. In *CVPR 2024*, pages 130–141, 2024. 2
- [40] Yahao Shi, Yanmin Wu, Chenming Wu, Xing Liu, Chen Zhao, Haocheng Feng, Jingtuo Liu, Liangjun Zhang, Jian Zhang, Bin Zhou, Errui Ding, and Jingdong Wang. GIR: 3d gaussian inverse rendering for relightable scene factorization. *CoRR*, abs/2312.05133, 2023. 2
- [41] Jingxiang Sun, Xuan Wang, Lizhen Wang, Xiaoyu Li, Yong Zhang, Hongwen Zhang, and Yebin Liu. Next3d: Generative neural texture rasterization for 3d-aware head avatars. In *CVPR 2023*, pages 20991–21002, 2023. 2
- [42] Tiancheng Sun, Kai-En Lin, Sai Bi, Zexiang Xu, and Ravi Ramamoorthi. Nelf: Neural light-transport field for portrait view synthesis and relighting. In *EGSR 2021*, pages 155–166, 2021. 2
- [43] Tiancheng Sun, Giljoo Nam, Carlos Aliaga, Christophe Hery, and Ravi Ramamoorthi. Human hair inverse rendering using multi-view photometric data. In *EGSR 2021*, pages 179–190, 2021. 1, 2, 4, 7
- [44] Zhentao Tan, Menglei Chai, Dongdong Chen, Jing Liao, Qi Chu, Lu Yuan, Sergey Tulyakov, and Nenghai Yu. Michigan: multi-input-conditioned hair image generation for portrait editing. *ACM Trans. Graph.*, 39(4):95, 2020. 2
- [45] Ayush Tewari, Justus Thies, Ben Mildenhall, Pratul P. Srinivasan, Edgar Tretschk, Yifan Wang, Christoph Lassner, Vincent Sitzmann, Ricardo Martin-Brualla, Stephen Lombardi, Tomas Simon, Christian Theobalt, Matthias Nießner, Jonathan T. Barron, Gordon Wetzstein, Michael Zollhöfer, and Vladislav Golyanik. Advances in neural rendering. *Comput. Graph. Forum*, 41(2):703–735, 2022. 2
- [46] Delio Vicini, Sébastien Speierer, and Wenzel Jakob. Path replay backpropagation: differentiating light paths using constant memory and linear time. *ACM Trans. Graph.*, 40(4):108:1–108:14, 2021. 5
- [47] Lvdi Wang, Yizhou Yu, Kun Zhou, and Baining Guo. Example-based hair geometry synthesis. *ACM Trans. Graph.*, 28(3):56, 2009. 4
- [48] Ziyang Wang, Giljoo Nam, Tuur Stuyck, Stephen Lombardi, Michael Zollhöfer, Jessica K. Hodgins, and Christoph Lassner. HVH: learning a hybrid neural volumetric representation for dynamic hair performance capture. In *CVPR 2022*, pages 6133–6144, 2022. 2
- [49] Lingyu Wei, Liwen Hu, Vladimir G. Kim, Ersin Yumer, and Hao Li. Real-time hair rendering using sequential adversarial networks. In *ECCV 2018*, pages 105–122, 2018. 2
- [50] Mengqi (Mandy) Xia, Bruce Walter, Eric Michielssen, David Bindel, and Steve Marschner. A wave optics based fiber scattering model. *ACM Trans. Graph.*, 39(6):252:1–252:16, 2020. 2
- [51] Mengqi (Mandy) Xia, Bruce Walter, Christophe Hery, Olivier Maury, Eric Michielssen, and Steve Marschner. A practical wave optics reflection model for hair and fur. *ACM Trans. Graph.*, 42(4):39:1–39:15, 2023. 2
- [52] Yuelang Xu, Bengwang Chen, Zhe Li, Hongwen Zhang, Lizhen Wang, Zerong Zheng, and Yebin Liu. Gaussian head avatar: Ultra high-fidelity head avatar via dynamic gaussians. In *CVPR 2024*, pages 1931–1941, 2024. 2
- [53] Ling-Qi Yan, Chi-Wei Tseng, Henrik Wann Jensen, and Ravi Ramamoorthi. Physically-accurate fur reflectance: modeling, measurement and rendering. *ACM Trans. Graph.*, 34(6):185:1–185:13, 2015. 2
- [54] Ling-Qi Yan, Henrik Wann Jensen, and Ravi Ramamoorthi. An efficient and practical near and far field fur reflectance model. *ACM Trans. Graph.*, 36(4):67:1–67:13, 2017. 2
- [55] 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*, pages 5453–5462, 2021. 2
- [56] Meng Zhang, Menglei Chai, Hongzhi Wu, Hao Yang, and Kun Zhou. A data-driven approach to four-view image-based hair modeling. *ACM Trans. Graph.*, 36(4):156:1–156:11, 2017. 2
- [57] Xiuming Zhang, Pratul P. Srinivasan, Boyang Deng, Paul E. Debevec, William T. Freeman, and Jonathan T. Barron. Nerfactor: neural factorization of shape and reflectance under an unknown illumination. *ACM Trans. Graph.*, 40(6):237:1–237:18, 2021. 2

- [58] Yufeng Zheng, Victoria Fernández Abrevaya, Marcel C. Bühler, Xu Chen, Michael J. Black, and Otmar Hilliges. I M avatar: Implicit morphable head avatars from videos. In *CVPR 2022*, pages 13535–13545, 2022. [2](#)
- [59] Yang Zheng, Qingqing Zhao, Guandao Yang, Wang Yifan, Donglai Xiang, Florian Dubost, Dmitry Lagun, Thabo Beeler, Federico Tombari, Leonidas J. Guibas, and Gordon Wetzstein. Physavatar: Learning the physics of dressed 3d avatars from visual observations. In *ECCV 2024*, pages 262–284, 2024. [2](#)
- [60] Yi Zhou, Liwen Hu, Jun Xing, Weikai Chen, Han-Wei Kung, Xin Tong, and Hao Li. Hairnet: Single-view hair reconstruction using convolutional neural networks. In *ECCV 2018*, pages 249–265, 2018. [4](#)
- [61] Yuxiao Zhou, Menglei Chai, Daoye Wang, Sebastian Winberg, Erroll Wood, Kripasindhu Sarkar, Markus Gross, and Thabo Beeler. Groomcap: High-fidelity prior-free hair capture. *ACM Trans. Graph.*, 43(6):254:1–254:15, 2024. [3](#), [4](#), [7](#)
- [62] Junqiu Zhu, Sizhe Zhao, Lu Wang, Yanning Xu, and Ling-Qi Yan. Practical level-of-detail aggregation of fur appearance. *ACM Trans. Graph.*, 41(4):47:1–47:17, 2022. [2](#)
- [63] Arno Zinke and Andreas Weber. Light scattering from filaments. *IEEE Trans. Vis. Comput. Graph.*, 13(2):342–356, 2007.
- [64] Arno Zinke, Martin Rump, Tomás Lay, Andreas Weber, Anton Andriyenko, and Reinhard Klein. A practical approach for photometric acquisition of hair color. *ACM Trans. Graph.*, 28(5):165, 2009. [2](#)