

# GroomLight: Hybrid Inverse Rendering for Relightable Human Hair Appearance Modeling

## Supplementary Material

### 1. Implementation Details

In this section, we provide additional implementation details about our method and the baselines.

#### 1.1. Our Method

Our inverse rendering pipeline consists of a two-stage optimization scheme. In *Stage 1*, we estimate the physical parameters of the extended hair BSDF model (Sec 3.2). We build upon the original hair BSDF [1], enabling albedo parameterization to support diverse hair colors and introducing rotation correctives to compensate for imperfect geometry input. Specifically, the albedo parameterization is implemented as a texture map, and rotation correctives are represented by a 2D rotation map similar to a normal map in common BRDF models. Both maps are defined in  $4096 \times 4096$  resolution. When evaluating a shading frame for the 3D point on the hair B-spline curves that intersects with the ray, we first find its rotation corrective using the rotation map and the point’s UV coordinates. We then apply the rotation to the shading frame to change the incident and exitant lighting directions, which are used for evaluating the hair BSDF. The entire process is differentiable, as implemented in Mitsuba 3 [2] based on path replay backpropagation [6]. Given multi-view OLAT images, we optimize the hair BSDF parameters using the L1 photometric loss and regularization losses (Sec 3.4). Note that we jointly optimize the lighting parameters, which we find leads to better performance. The optimization is performed on a single NVIDIA A100 GPU (40 GiB).

In *Stage 2*, we introduce a residual model based on the 3D Gaussian representation [3] with dual-level spherical harmonics to enhance rendering photorealism. We utilize eight NVIDIA A100 GPUs (40 GiB) for training and evaluation. During training, we randomly split the hair strands and equally distribute them across six A100 GPUs. We use the remaining two GPUs for splatting and computing gradients, respectively. In each training iteration, we randomly choose one split and render the residual map to compute the L2 photometric loss. Note that we employ the L2 loss instead of L1, as we find that the L1 loss leads to overly noisy results. During evaluation, we combine all the splits for rendering.

#### 1.2. Baselines

We compare our method to current state-of-the-art methods, including HairInverse [5] and GaussianHair [4]. Since no public code is available, we re-implement the methods

based on the algorithms described in the papers.

**HairInverse** [5] is the inverse rendering pipeline introduced in HairInverse consists of geometry reconstruction based on a synthetic setup and estimation of hair BSDF [1] parameters. However, the method only estimates the average color of the hair and struggles with real-world data due to failures in geometry reconstruction. To ensure fair comparisons, we use our reconstructed hair geometry as input and employ our optimized lighting and hair roughness parameters to render results for evaluation. The optimization process is performed using views under uniform lighting, following Algorithm 1 from the paper [5]. We build the training and testing pipelines in Mitsuba 3 [2].

**GaussianHair** [4] includes geometry and appearance reconstruction based on connected 3D Gaussians [3], similar to our chained cylinder Gaussian representation. To enable relighting, the method introduces a hair scattering function that employs optimized SH base colors and manually tuned material parameters. For simplicity, we leverage our Gaussian representation with dual-level spherical harmonics to evaluate the performance of relightable 3D Gaussians in modeling human hair appearance. We use the same inputs as our method (hair geometry and multi-view OLAT images) to train the model. Similarly, we solely optimize the spherical harmonics with other parameters (i.e., 3D Gaussian positions, covariance matrices and opacities) fixed.

### 2. Additional Results

We highly recommend watching the supplementary video for more results on relighting, appearance editing, dynamic rendering, and evaluations.

We present more qualitative comparison results with baselines in Fig. 4. In addition, we visualize the error maps of view synthesis evaluation, for both baseline comparisons (Fig. 3) and ablation study (Fig. 2). Consistently, our method achieves the best results both visually and quantitatively, against all baselines we have evaluated.

We also include additional visual results. In Fig. 4, we show relighting results under diverse environmental lighting and rotated hair geometry. Our results demonstrate high-fidelity rendering under novel lighting conditions and novel view configurations. In Fig. 1, we present more qualitative comparisons between our method and the baselines. GaussianHair [4] is overfitted to training lighting conditions, where the model learns baked-in highlights and thus generalizes poorly to novel test lighting. HairInverse [5] regresses

to an average color of the hair region, lacking the capacity to reconstruct diverse hair colors and highlight patterns. Our method estimates the base appearance in the first stage (*Ours Stage1*) and captures more details and highlights by leveraging the residual module (*Ours Full*).

## References

- [1] 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. [1](#)
- [2] 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>. [1](#)
- [3] 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. [1](#)
- [4] 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. [1](#)
- [5] 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](#)
- [6] 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. [1](#)



Figure 1. **More qualitative baseline comparisons** of novel testing views rendering under novel lighting conditions.

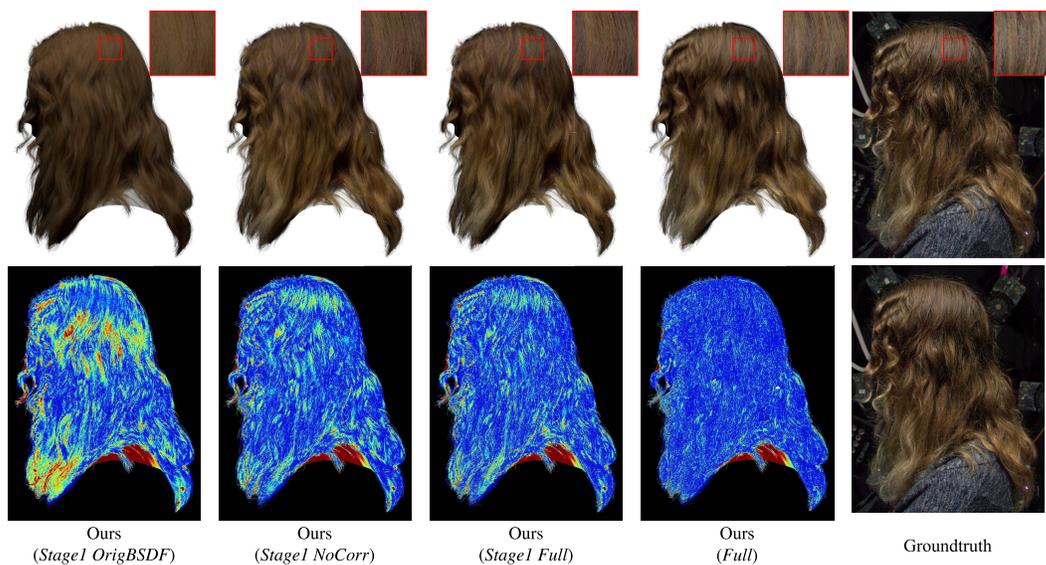


Figure 2. **Error maps of ablation study** on view synthesis under novel viewpoint and lighting condition.

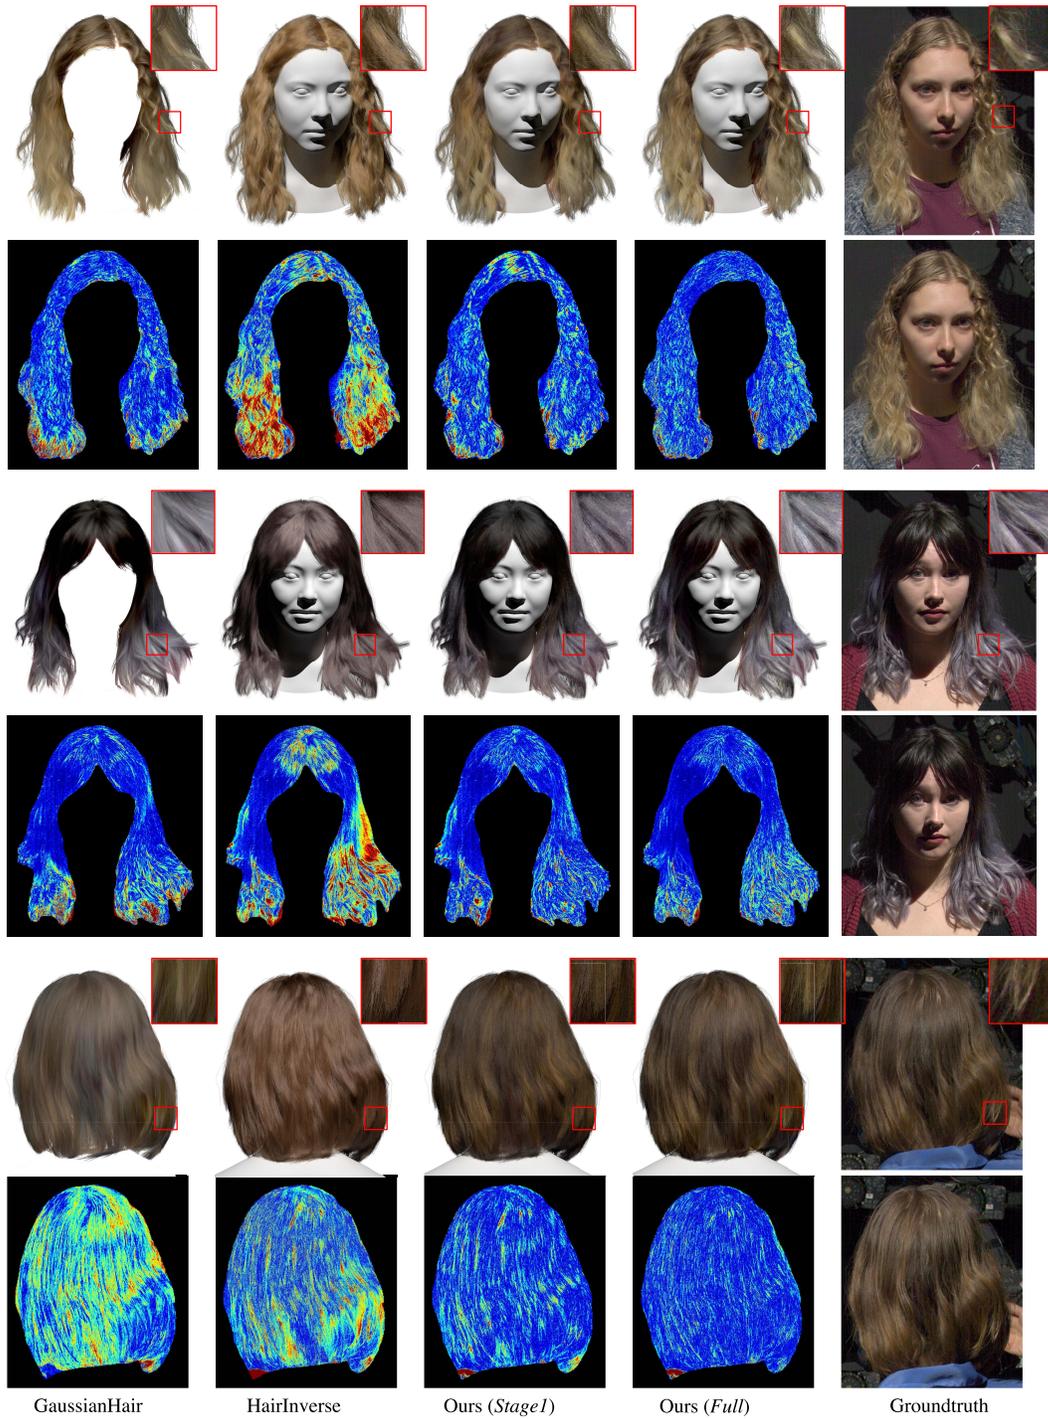


Figure 3. **Error maps of baseline comparisons** on view synthesis under novel viewpoint and lighting condition.



Figure 4. **More relighting results.** Each row shows relighting results with rotated environmental light (left) and rotated geometry (right).