

Dr.Hair: Reconstructing Scalp-Connected Hair Strands without Pre-Training via Differentiable Rendering of Line Segments

Yusuke Takimoto^{1*} Hikari Takehara^{1*} Hiroyuki Sato^{1*} Zihao Zhu^{1,2†} Bo Zheng¹

¹Huawei Technologies Japan K.K.

²Keio University

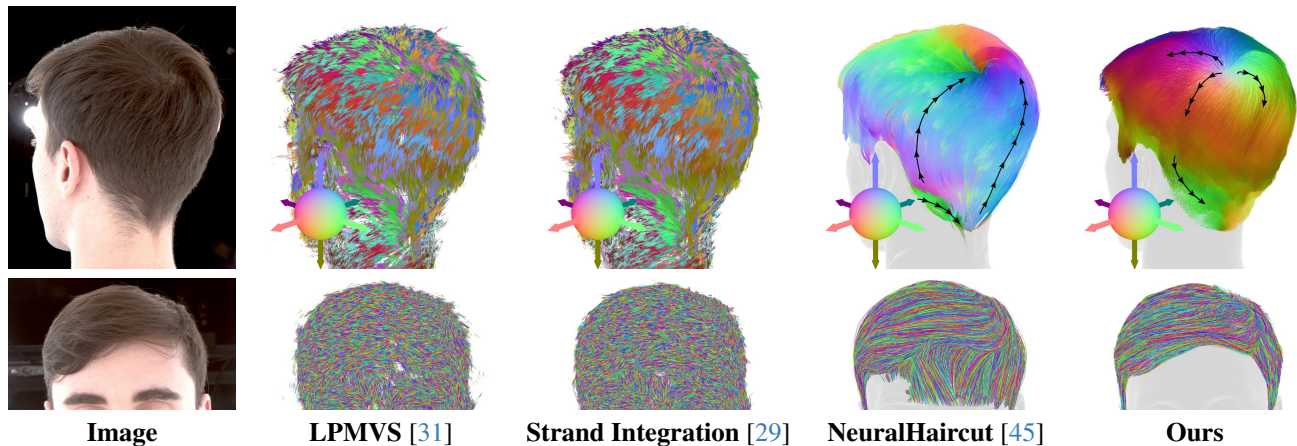


Figure 1. Results of existing strand-based 3D reconstruction methods and our method tested with the data captured by a multi-camera system. In the upper row, color and colored arrows represent 3D orientation of hair strands. The overlaid black arrows were drawn manually to visualize rough orientations. The lower row shows individual strands with random color. LPMVS and Strand Integration failed to estimate consistent direction, and their strands are too short not to connect to the scalp. The absolute orientation of strands estimated by NeuralHaircut is mostly the opposite of the actual hair orientation. Our method demonstrates better precision in reconstructing the directional flow of scalp-connected hair.

Abstract

In the film and gaming industries, achieving a realistic hair appearance typically involves the use of strands originating from the scalp. However, reconstructing these strands from observed surface images of hair presents significant challenges. The difficulty in acquiring Ground Truth (GT) data has led state-of-the-art learning-based methods to rely on pre-training with manually prepared synthetic CG data. This process is not only labor-intensive and costly but also introduces complications due to the domain gap when compared to real-world data. In this study, we propose an optimization-based approach that eliminates the need for pre-training. Our method represents hair strands as line segments growing from the scalp and optimizes them using a novel differentiable rendering algo-

ri thm. To robustly optimize a substantial number of slender explicit geometries, we introduce 3D orientation estimation utilizing global optimization, strand initialization based on Laplace’s equation, and reparameterization that leverages geometric connectivity and spatial proximity. Unlike existing optimization-based methods, our method is capable of reconstructing internal hair flow in an absolute direction. Our method exhibits robust and accurate inverse rendering, surpassing the quality of existing methods and significantly improving processing speed.

1. Introduction

High-quality 3D hair data is essential for depicting realistic human figures in movies, games, and metaverse. However, capturing real hair is notoriously difficult due to its intricate properties, including its elongated shape, overlapping strands, transparency, reflectivity, and uniformity, which

*Core authors

†Work done during an internship at Huawei Technologies Japan K.K.

challenge even the most advanced computer vision stereo techniques at a sub-pixel resolution. As a result, hair processing is one of the most formidable tasks in image-based 3D human modeling.

Hair-specific reconstructions methods have been studied for many years [15, 29, 31, 34, 35, 48, 54]. These approaches are characterized by applying a Gabor filter to the hair image to calculate the 2D orientation, which is then combined with 3D measurements for optimization. However, these methods require accurate calibration of the illumination and cameras, making it difficult to scale them to casual shooting environments. In addition, since only the hair’s surface can be measured, it is not easy to estimate the form of the hair connected to the scalp, which is commonly used in industry.

Data-driven reconstructions, particularly those using volumetric representations via neural networks, have recently gained traction in 3D hair modeling, and many studies on humans with hair [3, 24] are conducted. Among them, hair-specific methods using one or a few views [7, 14, 20, 42, 55, 60, 61, 64, 65] have been actively studied. The recently proposed methods [40, 45] use pre-trained priors and perform strand fitting through differentiable rendering [22, 39, 41, 57] against the multi-view images at runtime. However, creating the CG data used for pre-training is not only costly and requires manual work by artists but also has the problem of domain gaps.

In response, we propose Dr.Hair, an optimization-based pipeline that recovers individual strands connected to the scalp from multi-view images to address the above problems. We start with conventional hair representation used in real-time rendering and standard CG tools. After fitting a scalp to the raw hair mesh, we compute consistent 3D orientations from 2D orientation images. From the results, the guide strands are initialized based on a differential equation. Finally, a hierarchical relationship called *guide-child* is utilized for optimization based on differentiable rendering. To summarize, our contributions are:

- 3D orientation estimation using global optimization, estimating consistent surface orientation;
- Laplace’s equation-based strand initialization, filling interior hair flow smoothly from surface observations;
- Rasterization-based differentiable rendering algorithm for line segments, generating smooth gradient in image space while maintaining high-frequency detail;
- Reparameterization of strand shapes, propagating dense gradients throughout the geometry;
- An optimization framework using hierarchical relations of guide and child hair.

Finally, we validate the effectiveness of our method on synthetic and real data.

2. Related Work

2.1. Optimization-based methods

Since the early 2000s, to measure the strands on hair surface, optimization methods combining 3D geometry with 2D orientation extracted from multiple images have been explored [15, 34, 35, 54]. These methods utilized Gabor filters to address the high specularities arising from the cylindrical shape and semi-transparent material and have been extended to dynamic scenarios [26, 27]. Subsequently, Line-based PatchMatch MVS (LPMVS) [31] have enabled more accurate strand geometry acquisition. Strand Integration [29] further refines the strands of LPMVS. Joint measurement of material properties has been achieved [48]. However, these methods require controlled lighting conditions and struggle with the 180° ambiguity of 2D/3D orientation.

Some methods have addressed this ambiguity issue using user stroke input [6], Markov Random Field optimization [28], and pre-trained models [62] to estimate long strands connected to the scalp. Specialized techniques for braided hair [11] and for estimating simulation parameters [13] also exist. Although handling various hairstyles is challenging for these approaches, our global optimization robustly estimates consistent 3D orientations. To extrapolate internal hair flow from surface observations, second order differential equations [15, 35] and 3D PatchMatch [62] have been used. These interpolation techniques are valuable for seamlessly integrating hair strands onto the scalp.

Most existing methods stray from the conventional hair modeling process conducted by human artists employing tools such as XGen [1], Ornatrix [9], and Blender Hair Curves [2]. In these tools, hair manipulation commonly involves a hierarchical arrangement of guide and child strands. Artists primarily manipulate guides, while children are generated through interpolation. EnergyHair [63] facilitates image-based interactive hair modeling utilizing this hierarchical structure. Our pipeline brings this guide-child hierarchy into a fully automatic hair strand reconstruction.

Differentiable rendering (DR) has been attracting attention as a method for reconstructing 3D scenes. Not only implicit functions [23, 30, 51], but also methods for explicit geometric primitives such as meshes [8, 16, 21, 22, 32, 33, 39, 49] and point clouds [41, 57] have been studied extensively. However, DR of line primitives used for hair has not been well studied. We propose a DR framework for line segments to robustly optimize hair strands.

2.2. Learning-based methods

Data-driven hair strand reconstruction has been widely studied. An early work used simulated examples [10], and the field has gained popularity after the release of a synthetic dataset created by hand at significant cost, USC-HairSalon [12].

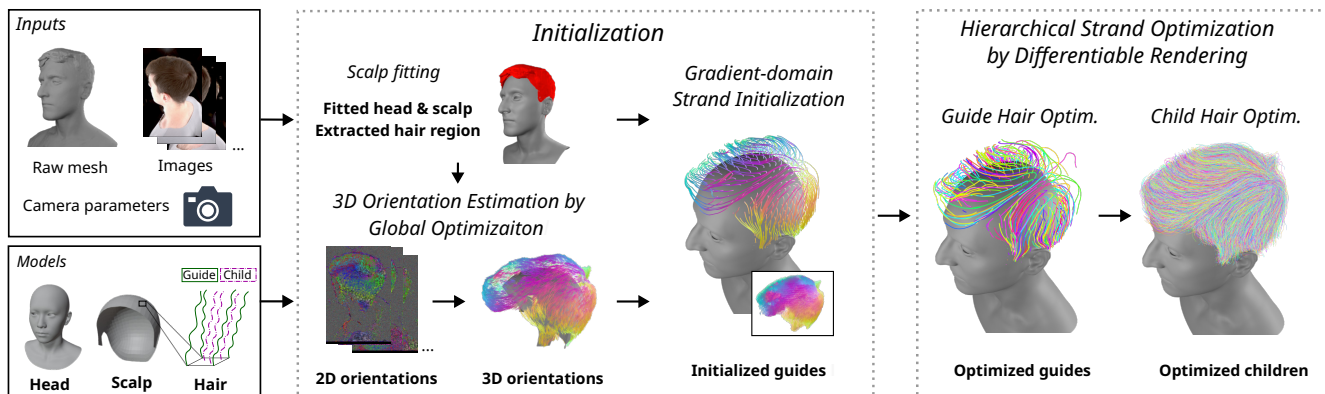


Figure 2. The overview of our pipeline. Our approach combines traditional real-time rendering techniques with recent advances in differentiable rendering. First, we fit a template to a raw mesh. Next, we compute consistent 3D orientations from 2D orientation images and initialize guide strands based on a differential equation. Finally, optimization based on differentiable rendering is applied by leveraging the hierarchical relationship between guides and children.




Figure 3. Surface 3D orientation estimation by global optimization: Hair regions are visualized on a gray scalp. (a) A raw mesh. (b) Surface 3D orientations before our optimization. The color stands for 3D orientation. Because of 180° ambiguity, orientations flip intermittently. (c) Minimum Spanning Tree (MST). The nodes are downsampled points, which are visualized as colored lines after propagation. Black lines represent edges. (d) Surface 3D orientations after our optimization. Consistent hair flow from the whorl to the ends is established.

Many methods attempt to overcome the ambiguity of thin hair strands through volumetric representation, particularly in simple setups like single-view [7, 14, 42, 55, 60, 64, 65], sparse views [20, 61], and sketches [43]. However, due to the small number of views, their 3D consistency is limited. Fine-tuning to deform strands against views [12, 20] is utilized to enhance high-frequency details. Neural-based volumetric representations [30] have been applied to head and hair reconstruction [3, 23, 24]. CT2Hair [44] reconstructs high-quality strands from wig CT data. Dynamic scenes are handled as well [52, 53, 56]. Generative models for hair strands have been recently proposed [46, 66].

In a multi-view setup, some methods train priors such as strand generators and perform the geometry texture [50] optimization against input images. NeuralStrands [40] uses surface 3D orientation [31] as a constraint and rasterizes strands generated from neural shape texture using point-based DR [41, 57]. It achieves photorealistic hair appear-

ance using neural appearance texture and a UNet-based neural renderer. However, manual annotation is needed to resolve the 180° orientation ambiguity. NeuralHaircut [45] uses volumetric reconstruction [51] as the first stage. Then, regularizing the geometry texture with a pre-trained diffusion model, broad gradient propagation by mesh-based soft rasterization [22, 39] is performed for sparse strands. It finally generates realistic images using UNet. These methods offer excellent rendering quality, but strand geometry accuracy is limited by blurred images of soft rasterization and domain gaps. Moreover, both pre-training and optimization are time-consuming. Recently proposed GaussianHair [25] utilizes gaussian splatting [18] along with a pre-trained strand decoder.

3. Method

Figure 2 illustrates the overview of the proposed method.

3.1. Initialization

In this step, the scalp and the hair strands connected to it are initialized. We use a head mesh and a separate scalp mesh with a different topology as templates as shown in Figure 2. A separate scalp with uniform vertex distribution and clear sideburn shapes is convenient for hair growth. The scalp mesh has a correspondence with the scalp region of the head mesh, and their vertex positions can be mutually transformed by linear interpolation.

The input comprises multi-view images, camera parameters, and a raw mesh. As a preprocessing step, the head is automatically fitted to the face region of the raw mesh by non-rigid ICP utilizing landmarks and segmentation. The scalp is then optimized to lay inside the hair region of the raw mesh, and the hair region is extracted. Details of this scalp fitting are provided in the supplementary material.

3D Orientation Estimation by Global Optimization

First, using Gabor filters [34], the 2D orientation and confidence are calculated from the color images of each view. Next, 3D-oriented points are reconstructed by using a modified version of LPMVS [31]. For fast computation, depth values were fixed by rendering the raw mesh from each viewpoint, and only the 3D orientation was optimized. Then, to reduce noise, mean-shift is applied to the points by following [31]. The resulting orientation contains a 180° ambiguity in the Euclidean space, so global consistency from the roots to the tips should be sought.

We make a graph structure that connects neighboring points using edge weights determined by the absolute value of the inner product of orientations. On this graph, an MST [37] is constructed. The orientation is then sequentially propagated from the initial point until all points are reached on the MST; if the inner product of neighboring orientations is negative, the destination is rotated by 180° . Since a single run may potentially lead to local optima, adding random perturbations to the edge weights of the graph, an MST creation and propagation are performed 100 times. We define the score of the graph as the sum of inner products of adjacent points' orientations and choose the best in the trials as the global optimal solution. At this stage, the orientation is globally consistent but uncertain in absolute terms. In other words, opposite directions, from the tips to the roots, may be estimated. Therefore, the heuristic that most hair strands should face the direction of gravity is used. The above process is applied to downsampled points, and the orientation is reflected back to the original resolution. Original points that differ much in orientation from the optimized points are removed as noise. Our global optimization process is shown in Figure 3.


Gradient-Domain Strand Initialization

We leverage gradient domain processing [17, 36] to estimate spatially smooth internal hair flow through volumetric representation [20, 60]. A visualization of this process is given in Figure 4.

Boolean operators extract the space Ω filled with hair, enclosed by the boundary of hair regions in the raw mesh and the scalp mesh. We consider filling Ω with a smooth hair flow field. $f_o(p) = (n_x, n_y, n_z)$ denotes the orientation at a position $p = (x, y, z) \in \Omega$. It should satisfy the following properties:

$$\begin{aligned} \nabla^2 f_o(p) &= 0 \quad \text{subject to} \quad \|f_o(p)\|_2 = 1, \\ f_o(p_h) &= H(p_h), \quad p_h \in \mathbb{H}, \quad f_o(p_s) = S(p_s), \quad p_s \in \mathbb{S} \end{aligned} \quad (1)$$

$f_o(p)$ follows a type of Laplace's equation, a special case of Poisson's equation, whose solution can be determined by boundary conditions. There are Dirichlet boundary conditions with multiple types: $\Omega = \{\mathbb{H}, \mathbb{S}, \mathbb{U}\}$. \mathbb{H} represents the



(a) Boundary cond. (b) Optimized volume (c) Extracted strands

Figure 4. Gradient-domain strand initialization: (a) Boundary condition. Colored region represents \mathbb{H} and \mathbb{S} , and black region represents \mathbb{U} . (b) Optimized volume sliced by certain planes. The interior is smoothly filled. (c) Extracted strands from the optimized volume.

hair surface boundary, and $H(p_h)$ is the estimated 3D orientation of the hair surface obtained in the previous step. \mathbb{S} and $S(p_s)$ are the scalp boundary and the orientation of the scalp surface, respectively. Based on the observation that hair on the top of the head grows upwards, but the hair on the back and sides tends to point downwards due to biological characteristics of scalp pores and gravity, we define $S(p_s)$ heuristically with a down vector d as follows:

$$S(p_s) = \text{normalize}(n_s(p_s) + d \min(n_s(p_s) \cdot d + 1, 1)) \quad (2)$$

\mathbb{U} is an undefined boundary without specific conditions. For example, it accompanies neck collision and no valid 3D orientation. We discretize $f_o(p)$ on a regular grid, initialize it by filling the interior space with zeros, and iteratively solve it for each element of XYZ with a successive over-relaxation method. To avoid instability, the norm constraint is enforced after the iterations.

Finally, we convert the 3D orientation field into guide strands. Starting from root vertices of \mathcal{V} on \mathbb{S} , we traverse the orientation of the voxels in sequence until reaching \mathbb{H} , generating guide strands.

3.2. Hierarchical Strand Optimization

We optimize hair line segments with a novel DR algorithm with reparameterization. Guide-child hierarchy is incorporated into our optimization framework.

Hair Strands as Line Segments

The representation of hair follows the common practice in real-time rendering [58]. Figure 5 displays our hair geometry. The geometry of hair $\mathcal{G} = \{\mathcal{V}, \mathcal{F}\}$ is a collection of line segments. Here, \mathcal{V} refers to the vertex positions corresponding to the division points, and \mathcal{F} represents the connectivity between upper and lower vertices, expressing spline curves [4]. During rendering, after tessellation, \mathcal{G} is further converted into camera-facing triangle strips $\mathcal{G} = \{\mathcal{V}, \mathcal{F}\}$ and rasterized. The tip becomes a single triangle. \mathcal{V} and \mathcal{F} denote the vertex positions and indices, respectively.




Figure 5. Line segments \mathcal{G} are subdivided and converted into billboard mesh \mathcal{G} for rasterization. The \mathcal{G} thickness is typically less than one pixel.




Figure 6. Adjacency used in our Laplacian ($k = 4$).

A two-stage structure of guide and child hair is employed to express hairstyles. Guide hair grows from each vertex of the scalp mesh, while child hair grows from sampled positions by Sobol sequence [47]. Child hair shape is linearly interpolated with the nearest four guide hairs. 653 guide strands and 50,000 child strands are used in this paper. For children, the geometry $\mathcal{G}_c = \{\mathcal{V}_c, \mathcal{F}_c\}$ are defined in the same way, and the same rasterization is applied.

Differentiable Rendering of Line Segments

After billboard mesh \mathcal{G} is generated, those triangles are rasterized with the help of hardware, and anti-aliasing (AA) for hair is applied to the rasterized screen-space buffer.

We designate the rasterized color at pixel position s as $c(s)$. N_8 and $s_n, n \in N_8$ denote the 8 neighboring pixels and their positions, respectively. $c_{bl}(s, s_n, \mathcal{G})$ represents the blended color. The function $\text{tri}(s)$ returns the triangle ID, while $0 \leq r(s, s_n, \mathcal{G}) \leq 1$ defines a function that returns the distance from pixel s to the edge of the triangle spanning s_n . If multiple edges cross a pixel boundary, the one with the closest depth is chosen. The screen-space gradient that can move \mathcal{V} is generated via this function that accesses \mathcal{G} . The AA color $c_{aa}(s, \mathcal{G})$ is the average of the blended colors.

$$c_{aa}(s, \mathcal{G}) = (c(s) + \sum_{n \in N_8} c_{bl}(s, s_n, \mathcal{G})) / (|N_8| + 1) \quad (3)$$

$$c_{bl}(s, s_n, \mathcal{G}) = \begin{cases} c_{bl'}(s, s_n, \mathcal{G}) & \text{if } \text{tri}(s) \neq \text{tri}(s_n) \\ c(s) & \text{otherwise} \end{cases} \quad (4)$$

$$c_{bl'}(s, s_n, \mathcal{G}) = r(s, s_n, \mathcal{G})c(s) + (1 - r(s, s_n, \mathcal{G}))c(s_n) \quad (5)$$

While we have referred to c_* as *colors* for convenience, this approach can be extended to handle any rasterized vertex attributes, such as silhouette and depth.

Our approach draws inspiration from nvdiffrast [21], a DR for meshes. The AA of nvdiffrast generates gradients only on the edge pixels where the occlusion actually occurs. Instead, all pixels with different IDs of adjacent triangles are updated for smoother gradients. AA can keep finer geometric details than soft rasterization [22, 39] used in Neural-Haircut [45]. As opposed to splatting [8], we leverage the

distance between pixels and geometric edges for stronger gradient propagation. The comparison of AAs is available in the supplementary material.

Reparameterization

Even if AA smooths the gradients in the screen-space, severe occlusions and non-deterministic rasterization of strands with less than one pixel width result in sparse propagation of gradients into the geometry. To address this issue, we propose a reparameterization of hair geometries as regularization, which is inspired by a mesh reparameterization [32]. We introduce a Laplacian matrix \mathbf{L} for line segments, considering both geometric connectivity and spatial proximity, to transform the parameter space for optimization. An example of the Laplacian we propose is shown in Figure 6.

$$(\mathbf{L})_{ij} = \begin{cases} -w_{ij} & \text{if } (i, j) \in \{\mathcal{F} \cup \mathcal{N}\} \\ \sum_{(i, k) \in \{\mathcal{F} \cup \mathcal{N}\}} w_{ik} & \text{if } i = j \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

Here, \mathcal{N} is a set of combinations of neighboring vertices obtained by searching for the k -Nearest Neighbors (k NN) based on Euclidean distance for each line segment. The number of neighbors k allows us to control the effect of spatial proximity relations. In all experiments presented in this paper, $w_{ij} = 1$ is used. Let \mathbf{x} be a matrix assembled from $x \in \mathcal{V}$. Using \mathbf{L} , we reparameterize the vertex positions \mathbf{x} in Cartesian coordinate into the values \mathbf{u} in the differential coordinate, where dense gradients are delivered. \mathbf{I} is the identity matrix, and the parameter λ controls regularization effect.

$$\mathbf{u} = (\mathbf{I} + \lambda \mathbf{L})\mathbf{x} \quad (7)$$

Guide Hair Optimization

We efficiently perform strand optimization following the Coarse-to-Fine strategy. First, the vertex positions of guide hair \mathcal{V} are optimized. Reparameterization with $k = K_g$ is applied for guide hair. We rasterize the strands at a thickness close to the actual hair, 0.2 mm. L_g is minimized using the Adam optimizer [19] in I_g iterations.

$$R_b = w_{stick} * R_{stick} + w_{root} * R_{root} + w_c * R_c \quad (8)$$

$$L_g = w_d * L_d + w_m * L_m + w_t * L_t + R_b \quad (9)$$

Here, R_b constitutes a base regularizer, and it consists of three parts. R_{stick} is an L1 regularizer to prevent the hair from penetrating the scalp, which is computed from the depth of the penetrated strands and that of the scalp. R_{root} is an L1 regularizer to ensure that the roots of the hair keep their initial positions as possible. R_c is a curvature regularizer for guide hair, representing the sum of the curvatures

formed by adjacent line segments. L_d is an L1 depth loss calculated between the depth values rendered by the raw mesh and the strands. L_m is an L1 mask loss computed between the hair mask extracted from the input image and one of the rendered strands. L_t is a 3D tangent loss computed as the sum of cosine losses between the estimated 3D orientation and the strands both rendered in screen-space. w_* are the weights for each loss.

Child Hair Optimization

In this step, we abandon the guide interpolation and optimize the vertex positions of child hair, \mathcal{V}_c , for finer alignment in the two stages. First, We apply reparameterization with $k = K_c$ for child hair and perform optimization in I_c^0 iterations. After the first stage, we relax the conditions by setting $k = 0$ and execute the final alignment in I_c^1 iterations. Both of these steps share a common loss term, L_c , which is minimized using the Adam optimizer.

$$L_c = w_d * L_d + w_m * L_m + w_o * L_o + R_b \quad (10)$$

L_o is a 2D orientation loss computed as the sum of absolute cosine losses between the 2D orientation extracted from the input image and that of the rendered strands.

4. Experiments

Qualitative and quantitative comparisons were carried out with state-of-the-art strand reconstruction methods that use multi-view images: **LPMVS** [31], **Strand Integration** [29], and **NeuralHaircut** [45]. An open CPU implementation made by the authors of Strand Integration was used for LPMVS, while the other methods were tested with official implementations. All methods require multi-view images with camera parameters as input. NeuralHaircut was trained for 300k iterations in Stage 1 (surface) and 200k iterations in Stage 2 (strand). The proposed method is implemented on the top of Blender, PyTorch, and Dressi-AD, a Vulkan-based DR framework with hardware rasterizer [49]. For Adam, learning rate was set to 0.001 with $\beta_1 = 0.9$ and $\beta_2 = 0.999$. Other hyperparameters were set as follows: $\lambda = 50$, $w_{root} = 1.0$, $w_{stick} = 0.1$, $w_d = 0.01$, $w_c = 0.01$, $w_m = 1$, $w_t = 1$, $w_o = 1$, $K_g = 4$, $K_c = 4$, $I_g = 2000$, $I_c^0 = 2000$, and $I_c^1 = 1000$. Our input raw mesh is reconstructed by OpenMVS [5] unless otherwise noted. Detailed experimental settings and more results are available in the supplementary material.

4.1. Synthetic data

We numerically evaluated our method on hair models prepared by Yuksel et al. [59]. The 58 images were ray-traced in Blender Cycles using a virtual camera on a hemisphere under uniform lighting. We did not use GTs as

Table 1. Speed comparison on synthetic data. We measured the time taken to generate strands starting from multiple image and camera parameter inputs on the same machine. The time taken for our surface reconstruction by OpenMVS is set to 1, and relative times are shown for the others.

Method	LPMVS	Strand Integ.	NeuralHaircut	Ours
Surface recon.	N/A	N/A	1080	1
Strand recon.	57	81	2160	25

input, except for camera parameters. We follow previous studies [31, 40, 45] that measure precision, recall, and F-score between reconstructed strands and GTs. Nevertheless, we chose the 3D correspondence to evaluate the internal strands. 3D space is searched per source sample to judge whether at least one destination sample is within the distance and angle error thresholds, and the quantities are computed. Moreover, while angular error evaluation in the previous studies accepts an ambiguity of 180° , we evaluated it in a range of 360° to assess absolute hair flow.

Quantitative comparison with existing methods and ablation study are shown in Table 2. Our full pipeline shows better values than the other methods in all criteria. In particular, the much higher recall values indicate that our method successfully recovers internal hair directions that were difficult to handle with existing methods. The effectiveness of each component in the proposed pipeline is also validated. For straight hair with simple internal flow, our strand initialization worked well and showed good scores even without DR. For complex curly hair, the influence of other modules, global optimization, reparameterization, and guide-child hierarchy becomes more pronounced.

The performance of each method was also compared. Table 1 shows the time taken to process Curly Hair. Typically, NeuralHaircut takes a couple of days, and LPMVS and Strand Integration require a few hours, but our pipeline finishes in less than one hour.

4.2. Real data

We show comparisons on H3DS dataset [38] in Figure 7. GT raw meshes and camera parameters are used for all methods but they are not very accurate. Our method is capable of reconstructing reasonable results despite the limited accuracy of the input data.

Figure 8 visualizes results on a monocular video sequence. In even worse calibrations, our method robustly reconstructs the strands. Hair length editing is also performed to prove that our method can reconstruct accurate hair flow.

We also conducted comparisons on a well-calibrated studio setup. The 58 images taken by cameras positioned on a hemisphere under uniform illumination are used. The results are shown in Figure 1. Ours reconstructs better strands than the existing methods in terms of direction. Figure 9 displays our other results that the challenging hairstyles

Table 2. Quantitative comparison with existing methods and ablation study on synthetic data. **P**, **R**, and **F1** denote precision, recall, and F1 score, respectively. Higher is better. The lower rows describe the values of our full pipeline and ours without individual modules. w/o DR: DR optimization is not applied, and the initialized strands are evaluated. w/o guide opt.: Child strands are optimized from the beginning of the DR step. w/o reparam.: Reparameterization is disabled. w/o \mathcal{N} : Only \mathcal{N} is abandoned in the reparameterization. w/o strand init.: Strands are initialized by straight lines parallel to the normal of the scalp. w/o global opt.: Only gravity heuristic is applied to the initial 3D orientation, and 180° ambiguity is accepted on the other steps.

Threshold Measure	Straight Hair									Curly Hair								
	1mm/10°			2mm/20°			3mm/30°			1mm/10°			2mm/20°			3mm/30°		
	P	R	F1	P	R	F1	P	R	F1	P	R	F1	P	R	F1	P	R	F1
LPMVS [31]	29.9	27.5	28.6	39.7	52.2	45.1	42.9	66.8	52.2	18.4	6.1	9.1	32.8	15.3	20.9	37.2	23.9	29.1
Strand Integration [29]	33.4	34.1	33.7	42.4	55.4	48.1	44.8	66.8	53.6	19.4	6.7	9.9	34.3	16.1	21.9	38.5	23.6	29.3
NeuralHaircut [45]	50.2	14.9	23.0	76.1	29.2	42.2	85.6	38.4	53.1	20.9	3.9	6.6	58.1	14.5	23.2	80.0	27.3	40.7
Ours	60.3	46.4	52.5	88.2	84.3	86.2	94.5	93.6	94.1	38.3	23.6	29.2	79.1	61.0	68.9	90.0	81.0	85.3
Ours (w/o DR)	65.4	41.8	51.0	88.6	78.8	83.4	93.3	88.6	90.9	22.1	15.7	18.4	59.2	56.1	57.6	75.8	82.0	78.8
Ours (w/o guide opt.)	61.1	46.8	53.0	86.8	86.1	86.5	92.7	95.0	93.9	36.8	22.9	28.2	77.4	60.6	68.0	88.8	80.7	84.6
Ours (w/o reparam.)	7.9	42.5	13.4	24.0	97.1	38.5	38.8	99.9	55.9	6.4	29.5	10.6	23.1	93.5	37.1	39.3	99.6	56.4
Ours (w/o \mathcal{N})	59.5	46.9	52.5	86.6	85.4	86.0	92.8	94.3	93.5	36.5	23.2	28.3	76.4	60.7	67.6	87.6	80.7	84.0
Ours (w/o strand init.)	5.3	1.0	1.6	23.4	7.0	10.7	37.7	18.6	24.9	9.3	4.9	6.4	22.3	23.8	23.0	32.4	48.0	38.7
Ours (w/o global opt.)	58.1	45.7	51.2	85.5	85.8	85.7	91.8	94.7	93.2	26.1	13.6	17.9	63.5	46.4	53.6	75.3	68.5	71.8

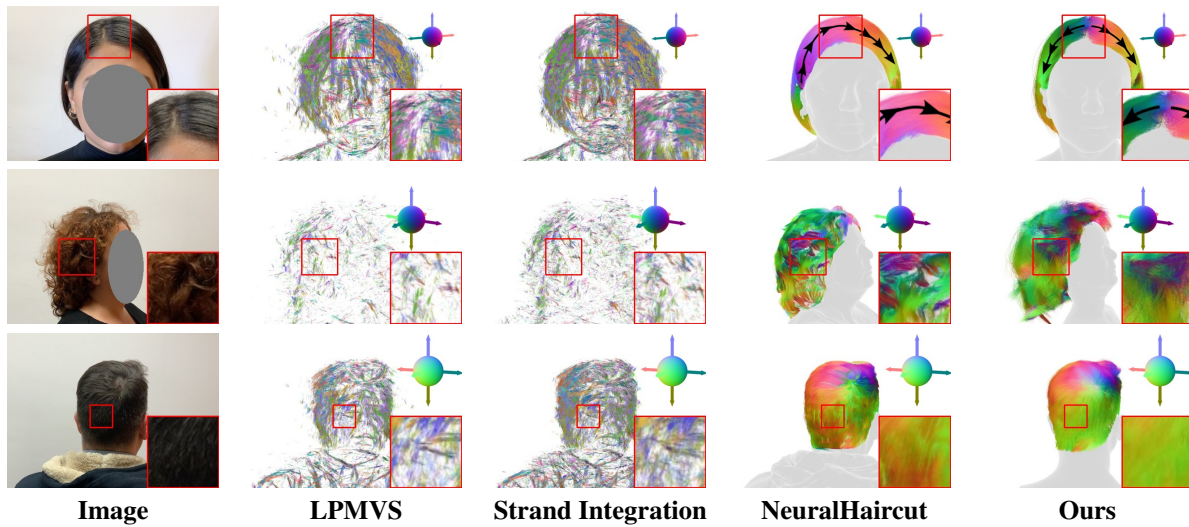


Figure 7. Comparison on a real-world multi-view dataset, H3DS [38]¹. Color and arrows represent 3D orientation of each strand. In the upper row, our method identifies the hair parting. In the middle row, ours reconstructs dense hair with no visible white scalp. In the bottom row, a smooth hair flow is estimated by ours. NeuralHaircut struggles in all cases. LPMVS and Strand Integration are prone to flying noise.

are successfully handled. Re-rendering comparison with reconstructed strands is shown in Figure 10. Our strands are shaded photorealistic, indicating that our method’s output is portable among rendering engines. Furthermore, the density distribution of our hair is more reasonable than NeuralHaircut’s. Figure 11 visualizes the comparison of physics simulation starting from the reconstructed strands. Although NeuralHaircut suffers from incorrect directions, ours exhibits reasonable behavior.

4.3. Limitations

Our method is affected by the input raw mesh quality. If the raw mesh shape has a notable difference from the subject, our strand reconstruction deteriorates, as shown in Figure 12. In addition, the assumption of a smooth flow makes it challenging to handle discontinuous hairstyles such as braids. Moreover, protruding strands can disrupt the overall flow, limiting our accuracy, particularly for highly curly or spiky hair. Finally, as shown in Figure 10, shaded hair color deviates from the actual one because our method does not recover material and lighting.

¹The images of this dataset are only used for testing and comparison with existing methods and not used for algorithm improvement or training.




Figure 8. Comparison on a hand-held monocular video captured by a smartphone [45]. On each method, the left image shows the original reconstruction, and the diagram on the right shows the hair length edited in half. Our method demonstrates robust reconstruction even under a severe capturing condition. Moreover, the half length image indicates our hair is editable thanks to the correct direction.




Figure 9. Robustness against challenging hairstyles. Our method can reconstruct (a) a half bald head and (b) tied hair.




Figure 10. Left and middle: Re-rendering comparison. The camera image is put on the top for hairstyle reference. The images are ray-traced by Blender Cycles, using the common hair material and lighting. Incorrect orientation of NeuralHaircut made artifacts at meeting points of opposite hair flow, which ours does not have. Right: Volumetric slice of the middle view with certain near and far planes. NeuralHaircut's hair is concentrated on the surface and scanty inside, whereas our method fills inside uniformly.

5. Conclusion

In this paper, we have introduced Dr.Hair, a novel method for reconstructing detailed human hair strands from multi-view images. Our approach recovers consistent surface orientations, estimates the internal flow using differential equations, and performs optimization based on differentiable rendering, leveraging the hierarchical relationship between guide and child strands. Our method's effective-




Figure 11. Physics simulation with gravity and certain hair stiffness applied to the same subject in Figure 1. The left image is the initial state given by reconstruction, and the right one is after simulation. NeuralHaircut's hairs are oriented from bottom to top, so the simulated result is severely affected by unnatural sagging. Thanks to the correct direction, our hair behaves naturally.




Figure 12. A limitation of our method. Raw mesh with inaccurate top geometry degenerates hair flow.

ness has been demonstrated through both qualitative and quantitative evaluations. Our method is capable of reconstructing a wide variety of hairstyles grown from a scalp without relying on priors trained on synthetic datasets, which are typically created through labor-intensive manual work. Moreover, our proposed method outperforms existing methods in terms of processing speed. We believe that our method can significantly contribute to the development of a cost-effective, photorealistic human digitization system.

References

- [1] Autodesk. Maya xgen, 2023. <https://www.autodesk.com/products/maya/>. 2
- [2] Blender. Hair nodes, 2023. https://docs.blender.org/manual/en/3.6/modeling/geometry_nodes/hair/index.html. 2
- [3] Chen Cao, Tomas Simon, Jin Kyu Kim, Gabe Schwartz, Michael Zollhoefer, Shun-Suke Saito, Stephen Lombardi, Shih-En Wei, Danielle Belko, Shou-I Yu, Yaser Sheikh, and Jason Saragih. Authentic volumetric avatars from a phone scan. *ACM Trans. Graph.*, 41(4), 2022. 2, 3
- [4] Edwin Catmull and Raphael Rom. A class of local interpolating splines. In *Computer aided geometric design*, pages 317–326. Elsevier, 1974. 4
- [5] Dan Cernea. OpenMVS: Multi-view stereo reconstruction library. <https://cdcseacave.github.io/openMVS>, 2020. 6
- [6] Menglei Chai, Lvdi Wang, Yanlin Weng, Xiaogang Jin, and Kun Zhou. Dynamic hair manipulation in images and videos. *ACM Trans. Graph.*, 32(4), 2013. 2
- [7] 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), 2016. 2, 3
- [8] Forrester Cole, Kyle Genova, Avneesh Sud, Daniel Vlasic, and Zhoutong Zhang. Differentiable surface rendering via non-differentiable sampling. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6088–6097, 2021. 2, 5
- [9] Ephere. Ornatrx for maya, 2023. <https://ephere.com/plugins/autodesk/maya/ornatrix/>. 2
- [10] Liwen Hu, Chongyang Ma, Linjie Luo, and Hao Li. Robust hair capture using simulated examples. *ACM Trans. Graph.*, 33(4), 2014. 2
- [11] Liwen Hu, Chongyang Ma, Linjie Luo, Li-Yi Wei, and Hao Li. Capturing braided hairstyles. *ACM Trans. Graph.*, 33(6), 2014. 2
- [12] Liwen Hu, Chongyang Ma, Linjie Luo, and Hao Li. Single-view hair modeling using a hairstyle database. *ACM Trans. Graph.*, 34(4), 2015. 2, 3
- [13] Liwen Hu, Derek Bradley, Hao Li, and Thabo Beeler. Simulation-Ready Hair Capture. *Computer Graphics Forum*, 2017. 2
- [14] Liwen Hu, Shunsuke Saito, Lingyu Wei, Koki Nagano, Jaewoo Seo, Jens Fursund, Iman Sadeghi, Carrie Sun, Yen-Chun Chen, and Hao Li. Avatar digitization from a single image for real-time rendering. *ACM Trans. Graph.*, 36(6), 2017. 2, 3
- [15] Wenzel Jakob, Jonathan T. Moon, and Steve Marschner. Capturing hair assemblies fiber by fiber. *ACM Trans. Graph.*, 28(5):1–9, 2009. 2
- [16] 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>. 2
- [17] Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. Poisson surface reconstruction. In *Proceedings of the fourth Eurographics symposium on Geometry processing*, page 0, 2006. 4
- [18] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4), 2023. 3
- [19] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5
- [20] Zhiyi Kuang, Yiyang Chen, Hongbo Fu, Kun Zhou, and Youyi Zheng. Deepmvshair: Deep hair modeling from sparse views. In *SIGGRAPH Asia 2022 Conference Papers*, New York, NY, USA, 2022. Association for Computing Machinery. 2, 3, 4
- [21] Samuli Laine, Janne Hellsten, Tero Karras, Yeongho Seol, Jaakko Lehtinen, and Timo Aila. Modular primitives for high-performance differentiable rendering. *ACM Transactions on Graphics*, 39(6), 2020. 2, 5
- [22] Shichen Liu, Tianye Li, Weikai Chen, and Hao Li. Soft rasterizer: A differentiable renderer for image-based 3d reasoning. *The IEEE International Conference on Computer Vision (ICCV)*, 2019. 2, 3, 5
- [23] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: Learning dynamic renderable volumes from images. *ACM Trans. Graph.*, 38(4):65:1–65:14, 2019. 2, 3
- [24] Stephen Lombardi, Tomas Simon, Gabriel Schwartz, Michael Zollhoefer, Yaser Sheikh, and Jason Saragih. Mixture of volumetric primitives for efficient neural rendering. *ACM Transactions on Graphics (ToG)*, 40(4):1–13, 2021. 2, 3
- [25] 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. *arXiv preprint arXiv:2402.10483*, 2024. 3
- [26] Linjie Luo, Hao Li, Thibaut Weise, Sylvain Paris, Mark Pauly, and Szymon Rusinkiewicz. Dynamic hair capture. *Rapp. tech. Princeton University*, 2011. 2
- [27] Linjie Luo, Hao Li, Sylvain Paris, Thibaut Weise, Mark Pauly, and Szymon Rusinkiewicz. Multi-view hair capture using orientation fields. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1490–1497. IEEE, 2012. 2
- [28] Linjie Luo, Hao Li, and Szymon Rusinkiewicz. Structure-aware hair capture. *ACM Trans. Graph.*, 32(4), 2013. 2
- [29] Ryota Maeda, Kenshi Takayama, and Takafumi Taketomi. Refinement of hair geometry by strand integration. *Computer Graphics Forum (proceedings of Pacific Graphics)*, 42(7), 2023. 1, 2, 6, 7
- [30] 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 *European Conference on Computer Vision*, pages 405–421. Springer, 2020. 2, 3
- [31] Giljoo Nam, Chenglei Wu, Min H. Kim, and Yaser Sheikh. Strand-accurate multi-view hair capture. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1, 2, 3, 4, 6, 7

- [32] Baptiste Nicolet, Alec Jacobson, and Wenzel Jakob. Large steps in inverse rendering of geometry. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)*, 40(6), 2021. [2](#), [5](#)
- [33] Merlin Nimier-David, Delio Vicini, Tizian Zeltner, and Wenzel Jakob. Mitsuba 2: A retargetable forward and inverse renderer. *Transactions on Graphics (Proceedings of SIGGRAPH Asia)*, 38(6), 2019. [2](#)
- [34] Sylvain Paris, Hector M Briceno, and François X Sillion. Capture of hair geometry from multiple images. *ACM transactions on graphics (TOG)*, 23(3):712–719, 2004. [2](#), [4](#)
- [35] 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):1–9, 2008. [2](#)
- [36] Patrick Pérez, Michel Gangnet, and Andrew Blake. Poisson image editing. *ACM Trans. Graph.*, 22(3):313–318, 2003. [4](#)
- [37] Robert Clay Prim. Shortest connection networks and some generalizations. *The Bell System Technical Journal*, 36(6):1389–1401, 1957. [4](#)
- [38] Eduard Ramon, Gil Triginer, Janna Escur, Albert Pumarola, Jaime Garcia, Xavier Giro-i Nieto, and Francesc Moreno-Noguer. H3d-net: Few-shot high-fidelity 3d head reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5620–5629, 2021. [6](#), [7](#)
- [39] Nikhila Ravi, Jeremy Reizenstein, David Novotny, Taylor Gordon, Wan-Yen Lo, Justin Johnson, and Georgia Gkioxari. Accelerating 3d deep learning with pytorch3d. *arXiv:2007.08501*, 2020. [2](#), [3](#), [5](#)
- [40] 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. [2](#), [3](#), [6](#)
- [41] Darius Rückert, Linus Franke, and Marc Stamminger. Adop: Approximate differentiable one-pixel point rendering. *ACM Trans. Graph.*, 41(4), 2022. [2](#), [3](#)
- [42] Shunsuke Saito, Liwen Hu, Chongyang Ma, Hikaru Ibayashi, Linjie Luo, and Hao Li. 3d hair synthesis using volumetric variational autoencoders. *ACM Trans. Graph.*, 37(6), 2018. [2](#), [3](#)
- [43] Yuefan Shen, Changgeng Zhang, Hongbo Fu, Kun Zhou, and Youyi Zheng. Deepsketchhair: Deep sketch-based 3d hair modeling. *IEEE transactions on visualization and computer graphics*, 27(7):3250–3263, 2020. [3](#)
- [44] Yuefan Shen, Shunsuke Saito, Ziyang Wang, Olivier Maury, Chenglei Wu, Jessica Hodgins, Youyi Zheng, and Giljoo Nam. Ct2hair: High-fidelity 3d hair modeling using computed tomography. *ACM Transactions on Graphics*, 42(4):1–13, 2023. [3](#)
- [45] Vanessa Sklyarova, Jenya Chelishev, Andreea Dogaru, Igor Medvedev, Victor Lempitsky, and Egor Zakharov. Neural haircut: Prior-guided strand-based hair reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 19762–19773, 2023. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#), [8](#)
- [46] Vanessa Sklyarova, Egor Zakharov, Otmar Hilliges, Michael J Black, and Justus Thies. Haar: Text-conditioned generative model of 3d strand-based human hairstyles. *ArXiv*, 2023. [3](#)
- [47] Il’ya Meerovich Sobol’. On the distribution of points in a cube and the approximate evaluation of integrals. *Zhurnal Vychislitel’noi Matematiki i Matematicheskoi Fiziki*, 7(4):784–802, 1967. [5](#)
- [48] Tiancheng Sun, Giljoo Nam, Carlos Aliaga, Christophe Hery, and Ravi Ramamoorthi. Human Hair Inverse Rendering using Multi-View Photometric data. In *Eurographics Symposium on Rendering - DL-only Track*. The Eurographics Association, 2021. [2](#)
- [49] Yusuke Takimoto, Hiroyuki Sato, Hikari Takehara, Keishiro Uragaki, Takehiro Tawara, Xiao Liang, Kentaro Oku, Wataru Kishimoto, and Bo Zheng. Dressi: A Hardware-Agnostic Differentiable Renderer with Reactive Shader Packing and Soft Rasterization. *Computer Graphics Forum*, 2022. [2](#), [6](#)
- [50] Lvdi Wang, Yizhou Yu, Kun Zhou, and Baining Guo. Example-based hair geometry synthesis. In *ACM SIGGRAPH 2009 Papers*, New York, NY, USA, 2009. Association for Computing Machinery. [3](#)
- [51] 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. [2](#), [3](#)
- [52] Ziyang Wang, Giljoo Nam, Tuur Stuyck, Stephen Lombardi, Michael Zollhöfer, Jessica Hodgins, and Christoph Lassner. Hvh: Learning a hybrid neural volumetric representation for dynamic hair performance capture. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6143–6154, 2022. [3](#)
- [53] Ziyang Wang, Giljoo Nam, Tuur Stuyck, Stephen Lombardi, Chen Cao, Jason Saragih, Michael Zollhöfer, Jessica Hodgins, and Christoph Lassner. Neuwigs: A neural dynamic model for volumetric hair capture and animation. In *CVPR*, pages 8641–8651, 2023. [3](#)
- [54] Yichen Wei, Eyal Ofek, Long Quan, and Heung-Yeung Shum. Modeling hair from multiple views. In *ACM SIGGRAPH 2005 Papers*, page 816–820, New York, NY, USA, 2005. Association for Computing Machinery. [2](#)
- [55] Keyu Wu, Yifan Ye, Lingchen Yang, Hongbo Fu, Kun Zhou, and Youyi Zheng. Neuralhdhair: Automatic high-fidelity hair modeling from a single image using implicit neural representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1526–1535, 2022. [2](#), [3](#)
- [56] Lingchen Yang, Zefeng Shi, Youyi Zheng, and Kun Zhou. Dynamic hair modeling from monocular videos using deep neural networks. *ACM Trans. Graph.*, 38(6), 2019. [3](#)
- [57] Wang Yifan, Felice Serena, Shihao Wu, Cengiz Öztireli, and Olga Sorkine-Hornung. Differentiable surface splatting for point-based geometry processing. *ACM Trans. Graph.*, 38(6), 2019. [2](#), [3](#)
- [58] Cem Yuksel and Sarah Tariq. Advanced techniques in real-time hair rendering and simulation. In *ACM SIGGRAPH 2010 Courses*, New York, NY, USA, 2010. Association for Computing Machinery. [4](#)

- [59] Cem Yuksel, Scott Schaefer, and John Keyser. Hair meshes. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia 2009)*, 28(5):166:1–166:7, 2009. 6
- [60] Meng Zhang and Youyi Zheng. Hair-gan: Recovering 3d hair structure from a single image using generative adversarial networks. *Visual Informatics*, 3(2):102–112, 2019. 2, 3, 4
- [61] 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), 2017. 2, 3
- [62] Meng Zhang, Pan Wu, Hongzhi Wu, Yanlin Weng, Youyi Zheng, and Kun Zhou. Modeling hair from an rgb-d camera. *ACM Trans. Graph.*, 37(6), 2018. 2
- [63] Yuanwei Zhang, Shinichi Kinuwaki, and Nobuyuki Umetani. Energyhair: Sketch-based interactive guide hair design using physics-inspired energy. In *Graphics Interface 2022*, 2022. 2
- [64] Yujian Zheng, Zirong Jin, Moran Li, Haibin Huang, Chongyang Ma, Shuguang Cui, and Xiaoguang Han. Hairstep: Transfer synthetic to real using strand and depth maps for single-view 3d hair modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12726–12735, 2023. 2, 3
- [65] 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 *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 235–251, 2018. 2, 3
- [66] Yuxiao Zhou, Menglei Chai, Alessandro Pepe, Markus Gross, and Thabo Beeler. Groomgen: A high-quality generative hair model using hierarchical latent representations. *ACM Trans. Graph.*, 42(6), 2023. 3