

# NeuralHDHair: Automatic High-fidelity Hair Modeling from a Single Image Using Implicit Neural Representations

Keyu Wu<sup>1\*</sup> Yifan Ye<sup>1\*</sup> Lingchen Yang<sup>2</sup> Hongbo Fu<sup>3</sup> Kun Zhou<sup>1</sup> Youyi Zheng<sup>1†</sup>

<sup>1</sup> Zhejiang University <sup>2</sup> ETH Zurich <sup>3</sup> City University of Hong Kong

## Abstract

Undoubtedly, high-fidelity 3D hair plays an indispensable role in digital humans. However, existing monocular hair modeling methods are either tricky to deploy in digital systems (e.g., due to their dependence on complex user interactions or large databases) or can produce only a coarse geometry. In this paper, we introduce NeuralHDHair, a flexible, fully automatic system for modeling high-fidelity hair from a single image. The key enablers of our system are two carefully designed neural networks: an IRHairNet (Implicit representation for hair using neural network) for inferring high-fidelity 3D hair geometric features (3D orientation field and 3D occupancy field) hierarchically and a GrowingNet (Growing hair strands using neural network) to efficiently generate 3D hair strands in parallel. Specifically, we perform a coarse-to-fine manner and propose a novel voxel-aligned implicit function (VIFu) to represent the global hair feature, which is further enhanced by the local details extracted from a hair luminance map. To improve the efficiency of a traditional hair growth algorithm, we adopt a local neural implicit function to grow strands based on the estimated 3D hair geometric features. Extensive experiments show that our method is capable of constructing a high-fidelity 3D hair model from a single image, both efficiently and effectively, and achieves the-state-of-the-art performance.

## 1. Introduction

As one of the most distinctive human characteristics, hair plays an indispensable role in digital humans [4, 10, 12, 15, 17]. Undoubtedly, a high-fidelity 3D hair model can significantly improve the realism of a virtual human. However, the existing single-view-based hair modeling methods [6, 11, 28, 36, 40] cannot sufficiently satisfy the require-

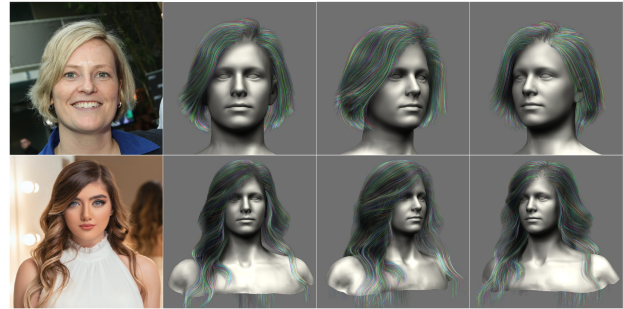


Figure 1. Given a single image, our NeuralHDHair reconstructs a high-fidelity 3D hair model.

ments of human digitalization in terms of flexibility, simplicity, and realism. On the one hand, data-driven methods [6, 11] could achieve high-fidelity results, but are complex and not very robust, e.g., entailing a sophisticated searching and matching process based on a large hair dataset. On the other hand, deep-learning based methods [28, 36, 40] are lightweight and flexible to deploy but could only achieve coarse results. Thus, in this work, we consider the problem of automatic high-fidelity 3D hair modeling from a single image utilizing a learning-based method.

Unlike other parts of the human body, the hair structure is more challenging to describe and extract due to remarkably intricate structures of interweaving strands, leading to extreme difficulty in reconstructing a high-fidelity 3D hair model only from a single view. Generally, almost all the existing methods tackle this problem in two steps: first estimating a 3D orientation field based on a 2D orientation map extracted from an input image and then synthesizing hair strands from the 3D orientation field. However, there are some problems with such a two-step mechanism. First, since a 2D orientation map is only a filtered version of the input hair growing information [24], using the 2D orientation map alone to bridge the domain gap between real and synthetic data would unavoidably lose hair details (e.g., the relationship of occlusive strands [33]). Second, the existing methods for inferring 3D orientation fields are either time-consuming due to the use of a complex searching and

\*The first two authors contributed equally. The authors are affiliated with the State Key Lab of CAD&CG. †Corresponding author: Youyi Zheng.

matching process based on a large hair dataset [6, 11], or liable to over-smoothness due to the use of deep networks to directly achieve image-to-voxel inference [36]. Third, the conventional hair growth algorithm [31, 39] to extract strands from the estimated 3D orientation field are inefficient and not conducive to one-shot hair modeling. Although Zhou et al. [40] have attempted to ignore the hair growth procedure by directly regressing 3D hair strands, their reconstruction results are generally unsatisfactory (see Sec. 4.2). Based on the above observations, we seek to build a fully automatic and efficient hair modeling method that can reconstruct a 3D hair model from a single image with fine-grained features (Fig. 1) while exhibiting high flexibility, e.g., reconstructing the hair model only needs one forward pass of the networks.

We found that implicit functions have excellent performance and great potential in representing [25] and inferring [23] 3D shapes. For example, Saito et al. [29] introduced PIFu to reconstruct a whole human body from a single image, including hair. However, the quality of their hair reconstruction results is less satisfactory. We observe that unlike human body modeling, which cares only the surface geometry of a human body, 3D hair modeling needs to consider both exterior shape and interior features, which are difficult to represent by a pixel-aligned implicit function [30]. To address this issue, we propose IRHairNet, which imposes a coarse-to-fine strategy to produce a high-fidelity 3D orientation field. Specifically, we introduce a novel voxel-aligned implicit function (VIFu) to extract global information from a 2D orientation map in the coarse module. Meanwhile, to supplement the lost local details in the 2D orientation map, we exploit a high-resolution luminance map to extract local feature and combine it with the global feature in the fine module for high-fidelity hair modeling.

To efficiently synthesize a hair-strand model from the 3D orientation field, we introduce GrowingNet, a deep learning-based hair growth method by leveraging a local implicit grid representation [13]. This is based on a key observation that although hair geometric shapes and growing directions vary globally, they share similar features at a specific local scale. Thus, we can extract a high-level latent code for each local 3D orientation patch, and then train a neural implicit function (a decoder) to grow strands inside it based on this latent code. After each growing step, new local patches centered at the ends of the strands will be used to proceed with the growing. After training, it is applicable to 3D orientation fields of arbitrary resolution.

IRHairNet and GrowingNet form the core of NeuralHdHair, a novel automatic monocular hair modeling method. We conduct extensive experiments, including comparison experiments and ablation study, and the results show that NeuralHdHair outperforms all existing monocular hair reconstruction methods [36, 40].

In summary, the main contributions of our work include:

- We introduce a novel fully-automatic monocular hair modeling framework, significantly outperforming the state-of-the-art methods.
- We introduce a coarse-to-fine hair modeling neural network (IRHairNet), where we use a novel voxel-aligned implicit function and a luminance map to enrich local details for high-quality hair modeling.
- We propose a novel hair growing network (GrowingNet) based on a local implicit function to efficiently generate strand models with arbitrary resolution, and is an order of magnitude faster than prior methods.

## 2. Related Work

**Hair Modeling from Single Images.** With the development of image sensing and computer graphics, 3D hair modeling [6, 8, 11, 16, 18, 24, 28, 39] has been extensively explored. Compared with multi-view-based techniques [16, 39], which are typically limited to carefully regulated environments and complex hardware setups, single-view hair modeling methods show their significant strength in feasibility, generality, and efficiency. The pioneering single-view-based methods rely on different kinds of priors such as layer boundary [7, 8] or shading cues [5] to reconstruct a hair model from a single image. They often require additional user interaction and cannot reasonably generate invisible regions. Subsequently, [6, 11] build a synthetic 3D hair database and produce impressive results from a single image based on data-driven methods. However, their reconstruction results depend highly on the quality and diversity of the database, and a large database will cause deployment difficulties. To this end, [28, 36, 40] introduce a lightweight technique for monocular hair modeling without the requirement of a pre-built database or complex hardware setups and a controlled environment. However, their methods focus on producing globally plausible models but ignore local details. In contrast, our end-to-end approach can fully automatically generate high-fidelity hair models with fine-grained features.

**Implicit Neural Representations.** Most recently, extensive studies [3, 20, 25, 26] have been conducted for representing 3D geometry in an implicit manner due to its simplicity and effectiveness. For example, Park et al. [25] represent 3D shapes by mapping the 3D coordinates to signed distance functions with MLPs. Mescheder et al. [19] introduce a new network to implicitly represent continuous surfaces as 3D occupancy fields for generating high-quality results at infinite resolution. However, these methods simply employ a global latent code to represent the total 3D shape, thus restricting to simple geometry and sacrificing robustness. To address these issues, several research studies [9, 21, 27, 29]

focus on combining local features to represent the corresponding 3D geometry instead of directly encoding an entire shape into a global latent code. PIFu proposed by Saito et al. [29] learns a per-pixel implicit representation from the pixel-aligned feature with the global context, and could preserve local details while producing high-fidelity reconstruction. Similarly, Jiang et al. [13] introduce a local implicit grid representation for arbitrary objects or scenes, and this representation greatly enhances their network’s generalization ability. In this work, we introduce a more expressive VIFu to represent intricate hair geometry and formulate the hair growth problem as an implicit function to improve the hair growth efficiency and achieve one-shot hair modeling.

**High-Resolution 3D Reconstruction.** Several recent research focus on the reconstruction of high-quality 3D texture or geometry based on the 2D cues. For example, [14, 32, 35] estimate geometric or color details using a texture map representation. [1, 37] explore the unwrapped UV space by regressing displacements to improve the reconstruction quality, and [2] deforms the high-fidelity models of humans based on a data-driven method to hallucinate plausible details. However, both approaches could not produce vivid details that match well with the truth. On the other hand, [30, 34] build two branches, the coarse one and the fine one, to fuse the global and local features. Their approach produces more detailed results while rarely imposing extra memory burden. We adopt a similar coarse-to-fine strategy to reconstruct a high-resolution hair model. However, unlike their task, where paired training data could be easy to obtain, we lack photo-realistic hair images paired with 3D hair models. Directly using rendered images would have a large domain gap with real images (e.g., due to the differences in color, illumination, texture, material, etc.). We thus utilize a 2D orientation map as the input to our coarse branch while a luminance map to our fine branch. The two maps could make up for most of the domain gap, and mutually contribute to each other.

### 3. Method

Fig. 2 shows the pipeline of NeuralHdHair. For a portrait image, we first calculate its 2D orientation map [24], and extract its luminance map [38]. Further, we automatically align them to the same reference bust model to get a bust depth map [36]. Then, these three maps are subsequently fed into our proposed IRHairNet. Specifically, it adopts a coarse-to-fine strategy and exploits a voxel-aligned implicit function (VIFu) to predict the initial 3D hair geometric features (3D orientation and 3D occupancy), which will be further enriched by the local features extracted from the hair luminance map (Sec. 3.1). Then, the enriched 3D orientation field is divided into many patches and fed into our GrowingNet (Sec. 3.2) to efficiently grow a complete hair-strand model in parallel.

### 3.1. IRHairNet

Our IRHairNet is designed to generate high-resolution 3D hair geometric features from a single image. The input to this network is comprised of a 2D orientation map, a luminance map and a fitted bust depth map, which are derived from the input portrait image. The output is a 3D orientation field, where each voxel inside contains a local growing direction, and a 3D occupancy field, where each voxel indicates whether there is any hair strand passing through (1) or not (0).

**Voxel-aligned Implicit Function.** Since the implicit function has shown impressive results in 3D shape representation, we consider defining an implicit function to represent 3D hair geometry. As discussed in Sec. 1, a pixel-wise implicit representation is not sufficient for representing complex 3D hair internal geometry and tends to produce over-smoothed results, directly ignoring the hair’s local details and spatial hierarchy (see Sec. 4). To this end, we propose an expressive voxel-aligned implicit function (VIFu) that targets each voxel instead of each pixel to represent the intricate 3D geometry of hair. Specifically, given a spatial point  $p$ , we can obtain the local orientation value  $O_p^{ori}$  and the occupancy value  $O_p^{occ}$  with two implicit functions  $f_c$  and  $f_r$  as follows:  $f_r(F(v_p), p) = O_p^{ori} \in \mathcal{R}^3$ , and  $f_c(F(v_p), p) = O_p^{occ} \in \mathcal{R}^1$ , where  $F(v_p)$  outputs the latent code of the voxel  $v$  that  $p$  locates. Since  $F(v_p)$  can vary within the entire volume, it is more reasonable to represent the interior features of complex hair than PIFu [29]. Furthermore, during testing, we can obtain the orientation and occupancy fields of higher resolution than during training by sampling more points inside each voxel when evaluating  $f$ . Importantly,  $F(\cdot)$  itself can be an implicit function, which maps the coordinate of the voxel  $v$  to the corresponding latent code.

**Implicit to Voxel Module with VIFu.** Inspired by [36], we use a variant U-Net architecture to achieve image-to-voxel task. The key idea is to build the skip connections between 2D and 3D features with a bridge module. Instead of simply increasing the number of 2D feature channels as in [36], where the depth information is not explicitly modeled, we propose an implicit to Voxel module to achieve this task, which employs the idea of our VIFu. Specifically, this module can be formulated as:

$$F(v) = \theta(I(x), Z(v)), \quad (1)$$

where for each voxel  $v$ ,  $x = \pi(v)$  is its 2D projection coordinate,  $I(x)$  is the 2D image feature locating at  $x$ , and  $Z(v)$  is the normalized depth value. This means each voxel’s latent code  $F(v)$  can be refined using its corresponding 2D feature  $I(x)$  and the normalized depth  $Z(v)$  through the MLP  $\theta$ . In practice, we first duplicate the 2D features along the depth dimension of the 3D features, then apply the MLP




Figure 2. The pipeline of NeuralHDHair. Given a single image, our IRHairNet employs a voxel-aligned implicit function to predict 3D hair geometric features based on a 2D orientation map and a high-resolution luminance map derived from the input image. Specifically, the coarse module produces a voxel-wise latent code with global context. The fine module extracts local details from the luminance map. We then use an MLP to decode the 3D hair geometric features from the voxel-wise 3D feature embeddings and local details. Afterwards, our GrowingNet divides the 3D orientation field into multiple patches and grows hair strands in individual patches in parallel. Finally, we remove the out-of-boundary hair strands using the 3D occupancy field to produce the final hair-strand model.

$\theta$  based on  $Z(v)$  and  $I(x)$ , as illustrated in Fig. 2. It has the following two advantages. First, it can refine a pixel-wise latent code to a voxel-wise latent code taking the depth into consideration, which is very important for inferring hair geometry. Second, this module can be inserted into every skip connection, thus suitable for progressive learning induced by the U-Net, e.g., the resolution of the 3D features will gradually increase (from  $6 \times 8 \times 8$  to  $96 \times 128 \times 128$ ) in the decoding process, which helps fuse multi-level features to learn the overall shape as well as the local details.

After the U-Net, we can obtain the latent code  $F(v_p)$  at any point  $p$  in the volume. Hence, the 3D hair geometric feature  $O_p$  can be decoded from  $F(v_p)$  as:

$$O_p = f(F(v_p), Z(p)), \quad (2)$$

where  $O_p$  represents an orientation or occupancy value at point  $p$  and  $f$  is an MLP representing  $f_c$  or  $f_r$  for simplicity.

**Coarse-to-fine Hair Modeling.** Although the proposed U-Net architecture combined with VIFu is able to achieve decent performance (see Sec. 4.1), plenty of hair details present in the image still cannot be captured. We thus utilize a coarse-to-fine strategy for high-quality reconstruction, as discussed in Sec. 2. For this, we design a coarse module and a fine module. The coarse module has the architecture of the proposed U-Net, which takes a 2D orientation map and a bust depth map as input, and outputs initial coarse 3D hair geometric features. The fine module is a Hourglass [22] network, extracting the local features from a high-resolution luminance map. The local features will be fused into the initial hair features to achieve high-resolution modeling ul-

timately, as follows:

$$O_p^H = f^H(\Omega(p), I^H(\pi(p)), Z(p)), \quad (3)$$

where  $\Omega(p)$  comes from the above coarse module (the second layer output of  $f$ ),  $I^H(\pi(p))$  is the 2D fine-grained local feature extracted from the luminance map, and  $f^H$  is an MLP.

Note that we convert the input image from the RGB space to the LAB space for obtaining a luminance map (L channel). The luminance map, highly disentangled with chrominance in the LAB space, makes it possible to describe the sophisticated hair structure. In addition, compared with the orientation map, which only contains the growth direction of the hair, the luminance map can capture more local details, such as illumination and depth information, making significant improvement in realism and high-fidelity.

**Loss Function.** The design of the loss function has a significant effect on the robustness of the trained model. Considering that forcing the network to fit occluded hair (especially the hair behind the head) might weaken the learning ability of the network, we assign small weights to invisible points. Similar to the previous work [36], we use the Binary Cross Entropy (BCE) loss for the occupancy field and the L1 loss for the orientation field. Thus, the final loss function is given by:

$$W_p = \begin{cases} 1 & , Z(p) - D(p) \geq \tau \\ 10 & , Z(p) - D(p) < \tau \end{cases} \quad (4)$$

$$\mathcal{L}_{Ori} = \sum W_p \cdot \|O_p^{ori} - \hat{O}_p^{ori}\|^1, \quad (5)$$

$$\begin{aligned} \mathcal{L}_{Occ} = & \sum W_p \cdot (\lambda \cdot \hat{O}_p^{occ} \log(O_p^{occ}) \\ & + (1 - \lambda)(1 - \hat{O}_p^{occ}) \log(1 - (O_p^{occ}))), \end{aligned} \quad (6)$$

where  $D(p)$  and  $Z(p)$  represent the corresponding depth value at the projected location  $\pi(p)$  on the depth image  $D$  and its actual z-coordinate, respectively. We regard that the points within  $\tau$  ( $\tau = 5$ ) voxels under the surface of the hair are visible since the hair is not flaky. Therefore, our model can perform reasonable reconstruction for invisible hair.

**Data.** Similar to the previous research [6, 31, 36], we collected 653 3D hair-strand models and aligned them to the same bust model within a boundary box. In addition, we also augment the data by horizontally flipping, scaling, and rotating. Then we voxelize the boundary box to prepare the training data, including 3D strand points, 3D orientation maps, and 2D orientation maps together with the luminance map. Note that we generate a luminance map from the rendered image of a 3D hair strand model at the training phase but produce it directly from an input portrait image at the testing phase. To train our IRHairNet, we randomly sample in the boundary box and sample points farther from the strands with Gaussian-decaying probabilities [25]. Finally, we perform trilinear interpolation to fill the discontinuous holes in the 3D orientation map for more robust training.

### 3.2. GrowingNet

Our GrowingNet is designed to efficiently generate a complete hair-strand model from the 3D orientation field and 3D occupancy field estimated by our IRHairNet, where the 3D occupancy field is used to limit the growth region of hair.

**Problem Formulation.** We regard the each hair strand as a continuous function in a high-dimensional space similar to the spatial surface. The derivative of the spatial point on the curve represents the local growing direction. If we set a fixed step size  $s$ , we can start from an arbitrary 3D point  $x_n$  and iteratively solve the next point  $x_{n+1}$  by computing the corresponding derivative to grow a complete strand. Specifically, our hair growth algorithm can be formulated as:  $x_{n+1} = x_n + s \cdot f_g(z, x_n)$ , where  $z = E(O^{ori})$  represents the latent code encoded from the whole 3D orientation field using the encoder  $E$ . We attempt to apply an implicit function  $f_g$  to regress the derivative of arbitrary point  $x$  in the space. Then, we can produce a complete hair model by iteratively solving the next point in parallel. In fact, since  $s$  is a constant, we force the network to learn the above formula and directly output the coordinates of the next point  $x_{n+1} = G(z, x_n)$ , where  $G$  is an MLP. Note that the tradition hair growing algorithm entails  $s$  is less than the voxel’s width while in our formulation, we can change  $s$  to satisfy different resolution requirements by training different MLPs.

**Local Implicit Hair Growing.** Since learning a global

implicit function is challenging, especially for complex structures, we consider learning a local implicit function (a decoder) to grow hair strands. Specifically, we divide the 3D orientation field into lots of independent patches, which share the same decoder to grow hair strands. Based on the above considerations, we first randomly sample a 3D point in the hair volume and query its corresponding patch in the world coordinate system. Then we decode the next point according to the following formula:

$$x_{n+1} = G(z_i, \frac{2}{d}(x_n - x_i)), \quad (7)$$

where  $z_i$  and  $x_i$  are the latent code and the center of the corresponding local patch  $i$  respectively, and  $d$  is the patch scale. It means that we unify all the local patches by converting the world coordinates to the local coordinate systems and normalizing the range to  $[-1, 1]$  before decoding.

**Hair Growth with Overlapping Latent Code.** Although our local implicit hair growing method has an excellent performance in capturing local details, we found that it would produce visible artifacts in the boundary of patches (see Sec. 4.1). To address this issue, we employ the overlapping latent patch scheme for any two adjacent patches overlapped by half of the patch scale. Specifically, given an arbitrary point  $x_n$  of a hair strand, we compute the next point  $x_{n+1}$  by applying trilinear interpolation to the implicit function values of all patches overlapped at the position of  $x_n$ :

$$\begin{aligned} x_{n+1} = & T(z_{\mathcal{N}}, G, x_n) \\ = & \sum_{j \in \mathcal{N}} W_j \cdot G(z_j, \frac{2}{d}(x_n - x_j)), \end{aligned} \quad (8)$$

where  $\mathcal{N}$  is the set of all overlapping patches at point  $x_n$ , and  $W_j$  is the trilinear interpolation weight corresponding to the local patch  $j$ . Finally, we generate a complete and smooth strand model by growing hair independently in each patch in parallel (e.g., growing 10,000 strands at the same time).

**Loss Function.** We train our GrowingNet by minimizing the following loss function:

$$\begin{aligned} \mathcal{L} = & \|T(z_{\mathcal{N}}, G, x_n) - \hat{x}_{n+1}\|^1 \\ & + \|T(z_{\mathcal{N}}, G_{Inv}, x_n) - \hat{x}_{n-1}\|^1, \end{aligned} \quad (9)$$

where  $G_{Inv}$  represents the decoder with the same architecture as  $G$ , but its output is the previous point  $x_{n-1}$ . This is because we bi-directionally grow a hair-strand from the sampled point in the local region to avoid the loss of entangled hair strands. Thus,  $\hat{x}_{n-1}$  and  $\hat{x}_{n+1}$  are the actual positions of the previous and next points of  $x_n$ , respectively.

**Data.** Similarly, we use the strand model as mentioned above to train our GrowingNet. The difference is that we

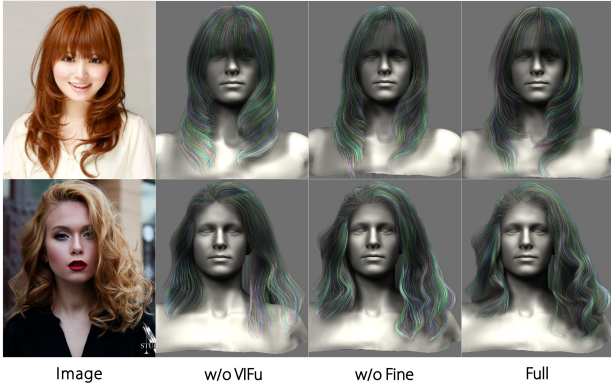


Figure 3. Qualitative evaluation of each key component of our IRHairNet. VIFu helps us represent complex hair structures, while the fine module contributes to supplement local details.

control the distance between two adjacent 3D points on one strand to be approximately equal (to achieve the same step size) using B-spline interpolation for more robust training.




Figure 4. Qualitative evaluation of GrowingNet. Our GrowingNet can efficiently achieve visually similar results to traditional methods. The employment of overlapping latent patch scheme helps produce smoother results across patches.

## 4. Experiments

In this section, we evaluate the effectiveness and necessity of each algorithmic component through ablation studies (Sec. 4.1), followed by comparing our approach with state-of-the-art methods [6, 28, 30, 36, 40] (Sec. 4.2). Besides, implementation details and more experimental results can be found in our supplementary materials.

### 4.1. Ablation Studies

**Evaluation of IRHairNet.** To evaluate the effectiveness and necessity of each key component of our IRHairNet in terms of geometric fidelity and realism, we compare our

Method	Traditional	Ours	w/o Overlap
Time(s)	10.57	1.21	0.25

Table 1. Our GrowingNet is more efficient than a traditional method. Sacrificing some precision (w/o Overlap) can save a lot of time.

Local size	4	8	16	32
Time(s)	5.75	1.21	1.13	1.04

Table 2. Comparison of time consumption under different local sizes.

full method with two simplified settings. One simplification removes the fine module to evaluate the effect of refinement (w/o Fine) while the other removes the implicit toVoxel module to evaluate the significance of VIFu (w/o VIFu).

The visual comparison results are shown in Fig. 3. When removing the overall implicit toVoxel module, our system severely suffers from the problem of over-smoothing: the results without VIFu only contain either a plausible shape with approximate growth direction or even an incorrect structure (second row in Fig. 3). This is because the representation ability of the network without VIFu would significantly deteriorate, making it difficult to describe intricate hair geometry. On the other hand, without the fine module, some local details are unavoidably lost, though the generated results still maintain most of the correct growth directions, especially for complex hairstyles. Employing the fine module can dramatically enhance the local details (e.g., strands’ hierarchical structure) by extracting fine-grained features from the high-resolution hair luminance map while increasing the resolution of 3D hair geometric features also improve the quality of the 3D hair model (3-th and 4-th columns of Fig. 3). The above experiments prove that the proposed VIFu has a strong representation ability for sophisticated hair geometry and such a coarse-to-fine approach can capture more local details while increasing the 3D hair model resolution.



Figure 5. Performance comparison of GrowingNet under different local sizes with different resolutions. The first row represents the low-resolution results ( $96 \times 128 \times 128$ ) and the second row represents the high-resolution one ( $384 \times 512 \times 512$ )




Figure 6. Compared to HairNet [40], our method leads to significantly better results in terms of shape and structure

**Evaluation of GrowingNet.** We also qualitatively and quantitatively evaluate the effectiveness of GrowingNet in terms of fidelity and efficiency. We first conduct three sets of experiments on synthetic data: i) traditional hair-growth algorithm [31], ii) GrowingNet without the overlapping latent patch scheme, iii) our full model. Fig. 4 and Tab. 1 show our GrowingNet has obvious advantages over traditional hair growth algorithm in time consumption, while maintaining the same growth performance in terms of visual quality. In addition, by comparing the third and fourth columns of Fig. 4, we can see that without the overlapping latent patch scheme, the hair strands at the patch boundary may be discontinuous and this issue is more serious where the hair strands’ growth directions change sharply. However, it is worth noting that this scheme greatly improves the efficiency at the cost of slightly reduced precision. Still, the improved efficiency is significant for the convenient and efficient application to human digitization.

We also evaluate the reconstruction performance of GrowingNet on different local sizes in different resolutions ( $96 \times 128 \times 128$  and  $384 \times 512 \times 512$ ). As shown in Fig. 5, as the input resolution increases, the quality of the result is higher. On the other hand, at the same resolution, the performance of modeling rises simultaneously with a smaller patch size, producing more accurate details, since decoding a small patch is much more simplified than a large one. It seems that we can reconstruct a high-fidelity model by simply reducing the local patch size. However, as shown in Tab. 2, when the local patch size is reduced to less than 8, the reconstruction time increases significantly with only a tiny improvement on the modeling quality, due to its time-consuming query process similar to the traditional hair growth algorithm. Therefore, we choose a local patch size of  $8 \times 8 \times 8$  to balance performance and efficiency.

## 4.2. Comparisons

To evaluate the effectiveness of our NeuralHdHair, we conduct several comparisons with the state-of-the-art methods [6, 28, 30, 36, 40], where Autohair [6] performs hair syn-

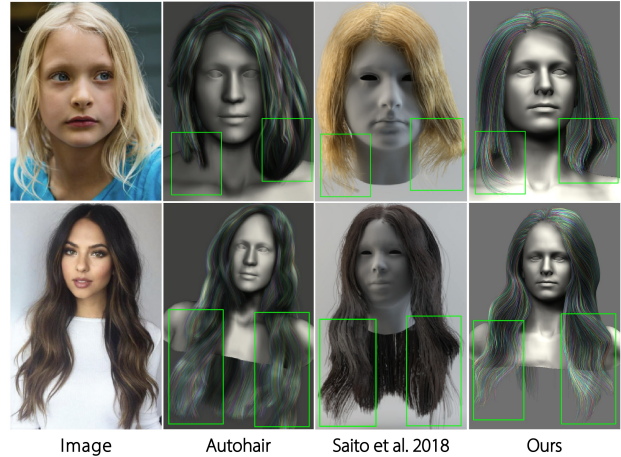


Figure 7. Comparison with Autohair [6] and Saito et al [28]. Their results are not consistent with the input image in shape or structure.

thesis based on a data-driven approach, and HairNet [40] ignores the hair growth procedure to achieve end-to-end hair modeling. In contrast, [28, 36] perform a two-step strategy that first estimates a 3D orientation field and then synthesizes hair strands from it. PIFuHD [30] is a state-of-the-art monocular high-resolution 3D modeling method based on the coarse-to-fine strategy, which can be reformulated for 3D hair modeling.

As shown in Fig. 6, it is evident that the result by HairNet looks plausible but the local details and even the overall shape are inconsistent with the hair in the input image. This is because they perform hair synthesis in a simple and crude manner by directly regressing the unordered hair strands from a single image. We also compared the reconstruction results with [6, 28]. As shown in Fig. 7, although Autohair can synthesize realistic results, its structure cannot match the input image well since the database contains limited hairstyles. On the other hand, the results of [28] lack local details, and the shape cannot be consistent with the input image. In contrast, our results better maintains the global structure and local details of the hair, while ensuring the consistency of the hair shape.

Similar to our task, PIFuHD [30] and Dynamic Hair [36] are committed to estimating high-fidelity 3D hair geometric features to produce realistic hair-strands models. Fig. 8 shows two representative comparison results. It can be seen that a pixel-wise implicit function adopted in PIFuHD cannot sufficiently describe the intricate hair, resulting in over-smooth results, without local details or even bad global structure. Although they attempt to supplement the local details to the coarse global features in their fine module, their generated results are still unsatisfactory due to the incorrect global structure. While [36] can produce more reasonable results with few details and the overall hair growth trend in their results can match the input image well, many local

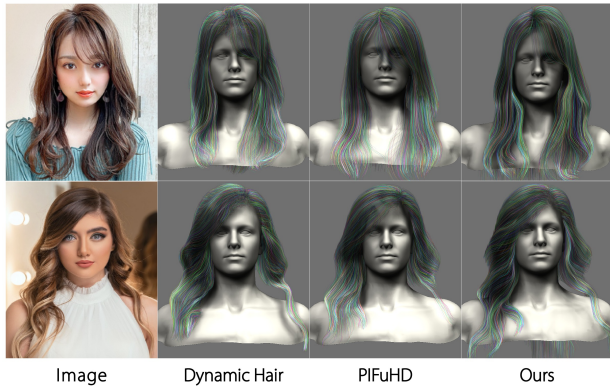


Figure 8. Comparison with PIFuHD [30] and Dynamic Hair [36]. Compared to our method, the two existing methods cannot sufficiently describe 3D hair geometry and generally produce over-smooth results, especially for complicated hair structure.

structural details (e.g., hierarchy) cannot be captured, especially for complex hairstyles. In contrast, with the coarse-to-fine strategy and the expressive implicit representation (VIFu), our method can adapt to diverse hairstyles, even extreme complex structures, and fully leverage the global features and local details to generate high-fidelity, high-resolution 3D hair models with more details.

## 5. Conclusion and Future Work

In this paper, we introduced an automatic learning-based monocular hair modeling system NeuralHDHair to produce a high-fidelity hair-strand model as shown in Fig. 9 (Please refer to the supplemental video<sup>1</sup> for more results). It consists of two carefully designed deep learning networks based on implicit functions: IRHairNet to provide high-resolution 3D hair geometric features. based on the proposed VIFu and high-resolution luminance map, and GrowingNet cleverly applies the local parallel growth strategy to provide an efficient hair synthesis procedure, which also supports us in producing a complete hair-strand model that only needs one forward pass of the networks. We also evaluate the effectiveness and necessity of each key component of our system through ablation studies and demonstrate that our approach achieves state-of-the-art performance among the monocular hair modeling methods.

As far as we know, all existing hair modeling methods are limited to the quality of the 2D orientation map. Although we attempt to use the luminance map to supplement local details, we found that it must be complementary with the 2D orientation map. Utilizing only a luminance map as input suffers from the problem of overfitting due to the insufficient diversity of our 3D hair models. Thus, enriching the 3D hair dataset or adopting more realistic rendering




Figure 9. Fully automatic hair modeling results for various hairstyles. Our NeuralHDHair can produce realistic, high-resolution 3D hair models with local details.

methods is essential to reconstructing a 3D hair model directly from a single image (or a luminance map) instead of an orientation map. In addition, for diversified portrait input, we employ landmark to automatically align it with a fixed bust model, which may cause the hair not to fit well with the head due to the different identities of the head. In the future, to significantly improve the visual quality of hair reconstruction, we may attempt to estimate the head pose and reconstruct the corresponding 3D face while reconstructing a 3D hair model.

**Acknowledgement.** This work was supported in part by National Key Research & Development Program of China (Grant No. 2018YFE0100900) and National Natural Science Foundation of China (Grant No. 62172363).

<sup>1</sup>[https://www.youtube.com/watch?v=g0wGRp\\_zKcI](https://www.youtube.com/watch?v=g0wGRp_zKcI)



## References

- [1] Thiemo Alldieck, Gerard Pons-Moll, Christian Theobalt, and Marcus Magnor. Tex2shape: Detailed full human body geometry from a single image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2293–2303, 2019. 2
- [2] Dragomir Anguelov, Praveen Srinivasan, Daphne Koller, Sebastian Thrun, Jim Rodgers, and James Davis. Scape: shape completion and animation of people. In *ACM SIGGRAPH 2005 Papers*, pages 408–416. 2005. 2
- [3] Matan Atzmon, Niv Haim, Lior Yariv, Ofer Israelov, Haggai Maron, and Yaron Lipman. Controlling neural level sets. *arXiv preprint arXiv:1905.11911*, 2019. 2
- [4] Xudong Cao, Yichen Wei, Fang Wen, and Jian Sun. Face alignment by explicit shape regression. *International journal of computer vision*, 107(2):177–190, 2014. 1
- [5] Menglei Chai, Linjie Luo, Kalyan Sunkavalli, Nathan Carr, Sunil Hadap, and Kun Zhou. High-quality hair modeling from a single portrait photo. *ACM Transactions on Graphics (TOG)*, 34(6):1–10, 2015. 2
- [6] Menglei Chai, Tianjia Shao, Hongzhi Wu, Yanlin Weng, and Kun Zhou. Autohair: Fully automatic hair modeling from a single image. *ACM Transactions on Graphics*, 35(4), 2016. 1, 1, 2, 3.1, 4, 4.2, 7
- [7] Menglei Chai, Lvdi Wang, Yanlin Weng, Xiaogang Jin, and Kun Zhou. Dynamic hair manipulation in images and videos. *ACM Transactions on Graphics (TOG)*, 32(4):1–8, 2013. 2
- [8] Menglei Chai, Lvdi Wang, Yanlin Weng, Yizhou Yu, Baining Guo, and Kun Zhou. Single-view hair modeling for portrait manipulation. *ACM Transactions on Graphics (TOG)*, 31(4):1–8, 2012. 2
- [9] Maximilian Denninger and Rudolph Triebel. 3d scene reconstruction from a single viewport. In *European Conference on Computer Vision*, pages 51–67. Springer, 2020. 2
- [10] Sunil Hadap, Marie-Paule Cani, Ming Lin, Tae-Yong Kim, Florence Bertails, Steve Marschner, Kelly Ward, and Zoran Kačić-Alesić. Strands and hair: modeling, animation, and rendering. In *ACM SIGGRAPH 2007 courses*, pages 1–150. 2007. 1
- [11] Liwen Hu, Chongyang Ma, Linjie Luo, and Hao Li. Single-view hair modeling using a hairstyle database. *ACM Transactions on Graphics (ToG)*, 34(4):1–9, 2015. 1, 1, 2
- [12] Liwen Hu, Shunsuke Saito, Lingyu Wei, Koki Nagano, Jae-woo Seo, Jens Fursund, Iman Sadeghi, Carrie Sun, Yen-Chun Chen, and Hao Li. Avatar digitization from a single image for real-time rendering. *ACM Transactions on Graphics (ToG)*, 36(6):1–14, 2017. 1
- [13] Chiyu Jiang, Avneesh Sud, Ameesh Makadia, Jingwei Huang, Matthias Nießner, Thomas Funkhouser, et al. Local implicit grid representations for 3d scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6001–6010, 2020. 1, 2
- [14] Verica Lazova, Eldar Insafutdinov, and Gerard Pons-Moll. 360-degree textures of people in clothing from a single image. In *2019 International Conference on 3D Vision (3DV)*, pages 643–653. IEEE, 2019. 2
- [15] Hao Li, Laura Trutoiu, Kyle Olszewski, Lingyu Wei, Tristan Trutna, Pei-Lun Hsieh, Aaron Nicholls, and Chongyang Ma. Facial performance sensing head-mounted display. *ACM Transactions on Graphics (ToG)*, 34(4):1–9, 2015. 1
- [16] 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
- [17] Linjie Luo, Hao Li, and Szymon Rusinkiewicz. Structure-aware hair capture. *ACM Transactions on Graphics (TOG)*, 32(4):1–12, 2013. 1
- [18] Linjie Luo, Cha Zhang, Zhengyou Zhang, and Szymon Rusinkiewicz. Wide-baseline hair capture using strand-based refinement. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 265–272, 2013. 2
- [19] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4460–4470, 2019. 2
- [20] Mateusz Michalkiewicz, Jhony K Pontes, Dominic Jack, Mahsa Baktashmotlagh, and Anders Eriksson. Implicit surface representations as layers in neural networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4743–4752, 2019. 2
- [21] Zak Murez, Tarrence van As, James Bartolozzi, Ayan Sinha, Vijay Badrinarayanan, and Andrew Rabinovich. Atlas: End-to-end 3d scene reconstruction from posed images. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VII 16*, pages 414–431. Springer, 2020. 2
- [22] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *European conference on computer vision*, pages 483–499. Springer, 2016. 3.1
- [23] Michael Oechsle, Songyou Peng, and Andreas Geiger. Unisurf: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. *arXiv preprint arXiv:2104.10078*, 2021. 1
- [24] 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. 1, 2, 3
- [25] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 165–174, 2019. 1, 2, 3.1
- [26] Songyou Peng, Michael Niemeyer, Lars Mescheder, Marc Pollefeys, and Andreas Geiger. Convolutional occupancy networks. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*, pages 523–540. Springer, 2020. 2
- [27] Sida Peng, Yuanqing Zhang, Yinghao Xu, Qianqian Wang, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Neural body:

- Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9054–9063, 2021. 2
- [28] Shunsuke Saito, Liwen Hu, Chongyang Ma, Hikaru Ibayashi, Linjie Luo, and Hao Li. 3d hair synthesis using volumetric variational autoencoders. *ACM Transactions on Graphics (TOG)*, 37(6):1–12, 2018. 1, 2, 4, 4.2, 7
- [29] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2304–2314, 2019. 1, 2, 3.1
- [30] Shunsuke Saito, Tomas Simon, Jason Saragih, and Hanbyul Joo. Pifuhd: Multi-level pixel-aligned implicit function for high-resolution 3d human digitization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 84–93, 2020. 1, 2, 4, 4.2, 8
- [31] 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. 1, 3.1, 4.1
- [32] Anh Tuan Tran, Tal Hassner, Iacopo Masi, Eran Paz, Yuval Nirkin, and Gérard G Medioni. Extreme 3d face reconstruction: Seeing through occlusions. In *CVPR*, pages 3935–3944, 2018. 2
- [33] Chufeng Xiao, Deng Yu, Xiaoguang Han, Youyi Zheng, and Hongbo Fu. Sketchhairsalon: Deep sketch-based hair image synthesis. *arXiv preprint arXiv:2109.07874*, 2021. 1
- [34] Qiangeng Xu, Weiyue Wang, Duygu Ceylan, Radomir Mech, and Ulrich Neumann. Disn: Deep implicit surface network for high-quality single-view 3d reconstruction. *arXiv preprint arXiv:1905.10711*, 2019. 2
- [35] Shugo Yamaguchi, Shunsuke Saito, Koki Nagano, Yajie Zhao, Weikai Chen, Kyle Olszewski, Shigeo Morishima, and Hao Li. High-fidelity facial reflectance and geometry inference from an unconstrained image. *ACM Transactions on Graphics (TOG)*, 37(4):1–14, 2018. 2
- [36] Lingchen Yang, Zefeng Shi, Youyi Zheng, and Kun Zhou. Dynamic hair modeling from monocular videos using deep neural networks. *ACM Transactions on Graphics (TOG)*, 38(6):1–12, 2019. 1, 1, 2, 3, 3.1, 3.1, 4, 4.2, 8
- [37] Xiaoxing Zeng, Xiaojiang Peng, and Yu Qiao. Df2net: A dense-fine-finer network for detailed 3d face reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2315–2324, 2019. 2
- [38] Bo Zhang, Mingming He, Jing Liao, Pedro V Sander, Lu Yuan, Amine Bermak, and Dong Chen. Deep exemplar-based video colorization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8052–8061, 2019. 3
- [39] Meng Zhang, Pan Wu, Hongzhi Wu, Yanlin Weng, Youyi Zheng, and Kun Zhou. Modeling hair from an rgb-d camera. *ACM Transactions on Graphics (TOG)*, 37(6):1–10, 2018. 1, 2
- [40] 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. 1, 1, 2, 4, 6, 4.2