Contents lists available at ScienceDirect

THE WAY CUMPERS



iournal homepage: www.elsevier.com/locate/visinf

Hair-GAN: Recovering 3D hair structure from a single image using generative adversarial networks

Meng Zhang, Youyi Zheng*

State Key Lab of CAD&CG, Zhejiang University, Hang Zhou 310058, China

ARTICLE INFO

ABSTRACT

Article history: Received 28 December 2018 Received in revised form 14 May 2019 Accepted 5 June 2019 Available online 26 June 2019

Keywords: Single-view hair modeling 3D volumetric structure Deep learning Generative adversarial networks We introduce *Hair-GAN*, an architecture of generative adversarial networks, to recover the 3D hair structure from a single image. The goal of our networks is to build a parametric transformation from 2D hair maps to 3D hair structure. The 3D hair structure is represented as a 3D volumetric field which encodes both the occupancy and the orientation information of the hair strands. Given a single hair image, we first align it with a bust model and extract a set of 2D maps encoding the hair orientation information in 2D, along with the bust depth map to feed into our *Hair-GAN*. With our generator network, we compute the 3D volumetric field as the structure guidance for the final hair synthesis. The modeling results not only resemble the hair in the input image but also possesses many vivid details in other views. The efficacy of our method is demonstrated by using a variety of hairstyles and comparing with the prior art.

© 2019 Zhejiang University and Zhejiang University Press. Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (http://creativecommons.org/licenses/by-nc-nd/4.0/).

1. Introduction

Techniques of 3D avatar modeling are becoming increasingly popular in today's emerging VR and AR applications. However, hair modeling, one of the most crucial tasks, still remains challenging due to the complexity and variety of hairstyles in the real world. Methods of single-view hair modeling (Chai et al., 2016; Hu et al., 2015; Chai et al., 2015) are considered as a much more user-friendly in comparison with those of multi-view modeling methods (Echevarria et al., 2014; Herrera et al., 2012; Hu et al., 2014; Jakob et al., 2009; Luo et al., 2013; Paris et al., 2008), which usually require specialized equipments in controlled studio environment and long processing cycles.

Due to the lack of information at views distant from the input one, single-view hair modeling techniques usually rely on a large database containing hundreds or thousands of synthetic hairstyles used as prior knowledge of hair shape, distribution, or complex structure. These data-driven methods (Chai et al., 2016; Hu et al., 2015, 2017) come with some problems. Firstly, the requirement of large storage for hair database restricts their application on resource-constrained platforms such as mobile devices. Second, the quality of modeling results relies on the retrieved hair exemplars which are searched from the quantity-limited database. Although post-refinement is introduced to improve the detail accuracy, the structure of the final result is

E-mail address: youyizheng@zju.edu.cn (Y. Zheng).

Peer review under responsibility of Zhejiang University and Zhejiang University Press.

still confined by the initially retrieved exemplar. In addition, the greedy searching process is slow and difficult to balance the selection criteria between the local detail similarity and the global shape similarity.

Recently, deep learning has been intensively exploited in many fields of researches. One of the most attractive characteristics of deep learning methods is their effectiveness in converting big datum into high-dimensional feature representations. Those learned features are sufficiently gualified to describe a new data and to set up a space mapping from the input to the target output, totally independently from those training datum. The method of Chai et al. (2016) and the concurrent works of Zhou et al. (2018) and Saito et al. (2018) introduce deep learning into single-view hair modeling. In Chai et al. (2016)'s method, they use CNN networks to detect hair regions and direction predictors from images as preprocessing for the subsequent hair modeling. Zhou et al. (2018) use an auto-encoder to directly learn the correspondence from 2D orientation map to hair stand assembly parameterized to a 2D scalp map of low resolution as in Wang et al. (2009). Saito et al. (2018) put an effort to achieve end-to-end 3D hair inference by leveraging the latent space of auto-encoders to build a bridge between the 2D image and the 3D hair structure.

In this paper, we introduce *Hair-GAN*, an architecture of generative adversarial networks to recover the 3D hair structure from a single input image. As studied in the vision communities, generative adversarial networks are capable of capturing the model distribution better than those of CNN and auto-encoder models because their adversarial training nature tends to learn more generic distance measures by themselves rather than handcoded (Goodfellow et al., 2014). In order to build the data for the

https://doi.org/10.1016/j.visinf.2019.06.001

2468-502X/© 2019 Zhejiang University and Zhejiang University Press. Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (http://creativecommons.org/licenses/by-nc-nd/4.0/).

^{*} Corresponding author.



Fig. 1. Given single-view inputs of a variety of hairstyles (the first column), our generator network recovers 3D hair structures represented as the rough shape and orientation field (the second column) to guide hair synthesis. Complete strand-level 3D hair models (the last three columns) are finally generated that resemble the hairs in input images.

network trainings, we use 3D artificially synthetic hair database. Like the deep learning methods of Zhou et al. (2018) and Saito et al. (2018), the large amount of database is put aside after the training of the network is completed.

A particular challenge of single-view hair modeling is that the input image only provides 2D information, lacking the cues along the depth direction in the 3D space. Therefore, a direct 2D-to -3D training may lead to unconstrained results. We thus involve the bust depth map into the input tensor of our network as a condition to our GAN, which provides depth prior for hairs growing around the human body. Meanwhile, we convert a succession of 2D features to a single channel of 3D features by a dimensional expansion layer to aggregate 3D knowledge from 2D convolutional neural networks. Inspired by the adversarial structure of GANs (Goodfellow et al., 2014), we minimize the objective loss of discriminator to enlarge the difference between the real and the fake, and competitively optimize the generator to reduce the distance between the output and the ground truth, not only voxel-wisely but also in the latent space of the discriminator. Our method can generate a high-quality 3D strand-level hair model with a single image as input, guided by the volume of occupancy and orientation field generated by our generator network. 3D hair structure recovered by our Hair-GAN takes care of high perceptual-quality details along the depth direction, rather than a smooth rough hair shape as in Saito et al. (2018). Fig. 1 shows the efficacy of our method for both shot and long, straight and curly hairstyles.

In summary, our contributions are:

- We introduce the architecture of GAN for single-view hair modeling. Our GAN transform 2D orientation maps into 3D volumetric field which encodes both the occupancy and orientation information of hair strands;
- We propose a dimension expansion layer into the design of our generator network which converts a succession of 2D features to a single channel of 3D features;
- We optimize the generator parameters by considering both the output and the latent features of the discriminator.

2. Related work

Hair is one of the vital components of digital characters in online games, virtual worlds, and virtual reality applications. The techniques of modeling high-quality 3D hair are extensively studied in computer graphics, which usually require professional skills and days of laborious manual work. Please refer to the survey (Ward et al., 2007) for a detailed discussion of seminal hair modeling methods.

Image-based hair modeling is a promising way to create compelling hair structures from captured hair images. According to the number of images required, methods of image-based hair modeling can be roughly separated into multi-view hair modeling and single-view hair modeling. Multi-view hair modeling methods (Echevarria et al., 2014; Herrera et al., 2012; Hu et al., 2014; Jakob et al., 2009; Luo et al., 2013; Paris et al., 2008) create high-quality 3D hair modeling from a number of views, which often require complex hard-ware setup, well-controlled environment, and long processing cycles. They are not consumerfriendly since these multi-view capture systems and professional skills are not easily accessible to average users. While singleview hair modeling methods are becoming increasingly popular and important as single-view, un-calibrated images are widely available on the Internet. Chai et al. (2012, 2013) first introduce the technique of single-view hair modeling by utilizing different kinds of prior knowledge, including layer boundary and occlusion, and shading cues (Chai et al., 2015). A major problem of their methods is the lack of control over the geometry at views distant from the input image.

Data-driven hair modeling methods provide a conceptually persuasive prior knowledge of the entire hairstyle, with 3D synthetic hairstyle database. Hu et al. (2015) assemble different hairstyles searched from the database by fitting with a few user strokes to reconstruct a complete hair shape. Chai et al. (2016) bring forward the model remixing step to precomputation stage. About 5-40 candidates are found from their enlarged database and afterward, they perform deformation on these candidates to achieve a model result with detail similarity. In order to enrich their 3D hair database to 40 K models, they cluster hair strands and recombine those cluster models. Zhang et al. (2017) only use the candidate searched by contour fitting for their four-view based hair modeling to build a smooth rough hair shape and the style details are introduced by texture melding and helixfitting. In Zhang et al. (2018), they introduce a local patch-based searching strategy to find candidates with local style patterns enough to guide hair synthesis, instead of to find those with the same style in global. All these data-driven methods need storage for hundreds or thousands of hairstyle database.

The recent success of deep learning also brings significant improvement in the fields of hair modeling. Chai et al. (2016) present a fully automated hair modeling method by replacing user interactions with deep convolutional neural networks to hair segmentation and hair growth direction estimation. Hu et al. (2017) introduce deep learning based hair attribute classifier to improve the candidate retrieval performance of the data-driven method. In order to get an end-to-end learning from 2D image knowledge to 3D hair representation, Zhou et al. (2018) use encoder-decoder architecture to generate hair strands represented as sequences of 3D points for 2D orientation fields as input. But their hair representation is parameterized to a lowresolution grid on the scalp, which leads the modeling result of a low quality. In a concurrent work, Saito et al. (2018) demonstrate that 3D occupancy field and the corresponding flow field with high resolution are easily handled by neural networks and compatible with traditional strand-based representation for highfidelity modeling and rendering. However, in their method, the occupancy field and flow field are decoded separately from the same volumetric latent space. They fine-tune the pre-trained network of ResNet-50 to encode input image to the hair coefficients which are aligned to the volumetric latent space by their trained embedding networks. Since there is compression during the processing of encoder and the information loss in latent coefficient alignment, their results lack corresponding details between the input image and the output hair structure. In comparison, with 2D information maps as input, our method is more direct to train *Hair-GAN* to predict the 3D volumetric field encoding both the occupancy and orientation information, with consideration of detail correspondence between the input image and the modeling result.

Generative adversarial networks (GANs) is introduced by Goodfellow et al. (2014) as a framework to construct a generative model that can mimic the target distribution. The goal of GANs is to train the generator model by iteratively training the discriminator and generator in turn. Conditional GANs (Mirza and Osindero, 2014) are a type of GANs using conditional information for the discriminator and generator, regarded as a promising tool on image domains, e.g., the conditional image synthesis (Odena et al., 2017), the generation of the images from text (Reed et al., 2016), and image to image translation (Zhu et al., 2017). We adopt GANs to recover 3D hair structure from 2D image information, taking advantage of the power of GANs to re-create the distributions of complex data sets. We make use of the latent space of discriminator to enforce the similarity of the ground truth and target output in distribution. Our Hair-GAN aims to learn a parametric transformation from 2D information maps to 3D volumetric occupancy and orientation field with no intermediate latent space.

3. Overview

We first clarify our unified model space where all synthetic hair database and the identical bust model are aligned. Based on the unified model space, we generate the training data of the ground-truth 3D volumetric field coupling with the corresponding 2D orientation and confidence maps (Section 4). Next, we introduce the architecture and loss functions of our *Hair-GAN* (Section 5). Similar to the original GANs, our networks are also composed by a discriminator and a generator. Given a real hair image as input, by using our trained *Hair-Generator*, we can recover the 3D hair structure based on the 2D orientation and confidence maps, along with the bust depth map (all of which are extracted from the image), and finally synthesis a high-quality 3D hair model (Section 6).

4. Data preparation

Inspired by Zhang et al. (2018), we regard a hairstyle as a fusion of local style patterns distributed around human body, in contrast to Chai et al. (2016), Hu et al. (2015) and Zhou et al. (2018), where different hairstyles are treated as different

combination of styled hair strands. Similar to those previous researches, we collect an original hair dataset with about 300 3D artificial hair models provided by Chai et al. (2016), which have already been aligned to an identical bust model. We define a unified model space (Section 4.1) to prepare our training data (Section 4.2) including both the ground-truth 3D volumetric field \mathcal{Y} along with the 2D hair information maps \mathcal{X} .

4.1. Unified model space

We define a bounding box as the boundary of our model space, where we generate the ground-truth 3D hair orientation volume and capture 2D hair orientation and confidence maps. Fig. 2 illustrates our defined bounding box and 2D capture as an example.

Bounding box. The model space is bounded by a bounding box defined in consideration of the bust model and all database hair models except some extreme long hairs (selected manually). Then a 3D volume with the resolution of $128 \times 128 \times 96$ is subdivided inside the bounding box ($H \times H \times D$).

2D capture. In order to get 2D information maps χ under the defined model space, we put a camera straight forward to the bust model. The center of the image plane coincides with the center of the bounding box. The 2D image is captured by orthogonal projection with a scale of 1024/H. Therefore, the size of the captured image is 1024×1024 .

4.2. Training data

Following Hu et al. (2015), we double the number of the database by simply flipping each model and remove the constrained hairstyles such as braids and bounds. There are 303 hairstyles in our database varying from short to long, straight to curly. We randomly rotate the hair around the center of the bounding box. The rotation ranges from -15° to 15° for X-axle, -30° to 30° for Y-axle, and -20° to 20° for Z-axle. Since all these database models are made up of polygon-strips, same as to Hu et al. (2015), we convert the polygon-strips to dense 3D orientation volume regarded as the ground truth y and grow strands afterwards. Then we render the hair strands to the 2D image at the camera view pose defined in Section 4.1. However, in order to remove the difference of real and synthetic hair image, we compute 2D orientation and confidence maps for the captured image using the iterative method of Chai et al. (2012). Considering the diverse quality of real images, the iteration for the database randomly ranges from 3 to 5. Usually, there is directional ambiguity in the orientation map and Chai et al. (2013) have confirmed that directional ambiguity should be removed to ensure the correct direction of hair growth. We can project model strand direction to the image plane to update the orientation map to avoid ambiguity. Then we diffuse the orientation with high confidence to obtain the final pixel-dense orientation map and encode direction vectors to color space. In addition, as mentioned earlier, the bust model should also be taken into account as a condition to our networks since hairs are grown from the scalp and distributed around the body. We compute the bust depth map by ray tracing pixel by pixel to get the distance from the bust to the camera, and divide the distance by D to range the value within [0, 1]. Finally our network input \mathcal{X} is generated made up of 2D orientation map, confidence map, and bust depth map.

All 2D maps are valued within [0, 1] and both 3D and 2D orientation vectors are encoded in color space. For each database model, we compute 12 pairs of \mathcal{X} and \mathcal{Y} . Therefore, we get 3636 pairs of training datum. An example of training data generation is shown in Fig. 3.



Fig. 2. Left: A bounding box set up with a camera pose for 3D volume definition and 2D capture. Right: An image plane aligned to the bounding box with an image rendered by the defined camera projection.



Fig. 3. Training data generation. The first row from left to right is from a polygon-strip model randomly posed in the defined bounding box to generate the ground-truth volumetric field y and to grow hair strands afterwards. Then the fourth of the second row is shown a 2D image captured. Based on the 2D image and the bust model, x are calculated made up of orientation, confidence, and bust depth maps (from left to right in the second row).

5. Hair-GAN

With the 2D maps and the bust depth map extracted from the input image, the goal of our *Hair-GAN* is to produce a 3D orientation volume encoding both the occupancy and orientation information to guide the hair synthesis. The input of the network is a 2D tensor \mathcal{X} with a size of 1024×1024 , composed of 4 feature channels that are captured in the unified model space: hair orientation map (the 2D direction vector *XY* encoded as the color of RG), confidence map (the confidence value as the color of gray), and the bust depth map (the depth value as the color of gray). The output is a 3D tensor \mathcal{Y} of size $128 \times 128 \times 96$, where the hair orientation vectors are encoded in color of RGB. We first describe the loss functions of our adversarial training networks (Section 5.1). Next, we describe the architecture (Section 5.2) and the training strategy of our *Hair-GAN* (Section 5.3).

5.1. Loss functions

The GANs (Goodfellow et al., 2014) is trained in a strategy of competition between two networks: the generator and the discriminator. We refer to the function form of WGAN-GP (Gulrajani

et al., 2017) in order to get an easy training. For our cases, the goal is to train a generator $G(\mathcal{X})$ that maps the input 2D tensor to a desired output 3D tensor $\tilde{\mathcal{Y}} : \tilde{\mathcal{Y}} = G(\mathcal{X})$. In the meantime, the discriminator maximizes the Wasserstein-1 distance between the generator distribution of $G(\mathcal{X})$ and the target distribution of \mathcal{Y} with a conditional latent projection $P(\mathcal{X})$.

Discriminator. The objective of our discriminator is to minimize the energy:

$$L_{D} = E[D(\widetilde{\mathcal{Y}}, P(\mathcal{X}))] - E[D(\mathcal{Y}, P(\mathcal{X}))] + \lambda E[(\|\nabla_{\hat{\mathcal{Y}}} D(\hat{\mathcal{Y}}, P(\mathcal{X}))\|_{2} - 1)^{2}]$$
(1)

Similar to Gulrajani et al. (2017), the third term in this function is the gradient penalty for random samples $\hat{\mathcal{Y}}$, that $\hat{\mathcal{Y}} \leftarrow \epsilon \mathcal{Y} + (1 - \epsilon)\mathcal{Y}$, and ϵ is a random number in [0, 1]. The coefficient λ is set to 10. $P(\cdot)$ is CNNs (denoted in Fig. 4) to map 2D tensor \mathcal{X} into a 3D latent space in order to be concatenated with \mathcal{Y} or $\hat{\mathcal{Y}}$ referred to the strategy introduced in Reed et al. (2016). And the parameters in $P(\cdot)$ are trained along with those of D.

Generator. Following the original WGAN-GP (Gulrajani et al., 2017), the energy function for generator is defined as

$$L_G = -E[D(\widetilde{\mathcal{Y}}, P(\mathcal{X}))] \tag{2}$$



Fig. 4. An overview of our Hair-GAN architecture. The generator and the discriminator are trained in conjunction.

However, in our experiment, we find that this function does not work well to optimize the generator, since the difference of the distribution between the real and the fake cannot be easily determined by the sign of plus and minus. Inspired by a previous work (Gatys et al., 2016) where they use selected layers of pre-trained networks VGG as feature representation to transfer texture style from a source image to a target one, here, we introduce the losses of style and content to our research, where the features are represented in the domains of selected discriminator layers. Thus, the objective to optimize the generator is to minimize the energy:

$$L_{G}^{*} = \alpha L_{content} + \beta L_{style} = \alpha \sum_{l} L_{content}^{l} + \beta \sum_{l} L_{style}^{l}$$
(3)

 α and β are the weighting factors. As described in Gatys et al. (2016), the content loss is defined by square-error loss between feature representations:

$$L_{content}^{l} = \frac{1}{2} \sum_{ik} [f_{ik}^{l}(\mathcal{Y}, P(\mathcal{X})) - f_{ik}^{l}(\widetilde{\mathcal{Y}}, P(\mathcal{X}))]^{2}$$
(4)

Here *l* is a selected layer, *i* is the is *i*th feature map, *k* is the index in feature tensor, and *f* is the discriminator features (denoted in Fig. 4). The style loss is defined by the mean-squared distance between the Gram matrices, where each element is calculated by the inner product between the vectorized feature maps *i* and *j*: $A_{ij}^{l} = \sum_{k} f_{ik}^{l} f_{jk}^{l}$. The objective is:

$$L_{style}^{l} = \frac{1}{4N_{l}^{2}M_{l}^{2}}\sum_{ij}[A_{ij}^{l}(\mathcal{Y}, P(\mathcal{X})) - A_{ij}^{l}(\widetilde{\mathcal{Y}}, P(\mathcal{X}))]^{2}$$
(5)

Here N_l is the number of feature maps and M_l is the size of feature tensors (e.g. if l = 0, $N_0 = 3$ and $M_0 = 128 \times 128 \times 96$).

5.2. Architecture

Fig. 4 and Table 1 describe the architecture of generator and discriminator in detail. To clearly specify the architecture of our *Hair-GAN*, we use the following notations: let *in*(resolution, feature channels) and *out*(resolution, feature channels) represent input and output information for operation blocks; *C*(input channels, output channels, strides) denotes the convolutional layer with a ReLU activation followed; ϖ is dimensional expansion layer, and ζ is fully connected node. We use + as the element-wise addition in the residual blocks constituted by *C*, and *I* denotes the input tensor current layer. For all 2D convolutional layers *C*₃. The operation blocks for *X*-, *Y*-, *Z*-info have the same form of scheme.

Table 1	
Architecture of the generator and the discriminator.	
Generator	
$\begin{array}{l} in(1024 \times 1024, 4) \Leftarrow \mathcal{X} \\ C_2(4, 16, 2) + [C_2(4, 8, 2), C_2(8, 16, 1)] \\ C_2(16, 64, 2) + [C_2(16, 32, 2), C_2(32, 64, 1)] \\ C_2(64, 256, 1) + [C_2(64, 128, 2), C_2(128, 256, 1)] \\ I + [C_2(256, 256, 1), C_2(256, 256, 1)] \\ out(128 \times 128, 256) \end{array}$	
X-, Y-, Z-blocks $in(128 \times 128, 256)$ $I + [C_2(256, 256, 1), C_2(256, 256, 1)]$ $I + [C_2(256, 256, 1), C_2(256, 256, 1)]$ $C_2(256, 128, 1)$ $C_2(128, 96, 1)$ $\overline{\omega}$ $uut(128 \times 128 \times 96, 1)$	
Concat. the outs from X-, Y-, Z-blocks in(128 × 128 × 96, 3) $I + [C_3(3, 3, 1), C_3(3, 3, 1)]$ $I + [C_3(3, 3, 1), C_3(3, 3, 1)]$ out(128 × 128 × 96, 3) $\Rightarrow \tilde{y}$	
Discriminator	
$ \begin{array}{c} \hline P(\cdot) \ block \\ in(1024 \times 1024, 4) \Leftarrow x \\ C_2(4, 32, 2) \\ C_2(32, 64, 2) \\ C_2(32, 64, 2) \\ C_2(128, 96, 1) \\ \hline m \\ \hline \end{array} $	Concat. $\tilde{\mathcal{Y}}/\mathcal{Y}$ with $P(\mathcal{X})$ in(128 × 128 × 96, 4) $C_3(4, 32, 2)$ $C_3(32, 64, 2)$ $C_3(64, 128, 2)$ $C_3(126, 256, 2)$ $C_2(256, 512, 2)$
$out(128 \times 128 \times 96, 1)$	ζ

Generator. We illustrate the generator as some blocks. The first block with χ as input, composed of 4 residual networks (Lim et al., 2017) element-wisely adding activation from earlier layers to later layers in order to get a residual correction from high-to low-level information, down-samples feature maps to a latent code from 1024 × 1024 to 128 × 128, along with the number of features increasing from 4 to 256. Then *X*-, *Y*-, and *Z*-blocks separately encode the latent code to features with the number of channels as 96, the resolution along *Z*-axis in the resulting volume. ϖ converts the succession of 2D features to a single channel of 3D features. Afterwards, we concatenate the output from *X*-, *Y*-, *Z*-blocks as input into the following 3D residual convolutional networks. More details please refer to Fig. 4 and Table 1.

Discriminator. Taking into consideration of the correspondence between the 2D input \mathcal{X} and the 3D desired output $\tilde{\mathcal{Y}}/\mathcal{Y}$, inspired by Mirza and Osindero (2014), we concatenate $\tilde{\mathcal{Y}}/\mathcal{Y}$ with $P(\mathcal{X})$, a feature map encoding \mathcal{X} to a 3D latent space with the same



Fig. 5. From left to right: input image, image aligned with the bust model fitted, and 2D maps as input for the generator network.



Fig. 6. Hair barbola. Modeling hair from single-view hair images created by cut-and-paste of different hairstyles. The left are input images of drastically different hairstyles. The middle is the hair barbola images. The right is 3D hair modeling results generated by our Hair-GAN with barbola images as input.

resolution as $\tilde{\mathcal{Y}}/\mathcal{Y}$. Then the concatenated 3D feature tensor is convoluted by a number of filters until the layer of ζ finally distinguishes the real and fake.

5.3. Training strategy

In Gulrajani et al. (2017), the generator updates after fivetimes training of the discriminator, which takes a lot of time. For time efficiency, we apply the two-timescale update rule (TTUR) (Heusel et al., 2017) to update the discriminator only once instead of five times. We employ the commonly used ADAM optimizer (Kingma and Ba, 2014) with $\beta_1 = 0$ and $\beta_2 = 0.9$ for training. The learning rate for the discriminator is set to 0.0003, and 0.0001 for the generator. Generally, our *Hair-GAN* are designed to generate a $128 \times 128 \times 96$ 3D volume encoded in both of the occupancy and orientation fields, using 2D maps as input with a size of 1024×1024 . The batch size for training is set to 4.

For the generator objective, the style and content weighting factors are set as: $\alpha = 1e - 2$, $\beta/\alpha = 5e + 2$. Selected layers for content loss are 0, 3, 6 and l = 0, 1, 2, 3, 4 for style loss. Specifically, when l = 0, $P(\mathcal{X})$ can be removed from $L_{content}^0$ and L_{style}^0 .

6. Hair synthesis

Given an input image, we first align it to our unified model space and generate 2D information maps and depth map (Section 6.1) as input to our generator network. After the 3D orientation volume is generated, we build a rough shape to confine hair synthesis and finally generate our strand-level hair model (Section 6.2).

6.1. Preprocessing

Similar to previous single-view methods (Chai et al., 2016; Hu et al., 2015; Chai et al., 2012), with an input image, we first run the face alignment algorithm (Cao et al., 2014) to fit the identical bust model to the facial landmark points detected in the image. Also, we segment the hair region and to generate the hair direction predictor to remove directional ambiguity for the orientation map computation afterwards (Chai et al., 2016, 2013). Specifically, assume that we have transformations: s (scaling) r (rotation) and t (translation) involved in the bust fitting, s and t are applied to the image and r is applied to the identical bust model (for correct depth map generation). After the alignment, we capture the image and hair mask by the projection described in Section 4.1 and generate maps of orientation, the confidence of orientation, and bust depth referred to Section 4.2. We also remove the directional ambiguity with the help of direction predictor from the orientation map. Fig. 5 shows an example of preprocessing.

6.2. Post-processing

We find that a direct trace along the 3D orientation field computed by our generator network could lead to non-smooth results for some complex hairstyles because the orientation is not smooth in the flow field. Therefore, we apply the marching cube



Fig. 7. Slices of volumetric fields. From left to right: the ground truth, our *Hair-GAN* output, the result learned using Eq. (2) as generator objective, a CNN generator output with only naive $L_{content}^0$ loss (no GAN training).



Fig. 8. Sampled strands tracing along the orientation field generated by *Hair-GAN* with (left) / without (right) the confidence map as input (cf. Fig. 5). There are more structural discontinuities in the tracing result generated without the confidence map input (the right blue square).



Fig. 9. Up to down in the first column: a synthetic image with blue-colored hair contour, comparison of hair mask (the blue colored is the mask of the input image, while the magenta colored is mask projected by the modeling result), orientation difference map (white color indicates there is a big difference between the input and the modeling result in orientation map). The right three columns: the first row is our hair modeling results and the second row is the ground-truth.

algorithm (Lorensen and Cline, 1987) on the occupancy field to compute a rough shape and smooth it by Poisson method (Kazhdan et al., 2006). We then smooth the exterior orientation field tangent to the surface of the rough shape. We also optionally refine the orientation field by warping the image orientation map to the rough shape surface. Constrained by the normal direction of the scalp and the image warping orientation field, we grow hair strands from the scalp uniformly distributed inside the rough shape following the previous method of hair synthesis (Chai et al., 2013; Hu et al., 2015). After that, in order to refine the local detail matching with the input image, we run the strands deformation according to the projected image orientation map, as in Hu et al. (2015).



Fig. 10. Given an input image (leflt), modeling hair results with (middle) / without (right) post-processing.



Fig. 11. Comparisons with Chai et al. (2016) (the top two rows), Zhou et al. (2018) (the middle two rows) and Saito et al. (2018) (the bottom three rows). For each comparison, from left to right: input image, our results (rough shape, orientation field, 3D hair models), the result of the previous method.

7. Results and discussion

The implementation of our hair structure recovery runs on a PC with an Intel Core i7-4790 CPU, 32G of memory and an NVIDIA

GeForce GTX 1080Ti. It took about 10 days to train our *Hair-GAN* with 200k generator iterations using our GPU. We use *tensorflow* frameworks to set up our deep learning networks. With a single-view input image as input, using our unoptimized pipeline, it



Fig. 12. Hair modeling results using our Hair-GAN with single-view images as input.

takes less than 3 s for data preparation by the automatic method of hair parsing (Chai et al., 2016), about 1 second to generate 3D volume field using our networks (after tensorflow python call), and less than 30 s to synthesis final hair model. Fig. 12 demonstrates the effectiveness of our method on a variety of hairstyles.

Hair-GAN can generate 3D models of hairstyles which do not exist in our hairstyle database. Different from Saito et al. (2018) to align the 3D hair structure with 2D image in a latent space, we build parametric transformation from 2D image to 3D structure guaranteed by the local respective property of convolutional neurons. Because we consider a hairstyle as a fusion of local style patterns as mentioned in Zhang et al. (2018), rather than a combination of different styled strands (cf. Chai et al., 2016).

Therefore, even with the image of an exaggerated hairstyle created by two drastically different hairs, our method can also create a strand-level 3D hair model as shown in Fig. 6.

Evaluation. To evaluate the efficacy of our generator objective function, we compare our volumetric field output with the ground truth in Fig. 7. Additionally, we also train our generator network separately by using the original function of WGAN-GP (Eq. (2)) and the naïve L_2 loss. The original function leads to much ambiguity and noise in both the fields of occupancy and orientation (Fig. 7, middle right). A naïve training of the generator just using the L_2 loss ($L_{content}^0$) results in a blurred output (Fig. 7, right), just like its usual effect on the super-resolution field (Xie et al., 2018). Our objective function of the generator considers

the distribution similarity of selected discriminator layers in both style and content, which leads to a higher perceptual quality.

To evaluate the impact of the confidence map on the 3D volumetric field result, we train another GAN without the confidence map as input. The *Hair-GAN* without confidence map can also generate 3D volumetric field to encode the hair structure globally fitted with the input image. However, there exit more local structural discontinuities in the orientation field than the result field computed with confidence map input, which are revealed by uniformly sampled strand tracing along the orientation as shown in Fig. 8. Thus, we prefer to involve the confidence map within the input maps which indicates reliable hair strands' trajectory.

We also evaluate the quality of our modeling results in Fig. 9. With a synthetic image as input, our *Hair-GAN* produces the modeling results similar to the ground-truth model even in views distant from the input one. The projected mask of our result can be well matched with the input mask. Moreover, the projected orientation map is similar to the input image in most of the hair region.

Additionally, the direct output of our generator network has already been able to provide a reasonable reconstruction of 3D volumetric field to guide hair synthesis (see Fig. 10). The postprocessing is optional for structural detail refinement.

Comparisons. We compare our method with the state-of-theart methods of Chai et al. (2016), Zhou et al. (2018) and Saito et al. (2018) in Fig. 11. The results of Chai et al. (2016) (Fig. 11, top two rows) rely on their database models which are used as the prior structures of hair modeling. Our results are generated by deep learning networks. Using our Hair-GAN, we can generate more faithful modeling results of hairstyles even they do not exist in our training data. Since the network of Saito et al. (2018) is trained with natural images, their method tends to recover hairstyles perceived in the original images rather than the modified ones in the case of hair barbola (the last row in Fig. 11). And our result is comparable to those by Saito et al. (2018) on those typical hairstyles (the 5th and 6th row in Fig. 11). Zhou et al. (2018) (Fig. 11, middle two rows) compute a feature map of low resolution parameterized on the scalp. Our results show more details in hair structure.

8. Conclusion

We have presented our *Hair-GAN*, an architecture of networks to learn 3D hair structure from a single-view input image. We use 2D CNNs to learn 3D cues and convert the succession of 2D features to a 3D feature map of a single channel by the dimensional expansion layer. The core of our method is to follow the training strategy of GANs, the competition between the generator and the discriminator. In order to produce a volumetric field of high perceptual quality, we design the generator objective function in consideration of the distribution similarity in all selected discriminator lays, inspired by the method of style-transfer. Our results on a variety hairstyles resemble the input images in both of the hair contour and hair texture. The network can also learn more details in the views distant from the input one.

Our architecture of *Hair-GAN* is subject to a number of limitations, which may inspire interesting future work. First, like almost all single-view hair modeling methods, our training datum only takes the frontal view image into account. Our method would fail when hair is partially occluded. Second, although our network can produce a considerable volumetric field for complex hairstyles, hair tracing method still merits some improvement on the quality of the final modeling result. Current efforts (Chai et al., 2016; Hu et al., 2014; Zhang et al., 2018) improve the hair strand geometry by using the guidance of database models. We hope in the future, there is a method independent of database to synthesis complex hairstyles, e.g., even the highly curly. Third, we hope to expand our training data dynamically to cover the ever-changing fashion of hairstyles. In addition, the resolution and size of the volumetric field have an impact on the quality of the final modeling results. Some complex hairstyles may need a larger and more subtilized volume. We believe the advances in efficient data structures for 3D-GANs could leverage this problem.

Acknowledgments

We would like to thank the anonymous reviewers for their constructive suggestions. This work is partially supported by the National Key Research & Development Program of China (2018YFE0100900), the China Young 1000 Talent Program, and the Fundamental Research Funds for the Central Universities.

Declaration of competing interest

There is no conflict of interests.

References

- Cao, X., Wei, Y., Wen, F., Sun, J., 2014. Face alignment by explicit shape regression. Int. J. Comput. Vis. 107 (2), 177–190.
- Chai, M., Luo, L., Sunkavalli, K., Carr, N., Hadap, S., Zhou, K., 2015. High-quality hair modeling from a single portrait photo. ACM Trans. Graph. 34 (6), 204:1–204:10.
- Chai, M., Shao, T., Wu, H., Weng, Y., Zhou, K., 2016. Autohair: Fully automatic hair modeling from a single image. ACM Trans. Graph. 35 (4), 116:1–116:12.
- Chai, M., Wang, L., Weng, Y., Jin, X., Zhou, K., 2013. Dynamic hair manipulation in images and videos. ACM Trans. Graph. 32 (4), 75:1–75:8.
- Chai, M., Wang, L., Weng, Y., Yu, Y., Guo, B., Zhou, K., 2012. Single-view hair modeling for portrait manipulation. ACM Trans. Graph. 31 (4), 116:1–116:8. Echevarria, J.I., Bradley, D., Gutierrez, D., Beeler, T., 2014. Capturing and stylizing
- hair for 3d fabrication. ACM Trans. Graph. 33 (4), 125:1–125:11.
- Gatys, L.A., Ecker, A.S., Bethge, M., 2016. Image style transfer using convolutional neural networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 2414–2423.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y., 2014. Generative adversarial nets. In: Advances in Neural Information Processing Systems, pp. 2672–2680.
- Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., Courville, A.C., 2017. Improved training of wasserstein gans. In: Advances in Neural Information Processing Systems. pp. 5767–5777.
- Herrera, T.L., Zinke, A., Weber, A., 2012. Lighting hair from the inside: A thermal approach to hair reconstruction. ACM Trans. Graph. 31 (6), 146:1–146:9.
- Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S., 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In: Advances in Neural Information Processing Systems. pp. 6626–6637.
- Hu, L., Ma, C., Luo, L., Li, H., 2014. Robust hair capture using simulated examples. ACM Trans. Graph. 33 (4), 126:1–126:10.
- Hu, L., Ma, C., Luo, L., Li, H., 2015. Single-view hair modeling using a hairstyle database. ACM Trans. Graph. 34 (4), 125:1–125:9.
- Hu, L., Saito, S., Wei, L., Nagano, K., Seo, J., Fursund, J., Sadeghi, I., Sun, C., Chen, Y.-C., Li, H., 2017. Avatar digitization from a single image for real-time rendering. ACM Trans. Graph. 36 (6), 195:1–195:14.
- Jakob, W., Moon, J.T., Marschner, S., 2009. Capturing hair assemblies fiber by fiber. ACM Trans. Graph. 28 (5), 164:1–164:9.
- Kazhdan, M., Bolitho, M., Hoppe, H., 2006. Poisson surface reconstruction. In: Proceedings of the Fourth Eurographics Symposium on Geometry Processing, pp. 61–70.
- Kingma, D.P., Ba, J., 2014. Adam: A method for stochastic optimization, arXiv preprint arXiv:1412.6980.
- Lim, B., Son, S., Kim, H., Nah, S., Lee, K.M., 2017. Enhanced deep residual networks for single image super-resolution. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops, vol. 1 (2), p. 4.
- Lorensen, W.E., Cline, H.E., 1987. Marching cubes: A high resolution 3d surface construction algorithm. ACM Trans. Graph. 21 (4), 163–169.
- Luo, L., Li, H., Rusinkiewicz, S., 2013. Structure-aware hair capture. ACM Trans. Graph. 32 (4), 76:1–76:12.
- Mirza, M., Osindero, S., 2014. Conditional Generative Adversarial Nets, arXiv preprint: abs/1411.1784.
- Odena, A., Olah, C., Shlens, J., 2017. Conditional image synthesis with auxiliary classifier GANs. In: Proceedings of the 34th International Conference on Machine Learning, vol. 70, pp. 2642–2651.

- Paris, S., Chang, W., Kozhushnyan, O.I., Jarosz, W., Matusik, W., Zwicker, M., Durand, F., 2008. Hair photobooth: geometric and photometric acquisition of real hairstyles. ACM Trans. Graph. 27 (3), 30:1–30:9.
- Reed, S., Akata, Z., Yan, X., Logeswaran, L., Schiele, B., Lee, H., 2016. Generative adversarial text-to-image synthesis. In: Proceedings of the 33rd International Conference on Machine Learning, pp. 1060–1069.
- Saito, S., Hu, L., Ma, C., Luo, L., Li, H., 2018. 3d hair synthesis using volumetric variational autoencoders. ACM Trans. Graph. 37 (6), 208:1–208:12.
- Wang, L., Yu, Y., Zhou, K., Guo, B., 2009. Example-based hair geometry synthesis. ACM Trans. Graph. 28 (3), 56:1–56:9.
- Ward, K., Bertails, F., Kim, T.-Y., Marschner, S.R., Cani, M.-P., Lin, M.C., 2007. A survey on hair modeling: Styling, simulation, and rendering. IEEE Trans. Vis. Comput. Graphics 13 (2), 213–234.
- Xie, Y., Franz, E., Chu, M., Thuerey, N., 2018. Tempogan: A temporally coherent, volumetric gan for super-resolution fluid flow. ACM Trans. Graph. 37 (4), 95.
- Zhang, M., Chai, M., Wu, H., Yang, H., Zhou, K., 2017. A data-driven approach to four-view image-based hair modeling. ACM Trans. Graph. 36 (4), 156:1–156:11.
- Zhang, M., Wu, P., Wu, H., Weng, Y., Zheng, Y., Zhou, K., 2018. Modeling hair from an rgb-d camera. ACM Trans. Graph. 37 (6), 205:1–205:10.
- Zhou, Y., Hu, L., Xing, J., Chen, W., Kung, H.-W., Tong, X., Li, H., 2018. HairNet: Single-View hair reconstruction using convolutional neural networks. In: Proceedings of the European Conference on Computer Vision, pp. 235–251.
- Zhu, J., Park, T., Isola, P., Efros, A.A., 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In: 2017 IEEE ICCV. pp. 2242–2251.