

Geometry-aware Single-image Full-body Human Relighting

Chaonan Ji¹, Tao Yu¹, Kaiwen Guo², Jingxin Liu³, and Yebin Liu¹

¹ Department of Automation, Tsinghua University, China

² Meta Reality Labs

³ Guangdong OPPO Mobile Telecommunications Corp., Ltd

Abstract. Single-image human relighting aims to relight a target human under new lighting conditions by decomposing the input image into albedo, shape and lighting. Although plausible relighting results can be achieved, previous methods suffer from both the entanglement between albedo and lighting and the lack of hard shadows, which significantly decrease the realism. To tackle these two problems, we propose a geometry-aware single-image human relighting framework that leverages single-image geometry reconstruction for joint deployment of traditional graphics rendering and neural rendering techniques. For the de-lighting, we explore the shortcomings of UNet architecture and propose a modified HRNet, achieving better disentanglement between albedo and lighting. For the relighting, we introduce a ray tracing-based per-pixel lighting representation that explicitly models high-frequency shadows and propose a learning-based shading refinement module to restore realistic shadows (including hard cast shadows) from the ray-traced shading maps. Our framework is able to generate photo-realistic high-frequency shadows such as cast shadows under challenging lighting conditions. Extensive experiments demonstrate that our proposed method outperforms previous methods on both synthetic and real images.

Keywords: single-image human relighting, single-image geometry reconstruction, ray tracing, neural rendering

1 Introduction

Human relighting aims to relight a target human with correct shadow effects under a desired illumination. Relighting can realize seamless replacement of the background while ensuring the light consistency of the foreground and the background and has a huge application prospect in film-making, video chat and Virtual Reality. Single-image human relighting for a general human person is convenient and promising for amateur photographers but also more challenging because the difficulty of decoupling lighting, geometry, and surface material from a single image.

Most existing single-image human relighting methods focus on portrait relighting [40, 50, 52, 37, 5, 16, 18, 33, 38, 43, 44, 47, 51, 32, 46, 59, 58, 22] and only a



Fig. 1. Given a single human image and an arbitrary high dynamic range lighting environment, our framework estimates a high-quality albedo map and generates photo-realistic relit images of the target human subject under the desired lighting conditions.

few works [28, 25, 11] focus on single-image full-body relighting. Since the style, texture and material of human clothes are widely varied and the geometry and poses of clothed humans are usually complex, decoupling the albedo and lighting is highly ill-posed. Moreover, mutual occlusion of limbs will produce self-shadows, which are not only difficult to remove for de-lighting tasks, but also challenging to generate under the target lighting conditions for relighting tasks. Previous full-body relighting methods [28, 25, 11] have attempted to address these issues, but the results still have the following drawbacks: (i) inability to disambiguate albedo and lighting (ii) inability to model high-frequency shadows due to reliance on Spherical harmonics representation of lighting.

We propose a novel geometry-aware single-image human relighting framework that leverages SOTA single-image geometry reconstruction [42, 19] for fully leveraging the traditional graphics rendering and neural rendering techniques simultaneously. First of all, accurate intrinsic decomposition for albedo estimation (de-lighting) is the cornerstone of producing high-quality relighting results. Previous methods [28, 25, 11] have used the UNet to infer albedo and suffer from severe entanglement between lighting and albedo. We discover that the skip connections in UNet is the culprit and found that HRNet[49] performs much better for de-lighting. In addition we further improve the de-lighting performance by: i) removing the aggressive down-sampling operations in the early stage, ii) eliminating skip connections and iii) fusing multi-scale features while maintaining high-resolution representations in the HRNet, and finally achieve better disentanglement between lighting and albedo on both synthetic and real images.

More importantly, even with high-quality albedo, we still need photo-realistic shadows to produce realistic relit results. Limited by the expressive ability of the spherical harmonics (SH) lighting model and the lack of 3D geometry information, previous methods [28, 25, 11] can only produce low-frequency shadows. Low-resolution environment maps [50, 52] and pre-filtered environment maps [40] have also been explored and struggle to generate hard cast shadows. With the development of single-image 3D human reconstruction technology, obtaining high-quality mannequins from a single image is possible, which can provide more 3D prior information for human relighting. Hence, we propose a geometry-

aware relighting method, which consists of a ray tracing-based per-pixel lighting representation that explicitly models high-frequency shadows and a learning-based shading refinement module to restore realistic shadows. The key idea of our method is that with the estimated 3D human model, we can render photo-realistic shading maps that encode full-band lighting information under target lighting conditions using ray tracing. This lighting representation is able to preserve high-frequency shadows and significantly improve the quality of the relit results which remains difficult for learning-based methods. However, the ray-traced shadows still suffer from artifacts due to the errors in geometry estimation, which prevents direct composition of the relit results. Thus, we further propose a learning-based refinement module that utilizes the ray-traced shading maps and the inferred ambient occlusion map as input to restore photo-realistic high-frequency shadows with rich local details and clear shadow boundaries. The proposed framework dramatically improves relighting performance and produces a photo-realistic relit image with high-frequency self-shadowing effects under the target lighting conditions.

To conclude, our main contributions are the following:

- A novel geometry-aware single-image human relighting framework that combines single-image 3D human reconstruction, ray tracing and neural rendering technologies.
- We demonstrate the UNet with skip connections is not suitable for the de-lighting task and propose a modified HRNet for achieving better disentanglement between lighting and albedo on both synthetic images and real photographs.
- We propose a ray tracing-based per-pixel lighting representation and a learning-based shading refinement module that utilizes the inferred ambient occlusion map as auxiliary input to restore photo-realistic shadows while preserving rich local details.

2 Related Work

2.1 Person-Specific Human Relighting

Turner Whitted [55] first used recursive ray tracing to simulate global illumination and render realistic images under target lighting conditions. However, the forward rendering technique requires a high-precision 3d model and corresponding PBR textures of the target object, which are unavailable or costly for real-world objects. Debevec *et al.*[13] proposed collecting OLAT(one-look-at-a-time) images using LightStage to synthesize specific person’s face under novel illuminations with no 3D model required. Subsequent works improved LightStage to capture higher-quality images and a larger range of human bodies[21, 14, 53, 12, 54, 8, 20, 57]. Guo *et al.*[20] attempted to combine image-based rendering technique with geometry and material estimated by LightStage and achieved unprecedented quality and photorealism for free viewpoint videos. However, collecting OLAT data, gradient data and object geometry requires professional

equipment and it is a time-consuming and complicated process for person-specific relighting. Li *et al.*[31] attempted to relight target human from multi-view video recorded under unknown illumination conditions and Imber *et al.*[23] extended this approach to scene relighting by introducing intrinsic textures. These two methods can produce high-quality relit results but multiview video must be collected for every target human or scene.

2.2 General Human Relighting

Deep learning makes general human relighting possible. Most existing general human relighting methods[40, 50, 52, 37, 5, 16, 18, 33, 38, 43, 44, 47, 51, 32, 46, 59, 58, 22, 35] focus on human portrait relighting. [50, 57, 40] achieved amazing single-image portrait relighting utilizing a dataset synthesized from OLAT images. For human full-body relighting, Meka *et al.*[36] combined traditional geometry pipelines with neural rendering to generate relit results using gradient images and estimated human geometry. Kanamori *et al.*[25] proposed an occlusion-aware single-image full-body human relighting method to infer albedo, geometry and illumination leveraging spherical harmonics (SH) lighting model. On the basis of [25], Lagunas *et al.*[28] added specular reflectance and light-dependent residual terms to explicitly handle highlights and Tajima *et al.*[11] used a residual network to restore neglected light effects. However, due to the limitation of the expressive ability of the SH lighting model, [25, 28, 11] could only produce low-frequency shadows and their method may fail under harsh illuminations.

2.3 Inverse Rendering

Ramachandran *et al.*[41] proposed the *shape-from-shading* method to estimate shape from shading given an input image with known lighting conditions. Subsequent works [10, 39, 34] assume a simple light source model such as directional and point light sources. Based on the Retinex theory [29], intrinsic images [4, 56, 7, 6, 27, 45, 30, 15] aims to decompose an input image into reflectance and shading. Recent single-image human relighting studies [25, 28, 11] have drawn on the idea of intrinsic images to decompose an input image into albedo, shape and illumination using UNet. We have further improved the intrinsic images decomposition results in the de-lighting stage utilizing the modified HRNet [49] and achieved better disentanglement between albedo and lighting.

3 Overview

Following current work [52], our framework consists of two stages: de-lighting and relighting. Fig. 2 shows the architecture of the whole framework. For the de-lighting stage, we first use two geometry networks (Geometry Module) to infer the per-pixel normal map \widehat{N} and ambient occlusion map \widehat{AO} separately. Similar to [40], the inferred normal map and the input image are concatenated as the input of Albedo Module to infer albedo \widehat{A} . For the relighting stage, we

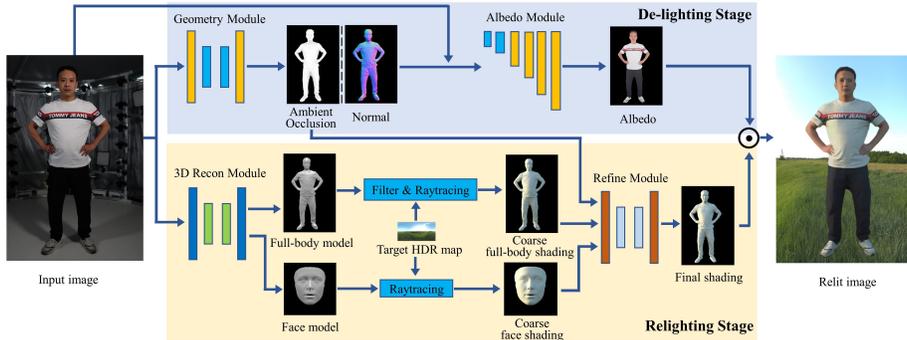


Fig. 2. Illustration of our framework architecture. There are two stages in our method: de-lighting and relighting. The de-lighting stage takes the input image and outputs estimated albedo \hat{A} (Section 4). For the relighting stage (Section 5), a full-body 3D model and a face 3D model are estimated by 3D Recon Module and then are sent to the renderer to render coarse shading maps (Section 5.1). The Refine Module takes the coarse shading maps and the inferred ambient occlusion map as input and produces the final shading map (Section 5.2).

first render the 3D human models estimated by [42] and [19] (3D Recon Module) using ray tracing to obtain the coarse full-body shading map S_{coarse}^{body} and coarse face shading map S_{coarse}^{face} . Then the Refine Module takes these coarse shading maps and the inferred ambient occlusion map as input to produce the refined full-body shading map S_{fine}^{body} and refined face shading map S_{fine}^{face} , which are composited together to obtain the final shading map \hat{S} . Details regarding the specific architecture and implementations of the networks are provided in the supplementary materials. Finally, the dot product of the inferred albedo and the final shading map is obtained to produce the relit image I_{relit} using the following formula:

$$I_{relit} = \hat{A} \odot \hat{S}$$

4 De-lighting

One of the most important factor for achieving better de-lighting results is the network architecture. Although the UNet architecture has been widely used in previous de-lighting networks [28, 25, 40, 52, 43], we find that it usually produces heavy lighting and albedo entangled results, possibly because of the fact that the lighting and albedo features remain coupled in the UNet architecture.

To guarantee that the lighting and albedo can be disentangled in the network architecture level, we modified the HRNet [49] architecture and achieved better delighting performance than UNet as shown in Fig. 3. Note that the vanilla HRNet can only produce oversmoothed albedo map due to the aggressive downsampling operations in the beginning stage. Even with skip connections



Fig. 3. Estimated albedo of a real image. Unet comes from RH [25] and is retrained on our dataset. We zoom in the area in the red anchor and place it at the right of corresponding image.

between the input downsampling features and output upsampling features, the results of vanilla HRNet remain poor and even make unrecovered shadows appear in the inferred albedo map, especially in outdoor scenarios. We assume that skip connections have a negative impact on the final de-lighting results because high-resolution features from the input still contain environmental lighting information. Based on this assumption, we modify the vanilla HRNet by removing the downsampling operations directly (which also avoid skip connections) to fuse multiscale features while maintaining high-resolution representations. Moreover, we modify HRNet to output lighting prediction at the layer with the lowest resolution of the final stage. This strategy guarantees a better decomposition between lighting and albedo at both the architecture and feature representation level. Details regarding the specific architecture and implementations are provided in the supplementary materials.

To estimate the albedo of an input human image, we employ losses as follows:

$$L_{DL} = \lambda_{input} \left\| \hat{A} - A \right\|_1 + \lambda_{vgg} Vgg(\hat{A}, A) \\ + \lambda_{light} \left\| w \odot \log(1 + \hat{I}_l) - \log(1 + I_l) \right\|_2^2$$

where \hat{A} is the inferred albedo and A is the ground truth albedo. Similar to portrait relighting [50], we use the weighted log- L_2 loss. \hat{I}_l and I_l are latitude-longitude representation of environmental illumination and are used to describe the estimated lighting and ground truth lighting, respectively. The ground truth lighting map is downsampled to $16 \times 32 \times 3$ using Gaussian pyramid. w is the solid angle of each “pixel”. The L_{DL} is the total loss for the de-lighting stage, and λ_{input} , λ_{vgg} and λ_{light} are the weight factors. Empirically, we find $\lambda_{input} = 500$, $\lambda_{vgg} = 100$, $\lambda_{light} = 0.025$ achieves the best performance.

5 Relighting

To render photorealistic shadows, we divide the relighting stage into two sub-stages: ray tracing and shading refinement. We first estimate 3D models of the

target human because the classical ray tracing algorithm requires a 3D model of target object.

For geometry estimation, given a single human image, we use PIFuHD [42] to estimate a complete 3D human model. Then we crop the human face from the input image and align it with the dense 3D face model following [19] by regressing the 3DMM parameters. Both the full-body 3D model and face 3D model have no texture because we only need to render shading maps of the target human.

5.1 Ray Tracing

For photorealistic rendering, we use the Cycles rendering engine in Blender [1] and a Principled BSDF shader. We set the camera mode to orthographic to ensure that the rendered shading map is pixel-aligned with the input image. Due to the limitations of the PIFuHD [42] network’s generation capacity and memory space, the surface of the estimated full-body 3D model is not smooth enough, which may produce checkered artifacts after raytracing. The schematic diagram of the artifacts is provided in the supplementary materials. To solve the problem, we adopt Laplacian smoothing [48] for the full-body 3D model and set smoothing steps to 10 with cotangent weight. Then the smoothed full-body model is sent to the renderer to render the full-body shading map under the target lighting environment using ray tracing. In addition, the generated 3D face model was directly sent to the renderer without smoothing for ray tracing.

In a word, we render a coarse full-body shading map S_{coarse}^{body} and coarse face shading map S_{coarse}^{face} under the target lighting environment using ray tracing. The rendered coarse shading maps are pixel-aligned with the input image.

5.2 Shading Refinement

After ray tracing under the target lighting environment, we obtained roughly correct shadows (especially for self-occlusions and hard shadows such as cast shadows) of the target human. However, the full-body 3D model reconstructed from a single image is not completely accurate and thus the ray-traced shading maps contain unnatural shadows and obvious geometry errors. To enhance the realism of the ray-traced full-body shading map, we introduce two refinement networks to compensate for shadow-rendering errors and restore a high-quality shading map. The first refinement network is designed to paint in the ray-traced full-body shading map and improve the overall quality of the shading details. The second refinement network is designed to refine facial shading details to ensure that we can fully leverage the geometry priors of human faces. Fig. 4 shows the entire shading refinement process.

Full-body Shading Refinement The full-body refinement network takes the coarse full-body shading S_{coarse}^{body} and the inferred ambient occlusion map \widehat{AO} as input and outputs the refined full-body shading residual. We choose an ambient occlusion map instead of a normal map for the following three reasons.

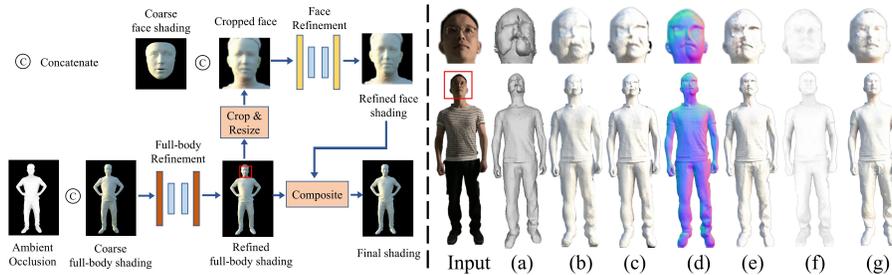


Fig. 4. Left: Illustration of Refine Module. A refined full-body shading map is inferred by the full-body refinement network. Then the cropped face from the refined full-body shading map and coarse face shading map are concatenated as the input of face refinement network, which outputs the refined face shading map. Finally, the refined face shading map and the refined full-body shading map are composited to generate the final relit shading map. **Right:** Refined shading maps. (a) 3D model estimated by PIFuHD [42] (b) Ray-traced shading map (c) Refined shading map without the inferred ambient occlusion map (d) Inferred normal map (e) Refined shading map with the inferred normal map (f) Inferred ambient occlusion map (g) Refined shading map with the inferred ambient occlusion map. Cropped faces of corresponding shading map are placed at the top.

First, the ambient occlusion map can supplement part of the self-shadows lost due to geometry prediction errors. Second, compared with inferring a normal map, inferring an ambient occlusion map is more robust under various lighting environments even with extreme lighting distributions. Third, existing 3D human reconstruction methods such as [42, 24] are highly dependent on the surface normal to predict 3D geometry surface details, which means that the normal prediction errors are consistent with the geometry errors and S_{coarse}^{body} cannot obtain extra correct geometry information from the normal map to compensate for existing shading errors. Fig. 4 shows the refined results with the normal and ambient occlusion maps as auxiliary inputs, respectively.

The architecture of full-body refinement network is similar to MIMO-UNet [9], and we deepen the network to 4 downsampling operations. The coarse full-body shading map S_{coarse}^{body} and the inferred ambient occlusion map \widehat{AO} are concatenated together as the original scale input and the downsampled \widehat{AO} is used as multiscale input. We use multiscale output for supervision. We add adversarial loss in the final highest-resolution output layer to help the network generate plausible shading effects and use PatchGAN as the discriminator. The training loss consists of content loss, fft loss and PatchGAN loss and is defined as follows:

$$\begin{aligned}
 L_{fb} = & \lambda_{content} \sum_{k=1}^K \left\| \widehat{S}_{fine}^{body^k} - S^k \right\|_1 + \left\| P(\widehat{S}_{fine}^{body}) - P(S) \right\|_2^2 \\
 & + \lambda_{fft} \sum_{k=1}^K \left\| F(\widehat{S}_{fine}^{body^k}) - F(S^k) \right\|_1
 \end{aligned}$$

where K is the number of levels and $K = 5$; $\widehat{S}_{fine}^{body^k}$ is the k_{th} level output and S^k is the k_{th} downsampled ground truth shading map. P is the PatchGAN discriminator, and $F()$ denotes the fast Fourier transform (FFT). $\lambda_{content}$ and λ_{fft} are weight factors and we empirically set $\lambda_{content} = 10$ and $\lambda_{fft}=0.001$.

Face Shading Refinement Although the above-mentioned refinement network produces realistic shading with rich details, the human eye is very sensitive to the details of the face and is able to distinguish small geometry and shadow errors. Therefore, we continue to refine the face region on the basis of the refined full-body shading map $\widehat{S}_{fine}^{body}$. The face refinement network takes S_{crop}^{face} cropped from $\widehat{S}_{fine}^{body}$ and S_{coarse}^{face} as the input and outputs the refined face shading residual. The training loss can be expressed as follows:

$$L_{ff} = \lambda_{face} \left\| \widehat{S}_{fine}^{face} - S^{face} \right\|_1 + \left\| L(\widehat{S}_{fine}^{face}) - L(S^{face}) \right\|_2^2$$

where L is the LSGAN discriminator, λ_{face} is the weight factor and is set to be $\lambda_{face} = 5$. S^{face} is the ground truth face shading map. We use UNet architecture as the backbone of face refinement network and details regarding the architecture are provided in the supplementary materials.

6 Implementation Details

We carefully select 811 scanned 3D human figures with good lighting conditions from Twindom [2], of which 700 figures are used for training and 111 figures for testing. The use of the dataset has been officially approved by Twindom. We collect 480 panoramic lighting environments sourced from www.HDRIHaven.com [3] and rotate them every 36 degrees to generate a total of 4800 HDR environment maps. We allocate 4600 HDR environment maps for training and 200 HDR environment maps for testing. To balance the amount of lighting in indoor and outdoor scenes, we add an extra 150 indoor HDR environment maps from the Laval Indoor Dataset [17] to the testing dataset. None of the test lighting conditions appear in the training dataset. Details regarding the specific data rendering, training and testing are provided in the supplementary materials.

7 Experiments

In this section, we first compare our method with previous state-of-the-art methods quantitatively and qualitatively to show that our method performs better on de-lighting task and produces more photorealistic relit results under challenging lighting conditions. Then we evaluate the key contributions of our proposed method and prove the effectiveness of the entire framework and each module. For quantitative comparisons, we adopt the metrics MSE, PSNR and SSIM to



Fig. 5. Relit results on synthetic and real images. The first column: synthetic images from our testing dataset. The second column: real images. For real images, "Reference" are the rendered images of a virtual 3D human model under the target lighting conditions and are used to indicate the position of shadows. The target HDR environment map is placed under the ground truth image or reference image.

compare the inferred albedo, relit shading and relit images with the corresponding ground truth images in our testing dataset. For qualitative comparisons, we show some intrinsic images decomposition results and relit results under target lighting conditions using both real images and synthetic images.

7.1 Comparisons with SOTA Methods

We compare our method with the state-of-the-art single-image human relighting methods RH [25], SFHR [28] and RHW [11]. All of them are based on intrinsic images decomposition and require only a single human image and target lighting for relighting as in our approach. Details regarding the specific comparison settings are provided in the supplementary materials.

Table 1. Quantitative comparisons of our single-image human relighting framework against prior works. Shading indicates the estimated shading map under target lighting condition. The value of MSE is scaled by multiplying 10^3 .

	Relight			Albedo			Shading		
	MSE ↓	SSIM ↑	PSNR ↑	MSE ↓	SSIM ↑	PSNR ↑	MSE ↓	SSIM ↑	PSNR ↑
RH [25]	8.501	0.9275	23.70	4.528	0.9521	24.69	18.96	0.9162	20.26
RH [25] (Re-trained)	3.781	0.9506	26.45	2.819	0.9559	26.73	8.333	0.9236	22.39
RHW [11]	10.90	0.9085	22.97	3.624	0.9586	25.28	26.18	0.9037	19.43
SFHR [28]	5.833	0.9349	24.66	4.185	0.9613	24.98	12.79	0.9155	21.41
Ours	1.311	0.9638	30.20	1.468	0.9814	30.27	0.729	0.9626	32.86

Comparison on Synthetic Data We first perform quantitative and qualitative comparisons on the testing dataset where we have ground truth images as a comparison. Tab. 1 shows quantitative comparison of intrinsic images decomposition performance and single-image human relighting quality. Our method outperforms the competitive methods on every metric for both tasks. To limit the comparison to decomposition and relighting quality only, all metrics are computed only on the foreground region for all methods.

Qualitative comparisons for relighting are shown in Fig. 5. To improve the visual effects, we use MODNet [26] to infer alpha channel of the input image and change the background of the relit images to the corresponding part of the environment lighting map. Benefiting from the combination of classical forward rendering and deep learning, our method can produce photorealistic shadows under arbitrary lighting conditions. As a comparison, RH [25], SFHR [28] and RHW [11] can only produce low-frequency shadows. Moreover, they also fail to relight images under outdoor lighting conditions and produce overly bright or dark relit results. As shown in Fig. 6, our method can produce plausible hard cast shadows and achieve better disentanglement between albedo and shading, which remain challenging for previous methods. Since all methods rely on albedo estimation at first for relighting, we also compare the de-lighting performance. As shown in Fig. 7, compared with RH [25], SFHR [28] and RHW [11], our method achieves better disentanglement between lighting and albedo and is able to remove large areas of shadows (even including self-shadows from concave regions on the real-world 3D shape of the human body, e.g., armpits, crotch, neck under the chin or folds on the clothing).

Comparison on Real-world Images Although our method is trained on synthetic datasets, it is generalizable to real data. The second column of Fig. 6 and Fig. 7 show shading estimation and de-lighting results on real images, respectively, and the second column of Fig. 5 shows the relit results of images photographed in the real world under arbitrary and complex illumination conditions. RH [25], SFHR [28] and RHW [11] still suffer from the entanglement of lighting and albedo and struggle to produce high-frequency shadows. By com-

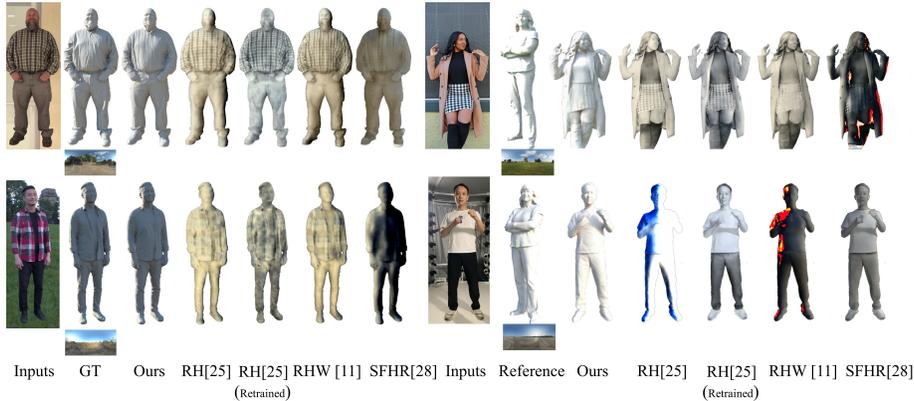


Fig. 6. Qualitative results for shading estimation under target lighting conditions. The first column: synthetic images from our testing dataset. The second column: real images.

Table 2. Quantitative results for ablation study of the full-body shading refinement.

	MSE \downarrow ($\times 10^{-3}$)	SSIM \uparrow	PSNR \uparrow
Raytracing	1.492	0.9499	29.84
Ours(w/o ambient)	1.250	0.9524	30.89
Ours(w normal)	0.886	0.9561	32.29
Ours	0.729	0.9626	32.86

parison, our method performs better at removing detailed shadows from the original images and generating photorealistic shadows on the relit images.

7.2 Ablation study

To demonstrate the effectiveness of our delighting network, full-body refinement network and face refinement network, we conduct comprehensive ablation studies both quantitatively and qualitatively.

First, we compare our delighting network with vanilla UNet, vanilla HRNet and vanilla HRNet with skip connections. Tab. 3 shows the quantitative results of different networks and Fig. 3 presents the de-lighting results on a real image. The vanilla HRNet is HRNet-W32 [49] has two transposed convolution layers with stride 2 to ensure that the output size and input size are the same. Based on the vanilla HRNet, the vanilla HRNet with skip connections further adds skip connections between the downsampled features of the first stage and the transposed convolution features of the output. “HRNet(w extra params)” means vanilla HRNet with skip connections, extra stages, modules and blocks. “Ours(5 stages)” indicates that the architecture of network is similar to that of our de-lighting network but only contains 5 stages. The proposed network “Ours” outperforms other networks on all metrics and achieves better disentanglement between lighting and albedo. By contrast, the vanilla HRNet fails to

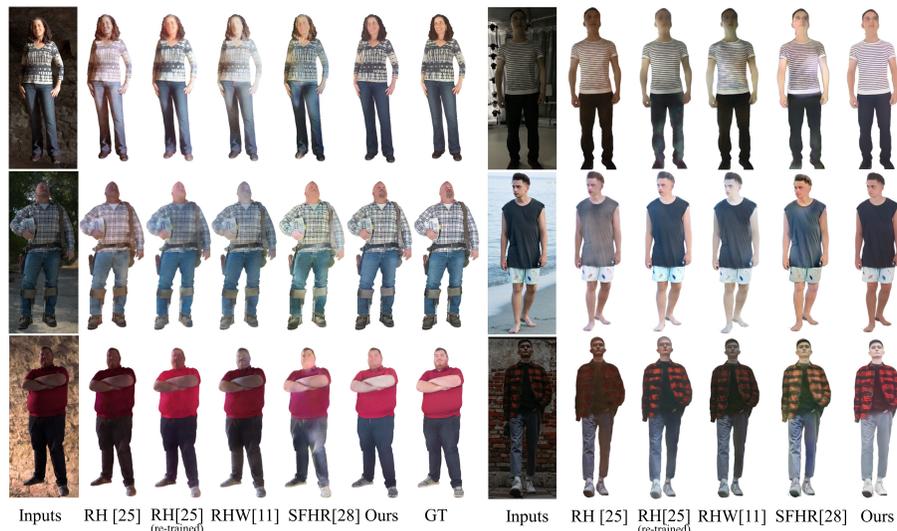


Fig. 7. De-lighting results on synthetic and real images. The first column: synthetic images from our testing dataset. The second column: real-world images.

Table 3. Quantitative results for ablation study of the de-lighting network. Vanilla Unet comes from RH [25] and is trained on our dataset.

	MSE \downarrow ($\times 10^{-3}$)	SSIM \uparrow	PSNR \uparrow
Vanilla Unet	2.819	0.9559	26.73
Vanilla HRNet	2.470	0.9550	27.21
Vanilla HRNet (w skip connections)	2.150	0.9767	28.53
Ours(5 stages)	1.728	0.9808	28.91
Ours	1.468	0.9814	30.27

produce high-resolution results and the removal of partial self-shadows such as those caused by clothing folds remains difficult, even with skip connections and extra parameters. Compared with UNet, our method greatly improves PSNR (increasing by 3) and MSE (dropping to half of the UNet’s).

Second, we verify the effectiveness of the ambient occlusion map used by the full-body refinement network. “Refinement Net(w/o ambient)” means that the refinement network only takes coarse full-body shading as input and removes SCM and FAM modules of MIMO-UNet[9]. “Refinement Net(w normal)” means that the refinement network takes the normal map as the auxiliary input rather than the ambient occlusion map. Tab. 2 shows quantitative results and Fig. 4 shows qualitative results. Without an ambient occlusion map, the refinement work cannot fill in missing geometry details and shadows. Moreover, when the input image is under extreme lighting conditions, the inference of normal map

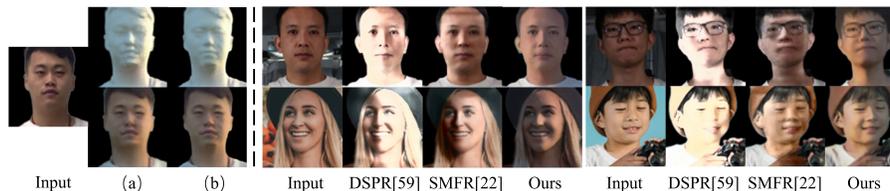


Fig. 8. Left: Ablation of face refinement module. (a) Top: cropped face from the refined full-body shading map Bottom: corresponding relit result (b) Top: refined face shading map Bottom: corresponding relit result. **Right:** Comparison with existing portrait relighting methods.

may fail around the boundary of the shadows. By contrast, the inferred ambient occlusion map is unaffected by shadows and able to recover more geometry and occlusion details, thus restoring better shading maps.

Finally, we evaluate the face refinement network. To highlight the role of this module, we present the cropped face shading for comparison. Fig. 8 shows the qualitative results for face refinement. The face without refinement contains unnatural facial details such as a twisted nose and asymmetric eyes. By contrast, thanks to the geometry priors provided by 3DMM templates, the refined face owns clearer and more natural facial features. Compared with DSPR [59] and SMFR [22], our method can generate plausible hard cast shadows, especially around the nose and neck, whereas SMFR [22] may produces patchy shadows and DSPR [59] may produces overexposed results.

8 Discussion

Conclusion. We propose a geometry-aware single-image human relighting framework that leverages 3D geometry prior information to produce higher-quality relit results. Our framework contains two stages: de-lighting and relighting. For the de-lighting stage, we use a modified HRNet as the de-lighting network and achieve better disentanglement between lighting and albedo. For the relighting stage, we use ray tracing to render the shading map of the target human, and further refine it utilizing learning-based refinement networks. The extensive results demonstrate that our framework can produce photorealistic high-frequency shadows with clear boundaries under challenging lighting conditions and outperforms the existing SOTA method on both synthetic images and real images.

Limitations. Due to the limitation of the dataset, we adopt the assumption of Lambertian materials for the clothed humans, which fails to produce specular reflectance in the relit results. For the same reason, our delighting network struggles to remove the highlights on the face. Moreover, inaccurate single-image geometry reconstruction may generate unnatural refined shading results.

Acknowledgement. This paper is supported by National Key R&D Program of China (2021ZD0113501) and the NSFC project No.62125107, No.62171255 and No.61827805.

References

1. <https://www.blender.org/>
2. <https://web.twindom.com/>
3. <https://polyhaven.com/hdri3>
4. Barron, J.T., Malik, J.: Color constancy, intrinsic images, and shape estimation. In: European Conference on Computer Vision. pp. 57–70. Springer (2012)
5. Barron, J.T., Malik, J.: Shape, illumination, and reflectance from shading. *IEEE transactions on pattern analysis and machine intelligence* **37**(8), 1670–1687 (2014)
6. Baslamisli, A.S., Le, H.A., Gevers, T.: Cnn based learning using reflection and retinex models for intrinsic image decomposition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6674–6683 (2018)
7. Bonneel, N., Kovacs, B., Paris, S., Bala, K.: Intrinsic decompositions for image editing. In: Computer Graphics Forum. vol. 36, pp. 593–609. Wiley Online Library (2017)
8. Chabert, C.F., Einarsson, P., Jones, A., Lamond, B., Ma, W.C., Sylwan, S., Hawkins, T., Debevec, P.: Relighting human locomotion with flowed reflectance fields. In: ACM SIGGRAPH 2006 Sketches, pp. 76–es (2006)
9. Cho, S.J., Ji, S.W., Hong, J.P., Jung, S.W., Ko, S.J.: Rethinking coarse-to-fine approach in single image deblurring. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4641–4650 (2021)
10. Christou, C.G., Koenderink, J.J.: Light source dependence in shape from shading. *Vision research* **37**(11), 1441–1449 (1997)
11. Daichi Tajima, Yoshihiro Kanamori, Y.E.: Relighting humans in the wild: Monocular full-body human relighting with domain adaptation. *Computer Graphics Forum (Proc. of Pacific Graphics 2021)* **40**(7), 205–216 (2021)
12. Debevec, P.: The light stages and their applications to photoreal digital actors. *SIGGRAPH Asia* **2**(4), 1–6 (2012)
13. Debevec, P., Hawkins, T., Tchou, C., Duiker, H.P., Sarokin, W., Sagar, M.: Acquiring the reflectance field of a human face. In: Proceedings of the 27th annual conference on Computer graphics and interactive techniques. pp. 145–156 (2000)
14. Debevec, P., Wenger, A., Tchou, C., Gardner, A., Waese, J., Hawkins, T.: A lighting reproduction approach to live-action compositing. *ACM Transactions on Graphics (TOG)* **21**(3), 547–556 (2002)
15. Ding, S., Sheng, B., Hou, X., Xie, Z., Ma, L.: Intrinsic image decomposition using multi-scale measurements and sparsity. In: Computer Graphics Forum. vol. 36, pp. 251–261. Wiley Online Library (2017)
16. Egger, B., Schönborn, S., Schneider, A., Kortylewski, A., Morel-Forster, A., Blumer, C., Vetter, T.: Occlusion-aware 3d morphable models and an illumination prior for face image analysis. *International Journal of Computer Vision* **126**(12), 1269–1287 (2018)
17. Gardner, M.A., Sunkavalli, K., Yumer, E., Shen, X., Gambaretto, E., Gagné, C., Lalonde, J.F.: Learning to predict indoor illumination from a single image. *ACM Transactions on Graphics (TOG)* **36**(6), 1–14 (2017)
18. Genova, K., Cole, F., Maschinot, A., Sarna, A., Vlasic, D., Freeman, W.T.: Unsupervised training for 3d morphable model regression. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 8377–8386 (2018)
19. Guo, J., Zhu, X., Yang, Y., Yang, F., Lei, Z., Li, S.Z.: Towards fast, accurate and stable 3d dense face alignment. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIX 16. pp. 152–168. Springer (2020)

20. Guo, K., Lincoln, P., Davidson, P., Busch, J., Yu, X., Whalen, M., Harvey, G., Orts-Escolano, S., Pandey, R., Dourgarian, J., et al.: The relightables: Volumetric performance capture of humans with realistic relighting. *ACM Transactions on Graphics (TOG)* **38**(6), 1–19 (2019)
21. Hawkins, T., Cohen, J., Debevec, P.: A photometric approach to digitizing cultural artifacts. In: *Proceedings of the 2001 conference on Virtual reality, archeology, and cultural heritage*. pp. 333–342 (2001)
22. Hou, A., Zhang, Z., Sarkis, M., Bi, N., Tong, Y., Liu, X.: Towards high fidelity face relighting with realistic shadows. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 14719–14728 (2021)
23. Imber, J., Guillemaut, J.Y., Hilton, A.: Intrinsic textures for relightable free-viewpoint video. In: *European Conference on Computer Vision*. pp. 392–407. Springer (2014)
24. Jafarian, Y., Park, H.S.: Learning high fidelity depths of dressed humans by watching social media dance videos. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 12753–12762 (2021)
25. Kanamori, Y., Endo, Y.: Relighting humans: occlusion-aware inverse rendering for full-body human images. *ACM Transactions on Graphics (TOG)* **37**(6), 1–11 (2018)
26. Ke, Z., Li, K., Zhou, Y., Wu, Q., Mao, X., Yan, Q., Lau, R.W.H.: Is a green screen really necessary for real-time portrait matting? (2020)
27. Laffont, P.Y., Bazin, J.C.: Intrinsic decomposition of image sequences from local temporal variations. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 433–441 (2015)
28. Lagunas, M., Sun, X., Yang, J., Villegas, R., Zhang, J., Shu, Z., Masiá, B., Gutierrez, D.: Single-image full-body human relighting. *CoRR* **abs/2107.07259** (2021), <https://arxiv.org/abs/2107.07259>
29. Land, E.H., McCann, J.J.: Lightness and retinex theory. *Josa* **61**(1), 1–11 (1971)
30. Li, C., Zhou, K., Wu, H.T., Lin, S.: Physically-based simulation of cosmetics via intrinsic image decomposition with facial priors. *IEEE transactions on pattern analysis and machine intelligence* **41**(6), 1455–1469 (2018)
31. Li, G., Wu, C., Stoll, C., Liu, Y., Varanasi, K., Dai, Q., Theobalt, C.: Capturing relightable human performances under general uncontrolled illumination. In: *Computer Graphics Forum*. vol. 32, pp. 275–284. Wiley Online Library (2013)
32. Li, Y., Liu, M.Y., Li, X., Yang, M.H., Kautz, J.: A closed-form solution to photorealistic image stylization. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 453–468 (2018)
33. Lin, J., Yuan, Y., Shao, T., Zhou, K.: Towards high-fidelity 3d face reconstruction from in-the-wild images using graph convolutional networks. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 5891–5900 (2020)
34. Lopez-Moreno, J., Hadap, S., Reinhard, E., Gutierrez, D.: Light source detection in photographs. In: *CEIG*. pp. 161–167 (2009)
35. Meka, A., Haene, C., Pandey, R., Zollhöfer, M., Fanello, S., Fyffe, G., Kowdle, A., Yu, X., Busch, J., Dourgarian, J., et al.: Deep reflectance fields: high-quality facial reflectance field inference from color gradient illumination. *ACM Transactions on Graphics (TOG)* **38**(4), 1–12 (2019)
36. Meka, A., Pandey, R., Haene, C., Orts-Escolano, S., Barnum, P., David-Son, P., Erickson, D., Zhang, Y., Taylor, J., Bouaziz, S., et al.: Deep relightable textures: volumetric performance capture with neural rendering. *ACM Transactions on Graphics (TOG)* **39**(6), 1–21 (2020)

37. Nagano, K., Luo, H., Wang, Z., Seo, J., Xing, J., Hu, L., Wei, L., Li, H.: Deep face normalization. *ACM Transactions on Graphics (TOG)* **38**(6), 1–16 (2019)
38. Nestmeyer, T., Lalonde, J.F., Matthews, I., Lehrmann, A.: Learning physics-guided face relighting under directional light. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 5124–5133 (2020)
39. Okatani, T., Deguchi, K.: Shape reconstruction from an endoscope image by shape from shading technique for a point light source at the projection center. *Computer vision and image understanding* **66**(2), 119–131 (1997)
40. Pandey, R., Escolano, S.O., Legendre, C., Haene, C., Bouaziz, S., Rhemann, C., Debevec, P., Fanello, S.: Total relighting: learning to relight portraits for background replacement. *ACM Transactions on Graphics (TOG)* **40**(4), 1–21 (2021)
41. Ramachandran, V.S.: Perception of shape from shading. *Nature* **331**(6152), 163–166 (1988)
42. Saito, S., Simon, T., Saragih, J., Joo, H.: 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*. pp. 84–93 (2020)
43. Sengupta, S., Kanazawa, A., Castillo, C.D., Jacobs, D.W.: Sfsnet: Learning shape, reflectance and illuminance of faces in the wild. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 6296–6305 (2018)
44. Shahlaei, D., Blanz, V.: Realistic inverse lighting from a single 2d image of a face, taken under unknown and complex lighting. In: *2015 11th IEEE international conference and workshops on automatic face and gesture recognition (FG)*. vol. 1, pp. 1–8. IEEE (2015)
45. Sheng, B., Li, P., Jin, Y., Tan, P., Lee, T.Y.: Intrinsic image decomposition with step and drift shading separation. *IEEE transactions on visualization and computer graphics* **26**(2), 1332–1346 (2018)
46. Shu, Z., Hadap, S., Shechtman, E., Sunkavalli, K., Paris, S., Samaras, D.: Portrait lighting transfer using a mass transport approach. *ACM Transactions on Graphics (TOG)* **36**(4), 1 (2017)
47. Shu, Z., Yumer, E., Hadap, S., Sunkavalli, K., Shechtman, E., Samaras, D.: Neural face editing with intrinsic image disentangling. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 5541–5550 (2017)
48. Sorkine, O.: Laplacian mesh processing. In: *Eurographics (State of the Art Reports)*. pp. 53–70. Citeseer (2005)
49. Sun, K., Xiao, B., Liu, D., Wang, J.: Deep high-resolution representation learning for human pose estimation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 5693–5703 (2019)
50. Sun, T., Barron, J.T., Tsai, Y.T., Xu, Z., Yu, X., Fyffe, G., Rhemann, C., Busch, J., Debevec, P.E., Ramamoorthi, R.: Single image portrait relighting. *ACM Trans. Graph.* **38**(4), 79–1 (2019)
51. Tewari, A., Zollhofer, M., Kim, H., Garrido, P., Bernard, F., Perez, P., Theobalt, C.: Mofa: Model-based deep convolutional face autoencoder for unsupervised monocular reconstruction. In: *Proceedings of the IEEE International Conference on Computer Vision Workshops*. pp. 1274–1283 (2017)
52. Wang, Z., Yu, X., Lu, M., Wang, Q., Qian, C., Xu, F.: Single image portrait relighting via explicit multiple reflectance channel modeling. *ACM Transactions on Graphics (TOG)* **39**(6), 1–13 (2020)
53. Wenger, A., Gardner, A., Tchou, C., Unger, J., Hawkins, T., Debevec, P.: Performance relighting and reflectance transformation with time-multiplexed illumination. *ACM Transactions on Graphics (TOG)* **24**(3), 756–764 (2005)

54. Weyrich, T., Matusik, W., Pfister, H., Bickel, B., Donner, C., Tu, C., McAndless, J., Lee, J., Ngan, A., Jensen, H.W., et al.: Analysis of human faces using a measurement-based skin reflectance model. *ACM Transactions on Graphics (ToG)* **25**(3), 1013–1024 (2006)
55. Whitted, T.: An improved illumination model for shaded display. In: *Proceedings of the 6th annual conference on Computer graphics and interactive techniques*. p. 14 (1979)
56. Ye, G., Garces, E., Liu, Y., Dai, Q., Gutierrez, D.: Intrinsic video and applications. *ACM Transactions on Graphics (ToG)* **33**(4), 1–11 (2014)
57. Zhang, L., Zhang, Q., Wu, M., Yu, J., Xu, L.: Neural video portrait relighting in real-time via consistency modeling. *arXiv preprint arXiv:2104.00484* (2021)
58. Zhang, X., Barron, J.T., Tsai, Y.T., Pandey, R., Zhang, X., Ng, R., Jacobs, D.E.: Portrait shadow manipulation. *ACM Transactions on Graphics (TOG)* **39**(4), 78–1 (2020)
59. Zhou, H., Hadap, S., Sunkavalli, K., Jacobs, D.W.: Deep single-image portrait relighting. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 7194–7202 (2019)