

Towards High Fidelity Face Relighting with Realistic Shadows

Andrew Hou¹, Ze Zhang¹, Michel Sarkis², Ning Bi², Yiying Tong¹, Xiaoming Liu¹
¹Michigan State University, ²Qualcomm Technologies Inc.

{houandr1, zhangze6, ytong, liuxm}@msu.edu, {msarkis, nbi}@qti.qualcomm.com

<https://github.com/andrewhou1/Shadow-Mask-Face-Relighting>




Figure 1: Given an input image and a target lighting, our model can produce a relit image with realistic hard cast shadows, such as those around the nose and eyes, which prior work struggles with. Images of SIPR [31] are provided by the authors.

Abstract

Existing face relighting methods often struggle with two problems: maintaining the local facial details of the subject and accurately removing and synthesizing shadows in the relit image, especially hard shadows. We propose a novel deep face relighting method that addresses both problems. Our method learns to predict the ratio (quotient) image between a source image and the target image with the desired lighting, allowing us to relight the image while maintaining the local facial details. During training, our model also learns to accurately modify shadows by using estimated shadow masks to emphasize on the high-contrast shadow borders. Furthermore, we introduce a method to use the shadow mask to estimate the ambient light intensity in an image, and are thus able to leverage multiple datasets during training with different global lighting intensities. With quantitative and qualitative evaluations on the Multi-PIE and FFHQ datasets, we demonstrate that our proposed method faithfully maintains the local facial details of the subject and can accurately handle hard shadows while achieving state-of-the-art face relighting performance.

1. Introduction

Face relighting is the problem of turning a source image of a human face into a new image of the same face under a desired illumination different from the original lighting. It has long been studied in computer vision and computer graphics and has a wide range of applications in face-related problems such as face recognition [10, 13, 23] and entertainment [22]. With the everlasting interest in consumer photography and photo editing, the ability to produce realistic relit face images will remain an important problem.

Many existing face relighting models utilize intrinsic decomposition of the image into face geometry, lighting, and reflectance [3, 5, 6, 13–16, 18, 20, 24, 25, 29, 32, 34, 36, 39]. The source image is then relit by rendering with a novel illumination. Other relighting methods employ image-to-image translation [20, 31, 41] or style transfer [17, 19, 27, 28].

For most face relighting applications, one important requirement is the preservation of the subject’s local facial details during relighting. Intrinsic decomposition methods often compromise high frequency details and can leave artifacts in the relit face images due to errors in the geometry or reflectance estimation. Another important feature of a

desirable face relighting model is proper shadow handling. For entertainment in particular, adding and removing shadows accurately is crucial in producing photorealistic results. Most existing relighting methods, however, do not model *hard self-cast shadows* caused by directional lights.

Our proposed method uses an hourglass network to formulate the relighting problem as a ratio (quotient) image [26] estimation problem. The ratio image estimated by our model can be multiplied with the source image to generate the target image under the new illumination. This approach allows our model to maintain the local facial details of the subject while adjusting the intensity of each pixel during relighting. We employ an estimation loss to enable ratio image learning, as well as a structural dissimilarity (DSSIM) loss based on the SSIM metric [37] to enhance the perceptual quality of relit faces. In addition, we incorporate PatchGAN [8] to further improve the plausibility.

During training, we generate and leverage shadow masks, which indicate estimated shadow regions for each image using the lighting direction and 3D shape from 3D Morphable Model (3DMM) [4] fitting. The shadow masks enable us to handle shadows through *weighted* ratio image estimation losses. We place higher emphasis on the pixels close to shadow borders in the source and target relighting images, with larger weights placed on borders of high-contrast cast shadows over soft ones. This simple strategy allows learning how to accurately add and remove both hard and soft shadows under various relighting scenarios.

Our training process can leverage images with both diffuse and directional lighting across multiple datasets, which improves our ability to handle diverse lighting and generalize to unseen data over methods that only train on a single dataset [20, 31, 41]. To enable this, we use our shadow masks to estimate the ambient lighting intensity in each image, and modify our lighting to account for differences in ambient lighting across images and datasets. Thus, our model accommodates for differences in the environment between images in controlled and in-the-wild settings.

Our proposed method has three main contributions:

- ◊ We propose a novel face relighting method that models both high-contrast cast shadows and soft shadows, while preserving the local facial details.
- ◊ Our technical approach involves single image based ratio image estimation to better preserve local details, shadow border reweighting to handle hard shadows, and ambient light compensation to account for dataset differences.
- ◊ Our approach achieves the state-of-the-art relighting results on two benchmarks quantitatively and qualitatively.

2. Related Work

Face Relighting. Among prior face relighting works, many relight via intrinsic decomposition and rendering [3, 5, 6, 13–

16, 18, 20, 24, 25, 29, 32–36, 39]: the source image is decomposed into face geometry, reflectance, and lighting, and render relit images via modified lighting. As decomposition generally relies heavily on single image face reconstruction, which remains an open problem, these methods tend to produce results that lack high frequency detail found in the source image and contain artifacts from geometry and reflectance estimation error. Our method bypasses this issue by avoiding intrinsic decomposition and estimating a ratio image instead. The ratio image only affects the intensity of each pixel in the source image in a smooth way aside from shadow borders, thus preserving the local facial details.

Other relighting methods do not perform explicit intrinsic decomposition. Sun *et al.* [31] use image-to-image translation to produce high quality relighting. However, their results deteriorate when the input image contains hard shadows or sharp specularities, or when presented with strong directional light. Zhou *et al.* [41] use a Spherical Harmonics (SH) lighting model and estimate the target luminance from the source luminance and a target lighting. However, their model is trained on images with primarily diffuse lighting and only handles soft shadows. We train our model on a mixture of images with diffuse and strong directional light. We also assign larger emphasis on learning the ratio image near the shadow borders of high-contrast cast shadows. Hence, our model handles shadows effectively.

Some methods relight using a style transfer approach [17, 19, 27, 28] by transferring the lighting conditions of a reference image as a style to a source image. However, since they require a reference image for each unique target lighting, they are less flexible to use in practice. Our model only requires a source image and a target lighting as input.

More recently, Zhang *et al.* [40] propose a method that can remove both foreign and facial shadows and can be considered a relighting method. While their method is excellent in removing facial shadows, it tends to not synthesize realistic hard shadows during relighting.

Ratio Images in Face Relighting. Prior face relighting works that incorporated ratio images often require multiple images as input [22, 26] or both source and target images [30], limiting their real-world applicability. Wen *et al.* [38] propose the first work on single image face relighting with ratio images, by estimating the ratio between the radiance environment maps. However, they use a fixed face geometry and only consider diffuse reflections. We instead estimate the ratio image between the source and target images and thus can directly produce non-diffuse relighting without resorting to a face mesh. Zhou *et al.* [41] use a ratio image based method to synthesize training data. Our work is the first ratio-image based face relighting method that can model non-diffuse relighting effects including shadows caused by strong directional lights, while requiring only one source image and a target lighting as input.

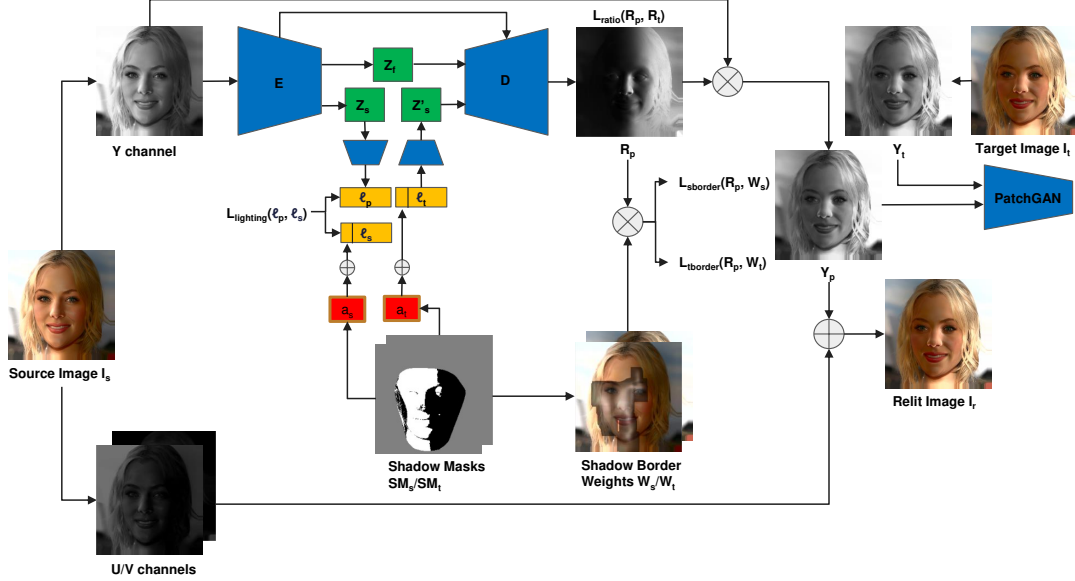


Figure 2: **Overview of our proposed method.** Given a source image I_s and a target lighting ℓ_t as input, our method estimates the source lighting ℓ_p and the ratio image \mathbf{R}_p in the luminance channel. In training, we assign high emphasis to learning the correct ratio image at pixels near shadow borders in both source and target images, helping the model add and remove shadows more accurately. We also estimate the ambient light intensity of the source image (\mathbf{a}_s) and of the target image (\mathbf{a}_t) and add them to the groundtruth source lighting ℓ_s and target lighting ℓ_t respectively to help our model adapt to images across different datasets with varying ambient light.

3. Proposed Method

3.1. Architecture

Our model takes a source image I_s and a target lighting ℓ_t as input and outputs the relit image I_r along with the estimated source lighting ℓ_p . Our lighting representation is the first 9 Spherical Harmonics (SH) coefficients. We adopt the hourglass network of [41], but rather than directly estimating the target luminance we estimate the ratio image \mathbf{R}_p between the input and target luminance. Similar to [41], our model only modifies the luminance channel: they use the *Lab* color space and estimate the target luminance before recombining with the source image’s *a* and *b* channels to generate the relit image, whereas we estimate \mathbf{R}_p for the *Y* channel of the *YUV* color space. The *Y* channel of the target image is computed by multiplying \mathbf{R}_p with the *Y* channel of the source image, which is then recombined with the *U* and *V* channels of the source image and converted to *RGB* to produce the final relit image (See Fig. 2).

3.2. Training Losses

We employ several loss functions to estimate ratio images that preserve high frequency details while capturing significant changes around shadow borders.

To directly supervise the ratio image learning, we employ the following ratio image estimation loss L_{ratio} :

$$L_{\text{ratio}} = \frac{1}{N} \|\log_{10}(\mathbf{R}_p) - \log_{10}(\mathbf{R}_t)\|_1. \quad (1)$$

Here, \mathbf{R}_p and \mathbf{R}_t are the predicted and ground truth ratio

images respectively, and N is the number of pixels in the image. Defining the loss in the log space ensures that ratio image values of r and $\frac{1}{r}$ receive equal weight in the loss.

We have two additional loss functions that place higher emphasis on the ratio image estimation near the shadow borders of the source and target images. The shadow border ratio image loss $L_{i,\text{border}}$ is defined as:

$$L_{i,\text{border}} = \frac{1}{N_i} \|\mathbf{W}_i \odot (\log_{10}(\mathbf{R}_p) - \log_{10}(\mathbf{R}_t))\|_1, \quad (2)$$

where i is s or t denoting the source or target respectively, and \odot is element-wise multiplication. \mathbf{W}_s and \mathbf{W}_t (See Sec. 3.3) are per-pixel weights that are element-wise multiplied with the per-pixel ratio image error, enabling our model to emphasize the ratio image estimation at or near shadow borders. N_s and N_t are the number of pixels with nonzero weights in \mathbf{W}_s and \mathbf{W}_t respectively. We denote the sum of L_{ratio} , $L_{s,\text{border}}$, and $L_{t,\text{border}}$ as L_{wratio} .

We also supervise the source lighting estimation using the loss term L_{lighting} defined as:

$$L_{\text{lighting}} = \|\ell_p - \ell_s\|_2^2, \quad (3)$$

where ℓ_p and ℓ_s are the predicted and ground truth source lighting, represented as the first 9 SH coefficients.

Similar to [41], we define a gradient consistency loss L_{gradient} to enforce that the image gradients of the predicted and target ratio images (\mathbf{R}_p and \mathbf{R}_t) are similar, and a face feature consistency loss L_{face} to ensure that images of the

same subject under different lighting conditions have the same face features. L_{gradient} preserves the image edges and avoids producing blurry relit images. L_{face} further preserves the local facial details of the subject during relighting.

To enhance the perceptual quality of our relit images, we employ two PatchGAN [8] discriminators: one operates on 70×70 and the other on 140×140 patches. We train the discriminators jointly with the hourglass network using the predicted luminance as fake samples and the target luminance as real samples. We denote this loss as $L_{\text{adversarial}}$.

Finally, similar to [20], we define a structural dissimilarity (DSSIM) loss L_{DSSIM} between \mathbf{R}_p and \mathbf{R}_t as:

$$L_{\text{DSSIM}} = \frac{1 - \text{SSIM}(\mathbf{R}_p, \mathbf{R}_t)}{2}. \quad (4)$$

Our final loss function L is the sum:

$$L = L_{\text{wratio}} + L_{\text{lighting}} + L_{\text{gradient}} + L_{\text{face}} + L_{\text{adversarial}} + L_{\text{DSSIM}}. \quad (5)$$

3.3. Shadow Border Reweighting

We reweight our ratio image estimation loss to assign larger weights to pixels near the shadow border. Our reweighting scheme incorporates two key ideas. First, higher-contrast hard shadows should have higher weight than lower-contrast soft shadows. Second, the reweighted neighborhood is larger for higher contrast shadows to provide a sufficiently wide buffer zone for handling hard shadows accurately using the ratio image. Since the contrast between the shadow and illuminated regions along the shadow border is high for hard shadows, larger weights along the border are needed for our model to learn this steep gradient.

The weights should peak at the shadow border and gradually decay as we move away from the border. To achieve this, we first compute a shadow mask for the image using the lighting direction and 3D shape estimated by 3DMM fitting [42, 43]. The shadow mask is a binary map that highlights shadow regions of the face, as detailed in Sec. 3.4.

We convolve the shadow mask with a 21×21 mean filter to get a smoothed mask $c(x, y)$. The smeared border can be found where the smoothed mask value is between two thresholds, $\tau_1 < c(x, y) < \tau_2$, as a border pixel’s neighborhood contains a mixture of pixels inside and outside shadows. For each border pixel, we then compute its local contrast. Instead of using the more costly root mean squared contrast in its neighborhood, we find the sum of absolute values of the four directional derivatives along the horizontal, vertical, and two diagonal directions efficient and effective. The derivatives are obtained through 21×21 filters, and their sum is stored as $t(x, y)$ for each border pixel (x, y) .

Given the local contrast $t(u, v)$ for each shadow border pixel (u, v) , we model its contribution to the shadow border weight at pixel (x, y) as a Gaussian $g_{(u,v)}(x, y)$, which




Figure 3: **Shadow Masks and Weights.** (a) input image, (b) shadow mask, (c) shadow border weights (intensity proportional to weights), (d) shadow border weights overlaid on the input image. For each image, we use the source lighting direction and estimated 3D shape to generate shadow masks indicating the shadowed regions. With the shadow mask, we assign high weights to learning the ratio image along the shadow border. Note that our reweighting scheme appropriately assigns high weights to high contrast shadows including hard cast shadows around the nose.

peaks at the center of the border and smoothly decays with distance away from the border:

$$g_{(u,v)}(x, y) = \frac{1}{\sigma_{max} \sqrt{2\pi}} e^{-\frac{(c(x,y) - \mu_c)^2}{2\sigma_{t(u,v)}^2}}. \quad (6)$$

Here, μ_c is the mean of $c(x, y)$ and $\sigma_{t(u,v)} = \frac{t(u,v)}{t_{max}} \sigma_{max}$ is the deviation controlling the decay speed, where t_{max} is the highest contrast in the image, and σ_{max} is the preset maximum deviation. We tally the contribution of $g_{(u,v)}$ to pixel (x, y) for every border pixel (u, v) whose neighborhood $N(u, v)$ contains (x, y) with the neighborhood size proportional to $t(u, v)$, and obtain the shadow border weight:

$$\mathbf{W}(x, y) = \sum_{N(u,v) \ni (x,y)} g_{(u,v)}(x, y). \quad (7)$$

As a final step, we linearly normalize \mathbf{W} to the range $[0, 10]$.

Note that the amplitudes and standard deviations of $g_{(u,v)}$ are larger for high contrast shadow borders (hard shadows) and smaller for low contrast shadows (soft shadows). Our reweighting thus assigns the highest weights to hard shadows and their neighborhoods (See Fig. 3) — crucial in handling directional-light-caused shadows.

During training, we apply this reweighting scheme to both the source and target image to generate \mathbf{W}_s and \mathbf{W}_t respectively. This allows us to place higher emphasis on any shadow border regions in either the source or target image when predicting the ratio image.




Figure 4: **Ambient Light Distributions.** We estimate the ambient light of our training images using shadow masks. Clearly, there is a major difference between the ambient light distributions of the images in the Extended Yale Face Database B [7] and the images in DPR [41], and also a wide distribution within each dataset.

3.4. Shadow Masks

We create shadow masks for all training images using the lighting direction and the 3D face mesh offered by 3DMM fitting [42, 43]. We estimate the rotation, translation, and scale of the shape and align the 3D face to the 2D image.

To generate the shadow mask, we cast parallel rays along the z -axis towards the mesh. If the ray hits the face, we check for two kinds of shadows at the intersection point: self shadows and cast shadows. Portions of the face between the light source and the intersection point will block the light and cast a shadow. To determine if the point lies in a cast shadow, we cast a “shadow feeler” ray from the intersection point to the light source [2]. If the ray hits the surface, the intersection point is in a cast shadow. To determine if the point lies in self shadow, we compute the angle between the light source’s direction and the surface normal of the intersection point. If the angle is obtuse, the light is from the back of the surface and the point is in self shadow. If the intersection point is either in a cast shadow or self shadow, we assign 0 to the corresponding mask pixel. Otherwise, the pixel is illuminated and is 1 in the shadow mask.

3.5. Estimating Ambient Light

Since ambient light is inconsistent across images, especially between light stage and in-the-wild images, we introduce a method to estimate ambient light using the shadow mask. Since only ambient light contributes to shadowed regions, we use the average intensity of the shadow pixels as an estimate of the ambient light intensity in the image.

To sum the contributions of directional and ambient light, we first model each image’s directional light as a point light and estimate the corresponding 9 SH coefficients. We then add the estimated ambient light intensity to the 0th order SH coefficient, which represents overall light intensity.

During training, we use the source image’s shadow mask to estimate its ambient light intensity and add that to groundtruth source lighting ℓ_s . We do the same using the target image’s shadow mask for target lighting ℓ_t .

3.6. Implementation Details

We train our model in Tensorflow [1] for 51 epochs using a learning rate of 0.0001 and the Adam optimizer [12] with default parameters. Similar to [41], we initially train for 8 epochs with no skip connections, and then add them one at a time at epochs 9, 11, 13, and 15. We only apply L_{face} after epoch 16, whereas all other losses are always active.

Since cameras often do not correctly capture luminance in the physical world, we also apply gamma correction to the input and target luminance during training ($\gamma = \frac{1}{2.2}$) and learn the ratio image in this corrected space.

4. Experiments

4.1. Training Data

We train our model using images from two datasets: one with mostly diffuse lighting and one with strong directional lights. Our diffuse lighting dataset is the Deep Portrait Relighting (DPR) dataset [41], where we use the same training images as [41]. It contains 138,135 images of 27,627 subjects with 5 lightings each. Our dataset with strong directional lighting is the Extended Yale Face Database B [7], which contains 16,380 images of 28 subjects with 9 poses and 65 illuminations (64 distinct lighting directions and 1 ambient lighting). To compensate for the large differences in ambient lighting intensities between the two datasets and within each dataset (See Fig. 4), we use our method in Sec. 3.5 to produce the groundtruth source and target lightings. All quantitative and qualitative evaluations are performed using this model trained on these two databases. For training and for all evaluations, we use 256×256 images.

4.2. Quantitative Evaluations

We compare our model’s relighting performance with prior work on the Multi-PIE dataset [9], where each subject is illuminated under 19 lighting conditions (18 images with directional lights, one image with only ambient lighting).

Relighting Using Target Lightings. We compare against prior relighting methods [20, 24, 41] that can accept a target lighting as input. For each Multi-PIE subject and each session, we randomly select one of the 19 images as the source image and one as the target image, which serves as the relighting groundtruth. The target image’s lighting is then used to relight the source image. This leads to a total of 921 relit images. For prior methods, we use the pretrained models provided by the authors. We evaluate the performance using 3 error metrics: Si-MSE [41], MSE, and DSSIM, as shown in Tab. 1. Unlike traditional MSE, Si-MSE computes the MSE between the target image and the relit image multiplied by the scale that yields the minimum error for the image. This is meant to resolve scale ambiguities in the relit image caused by the target lighting. We don’t compare quantitatively with SIPR [31] since their method masks




Figure 5: **Relighting on Multi-PIE Using Target Lighting.** All methods take a source image and a target lighting as input. Images for SIPR [31] are provided by the authors. Notice that our model produces significantly better cast shadows, especially around the nose.




Figure 6: **Lighting Transfer on Multi-PIE.** All methods estimate the target lighting from the reference. Images of [28] are provided by the authors. Our model produces the correct lightings, whereas [27] does not. Unlike ours, both baselines strongly alter the skin tones.

Method	Si-MSE	MSE	DSSIM
SfsNet [24]	0.0545 ± 0.0201	0.1330 ± 0.0531	0.3151 ± 0.0523
DPR [41]	0.0282 ± 0.0125	0.0702 ± 0.0361	0.1818 ± 0.0490
Nestmeyer [20]	0.0484 ± 0.0318	0.0584 ± 0.0335	0.2722 ± 0.0950
Our model	0.0220 ± 0.0078	0.0292 ± 0.0148	0.1605 ± 0.0487

Table 1: **Relighting Results (mean ± standard deviation) Using Target Lighting on Multi-PIE.** All methods take a source image and a target lighting as input. Our method outperforms prior works across all metrics, and also has the lowest standard deviation.

out the background. Most notably, our model outperforms DPR [41] while using their model architecture, showing the benefits of the ratio image and our shadow border losses.

Relighting Using Reference Images. We also compare with relighting methods that require both a source and a reference image as input [27, 28], which relight by transferring the reference’s lighting to the source. For each Multi-PIE image, we randomly select a reference image across all subjects of the dataset and estimate the target lighting from the reference. We use the available code to compare against Shih *et al.* [27] and the results for Shu *et al.* [28] are provided by the authors. The evaluation is shown in Tab. 2.

Facial Detail Preservation. To compare our model’s ability to preserve the subject’s facial details during relighting


Method	Si-MSE	MSE	DSSIM
Shih [27]	0.0374 ± 0.0190	0.0455 ± 0.0203	0.2260 ± 0.0443
Shu [28]	0.0162 ± 0.0102	0.0243 ± 0.0170	0.1383 ± 0.0499
Our model	0.0148 ± 0.0096	0.0204 ± 0.0153	0.1150 ± 0.0404

Table 2: **Lighting Transfer Results (mean ± standard deviation) on Multi-PIE.** Each input image is assigned a random reference image. Our model outperforms both methods in all metrics.

with prior work, we compute the average cosine similarity of the VGG-Face [21] features of the relit and groundtruth Multi-PIE images across different layers of VGG-Face. In particular, we use the layer before each max-pooling layer in VGG-Face. As shown in Fig. 8, our model achieves noticeably higher cosine similarity than prior work in earlier layers and is comparable or slightly better in later layers, indicating that ours is better at preserving local facial details.

4.3. Qualitative Evaluations

We qualitatively compare relighting results on Multi-PIE [9] and FFHQ [11]. On Multi-PIE, we include relit images produced from target lightings (See Fig. 5) and from lighting transfer (See Fig. 6). When applying target lighting, our model is able to produce noticeably better cast shadows than DPR [41], SIPR [31], SfsNet [24], and



(a) Source (b) Target (c) Our model (d) DPR [41] (e) SIPR [31] (f) SfSNet [24]

Figure 7: **Relighting on FFHQ dataset Using Target Lighting.** Images for SIPR [31] are provided by the authors. Across all lighting conditions, our model produces better cast shadows than prior work, especially around the nose and eyes.





Figure 8: **Facial Detail Preservation on Multi-PIE.** We compute the cosine similarity between the VGG-Face [21] features of each method’s relit images and groundtruth target images to measure their ability to preserve local and global facial details during relighting. Our method is consistently the best at preserving local details based on the cosine similarity in the *earlier* layers.

Nestmeyer *et al.* [20]. Our model also avoids overlighting the image, whereas [41] often produces images that appear overexposed. When performing lighting transfer, Shih *et al.* [27] is unable to produce the correct lighting whereas ours can transfer the correct target lighting. Both baselines also greatly alter the subjects’ skin tones, whereas our model does not.

On FFHQ, we perform relighting using target lighting (See Fig. 7). Our approach handles cast shadows better than prior work, as seen by the shadows around the nose, eyes,



(a) Input (b) Target (c) No $L_{s/tborder}$ (d) Proposed

Figure 9: **Shadow Border Ablation.** The shadow border losses improve the quality of cast shadows, especially around the nose.

and other regions, while also maintaining similar or better visual quality in the relit images. When evaluating Nestmeyer *et al.* [20], we find that it does not generalize well to the in-the-wild images in FFHQ due to their training dataset and thus we do not include it in Fig. 7. We discuss this further in the supplementary materials.

4.4. Ablation Studies

Shadow Border Reweighting. To demonstrate the effectiveness of our reweighting of the ratio image estimation loss, we train an additional model without $L_{sborder}$ and $L_{tborder}$. Our proposed model is able to produce signifi-

Method	Si-MSE	MSE	DSSIM
No shadow border	0.0212	0.0290	0.1528
No ratio image	0.0230	0.0343	0.1617
Proposed	0.0220	0.0292	0.1605

Table 3: **Ablation on Multi-PIE.** We ablate on our Multi-PIE test set against our model with no shadow border losses and our model that estimates the luminance directly rather than the ratio image.

cantly better cast shadows, especially around the nose region, whereas the model trained without $L_{\text{shadow border}}$ and L_{border} often cannot produce cast shadows (see Fig. 9). This highlights the contribution of our shadow border losses. We also quantitatively evaluate both models on Multi-PIE, as shown in Tab. 3. The model with no shadow border weights appears to slightly outperform the proposed model quantitatively, however we’ve shown in Fig. 9 that it cannot produce cast shadows. The small difference is likely due to the emphasis that the border weights place on the shadow borders over other pixels in the image, resulting in a tradeoff between high quality cast shadows and overall visual quality.

Ratio Image Estimation. We perform a second ablation to show the benefits of modeling the relighting problem as a ratio image estimation problem. We train a second model that estimates the luminance directly rather than the ratio image. As shown in Fig. 10, our proposed model is able to better preserve the local facial details of the subject in several regions, particularly around the mouth. We also quantitatively evaluate both models on Multi-PIE, as shown in Tab. 3. Our proposed model performs better across all metrics. This validates that estimating the ratio image rather than the luminance improves the fidelity of the relit images.

Ambient Lighting Estimation. We ablate the benefits of modeling the different ambient lighting intensities across different datasets, which enables our model to train with multiple datasets. We train an additional model that does not account for the ambient lighting differences between our two datasets and instead only uses the SH lighting estimate from the point light source. Our proposed model produces more accurate global lighting intensities on Multi-PIE, especially around the face, showing the benefit of our ambient lighting estimation (See Fig. 11). To demonstrate the performance quantitatively, we compute the MSE in the face region using portrait masks. Our proposed model achieves a lower MSE of 0.0233 compared to our model trained with no ambient lighting estimation, which achieves 0.0242.

5. Conclusion

We have proposed a novel face relighting method that creates realistic hard shadows. We introduce the idea of improving shadow modeling by increasing the emphasis on correctly relighting the pixels near shadow borders, and provide an algorithm to locate the borders and reweight




Figure 10: **Ratio Image Ablation.** d) better preserves the local facial details of the subject, most notably around the mouth. In the case of the bottom subject, c) seems to be gazing to the left instead of forward which is fixed by d). Best viewed if enlarged.




Figure 11: **Ambient Light Ablation.** d) produces a more accurate global lighting intensity than c), especially around the face, which shows the benefit of our ambient light estimation method.

losses around them using shadow masks. We also introduce the first ratio-image based face relighting method that both models the effects of strong directional lights and requires only a single source image and a target lighting as input, and we demonstrate our model’s ability to better preserve the local facial details of the subject as a result of using the ratio image. We further introduce a method to estimate the ambient lighting intensity using an image’s shadow mask, allowing our model to account for differences in environmental lighting across datasets. Our model can thus benefit from multiple datasets during training, which improves its generalizability to unseen data. We show that our shadow border reweighting scheme, our use of the ratio image, and our ambient light estimation all improve the relighting performance. We hope that this work will inspire future relighting methods, especially with respect to shadow handling.

Acknowledgement This work was partially funded by Qualcomm Technologies Inc. The authors thank Dr. Zhixin Shu and Tiancheng Sun for kindly providing the results for their methods.

References

- [1] Martin Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, et al. Tensorflow: A system for large-scale machine learning. In *OSDI*, 2016. 5
- [2] Arthur Appel. Some techniques for shading machine renderings of solids. In *AFIPS*, 1968. 5
- [3] Jonathan T Barron and Jitendra Malik. Shape, illumination, and reflectance from shading. *PAMI*, 2015. 1, 2
- [4] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3D faces. In *SIGGRAPH*, 1999. 2
- [5] Bernhard Egger, Sandro Schonborn, Andreas Schneider, Adam Kortylewski, Andreas Morel-Forseter, Clemens Blumer, and Thomas Vetter. Occlusion-aware 3D morphable models and an illumination prior for face image analysis. *IJCV*, 2018. 1, 2
- [6] Kyle Genova, Forrester Cole, Aaron Maschinot, Aaron Sarna, Daniel Vlasic, and William T. Freeman. Unsupervised training for 3D morphable model regression. In *CVPR*, 2018. 1, 2
- [7] Athinodoros Georgiades, Peter Belhumeur, and David Kriegman. From Few to Many: Illumination Cone Models for Face Recognition under Variable Lighting and Pose. *PAMI*, 2001. 5
- [8] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NeurIPS*, 2014. 2, 4
- [9] Ralph Gross, Iain Matthews, Jeffrey Cohn, Takeo Kanade, and Simon Baker. Multi-PIE. *Image and Vision Computing*, 2010. 5, 6
- [10] Yuge Huang, Pengcheng Shen, Ying Tai, Shaoxin Li, Xiaoming Liu, Jilin Li, Feiyue Huang, and Rongrong Ji. Improving face recognition from hard samples via distribution distillation loss. In *ECCV*, 2020. 1
- [11] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, 2019. 6
- [12] Diederik Kingma and Jimmy Ba. Adam. A method for stochastic optimization. In *ICLR*, 2014. 5
- [13] Ha Le and Ioannis Kakadiaris. Illumination-invariant face recognition with deep relit face images. In *WACV*, 2019. 1, 2
- [14] Gun-Hee Lee and Seong-Whan Lee. Uncertainty-Aware Mesh Decoder for High Fidelity 3D Face Reconstruction. In *CVPR*, 2020. 1, 2
- [15] Jinho Lee, Raghu Machiraju, Baback Moghaddam, and Hanspeter Pfister. Estimation of 3D faces and illumination from single photographs using a bilinear illumination model. In *EGSR*, 2005. 1, 2
- [16] Chen Li, Kun Zhou, and Stephen Lin. Intrinsic face image decomposition with human face priors. In *ECCV*, 2014. 1, 2
- [17] Yijun Li, Ming-Yu Liu, Xueting Li, Ming-Hsuan Yang, and Jan Kautz. A closed-form solution to photorealistic image stylization. In *ECCV*, 2018. 1, 2
- [18] Jiangke Lin, Yi Yuan, Tianjia Shao, and Kun Zhou. Towards High-Fidelity 3D Face Reconstruction from In-the-Wild Images Using Graph Convolutional Networks. In *CVPR*, 2020. 1, 2
- [19] Fujun Luan, Sylvain Paris, Eli Shechtman, and Kavita Bala. Deep photo style transfer. In *CVPR*, 2017. 1, 2
- [20] Thomas Nestmeyer, Jean-Francois Lalonde, Iain Matthews, and Andreas Lehrmann. Learning Physics-guided Face Relighting under Directional Light. In *CVPR*, 2020. 1, 2, 4, 5, 6, 7
- [21] O.M. Parkhi, A. Vedaldi, and A. Zisserman. Deep face recognition. In *BMVC*, 2015. 6, 7
- [22] Pieter Peers, Naoki Tamura, Wojciech Matusik, and Paul Debevec. Post-production facial performance relighting using reflectance transfer. In *SIGGRAPH*, 2007. 1, 2
- [23] Laiyun Qing, Shiguang Shan, and Xilin Chen. Face relighting for face recognition under generic illumination. In *ICASSP*, 2004. 1
- [24] Soumyadip Sengupta, Angjoo Kanazawa, Carlos D. Castillo, and David W. Jacobs. SfSNet: Learning shape, reflectance and illuminance of faces in the wild. In *CVPR*, 2018. 1, 2, 5, 6, 7
- [25] Davoud Shahlai and Volker Blanz. Realistic inverse lighting from a single 2D image of a face, taken under unknown and complex lighting. In *FG*, 2015. 1, 2
- [26] Amnon Shashua and Tammy Riklin-Raviv. The quotient image: Class-based re-rendering and recognition with varying illuminations. *PAMI*, 2001. 2
- [27] YiChang Shih, Sylvain Paris, Connelly Barnes, William T. Freeman, and Fredo Durand. Style transfer for headshot portraits. In *SIGGRAPH*, 2014. 1, 2, 6, 7
- [28] Zhixin Shu, Sunil Hadap, Eli Shechtman, Kalyan Sunkavalli, Sylvain Paris, and Dimitris Samaras. Portrait lighting transfer using a mass transport approach. *TOG*, 2017. 1, 2, 6
- [29] Zhixin Shu, Ersin Yumer, Sunil Hadap, Kalyan Sunkavalli, Eli Shechtman, and Dimitris Samaras. Neural face editing with intrinsic image disentangling. In *CVPR*, 2017. 1, 2
- [30] Arne Stoschek. Image-based re-rendering of faces for continuous pose and illumination directions. In *CVPR*, 2000. 2
- [31] Tiancheng Sun, Jonathan T Barron, Yun-Ta Tsai, Zexiang Xu, Xueming Yu, Graham Fyffe, Christoph Rhemann, Jay Busch, Paul Debevec, and Ravi Ramamoorthi. Single image portrait relighting. In *SIGGRAPH*, 2019. 1, 2, 5, 6, 7
- [32] Ayush Tewari, Michael Zollhofer, Hyeonwoo Kim, Pablo Garrido, Florian Bernard, Patrick Perez, and Theobalt Christian. Mofa: Model-based Deep Convolutional Face Autoencoder for Unsupervised Monocular Reconstruction. In *ICCV*, 2017. 1, 2
- [33] Luan Tran, Feng Liu, and Xiaoming Liu. Towards high-fidelity nonlinear 3D face morphable model. In *CVPR*, 2019. 2
- [34] Luan Tran and Xiaoming Liu. Nonlinear 3D face morphable model. In *CVPR*, 2018. 1, 2
- [35] Luan Tran and Xiaoming Liu. On learning 3D face morphable model from in-the-wild images. *PAMI*, 2019. 2
- [36] Yang Wang, Lei Zhang, Zicheng Liu, Gang Hua, Zhen Wen, Zhengyou Zhang, and Dimitris Samaras. Face relighting from a single image under arbitrary unknown lighting conditions. *PAMI*, 2009. 1, 2

- [37] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *TIP*, 2004. [2](#)
- [38] Zhen Wen, Zicheng Liu, and Tomas Huang. Face Relighting with Radiance Environment Maps. In *CVPR*, 2003. [2](#)
- [39] Shuco 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. In *SIGGRAPH*, 2018. [1](#), [2](#)
- [40] Xuaner Zhang, Jonathan T. Barron, Yun-Ta Tsai, Rohit Pandey, Xiuming Zhang, Ren Ng, and David E. Jacobs. Portrait Shadow Manipulation. *TOG*, 2020. [2](#)
- [41] Hao Zhou, Sunil Hadap, Kalyan Sunkavalli, and David W Jacobs. Deep single-image portrait relighting. In *ICCV*, 2019. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#)
- [42] Xiangyu Zhu, Zhen Lei, Xiaoming Liu, Hailin Shi, and Stan Z. Li. Face Alignment Across Large Poses: A 3D Solution. In *CVPR*, 2017. [4](#), [5](#)
- [43] Xiangyu Zhu, Xiaoming Liu, Zhen Lei, and Stan Li. Face alignment in full pose range: A 3D total solution. *PAMI*, 2017. [4](#), [5](#)