

# A Visual Representation for Editing Face Images

Jiajun Lu<sup>1</sup>, Kalyan Sunkavalli<sup>2</sup>, Nathan Carr<sup>2</sup>, Sunil Hadap<sup>2</sup>, David Forsyth<sup>1</sup>

<sup>1</sup>University of Illinois at Urbana Champaign, <sup>2</sup>Adobe Research  
<sup>1</sup>{jlu23, daf}@illinois.edu, <sup>2</sup>{sunkaval, ncarr, hadap}@adobe.com

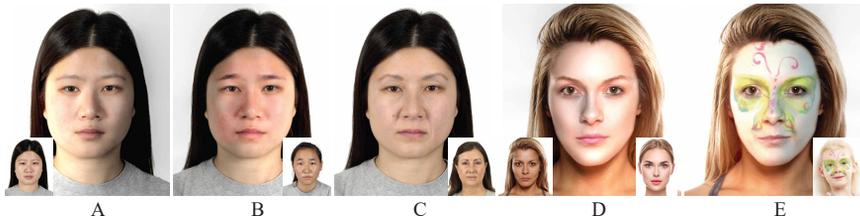
**Abstract.** We propose a new approach for editing face images, which enables numerous exciting applications including face relighting, makeup transfer and face detail editing. Our face edits are based on a visual representation, which includes geometry, face segmentation, albedo, illumination and detail map. To recover our visual representation, we start by estimating geometry using a morphable face model, then decompose the face image to recover the albedo, and then shade the geometry with the albedo and illumination. The residual between our shaded geometry and the input image produces our detail map, which carries high frequency information that is either insufficiently or incorrectly captured by our shading process. By manipulating the detail map, we can edit face images with reality and identity preserved. Our representation allows various applications. First, it allows a user to directly manipulate various illumination. Second, it allows non-parametric makeup transfer with input face’s distinctive identity features preserved. Third, it allows non-parametric modifications to the face appearance by transferring details. For face relighting and detail editing, we evaluate via a user study and our method outperforms other methods. For makeup transfer, we evaluate via an online attractiveness evaluation system, and can reliably make people look younger and more attractive. We also show extensive qualitative comparisons to existing methods, and have significant improvements over previous techniques.

**Keywords:** face visual representation, face editing, face relighting, makeup, appearance editing, transfer, detail map

## 1 Introduction

Post-capture editing of lighting, makeup and face appearance is an important component of photographic retouching, visual effects, and image compositing. However, it is a very challenging problem because realistic editing usually requires precise estimates of scene geometry, material properties, and lighting. This is especially true for faces where even small artifacts lead to perceptually implausible results.

Our representation exposes natural “hooks” for image editing. We align a morphable face model to the face in an image, then recover albedo and shading maps for the image, which are used in conjunction with the registered face model to estimate scene illumination. We render the appearance of the face given the



**Fig. 1.** Our paper proposes a visual representation to automatically relight a single face image with face details preserved (A), edit face detail by transferring detail maps from a different subject (B,C), and transfer face makeup from a reference (lower right inset) image to an input (lower left inset) image (D,E). Our method keeps the image’s realism during all three operations, and preserves the input face’s identity (detail around nose, eyes and mouth) while transferring makeup. *Best viewed at high resolution, in color.*

estimated geometry, lighting, and albedo using Lambertian shading. However, current technology does not allow recovering a face geometry up to the scale of wrinkles, pores, etc. These small details are important parts of what make a face individual (for example, careful preservation of detail around eyes, nose and mouth ensures that changing face makeup by editing albedo does not change the input’s identity). Thus we compute a detail map containing image information not represented by recovered albedo, shading and geometry. The idea behind our method is that editing operations enabled by such a detail map can be made reasonably robust against the accuracy of the models used to produce the detail map. We demonstrate three types of face editing tasks: relighting, makeup transfer (we take makeup from a reference and place it on an input) and detail editing (we change surface detail on a face), see Fig. 1.

We evaluate our method by: qualitative evaluations of result images; qualitative comparisons to the results of existing methods; quantitative results of a user study, where we compare the ability of persons to distinguish between edited faces and untouched images; quantitative results of an online system to evaluate the effect of makeup on age and attractiveness. These experiments show that our technique is able to produce results that have significant improvements over previous techniques in terms of perceptual realism. For relighting and detail transfer, our results are hardest to distinguish. For makeup transfer, we can reasonably make people look younger and more attractive.

## Contributions:

1. We propose an effective visual representation for face (with face detail maps related to face realism and face identity) that can realistically and reliably perform multiple tasks automatically, including face relighting, makeup transfer and face detail editing. Previous methods are usually effective only on a single task.
2. Our model is the first model that allows interactive relighting of a whole single face image. We propose the first visual quality based quantitative way to evaluate face relighting results, and our method produces state of the art face relighting results.

3. Our makeup transfer preserves the person’s identity and creates realistic results with various makeup images ( such as normal makeup, unusual makeup, etc), which previous methods have problems dealing with. Even for normal makeup, we create state of the art results.
4. Our model also enables us to edit face details to achieve various effects, such as making the person look younger or older, changing skin properties and fine-scale geometric details. By grouping and choosing appropriate source images, we can control tons of editing operations to create desired images. Compared to previous method, we support a larger variety of tasks, and perform as well or better.

## 2 Background

**Face relighting:** Ratio image [24] [30] is a standard method to relight face images. Numerous variants are available to: improve performance under extreme lighting conditions [34]; sew together local ratio estimates to resolve difficulties created by small errors in registration [7]; relax the requirement that reference images are registered [32]; allow expression transfer [23]; relight temporal sequences, using a large database of reference images [25].

The requirement for two reference face images can be relaxed in ratio image methods. Wen et al. [36] describe a method that replaces the reference face images with reference spheres. Chen et al. [7] describe a method that requires one reference face, in the target illumination. Li et al. [16] estimate low spatial frequency components using logarithmic total variation. Wang et al. [35] relight single images by recovering face geometry. Each method is demonstrated only on masked faces. Chai et al. [6] focus on accurate geometric reconstruction of face and hair, which is used to relight portraits. In contrast to our method, there is no explicit detail map representation that captures rendering errors for the inferred model.

For the methods described, evaluation is by qualitative comparison, and (for [35] [34]) by quantitative studies of face recognition results (the percentage recognition accuracy improved by relighting). In contrast, our quantitative evaluation tests whether people are fooled into thinking that relit faces are real images.

**Makeup transfer:** Analyzing and modifying images to improve perceived beauty is an active area (review in [22]). A variety of methods allow interactively applying makeup to the input face. By recovering and then editing multiple reflectance layers, makeup can be successfully simulated [15]. TAAZ allows users to apply chosen makeup components to face images interactively ([www.taaz.com](http://www.taaz.com)). In contrast, our method allows automatic transfer of makeup from reference to input image.

Automated makeup transfer methods typically use 2D face models. Tong et al. [33] transfer makeup using a ratio image. Guo et al. [9] transfer makeup by decomposing images into several layers, then blend layers and recomposite.

Their detail layer is obtained by image filtering. Liu et al. [21] register facial components independently, and synthesize made-up faces using database. Liang et al. [17] enhance facial skin by suppressing a detail layer, estimated using image filtering. Shih et al. [31] adjust face image subband energies to match the makeup. Scherbaum et al. [29] suggest makeup for a 3D morphable face model based on a collection of 3D example morphable face models. In contrast, our detail layer is obtained from rendering residual, and contains material and identity information; we only need 2D images and allow relighting; we offer various makeup types and larger number of evaluation images.

**Appearance editing:** Face expression editing is possible using ratio images, and produces good results [23]. Bitouk et al. achieve expression and illumination transfer by replacing a face image with a similar image (with appropriate expression and illumination parameters) from an enormous collection [2]. Yang et al. warp local face components from other images to transfer expression [37]. Kemelmacher-Shlizerman et al. combine models of illumination subspaces of faces at various ages [12]. In contrast, our editing is achieved by warping and applying detail maps.

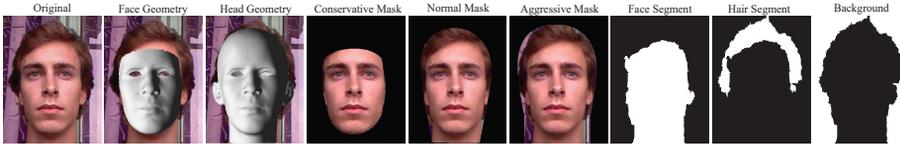
**Detail maps:** Fine relief on surfaces creates small shadows that appear to be strong material cues to human viewers (eg. [26]). Representations of this detail can be used to slightly improve albedo estimates and material classifiers [19]. Liao et al. [18] demonstrated that residual map can help relight some objects, because convincing material appearance is more important to human subjects than physical accurate shading fields. Boyadzhev et al. [4] edit surface material (e.g. skin oilness or blemishes on faces) by adjusting selected image bands typically at high spatial frequencies. In contrast, for relighting we preserve those bands and adjust low spatial frequencies; for makeup and detail transfer we use a non-parametric model of those bands.

### 3 Recovering a Visual Representation

We need to recover our visual representation of faces to build our various face editing pipelines. Let  $F$  define the aligned 3D face geometry, which allows us to build coherent face coordinates, and estimate face normals. Let  $s(x)$  and  $\rho(x)$  define image shading and albedo, which are recovered by some intrinsic image process. Let  $\theta$  represent the image illumination, which can include luminaires such as spherical lighting and point light source. Finally, we define a detail map  $I_d(x)$ , which carries identity and finer scale geometry. We render a representation  $(\rho, F, \theta, I_d)$  to produce a new image by the rule

$$I(x) = \text{render}(\rho(x), F, \theta) + I_d(x) \quad (1)$$

where  $\text{render}()$  denotes any reasonable modern 3D rendering process. Next, we will talk about recovering each component in detail, and example components can be found in makeup transfer pipeline in Fig. 3.



**Fig. 2.** The reconstructed face geometry registered to images, and face masks in face segmentation.

**Face geometry:** We reconstruct the geometry of the face in the input image using a morphable face model [3] built from the FaceWarehouse data [5]. The model produces a 3D face mesh  $\mathbf{F}(\omega_i, \omega_e)$  that is a function of identity parameters,  $\omega_i$ , and expression parameters,  $\omega_e$ . FaceTracker [28] is used to detect landmarks  $l_i$  on the face, then we recover parameters and pose of the face mesh by minimizing the distance between the projected landmark vertices and their corresponding landmark locations on the image plane. We propose a Gaussian Mixture Model as a regularizer. The GMM fits to the values of these parameters for the FaceWarehouse data, and a scaled negative probability density value is used as our regularizer. Our morphable face model gives us vertex-to-vertex correspondences between all face meshes. By projecting image pixels to the face meshes (via barycentric coordinates), we get pixel-to-pixel correspondence between all images, see Fig. 2.

**Face segmentation:** Accurate face segmentation masks are essential. We use the aligned face mesh as a cue to segment the input image into hair, face (including neck and ears) and background regions ( $\Omega_h$ ,  $\Omega_f$ , and  $\Omega_b$  respectively). We first project the aligned face geometry onto the image plane and construct three masks – a conservative face mask (strictly corresponding to face), a normal face mask (approximately face and head) and an aggressive face mask (expansion of the normal mask), see Fig. 2.

Hair is typically present at the top of the head; therefore we train a GMM-based hair detector from the pixels at the top of the normal face mask. The hair detector response, intersected with the normal face mask, yields a hair region. Pixels inside the conservative face mask are assigned the face label. Pixels outside the aggressive face mask and are not detected by the hair detector, are initially given the background label. There might be pixels in the image that are not assigned any label. For each of these labels, we learn a GMM-based classifier specific to that input image using color and MR8 features, and re-classify the image. Finally, we use matting Laplacian [14] to correct the face mask.

**Intrinsic images:** The intrinsic image algorithm has a significant qualitative effect on our results. We estimate  $s(x) = \text{Intensity}(x)$  and  $\rho(x) = \mathcal{I}(x)/s(x)$ . Our approach (assume small dark patches are likely shadow) will result in significant errors on current evaluation protocols, but for the tasks we focused on, our method reliably produces better results. A much better albedo recovery algorithm, such as Bell et al. [1] produces significantly worse makeup transfer (Fig. 8, discussion below). Similar problems occur with Retinex [13].

**Face illumination:** We represent the illumination in the image using a Spherical Harmonic (SH) model [27]. Per-pixel normal  $n(x)$  is got by projecting the face mesh onto the image. We then estimate the first-order SH light coefficients  $\theta_s$ , as:

$$\theta_s = \arg \min_{\theta} \sum_{x \in \Omega_f} (s(x) - \text{SH}(n(x), \theta))^2, \quad (2)$$

where  $s(x)$  is the shading computed using our intrinsic method,  $\text{SH}()$  denotes the SH-based rendering of the normals  $n(x)$  with the lighting coefficients  $\theta$ , and  $x \in \Omega_f$  restricts this optimization to the face pixels.

**Face detail map:** Given our estimates of face geometry, albedo, and scene illumination, we can reconstruct the image as:

$$I_e(x) = \text{render}(\rho(x), F, \theta_s) = \rho(x) \text{SH}(n(x), \theta_s) \quad (3)$$

Under ideal circumstances, this reconstruction will approximate the input image, i.e.,  $I_e(x) \approx I(x), \forall x$ . However, the geometric model cannot represent fine-scale spatial structure on the face (for example: detailed folds on eyelids; small grooves and wrinkles; facial blemishes; the wrinkles on lip tissue; the folds around the outside of the lower end of the nose; and so on). These details are typically the results of shading effects. The intrinsic image algorithm we used produces an albedo map that cannot account for these details, but represents the albedo as large blotches of constant color. As a result,  $I_e(x)$  lacks these details. These details are represented by computing a *face detail map*,  $I_d$ :

$$I_d(x) = I(x) - \text{render}(\rho(x), F, \theta_s), \quad (4)$$

## 4 Face Editing

As can be seen from Eqn. 1, our model allows us to manipulate face appearance in three ways. We can change the illumination in the scene by rendering the face under a new set of lights; we can edit the face appearance by editing the albedo layer; we can also edit the fine-scale details of the face by manipulating the detail map. The novel face image is then given by:

$$I^{(n)}(x) = \text{render}(\rho^{(n)}(x), F, \theta^{(n)}) + I_d^{(n)}(x), \quad (5)$$

where  $\theta^{(n)}$  represents the new illumination,  $I_d^{(n)}(x)$  is a (potentially) new detail map, and  $\rho^{(n)}(x)$  represents new albedo.

### 4.1 Face Relighting

Since we have already reconstructed the face albedo and geometry, we can render the face with different lighting effects. Our current implementation supports three distinct types of lights: SH lighting, directional lights, and colored spot lights. SH lighting is our primary source of light, and it enables soft lighting

effects [27]. For other light types (i.e., directional and spotlight) we support casting shadows by leveraging the geometry stored in a depth buffer. For colored spot lights, we separate shading into three distinct channels to allow for finer grain control of light color.

**Rendering and compositing non-face pixels:** While we estimate the albedo at every pixel of the image, we reconstruct geometry only on face pixels. This means that we can only render new shading inside face region, and we must extend this information to cover other regions of the image such as neck, ears and background. We assign each pixel outside the face mask to its nearest shading value inside the face mask, and then smooth the shading.

This gives us the new shading field  $s^{(n)}(x)$  across the full image. In addition, we also have the original reconstructed shading,  $s(x)$ . We blend the two shading fields differently according to different labels. For face region, new shading  $s^{(n)}(x)$  is used. Hair has a characteristic specular appearance, and in order to preserve it, we retains more of the original shading at bright pixels (which are more likely to be specular highlights). The shading for the background is a constant mixture. The final shading is smoothly blended from the three areas.

**Scattering:** There are subsurface scattering effects [10] on human faces, but this is expensive to simulate exactly. We simulate this effect by smoothing the rendered shading field. Comparison figures both with and without smoothing is provided in the supplementary materials.

## 4.2 Makeup Transfer

We transfer makeup effects from a reference face to an input face. We want our results to meet these important qualitative criteria: **Faithfulness:** result images should preserve the input’s distinctive facial features, particularly around the eyes, nose, and lips. **Realism:** the transferred makeup should look like real makeup, so that caking, finger smears, powder texture and so on should be preserved; and relit images should look like pictures taken under a real illumination field. **Predictability:** the pattern, color and qualitative appearance of the transferred makeup should look like the reference.

To transfer colored patterns, we need to replace input image albedo  $\rho^{(i)}(x)$  with reference albedo  $\rho^{(r)}(x)$  for most regions of the face. Define  $I_d^{b(i,r)}$  as a detail map somehow blended by input face detail map  $I_d^{(i)}(x)$  and reference detail map  $I_d^{(r)}(x)$ . Our transfer process is represented as

$$I^{(n)}(x) = \text{render}(\rho^{(r)}(x), F^{(i)}, \theta^{(n)}) + I_d^{b(i,r)}(x) \quad (6)$$

where  $\theta^{(n)}$  could be either  $\theta^{(i)}$  or some new illumination (for relighting).

Our algorithm proceeds as follows. We transfer the whole albedo layer from reference to input face. We construct a new detail map  $I_d^{b(i,r)}$ , which is a blending of reference and input detail maps. The blending preserves the input detail map around eyes, nose and mouth, and reference details elsewhere (Fig. 3 shows our



**Fig. 3.** A makeup transfer example, which includes albedo  $\rho(x)$ , rendered shading  $\text{render}(\rho(x), F, \theta)$ , detail map  $I_d(x)$ , the masks used to blend detail maps, and results.

detail transfer process, and examples of intermediate results). We then choose some illumination  $\theta^{(n)}$  (which might be just the input illumination  $\theta_s$ ) and use Eqn. 6 to composite. Experiment shows  $I_d^{b(i,r)}$  is crucial to the success of our approach. If other intrinsic image methods are used (such as Bell et al. [1]), facial shadows etc. will mix into albedo (rather than only in detail map), and when transfer is performed, the input face appears to have lost distinctive facial features (Fig. 8).

**Boundary processing:** Boundaries need to be processed more carefully for makeup transfer because portraits have a variety of face shapes, face sizes and hairlines, and misclassifying hair as face will create problems. We shrink the face boundary to remove potential hair pixels, then use mirror mapping along the boundary to expand the image (we copy a 30 pixel width ring back and forth). Our approach is fully automatic, however, some level of user interaction could produce even better results.

### 4.3 Face Detail Transfer

This section focuses on transferring face details from a reference image  $I^{(r)}$  to an input image  $I^{(i)}$ , while preserving the general lighting of the input photograph. Different parts of the face have different material properties and fine-scale geometric variation. We account for this by transferring the detail map in nine standard components: left eyebrow, right eyebrow, left eye, right eye, forehead, nose, mouth, left cheek and right cheek. We can preserve the base face shape, and edit face details by transferring one component to all components.

In order to transfer detail maps between  $I^{(r)}$  and  $I^{(i)}$ , we need pixel-accurate alignment between two images. We already have an initial alignment in the form of the aligned morphable face model. In some cases, there might still be some small misalignments, and we fine tune the detail map alignment with optical flow based warping [20]. Makeup transfer also takes this approach to improve correspondence.

A Figure in the supplementary materials shows an example of a shiny nose (caused by specular highlights) being transferred from one source image to a variety of target faces. The alignment and warping ensure that this transfer produces reasonably realistic results.



**Fig. 4. First Row:** Baseline rendering methods compared to ours. Removing detail maps produces overly smooth faces, and using multiplicative details loses detail in dark or bright areas. **Second Row:** This figure compares harsh light relighting and detail transfer results for the average face and our methods. The average face leads to obvious failures especially when lighting is harsh or the face has a strong expression.

## 5 Results

For all face edits, we evaluate our methods qualitatively by direct image comparisons. Previous work has mostly evaluated qualitatively, but Wang et al. report face recognition rates for face relighting [35]. These do not directly evaluate whether our results convince humans. We therefore evaluate quantitatively with a user study for face relighting and detail transfer. For makeup transfer, we use an online attractiveness evaluation system to quantitatively evaluate how much we can make people look more attractive. All the results in this paper are best viewed at high resolution. Please refer to the supplementary materials for more examples and comparisons.

### 5.1 Relighting

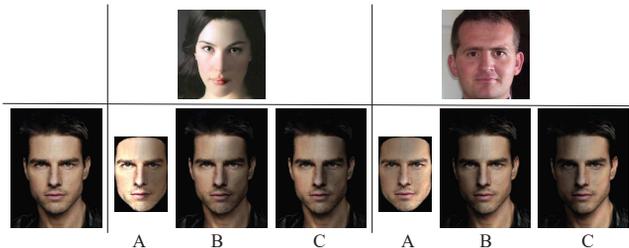
We compare our method to various baselines in Fig. 4 and in the supplementary materials.

**Average Face:** We register the average face mesh to the image, but all other steps in our system remain unchanged, allowing us to evaluate the importance of our morphable face model. This baseline performs poorly, because not morphing the geometry leads to poor registration of nose and mouth (Fig. 4).

**No Detail Map:** When the detail map is not applied, the results have too little spatial detail. This is due to our approximate (and overly smooth) geometry, albedo, and shading estimates.

**Multiplicative Detail (Shading Ratio):** A natural variant of our method is to compute a multiplicative detail map, rendering the new image as

$$I^n(x) = I(x) \frac{\text{render}(\rho(x), F, \theta^{(n)})}{\text{render}(\rho(x), F, \theta_s)} \quad (7)$$



**Fig. 5.** A qualitative relighting comparison between our method and Chen et al. [8] for one source image and two different targets. In each case, A comes from [8] (lighting transfer). B is our result with a manual choice of spherical harmonic light to match target lighting, and C is our method with manually choice of directional light to match target lighting.



**Fig. 6.** Face relighting. Each example shows the original image, our results with minor/major changes to lighting, a result from Portrait Pro 15, and our addition of spot lights. Best viewed at high resolution in color.

and so allowing a comparison with a form of ratio image construction. As Fig. 4 shows, this approach creates odd-looking materials and fails frequently, especially in areas where the ratio is large or small.

**Portrait Pro 15:** It is the best available software for face relighting, but allows only minor changes in lighting before suffering from significant artifacts.

**Our relighting** We have designed an interface to interactively edit spherical harmonics lights, directional lights and spot lights. We relight images using: (a) minor changes to spherical harmonic lighting; (b) significant changes to spherical harmonic lighting; and (c) introducing spotlights. In Fig. 6, we compare our results with the commercial software Portrait Pro 15. Even small lighting changes can cause the Portrait Pro 15 results to have noticeable artifacts; facial skin often appears unnatural. In contrast, even complex lighting edits produced with our method still look natural. This is because our detail map preserves perceived material properties of skin. We also compare qualitatively with a recent lighting transfer method [8], see Fig. 5. Quantitative comparisons show in the user study.

## 5.2 Makeup Transfer

Our makeup transfer is fully automatic, and works well on both normal makeup and unusual (wild) makeup. We evaluate results qualitatively by faithfulness



**Fig. 7.** Our makeup transfer results for various situations: same gender, across gender, from multiple references to one input, and from one reference to multiple inputs.

(does the input face still look like themselves?), realism (does the input face look like they are wearing real makeup?) and predictability (does the makeup on the input look like the makeup on the reference?). For normal makeup, we can compare with previous methods qualitatively; for wild makeup, there are no previous methods. We evaluate quantitatively by investigating the effect of makeup on online services that predict age and attractiveness from face images. Fig. 7 shows makeup transfer of one reference to multiple inputs (to evaluate predictability) and from multiple references to one input (to evaluate faithfulness). It also shows we can transfer makeup across gender.

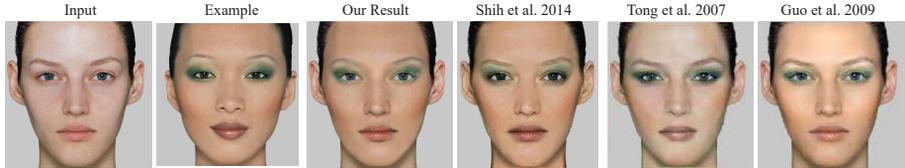
**Different intrinsic images:** Fig. 8 shows the effect of using a different intrinsic image method (a state-of-the-art method [1]). Bell et al.’s method allocates small shadows to the albedo map, resulting in dark lips and nose region in both reference and input’s albedo map. As a result, these shadows are missing from the detail map, so the transferred image result from Bell’s method keeps some of the reference’s facial structure (the shape of the input’s mouth, the shadows in the input’s nostrils, and the shape at the end of the input’s nose). Video contains more intrinsic image comparisons.

**Comparisons:** Fig. 9 shows a comparison to recent methods with normal light makeup. Shih et al. [31] is generally successful, but loses eye makeup, and the skin around cheek is a little unnatural. Their result also suffers ringing artifacts at nostrils, chin and around the eyes. Earlier methods (Tong et al. [33] and Guo et al. [9]) are less successful. These methods are not likely to work with stronger makeup and wild makeup, and they also don’t support multiple face editing operations. The supplementary material contains more results and further comparisons with other methods [15], etc.

**Makeup and attractiveness:** Several websites predict age and attractiveness for a face image; we use <http://howhot.io>, which predicts age and evaluates attractiveness from 1 (least) to 6 (most). Table 1 summarizes the results, and



**Fig. 8.** The albedo produced by Bell et al. [1] method, compared with that produced by our method. Their method retains some face detail in the albedo map (small shadows around nostril and lips); as a result, in the transferred image, the shape of the input’s nostrils and mouth in the transfer “look like” those features in the reference. The effect is further explored in the video submission.



**Fig. 9.** A comparison to the methods of Shih et al., Tong et al., and Guo et al. Our face reference points were adjusted by hand to give the best results (Shih et al. allow manual adjustment of reference points and masks). Results of Tong et al. and Guo et al. are far behind and even don’t get the makeup color right. The skin of Shih et al’s result looks fake, especially around cheek. There are ringing artifacts at chin, left eyelid and nose in their result, and their eye makeup is lost during transfer. Our results looks more natural, eye makeup is preserved, and our method is numerically stable, so does not have numerical artifacts. *Best viewed at high resolution, in color.*

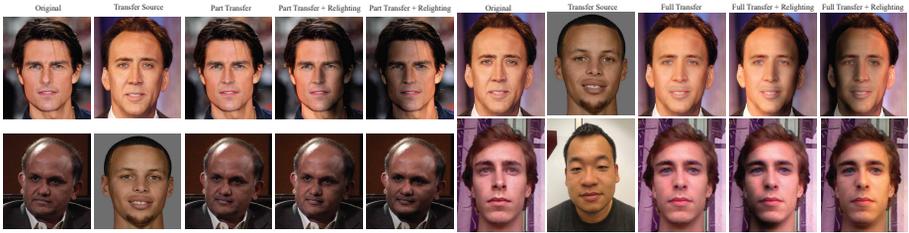
	Normal Makeup		Wild Makeup	
	Age (year)	Attractive (1-6)	Age (year)	Attractive (1-6)
Men	-3.8	0.6	-1.6	-0.2
Women	-2.3	1.5	-1.6	0

**Table 1.** Age and attractive score change after applying our makeup transfer algorithm. Age change is in years, men with normal makeup looks 3.8 years younger; attractiveness has discrete values range from 1 (least) to 6 (most), women with normal makeup look 1.5 levels more attractive. We care about the best improvement for each person using a small number of references. The number in the table is the average of the best improvement. Detailed table see supplementary materials.

detailed table is in supplementary materials. Transfer of naturalistic makeup tends to make a person look younger and more attractive; while wild makeup’s effects depend on images, and on average lower the attractiveness score. Our approach quite reliably yields a substantial, automatic improvement in attractiveness score by allowing an input to search a small set of references for the transfer that yields the best improvement. We anticipate building a simple app that allows users to manipulate their photo profiles.

### 5.3 Detail Transfer

Our method can transfer face details like skin material attributes, wrinkles, scars, small facial expressions and even the shape of nose and mouth, while retaining the original geometric model of the face. Fig. 10 shows detail transfer for areas of



**Fig. 10.** We show results of combining relighting with detail transfer either in one of nine regions or the entire face, more examples and marked transfer regions are available in supplementary materials.



**Fig. 11.** Each example shows a series of detail map transfers using our methods, illustrating how these transfers change the material attributes of skin. The original image in the second example comes from [4], who also edits material attributes by per image interaction (rather than non-parametric transfer). Comparable images appear in that paper, and are duplicated in the supplementary materials. Supplementary materials also contain our reference images and more examples.

the face, and entire faces. By grouping and using attributes of source images, we can reasonably control the detail transfer process, see supplementary materials. In comparison to [4], which offers users the ability to interactively edit face material attributes, our method offers a non-parametric transfer of a face region or whole face, see Fig. 11.

## 5.4 User Study

Relighting methods do not admit comparisons with ground truth images (we do not have ground truth; the model is never physically correct; and the key issue is whether the error is discernible to users). We performed a user study to see how natural our results are. Users are presented with two images, and must decide which of two images has not been edited. We use the error rate in this task to judge the method; a 50% error rate means the method fools users perfectly. To make our comparison fair, we mask the unedited image when the edited image is masked. This likely confuses users and somewhat overestimates the effectiveness of these methods.

There are two groups of users. Group 1 has 29 people with some image editing background, and each finished all tasks. Group 2 has around 100 people and is crowdsourced (CrowdFlower); we assume these are naive users. The error rates are listed in Table 2. In each case of: minor change; major change; detail transfer; relight and detail transfer, we have five image pairs and we average the error rate over image pairs and users. For the comparison methods (listed in Sec. 5.1), we have two image pairs for each task. Tasks are presented in random order. For all the image pairs in the user study, refer to our supplementary materials. For reference, we supply error rates for object relighting methods of [11, 18]. Despite

Group	Methods	Skilled	Crowd Source
Main	Our Minor Light	39.3%	48.9%
	Our Major Light	26.2%	44.1%
	Our Transfer	46.9%	47.1%
	Our Light + Detail	33.1%	37.6%
	Our Average	36.4%	44.4%
Ablation	Average Face Relight	13.8%	43.7%
	Average Face Transfer	13.8%	8.0%
	No Detail Map	3.5%	9.7%
	Multiplicative Detail	3.5%	10.5%
Baseline	Chen 11 [8]	19.0%	20.0%
	Wang 09 [35]	22.4%	14.7%
	Chen 10 [7]	19.0%	10.6%
	Portrait Pro 15	25.9%	32.3%
Objects(graphics)	Karsch [11]	35.5%	
Objects(real)	Liao [18]	44.0%	

**Table 2.** Error rates for our user study, higher is better, see Sec. 5.4.

reasonable concerns about uncanny valley effects, our methods can produce error rates that match or exceed theirs.

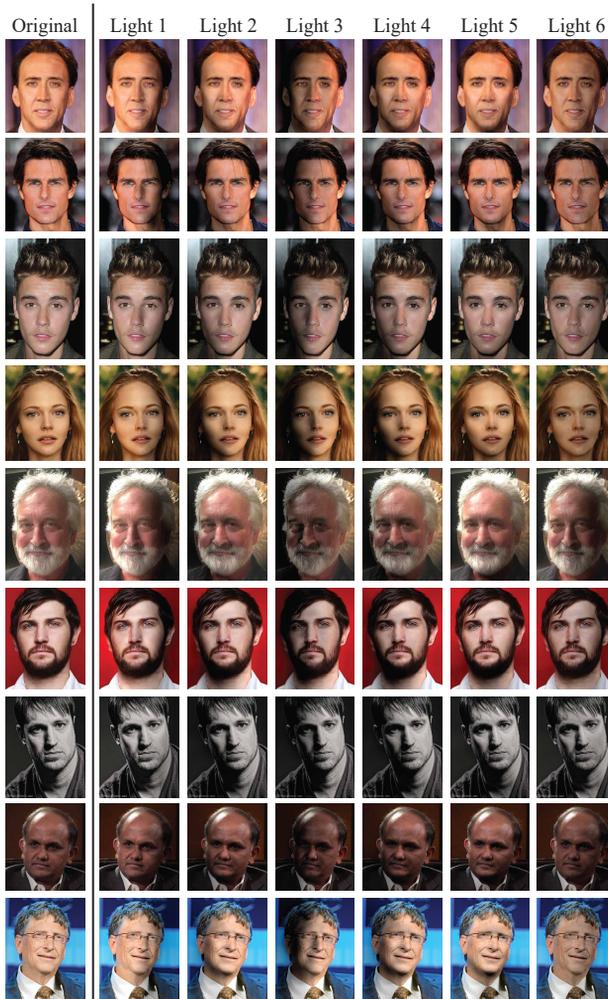
Our study supports the conclusions: 1) users attended to tasks, because the error rate for some tasks is low; 2) both our detail transfer and relighting methods have high error rates, meaning our method produces results people can not distinguish from natural images; 3) it’s important to use morphable models and detail maps, because the error rate drops in the absence of either strategy; 4) the multiplicative detail map (or shading ratio) works poorly; 5) methods that mask faces solve an easier task (because they need not relight hair, ears, neck and background) but still fool users less often than our method; 6) Portrait Pro 15 is relatively effective, but cannot change lighting much (supplementary materials); 7) unlike the experience of [11, 18], skilled users significantly outperform naive users, but can still be fooled by our method at useful rates.

## 6 Conclusion

We propose a visual representation for face editing that combines an approximate shading model based on coarse estimates of geometry, albedo, and lighting, with a non-parametric face detail map that captures deviations like non-Lambertian reflectance, complex shading, fine-scale albedo and geometry variations. Our model can be used to solve multiple tasks: relight faces by changing the lighting parameters, transfer makeup by transferring albedo and blending detail maps, and edit face details by only transferring detail maps. We present a quantitative evaluation of makeup transfer by tracking the age and attractiveness from an online evaluation system. Results show that we can reasonably make people look younger and more attractive by our makeup transfer. We also present the first quantitative perceptual evaluation of face relighting and detail editing via a user study; this study shows that, to a large extent, our method produces results that users are unable to distinguish from real images. Our approach creates state of the art results on multiple face editing tasks, and in the future, we would like to model the face detail map more thoroughly and extend this work to more face editing applications and also applications in other fields, such as building appearance editing, etc.

## 7 Appendix

Refer to our support materials on project website for additional results and explanations.



**Fig. 12.** Relighting Table. Relighting the image according to several lighting conditions. Each row is one subject, each column is one lighting condition.



**Fig. 13.** Relighted makeup transfers from reference (lower right inset) image to input (lower left inset) image, for both genders *Best viewed at high resolution, in color.* Further examples.



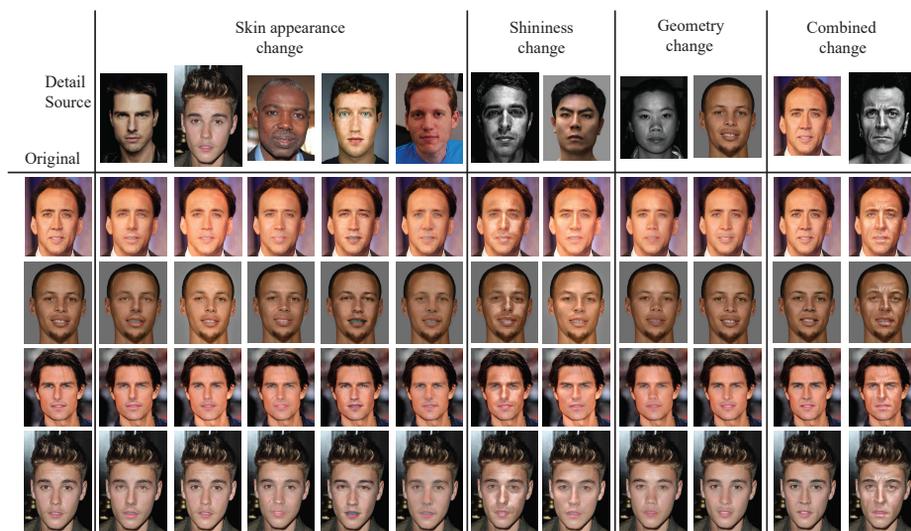
**Fig. 14.** Across gender transfers are usually successful. *Best viewed at high resolution, in color.* Further examples.



**Fig. 15.** One reference transferred to multiple inputs. The input retains their face shape, but appears to wear the reference's makeup. *Best viewed at high resolution, in color.* Further examples.



**Fig. 16.** Multiple references transferred to one input. The input retains their face shape, but appears to wear the reference’s makeup. *Best viewed at high resolution, in color.* Further examples.



**Fig. 17.** Transfer Table. We transfer all of the nine parts of the face detail map (full transfer) from the subjects in the detail source row (first row) to the subjects in the original column (first column). Each row is one subject, and each column shares the same face detail map from the detail source person in that column. We grouped the transfer effects into skin property change, shininess change, geometry change and combined change.

## References

1. Bell, S., Bala, K., Snavely, N.: Intrinsic images in the wild. *ACM Trans. on Graphics (SIGGRAPH)* 33(4) (2014)
2. Bitouk, D., Kumar, N., Dhillon, S., Belhumeur, P., Nayar, S.K.: Face swapping: automatically replacing faces in photographs. *TOG* 27(3), 39 (2008)
3. Blanz, V., Vetter, T.: A morphable model for the synthesis of 3D faces. In: *SIGGRAPH*. pp. 187–194. *ACM* (1999)
4. Boyadzhiev, I., Bala, K., Paris, S., Adelson, E.: Band-sifting decomposition for image-based material editing. *ACM Trans. Graph.* (2015)
5. Cao, C., Weng, Y., Zhou, S., Tong, Y., Zhou, K.: Facewarehouse: A 3d facial expression database for visual computing. *IEEE Transactions on Visualization and Computer Graphics* 20(3), 413–425 (Mar 2014), <http://dx.doi.org/10.1109/TVCG.2013.249>
6. Chai, M., Luo, L., Sunkavalli, K., Carr, N., Hadap, S., Zhou, K.: High-quality hair modeling from a single portrait photo. *ACM Trans. Graph.* 34(6), 204:1–204:10 (Oct 2015)
7. Chen, J., Su, G., He, J., Ben, S.: Face image relighting using locally constrained global optimization. In: *ECCV*, pp. 44–57. *Springer* (2010)
8. Chen, X., Chen, M., Jin, X., Zha, Q.: Face illumination transfer through edge-preserving filters. *CVPR*, *IEEE Computer Society* pp. 281–287 (2011)
9. Guo, D., Sim, T.: Digital face makeup by example. In: *CVPR* 2009
10. Jimenez, J., Sundstedt, V., Gutierrez, D.: Screen-space perceptual rendering of human skin. *ACM Transactions on Applied Perception* 6(4), 23:1–23:15 (2009)
11. Karsch, K., Sunkavalli, K., Hadap, S., Carr, N., Jin, H., Fonte, R., Sittig, M., Forsyth, D.: Automatic scene inference for 3d object compositing. *ACM Trans. Graph.* 33(3) (June 2014)
12. Kemelmacher-Shlizerman, I., Suwajanakorn, S., Seitz, S.M.: Illumination-aware age progression. *CVPR* (2014)
13. Land, E.H., John, Mccann, J.: Lightness and retinex theory. *Journal of the Optical Society of America* pp. 1–11 (1971)
14. Levin, A., Lischinski, D., Weiss, Y.: A closed-form solution to natural image matting. *TPAMI* 30(2) (2008)
15. Li, C., Zhou, K., Lin, S.: Simulating makeup through physics-based manipulation of intrinsic image layers. *CVPR* (2015)
16. Li, Q., Yin, W., Deng, Z.: Image-based face illumination transferring using logarithmic total variation models. *The visual computer* 26(1), 41–49 (2010)
17. Liang, L., Jin, L., Li, X.: Facial skin beautification using adaptive region-aware masks. *Cybernetics*, *IEEE Transactions on* 44(12), 2600–2612 (2014)
18. Liao, Z., Karsch, K., Forsyth, D.: An approximate shading model for object relighting. *CVPR* (2015)
19. Liao, Z., Rock, J., Wang, Y., Forsyth, D.: Non-parametric filtering for geometric detail extraction and material representation. In: *CVPR*. pp. 963–970. *IEEE* (2013)
20. Liu, C.: Beyond pixels: Exploring new representations and applications for motion analysis. *Doctoral Thesis*. *Massachusetts Institute of Technology* (May 2009)
21. Liu, L., Xing, J., Liu, S., Xu, H., Zhou, X., Yan, S.: Wow! you are so beautiful today! *TOMM* 11(1s), 20 (2014)
22. Liu, S., Liu, L., Yan, S.: Fashion analysis: Current techniques and future directions. *MultiMedia*, *IEEE* 21(2), 72–79 (2014)

23. Liu, Z., Shan, Y., Zhang, Z.: Expressive expression mapping with ratio images. In: Proceedings of the 28th annual conference on Computer graphics and interactive techniques. pp. 271–276. ACM (2001)
24. Marschner, S.R., Greenberg, D.P.: Inverse lighting for photography. In: Color and Imaging Conference. vol. 1997, pp. 262–265. Society for Imaging Science and Technology (1997)
25. Peers, P., Tamura, N., Matusik, W., Debevec, P.: Post-production facial performance relighting using reflectance transfer. *ACM Transactions on Graphics (TOG)* 26(3), 52 (2007)
26. Pont, S.C.: Material–illumination ambiguities and the perception of solid objects. *Perception* 35(10), 1331 (2006)
27. Ramamoorthi, R., Hanrahan, P.: An efficient representation for irradiance environment maps. *SIGGRAPH* 2001
28. Saragih, J., McDonald, K.: Facetracker. <https://github.com/kylemcdonald/FaceTracker>
29. Scherbaum, K., Ritschel, T., Hullin, M., Thormählen, T., Blanz, V., Seidel, H.P.: Computer-suggested facial makeup. In: *CGF*. vol. 30, pp. 485–492. Wiley Online Library (2011)
30. Shashua, A., Riklin-Raviv, T.: The quotient image: Class-based re-rendering and recognition with varying illuminations. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 23(2), 129–139 (2001)
31. Shih, Y., Paris, S., Barnes, C., Freeman, W.T., Durand, F.: Style transfer for headshot portraits. *ACM Trans. Graph.* (2014)
32. Stoschek, A.: Image-based re-rendering of faces for continuous pose and illumination directions. In: *Computer Vision and Pattern Recognition, 2000. Proceedings. IEEE Conference on*. vol. 1, pp. 582–587. IEEE (2000)
33. Tong, W.S., Tang, C.K., Brown, M.S., Xu, Y.Q.: Example-based cosmetic transfer. In: *PG'07*. pp. 211–218. IEEE (2007)
34. Wang, Y., Liu, Z., Hua, G., Wen, Z., Zhang, Z., Samaras, D.: Face re-lighting from a single image under harsh lighting conditions. In: *CVPR*. pp. 1–8. IEEE (2007)
35. Wang, Y., Zhang, L., Liu, Z., Hua, G., Wen, Z., Zhang, Z., Samaras, D.: Face relighting from a single image under arbitrary unknown lighting conditions. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 31(11), 1968–1984 (2009)
36. Wen, Z., Liu, Z., Huang, T.S.: Face relighting with radiance environment maps. In: *CVPR 2003*. vol. 2, pp. II–158. IEEE (2003)
37. Yang, F., Wang, J., Shechtman, E., Bourdev, L., Metaxas, D.: Expression flow for 3d-aware face component transfer. *TOG* 30(4), 60 (2011)