

Optimizing Portrait Lighting at Capture-Time Using a 360 Camera as a Light Probe

Jane L. E
Stanford University
ejane@stanford.edu

Ohad Fried
Stanford University
ohad@stanford.edu

Maneesh Agrawala
Stanford University
maneesh@cs.stanford.edu




Figure 1. In a fixed lighting environment like this *arch walkway*, photographers can produce many different lighting styles (e.g. butterfly, right loop, right split, and right rim) just by rotating the subject in place without changing their location. Given an HDR environment map from a 360 camera at some initial orientation (Figure 2) and a target lighting style (bottom left), our tool automatically identifies the optimal angle for reorienting the subject to match the desired lighting – e.g. 90° for butterfly lighting. We use a precomputed radiance transfer-based method on a generic head, skin, and camera model for efficiently optimizing lighting orientation and for visualizing the best orientation match (bottom right).

ABSTRACT

We present a capture-time tool designed to help casual photographers orient their subject to achieve a user-specified target facial appearance. The inputs to our tool are an HDR environment map of the scene captured using a 360 camera, and a target facial appearance, selected from a gallery of common studio lighting styles. Our tool computes the optimal orientation for the subject to achieve the target lighting using a computationally efficient precomputed radiance transfer-based approach. It then tells the photographer how far to rotate about the subject. Optionally, our tool can suggest how to orient a secondary external light source (e.g. a phone screen) about the subject’s face to further improve the match to the target lighting. We demonstrate the effectiveness of our approach in a variety of indoor and outdoor scenes using many different subjects to achieve a variety of looks. A user evaluation suggests that our tool reduces the mental effort required by photographers to produce well-lit portraits.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

UIST’19, October 20–23, 2019, New Orleans, LA, USA

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM. ISBN 978-1-4503-6816-2/19/10...\$15.00

DOI: <https://doi.org/10.1145/3332165.3347893>

Author Keywords

portrait lighting; environment map; 360 camera

CCS Concepts

•Human-centered computing → Graphical user interfaces; •Computing methodologies → Graphics systems and interfaces; Computational photography

INTRODUCTION

“We can do everything else beautifully, but if our lighting is bad, our portrait will be bad. It is that simple.”

Light Science & Magic [22]

In portrait photography, lighting is one of the most important elements for establishing the overall look and mood of the image [8, 22]. Professional portrait photographers working in a studio, typically place one or more lights around the subject to carefully control the distribution of bright and dark regions on the subject’s face. Outside the studio, where they cannot control light placement, professional photographers instead focus on orienting the subject with respect to the lights in the environment to adjust the distribution of light falling on the face. While the fixed light placement makes it impossible to achieve all the different looks that are possible in a studio, good photographers can often produce several distinctly different appearances just by reorienting the subject (Figure 1). Yet, most casual photographers primarily focus on maintaining



Figure 2. Photographers first capture an HDR environment map at the expected location of the subject’s head using a 360 camera (left). The subject then moves to the location, and our tool suggests how far to rotate the subject to achieve different lighting styles (middle). After re-orientation, the photographer shoots the final portrait using a primary portrait camera (right). Note that our approach also allows for selfies, where the subject is the photographer.

even brightness on the face [29], rarely considering all the variations in facial appearance that are possible via reorientation.

In this paper, we present a capture-time tool that suggests how portrait photographers should orient their subject within an environment to best achieve a user-specified target facial appearance. Photographers start by using an off-the-shelf 360 camera as a light probe to capture an HDR environment map of the light in a scene (Figure 2). They then select a target facial appearance from a pre-designed gallery of common studio lighting styles that the tool determines are most feasible in the current environment. Photographers can optionally paint weights to adjust the distribution of bright and dark regions in the target appearance. Our tool then tells the photographer how far the subject should rotate (clockwise or counterclockwise) to produce the best match to the target appearance and provides real-time feedback showing how close the current camera view is to the view at the target orientation. If the photographer has a secondary source of light (e.g. a phone screen, a flash, etc.), our tool can also suggest how to orient it about the subject’s face to further improve the match to the target lighting. After reorientation, the photographer can shoot the portrait using any available primary camera (e.g. a phone, a DSLR, etc.).

Our prototype implementation runs on a laptop, uses a 360 camera to capture the environment, and a webcam to determine the current view to aid in reorientation. But as 360 cameras become more widespread, we envision that they will become a standard tool for capturing environment maps and that all of our prototype hardware would be built into the primary portrait camera. Already it is possible to buy a 360 camera that attaches to a cell phone [23]. Thus, we have designed our implementation to be computationally lightweight enough to easily run on a phone or a DSLR, while remaining performant enough to be used at capture-time.

The key idea of our approach is to use an efficient precomputed radiance transfer (PRT) method [9, 36, 49] with a generic head, skin, and camera model, to compute the appearance of the face under different orientations of the environment. We show how to formulate a target facial appearance as a weighting function

in the image domain, and pre-integrate it against a light transport matrix for the generic head to significantly accelerate the search for the orientation of the lighting environment that best achieves the target appearance. We also introduce a method for optimally placing a secondary light source in the environment to best achieve a target facial appearance. Our prototype implementation running in Matlab takes 0.19 milliseconds (5263 fps) to compute the optimal orientation. A user evaluation suggests that our tool is very useful and reduces the mental effort required to produce well-lit portraits. We discuss why our approach may be suitable for implementation directly in phones and/or DSLR hardware.

RELATED WORK

Human faces are one of the most popular subjects in photography. As a result, researchers have developed a number of techniques for enhancing portraits via perspective correction [17], transferring makeup [20, 54], improving attractiveness [28], transfiguring appearance based on Internet photos [24], and bringing still portraits to life [4]. Lighting design is another well studied problem, and here we focus on the subset of these methods that are most related to our work on portrait lighting.

Computational image relighting. Debevec et al. [14] were the first to introduce the approach of capturing multiple images of a face (or scene) from the same viewpoint, but under different lighting conditions, and then letting users composite these basis images to produce the image under novel illuminations. Extensions to this approach let users specify higher level design objectives (e.g. emphasize contour, remove shadows), and then automatically find the set of input images that should be composited to achieve the desired result [2, 3, 10, 11]. However, all of these methods require 10s to 100s of input images under varying lighting conditions and some require specialized lighting hardware, which can make it difficult for subjects to remain still throughout the capture process.

Others have focused on computational relighting using a smaller number of input images. Quotient image methods [30, 32, 39, 44, 57] require two images of the same face under different lighting conditions A and B and compute the ratio of the pixel value. Then given an image of a new face under lighting condition A, these methods can generate the new face under lighting condition B. Portrait style transfer [48] focuses on transferring image statistics from a user-chosen style exemplar to an input portrait. While this method does not focus on relighting, the resulting changes to the image usually affect local contrast and can appear to change the overall illumination falling on the face.

While these computational relighting methods are powerful, they are designed to post-process input images after capture, and all require pixel-level alignment between the faces in each input image to ensure artifact-free results. Moreover, some of them can generate unrealistic lighting as they may fuse together pixels or image statistics from portions of different images. In contrast, we focus on directing the photographer at capture-time to shoot the single photograph that best achieves the desired lighting. Therefore, our approach cannot suffer from pixel-level artifacts or unrealistic lighting.




Figure 3. Our tool includes a set of 16 lighting styles (top), selected from commonly used studio lighting styles to represent the diversity of placement of the main light source(s). Only 11 styles are shown here as each named style that starts with the word “right” also has a “left” version in our tool. We render these in PBRT [42] using the generic head and skin model, manually placing area light sources around the head based on studio lighting setups shown in portrait photography books [8, 22]. We convert a lighting style into a target facial appearance weighting function f' by rescaling it to lie between $[0, 1]$ and then shifting it to put the mean brightness of the facial pixels at 0 (middle). We then pre-integrate the weighting function against the light transport matrix T to form the pre-integrated target in the lighting domain F' . Orange represents positive values where the image should be brighter and blue represents negative values where the image should be darker.

Automated lighting design for synthetic scenes. Lighting plays a central role in the perception of scenes. Based on this fact, researchers have developed lighting optimization techniques designed to place highlights and shadows [25, 40] to enhance shape [19, 26, 45, 46, 56], and to improve materials perception [9] in synthetic scenes. Our work is inspired by these methods, but focuses on optimizing the lighting for depicting faces in real world environments.

Capture-time lighting optimization. A few research groups have developed capture-time lighting optimization techniques that focus on adding light to a scene at capture-time. Adelsberger et al. [1] develop custom flash hardware that uses a depth image to spatially adapt the illumination so that the scene is evenly lit despite variations in depth. Srikanth et al. [51] automatically position rim lights around the subject using a robotic light-carrying drone to create dramatic contour lighting. Murmann et al. [34] develop custom bounce flash hardware that automatically reorients the flash to bounce off nearby surfaces to ensure that the subject’s face is well lit. While these projects are similar to our work, they require customized hardware (e.g. depth sensors, drones, reorientation motors) making them inaccessible to photographers today. In contrast, our approach focuses on using off-the-shelf hardware so that anyone with a 360 camera and primary portrait camera could use it today. Li and Vogel [29] present a smartphone application that can guide users towards an evenly lit portrait at a fixed location. Their work focuses on scenes with one main directional light source whereas we account for all lighting in the scene.

PORTRAIT LIGHTING OPTIMIZATION

Our lighting optimization tool takes two inputs: (1) an HDR environment map, representing the illumination falling on the subject’s face, and (2) a user-specified target facial appearance that represents the desired distribution of bright and dark regions on the subject’s face. The tool then suggests how to reorient the subject with respect to the environment to best produce the target appearance.

We have analyzed a number of books on portrait photography [8, 21, 22, 31, 33] and they all suggest that because skin is a diffuse reflector, photographers should focus on the broad, low-frequency distributions of bright and dark regions on the face rather than on high-frequency highlights and details. They describe studio lighting setups (placement and orientation) that produce different looks by changing the bright and dark facial regions (Figure 3 top row). Outside the studio, they suggest orienting the subject with respect to the lights in the environment to similarly control facial lighting. In short, the orientation of the face with respect to the environment has a much greater impact on facial appearance due to lighting, than either the individual differences in facial features (e.g. geometry, skin tone, etc.), or perspective effects from the camera, and a low-resolution environment map is enough to capture the light falling on the subject.

These guidelines allow us to use three key approximations in rendering a suitable proxy for evaluating the lighting on the face: we use (1) a generic head and skin model, (2) a camera with a default field of view placed at a fixed distance from the head, and render it with (3) a significantly downsampled version of our input HDR environment map. Because the scene and camera models are fixed, we can adapt Bousseau et al.’s [9] pre-computed radiance transfer approach to our problem and efficiently find the orientation of the face that best matches the target lighting. In the next four subsections, we first summarize how we adapt Bousseau et al.’s lighting optimization approach, then describe how the target facial appearances are specified for our problem, next show how our approach can be extended to place a secondary light source, and finally provide implementation details.

Lighting Optimization with Pre-Integration

Bousseau et al.’s [9] lighting optimization approach builds on precomputed radiance transfer techniques [36, 49]. In matrix form, the image \mathbf{B} of the scene lit by an environment map \mathbf{L} is given by

$$\mathbf{B} = \mathbf{T}\mathbf{L}. \quad (1)$$

Both \mathbf{B} and \mathbf{L} are vectors, of the image and environment respectively, while \mathbf{T} is the light transport matrix that represents how light from the environment map is transported through the scene to form the image. The columns of \mathbf{T} are each a rendered image of the scene as lit by a single pixel of the environment map.

Suppose $f(\mathbf{x})$ is a weighting function that specifies the target facial appearance at each image pixel \mathbf{x} , using positive values where the face should be brighter and negative values where the face should be darker (Figure 3). Then we can define an image quality metric as the integral over all the pixels in the image of the pixel-wise product of the weight function and the image. In matrix form, the image quality metric C is

$$C = \mathbf{f}'\mathbf{B} = \mathbf{f}'\mathbf{T}\mathbf{L}. \quad (2)$$

Once the target weighting function \mathbf{f} is fixed, we can pre-integrate it over the image domain

$$\mathbf{F}' = \mathbf{f}'\mathbf{T}, \quad (3)$$

and the image quality metric becomes

$$C = \mathbf{F}'\mathbf{L}. \quad (4)$$

The pre-integration allows us to efficiently compute the image quality C under any lighting \mathbf{L} as a dot product of two vectors that are the size of the environment map.

In order to compute the optimal orientation for the lighting \mathbf{L} , we find the rotation \mathbf{R} that maximizes

$$C(\mathbf{R}) = \mathbf{F}'\mathbf{R}(\mathbf{L}). \quad (5)$$

In practice, since most subjects are standing or sitting up vertically with respect to the environment, we limit the search to a set of 1D rotation angles about the vertical axis (azimuth).

Specifying Target Facial Appearance

The target facial appearance weighting function $f(\mathbf{x})$ specifies which pixels in the final image should be brighter (positive values) and which should be darker (negative values). A weight value of $f(\mathbf{x}) = 0$ implies that pixel \mathbf{x} should be ignored when evaluating the image quality metric C .

Our lighting optimization tool provides a built-in set of such weighting functions that are based on common studio lighting styles (Figure 3). Suppose \mathcal{T} is a rendered image of our generic head model in one of these lighting styles. We mask out the pixels falling outside of the rendered face (i.e. background pixels) as they are irrelevant to the target facial appearance. We then linearly rescale the remaining pixels of \mathcal{T} to lie between $[0, 1]$. Finally we set

$$f(\mathbf{x}) = \begin{cases} \mathcal{T}(\mathbf{x}) - \text{mean}(\mathcal{T}) & \text{if } \mathbf{x} \in \text{Face} \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

where \mathcal{T} is the masked and re-scaled version of \mathcal{T} , and $\text{mean}(\mathcal{T})$ is the mean pixel value of \mathcal{T} across the facial pixels. By setting our target weighting function $f(\mathbf{x})$ in this way, our optimization tries to match the brightness at each point on the face relative to the average brightness of the face. In practice, we have found that this approach generally finds a good match in terms of the overall distribution of bright and

dark regions on face and is relatively robust to differences in the total amount of light in the environment – i.e. we can use the same target in both dark and bright settings. For additional control, our tool lets users manually adjust the built-in target weighting functions using a painting interface as described in Step 2 of the Interaction Section.

Placing an Additional Light Source

Photographers may sometimes have an additional light source (e.g. a phone screen, a flash unit, etc.) that they can manually orient around the face to further improve the match to the target lighting. We can adapt Equation 5 to compute the optimal orientation for this additional light source as follows

$$C(\mathbf{R}_{\text{add}}, \mathbf{R}) = \mathbf{F}'(\mathbf{R}_{\text{add}}(\mathbf{L}_{\text{add}}) \mathbf{R}(\mathbf{L})), \quad (7)$$

where \mathbf{L}_{add} is an environment map capturing just the additional light source with an accompanying alpha mask set to 1 within the light source and 0 outside it, and \mathbf{R} is the *over* image compositing operator. Our optimization searches over the space of rotations for the additional light source \mathbf{R}_{add} as well as the primary scene lighting \mathbf{R} and considers the quality metric C for the composite of these two environment maps. We use the over operator under the assumption that the additional light source is closer to the face and therefore occludes the scene lighting behind it. Thus, our optimization finds the pair of orientations that best produce the target appearance. Note that because the additional light source is often handheld we search a 2D band of rotation angles that vary both the light source’s azimuth and altitude angles.

Implementation

Our optimization algorithm requires an environment map \mathbf{L} and a target facial appearance weighting function $f(\mathbf{x})$ as input. While current 360 cameras can generate high resolution (e.g. 4K) environment maps, we have found that resolutions of 64×32 for \mathbf{L} and 100×100 for the image are sufficient for our optimization (see Algorithmic Evaluation Section). Additional resolution in the lighting or the image makes little visual difference in the appearance of the rendered head model, as the subsurface skin material we use is relatively diffuse.

We precompute the transport matrix \mathbf{T} using the PBRT ray-tracer [42] treating the generic PBRT head with the built-in subsurface scattering skin material model (`kdsurface`) as our scene. Each column of \mathbf{T} is an image of the scene as lit by a single pixel of the environment map \mathbf{L} . For our image and lighting resolutions, serially computing \mathbf{T} on a single machine takes 17.1 hours. The resulting three-channel RGB transport matrix \mathbf{T} is 246 MB uncompressed, and after precomputing \mathbf{T} , we can generate a color rendering of the scene under any lighting \mathbf{L} using Equation 1 in 13 milliseconds (77 fps). In practice however, we only use the three-channel transport matrix to render the visualization of the lighting on the generic head model at the optimal orientation which we call the *PRT render match* (Figures 1 and 4). For all other computations, since the distributions of bright and dark regions matter more than color, we use a grayscale transport matrix \mathbf{T} of size 81.9 MB uncompressed and also convert the lighting \mathbf{L} to grayscale.

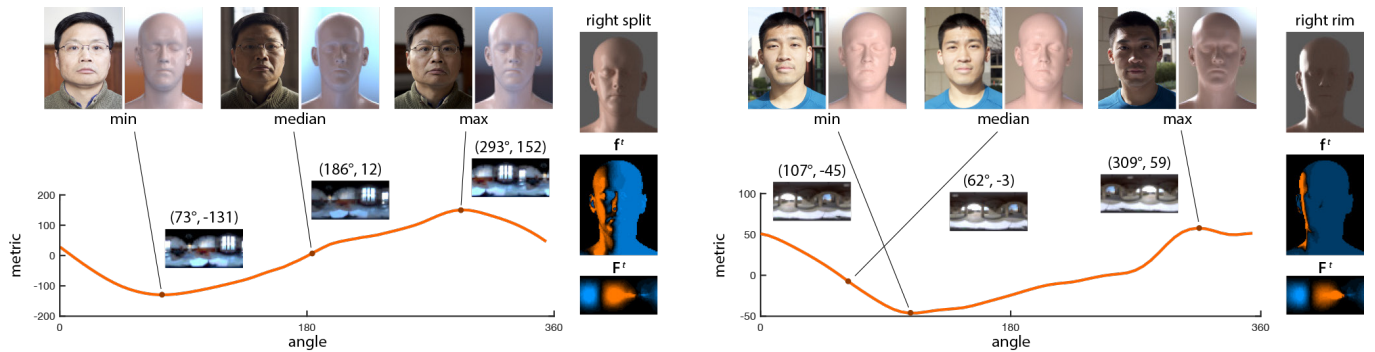


Figure 4. The image quality metric varies as we rotate the environment. Maximizing the metric produces a good match (both photo and PRT render) to the target facial appearance and the bright/dark areas in the environment map match those in the pre-integrated target weighting function F^l . At the median and minimum metric values the matches are not as good.

Pre-integrating the target weighting function f^l against the single-channel transport matrix T to compute F^l as in Equation 3 takes 3.5 milliseconds (286 fps). After this pre-integration, computing our image quality metric is very fast as it requires computing a dot product between two vectors F^l and L and takes 0.003 milliseconds (333333 fps). Given the pre-integrated target weighting function F^l , we can identify the lighting rotation R that maximizes our image quality metric C using Equation 5 (Figure 4). We have found that it is sufficient to compute the image quality metric C at 5.625° rotation increments about the vertical axis. Finer sampling is unnecessary in practice as photographers usually only reorient to within a few degrees of the target orientation (see User Evaluation Section). Computing the optimal orientation for the environment lighting takes 0.19 milliseconds (5263 fps) for a single lighting target and is linearly related to the number of rotation increments we sample (in our case 64).

To optimize placement of an additional light source, we sample a 2D band of rotation angles ranging from 0° to 360° at 5.625° increments around the vertical axis (azimuth) and ranging from -67.5° to 67.5° about the horizontal axis (altitude) at 16.875° increments. Computing the orientation for the environment as well as an additional light source, then takes 110 milliseconds (9 fps) for a single target and is again linearly related to the number of rotation increments we sample (38864 in our case; 64 for the environment times 64×9 for the additional light). We have experimented with a greedy approach in which we first find the optimal rotation for the environment, keep that rotation fixed, and then search for the optimal rotation for the additional light source. This approach significantly reduces the number of rotation samples ($640 = 64 + 64 \times 9$), and therefore our overall cost to 1.9 milliseconds (526 fps). While it generally produces similar results to the full optimization, they are not always the same. In this paper we report all results using the full optimization and leave it to future work to further improve the efficiency of placing an additional light source.

All of our timings are computed using our Matlab implementation running on a MacBook Pro laptop (2017) with a 3.1 GHz Intel Core i7 processor. We expect a natively compiled implementation running on a GPU or even a CPU would run faster and could easily run on modern smartphones.

INTERACTION

For a photographer, using our portrait lighting optimization tool involves four steps: (1) capturing an HDR environment map of the location, (2) specifying a target facial appearance, (3) following our tool’s guidance to reorient self and subject, and (4) taking the final portrait. Note that our tool can be used without modification to take selfies in which the photographer is the subject. We detail how the photographer interacts with our tool in the first three of these steps.

Step 1: Capture Environment Map

The photographer starts by capturing an HDR environment map in the location they would like the subject’s face to appear in the final photograph using a 360 camera in bracketed exposure mode (we use a Ricoh Theta V). The photographer should ensure that nothing is obstructing the camera’s view of the light sources in the environment. For handheld capture, we recommend that the photographer kneel under the 360 camera to minimize occlusions.

In bracketing mode, we take 3 exposures at fixed aperture and ISO – one exposed at the auto-exposure settings, one two stops brighter, and one two stops darker. The Ricoh takes about 10 seconds to capture the 3 images and outputs gamma-corrected sRGB at 5376×2688 resolution in JPG format. We provide a pre-processing tool that applies Ebner’s [15] linearization algorithm via Matlab’s `rgb2Lin` function and then uses Reinhard et al.’s [43] algorithm via Matlab’s `makehdr` function to convert the exposures into a single HDR environment map. Finally, we downsample the environment map to 64×32 resolution for use in our optimization procedure. This pre-processing takes about 16.2 seconds.

While the HDR capture, transfer to the laptop, and pre-processing, can take a few minutes in our prototype implementation, the vast majority of the time is spent in the transfer. We believe that in the fully integrated tool we envision, where the tool runs on a single device that captures the environment at low resolution, this time would be significantly reduced (see Discussion Section). To place an additional light source the photographer must a priori capture an HDR environment map of the additional light (L_{add} in Equation 7) in a dark (ideally blackout) room, using the same procedure and camera settings as used to capture the scene. We set the light (a phone screen) about one foot from the 360 camera in our experiments. In

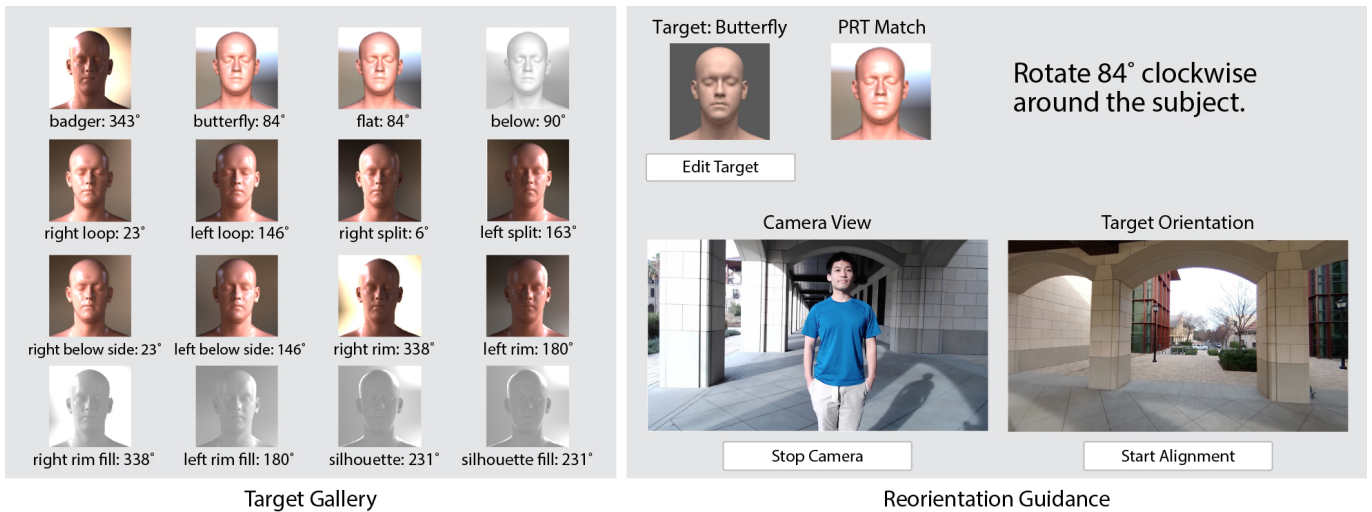


Figure 5. Upon loading an HDR environment map, the interface of our portrait lighting tool computes the optimal orientation for each target lighting style (Figure 3) and displays a gallery of target appearances (left). The tool grays out *unattainable* targets. Selecting a target brings up the reorientation guidance screen, which displays how far the photographer should rotate clockwise or counterclockwise about the subject (right). It also shows the background at the target orientation and the current view from a webcam at the current location of the photographer. As the photographer gets close to the target orientation, clicking the *Start Alignment* button overlays the current view on the target orientation view (see supplemental material).

general we expect photographers would choose this offset distance based on the power of the light source and how close to the face they are willing to place the additional light. The photographer must capture a new HDR environment map for each additional light source and offset distance, but once captured they can be reused in any setting. In this case, the main environment map should be captured with matching camera exposure settings.

Step 2: Specify Target Facial Appearance

Our optimization tool includes a pre-designed gallery of target facial appearance weighting functions that are based on common studio lighting styles (Figure 3). Because these targets are known in advance, we pre-integrate them against the transport matrix \mathbf{T} a priori. As soon as the photographer loads the environment map, our tool applies the optimization procedure (see Lighting Optimization with Pre-Integration Section) to identify the optimal orientation angle for each target and displays all of these results gallery of target appearances (Figure 5).

Not all target appearances are achievable in every environment. In such cases, the image quality metric is relatively low for all orientations and the quality of the maximizing orientation is close to the mean quality across all orientations. Therefore, our tool computes the maximum image quality across all orientations and if it is less than an absolute threshold ϵ , our tool marks the target appearance as *unattainable*. We empirically set $\epsilon = 45$. In the gallery of target appearances, all unattainable targets are grayed out. Additionally, photographers can manually adjust any of the targets, including those that are grayed out, using a painting interface (see supplemental materials). Our tool re-computes the pre-integration of the adjusted target weighting \mathbf{f}' against the transport matrix \mathbf{T} , and then runs the optimization using the adjusted \mathbf{F}' .

Step 3: Follow Guidance to Reorient Self and Subject

Once the photographer chooses a specific target facial appearance, our lighting optimization tool tells the photographer how

far to rotate (clockwise or counterclockwise) about the subject to achieve the desired appearance (Figure 5). In addition to the rotation angle, it provides the current view from the primary camera, and also displays the portion of the environment that should be visible after the rotation. The goal of the photographer is to match these two views. To further aid the reorientation process, our tool optionally provides real-time feedback showing how well the current and desired views match, by overlaying the current view onto the desired view after transformation by the best-fit homography between them (see supplemental materials). If the photographer is placing an additional light source, our tool next describes how to orient the light with respect to the subject’s face as a pair of rotation angles (azimuth and altitude) for the center of the light, and provides a schematic diagram illustrating the relative orientation between the light and the face (Figure 8) after the subject has been oriented with respect to the environment. In practice, either the photographer, the subject, a third person or a tripod needs to hold the additional light source in place.

RESULTS

We have tested our portrait lighting tool in a variety of locations with many different subjects – e.g. of varying skin tones, hair styles, facial hair/accessories, etc. In each case, we captured the optimal orientation as reported by our tool for each of our built-in targets, including those that the interface suggested were unattainable in the environment. Figures 1, 6, and 7 show many of these results and the complete set can be found in our supplemental materials. These examples include a range of different lighting conditions – e.g. indoors with windows and lamps, and outdoors during the day and night with sun, overcast, streetlights, partly covered, etc.

As shown in Figures 1 and 6, our tool is usually able to find good matches for the butterfly, loop (left and right), and split (left and right) lighting styles as these include a strong primary light source located above and either in front of the face (butterfly) or to one side of it (loop, split). In many of these




Figure 6. Portraits captured in different environments using our tools to achieve variety of target lighting styles. These lighting styles typically contain a strong primary light source in front of and above the face (butterfly) or to the side of the face (loop, split, rim). Our tool is able to emphasize the evenness of butterfly lighting and the key triangle in loop lighting. Rim lighting reproduces the strong highlight at the contour of the face in well-lit environments, but in the darker *library* and *quad night* scenes the light is not quite strong enough to produce a bright contour.

cases, the strongest light is the sun either directly illuminating the face or passing through a window; but in some indoor and nighttime scenes like *library* and *quad night*, the primary light is due to lamps. Our tool is generally able to find orientations that emphasize the key triangle on the subject’s cheek characteristic of loop lighting (i.e. in left loop, the cheek under the subject’s right eye is bright relative to the right side of the face and vice versa for right loop), and reproduce the evenness of butterfly lighting. Rim lighting places a strong primary light slightly behind the subject’s head to one side. In most of these environments, our tool finds a good orientation match for this lighting style (*arch walkway*, *gothic walkway*, *apartment*, *balcony*), but in darker environments like *library* and *quad night*, the primary light is not always strong enough to brightly illuminate the contour of the face.

Lighting styles that require two relatively strong lights facing one another (badger) or that require lighting from below (below, below side) are more difficult to achieve in most fixed lighting environments. Nevertheless as shown in Figure 7, our tool can sometimes find orientations that match these styles as we see for *academic building*, *columns night*, and *windowed hall*. In other cases like *modern quad* and *walkway night*, it matches the badger lighting on only one side of the face since there is only one strong main light. Similarly, in these two scenes, for the below lighting styles our tool suggests an orientation in which the subjects are lit a bit more from the side than from below.

Alternatively, users can use an additional light source to better achieve targets that require light from multiple directions.

For example, Figure 8 shows how our tool suggests placing a phone screen to capture targets that are unattainable in the fixed lighting of the environment. In the target combining loop and rim lighting (top row), the additional light generates the bright rim highlight on the right contour of the face opposite the side lit by the primary light in the environment. In the badger example (bottom row), the additional light acts as a second key light brightening the left side of the face and approximately mirroring the primary light in the environment. More results for placing an additional light source and using our painting interface to adjust lighting targets are in the supplemental materials.

Algorithmic Evaluation

To evaluate the robustness of our optimization algorithm, we consider how several parameters affect the image quality metric and the resulting optimal orientation angle.

Resolution of lighting and rendered image. To check how the resolution of the lighting and the image affect the image quality metric, we varied the lighting and image resolutions and generated image quality graphs for 192 (scene, target lighting) pairs as in Figure 9 (12 scenes and 16 lighting targets). For each such graph, we compute the difference between the optimal angle that maximizes image quality, and the corresponding optimal angle for the highest resolution lighting or image respectively. Averaged across all 192 (scene, target lighting) pairs, we obtain the following differences in optimal angles for lighting: 32×16 (diff: 6.15°), 64×32 (diff: 1.9°), 128×64 (diff: 0°); and for images: 50×50 (diff: 4.95°),

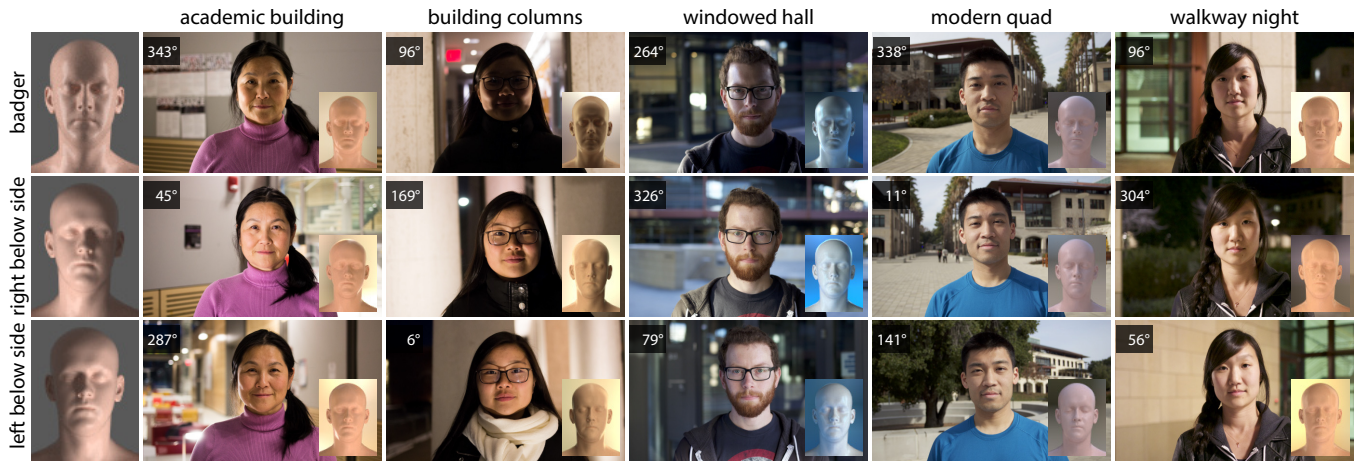


Figure 7. Some lighting styles are difficult to achieve in most fixed light environments. Badger requires two main lights facing one another with a dark center while the below lighting styles require strong illumination from below. However, our tool can sometimes find orientations that match such difficult target styles as we see here for the *academic building*, *columns night*, and *windowed hall*. For the other environments, *modern quad* and *walkway night*, it can only match the badger lighting on one side of the face and the below lighting appears a little more like side lighting than lighting from below.

100×100 (diff: 2.02°), 200×200 (diff: 0°). While the differences in angles are relatively small – within a few degrees – on close inspection it is possible to see small differences in some of the renderings. Overall, we find that a lighting resolution of 64×32 and an image resolution of 100×100 provides minimal loss in quality, while maintaining a relatively small sized transport matrix.

Distance between 360 camera and subject’s head. We checked the robustness of our algorithm to using lighting captured by a 360 camera at various distances from the subject’s head. As shown in Figure 10, lighting taken from 1 to 2 feet away from the subject produce similar image quality metric graphs, with maxima that are relatively close to one another. We find that the optimal angles differ from the optimal angle for the centered lighting by an average of 6.43° across all other lighting locations and lighting targets. These results suggest that as long as the lighting is relatively far away from the viewer, the environment map does not need to be taken exactly at the location of the subject.

Width of optimal image quality peak. Our optimization is designed to find the peak in the image quality metric graph. For each such peak, we compute the range of angles on either side of the peak for which the image quality metric remains within 3% of the maximum value. Averaged across all of our 192 (scene, target lighting) pairs, we find that this range is $17.05^\circ \sigma = 9.73^\circ$, or $\pm 8.5^\circ$ on either side of the peak. As long as photographers are within this range of the optimal orientation computed by our tool, they will produce an image with lighting that is very similar to the rendered target image (within 3% in terms of the image quality metric). As we show in User Evaluation, we find that users typically match the optimal orientation to within $\pm 9^\circ$ of the peak.

User Evaluation

To better understand how our tool helps users, we asked 28 people (20 amateur, 8 experienced) to complete three tasks. (1) First, we asked them to position and orient their subject (half used themselves as the subject and took selfies, while the other




Figure 8. Our tool can suggest how to place a phone screen to better achieve lighting targets such as these that are not easily attainable in the fixed lighting environment. For each of these portraits, the environment provides light on one side of the face (L only), and the additional light is able to provide the light on the other side of the face (L and L_{add}). The schematic shows how our interface conveys the optimal orientation for the additional light source relative to the face after orienting the subject to the environment.

half used a friend) to take a nice, well-lit portrait. (2) We next showed them images of 16 pre-defined studio lighting styles (Figure 3 top row) and asked them to use one as reference and take another well-lit portrait matching it. (3) Finally, we asked them to use our interface to select an attainable target lighting style and capture it.

After each task, we asked participants to rate the usefulness of the method they used to produce a well-lit portrait on a Likert scale running from 1-strongly disagree to 7-strongly agree. We found a significant difference (Friedman test) in usefulness between the three tasks [$\chi^2(2) = 18.6$, $p < 0.001$]. All three pairwise comparisons (Wilcoxon signed-rank test) were significant; in particular, participants rated our tool in task 3 as more useful ($\mu = 6.5$, $\sigma = 0.7$) than simply seeing a reference lighting style image as in task 2 ($\mu = 5.5$, $\sigma = 1.1$, $p < 0.025$) or having no guidance as in task 1 ($\mu = 4.5$, $\sigma = 1.2$, $p < 0.005$). We similarly asked them to rate how well-lit the images were in each condition on a 7-point Likert scale and found a significant difference [$\chi^2(2) = 7.0$, $p < 0.03$].

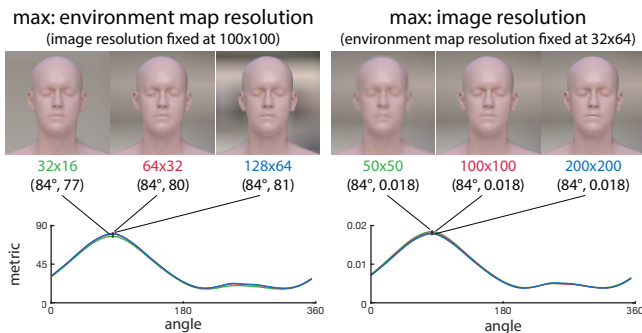


Figure 9. Effects of varying environment map and image resolution on image quality graphs for the *modern quad* scene and *butterfly* target. We have tested 3 lighting resolutions: 32×16 (green), 64×32 (red), 128×64 (blue), and 3 image resolutions: 50×50 (green), 100×100 (red), and 200×200 (blue). The corresponding image quality graphs sit on top of one another indicating very little difference between them (readers should zoom in to see the red and green curves). It is possible to see some small differences in the renderings, however for this (scene, lighting target) pair the optimal angles are the same across all the resolutions.

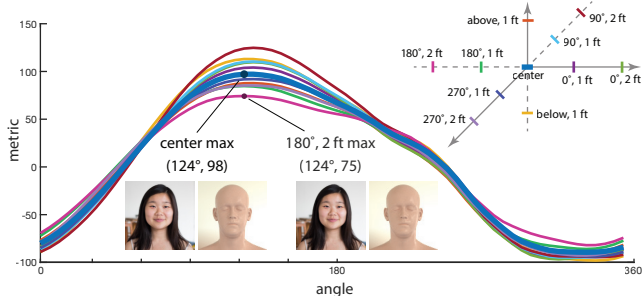


Figure 10. Effects of varying the position of lighting with respect to the subject’s head for the *art studio* scene and *left loop* target. We vary the position of the 360 camera around the subject’s head (located at the center of coordinate frame as in inset). Despite the differences in the graphs, in this case there is no difference in the optimal angle for the centered lighting versus the optimal angle for the lighting 2 feet away at 180° . We find that the optimal angles differ from the optimal angle for the centered lighting by an average of 6.43° across all other lighting locations and lighting targets.

The images produced using our interface were more well-lit ($\mu = 6.3$, $\sigma = 0.6$) than using only a reference target ($\mu = 5.5$, $\sigma = 1.3$, $p < 0.025$) or having no guidance ($\mu = 5.5$, $\sigma = 1.3$, $p < 0.025$).

In a post-study interview, our 20 amateur participants consistently said that our tool reduced the mental effort required to produce well-lit portraits. They attributed this reduction to the separation between choosing the target lighting style and the guidance to achieve the chosen style. One wrote, “*The thinking was concentrated in the stage when I chose the target... After I chose what I wanted, the tool made it pretty simple to get [it]*” (P7). We used a NASA-TLX questionnaire (excluding physical demand) to further assess differences in cognitive load between tasks 2 and 3 for amateur participants. We found a significant difference [Wilcoxon $V(2) = 64.5$, $p < 0.05$], suggesting that our tool makes it mentally easier to achieve desired lighting styles compared to simply using a reference target image. Directly asking the amateurs to rate the ease of orienting the subject on a 7-point Likert scale, we also found a significant difference between using a reference target in task 2 ($\mu = 3.8$, $\sigma = 1.6$) and using our interface in task 3 ($\mu = 6.1$, $\sigma = 0.9$) [$V(2) = 153$, $p < .001$].

We checked how well all the participants matched the orientation angle proposed by our tool in task 3, by comparing the background of their portraits to the corresponding environment map. We found that they were able to reorient to within an average of 9° ($\sigma = 7.5$), on either side of the target orientation. When manually orienting the subject for task 2, participants were only able to orient to within an average of 49° ($\sigma = 50.1$) [$t(27) = 4.3$, $p < .001$]. In this case, only 3 of 20 amateur participants and 5 of 8 experts were able to reorient to within 10° of the optimal target orientation.

Finally, interviews with our expert photographers revealed that our interface made them more confident and deliberate in their lighting choices. They said they were able to achieve better lighting “*without as much mental demand*” (P5). They similarly appreciated that “*it separated the choice of lighting from the rest of the photographic process*” (P14), encouraging them to “*branch out*” (P14) and “*explore*” (P2). A professional New York Times photographer said that our tool would also be helpful in scenarios where he has limited time in a location with a subject. See supplemental materials for additional user study details.

LIMITATIONS AND FUTURE WORK

While our approach for optimizing portrait lighting at capture-time can often suggest how to achieve a variety of lighting styles, it does have a few limitations that suggest directions for future work.

Fixed lighting. The main assumption of our approach is that lighting in the environment is fixed. As noted in Figure 7, some environments may not contain lights at the locations necessary to achieve certain lighting styles. However, our tool does mark lighting styles as unattainable in such situations so that the photographer is aware of the problem. Our tool can also help the photographer place an additional light source to achieve the target lighting. However this requires that either the photographer, the subject, a third person or a tripod hold the additional light in place.

Generic head and skin models. Using generic head and skin reflectance models allows our tool to use an efficient PRT-based approach to compute the optimal lighting orientation. As a result, our approach cannot account for individual differences in the geometric shape of the face or reflectance properties of the skin (due to skin tone, skin type, facial hair, etc.). We argue that the overall distribution of light and dark regions on the face is less affected by such individual differences. Nevertheless, these can have some effect on the appearance and shape of light and dark regions. We also do not consider how hair can occlude the light falling on the face. One approach may be to build a subject-specific head model using recent image-based facial geometry acquisition methods [4, 17] or RGB-D cameras. But even given the geometry, efficiently optimizing the lighting orientation without using our PRT-based approach is an open challenge.

Optimizing subject position. Our tool assumes distant lighting and only considers the orientation of the subject relative to the lighting environment. In many locations, changing the sub-

ject’s position, even by a few feet, can significantly alter the light falling on their face, especially when there are nearby occluders. Optimizing position in addition to orientation would enable much more control over the lighting, but likely require more significant modeling of geometry and light transport in the environment.

Alternative approach. An alternative approach for producing well-lit portraits, might involve first recording a video centered on the subject’s face as the subject rotates in place by 360 degrees, and then analyzing the video to find the frame that best matches each target lighting style. Our experience is that capturing such video at high-quality (without motion blur, while maintaining a nice facial expression, etc.) is very difficult. Nevertheless, assuming it is possible to capture such a video at high-quality, an open direction for future work is to develop an automatic algorithm for matching frames of the video to a target lighting style.

DISCUSSION

Our goals in designing a capture-time lighting optimizing tool were to provide a computationally simple and efficient technique that could be used with off-the-shelf hardware today and could be fully incorporated into a single camera system in the future. Our prototype implementation achieves these goals using two complementary ideas. The first idea is to treat a 360 camera as a light probe that can quickly capture the lighting in the environment. We believe that such light probes could eventually be built into cameras as just another sensor akin to light meters and focus sensors today. The second idea is to efficiently find the optimal lighting orientation by evaluating the appearance of a generic face model under different rotations of the captured lighting using a PRT-based approach. In fact, we show that evaluating each orientation requires computing a single dot product between vectors that are the size of the environment map, and that a low-resolution environment map of size 64×32 is sufficient for optimizing portrait lighting. One of the main implications of our work is that incorporating even a very low-resolution 360 light probe into a camera system would be useful for our application. Moreover, because our optimization is extremely efficient, requiring 0.19 milliseconds, it is well suited for running on a single integrated device. This would further eliminate the cumbersome switching between operating a 360 camera and a laptop and a primary portrait camera as in our prototype.

A concern when using PRT-based techniques is that the size of the transport matrix \mathbf{T} (246 MB uncompressed in our case) can be prohibitive for devices with limited memory. However, if the target appearance weighting functions \mathbf{f}^t are known a priori, we only have to store the pre-integrated lighting domain vectors \mathbf{F}^t to evaluate the image quality, and it is unnecessary to store \mathbf{T} . In our application, we use \mathbf{T} to render the PRT matches for the optimal orientation angles we identify for each target. These visualizations can help the user understand what the face will look like in the target orientation, but they are not essential for using our system. We also use the transport matrix \mathbf{T} to pre-integrate the target weighting function when the user manually adjusts it via our painting interface. In practice however, we believe most casual users would focus

on the built-in set of target lighting styles and would not need the additional control afforded by the painting interface. Thus, it may be possible to offer most of the benefits of our approach without the transport matrix in the end-user interface.

CONCLUSION

Lighting is a crucial element of portrait photography. But choosing how to orient a subject with respect to the environment has traditionally required paying careful attention to the available light, as well as understanding of how that light might fall on the face to produce different looks. We have demonstrated a capture-time tool that uses a 360 camera as a light probe and suggests how to orient the subject with respect to the surrounding environment to best achieve a user-specified target appearance. We believe that such a tool can make it easier for photographers to consider a variety of different looks that may be possible in an environment.

ACKNOWLEDGEMENTS

We thank Parastoo Abtahi, Ali Alkhatib, Dorothy Ahn, Cody Coleman, Matthew Cong, Griffin Dietz, Ilene E, Weinan E, Dan Fu, Hongjun Li, Winnie Lin, Sean Liu, Josh Newman, Ante Qu, Ed Quigley, Anh Truong, Xinwei Yao, and Yeshuang Zhu for posing for our results photos. We additionally thank Anh Truong for her help with figures and supplemental materials, and Matthew Cong for his help with the video. Finally we thank James Landay and Pat Hanrahan for their valuable advice and feedback. Our research is supported by Brown Institute for Media Innovation and the Microsoft Research Dissertation Grant.

REFERENCES

- [1] Rolf Adelsberger, Remo Ziegler, Marc Levoy, and Markus Gross. 2008. *Spatially adaptive photographic flash*. ETH, Department of Computer Science.
- [2] Aseem Agarwala, Mira Dontcheva, Maneesh Agrawala, Steven Drucker, Alex Colburn, Brian Curless, David Salesin, and Michael Cohen. 2004. Interactive digital photomontage. In *ACM Transactions on Graphics (TOG)*, Vol. 23. ACM, 294–302.
- [3] David Akers, Frank Losasso, Jeff Klingner, Maneesh Agrawala, John Rick, and Pat Hanrahan. 2003. Conveying shape and features with image-based relighting. In *Proceedings of the 14th IEEE Visualization 2003 (VIS'03)*. IEEE Computer Society, 46.
- [4] Hadar Averbuch-Elor, Daniel Cohen-Or, Johannes Kopf, and Michael F. Cohen. 2017. Bringing Portraits to Life. *ACM Trans. Graph.* 36, 6, Article 196 (Nov. 2017), 13 pages. DOI : <http://dx.doi.org/10.1145/3130800.3130818>
- [5] Soonmin Bae, Aseem Agarwala, and Frédo Durand. 2010. Computational rephotography. *ACM Trans. Graph.* 29, 3 (2010), 24–1.
- [6] Patrick Baudisch, Desney Tan, Drew Steedly, Eric Rudolph, Matt Uyttendaele, Chris Pal, and Richard Szeliski. 2005. Panoramic viewfinder: providing a real-time preview to help users avoid flaws in panoramic pictures. In *Proceedings of the 17th Australia conference*

- on *Computer-Human Interaction: Citizens Online: Considerations for Today and the Future*. Computer-Human Interaction Special Interest Group (CHISIG) of Australia, 1–10.
- [7] Herbert Bay, Tinne Tuytelaars, and Luc Van Gool. 2006. Surf: Speeded up robust features. *Computer vision–ECCV 2006* (2006), 404–417.
- [8] Edward S Bomback. 1971. *Manual of photographic lighting*. Fountain Press, Limited.
- [9] Adrien Bousseau, Emmanuelle Chapoulie, Ravi Ramamoorthi, and Maneesh Agrawala. 2011. Optimizing environment maps for material depiction. In *Computer graphics forum*, Vol. 30. Wiley Online Library, 1171–1180.
- [10] Ivaylo Boyadzhiev, Jiawen Chen, Sylvain Paris, and Kavita Bala. 2016. Do-it-yourself lighting design for product videography. In *Computational Photography (ICCP), 2016 IEEE International Conference on*. IEEE, 1–9.
- [11] Ivaylo Boyadzhiev, Sylvain Paris, and Kavita Bala. 2013. User-assisted image compositing for photographic lighting. *ACM Trans. Graph.* 32, 4 (2013), 36–1.
- [12] Ewen Cheslack-Postava, Nolan Goodnight, Ren Ng, Ravi Ramamoorthi, and Greg Humphreys. 2007. 4D Compression and Relighting with High-resolution Light Transport Matrices. In *Proceedings of the 2007 Symposium on Interactive 3D Graphics and Games (I3D '07)*. ACM, New York, NY, USA, 81–88. DOI: <http://dx.doi.org/10.1145/1230100.1230115>
- [13] Wayne W Daniel. 1990. Kruskal–Wallis one-way analysis of variance by ranks. *Applied Nonparametric Statistics* (1990), 226–234.
- [14] Paul Debevec, Tim Hawkins, Chris Tchou, Haarm-Pieter Duiker, Westley Sarokin, and Mark Sagar. 2000. Acquiring the reflectance field of a human face. In *Proc. of SIGGRAPH*. 145–156.
- [15] Marc Ebner. 2007. *Color constancy*. Vol. 6. John Wiley & Sons.
- [16] Elmar Eisemann and Frédo Durand. 2004. Flash photography enhancement via intrinsic relighting. *ACM transactions on graphics (TOG)* 23, 3 (2004), 673–678.
- [17] Ohad Fried, Eli Shechtman, Dan B Goldman, and Adam Finkelstein. 2016. Perspective-aware manipulation of portrait photos. *ACM Transactions on Graphics (TOG)* 35, 4 (2016), 128.
- [18] Marc-André Gardner, Kalyan Sunkavalli, Ersin Yumer, Xiaohui Shen, Emiliano Gambaretto, Christian Gagné, and Jean-François Lalonde. 2017. Learning to Predict Indoor Illumination from a Single Image. *arXiv preprint arXiv:1704.00090* (2017).
- [19] Stefan Gumhold. 2002. Maximum entropy light source placement. In *Visualization, 2002. VIS 2002. IEEE*. IEEE, 275–282.
- [20] Dong Guo and Terence Sim. 2009. Digital face makeup by example. In *Computer Vision and Pattern Recognition*. IEEE, 73–79.
- [21] Barry Huggins. 2005. *Creative photoshop lighting techniques*. Lark Books.
- [22] Fil Hunter, Steven Biver, and Paul Fuqua. 2015. *Light Science & Magic: An Introduction to Photographic Lighting*. CRC Press.
- [23] Insta360. 2018. Insta360 Nano. (2018). <https://www.insta360.com/product/insta360-nano/>
- [24] Ira Kemelmacher-Shlizerman. 2016. Transfiguring portraits. *ACM Transactions on Graphics (TOG)* 35, 4 (2016), 94.
- [25] William B Kerr and Fabio Pellacini. 2009. Toward evaluating lighting design interface paradigms for novice users. In *ACM Trans. on Graphics (TOG)*, Vol. 28. ACM, 26.
- [26] Chang Ha Lee, Xuejun Hao, and Amitabh Varshney. 2006. Geometry-dependent lighting. *IEEE TVCG* 12, 2 (2006), 197–207.
- [27] Christian Lessig, Tyler de Witt, and Eugene Fiume. 2012. Efficient and accurate rotation of finite spherical harmonics expansions. *J. Comput. Phys.* 231, 2 (2012), 243–250.
- [28] Tommer Leyvand, Daniel Cohen-Or, Gideon Dror, and Dani Lischinski. 2008. Data-driven enhancement of facial attractiveness. In *ACM Transactions on Graphics (TOG)*, Vol. 27. ACM, 38.
- [29] Qifan Li and Daniel Vogel. 2017. Guided Selfies using Models of Portrait Aesthetics. In *Proceedings of the 2017 Conference on Designing Interactive Systems*. ACM, 179–190.
- [30] Zicheng Liu, Ying Shan, and Zhengyou Zhang. 2001. Expressive expression mapping with ratio images. In *Proc. of SIGGRAPH*. ACM, 271–276.
- [31] Barbara London, Jim Stone, and John Upton. 2011. *Photography*. Prentice Hall.
- [32] Stephen R Marschner and Donald P Greenberg. 1997. Inverse lighting for photography. In *Color and Imaging Conference*, Vol. 1997. Society for Imaging Science and Technology, 262–265.
- [33] Dave Montizambert. 2003. *Creative lighting techniques for studio photographers*. Amherst Media, Inc.
- [34] Lukas Murmann, Abe Davis, Jan Kautz, and Frédo Durand. 2016. Computational bounce flash for indoor portraits. *ACM Transactions on Graphics (TOG)* 35, 6 (2016), 190.
- [35] Markus Neuhäuser. 2011. Wilcoxon–mann–whitney test. In *International encyclopedia of statistical science*. Springer, 1656–1658.

- [36] Ren Ng, Ravi Ramamoorthi, and Pat Hanrahan. 2003. All-frequency shadows using non-linear wavelet lighting approximation. In *ACM Transactions on Graphics (TOG)*, Vol. 22. ACM, 376–381.
- [37] Ren Ng, Ravi Ramamoorthi, and Pat Hanrahan. 2004. Triple product wavelet integrals for all-frequency relighting. *ACM Transactions on Graphics (TOG)* 23, 3 (2004), 477–487.
- [38] Derek Nowrouzezahrai, Patricio Simari, and Eugene Fiume. 2012. Sparse zonal harmonic factorization for efficient SH rotation. *ACM Transactions on Graphics (TOG)* 31, 3 (2012), 23.
- [39] Pieter Peers, Naoki Tamura, Wojciech Matusik, and Paul Debevec. 2007. Post-production facial performance relighting using reflectance transfer. In *ACM Transactions on Graphics (TOG)*, Vol. 26. ACM, 52.
- [40] Fabio Pellacini. 2010. envyLight: an interface for editing natural illumination. *ACM Transactions on Graphics (TOG)* 29, 4 (2010), 34.
- [41] Georg Petschnigg, Maneesh Agrawala, Hugues Hoppe, Richard Szeliski, Michael Cohen, and Kentaro Toyama. 2004. Digital Photography with Flash and No-flash Image Pairs. *ACM Trans. Graph.* 23, 3 (Aug. 2004), 664–672. DOI : <http://dx.doi.org/10.1145/1015706.1015777>
- [42] Matt Pharr, Wenzel Jakob, and Greg Humphreys. 2016. *Physically based rendering: From theory to implementation*. Morgan Kaufmann.
- [43] Erik Reinhard, Wolfgang Heidrich, Paul Debevec, Sumanta Pattanaik, Greg Ward, and Karol Myszkowski. 2010. *High dynamic range imaging: acquisition, display, and image-based lighting*. Morgan Kaufmann.
- [44] Tammy Riklin-Raviv and Amnon Shashua. 1999. The quotient image: Class based recognition and synthesis under varying illumination conditions. In *Computer Vision and Pattern Recognition*, Vol. 2. IEEE, 566–571.
- [45] Szymon Rusinkiewicz, Michael Burns, and Doug DeCarlo. 2006. Exaggerated shading for depicting shape and detail. In *ACM Transactions on Graphics (TOG)*, Vol. 25. ACM, 1199–1205.
- [46] Ram Shacked and Dani Lischinski. 2001. Automatic lighting design using a perceptual quality metric. In *Computer graphics forum*, Vol. 20. Wiley Online Library, 215–227.
- [47] Xiaoyong Shen, Aaron Hertzmann, Jiaya Jia, Sylvain Paris, Brian Price, Eli Shechtman, and Ian Sachs. 2016. Automatic portrait segmentation for image stylization. In *Computer Graphics Forum*, Vol. 35. Wiley Online Library, 93–102.
- [48] YiChang Shih, Sylvain Paris, Connelly Barnes, William T Freeman, and Frédo Durand. 2014. Style transfer for headshot portraits. (2014).
- [49] Peter-Pike Sloan, Jan Kautz, and John Snyder. 2002. Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. In *ACM Transactions on Graphics (TOG)*, Vol. 21. ACM, 527–536.
- [50] sRGB Wikipedia Page. 2017. sRGB — Wikipedia, The Free Encyclopedia. (2017). <https://en.wikipedia.org/w/index.php?title=sRGB&oldid=816935716> [Online; accessed 14-January-2018].
- [51] Manohar Srikanth, Kavita Bala, and Frédo Durand. 2014. Computational rim illumination with aerial robots. In *Proceedings of the Workshop on Computational Aesthetics*. ACM, 57–66.
- [52] Inc. The MathWorks. 2007-2014a. makehdr. (2007-2014). <https://www.mathworks.com/help/images/ref/makehdr.html>
- [53] Inc. The MathWorks. 2007-2014b. rgb2lin. (2007-2014). <https://www.mathworks.com/help/images/ref/rgb2lin.html>
- [54] Wai-Shun Tong, Chi-Keung Tang, Michael S Brown, and Ying-Qing Xu. 2007. Example-based cosmetic transfer. In *Computer Graphics and Applications, 2007. PG'07. 15th Pacific Conference on*. IEEE, 211–218.
- [55] Philip HS Torr and Andrew Zisserman. 2000. MLESAC: A new robust estimator with application to estimating image geometry. *Computer Vision and Image Understanding* 78, 1 (2000), 138–156.
- [56] Romain Vergne, Romain Pacanowski, Pascal Barla, Xavier Granier, and Christophe Schlick. 2009. Light warping for enhanced surface depiction. In *ACM Transactions on Graphics (TOG)*, Vol. 28. ACM, 25.
- [57] Zhen Wen, Zicheng Liu, and Thomas S Huang. 2003. Face relighting with radiance environment maps. In *Computer Vision and Pattern Recognition*, Vol. 2. IEEE, II–158.