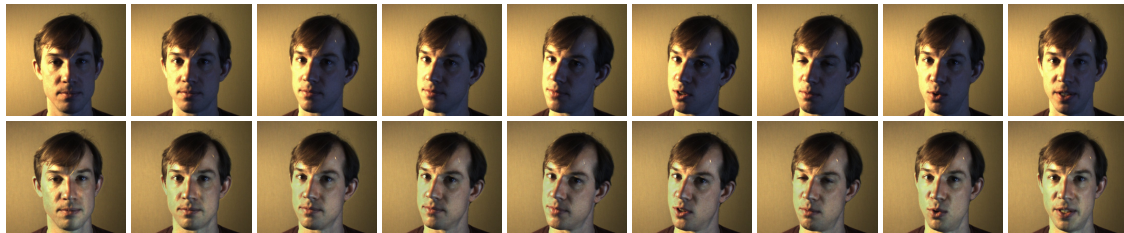


Video Relighting Using Infrared Illumination

Oliver Wang and James Davis and Erika Chuang and Ian Rickard and Krystle de Mesa and Chirag Dave

Department of Computer Science
University of California, Santa Cruz



Abstract

Inappropriate lighting is often responsible for poor quality video. In most offices and homes, lighting is not designed for video conferencing. This can result in unevenly lit faces, distracting shadows, and unnatural colors. We present a method for relighting faces that reduces the effects of uneven lighting and color. Our setup consists of a compact lighting rig and a camera that is both inexpensive and inconspicuous to the user. We use unperceivable infrared (IR) lights to obtain an illumination bases of the scene. Our algorithm computes an optimally weighted combination of IR bases to minimize lighting inconsistencies in foreground areas and reduce the effects of colored monitor light. However, IR relighting alone results in images with an unnatural ghostly appearance, thus a retargeting technique is presented which removes the unnatural IR effects and produces videos that have substantially more balanced intensity and color than the original video.

Categories and Subject Descriptors (according to ACM CCS): I.3.6 [Computer Graphics]: Methodology and Techniques

1. Introduction

While video conferencing is an increasingly used method of communication, problems exist that prevent a more wide scale adoption of the technology. Some of the difficulties, such as compression and latency artifacts, dominate the current user experience and are due primarily to networks with insufficient bandwidth. Fortunately, current technology trends will provide sufficient bandwidth to most home and office users within a few years. Looking forward to that point in time, we can begin to ask what quality issues will remain? This paper addresses one of the most important of these, relighting of faces.

Existing lighting setups in offices and homes have not been designed properly for photography, and therefore lead to poorly lit faces and distracting images. These lighting effects not only cause aesthetically displeasing faces, but can

also make communication more difficult. Non-verbal communication, especially in the form of facial expressions, is thought to convey meaning more quickly than speech [VP06]. Bright office windows, harsh desk lights, backlighting and monitor glow all make the perception of these facial expressions difficult. Movie and TV studios have lighting crews precisely because everyday environments do not provide adequate lighting.

Portrait photographers routinely create proper lighting in studio setups. The most flattering light is generally thought to be off center, white, and very diffuse [Gre04]. This has the effect of de-emphasizing facial features and lessening skin imperfections. Furthermore, the human visual system is very sensitive to color discrepancies across human faces. Therefore, using lights with a non-white hue is usually avoided.

An obvious approach to fixing lighting would be to sim-

ply construct studio quality lighting setups. There are some commercial services that do attempt to solve the lighting problems of video conferencing using this approach. However, these services are often prohibitively expensive. For example, the Halo system, developed at HP, costs around \$500,000 [HP]. In this case, the lighting and materials in a room with no windows are chosen carefully so as to avoid any problems with uneven lighting. This solution is clearly impractical to implement in every home and cubicle.

There have been many attempts to digitally relight images, but few are suitable for desktop video-conferencing. Existing solutions have often addressed general relighting, which can recreate arbitrary lighting conditions. These solutions have made use of large complex acquisition devices, careful calibration, offline 3D model generation, non-real-time algorithms, and intrusively modulated visible lighting. In contrast we present a method which does not provide general relighting capability, but rather a very specific kind of lighting improvement. This simplification in output requirement comes at the cost of stricter implementation requirements: a compact uncalibrated desktop device, real-time capture and processing, and no visually disturbing lighting change.

We attempt to emulate lighting conditions that closely match desired portrait photography lighting and allow for an easier interpretation of facial expressions. Our method also removes unnatural hues cast by colored monitor light. We relight a video sequence using additional information acquired by an inexpensive, uncalibrated IR lighting setup. We chose to use IR lighting to compute an illumination basis as it is invisible to the human optical system and therefore can be placed in interactive environments without distracting the user. Because room lighting can change during the course of a video, our method computes an optimal relighting based on the existing lighting at each frame. We generate relit images using a weighted linear combination of IR basis images. However, because we are using IR illumination, we cannot directly insert these relit images into the original video as they have a different appearance than what we expect to see in the visible spectrum. Instead, we take advantage of the similarity of the shape of an object's BRDF across different wavelengths to retarget the information into the visible domain. We also reduce the colored glow of monitors by solving for different ideal images separately across the RGB channels.

The primary contribution of our work is a novel method for relighting faces which is non-intrusive, compact enough for office desks, requires no user calibration, and is sufficiently simple for real-time implementation. By using these ratio images, we are able to correct both intensity and hue inconsistencies in the original video using unobtrusive IR information while still preserving its visible domain appearance. The resulting video has substantially reduced overexposed regions, shadowing, and color imbalance.

2. Related Work

Image relighting is a widely investigated topic used in areas as diverse as face recognition and movie post-production. Here we discuss the relighting methods most closely related to those presented in this work.

Geometric Based Relighting A traditional graphics based solution to generating new images is to render the object under artificial lighting conditions using a 3D representation of the scene and some estimation of its reflectance properties. There is a body of work which attempts to capture the complicated reflectance characteristics of human skin and uses it in combination with 3D scanners to render a human face [WMP*06] [MLTG99] [YDMH99] [FLS05]. Additional methods exist that use only a rough 3D model to generate new lighting conditions, and then apply that information to adjust the lighting in images. This reduces the requirement for extremely complex geometry and the necessity for accurate reflectance functions. Other methods generate geometry at run time by integrating normals acquired from images of varying lighting conditions [LPMM05] [GBK99] [NN04]. Other methods use a prior 3D model acquired by some kind of 3D scanner [MG97] [PSS99] [LKG*03] or some use a simple geometric representation such as an ellipse or generic face model fit to an image [BGG*01] [WLH03]. These methods composite the rendered light from these geometric approximations with the original video.

The primary difficulty with these approaches is that they all require some kind of 3D geometry. This geometry must either be approximated by image based methods such as integrating normals, known beforehand by using a 3D scanner, or assumed to be some preexisting model. It is difficult to acquire accurate geometry from a purely image based approach, and for our purposes, it is unreasonable to require 3D head scans for each separate individual that we wish to relight. If the precision of the model is reduced, alignment errors can cause artifacts when relighting the images, a problem complicated by non-rigid deformation in the face.

Image Based Relighting Another way to create new lighting conditions for a scene is to use image based relighting methods. These methods do not use a geometric representation of the object, but rather use one or multiple images to generate new scenes. Debevec et al. [DHT*00] introduced a complex gantry capable of capturing a dense basis space of images which were combined to produce new lighting conditions. This gantry was extended to very high frame rates for real time video applications [WGT*05]. Further work showed that you could reduce the complexity of capturing a suitable basis for image relighting by using projected light patterns and solving for an optimal lighting basis for each scene [MNNB05]. Other work captured lighting information of objects, stored the information as polynomials and then interpolated between them to generate new images [MGW01]. These methods use well defined basis

spaces and therefore require complex capture setups that have to be precisely calibrated. In addition, all of these methods relight by using modulated lighting in the visible spectrum, which would be distracting for the user in a real time application.

Image Combinations A different way to solve the lighting problem would be to combine separate parts from multiple images with varying lighting conditions to form one final image. This method has been applied both to user guided lighting design as well as more unconstrained image recombination [ALK*03] [ADA*04]. However, these applications often rely on user annotation of the image, and require roughly normalized lighting conditions across images to prevent visible seams. Both of these requirements would be unreasonable to require for our application.

Ratio Images The idea of using ratio images (also called quotient images) to encode the relation between two domains is not new. Liu et al. [LSZ01] used ratio images to map the expression of one person onto another. Marschner et al. [MG97] used ratio images to preserve geometric details not present in relit graphics renderings. Rinklin-Ravin et al. [RRS01] used the ratio image between different people to apply lighting conditions to new faces assuming Lambertian reflectance and fixed viewpoint. Recently, work was presented to use ratio images to map the illumination from one face to another [PTMD07]. This method allowed for the capture of a complete reflectance field for one face, and apply it on new faces. Our paper approaches a different problem of capturing the reflectance field in a different spectrum, in real time. This allows for more unpredicted changes, as we do not require that pre-captured data must account for all possible poses.

3. Research and Methods

Our main motivations in relighting video conferencing scenes are to reduce the effects of harsh and uneven lighting and to correct for unnatural colored light cast by monitor glow. In order to gain additional information about a scene, we capture a series of IR images along with the original color video. We use this IR basis to compute a well lit IR image and then perform a retargeting into the visible domain.

3.1. Image Acquisition

Our setup consists of an array of IR lights, two cameras, a beam splitter, and a computer. We use two PointGray DragonFly 2 Cameras, one that is sensitive to IR and visible wavelengths and one that is sensitive to just the visible spectrum. The IR LEDs are inset in bars on both sides of the monitor. This setup is shown in Figure 1 with a top-down view of the beam splitter and cameras inset in the image. Our IR LEDs are controlled by a custom control board with a serial link to our computer. We used Osram surface mount IR LEDs, which have a radiant intensity of 200mW/sr, 20.0

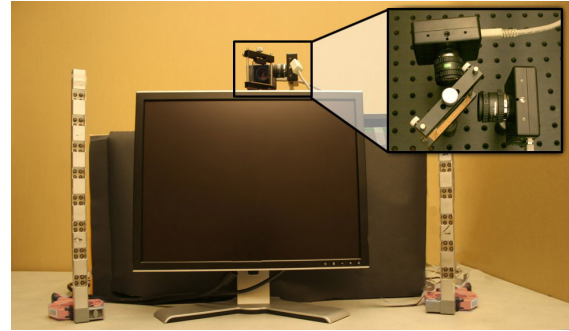


Figure 1: *Our Infrared Lighting Setup. A beam splitter and two cameras are situated above the monitor and IR LEDs are mounted in the columns along the sides. The inset shows a top down image of our cameras and beam splitter*

nanosecond rise and fall time, 940.0nm peak emission wavelength, and 60.0° half angle. We note that we are operating well within the safety guidelines for near-IR light set by the European Standard EN 60825-1 [Eur]. In order to make the IR LED's less directional, we place diffusers over the lights. These have been removed in the picture for the purpose of clarity. We use a cold mirror as our beam splitter which reflects visible spectrum light while passing light with a wavelength longer than 780nm. We have experimented with using different numbers of basis images for each frame, and have settled on using eight. Our color camera runs at a rate of 15fps while our IR camera captures eight frames for each color frame, running at 120fps. We perform a calibration step to align the two cameras by warping the images with an affine transformation.

Many lighting sources emit light both in the IR and visible spectrum. We must therefore preprocess the IR images by removing all ambient IR room lighting, which we capture as one of our basis images. Figure 2 shows some of these input images. The brightness has been increased for viewing purposes. In our acquisition setup, we used standard fluorescent overhead office lights as well as fluorescent and incandescent desk lamps in various positions.

Because we cannot control the intensity of the room lighting, we must deal with the possibility of a low signal-to-noise ratio in our IR images. We smooth image noise using a joint bilateral filter as described in Petschnigg et al. [PSA*04]. We use the color image to calculate the weights for the bilateral filter, which allows us to smooth the IR images without crossing edges. Figure 3 shows before and after the denoising step.

Our goal is to minimize lighting variation across the face region of an image. We use a foreground mask to determine which parts of the image we use in computing the image weights. Because we already have an active illumination setup and a controlled scene, we can compute a rough fore-

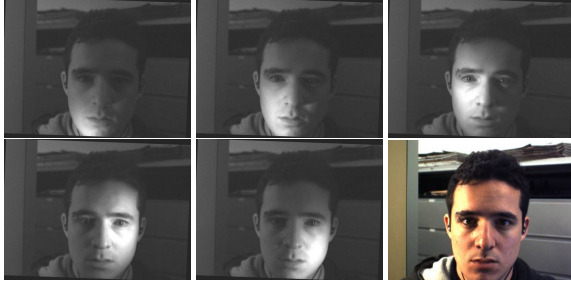


Figure 2: A subset of the images captured at each frame showing some IR lighting conditions and the color image.

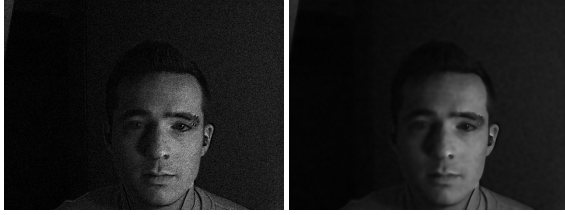


Figure 3: Our noise removal step. We use a bilateral filter, which yields smooth results without blurring across edges.

ground mask using the $\frac{1}{R^2}$ falloff of light, similar to the technique shown in Flash Matting [SLKS06]. In general, objects closer to our IR lights will be illuminated more than those farther away. We look at a difference between two IR images, one with all the lights on, and the other with all the lights off. Areas where the intensity change is greatest are assumed to be foreground. We can see the results of a simple threshold on this difference in Figure 4. Because the map is only used for computing weights, which will then be applied evenly to the whole image, it is unnecessary to have pixel perfect edges, which can be a difficult problem.

3.2. Image Relighting

It has been shown that realistic looking relighting can be accomplished by combining lighting basis images [DHT*00]



Figure 4: A visible spectrum intensity image and a foreground region mask created by thresholding the difference between images with the IR lights on and off.

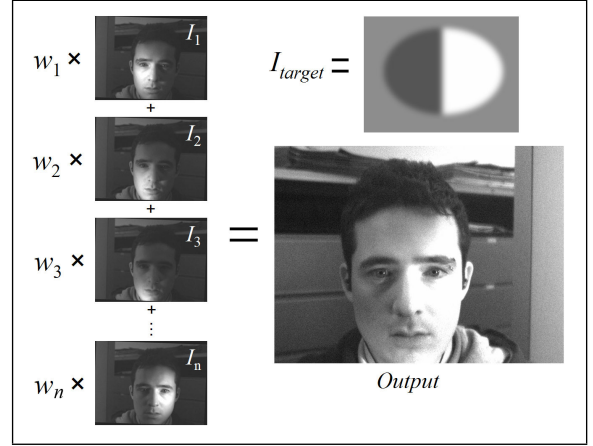


Figure 5: A linear combination of lighting bases is computed so that the output image closely matches the target, I_{target} . Notice that the output is a valid image of a face with approximately the desired lighting, even though the supplied target image does not resemble a face.

[FBS05] [WGT*05]. Existing work often relies on carefully calibrated and a largely complete basis space. We use a similar technique but with a much sparser basis space. We show that even with this sparse and uncalibrated lighting set, we can still find a relit image that closely matches target lighting conditions. We compute a weight vector for the images using the equation:

$$A\bar{w} \approx I_{target} \quad (1)$$

Where A is a matrix with each column corresponding to an image in the basis space, I_{target} is the target image to solve for, and \bar{w} is the vector of image weights. After we compute \bar{w} , we can create the new image using the following equation:

$$I_{new} = \sum_{i=1}^n w_i \cdot I_i \quad (2)$$

Where n is the number of images in our basis, I_k corresponds to the image under the k th lighting condition, and w_i is the i th image in the weight vector.

Figure 5 shows graphically how this step works. It is important to note that while we cannot achieve the target image exactly, the output image appears to provide convincing lighting configurations. This technique is not new, and has been used frequently in relighting research. Often, the target images are constructed manually to obtain desirable lighting effects [ALK*03] [MTB*06] [AD04].

We can use any image I_{target} and solve for an approxima-



Figure 6: The visible and IR spectrum reflectance of an eye. The skin and iris can appear lighter in IR, giving the IR image an unnatural ghostly appearance.



Figure 7: The original visible spectrum image with uneven lighting and its intensity channel are shown on top. An IR image with more balanced lighting is shown on the bottom left. Notice the ghostly appearance in the IR intensity image. If this image is merely substituted into the intensity channel of the visible spectrum image, obvious artifacts occur, as shown on the bottom right.

tion of it as a combination of our bases. Given that we have poor lighting in the visible spectrum and arbitrary control over our IR illumination, a simple approach may be to just replace the poorly lit intensity channel with a well lit IR image. However, this does not produce a satisfactory result as we are not used to seeing objects in the IR spectrum. This is especially true of human faces, which can appear ghostly when viewed under IR illumination, as shown in Figure 6. Specifically, the iris and skin appear lighter and unrealistic. We show how it would look if we just replaced the visible spectrum with information from the IR bases in Figure 7.

3.3. Spectral Consistency

Objects appear unnatural when viewed in the IR spectrum because many materials do not have the same reflectance across different spectral bands. We hypothesize however, that while objects may have different relative intensities at different wavelengths, the general shape of their reflectance

functions is more likely to be closely related across these bands. In other words, if an arbitrary change in the scene illumination makes an object brighter in the IR spectrum by some amount, it is also likely to make the object brighter in the visible spectrum by a similar amount. We should note that this is not the same observation as the linearity of light, as we are also accounting for changes in the position of lights not only their intensity.

We verify this hypothesis empirically by analyzing data from numerous different scenes. If we look at the intensity of a pixel in the IR spectrum and the same pixel in the visible spectrum, we notice that there is not an obvious relationship, some objects are bright in IR and dark in the visible spectrum while other objects have opposite characteristics. This can be seen in the top plot in Figure 8. However, because this plot does not show the density of the points very well, it is hard to interpret this data. Therefore, we plot the same data but instead of looking at the relative intensities, we look at a histogram of the angle θ of each point in polar coordinates. If the intensity were the same in both spectral bands equally, we would expect all points to lie on the line $y = x$, (or in other words, $\theta = 45^\circ$). The middle figure in Figure 8 shows that this is not the case. We then plot the angle θ of the ratio of a pixel's intensity between two scenes with the same viewpoint but different lighting directions in both the IR and visible spectrums. The bottom plot in Figure 8 shows the histogram of a series of points where each point represents a pixel whose coordinates are: $(\frac{IR_1}{IR_2}, \frac{Visible_1}{Visible_2})$. We can see that in this plot, there is in fact a very strong correspondence, with most points laying along the $y = x$, or 45° line. The intuition is that for most objects, the *shape* of their BRDFs are not nearly as dependent on wavelength as their relative intensities are.

3.4. IR Retargeting

Based on the results of our analysis, we would like to relight our scene not by substituting a well lit IR image directly into our final image, but by using it as an indicator of how much we want to change our visibly lit image. If an ideally lit IR image can be obtained from a poorly lit IR image by the equation $I_{ideal} = (\% \text{ change IR}) * I_{observed}$. Then Section 3.3 indicates that a visible light image can also be corrected via $V_{ideal} \approx (\% \text{ change IR}) * V_{observed}$. Combining these equations gives us the "ratio equation" 3.

$$\frac{I_{ideal}}{I_{observed}} \approx \frac{V_{ideal}}{V_{observed}} \quad (3)$$

We are faced with a problem where we would like to compute V_{ideal} , the ideal intensity of the visible spectrum. The only information that we have is the measured intensity of the visible spectrum ($V_{observed}$), and an IR basis space $I_1 \dots I_n$. In order to find V_{ideal} we must approximate both $I_{observed}$ and I_{ideal} .

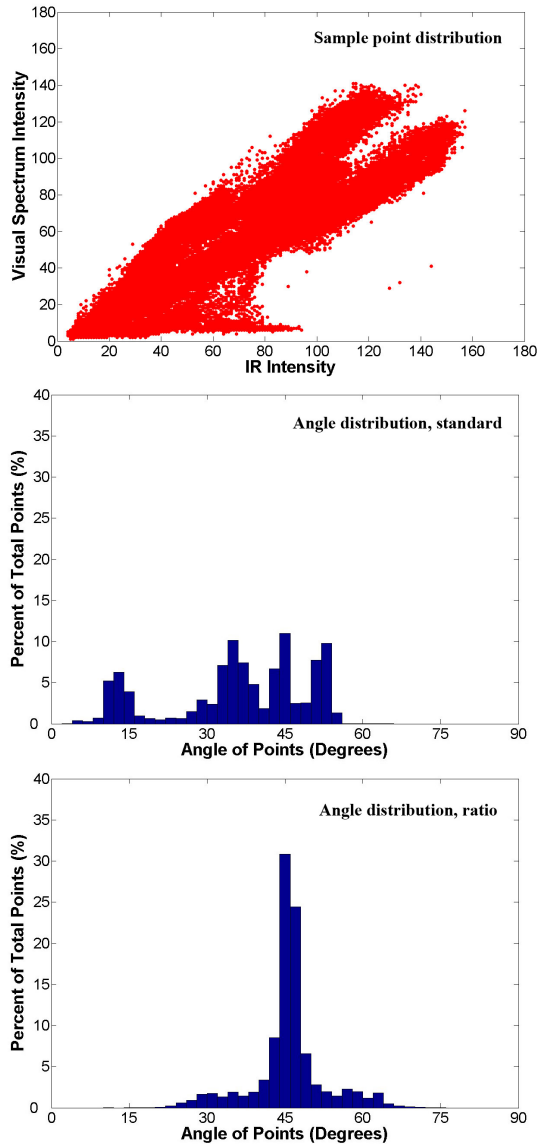


Figure 8: Object albedo in the visible and IR spectrums are not necessarily correlated. Each point in the top plot represents a single pixel's IR intensity plotted against its visible intensity. Notice that these values do not fall along a straight line. The middle plot shows a different view of the same data, a histogram of the θ value for each point in polar coordinates. Note that the data is not well correlated so there is a wide distribution of θ values. In contrast, ratio images are well correlated between the IR and visible spectrum. The bottom plot shows a histogram of the angles formed by the points $(IR_1/IR_2, Vis_1/Vis_2)$ where lighting conditions 1 and 2 are images of the same scene with the lights placed in different physical locations. These points show a strong correlation and lie close to the 45° line.



Figure 9: The ratio image as formed by dividing the ideal image with balanced lighting, I_{ideal} , by the image with observed lighting, $I_{observed}$.

Ideally, we would like the image $I_{observed}$ to be an image with the same position and intensity of lights as in $V_{observed}$ only in IR. However, we cannot capture this image, as it is not guaranteed that all the lights in the scene will have the same brightness in IR as in the visible spectrum. Notably, fluorescent lights tend to have very little intensity in the IR spectrum, while incandescent lights are roughly equivalent in the IR and visible bands. Therefore, we project the image $V_{observed}$ onto the space of IR images defined by our basis. We do this by simply solving for a combination of IR bases that most closely matches $V_{observed}$ using the method described in Section 3.2.

We then must then compute I_{ideal} . Our goal is to minimize lighting variation across the face. Therefore, we define I_{ideal} to be the projection of a constant intensity image in the foreground face region into our IR bases. This is similar to the sketch guided lighting optimization in the work in Mohan et al. [MTB*06]. Here, this constant image serves as the target image in our solution. We now have the ratio image, which we achieve by dividing these two IR images. The formulation and ratio image is shown in Figure 9. We can now use equation 3 to approximate V_{ideal} by multiplying the ratio image with $V_{observed}$. The result of our algorithm on one frame of our video is shown in Figure 10. Notice that the intensity difference is greatly improved in the resulting image.

3.5. Chroma Correction

We note that in addition to the uneven intensity of lighting in the face, there may also be significant chroma variation due to colored light sources in the scene. We can correct for this uneven chroma by applying our technique separately to the different color channels. We first determine a suitable face color heuristically by finding the median color in the face region. We can then solve for each color channel separately targeting the median face color in that channel. By adjust-

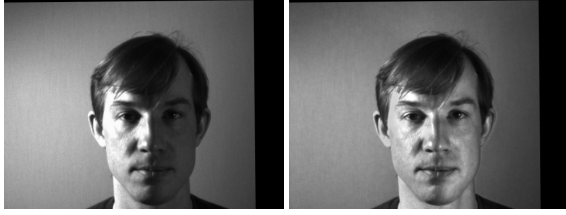


Figure 10: The original intensity channel and the fixed intensity channel using our method. Notice that the lighting is much more balanced, revealing details in the previously shadowed portion of the face.

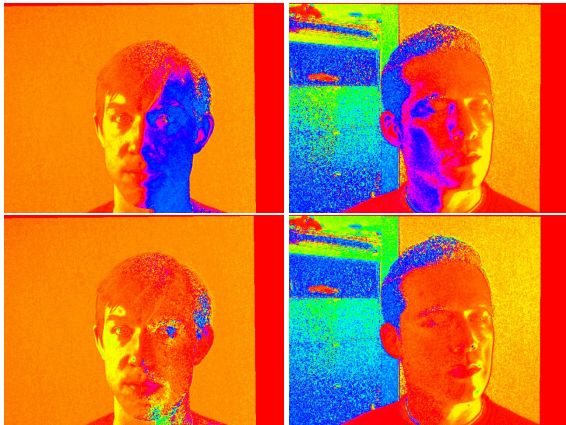


Figure 11: The images on top show the hue channels of two untouched color images with unbalanced lighting, the blue light from the monitor is especially visible in these images. The images on the bottom show the hue channels after applying our method. Much of the color cast has been removed. Remaining errors are primarily in areas of low saturation.

ing color channels separately, we are actually changing the relative intensities in different wavelengths, which fixes the inconsistency in the color channels better than would be possible with simple color balancing.

The effect that applying our method to the individual color channels has on the hue of the images can be seen in Figure 11. The improvement of the hue values is clearly visible in this space. It is important to note that for the purpose of visualization, saturation is not factored in these images, but it plays an important role in the perceptual closeness of colors. This is the reason for the remaining noise in the output images, as these points are actually close to the ideal color. However, at very low saturation values, slight differences in hue cause a noisy projection into hue space, but do not create visible artifacts.

4. Results

We implemented our algorithm in C++ using Intel's OpenCV library. Using a Pentium 4 3.00 Ghz. machine, we were able to process frames at 13fps at 320x240 resolution. After applying our algorithm, we observe that the corrected images contain more balanced lighting than the original image. In addition, the blue glow from the monitor was largely reduced in the output. As a result, details in the face are much easier to make out than in the original images as more information is filled in to shadowed and specular regions. These results are shown from two videos in Figure 12.

5. Limitations and Future Work

While our method is able to correct for some degree of lighting problems in videoconferencing there are still several limitations. For one, excessive movement between frames can cause artifacts, since the IR and color images are not in perfect temporal alignment. This could be partially addressed using an optical flow technique similar to [WGT*05].

In addition, areas such as deep shadows or specular highlights contain very little color information and thus it is not possible to accurately correct the light intensity in these regions without producing chroma artifacts. Our method also assumes specular highlights to be the same color as the target skin color, which is untrue. It may be possible to use a specular highlight mask when relighting the face to preserve their appearance, or conversely, to remove them with a detail transfer method.

While our prototype implementation makes use of two cameras and a beam splitter. We are interested in the possibility of widely available deployment so have paid careful attention to cost. In the future we believe that it would be possible to replace the standard RGB Bayer pattern in a camera with one that includes an IR sensitive color channel in addition to RGB. One could then imagine a system where the IR lights and a single modified camera are built into a monitor frame for a cost approximately the same as current implementations. We believe that it is also possible to adaptively change the IR bases so that fewer images need to be captured per frame. This could relax the need for a high speed camera as well as reduce errors from temporal misalignment.

6. Conclusion

We have introduced a system that is inexpensive, compact, and non-invasive. With our method, we are able to fix much of the poor lighting that noticeably degrades typical video conferencing. Our setup is uncalibrated and robust, the light bars can be positioned by hand and even moved during capture.

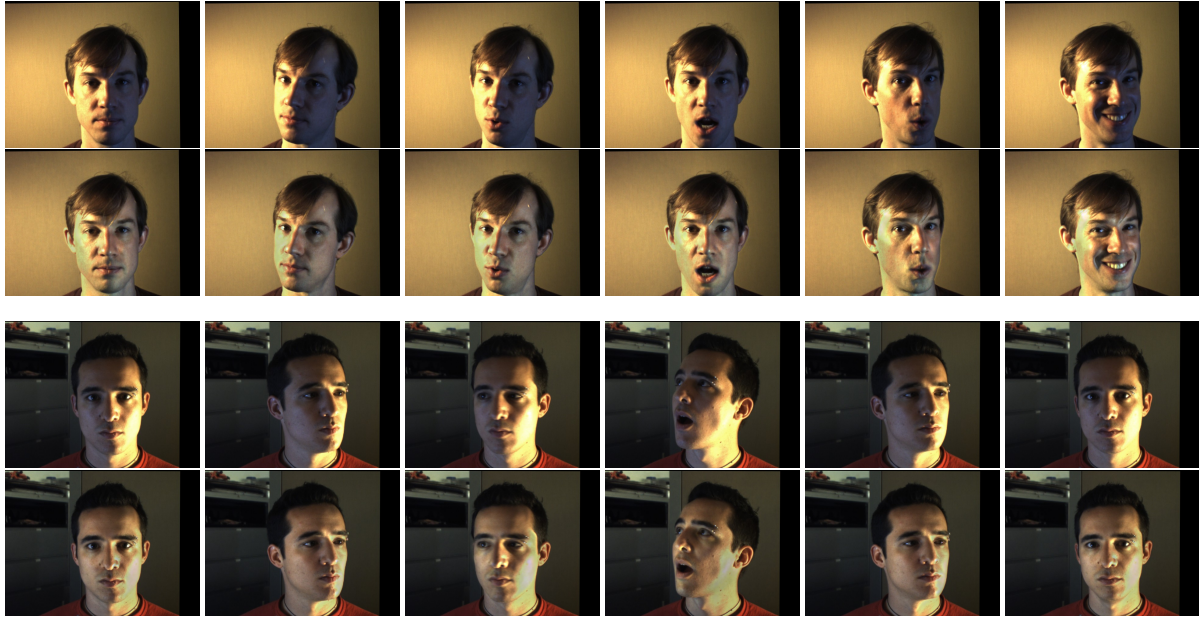


Figure 12: The original poorly lit frames of a video sequence can be corrected using our method. The results have more balanced lighting that better reveals facial details.

7. Acknowledgments

We would like to thank Jeff Mulligan and NASA for kindly donating the IR light rig and control board. We would also like to thank Hewlett-Packard for partial funding, and all the reviewers for their helpful comments and suggestions.

References

- [AD04] ANRYS F., DUTRE P.: Image based lighting design. In *4th IASTED International Conference on Visualization, Imaging, and Image Processing* (2004).
- [ADA*04] AGARWALA A., DONTCHEVA M., AGRAWALA M., DRUCKER S., COLBURN A., CURLESS B., SALESIN D., COHEN M.: Interactive digital photomontage. *ACM Trans. Graph.* 23, 3 (2004), 294–302.
- [ALK*03] AKERS D., LOSASSO F., KLINGNER J., AGRAWALA M., RICK J., HANRAHAN P.: Conveying shape and features with image-based relighting. In *VIS '03: Proceedings of the 14th IEEE Visualization 2003 (VIS'03)* (Washington, DC, USA, 2003), IEEE Computer Society, p. 46.
- [BGG*01] BASSO A., GRAF H., GIBBON D., COSATTO E., LIU S.: Virtual light: Digitally-generated lighting for video conferencing applications. *International Conference on Image Processing* (2001).
- [DHT*00] DEBEVEC P., HAWKINS T., TCHOU C., DUIKER H.-P., SAROKIN W., SAGAR M.: Acquiring the reflectance field of a human face. In *SIGGRAPH '00: Proceedings of the 27th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 2000), ACM Press/Addison-Wesley Publishing Co., pp. 145–156.
- [Eur] Safety of laser products. equipment classification, requirements and user's guide, EN 60825-1, 1994-1996.
- [FBS05] FUCHS M., BLANZ V., SEIDEL H.: Bayesian relighting. In *Rendering Techniques 2005: Eurographics Symposium on Rendering* (Konstanz, Germany, July 2005), Deussen O., Keller A., Bala K., Dutré P., Fellner D. W., Spencer S. N., (Eds.), *Rendering Techniques*, Co-sponsored by ACM SIGGRAPH and EUROGRAPHICS Association, Eurographics, pp. 157–164.
- [FLS05] FUCHS M., LENSCH H., SEIDEL H.-P.: Reflectance from images: A model-based approach for human faces. *IEEE Transactions on Visualization and Computer Graphics* 11, 3 (2005), 296–305. Member-Volker Blanz.
- [GBK99] GEORGHIADES A., BELHUMEUR P., KRIEGSMAN D.: Illumination-based image synthesis: Creating novel images of human faces under differing pose and lighting, 1999.
- [Gre04] GREY C.: *Master Lighting Guide for Portrait Photographers*. Amherst Media, 2004.
- [HP] HEWLETT-PACKARD: <http://www.hp.com/halo/index.html>.

- [LKG*03] LENSCH H., KAUTZ J., GOESELE M., HEIDRICH W., SEIDEL H.-P.: Image-based reconstruction of spatial appearance and geometric detail. *ACM Trans. Graph.* 22, 2 (2003), 234–257.
- [LPMM05] LEE J., PFISTER H., MOGHADDAM B., MACHIRAJU R.: Estimation of 3d faces and illumination from single photographs using a bilinear illumination model. In *Rendering Techniques* (2005), pp. 73–82.
- [LSZ01] LIU Z., SHAN Y., ZHANG Z.: Expressive expression mapping with ratio images. In *SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 2001), ACM Press, pp. 271–276.
- [MG97] MARSCHNER S. R., GREENBERG D. P.: Inverse lighting for photography. In *Proceedings of the Fifth Color Imaging Conference, Society for Imaging Science and Technology* (1997).
- [MGW01] MALZBENDER T., GELB D., WOLTERS H.: Polynomial texture maps. In *SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 2001), ACM Press, pp. 519–528.
- [MLTG99] MARSCHNER S., LAFORTUNE E., TORRANCE K., GREENBERG D.: Image-based brdf measurement including human skin. *10th Eurographics Workshop on Rendering* (1999), 139–152.
- [MNNB05] MORENO-NOGUER F., NAYAR S. K., BELHUMEUR P. N.: Optimal illumination for image and video relighting. In *SIGGRAPH '05: ACM SIGGRAPH 2005 Sketches* (New York, NY, USA, 2005), ACM Press, p. 75.
- [MTB*06] MOHAN A., TUMBLIN J., BODENHEIMER B., GRIMM C., BAILEY R.: Table-top computed lighting for practical digital photography. In *SIGGRAPH '06: ACM SIGGRAPH 2006 Courses* (New York, NY, USA, 2006), ACM Press, p. 3.
- [NN04] NISHINO K., NAYAR S. K.: Eyes for relighting. In *SIGGRAPH '04: ACM SIGGRAPH 2004 Papers* (New York, NY, USA, 2004), ACM Press, pp. 704–711.
- [PSA*04] PETSCHNIG G., SZELISKI R., AGRAWALA M., COHEN M., HOPPE H., TOYAMA K.: Digital photography with flash and no-flash image pairs. In *SIGGRAPH '04: ACM SIGGRAPH 2004 Papers* (New York, NY, USA, 2004), ACM Press, pp. 664–672.
- [PSS99] PIGHIN F. H., SZELISKI R., SALESIN D.: Resynthesizing facial animation through 3d model-based tracking. In *ICCV (1)* (1999), pp. 143–150.
- [PTMD07] PEERS P., TAMURA N., MATUSIK W., DEBEVEC P.: Post-production facial performance relighting using reflectance transfer. *ACM Trans. Graph.* 26, 3 (2007), 52.
- [RRS01] RIKLIN-RAVIV T., SHASHUA A.: The quotient image: Class-based re-rendering and recognition with varying illuminations. *IEEE Trans. Pattern Anal. Mach. Intell.* 23, 2 (2001), 129–139.
- [SLKS06] SUN J., LI Y., KANG S., SHUM H.: Flash matting. In *SIGGRAPH '06: ACM SIGGRAPH 2006 Papers* (New York, NY, USA, 2006), ACM Press, pp. 772–778.
- [VP06] VUILLEUMIER P., POURTOIS G.: Distributed and interactive brain mechanisms during emotion face perception: Evidence from functional neuroimaging. *Neuropsychologia* 45 (2006), 174–194.
- [WGT*05] WENGER A., GARDNER A., TCHOU C., UNGER J., HAWKINS T., DEBEVEC P.: Performance relighting and reflectance transformation with time-multiplexed illumination. In *SIGGRAPH '05: ACM SIGGRAPH 2005 Papers* (New York, NY, USA, 2005), ACM Press, pp. 756–764.
- [WLH03] WEN Z., LIU Z., HUANG T.: Face relighting with radiance environment maps. *CVPR 02* (2003), 158.
- [WMP*06] WEYRICH T., MATUSIK W., PFISTER H., BICKEL B., DONNER C., TU C., MCANDLESS J., LEE J., NGAN A., JENSEN H. W., GROSS M.: Analysis of human faces using a measurement-based skin reflectance model. In *ACM Transactions on Graphics* (July 2006), vol. 25(3).
- [YDMH99] YU Y., DEBEVEC P., MALIK J., HAWKINS T.: Inverse global illumination: recovering reflectance models of real scenes from photographs. In *SIGGRAPH '99: Proceedings of the 26th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 1999), ACM Press/Addison-Wesley Publishing Co., pp. 215–224.