

A Model-Based Approach to Image Relighting with a Potential for Real-Time Implementation

Claus B. Madsen and Rune Laursen

Lab. of Computer Vision and Media Technology
Aalborg University
Aalborg, Denmark

Abstract

Image relighting is a very unique special visual effect which promises to have many important practical applications. Image relighting is essentially the process of, given one or more images of some scene, computing what that scene would look like under some other (arbitrary) lighting conditions, e.g., changing positions and colors of light sources. Image relighting can for example be used for interior light design. This paper describes an approach to image relighting which can be implemented to run in real-time by utilizing graphics hardware, as opposed to other state-of-the-art approaches which at best run at a few frames per second.

Categories and Subject Descriptors (according to ACM CCS): I.3.4 [Computer Graphics]: Three-Dimensional Graphics and Realism

1. Introduction

This paper addresses the subject of developing relighting techniques, i.e. techniques allowing a user to completely alter the lighting conditions in an image of a real scene. Specifically, the paper focuses on techniques providing *real-time* relighting functionalities, thus enabling the user to interactively change lighting conditions and get "instant" visual feedback. Figure 1 provides an example of the kind of relighting this paper addresses. It should be noted that the work presented here presumes the availability of three things: 1) an original image of the scene, 2) a 3D model of the scene, and 3) a model of the lighting conditions in the scene at the time the original image is acquired. We will return to ways in which the two last pieces of knowledge can be obtained.

Conceptually image relighting in this manner is a two step process. In the first step all effects of the original lighting conditions are *removed*, e.g., highlights, shadows, and differences in shading across surfaces due to varying light incidence angles. In the second step the scene is subjected to some arbitrary new lighting conditions and the appearance of the scene in these conditions is computed. The second step thus "adds" new highlights, shadows and shading etc. Of these two steps the former is the tricky one, while the latter can be performed using any preferred rendering tech-

nique, e.g., ray tracing, radiosity or standard hardware accelerated approaches. Which rendering technique is employed depends on the preferred balance between rendering speed and accuracy in handling various lighting phenomena. In order to achieve true real-time performance we have chosen to use a hardware accelerated approach for step 2, thus sacrificing certain global illumination phenomena.

In our approach step 1 is achieved by a *computational* approach which requires, as stated above, a 3D model of the scene and a model of the original lighting conditions. Alternatively, one could in principle acquire fronto-parallel digital images of the surfaces in the scene under perfectly diffuse, white-balanced lighting conditions and use these images as textures on the 3D model, which is subsequently rendered under novel lighting conditions (step 2). This would be a *mechanical* or *image acquisition* approach to step 1, but in reality acquiring such 'clean' textures devoid of lighting effects is not practical for general scenes.

The contributions in this work lie in the specific manner in which the operations performed in steps 1 and 2 are carried out. With the approach described here the two steps can be combined such that the image relighting becomes a matter of modulating the original image on a pixel by pixel basis with a "relighting map". The relighting map can be computed in



Figure 1: Left: original image acquired outdoors on a sunny day. Right: same scene as left but this image is a simulation of what the scene would look like if the position of the sun were different. The work in this paper enables such changes in lighting conditions to be performed in real-time.

real-time using standard techniques, and the modulation can also easily be performed in real-time. Thus, our approach has two advantages: 1) it is directly designed for real-time performance, and 2) the original image is used directly and therefore the final image is not subject to filtering and/or aliasing effects involved in doing reprojections of textures mapped to a 3D model of the scene.

This paper represents the current state of work in progress and will only address the problem for scenes with perfectly diffuse reflectance properties. The paper is organized as follows. In section 2, we present an overview of our approach and show how the relighting effects are achieved. Section 3 then describes related work. Section 4 describes our approach in more detail, followed by section 5 giving some practical details behind the initial experiments we have performed to validate the proposed approach. Section 6 discusses central aspects of the work and points to future research. Finally section 7 offers conclusions.

2. Overview of approach

Prior to describing the proposed approach in a more technically rigorous manner this section attempts to provide the reader with an intuitive understanding of the issues involved and of the process behind our technique.

The approach requires different types of input information. First of all an image of the original scene is required. Secondly, a 3D model of the scene must be available, and the original image must be calibrated to the 3D model such that every pixel corresponds to a known 3D point in the scene model. Third, the original lighting conditions in the scene must be known, i.e., we need to know the sources of light in the scene, and their relative intensities.

The 3D model can be obtained in many different ways [Oli02], e.g., by reconstruction from multiple images using approaches such as [PKV99], by Image-Based Modelling, e.g. [DTM96], or by laser range scanning. Alternative the scene can be measured and a model constructed manually. The latter is the approach employed for our experimental results, i.e., we have measured the scene, constructed crude polygonal models of the objects, and then calibrated the camera to the 3D model using manually established 2D to 3D point correspondences. Figure 2 shows the scene model used for the relighting illustrated in figure 1.

The required knowledge of the original lighting conditions can most easily be acquired using the popular light probe approach, i.e., by taking high dynamic range images of a reflective sphere placed in the scene, [Deb98, GCHH03]. Alternatively, light source positions, sizes and power can be measured manually as done in [LDR00] or semi-automatically using multiple images as in [YDMH99]. For the experimental results in this paper we have done it manually in a manner described in section 5.

In this paper we will limit ourselves to discussing the case of scenes containing surfaces with perfectly diffuse reflectance properties.

Each pixel in the original image is a measurement of the radiance (in the three RGB bands) from a unique 3D point in the scene in the direction of the viewpoint. Thus the original image is a 2D **radiance map**, $L_o(u, v)$, where u and v are the image coordinates. Because we have the 3D model and knowledge of the original lighting conditions it is trivial[†] to

[†] Irradiance computation is trivial provided global illumination issues (irradiance contributions from diffuse reflections) are disre-




Figure 2: 3D model corresponding to scene shown in figure 1. The image illustrates the model as a depth map, i.e., intensity is a measure of distance from the camera. We have used one plane for the tiled ground plane, a plane for the brick wall on the right, six quadrilaterals for the speaker on the left, and two quads and one triangle for for calibration object in the center.

compute the amount of light arriving at the same 3D points in the scene, i.e., it is possible to construct an **irradiance map**, $I(u, v)$. When the irradiance map is computed using the known original scene lighting conditions we will call it $I_o(u, v)$. Conversely, when the irradiance map is computed using some arbitrary different *relighting* conditions it will be denoted $I_r(u, v)$.

For purely diffuse Bidirectional Reflectance Distribution Functions (BRDFs) there is a linear relationship between radiance and irradiance (radiance is proportional to diffuse albedo times irradiance). Therefore diffuse scenes can be very simply relit by dividing the radiance map with the original irradiance map, and then multiplying with the relighting irradiance map. Using the introduced terminology $L_r(u, v) = L_o(u, v) \cdot I_r(u, v) / I_o(u, v)$.

Figure 3 shows original and relighting irradiance maps corresponding to the relighting example in figure 1.

Section 1 presented the general concept of relighting as a two step process: 1) removing original light from the image, and 2) adding new light. In the described diffuse case step 1 is represented by the $L_o(u, v) / I_o(u, v)$ operation, whereas step 2 is performed by multiplying the result from step 1 by the relight irradiance, $I_r(u, v)$. Step 1 is a once-only process as it only involves elements that do not change over time (original image and original irradiance map). Step 2 has to be

garded. If this is not a fair assumption the work by Yu et al., [YDMH99], can be used to compute these contributions.

re-performed constantly in response to the user's changes in the desired lighting conditions for the relit scene, so $I_r(u, v)$ must in general be re-computed for every frame.

Step 1 could be pre-computed and the result stored in an image which is then subsequently modulated by the real-time computed relighting irradiance map. Our approach does not do this; we have designed a solution which embeds the normalization with the original irradiance into the computation of the relighting irradiance map, $I_r(u, v)$. Thus, at run-time we actually compute the relight/original light ratio $I_r(u, v) / I_o(u, v)$ directly. This "ratio map" is then used for modulating the original image. By doing this we avoid the non-trivial implementation of a real-time texture division operation.

Computing the ratio map is described in detail in section 4. The idea is based on the observation that if one renders a radiance image of an all-white perfectly diffuse 3D model of the scene under the chosen relighting conditions, this image is then identical to the required irradiance map, $I_r(u, v)$. If instead we set the reflectances of the 3D model to be proportional to the inverse of the original irradiances then the rendered image automatically becomes the desired relight/original irradiance ratio map.

3. Related work

In the area of image relighting there are basically two different paradigms or approaches: image-based relighting and model-based relighting. The work in this paper falls in the model-based category, and model-based in this context refers to an assumption that a 3D model of the scene is available. Image-based relighting on the other hand typically involves multiple images acquired under different, often highly controlled, illumination conditions (light direction, special light patterns, etc.). Recent work in this category include [ALK*03, SC04, SC05].

There is a small amount of related work in the literature over the model-based approaches, which as stated above is the category in which this work falls. First of all Yu et al., [YM98], demonstrated how they could acquire reflectance properties of architectural scenes by taking at least two images of each surface of the objects under different lighting conditions. The recovered 3D model combined with the estimated reflectance parameters could then be used to render the scene under changing lighting conditions. The focus of this work is entirely on parameter recovery and relighting is by no means done in real-time.

Similarly, inverse global illumination was proposed by Yu et al., [YDMH99] for recovery of reflectance parameters for indoor scenes using multiple images of each surface from different viewpoints. Again this work focuses on reflectance parameter estimation and relighting is done using RADIANCE, [War94], which again is far from real-time.

The most closely related work is that of Loscos et al.,




Figure 3: Left: irradiance map, $I_o(u, v)$, corresponding to original lighting conditions of the scene shown in figure 1. Center: irradiance map, $I_r(u, v)$, corresponding to lighting conditions where the dominant light source, the sun, has changed location relative to the scene. These are the lighting conditions valid for the relit image in figure 1. Right: illustrates the "relighting factor", i.e., the ratio of original irradiance to relighting irradiance, I_r/I_o . Every pixel in the original radiance map is modulated with its corresponding pixel in this map. Dark means "light is removed" from original pixels, and bright means "light is added" to original pixels.

[LDR00]. This work also enables a user to change the lighting conditions in an image of a scene in an interactive manner, but this work is centered around a radiosity method for irradiance computations. Therefore, the method performs at a few frames per second when the only lighting changes performed are intensity adjustments. If the number of sources or their positions are changed updating takes on the order of 10 seconds.

The work in [LDR00] also employs texture modulation for efficient relighting, and the modulation texture (irradiance map) is computed using radiosity. We have chosen to focus specifically on true real-time performance and therefore the computation of the relighting irradiance maps does not account for global illumination phenomena such as color bleeding. Nevertheless, with the work currently being done in Pre-computed Radiance Transfer and Photon Mapping, real-time global illumination is coming closer and closer to reality, and our approach can readily be combined with such efficient global illumination techniques.

4. The perfectly diffuse case

For perfectly diffuse reflectors the relationship between incident irradiance, I , and radiance, L , in any direction is given by the diffuse albedo, ρ_d :

$$L = \frac{\rho_d}{\pi} I \quad (1)$$

The original image (radiance map), $L_o(u, v)$, provides us with measured radiances from a dense set of 3D points in the scene, and these points are known since we assume the camera is calibrated to the scene. In general the diffuse albedo and the irradiance vary for every point in the scene, so the relationship between radiance maps and irradiance maps becomes:

$$L_o(u, v) = \frac{\rho_d(u, v)}{\pi} I_o(u, v) \quad (2)$$

Here, $\rho_d(u, v)$ is the "albedo map". When doing relighting the albedo stays constant; the only thing that changes is the irradiance at each scene point. Therefore, the radiance map of the relit image/scene can be expressed as:

$$\begin{aligned} L_r(u, v) &= \frac{\rho_d(u, v)}{\pi} I_r(u, v) \\ &= (L_o(u, v)/I_o(u, v)) \cdot I_r(u, v) \\ &= L_o(u, v) \cdot (I_r(u, v)/I_o(u, v)) \end{aligned} \quad (3)$$

Eq. 3 simply shows that the relit image can be computed by modulating the original image with a ratio of two irradiance maps: the relight irradiance map, $I_r(u, v)$, corresponding to the user's desired scene (new) lighting conditions, and the original irradiance map, $I_o(u, v)$. The key element in our approach is a technique for computing this map in real-time and using it for modulation of the original image.

4.1. Computing the irradiance ratio map

How can we efficiently compute the irradiance ratio map? First, let us describe how simply the relighting irradiance map can be computed using standard local illumination techniques (specifically we will use the Phong lighting model of OpenGL, a description of which may be found in books such as [Bus03, WP01]). Rendering an image of a scene using the Phong lighting model results in a radiance from a 3D point which can be formulated as (disregarding specular reflection):

$$L = \rho_a I_a + \rho_d \sum_{i=1}^k I_i \cos(\theta_i) \quad (4)$$

Here L is the radiance from a 3D point in the direction of the viewpoint. ρ_a and ρ_d are the ambient and diffuse reflectances, respectively, (eq. 4 is to be evaluated for each of three RGB colors). I_a is the ambient irradiance at the 3D

point. $I_i \cos(\theta_i)$ is the irradiance at the point caused by the i th point light source, $(\theta_i(u, v))$ is the angle between the surface normal and the direction vector to the i th light source at the 3D point). k is the number of light sources.

If we set the ambient and diffuse reflectances equal, eq. 4 changes to:

$$L = \rho_d \left(I_a + \sum_{i=1}^k I_i \cos(\theta_i) \right) \quad (5)$$

Eq. 5 states that rendering with OpenGL Phong lighting the radiance from a point equals the reflectance at the point times the total irradiance (ambient plus sum of individual point source contributions) at the point. Thus, by setting unit reflectances, the *radiance equals the irradiance*. This may be self-evident but is important because it shows that we can use the graphics card's efficient lighting computation capabilities to produce irradiance maps needed for relighting.

That is, if we render the 3D scene model from a viewpoint corresponding to the original image, and if all surfaces in the rendered 3D model have unit reflectances, then the resulting image is an irradiance map. This means that relighting irradiance maps, $I_r(u, v)$, for any user desired lighting conditions can be rendered simply by rendering a diffuse, all-white 3D scene model under the chosen lighting conditions.

To actually do relighting we not only needed real-time computation of $I_r(u, v)$, but we needed the lighting ratio map, $I_r(u, v)/I_o(u, v)$. This is accomplished by setting the reflectances of points in the 3D model to the inverse of the original irradiance at that point, $\rho_d = 1/I_o$.

To summarize the light ratio maps are generated by doing the following rendering using hardware acceleration:

1. upload 3D scene model to graphics card
2. set ambient and diffuse RGB reflectances of all vertices to the inverse of the original irradiance at that 3D point
3. setup the desired lighting conditions consisting of ambient and point source contributions
4. render the model to a texture using a viewport corresponding to the camera in the original image

4.2. Practical issues

In the previous section we described how to use hardware accelerated local lighting rendering to produce irradiance ratio maps for modulating the original image. With this approach there is really no limits to how much the lighting conditions in the scene can be altered.

We are presently implementing the proposed technique but all images in this paper were produced by a non-real-time simulation of the presented approach. Figure 4 shows what the original scene looks like with a (non-existing) light source in the very center of the scene.

For the ongoing implementation of the real-time version

the only real issue to contemplate is the resolution of the 3D model of the scene. In order to properly capture gradients in the original irradiances of the scene the resolution of the 3D model has to be high at such gradients. We are working on designing methods for adaptive subdivision based on evaluating irradiance differences between 3D model vertex locations. If differences are too high the surface is subdivided.

Computing original irradiances to be used inversely as reflectances of the (subdivided) 3D scene model is an off-line process which can be done using any preferred rendering technique, for example Monte Carlo ray tracing to enable proper handling of area light sources. This is especially possible if a high dynamic range light probe image of the scene is available, because then an Image-Based Lighting approach, [Deb98, Deb02], can be used to compute accurate irradiances which properly handle soft shadows in the original image.

In the on-line stage, when rendering the 3D model with the assigned reflectances, cast shadows are important for proper irradiance computation. For this we propose to use a shadow volume approach to detecting shadowed areas.

5. Experiments

As mentioned previously the images shown in this paper are produced using a non-real-time version of the presented technique. The original image was acquired with a standard 5 mega pixel digital camera. The scene was measured manually and a simple 3D model of it was constructed (as described in section 2).

The camera was calibrated to the 3D model using manually established 2D to 3D point correspondences, and the estimation of internal and external camera parameters was done using an approach from [TV98].

The original lighting conditions were modelled as a combination of a point light source (the sun) and an ambient term (the sky). The position of the sun relative to the scene model was determined by orienting the calibration object such that sun rays were parallel to the xz -plane, thus fixing the sun's y -coordinate to zero. The x and z coordinates were then found by measuring the length of a shadow cast by an object of known height. The RGB intensities of the blue ambient sky light was determined from the image colors of the white paper of the calibration object in areas not exposed to sunlight. By comparing RGB values of calibration object cardboard in shadow and in direct light the relative intensities between ambient and sun light were determined, (taking the cosine fall-off for diffuse reflection into account for the sun point source).

It should be made very clear that the original lighting conditions modelled as described above are extremely crude and this was only done to get quick working results. A much more precise method is to use light probe images, and the




Figure 4: Left: irradiance map corresponding to a user defined lighting environment with a weak ambient term and a point light source in the middle of the scene. Right: resulting relit image.

last two weeks we have been waiting for nice weather to capture better test images, but the deadline ran out before we could redo the experiments.

For computing the original irradiance map a simple ray tracing approach was implemented which considers local illumination only, by casting primary rays plus shadow feelers. The original radiance map was computed in image resolution. The relighting examples given in the paper basically involve changing the location of the sun source. Given some desired sun position the simple raytracer was used to render a relighting irradiance map, again in image resolution. The relighting irradiance map was divided by the original irradiance map and the result multiplied with the original image to complete the diffuse relighting process.

6. Discussion

In this section we will briefly discuss some important points in relation to our proposed method.

Using our approach it is possible to employ arbitrarily complex and accurate computations of the original irradiance map. This is an off-line, once-only computation the results of which are used to set the reflectances of the 3D scene model subsequently used for relighting. We believe handling area light sources to be very important, even for outdoor images, since shadows due to sun light actually do have noticeable penumbra regions. Similarly, we believe taking global illumination phenomena (indirect light) into account is important, especially for indoor scenes, where reflections from other surfaces may be a significant irradiance contribution for a given surface.

Conversely, for the actual on-line, real-time rendering of irradiances during interactive relighting we have here proposed a straight forward local illumination approach. Yet,

the basic approach with using the 3D scene model, normalized with original irradiances, can be used in conjunction with any lighting algorithm depending on how accurate one desires the result should be.

Throughout this paper we have assumed scenes to consist entirely of Lambertian materials. Our approach actually does generalize nicely to scenes with glossy BRDFs. It requires an additional rendering pass in the real-time relighting process, in order first to modulate the original image with the diffuse part of the relighting/original irradiance ratio map and subsequently add the specular radiance part. Figure 5 demonstrates the effect of adding a specular component to the surfaces during relighting.

7. Conclusions

We have described an approach to image/scene relight which based on a 3D model of the scene and knowledge of the original lighting conditions can compute the appearance of the scene under any arbitrary new lighting conditions, including changing the number of light sources, their positions and radiant powers.

The main contribution of the work is the fact that the approach is directly designed for real-time performance, enabling a user to get instant visual feedback upon having changed the parameters of the lighting environment. A smaller contribution lies in the idea of performing the normalization with the original irradiance by appropriately setting the reflectances of the 3D model used for real-time irradiance computations. This allows the approach to operate directly on the original image, rather than computing the albedo map off-line and modulating it at run-time.

An important aspect of the proposed approach is that relighting is performed as a modulation of the original image.



Figure 5: A specular reflection component has been added to each surface during relighting to illustrate the possibility of playing with the reflectance properties.

It is believed that doing relighting in this manner is superior to an approach where textures extracted from the image are mapped to the scene geometry and reprojected at run-time, because the multiple re-sampling steps involved will cause the resulting image to be blurred.

Acknowledgments

This research is funded in part by the BENOGO project under the European Commission IST program (IST-2001-39184), and in part by the ARTHUR project (IST-2000-28559). This support is gratefully acknowledged.

References

[ALK*03] AKERS D., LOSASSO F., KLINGNER J., AGRAWALA M., RICK J., HANRAHAN P.: Conveying shape and features with image-based relighting. In *Proceedings: IEEE Visualization 2003* (2003), pp. 349 – 354.

[Bus03] BUSS S. R.: *3-D Computer Graphics – A Mathematical Introduction with OpenGL*. Cambridge University Press, 2003.

[Deb98] DEBEVEC P.: Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *Proceedings: SIGGRAPH 1998, Orlando, Florida, USA* (July 1998).

[Deb02] DEBEVEC P.: Tutorial: Image-based lighting. *IEEE Computer Graphics and Applications* (March/April 2002), 26 – 34.

[DTM96] DEBEVEC P., TAYLOR C., MALIK J.: Modelling and rendering architecture from photographs: a hybrid geometry and image-based approach. In *Proceed-*

ings: SIGGRAPH 1996, New Orleans, LA, USA (August 1996), pp. 11 – 20.

[GCHH03] GIBSON S., COOK J., HOWARD T., HUBBOLD R.: Rapid shadow generation in real-world lighting environments. In *Proceedings: EuroGraphics Symposium on Rendering, Leuven, Belgium* (June 2003).

[LDR00] LOSCOS C., DRETTAKIS G., ROBERT L.: Interactive virtual relighting of real scenes. *IEEE Transactions on Visualization and Computer Graphics* 6, 4 (October-December 2000), 289 – 305.

[Oli02] OLIVEIRA M. M.: Image-based modelling and rendering: A survey. *RITA - Revista de Informatica Teorica e Aplicada* 9, 2 (October 2002), 37 – 66. Brazilian journal, but paper is in English.

[PKV99] POLLEFEYS M., KOCH R., VANGOOL L.: Self-calibration and metric reconstruction in spite of varying and unknown internal camera parameters. *International Journal of Computer Vision* 32, 1 (1999), 7 – 25.

[SC04] SHIM K. H., CHEN T.: Efficient representation of lighting patterns for image-based relighting. In *Proceedings: Picture Coding Symposium* (2004).

[SC05] SHIM K. H., CHEN T.: A statistical framework for image-based relighting. In *Proceedings: IEEE International Conference on Acoustics, Speech, and Signal Processing, Philadelphia, Pennsylvania* (2005).

[TV98] TRUCCO E., VERRI A.: *Introductory Techniques for 3D Computer Vision*. Prentice Hall, 1998.

[War94] WARD G. J.: The radiance lighting simulation and rendering system. In *Proceedings: SIGGRAPH 1994* (July 1994), pp. 459 – 472.

[WP01] WATT A., POLICARPO F.: *3D Games: Real-Time Rendering and Software Technology*, vol. 1. Addison-Wesley, 2001.

[YDMH99] YU Y., DEBEVEC P., MALIK J., HAWKINS T.: Inverse global illumination: Recovering reflectance models of real scenes from photographs. In *Proceedings: SIGGRAPH 1999, Los Angeles, California, USA* (August 1999), pp. 215 – 224.

[YM98] YU Y., MALIK J.: Recovering photometric properties of architectural scenes from photographs. In *Proceedings: SIGGRAPH 1998, Orlando, Florida, USA* (July 1998), pp. 207 – 217.