Real-time Illumination Estimation from Image Sequences

Claus B. Madsen, Moritz Störring,
Tommy Jensen, Mikkel S. Andersen, and Morten F. Christensen
Laboratory of Computer Vision and Media Technology
Aalborg University, Aalborg, Denmark
cbm/mst@cvmt.aau.dk
www.cospe.dk

Abstract

Knowledge about the illumination conditions in a real world scene has many applications among them Augmented Reality which aims at placing virtual objects in the real world. An important factor for convincing augmentations is to use the illumination of the real world when rendering the virtual objects so they are shaded consistently and cast consistent shadows.

This paper proposes two approaches to continuously estimate the illumination conditions in a static outdoor scene based on images from a single viewpoint of that scene while using the scene itself as light probe. Thus, no additional calibration objects are required. Experimental results show that the proposed illumination estimation is sufficient for Augmented Reality applications.

1 Introduction

Images are formed as a result of light interacting with surfaces. The radiation emitted by a light source hits a material's surface under a certain angle where it is then reflected, absorbed, and transmitted depending on the material's properties. The reflected light may hit other objects causing interreflections, and one object may occlude another object's reflections or a light source resulting in shadowing. When images are synthesized using computer graphics techniques it is important to have good models of these interactions in order to achieve realism. Similarly, when the images are real images acquired with some form of camera it is paramount to understand how the image was formed in order to analyze it using computer vision techniques. Generally, three different elements come together in forming images: 1) the 3D geometry of the scene, 2) the reflectance properties of the surfaces in the scene, and 3) the illumination conditions in the scene. Given a model of all three elements it is possible both to render synthetic images and to design robust computer vision techniques for analyzing images of the scene.

The Laboratory of Computer Vision and Media Technology at Aalborg University, Denmark (CVMT/AAU) has recently initiated a research project (CoSPE: Computer Vision-Based Scene Parameter Estimation) which lies on the border between computer vision and computer graphics. The project focuses on estimating the reflectance properties and the illumination conditions in scenes based on images. For more information about the project please visit the project's web-site www.cospe.dk

In this paper we present some initial results of this research, namely two approaches to the same problem: to continuously estimate the illumination conditions in a static scene based on a sequence of images from a single viewpoint. So far, the most commonly used approach to scene illumination measurement/estimation has been the so-called light probe, which is a reflective sphere placed in the scene and photographed with a camera to get an omni-directional measurement of light, [5, 8, 10]. None of the approaches presented in this paper require any special purpose radiometric calibration objects to be present in the scene. In fact one could say we are proposing techniques that allow the scene to act as its own light probe.

Real-time, continuous estimation of scene illumination conditions is really important for Augmented Reality (AR) systems. Figure 1 shows an example of an AR system where a virtual object has been rendered into a real scene. The virtual object is rendered with illumination conditions corresponding to the illumination condition that are estimated for the real scene, so the virtual object is shaded consistently with the scene, and it also casts a consistent shadow on surfaces in the real scene.

The application scenario we are targeting is a system, be it an AR or a vision system, which needs to continuously update its internal model of the illumination in an outdoor scenario. Consider for example a computer screen mounted on a pole at an archaeological site allowing visitors to view the real scene (filmed with a video camera) augmented with visualization of virtual 3D buildings that no longer exist. For such a system the illumination conditions constantly change due to the passing of time causing the sun to travel across the sky, clouds causing partial or complete blockage of the direct light from the sun, and changing the illumination from the sky.

The approaches presented in this paper can estimate the intensities and color of the direct sunlight and of the

Figure 1: Two images taken at different times (approximately one hour apart) on a sunny day with partial cloud cover causing constant changes in the illumination conditions. With one of the methods proposed in this paper we have automatically estimated the current illumination conditions and used this illumination estimate to render a virtual sculpture into the scene.

indirect skylight. Additionally one method is able to estimate the direction of the sunlight relative to the scene, whereas the other technique assumes that the system knows the direction of the sunlight using date, time and position information. The latter approach is more robust for cloudy conditions, whereas the former is more readily applicable to a scene as there is less positional and orientational calibration to carry out.

Both approaches involve an "off-line" photometric calibration phase where the reflectances (albedo) of diffuse surfaces in the scene are estimated. After the once-only reflectance calibration the approaches enable continuous "on-line" illumination estimation.

This paper is organized as follows. In section 2 we give an overview of related work. Section 3 then lists the assumptions behind the presented techniques, and presents the illumination model used by both approaches (both approaches estimate the values of parameters in this model). In section 4 we then present the approach which assumes availability of sunlight direction information, whereas the approach which also estimates sunlight direction is presented in section 5. Conclusions are given in section 6.

2 State-of-the-art

Estimating scene illumination conditions from images is the dual problem of estimating surface reflectance properties, because the image represents light reflected off surfaces, and this reflection is governed by the illumination and the reflectances. Therefore illumination estimation cannot be performed without knowledge of surface reflectance. This is the reason all related work is based on placing some kind of special purpose object with a priori known reflectance properties in the scene. For continuously operating AR or vision systems performing illumination estimation it is not a viable approach to be forced to have calibration objects in the scene. Therefore we have developed and tested two approaches to estimate dynamic illumination conditions based on the surfaces naturally present in the scene. Subsequently we briefly describe some of the most closely related work. Recent surveys on illumination estimation may be found in [9, 13].

One group of related work has a somewhat different focus, namely that of estimating scene reflectances. Yu and Malik proposed estimation of pseudo BRDFs[1] for outdoor scenes, [17]. The approach requires multiple images of the outdoor scene taken from different viewpoints and under differing illumination conditions. The goal is to be able to re-render the scene under arbitrary novel illumination conditions. Knowledge of scene illumination is obtained by combining a parameterized outdoor skylight model with light probe images.

Yu and Debevec proposed an inverse global illumination rendering approach, [16]. By using multiple images of all surfaces and a complete 3D model of an entire indoor scenario, and by using knowledge of the illumination conditions they demonstrate that it is possible to estimate glossy BRDFs for all surfaces. Knowledge of illumination conditions is obtained by manually measuring the positions and emittances of all light sources.

Loscos and Drettakis proposed a system for interactive re-lighting of indoor scenarios, [11]. Using a single image of the scene, combined with a complete 3D model of the entire room, and knowledge of the original illumination conditions they are able to re-render the scene under arbitrary novel illumination conditions. The knowledge of the original scene illumination is obtained by manual measurement of the positions and emittances of light sources.

Boivin and Gagalowicz proposes an iterative global

[1]BRDF: Bi-directional Reflectance Distribution Function is defined as the ratio of the reflected radiation to the incident radiation on a surface.

illumination approach to estimating surface reflectance parameters, [3, 4]. This work is also based on a single image of an indoor scene, and assumes that the positions and emittances of the light sources are measured manually.

Masselus and Dutre proposed an approach to image-based modeling of surface reflectances with the aim of being able to re-render under novel illumination conditions, [12]. By acquiring multiple images from a single viewpoint of a scene illuminated with a manually moved single light source they were able to model the reflectance field for re-lighting. The location of the moving light source is computed for each image by a triangulation technique based on the shading of four diffuse spheres present at known locations in the scene.

Sato and Sato proposed a technique for estimation of complex illumination environments, [15]. The technique requires that a known object is casting shadows on a surface in the scene, and the reflectance of the shadow receiver must be known. If this information is not available the method requires an image of the scene without the shadow casting object. In this case the method cannot be applied to scenes with changing illumination conditions.

Kanbara and Yokoya designed an approach to automatic, real-time estimation of scene lighting for augmented reality, [10]. The approach involves placing a reflective sphere which is always in the camera's field of view. The dynamic scene illumination conditions are estimated from the environment's reflection in this special purpose sphere.

Using reflective spheres has for several years been the standard approach to acquiring omni-directional knowledge of scene illumination. The approach has been pioneered by Debevec and taken up by several other for various purposes, including real-time AR systems [7, 5, 6, 8]. The problem with using this approach for continuously operating systems is that it requires high resolution images of the reflective sphere, which has to be placed in the scene.

As seen from the above brief review the standard approaches to determining scene illumination conditions are either to manually measure the light sources, or to photograph a reflective sphere placed in the scene. As stated our goal is to investigate whether images of surfaces naturally present in the scene can be used for estimating illumination, i.e., to determine if changing illumination can be detected from a video sequence.

3 Background

This work is based on a number of assumptions, which we will list together here. First of all our approaches are targeted at daytime outdoor scenarios, allowing us to assume that the illumination conditions are in effect completely governed by light from a directional source (the sun) and light from the sky hemisphere. In addition we assume that the imaged scene is static, that a complete 3D model of the scene is available, and that the camera is internally and externally calibrated, such that each pixel corresponds to a ray that can be traced to a unique 3D point in the scene.

Additionally the presented techniques assume that the scenes contain diffusely reflecting surfaces and that different normal directions are represented by these diffuse surfaces. We use the approach that the 3D model of the scene is manually annotated with information about which surfaces can be considered diffuse reflectors. As described in section 1 the techniques involve a reflectance calibration phase, and it is assumed that surface reflectances do not change after this calibration phase. This means that precipitation is not allowed, i.e., it is not allowed to rain or snow after reflectance calibration.

Both presented techniques are based on an assumption that the Phong Illumination Model can be used as a reasonable approximation to outdoor illumination conditions, and finally one of the techniques further assumes that the direction of the sun light is known at all times, computed automatically based on knowledge of date, time, and the camera's position in latitude and longitude.

In order to estimate the illumination conditions of a scene from 2D images of that scene a model of the image formation process is needed that describes the interactions between light and surfaces. Such a model requires the reflectance properties of the surfaces in the scene as well as a 3D description of the scene. Given a sufficient number of surfaces of different orientations it is then possible to set up a system of equations and solve for the variables describing the illumination conditions.

The remainder of this section describes an illumination model and the acquisition of reflectance properties. As stated we assume that the scene can be measured and modeled manually, using for example 3D Studio Max or similar 3D modeling software.

3.1 Phong Illumination Model

The Phong Illumination Model [14] is a local illumination model that is often used in computer graphics because it is fast to compute and gives reasonably realistic results although it is largely an empirical model. It is called a local illumination model because interreflections between surfaces are not considered. Interreflections – also known as global illumination effects – are approximated with an ambient term that allows for a global control of brightness in a scene. Besides the ambient term the Phong Model is composed of two reflection components that are due to direct illumination on a surface: a diffuse and a specular term. Diffuse reflections scatter light equally in all directions, i.e., the intensity at a point on a surface does not depend on the viewing direction. The diffuse reflections are modeled with Lambert's cosine law which states that the reflected light is proportional to the cosine of the angle between the surface normal and the incident light θ_i. The specular reflections depend on both the incident angle θ_i and the viewing angle θ_r, and may be modeled as proportional to the cosine of the angle α, see figure 2.

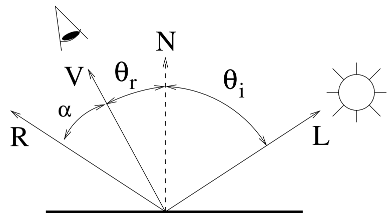

Figure 2: Image formation components of the Phong Illumination Model. L is a unit vector in the direction to the light source, N is the unit surface normal, V is the viewing direction, and R is the mirror-like reflected light.

The illumination of an outdoor scene may be modeled by one direct light source, the sun, and ambient light representing the skylight. The reflected light I_r approximated with Phong's Illumination Model is then given by the following equation:

$$I_{r,l} = k_{a,l} \cdot I_{a,l} + I_{i,l} \cdot (k_{d,l} \cdot \cos(\theta_i) + k_{s,l} \cdot \cos(\alpha)^m) \quad (1)$$

where k_a, k_d, and k_s are the reflection coefficients for the ambient, diffuse, and specular components, respectively, also called albedos which are described in the next subsection. I_a is the ambient illumination, I_i the direct light source, and m is a factor controlling the shininess of the surface. To handle color there is a separate equation for red, green, and blue, therefore the subscript $l \in \{R, G, B\}$.

In the following we assume pure diffuse surfaces and that $k_a = k_d$. Equation 1 then becomes:

$$I_{r,l} = k_{d,l} \cdot (I_{a,l} + I_{i,l} \cdot \cos(\theta_i)) \quad (2)$$

or using the vectors from figure 2:

$$I_{r,l} = k_{d,l} \cdot (I_{a,l} + I_{i,l} \cdot (\vec{N} \bullet \vec{L})) \quad (3)$$

3.2 Reflectance Properties

The reflectance properties of the surfaces are needed when using the scene as light probe. In the Phong Illumination Model (eq. 3) the reflectance properties are modeled with the scalar k_d. This is often called the *albedo* which is the reflectivity of a surface, or in other words the ratio of radiation reflected to the amount incident upon it. The reflected radiation may be expressed by the radiometric term *radiance* L_e which is the power leaving a surface per unit solid angle[2] and per unit surface area. The radiance can be measured using an image of a scene.

[2]The solid angle is the angle that, seen from the center of a sphere, includes a given area on the surface of that sphere. The value of the solid angle is numerically equal to the size of that area divided by the square of the radius of the sphere. It is measured in steradians [sr].

The radiometric term describing the received power per unit area, i.e., the power falling onto a surface, is the Irradiance E_e. For pure diffuse surfaces the albedo is then:

$$k_d = \frac{L_e}{E_e} \quad (4)$$

While it is rather easy to obtain the radiance from a scene using an image, the irrandiance requires knowledge of a 3D model of the scene and the light sources. One way to calculate the irrandiances for every pixel is then to synthesize (render) an image using the 3D model and setting all surface albedos to one.

4 Illumination Estimation under known Sun Position

In this first approach we take the illumination model presented in the previous paragraph and use it to model the measured pixel intensities from an image of the scene. If it is assumed that the system continuously can compute the unit direction vector to the light source relative to the scene coordinate system, then we arrive at a set of equations, one for each color channel for each pixel. These equations are linear in the ambient and the direct light, $I_{a,l}$ and $I_{i,l}$, respectively.

4.1 Approach

Let subscript j refer to the jth 3D point in the scene. Some points will in fact be in shadow and not receive direct light from the sun. Let S_j be a boolean parameter of value 1 if the jth point is in direct light, and 0 if it is in shadow. Furthermore, let C_j be a real number between 0 and 1, with the value of 1 if the jth point receives light from the entire hemi-spherical sky, and 0 if the sky is completely occluded seen from the jth point. The reflected light from the jth point can then be written as (the scene is small compared to the distance to the sun, so the unit direction vector to the sun is the same for all points in the scene):

$$I_{r,j,l} = k_{d,j,l} \cdot (C_j \cdot I_{a,l} + S_j \cdot I_{i,l} \cdot (\vec{N}_j \bullet \vec{L})) \quad (5)$$

The ambient occlusion factor, C_j, can be computed a priori for all points in the scene. Given knowledge of the sun's position the shadow masking parameter S_j can be computed at run-time for all points in the scene. From the offline reflectance calibration we know the albedos, $k_{d,j,l}$, of all diffusely reflecting points. The surface normal for all points, N_j is known from the 3D model of the scene. The direction vector to the sun, L, can be computed given: 1) the date, 2) the time, 3) the Earth position in latitude and longitude of the scene origin, 4) and the direction of North in the scene, [1]. The only unknowns in eq. 5 are the 6 parameters for ambient and direct light, $I_{a,l}$ and $I_{i,l}$.

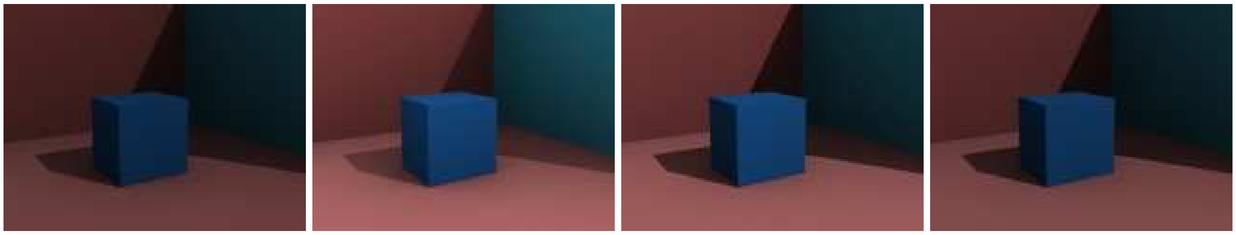

Figure 3: Frames 0, 29, 59 and 89 from synthetic test sequence with known illumination changes. The direction of the direct light source is not changing but the ambient and source emittances both change over the sequence.

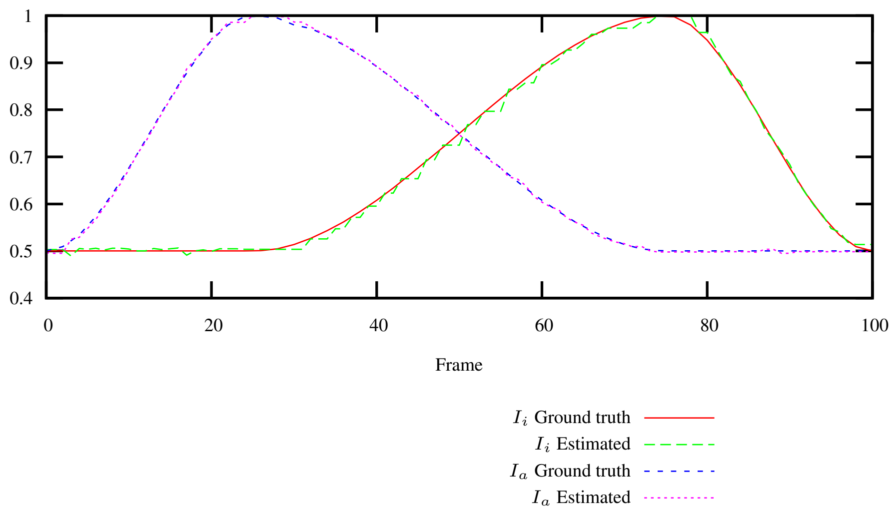

Figure 4: Comparison of estimated and true values for direct and ambient radiances for the synthetic test sequence shown in figure 3. All values are normalized to a maximum of 1.

Using the assumption that we know what surfaces in the scene can be considered diffuse and that the camera is calibrated to the scene we can also find those pixels that correspond to diffuse surfaces. The RGB pixel values of such a pixel is denoted $P_{j,l}$. If the camera is radiometrically linear the measured pixel values are some camera constant K times the reflected radiance from the corresponding scene point:

$$P_{j,l} = K \cdot I_{r,j,l} \quad (6)$$

Thus, by picking pixels from the image we can set up a system of linear equations in $I_{a,l}$ and $I_{d,l}$ of the form:

$$1/K \cdot P_{j,l} = k_{d,j,l} \cdot (C_j \cdot I_{a,l} + S_j \cdot I_{i,l} \cdot (\vec{N}_j \bullet \vec{L})) \quad (7)$$

In our implementation of this framework we at random select on the order of a few hundred pixels evenly distributed across the image (among those pixels that correspond to diffuse surfaces). It is important that the pixel population represents both areas in shadow (only ambient light) and in direct light (both ambient and direct light). The camera scene radiance to pixel value scaling factor K is of course unknown, but a system of equation of the form of eq. 7 allows us to estimate scene illumination up to a scaling factor.

4.2 Experiments and Results

The presented framework has been tested extensively on both synthetic and real images. Figure 3 shows a few frames from a synthetically generated sequence where a simple scene has been rendered with known ambient and direct intensities. Correspondingly, figure 4.1 shows the estimated intensities. As seen synthetic data results in near perfect estimations. The same scene has been tested with generating a sequence where a yellow ball is falling into the scene and bouncing out again in order to test how the illumination estimation procedure reacts to dynamic objects in the scene, thus violating the static scene assumption. The estimation results from this scenario is not shown, but due to the extraction of a large number of sample points across the entire image the illumination estimation is very stable and only in minor degree affected by the dynamic object.

To test the approach on real data a 2 hour time-lapse sequence has been acquired with one frame every 20 seconds. Figure 5 shows select frames from the sequence. Naturally, we do not have ground truth data for the illumination conditions in this real scene, but figures 4.2 and 4.2 show the estimated ambient and direct light. The estimated illumination has been verified

Figure 5: Frames 1, 72, 134, 201, and 259 from real test sequence covering approximately 2 hours of moving sun and changing cloud cover.

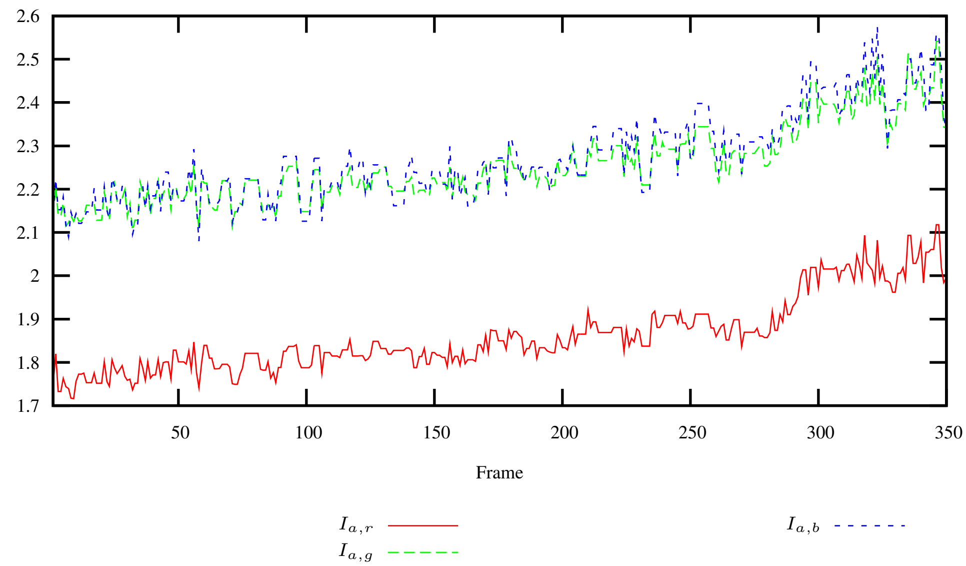

Figure 6: Estimated values for ambient radiance for the real test sequence shown in figure 5. Notice that the ambient light in the scene has a relatively low red (R) component as is expected for a sky with only partial cloud cover.

quantatively by taking the known albedos of surfaces in the scene, illuminating these surfaces with the estimated light and comparing these values with the real sequence pixel value for the same surface. These tests (not shown) demonstrate that the estimated illumination follows the real scene illumination quite accurately apart from a tendency to over-estimate the red component of the direct light with approximately 10%. This may be caused by mis-estimating the albedo of the dominant red brick wall due to an in-accurate determation of the illumination conditions at time the image was acquired for albedo estimation.

In addition to quantitative tests on real data qualitative tests of the estimated illumination has been evaluated by rendering virtual objects into the scenes and visually judging the quality of the virtual shading. Especially for sequences with very dynamic lighting conditions it is clearly seen that the estimated light results in consistent shading of virtual objects. Two frames from such a test were shown in figure 1.

The current implementation of this estimation technique can run the estimation at about 10 frames per second and is thus easily able to respond to the illumination condition changes an outdoor AR system would experience.

5 Illumination Estimation under unknown Sun Position

This section describes the estimation of the illumination conditions including the sun direction from an image of a scene given a 3D model of the scene and the albedos of the surfaces in the scene, and assuming that the illumination model in equation 2 can be used to approximate outdoor illumination conditions.

5.1 Approach

The illumination conditions may be estimated by setting up a sufficient number of equation 3, and solving this non-linear system of equations for the unknowns $\vec{L}$, $I_{a,l}$, and $I_{d,l}$ ($l \in \{R, G, B\}$). Thus, there are nine unknowns. Assuming that the distance r to the sun is known the estimation of $\vec{L}$ reduces to the two angles (azimuth, φ, and zenith, θ). In order to solve a non-linear system of equations one may formulate it as a least squares problem and then use a numerical optimization. Let $f_i(x)$ be the minimization function that should converge to zero. The minimization function for the Phong Illumination Model is given in equation 8:

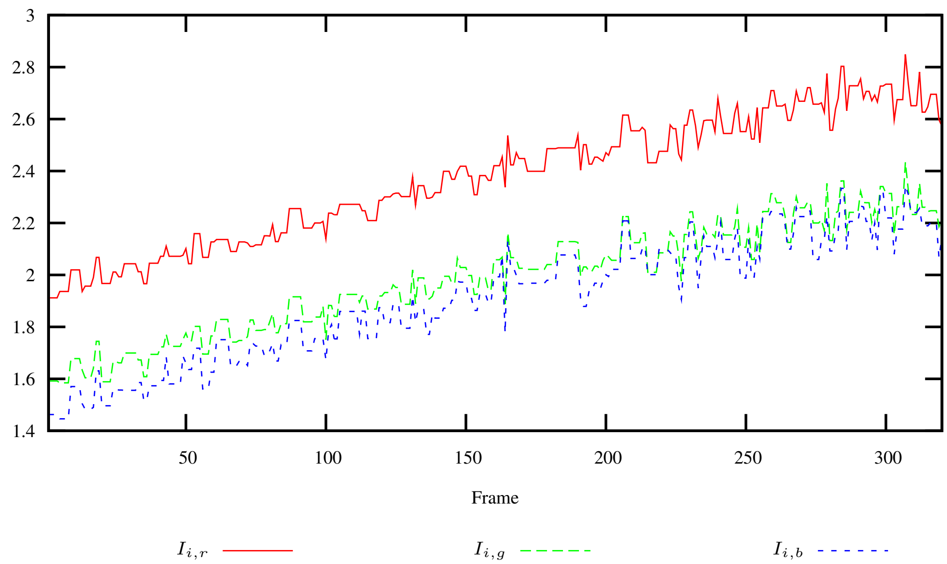

Figure 7: Estimated values for direct radiance for the real test sequence shown in figure 5.

$$f_{j,l}(x) = k_{d,j,l}(I_{a,j,l} + I_{d,j,l}\overline{\cos\theta_j}) - I_{r,j,l} = 0 \quad (8)$$

where

$$\overline{\cos\theta}_j = \begin{cases} 0 & , \quad \theta > \pi/2 \\ \dfrac{\vec{N}_j \bullet \vec{L}}{\|\vec{N}_j\| \cdot \|\vec{L}\|} & , \quad 0 \le \theta \le \pi/2 \end{cases} \quad (9)$$

and

$$\vec{L} = \begin{pmatrix} r \cdot \cos\varphi \cdot \sin\theta \\ r \cdot \sin\varphi \cdot \sin\theta \\ r \cdot \cos\theta \end{pmatrix} \quad (10)$$

5.2 Experiments and Results

The evaluation of computer vision methods using real image data is often difficult due to the lack of ground truth. In this work the estimation of a 3D model and of the albedos introduces an error which makes the evaluation of the illumination estimation inaccurate. Therefore the main evaluation was done using synthetic image data that were generated from a 3D scene description including light positions and object reflectances. They were rendered using a ray-tracer (Radiance [2]) that generates radiometric correct images including global illumination effects. Radiance supports all kinds of light sources among them a daylight model to create realistic illumination in outdoor scenes with sun and skylight. Figure 8 shows an example of a rendered image that was used for evaluation.

All estimations were done using MATLAB's *lsqnonlin* for solving non-linear least squares problems.

Using the daylight model images were rendered for illumination conditions from sunrise to sunset. The estimation errors are shown in figure 9. In most of the estimations the error lies around a few degrees $[0.6-2.5°]$ except from three estimations, which are at 9, 9:30 and 17 o'clock. With a look at the image of these time periods it can be seen that some of the measurement pixels were occluded resulting in no direct illumination.

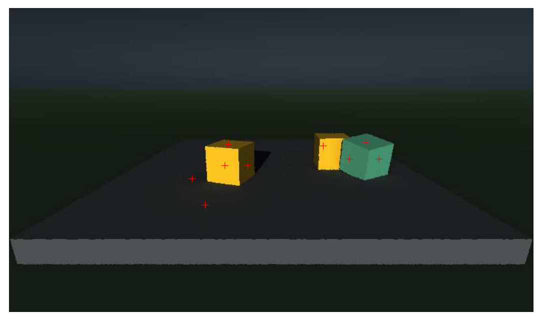

Figure 8: Synthesized image using the daylight model.

These estimation tests have been extended with sun positions across the entire hemisphere. The azimuth ranges $[0 \rightarrow 2\pi]$ and the zenith ranges $[0 \rightarrow \frac{1}{2}\pi]$. All in all 339 images have been rendered in Radiance. The scene illustrated in figure 8 has also been used in this experiment.

The error in these experiments are given as the total angle between the estimated and actual angles. Total error means the two actual angles (azimuth φ and zenith θ) seen in relation to the two estimated, which has been calculated from equation 11.

$$\Delta E = \arccos\left(\frac{\vec{a} \bullet \vec{e}}{|\vec{a}||\vec{e}|}\right) \quad (11)$$

The two vectors for actual ($\vec{a}$) and estimated ($\vec{e}$) are calculated from equation 12, where $r = 1$ since it is only the angular difference that is of interest.

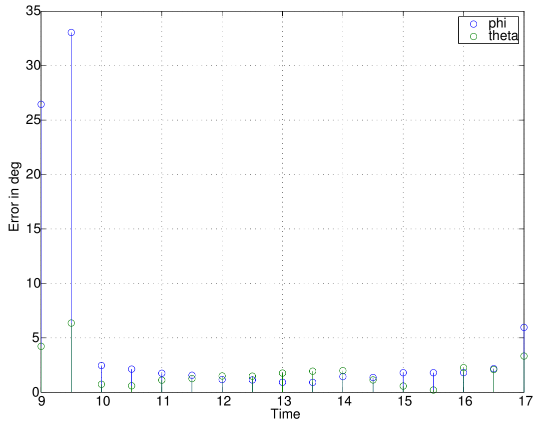

Figure 9: Estimation error in degrees as a function of the time for azimuth (φ) and zenith (θ) respectively.

$$\vec{a} = \begin{pmatrix} r \cdot \cos\varphi \cdot \sin\theta \\ r \cdot \sin\varphi \cdot \sin\theta \\ r \cdot \cos\theta \end{pmatrix} \quad (12)$$

The total angular error for all azimuth and zenith angles of the sun is shown in figure 10 where the error is indicated by the height of the small circles. The azimuth angle is given as the angle in the plan, and the zenith is zero in the center of the plot and increasing with distance to the center.

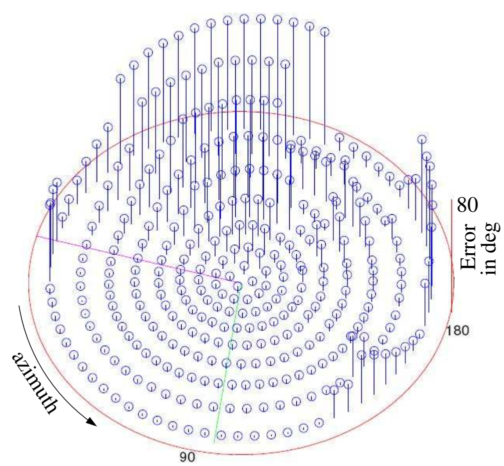

Figure 10: Total angular error over all azimuth and zenith angles.

The camera viewing direction is indicated by the green line around the an azimuth angle of $90°$. It can be seen that the error is rather low when the illumination direction is close to the camera direction, whereas the error increased significantly (up to $80°$) when the illumination is opposite to the camera. This is due to the reduced number of directly illuminated surfaces.

Besides simulated image data we tested the method on real images of building bricks. These bricks have rather diffuse reflectance properties. Figure 11 shows an image of a real scene that was used for illumination estimation, and figure 12 show a part of that scene with a virtual shadow that was simulated using the estimated illumination direction.

Figure 11: Real image used for illumination estimation.

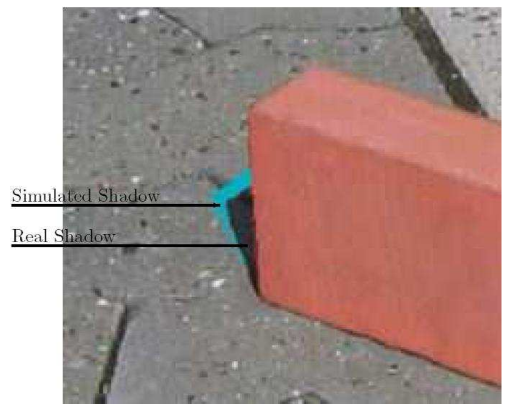

Figure 12: Real scene showing both, the real shadow casted by the brick and a virtual shadow casted by the brick. The virtual shadow is rendered using the estimated sun position.

Figure 13 show an example application where the estimated illumination direction was used to augment the scene with four virtual vases.

6 Conclusions

In this paper two methods were proposed and tested to estimate the illumination conditions of a real outdoor scene while using the scene itself as a light probe, i.e., no additional light probe has to be placed into the scene. The methods work on single view images and require a 3D model of the scene as well as the reflectance properties of the surfaces present in the scene. The preliminary results show their applicability to Augmented Reality.

In future work we aim at combining several complementary methods in order to achieve more robust illumination estimation. Furthermore, we will look into possibilities to reduce the offline calibration, e.g., using reflectance models for common everyday outdoor objects. These objects may be recognized by some automatic object recognition and be used as light probes.

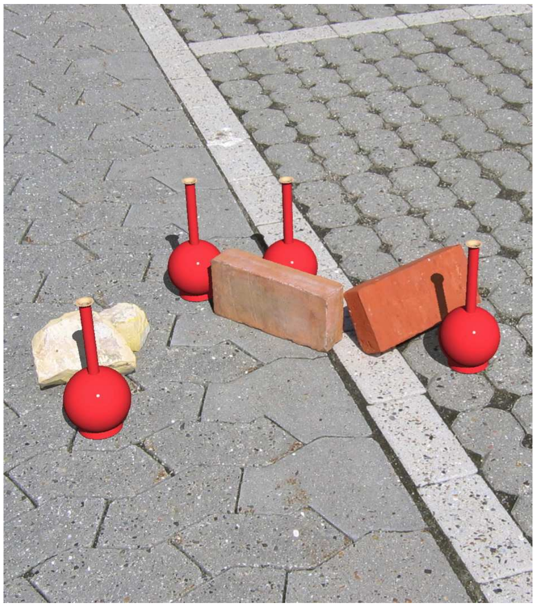

Figure 13: Real scene augmented with virtual vases.

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 CoSPE project (26-04-0171) under the Danish Research Agency. This support is gratefully acknowledged.

References

[1] Computing planetary positions – a tutorial with worked examples. http://www.stjarnhimlen.se/comp/tutorial.html.

[2] Radiance synthetic imaging system homepage. http://radsite.lbl.gov/radiance/HOME.html.

[3] Samuel Boivin and André Gagalowicz. Image-based rendering of diffuse, specular and glossy surfaces from a single image. In *Proceedings of ACM SIGGRAPH 2001*, Computer Graphics Proceedings, Annual Conference Series, pages 107–116, August 2001.

[4] Samuel Boivin and André Gagalowicz. Inverse rendering from a single image. In *IS&T's First Europ. Conf. on Color in Graphics, Images and Vision*, pages 268–277, Poitiers, France, April 2002.

[5] P. Debevec. 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.

[6] P. Debevec. Tutorial: Image-based lighting. *IEEE Computer Graphics and Applications*, pages 26 – 34, March/April 2002.

[7] P. Debevec and J. Malik. Recovering high dynamic range radiance maps from photographs. In *Proceedings: SIGGRAPH 1997, Los Angeles, CA, USA*, August 1997.

[8] S. Gibson, J. Cook, T. Howard, and R. Hubbold. Rapic shadow generation in real-world lighting environments. In *Proceedings: EuroGraphics Symposium on Rendering, Leuwen, Belgium*, June 2003.

[9] K. Jacobs and C. Loscos. State of the art report on classification of illumination methods for mixed reality. In *EUROGRAPHICS*, Grenoble, France, September 2004.

[10] M. Kanbara and N. Yokoya. Real-time estimation of light source environment for photorealistic augmented reality. In *Proceedings of the 17th International Conference on Pattern Recognition, Cambridge, United Kingdom*, pages 911–914, August 2004.

[11] Celine Loscos, George Drettakis, and Luc Robert. Interactive virtual relighting of real scenes. *IEEE Transactions on Visualisation and Computer Graphics*, 6(3), July 2000.

[12] V. Masselus, P. Dutre, and F. Anrys. The free-form light stage. In *Proceedings of the 13th Eurographics Workshop on Rendering*, pages 247–256, Pisa, Italy, June 2002.

[13] G. Patow and X. Pueyo. A survey of inverse rendering problems. *COMPUTER GRAPHICS forum*, 22(4):663–687, 2003.

[14] Bui Tuong Phong. Illumination for computer generated pictures. *Communications of the ACM*, 18(6):311–317, June 1975.

[15] I. Sato, Y. Sato, and K. Ikeuchi. Illumination from shadows. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 25(3):290–300, March 2003.

[16] Yizhou Yu, Paul Debevec, Jitendra Malik, and Tim Hawkins. Inverse global illumination: Recovering reflectance models of real scenes from photographs. In *Proceedings of SIGGRAPH 99, Computer Graphics Proceedings, Annual Conference Series*, pages 215–224, Los Angeles, California, USA, August 1999. Addison Wesley Longman.

[17] Yizhou Yu and Jitendra Malik. Recovering photometric properties of architectural scenes from photographs. In *Proc. of SIGGRAPH 98*, pages 207–217, Orlando, Florida, USA, July 1998.