

# Real Time Augmentation of Omni-directional Images with Consistent Lighting

Thorsten Grosch  
Institut für Computervisualistik  
Universität Koblenz-Landau  
56070 Koblenz  
Tel.: +49 (0)261 287 2734  
Fax: +49 (0)261 287 2735  
E-Mail: grosch@uni-koblenz.de

**Abstract:** Panoramic images can be used as an image-based method for navigating in a real environment. Recent publications have shown that it is possible to augment photographs with correct lighting. We combine these two approaches to an interactive tool with real-time modifications of the real scene maintaining consistent lighting. The possible changes are: Insertion of virtual objects, new light sources and changes of materials. The only user input is a high-dynamic-range photograph of a light probe. A 3D model is reconstructed from the light probe image using object-based reconstruction methods. The differential rendering for all user manipulations is performed with the latest generation of graphics hardware. Finally, we remove the restriction to a fixed viewpoint by projecting the environment map to the geometry in the fragment program and we present a hole-filling algorithm for the parts of the geometry which were not visible in the light probe image

**Keywords:** Augmented Image Synthesis, Augmented Reality, Virtual Reality, Radiosity

## 1 Introduction

Augmenting real images with correct lighting conditions becomes a more and more important area of research. Example applications are virtual furniture in a real room or new virtual buildings in a real landscape. Also, virtual characters in a movie do not look believable if the illumination is wrong. The goal of all these changes is to modify the image with changes that can not be distinguished from the real parts. This means that all changes in lighting must be calculated. Therefore, a time-consuming lighting simulation must be performed. On the other hand, the user would like to interactively modify the real photograph without a long response time. The graphics hardware has rapidly developed in the last few years and can be used to display the changes in real-time. Many approaches augment a single photograph which shows only a small part of the whole environment. We therefore decided to augment an environment

map and use a panoramic image viewer to display the augmented photograph. This enables the user to look around and manipulate the whole surrounding.

## 2 Previous Work

Adding virtual objects in photographs with correct lighting conditions was introduced by Fournier et al. [FGR93]. Drettakis et al. [DRB97] improved the update rates when moving a virtual object and Loscos et al. [LFD99] separated direct and indirect light for better display quality. The first system directly using the graphics hardware for differential rendering was presented by Gibson and Murta [GM00] who extracted a set of directional light sources from a light probe image. The shadows of the virtual objects are displayed with shadow mapping in a multi-pass algorithm. The most recent approach was presented by Gibson et al. [GCH03]. Here, virtual objects with spatially varying light and shadows are illuminated correctly in real-time. Using high-dynamic-range images [DM97] of a light probe was introduced by Debevec [Deb98]. The light probe is a mirrored ball that captures the illumination from all incoming directions. This omni-directional image can be used for illuminating synthetic objects under real lighting conditions. Another alternative was presented by Sato et al. [SSI99], they use a fish-eye lens to capture the illumination from the upper hemisphere in one photograph. An image-based system for navigation in virtual environments was introduced by Chen [Che95]. A set of photographs is taken from a fixed viewpoint in different directions and stitched together to a cylindrical/spherical environment texture. The user can pan and zoom the camera, moving is only possible by switching between different panorama points or special object movies for hot spots. Yu et al. [YDM99] showed augmented panoramic images with correctly illuminated virtual objects and new lighting conditions, calculated with a complex lighting simulation. Wong et al. [WOF03] relighted a synthetic panorama image by precomputing basis-panorama images with a directional light source and combining the basis images.

In this paper, we develop a system with more user manipulations than rotating the camera. The user can move virtual objects, add new virtual lights, change existing materials and move the camera. The rest of the paper is organized as follows: In section 3 we explain the necessary preparations for the augmentation, section 4 describes the details about all possible user manipulations and in section 5 we conclude and give some ideas about future work.

## 3 Preparation

The first step is to take a high-dynamic-range photograph of a light probe using the algorithm described in [DM97]. Next, we convert the mirror ball image to a latitude/longitude environment map and use it as a texture for a sphere. We use a fragment program with a simple gamma correction as a tone mapper to display the floating point texture pixels. Usually, a high resolution of the environment map is necessary, otherwise the resulting image looks blurry. To further

improve image quality, we use a sharpening filter to create a laplacian-enhanced image of the HDR environment map.

We employ the latest generation of graphics hardware because of the support of floating point textures. All calculations with input data can be done with floating-point precision and only the result is mapped to the [0,1] range. Because many of the useful texture functions like filtering, cube mapping or blending do not work with floating-point textures, we implemented these functions in our own vertex and fragment programs. Our system was developed using C++ with OpenGL and the RenderTexture library [Har04] for the floating point textures. The vertex and fragment programs were created with Cg. All timing values were measured with a dual Xeon PC, running at 2.8 MHz, equipped with a GeForceFX 5950 graphics card. The environment map resolution is 2048 x 1024 pixels, all snapshots were made with a screen resolution of 400 x 400 pixels.

### 3.1 Geometry Reconstruction

The first step of our algorithm is the geometry reconstruction from a single light probe image. Reconstruction from a single image is a well-known problem: [GH00], [CRZ99], [WBS01], [Heu98], [SM99] present object-based algorithms for geometry reconstruction from a single image. The intrinsic parameters, especially the focal length can be determined from the position of the vanishing points in the image. With the known focal length and direction vectors for the three main axes, simple geometric primitives like rectangles or boxes can be placed in the scene using constraints like coplanarity, parallelism, orthogonality and symmetry [Heu98]. Shum et al. [SHS98] describe how to reconstruct geometry from a panoramic image mosaic. The extension for reconstructing geometry from a light probe image is straight forward: In our implementation, the image is first converted to the latitude/longitude environment map format (Fig. 1) and mapped as a texture to a sphere. The camera is placed in the center of the sphere with a perspective projection. An initial viewing direction is set, so that at least two vanishing points are visible. The resulting image is used as the source image for the reconstruction. When all visible objects for the initial viewing direction are reconstructed, the camera is rotated and the reconstruction can be continued using all the information already computed like the plane equations for floor and walls. Fig. 1 shows the geometry reconstruction: The coordinate system is defined by picking parallel lines and assigning them to vanishing points. Reconstruction of the focal length is not necessary because the camera angle was set in the projection of the perspective camera. With the known focal length  $f$  and two vanishing points, the direction vectors of the world coordinate system can be calculated as [Gro01]:

$$\bar{x} = \begin{pmatrix} -x_{\infty} / f \\ -y_{\infty} / f \\ 1 \end{pmatrix} \quad \bar{z} = \begin{pmatrix} -x'_{\infty} / f \\ -y'_{\infty} / f \\ 1 \end{pmatrix} \quad \bar{y} = \bar{z} \times \bar{x}$$

The user defines two points with a known distance to determine the unknown scaling factor of the geometry. Usually, we take the distance between floor and ceiling to determine the correct size of the 3D geometry. Now, simple geometric primitives can be placed in the coordinate system. A whole room reconstructed from a light probe image is shown in Fig. 1. In this example, the reconstruction process took about 30 minutes.

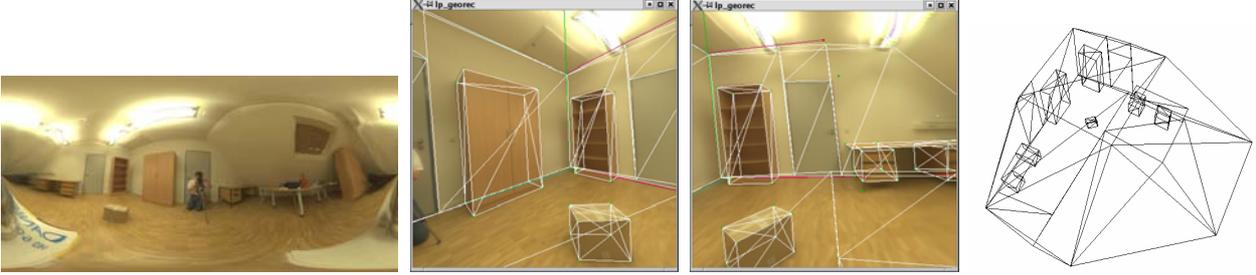


Figure 1: Left image: The latitude/longitude environment map. The second image shows the camera calibration and the reconstructed geometry, in the third image the camera is rotated and the new geometry is reconstructed using all the information from the first view. The whole geometry reconstructed from the environment map image is shown on the right image.

### 3.2 Reflectance Estimation

For correct differential rendering we need to subtract radiance of the light sources in the new shadow regions [GCH03]. Therefore, we need to know the reflectance values of all geometry objects. In our implementation, all polygons are subdivided into patches. If we assume a diffuse environment, the reflectance value of a non-emitting patch is simply its radiosity divided by the irradiance

$$\rho = B / E \quad (\text{EQ1})$$

The radiosity values can be obtained from the HDR image ( $B = \pi L$ ), where L is the average radiance of all pixels inside the patch. If we assume that all important light sources are visible in the light probe image, we can also calculate the irradiance

$$E_e = \sum_s B_s F_{es} \quad (\text{EQ2})$$

and the reflectance value for patch e can be determined as [Deb98]

$$\rho_e = B_e / \sum_s B_s F_{es} \quad (\text{EQ3})$$

During geometry reconstruction, we assign polygons to material groups. The reflectance value is calculated for each patch, and afterwards each material group gets the average reflectance value of all patches with the same material. The reflectance value is correct for patches near the light probe position, for other positions we do not gather the complete irradiance. Instead we gather

irradiance from some hidden regions which are not visible in the light probe image. We use a simple iteration scheme to fill these invisible regions with useful values and to improve the reflectance estimate:

1. Calculate reflectance estimates for all patches which are visible in the light probe image by using (EQ3)
2. Calculate an average reflectance value for each material
3. Copy the average reflectance values to all patches in hidden regions
4. Calculate irradiance values  $E$  for all patches in hidden regions with (EQ2) and compute a radiance estimate  $L = \rho E / \pi$
5. Go back to step 1 if not converged

Tab. 1 shows the result of the diffuse reflectance estimation for the three color channels. In our case, there are no significant changes after two iterations of our algorithm. The right column shows the real reflectance values, calculated by measuring radiance and irradiance at several locations and calculating the average of all measurements.

	First Iteration	Second Iteration	Real
Wall	0.98 0.95 0.94	0.91 0.88 0.86	0.79
Door	0.52 0.49 0.45	0.51 0.49 0.44	0.49
Cupboard	0.53 0.32 0.19	0.53 0.32 0.19	0.37
Floor	0.46 0.31 0.17	0.46 0.31 0.16	0.38
Table	0.57 0.43 0.28	0.53 0.39 0.25	0.38
Box	0.64 0.48 0.29	0.64 0.48 0.29	0.29

Table 1: Reflectance Values

### 3.3 Finding Important Light Sources

Currently, we subdivide the geometry in patches based on an area threshold. For each patch, we calculate the average radiance and then we select  $N$  patches with the highest luminous flux  $\phi = BA$ , like in progressive refinement radiosity [CCW88]. This set of patches is considered as *important* light. The shadows from virtual objects are generated from these patches and the direct lighting from these patches is recomputed when the virtual objects are moved.

## 4 Image Manipulations

### 4.1 Virtual Objects

For displaying virtual objects we use a modification of the method described in [GCH03]: We first precompute an irradiance volume [GSH98] to capture the irradiance for the whole scene. Specular materials are displayed with floating-point reflection maps. The radiance of the important light sources is stored at the receiver patch vertices and interpolated. In shadow regions it is subtracted from the background. No global illumination simulation is used here, the contributions of the important light sources  $i = 1..N$  are subtracted from the real brightness values from the photograph [GCH03]:

$$\Delta L_i = \frac{\rho}{\pi} E_i \quad L'_{\text{pixel}} = L_{\text{pixel}} - \Delta L_i$$

We use shadow volumes [Cro77] to display the shadows cast by virtual objects. In contrast to shadow maps, shadow volumes have no aliasing effects at the shadow borders and the light sources are not restricted to spotlights.

When using the irradiance volume, the problem of light and shadow leaks arises: This means that shadow regions are interpolated in neighbouring regions without shadows. Therefore, we separate the *important* light from the *unimportant* light. For each grid point, two values are precomputed and stored: The first value is the normal irradiance gathered from all directions, the second value is the irradiance gathered only from unimportant sources of light. The important sources are all patches selected as described in (3.3), the unimportant sources are all remaining patches. We assume that the resulting irradiance gathered in the second value is slowly varying and the error due to linear interpolation is small. When moving the virtual object, we use two different ways to display the virtual object. The first version is fast: The virtual object is displayed by interpolating the first value of the irradiance volume. The shadows are created by building the shadow volume with all silhouette edges of the virtual object and the precomputed radiance is subtracted from the background image in the shadow regions. For a better, but more time consuming display, the second value is used and the important light is recomputed for all important sources in the vertex program. When displaying the shadows we distinguish between shadows cast on the virtual object and shadows cast on real objects. Shadows on the virtual object are determined by creating the shadow volume with all silhouette edges of real and virtual geometry. The resulting shadow is a mask for the direct light computation in the vertex program. Shadows on real objects are determined by creating two shadow images: The first image contains the shadows of the real geometry, the second image contains the shadows of the real scene with the virtual object. When subtracting the two shadow images, we get a mask for the new shadows cast by the virtual object on real objects [GM00]. In these shadow regions we subtract the precomputed radiance values from the HDR background image. Fig. 3 shows the difference between the two versions.



Figure 2: A virtual armchair, illuminated by 32 important light patches, is inserted. The second image shows the armchair with a floating-point reflection map. The images on the right show the virtual object and the light sources. The chair can be moved at 3..12 fps (32..2 important lights)



Figure 3: The armchair is moved under the desk. Note that the shadows on the armchair fit seamlessly to the real shadows. 32 important light patches are used for direct lighting and shadows. The right image shows the fast image display. The chair can be moved at 3.0 fps with the simple algorithm and at 1.0 fps with the improved display.

Anti-aliasing techniques significantly improve the realism of the virtual object as shown in [GMK03]. We use a simple gaussian blur filter in the fragment program, applied to all virtual pixels in the composited image.

## 4.2 New Light Sources

Beside virtual objects, the user can insert new light sources in the virtual environment. We use light goniometrics [Sch00] for the description of new lights since this is one of the most realistic approximations of a real light source and better than simple point or area lights often used in computer graphics. A light goniometric describes the luminous intensity of a light source for any outgoing direction  $(\theta, \varphi)$ . If the data for the desired light is not available, it can be reconstructed from a set of photographs [GMK03]. The irradiance  $E$  of a light goniometric is

$$E = \frac{I(\theta, \varphi) \cdot \cos \theta_e}{d^2} \quad (\text{EQ5})$$

where  $d$  is the distance between receiver and light source and  $\theta_e$  is the angle between incoming direction and surface normal. To display the direct light of a light goniometric in real-time, we convert the luminous intensity values  $I(\theta, \varphi)$  to a cube-map in a pre-process. The 3D position of a pixel is determined by copying the world coordinate to a texture coordinate in the vertex program, because it is interpolated for the fragment program. The fragment program calculates

the light vector by subtracting the pixel position from the light position and uses the standard cube mapping equations to select the corresponding pixel from one of the sides of the cube map. To handle light sources with arbitrary brightness, we use floating point textures for storing the luminous intensity values in the pixels.



Figure 4: A new light is inserted and can be moved and rotated interactively by the user at 12 fps. Note the shadows of real objects on real objects. The right image shows an artificially generated light goniometric with different colors, stored in one cube map.

The fragment program calculates the irradiance for the current position using (EQ5). We calculate the shadow volume for the whole scene with virtual objects for the current light position and draw the resulting shadow with black on a white background. The shadow image is rendered in a texture and used as a mask for the transfer function. A mask of 1 describes a pixel which is lit, a mask value of 0 describes a pixel in shadow. We multiply this mask with the texture containing the unoccluded light goniometric radiance to obtain light and shadows of the new light source (Fig. 4).

### 4.3 Changing Materials

The third possible user manipulation is the adjustment of the reflectance values. The user can point to an object and change the computed diffuse reflectance value of the corresponding material group. When changing the diffuse material from  $\rho$  to  $\rho'$ , the radiance is changing from  $L$  to  $L'$ :

$$L = \frac{\rho}{\pi} E \quad L' = \frac{\rho'}{\pi} E = \frac{\rho'}{\rho} L \quad (\text{EQ6})$$

The direct material changes can be visualized immediately, we simply draw the reconstructed geometry with colors set to  $(\rho'_R / \rho_R, \rho'_G / \rho_G, \rho'_B / \rho_B)$  where  $\rho'$  is the user defined reflectance and  $\rho$  is the reconstructed reflectance value. These relative reflectance values are rendered into a float texture and the HDR background image is scaled with this texture in the fragment program. To simulate color bleeding of the new surface color, we calculate delta radiosities  $\Delta B = B' - B$  for all patches with a changed material and send them to all other patches. This process is similar to dynamic radiosity approaches [Che90][DSG90][MS94]. A one-bounce approximation gives good results in most cases, but additional bounces can be computed.



Figure 5: First image: The material of door and box are changed. The resulting color bleeding can be calculated in 0.3 - 1.0 seconds. Second image: The camera is moved to a new position. The environment texture is projected onto the geometry in the fragment program. The top of the furniture, parts of the wall and the area behind the box on the floor are not visible in the light probe image. These areas are filled with the algorithm described in (3.2). Third image: A combination of all possible manipulations. Right image: Omni-directional depth map.

#### 4.4 Camera Movement

Omni-directional images allow the user to look in every possible direction but it is not possible to move the camera. Because we reconstructed a 3D model and filled the invisible regions, we can remove this limitation by projecting the environment map onto the reconstructed geometry. This is usually done by extracting a texture for each polygon, which is rectified and mapped on the polygon [GMM00]. We directly use the latitude/longitude image to project the background image to the geometry. The fragment program gets the interpolated 3D position in a texture coordinate, converts this direction vector in  $(\theta, \varphi)$  coordinates and fetches the texture pixel from the environment map. For the invisible parts of the geometry we display the gathered radiance estimate as described in (3.2). This distinction is done in the fragment program by using an omni-directional depth-map storing the z-values for each direction  $(\theta, \varphi)$  (Fig. 5). If the z-value in the texture is smaller than the real z-value in the fragment program, we use the radiance estimate. Otherwise, we take the pixel out of the environment texture. This algorithm is similar to shadow mapping

## 5 Conclusions and Future Work

We presented a new algorithm for real-time augmentation of omni-directional images with consistent lighting. Three types of changes are possible: Addition of virtual objects, new light sources and material changes. All manipulations can be done in real-time using the latest generation of graphics hardware. As future work we plan to improve the detection of important light sources and to include hardware soft-shadow algorithms to reduce the number of important light sources. The hole-filling algorithm can be improved by analyzing the texture in visible regions and modulating the radiance estimate with the texture in the hidden regions [LFD99]. The hole-filling algorithm can also be used when real objects are removed or moved to a new position. Another important topic is the rapid display of changes in indirect light of new light

sources. Finally, we would like to improve the appearance of the virtual objects by using better reflection models which can be calculated with the graphics hardware.

## Acknowledgements

We would like to thank Markus Geimer and Angelika Bräunig de Colan for the corrections. Many thanks go to Oliver Abert for creating the videos.

## 6 References

- [CCW88] COHEN M.F., CHEN S.E., WALLACE J.R., GREENBERG D.P.: A Progressive Refinement Approach to Fast Radiosity Image Generation, Computer Graphics (SIGGRAPH 88 Proceedings), Vol22, No.4, August 1988, pp.75-84
- [Che90] CHEN S.E.: Incremental Radiosity: An Extension of Progressive Radiosity to an Interactive Image Synthesis System, ACM Computer Graphics (SIGGRAPH 1990 Proceedings) Vol 24, No.4, pp.135-144, August 1990
- [Che95] CHEN S.: Quicktime VR: An image-based approach to virtual environment navigation. Proceedings of SIGGRAPH 1995, pp 29-38
- [Cro77] CROW F.: Shadow Algorithms for Computer Graphics, Proceedings of ACM SIGGRAPH 1977, pp. 242-248
- [CRZ99] CRIMINISI A., REID I., ZISSERMAN A.: Single view metrology, ICCV 1999, pp. 434-441
- [Deb98] DEBEVEC P.: Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. Proceedings of SIGGRAPH 98 (July 1998), 189-198
- [DM97] DEBEVEC P.E., MALIK J.: Recovering high dynamic range radiance maps from photographs. Proceedings of SIGGRAPH 97 (August 1997), 369-378
- [DRB97] DRETTAKIS G., ROBERT L.; BOUGNOUX S.: Interactive Common Illumination for Computer Augmented Reality, Eighth Eurographics Workshop on Rendering (1997). In: Dorsey, J.; Slusallek, P. (Hrsg.): Rendering Techniques '97. Springer Verlag, 1997, pp. 45-56
- [FGR93] FOURNIER A., GUNAVAN A.S., ROMANZIN C.: Common Illumination between real and computer generated scenes. In Proceedings of Graphics Interface '93 (San Francisco, CA, May 1993), Morgan Kaufmann, pp. 254-262

- [GH00] GIBSON S., HOWARD T.: Interactive Reconstruction of Virtual Environments from Photographs, with Application to Scene-of-Crime Analysis, Proc. of ACM Symposium on Virtual Reality Software and Technology 2000, Seoul, Korea, October 2000, pp. 41-48
- [GM00] GIBSON S.; MURTA A.: Interactive Rendering with Real-World Illumination, Eleventh Eurographics Workshop on Rendering. In: Péroche, B.; Rushmeier, H. (Hrsg.): Rendering Techniques 2000. Springer Verlag, 2000, pp. 365–376
- [GCH03] GIBSON S., COOK J., HOWARD T., HUBBOLD R.: Rapid Shadow Generation in Real-World Lighting Environments, Eurographics Symposium on Rendering 2003
- [GMK03] GROSCH T., MUELLER S., KRESSE W.: Goniometric Light Reconstruction for Augmented Reality Image Synthesis, Graphiktag im Rahmen der GI Jahrestagung, Frankfurt am Main, Germany, September 2003
- [Gro01] GROSCH T.: Interaktive Erweiterung und Beleuchtungsrekonstruktion von Fotos mit Radiosity, Diploma Thesis, Technische Universität Darmstadt 2001
- [GSD90] GEORGE D.W., SILLION F.X., GREENBERG D.P.: Radiosity Redistribution for Dynamic Environments, IEEE Computer Graphics & Applications, Vol.10, No.4, July 1990, pp.26-34
- [GSH98] GREGER G., SHIRLEY P., HUBBARD P.M., GREENBERG D.: The irradiance volume, IEEE Computer Graphics and Applications, 18(2): 32-43, March 1998
- [GMM00] GULLIOU E., MENEVAUX D., MAISEL E., BOUATOUCH K.: Using vanishing points for camera calibration and coarse 3D reconstruction from a single image. In Visual Computer 16 (2000) 7, 396-410
- [Har04] [www.markmark.net/misc/rendertexture.html](http://www.markmark.net/misc/rendertexture.html)
- [Heu98] VAN DEN HEUVEL F.: 3D reconstruction from a single image using geometric constraints. J. Photogrammetry Remote Sensing 53:354-368
- [LDR00] LOSCOS C., DRETTAKIS G., ROBERT L.: Interactive Virtual Relighting of Real Scenes, IEEE Transactions on Visualization and Computer Graphics, 6(4):289–305, 2000.
- [LFD99] LOSCOS C., FRASSON M.-C., DRETTAKIS G., WALTER B., GRANIER X., POULIN P.: Interactive Virtual Relighting and Remodeling of Real Scenes, Tenth Eurographics Workshop on Rendering. In: Lischinski, D.; Ward-Larson, G. (Hrsg.): Rendering Techniques '99. Springer Verlag, 1999, pp. 329–340

- [NHI86] Nakamae, E., Harada, K., Ishizaki, T., and Nishita, T.: A montage method: The overlaying of the computer generated images onto a background photograph. *Computer Graphics (Proceedings of SIGGRAPH 86)* 20, 4 (August 1986), 207-214
- [MS94] MÜLLER S., SCHÖFFEL F.: Fast radiosity repropagation for interactive virtual environments using a Shadow-Form-Factor-List, *Fifth Eurographics Workshop on Rendering (Darmstadt, Germany, May 1994)* pp.325-342
- [SHS98] SHUM H.Y., HAN M., SZELISKI R.: Interactive reconstruction of 3D models from Panoramic Mosaics. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pages 427-433, Santa Barbara, June 1998
- [SM99] STURM P., MAYBANK S.J.: A Method for Interactive 3D Reconstruction of Piecewise Planar Objects from Single Images, *BMVC England*, pp. 265-274, September 1999
- [SSI99] SATO I., SATO Y., IKEUCHI K.: Acquiring a radiance distribution to superimpose virtual objects onto a real scene. *IEEE Transactions on Visualization and Computer Graphics* 5, 1 (January - March 1999), 1 - 12
- [Sch00] Schmidt-Clausen, H.J.: *Grundlagen der Lichttechnik, Unterlagen zur Vorlesung 2000*, TU Darmstadt
- [WBS01] WILCZKOWIAK M., BOYER E., STURM P.: Camera Calibration and 3D Reconstruction from Single Images Using Parallelepipeds, *ICCV 2001*, pp. 142-148
- [WOF03] WONG T.T., OR S.H., FU C.W.: *Real-Time Relighting of Compressed Panoramas, Graphics Programming Methods (edited by Jeff Lander)*, Charles River Media 2003, pp 375-388
- [YDM99] Yu, Y., Debevec, P., Malik, J., and Hawkins, T.: Inverse global illumination : Recovering reflectance models of real scenes from photographs. In A. Rockford, editor, *Computer Graphics (Proceedings of SIGGRAPH 99)*, volume 19, pages 215-224. Addison Wesley Longman, August 1999