

Low-cost Photometric Calibration for Interactive Relighting

Céline Loscos^{1,2}

George Drettakis¹

¹ University College London
Department of Computer Science
Gower Street, London WC1E 6BT, UK
<http://www.cs.ucl.ac.uk/>

² iMAGIS¹ – GRAVIR/IMAG–INRIA
ZIRST, 655 avenue de l’Europe
38330 Montbonnot Saint Martin, FRANCE
<http://www-imagis.imag.fr/>

Abstract

Computer augmented reality is a rapidly emerging field allowing users to mix virtual and real worlds. Our interest is to allow relighting and remodelling of real scenes, using a reflectance estimation method. Most previous work focused on the quality of the results without considering the expense in computation and the price of acquisition equipment. In this paper, we present a low-cost photometric calibration method which improves the reflectance estimate of real scenes. This is achieved by adapting high-dynamic range image creation to a low-cost camera, and an iterative approach to correct reflectance estimation using a radiosity algorithm for indirect light calculation.

keywords: High-Dynamic Range Images, Inverse Illumination, Relighting, Remodelling, Augmented Reality, Common Illumination.

1 Introduction

Computer augmented reality consists in mixing virtual worlds with real worlds; numerous applications have recently emerged such as video games, training systems or film/video productions. Most existing systems choose to ignore the interaction of light between real and virtual worlds, mainly because realistic lighting simulation is computationally expensive, and incompatible with the real time constraints of these applications. Consistent lighting between virtual and real is however important for some applications, e.g., simulating new lighting conditions in existing real environments for architectural design.

Previous work has been published [SHC⁺96, FGR93, JNP⁺95, Deb98, SSI99] treating consistent lighting (or *common illumination*) between real and virtual objects, without modifying real-world lighting. To be able to modify real-world illumination virtually (a process we will call *virtual relighting*), several preprocessing steps are necessary to acquire knowledge of the real scene. First, several systems, such as 3D laser scanners [MNP⁺99], or image-based modelling [BR97, DTM96, POF98], can be used to build a simple 3D-model of a real scene. Then textures must be extracted

¹iMAGIS is a joint project of CNRS, INRIA, INPG and UJF

from input photographs and mapped onto previously built 3D-polygons. Finally, *photometric* information on real light sources (geometry, emittance) and reflectance need to be retrieved to be able to deal with common illumination.

The main purpose of this paper is to improve the estimated reflectance while using non-specific, low-cost equipment and a simple input process. In previous work, we have developed a complete common illumination system [LFD⁺99], using a simple, low-cost input procedure. Photometric acquisition was performed from photographs, which has proven to be inadequate in certain circumstances. In this paper, we show how to achieve higher quality photometric calibration while maintaining the requirement for low-cost input. The resulting quality of this calibration allows us to improve the reflectance estimation of [LFD⁺99], by introducing an iterative algorithm.

Inverse illumination

Virtual relighting of real-world environments is a very hard problem. Real lighting is already included in the textures extracted directly from real photographs. If we want to virtually switch off real light sources and perform relighting, we need to know the reflectance of the objects of a real scene. This can be seen as having a description of the material properties of a scene independently of the given real-world lighting conditions.

Previous work has resulted in solutions to *inverse illumination*, i.e., to extract reflectance parameters from an already lit environment. Some of these papers deal only with a single object (e.g., [SWI97, Mar98, MWL⁺99]), whereas others (e.g., [YM98, YDMH99, LFD⁺99]) extract parameters for a complete environment. For the case of general complete environments, accurate reflectance parameters cannot be extracted for every object, because the reconstructed geometry is often simplified (e.g., a desk is modelled as a set of boxes, ignoring the details of drawers etc.). Most relighting algorithms have focused on the quality of the results, and are not designed with interactivity in mind. Moreover, significant amounts of data need to be acquired. In previous work of ours [LDR98, LDR00, LFD⁺99] interactive relighting and remodelling are performed with a low-cost and simple input procedure. To achieve this goal, two simplifying assumptions are made: the recovered reflectance is assumed diffuse, and the scene is viewed from a single viewpoint.

2 Paper overview

In this paper, we improve the method presented in [LFD⁺99] by providing photometric calibration of the input data. In combination with an iterative approach, this results in the acquisition of more precise diffuse reflectance. We begin by overviewing the details of the original method [LFD⁺99].

The Interactive Relighting System

In previous work [LFD⁺99], we developed a solution to retrieve reflectance parameters for an indoors environment, using a simple 3D-model of the real scene. We use a semi-automatic digital camera¹ to take photographs of the real scene. We use two sets of photographs, one used to construct a 3D model of the real scene using the Rekon image-based modeller [POF98], and one to recover the reflectance.

The second set consists of photographs taken from the same viewpoint but under different lighting conditions controlled by moving a single, small and portable light source (a garden lamp) in several positions. We call the resulting photos *radiance images*. Examples of radiance images are shown in figure 1 (a), (b), (c) (notice that the shadows of the desk are in different positions for each image). During the geometry reconstruction process, objects visible from the chosen viewpoint for the radiance images are modelled more accurately than objects outside the view frustum.

¹Kodak DC260

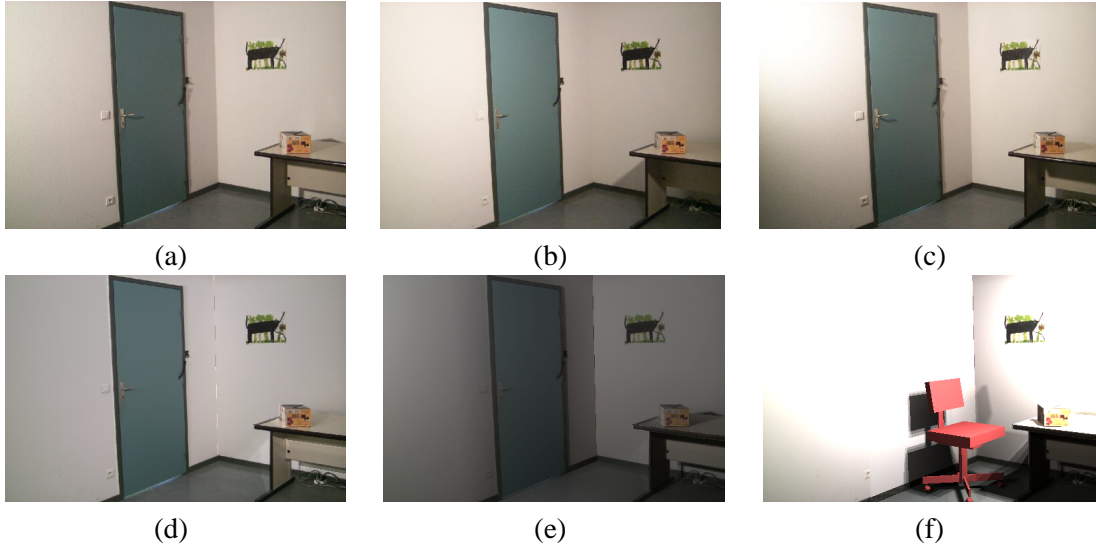


Figure 1: (a), (b), (c) Three of the seven input radiance images used. (d) Recovered reflectance. (e) Simulated lighting using the reflectance in (d). (f) The real door has been removed and a virtual chair and light source have been added.

The acquisition process of radiance images has several advantages. First, light comes from a single source. It is therefore easier to simulate direct lighting (i.e., light arriving directly from a light source), since the geometry of the source and the receiving objects is known. Second, we choose the light positions so that areas in shadows in some pictures are directly illuminated in others. Since direct illumination is easier to simulate than indirect lighting (i.e., light which has been reflected one or more times by non-light source surfaces), it is easier to estimate reflectance in directly illuminated areas. The reflectance estimation is performed pixel by pixel in 3 steps:

- For each pixel i of each radiance image with the associated light source position s , the computation of the diffuse reflectance ρ_{is} is done using the equation:

$$\rho_{is} = \frac{B_{is}}{(F_{is}V_{is}E_s + I_s)}, \quad (1)$$

where B_{is} is the radiosity of pixel i in image s , F_{is} is the form factor, V_{is} is the visibility factor between the 3D point associated to pixel i and the light source s , and E_s is the emittance of source s (the same for every image). I_s represents indirect illumination.

Since we do not have enough information to estimate the indirect light I_s , we initially use an ambient term which is the same for every pixel of a radiance image s , and is computed as the average value of pixel colours divided by an average reflectance given by the user. For B_{is} , we initially use the colour of the pixel in the radiance image. As we shall see, the estimates of both I_s and B_{is} will be improved using the calibration techniques presented in this paper.

- We associate a *confidence* value to each pixel of each radiance image: this value must reflect the confidence we have in the reflectance estimation for each pixel and is in the range $[0..1]$, where 0 is lowest confidence, and 1 highest. Evidently, it must be low for regions not directly lit in the radiance image (since we have a very coarse approximation to indirect light). We then filter the result to compensate for inaccuracies.
- We combine reflectance values from each individual radiance image to get a single reflectance value per pixel: a new reflectance value is computed for each pixel, and is equal to the average reflectance in each radiance image weighted by the confidence values.

This estimation process allows us to use a simple model of the real scene, with a simple geometric model and an approximate lighting model, and to get the resulting reflectance images with few artifacts. The resulted estimated reflectance is shown in figure 1 (d). Although the estimated reflectance is not perfect, it allows convincing interactive relighting². The complete estimation process takes a few minutes, depending on the complexity of the model for the visibility computation and the number of pixels.

The estimated reflectance is used to initialise a lighting system. To simulate the illumination, we combine a per-pixel computation for direct lighting with a hierarchical radiosity solution [Sil95] for indirect illumination. The resulting simulation with the original lighting conditions is shown in figure 1 (e). Compare it to the original photograph in (a): The two images are not exactly the same, but as expected they are similar.

Once we have initialised the lighting system, we are able to interactively relight and remodel the scene, using local updates of the lighting solution. An example is shown in figure 1 (f), where a real door was virtually removed in 3.3 sec., a virtual chair was added in 5.37 sec., and a virtual light was inserted in 7.6 sec (on a SGI R10000 workstation at 195Mhz).

Limitations and suggested improvements of this method

This method however suffered from limitations due to the quality of the input data. If we closely examine the estimated reflectance, we can see that there are discontinuities in the reflectance estimate for the same material (see the wall in figure 2 (a)), and the reflectance is incorrectly estimated in indirect lighting regions (see the region under the desk in figure 2(b)). If we plot the reflectance values extracted for a scanline in the reflectance image (shown in black in figure 2 (c)), we can also see that the reflectance estimate for a same material (e.g., for the white wall) is not as homogeneous as we would expect. This can be seen with the three curves in figure 2 (d) that correspond to the pixel values in RGB for the scanline in black in (c). These limitations are due to several problems involving the way the camera interprets pixels, the input photographs, and the approximation made for indirect lighting in our reflectance estimation algorithm.




Figure 2: Limitations in the reflectance estimate as presented in [LFD⁺99]. See Figures 3 (a), (a') and 4 (d) for larger images.

As stated in the introduction, our goal is to use low-cost, non-specific equipment. As a result, we use a semi-automatic camera to capture the radiance images. The camera chooses the exposure time automatically. As lighting changes in each photograph (because the real lamp is moved to another position), the camera chooses a different exposure time for each image. If the lighting is very different in the same image, there will be saturated areas because the exposure time chosen will not correctly compensate. Moreover, the input photographs are “standard” RGB images. By this we mean that colours are encoded in the integer 0 to 255 range. The camera thus transforms the

²See <http://www-imagis.imag.fr/Publications/loscos/relight.html> for Quicktime movie examples of interactive sessions.

perceived image to reduce the range of colours. Another reason for the inhomogeneous reflectance estimate for the same material is that the camera also transforms the colour. These transformations are described by the camera's response function. Finally, the approximate geometric model, the use of ambient term as the indirect illumination, and geometric factors such as the form factors and the visibility computation between patches, involve estimation errors that are not compensated by the combined and filtered estimations.

It is also often difficult to ensure that we have a directly lit region for every pixel, in at least one radiance image. Regions which are exclusively lit by indirect light often have very low values, which are inaccurately represented. Moreover, since in the original method we use an ambient term to estimate the reflectance in such areas, the result is much less precise than for directly illuminated regions.

We present three improvements for reflectance estimation:

- We create high–dynamic range images from a semi–automatic camera, a process which provides the camera–response function and limits the problem of saturation and inaccurate representation of colour values.
- We correct the pixel values to make them consistent with each other even if the lighting conditions change.
- We compute a better approximation of the indirect illumination, which in turn iteratively improves the reflectance estimate, especially in regions lit exclusively by indirect light.

3 Low–cost high–dynamic range images

Images provided by a digital camera or scanned images are typically represented in a restricted range of 0 to 255, for each colour channel (red, green, blue). Each camera has a response function, that transforms real colours to its internal digital representation. Moreover, if the lighting of the scene has a large luminance range, dark and bright areas of the image will be inaccurately represented. The solution to this is the creation of high–dynamic range images as presented by Debevec and Malik [DM97]. Their approach uses multiple photographs at different exposure times. A first step is the extraction of the response function. To extract more precise values, we need to apply the inverse of this response function to the images acquired directly from the camera. The photographs are then assembled to create radiance images.

The new radiance images use RGB values stored as floating point numbers. The full range from low values (dark colours) to high values (bright colours) is well represented, and saturation problems are eliminated. Using the initial approach the camera has to allow manual control of the exposure time for each image.

Using the EV function of the camera

The camera³ we are using is semi–automatic and costs less than 1,000 Euros, which is currently ten times less than a professional digital camera with complete manual exposure control. As the aperture is controllable, we set it to a very small value to avoid radial falloff problems. However, the exposure time is not directly controllable. To overcome this problem, we use the *EV* parameter, which is a standard feature available on low–cost cameras which implicitly controls the exposure time. When *EV* equals zero, the camera automatically chooses the parameters (aperture and shutter speed) that fit the lighting. When *EV* is negative, the camera chooses faster shutter speeds, and when it is positive, it chooses slower shutter speeds. We adapt the Debevec and Malik algorithm

³Kodak DC260

for semi-automatic cameras by converting each possible EV value to a time exposure according to a reference time t , set for $EV = 0$. We based our assumptions on traditional F-stop ranges (2^{EV}).

EV	-2	-1.5	-1	-0.5	0	0.5	1	1.5	2
exposure time	$t/4$	$t/2\sqrt{2}$	$t/2$	$t/\sqrt{2}$	t	$\sqrt{2}t$	$2t$	$2\sqrt{2}t$	$4t$

The results of this adapted algorithm are quite satisfactory. The camera response function extracted seems reasonable, and we avoid the problems of saturation and low dynamic range of traditional RGB images. Using this algorithm, we get high-dynamic range radiance values free from a camera transfer function, for each light position.

We also take a image of the light itself to get its emittance spectrum (sampled in RGB). While we cannot get an absolute value due to the automatic adjustment of the shutter speed, the relative values give us information on the colour of the source.

4 Computing consistent radiance values

The reference exposure time t chosen by the camera for each radiance image is different for each image. Since we do not know the reference time chosen by the camera, we arbitrarily fix it to create each radiance image. Since the lighting conditions change as the source moves, the resulting radiance values are non necessarily consistent between each radiance image. To calibrate these radiance values, we developed a new algorithm based on the radiosity equation.

The principle is to compute a reference reflectance value for each pixel of a chosen radiance image, say image $s = 1$, and to use it to compute the error in radiance in the other images. This error gives a scaling factor, subsequently applied to the image.

As before, for every image s and for every pixel i , we have a relationship between the radiance L_{is} , the reflectance ρ_{is} , the form factor F_{is} , the visibility V_{is} , the emittance E_s , and an indirect radiosity value B_s :

$$L_{is} = \rho_{is}F_{is}V_{is}E_s + B_s. \quad (2)$$

We first compute reference reflectance values ρ_{i1}^{ref} using this equation for the chosen image $s = 1$. If radiance values were consistent and all parameters were accurate, the reflectance ρ_{is} would be the same for each image s . This however, is not the case, due to the aforementioned acquisition problems, and the approximations in the radiosity computations.

To overcome this problem, we use the difference between an expected radiance and the registered radiance in the image to evaluate the error. We compute an expected radiance value L_{is}^{ref} with the reference reflectance ρ_{i1}^{ref} . Since we only have a very coarse approximation of the indirect illumination, we only consider pixels for which $V_{is} = 1$ for both the chosen image and the current image (i.e., directly lit pixels). For each radiance image, we then solve a least squares system:

$$\varepsilon = \sum_i \|a_s L_{is} - L_{is}^{\text{ref}}\|^2. \quad (3)$$

The computed value a_s is the error of the radiance registered in the image with respect to the expected radiance. We then apply a_s as a scale factor to each pixel of each image s , making the values consistent.

5 An Iterative Method to Improve Reflectance Estimation

We now have calibrated radiance images, which we will use to improve the reflectance estimation in [LFD⁺99]. Initially, we use the same algorithm as before [LFD⁺99] to compute reflectance values, but using the high-dynamic range images instead of the raw camera image data. As can be

seen in figure 3 (b) and (b'), these values are still inaccurate due to the coarse approximation used for indirect illumination. Note in particular the discontinuities in the reflectance on the wall and floor.

To improve the estimation, we compute a radiosity solution with the initial reflectance values. This results in direct and indirect illumination values on patches corresponding to subdivided polygons of the reconstructed scene. Using this information, we extract an indirect illumination value for each pixel at each light position. This is done by first extrapolating radiosity to the vertices of each patch and then interpolating. The point corresponding to the pixel is found, and the interpolated indirect radiosity value computed.

Using these radiosity values, we estimate a new reflectance value by Eq. (1), repeated here:

$$\rho_{is} = \frac{B_{is}}{(F_{is}V_{is}E_s + I_{is})}, \quad (4)$$

Note that the ambient term I_s (equal for all pixels) is now replaced by the indirect light I_{is} computed by the radiosity system, which is different and more accurate for each pixel. As a result the quality of the reflectance estimated is improved, since the approximation using the ambient term introduces significant error. We iterate until the reflectance values are stable. In our experiments, the reflectance values converge after 3 or 4 iterations.

Although the indirect illumination is approximate due to the simple geometric model and the even coarser model for objects not in the view frustum, we will show in the next section that the reflectance estimated using this iterative algorithm is more homogeneous than the non-iterative approach and has fewer artifacts.

6 Results

We took a set of pictures, at different EV values, and converted them to radiance images using the algorithm described in section 3. In the computed radiance images, the effect of the camera's response function is removed and there is no saturation. We then used the algorithm described in section 4, to make the radiance values consistent for different lighting conditions. We use these radiance images to compute a reflectance image, using the initial algorithm [LFD⁺99]. But now B_{is} in the Eq. (1) is a high-dynamic radiance value instead of a colour as before. The resulting reflectance is shown in figure 3 (b). In (b'), the pixel values were plotted for the scanline in black, in the three colour channels red, green and blue. The curves for a reflectance computed with "standard" RGB values are shown in (a'). In this figure, the blue channel is saturated. The red channel curve has a variation of ≈ 0.2 in the reflectance values, for the same material (the white wall). With the newly computed reflectance shown in (b), the reflectance values, shown in (b') for the line in black, are all under 0.6. The variation of the reflectance values are of ≈ 0.15 for a same material, which is less than the variation in (a'). Moreover, the neighbourhood values are more homogeneous in (b') than in (a'). This shows that the reflectance estimate has been improved by the use of radiance images even before the iterative process. After the iterative process, we get the reflectance image shown in (c). For the scanline shown, the values shown in (c') are all under 0.4. For the same material, values are more homogeneous, with a variation of ≈ 0.08 .

We have also tested our calibration algorithm on another example. For this example, the inverse camera response function is applied, but only a single exposure was available for each lighting position. We applied the algorithm presented in section 4 to make the radiance values consistent. With the radiance images, we compute a new reflectance image shown in figure 4 (b). We then compute a new reflectance image after the iterative process shown in (c). Because we do not have different exposure times, the quality of the results is not completely satisfactory. However, if we look

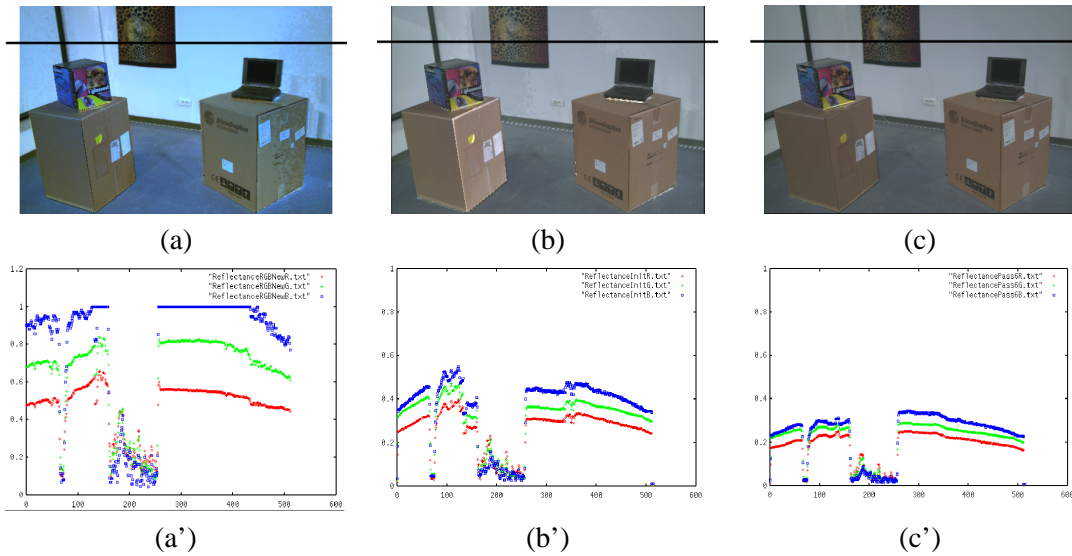


Figure 3: Comparison of the evolution of the computed reflectance. (a) Reflectance computed from RGB images. (b) Reflectance computed from calibrated radiance images. (c) Reflectance computed from calibrated radiance images, after iterations for indirect illumination values. (a'), (b'), (c') RGB reflectance values for a same scanline (shown in black), associated respectively to images (a), (b), (c).

more carefully, we can see that removing the camera response function improves the quality of the reflectance obtained. As can be seen in figure 4 (e) for the initial reflectance with radiance images and in (f) after the iterative process, the discontinuities on the wall in (d) are less pronounced. Another area of interest is under the desk, which is lit exclusively by indirect light in all images. We can see that the reflectance computation improves in (f) after the iteration on the indirect illumination. We believe that the use of multi-exposure input images would further improve these results.

7 Conclusion

In this paper we have presented a new photometric calibration algorithm which improves the reflectance estimate of [LFD⁺99]. This algorithm improves the quality of the computed reflectance, while maintaining the initial goals of a low-cost and simple input process. Using a semi-automatic camera, we capture pictures at different shutter speeds using the *EV* parameter provided by the camera. The algorithm of Debevec and Malik [DM97] is adapted to our data to create radiance images. In these images, the effects of the camera's response function are removed, and the range of the values is larger, avoiding colour saturation and compression. Since the *EV* parameter does not provide the shutter speed, the computed radiance values may not be consistent for different lighting conditions. We have developed an algorithm to correct this inconsistency using the radiance provided in the image pixels and a computed expected radiance value. With the radiance images obtained after these two process, we estimate reflectance values using the previous method [LFD⁺99].

The precision of the estimated reflectance values is then improved using an iterative algorithm. In our approach we simulate indirect lighting with a hierarchical radiosity solution. As shown in the results, we achieve our goal, since we obtain better reflectance estimates, which are more homogeneous for the same material.

In order to improve the visualisation of the relighting, tone mapping [Sch94, LRP97] and the camera's response function could be applied for final rendering. In future work, we would like to

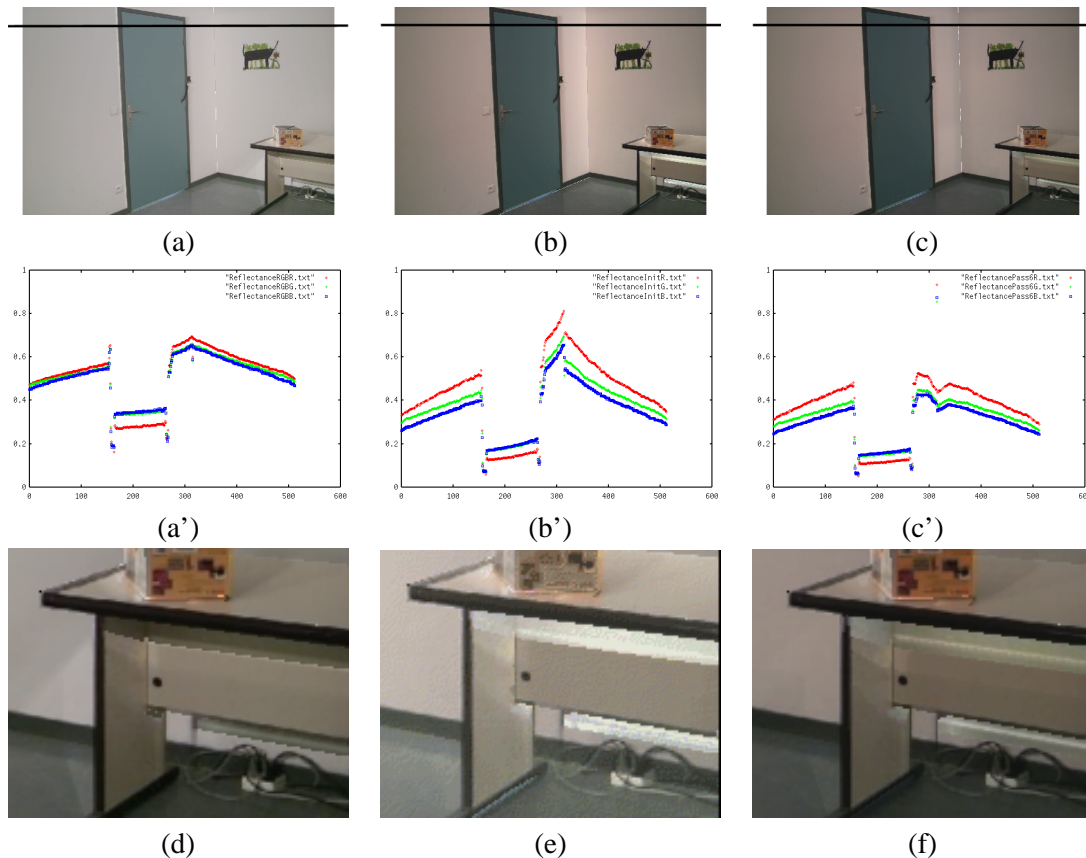


Figure 4: Comparison of the evolution of the computed reflectance. (a) Reflectance computed from low-range RGB images. (b) Reflectance computed from calibrated radiance images. (c) Reflectance computed from calibrated radiance images, with iterations on the indirect illumination values. (a'), (b'), (c') RGB reflectance values for a same scanline (shown in black), associated respectively to images (a), (b), (c). (d), (e), (f) Zoomed version of images shown in figure (a), (b) and (c): A zoom on the wall left to the desk and under the desk. The region under the desk is lit only by indirect illumination in all radiance images. (d) Reflectance values computed from low-range RGB images. (e) Reflectance values computed from calibrated radiance images. (f) Reflectance values computed from calibrated radiance images, with iterations on the indirect illumination values.

further simplify the capture process and remove the major restrictions of the algorithm, which are the diffuse-only assumption and the limitation to a fixed viewpoint.

References

- [BR97] S. Bougnoux and L. Robert. TotalCalib: a fast and reliable system for off-line calibration of images sequences. In *Proceedings of International Conference on Computer Vision and Pattern Recognition*, 1997. The Demo Session.
- [Deb98] P.E. Debevec. Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *SIGGRAPH '98 Conference Proceedings*, Annual Conference Series, pages 189–198, July 1998.
- [DM97] P.E. Debevec and J. Malik. Recovering high dynamic range radiance maps from photographs. In *SIGGRAPH '97 Conference Proceedings*, Annual Conference Series, pages 369–378, August 1997.
- [DTM96] P.E. Debevec, C.J. Taylor, and J. Malik. Modeling and rendering architecture from photographs: A hybrid geometry- and image-based approach. In *SIGGRAPH '96 Conference Proceedings*, Annual Conference Series, pages 11–20, July 1996.

- [FGR93] A. Fournier, A.S. Gunawan, and C. Romanzin. Common illumination between real and computer generated scenes. In *Proc. of Graphics Interface '93*, pages 254–262, May 1993.
- [JNP⁺95] P. Jancène, F. Neyret, X. Provot, J-P. Tarel, J-M. Vézien, C. Meilhac, and A. Verroust. Res: computing the interactions between real and virtual objects in video sequences. In *Second IEEE Workshop on networked Realities*, pages 27–40, Boston, Massachusetts (USA), October 1995. <http://www-rocq.inria.fr/syntim/textes/nr95-eng.html>.
- [LDR98] C. Loscos, G. Drettakis, and L. Robert. Interactive modification of real and virtual lights for augmented reality. In *SIGGRAPH '98 Technical Sketch (Visual Proceedings)*, July 1998.
- [LDR00] C. Loscos, G. Drettakis, and L. Robert. Interactive virtual relighting of real scenes. *IEEE Transactions on Visualization and Computer Graphics*, To appear 2000. <http://www-imagis.imag.fr/Publications/loscos/TVCG00/>.
- [LFD⁺99] C. Loscos, M.-C. Frasson, G. Drettakis, B. Walter, X. Granier, and P. Poulin. Interactive virtual relighting and remodeling of real scenes. In *Rendering Techniques '99 (10th Eurographics Workshop on Rendering)*, pages 329–340. Springer-Verlag, June 1999.
- [LRP97] G. Ward Larson, H. Rushmeier, and C. Piatko. A visibility matching tone reproduction operator for high dynamic range scenes. *IEEE Transactions on Visualization and Computer Graphics*, 3(4):291–306, October–December 1997.
- [Mar98] S. R. Marschner. *Inverse Rendering in Computer Graphics*. PhD thesis, Program of Computer Graphics, Cornell University, Ithaca, NY, 1998.
- [MNP⁺99] D.K. McAllister, L. Nyland, V. Popescu, A. Larsa, and C. McCue. Real-time rendering of real world environments. In *Rendering Techniques '99*, pages 145–160, New York, NY, June 1999. Springer Wien.
- [MWL⁺99] S. R. Marschner, S. H. Westin, E. P. F. Lafortune, K. E. Torrance, and D. P. Greenberg. Image-based BRDF measurement including human skin. In *10th Eurographics Workshop on Rendering*, pages 131–144, June 1999.
- [POF98] P. Poulin, M. Ouimet, and M.-C. Frasson. Interactively modeling with photogrammetry. In *Rendering Techniques '98 (9th Eurographics Workshop on Rendering)*, pages 93–104. Springer-Verlag, June 1998.
- [Sch94] C. Schlick. Quantization techniques for visualization of high dynamic range pictures. In *Fifth Eurographics Workshop on Rendering*, pages 7–18, Darmstadt, Germany, June 1994.
- [SHC⁺96] A. State, G. Hirota, D.T. Chen, B. Garrett, and M. Livingston. Superior augmented reality registration by integrating landmark tracking and magnetic tracking. In Holly Rushmeier, editor, *SIGGRAPH 96 Conference Proceedings*, Annual Conference Series, pages 429–438, August 1996.
- [Sil95] F. X. Sillion. A Unified Hierarchical Algorithm for Global Illumination with Scattering Volumes and Object Clusters. *IEEE Transactions on Visualization and Computer Graphics*, 1(3), September 1995.
- [SSI99] I. Sato, Y. Sato, and K. Ikeuchi. Acquiring a radiance distribution to superimpose virtual objects onto a real scene. *IEEE Transactions on Visualization and Computer Graphics*, 5(1):240–254, January–March 1999. ISSN 1077-2626.
- [SWI97] Y. Sato, M.D. Wheeler, and K. Ikeuchi. Object shape and reflectance modeling from observation. In *SIGGRAPH '97 Conference Proceedings*, Annual Conference Series, pages 379–387, August 1997.
- [YDMH99] Y. Yu, P.E. Debevec, J. Malik, and T. Hawkins. Inverse global illumination: Recovering reflectance models of real scenes from photographs. In *SIGGRAPH '99*, 1999.
- [YM98] Y. Yu and J. Malik. Recovering photometric properties of architectural scenes from photographs. In *SIGGRAPH '98 Conference Proceedings*, Annual Conference Series, pages 207–218, July 1998.