

Image-based Participating Media

Jorge Lopez-Moreno, Angel Cabanes and Diego Gutierrez

Universidad de Zaragoza, Spain

Abstract

Light transport inside participating media, like fog or water, involves complex interaction phenomena, which make traditional 3D rendering approaches challenging and computationally expensive. To circumvent this, we propose an image-based method which adds perceptually plausible participating media effects to a single, clean high dynamic range image. We impose no prior requirements on the input image, and show that the underconstrained nature of the problem (where variables like depth or reflectance properties of the objects are obviously unknown) can be overcome with relatively little unskilled user input, similar to other image-editing techniques. We additionally validate the visual correctness of the results by means of psychophysical tests.

Categories and Subject Descriptors (according to ACM CCS): I.3.4 [Computer Graphics]: Graphics Utilities

1. Introduction

Participating media like fog or smoke have a great influence in the light transport of a scene. They imply a series of complex phenomena which greatly affect how objects are perceived. More precisely, light transport through participating media is affected by the following phenomena [SKSU05] (fig. 1):

- Emission: Radiance is increased by the photons emitted by the participating medium itself.
- Absorption: Radiance decreases when photons are absorbed by the particles composing the participating medium.
- In-scattering: Radiance increases due to photons scattered in the direction of the considered path.
- Out-scattering: Radiance decreases due to photons scattered out of the direction of the considered path.




Figure 1: The four types of interaction of light in participating media (after [PPS97]).

All this makes simulating light transport in participating media a computationally expensive process which requires previous 3D knowledge of the scene. Rather than attempting to provide a physically-based simulation, which would require complete knowledge of the scene's properties, such as dimensions, optical thickness of the medium or reflectance properties of the objects, we aim to simulate its effects in image-space, starting with a single high dynamic range (HDR) image as input. Given the underconstrained nature of the problem, a physically accurate solution is obviously impossible to achieve. However, Ramanarayanan and colleagues [RFWB07] showed that physical inaccuracies in an image come largely undetected in some situations, and therefore a perceptually plausible solution can be achieved that will be perceived as correct by a human observer. In that regard, our work is similar in spirit to the work of Khan et al. [KRFB06], which provides an algorithm for image-based material editing which exploits the limitations of the human visual system. The purpose of this research is to extend the current capabilities of image-editing tools (such as PhotoshopTM), for which we need easy interactivity and short computational times.

The rest of the paper is structured as follows: In section 2 we discuss previous work similar to ours. In section 3, first we analyze the natural process: the physical interpretation and its influence on the perception of scenes. Second, we implement a processing pipeline capable of simu-

lating the presence of participating media in a single HDR image yielding visually realistic results. Finally we validate our results in section 4 by means of two psychophysical tests, and by measuring our rendering times against several artists' renditions using commercial image editing software (like PhotoshopTM).

2. Previous Work

Our research is closely related to the image-based work by Nayar and colleagues [SNN01], [NN01], [NN03b], [NN03a]. In [SNN01] the authors present a method to remove haze based on the partial polarization of airlight (defined as the ambient light scattered towards the viewer). The method requires two images taken with polarization filters, preferably at parallel and perpendicular orientations. Contrast is partially restored in [NN01]; although the method does not require any prior information, it also needs several images as input. This restriction is subsequently lifted in [NN03b], in exchange of some user input. The results in all these works are quite impressive. They contribute with a solution to the problem of removing undesirable effects from images, caused by light transport in participating media. However, it is not clear if the processes could be reversed to *add* those effects to clean input images.

By contrast, in [NN03a] multiple scattering of single point light sources is simulated by means of a point-spread function. Unfortunately with this method, only the light transport originated by the light sources which are visible in the original image can be simulated, and the results are constrained to a limited subset of cases: light sources in almost completely dark scenes (e.g. lamps in a misty night). We overcome these limitations by presenting a method to simulate participating media in images relying in image processing techniques. We show that the underconstrained nature of the problem can be overcome with little unskilled user input. The algorithms use a single HDR image as input, and no previous knowledge of the scene is required.

An important difference with the discussed previous work is our overall goal: that whilst Nayar and co-workers focus in computer vision related problems, our aim is to extend the available repertoire of image editing tools. We believe that the progressive establishment of an HDR imaging pipeline opens up new possibilities for these kind of applications, creating image editing techniques that were not possible before due to the quantization and loss of data in traditional low dynamic range images. An example of this is the work by Khan et al. [KRFB06], which shows how extreme material edits can be performed in perceptual space, leveraging the wealth of information available in HDR format. We thus aim at producing simulated participating media that can be seen as perceptually plausible, a claim we support by means of psychophysical validations. Sundstedt and colleagues [SGA⁺07] already proved the convenience of such an approach for participating media: the authors were able to

drastically cut down rendering times by taking advantage of the limitations of human perception, producing images indistinguishable from ground-truth, Monte Carlo based renderings at a fraction of the time. While the aforementioned authors used a traditional 3D rendering approach (with a complete description of the scene and the medium), we work on a single HDR image instead.

3. Light in participating media

In this section we analyze light-medium interactions and how they influence the visual perception of the image (Sections 3.1 and 3.3). Then, we simplify the physical model, thus circumventing the underconstrained nature of working in image space while still producing plausible results (Section 3.2). Finally we present a processing pipeline capable of simulating the complex interactions inside the participating medium by means of a sequence of simple image filters (Section 3.4).

3.1. Assumptions

Starting with a single HDR image, we follow the approach introduced by Debevec and co-workers for image-based lighting [Deb98] and interpret every pixel as a light source. We limit ourselves to isotropic light sources and homogeneous participating media. Further, we will show how non-homogeneous media can be simulated by simply adding Perlin noise to our algorithm.

3.2. Simplifying the physical model

We describe the physical process in terms of the Radiance Transfer Equation (RTE) [Gla95]. Marching along a ray, we can find the total change in radiance per unit distance t as follows:

$$\begin{aligned} (\vec{w} \cdot \nabla)L(t, \vec{w}) &= \alpha(t)L_e(t, \vec{w}) \\ &+ \sigma(t) \int_{\Omega} p(t, \vec{w}', \vec{w})L_i(t, \vec{w}')d\vec{w}' \\ &- k(t)L(t, \vec{w}) \end{aligned} \quad (1)$$

Where the term $\alpha(t)L_e(t, \vec{w})$ adds energy due to emission, $\sigma(t) \int_{\Omega} p(t, \vec{w}', \vec{w})L_i(t, \vec{w}')d\vec{w}'$ represents in-scattering events and $k(t)L(t, \vec{w})$ subtracts energy due to absorption and out-scattering. $\alpha(t)$, $\sigma(t)$ and $k(t)$ are the emission, in-scattering and extinction coefficients respectively. We can simplify this equation by assuming a homogeneous medium, therefore the three coefficients become constant values for each differential step of t : α , σ and k . We further consider an isotropic L_i term, therefore $L_i(t, \vec{w}')$ can be reduced to $L_i(t)$. Furthermore, in many cases we can dismiss the term representing the emission of light, as most of the participating media, like fog, do not have light-emitting particles.

Given that the phase function is considered to be isotropic and the medium is homogeneous, we get $p(t, \vec{w}', \vec{w}) = 1/4\pi$,

and we can further simplify the in-scattering integral term. Equation 1 now can be written as:

$$(\vec{w} \cdot \nabla)L(t, \vec{w}) = \sigma \frac{L_i(t)}{4\pi} \int d\vec{w}' - kL(t, \vec{w}) \quad (2)$$

By integrating the in-scattering term in the whole sphere Ω we substitute it by a constant value *InScat*. This value will be given by the user as a parameter:

$$(\vec{w} \cdot \nabla)L(t, \vec{w}) = \sigma InScat - kL(t, \vec{w}) \quad (3)$$

Given that we are working in image space, we need to integrate the radiance along each ray, thus obtaining:

$$L(x_s, y_s) = \sigma InScat \Delta t + L_0 e^{-k\Delta t} \quad (4)$$

where (x_s, y_s) represent pixel coordinates after integration in t , L_0 is the original luminance value without participating media (given by each pixel value in the image) and Δt is the estimated per-pixel depth of the scene (given by the user as discussed in Section 3.4). In the following, we show how using this simplified equation, coupled with some unskilled user input, allows us to simulate in-scattering, out-scattering and extinction phenomena in images.

3.3. Perception of the natural process

Visual inspection of images with participating media (see figures 2 and 3) allows us to identify the main telltale cues that reveal the presence of participating media from a perceptual perspective. In the absence of existing literature on this topic, we propose the following, which will work well enough for our purposes:

- **De-saturation**, contrast reduction and loss of apparent volume. When surrounded by participating media, shadows are softened and colors lose intensity, due to multiple scattering phenomena.
- **Attenuation of highlights** both from light sources and in the objects of the scene, also due to multiple scattering phenomena.
- **Airlight**. Added luminosity due to in-scattering effect originated by light sources located inside or outside of the medium (e.g.: the sun).
- **Extinction** of the original pixel luminosities, due to out-scattering and absorption in the medium.
- **Blur and detail loss** due to multiple scattering.

We now show how our method applies these visual cues to simulate the effects of participating media in an image.

3.4. Image processing

From the aforementioned observations of the natural processes we derive an image processing pipeline: First the original image is "relit" as if it were inside the participating media. As we cannot perform an actual relighting of the scene



Figure 2: Some photographs of real participating media, showing combinations of the perceptual cues enumerated in Section 3.3.



Figure 3: Highlights attenuation due to scattering of light sources.

(unknown geometry and reflectance properties), we simulate its effects of the medium through simple filters: desaturation of colors, contrast reduction, increased luminance in shadows, and highlight attenuation. This is achieved by means of histogram manipulation. Extinction and airlight effects are subsequently simulated by following equation (4). Finally, blur and detail loss is simulated by defining a point spread function for the image.

3.4.1. Depth estimation

Before going into detail describing the process, we will discuss one of the main user inputs to the image processing pipeline: the approximated depth information of the scene. Given that our input is a single image, we lack any depth information associated to its pixels but the perception of the objects in a scene with participating media is highly dependent on how far they are with respect to the sensor, and thus this depth information needs to be approximated somehow. For our purposes there is no need for great accuracy: the human visual system is not a perfect light meter and thus great discrepancies from a physically accurate solution go undetected (as our psychophysical tests will show). We propose two different approaches:

1. **Depth simplification:** In terms of composition, almost any image can be decomposed in up to three planes: close-up plane, middle plane (optional) and background. Therefore, an user-made segmentation of the scene is enough to create a discrete depth map, capable of creating visually plausible depth perceptions in our media simulations. However, in certain cases perspective issues make this discrete plane segmentation not good enough (a ground receding into the distance, for instance). In those cases it is necessary to create a smooth depth gradient instead. We generate those gradients by hand simply by dragging the mouse over the region of interest and defining a perspective plane as in the work by Oh and colleagues [OCDD01], as illustrated by Figure 4.
2. **Shape from shading techniques:** In cases where the surface of an object is too complex, Shape From Shading (SFS) techniques [ZTCS99, EP06] can be used. SFS attempts to recover the shape of an object in an image by analyzing shading variations across its projected surface. As it is discussed in [KRFB06], the drawback of these algorithms is that they are greatly constrained to the conditions where the image was taken (presence of textures, self shadows, highlights,...) and usually show poor results for arbitrary images. Other techniques [Kan98], [OCDD01] depend greatly on the quality and amount of user input to infer depth in the scene. To avoid these problems, we again leverage the limitations of human perception to obtain an approximation which suffices for our purposes. In particular, we follow the approach by Khan and co-workers [KRFB06], which is based on a surprisingly simple assumption: the brighter a pixel, the closer it is to the camera. Thus, darker values are interpreted as points far from the observer. This is clearly not true for a vast amount of cases, but it has been proved to be one of the basic assumptions of the human visual system [KvDS96], [LB00]. We use this SFS-based segmentation approach when required by the complexity of the scene. Although the ideal would be to perform this step without any user input, we find that hand-made segmentation of the regions of interest is still the best option: constraining the users to an algorithm default outcome would reduce artistic criteria, thus limiting its flexibility as an image editing tool. Furthermore, the mental framework of the artists usually includes the concept of layers to separate image areas in depth by its visual importance[†].

3.4.2. Image processing pipeline

The image processing pipeline starts with the transformation of the color space from RGB space to HLS space. We subsequently pre-process the image, simulating the light transport from the light sources to the pixels of the image. This is done in two steps:




Figure 4: Binary depth information (left) and detailed depth information using a multiple level z-buffer image (right).

1. Highlight detection by finding the minimum in the derivative of the image histogram. This minimum is usually a reasonable start for a highlight as shown in [KRFB06].
2. Per-pixel attenuation of the luminance channel L for the corresponding pixel containing highlights. Given the high dynamic range of the input image, it is possible to recover the original colors in most of the cases (see Figure 5).



Figure 5: Example of highlight attenuation by histogram analysis and manipulation.

Once the highlights have been processed, we simulate the rest of the light transport in the scene by directly increasing the luminance in shadowed areas (defined as pixels below five f-stops in the histogram) and modifying the hue and saturation values in proportion to the σ and k parameters. We illustrate the process in Algorithm 1:

Data: I_o : Original image, I_f : Final image, Z : z-buffer, σ : attenuation index, $InScat_H$: Airlight hue, S_R : Shadow reduction coefficient (0...1, default=0.5)

for each pixel x_s, y_s **in the image do**

$$I_f(x_s, y_s)_S \Leftarrow [I_o(x_s, y_s)_S \cdot (1 - k)];$$

$$I_f(x_s, y_s)_H \Leftarrow [I_o(x_s, y_s)_H \cdot (1 - \sigma) + InScat_H \cdot \sigma];$$

if $I_o(x_s, y_s)_L < (Max_L - Min_L)/5$ **then**

$$I_f(x_s, y_s)_L \Leftarrow [I_o(x_s, y_s)_L + \frac{(Max_L - Min_L) \cdot S_R}{5}];$$

end

end

Algorithm 1: Preprocessed image "relighting".

[†] As confirmed by interviews with artists using our system.

Attenuation due to extinction and out-scattering is computed again manipulating the luminance channel, following the second term of equation (4), as shown in Algorithm 2:

Data: I_i : Input image, I_f : Final image, Z : depth z-buffer, K : extinction index
for each pixel x_s, y_s **in the image do**
 $I_f(x_s, y_s)_L \leftarrow [1 - I_i(x_s, y_s)_L \cdot e^{-K \cdot Z(x_s, y_s)}];$
end

Algorithm 2: Extinction due to out-scattering.

Next, we add the airlight typical of a participating medium, by increasing the luminance and hue channels for each pixel (see Algorithm 3), as suggested in equation (4). For implementation purposes, algorithms 2 and 3 are actually computed in the same loop.

Data: I_i : Input image, I_f : Final image, Z : z-buffer, σ : attenuation index, $InScat_L$: airlight luminance
for each pixel x_s, y_s **in the image do**
 $I_f(x_s, y_s) \leftarrow [I_i(x_s, y_s) + \sigma \cdot InScat_L \cdot Z(x_s, y_s)];$
end

Algorithm 3: Airlight due to in-scattering.



Figure 6: Images obtained varying the scattering and absorption coefficients (σ, k): (a) = (0.2, 0.2), (b) = (0.44, 0.5) y (c) = (0.75, 0.9.)

We simulate the blur and detail loss due to the multiple scattering of the light by means of an atmospheric point spread function (APSF). The concept of the APSF was first described in [NN03a] to simulate the effects of multiple light scattering without the cost of ray tracing techniques. It can then be seen as an extension of traditional point spread functions, which model the response of any optical system in the presence of a point light source. For our functions, we use the same values as described in [NN03a], shown in Figure 7. By applying a convolution of the luminance channel of the image with the APSF, we simulate the characteristic blur of a participative media.

When multiple levels of depth are present in the scene, we cannot apply the same APSF to all of them, as the kernel size of the function is determined by the very nature of the participating medium and it is fixed to a determined depth. Therefore we need to resize it according to the user's previous segmentation of the image. We simply interpolate the values of each kernel for intermediate distances as shown in Figure 8.

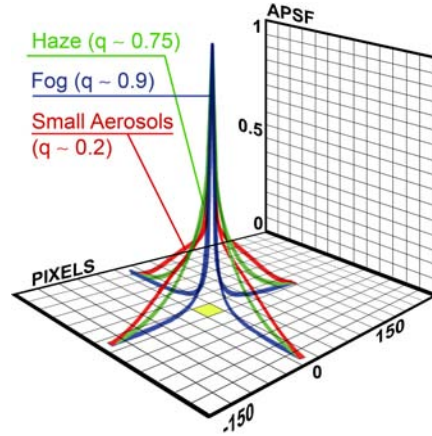


Figure 7: Graphical representations of some APSFs. After Nayar et al. [NN03a]

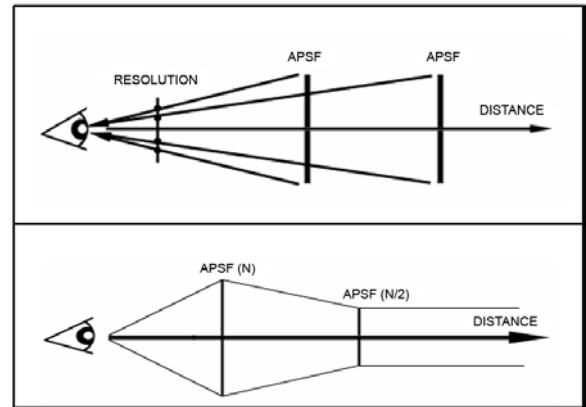


Figure 8: We can see how maintaining the kernel size of the APSF constant requires decreasing the size of the kernel of its projection in the screen(top). We therefore interpolate between decreasing kernel sizes at three different depths.(bottom)

Finally, Perlin noise can be added to achieve the visual appearance characteristic of non-homogeneous media. Perlin noise is a procedural pseudo-random noise which takes two parameters as input for 2D images: Period and number of octaves (amplitude is equal to 1 in our case). An octave is the number of noise functions which, when added together, yield the final noise. Each noise function doubles the frequency of its predecessor (see fig. 9). Perlin noise divides




Figure 9: 2D noise functions with different frequency and amplitude. The last one is the total sum: Perlin noise.

the image into a grid with the size of the cell side equal to *period*. If the pixel corresponds to a vertex of the grid, the value of the noise function is returned. Otherwise, an interpolation of its 4-neighbors is performed. Then for each octave the process is repeated dividing the period by two. The final Perlin noise image is used as a density function for the participating media simulation. To summarize, the user input required by our method is as follows:

- **Approximate Z-buffer:** depth information for each pixel.
- **Scattering and absorption coefficients,** which characterize the medium.
- **InScat value:** which defines airlight intensity.
- **APSF Threshold:** Used to divide the depth of the scene in three regions, as in Figure 4.
- **Perlin noise parameters:** Octaves and period. Used to configure the number of signals to shape the noise function.

4. Validation

Our motivation is to develop a method capable of computing a visually plausible simulation of participating media. As there is no accurate method to measure the degree of realism achieved by our system, we rely on the psychophysical analysis of the perceptions induced by our results. First (Section 4.1), we add participating media to a photograph with our method (we restrict ourselves to fog); then we ask the participants to recreate similar media in other images, using commercial image editing software (we use PhotoshopTM). Second (Section 4.2), we compare the outcome of our system with the artists renditions from a perceptual point of view.

4.1. Adding participating media

6 individuals took part in this part of the experiment. Two of them, from now on called 'average users', had low or medium-low expertise with the tool. The remaining four, henceforth called 'artists' had a great knowledge of both the tool and its potential applications.



Figure 10: (a) Original image used by participants as input. (b) Image with fog computed by our method, used as sample.

The participants were asked to work with a clean image (fig. 10 (a)), and manipulate it so that it looked as if contained the same kind of fog as the image generated by our method (fig. 3.1 (b)). They were instructed that of course the meaning of *same kind of fog* is subjective, and thus artistic expression was not hindered. The participants were given

unlimited time to finish the images. The images generated by the participants are shown in Figure 12. These images, together with the result of our method (using the same parameters for the fog as in the sample image; See fig. 11) will be used as input for the psychophysical test. Now, we analyze Figure 12 qualitatively: we observe that both average users (u1, u2) have not considered how depth affects the opacity and luminance of the fog. Furthermore (u2) has tried to simulate the non homogeneity of the medium adding a disproportionate amount of noise. The artists (a1, a2, a3, a4) have taken into account depth attenuation effects, although some seem to have added too much noise to the fog (a1,a2). (a3) created the closer outcome to our image although extinction effects are not visible, producing a global over-illumination of the image. Finally (u4) overexposes certain areas too much, obtaining a very artificial finish.



Figure 11: Result generated by our model.



Figure 12: Images created by average users (u1,u2) and artists (a1...a4)

In Figure 13 we show the times needed by each participant and our method. The data show that our model generates participating media simulation in less than a quarter of the time needed in average by any user. Furthermore, our model does take into account all the telltale visual cues from participating media, whereas as we have seen, most of the participants failed to capture at least some of them (qualitatively speaking).

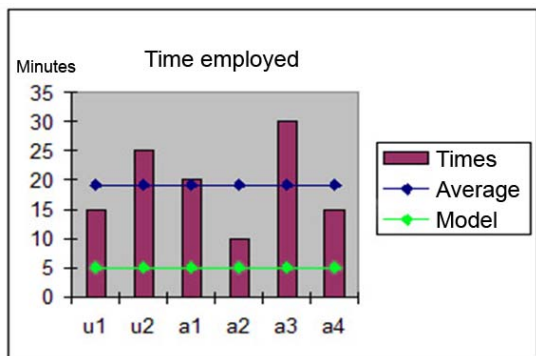


Figure 13: Images created by average users (u1,u2) and artists (a1...a4)

4.2. Psychophysical test

Once we have shown that our algorithm produces faster results, we now want to shed some light on two other important questions: does our simulation look real? And, is it comparable to what an artist’s rendition? The image set used in the test is composed by the five best renditions of the previous step (we discarded the image by (u2)) and the outcome of our model (image 4 ni Figure 14). A gender-balanced total of 20 individuals took part in the experiment, all of them having reported normal or corrected-to-normal vision. About half of them had previous knowledge of computer graphics in general.

The images were shown in a 22" LCD DELL monitor. The test had two parts. First, each individual was exposed to a random sequence of the 6 images (fig. 14) without any possible user interaction. They were simply asked the following question:

"Please indicate if the fog in the image corresponds to a real photograph or if it was digitally processed."

Participants could give only a yes-or-no answer. Time to observe and answer was limited to 20 seconds per image.



Figure 14: Images shown in the test.

Although useful to detect preference trends in the participants this question might introduce a bias. In order to disambiguate and measure this degree of preference we per-

formed a second experiment showing the 6 images at the same time (a 'stimulus sextuple') while asking the participants to rank them (1: less realistic,... 6: most realistic), as suggested in [MMS06]. The display was the same as in the previous experiment, but no time limits were imposed for this task.

The results are shown in Figure 15. (a) shows that the result from our algorithm has the highest ratio of high scores (5 or 6 on a 6-point scale) assigned by the users (52,94%). (b) shows the average scores, where our algorithmic result competes with images 3 and 5 (but has been generated at a fraction of the time as shown before). In particular, our image obtained an average score of 4, only 0.05 points below Image 3 and 0.05 points over Image 5. Finally, (c) shows that 70,59% of the participants perceived our result as a real, photographed scene without digital editing, a percentage not equaled by any of the artists’ renditions.

5. Conclusions and Future Work

We have presented an image based method to simulate plausible participating media in 2D images, using an HDR image as input. The underconstrained nature of the problem is circumvented by means of unskilled user input. We believe the amount of user input is reasonable, given that our results are four times faster than the average artist’s time using a conventional, image editing tool. We show our results on some images and the parameters used in Figure 16 and Table 1.

After analyzing the results of our psychophysical tests, we can state the following:

- Our model is able to simulate participating media with, at least, the same degree of realism and accuracy as an artist using an image editing tool like Photoshop™.
- Observers tend to show preference for our image in detriment to artist paintings.
- Our model generates an image in less than 5 minutes against the 20 minutes needed in average by any user. We have to remark that most of those 5 minutes are processing time without user input. Even unskilled users could match artist renditions in a quarter of the time.

Thus far, in terms of time, the amount of user input needed to create the depth maps is the greatest bottleneck of our method. We believe that our method would be highly improved if the users were provided with computational tools in order to set (with feedback) the depth parameters. Some interesting future work lies ahead: recently, Saxena et al. [SCN08] introduced a new depth acquisition method based in machine learning through the analysis of multiple sets of images and their corresponding actual depths, obtained by laser scan. In this manner, the system is able to infer a likely depth for each pixel based on its previous experience. Although it lacks of great accuracy, the flexibility and robustness of this method makes it a very good candidate to feed depth information to our algorithm, ameliorating the need for user-guided segmentation of the scene.




Figure 15: Graphs representing: (a) Percentage of participants evaluating the image with high scores (5 or 6). (b) Average score assigned to each image (1..6). (c) Percentage of individuals who considered the image as real.

6. Acknowledgments

The authors would like to thank Adolfo Muñoz and the anonymous reviewers for their valuable comments. This research has been funded by the projects UZ2007-TEC06 (University of Zaragoza) and TIN2007-63025 (Spanish Ministry of Science and Technology). Diego Gutiérrez was additionally supported by a mobility grant by the Gobierno de Aragón (Ref: MI019/2007).

References

- [Deb98] Paul E. Debevec. Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *SIGGRAPH*, pages 189–198, 1998.
- [EP06] Olivier Faugeras Emmanuel Prados. *Handbook of Mathematical Models in Computer Vision*, chapter Shape From Shading, pages 375–388. Springer, 2006.
- [Gla95] A. Glassner. *Principles of digital image synthesis*. Morgan Kaufmann, San Francisco, California, 1995.
- [Kan98] S. Kang. Depth painting for image-based rendering applications, 1998.
- [KRFB06] Erum Arif Khan, Erik Reinhard, Roland Fleming, and Heinrich Buehler. Image-based material editing. *ACM Transactions on Graphics, Proceedings of SIGGRAPH 06*, 25(3):654–663, 2006.
- [KvDS96] Jan J. Koenderink, Andrea J. van Doorn, and Marigo Stavridi. Bidirectional reflection distribution function expressed in terms of surface scattering modes. In *ECCV '96: Proceedings of the 4th European Conference on Computer Vision-Volume II*, pages 28–39, London, UK, 1996. Springer-Verlag.
- [LB00] Michael S. Langer and Heinrich H. Bulthoff. Depth discrimination from shading under diffuse lighting. *Perception*, 29:649–660, 2000.
- [MMS06] Rafal Mantiuk, Karol Myszkowski, and Hans-Peter Seidel. A perceptual framework for contrast processing of high dynamic range images. *ACM Trans. Appl. Percept.*, 3(3):286–308, 2006.
- [NN01] Srinivasa G Narasimhan and Shree K Nayar. Removing weather effects from monochrome images. In *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 2, pages 186 – 193, June 2001.
- [NN03a] S.G. Narasimhan and S.K. Nayar. Shedding Light on the Weather. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume I, pages 665–672, Jun 2003.
- [NN03b] Srinivasa G Narasimhan and Shree Nayar. Interactive deweathering of an image using physical models. In *IEEE Workshop on Color and Photometric Methods in Computer Vision, In Conjunction with ICCV*, October 2003.
- [OCDD01] Byong Mok Oh, Max Chen, Julie Dorsey, and Frédo Durand. Image-based modeling and photo editing. In Eugene Fiume, editor, *SIGGRAPH 2001, Computer Graphics Proceedings*, pages 433–442. ACM, 2001.
- [PPS97] Frederic Perez, Xavier Pueyo, and Francois X. Sillion. Global illumination techniques for the simulation of participating media. In Julie Dorsey and Phillip Slusallek, editors, *Rendering Techniques '97 (Proceedings of the 8th Eurographics Workshop on Rendering)*, pages 309–320, NY, 1997. Springer Wien.
- [RFBW07] Ganesh Ramanarayanan, James Ferwerda, Bruce Walter, and Kavita Bala. Visual equivalence: towards a new standard for image fidelity. In *SIGGRAPH '07: ACM SIGGRAPH 2007 papers*, page 76, NY, USA, 2007. ACM.
- [SCN08] Ashutosh Saxena, Sung H. Chung, and Andrew Y. Ng. 3-d depth reconstruction from a single still image. *Int. J. Comput. Vision*, 76(1):53–69, 2008.
- [SGA⁺07] V. Sundstedt, D. Gutierrez, O. Anson, F. Banterle, and A. Chalmers. Perceptual rendering of participating media. *ACM Transactions on Applied Perception*, 4(3):1–22, 2007.
- [SKSU05] László Szirmay-Kalos, Mateu Sbert, and Tamás Umenhoffer. Real-time multiple scattering in participating media with illumination networks. In *Rendering Techniques*, pages 277–282, 2005.
- [SNN01] Yoav Y Schechner, Srinivasa G Narasimhan, and Shree K Nayar. Instant dehazing of images using polarization. In *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 1, pages 325 – 332, June 2001.
- [ZTCS99] Ruo Zhang, Ping-Sing Tsai, James Edwin Cryer, and Mubarak Shah. Shape from shading: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(8):690–706, 1999.



Figure 16: First column: Original images. Second, third and fourth columns: Processed images using the parameters shown in table 1.

Table 1: Parameter data of images shown in figure 16. $InScat_H$ is constant.

Image	Parameters					
	k	σ	$InScat_L$	$InScat_S$	Octaves	Period
House (left)	1.2	0.8	0.6	0.2	5	250
House (center)	0.6	0.35	0.25	0.2	6	450
House (right)	0.7	0.5	0.5	0.3	3	300
Forest (left)	1.0	0.8	0.4	0.35	6	250
Forest (center)	0.5	0.35	0.15	0.35	5	400
Forest (right)	0.9	0.65	0.5	0.35	4	400
Statue (left)	0.5	0.3	0.01	0.025	6	900
Statue (center)	0.65	0.45	0.15	0.025	4	1500
Statue (right)	0.75	0.5	0.025	0.04	5	750