

Perception of Lighting Errors in Image Compositing

James A. Ferwerda¹, Jeremy Selan² and Fabio Pellacini³

¹Munsell Color Science Laboratory, Chester F. Carlson Center for Imaging Science, Rochester Institute of Technology, Rochester, NY / USA, ²Program of Computer Graphics, Cornell University, Ithaca, NY/USA, ³Department of Computer Science, Dartmouth College, Hanover, NH/USA

Abstract

Image compositing is a standard computer graphics technique used to merge independently created visual elements into a single image. However, errors in the source material or the compositing process can destroy the realism of the result. In this paper we describe a series of experiments that investigate the role that lighting errors in the source material have on the realism of composite images. We study two classes of errors: pixel errors - differences in illumination intensity or color temperature; and cue errors - differences in illumination or shadow direction. The results show that the sensitivity to errors varies with the type of error. While observers are reasonably sensitive to pixel errors that change the statistics of the foreground and background image regions, they are less sensitive to cue errors that leave the image statistics unchanged but introduce conflicting information about scene lighting conditions. These trends are modulated by the scene content. These studies represent some first steps towards developing perceptual metrics for error tolerance in image compositing that can be used to improve the fidelity and efficiency of compositing process.

Introduction

Image compositing is a technique in computer graphics that is used to merge independently created visual elements into a single image. The origins of compositing date back to 1857, when Oscar G. Rejlander created a single image by combining different regions of 32 photographs [5]. Digital image compositing has become ubiquitous in filmmaking [1,8], television [22], virtual sets [11], and augmented reality [7].

When used successfully, compositing creates the illusion that the composited elements are part of a single, cohesive scene. However, errors in either the compositing process or the source material can destroy the illusion. For example, a standard compositing technique is to photograph an object against a uniformly colored screen, and then to superimpose this object on a new background. Figure 1 (left) shows that errors in the compositing process, such as a colored silhouette around the object caused by poor segmentation, clearly indicate that the image is not real. Recent advances in compositing technology [22,16] have eliminated the majority of these *process errors*. Unfortunately, even scenes that are “process perfect” may not appear realistic because of lighting errors in the source material. An example of a common *source error* is shown in Figure 1 (right), where the composited object’s lighting properties do not match the target environment’s lighting. Although in the compositing industry there are rules of thumb for creating realistic results, there have been no systematic studies of the tolerance for such errors. An understanding of how these source errors affect the realism of composite images would greatly facilitate the compositing process.

In this paper, we study the visual effects of source errors in



Figure 1. Errors in image compositing. Left: Composite process error (improper segmentation), Right: Image source error (differences in lighting color and direction).

image compositing. In a series of psychophysical experiments we measure visual sensitivity to four common *source lighting errors*: brightness errors, color errors, illumination direction errors and shadow direction errors. We use two scenes that are representative of common compositing situations: a tabletop still life, and a television “talking heads” scene with two subjects sitting behind a desk. Our results show that different lighting errors are not equally detectable, and also that the detectability of errors depends in part on the scene content.

Background

In 1984 Porter and Duff [17] introduced a digital matting algebra that is now widely used in compositing applications, that allowed two images to be merged by employing an extra image “alpha” channel that contains transparency information. This operation allowed objects photographed against a uniform “blue screen” to be segmented and superimposed over selected backgrounds. Brinkmann [5] and Rickitt [20] offer excellent surveys of the history and development of digital image compositing.

Process errors in compositing

Recent advances in blue screen technology [22,16] have solved many of the problems traditionally associated with the compositing process. For example, one early problem was how to correct for the blue (or green) light that would reflect from the background and would produce halos around composited objects. This problem is now solved automatically by most systems (known as spill suppression).

Camera errors have been another source of problems in image compositing. In standard cameras, the finite depth of field causes image focus differences. Film grain or CCD noise is imparted to the images. Motion blur is recorded due to finite exposure times. Perspective distortion is added because of the nature of lenses. Color imbalances may be present in the media. Finally, the CCD or film stock will have a limited dynamic range that will impart some transfer function onto the recorded images. To effectively

composite images created through different processes, one must understand the limitations inherent in each process and compensate for them. Fortunately, for the most part in modern compositing systems, these “process” problems have been solved [5].

Source errors in image compositing

Unfortunately, the focus on the technical “process” of compositing, ignores the role that “source” factors have on the realism of the result. However, manufacturers of compositing systems are beginning to realize that creating better source material is a critical issue. For example, in a technical bulletin [21], Ultimatte makes the following plea to users: “Proper lighting is the key to realism in image compositing. Not only will it be difficult to make the Ultimatte function properly if the lighting is not right, but even a technically perfect composite will look phony with bad lighting”. Books and papers on compositing techniques [21,5] have begun to note rules of thumb to guide users towards better results:

- Composite scenes that are matching exterior shots should be filmed outside.
- Discrepancies in skin tones should be avoided.
- Matching the location of light sources is important for scenes with strongly cast shadows.
- Color filters (gels) should be used to recreate the lighting color for scenes.

While these rules of thumb are widely known, guidelines for applying them to achieve visually realistic results are lacking. Studies that link compositing practice to knowledge of the requirements and limitations of human vision would be of great value.

Source lighting errors

Lighting errors in image compositing can be placed into two categories: *pixel errors* and *cue errors*. Pixel errors include brightness and/or color differences between the composite layers. These are typically produced by differences in the intensity or color temperature of the illumination in the scenes that are being composited. Cue errors on the other hand, are inconsistencies in the spatial properties of illumination in the composited scenes. Sources of cue errors include discrepancies in the direction of illumination (surface shading) and differences in the directions of cast shadows.

With pixel errors, the image statistics of the composited element and the scene context are different, according to a uniform, histogram-modifying rule. In contrast, cue errors preserve the image statistics of the composited region and the scene context, but the visual information about direction and qualities of illumination are contradictory.

In the experiments that follow we study four lighting errors that are considered most likely to affect the realism of image composites. With respect to pixel errors we study the effects of source brightness and color temperature differences. With respect to cue errors we study the effects of differences in shading and shadows caused variation in the spatial properties of lighting. The errors introduced by these manipulations are representative of many real-world compositing situations and therefore the results of

the studies should be directly applicable to practice.

Contributions from vision and graphics

In the vision literature, there are many studies on the perception of objects and scene illumination [3,2,10,12,15], but little of this work is directly applicable to the analysis of lighting errors in compositing. One exception is the work of Ostrovsky et al. [14], which directly addresses the issue of lighting direction inconsistencies in images. They found that observers were relatively insensitive to even large lighting inconsistencies, however, they did not study the errors parametrically, which limits the applicability of the results.

In the computer vision literature there have recently been efforts [26,27,28] to develop algorithms for automatically detecting image forgeries and composites on the basis of errors in shading, shadows, and perspective. These algorithms have focused on taking advantage of photometric and projective inconsistencies and have not focused on human abilities for detecting errors. One exception is the work of Farid and Bravo [29] who have recently started to investigate the perception of shading, shadow, and perspective errors.

In the computer graphics literature, studies by a number of researchers [4,13,19,18,23] have focused on the promise of perceptually-based rendering, in which models of human vision are used to improve the efficiency and fidelity of the image synthesis process. One recent project that is particularly relevant is Ramanarayan et al.’s [25] work on the concept of “visual equivalence” that quantifies how changes in the light field in a scene affects the appearance of objects, and develops a metric that can predict when two different images are equally realistic as representations of a scene.

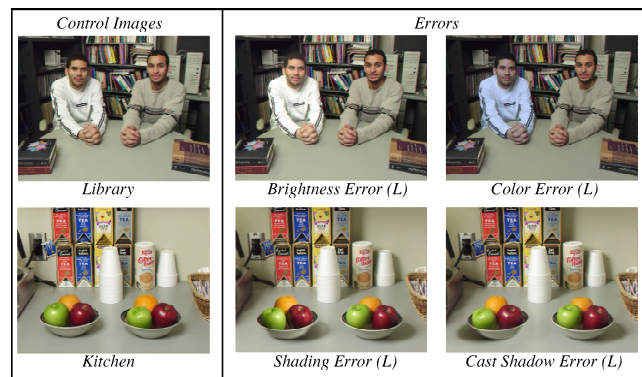


Figure 2. Test scenes and lighting errors. Rows show the library (top) and kitchen (bottom) scenes. Left column shows the reference (control) images. Middle and right columns show representative brightness, color, shading, and shadow errors (reading order)

Experiments

We have conducted a series of psychophysical studies to measure visual sensitivity to lighting errors in image compositing. In each study we ask subjects to judge which of two objects has been composited into an image. Two scenes are tested and the types and magnitudes of lighting errors are varied. The studies yield a set of psychometric curves that describe an average observer’s ability to detect each type of compositing error.

Test scenes

We studied the detectability of lighting errors in two scenes. Representative images of these two scenes are shown in Figure 2. The first scene (kitchen) was a simple tabletop still life consisting of two bowls of fruit against a kitchen backdrop. The second scene (library) approximated a standard television “talking heads” context with two people seated behind a desk.

Each scene consisted of two objects (bowls and people respectively) side by side within the background context. We created these arrangements for three reasons. First we wanted to be able to use a standard two-alternative forced choice (2AFC) procedure to avoid bias and facilitate data analysis. Second we wanted observers to make their judgments on the basis of how well the objects “fit” into the scenes rather than on simple pixel-to-pixel comparisons. Note that the two test objects are completely different pixel-for-pixel though they are similar as objects and therefore have many of the same visual features and fit equally well within the background contexts. Finally, to maximize the detectability of errors, we wanted scenes where the errors were centrally located in the images and were associated with the “subjects” of the scenes.

Stimulus images

In producing the image sets used in the experiments, our goal was to vary each of the error parameters to span its threshold of detectability. Appropriate ranges for each type of error were determined in pre-testing. Seven points along each range were selected for further testing. The specifics of the image sets for each error type are described below.

Pixel errors:

Brightness errors: To simulate modifying light source brightness, seven stimulus images were created by adding or subtracting offsets to the black and white points of a reference image. The increments were set to be -11.7%, -7.8%, -3.9%, 0 (no change), 3.9%, 7.8%, 11.7%. This range was chosen to span the expected threshold values.

Color temperature errors: To simulate changes in light source color temperature, we modified a reference image by shifting pixel chromaticities according to changes in the image’s white point along the black body locus. Seven stimulus images were generated using the following white points: 2869K, 2972K, 3082K, 3200K (defined as reference white), 3328K, 3466K, and 3617K.

Cue errors:

Shading errors: To modify object shading we varied the scene illumination direction. The light source was a Lowel Totalight with a tungsten halogen bulb rated at 3200K and a two foot diameter bounce umbrella. Starting with reference lighting set approximately 30° to the left of the camera viewpoint, we varied the illumination direction in six steps between 0° and 105° relative to the original light axis to generate seven stimulus images (see Figure 3).

Shadow errors: Stimulus images for the shadow errors set were created using the images generated for the shading set. To create images with inconsistent shadows, shadow regions in the reference image were selected and altered using Adobe Photoshop.

The seven shadow error directions created corresponded with the lighting directions used in the shading set.

Images were acquired at 1344x1024 pixels using a Sony DSC-D770 digital camera with lens focal length set to 50mm, aperture 5.6 and ISO 50. All camera settings were fixed for the duration of the capture session. Images were transferred from the camera as uncompressed TIFFs. Adobe Photoshop was used for all histogram adjustments and image modifications.

Eight image sets were created, (2 environments x 4 error dimensions), with each set of images containing both a left error and right error variant. With 7 samples/dimension, and an extra image (with the objects swapped) to serve as a L/R counterbalance the entire stimulus image set consisted of 128 images. Selected examples are shown in Figure 3.



Brightness errors, library scene, left subject, percent changes left to right : -11.7%, 0% control, +11.7%.



Color temperature errors, library scene, left subject, temp. changes (left to right): 2869K, 3200K (control), 3617K.



Illumination direction (shading) errors, kitchen scene, right subject, angular difference in illumination (left to right): 0° (control), 15° , 33° , 61° , 73° , 87° , 105° .



Shadow direction errors, kitchen scene, right subject, angular difference in shadows (left to right): 0° (control), 61° , 105° .

Figure 3. Example image error ranges.

Procedure

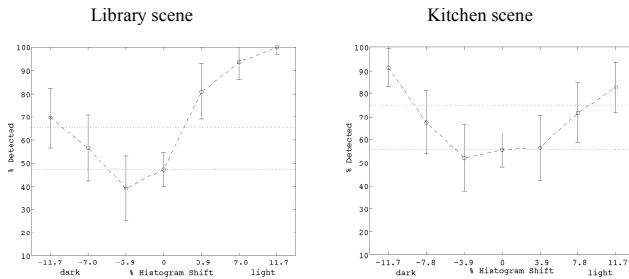
All stimulus images were printed at 300 dpi on glossy paper using a Kodak XLS 8600 dye sublimation printer. Image size was at 14 x 10.7 cm.. Each image subtended a horizontal viewing angle of approximately 16 degrees. Viewing took place under normal office lighting conditions.

The experiments used a two-alternative forced-choice procedure (2AFC), with the two test objects seen next to each other in a common image. For each image the observer was asked to indicate which object appeared to be more realistic with respect to the scene context. Observers sat at desks while taking the tests and viewed the images one at a time without time or spatial constraint.


Image presentation order was randomized across subjects and

across all error dimensions. However, the order was constrained so that observers never saw two images from the same environment in a row. For each error condition, two test images were created, one with the error on the right, and one with the error on the left. Each subject was presented with a random set of error left/error right images, so that the average was 50/50 left vs. right.

Forty-six observers were tested. All were college students, age 19-23. All were naïve to the purposes of the experiment, had no knowledge of compositing or perception psychology, and were generally in non-technical majors. All had normal or corrected to normal vision. Each experimental session lasted approximately 15 minutes.



Figures 4 and 5. Sensitivity to brightness errors. Upper dotted line indicates the correct detection rate at threshold. Lower dotted line is correct detection rate for the control stimulus.



Figures 6 and 7. Sensitivity to color temperature errors. Upper dotted line indicates the correct detection rate at threshold. Lower dotted line is correct detection rate for the control stimulus.

Results and discussion

The following sections summarize the results of our experiments. The data from each of the conditions are summarized in Figures 4 through 13. In each graph the abscissa indicates the magnitude of the particular lighting error, and the ordinate indicates the percentage of trials on which the observers correctly detected the composited object. Detection rates range from 50% (pure chance in a two-alternative forced choice procedure) to 100% (perfect detection).

Logistic regression methods were used to fit psychometric functions to the data. The Chi-Square [30] statistic was used in all tests of significance. The Yates' correction for continuity was applied because particular data categories sometimes had low (<5) numbers of entries.

The detection threshold for each type of error is indicated by the upper dotted line in each graph and was determined by testing for the smallest significant difference in the psychometric function, using the Yates Chi-Square measure ($p < 0.05$). The lower dotted line denotes the detection rate for the null hypothesis (no visible difference).

Brightness errors: Figures 4 and 5 show the results for errors in the brightness of the composited object. In the library scene (Figure 4), observers were able to reliably detect the composited element when the brightness was increased by one step (+3.9%), however the effect was asymmetric, and brightness could be decreased by three steps (-11.7%) before the error was detectable. These effects were significant at the ($p < 0.001$) and ($p < 0.01$) levels respectively. Thresholds calculated from the psychometric function fall between the discrete step levels.

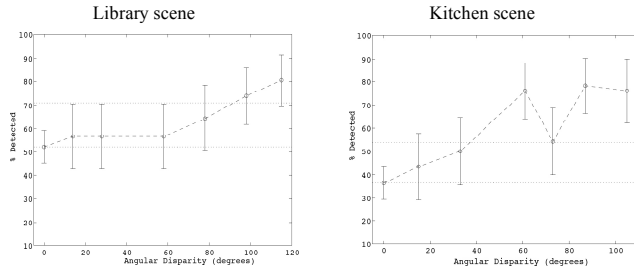
Sensitivity to brightness errors was somewhat lower in the kitchen scene (Figure 5). Here brightness had to be increased by three steps to be detected (+11.7%). Sensitivity to brightness decrements was the same as for the library scene (-11.7%). These effects were both significant at the ($p < 0.001$) levels.

These results indicate that brightness errors are detectable over the range of magnitudes we studied. Further, the asymmetry in sensitivity in the brighter and darker directions found in the library scene suggests that we may be more sensitive to positive brightness errors. However other factors may also be contributing, so further investigation is required before firm conclusions can be drawn.

Color temperature errors: Figures 6 and 7 show the results for errors in the chromaticities of the composited object. In the library scene (Figure 6), observers were able to reliably detect the composited element when the color temperature was shifted by one step toward the blue (temp > 3328K), however the color temperature had to be shifted by three steps toward the red to be detectable. These effects were significant at the ($p < 0.001$) and ($p < 0.01$) levels respectively. Figure 7 shows that no such effects were found in the kitchen scene. Here, over the full range of color shifts tested, subjects were never reliably able to detect the composited object ($p = 0.15$). This occurred despite the fact that the fruits were highly saturated in color.

These results suggest that observers' sensitivity to chromaticity errors in lighting composites, depends in part upon the subject matter in the scene. The bias toward lower detectability of red shifts in the library scene may be due to the fact that shifts in this direction are within the acceptable range of human skin tones while blue shifts are not. It may also indicate that observers are misestimating the subject's true skin colors or the scene illumination (as indicated by the fact that the observers actually mistake the subject with a slight red shift (3082K) as the uncomposited element. The overall lower sensitivity for color shifts in the kitchen scene may also reflect a greater sensitivity for changes in relatively the neutral skin tones over other object colors. That said, it is likely that color shifts of greater magnitude in the kitchen scene would eventually be detectable and that the data in Figure 5 actually shows the central section of a shallowly sloped bi-directional sigmoid psychometric function.

Shading errors: Figures 8 and 9 show the results for errors in the illumination direction of the composited object. In the library scene (Figure 8), observers were able to reliably detect the composited element when the illumination direction error increased by 5 steps ($\theta > 98^\circ$). This effect was significant at the ($p < 0.01$) level. Figure 9 shows that similar effects were found in the kitchen scene. Here, the observers were able to reliably detect the composited element when the illumination direction error increased by 3 steps ($\theta > 61^\circ$). This effect was significant at the ($p < 0.001$) level.



Figures 8 and 9. Sensitivity to illumination direction (shading) errors. Upper dotted line indicates the correct detection rate at threshold. Lower dotted line is correct detection rate for the control stimulus.

These results suggest that observers are relatively insensitive to errors in the direction of illumination in composited objects. Note that even though detection levels reached significance with respect to the psychometric functions, absolute levels of detection only barely exceeded 75%. This finding is in concert with previous studies of illumination perception [14], and is an important result for the field of image compositing since it suggests that careful matching of lighting angles may not be necessary to achieve visually acceptable composites. However further studies of these effects across a range of scenes and subject classes should be done before any firm conclusions are drawn.

It should also be noted that in the kitchen scene (Figure 9) a stimulus bias was found. This can be seen in the fact that the detection rate for the control condition is below 50%. This is surprising since from a visual inspection, the bowls look identical. This may reflect a directional bias in illumination estimation, context effects, or other high level factors. We were able to remove this effect from our statistical tests by comparing the null condition rates for counterbalanced (left, right switched) sets.

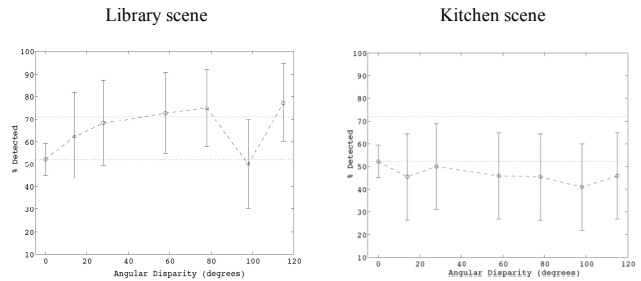
Shadow errors: Figures 10, 11, 12, and 13 show the results for shadow direction errors. In contrast to the other errors studied, significant differences in performance were found when the composited object was on the right vs. the left side so we could not combine the data across conditions. One reason for this difference is because errors on the left cause the real and composited shadows to diverge, while errors on the right cause the shadows to converge.

Figures 10 and 11 show that in the library scene, overall sensitivity to shadow direction errors was low. Regardless of whether the composited object was on the left (diverging shadows, Figure 10) or the right (converging shadows, Figure 11) shadow direction errors were never reliably detected.


We found similar results for the kitchen scene when the errors caused diverging shadows (Figure 12). Once again shadow direction errors were never reliably detected. However the situation was quite different when the errors caused converging shadows (Figure 13). Here, shadow direction errors were reliably detected when the angular disparity exceeded 1 step ($\theta > 15^\circ$). This effect was significant at the ($p < 0.05$) level. However the significance of this result should be weighed against the observations that the data is noisy, variance is high, and overall levels of detectability are relatively low, only reaching 80% in one case.

Overall the results suggest that sensitivity to shadow direction errors in composites is relatively poor. In the library scene

direction error were never detectable. In the kitchen scene (where arguably the shadows were more salient), the differences in performance for the diverging and converging shadow conditions can be explained by the fact that while diverging shadows are quite common in the real world, and are created from singular or clustered light sources, converging shadows are much less likely, and for closely spaced objects can only be created using multiple carefully balanced light sources. The greater sensitivity the observers showed in the converging shadow condition in the kitchen may be because the error is more detectable because the illumination conditions are less likely. Further investigation of this issue is warranted.



Figures 10 and 11. Sensitivity to shadow errors in the library scene. (Left: diverging shadows, Right: converging shadows) Upper dotted line indicates the correct detection rate at threshold. Lower dotted line is correct detection rate for the control stimulus.



Figures 12 and 13. Sensitivity to shadow errors in the kitchen scene. (Left: diverging shadows, Right: converging shadows) Upper dotted line indicates the correct detection rate at threshold. Lower dotted line is correct detection rate for the control stimulus.

Conclusions and Future Work

In this paper we presented the results of a series of psychophysical experiments to measure visual sensitivity to four kinds of lighting errors that occur in image compositing. We found threshold measures for the detectability of the different classes of errors in two representative scenes. The results show that the sensitivity to errors varies both with the type of error and the subject matter of the scene. Also, while observers are reasonably sensitive to discrepancies in the image statistics of the composited and context regions of the images (pixel errors: brightness, color temperature) observers appear to be less sensitive to cue errors (shading/shadow direction) that leave the image statistics unchanged but introduce conflicting information for the lighting conditions in the scene. These studies represent some first steps towards developing perceptual metrics of illumination errors in image compositing that can be used to facilitate the compositing process.

While the results are interesting, it should be emphasized that

this work is preliminary, and that the primary goals of this paper are 1) to raise consciousness about the different categories of lighting errors in image compositing; 2) to show that different perceptual mechanisms with different sensitivities are involved in their processing; and 3) to suggest that different perceptual metrics are needed to quantify the different classes of errors. However much more work needs to be done before generally applicable perceptual metrics can be developed, and caution should be used in applying the specific findings of our experiments in practice.

There are many avenues of further exploration within this topic. First, we have only explored a subset of lighting errors in image compositing. At least two other lighting properties, such as the area of the light source, and the key-to-fill ratio (approximately the ratio of direct to indirect illumination) will likely have similar impacts on the realism of the compositing process. Second, testing over a larger number of scenes and objects should allow stronger conclusions about how scene content affects tolerance for lighting errors in compositing. Finally since compositing is widely used in filmmaking, scene dynamics will likely play an important role in the realism of composites, and understanding when dynamics hide and highlight errors would be very valuable to practitioners. These findings would add to our basic understanding of the limits and capabilities of visual perception and contribute the high level goal of developing sound perceptual metrics to improve the fidelity and efficiency of the compositing process.

References

- [1] J. Barney, J.K. Redd, Stewart Little: A Tale of Fur, Costumes, and Integration. Siggraph 2000, Course 14 notes, (1989).
- [2] J. Beck, Stimulus correlates of judged illumination of a surface. *Journal of Experimental Psychology*, 58:267-274, (1959).
- [3] A. Blake, and H.H. Bulthoff, Does the brain know the physics of specular reflection? *Nature*, 343:165-168, (1990).
- [4] P. Debevec. A Perceptually Based Adaptive Sampling Algorithm. *ACM Trans. on Computer Graphics (Proc. SIGGRAPH '98)*, (1998).
- [5] R. Brinkman. *The Art and Science of Digital Compositing*. Morgan Kaufmann, (1999).
- [6] V. Bruce and A. Young, *In the eye of the beholder: The science of face perception*. Oxford University Press, (1998).
- [7] P. Debevec. Rendering Synthetic Objects Into Real Scenes. *ACM Computer Graphics (Proc. of SIGGRAPH '98)*, 1998.
- [8] J. Egstad. Nuke Case Study- Titanic *ACM Computer Graphics (Course Notes. of SIGGRAPH '98)*, 40:51-56, (1998).
- [9] J. Foley, A. van Dam, S. Feiner, J. Hugues, and R. Phillips. *Introduction to Computer Graphics*. Addison Wesley, (1993).
- [10] A.L. Gilchrist, Perceived lightness depends on perceived spatial arrangement. *Science*, 195:185-187, 1977.
- [11] M. Hayashi. Image Compositing Based on Virtual Cameras. *IEEE Multimedia.*, 5:36-48, March (1998).
- [12] P. Mamassian, D. Knill and D. Kersten. Perception of Cast Shadows. *Trends In Cognitive Science*, 2, August (1998).
- [13] K. Myszkowski, T. Tawara, H. Akamine, and, H. Seidel, Perception-Guided Global Illumination Solution for Animation Rendering. *ACM Trans. on Computer Graphics (Proc. SIGGRAPH '01)*, (2001).
- [14] Y. Ostrovsky, P. Cavanagh and P. Sinha, Perceiving illumination inconsistencies in scenes, *Perception* 34, 1301-1314 (2005).
- [15] A.P. Pentland, Finding the direction of illumination. *Journal of the Optical Society of America*, 72:448-455, (1982).
- [16] Primatte S-100 Manual/Tutorial. www.photron.com. 1996.
- [17] T. Porter and T. Duff. Compositing digital images. *ACM Computer Graphics (Proc. of SIGGRAPH '84)*, 18:253-259, 1984.
- [18] P. Rademacher, J. Lengyel, E. Cutrell and T. Whitted, T. Measuring the perception of Visual Realism in Images. *Proc. Eurographics Rendering Workshop*, 235-248, (2001).
- [19] M. Ramasubramanian, S. Pattanaik and D.P. Greenberg, A Perceptually Based Physical Error Metric for Realistic Image Synthesis. *ACM Trans on Graphics (Proc. SIGGRAPH '99)*, (1999).
- [20] R. Rickett. *Special Effects, the history and techniques*. Billboard Books, (2000).
- [21] Ultimatte Technical Bulletin:#1 Matching Foreground and Background in a Composite, www.ultimatte.com. (2002).
- [22] Ultimatte Overview Manual www.ultimatte.com. (2002).
- [23] B. Watson, A. Friedman and A.A. McGaffey, Measuring and Predicting Visual Fidelity. *ACM Trans. on Graphics (Proc. SIGGRAPH '01)*, (2001).
- [24] W. Zongker, B. Curless and D. Salesin. Environment Matting and Compositing. *ACM Trans. on Graphics (Proc. SIGGRAPH '99)*, 205-214, (1999).
- [25] G. Ramanarayanan, J.A. Ferwerda, B.J. Walter and K. Bala, Visual Equivalence: towards a new standard for image fidelity. *ACM Transactions on Graphics*, 26(3), (SIGGRAPH '07), 1-11. (2007).
- [26] W. Zhang, X. Cao, J. Zhang, J. Zhu, and P. Wang, Detecting photographic composites using shadows, *Proc. IEEE International Conference on Multimedia*, 1042-1045 (2009).
- [27] M. Johnson, and H. Farid, Exposing digital forgeries by detecting inconsistencies in lighting, *ACM Multimedia and Security Workshop*, (2005).
- [28] M. Johnson, and H. Farid, Exposing digital forgeries in complex lighting environments, *IEEE Transactions on Information Forensics and Security* 3(2), 450-461 (2007).
- [29] H. Farid and M.J. Bravo, Image forensic analyses that elude the human visual system, *SPIE Symposium on Electronic Imaging*, (2010).
- [30] Engeldrum, P.G. *Psychometric scaling: A Toolkit for Imaging Systems Development*. Imcotek Press, (2000).