

The Importance of Shadows in Augmented Reality

Abstract

In Augmented Reality (AR) virtual objects are visually added to real scenery, i.e., the real scene is *augmented* with virtual objects. In some AR applications it is a goal in itself that the virtual objects are indistinguishable from real objects. A major difficulty in achieving this realism is to ensure that the virtual objects are rendered using lighting conditions which are consistent with the lighting conditions in the real scene. If this difficulty can be overcome the virtual objects can be shaded, and can cast shadows, just as if they had actually been part of the real scene. Our work focuses on developing methods for determining such lighting conditions and results show that adding the correct shading and shadowing to virtual objects substantially enhances the visual quality of the augmented scene, and makes it much easier for people to perceive the spatial relations between real and virtual objects.

1 Introduction

The concept of Augmented Reality (AR), where real scenes are augmented with visualizations of virtual objects, is well-known, and over the past decade or so substantial progress has been made in terms of hardware supporting the concept, and methodologies enabling credible augmentations, [2, 3].

Our work is aimed at *realistic* Augmented Reality, i.e., the virtual objects should appear as if they were actually part of the real scene. This is opposed to some AR applications, where visual realism of the virtual objects is not of particular importance, for example when augmenting with hidden structures such as wires and cables inside walls or other opaque elements.

One popular example of the realistic AR concept is special effects in film where visualizations of e.g., extinct animals are augmented into real scenes using advanced

computer graphics techniques. Film is not AR per se, as definitions require AR to be interactive, but without doubt big Hollywood productions have demonstrated that virtual objects can indeed be realistically augmented into real scenes. One key element to achieving this realism is that the virtual objects are lit and cast shadows just as if they were actually part of the real scene.

Two things make it difficult to transfer methodologies from movie production to realistic, interactive AR. First of all lighting and shadow casting from virtual objects is manually tuned and optimized when rendering a scene for a film. This is a very complex and time consuming task requiring great skill and experience. Secondly, the computer graphics techniques used for rendering for film are not suitable for real-time performance.

In our work we have developed techniques which enable us to achieve quite convincing results in real-time with a method which automatically estimates the dominant characteristics of the lighting conditions in the real scene, and uses these characteristics when rendering the virtual objects. Figure 1 demonstrates the performance of our techniques.

Our work is primarily technical in nature, aiming at developing methods for estimating lighting conditions in real scenes for use in realistic AR. But in the subsequent section we will give a brief introduction to work on the role of shadows in visual perception of scenes.

2 Shadows and perception

In realistic AR the goal is to make virtual objects indistinguishable from real objects, and it is extremely relevant to discuss AR in the context of *(tele-)presence*, i.e., the sense of presence experienced by a person sensing a (partly) technologically created/mediated experience.

In AR a person is sensing the real world "as usual". I.e., the person is seeing the real world with her own eyes,




Figure 1: Top: a real indoor scene with a few objects on a table top and two lightsources (lamps) outside the field of view. Bottom row, left: some virtual objects have been added to the scene taking occlusions into account. Bottom row, right: same scene, but with shadows cast by virtual objects added. The lighting of the virtual objects, including the virtual shadows, are geometrically and spectrally consistent with the lighting in the real scene, leading to greatly enhanced realism. The bottom row figures are screen shots from a program running at 30 frames per second on a standard PC with a GeForce 2 graphics acceleration card, and the virtual objects can move around the scene. Notice how the addition of cast shadows makes perceiving the actual position of the left sphere much easier; in fact adding the shadows seems to pull the sphere down towards the table top, but in reality the sphere does not move.




Figure 2: Virtual ball-in-box scenario used in the experiments performed by Kersten et al., [5], leading to the conclusion that shadows are the dominant cue with respect to determining the spatial relationships between objects. In the middle and the right images the ball is in exactly the same place, but the different positions of the cast shadow cause completely different interpretations regarding the position of the ball relative to the surface.

either directly with see-through Head Mounted Displays (HMDs), or by seeing a video image of the real world on a (Head Mounted) display. Therefore, in the context of AR applications there is already a real world, in which the person experiences some level of presence, and if the level of presence is not high, it is in principle not due to a technical imperfection (although wearing HMDs can be distracting and cumbersome, and the relatively small Field-of-View of HMDs is also a factor).

But as a concept, AR starts with an existing reality, which the user/person senses as a "first order" mediated experience, ([1]). AR augments virtual objects onto/into this reality, creating a revised reality, which is visually different from the original one. Typically AR *adds* virtual objects to the scene, but in principle the AR concept can just as well be used to remove objects from the scene.

Our goal is to add virtual objects, but to do it so realistically that they do not detract from the perceived realism. We do not dare to state that visual realism equals experience of presence, but we do dare to state that if the augmented objects do not appear realistic, the credibility of the whole scene breaks down. Slater et al., *slater95*, state "... presence is a function of immersion, and immersion requires vividness.". Here Slater et al. are equating frame rate with vividness, and similarly we firmly believe that visual realism supports immersion, which again supports the feeling of presence. (That shadows are important for visual realism should be an obvious fact, but it is substantiated by recent work on establishing quantitative measures for visual realism, [7]).

We have found some very interesting studies on the role of lighting issues in perception, more specifically the role

of shadows in perception. If accepted, our presentation will discuss our work in the context of some of that work. Below we briefly mention some of the more prominent results which have a clear bearing on our work with making AR more realistic by ensuring the virtual objects are subjected to lighting conditions which are consistent with the real scene lighting.

Lighting, shading and shadowing are three interrelated subjects. Lighting conditions cause a particular shading of objects, and the role of shading in biological perception of object shape has been acknowledged for many years, [4]. With a slight simplification of matters it can be said that the role of shading is in shape perception of individual objects, whereas the role of shadows is a very strong cue for the relative disposition of objects in a scene, [6]. I.e., shadows primarily assist in the perception of the spatial relationships between objects in a scene.

In fact studies strongly indicate that shadows are the dominant cue in this task. I.e., shadow information dominates over for example stereoscopic disparities, if the two cues are contradictory concerning the spatial relationships, [5]. This is particularly true for dynamically moving objects and shadows. Their studies also reveal that shadow information dominates over other strong cues for depth information, such as change in projected size. This famous study was performed using the the completely virtual "ball-in-box" scenario shown in figure 2.

A study has once been performed, which directly measured the level of presence induced by adding shadow information to completely virtual environments, [8]. At the time of the study hardware and software did not allow fast generation of virtual shadows, so the frame rate suffered

severely, but the study did show that shadows led to an increase in the level of experienced presence for the visually dominated test persons.

3 Our work so far

The core of our work so far is quite technical in nature and it is beyond the scope of this abstract to give any detail on this. The gist of the approach will though be described in the presentation if accepted. Essentially the developed method allows us to determine the lighting conditions in the real scene and recreate these conditions in the software performing the visualization of virtual objects, including rendering the shadows cast by virtual objects. The result is that the virtual objects appear more realistically embedded into the scene, as figure 1 showed.

4 Future work

With our current level of performance from the developed technique we believe we have an extraordinarily good starting point for further experimentation. Slater's study on whether shadows led to increase in presence level, [8], was hampered by poor performance of algorithms, and was done in completely virtual environments. With our techniques we can perform similar experiments but in a much more realistic setting.

Similarly, the study by Kersten et al., [5], could be repeated in a more realistic scenario than the "ball in box" scenario shown in figure 2.

Lastly, we expect to benefit from some of more detailed findings (not mentioned here) in Kersten's studies, allowing us to concentrate computational efforts on aspects of shadows which are more important. Because for very complex scenes it is still a computational problem to create realistic shadows in real time.

5 Conclusion

We have developed a method to estimate real scene lighting conditions and to transfer them to visualization of virtual objects. This enables realistic real-time augmentation of real scenes with animated virtual objects. We have

found several studies on the perceptual aspects of shadows and within the scope of our research projects we see a great potential for exploring how our work can be used in more experimentation on perception of shadows. Furthermore we feel that our future work can benefit from a better understanding of how shadows are perceived.

References

- [1] What is presence? www.temple.edu/ispr/explicat.htm, April 2000.
- [2] R. T. Azuma. A survey of augmented reality. *Presence: Teleoperators and Virtual Environments*, 6(4):355 – 385, August 1997.
- [3] R. T. Azuma, Y. Baillet, R. Behringer, S. Feiner, S. Julier, and B. MacIntyre. Recent advances in augmented reality. *IEEE Transactions on Computer Graphics and Applications*, 21(6):34 – 47, Nov/Dec 2001.
- [4] V. Bruce and P. R. Green. *Visual Perception - Perception, Psychology and Ecology*. Lawrence Erlbaum Associates Ltd., 2nd edition, 1990.
- [5] D. Kersten, P. Mamassian, and D. C. Knill. Moving cast shadows induce apparent motion in depth. *Perception*, 26:171 – 192, 1997.
- [6] P. Mamassian, D. C. Knill, and D. Kersten. The perception of cast shadows. *Trends in Cognitive Sciences*, 2(8):288 – 295, August 1998.
- [7] P. Rademacher, J. Lengyel, E. Cutrell, and T. Whitted. Measuring the perception of visual realism in images. In S. J. Gortler and K. Myszkowski, editors, *Proceedings: EUROGRAPHICS Workshop on Rendering Techniques, London, United Kingdom*, pages 235 – 248. Springer, June 2001.
- [8] M. Slater, M. Usoh, and Y. Chrysantou. The influence of dynamic shadows on presence in immersive virtual environments. In M. Goebel, editor, *Proceedings: Virtual Environments*. Springer, July 1995.