Visual Glue

William B. Thompson *Peter Shirley* *Brian Smits*
Computer Science Department
University of Utah

Daniel J. Kersten *Cindee Madison*
Psychology Department
University of Minnesota

University of Utah Technical Report UUCS-98-007
Also available (with better pictures) in http://www.cs.utah.edu/projects/sim/glue/

March 12, 1998

Abstract

One key function of graphics systems is to present information about the 3-D structure of modeled environments. For real-time simulations, conveying a sense of contact between touching surfaces and relative position and motion between proximate objects is particularly critical. Neither stereo nor occlusion cues are completely effective for such fine judgments. Conventional wisdom often argues that shadows play a critical role. Less often, it is argued that interreflection also contributes to the sense that two surfaces are touching. This paper explores the actual utility of shadows and interreflection in signaling contact and suggests how this information can be exploited in real-time rendering systems to glue objects to surfaces.

1 Introduction

Every opaque body is surrounded and its whole surface enveloped in shadow and light.

– Leonardo da Vinci [1]

Leonardo wrote at length about the relationships between light sources, objects, and shadows. He also described the phenomenon we now call interreflection. Subsequent work on the part of many researchers has yielded an effectively complete understanding of the physics of light transport as it relates to shadows and indirect illumination. More recently, methods have been developed for computationally simulating light transport with sufficient fidelity to produce physically correct images of synthetic scenes [2].

While the physics of shadows and interreflections are now well understood, we know much less about their perceptual effects. One reason for this is the difficulty of constructing controlled experiments involving shadows and indirect illumination. Gilchrist, in discussing how to analyze how the vision system might decouple illumination, albedo, and luminance, muses about "some sort of magical filter that could filter out all light that had been reflected off a surface . . . more than once" [3]. The graphics community now has the tools to provide exactly these sorts of manipulations.

These tools make it feasible to learn directly about the effects of shadows and interreflections on perception.

There has been extensive work on geometric cues for spatial perception (e.g., [4, 5]), but much less work has been done on illumination cues. Research on the perception of shadows deals almost exclusively with how detached shadows act as a 3–D position cue, locating the shadow-generating object within a larger three-dimensional environment [6–9], or with shadows as a cue for object shape [10–12]. Within the perception community, the little work that has been done on indirect lighting has dealt with the perception of albedo [3, 13]. Computational analyses of the information about surface shape conveyed by interreflections have been done [14–17], but these results do not directly address the determination of spatial organization.

In this paper, we address a specific aspect of visually determining spatial structure: how do shadows and interreflection provide a sense of contact between touching surfaces. Rendered images involving objects in contact with an extended surface often have a "cookie cutter" appearance in which the object looks as if it is one image pasted onto another. This effect is common in real-time applications and can even occur in some realistic rendering algorithms such as radiosity. Current solutions to this problem rely on explicitly adding fairly accurate shadows, either through real-time shadow techniques [18] or discontinuity meshing for radiosity [19]. These approaches tend to be computationally expensive.

We show that interreflections as well as shadows can be used to significantly reduce the problem of floating objects by gluing them to the surface they lie on. We also show that a wide range of manipulation of the cues still results in effective perception of spatial organization. We use this knowledge to make a real-time program that includes both shadows and interreflection on a low-end workstation. This program serves as a proof-of-concept that coarse approximations of complex illumination effects are sufficient for establishing contact and conveying spatial relationships.

2 Surfaces in Contact

In this section we show that shadows and interreflection provide powerful perceptual cues for physical contact between objects and surfaces. As a result, effective rendering of shadowing and indirect lighting can provide benefits even in applications such as VR, where realism in and of itself is not the primary goal.

Shadows come in two distinct types. *Self shadows*, sometimes called *intrinsic shadows*, occur on those parts of the surface of a shadowing object that face away from an illumination source. *Cast shadows*, also called *extrinsic shadows*, occur on surfaces oriented towards an illumination source but occluded from that source by a shadow-generating object. Cast shadows can be *attached* or *detached*, depending on whether or not they are touching the self shadow of the generating object. Cast shadows can signal that the generating object is in contact with the surface on which the shadow is cast. Extended illumination sources cause shadows with penumbra (called *soft shadows* in the computer graphics literature). The "softness" will be evident for any portion of the boundary of a cast shadow not in contact with the generating object. The photometric details of these shadows have been studied extensively in the case of diffuse luminaires illuminating diffuse reflectors [19].

Corresponding terminology has not developed for interreflection effects. For interreflections involving light diffusely reflecting from an object surface onto another, extended surface, contact is signaled by the lightening associated with the interreflection being adjacent to the directly illuminated surface. The boundaries of interreflection patterns associated with diffuse surfaces are always "soft," whatever the nature of the illumination. Unlike shadows, however, the shape of the interreflection pattern is not a function of the direction of illumination.

For objects in contact with an extended surface, diffuse illumination produces a luminance pattern on the extended surface that is a hybrid of that associated with shadows and with interreflections. Near the contacting object, light from some directions will be blocked. This will produce a darkening of the nearby surface, similar to a conventional shadow due to a compact illumination source, but the shape of the pattern will be similar to the shape of the brightening due to interreflection. The effect is most apparent for dark objects, where it is not masked by secondary illumination. We will refer to this effect as *diffuse shadows*[1]. We note that neither the graphics nor the psychology literature has terminology related to diffuse shadows, which emphasizes the lack of attention they have received from either research community.

Illumination generates powerful perceptual cues indicating that two objects are in physical contact. Figure 1a shows a block sitting on a flat surface, rendered with a mix of direct and diffuse illumination but without any consideration of shadows or interreflections. Figure 1b shows the same configuration, rendered with the addition of a soft shadow and interreflections. There is now a strong sense of contact not apparent in the previous figure. Figure 1c shows a rendering with a shadow but no interreflections. Note that the sense of the front of the object being in contact is diminished compared to Figure 1b. Figure 1d shows the same configuration, this time rendered with interreflections but not a shadow. Note that the interreflections are not at all prominent, but they have a large influence on the perception of contact when compared to Figure 1a. Also note that in Figure 1d that the interreflections "glues" the front of the object just as well as the shadow glues the side of the object in Figure 1c. Figures 2a-2d show the same effects for an object positioned just above an extended surface. Figures 1b, 1d and 2b, 2d are effective in portraying whether the object is in contact. This contradicts the idea that indirect lighting is an esoteric effect largely of interest only when photorealistic realism is required. In fact, indirect lighting clearly complements shadowing in establishing a sense of contact. When used on its own as a contact cue, indirect lighting is of comparable effectiveness to shadows, despite being less visually prominent.

Strong shadow cues disappear under diffuse illumination. Figure 3a shows a block illuminated by lighting that is approximately uniform, without any account being taken of the interactions between the lighting, object, and surface. Figure 3b shows the same arrangement of light source and scene geometry, this time rendered with accurate light transport. While the effect is subtle, it clearly signals contact between object and surface. The information provided about spatial proximity is even more apparent in Figure 3c, in which the block is lifted slightly off of the surface. Comparing Figures 3a and 3b emphasize that a lighting effect whose presence may not be obvious in isolation can still be an important visual cue for spatial organization.

[1]The convention in both vision and graphics is to consider lighting to be broken into effects such as shadows and indirect light. In reality there is a continuum of effects that depends on the configuration of incoming light over the hemisphere (field-radiance), and the standard effects are merely extremes in this continuum. However, it is convenient to have a specific term (diffuse shadows) for the object interactions that occur for nearly uniform field-radiance.

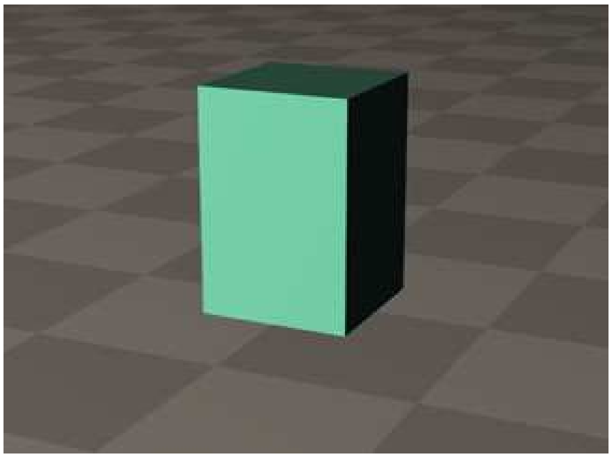

a. No shadows or interreflections.

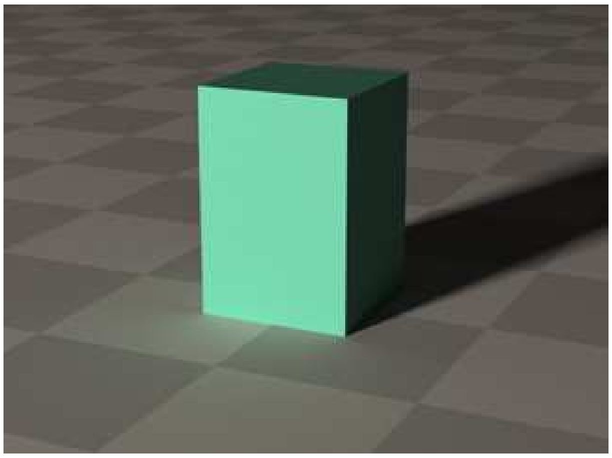

b. Shadows and interreflections added to **a**.

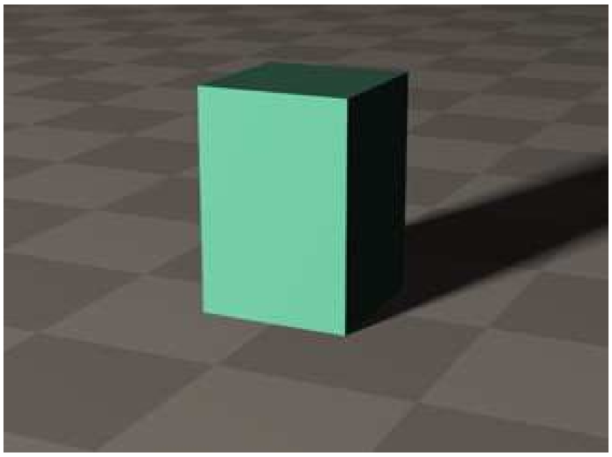

c. Shadows added to a.

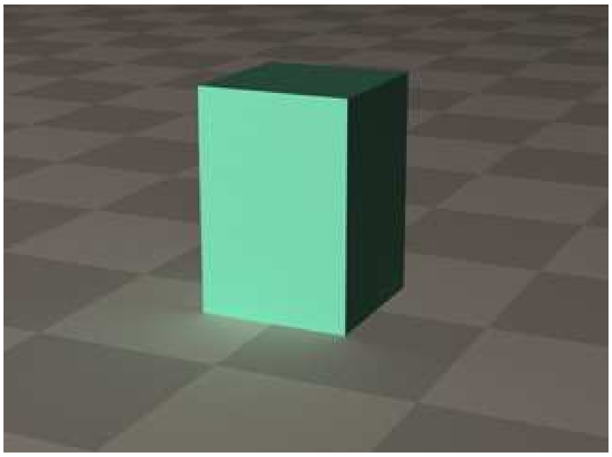

d. Only interreflections added to textbfa.

Figure 1: Objects in contact with extended surface.

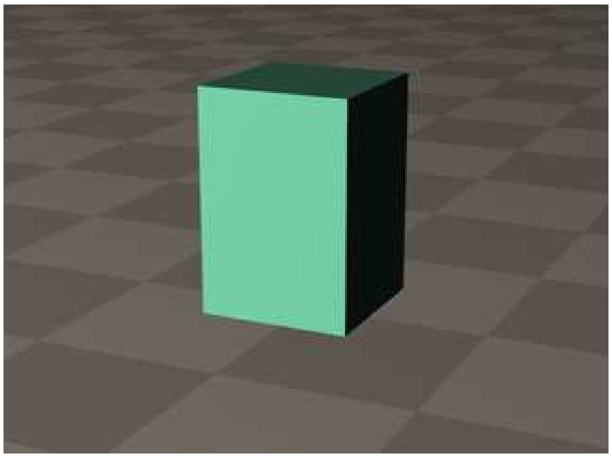

a. No shadows or interreflections.

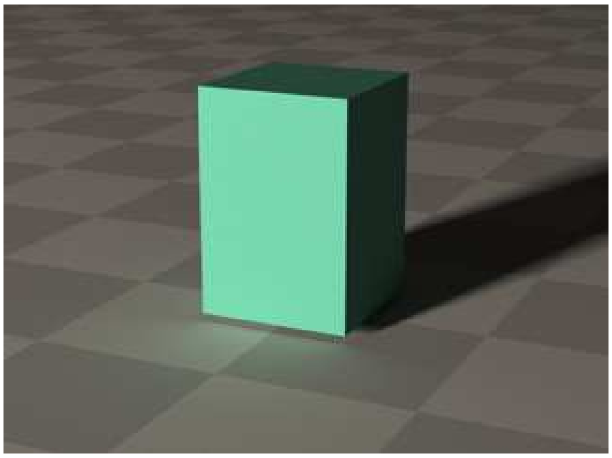

b. Shadows and interreflections added to **a**.

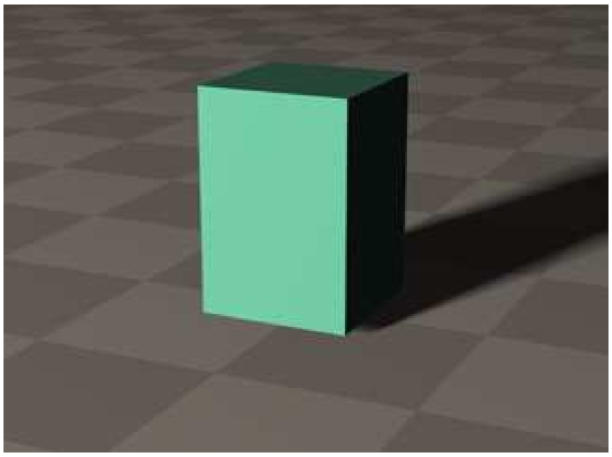

c. Shadows added to a.

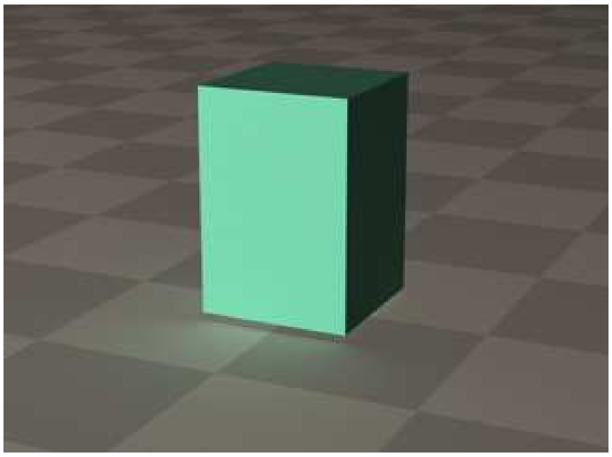

d. Only interreflections added to **a**.

Figure 2: Objects just above extended surface.

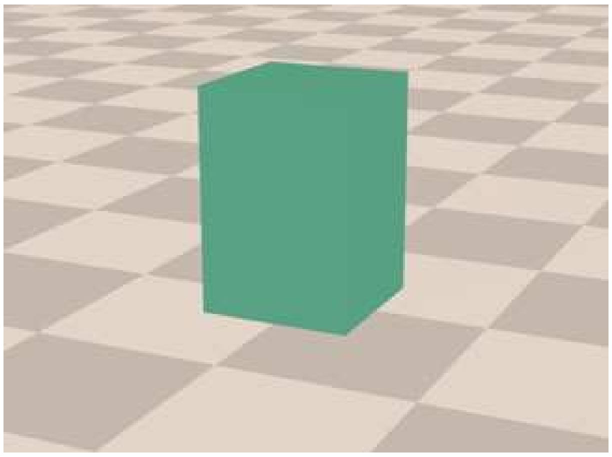

a. No interaction between lighting, block, and surface.

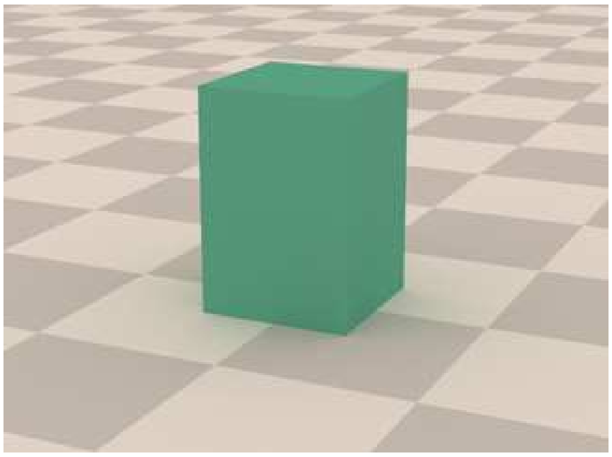

b. "Diffuse shadows" added.

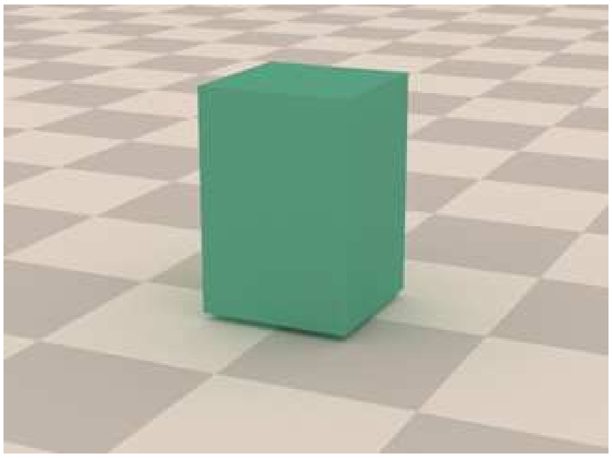

c. Same as **b** but with raised block.

Figure 3: Diffuse shadows also act as a cue for contact.

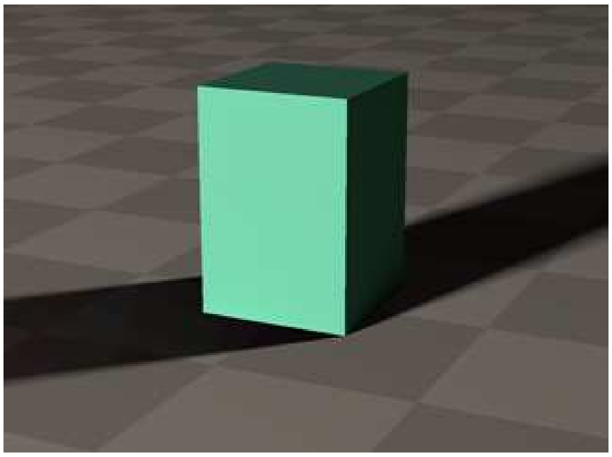

a. Two shadows.

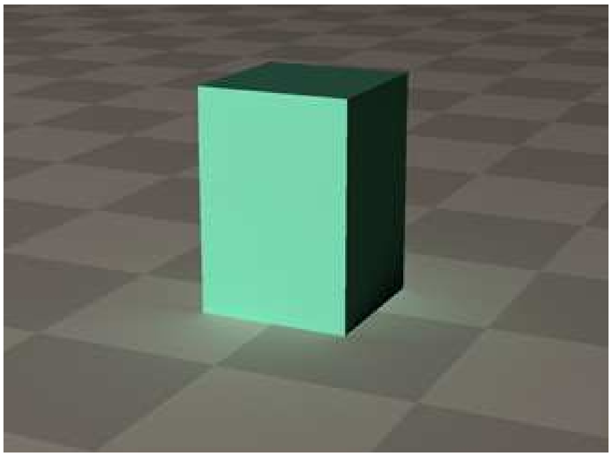

b. Two interreflections.

Figure 4: Wrap-around glue.

Applications involving real-time interaction with complex virtual environments need to present users with a clear sense of spatial organization and contact. Based on what we have just shown, shadows and indirect illumination are likely to play an important role in achieving this objective. This is a disquieting observation, since accurate rendering of interreflection is extremely costly. Shadows, while computationally less complex, still require substantial resources to render. This is particularly true when soft shadowing is done. As a result, it is important to probe more deeply into the role shadows and interreflection play in signaling contact between one object and another.

3 How accurate do we need to be?

Limitations of displays and computational power necessarily require that graphical renderings at best only approximate what would be seen in a physical environment. Thus, one of the goals of graphics research is to discover computable and displayable approximations to light transport that generate appropriate visual cues. It is particularly important to understand shadows and interreflections within this context, since both require substantial computational effort to render accurately and even then, limitations on dynamic range preclude faithful displays.

The first question to address is whether there is really a synergistic interaction between shadows and interreflections, as is suggested by Figures 1b and 2b. As shown in Figure 4, the answer is something of a surprise. Figure 4a shows a rendering comparable to Figure 1b except that the interreflection has been replaced by a second shadow added to the extended surface. Figure 4b is similar, except that this time the shadow has been replaced by a second interreflection. Comparing these images with Figures 1b and 2b shows that two shadows, two interreflections, or correct shadows and interreflection are all approximately the same in their ability to signal contact. This leads to the conclusion that it matters more that cues to contact be present along the entire visible perimeter of the object-object contact boundary than whether the cues are associated with shadows, interreflections, or both. The interchangability of shadows and interreflection also suggests that visual prominence, which is large for shadows and small for interreflections, may not be related to strength as a contact cue.

The next question to ask deals with the fidelity required of shadow or interreflection cues in order to signal contact. The answer to this too will come as something of a surprise.

The accuracy with which shadows need to be rendered in order to convey a sense of position in a larger spatial environment is highly variable, depending on the specific circumstances. Sometimes, soft shadows are important [8], other times they are not [9, 20]. When common motion of object and shadow is present, the actual appearance of the shadow seems almost irrelevant, with patterns that would never be confused for a shadow on their own sufficing perfectly well [9]. This is an effect that is commonly used in video games and animation to convey a sense of height. On the other hand, film requires more accurate shadows when it uses composited special-effects images that must touch the background, which may have as much to do with visual realism as it does with enhancing the sense of contact.

In the case of contact, we know that the physically realizable patterns of light are themselves highly variable, ranging from shadows to interreflections to diffuse shadows. The vision system could conceivably deal with this variability in at least two ways. One possibility is that the vision system is able to detect and interpret each type of cue on its own. A second possibility is that the vision system uses approximations sensitive to any of the contact cues. The tools of photorealistic rendering can be modified in ways that let us test which of these hypotheses is more likely. As with 3–D localization, the key is to discover whether or not approximate imagery can still convey the same sense of spatial organization.

In Figure 5, we have manipulated the shadow and interreflection patterns first presented in Figure 1c. Figure 5a is rendered with no shadows or interreflections. In Figure 5b, photorealistic shadows and interreflections have been added. In Figure 5c, the contrast of the indirect illumination effects has been increased, while the contrast of the shadow has been halved. Figure 5d was produced by subjecting the interreflection to a significant hue shift while reversing the contrast of the shadow. Even when the intensity and hue of shadows and indirect lighting are grossly distorted, the impression that the objects are in contact is not significantly diminished.[2]

4 Relevance to interactive rendering

In the previous section we showed that the ability for illumination cues to "glue" an object to another surface does not depend strongly on the details of the illumination cues, provided they are present. For applications where accurate spatial perception is more important than subjective realism, even coarse approximations to indirect illumination, conventional shadows, and diffuse shadows are able to indicate contact in a rendered image. To show the value of this in applications where real-time performance is of critical concern, we have developed an algorithm which is capable of generating perceptually effective contact cues using very crude approximations to shadows and interreflection.[3] This algorithm is simpler and faster than methods for accurately

[2]This insensitivity to hue does *not* carry over into other perceptual effects involving the interaction of indirect illumination and scene geometry [13].

[3]The algorithm modifies the intensities of the area using the multi-pass rendering and blending functionality of OpenGL. A dark projected texture is subtracted from the base intensity, and a separate light projected texture is added

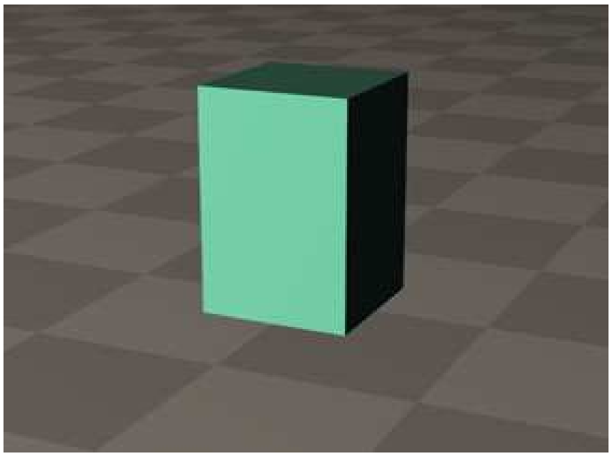

a. No shadows or interreflections.

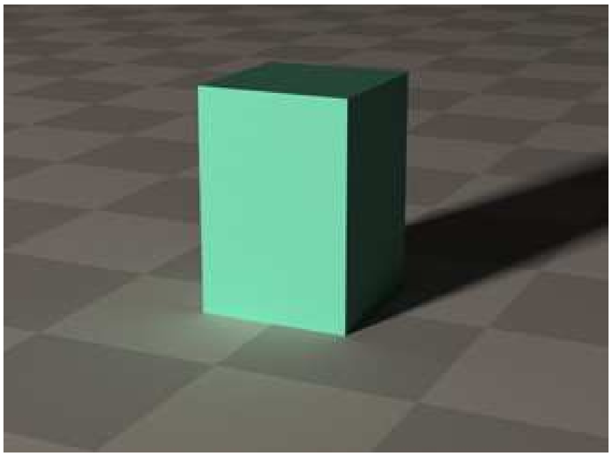

b. “Correct” shadow and interreflections.

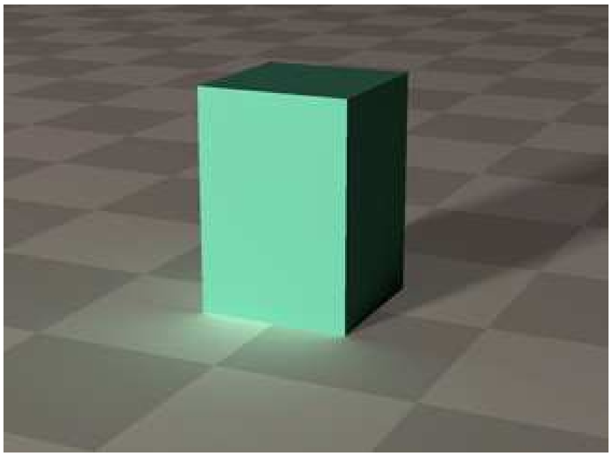

c. Bright interreflections, less dark shadows.

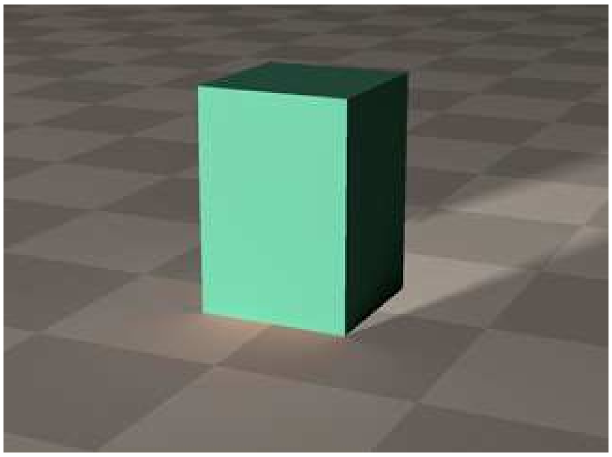

d. Hue shifed interreflections,
reverse contrast shadows.

Figure 5: Manipulating shadows and interreflection cues.

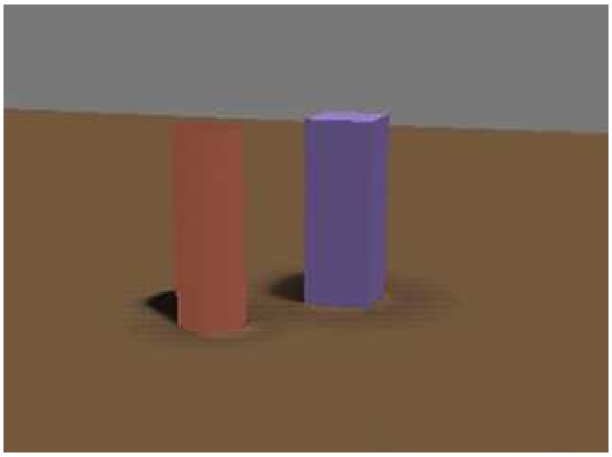

a. Simple approximation to shadows
and interreflections.

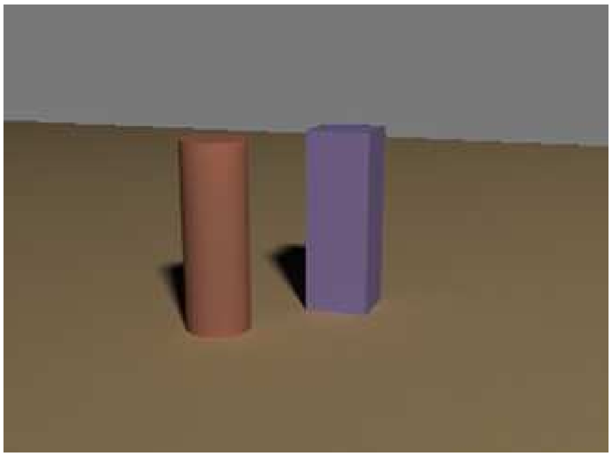

b. Accurate shadows and interreflections.

Figure 6: Coarse approximations are sufficient to convey a sense of contact.

rendering shadows in real-time (e.g., [21]) and is orders of magnitude more efficient than current methods for generating real-time indirect lighting effects. Because it uses projected textures, it also works for irregular ground planes. Figure 6a shows an image generated on a low-end workstation at interactive rates using this technique. Figure 6b shows the same scene rendered with a highly accurate renderer, running on a high-end processor and requiring close to an hour of CPU time. Although the second image is subjectively more realistic, it does not visually imply contact much more effectively than the approximation.

5 Conclusion

Four fundamental points are made above:

- *Both* interreflections and shadows – either alone or in combination – can serve as cues for contact.
- The visual prominence of shadows and interreflections is not an indication of their effectiveness in conveying information about spatial organization.
- The presence of contact cues along the whole of the line of contact is more important than whether the cues involve shadows or indirect illumination.
- Crude approximations to shadows and interreflections are sufficient to establish a sense of contact, even when the subjective sense of realism is seriously compromised.

While the graphics community has known that shadows are useful, indirect illumination has often been considered only important for applications where high subjective realism is valued. This historical down-grading of indirect illumination relative to shadows is probably due to the difference in visual prominence, which we have argued is not directly related to their effectiveness as contact cues. Our results imply that even VR systems where realism is not the primary goal might well benefit from at least approximating indirect lighting, particularly since in such cases simple techniques are likely to be sufficient for conveying an adequate sense of spatial organization. A further use is for visualization applications where demonstrating contact may be important, but shadows obscure too much information. The ability to use indirect lighting cues gives another choice to the designer of such applications.

Haddon and Forsyth, commenting on a computational analysis of the information available about scene geometry in shading patterns due to interreflections, observe that "the best prospect for extracting shape information from shading is to construct programs that observe stylized properties of shading and associate those properties with shape primitives or their properties" [22]. We reach the same conclusion approaching a closely related question from a human vision perspective: stylized patterns of lightness and darkness are sufficient to signal perceptions associated with shadows and interreflections in a way that is almost invariant to the actual radiance values that are present.

---

to the base intensity. These operations are accomplished using `glBlendFunc(GL_DST_COLOR, GL_ZERO)` and `glBlendFunc(GL_DST_COLOR, GL_ONE)` for shadows and indirect illumination respectively.

There are several important perceptual issues that are not addressed in this paper. We have demonstrated that in static scenes shadow and indirect illumination cues can establish object contact. When moving objects are brought into proximity with other objects, the dynamic cues may have different characteristics than the static cues. Also, non-diffuse effects may raise issues that do not arise in the diffuse case.

Finally, this paper is an example of the symbiotic relationship that is naturally arising between the graphics and psychology communities. Graphics researchers can provide variation to the optical behavior of the world which enable new sorts of programmable stimuli, and psychology researchers can help the graphics community prioritize what is rendered, so that efficient and effective algorithms can be developed.

References

[1] L. da Vinci, *The Notebooks of Leonardo da Vinci*, vol. 1. NY: Dover Press, 1970.

[2] D. P. Greenberg, K. E. Torrance, P. Shirley, J. Arvo, J. A. Ferwerda, S. Pattanaik, E. P. F. Lafortune, B. Walter, S.-C. Foo, and B. Trumbore, "A framework for realistic image synthesis," in *SIGGRAPH 97 Conference Proceedings*, pp. 477–494, Aug. 1997.

[3] A. Gilchrist and A. Jacobsen, "Perception of lightness and illumination in a world of one reflectance," *Perception*, vol. 13, pp. 5–19, 1984.

[4] H. A. Sedgwick, "Space perception," in *Handbook of Perception and Performance* (K. R. Boff, L. Kaufman, and J. P. Thomas, eds.), Wiley - Interscience, 1986.

[5] J. E. Cutting and P. M. Vishton, "Perceiving layout and knowing distances: The integration, relative potency, and contextual use of different information about depth," in *Perception of Space and Motion* (W. Epstein and S. Rogers, eds.), pp. 69–117, San Diego: Academic Press, 1995.

[6] A. Yonas, L. T. Goldsmith, and J. L. Hallstrom, "Development of sensitivity to information provided by cast shadows in pictures," *Perception*, vol. 7, pp. 333–341, 1978.

[7] L. R. Wanger, J. A. Ferwerda, and D. P. Greenberg, "Perceiving spatial relationships in computer-generated images," *IEEE Computer Graphics & Applications*, pp. 44–58, May 1992.

[8] D. Kersten, D. C. Knill, P. Mamassian, and I. Bülthoff, "Illusory motion from shadows," *Nature*, vol. 379, p. 31, 1996.

[9] D. Kersten, P. Mamassian, and D. C. Knill, "Moving cast shadows induce apparent motion in depth," *Perception*, vol. 26, no. 2, pp. 171–192, 1997.

[10] S. A. Shafer, *Shadows and Silhouettes in Computer Vision*. Kluwer Academic Publishers, 1985.

[11] J. R. Kender and E. M. Smith, "Shape from darkness: Deriving surface information from dynamic shadows," in *Proc. IEEE First International Conference on Computer Vision*, pp. 539–546, 1987.

[12] D. C. Knill, D. Kersten, and P. Mamassian, "Geometry of shadows," *Journal of the Optical Society of America A*, vol. 14, no. 12, pp. 3216–3232, 1997.

[13] D. Kersten and A. Hurlbert, "Discounting the color of mutual illumination: A 3D shape-induced color phenomenon," *Investigative Ophthalmology and Visual Science*, vol. 37, no. 3, 1996.

[14] S. K. Nayar, K. Ikeuchi, and T. Kanade, "Shape from interrefletion," *International Journal of Computer Vision*, vol. 6, no. 3, pp. 173–195, 1991.

[15] D. Forsyth and A. Zisserman, "Reflections on shading," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 13, no. 7, pp. 671–679, 1991.

[16] J. J. Koenderink and A. J. van Doorn, "Geometrical modes as a method to treat diffuse interreflections in radiometry," *J. Opt. Soc. Am.*, vol. 73, no. 6, pp. 843–850, 1983.

[17] C. Christou and A. Parker, "Virtual realism and virtual reality: A psychological perspective," in *Simulated and Virtual Realities: Elements of Perception*, pp. 53–84, Taylor and Francis, 1995.

[18] M. Segal, C. Korobkin, R. van Widenfilt, J. Foran, and P. Haeberli, "Fast shadows and lighting effects using texture mapping," *Computer Graphics*, vol. 26, pp. 249–252, July 1992. ACM Siggraph '92 Conference Proceedings.

[19] D. Lischinski, F. Tampieri, and D. P. Greenberg, "Discontinuity meshing for accurate radiosity," *IEEE Computer Graphics and Applications*, vol. 12, pp. 25–39, Nov. 1992.

[20] L. Wanger, "The effects of shadow quality on the perception of spatial relationships in computer generated imagery," in *Proc. Symposium on Interactive 3D Graphics*, pp. 39–42, 1987.

[21] P. Heckbert and M. Herf, "Simulating soft shadows with graphics hardware," Tech. Rep. CMU-CS-97-104, Carnegie Mellon University, January 1997.

[22] J. Haddon and D. Forsyth, "Shading primitives: Finding folds and shallow grooves," in *Proc. Int. Conf. on Computer Vision*, January 1998.