

Cartoon Dioramas in Motion

Ramesh Raskar, Remo Ziegler, Thomas Willwacher

TR2006-017 July 2006

Abstract

Cartoon animations delight the audience with moving characters but they remain on a flat 2D screen. The cartoon dioramas, on the other hand, are detailed, three-dimensional and allow physical interaction but they are static. We present techniques to combine the two in some limited cases. We illuminate static physical models with projectors. The images are generated with real time three dimensional computer graphics. We describe a system to demonstrate various visual effects such as non-photorealistic shading, apparent motion and virtual lighting on a toy-car model.

ACM SIGGRAPH 2006

This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

Cartoon Dioramas in Motion

Ramesh Raskar

Remo Ziegler

Thomas Willwacher

MERL, Mitsubishi Electric Research Labs

{raskar, ziegler, thomasw}@merl.com

Abstract

Cartoon animations delight the audience with moving characters but they remain on a flat 2D screen. The cartoon dioramas, on the other hand, are detailed, three-dimensional and allow physical interaction but they are static. We present techniques to combine the two in some limited cases. We illuminate static physical models with projectors. The images are generated with real time three dimensional computer graphics. We describe a system to demonstrate various visual effects such as non-photorealistic shading, apparent motion and virtual lighting on a toy-car model.

CR Categories: I.3.7 [Computing Methodologies]: Computer Graphics—3D Graphics

Keywords: perception, immersive environments, virtual reality, augmented reality, non-photorealistic rendering.

1 Introduction

Cartoon animations are made from a series of drawings simulating motion by means of small progressive changes in the drawings. The gross movement is achieved using relative displacement between layers of characters and backgrounds. The stored animation then can be seen at a desired frame rate as a video. Life-sized or scaled versions of cartoon characters and environments, such as funny shaped vehicles and buildings are also seen in shops, amusements parks and design studios. Inspired by recent advances in augmented reality, our goal is to combine the two mediums: two-dimensional cartoon animation and static three dimensional dioramas of physical models into lively non-realistic setups and non-realistic movement. We substitute pre-recorded sequence of images with real-time three dimensional computer graphics renderings to create an animated physical world. In this paper we describe a system to show some limited effects on a static toy-car model and present techniques that can be used in similar setups. Our focus is on creating apparent motion for animation.

1.1 Motivation

Our approach is based on techniques in two very different fields, projector-based augmentation and illusions for creating apparent motion.




Figure 1: (Top-left) The diorama of a car and approximate background. (Bottom) A cartoon-like environment created by image projection in which the car appears to move (Top-right) A night-time simulation using virtual headlights.

1.1.1 Augmenting Physical Models

Projector can be used to change the surface appearance of neutral colored physical models by illuminating them with rendered images. The projectors are treated as ‘Shader Lamps’ [Raskar et al. 2001]. In this type of projector-based augmentation, complete illumination of complex 3D shapes is made possible by rendering geometric models of those shapes from multiple overlapping projectors. The geometric models are pre-authored with detailed color, texture and material properties to reproduce desired surface appearance. We use the same paradigm to create a cartoon-like appearance. Although, reproduction of surface reflectance is well studied, simulating apparent motion from static models remains unexplored.

1.1.2 Apparent Motion

The motion analysis by humans is divided into three levels of processing: retinal motion sensing, 2D integration and 3D interpretation. The need to sense retinal motion, and analyze it as quickly as possible, places great demands on the visual system. A great deal of perceptual and physiological research has been conducted to discover the properties of these mechanisms. With several (retinal and other) cues, some potentially conflicting, the visual system attempts to integrate them into a meaningful ‘best’ solution, which may actually be incorrect [Mather et al 1998]. The illusions that may result in such cases provide us an opportunity.

We are specifically interested in exploiting the perception of motion when there is no corresponding physical displacement in space. This perception of motion is known as *apparent motion* [Smith and Snowden 1994]. There are four well-known types of apparent motion: motion after effect (MAE) (e.g. upward motion of objects near a waterfall), phi phenomenon (e.g. filled in motion in a movie of flickering animation stills), induced motion (e.g. perceived movement due to a slowly moving train on the next track) and autokinetic movement (e.g. perturbation seen at small light in a dark room). In addition, one type of physical motion may result in some other types of apparent motion due to either 2D integration or 3D interpretation error. A good example is moving shadows which act as a strong cue to depth change.

Using 3D computer graphics, in this paper, we mainly exploit the induced motion effect and the 3D interpretation error for some very constrained cases. To induce motion, we segment a continuous static physical scene into sub-parts. We then illuminate each sub-part with rendered animations so that the multiple movements appear compatible. We also force temporary and invalid 3D interpretation using shadows, lighting and texture movements.

Motion perception is, however, a complex topic. New models in understanding apparent motion are still evolving. There may be many other ways of using the errors in human motion analysis.

1.2 Previous Work

In cartoon and in live action films, characters are made to appear in motion relative to the background by moving the background in the opposite direction. However, there are surprisingly few efforts at generating apparent and sustained motion in the real world. A well-known example is the singing heads at Disney's "Haunted Mansion" [Liljegren 1990]. Pre-recorded video, synchronized with sound, is projected on semi-transparent busts of men to create a compelling effect that the busts are changing shapes corresponding to facial expressions.

Based on the underlying complexity, the previous work in using projectors for augmented reality can be divided into two main groups (i) projection of useful information on a planar surface, and (ii) more complex insertion of 3D shapes and attributes into the real world. There are several examples of 2D augmentation - interactive 2D images projected on flat surfaces in a room to enhance the user's environment, Luminous room [UnderKoffler et al 1997], digital desk [Wellner 1993], and smart whiteboards. The 3D augmentation examples are fewer, Shader Lamps [Raskar et al. 2001][Bandyopadhyay et al. 2001] and the related project Being There [Low et al. 2001].

1.3 Contribution

Our main contribution is a general set of techniques to allow a limited combination of non-photorealistic images with physical models to generate non-realistic appearance and movement. We demonstrate the ideas in the context of a system, discuss methods to make them practical and describe user experiences. We hope to contribute to the body of research in immersive non-realistic animation and spatially augmented reality.

2 Effects

In this paper, our demonstration system includes augmentation of a diorama made up of a toy-car model and its surroundings. In

that context, we discuss and describe new rendering techniques that take into account the relationships between the user position, real objects, virtual objects, light projectors and the surface attributes of the virtual objects. The techniques can also be applied to other types of dioramas.

There are effectively two steps: augmenting objects so that they have cartoon appearance and simulating apparent motion.

2.1 Spatial Augmentation

Many physical attributes can be effectively incorporated into the light source to achieve a perceptually equivalent effect on a neutral object. While the results are theoretically equivalent for only a limited class of surfaces and attributes, our experience is that they are quite convincing and compelling for a broad range of applications. The task is easier when a non-realistic appearance is desired. This subsection describes a technique very similar to the one used in [Raskar et al. 01]. The approach for a general projector-based augmentation system is summarized below.

During pre-processing

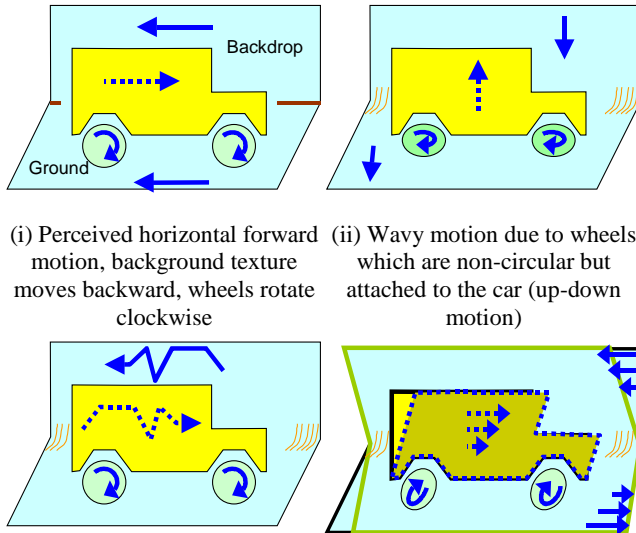
- Create 3D graphics model, G , of physical object
- Create 3D graphics model, B , of background
- Approximately position the projector
- Find perspective pose, P , of the projector wrt the physical object

During run-time

- Get user location, U
- Get animation transformation, T
- Modify G 's surface attributes
- Render G using the pose P , and user location U
- Transform B using T^{-1} , B'
- Render B' using the pose P , and user location U
- Modify image intensity to compensate for surface orientation

To reproduce purely view-independent surface appearance (diffuse reflectance), user location is not required. For view-dependent effects such as specular highlights, approximate user location is necessary. The user's location, U , can be tracked using magnetic or optical tracking technology. In our case, we assume the user is at a sweet-spot standing in front of the car and do not track the user.

The projector projection matrix, P , is obtained using an off-line calibration process similar to the technique used for finding internal and external parameters of a camera [Faugeras 1993]. We take a set of fiducials with known 3D locations on the physical object and find the corresponding projector pixels that illuminate them. This allows us to compute a 3x4 perspective projection matrix, up to scale, which is decomposed to find the internal and the external parameters of the projector. The internal parameters represent the focal length and principal point, while the external parameters represent the rigid transformation (translation and rotation) with respect to the coordinate system of the physical object. The rendering process uses the same internal and external parameters to render the pre-authored geometric model, G , so that the projected images are registered with the physical objects. During run-time, instead of the object, G , the background, B , is transformed to create the apparent motion. At each frame, an intensity correction stage pre-multiplies the projected image with intensity weights to compensate for the local surface orientation. Otherwise, surfaces normal to the incident light will appear



(i) Perceived horizontal forward motion, background texture moves backward, wheels rotate clockwise

(ii) Wavy motion due to wheels which are non-circular but attached to the car (up-down motion)

(iii) Bumpy movement on a rough surface

(iv) Shear during acceleration

Figure 2: Types of apparent motion for a car. Rendered motion is shown with blue arrow. The resultant apparent motion is shown with dashed blue arrow.

brighter than surfaces illuminated obliquely due to the cosine fall-off.

2.2 Cartoon Shading

Our focus is mainly creating cartoon-like (apparent) movements. To enhance the cartoon-feel, however, we implemented well-known non-photorealistic shading techniques with some modifications. The scene is made up of a toy-car model, with a simple background of a horizontal and a vertical surface. The horizontal surface coarsely represents the ground surface, either a road or rough terrain. The vertical surface is the backdrop which represents everything else including bushes, trees and sky. To improve the visual fidelity, the background can be made arbitrarily complex as long it has a constant profile along any plane perpendicular to the direction of apparent motion..

We use the technique proposed by [Lake et al. 2000] to render a flat-shaded look. As described below, we modify the color quantization procedure to take into account the cut-off angle of the spot-beams for street lights. For static scenes, we also highlight sharp edges as seen in the traditional 2D cartoon animation stills [Raskar 2001]. However, we realized that for objects with apparent motion, the sharp features are distracting because they work against the notion of motion blur. For dynamic three dimensional scenes, human visual system segments scenes based on grouping of objects in motion [Mather 1998]. Hence, we speculate that, in the retinal image, the boundaries between regions with different optical flow are more important than edges corresponding to discontinuity in surface normal (i.e. sharp features). However, this aspect needs more investigation. Edges corresponding to discontinuity in depth (i.e. silhouettes) are generally a subset of the boundaries between regions with different directions or rates of optical flow.



Figure 3: Vertical displacement by shifting the background and shadows. These two pictures are taken from the same camera position, so the car in both images is at the same location. Note (i) the change in position of the virtual shadows and (ii) the parallax for the edge between ground and backdrop with respect to the right headlight.

2.3 Motion Effects

The car is simulated as driving along a road, on a rough surface, or in various other environments. To create apparent motion, we illuminate wheels with images of rotating wheels. The images of the background (made up of the backdrop and ground) move in a direction opposite to the intended car movement. In the simplest case, as seen in the video, the car appears to move forward i.e. left to right. To create this effect, the wheels rotate clockwise and the background moves right to left (Figure 2(i)). The crucial task for a believable movement is maintaining consistency between the angular wheel movement and corresponding translation of the ground (and the backdrop). For any duration,

$$\int (\text{wheel perimeter arc length}) = |\text{displacement of the bkgnd}|$$

This ensures that the wheels are not sliding while the car is in motion. A small amount of motion blur is added to the wheels. Along the direction of the motion, the background geometry with the associated texture maps is infinitely long, and is implemented using a simple sink and a source.

We experimented with many types of cartoon or non-realistic motions. Two additional factors that add to the effect are sound and removal of (physical) sharp edges in the background.

2.3.1 Wobbly Motion

Slow moving cartoon cars usually have a wavy or bumpy movement resulting in a small periodic or random displacement along the vertical direction. This sometimes is emphasized by non-circular wheels. In our case, the car is static. Hence, to create the desired apparent vertical shift while the car is in motion, we instead translate the background in the opposite (vertical) direction (Figure 3). During rendering, the 3D geometric model of the backdrop as well as the ground is translated vertically. The amount of translation, in the case of non-circular wheels, is determined by distance between the point of contact of the wheels from wheel axis (Figure 2(ii)). The distance traveled is again consistent with integration of the arc length of wheel perimeter. We encourage the reader to watch the accompanying video to view both types of motions. (The video gives a good idea of the effect of the transitions, but it is stronger when experienced in person with all the three dimensional cues. Please see the subsection on user reactions.)

It is important to note that to create apparent motion, sometimes, the underlying physical geometry does not match up with the virtual model. For example, with vertical displacement, the boundary between vertical and horizontal surface (shown as



Figure 4: Shading due to street lights as spot lights.

brown sharp edge in Figure 2(i)) does not correspond to the projection of the boundary between the backdrop and ground. Hence, first we need to eliminate any sharp edges that the user may be able to use as a frame of reference. A conflicting cue affects the perception of vertical component of the motion. Second, we need to represent the surface with the approximate average of corresponding variation in the geometry. The solution is to the smooth the edge using a paper with small curvature (see brown lines in Figure 2(ii), (iii)).

2.3.2 Lights and Shadows

Shadows provide a very important cue in apparent motion [Kersten et al. 1997]. For example, in a two-dimensional image sequence (i.e. without any other depth cues), a ball moving diagonally upwards across the screen can be made to appear as if it is gradually moving away from the viewer by rendering a shadow that appears to stay with and under the ball. The cue overrides contradictory information regarding the unchanging size of the ball. The strength of the effect does depend on the properties of the shadow region. For example, shadows with a penumbra, and which fall below the object, work best.

According to [Kersten et al. 1997] the existence of the illusion shows that our perception of the spatial layout of scenes relies on the assumption that the light source responsible for the shadow is stationary, so that any motion of the shadow is attributed to the object casting it.

We enhance the wobbly motion effect by rendering virtual shadows from directional light source (sun) or local lights (street lights). The change in shadow position creates the illusion that the change is a result of changing vertical distance between the car and the ground. We noticed that, it is not necessary to use the same light position to calculate shading and shadows! Further, perceptually, the movement of virtual shadows (Figure 3) is not affected by the fixed real shadows of the physical car on the background.

For night time simulation, the only cues are headlight beams and shading due to the street lights. The spread of the parabolic projection of the headlights and change in overlap between the two beams indicates the vertical shift. The color of the two beams is intentionally chosen slightly different, so that the lateral shift as the car moves up and down is clearly seen (Figure 1 and 7). We also exploit spot lights from the street lights. The basic cartoon shading equation is modified to take into consideration the spot-cut off angle (Figure 4).

2.3.3 Acceleration dependent modifications

Cartoon animators emphasize acceleration, such as a sudden start or a screeching halt, by geometric deformation of the object. For example, a car appears to ‘work hard’ to move forward while starting when the top of the car is sheared in the direction of the acceleration. Similarly a hard brake and stop is indicated by

shearing the top backwards. Since we cannot shear the physical model, we enhance this effect using two tricks.

First, we implement a shear in the background that is opposite of the shear expected in the car (Figure 2(iv)). The shear is along one dimension, along the vertical axis. Hence, for example, during a sudden start, the background shears backward and the shear at a point is proportional to the vertical distance from the center of the car. The points above the vertical center of the car translate backwards while points below translate forward (Figure 5). Without a loss of generality, let's assume that the car center is at the origin, the vertical direction is parallel to the z-axis, and the forward direction is parallel to the x-axis. Then the shear at a given frame is achieved using a simple transformation matrix $[1, 0, -a, 0; 0, 1, 0, 0; 0, 0, 1, 0; 0, 0, 0, 1]$. Where a is proportional to the acceleration at that frame.

Since the acceleration is positive during starting, negative during braking, and zero during constant speed and velocity, the same shear equation can be used throughout the animation.

For the second trick, we observe that rapid acceleration also means relative slide between the ground and the wheels. Hence, for example, a sudden brake results in halt in rotation of the wheels, but the background continues to move (Please see the video).

2.3.4 Tracked Illumination

We synchronize and update the rendering of the objects so that the surface textures appear glued to the objects even as the objects move. In this case, as seen in the video, we rotate the car (along with the background) on a turntable. Keeping the virtual light sources fixed, we see corresponding shading changes.

2.3.5 Issues

The type of effects possible are obviously limited. For example, we played with motion blur and motion lines typically seen in cartoon animations, but without success. A simple problem is lack of corresponding physical surface to show motion lines behind the car. We also could not bring in effects such as fog, smoke or fumes. Displaying fog simply washes out the image on the model. During the night time simulation, the backdrop, which is supposed to be at a large depth and hence dark, gets illuminated by the secondary scattering of headlight beams.

3 Other Applications

A setup to introduce non-realistic appearance and motion can be used for many other applications. Some examples are educational




Figure 5: Shearing of the background during acceleration.

interactive setups (e.g. physics experiments), science and art museums, stage drama shows, showroom displays, virtual sets and physical mock-ups for special effects. We assumed the user to be at a sweet-spot, but with head-tracking many view-dependent effects can be incorporated. They include simulating complete (possibly non-realistic) reflectance properties (BRDF), inserting animated floating 3D objects close to the physical surface (e.g. crawling ants or flowing water) or advanced real-time non-photorealistic shading effects [Praun et al. 2001].

We used a simple turn-table to demonstrate illumination of a moving physical object. This can be extended further in many ways depending on the degrees of freedom of the movement. For example, current car motion is mostly uniform. By adding small and sudden turntable rotations in and out of the plane, the car can be made to appear to be taking turns around a corner or avoiding an obstacle in the middle of the road.

While there is a growing body of physiological research and studies of apparent motion, the work so far is limited to two-dimensional experiments and in many cases limited to computer screen simulations. It will be interesting to use the proposed setup to study other type of stimulation for the human visual system and better understand the model of motion analysis.

4 Benefits and limitations

A key benefit of this type of spatial augmentation is that a visually rich, large field-of-view diorama can be generated with greater amount of integration of virtual (surface and motion) attributes with the real world. Multiple people can enjoy the enhanced diorama just as in the traditional dioramas.

A crucial problem is the dependence on display surface properties. A light colored diffuse object with smooth geometry is ideal. It is practically impossible to render vivid images on highly specular, low reflectance or dark surfaces. The ambient lighting as well as secondary scattering from diffuse surfaces can also affect the contrast in the images. This limits the applications to controlled lighting environments with restrictions on the type of objects with which virtual objects will be registered. Shadows can also be a problem if close human interaction is desired.

5 Implementation and lessons

Currently we use a Mitsubishi X80U LCD projector, 1000 lumen, 1024x768 resolution. The toy-car (Barbie Beetle, bought at a toy-store) is 50cm x 25cm x 20cm (Figure 6). The 3D model is generated from multiple photographs using PhotoModelerPro and then authored in 3DSMax. The model is very coarse, 602 triangles. We struggled with creating smoothly varying normals across the car surface, resulting in shading artifacts. A better 3D scanning method should overcome those problems. The animations are rendered at approx. 15-20 fps. To determine projector pose during calibration, we used 15 pairs of pre-determined 3D points on the car surface and the corresponding projector pixels that illuminate them. The projectors pixels are interactively found by projecting a cross-hair and moving the mouse till the cross-hair lines up with one of the 3D feature point. This process takes only two minutes. The resulting re-projection error is less than one pixel allowing near-perfect static registration. To find the axis of rotation of the turn table, a second set of projector pixels for the same 3D points, after angular displacement, is used.



Figure 6: Setup with a projector and simple diorama.

5.1 Use Reactions

We do not have systematic user experience data but the system has been seen by over two thousand people. The working system has been demonstrated to visitors of the lab, at an art show, during UIST'01 conference, and to dozens of Disney researchers in Orlando, Florida. Most users get excited at the following stages: (i) when the car *starts*, (ii) when the car motion *switches* to bumpy motion, (iii) when the car is rotated on the turn table while maintaining its augmented surface appearance. As mentioned earlier, the effect is most noticeable during any changes in apparent motion and it lasts for a short time. The human visual motion detection system quickly becomes balanced again. Hence, it is important to constantly change the motion parameters. The synchronized car engine (or brake) sound seems to have a huge impact, it generates anticipation before the car starts and reinforces the notion of a car movement (as opposed to movement of the background). We implemented a photo-realistic version as well as the cartoon-like version (Figure 7). The cartoon-like version appeals more when the car is in motion. Viewers, not familiar with the technique, are unimpressed to see the static cartoon version. This is probably because we are all used to looking at much higher quality cartoon-like setups. But, the



Figure 7: Photo-realistic diorama (left column) vs. cartoon diorama (right column). Note the headlight beams, specular appearance and textures.

interest increased when we interactively changed the surface appearance of the car or the background. Many viewers, even after explaining that we are projecting on white surfaces, are very surprised when the projector light is blocked to show the underlying dull setup of simple static models. Comparing figure 6 and figure 7, it is surprising to see the amount of fine detail possible with purely projected images.

So far, two viewers (both women in their 30s) have complained of feeling dizzy (a well known side effect of VR/VE worlds). Children show a tendency to run and catch the moving car. Advanced technical viewers after careful observation sometimes notice the lack of complete illumination (because we are using only one projector and, for example, shadow of the car on the backdrop is not filled in by multiple projectors). Please view the video and more images at <http://www.shaderlamps.com>.

6 Conclusion and Future work

We have presented a system and associated techniques to show apparent cartoon-like motion on a simple diorama. The techniques allow only a limited combination of rendered animation and static neutral colored physical objects. But, the proposed methods are sufficiently general and can be used in a variety of applications.

Similar to *non-realistic* rendering, there is an interesting unexplored area of *non-realistic* motion simulation. We have presented some preliminary ideas and a prototype system. There are many possible directions for combining apparent motion effects with real time 3D rendering in projector-based augmentation. For example, despite the visually compelling results, the visual adaptation effect lasts for a short time. How can we make it last longer? For most effects, the motion detection system quickly becomes balanced and the apparent movement becomes weaker or stops. However, the four-stroke apparent motion [Mather et al. 1998] is a wonderful illusion where the motion appears to continue forever in the same direction. The animation involves a repeating cycle of four different frames. e.g. a car moves forward and then backward, but the backward step is accompanied by a reversal in intensities. The effect is to reverse the direction of perceived motion, resulting in forward and then again forward motion. Can we achieve similar effects on three dimensional objects?

We currently limit effects to static models. However, one can imagine using phi phenomenon to create apparently morphing shapes out of two or more static models. Shape change over time presents the visual system with a correspondence (matching) problem. The 'best' solution is interpreted based on a few geometric properties. Thus, successively illuminating different shapes creates the apparent and smooth interpolating morph. This has been explored in 2D computer simulations and may be extended to 3D objects by changing focus of attention from one static shape to the next with projected light.

Lastly, our use of synchronized realistic sound was limited to enhancing the visual effects. But, as noted recently by [Shams et al. 2000] sound can create virtual motion perception. It should be worth exploring the use of realistic or even non-realistic spatialized sound to further intensify the effects.

Acknowledgements

We would like to thank members of the ShaderLamps project, Greg Welch, Kok-lim Low and Deepak Bandyopadhyay for useful discussions.

References

- BANDYOPADHYAY D., RASKAR R., FUCHS H. 2001. Dynamic Shader Lamps: Painting on Movable Objects. In *Proceedings of Int. Symp. On Augmented Reality*.
- FAUGERAS, O. 1993 *Three-Dimensional Computer Vision: A Geometric Viewpoint*. MIT Press, Cambridge, Massachusetts.
- KERSTEN, D., MAMASSIAN P., AND KNILL D. 1997. Moving cast shadows induce apparent motion in depth. *Perception*, 26, 171-192.
- LAKE, A., MARSHALL, C., HARRIS, M. AND BLACKSTEIN, M. 2000. Stylized Rendering Techniques for Scalable Real-Time 3D Animation, In *Proceedings of Non-Photorealistic Animation and Rendering*, Annecy, France, June 5-7.
- LILJEGREN, G. E. AND FOSTER, E. L. 1990. Figure with Back Projected Image Using Fiber Optics. US Patent # 4,978,216, Walt Disney Company, USA.
- LOW, K., WELCH, G., LASTRA, A., FUCHS H. 2001. Life-Sized Projector-Based Dioramas. *Symposium on Virtual Reality Software and Technology*.
- MATHER, G., VERSTRATEN, F., AND ANSTIS, S. 1998. *The Motion Aftereffect: a Modern Perspective*. MIT Press, Cambridge, Massachusetts.
(http://www.biols.susx.ac.uk/home/George_Mather/Motion/)
- MILGRAM, P., TAKEMURA, H., UTSUMI, A., AND KISHINO, F. 1994. Augmented Reality: A class of displays on the reality-virtuality continuum. SPIE Vol. 2351-34, Telem manipulator and Telepresence Technologies.
- PRAUN, E., HOPPE, H., WEBB, M., FINKELSTEIN, A. 2001 Real-Time Hatching. In *Proceedings of ACM SIGGRAPH 2001*, ACM Press / ACM SIGGRAPH, E. Fiume, Ed., Computer Graphics Proceedings, Annual Conference Series, ACM.
- RASKAR, R., WELCH, G., LOW, K., BANDYOPADHYAY D. 2001. Shader Lamps, Animating Real Objects with Image Based Illumination. In *Proceedings of the 12th Eurographics Workshop on Rendering*.
- RASKAR, R. 2001. Hardware Support for Non-photorealistic Rendering. In *Proceedings of the ACM/Eurographics Workshop on Graphics Hardware*.
- SHAMS, L., KAMITANI Y. AND SHIMOJO, S. 2000. What you see is what you hear. *Nature*, pp 788. Dec 14, 2000. (<http://neuro.caltech.edu/publications/shams.shtml>)
- SMITH A, AND SNOWDEN R 1994. *Visual Detection of Motion*. London: Academic Press.
- UNDERKOFFLER, J. 97. A View From the Luminous Room. Springer-Verlag London Ltd., *Personal Technologies* (1997) 1:49-59.
- WELLNER, P. 1993. Interacting with paper on the DigitalDesk. *Communications of the ACM*, 36(7):87-96. July

<http://www.shaderlamps.com>