Max-Planck-Institut
für biologische Kybernetik
Arbeitsgruppe Bülthoff

Technical Report No. 6 June 1994

Moving Cast Shadows and the Perception of Relative Depth

Daniel Kersten, Pascal Mamassian & David C. Knill

Abstract

We describe a number of visual illusions of motion in depth in which the motion of an object's cast shadow determines the perceived 3D motion of the object. The illusory percepts are phenomenally very strong. We analyze the information which cast shadow motion provides for the inference of 3D object motion and experimentally measure human observers' use of this information. The experimental results show that cast shadow information overrides a number of other strong perceptual constraints, including viewers' assumptions of constant object size and a general viewpoint. Moreover, they support the hypothesis that the human visual system incorporates a stationary light source constraint in the perceptual processing of shadow motion. The system imposes the constraint even when image information suggests a moving light source.

DK and PM were supported by the National Science Foundation (BNS-9109514) and the Max Planck Society. DCK was supported by the Air Force Office for Scientific Research (AFOSR 90-2074) and NIH (EY09383-01A1). We thank Albert Yonas and Deborah Rossen for their comments and suggestions. Correspondence should be sent to: Daniel Kersten, N218 Elliott Hall, Psychology Department, 75 East River Road, Minneapolis, MN 55455, U.S.A.. Email: kersten@eye.psych.umn.edu.

This document is available as /pub/mpi-memos/TR-6.ps.Z via anonymous ftp from ftp.mpik-tueb.mpg.de or by writing to the Max-Planck-Institut für biologische Kybernetik, Spemannstr. 38, 72076 Tübingen, Germany.

1.0 Introduction

The relative displacement between an object and its cast shadow in an image provides an important source of visual information about the spatial layout of objects. Leonardo da Vinci elucidated the principle relating shadow displacement and the perception of relative depth in his notebooks: "...when representing objects above the eye and on one side--if you wish them to looked detached from the wall--show, between the shadow on the object and the shadow it casts, a middle light, so that the body will appear to stand away from the wall." (da Vinci, 1970) Artists regularly exploit this principle in static drawings and paintings of 3D scenes, and psychophysical research has shown the salience of static cast shadow information for judgments of depth (Yonas, 1978). Yonas et al. (1978) were able to show that the location of a cast shadow was able to influence the judged depth and height of an object above a ground plane in observers as young as three years old. The role of dynamic shadows in human perception, however, has received no scientific study. Because movement due to shadow boundaries is almost always present in the retinal image, understanding how the visual system processes shadow motion is a fundamental issue in vision. In this paper, we report a set of controlled experiments and phenomenal demonstrations which show:

- the relative motions of objects and their cast shadows in an image can produce remarkably strong percepts of 3D motion
- information provided by the motion of an object's shadow overrides other strong sources of information and perceptual biases, such as the assumption of constant object size and a general viewpoint
- image features such as shadow darkness can be utilized, but are not necessary for the perception of depth from moving cast shadows
- support for a prior assumption of a *stationary light source constraint* by the visual system.

2.0 The Phenomenon

2.1 Experiment 1: Cast shadow motion is sufficient for the perception of motion in depth.

The first question is whether shadow motion is in fact used for the perception of relative motion in depth. Although it is reasonable to assume that an affirmative answer would follow given the evidence from judging static shadows in pictures, it is not necessarily the case for at least three reasons. First, the fact that a pictorial cue is useful for judgments of depth does not necessarily imply that variations of that cue will produce the perception of motion in depth. The reason is that judgments based on static cues with long viewing times can involve conscious reasoning as well as perceptual processing. Second, the computational problem of identifying shadows is known to be very difficult. The real-time requirements of identifying shadows in motion may be even harder. Although processing of static shadows has received some study in computer vision (Waltz, 1972; Shafer, 1985), with few exceptions (Kender, J. R., & Smith, E. M., 1987) computer vision has ignored moving cast shadows. Third, if vision's primary function is to determine the identity and spatial layout of surfaces and objects, one could argue that variation of intensity in the image due to illumination might be discounted early given the processing overhead required. A related argument that the visual system discounts variations in illumination in order to determine surface color has been discussed since Helmholtz.

The computational difficulty lies in the fact that optic flow is determined by a complex interaction of causes. The form and evolution of optic flow is influenced by changes in the viewpoint of the observer, positions and shapes of the objects, and the illumination. Unlike the effect of shape, the effect of illumination on the image is not just local. Shadow boundaries are determined by the illumination, the casting object, the receiving object and the viewpoint. Unfortunately, there is no unambiguous local cue for a shadow edge. Nevertheless for human shape perception, static cast shadow boundary is useful for object shape perception as well as depth perception (Cavanagh, & Leclerc, 1989). How are shadows identified? Cavanagh (1991) argues, based on work with images of faces, that the identification of shadow boundaries and utilization of shadow information may in fact *follow* the recognition of the category that a shape belongs to. From this point of view, it is not unreasonable to suppose that judgments involving static shadows may require processes that are too slow to be useful in processing dynamic shadows for depth information. Yet, moving cast shadows are used routinely in cartoon animations and in video games; but does this merely enhance the realism of the pictures, or is this information useful for depth?

Figure 1 illustrates the well-known effect of shadow displacement on the perception of relative depth in static images: the closer an object is to its cast shadow in an image, the closer it appears in depth to the background surface. We created a motion analog of this demonstration, in which the shadow cast by a stationary square moves back and forth relative to the square (figure 2). We then ran a simple psychophysical experiment (Experiment 1) to test whether subjects would see the square move in depth (see figure 2 caption for details). When the shadow was rendered real-

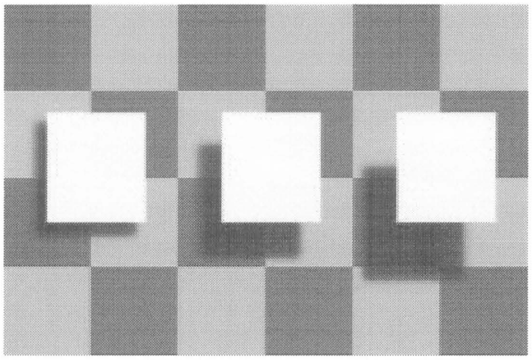

Fig. 1. Increasing the displacement between the cast shadows and the three foreground squares tends to produce an impression of increasing depth (from left to right) relative to the background checkerboard.

istically dark, subjects reported seeing the square move toward and away from the background surface 78% of the time. When the shadow was implausibly lighter than its background, subjects only reported seeing the square move in depth 40% of the time. Subjects who perceived the motion reported that the percept was phenomenally strong and immediate.

The result clearly shows an effect of cast shadow motion on observers' perception of 3D motion of an object. Moreover, a close analysis of the experimental stimuli reveals that for the observers who saw the motion in depth, the motion of the shadow overrode a number of conflicting cues which suggested that the square was stationary: the lack of any change in size of the square, and the lack of any 2D motion of the square in the image. That these features of the stimulus would suggest object stationarity results from the human visual system's bias to assume, first, that objects do not change size over time (related to object size constancy, cf. Gogel, Hartman, & Harker, 1957) , and second, that the viewer is viewing the scene from a non-accidental, or general viewpoint (Biederman, 1985; Nakayama, & Shimojo, 1992) . The assumption of object size constancy would lead the visual system to interpret the non-changing size of the square as information that the square was stationary, since any change in depth of a rigid object would lead to a correlated change in the size of the object's image. The general viewpoint assumption would lead the system to interpret the lack of any 2D motion of the square also as information for stationarity, since for almost all viewpoints (except one "accidental" view in which the viewer is looking along the direction of motion), motion in depth of an object would cause a correlated 2D motion of the object's image. The cues for stationarity could well have led to the result that on 22% of the trials with dark shadows, subjects did not see the square move in depth. This raises the possibility that elimination of the stationarity cues would lead to greater effects of cast shadow motion on observers' percepts of 3D motion. Unfortunately, one cannot remove the effect of the size constancy constraint from an experiment, since the image size of an object is an inherent property of a stimulus. What is possible, however, is to remove the effect of the general viewpoint constraint by simply moving the object, as well as its cast shadow, in the image plane.

2.2 Demonstration 1: Phenomenally strong illusion of motion in depth with accidental view removed.

We generated a 3D graphics simulation which we call the ball-in-a-box animation (figure 3), in which we simulated a ball moving inside a box in such a way that it followed a diagonal trajectory in the image plane. As in Experiment 1, the size of the object's

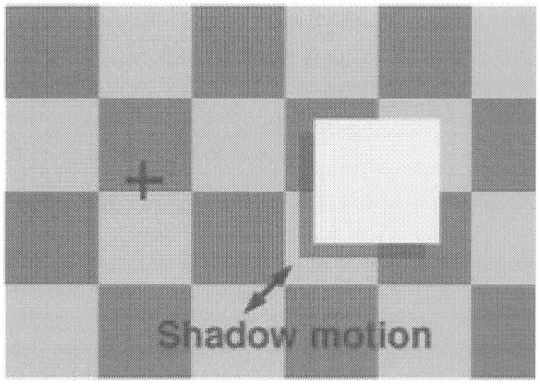

Fig. 2. Observers were asked to look at a fixation mark (+) placed on a checkerboard plane which subtended 6.6 x 10° of visual angle. Viewing distance was 500 mm. At a position 4.1° to the right of the fixation point, a foreground square was superimposed over a sharp shadow of the same size as the square. In a 500 msec. animated sequence, the shadow oscillated for one cycle through a 0.34° displacement from the foreground square. The foreground square remained stationary throughout the animation. Observers were asked to indicate whether the foreground square appeared to oscillate in depth or appeared to be stationary. Six different types of shadow were used for the experiment: three "dark" shadows simulated as film transparencies with transmittances of 12, 16, and 36%; and three physically implausible "light" shadows corresponding to transmittances of 180, 284, and 394% (i.e. light was added within the shadow). The background checkerboard had a mean luminance of 17.4 cd / m^2 with an 82% contrast between dark and light squares. Subjects were split into two groups of ten. The order of presentation of different shadow conditions for one group, in terms of effective transmittance, was: 16, 284, 12, 394, 36, and 180%. The other group saw the stimuli in the order 284, 16, 394, 12, 180 and 36%. Each subject viewed three series of presentations, making a total of 18 trials. On 78% of the trials using dark shadows, observers reported seeing the foreground square as oscillating in depth--toward and away from the viewer. On only 40% of the trials using light shadows did subjects report seeing the square oscillating in depth (A Wilcoxon signed rank order test on the difference between light and dark shadows gave p= 0.001).

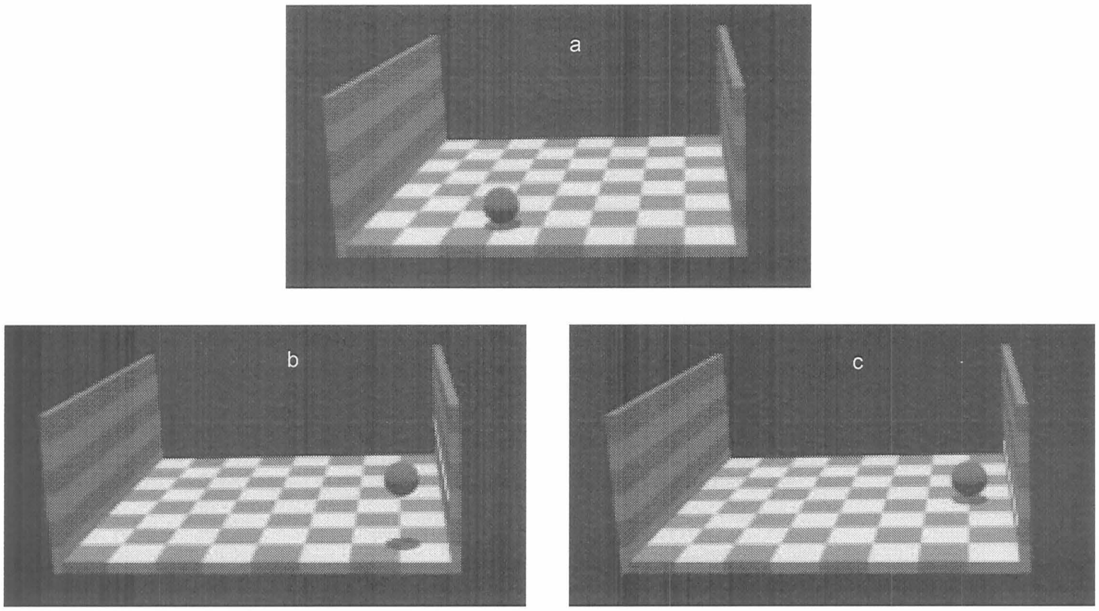

Fig. 3. Three frames from animations made with the ball-in-a-box simulation. In a simulated world, a ball was placed in a small 132 x 132 mm box and viewed from a point 355 mm from the center of the box with an elevation of 21.8˚ relative to the floor of the box. The viewpoint was offset slightly to the right, as shown. Each animation was created in two stages: first, we rendered a scene with a moving ball without cast shadows. Second, we independently added the ball's cast shadow to the images in an animation, so that we could manipulate the motion of the shadow independently of the ball's motion. The shading on the ball and in the room for all the animations, except those used in Experiment 3, was generated by simulating a light source at infinity with a slant of 63.4˚ degrees relative to the floor of the box. In Experiment 3, we manipulated the shading on the ball as an independent variable. In all the animations, the ball moved in a linear trajectory in the image at an angle tilted by 21.8˚ from the horizontal. Its velocity varied sinusoidally (period = 4 sec), so that the ball repeated its motion back and forth between its left- and right-most positions in the image. The shadow moved so that it remained vertically below the ball in the image. Only the distance between the shadow and the ball varied as the shadow and ball moved. The images shown here are copies of those used in the two animations for Demonstration 1. Figure 3a shows the left-most positions of the ball and shadow in both animations. Figure 3b shows the right-most positions in one of the animations and figure 3c shows the right-most position in the other. The demonstration animations were recorded on videotape, and observers were shown the taped animations. For the experiments (Experiments 2 and 3), however, the animations were shown on the screen of a Stardent GS2000 graphics computer. Subjects were given the task of adjusting a line along the right wall (shown in 3b and c) to match the apparent height of the middle of the ball at the right-most point of its trajectory. Subjects adjusted the height of the line by moving the computer's mouse and indicated a match by pressing the mouse button. The motion of the ball and its shadow continued throughout the course of a trial.

image, in this case that of a ball, remained fixed throughout the animation.

The first demonstration using this simulation (Demonstration 1) consisted of two different animation sequences: In the first, the ball's cast shadow followed a horizontal trajectory in the image (ending up at the position shown in figure 3b); in the second, it followed a diagonal trajectory identical to that of the ball's image (ending up at the position shown in figure 3c). Despite the fact that the ball's image remained the same size and had an identical trajectory in the image plane in both animations, all observers reported the striking percept of seeing the ball rise above the checkerboard floor when the shadow trajectory was horizontal, and recede smoothly in depth along the floor when the slope of the shadow trajectory matched that of the ball. Because the size of the ball's image remained fixed, it is clear that the apparent depth from the moving cast shadow was sufficient to override the constant size constraint in this experiment.

2.3 Demonstration 2: Apparent depth produced by cast shadows induces apparent size change.

If observers have an implicit perceptual assumption that objects do not change physical size, one would predict that when the slope of the shadow trajectory matched the ball, the ball would appear to grow in size as it recedes in depth. Indeed, several of our observers reported this perception. In Demonstration 2, everything was as with Demonstration 1, except that we tripled the length of the box in world coordinates (figure 4). For constant ball size, the image should decrease in size by about 50% if it were indeed receding to the back of the box. However, as before, the image of the ball was kept constant. The ball made a full excursion (in the image) from the lower left corner of the box to the upper right corner. All of our observers reported seeing the ball apparently inflating and shrinking when the trajectory of the shadow matched the ball, but remaining fixed in size when the shadow trajectory was horizontal. In another study, we explicitly varied the image size of the ball together with the shadow trajectory slope and found a non-linear integration of the two sources of information in the perception of the relative position of the ball (Mamassian, Kersten, and Knill, 1992).

2.4 Demonstration 3: Moving cast shadow can produce the illusion of a non-linear object trajectory.

A third demonstration (Demonstration 3) further shows the sophistication of human 3D motion perception from relative shadow motion. We modified the animations used for Demonstration 1 in the following way: the shadow was given a non-linear motion trajectory in which it initially touched the ball's image, moved towards the front of the box, at mid-trajectory returned to touch the ball's image, and then swung to the front again (see figure 5a). The ball's image moved in the same straight, diagonal trajectory as before. All observers reported seeing the ball as moving in a non-linear 3D trajectory in which the ball appeared to come forward, retreat in depth, and then come forward again, as it moved from left to right in the box. Moreover, the observers reported seeing a singularity, or bounce, in the path of the ball when the shadow touched the ball's image and changed direction. Observers saw the bounce despite the fact that the ball's motion in the image was smooth at that point

3.0 The Stationary Light Source Constraint

Like many other monocular cues, the relative displacement of an object's image and its cast shadow provides theoretically ambiguous information for spatial layout. In order to interpret the cues, the visual system must use other information about the scene and make prior assumptions about the world. Since cast shadow displacement is a function of both object position and light source position (figure 6), the visual system must make implicit assumptions, or inferences from image data, about the position of the light source creating the shadows in order to infer the spatial positions of the casting objects. In this section, we present experimental data and phenomenal demonstrations which reveal the nature of the information and prior assumptions about light source position which the visual system brings to bear on the interpretation of cast shadow motion..

For static images of objects with cast shadows, the visual system must either assume a single light source illuminating all the objects in a scene or estimate the positions of different light sources illuminating the different objects. The phenomenal demonstration in

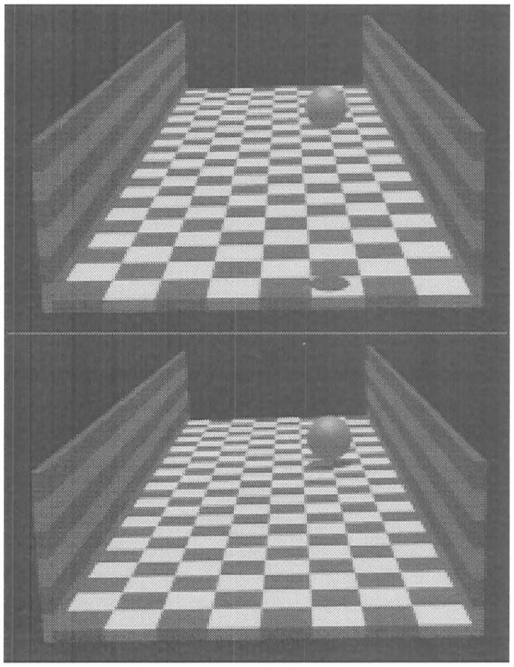

Fig. 4. The top and bottom panels show the extreme right position of the ball for the horizontal and diagonal shadow trajectories, respectively. In these static images, the effect of the shadow on the apparent size of the ball is small, but noticeable. In the dynamic case with diagonal trajectory, the ball has the striking appearance of inflating as it moves from left to right. For the horizontal trajectory, the ball appears to remain the same size.

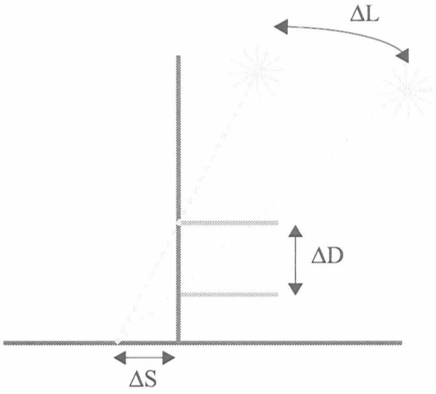

Fig. 6. A displacement ΔS between an object and its shadow can be produced either by a change in light source position, ΔL or by a change in depth of the object, ΔD.

DOCKET ALARM

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.

WHAT WILL YOU BUILD? | **sales@docketalarm.com** | **1-866-77-FASTCASE**

fastcase®
Smarter legal research.