

Cast Shadow Illusions

Daniel Kersten^{1,2} and Pascal Mamassian³

¹Department of Psychology, University of Minnesota

²Department of Brain and Cognitive Engineering, Korea University, Seoul 136-713, South Korea

³CNRS & Université Paris Descartes, France

Introduction

Shadows are everywhere, but they usually go unnoticed. We do, however, use their consequences all the time. Figure 1A shows two pictures each with a green square in front of a checkerboard background. The two pictures differ only in the cast shadows of the square. Because of these differences, the right square appears to be further from the background than the left square. By adding motion to the cast shadow, the effect on the perception of depth can be made quite striking. Moving the shadow, leaving the square fixed in the image, results in the square appearing to move in depth (see: <http://youtu.be/ig2CV1TLn5A>). It is only with scrutiny that one is convinced that the square is not moving at all within the picture frame. In particular, the square is not changing size, as it should if it was indeed moving in depth. Shadows seem to create a paradox for the visual system. While on the one hand they can be quite useful as cues for depth, their unobtrusiveness suggests that there may be good reasons for a visual system to get rid of them.¹

¹Cast shadows can be distinguished from “attached” or “form” shadows (see Mamassian et al., 1998; Castiello, 2001). An attached shadow on an object is the darkened surface that faces away from the light source. A cast shadow on a surface is the dark area created by an object blocking light onto that surface. A cast shadow can also be created when part of an object blocks light onto another part of the same object. For simplicity, in this chapter, the word “shadow” will usually refer to “cast shadow”.

To appear in: Oxford Compendium of Visual Illusions. Arthur G. Shapiro and Dejan Todorovic, Editors.

Comments may be sent to the author at kersten@umn.edu. D.K. was supported by the WCU (World Class University) program funded by the Ministry of Education, Science and Technology through the National Research Foundation of Korea (R31-10008).

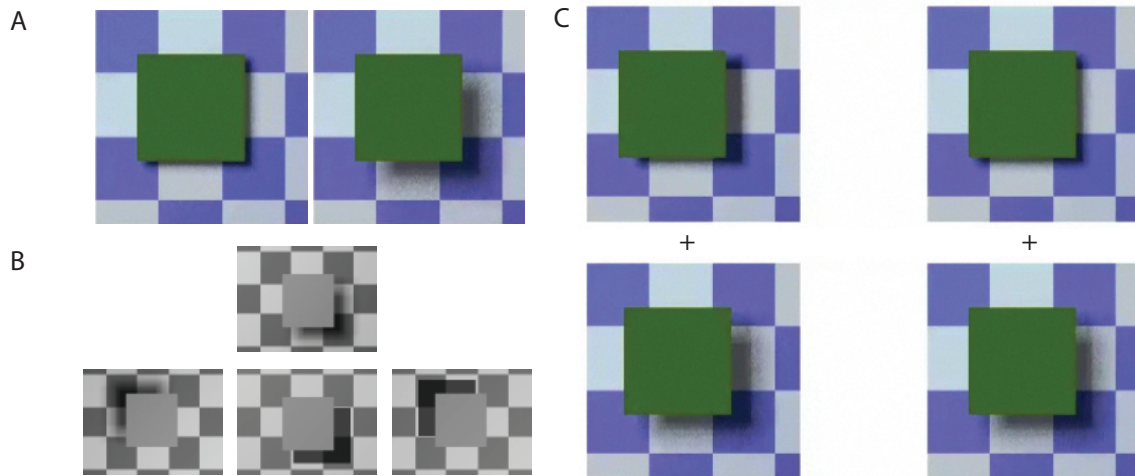


Figure 1. **A.** A cast shadow provides useful information for the depth of an object relative to the background surface. One can produce the appearance of motion in depth of the unmoving square by moving the shadow over a sequence of frames, see: <http://youtu.be/ig2CV1TLn5A>. **B.** Cast shadows are most effective at conveying depth when they have typical properties (top panel). If the shadows are above the object, or are sharp, they tend to be less effective (bottom three panels) (Kersten et al., 1996). **C.** Information about depth from shadows can be placed in conflict with information from stereo disparity. The stereo disparity provides information indicating that the top square is further from the background, whereas the shadow indicates it is the bottom square that is further away from the background. There are differences between people, and readers may be able to judge for themselves which cue wins. To experience the stereo effect, the reader should try to cross the two eyes so that the left and right images are seen by the right and left eyes, respectively. An animated version of this demonstration, with moving shadows, can be viewed at: <http://youtu.be/d4KMLLSs1C0>.

Shadows: Use them or lose them?

As visual animals, we typically need information about objects for the job at hand, whether that be describing, handling, navigating around, or tracking them. It would make sense that shadows go unnoticed because our attention is primarily directed towards the tangible objects we need rather than towards their shadows. Depending on the task, it might even make sense for the visual system to filter out shadows completely; for example, when they fall on the very object of interest, making it difficult to detect the features specific to the object (see Figure 2A; Cavanagh 1991; Lovell et al. 2009). The problem of detecting object-specific features is part of a more general problem that the brain must solve to reliably recognize objects. The images of an object (called its “appearances”) vary

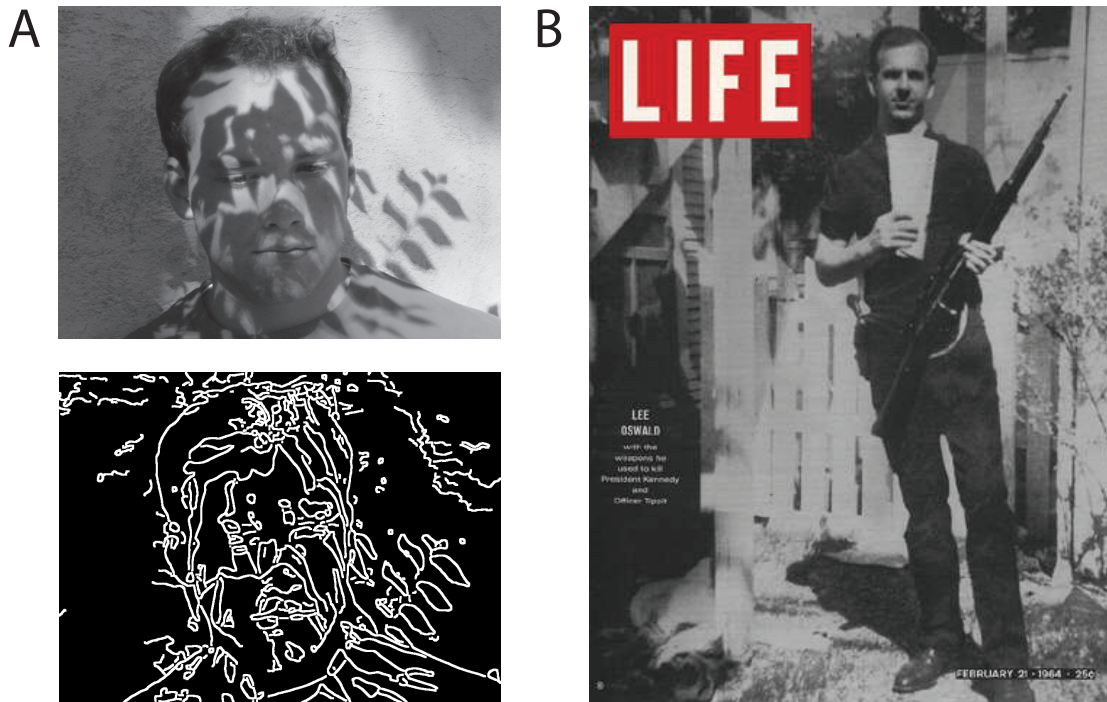


Figure 2. **A.** Shadows cast from other objects create problems for visual recognition. Upper and lower panels show the original gray level image and the output of a standard edge detector, respectively. While cast shadow edges can, in theory, provide some information regarding shape, they tend to be more trouble than they are worth. **B.** It isn't easy to tell whether cast shadows are physically correct, as illustrated by a photograph of Lee Harvey Oswald, the alleged assassin of John F. Kennedy (LIFE Magazine, February 21, 1964). Conspiracy theorists have argued that Oswald's claims that the photographs were doctored were supported by the supposed "fact" that shadows cast by his nose and body, were not consistent with a single light source (Farid, 2009).

greatly from one instance to the next because of changes in both viewpoint and illumination. Illumination gives rise to shading patterns that are affected by the object's orientation relative to the illumination sources and the viewpoint. And shadows, in particular, can have a strong effect on appearances.

Given the confounding nature of shadows, it may not be surprising that we seldom notice shadow inconsistencies in paintings or when artificially created in photographs (Mamassian, 2004; Ostrovsky et al., 2005; Cavanagh, 2005). Even with careful inspection we are not very good at deciding whether a cast shadow is faithful to the laws of optics or is faked (for example see the analysis of the Lee Harvey Oswald backyard pictures, Farid 2009; Figure 2B). Further, pictorial artists often eliminate shadows that might interfere with interpretation (Jacobson & Werner, 2004; Gombrich, 1995). Human visual recognition can suffer in the presence of shadows (Tarr et al., 1998; Braje et al., 2000, 1998), and considerable computer vision research has gone into detecting shadows in order to remove them (G. Finlayson et al., 2006; Jacobson & Werner, 2004). Does their confounding effect on recognition (Porter et al., 2010), together with the need to attend to objects imply that shadows are more trouble than they are worth, and that the visual system filters them out? Our illustration of depth from cast shadows suggests obviously not.

In fact, artists, photographers, animators, and lighting designers have long known that shadows have positive as well as negative consequences for perception—their judicious placement adds a sense of depth, realism, and evokes mood, whereas injudicious placement of even optically correct shadows can be misleading or distracting. So while shadows don't have a starring role to play on the stage of conscious perception, the illusory motion demonstration described above does suggest that they have a strong supporting role in conveying information about depth. What kinds of parts do shadows play in depth perception and just how strong are their roles? And what information does the human visual system extract in order to use shadows?

The perceptual use of shadows

Because of their role in providing depth information, shadows can be useful for diverse tasks such as deciding whether two surfaces are touching (Thompson et al., 1998), estimating the future trajectory of a flying object (Taya & Miura, 2010), determining the structural form of patterns (Khuu et al., 2012), deciding whether one surface occludes another (Tomonaga & Imura, 2010), and detecting differences in shape (Rensink & Cavanagh, 2004). Infants as young as 5-7 months use cast shadows for depth (Yonas & Granrud, 2006; Imura et al., 2006). Cast shadows improve perception of relative depth in chimpanzees (Imura & Tomonaga, 2009). Shadow features, unlike specular highlights, provide reliable surface features for stereo disparity, which in turn provides information for



Figure 3. Left panel. The absence of a strong true shadow of the man, together with a “good-enough” substitute—the dark spot on the road— produces apparent levitation. Right panel. The shadow of a flag appears to belong to the platform, creating the illusion that the speaker is elevated with no visible means of support.

3D shape (Puerta, 1989). (For a general review of the role of cast shadows in perception, see Dee & Santos (2011). For a historical and cultural treatment of shadows, see Casati 2004b.)

The role of cast shadows as a means to convey depth in drawings was recognized long ago by Leonardo da Vinci (Fiorani, 2008). One aspect of his advice—the ability to add “drop shadows”— is used everyday in graphics presentation and layout software (e.g. as a text format option in Microsoft’s *Powerpoint* and Apple’s *Keynote* presentation software). The strength of real-life drop shadows is illustrated below by putting them in conflict with our daily experience that objects shouldn’t defy gravity.

How strong are cast shadow cues for depth? A critical visual function is to know where an object is on a ground plane. Cast shadows provide pervasive and strong cues for contact between one surface and another, acting as a kind of “visual glue” (Thompson et al., 1998). The information lies in the proximity of the image of an object and its cast shadow, the sharpness of the penumbra, together with how local shadow and object contours meet (Madison et al., 2001). When there is no visual glue, the depth of an object becomes ambiguous. As an object rises above a plane, the shadow gets further from the

casting object, and the size of the penumbra increases.²

Apparent Levitation. Cast shadows are sufficiently powerful as a depth cue, that lack of contact information can inveigle perception to believing that objects or people can defy gravity, a trick used by street magicians. Here the magician raises one foot creating a displaced cast shadow while using the same foot to hide the fact that the other foot is holding his body up tiptoed (cf. http://en.wikipedia.org/wiki/Balducci_levitation).

Another trick is to take advantage of a gray day when the contact cue is weak, increasing depth ambiguity. Viewing from just the right angle allows a distant wet spot to assume the role of cast shadow, resulting in apparent levitation (left panel of Figure 3). Even on a bright day, an overhead sun can make an object's true shadow hard to see, reducing its role as visual glue, allowing a visually more salient shadow from another object to become associated with the object of interest making it appear to float mysteriously above the ground (right panel of Figure 3).

By creating variations of the green-square animation described earlier, one can show that the strength of the apparent change in depth is stronger if 1) the light source appears to come from above, 2) the cast shadow has a visible penumbra whose fuzziness grows with relative depth, and 3) the shadow is dark rather than light (which is physically impossible). These factors all suggest that the human visual system is sensitive to naturally and frequently occurring properties of cast shadows in images (see Kersten et al., 1996; Figure 1B). Recognizing these factors has practical consequences; for example, illumination from above the camera improves the utility of cast shadows in endoscopy (Mishra et al., 2004).

It is well-known that in addition to shadows, the human visual system uses multiple sources of depth information, including stereo disparity, motion parallax, relative brightness, and relative image size to produce estimates of relative depth (Bülthoff & Mallot, 1988; Landy et al., 1995; Cutting & Vishton, 1995). One can use computer graphics to place different sources of information, or cues, in opposition. For example, Figure 1C shows stereo disparity competing with cast shadows. Which cue wins depends on the weight an individual's visual system gives to the stereo over the shadow information.

In many cases, the percept experienced is the consequence of how these cues get weighted based on their respective reliabilities (Jacobs, 1999; Ernst & Bülthoff, 2004; Clark & Yuille, 1990). This suggests that as we make a visual cue increasingly strong and more

²The size of the penumbra also depends on the angular size of the light source. On a sunny day, the penumbra caused by the sun is quite small, the sun being far away, and only $1/2^\circ$ in diameter. On a cloudy day, the effective diameter approaches 180° , and penumbras are all but invisible except for small depth differences. And when there are multiple light sources, e.g. light filtering through the leaves of a tree, there can be multiple shadows that have characteristic transparency-like properties where they overlap.

diagnostic of a relative depth change (vs. some other cause), one could make visual illusions in which a stronger cue (e.g. the fuzziness of a shadow penumbra) overrides cues that say otherwise. We've already seen this with the green-square demonstration in that the size of the image of the green square doesn't increase as it approaches the viewpoint, as it normally would under perspective projection.

One way to increase the perceptual strength of shadows for depth is to move the object in the image, while keeping the illumination and viewpoint constant. The left top panel of Figure 4A shows the first frame of an animation sequence, and the middle and right panels show two alternative endpoints of the ball's shadow trajectory. While a comparison of the middle and right-hand panels provides some sense of depth differences, the animated versions are considerably more striking (see <http://youtu.be/hdFCJepvJXU>). When the shadow travels horizontally in the image, the ball appears to rise above the checkerboard floor, and when the shadow travels diagonally, the ball appears to recede in depth towards the back of the box. The ball's size, appearance, and trajectory are identical in the two animation conditions. The lack of a change in size of the image of the ball presents a cue conflict, which the visual system largely ignores. The bottom row of Figure 4A shows a similar sequence, except that a yellow, solid slab is substituted for the shadow. As long as the slab appears to stay on the floor, viewers typically see the ball rise above the floor for the animation whose shadow endpoint corresponds to the bottom, middle panel; but, remain on the floor for the end-point illustrated by the bottom right panel. One explanation for the perception of the different trajectories with the yellow slab is that human vision is sensitive to the correlated motion typical of objects and their shadows resulting from light sources that do not move. This is called the "stationary light source" constraint (Kersten et al., 1997). So even if the slab, ersatz shadow, lacks most of the properties of a shadow, the visual system is equipped to perceive different trajectories as a consequence of the differences in correlated motion. Other factors, such as matching sizes between casting object and shadow, also play a role (Ni et al., 2004), but not as much as the correlated motion.

Another demonstration of the effect of the stationary light source constraint, and its impact on correlated motion, is illustrated in Figure 4B. The top row shows samples of several successive frames in which a ball and its shadow alternately accelerate and decelerate consistent with fixed light source from above. When all frames are shown in a movie, the ball appears to bounce (see <http://youtu.be/6Wbbx6Zi3g4>). However, if the ball has constant velocity (samples shown in bottom row of Figure 4B), the perceived coupling between the ball and its shadow is lost, the shadow seems to have a life of its own, and the ball no longer appears to bounce.




Figure 4. **A.** The images in the left column show the first frame, and the middle and right frames show alternative endpoints of the ball's shadow. The trajectory of the ball is the same in all trajectories. Top panel. When the trajectory of the shadow is horizontal, the ball appears to rise above the floor; when the shadow trajectory is diagonal, the ball appears to travel towards the back of the box. See <http://youtu.be/hdFCJepvJXU>. The perception of the different trajectories in depth relies on the visual system's ability to detect correlated motion (consistent with a stationary light source). Bottom panel: If the shadow is replaced by a yellow slab, the apparent difference in trajectories can still be seen. In addition to correlated motion, treating the slab as a substitute shadow relies on perceiving the slab to be on the floor. The coincidental alignment of one of the primary axes of the slab with that of the floor encourages this. See: <http://youtu.be/TGX7153-Ccw>. Figure from Kersten et al. (1997). **B.** See <http://youtu.be/6Wbbx6Zi3g4> and main text for an explanation.

The future of shadows

Despite their ubiquity, there are substantial computational, perceptual, and neuroscience challenges to our understanding of the visual processing of shadows.

Computational challenges. While the potential usefulness of cast shadows in computer vision was recognized many years ago (Waltz, 1972), the application to depth estimation has largely been neglected. This stands in contrast to substantial progress in understanding and using other cues such as stereo disparity and motion parallax. To use cast shadows, a visual system must effectively segment a scene into objects and shadows, determine which object (or object part) goes with which shadow (or shadow part; i.e., the shadow “correspondence problem”, Mamassian 2004; Ni et al. 2004; Casati 2008), and use this information in the context of the overall spatial layout (Figure 5). While there have been a number of analyses of the information provided by shadows (Casati, 2004a; Dee & Santos, 2011; Knill et al., 1997), the computational challenge will be to turn these observations into useful algorithms.

Statistical constraints on shadow photometry can be used for shadow recognition. For example, the color change across a shadow boundary is largely a change in luminance rather than chromaticity (Kingdom et al., 2004; G. D. Finlayson et al., 2009). Similarly, texture statistics tend to remain similar across shadow boundaries (Zhu et al., 2010). The global shape of shadow contours have statistical regularities consistent with the shadow patch laying on the surface receiving it. More specifically, shadow shapes follow the geometry of double-projection: the shadow on a surface is the projection of the casting object by the light source on to an oriented surface, and shadow’s image is the result of a projection of the surface shadow onto the retina. One consequence of this is that it is not unusual to have a long thin shadow in an image cast by a compact object (e.g. with the highly oblique sun angles during sunrise or sunset). In contrast, a long, thin image is less likely due to a compact object in the world (Brady & Yuille, 1984).

Perceptual challenges. A major problem in visual perception is understanding how the visual system combines multiple cues for a task such as depth perception. One problem in combining shadows with other depth cues is the commensurability of their information. Some cues, like “looming”, the expansion of the size of an object’s image on the retina, provide the visual system with information for distance between an object and the observer. Rather than depth from the observer, cast shadows more directly provide an estimate of the distance between two objects (the object casting the shadow and the receiving surface; Schrater & Kersten, 2000). In addition, this distance is dependent on usually invisible factors like the light source position. Thus, shadows should *prima facie* not be

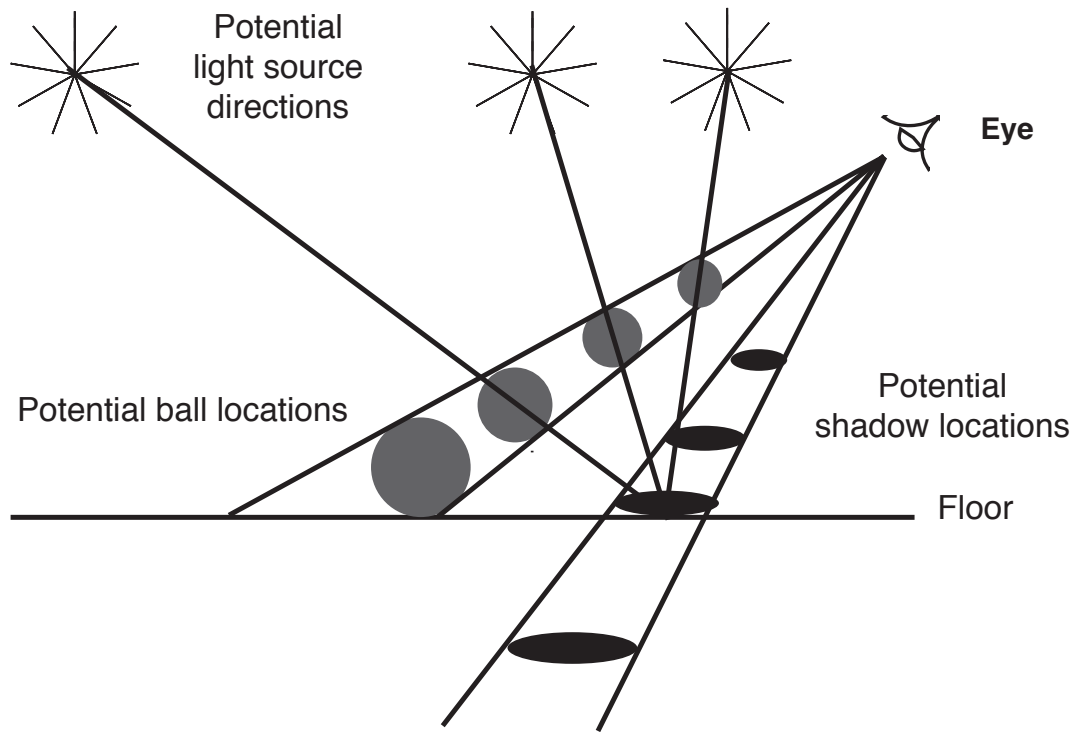


Figure 5. The human visual system solves several problems when estimating depth from cast shadows. Even if the “patches” representing an object and shadow are identified and paired, in order to determine relative depth, their locations need to be localized in an appropriate frame of reference (Schrater & Kersten, 2000). For example, while it is obvious that the shadow should be on the ground, an algorithm needs also to determine the ground plane. From (Kersten, 1997).

considered as a metric depth cue like binocular disparity, and this makes it more difficult to understand how these two cues can interact (e.g. Figure 1C). However, it is conceivable that, given prior assumptions on the statistical structure of the environment, shadows could be “promoted” to a metric depth cue. Such a reliance on the statistical structure of natural scenes has been shown to promote the figure-ground convexity cue (an object with a convex boundary tends to be seen in front) to a full metric depth cue (Burge et al., 2010).

Neuroscience challenges. The dual nature of shadows reflects, in large part, two different functional requirements of vision. A basic architectural aspect of the visual brain is the division into two main pathways: the ventral (or temporal) and dorsal (or parietal) streams, believed to be specialized for processing intrinsic object properties, and viewer-

object spatial relationships, respectively (Milner & Goodale, 2006). As discussed above, on the one hand, to recognize objects, the visual system needs to filter out image features, including those due to shadows, that are not intrinsic to the object. On the other hand, to determine where objects are, we've seen that the visual system uses cast shadows in addition to other depth cues. Bonfiglioli et al. (2004) argue that this functional distinction underlies their finding that incongruent shadows had a significant impact on motor task, but not on recognition performance. However, it has also been shown that cast shadows can be relatively ineffective in the online control of pointing movements in depth (Hu & Knill, 2011). Consistent with a role for the dorsal stream in processing depth from shadows, neuroimaging results have found that the posterior portion of the medial parietal cortex is activated when observers perceive depth from cast shadows (Katsuyama et al., 2011).

Shadows have been studied for millennia, yet there are always new and fascinating observations to be made (Casati, 2012, 2010). Shadow illusions, in particular, raise significant questions about how the brain processes them. Attempting to answer these questions will insure that shadows continue to captivate artists and scientists alike for years to come.

References

- Bonfiglioli, C., Pavani, F., & Castiello, U. (2004). Differential effects of cast shadows on perception and action. *Perception*, *33*(11), 1291–1304.
- Brady, M., & Yuille, A. (1984, March). An extremum principle for shape from contour. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, *6*(3), 288–301.
- Braje, W. L., Kersten, D., Tarr, M. J., & Troje, N. F. (1998). Illumination Effects in Face Recognition. *Psychobiology*, *26*(4), 371–380.
- Braje, W. L., Legge, G. E., & Kersten, D. (2000). Invariant recognition of natural objects in the presence of shadows. *Perception*, *29*(4), 383–398.
- Bülthoff, H. H., & Mallot, H. A. (1988, October). Integration of depth modules: stereo and shading. *Journal of the Optical Society of America A*, *5*(10), 1749–1758.
- Burge, J., Fowlkes, C. C., & Banks, M. S. (2010, May). Natural-Scene Statistics Predict How the Figure-Ground Cue of Convexity Affects Human Depth Perception. *Journal of Neuroscience*, *30*(21), 7269–7280.
- Casati, R. (2004a). The shadow knows: a primer on the informational structure of cast shadows. *Perception*, *33*(1), 1385–1396.
- Casati, R. (2004b). *Shadows: Unlocking Their Secrets, from Plato to Our Time*. Vintage.
- Casati, R. (2008). The copycat solution to the shadow correspondence problem. *Perception*, *37*(4), 495–503.

- Casati, R. (2010). Impossible Cast Shadows in Ukyo-e Paintings. *CARLS Series of Advanced Study of Logic and Sensibility, Vol. 3*, 265–268.
- Casati, R. (2012). Some varieties of shadow illusions: Split shadows, occluded shadows, stolen shadows, and shadows of shadows. *Perception, 41*(3), 357–360.
- Castiello, U. (2001, August). Implicit processing of shadows. *Vision Research, 41*(18), 2305–2309.
- Cavanagh, P. (1991). What's up in top-down processing. In A. Gorea (Ed.), (pp. 295–304). *Representations of vision: Trends and tacit assumptions in vision research*.
- Cavanagh, P. (2005, March). The artist as neuroscientist. *Nature, 434*(7031), 301–307.
- Clark, J. J., & Yuille, A. L. (1990). *Data Fusion for Sensory Information Processing Systems*. Springer.
- Cutting, J. E., & Vishton, P. M. (1995, June). Perceiving layout and knowing distances: The integration, relative potency, and contextual use of different information about depth. In W. Epstein & S. Rogers (Eds.), *Handbook of perception and cognition* (pp. 69–117). San Diego: Academic Press.
- Dee, H. M., & Santos, P. E. (2011, August). The Perception and Content of Cast Shadows: An Interdisciplinary Review. *Spatial Cognition & Computation, 11*(3), 226–253.
- Ernst, M. O., & Bühlhoff, H. H. (2004, April). Merging the senses into a robust percept. *Trends in Cognitive Sciences, 8*(4), 162–169.
- Farid, H. (2009). The Lee Harvey Oswald backyard photos: real or fake? *Perception, 38*(11), 1731–1734.
- Finlayson, G., Hordley, S., & Drew, M. (2006). Removing shadows from images. *Computer Vision—ECCV 2002*, 129–132.
- Finlayson, G. D., Drew, M. S., & Lu, C. (2009, May). Entropy Minimization for Shadow Removal. *International Journal of Computer Vision, 85*(1), 35–57.
- Fiorani, F. (2008). The colors of Leonardo's shadows. *Leonardo, 41*(3), 271–278.
- Gombrich, E. E. H. (1995). *Shadows : the depiction of cast shadows in Western art*. London : National Gallery Publications ; [New Haven] : Distributed by Yale University Press.
- Hu, B., & Knill, D. C. (2011, June). Binocular and monocular depth cues in online feedback control of 3D pointing movement. *Journal of Vision, 11*(7), 23–23.
- Imura, T., & Tomonaga, M. (2009). Moving shadows contribute to the corridor illusion in a chimpanzee (*Pan troglodytes*). *Journal of Comparative Psychology, 123*(3), 280–286.

- Imura, T., Yamaguchi, M. K., Kanazawa, S., Shirai, N., Otsuka, Y., Tomonaga, M., et al. (2006, March). Perception of motion trajectory of object from the moving cast shadow in infants. *Vision Research*, *46*(5), 652–657.
- Jacobs, R. (1999, October). Optimal integration of texture and motion cues to depth. *Vision Research*, *39*(21), 3621–3629.
- Jacobson, J., & Werner, S. (2004). Why cast shadows are expendable: Insensitivity of human observers and the inherent ambiguity of cast shadows in pictorial art. *Perception*, *33*(11), 1369–1383.
- Katsuyama, N., Usui, N., Nose, I., & Taira, M. (2011, January). Perception of object motion in three-dimensional space induced by cast shadows. *NeuroImage*, *54*(1), 485–494.
- Kersten, D. (1997, August). Perceptual Categories for Spatial Layout. *Philosophical Transactions: Biological Sciences*, *352*(1358), 1155–1163.
- Kersten, D., Knill, D. C., Mamassian, P., & Bühlhoff, I. (1996, January). Illusory motion from shadows. *Nature*, *379*(6560), 31.
- Kersten, D., Mamassian, P., & Knill, D. C. (1997). Moving cast shadows induce apparent motion in depth. *PERCEPTION-LONDON-*, *26*(2), 171–192.
- Khuu, S., Khambiye, S., & Phu, J. (2012, October). Detecting the structural form of cast shadows patterns. *Journal of Vision*, *12*(11), 10–10.
- Kingdom, F. A. A., Beauce, C., & Hunter, L. (2004). Colour vision brings clarity to shadows. *Perception*, *33*(8), 907–914.
- Knill, D. C., Mamassian, P., & Kersten, D. (1997). Geometry of shadows. *JOSA A*, *14*(12), 3216–3232.
- Landy, M. S., Maloney, L. T., Johnston, E. B., & Young, M. (1995). Measurement and modeling of depth cue combination: In defense of weak fusion. *Vision Research*, *35*(3), 389–412.
- Lovell, P. G., Gilchrist, I. D., Tolhurst, D. J., & Troscianko, T. (2009). Search for gross illumination discrepancies in images of natural objects. *Journal of Vision*, *9*(1), 37.1–14.
- Madison, C., Thompson, W., Kersten, D., Shirley, P., & Smits, B. (2001, February). Use of interreflection and shadow for surface contact. *Percept Psychophys*, *63*(2), 187–194.
- Mamassian, P. (2004). Impossible shadows and the shadow correspondence problem. *Perception*, *33*(11), 1279–1290.
- Mamassian, P., Knill, D. C., & Kersten, D. (1998). The perception of cast shadows. *Trends in Cognitive Sciences*.

- Milner, D., & Goodale, M. (2006). *The Visual Brain in Action (Oxford Psychology Series)* (2nd ed.). Oxford University Press, USA.
- Mishra, R. K., Hanna, G. B., Brown, S. I., & Cuschieri, A. (2004, August). Optimum shadow-casting illumination for endoscopic task performance. *Archives of surgery (Chicago, Ill. : 1960)*, *139*(8), 889–892.
- Ni, R., Braunstein, M. L., & Andersen, G. J. (2004). Perception of scene layout from optical contact, shadows, and motion. *Perception*, *33*(11), 1305–1318.
- Ostrovsky, Y., Cavanagh, P., & Sinha, P. (2005). Perceiving illumination inconsistencies in scenes. *PERCEPTION-LONDON-*, *34*(11), 1301–1314.
- Porter, G., Tales, A., & Leonards, U. (2010). What makes cast shadows hard to see? *Journal of Vision*, *10*(3), 1–18.
- Puerta, A. M. (1989). The power of shadows: shadow stereopsis. *JOSA A*, *6*(2), 309–311.
- Rensink, R. A., & Cavanagh, P. (2004). The influence of cast shadows on visual search. *Perception*, *33*(11), 1339–1358.
- Schrater, P., & Kersten, D. (2000). How optimal depth cue integration depends on the task. *International Journal of Computer Vision*, *40*(1), 71–89.
- Tarr, M., Kersten, D., & Bülthoff, H. H. (1998, August). Why the visual recognition system might encode the effects of illumination. *Vision Research*, *38*(15-16), 2259–2275.
- Taya, S., & Miura, K. (2010, October). Cast shadow can modulate the judged final position of a moving target. *Attention, Perception, & Psychophysics*, *72*(7), 1930–1937.
- Thompson, W. B., Shirley, P., Smits, B., Kersten, D. J., & Madison, C. (1998). Visual glue. *University of Utah Technical Report UUCS-98-007*.
- Tomonaga, M., & Imura, T. (2010). Pacman in the sky with shadows: the effect of cast shadows on the perceptual completion of occluded figures by chimpanzees and humans. *Behavioral and Brain Functions*, *6*(1), 38.
- Waltz, D. L. (1972). *Generating semantic descriptions from drawings of scenes with shadows* (Tech. Rep.).
- Yonas, A., & Granrud, C. E. (2006). Infants' perception of depth from cast shadows. *Attention, Perception, & Psychophysics*, *68*(1), 154–160.
- Zhu, J., Samuel, K. G., Masood, S. Z., & Tappen, M. F. (2010). Learning to recognize shadows in monochromatic natural images. In *Computer vision and pattern recognition (cvpr), 2010 ieee* (pp. 223–230). IEEE.