

Change blindness for cast shadows in natural scenes: Even informative shadow changes are missed

Krista A. Ehinger^{1,2} · Kala Allen³ · Jeremy M. Wolfe^{1,2}

Published online: 4 February 2016
© The Psychonomic Society, Inc. 2016

Abstract Previous work has shown that human observers discount or neglect cast shadows in natural and artificial scenes across a range of visual tasks. This is a reasonable strategy for a visual system designed to recognize objects under a range of lighting conditions, since cast shadows are not intrinsic properties of the scene—they look different (or disappear entirely) under different lighting conditions. However, cast shadows can convey useful information about the three-dimensional shapes of objects and their spatial relations. In this study, we investigated how well people detect changes to cast shadows, presented in natural scenes in a change blindness paradigm, and whether shadow changes that imply the movement or disappearance of an object are more easily noticed than shadow changes that imply a change in lighting. In Experiment 1, a critical object's shadow was removed, rotated to another direction, or shifted down to suggest that the object was floating. All of these shadow changes were noticed less often than changes to physical objects or surfaces in the scene, and there was no difference in the detection rates for the three types of changes. In Experiment 2, the shadows of visible or occluded objects were removed from the scenes. Although removing the cast shadow of an occluded object could be seen as an object deletion, both types of shadow changes were noticed less often than deletions of the visible, physical objects in the scene. These results show that even

informative shadow changes are missed, suggesting that cast shadows are discounted fairly early in the processing of natural scenes.

Keywords Change blindness · Attention · Scene perception · Shadows

Shadows are ubiquitous features of everyday scenes, but we rarely give them much attention, except in some rare circumstances. For example, we may notice a particularly dramatic long shadow at sunset, or actively search for shadows when looking for a shady spot to sit on a hot day. But in general we have little reason to attend to shadows in daily life, and we don't seem to encode much information about shadows: For example, we are generally poor at recognizing impossible shadows in scenes (Ostrovsky, Cavanagh, & Sinha, 2005), and worse at detecting changes to shadows than at detecting changes to objects (Wright, 2005). However, most of this previous work had looked at scenes in which the shadow information was largely redundant, and changing a shadow would not have changed anything about the meaning of the scene. Here we looked at whether people neglect shadow changes that do affect the meaning of the scene—for example, by changing the number of objects or their spatial relations.

In these experiments, we focused on *cast shadows*, which are the shadows that light-occluding objects cast onto an external surface, such as the ground. These should be distinguished from *attached shadows*, which are the gradations in shading on the surface of an object, and *self shadows*, which are shadows cast by one part of an object onto another part of its own surface. All shadows are produced when light emanating from a source interacts with objects in its path. The

✉ Krista A. Ehinger
k.a.ehinger@gmail.com

¹ Harvard Medical School, Visual Attention Lab, 64 Sidney St, Suite 170, Boston, MA, USA

² Brigham & Women's Hospital, Boston, MA, USA

³ Miami University, Oxford, OH, USA

brightness (or *luminance*) that we perceive when looking at an object or surface is the product of the amount of light hitting that surface (*illuminance*) and the actual color of the surface (*reflectance*).

In order to recognize objects under a variety of lighting conditions, a visual system must be able to subtract out the effects of illumination in order to perceive the true surface colors of objects (Adelson, 2000; DiCarlo & Cox, 2007; Land & McCann, 1971; Rolls & Stringer, 2006). Cast shadows, in particular, should be discarded because they are produced by a change in illumination only: The surface appears darker because it receives less light, not because there is any actual difference in the surface reflectance. Experiments have shown that people discount shadow information across a variety of visual tasks, although there is some debate about the mechanisms involved in this discounting process. People are poor at identifying lighting inconsistencies in complex object arrays or detecting incorrect cast shadows in natural scenes (Ostrovsky et al., 2005). In change blindness tasks, people are slower to detect a disappearing cast shadow than a disappearing object (Wright, 2005). People seem to discount cast shadows in search tasks, in that they are slower to detect an odd shadow in an array than an equivalent odd object (Rensink & Cavanagh, 2004). This could be taken as evidence for an early visual processing stage that identifies cast shadows in images and corrects for them, but later experiments have suggested that cast shadows are retained in early vision and that the discounting occurs later, as part of the object recognition process (Porter, Tales, & Leonards, 2010). Furthermore, other studies have shown rapid, “pop-out” search for odd shadows that suggest a different lighting direction or a change in object depth, which provides more evidence that cast shadow information is available in early visual processing (Elder et al., 2004). The presence or absence of cast shadows seems to have no effect on object recognition from photographs (Braje, Legge, & Kersten, 2000), but later experiments using a more varied set of computer-generated objects showed that recognition is slightly slowed when cast shadows are absent or incorrect (Castiello, 2001). Experiments using novel objects have revealed a small recognition penalty when the cast shadow information changes between learning and test (Leek, Davitt, & Cristino, 2015; Tarr, Kersten, & Bühlhoff, 1998). This suggests that some cast shadow information is retained during object processing, though it may only be useful for identifying unfamiliar, artificial objects.

Although cast shadows give incorrect information about the surface on which they are cast, they can give very useful information about the casting objects. In particular, cast shadows can provide 3-D information that is not otherwise available in the scene, and can be used to disambiguate the positions of objects in depth (Mamassian, Knill, & Kersten, 1998). For example, cast shadows can be used to determine

whether an object is resting on a surface or floating above it, or to determine which of two ambiguous surfaces is supporting the other—assuming light from above, an object can cast a shadow onto a supporting surface, but not vice versa. People’s perception of the depth and spatial position of objects in scenes seems to be very dependent on cast shadows, and people perceive the movement of an object’s cast shadow as movement of that object in depth, even when other cues in the scene contradict that interpretation (Kersten, Mamassian, & Knill, 1997).

Cast shadows are potentially useful to the visual system because they provide a second, 2-D projection of the scene (Casati, 2004; Dee & Santos, 2011). They depict the shapes of objects as “seen” from the point of view of the light source, which must necessarily be a different point of view than the one observed by the eye (Vinci, 1888). A second projection of the scene can provide useful 3-D information, which is why shadows are so valuable for disambiguating object depth and establishing spatial relationships between objects and surfaces. In this sense, shadows are somewhat similar to reflections—they provide incorrect information about the surface on which they are cast (i.e., a shadowed surface is not really dark, and a mirror does not actually contain depth), but provide a second view of the scene that may contain information that is not otherwise visible.

In this study, we investigated the discounting of shadows in natural scenes using a change blindness paradigm similar to one previously used to study reflections in natural scenes (Sareen, Ehinger, & Wolfe, 2015). The experiments make use of the “flicker paradigm” of Rensink, O’Regan, and Clark (1997), in which two versions of a scene alternate with a brief blank period between presentations. The observer attempts to locate the difference between the two frames. Under these conditions, even quite large changes in a scene can go unnoticed for many seconds. For change to be noticed, the item being changed must be attended before and after the change. With very simple displays like arrays of colored squares, observers are able to keep track of the state of three or four items from frame to frame (Luck & Vogel, 1997, 2013). With scenes, typically many more items could be the focus of attention, so change blindness becomes a way of assessing what attracts attention in a scene, at least when the observer is looking for a change. In detecting change, low-level salience does not appear to be as important as the meaning of items in a scene (Stirk & Underwood, 2007), and changes in the existence or position of an object are more readily detected than changes in surface properties such as color (Aginsky & Tarr, 2000). Wright (2005) looked at change blindness for the deletion of cast shadows and found that observers were not good at detecting them. Apparently, having a shadow go out of existence is not like having an object go out of existence. People may be particularly insensitive to this type of shadow change because it doesn’t change the

meaning of the scene. It does add a lighting inconsistency, but people are poor at recognizing lighting inconsistencies in natural scenes (Ostrovsky et al., 2005). In the present study, we considered cases in which a change to a cast shadow affected the spatial relations of objects in the scene (Exp. 1) or effectively deleted an implied object from the scene (Exp. 2). We compared these types of informative shadow changes to object changes and to less informative shadow changes, such as the deletion or rotation of cast shadows, which add a lighting inconsistency but do not alter the meaning of the scene.

Experiment 1

We first looked at whether participants had different detection rates for different types of cast shadow changes in scenes. We compared two types of shadow changes that should not significantly affect the meaning of the scene—shadow deletion and rotation—to a change that arguably does change the gist of the scene—shifting a cast shadow downward to suggest that the casting object is floating. This turns an ordinary scene into an impossible scene with an object that violates the laws of physics. These shadow manipulations were compared to two control conditions in which changes occurred to objects or surfaces in the scene.

Methods

Participants A total of 21 people participated in an online experiment on Amazon’s Mechanical Turk. The participants were based in the U.S. and had a good track record on the Mechanical Turk site (at least 100 HITs completed and an acceptance rate of at least 95%). All participants gave informed consent before starting the task. Participants were paid \$1.00 upon completing the task and an additional \$0.10 for each correct response, resulting in an average total payment of about \$14.00.

Stimuli The stimulus images were photographs of everyday scenes. They were 24-bit color and scaled to a resolution of 1,024 pixels wide by 768 pixels high (a few images had a different aspect ratio but were scaled to the maximum possible size within this box). Sixty images were used in the experimental trials, and an additional 150 images were used as “fillers” to disguise the fact that the main experimental manipulation involved shadows.

We created a second version of each image by changing one object or region using Photoshop. There were five change conditions for the experimental images; each image appeared in all five conditions, but an individual participant would only see one of the five versions. The five conditions for a single example image are shown in Fig. 1. On *object change* trials, a random object in the scene would either

disappear or change color (each type of change occurred on half of the experimental trials). On *shadow deletion* trials, one object’s shadow would disappear. On *shadow rotation* trials, the same shadow would be deleted and redrawn as though the object were lit from another direction. On *floating shadow* trials, the shadow was shifted downward to suggest that the object was floating above its supporting surface. On *surface change* trials, a patch with the same size, shape, and color values as the original shadow was added to a flat surface in the scene. This patch was located away from objects so it would not be mistaken for a cast shadow, and was blurred to match the original shadow’s penumbra. The images used in this experiment (and in Exp. 2) are available online as part of the Change Blindness (CB) database at <http://search.bwh.harvard.edu/new/CBDatabase.html> (Sareen, Ehinger, & Wolfe, 2016).

Design Each participant saw six images from each of the five experimental conditions; the image selection was random and counterbalanced across subjects, so each image occurred equally often in each condition across the experiment. These 30 experimental trials were randomly interleaved with 150 filler trials. The filler trials were the same for every participant and included 114 trials with object changes (half were object deletions, and the other half were object color changes) and 36 trials with no change. The trial order was randomized for each participant. Across the experiment, 80% of trials had a change, and 16.67% of those changes involved shadows.

Procedure Participants were told that the experiment would involve finding changes in scenes, and they were given an example of a change blindness trial with an object deletion. The images were presented in a Web browser using jsPsych (de Leeuw, 2015). On each trial, participants saw a pair of images, presented alternately one after the other for 1 s each, with a 250-ms blank in between. On each trial, the starting image was chosen at random. This sequence repeated 24 times (1 min) or until the participant pressed a response key on the keyboard: “Y” if the participant saw a change, or “N” if the participant thought there was no change. Keypress responses and response times were recorded.

Next, participants were shown one of the two images from the change blindness sequence. This probe image was always the original image from the pair (the one with the original shadow, or the original object present), except on surface change trials, which showed the modified image with the extra surface feature present. Participants were asked to click on the location of the change. If participants had not detected a change, they were asked to skip this step. There was no time limit on this step, and response times were not recorded.

After each trial, participants were shown the number of trials remaining. They could press a key to start the next trial, or quit the experiment and return to it later. Participants were

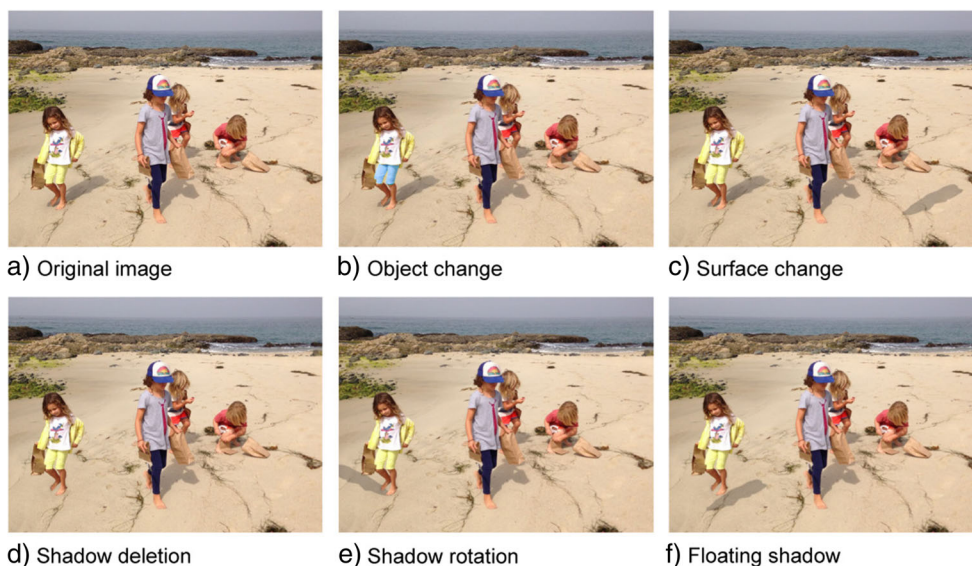


Fig. 1 An example image (a) and the change conditions (b–f) used in Experiment 1. In the flicker change blindness paradigm, the original image is shown alternately with one of these modified versions, with a brief blank screen in between. The changes in this image are: (b) the girl’s

shorts change color; (c) a dark patch appears in the sand near the bottom right; (d) the girl’s shadow is deleted; (e) the girl’s shadow changes direction; and (f) the girl’s shadow is moved so that she seems to float above the sand

encouraged to complete the trials in their own time and to take breaks whenever they wished. There was no feedback after the trials.

Results

One participant who detected only a single change during the experiment was dropped and replaced. The data were screened for very short response times (<250 ms) and very long response times (>3 SDs above the mean response time, not including trials that timed out), but no trials were dropped due to these criteria.

We compared accuracy, response times, and the numbers of timed-out trials across the five experimental conditions, shown in Fig. 2. Experimental trials were marked “correct”

if the participant responded “yes” and correctly marked the location of the change, and “incorrect” if the participant responded “no” or failed to mark the change location, or if the trial timed out with no response. A one-way repeated measures analysis of variance (ANOVA) showed a significant effect of condition on accuracy [$F(4, 19) = 10.89, p < .001$]. Bonferroni-corrected pairwise *t* tests showed significant differences between each of the two most accurate conditions (object change and surface change) and each of the three least accurate conditions (shadow deletion, shadow rotation, and floating shadow), but no significant differences within either of these groups. A similar analysis was performed to compare response times on correct trials (four participants were dropped from this analysis because they had no correct trials in at least one condition). A one-way ANOVA showed no




Fig. 2 Response times (left), accuracy (center), and percentages of trials that timed out across the five types of changes in Experiment 1. Error bars represent standard errors of the means

significant difference in response times across conditions [$F(4, 15) = 1.88, p = .125$]. We also observed no significant difference in number of trials that timed out across the experimental conditions [$F(4, 19) = 0.41, p = .802$].

Most of the errors were failures to detect the changes. These were trials on which participants either said that no change had occurred or the trial timed out before the change was detected. Only 2.6% of the errors (28 trials) were location errors, in which people said there was a change but clicked on a wrong object or region. We checked whether participants had misinterpreted the shadow change as an object change in any trials and clicked the object, but there were no errors of this type. The low incidence of location errors is probably due to the reward scheme in our task: Participants received bonuses for correct localization or correctly rejecting a trial as a “no change” trial. There was no bonus for correctly detecting a change but marking the wrong location, which probably encouraged participants to guess “no change” if they thought there was a change but could not identify what was changing.

We also looked at whether any effects emerged of the number of image pixels affected by the change or the spatial location of the changes across the five experimental conditions. Although we tried to keep the sizes and eccentricities of the changes similar across the conditions, these could not be perfectly matched in natural scenes. For each condition, we created a mask image in which the changing pixels were set to 1 and all other pixels in the image were set to 0. For the conditions with moving shadows (shadow rotation and floating shadow), the mask included all pixels that were part of the shadow in one version of the changing image but not the other. To measure the area of the image that had changed in each condition, we computed the number of pixels in each mask. To measure the eccentricity of the changing region, we computed its center of mass from the mask and took the distance in pixels from that point to the center of the image. The average areas and eccentricities of changes in the experimental conditions are shown in Fig. 3. One-way repeated measures ANOVAs showed significant differences in both the area [$F(4, 59) = 13.48, p < .001$] and eccentricity [$F(4, 59) = 3.48, p < .01$] of changes across the experimental conditions. However, these differences do not seem to predict the observed differences in accuracy: The object and surface changes were detected most frequently, but were generally smaller in area than changes that were detected less frequently (shadow rotation and floating shadow). The object changes were also, on average, less central than changes in the other conditions.

Twelve participants (60%) mentioned that they had noticed or been looking for shadow changes when they were debriefed at the end of the experiment. Although none of them thought the experiment was primarily about shadow perception, this does suggest that at least some participants may have been paying extra attention to shadows. Figure 4 shows the

accuracies across conditions for the participants who mentioned shadow changes in their debriefing versus those who did not. The former group seems to have noticed more of the shadow changes, but both subsets of participants showed a significant main effect of condition on accuracy in separate ANOVA analyses [among participants who reported looking for shadows, $F(4, 44) = 3.65, p < .05$; among participants who did not, $F(4, 28) = 9.88, p < .01$]. This suggests that even participants who were aware of or actively looking for shadow changes showed some effect of the experimental manipulation.

Experiment 2

In a second experiment, we looked at whether cast shadow deletions that imply an object deletion are more noticeable than shadow deletions that do not affect objects.

Method

Participants A total of 24 people participated in the online experiment on Amazon’s Mechanical Turk. Participants were based in the U.S. and had a good track record on the Mechanical Turk site (at least 100 HITs completed and an acceptance rate of at least 95%). All participants gave informed consent before starting the task. Participants were paid \$1.00 upon completing the task and an additional \$0.10 for each correct response, resulting in an average total payment of about \$14.00.

Stimuli In all, 36 new stimulus images were created for Experiment 2. As in Experiment 1, the stimulus images were photographs of everyday scenes in 24-bit color with a resolution of 1,024 pixels wide by 768 pixels high. Three variations of each image were created in Photoshop; these are shown in Fig. 5. For the *object deletion* trials, a single object was removed from the scene. For the *shadow deletion* trials, the shadow of a visible object was removed from the scene. For the *hidden object shadow deletion* trials, the shadow of an occluded or out-of-frame object was deleted from the scene. In each scene, the two critical shadows were cast by the same type of object (one visible and one occluded).

Experiment 2 included 162 filler images. Most of the filler images from Experiment 1 were reused, except for a few that had been repurposed as experimental images or dropped because we determined that they were near-duplicates (reflections or differently cropped versions) of other filler images.

Design Each participant saw six images from each of the three experimental conditions; the image selection was random and counterbalanced across participants so that each image occurred equally often in each condition across the experiment.




Fig. 3 Comparison of the sizes and positions of the changes in Experiment 1. Error bars represent standard errors of the means

These 18 experimental trials were randomly interleaved with 162 filler trials that were the same for every participant and included 63 object deletions, 63 object color changes, and 36 trials with no change. Trial order was randomized for each participant. As in Experiment 1, 80% of the trials in Experiment 2 had a change. Shadow changes were half as common as in Experiment 1, now occurring on 8.3% of trials.

Procedure The instructions and experimental procedure were exactly the same as in Experiment 1.

Results

One trial was dropped due to a very long response time (>3 SDs above the mean, not including trials that timed out). We compared accuracies, response times, and the numbers of timed-out trials across the three experimental conditions, shown in Fig. 6. No significant difference in response time emerged across the three conditions [one-way repeated measures ANOVA: $F(2, 23) = 0.75, p = .49$], but there was a massive difference in accuracy [$F(2, 23) = 25.2, p < .001$]. Bonferroni-corrected pairwise t tests indicated that people

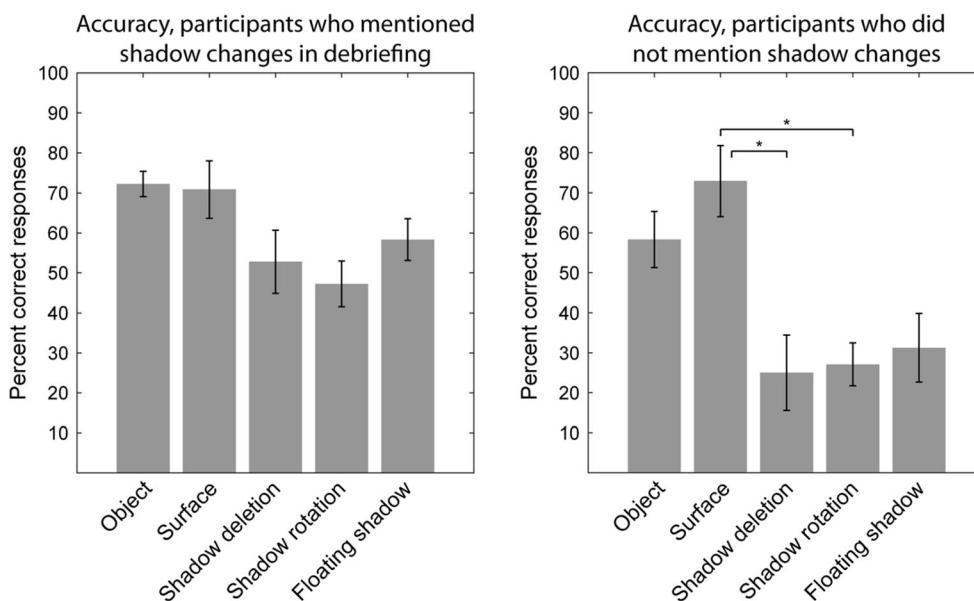


Fig. 4 Change detection performance of participants who mentioned shadow changes when debriefed after Experiment 1 (left), versus participants who did not mention shadow changes (right). Error bars represent standard errors of the means



Fig. 5 Example image (a) and the change conditions (b–d) used in Experiment 2. In this image, the changes are (b) silverware by the leftmost plate is removed; (c) the shadow of the glass on the left is removed; and (d) the shadow of the glass behind the flowers is removed

were significantly better at detecting object changes than at detecting either type of shadow change, but we found no significant difference in accuracy across the two types of shadow changes. There was also a significant difference in the percentages of trials that timed out across the three experimental conditions [$F(2, 23) = 4.35, p < .05$]. Object change trials timed out less frequently than either type of shadow change trial, but only the difference between the object change and shadow change (visible object) conditions was significant.

As in Experiment 1, most of the errors were failures to detect the changes: People said that there was no change, or the trial timed out before they detected the change. Only 1.9% of the errors (19 trials) were location errors, and these appeared to be random guesses.

The average sizes and eccentricities (distance from the center of the image) of the changes across the three experimental conditions are shown in Fig. 7. A one-way repeated measures ANOVA showed no significant differences in sizes [$F(2, 35) =$

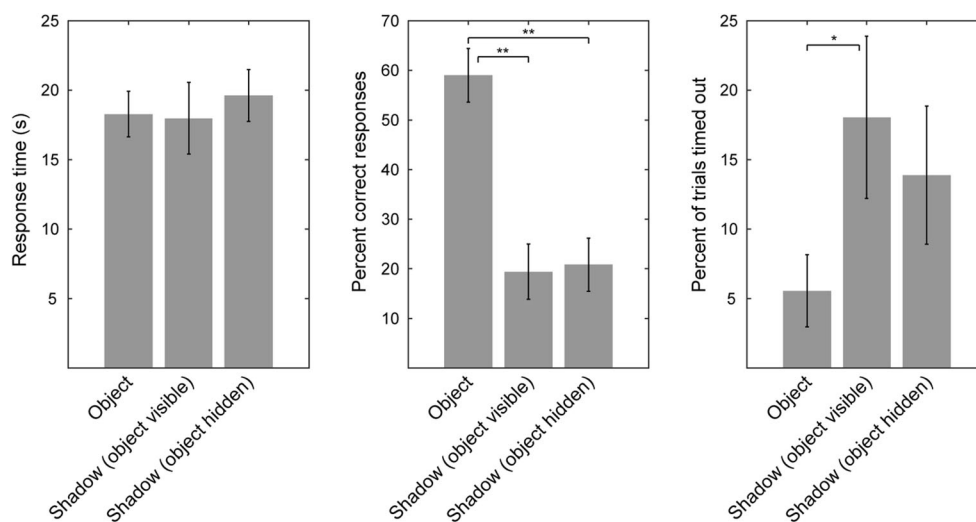


Fig. 6 Response times (left), accuracy (center), and percentages of trials that timed out across the three types of changes in Experiment 2. Error bars represent standard errors of the means




Fig. 7 Comparison of the sizes and positions of the changes in Experiment 2. Error bars represent standard errors of the means

0.24, $p = .78$] or eccentricities [$F(2, 35) = 2.38, p = .10$] across conditions. There was a tendency for the object deletion changes to be smaller and farther from the image center than the shadow changes, even though they were the most easily detected, so size and eccentricity differences probably do not explain our results.

As in Experiment 1, participants were asked at the end of the experiment if they had noticed any patterns to the changes in the scenes. Only one person mentioned shadow changes.

Discussion

In two experiments, we found that people are worse at detecting changes to cast shadows than at detecting changes to the objects or surfaces in scenes. We expected this would be true for certain types of changes, such as the deletion of a cast shadow or the rotation of a shadow to suggest that an object was lit from the wrong direction. In these cases, the change in the shadow does not really alter the meaning of the scene; it introduces a lighting inconsistency, but people are known to be poor at detecting wrong shadows in scenes (Ostrovsky et al., 2005). However, people are no better at detecting shadow movement that implies that an object is floating impossibly above a surface, or shadow deletions that imply an object has been deleted from a scene. These types of changes could be interpreted as object or scene gist changes, but they are significantly less noticeable than changes to the physical objects in the scene.

People may fail to notice changes to shadows because shadows are less salient than other parts of a scene.

However, the surface changes in Experiment 1 had similar local contrast to the shadow changes, so they should have been equally nonsalient. Nevertheless, they were noticed significantly more often than shadow changes. In addition, it is unclear that low-level feature saliency alone drives change detection in natural scenes: Higher-level scene semantics may play a more important role (Stirk & Underwood, 2007).

The results of this study are similar to previous findings by Sareen, Ehinger, and Wolfe (2015), who compared detection rates in a change blindness task when changes occurred to physical objects in the scene or to the same objects' reflections in mirrors. People were significantly worse at detecting changes to reflections, even when the changing object was only visible in the reflection. This is also somewhat surprising, given that other studies have shown that objects visible only in a reflection are treated as being more “real” than the reflections of objects that are also visible in the scene; for example, when counting the number of a given object in a scene, people will include reflections if the object is not visible in the scene, but they never count reflections otherwise (Chesney & Gelman, 2015).

Both cast shadows and reflections are redundant when the object casting them is clearly visible in the scene. Although there are cases in which a cast shadow or reflection provides useful information about an object, these may be fairly rare in everyday life, and it may be more efficient for the visual system to ignore this information and focus on processing the physical objects and surfaces in the scene. This processing should discount shadows and reflections, since these are produced by lighting alone, and do not represent the true color or shape of the surface.

Even when cast shadows are informative, using that information may require some prior processing of the objects in the scene. In order to correctly interpret shadow information, the visual system must be able to determine which shadows are cast by which objects, a problem known as the *shadow correspondence problem* (Dee & Santos, 2011). In this experiment, people may not have noticed the scene-altering shadow changes because they were not able to process the scenes deeply enough to solve this problem. For example, to recognize a shadow change that deletes a hidden object as an “object deletion,” people may need to process both the objects and shadows in the scene and then compare them to realize that, some of the time, an extra shadow is present that does not correspond to any visible object. It is also possible that people do process shadows unconsciously, but for whatever reason this information does not become consciously available for change detection, while information about objects in the scene is consciously available.

Another way to think about these results is to note that attention is preferentially directed to objects (Egly, Driver, & Rafal, 1994; Fiebelkorn, Saalman, & Kastner, 2013), or at least to “proto-objects” (Rensink, 2000; Russell, Mihalaş, von der Heydt, Niebur, & Etienne-Cummings, 2014). Shadows are not objects, and may attract less attention. Thus, changes to shadows may be detected more slowly and/or less successfully. Displaced shadows, acting as surface markings, may be treated more like objects, and thus changes to those markings are found more readily.

Taken with the previous study by Sareen and colleagues (2015), these change blindness results support the idea that cast shadows and reflections are discounted in early visual processing in favor of “real,” physical objects and surfaces. Deeper processing of shadows and reflections may occur somewhat later, after the visual system has already built a rough sketch of the scene. This later processing may reconcile hypotheses about the physical objects in the scene with the information available in shadows and reflections, in order to resolve ambiguities such as a mismatch in the number of objects shown and the number of cast shadows, or inconsistencies in light direction across objects and shadows.

Author note This work was funded by a Center of Excellence for Learning in Education, Science and Technology Grant, No. SBE-0354378, and an Office of Naval Research Grant, No. N000141010278, to J.M.W.

References

- Adelson, E. H. (2000). Lightness perception and lightness illusions. In M. Gazzaniga (Ed.), *The new cognitive neurosciences* (2nd ed., pp. 339–351). Cambridge: MIT Press.
- Aginsky, V., & Tarr, M. J. (2000). How are different properties of a scene encoded in visual memory? *Visual Cognition*, 7, 147–162.
- Braje, W. L., Legge, G. E., & Kersten, D. (2000). Invariant recognition of natural objects in the presence of shadows. *Perception*, 29, 383–398.
- Casati, R. (2004). The shadow knows: A primer on the informational structure of cast shadows. *Perception*, 33, 1385–1396.
- Castiello, U. (2001). Implicit processing of shadows. *Vision Research*, 41, 2305–2309.
- Chesney, D. L., & Gelman, R. (2015). What counts? Visual and verbal cues interact to influence what is considered a countable thing. *Memory & Cognition*, 43, 798–810. doi:10.3758/s13421-015-0505-7
- de Leeuw, J. R. (2015). jsPsych: A JavaScript library for creating behavioral experiments in a Web browser. *Behavior Research Methods*, 47, 1–12. doi:10.3758/s13428-014-0458-y
- Dee, H. M., & Santos, P. E. (2011). The perception and content of cast shadows: An interdisciplinary review. *Spatial Cognition and Computation*, 11, 226–253. doi:10.1080/13875868.2011.565396
- DiCarlo, J. J., & Cox, D. D. (2007). Untangling invariant object recognition. *Trends in Cognitive Sciences*, 11, 333–341.
- Egly, R., Driver, J., & Rafal, R. D. (1994). Shifting visual attention between objects and locations: Evidence from normal and parietal lesion subjects. *Journal of Experimental Psychology: General*, 123, 161–177. doi:10.1037/0096-3445.123.2.161
- Elder, J. H., Trithart, S., Pintilie, G., & MacLean, D. (2004). Rapid processing of cast and attached shadows. *Perception*, 33, 1319–1338.
- Fiebelkorn, I., Saalman, Y., & Kastner, S. (2013). Rhythmic sampling within and between objects despite sustained attention at a cued location. *Current Biology*, 23, 2553–2558.
- Kersten, D., Mamassian, P., & Knill, D. C. (1997). Moving cast shadows induce apparent motion in depth. *Perception*, 26, 171–192.
- Land, E. H., & McCann, J. J. (1971). Lightness and Retinex theory. *Journal of the Optical Society of America*, 61, 1–11.
- Leek, E. C., Davitt, L. I., & Cristino, F. (2015). Implicit encoding of extrinsic object properties in stored representations mediating recognition: Evidence from shadow-specific repetition priming. *Vision Research*, 108, 49–55.
- Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, 390, 279–281. doi:10.1038/36846
- Luck, S. J., & Vogel, E. K. (2013). Visual working memory capacity: From psychophysics and neurobiology to individual differences. *Trends in Cognitive Sciences*, 17, 391–400. doi:10.1016/j.tics.2013.06.006
- Mamassian, P., Knill, D. C., & Kersten, D. (1998). The perception of cast shadows. *Trends in Cognitive Sciences*, 2, 288–295.
- Ostrovsky, Y., Cavanagh, P., & Sinha, P. (2005). Perceiving illumination inconsistencies in scenes. *Perception*, 34, 1301–1314.
- Porter, G., Tales, A., & Leonards, U. (2010). What makes cast shadows hard to see? *Journal of Vision*, 10(3), 13:1–18. doi:10.1167/10.3.13
- Rensink, R. A. (2000). Seeing, sensing, and scrutinizing. *Vision Research*, 40, 1469–1487.
- Rensink, R. A., & Cavanagh, P. (2004). The influence of cast shadows on visual search. *Perception*, 33, 1339–1358. doi:10.1068/p5322
- Rensink, R. A., O’Regan, J. K., & Clark, J. J. (1997). To see or not to see: The need for attention to perceive changes in scenes. *Psychological Science*, 8, 368–373. doi:10.1111/j.1467-9280.1997.tb00427.x
- Rolls, E. T., & Stringer, S. M. (2006). Invariant visual object recognition: A model, with lighting invariance. *Journal of Physiology*, 100, 43–62.
- Russell, A. F., Mihalaş, S., von der Heydt, R., Niebur, E., & Etienne-Cummings, R. (2014). A model of proto-object based saliency. *Vision Research*, 94, 1–15. doi:10.1016/j.visres.2013.10.005
- Sareen, P., Ehinger, K. A., & Wolfe, J. M. (2015). Through the looking-glass: Objects in the mirror are less real. *Psychonomic Bulletin & Review*, 22, 980–986. doi:10.3758/s13423-014-0761-8

- Sareen, P., Ehinger, K. A., & Wolfe, J. M. (2016). CB database: A change blindness database for objects in natural indoor scenes. *Behavior Research Methods*. doi:10.3758/s13428-015-0640-x
- Stirk, J. A., & Underwood, G. (2007). Low-level visual saliency does not predict change detection in natural scenes. *Journal of Vision*, 7(10), 3:1–10. doi:10.1167/7.10.3
- Tarr, M. J., Kersten, D., & Bülthoff, H. H. (1998). Why the visual recognition system might encode the effects of illumination. *Vision Research*, 38, 2259–2275.
- Vinci, L. da. (1888). Notebooks of leonardo da vinci. Project Gutenberg. (J. P. Richter translation, <http://www.gutenberg.org/etext/5000>).
- Wright, M. J. (2005). Saliency predicts change detection in pictures of natural scenes. *Spatial Vision*, 18, 413–430.