Shuffling your way out of change blindness

Emilie Josephs¹ · Trafton Drew² · Jeremy Wolfe^{1,3}

Published online: 24 June 2015
© Psychonomic Society, Inc. 2015

Abstract Change detection is typically discussed in the literature as a 2-state phenomenon. Small differences between otherwise identical images are easy to detect when the images are superimposed in space and alternated in time (“shuffled”). However, change blindness results from any disruption that prevents the critical change from generating the sole salient transient. Here we show that different presentation strategies produce different degrees of change blindness for the same change. Specifically, shuffling the images supports faster change detection than viewing the same 2 images side by side, even when the images contain a number of distracting differences. In Experiment 1, pairs of photographs having a 50 % chance of containing a difference were viewed either in alternation, in a “Shuffle” condition, or simultaneously, in a “Side-by-Side” condition. Change detection was about 6 seconds faster when the images were viewed in the “Shuffle” mode. In Experiment 2, we presented two images that were slightly laterally shifted relative to each other (0–48 pixels). The RT benefit for the Shuffle viewing mode was very strong when the relative shift was small, to insignificant when there was a large difference between the two images. However, at large shifts, Shuffle maintained an accuracy advantage. It seems that changes are easier to detect when comparing images in a Shuffle condition rather than Side-by-Side. This has important implications for real world tasks like radiology where detection of change is critical.

Keywords Visual search · Change blindness · Eye-tracking · Scene perception

Suppose that you want to determine if there are any changes between two otherwise identical pictures. If you place the pictures side by side, this task is hard enough to be a classic children’s game (see <http://www.highlightskids.com/double-check>). If, however, the images are perfectly superimposed in space, and alternated in time (henceforth “shuffling”), then the task is trivial. Shuffling has been used for many years in astronomy, where the stars in the night sky are fixed and the two images of those stars can be readily aligned. Shuffling the two images will reveal the presence of a planet or satellite because it, alone, will appear to move against the fixed background. Indeed, the planet (or ex-planet), Pluto, was discovered by the use of a “blink comparator” that worked on this principle (Tombaugh, 1946). Unfortunately, in a variety of other important tasks that rely on image comparison, the situation is not so simple. In tasks like satellite surveillance or mammographic screening for breast cancer, the observer wants to know if anything significant has changed since the last image was acquired. Unfortunately, elements of the background, though similar from one image to the next, have also changed. Cloud shadows, traffic patterns, and so forth produce irrelevant but distracting differences between two satellite images. The inability to compress a breast in exactly the same way twice alters the appearance of the breast tissue between two mammograms. The result of these relatively modest differences between scene backgrounds is that the critical change becomes much harder to locate, even when the images are shuffled, a phenomenon known as “change blindness.”

Change blindness, the failure to detect even quite substantial differences between two successive stimuli (Simons &

✉ Emilie Josephs
emilie.josephs.1@gmail.com

¹ Brigham and Women’s Hospital, Boston, MA, USA

² University of Utah, Salt Lake City, UT, USA

³ Harvard Medical School, Boston, MA, USA

Levin 1997), can be produced by introducing a brief blank interstimulus interval between consecutive presentations of images that are identical except for the critical change. Changes introduced during saccades (Bridgeman, Hendry, & Stark, 1975), blinks, or camera cuts in movies similarly often go unnoticed (O'Regan, Deubel, Clark, & Rensink, 2000). A similar difficulty is also evident when observers are asked to find the difference between two images presented side by side, forcing a saccade between the two versions of the picture (Scott-Brown, Baker, & Orbach, 2000). Even relatively modest, local changes can produce change blindness, which is why the irrelevant changes in a mammogram or satellite image may disrupt change detection. O'Regan, Rensink, and Clark (1999) produced strong change blindness by simply adding a “mudsplash,” a relatively small, irrelevant, but very salient change that flashed in the images without obscuring the target change. Turatto, Bettella, Umiltà, and Bridgeman (2003) showed that briefly reversing the polarity of the image when the change was introduced also caused change blindness. Finally, Wolfe, Reinecke, and Brawn (2006) showed that it was almost impossible to determine if a spot had changed color if it changed luminance at the same time. These results suggest that our impressive ability to locate changes in consecutive, otherwise identical images is easily disrupted by other irrelevant changes.

The previous work outlined above suggests a two-state description of change detection. If two images are identical except for the critical change, and the images shuffled, change detection is trivial. Any other situation, from side-by-side viewing, to imperfect alignment of the two shuffled images, to the introduction of incidental differences between the images, produces change blindness. Here we ask if all such conditions are equally susceptible to change blindness. Specifically, we test the hypothesis that shuffling two images, even if they are imperfectly matched, will produce faster or more accurate change detection than showing them side by side. In Experiment 1, observers were shown two pictures that were taken from slightly different angles so that they were not perfectly aligned (see Fig. 1), and instructed to find the difference between them. On half of the trials, a single object was removed from one of the pictures. In the Shuffle condition, we superimposed the images in space and allowed observers to alternate between them at will. In the Side-by-Side condition, we presented the images side by side. In Experiment 2, we systematically manipulated the degree of alignment between the images being compared to better understand how reaction times in each viewing condition were influenced by this factor. Beyond our basic interest in change detection, this is a question of practical importance because tasks like mammography often involve comparisons of two similar but not identical images. If one mode of comparison saved time or improved accuracy, even by a small amount, this could be of substantial real-world benefit.

Fig. 1 (a) Can you find the change between these images? This is an example of a typical change-present pair of images. There are a number of incidental differences between the two images (caused by tilting the camera between shots), such as the position of the piano bench relative to the edges of the image. The target is the red dot on the floor that is in one of the images but not the other. (b) The displays used in Side-by-Side (left) and Shuffle trials (right). In the Shuffle condition, observers shifted between the two images using the up and down arrow keys. In the Side-by-Side condition, images were presented simultaneously

Experiment 1

Observers

With no idea of the effect size, we tested roughly double the 6–12 participants tested in typical change blindness experiments (Rensink, O'Regan, & Clark, 1997; Rensink, O'Regan, & Clark, 2000; Saiki & Holcombe, 2012). All observers in Experiment 1 ($N = 23$, mean age 24) gave informed consent, according to the procedures at Brigham and Women's Hospital, and were paid \$10/hour to participate. All were recruited from the general population, had normal or corrected to normal vision, passed the Ishihara color-blindness test, and had no history of eye or muscular disorders.

Apparatus

The experiment was designed in MATLAB version 7.10, using the Psychophysics Toolbox version 3.0.9 (Brainard, 1997; Pelli, 1997). Observers viewed stimuli on a 19-in. Mitsubishi Diamond Pro 91TXM CRT monitor with their heads restrained in a chin rest at a viewing distance of 65 cm. The monitor subtended visual angles of 37° horizontally and 30° vertically, and had a resolution of 1024 × 768. We monitored eye movements using an SR Research EyeLink1000 desktop eye tracker.

Method

The images for this experiment, taken with a digital camera, were predominantly pictures of indoor scenes (apartment,

office, retail). Three images were generated for each scene: one “original” and two “comparison” images. In one comparison image, no changes were made to the scene, while in the other, one object was removed. Both comparison images were taken from a slightly different angle than the original picture. This was achieved by tilting the camera roughly 1 inch down and to the side between pictures. This slight shift introduced small changes between the otherwise identical images. This emulates the small differences between images taken at different times in applications such as radiology and satellite surveillance. Care was taken to ensure that lighting conditions remained the same for all three pictures. All images were resized to 600×450 pixels in Adobe Photoshop CS5 ($15.2^\circ \times 11.8^\circ$ of visual angle).

Each trial consisted of the original image and one comparison image. Observers were instructed to find the change as quickly as possible and click on its location in the scene or, if there was no change, to click on a “no change” box at the bottom of the screen. As the two photos were taken from slightly different angles, there were a number of incidental differences near the edges of the photographs, so “a change” was operationally defined as a place where an object was present in one photograph and absent in the other. Participants were warned about the existence of such incidental differences and were told to be careful in their responses to avoid them. Fifty percent of the trials were change-present trials.

There were two viewing conditions (Fig. 1). In the “Side-by-Side” condition, the images were located next to each other. In the “Shuffle” condition, only one image was visible at a time, positioned in the center of the screen. Observers could toggle back and forth between the two images by pressing the up and down arrow keys of the keyboard. In both conditions, observers could click on the location of the change in either of the images. We counterbalanced viewing conditions and target presence across observers. Following their response, observers were given feedback on their accuracy and the location

of the change, if one was present. Observers then pressed a button to move on to the next trial. There were 88 trials in this experiment, including 4 trials for practice. The practice trials were designed to clarify our definition of a “change” and to illustrate all four combinations of change, no change, Shuffle and Side-by-Side.

Results

To simulate real-world change-detection tasks, observers were given unlimited time to complete each trial and were told to be as accurate as possible. As a result, our primary variable of interest was response time rather than accuracy. We limited our reaction time analyses to correct trials since we are primarily interested in the metrics of accurate performance. As shown in Fig. 2, correct Shuffle trials were completed approximately 6 seconds faster than correct Side-by-Side trials, $t(23) = 5.904$, $p < .0001$, Cohen's $d = 0.92$). This large effect size was comparable for absent trials, $t(23) = 6.138$, $p < .0001$, $d = 0.80$, and present trials, $t(23) = 3.1$, $p = .0044$, $d = 0.83$.

All accuracy measurements are reported in the text and in graphs as the proportion of incorrect trials, but these values were arcsine transformed for significance testing to mitigate the fact that the distribution of proportion data is binomial rather than normal. There was no difference in overall performance as a function of viewing mode: Shuffle 84.52 % correct, Side-by-Side 83.12 %, $t(23) = .439$, $p = .697$. However, as shown in Fig. 2, while there was no difference in error rates when the target was present—Shuffle 74.80 % correct, Side-by-Side 73.02 %, $t(23) = .84$, $p = .409$ —the Shuffle viewing mode led to a small (<3 %), but reliable increase in false alarm rates for target-absent trials, Shuffle 94.25 % correct, Side-by-Side 97.22 %, $t(23) = 2.43$, $p = .023$, $d = 0.46$. False-alarm errors most likely occur when observers click on an incidental change near the margin of the image. It is unlikely that these error rates are a product of a speed-accuracy trade-off, as there

Fig. 2 Reaction time and error rate data from Experiment 1, separated by target-present and target-absent trials. Error bars here and throughout the paper represent standard error of the mean

is no reason to imagine that such a small difference in the number of false alarm errors (~ 1 error per observer) would have a 6 second effect on the time required to correctly detect a change, especially when no such trade off is seen in target-present trials, which also had faster reaction times in Shuffle trials. Overall, the Shuffle conditions led to substantially faster decisions with little influence on overall accuracy.

To better understand the cause of the differences in reaction times, we compared the duration of two periods of interest in the eye-movement data: the time from the start of the trial until the observers' first fixation on the location of the change (time to first fixation) and the time between the first fixation and the observer's response (decision time; see Fig. 3). For these variables, we focused on change-present trials, as change-absent trials lack any specific region of interest. While there was no difference in the time to first fixation, $t(23) = 0.76$, $p = .46$, decision time was decreased, $t(23) = 3.377$, $p = .0026$, $d = 0.85$ in Shuffle trials. We also compared the deployment of eye movements between the two conditions, measuring the number of fixations, the average duration of fixations, and the saccade amplitude. For these measures, we included target-present and target-absent trials, since these are metrics of overall strategies differences between the two conditions. Observers made fewer fixations in Shuffle trials, $t(23) = 10.31$, $p < .0001$, $d = 1.87$, but each fixation lasted nearly twice as long as in Side-by-Side trials, $t(23) = 14.08$, $p < .0001$, $d = 2.77$. Saccades within the same image (i.e., excluding saccades from one image to another) were shorter in the Side-by-Side condition than in the Shuffle condition, $t(23) = 4.068$, $p = .0005$, $d = 0.77$.

Discussion

As expected, our modified change blindness task was quite difficult, as evidenced by the long reaction times in both the Side-by-Side (24 seconds) and the Shuffle condition (18 seconds). However, our data suggest that the Shuffle viewing mode is less susceptible to change blindness: Shuffle trials were completed an average of 6 seconds faster than trials in the Side-by-Side condition. Data from eye movements suggest a possible mechanism for this effect: The time between the first fixation on the target and the participant's response was 4 seconds faster in Shuffle trials while time to first fixation was unaffected. This suggests that the Shuffle viewing mode allowed observers, having found a potential change in the scene, to confirm it more quickly. Eye movement data also revealed different patterns in saccadic amplitude, fixation count, and average fixation duration in the different viewing modes, suggesting the recruitment of different search strategies, which may have contributed to the observed benefit for the Shuffle technique. Taken together, these results reveal that the different viewing modes led to substantially different search strategies, which may have contributed to the observed benefit for the Shuffle technique.

The observed reaction time benefit for the Shuffle viewing mode may be relevant for a number of tasks that require comparing images to find changes between them. Furthermore, the finding that this benefit is robust to a moderate amount of misalignment between the images (created by tilting the camera between shots) is important given that images in these tasks often also contain incidental differences. However, pairs

Fig. 3 Eye-tracking metrics from Experiment 1: Time to first fixation, decision time, average fixation duration, saccade amplitude, and fixation count. Observers first fixated the change equally quickly in both conditions but confirmed it faster ($p = .029$) on the Shuffle trials.

Observers made longer saccades in Shuffle than in Side-by-Side trials ($p = .0005$). They also made fewer fixations ($p < .0001$), but those fixations lasted longer ($p < .0001$)

of images in these real-life tasks can span a broad range of misalignment, while our test images all contained approximately the same degree of shift. In Experiment 2, we varied the range of displacement in our test stimuli in a systematic manner to test whether Shuffle trials remain faster than Side-by-Side trials across a broad range of differences.

Experiment 2

Experiment 1 showed that observers perform the change-detection task more quickly when images are viewed in a Shuffle mode than in a Side-by-Side mode, even when there was a slight misalignment between the two images. Experiment 2 systematically varied the degree of displacement between the two images to understand how this influenced performance in each viewing condition.

Observers

There were 21 participants in Experiment 2. Two participants were excluded for error rates exceeding 25 %, and one was excluded because the program failed in the middle of the experiment. All observers in Experiment 2 (mean age = 28) gave informed consent according to the procedures at Brigham and Women's Hospital and were paid \$10/hour to participate. All were recruited from the general population, had normal or corrected to normal vision, passed the Ishihara color-blindness test, and had no history of eye or muscular disorders.

Apparatus

The experiment was designed in MATLAB version 7.10, using the Psychophysics Toolbox version 3.0.9. Observers viewed the stimuli on 19-in. Viewsonic NX1932w monitors with 1440×900 resolution, which subtended approximately $35.4^\circ \times 24.5^\circ$ of visual angle. Observers were seated approximately 60 cm from the screen, but unlike Experiment 1, heads were unrestrained.

Design

Experiment 2 consisted of the same task as Experiment 1: Participants were instructed to find the difference between two images, with the difference defined as the object that was present in one image and removed from the other. Unlike Experiment 1, all trials in Experiment 2 were target-present trials. Pilot studies performed in our lab suggested that when there is a broad range of search difficulty, as in this case, the presence of both target-present and target-absent trials may lead to criterion effects in which participants end the search prematurely by guessing “target absent” in difficult trials

rather than taking the time to thoroughly search the scene. Participants had unlimited time to search the scenes, and if they could not find the change, there was a region at the bottom of the screen labeled “Give up” that the participants could click to go to the next trial.

In order to test whether the reaction-time advantage is robust to displacement between the images, we slightly modified the method of stimulus creation from Experiment 1. We obtained a new set of images that were exactly the same save for the change in the target object. One hundred and three images were obtained by using a tripod to take two pictures from the same location while physically removing an object from the scene between pictures. An additional 84 image pairs were obtained from a set created by Sareen, Ehinger, and Wolfe (2014), in which scenes were modified in Adobe Photoshop CS4 to digitally remove one object. All images were then resized to 600×450 pixels. To create slightly different images, emulating conditions commonly encountered during satellite image analysis or mammography, we used 552×450 pixels ($13.8^\circ \times 11.8^\circ$ of visual angle at 60 cm) subsets of the larger image. Subset images were created that differed by 0, 6, 12, 24, or 48 pixels in the horizontal direction (0.00 %, 1.09 %, 2.17 %, 4.35 %, and 8.70 % of the total image size). Thus, one member of each image pair (except for the 0 pixel condition) had a unique vertical region at the left margin of the image while the other had a unique region on the right (Fig. 4). As in Experiment 1, this shifting of the images caused a number of incidental differences at the edges of the photographs.

Fig. 4 Figure 4 demonstrates how the image pairs were created for Experiment 2. The original image size was 600×450 pixels. One of the stimulus images was a 552×450 section of the original image, and the second was a section of the same size, taken from a laterally displaced segment of the original image. The two segments could be displaced, relative to each other, by 0, 6, 12, 24, or 48 pixels. The only other difference between the two images was the removal of a single object from one of the images (in this case, a tube of hand cream is present on the desk in the picture on the left, but not on the right)

Fig. 5 Average reaction times in seconds for Experiment 2, by View Type and Shift Size. The displacement between the images was 0, 6, 12, 24, or 48 pixels, here displayed in the *x*-axis as the percentage of total image size

The target change was again operationally defined to the observers as the object that was intentionally removed from the scene. There were 17 trials at each Shift Size for both the Shuffle and Side-by-Side trials, leading to a total of 170 trials. In addition, there were four practice trials at the beginning that served to introduce the image viewing modes and our

definition of a change. Because of program failures, two people completed fewer than the 170 trials (150 and 142 trials).

Results

Reaction-time data (see Fig. 5) were submitted to a two-way ANOVA with factors of View Type (Shuffle or Side by Side) and Shift Size (0.00 %, 1.09 %, 2.17 %, 4.35 %, or 8.70 % of the total image). This revealed a main effect of View Type, $F(1, 20) = 85.37, p < .0001, \mu^2 = 0.27$, and a main effect of Shift Size, $F(4, 80) = 23.94, p < .0001, \mu^2 = 0.16$. There was also a significant interaction, $F(4, 80) = 15.91, p < .0001, \mu^2 = 0.09$, with Shuffle trials being faster than Side-by-Side trials at smaller Shift Sizes. Pairwise comparisons using Sidak's multiple comparisons test showed that Shuffle was faster for Shift Sizes of 0.00 %, 1.09 %, 2.17 %, and 4.35 %, $t(21) = 10.47, t(21) = 9.63, t(21) = 4.58$, and $t(21) = 4.69$, respectively, $p < .05$ for all. At the highest Shift Size, there was no difference in reaction time, $t(21) = 0.83, p > .05$. Pearson's correlation coefficients were computed to assess how Shift Size influenced reaction times in each condition. While Shift Size did not

Fig. 6 Average error rates for Experiment 2, by error type. The displacement between the images is displayed in the *x*-axis as the percentage of total image size. (a) Average rate of "Give-up" responses. (b) Average rate of False Alarms

predict performance in the Side-by-Side condition, $r^2 = .24$, $p > .05$. Shuffle performance was very strongly predicted by the amount of displacement, $r^2 = .92$, $p = .0099$.

Error rates were divided by type: Give-up errors (see Fig. 6a) meant the participants had passed on to the next trial without making a guess. Trials where participants clicked on portions of the image that did not include the change were False Alarms (see Fig. 6b). Give-up errors were not modulated significantly by View Type, $F(1, 20) = 1.59$, $p = .22$, Shift Size, $F(4, 80) = 0.0876$, $p = .088$, or the interaction thereof, $F(4, 80) = 1.55$, $p = .20$. False alarms, however, showed a main effect of View Type, $F(1, 20) = 30.86$, $p < .0001$, $\mu^2 = 0.07$, a main effect of Shift Size, $F(4, 80) = 8.29$, $p < .0001$, $\mu^2 = 0.12$, and a significant interaction, $F(4, 80) = 3.91$, $p = .006$, $\mu^2 = 0.04$. Note that in Experiment 2, false-alarm errors are markedly more common in the Side-by-Side condition. The number of false-alarm errors in Side-by-Side trials was very strongly predicted by Shift Size, $r^2 = .99$, $p = .0006$, but errors in Shuffle trials were not correlated with Shift Size, $r^2 = .54$, $p = .16$. This pattern of errors is different from Experiment 1, where there was no difference in errors for target present trials. However, Experiment 1 did not have a separate “Give-up” category, so error analyses in that experiment reflect all errors (false alarms, misses, and give up), rather than just false alarms.

Discussion

Experiment 2 shows that reaction times were shorter in Shuffle viewing mode over a wide range of displacement sizes. At the highest amount of displacement between the two images, 8.70 % of the image size, the reaction times were not significantly different between Shuffle and Side-by-Side trials. However, even in this case, false alarm rates were markedly lower in Shuffle trials, suggesting an advantage for Shuffle in this condition as well. False alarm rates increase with Shift Size, presumably because there are more objects incidentally disappearing and reappearing at the margins. While participants were instructed to ignore changes caused by the shift in the image, it appears that they were not always successful, especially in Side-by-Side trials. In the Shuffle condition, the images were alternating in the same location, making it is easy to determine the size of the shift, and easy to decide whether the disappearance of a given object was caused by this shift or not. In contrast, in Side-by-Side trials, where both images were present on different sides of the screen, it is phenomenologically more difficult to determine how much the image has shifted, leading to more confusion.

General discussion

In this study, we compared two viewing modes to evaluate their influence on the speed and accuracy of change detection in

scenes that were slightly misaligned. We found that shuffling between images in the same position on the screen led to faster change detection than showing the images simultaneously. This effect was evident for both present and absent trials in Experiment 1. In Experiment 2, we took two snapshots of the same scene from the same position and systematically varied the horizontal displacement between two images to create a range of image misalignment. We found that shuffling the images remained faster if the images were displaced by 0 % to 4.35 %. The reaction time advantage for shuffling was not significant when images were displaced by 8.7 %, though the false-alarm rate was lower, indicating that shuffling might still be the more effective method for finding changes. Our findings are broadly consistent with Riley, Simpson, Bochud, Steel, and Porter (2013), who found that detection of a Gaussian blob in noise was better when the images were viewed successively in the same location as compared to side by side, so long as the noise in the two images was highly correlated.

The present results suggest that image analysts, looking for changes between images, would benefit from shuffling the images rather than viewing them side by side, even if two shuffled images are not perfectly aligned. With modest differences between images, we observed a reliable benefit in reaction time or error rates for viewing the images in the Shuffle mode. We suspect that much of this effect is due to the different oculomotor demands of the two situations, and our eye-tracking data provide some evidence in favor of this assertion. If two images are alternated in time, there is no need for a series of eye movements between corresponding points in the two images. Thus, Shuffling reduces the demands on both the memory and oculomotor systems. Traditional analog viewing techniques in radiology and satellite image analysis would have made the use of shuffle viewing modes difficult (though the astronomers, hunting for Pluto, managed). The advent of modern, digital workstations could make shuffling very easy. The digital era has brought a vast increase in the volume of imagery (e.g., Andriole et al., 2011; Skaane et al., 2013; Buist et al., 2011) so a technique that reduced comparison time by a few seconds could be valuable when those seconds are aggregated over thousands of comparisons. To document a real benefit, future work will need to replicate these results with experts looking for changes in images in their area of expertise.

Acknowledgments We thank Raymond Farmer, Cameron Douglas, Eric Chan, and Alexander Hall for creating the stimulus images

References

- Andriole, K. P., Wolfe, J. M., Khorasani, R., Treves, S. T., Getty, D. J., Jacobson, F. L., ... Seltzer, S. E. (2011). Optimizing analysis, visualization, and navigation of large image data sets: One 5000-section CT scan can ruin your whole day. *Radiology*, 259(2), 346.

- Brainard, D. H. (1997) The psychophysics toolbox, *Spatial vision* 10, 433–436.
- Bridgeman, B., Hendry, D., & Stark, L. (1975). Failure to detect displacement of the visual world during saccadic eye movements. *Vision Research*, 15(6), 719–722.
- Buist, D. S., Anderson, M. L., Haneuse, S. J. P.A., Sickles, E. A., Smith, R. A., Carney, P. A., ... Miglioretti, D. L. (2011). Influence of annual interpretive volume on screening mammography performance in the United States. *Radiology*, 259(1), 72.
- O'Regan, J. K., Deubel, H., Clark, J. J., & Rensink, R. A. (2000). Picture changes during blinks: Looking without seeing and seeing without looking. *Visual Cognition*, 7(1/3), 191–211.
- O'Regan, J. K., Rensink, R. A., & Clark, J. J. (1999). Change blindness as a result of mudsplashes. *Nature*, 398(6722), 34.
- Pelli, D. G. (1997) The video toolbox software for visual psychophysics: Transforming numbers into movies, *Spatial Vision* 10, 437–442.
- 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(5), 368–373.
- Rensink, R. A., O'Regan, J. k., & Clark, J. J. (2000). On the failure to detect changes in scenes across brief interruptions. *Visual cognition*, 7(1–3), 127–145.
- Riley, G., Simpson, W. A., Bochud, F. O., Steel, J., & Porter, G. (2013, August 14). *Animated mammogram displays for improved breast cancer detection*. Washington, DC: Talk presented at Medical Image Perception Society.
- Saiki, J., & Holcombe, A. O. (2012). Blindness to a simultaneous change of all elements in a scene, unless there is a change in summary statistics. *Journal of vision*, 12(3), 2.
- Sareen, P., Ehinger, K. A., & Wolfe, J. M. (2014). Through the looking glass: Objects in mirrors are less real. *Psychonomic Bulletin & Review*. doi:10.3758/s13423-014-0761-8.
- Scott-Brown, K. C., Baker, M. R., & Orbach, H. S. (2000). Comparison blindness. *Visual Cognition*, 7(1/3), 253–267.
- Simons, D. J., & Levin, D. T. (1997). Change blindness. *Trends in Cognitive Sciences*, 1(7), 261–267.
- Skaane, P., Bandos, A. I., Gullien, R., Eben, E. B., Ekseth, U., Haakenaasen, U., ... Gur, D. (2013). Comparison of digital mammography alone and digital mammography plus tomosynthesis in a population-based screening program. *Radiology*, 267(1), 47–56.
- Tombaugh, C. W. (1946). The search for the ninth planet, Pluto. *Astronomical Society of the Pacific Leaflets*, 5(209), 73.
- Turatto, M., Bettella, S., Umiltà, C., & Bridgeman, B. (2003). Perceptual conditions necessary to induce change blindness. *Visual Cognition*, 10(2), 233–255.
- Wolfe, J. M., Reinecke, A., & Brawn, P. (2006). Why don't we see changes? The role of attentional bottlenecks and limited visual memory. *Visual Cognition*, 19(4/8), 749–780.