CHANGE DETECTION BLINDNESS IN MOVING VISUAL SCENES

Michael Coppola, Jillian DeMair, Ashley Feneran, Kelly Huang, Janet Lee, Deepali Malhotra, Joshua Min, Raina Shivashankar, Alice Xie, Grace Yeung

Advisor: Patrick Dolan
Assistant: Anna Labowsky

ABSTRACT

Experience shows that an individual, after viewing a rich visual scene, is left with a detailed perceptual trace of the properties and locations of objects in that scene. At the same time, there are many classic instances of unnoticed editing mistakes in movies. Even though some changes within the clip are obvious, the observer surprisingly cannot detect them. When Dorothy meets the scarecrow in "The Wizard of Oz," for example, her pigtails are at first short (above her shoulders) but become longer (below her shoulders) across scene cuts. Her hair length, as a matter of fact, changes five times in one scene, yet the viewer typically fails to notice these alterations. It has been suggested that these changes may go unnoticed because the viewers' attention is directed towards the plot, scenery, and dialogue of the film. They do not expect these changes to occur, especially since the modifications to Dorothy's pigtails are small and subtle. To study this pattern of lack of perception and attention, we created a video with both obvious and minor changes across scenes. In addition, we manipulated whether or not viewers expected changes to occur by having an instructed and non-instructed group. Participants of the experiment watched the clip and reported the changes that they observed. Results indicated that attention to changes did increased the viewers' ability to detect them. Furthermore, it was found that changes to more prominent objects in the visual scene were more noticeable than others.

INTRODUCTION

Imagine watching a video clip in which three actors are gossiping about their mundane lives. While listening to their banal conversation, you would think that you are attentive to the visual scenery in the film. Yet, you fail to notice subtle changes that have taken place across scenes. Your inability to discern these disparities is caused by a condition known as change detection blindness.

Change detection blindness, or the inability to notice changes in our surroundings, occurs in our everyday lives. Since we expect a constant environment around us, we do not necessarily notice subtle changes that can occur. Several current experiments have been conducted to observe change detection blindness. Studies have varied in the type of stimuli and/or tasks required of subjects, but all are based on a fundamental design. Every experiment incorporates the same elements: presenting subjects with a stimulus, changing the stimulus, and observing a subject's response to the change.

The extent to which the observer notices a modification in a given scene depends on the degree of change. For example, a simple type of change involves the existence (addition/deletion) of the object and its spatial arrangement (layout). The three changes that are used in most studies of this type include the following: change to an object with a property unique in the scene, change to a nonessential object, and a switch in properties of two or more objects. In previous change detection blindness studies, changes made to the displays were minor and did not drastically alter the overall appearance of the scene. These changes can be as subtle or as apparent as desired by the conductor of the experiment. In our experiment, we varied the degree of changes involving the properties of the objects, such as size, shape, orientation, and color.

A study conducted by psychologist Ronald A. Rensink [1] demonstrated that people are more likely to detect changes that deal with objects vital to the scene's meaning or plot rather than changes that deal with the periphery. It is evident that viewers are more likely to encode or attend to a difference in a central object than they are to an object or setting of secondary importance. In our experiment, we implemented the cut-contingent method in which objects were altered when the camera angle changed. Rensink maintained that changes made during these camera cuts cut-contingent were difficult for subjects to detect. The viewer does not notice these differences because his or her attention is focused on the central objects.

An important factor in the results of a change detection blindness experiment is the intention of the observer. Intention can be defined as the extent to which the observer is expecting a change in a given scene. For example, previous studies have used the intentional approach. In this type of experiment, the viewer is informed about the changes, and he or she concentrates more on detecting those changes. However, prior studies have shown that under these circumstances change detection blindness is likely to occur. Even though the observer is told to pay attention to possible changes, he or she will tend to miss the alterations. Another means of testing the observer is the incidental approach in which the observers are not informed of any changes. As expected, the degree of blindness found in this type of study is usually much greater than that found in the intentional approach. In our experiment, one of the main purposes was to examine these differences between informed and uninformed viewers.

The ability to detect changes pertains to an individual's tendency to focus his or her attention on both central and detailed events. For example, observers, while watching a film, must be aware of several stimuli in sight and sound. The complication of many stimuli at once distracts the observers from noticing slight modifications, especially when the objects appear appropriate to the environment. In order to create a sense of comfortable realism in our experiment, we chose to use a dynamic display, or a visual scene rich with detail and plot. Since our video clip emulated a real life situation, with rich dialogue, scenery, and characters, we predicted that it would be particularly difficult for viewers to detect the minor changes.

Several past experiments have been important to our study of change detection blindness. In 1 study by Levin and Simon, [5] their first experiment involved a video clip of two actors conversing across a table in which at least one error could be found between cuts. For example, the plates on the table would change from red to white. In addition the bright scarf worn by an actor in one scene disappeared in the next. When the video clip was shown to observers, both the incidental and intentional approaches were taken. For the incidental approach, it was found that 89 of the 90 changes in the clip went unnoticed. Even when the intentional approach was used, the observers only noticed twenty out of the ninety changes.

In the same paper, Experiment 2 involved changing the central actor rather than that of any minor items. In this experiment, an actor sitting at a desk hears the telephone ringing in a hallway. As he gets up from his seat to leave the room, the camera cuts to the next shot and a different actor is seen answering the telephone. Data were gathered from a group of viewers, half of which were forewarned about the actor change. As expected, viewers who were told about the change had little trouble differentiating between the actors, while uninformed viewers beforehand were unaware of the change.

We utilized a combination of elements from the aforementioned studies to design and conduct our experiment on change detection blindness. Our study includes both object changes and an actor change. In addition, our experiment uses the intentional as well as the incidental approach. We assumed that viewers would notice the apparent changes (i.e., the picture) more easily than the inconspicuous changes (i.e., the necklace). For example, the necklace that disappears in one scene is much more subtle compared to the picture that is turned upside down in another scene. We also predicted that viewers would find it easier to detect changes when they were involved in the intentional approach as opposed to the incidental approach.

METHODS

Video

Using a digital video camera, we filmed a two-minute video clip of three actors having an inane conversation at a table. The footage was edited and converted to a computer QuickTime file. The actors, one male and two females, were seated side-by-side with the male in the center. In the background, there was a window with a shade and on the window sill were two plants. There also was a portrait hanging on the wall. On the table in front of the actors were several props, such as a fruit basket, cups, and plates (series of still photos in Appendix A).

There were six scenes, each with a different camera angle. Between the camera changes, a total of eleven alterations were made to the scenery. In the transition from scene one to scene two, the shade on the window was lowered, and a plate with fruit on it switched positions with an empty plate. From scene two to scene three, a cup that was crushed and thrown in the previous scene reappeared, the portrait on the wall was inverted, and the necklace on the female actor to the viewer's left was removed. In scene four, a water bottle on the table was changed into a Dr. Pepper bottle, and the watch on the wrist of the female actor to the right disappeared. In the following scene, the plants on the windowsill swapped positions, the fruit in the basket was rearranged, and the female actor on the viewer's right was replaced by another similar-looking actor with similar but different clothes and hair (Appendix B). In the final scene, the male actor's tie changed from silver to black. Each alteration that was made across scenes was reversed prior to the next scene, with the exception of the watch, the fruit basket, and the actor. Since we had to re-film the last two scenes on a separate occasion due to technical difficulties, some negligible details could not be replicated exactly. For example, the male actor wore two different-colored pairs of shorts and the Coke cans were not exactly the same. These discrepancies were likely invisible to the viewer and unlikely to have affected our results.

Experimental Protocol

For our experiment, the subjects had to be 18 years of age or older. Twenty people, some given candy as an incentive, volunteered to participate. These subjects were tested individually at Drew University, either in a secluded area of the cafeteria in the University Commons or at the psychology department of the Hall of Sciences. Each subject was first given a consent form (Appendix C) to read and sign, and then received either one of two instructions. Ten individuals, five males and five females with an average age of 32.2, were told to simply watch the video clip (uninstructed). The other ten individuals, with the same 1:1 gender ratio and an average age of 32.6, were told to pay close attention and watch for changes between scenes (instructed). Most of the subjects viewed the video clip on a laptop screen with headphones, while others watched the clip on desktop monitors with regular speakers. There were several instances of technical difficulties, including poor sound quality and a computer crash between viewings. These problems did not appear to adversely affect our results. Subjects were tested one at a time to ensure that their answers were their own and that they did not confer with other participants. On two separate occasions, pairs of subjects viewed the video clip on a projection screen. In one instance, two subjects who were shown the video together consulted each other during the video clip. This did not seem to significantly affect their performances on the questionnaire.

After the subjects finished watching the video clip, they were given a questionnaire and asked to answer question #1 (Appendix D). They were then shown the video clip a second time, and regardless of which group they were initially in, they were told to watch for changes. After the second viewing, they were instructed to answer question #2, which was identical to question #1. Participants were then asked to provide their gender and age. For the last part of the experiment, subjects were told that at some point in the clip there was an actor change and were shown two photos that resembled (but were not actually from) scenes in the video clip. They were instructed, upon a comparison of the actors in the two photos, to identify the one who appeared in the latter portion of the clip. The subjects were told to base their decisions solely on the physical appearance of the actor that had been substituted and were discouraged to use any other elements of the photographs in making their choice. At the conclusion of the experiment, subjects were given a debriefing form (see Appendix E), which explained the intent and details of our project and provided contact information to answer any of their remaining questions.

RESULTS

Our video contained ten scenery changes and one actor change. We first evaluated how well the instructed and uninstructed groups compared in their overall ability to detect the changes. Any differences between the two groups would skew any further analyses. Data were combined across the two viewings such that both groups were given the opportunity to view the film at least once while looking for changes. The two groups performed very similarly, noticing an average of 50% of the changes (48% for the instructed group and 51% for the uninstructed group; see Figure 1 below). Both groups of viewers tended to notice very few changes because they often focused their attention on the engaging dialogue, rather than the changing visual scenes. The visual scene was too complex for the audience to notice all the minute changes.

The effect of instruction was the first aspect we examined. We calculated the average percentage of changes detected during the first viewing for both groups. The uninstructed group noticed only 16% of the changes whereas the instructed group noticed 22% of them. This reflects a 37.5% increase in identification of changes above the uninstructed group. The limited number of subjects in our experiment resulted in low statistical power that influenced all statistical analyses on the results reported below. Upon viewing the video a second time, the uninstructed group was able to identify more changes, bringing up their total to levels similar to the instructed group.

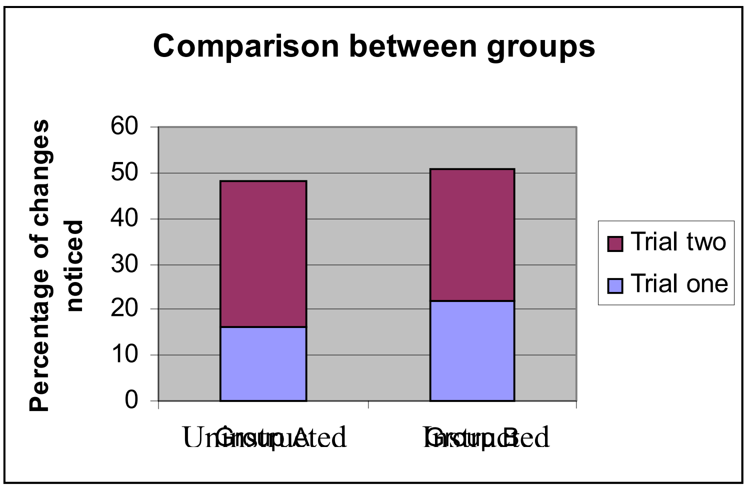

Figure 1.

Looking at the differences in identification for the individual changes, we found considerable variability. For example, when combining the results of the two consecutive viewings, the inverted picture (Appendix A Scene 3) was identified by 100% of the viewers and the moved plate of fruit (Appendix A Scene 2) was identified by 75% of the viewers. On the other hand, the re-arranged fruit basket (Appendix A Scene 5) was only identified by 25% of the viewers and the lowered window shade (Appendix A Scene 2) was noticed by 15% of the viewers.

After tabulating the identification of the individual items for the first trial, we found a number of interesting results. Among the ten object changes, four of these (the window shade, the cup, the necklace, and the watch) went unnoticed in the first trial for both groups of subjects, while a second viewing was required for others to observe these differences (see Figure 2 below). These changes may not have been noticed in the first trial because they were incidental and did not change the overall visual scene.

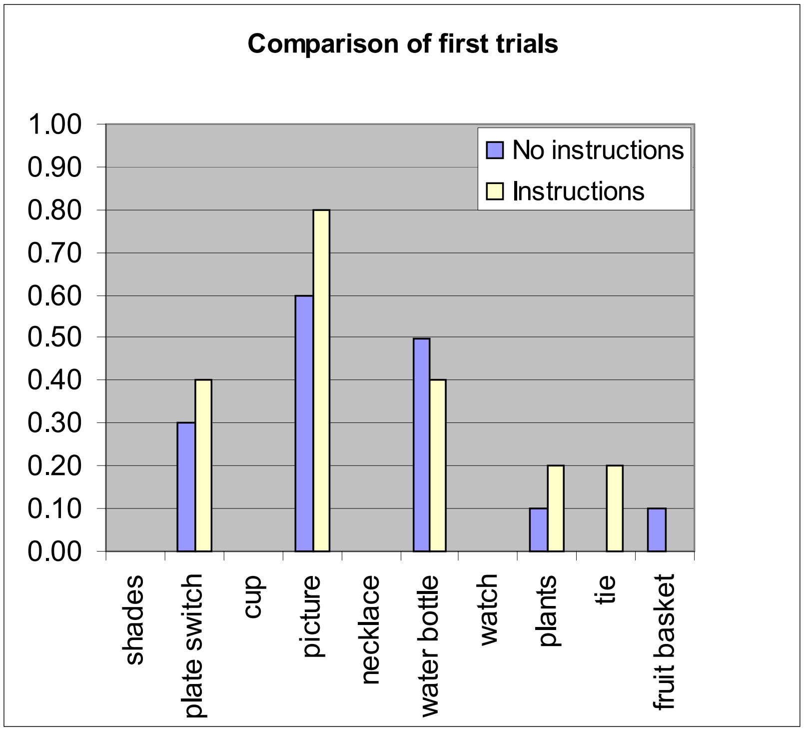

Figure 2.

We tried to determine what caused some changes to be noticed and others to be overlooked. We hypothesized that the duration for which each change was present affected the viewers' ability to perceive it. In support of our theory, a positive correlation was found between the number of times the uninstructed group identified the change and the amount of time each change was on the screen (see Table 1). Our correlation coefficient (r) was calculated to be positive .28. This indicates that a moderate relationship exists between the elapsed time and the proportion of objects identified by the test subjects.

Comparison of Changes Noticed and their Duration

Change	shades	plate switch	picture	necklace	water bottle	watch	plants	tie	fruit basket
number of seconds	9	9	45	45	19	19	16	13	29
proportion identified	0	0.3	0.6	0	0.5	0	0.1	0	0.1

Table 1.

Since the differences between the instructed and uninstructed groups were subtle and never reached statistical significance, we collapsed across both groups for the following comparisons. Turning first to sex of the viewers, we had a total of 10 females and 10 males. On average, females, after viewing the video clip twice, noticed 53% of the changes and males noticed 46% (see Figure 3). Similarly, this also held true for all subjects after watching the clip only once, with females noticing 22% of the changes and males noticing only 16%.

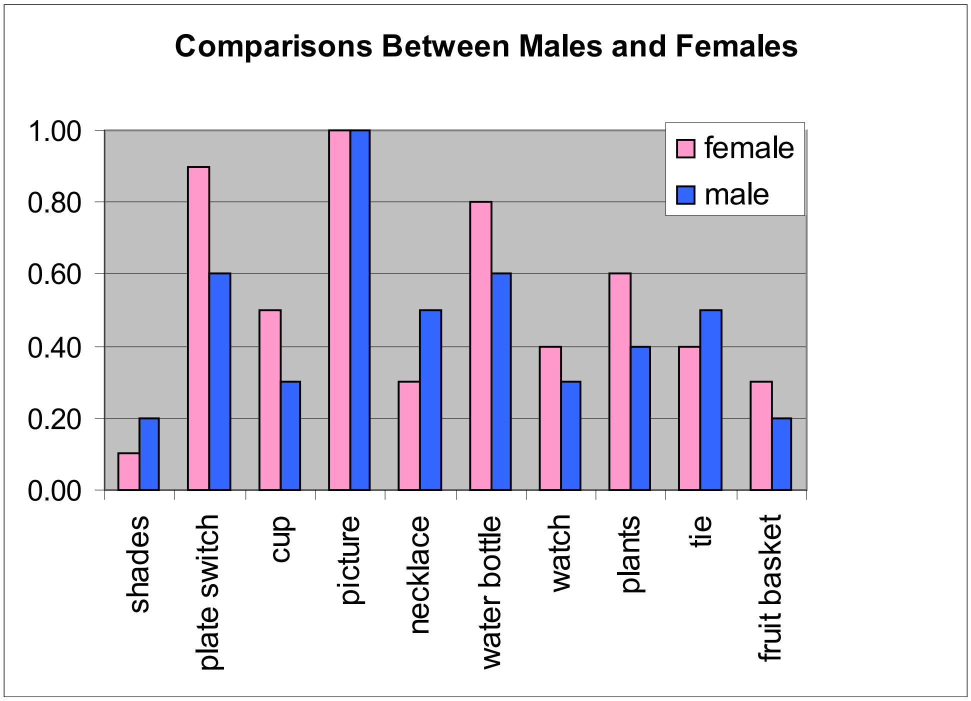

Figure 3.

Additionally, comparisons were made based on age. A median split was performed on the age of our 20 subjects, resulting in the younger half averaging 24 years of age (range = 20-32) and the older half averaging 42 years of age (range = 32-58). Generally, the younger group noticed 47% of the changes while the older group noticed 52% of the changes, thus indicating that the two age groups performed similarly. An important finding to note is the presence of false memories (the situation when subjects report changes that did not actually occur) between our two age groups. For example, one test subject indicated that the male actor in the center had removed his watch between cuts when in fact he never wore a watch. These errors were more prevalent in the older group in comparison to the younger group. After viewing the film twice, 30% of the younger subjects and 60% of the older subjects reported false memories. This finding is consistent with the literature on memory and aging that frequently reports large increases in false memories among older adults [3].

We also observed that after the first viewing, 60% of the uninstructed group reported false memories in comparison to 30% for the instructed group. These results were appropriate because the uninstructed group was faced with a task for which they were unprepared, and thus more likely to generate potential changes. In contrast, the instructed group initially watched the video with the intention of looking for changes and may have been less likely to make guesses.

The most significant change we included in our video was an actor change in the last two scenes. Surprisingly, not a single participant noticed this change, even after viewing the film twice and being fully informed of potential changes. At the conclusion of the experiment, the subjects were shown two pictures of the actors. They were informed of the actor change and asked to identify the second actor. Only 15% were able to identify the actor based on her face, and another 10% happened to guess correctly. These percentages are comparable to the 25%

who guessed incorrectly, indicating that the subjects were performing at chance levels. The remaining 50% were able to correctly identify the actor based on some other detail (shirt, watch, etc.) (see Figure 4).

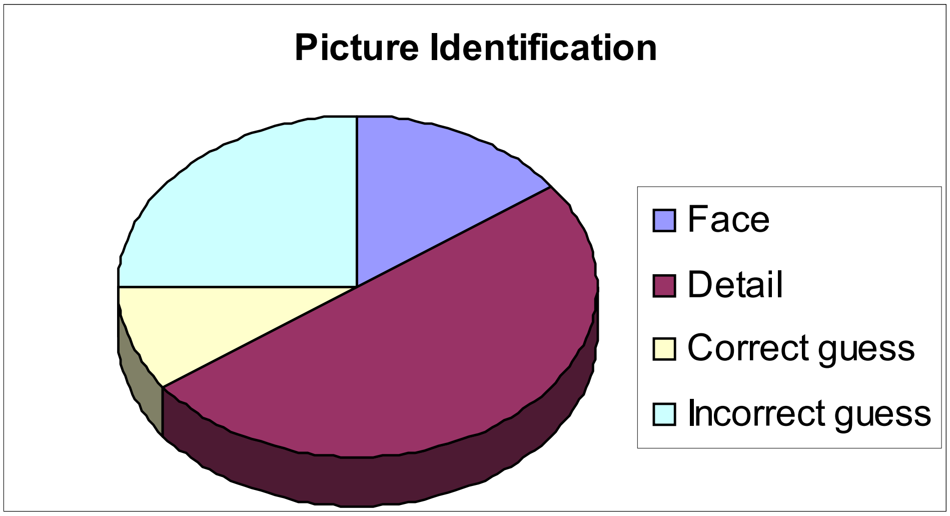

Figure 4.

DISCUSSION

This research project explored the effects of instruction and varying size of detail change on the phenomenon of change detection blindness [4]. We hypothesized that in the first trial, the instructed subjects would significantly outperform those without instruction. We also believed that an object's role in changing the overall appearance of a scene would determine its perceptibility. For example, every subject noted the change in position of the picture on the wall. Inversely, very few subjects noticed the removal of the necklace from the character's neck. As the picture was a large object with extreme color contrast to the wall, its inversion was more conspicuous than the white necklace's disappearance and reappearance on the character's neck.

When analyzing data, it is important that we consider the knowledge of the observer. If the observer knows which aspects of a scene are important, attention will be paid to the appropriate items at the appropriate time. For instance, a study was done on the ability of observers to detect changes in scenes from American football games [2]. Upon comparing the performance of experts and non-experts, it became evident that the experts were able to discern the changes in the scenery with much more precision and efficiency than non-experts. With this in mind, the results of testing may be skewed by the level of education, expertise, and experience one might have in a specific subject matter.

The results of our first trial yielded data that supported both the above study and our initial hypothesis. The uninstructed group encoded only 16% of the changes made, while the instructed individuals only perceived 22%. At first glance, these figures seem insignificant; however, the instructed group was able to attend to 37.5% more alterations than the uninstructed group. Thus, it is evident that instruction plays an important role in the study of change-detection blindness.

In the second trial, the instructed subjects noticed an additional 29% of the changes, while the uninstructed individuals noticed 32% more changes. These data are conclusive with the fact that the uninstructed group had more changes to find in the second viewing. However, when we sum the percentages of each trial, the instructed group noticed 51% of all changes, while the uninstructed subjects ascertained only 48% of the changes, a difference of only 6.25%. This close proximity of results may be due to the complexity and richness of the plot and scenery and the small sample size.

In a past study, Rensink proposed that it is more likely that subjects will notice changes to central figures and objects that are essential to a scene's plot [1]. Rensink found that people were less likely to acknowledge changes that occurred in the surroundings. In contrast to these data, our results suggested that the subjects had difficulty in identifying changes to central figures. For example, the actor change went entirely undetected in each viewing, and subjects were at chance performance in identifying the correct actor when provided with 2 photos. Meanwhile, participants readily noticed changes to the periphery, in particular the inverted picture. A possible explanation for the above dilemma revolves around the focus of participants and involves various studies of gap-contingent change, which suggest that only four items can be monitored at a given time. If the focus of the observer is on the surroundings (as in the instructed group), it is likely that they will miss changes made to the main characters. Therefore, if one watches a complex, realistic visual scene, it is likely that details will not be noticed.

As mentioned earlier, sample size plays an important role in data collection. In this study, our sample size not only reduced our power to detect significant differences between means, it also may have been a source of error, as the examination consisted of a total of 20 participants. Each group was composed of only 10 volunteers. Ideally, we would have liked to have a minimum of 40 participants (20 in each test group). Due to the small sample size, there may have been biases, and the data may not have accurately represented the population. Therefore, we are not able to extend our conclusions from this study, such as women performing better then men or the older generation being more perceptive than the younger, to the general population.

The most significant change in the video clip was the actor change. In another portion of this experiment, viewers were presented with images of the first and second actors. Despite the fact that none of the test subjects were able to distinguish between the two during the media clip, 10% of the participants correctly guessed which image represented the second actor. This supports an interesting point that has been studied in the past, "forced-choice guessing" [1]. In previous studies, subjects were able to guess the location of a change more often than not. This occurred even when there was no awareness that a change occurred. Therefore, participants in this study were able to identify the second actor largely due to chance [5].

A question then arises: do we, as humans, have a poorly designed visual system? Our findings report that only 50% of the total changes in the media clip were acknowledged by viewers, even after 2 viewings. How could participants miss a significant disparity in the media clip, such as an actor change? Despite these seemingly huge deficiencies, it is possible that we have a well defined visual system. The problem may not lie in our ability to physically see change, but rather our perception of change. We have come to expect consistency in our lives and therefore have no reason to suspect such irregularities as scenery changes. Consequently, it is likely that our environment has influenced our ability to detect certain types of changes.

ACKNOWLEDGEMENTS

We would like to extend our gratitude to Professor Dolan for trusting us with his new digital camera and helping us with our team project. Special thanks to Professor Surace for finding us test subjects and to all those who participated in our study. Last but not least, we would like to thank Anna Labowsky for dedicating her time and effort in helping us finish the paper.

REFERENCES

[1] Rensink, R. A., O'Regan, J. K., & Clark, J. J. (1996). To see or not to see: The need for attention to perceive changes in the scenes. *Investigative Ophthalmology & Visual Science*, **37**, S213.

[2] Rensink, Ronald A. (2002). Change Detection. *Annual Reviews of Psychology*

[3] Balota, D. A., Cortese, M. J., Duchek, J. M., Adams, D., Roediger, H. L., McDermott, K. B., and Yerys, B. E. (1999). Veridical and false memories in healthy older adults and in dementia of the Alzheimer's type. *Cognitive Neuropsychology*, **16**, 361-384.

[4] http://www.cogsci.msu.edu/sigma/documents/cogsci_abstract1.pdf Scene Context and Change Blindness: Memory Mediates Change Detection. accessed 8/8/02.

[5] Levin, D.T., & Simons, D.J. (1997). Failure to detect changes to attended objects in motion pictures, *Psychonomic Bulletin and Review*, **4**, 501-506.

APPENDICES:

Appendix A: The following pictures are scenes captured from our video clip.

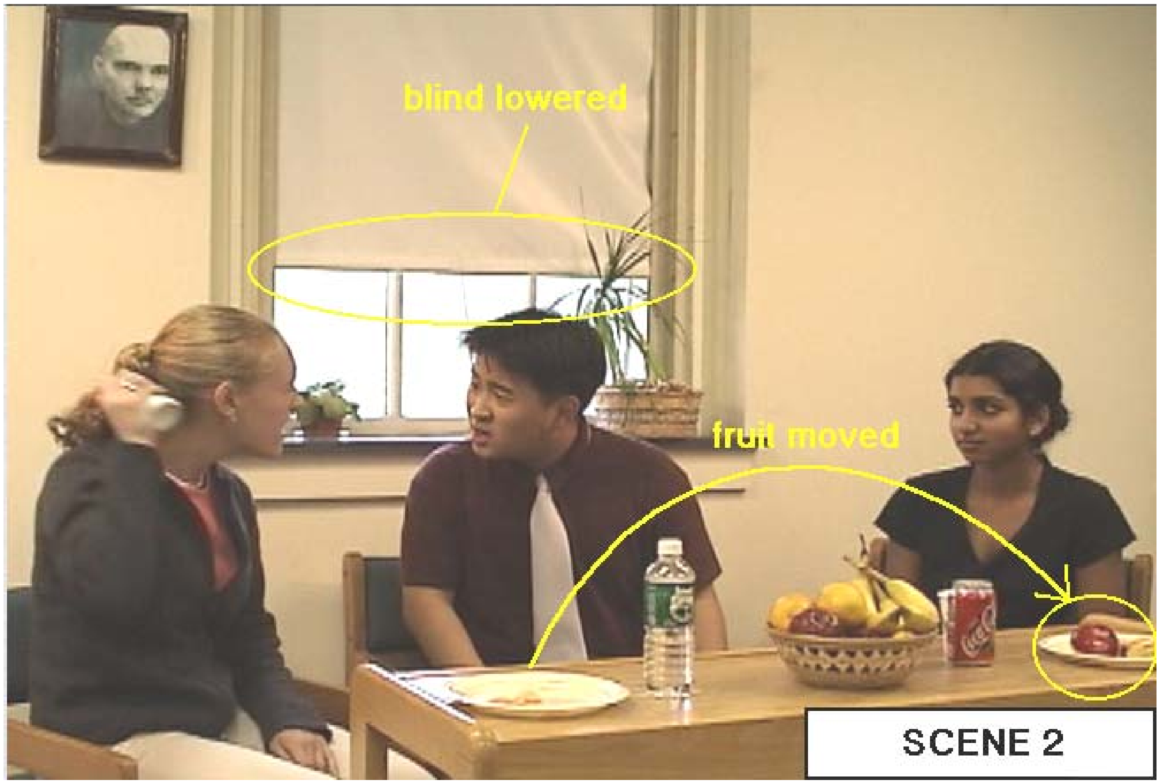
blind lowered
fruit moved
SCENE 2

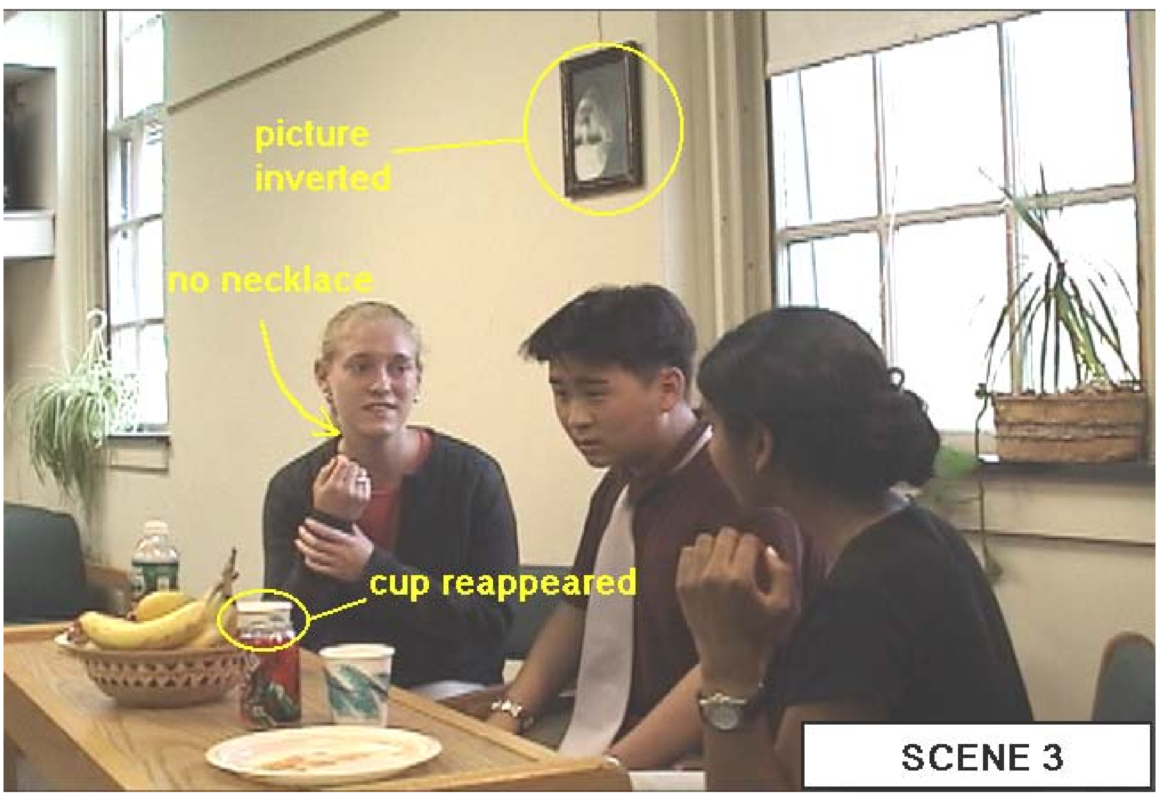
picture
inverted
no necklace
cup reappeared
SCENE 3

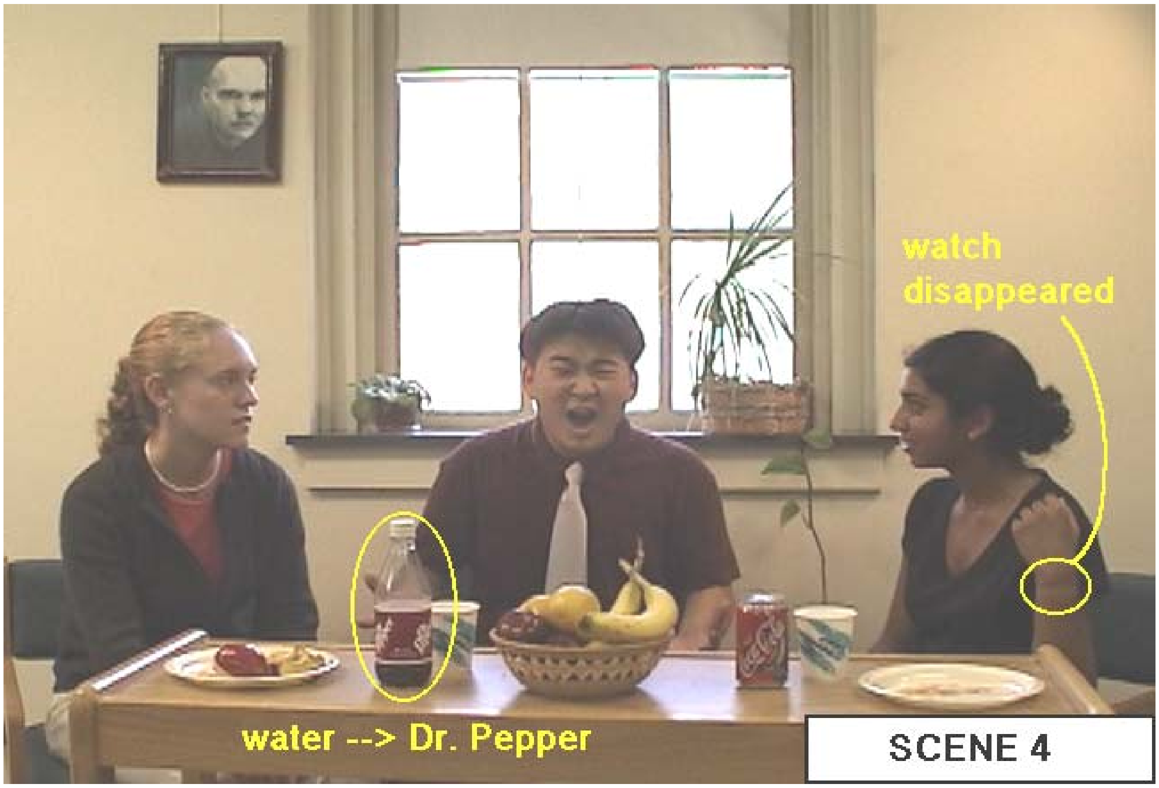
watch
disappeared
water --> Dr. Pepper
SCENE 4


actress change
plant moved
fruit
rearranged
SCENE 5

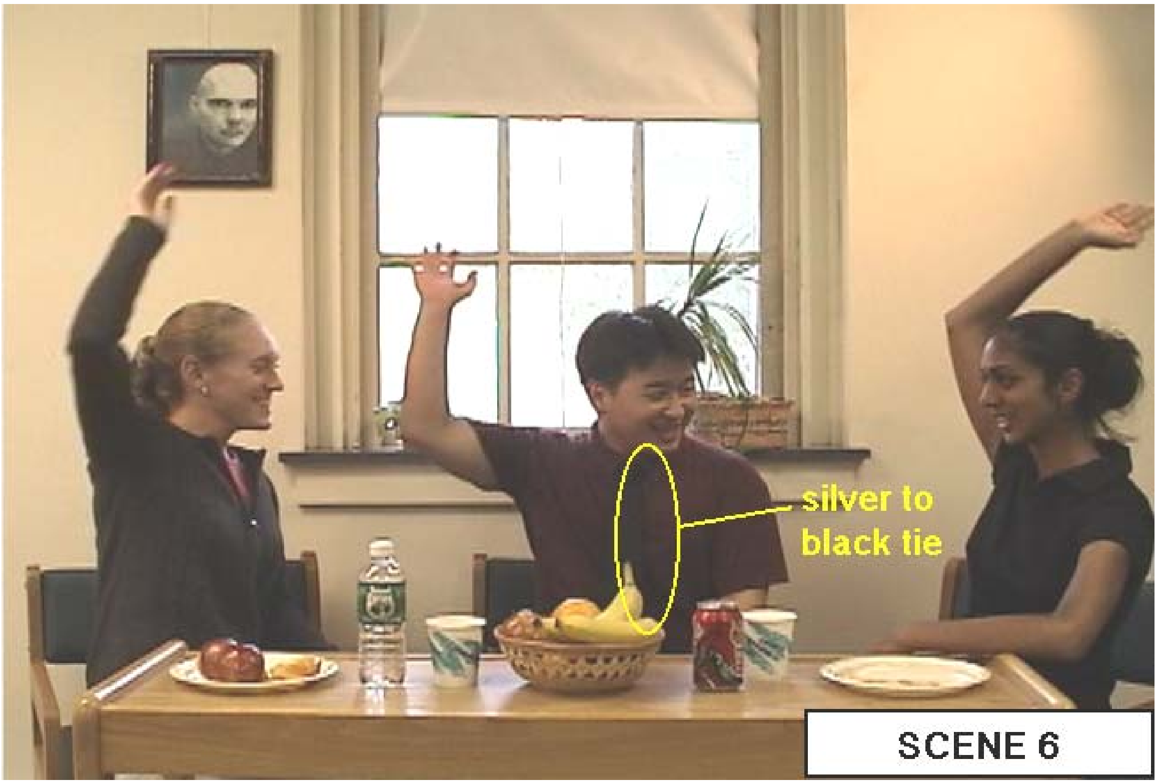
silver to
black tie
SCENE 6

Appendix B: The pictures below are used to differentiate between the two actors who changed in our clip.

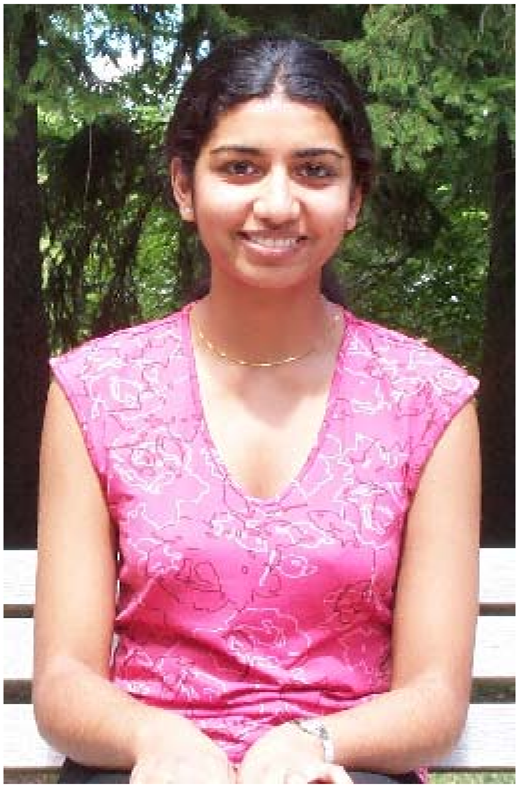

First Actor

Second Actor

Appendix C:

Consent Form

I state that I wish to participate in a program of research being conducted by the psychology students of the 2002 NJ Governor's School of the Sciences under the supervision of Professor Patrick Dolan at Drew University, Department of Psychology.

The purpose of the research is to study perception and attention. In the remainder of the study, I will be asked to view a brief video clip from a computer screen. Immediately following the viewing, I will be asked a series of questions about the film and I will answer them to the best of my ability. The research will last approximately ten to fifteen minutes.

I consent to allow my responses in the study to be recorded. I understand that my results will be used for research purposes only, and no one other than the experimenters and their immediate partners will have access to the results.

I also understand that code numbers will be used so that my name will not be connected to my responses.

I understand that the experiment is not designed to help me personally, but that the researchers hope to learn more about visual continuity.

I understand that my participation is voluntary, and that I have the right to withdraw at any time during the study without any consequences.

I understand that I am entitled to an explanation of the study upon completion of the experiment. I understand that any questions I may have at the end of the experiment will be answered. If I have any additional questions I may contact Professor Dolan.

Professor Patrick Dolan
Department of Psychology
Hall of Sciences 56
(973) 408-3558

Signature of Participant ____

Printed Name ____

Date _____________

Signature of experimenter ____

Appendix D:

Post-Viewing Questionnaire

1. In the video, did you notice any unusual differences from one shot to the next where objects, scenery, clothing, or people suddenly changed? If yes, please describe the changes you noticed.

2. Now that you have viewed the video for a second time, did you notice any unusual differences from one shot to the next where objects, scenery, clothing, or people suddenly changed? If yes, please describe the changes you noticed.

Personal Information Gender: _____ Age: _____

For Office Use Only: Correct Picture: Y / N

Appendix E:

Debriefing form

"Change Detection Blindness"

Professor Patrick Dolan,	Governor's School of the Sciences	Anna Labowsky
Psychology Department	2002 students	GSS Counselor
Drew University	Group Project: Psychology	
(973) 408-3558	Drew University	

Thank you for watching the video clip and participating in our psychology project. The study you took part in was designed to examine a subject's ability to detect subtle changes between frames of a short video. During the test, setting and characters were manipulated in a way to test the audience's response to these minor changes. For example, a coke can in one scene was changed into a vanilla coke can in the following scene. In another scene, an actress was replaced with a different but similar-looking actress. In addition to differences presented in the video, we also examined how subjects react to different instructions given at the beginning of the experiment. We hope to observe audience response to minor and major changes and to note any similarities between subjects' reactions.

If you would like more information on this topic, the following references may be of interest:

- Levin, Daniel T., and Simons, Daniel J. (1997). Failure to detect changes to attended objects in motion pictures. Psychonomic Bulletin and Review, 4 (4), 501-506.
- http://www.wjh.harvard.edu/~dsimons

If you are interested in the results of this study, you may contact Patrick Dolan (973-408-3558) at the completion of this study (August 2002). Please note that only overall results, not individual results, will be available.

If you have any questions or concerns about this study, please contact Patrick Dolan (973-408-3558).

Thank you for your participation.