EFFECT OF ISOLATED FACIAL FEATURE TRANSFORMATIONS IN A CHANGE BLINDNESS EXPERIMENT INVOLVING A PERSON AS THE OBJECT OF CHANGE

Hadar Kadosh

A dissertation submitted to the Faculty of Humanities, University of the Witwatersrand, Johannesburg, in fulfilment of the requirements for the degree of Master of Arts in Psychology by Dissertation

Johannesburg, 2007

ABSTRACT

Research has shown that people often fail to notice changes to visual scenes. This phenomenon is known as change blindness. This study investigated the effect of facial feature transformations on change blindness using change detection tasks involving a person as the object of change. 301 participants viewed a photo-story comprised of a few still frames. In the final frame, a selected facial feature of a character in the story was altered. Four different photo-stories were used, each utilising a different alteration. Questionnaires designed to determine whether the change was detected were administered. Results showed that changes to facial features considered to be more salient produced higher levels of change detection. A flicker test using the same images from the photo-story was administered to a further 75 participants and showed a similar pattern of results. It was concluded that in order to detect change, the changing stimuli have to be both salient and meaningful.

Keywords: change blindness; change detection; visual perception; attention; memory; facial features; cue salience

DECLARATION

I declare that this dissertation is my own unaided work. It is submitted for the degree of Master of Arts in Psychology by Dissertation in the University of the Witwatersrand, Johannesburg. It has not been submitted before for any other degree or examination in any other university.

____

Hadar Kadosh

________day of ______________________, 2007.

ACKNOWLEDGEMENTS

I would like to express my deepest and most heartfelt thanks to the following: The Department of Psychology at the University of the Witwatersrand for approving my proposal and allowing me to pursue this research; Cameron Peter Tandy and Gary David Chenick for allowing me to use their photographs in my experimental stimuli; Natalie Urban for working her magic on my photo transformations; Sharon Shamir, Nicky Israel, Natalie Zahringer, Denise Newfield, James Braid and Dr. Phil Ferrer for generously assisting me in my quest for participants; Tomer Kadosh, who truly went beyond the call of brotherly duty, for helping me facilitate all my experiments; Mike Greyling for all his help with the statistical analyses; a special thanks to my supervisor, Michael Pitman, for all his time, guidance and insight; and last but not least, my thanks to the kind and obliging students of the University of the Witwatersrand for participating in this study.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1 OVERVIEW

This study aims to investigate the effect of facial feature salience on the change blindness phenomenon in a change detection task involving a person as the object of change. The change blindness phenomenon was selected for investigation due to its striking and surprising nature, its consistent popularity in the field of cognitive research, and the fact that while change blindness experiments have been extensive and able to demonstrate the robustness of the phenomenon across a wide variety of scenarios and tasks, a number of issues remain open to investigation. Among these is the question of what draws attention to some scene elements and not others in change detection tasks. The primary motivation for the present study was to set up an experiment that may begin to address this question.

In Chapter 2, change blindness is defined, with particular reference to the various definitions of change and how these can guide the investigation and understanding of change blindness, both theoretically and operationally. A distinction is drawn between change blindness and inattentional blindness, as these two terms are sometimes used interchangeably in the literature.

This is followed by a review of the various studies conducted on change blindness, with its simple-stimuli centred beginnings in laboratory experiments in the 1950s to the more elaborate and naturalistic experiments used in contemporary change blindness research. This review shows that change blindness can be demonstrated in a broad array of scenarios and with varying stimuli, fuelling interest in the phenomenon as a possible platform for deepening an understanding of visual processes, attention and memory. The role of attention and memory in change blindness is explored in the section that follows, and various theories and studies that attempt to explain the change blindness phenomenon are discussed.

Chapter 3 delineates the motivation for the use of faces as the chosen stimulus in this study. The nature of the face as an object likely to receive preferential attention is discussed, and differences between the stimuli employed in this study and those used in other change blindness experiments are outlined. This is accompanied by an overview of the mechanisms involved in facial recognition, and how these may differ from mechanisms involved in the recognition and processing of non-face objects. Given that this study isolates particular facial features in the hope of establishing a relationship between feature salience and the resulting level of change detection, a section of this chapter addresses experimental and theoretical discussions suggesting that facial features are a legitimate and sensible unit of perception that can be investigated operationally. A review of studies on the salience of isolated facial features is also included.

Chapter 4 outlines the motivations for and aims of this research. Chapter 5 describes the methods employed by the present study, including a detailed description of the designs, samples, instruments, procedures, and analyses used in the experiments. In Chapter 6 the results of the experiments are presented.

Finally, Chapter 7 contains a discussion of the findings outlined in Chapter 6 and the relevance of these in view of the arguments and theoretical issues highlighted in the literature review. The various limitations of the study are discussed and recommendations for future research which may arise from these limitations are included.

2 CHANGE BLINDNESS

2.1 Change Blindness Defined

Change Blindness can be defined as individuals' inability to detect changes that occur between views in their visual fields (Levin, 2002). Studies have shown that the change blindness phenomenon occurs in a broad variety of situations, ranging from the inability to detect a changed object (such as letters, numbers or pictures of objects) in a visual array of a few items (Pashler, 1988; Simons, 1996), delays in detecting changes in still photos that alternate between the original and a slightly changed image, at short intervals (Rensink, O'Regan & Clark, 1997), failure to detect between-view changes in motion pictures (Levin & Simons, 1997) to the inability to detect between-view changes in real world situations (Simons & Levin, 1998; Levin, Simons, Angelone & Chabris, 2002).

2.1.1 What is Change?

Rensink (2002) emphasises the importance of identifying distinctions in the definition of change that not only guide designs of experiments in change blindness, but may also point to the different underlying neural processes involved in both change detection and the inability to detect certain changes. The first distinction is between change and motion. Change can be defined as the transformation of a well defined, enduring structure, while motion refers to a change of position over time. In other words, motion can be thought of as variation referenced to location, whereas change is referenced to structure. The second distinction involves distinguishing between change and difference. If change, as defined above, refers to a single structure which undergoes some transformation over time, then difference refers to a lack of similarity in the properties of two structures (Rensink, 2002). Thus, in order to say that a change has been detected (or not detected), it is imperative that the phenomenological experience of the observer is that only a single structure, or single set of structures were involved. The third distinction is between detecting dynamic change, which involves seeing the change in progress, and detecting completed change, that is, seeing that something has changed. Watching a person unscrew and remove the lid of a jar, for example, is seeing dynamic change: the entire process of transformation from lidded to open jar is observed. Observing a closed jar on the counter top, leaving the kitchen, and returning later to find the same jar without its lid is an example of detecting completed change.

As noted, these distinctions can help guide the design of change blindness experiments, as the design has to ensure that the detection of change is not the result of the detection of motion, nor is it phenomenologically perceived as the detection of difference by the subject. Decoupling change and motion, Rensink (2002) asserts, is achieved by making the change contingent on some event. Various contingencies of change can be employed and some are discussed in the following sections. This study employs a 'cut-contingent' method, which generally refers to motion pictures, and where the change is made during a cut from one camera angle to another. In this particular study, still images are employed as opposed to motion pictures, but the flowing narrative of the images and the alternating camera angles mimic those normally used in films, thus allowing the type of change used to be labelled cut-contingent change.

2.1.2 Change Blindness vs. Inattentional Blindness

An important distinction worth highlighting is that between change blindness and inattentional blindness. Rensink (2000a) points out that change blindness (i) denotes the failure to report the presence of significant changes in the visual input under particular experimental conditions, (ii) concerns itself with transitions between quantities, as opposed to the simple presence of quantities, and (iii) pertains only to subjective visual experience (or qualia) as opposed to other aspects of seeing, such as visuomotor response or physiological response (for example, differential activity of various brain regions). In contrast, inattentional blindness pertains to not seeing the presence of an object/stimulus at any moment in time, and not necessarily the perception of change. An often quoted example of inattentional blindness is the 'gorilla' experiment conducted by Simons & Chabris (1999). Participants viewed a video where two small teams of people passed a basketball to one another. Participants were asked to mentally (silently) count the number of basketball passes that were made. In the middle of the video a woman dressed in a gorilla suit walked onto the scene, beat her chest and walked off. Participants were then asked a series of questions designed to determine whether they noticed this had occurred. Forty-six per cent of participants did not notice the unexpected event.

Although inattentional blindness may be found to occur with a changing stimulus input, the input does not need to be experienced as changing. The input can simply be experienced as different images shown in succession, without a 'history' or narrative to connect them. Simply put, inattentional blindness refers to the failure of perception of an input in the here and now, whereas change detection requires memory or comparison operations. This, Rensink (2000a) argues, shows that while these two phenomena may be related, there may be different mechanisms responsible for each.

2.2 Change Blindness Experiments: Past and Present

As early as the 1950s, researchers began developing an interest in change detection tasks and change blindness (Simons & Ambinder, 2005) with studies revealing people's failure to detect changes to simple arrays of dots and letters when these were presented sequentially and separated by a brief blank interval. During the 1970s and 1980s this trend continued with research by Pashler (1988) for example, which demonstrated people's inability to detect a changed object (letters, numbers or pictures of objects) in a visual array of a few items. In the 1990s change blindness was further demonstrated in a series of experiments by Ron Rensink and his colleagues (Rensink, O'Regan & Clark, 1997). Some of these tasks were inspired by findings that changes introduced during eye movements often go undetected, referred to as saccade-contingent change blindness. Other tasks used still photos that alternated between the original and the slightly changed image at short intervals, such as photos of a jumbo jet which alternate between an image of the plane with an engine and without an engine. In such tasks participants showed significant delays in detecting the changes (Rensink et al., 1997).

While it may not be very surprising that changes to unattended objects in the visual field may go undetected – most people are aware that so called 'continuity errors' in films are often not noticed by the audience – it is interesting that the change blindness phenomenon extends to attended objects as well. Levin and Simons (1997) conducted a study where participants viewed a short video depicting an actor who heard a ringing phone and stood up to answer it. As the actor approached the phone, the camera angle changed and the actor was replaced with another actor. Two thirds of participants failed to notice the switch.

The change blindness phenomenon was even demonstrated as taking place with attended objects in real-world scenarios. In the famous 'door experiment' by Simons and Levin (1998) an experimenter approached individuals on the campus of an American university and asked for directions to a building on the campus. While the participant was giving directions, two more experimenters carrying a door stepped between the first experimenter and the participant, and the first experimenter switched positions with one of the door-carriers who stayed behind to continue the conversation. About half of the participants did not notice that their conversation partner had changed. Simons and Levin (1998) admitted that part of this failure to detect the change may have been due to the fact that participants were concentrating on a task, giving directions, that involved looking at a map and/or at the direction of the building, and the fact that an unnatural disruption took place – the door-carrying experimenters chose to walk directly between two people who were clearly conversing.

To address such concerns, Levin, Simons, Angelone and Chabris (2002) devised a study where the participants would not be distracted by a task while the change took place and the change would occur in a less disruptive and more natural way.

Individuals were asked to participate in a study and they approached a counter in a building behind which stood an experimenter who presented them with an informed consent sheet to sign. When participants returned the sheet to the experimenter, he crouched down behind the counter to “get them some forms”, and another experimenter stood up in the same place the first experimenter had crouched down. The second experimenter handed the participant a questionnaire designed to determine whether the switch was detected. Seventy five per cent of participants failed to detect the change.

Such findings go against most people’s intuitive belief that they would be able to detect such changes (Levin, Momen, Drivdahl, & Simons, 2000; Levin, 2002). When a number of possible visual change scenarios were described to a group of psychology students, 90.5% of those questioned thought they would detect an actor’s changing scarf across shots in a video, 69.5% felt they would detect a change of actors across views in a video and 97.6% thought they would detect a change of conversation partner in a real-world scenario (Levin et al., 2000). These estimations exceed by far the change detection by participants in experiments. Levin (2002) refers to this intuitive inclination to think one would detect such changes as change blindness blindness.

It also appears as though knowledge of the change blindness phenomenon, and specifically knowledge of person-change studies, has no impact on participants’ ability to detect a switch in conversation partners. In the counter experiment by Levin et al. (2002) described above, six of the participants recruited for the study had, less than one hour earlier, (unbeknownst to the experimenters) attended a psychology seminar where the door study (Simons & Levin, 1998) had been described and discussed. Four out of the six still failed to detect the change.

Despite their variations, what all the above-mentioned experiments have in common is that the change signals, or transients associated with change, are masked or otherwise impaired (Simons and Rensink, 2005). This is significant and is explored in further detail in the following section.

2.3 Change Blindness Explained?

Various possible explanations of the change blindness phenomenon have been put forward, and these often depend on the type of change detection task being explored. For example, in saccade-contingent change detection tasks, some explanations suggest that those mechanisms that generate a movement of the eyes actively suppress perception during the eye movement (Simons & Ambinder, 2005), while other explanations refer to the observation that rapid movement of the eye produces a blur on the retina, and that this blurring may mask the change signal (Simons & Ambinder, 2005). Explanations pertaining to the type of change blindness experiments that are the focus of this study include the possibility that people fail to represent sufficient information about pre- and post-change objects or people (Levin et al., 2002), that people do form detailed representations of pre and post change objects or people, but are unable to access them consciously (Levin et al., 2002), or that representations are both effective and consciously accessible, but people simply fail to make cross-view comparisons and therefore fail to detect the changes (Levin et al., 2002; Mitroff, Simons, & Levin, 2004).

2.3.1 Visual Experience

All the abovementioned experiments fuelled the idea that our visual experiences are not as richly detailed as we might assume. In fact, such findings led some theorists to assert that our visual world is in fact a 'Grand Illusion' (for example, Dennett, 1991 and Ballard, 2002). Noe (2002) attempts to summarise the grand illusion theory by saying that according to this theory our beliefs about what our perceptual experience is like are radically false, and that in fact our perceptual consciousness is a sort of confabulation. Clark (2002) extends this explanation of the grand illusion by expanding on what this confabulated consciousness consists of. Namely, that we experience a rich stream of highly detailed visual information covering an extensive part of our surroundings at any given moment in time. According to the theory, however, this is not the case. Our brains do not create such richly detailed models of visual scenes, so the 'seemingness' described above is an illusion.

A number of theorists have suggested that, although some aspects of the notion of the grand illusion contain convincing kernels of truth, the grand illusion position is too extreme. The discussion now turns to arguments presented by three such theorists, Alva Noe, Andy Clark, and Arien Mack. Their positions are described, compared and contrasted, with specific reference to how these might shed light on the causes of change blindness.

Noe (2002) suggests that the grand illusion theory emerged as a response to the more orthodox conception of experience which holds that it is understood that the sensory visual information we have access to via the eye is limited. This is based on the observation that our visual input is made up of two tiny, distorted and upside-down retinal images, that our eyes have a foveal region that is relatively small, outside of which the retina is somewhat colour-blind, and that the eye is in nearly constant motion. The orthodox conception holds that despite these limitations, our brain is able to integrate all this sketchy information into a stable and detailed representation. However, work in perceptual psychology led to a rethinking of this orthodox conception, and the grand illusion theory was born. As an example of the new kind of thinking, Noe (2002) refers to Dennett's (1991) discussion of the blind spot (the place on each retina where there are no photoreceptors). When looking at an image or scene or picture through only one eye, a part of this visual scene inevitably falls on the blind spot. So, in essence there is a gap in our visual experience. The fact that we don't experience this gap as a gap – the fact that our visual experience appears to be continuous even though it is not – Dennett argues, shows that we are at the mercy of an illusion of visual consciousness. For a demonstration of the effect of the blind spot, see Figure 2-1 below. Cover your right eye with your hand and focus your gaze on the smiley face. Hold the page approximately twenty centimetres from your face. The heart should disappear, as its image has been forced to fall on your blind spot.

Figure 2-1: Blind Spot Demonstration

♥ ☻

Noe (2002) counters that while it is true that we do not experience the gap, it does not follow that we are under the illusion that we are experiencing the entire visual scene all at once. He gives the example of shutting one eye and looking at a wall.

Even though it seems to one that there is an unbroken expanse of wall, it does not seem as though one could experience the whole of the wall's surface with a single fixation. Rather, the wall's presence appears to be constant, but the details of the surface of the wall are accessed by looking at one bit of the wall or another. According to Noe (2002), we do not take our experiences to be as full as grand illusion theorists might suggest we do. It does however seem to us, Noe (2002) argues, that we have perceptual access to a world that is richly detailed and complete, because we do in fact have this access. This forms the basis of Noe's sensorimotor approach to perception (O'Regan & Noe, 2001).

The basis of this perceptual access, according to Noe (2002), is the possession of sensorimotor skills. These skills consist in the knowledge of how one relates, via the rules of sensorimotor dependence, to one's immediate environment. So, returning to the question of change blindness, according to this sensorimotor skills theory, we know that we possess both the ability and the know-how to attend to particular aspects of a scene, at will. Change blindness results when we attend to a certain aspect, and therefore only see that aspect of the scene. But Noe (2002) argues that it does not stop there, since, phenomenologically speaking, we possess a perceptual awareness of some unattended aspects of the scene as well.

This, too, Noe (2002) claims, can be accounted for by the sensorimotor skills theory, if one takes into account the notion of perceptual presence. Perceptual presence is illustrated using the example of a cat sitting behind a picket fence. One senses the presence of the cat (the whole cat) despite seeing only those parts of the cat not hidden behind the fence. This is not just a thought or inference that the entire cat is present, but rather the perceptual experience that the cat is present. This is a form of amodal perception, that is, genuinely perceiving what is in fact out of view. This mode of perceiving results from two types of knowledge: the knowledge of what cats are like, and the (sensorimotor) knowledge that if you move your head or walk towards the cat, the parts previously hidden will come into view.

In his discussion of a number of theories that arose, either in response to, or containing explanations for, the findings of work on change blindness and inattentional blindness, Clark (2002) addresses both the question of the grand illusion, and the senorimotor skills theory described above. Like Noe, Clark (2002) argues that the 'seemingness' of seeing is not the grand illusion that some theorists would have us believe. Clark (2002) commends the sensorimotor skills theory for successfully and elegantly dispelling the "miasma of Grand Illusion" (Clark, 2002, p. 194). He argues, however, that the theory developed by O'Regan and Noe (2001) ties the nature of conscious visual experience far too closely with specifically sensory engagement with the world, and terms this strong view 'Sensory Chauvinism'.

Clark (2002) suggests a different way of unpacking the theory and its critique of the grand illusion argument. This entails placing less weight on the dependence of conscious experience on sensorimotor contingencies, and more weight on the role of semantic and episodic memory in conscious experience. In his discussion of the reasoning behind this argument, Clark (2002) refers to the work of Milner and Goodale (1998). Milner and Goodale's (1998) claim is that, in the organisational structure of the brain, there exist two structurally and functionally distinct kinds of vision: a visual system for action, and a visual system for perception. The former, which represents a much earlier evolutionary development, concerns the visually-based control of fine motor action in the here and now, and is identified with the dorsal visual-processing stream which leads to the posterior parietal cortex, and the latter, a far more recent evolutionary development, concerns the knowledge and memory-based aspects of deliberate actions relating to semantics, planning and communication, and is identified with the ventral visual-processing stream which leads to the infero-temporal cortex.

While Clark (2002) comments that Milner and Goodale's discussion of the two visual streams perhaps overstates the extent to which these work in isolation from one another, and discusses a few slightly weakened accounts of the dual visual systems hypothesis, he does use it as a springboard for the argument that the link between consciousness and action involves memory systems just as much as it involves motor systems. Conscious experience, Clark (2002) argues, forms the base for reasoned action selection, and thus requires memory systems to mediate between sensory input and action.

Simply put, O'Regan and Noe's (2001) theory of sensorimotor skills pitches implicit knowledge of a set of possible actions and results at the level concerning head and eye movements and the possible resulting distortions of the retinal image, but Clark (2002) argues that in the case of conscious visual perception this implicit knowledge is pitched at the level of types of action and types of objects (that is, the semantic categorisation of these). It is the latter that appear to underpin visual experience, Clark (2002) explains, as the errors demonstrated by phenomena such as change blindness help to show. It must be noted, however, that while the emphasis of the sensorimotor skills theory may be on sensorimotor contingencies, Noe (2002) still refers to the important role of semantic knowledge in conscious visual experience as mentioned in the cat behind a fence example above (that our perception of the semi-obscured cat depends on our understanding of cats, what they are and how they are).

Mack (2002), like Noe and Clark, dismisses the grand illusion theory, suggesting that the real issue in question is not whether all that we think we see is an illusion, but rather, how much of what is before us do we see and why we have the phenomenological impression that we see it all. Mack (2002) points to the fact that the findings of explorations of change and inattentional blindness have highlighted the need for such questioning. Additionally, like Noe and Clark, Mack (2002) concedes that it appears that attention is needed in order to see and that we do not attend to everything all at once, therefore we cannot see everything all at once. The reason we get the sense that we see much of what there is to be seen, Mack (2002) argues, is because our attention is often broadly distributed, and not narrowly focused (as it is in inattentional blindness experiments of the sort described in the previous section).

At this stage, Mack's (2002) argument appears to fit in quite neatly with that of Noe (2002) and O'Regan and Noe (2001), in that we get a sense of seeing a large portion of what there is to be seen at any given moment because under normal circumstances our attention is distributed in such a way as to give us the gist of the scene – and with it the impression that we processed the whole scene. Where these theorists begin to part ways is with Mack (2002) raising the problem of processing and storing information that we are not consciously aware of. Mack (2002) argues that O'Regan and Noe's (2001) theory is not able to explain why it is that a lot of what we neither attend to nor see is still represented in memory (although it cannot be recalled consciously). Evidence supporting this comes from demonstrations of priming in change blindness and inattentional blindness experiments. Mack (2002) suggests that the reason we process and store unattended input is so that we can access its meaning and attend to it if necessary. To support this claim, Mack (2002) uses as an illustration the findings of Mack and Rock (1998, in Mack, 2002) in their analysis of inattentional blindness which shows that when attention is strongly focused on an aspect of a scene, the observer will be blind to everything else in the scene unless there is something present in the unattended aspects of the scene that is both very salient and meaningful which will capture attention and cause its detection. Interestingly, Mack (2002) points out that the stimuli most capable of achieving this include the observer's own name and the iconic representation of a happy face.

Taken together, the claims of these three theorists present a picture of visual experience that suggests: (i) that we are not subject to an illusion that we believe we see a full and richly detailed, snap-shot like picture of the world when in fact we do not, (ii) that we have perceptual access to an entire visual scene since we possess knowledge, in the form of sensorimotor skills, that can bring the various aspects of a scene into conscious awareness at will (Noe, 2002), (iii) that in addition to sensorimotor contingencies, we also have memory-based, semantic knowledge of types of actions and types of knowledge that inform our visual consciousness (Clark, 2002), and (iv) that while our visual attention is mediated by both these forms of knowledge, we also process and encode unattended information, so that we may attend to it if we deem it both salient and meaningful enough. These descriptions of the nature of visual experience manage to explain both why change blindness occurs - we only see attended aspects of a scene and only attend to some aspects of a scene at once – while still allowing for the fact that we have the experience of seeing a lot more than we do.

The discussion will now turn to the more specific role of attention and its mechanisms in the context of change blindness. This relationship is delineated explicitly by Rensink's (2000b) coherence theory of focused attention.

2.3.2 Coherence Theory

Rensink, O'Regan and Clark (1997) assert that focused attention is needed to see change. Normally in real world situations, the argument goes; change is accompanied by a motion signal which will attract attention to the location of the change. If the motion signal is too weak, or some other factor interferes with or masks the motion signal (such as the eye saccades, photo flickers or motion picture cuts discussed above) insufficient attention will be drawn to the signal resulting in change blindness. This forms the basis of Rensink's (2000b) coherence theory of focused attention.

Within this theory, change blindness is understood to result from the masking or neutralisation of the motion signal which accompanies a particular change, which would otherwise draw attention to the location of the change (Rensink, 2000b). It asserts that attention cannot integrate visual features into a detailed and long lasting coherent representation, but rather that a structure/object's coherence lasts only as long as focused attention is directed towards it. Coherence in this sense refers to spatial and temporal consistency, that is, that representations at different spatial locations are understood to refer to the same object and that representations at different times are understood to refer to the same object.

Rensink (2000b, 2001, 2002) breaks the theory down as follows: (i) Before attention is focused, structures known as proto-objects form rapidly and in parallel across the visual field, with limited spatial and temporal coherence. These structures are volatile and constantly regenerated, which means they are easily replaced by any new stimuli appearing in the same location on the retina. (ii) Focused attention is described as a metaphorical hand that grasps a few of these proto-objects, allowing them to be perceived as distinct objects with a high degree of spatial and temporal coherence. This coherence allows the object to retain its identity across interruptions, which means it is understood to transform into, as opposed to being replaced by, any new stimulus that arrives in the same retinal location. (iii) The moment focused attention is averted however, the object loses its coherence and reverts back to being a proto-object.

If no more than a few coherent object representations can exist at any one time, one might ask how complex scenes can be represented and how it is that people generally have the impression of observing a large number of coherent objects simultaneously. Rensink (2000b, 2001, 2002) responds to such questions by asserting that instead of a detailed and coherent representation of all objects in a scene being created, such representations are only created for those few objects needed for a given task. Whenever requested, attention will form a coherent representation of an object or any of its aspects, resulting in a 'virtual representation' of a scene that appears as if all objects have detailed and coherent representations simultaneously.

The above discussion pertains mainly to the theoretical and philosophical discussions surrounding visual consciousness and perception, which may inform an understanding of the change blindness phenomenon at a more general, theoretical level. There have, however, been a number of empirical attempts to uncover and tease out, more specifically, the potential causes and underlying mechanisms of perception that may explain the occurrence of change blindness. The present study is conceived of as belonging to this latter category.

2.3.3 Experimental Attempts to Uncover the Causes of Change Blindness

Contemporary research includes a number of experiments that attempt to uncover some of the underlying causes of change blindness, such as a study by Levin, Simons, Angelone and Chabris (2002). The study replicated earlier door studies, but in addition to establishing levels of change detection, participants were also required to try and identify the pre-change experimenter from a photo line-up. Line-up accuracy was found to be poor for participants who failed to detect the change, but significantly greater and above chance for those participants that did notice the change.

Noting that while this may serve as evidence supporting the possibility that the pre-change object was insufficiently represented, it does not rule out the possibility that while information about the pre-change experimenter *was* sufficiently represented, it was then *overwritten* as a result of exposure to the post-change experimenter. Levin, Simons, Angelone and Chabris (2002) included a further experiment to investigate this possibility. As before, some classic change blindness scenarios were carried out, but in this instance participants were asked to select both the pre- and post-change experimenters from photo line-ups. If overwriting had occurred, one would expect that participants would be better at selecting the second experimenter than at selecting the first. Results showed that this was not the case: identification frequencies of pre- and post-change experimenters were similar.

A study by Mitroff, Simons, & Levin (2004) was conducted to strengthen the claim that change blindness can result from a failure to compare retained representations of both pre- and post-change objects. A series of experiments were set up where participants were shown an array of objects on a computer screen, which then disappeared and reappeared with one of the objects potentially replaced by a new one. In addition to reporting on whether or not they detected a change, participants were required to answer a series of two-alternative forced-choice questions, where two objects were shown and the participants had to select the one they thought had been displayed in the trial. The questions included choices between the pre-change object and a novel object, the post-change object and a novel object, and an unchanged object and a novel object. These three were presented in random order, and in trials where there was no change, all three questions included an unchanged object and a novel object. The experimenters reasoned that in cases where no change was detected, if participants were able to pick out both the pre- and post-change objects at levels greater than chance, this would support the comparison failure hypothesis. Results did indicate that retained representations of pre- and post-change objects were present even under conditions of change blindness (Mitroff et al., 2004). They noted, however, that this does not mean that the representations formed were necessarily perfect or complete, but rather that they contained sufficient information to potentially support change detection.

Similarly, a series of experiments by Varakin and Levin (2006) were carried out to test the relationship between change blindness and visual memory: specifically, whether good visual memory for stimuli can be demonstrated in the context of poor change detection. For example, participants viewed arrays of line drawn objects on a computer screen. A blank screen flicker appeared and the arrays were displayed again, with one of the objects potentially changed. Both before and after the screen flicker, all the objects in the array were cued in random order: each object was highlighted by two red bars appearing above and below it, in order to ensure that attention had been focused on both pre- and post-change objects. Following the post-flicker cuing sequence, forced choice change detection questions appeared on the screen. Participants responded 'yes' or 'no' to the question of whether they had noticed any objects change and had to select the location on the screen where the change had taken place. These trials were followed by a surprise recognition test, where two versions of arrays that were used in the no-change trials were presented simultaneously on the screen, one of which contained one token-level substitution of an object. Participants had to select the array they had seen previously. For the tasks involving incidental (that is, unexpected) change detection, change detection levels were equivalent to chance, whereas long term recognition levels were found to be above chance. This led Varakin and Levin (2006) to conclude that change blindness can be demonstrated in situations which lead to good visual memory.

While it may seem contradictory that change blindness can occur even when we are able to form lasting visual representations, Varakin and Levin (2006) argue that it is important to distinguish between the visual system's ability to store representations in memory and its tendency to use such representations to monitor visual information's consistency over time. As Varakin and Levin (2006) put it: "it is important to distinguish between what people *can* do and what they *generally* do" (Varakin & Levin, 2006, p.52). Thus, while the capacity exists to retain a great deal of visual information and, depending upon the nature of the task or the intention of the observer, this capacity can be demonstrated, we do not necessarily utilise this capacity to detect changes to visual scenes. This conclusion seems to implicitly support the comparison failure hypothesis, since the results demonstrated good long term memory for both pre- and post- flicker arrays, and the researchers used visual cues in the form of highlighting objects within arrays to ensure that the changing objects would receive attention. One criticism of the Varakin and Levin (2006) study is that all the arrays used in the recognition tests were selected exclusively from the no-change trials. Although Varakin and Levin (2006) were careful to conclude that good memory for objects can be demonstrated in the same *scenarios* that lead to change blindness, as opposed to claiming that good memory can be demonstrated for pre- and post-change objects, it is possible that the good memory demonstrated for the arrays in the recognition tests was the result of longer visual exposure, since in no-change trials the arrays were present on the screen twice as long as the pre- and post-flicker arrays in the trials where an object was changed.

The various theories of visual experience described above (O'Regan & Noe, 2001, Noe, 2002, Clark, 2002, & Mack, 2002), Rensink's (2000b) coherence theory, and the evidence from studies such as the one by Levin et al. (2002) – which were driven wholly, or in part, by the findings of the various change blindness experiments outlined in section 2.2 – can be taken together to begin to form a better understanding of why the change blindness phenomenon occurs, and what this tells us about normal vision.

The present study fits into the body of work relating to the empirical investigation of the possible explanations of change blindness, and could contribute to an understanding of the phenomenon in a number of ways. Firstly, at the simplest level, the findings of this study may re-demonstrate the change blindness phenomenon itself, thus contributing to the existing body of work that shows the robustness of this phenomenon. Secondly, it was noted above that there exists the general understanding that attention is needed in order to see change (Rensink, 2000b, O'Regan & Noe, 2001, Noe, 2002, Clark, 2002, & Mack, 2002). What remains open to investigation, however, is what aspects of a particular scene are likely to draw attention in a change detection task. The present study aims to address this question empirically, via specific manipulation of selected experimental stimuli. It is hypothesised here that the more salient aspects of a visual scene will attract more attention, and as a result, just as coherence theory suggests (Rensink, 2000b), changes to these aspects or parts will be more readily detected by observers. The choice of faces as an experimental stimulus (discussed in detail in the chapter that follows) is driven by this hypothesis, as it is reasoned that the face as a stimulus is particularly salient, and that the face's constituent parts, that is its features, can be considered sensibly as unique and distinct parts, with some features being more salient than others. Finally, the findings yielded by the present study could go some way towards challenging the comparison failure hypothesis posited by Mitroff, et al. (2004). It is hypothesised here that some facial features, when altered, will be more likely to produce higher levels of change detection than others. If change blindness results from a failure to compare pre- and post-change objects, despite these objects being retained in memory, then levels of change detection should be similar regardless of which particular object (or part of the stimulus image) undergoes the change.

3 THE FACE AND ITS USE AS AN EXPERIMENTAL STIMULUS

3.1 Nature of Stimuli in Past Change Blindness Experiments

Before addressing the questions surrounding the choice of stimulus materials in the present study, it is important to take a closer look at the type of stimuli used in change blindness experiments in the past, and delineate some significant distinctions between these. The type of experimental stimuli used also produce different 'types' of change detection, and distinguishing between these can have important implications for one's understanding of the underlying mechanisms and causes of change blindness.

Table 3-1 shows a summary of a number of change blindness experiments and the resulting change detection rates of each. While the literature often groups these experiments as a collective body of work which demonstrates change blindness, some important distinctions need to be made. The first of these is the difference between what will be referred to here as changes that occur within a 'narrative paradigm' and changes that occur within a 'flicker paradigm'. The narrative paradigm refers to changes occurring towards the end of a series of events that can be understood by the perceiver as a related, story-like narrative sequence. For example, giving directions to a person in the street, being interrupted by two people carrying a door who pass between you and your conversation partner, and then continuing the conversation once they have passed is a narrative sequence of events. The experiments by Levin and Simons (1997), Simons and Levin (1998),

Table 3-1: Summary of Selected Change Blindness Experiments and Detection Rates

YEAR	RESEARCHERS	STIMULUS MATERIAL	NATURE OF CHANGE	CHANGE DETECTION
1997	Rensink, O'Regan & Clark	Photographic images, alternating between original & changed variation	Changes to aspects of photographs thought to be of marginal interest	An average of 17.1 alternations were required before observers detected the change
1997	Rensink, O'Regan & Clark	Photographic images, alternating between original & changed variation	Changes to aspects of photographs thought to be of central interest	An average of 7.3 alternations were required before observers detected the change
1997	Levin & Simons	Motion pictures	Actor switch across a movie cut	Approximately 34% of viewers noticed a change had occurred
1998	Simons & Levin	Real world interaction	Person switch occurring when a door was carried between experimenter and participant	Approximately 47% of participants noticed a change had occurred
2002	Levin, Simons, Angelone & Chabris	Real world interaction	Person switch occurring when experimenter bent down behind a counter and another experimenter stood up in his place	Approximately 25% of participants noticed a change had occurred
2002	Levin, Simons, Angelone & Chabris	Real world interaction	Person switch occurring when a door was carried between experimenter and participant	Approximately 62% of participants noticed a change had occurred
2002	Levin, Simons, Angelone & Chabris	Real world interaction	Person switch occurring when a cardboard screen was carried between experimenter and participant while participant was attempting to photograph the experimenter	Approximately 47% of participants noticed a change had occurred
2004	Mitroff, Simons & Levin	A group of 4-8 line-drawn objects on a computer screen	One of the objects on the screen was replaced with a new object across a blank screen interruption	Observers reported approximately 66% of the changes
2005	Kadosh (Unpublished Honours Research Essay)	Motion Pictures	Person switch occurring across a movie cut	Approximately 36% of participants noticed a change had occurred

Levin et al. (2002) and Kadosh (2005) fit into the narrative paradigm. Alternatively, the flicker paradigm refers to changes to photographic or illustrated scenes that repeatedly alternate between an original image and a slightly changed variation, with gaps in between each alternation, used to mask the change transient. The alternating images do not form a story-like narrative sequence, but are simply the presentation of an image, then the image with a change, then the original image, and so on. Rensink et al.'s (1997) study fits into this paradigm.

Another important difference that results from the different stimulus paradigms is what exactly is meant by change detection. In change blindness experiments belonging to the narrative paradigm, change detection generally refers to the number of participants that reported detecting the change. For example, in the door study by Simons and Levin (1998), about 47 per cent of participants reported noticing the person switch. On the other hand, within the flicker paradigm, change detection refers to the number of alternations required for a participant to accurately detect the change (i.e. to report that the change was noticed and to point out what had changed). Thus, the main difference is that within the flicker paradigm, change is always detected, and the question being asked is how long it will take for the change to be detected, whereas within the narrative paradigm, the question being asked is whether or not the change will be detected at all.

A further, important distinction is that within the flicker paradigm, participants are informed that there will be a change, and actively search for this change. In the narrative paradigm, participants are naïve, that is, they do not know that a change is about to occur.

It is important to note that despite the differences between narrative and flicker paradigms discussed above, not all change blindness experiments fit neatly into one category or the other. In the experiments conducted by Mitroff et al. (2004), participants were neither involved in observing a narrative sequence of events, nor were they actively searching for change in flickering images. The experimental tasks involved viewing arrays of line-drawn objects, then a mask (a blank screen), then the same array with an object potentially changed. For every array there was a single pre- and post- change exposure, unlike the flicker tests where the images alternated until the change was detected. Thus, the nature of the stimulus was similar to those in the flicker paradigm (image – mask – altered image), but the operational definition of change was similar to those in the narrative paradigm (the question of whether change was detected was asked, as opposed to how long it took to detect change).

It is important to draw distinctions between the various types of stimulus used in change blindness experiments because it is possible that different types of experiments reveal different facets of the underlying causes and mechanisms involved in change detection. In the flicker paradigm, for example, it is sensible to assume that delays in change detection result from failures relating to perceptual mechanisms, since the task does not necessarily call for a semantic or meaningful interpretation of the visual scene, but rather involves an active visual search for change of any kind. On the other hand, experiments relating to the narrative paradigm may demonstrate that change blindness results from the subordination of perceptual mechanisms to the semantic content of visual scenes and sequences. Returning again to the famous 'door studies' (Simons and Levin, 1998, Levin et al., 2002), the participant may have detected that her conversation partner had changed at a perceptual level, but because she was not expecting such an unusual and far fetched event, and because experience in real world interaction had taught her that conversation partners usually remain stable for the duration of the conversation, she was unable to consciously access, and therefore remember and report, that the man standing before her at the end of the conversation was not the same person who had approached her to begin with.

The present study addresses the operational and theoretical complexities relating to the nature of stimuli in change blindness tasks by using the same set of images (photographic images of faces) in two separate experiments, one designed to fit into the narrative paradigm, and one designed to fit into the flicker paradigm. It is reasoned that if the same set of images is used in two operationally distinct tasks, the possible effect of the particular experimental design on change detection can be teased out more sensibly, thus contributing to a more thoughtful consideration of the possible mechanisms underlying the change blindness phenomenon. The choice of faces as an experimental stimulus in the present study may also contribute toward this end, as discussed in the following section.

3.2 Why Faces?

Change blindness experiments involving a person as the object of change (e.g. Levin & Simons, 1997, Simons & Levin, 1998, Levin, Simons, Angelone & Chabris, 2002) marked a move from the simplistic stimuli often used in laboratory experiments, to more complex and naturalistic stimuli, the likes of which are more likely to be encountered in real world situations. This is significant as it supports the view that change blindness is not an experimental artefact produced by artificial disruptions but rather a genuine and general failure to either retain and/or compare visual information from one moment to the next (Simons and Rensink, 2005).

In addition to this, the role of attention in change detection can be more thoroughly explored in experimental scenarios where the object of change will be more likely to receive preferential attention. It is perhaps not very surprising, at least intuitively, that unattended objects or events may go unnoticed. But when attending to visual scenes or images, it is more surprising that changes to the scene or image might not be noticed. When asking the question of what aspects of *attended* visual objects go unnoticed, one might uncover clues pertaining to facets of visual perception such as encoding and visual comparisons. Using people as the changing input stimulus can go some way towards guaranteeing that the stimulus will receive attention, since attention is normally drawn to semantically central items (Rensink, O'Regan, & Clark, 1997). More specifically, in images involving people, faces are found to be of particular interest, and draw a lot of attention. Studies involving the scanning of eye movements show that people tend to fixate on faces when viewing images containing people, objects and scenery (Kanwisher, 2001).

The emphasis on faces in this study is further driven by the observation that in a large number of change blindness experiments, particularly those involving the flicker paradigm (e.g. Rensink et al., 1997) the visual stimuli often contain changes to objects of marginal interest. That is, objects that when changed, have little effect on the semantic interpretation of the entire image, and changes to objects in very ‘busy’ images – in other words images that contain a large number of objects and details. Rensink et al. (1997) pointed out that when the change occurs in an object of central interest, the change is detected far more rapidly than for an object of marginal interest. For example, when an image of a man and woman dining was used – containing a change where the railing behind them was shifted further up to a new location – an average of 16.2 alternations was required to identify the change. However, when an image of a cock-pit windscreen overlooking a helicopter flying outside was used – containing a change where the location of the helicopter was altered – an average of only 4.0 alternations was required to identify the change. The centrality or marginality of objects in the photographs was determined by asking five naïve observers to provide a brief verbal description of the scene. Objects of central interest were defined as those parts of the image which were mentioned by three or more observers and objects of marginal interest were defined as those mentioned by none (Rensink et al., 1997). This demonstrates the importance of centrality as opposed to marginality of the change in change detection. However, it fails to address the issue of the semantic importance of the changing object in the meaningful interpretation of the image, and the relative ‘busyness’ or sparseness of details in the image and how these may affect change detection.

This study may go some way towards addressing questions regarding the effect of the types of stimuli used in change blindness experiments, since the stimulus materials used in this study can be contrasted with stimuli used in experiments such as the one discussed above (Rensink et al., 1997) – the images used in this study are sparse and uncluttered, with a single object in the foreground – while still being similar enough at the level of experimental methods to allow for meaningful comparisons.

3.2.1 Faces vs. Objects: Possible Implications of Different Processing Mechanisms

Given that the change blindness phenomenon has been demonstrated in both scenarios involving changing objects and changing people, and that this study involves faces as the changing stimulus, it is important to point out any possible differences in the processing mechanisms which lead to the recognition of objects and faces under normal circumstances.

How we perceive, recognize and remember faces has been the topic of much and varied research. Currently, holistic and configural processing of faces appears to be the dominant view (Bartlett, Searcy, & Abdi, 2003). This view holds that individual facial features, especially of upright faces, are not encoded independently of each other, but rather that the processing of faces involves the encoding of interfeatural relations and also interactive encoding across broad regions of a face (Bartlett et al., 2003). Even before one explores the experimental data which supports this view, it is intuitively sensible. After all, faces in nature exist in their entirety: one is rarely, if ever, faced with the task of recognizing a neighbour's eyes without having the neighbour's whole face present.

Tanaka and Farah (2003) argue that because all human faces share the same basic features which are arranged in the same basic configuration, recognition, identification and the ability to distinguish between particular faces must rely on a very detailed analysis of both the face's featural *and* configural information. They describe research where face components such as eyes, noses and mouths were better recognised as belonging to a particular face when they were presented within the context of the whole face than when they were presented in isolation. This was not the case, however, when the component parts used belonged to non-face objects, scrambled faces or inverted faces.

On the other hand, object recognition is generally thought to be achieved using part-based processing. Some examples of theories that posit that objects are recognised on the basis of the nature and location of their parts include Biederman's (1987, in Bruce and Humphreys, 1994) theory, that asserts that objects are perceived as being comprised of a limited set of volumetric shapes named 'geons'. Recognition of objects is based on the particular relationship between these geons. Biederman (1987, in Bruce and Humphreys, 1994) showed that when individuals were presented with line drawings of objects which had parts of their contours removed, recognition of the object was more accurate when the deleted contours maintained the underlying part structures than when they destroyed the part structures.

There also appears to be neuroanatomical evidence for the separate processing of objects and faces. In the case of neurological impairments, the dissociation of deficits in object processing, known as agnosia, and face processing, known as prosopagnosia, seems to suggest separate neurological processing for faces and objects (Bruce and Humphreys, 1994).

Although such evidence seems to suggest different mechanisms for the recognition of faces and objects, Tanaka and Farah, (2003) suggest that perhaps face recognition may not be that distinct from object recognition, and that the holistic processing most often attributed to faces may also be employed for other stimuli, particularly by 'experts' whose knowledge of particular objects or animals is quite extensive. Human beings, the argument goes, are experts at facial recognition, since faces are arguably the most important visual stimulus we encounter. If one possesses a certain degree of expertise in the recognition of non-face objects, Tanaka and Farah (2003) hypothesised, then the type of processing used in the recognition of such an object might be similar to the face recognition processes. They refer to a study in which participants were trained to become experts at the recognition of artificial objects termed 'Greebles', which, like faces, all had the same basic features arranged in a canonical configuration. The experts were found to be faster than non-experts at recognising Greeble parts when these were shown in an original configuration than when they were shown in a new configuration, indicating that objects other than faces make use of holistic recognition mechanisms.

With facial stimuli being the focus of this study, and given the discussion above which points to the possibility of different processing mechanisms for the recognition of objects and faces, it is important to note that the findings of this study may not be directly applicable to the role of objects within the change blindness phenomenon.

3.3 Should facial features be considered in isolation?

The discussion above referred to the generally accepted understanding that faces are recognised and processed holistically or configurally. Given this, can facial features be sensibly considered in isolation, as in this study? Evidence from experiments involving memory conjunction errors and inversion effects shows that featural (or piecemeal) information plays a unique role in facial recognition, prompting Bartlett et al. (2003) to suggest dual-route hypothesis of face recognition and perception.

Facial inversion studies have shown that inversion affects the processing of features less than the processing of configural information (for example spatial relations between features). A study by Freire, Lee and Symons (2000) showed that participants' accuracy levels for detecting spatial differences among faces were significantly lower when faces were presented inverted than when presenting upright faces, but when the differences were featural, the accuracy of detection was unaffected by inversion. A series of experiments by Searcy and Bartlett (1996) showed similar findings. Examining the perceived grotesqueness of faces that had been altered either spatially/configurally (for example moving features farther apart or closer together) or distorted at the feature level (for example blackening some teeth), they found that while the perceived grotesqueness of the former was significantly reduced by inversion, inversion did not reduce the perceived grotesqueness of the featurally altered faces. Searcy and Bartlett (1996) also showed that speed of detection in face-matching studies is slowed less by inversion of faces with featural differences than by inversion of faces with spatial differences.

Such evidence would seem to suggest a qualitatively different process for recognizing and perceiving facial features than for processing spatial-relational information. This is further supported by findings from studies demonstrating memory conjunction errors. These errors are a phenomenon that occurs when people combine parts of separately perceived experiences and 'remember', falsely, a conjunction that does not correspond to a single previously experienced event (Reinitz, Lammers, & Cochran, 1992). Reinitz, Lammers, and Cochran (1992) conducted an experiment where participants were shown a series of faces, first in a study condition, then in a test condition. In the test condition participants were required to respond "yes" or "no" depending on whether they thought they had seen the face before, during the study condition. The tested faces were either 'old' faces, that is faces shown during the study condition, 'conjunction' faces, made up of combining the hair and mouth of one studied face with the eyes and nose of another, or entirely new faces that had not been seen in the study condition. Conjunction faces produced many false alarms, with participants responding that they had seen the face before, leading Reintiz et al. (1992) to conclude that featural information is encoded and retrieved more readily than configural information. Other studies yielded similar findings (for example Searcy, Bartlett, & Memon, 1999). McKone and Peh (2006) summarize the logic of memory conjunction experiments as follows:

> If performance is based entirely on matching the whole studied item (e.g. because no part decomposition takes place in either perception or memory), the

> pattern old>conjunction=new will be obtained. Alternatively, if performance is based purely on part matches (e.g., because there is no binding of parts into wholes), the pattern will be old=conjunction>new.

McKone and Peh (2006) argue, however, that the high levels of false alarm rates for conjunction faces in studies such as the one described above may have been influenced by a number of other factors. The first is that these studies employed as their stimuli schematic line drawings that showed ‘unnatural’ representations of faces, possibly causing participants to pay more attention to parts. The second is that in studies which used real photographs, the photographs included hair, considered to be a particularly salient recognition cue (discussed further below). To overcome such effects, McKone and Peh (2006) conducted a memory conjunction experiment using natural-appearing composite photographs with no hair, with conjunction stimuli made entirely of internal facial features: combining the eyes and eyebrows region of one face with the nose and mouth region of another (these feature combinations were chosen based on results of free scanning tasks which show that subjects generally fixate more on the eye region than either the nose or mouth – such findings are briefly discussed below). Results showed that conjunction faces were judged as old some of the time, producing the pattern old>conjunction>new. McKone and Peh (2006) concluded that these results demonstrate that configural or holistic processing does take place for upright faces, which explains the higher levels of recognition responses to the ‘old’ condition, but also that the results indicate memory for local feature sets, shown by the higher levels of recognition responses to conjunctions than to new faces.

The above discussion shows that as Bartlett et al. (2003) suggested, both configural/holistic processing and piecemeal processing appear to play a role in facial recognition, that there is experimental evidence supporting the claim that facial features can be unique units of perception at the level of encoding and/or retrieval, and that featural information plays an important role in long term memory tasks.

3.4 Salience of facial features in recognition and memory tasks

Shepherd, Davies, and Ellis (1981) discuss various experiments which suggest that particular attributes or elements of the face as a stimulus seem to be more critical for identification purposes than others. They divide these experiments into different groups such as studies relying upon subjective reports or verbal descriptions, experiments investigating participants' ability to detect changes in facial cues and how such changes might affect recognition of faces, and experiments investigating participants' eye movements while examining faces.

Results across the various types of studies, according to Shepherd et al. (1981), were not always homogeneous or consistent. Studies which employed the method of asking participants to report which cues they felt were important in recognizing a face showed that eyes and facial structure were listed as most subjectively salient, while the hair and mouth and chin regions were considered less salient. Studies which used a method of asking participants to describe a face (either from memory or from a present stimulus) found that the features most frequently mentioned, in descending order, were hair, eyes, nose, face structure, eyebrows, chin, lips, and others. Studies into the eye movements of participants (while examining a face) generally showed that people spent more time fixating their gaze on the central features of the face, namely nose, eyes and mouth, and less time on peripheral features such as hair (for example Stacey, Walker & Underwood, 2005).

Perhaps most relevant to this study are those experiments discussed by Shepherd et al. (1981) which measured the effect of distorting or altering certain features of a face in face recognition tasks. These experiments varied in their methods, with some having participants make same/different judgements on simultaneously presented pairs of faces constructed from forensic Identikit tools and measuring reaction times involved, while others had participants make same/different judgements on successively presented pairs of faces constructed from Photofit composites (for example Davies, Ellis & Shepherd, 1977). The various experiments did not always produce consistent results: the simultaneous, Identikit studies showed changes to the hair, eyes, and chin region to be most salient, with nose and mouth changes showing slower reactions, whilst the successive Photofit studies showed that while hair changes were most rapidly detected, eyes, nose and mouth regions showed more rapid detection than changes to the chin.

Despite the seemingly inconsistent salience data derived from the various studies involving facial recognition tasks, there do appear to be some overarching trends, according to Shepherd et al. (1981). Firstly, hair emerged as the most important, and by inference salient, single feature in the majority of the studies of cue salience. Secondly, all studies that posed the questions, "What parts of the face are most important for allowing an individual to identify a face as one s/he has seen before?" or "What parts of the face are most important in judging two faces to be similar or different?" found that some features had a more significant role to play than others.

Taken together with the evidence marshalled above which demonstrates the important role of features in visual processing and memory tasks involving faces, such findings may contribute to investigations into the phenomenon of change blindness, specifically those change detection tasks involving people as the object of change. The above findings are used in the present study as a guide, in order to construct an experiment where changes are made to selected aspects of a face, and it is expected that some of these changes are more salient than others. The aim of the study is then to test whether the more salient changes produce higher levels of change detection. The findings yielded by the present study could contribute to the understanding of change blindness in a number of ways. First, they may begin to answer the question posed above about what draws attention to some scene elements and not others in a change detection task. Second, they may provide evidence against the hypothesis that change blindness results from a failure to make between-view comparisons. Third, if this study yields findings which demonstrate that feature salience plays a role in change detection/blindness; this may provide a motivation for the investigation of cue salience in change detection tasks involving stimuli other than faces.

In addition, deepening our understanding of the change blindness phenomenon could have a number of applications in 'real world' settings. For example, as discussed in the previous chapter, change blindness tends to occur when the motion signal which normally accompanies change is masked or disrupted. For a number of ‘real world’ activities that require sensitive visual alertness and vigilance, such as driving, attending to closed-circuit television security images, or radar observation, a more thorough understanding of the processes involved in change detection and its ‘pitfalls’ could influence the way in which such activities are carried out, as well as informing technological and ergonomic development in such fields.

More specifically, the focus on faces and people in the context of change blindness could also have important ‘real world’ relevance. In Kadosh’s (2005) research investigating the effect of race in a change blindness experiment involving a person as the object of change, it was found that individuals were less likely to detect a change involving an actor switch in a motion picture if the actor pair were of a different race than the observer, and the implications of this for cross-race eyewitness identification were pointed out. Similarly, the effects of isolated facial feature transformations and change blindness could also contribute to our understanding of eye-witness testimony and its associated applications, since isolated facial features are often used in the construction of suspect composites, and feature-based retrieval techniques are used in order to narrow down the number of offender photographs shown to potential witnesses during criminal investigations (Lee, Whalen, Sakalauskas, & Baigent, 2004).

While eye-witness identification tasks are often quite different from change detection tasks like the ones employed in change blindness experiments, the potential relationship between change detection and eyewitness identification was illustrated using a hypothetical scenario in Kadosh's (2005) discussion. Imagine an old man sitting on a park bench reading a newspaper, when a young man sitting on the adjacent bench moves closer and asks him for the time. The old man responds with the time and goes back to reading his newspaper, until he decides to go get a cup of tea from a nearby kiosk. Before he gets up he reaches into his jacket pocket where he knows he should find some money, but he finds none, and, knowing that the money was present in his pocket when he first sat on the bench, he turns to the man on the adjacent bench and accuses him of picking his pocket when he came closer to ask for the time. Of course, the man on the bench would be a little confused because he had never asked the old man for the time; he sat down on the bench moments earlier when the other young man who had been sitting on it got up and left. This is a very unlikely scenario but it does demonstrate the possibility of failure of change detection in a real-life situation with relevance to eyewitness identification.

4 THE PRESENT STUDY: MOTIVATION AND AIMS

4.1 Research Rationale

The change blindness phenomenon has been the topic of a great deal of research in recent years (Pashler, 1988; Simons, 1996; Rensink, O'Regan, & Clark, 1997; Levin & Simons, 1997; Simons & Levin, 1998; Levin, Momen, Drivdahl, & Simons, 2000; Levin, Simons, Angelone, & Chabris, 2002; Mitroff, Simons, & Levin, 2004). While results yielded by such investigations seem to have important implications in areas such as visual perception and attention, with the change blindness literature converging on a relatively homogeneous set of findings, some questions remain open to investigation.

For example, it has not been clearly established what draws attention to some scene elements in a change detection task and not others. As discussed in the previous chapters, change blindness experiments which utilise more naturalistic scenes, such as photographs, as their visual stimuli are often fairly 'busy' images containing a large number of visual items, or contain changes to semantically or physically marginal items. This study utilises visual stimuli where the changing object is central to the scene, both in terms of meaning, and physically, with few other visual distracters present. This could not only help to test the robustness of the change blindness phenomenon in general, but could go some way towards uncovering which elements of a visual scene are most salient and thus less likely, when changed, to result in the failure of an observer to detect the change. Faces are a logical choice as stimulus material as they have already been found to receive preferential visual attention, as discussed in the previous chapter.

Furthermore, while a number of experiments have been designed to try and determine the cognitive processes which may be responsible for change blindness, these have not always yielded consistent findings, and the issue remains open to the possibility of further investigation. The previous chapters included a discussion of the comparison-failure hypothesis, and it was noted that this study's findings, if consistent with its hypothesis that some facial features are more salient and therefore more likely to elicit change detection, could go some way towards challenging the comparison-failure hypothesis. This forms a secondary motivation for the present study.

In short, by incorporating findings from the study of the salience of facial features in facial recognition tasks into a change blindness experiment, this study hopes to further the understanding of what aspects of a visual scene are most important in the context of a change detection task, and what such findings may reveal about the cognitive mechanisms underlying such change detection (or lack of detection).

4.2 Research Aims

The study investigated the effect of facial feature salience on the change blindness phenomenon in a change detection task involving a person as the object of change. The study utilised two experimental paradigms. The first was the narrative paradigm, where participants looked at a short photo-story, comprised of a few still frames, depicting individuals engaging in an everyday activity. In the final frame, a selected facial feature of a character in the story was altered. Four different photo-stories were used, each utilising a different alteration. A short questionnaire designed to determine whether the change was detected was administered. The second was the flicker paradigm, where a flicker test was carried out using three sets of alternating photograph pairs. Participants watched the continuously alternating photographs on a screen and were timed to determine how long it took them to spot the difference between the photographs.

4.3 Research Questions

1. What levels of change detection will result from experiments 1 and 2, and how do these findings compare to similar studies utilising narrative and flicker paradigms in change detection tasks?

2. Will particular facial features, when altered, produce higher rates of change detection than others in change blindness experiments involving a person as the object of change?

3. What are the implications of the findings for questions 1 and 2 for an understanding of the change blindness phenomenon and its underlying mechanisms?

5 METHODS

5.1 Experiment 1: The Photo Story

5.1.1 Design

The design of this study was pre-experimental: there was manipulation of the independent variable, but no control groups or random assignment. The dependent measure was the number of people who failed to notice the change in the photo story.

It is important to note here, however, that while there was no control group per se – the nature of the experiment is such that there could be no sensible absence condition – a theoretical baseline condition was used. The different person (DP) condition was selected as a baseline since it was reasonable to expect that people would notice a change to the entire object. More specifically, it was reasoned that if people were to notice a change to part of an image, they would certainly notice a change to the entire image.

Given the aim of the study; to determine whether some facial features, when changed, produce higher rates of change detection, this design, while pre-experimental, could allow the researcher to touch on causality but relative to a baseline that is not an absence condition.

5.1.2 Sample

The sample consisted of Wits University students (N= 301) with ages ranging from 15 to 50 years (mean age = 20.14 years, standard deviation = 2.968). Table 5-1 shows a break-down of some demographic characteristics of the sample.

Ninety participants (n=90) were tested on the different face condition, seventy-one participants (n=71) were tested on the different hair condition, sixty-nine participants (n=69) were tested on the different eyes and eyebrows condition, and seventy-one (n=71) participants were tested on the different nose and mouth condition. Table 5-2 shows a break-down of experimental condition by gender.

Table 5-1: Demographic Characteristics of Sample: Age by Gender

AGE	FEMALE	MALE	TOTAL
15 – 19	76	75	151
20 – 24	87	52	139
25 – 29	1	6	7
30 +	2	-	2
Unknown	1	1	2
TOTAL	167	134	301

A non-probability sampling technique was employed to recruit participants. All 301 participants volunteered to take part in the study, and were selected on the basis of convenience. This is discussed in more detail in the procedure section below.

Table 5-2: Sample: Experimental Condition by Gender

EXPERIMENTAL CONDITION	FEMALE	MALE	TOTAL
Different Person	58	32	90
Different Hair	31	40	71
Different Eyes & Eyebrows	29	40	69
Different Nose & Mouth	48	23	71
TOTAL	166	135	301

5.1.3 Instruments

Four short photo-stories or photo-comics were created for the purposes of the study. Each depicted a person engaging in an everyday activity, and was comprised of 7 still frames, creating a narrative sequence. All four photo stories were identical in every respect except that in the last frame an alteration was made to the appearance of the main character in the photo-story. The alterations fell into four categories: (i) the actor in the photo-story was replaced by a different actor altogether, (ii) the shape and colour of the hair was changed, (iii) the shape and colour of the eyes and eyebrows was altered, and (iv) the nose and mouth were changed. Photo-manipulation software was employed to achieve the effects in conditions 2-4. Figure 5-1 shows the original image and the four alternative images used in the photo-stories.

The photo story was presented one frame at a time, in sequence, with the exposure for each frame lasting between 3-6 seconds. The images were accompanied by an audio narrative which, coupled with the images, ensured the creation of a narrative flow which made it clear that the person in the final frame, the altered frame, was expected to be the same character seen in previous frames in the story. A single voice was used to narrate all the parts in the photo story. For a copy of the entire photo story and accompanying narrative please see Appendix A.

Original Image

Different Person

Different Hair

Different Eyes

Different Nose and Mouth

Figure 5-1: Photo story images

As an experimental stimulus, this photo story compares to the type of stimulus materials used in change blindness experiments that fit into the narrative paradigm, as discussed in the literature review, and introduces change in a surprising and unexpected way. Just as one would not expect one's conversation partner to change into a different person mid-way through an interaction (Simons and Levin, 1998, Levin et al., 2002), one would not expect a photo story character's hair to change from red to black from one moment to the next. Unlike previous studies, however, which used either real world interaction or motion pictures where a person was the object of change, the instruments here are constructed using photographic images. This may reduce the realism of the instruments to some degree, but it is worth noting that photographic images represent a more naturalistic stimulus than line drawings or illustrations (for example as used in experiments by Mitroff et al., 2004).

Since in real world interactions the assumption is that a person would be looking directly at the person with whom one is interacting – which is imperative in the investigation of change blindness, as the concept of change blindness refers to looking at but not noticing change – it was important to ensure that participants were actively looking at the images in the photo story, and they were verbally instructed to do so before the experiment began. This helped ensure that if an observer failed to report the change, this was not due to not having watched the photo story.

A questionnaire was also used in the study. The questionnaire comprised of two parts, the first containing demographic questions to record the participants' age and gender, the second containing questions designed to determine whether the change in the photo story was detected. No identifying information was included in the demographic questionnaire, so as to maintain the anonymity of the participants.

The question of change detection was not addressed directly. In other words, participants were not asked “Did you notice that in the last frame the main character’s hair changed colour?” for example. Instead, questions were similar to questions used in previous change blindness studies (Levin et al., 2002) so as to allow for more honest responses. Participants were first asked to describe the events in the photo story as they remembered them, and were allowed space for an open response. The second question asked participants whether they noticed anything unusual in the photo story with closed-ended Yes/No selections followed by an open-ended, “If yes, please specify” section. The inclusion of the first question was based on the reasoning that if a given participant did notice the change, but perhaps did not consider this to be particularly unusual; s/he may have described the change in answering the first question, but not in the second one. These questions were the same as the ones used in Levin et al.’s (2002) experiments involving person changes during real world interactions. The questionnaire took between 3 and 7 minutes to complete, depending upon how much detail participants chose to include in their responses. A copy of the questionnaire can be found in Appendix B.

Given the relative simplicity of the concept of change blindness, and its proven robustness in previous studies on which the present study was operationally modelled, it is sensible to assume the reliability and validity of the instruments.

5.1.4 Procedure

Data collection took place over several days, and the experiment was conducted in one of two variations. In the first variation, large groups of students (for example an entire lecture audience/class) were addressed en masse at various lecture theatres on campus at the University of the Witwatersrand, with the experimenter asking for their participation in the study. The students were informed, both verbally and in writing, that participation was purely voluntary and that electing to participate or not would neither advantage nor disadvantage any person in any way (academically or otherwise). Students were also briefly informed of what participation would entail, in terms of the type of tasks involved and the time it would take to complete the tasks. Those electing to participate signed an informed consent form and remained at the venue to take part in the experiment which immediately followed.

Volunteers were issued a series of verbal instructions by the experimenter. They were asked to pay close attention to the images in the photo story, and not to discuss the photo story or their responses in the questionnaire until the experiment was completed and all questionnaires collected. The photo story was then projected onto a large screen at the front of the lecture hall, accompanied by audio narration. The questionnaires were then administered, volunteers completed the questionnaires, and these were collected by the experimenter. A debriefing session followed where the nature of the experiment was described to the participants.

The second variation of data collection was similar to the first in every respect except that potential participants were approached in much smaller groups (between 1 and 6 individuals) at various locations on campus, and watched the photo story on a smaller computer screen.

It is important to note that any one given participant only watched one variation/condition of the photo story.

5.1.5 Analysis

Descriptive statistics were used to answer question one. The levels of change detection were determined using frequencies, that is the number of people who reported noticing, or failed to notice, the change. This form of descriptive statistics is the same as those used in similar change blindness experiments (for example Simons and Levin, 1998), and so can be easily used as the basis for qualitative comparison to other studies.

A logistic regression analysis was employed in experiment one in order to test the relationship between the various experimental conditions, that is which photo story any given participant watched, and the likelihood of detecting the change. Logistic regression, which is a regression model, is appropriate for this analysis, firstly, as the research question being posed here is merely one of association, and secondly, as the dependent variable involved in the analysis is nominal. Statistical tests of the effects of other variables in the study, such as gender and age, on change detection were also included.

5.2 Experiment 2: The Flicker Test

5.2.1 Design

The design of this study was pre-experimental: there was manipulation of the independent variable, but no control groups or random assignment. The dependent measure was the number of alternations taken to detect the change.

As in experiment one, there was no sensible absence condition that could be used as a control, but the theoretical baseline (purely theoretical in this instance, while experiment one had a theoretical and operational baseline) is the reasoning that a change to the entire image would be noticed immediately. Thus, while the design is pre-experimental, the question of causality can be touched on relative to this theoretical baseline.

The aim of this experiment, as in experiment one, is to determine whether some facial features when changed produce higher rates of change detection. But in this case the logic of change detection is slightly different as this experiment fits into the flicker paradigm, so change detection refers to how long it takes an observer to notice the change, and not whether or not the change was noticed.

5.2.2 Sample

A sample consisting of 75 students from the University of the Witwatersrand was used in this experiment, 40 females and 35 males. Twenty-five participants (n=25) were tested in the different hair condition, twenty-five (n=25) in the different eyes and eyebrows condition, and twenty-five (n=25) in the different nose and mouth condition. Table 5-3 shows a break-down of the sample in terms of experimental condition and gender.

Table 5-3: Sample: Experimental Condition by Gender

EXPERIMENTAL CONDITION	FEMALE	MALE	TOTAL
Different Hair	13	12	25
Different Eyes & Eyebrows	12	13	25
Different Nose & Mouth	12	13	25
TOTAL	37	38	75

A non-probability sampling technique was employed to recruit participants. All 75 participants volunteered to take part in the study, and were selected on the basis of convenience.

5.2.3 Instruments

A simple computer programme was created for the purpose of this study. Photographic images on a computer screen, alternating between the original (A) and a changed variation (A') were presented at an exposure rate of 1000ms, with a blank grey screen being presented in between at an exposure rate of 100ms (this is represented graphically in Figure 5-2). The purpose of the blank grey screen is to mask the change transient, which, as noted in the literature review, is always masked or otherwise interrupted in change blindness experiments, as the motion signal associated with the change transient instantly draws attention to the location of change. This cycle continues until a button is pressed. The number of alternations is recorded by the programme, counting from the first appearance of the altered variation up to and including the photographic image which appears at the time the button is pressed. This programme was used with three different image conditions, pictured in Figure 5-3.

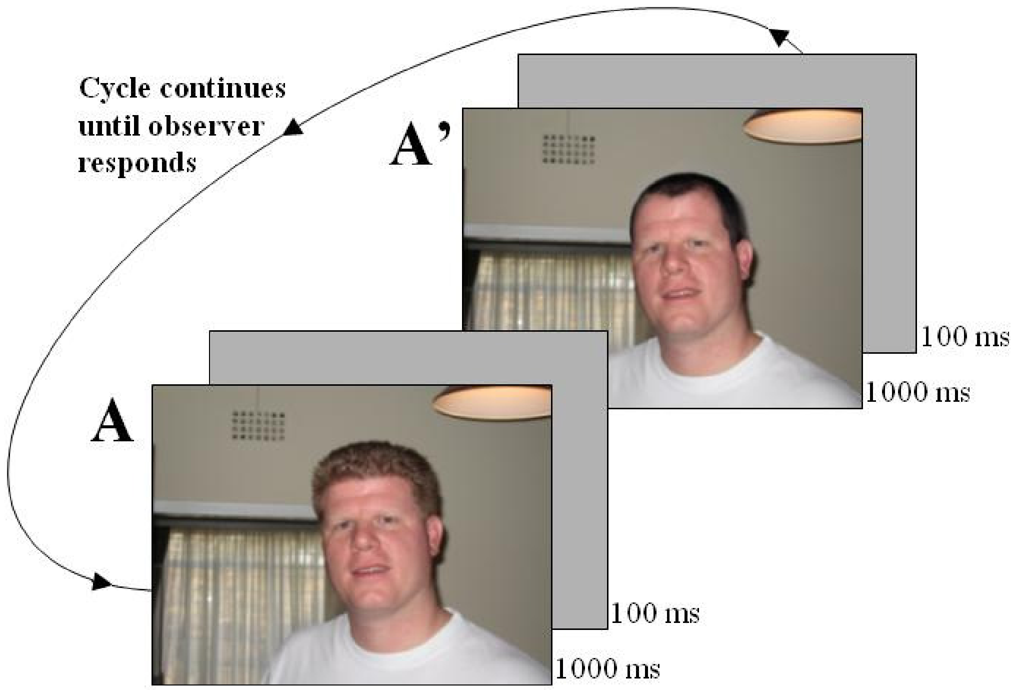

Figure 5-2: Flicker test cycle

The design of this experimental stimulus was modelled on the flicker tests conducted by Rensink et al. (1997). The verbal instructions for taking part in the flicker test were also like the ones used in the Rensink experiments, and are detailed in the following section. The similarity of the instruments to those used in previous studies goes some way towards ensuring their reliability and validity.

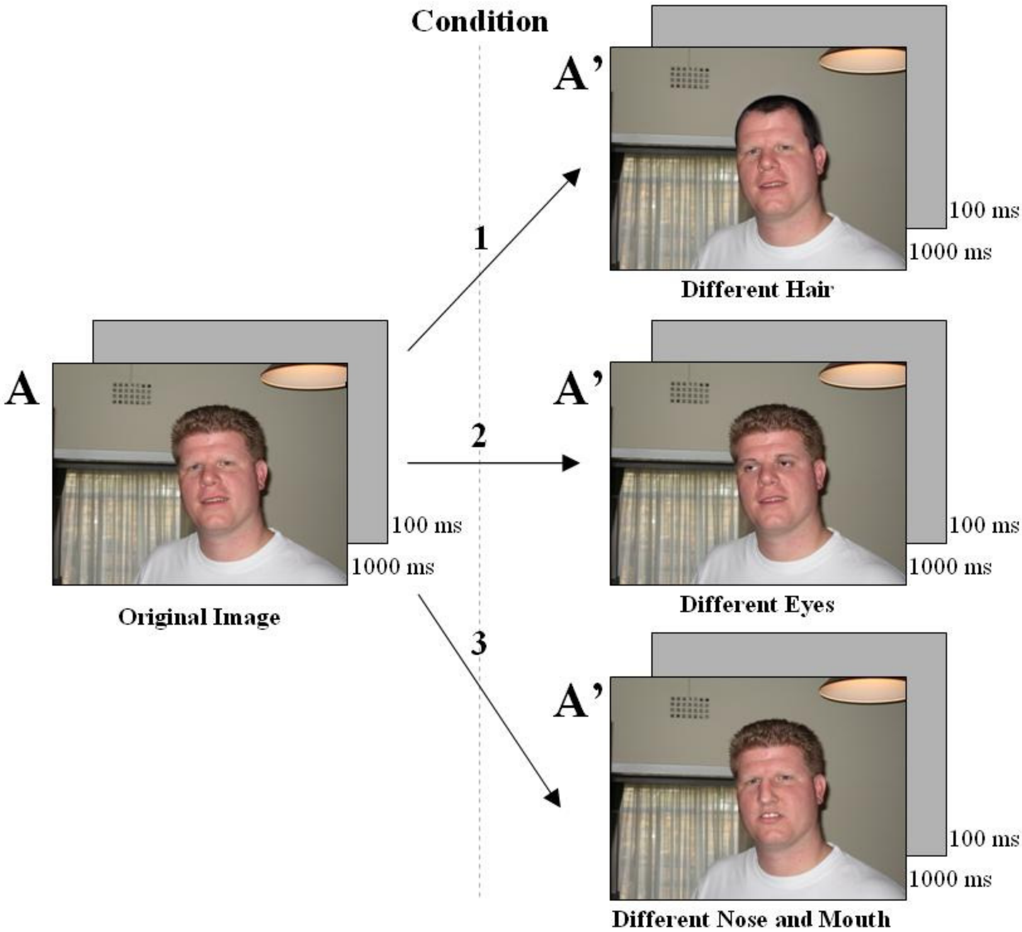

Figure 5-3: Flicker test conditions

5.2.4 Procedure

Potential participants were approached individually at various locations on the campus of the University of the Witwatersrand and asked to participate in the study. They were verbally informed of the nature of the study and what participation would entail. They were told that participation is voluntary, and that electing to take part or not would neither advantage nor disadvantage them in any way. Those electing to participate were tested on the spot.

A laptop computer was placed before them, and they were asked to attend to the alternating photographs on the screen, which contained a change, and to press a button the moment they noticed the change. The experimenter then recorded the number of alternations it took for the participant to notice the change, and asked the participant to verbally report what change had taken place, to ensure it was accurately detected.

This was carried out using three different experimental conditions, with each participant taking part in only one of these.

5.2.5 Analysis

A Kruskal-Wallis test of statistical significance was used in experiment two in order to test the relationship between rate of change detection and experimental condition. Analysis of variance was selected in this case as there were three experimental conditions and the dependent measure was not nominal (number of alternations needed to detect the change). A nonparametric analysis of variance was appropriate as the results were not normally distributed.

5.3 Ethical Considerations

The study relied on volunteers, and every potential participant was issued an information sheet which assured them that the choice to participate in the study or not was their own, and gave them some indication of what the study entails, as well as the details of the experimenter, should they wish to contact her with questions in the future. Every potential participant was also issued an informed consent form which they signed and returned to the experimenter upon electing to participate in the study. Anonymity of individuals' responses was assured, so as to make participants feel as comfortable as possible to answer the questions honestly and openly. The debriefing following the experiment ensured that participants had knowledge of the nature of the study they participated in, and any feelings of inadequacy at failing to spot the change were alleviated by the experimenter informing participants that this is a common phenomenon. The experimenter was available to answer specific questions and address specific concerns of participants after the experiment had taken place. Ethics clearance was obtained prior to the commencement of the study. For a copy of the ethics clearance certificate, see Appendix D. For a copy of the information sheet and informed consent form please see Appendix C.

6 RESULTS

6.1 Experiment 1: The Photo Story

6.1.1 Change Detection

Table 6-1 shows the levels of change detection for the different conditions in the photo story experiment. The different person condition showed the highest levels of change detection, with 70 per cent of participants noticing the change. Change detection dropped markedly in the remaining conditions, with only 24 per cent, 14 per cent and 13 per cent of participants noticing the change in the different hair, eyes and eyebrows and nose and mouth conditions respectively.

Table 6-1: Photo Story Experiment Results

EXPERIMENTAL CONDITION	NOTICED CHANGE	DID NOT NOTICE CHANGE	TOTAL
Different Person	63 (70%)	27 (30%)	90
Different Hair	17 (24%)	54 (76%)	71
Different Eyes & Eyebrows	10 (14%)	59 (86%)	69
Different Nose & Mouth	9 (13%)	62 (87%)	71
TOTAL	99	202	301

Change detection results also showed that a greater percentage of female participants than male participants detected the change across all the experimental conditions. This is represented in Table 6-2.

Table 6-2: Change Detection by Gender

GENDER	NOTICED CHANGE	DID NOT NOTICE CHANGE	TOTAL
Female	65 (39%)	101 (61%)	166
Male	34 (25%)	101 (75%)	135
TOTAL	99	202	301

6.1.2 Logistic Regression Analysis

In a global analysis of the effects of the different variables in the study on change detection, the effect of experimental condition, that is which photo story was viewed, was found to be statistically significant at the 5% level of significance [Wald chi-square = 58.934, p <.0001], as was the effect of gender [Wald chi-square = 4.7518, p = 0.0293]. Age, however, was not found to be statistically significant [Wald chi-square = 0.9746, p = 0.3235]. There was also no statistically significant interaction between gender and experimental condition [Wald chi-square = 4.2753, p = 0.2332].

According to Simonoff (1997), the logistic regression model is appropriate for cases where the dependent variable is nominal or dichotomous. In this case, the dependent variable is change detection, and it takes on two possible 'values': noticing, or not noticing the change. Other variables, including other categorical variables, can act as predictors for the dependent variable. In this instance, the various experimental conditions act as the predictor variables. This is appropriate as the question being asked is whether changes to particular facial features are more likely to facilitate change detection than others. The logistic regression model generates odds ratios, which are the ratio of odds of the predicted events under different conditions.

Odds ratios were generated using the logistic regression procedure for the different hair (H), different eyes and eyebrows (EE) and different nose and mouth (NM) conditions with the different person (DP) condition acting as a baseline for all of these. Although statistically the choice of baseline is arbitrary, the different face condition was the obvious choice in this case as it was reasoned that if a given observer is likely to detect a change to a single facial feature then it is expected that this observer would certainly detect a change to all facial features.

The odds ratios for the conditions were as follows: 0.159 for H (different hair), 0.076 for EE (different eyes and eyebrows), and 0.054 for NM (different nose and mouth). Simply put this means that the odds of noticing the change in condition H were approximately one sixth the odds of noticing the change in condition DP, the odds for EE were approximately one thirteenth the odds for condition DP, and the odds for NM were approximately one eighteenth the odds for condition DP. This produces a pattern of change detection likelihood where change detection is most likely to occur for observers viewing the different face condition, followed by the hair, eyes and eyebrows, and nose and mouth conditions, in that order.

The relationships between individual pairs of conditions were also tested. The results are presented in Table 6-3. The probability of noticing the change in condition DP was found to be significantly different from all three other conditions at the 5% level of significance. The difference between the H and NM conditions was also found to be statistically significant at the 5% level of significance. The difference between conditions EE and H, and conditions EE and NM, however, were not found to be statistically significant at the 5% level. All the statistical results generated for this experiment can be found in Appendix E.

Table 6-3: Contrast Test Results: Condition Comparisons

CONTRAST	DF	WALD CHI-SQUARE	Pr > CHI-SQUARE
DP vs. H	1	53.8074	<.0001
DP vs. E	1	56.7782	<.0001
DP vs. NM	1	55.7503	<.0001
H vs. E	1	2.4533	0.1173
H vs. NM	1	4.2966	0.0382
E vs. NM	1	0.3616	0.5476

Key: **DP =** Different Person, **H** = Hair, **E** = Eyes & Eyebrows, **NM** = Nose & Mouth

6.2 Experiment 2: The Flicker Test

The results of the flicker test experiment are summarised in Table 6-4. These reflect the number of alternations required on average before participants detected the change in the various conditions. The H condition produced the fastest rates of detection, with a mean of 1.84 (and a median of 2) alternations needed to detect the change. This was followed by the EE condition with a mean of 2.72 (and a median of 3) alternations and finally the NM condition producing the slowest rates of detection with a mean of 4 (and a median of 2) alternations. Results also showed that the range of scores (measured in number of alternations) varied in amplitude between the various conditions. The range of scores for the different hair, different eyes and eyebrows and different nose and mouth conditions were between 1 and 5 alternations, 1 and 12 alternations and 1 and 19 alternations respectively.

Table 6-4: Flicker Test Results

EXPERIMENTAL CONDITION	NUMBER OF PARTICIPANTS	MEAN (and MEDIAN) NUMBER OF ALTERNATIONS	STANDARD DEVIATION
Different Hair	25	1.84 (2)	0.99
Different Eyes & Eyebrows	25	2.72 (3)	2.28
Different Nose & Mouth	25	4 (2)	4.09

6.2.1 Analysis of Variance

A histogram generated for results of the flicker test showed that these were not normally distributed. Results also showed unequal variances – see the standard deviations in Table 6-4 – and so a nonparametric analysis of variance was used (For a printout of the histogram please see Appendix F).The Kruskal-Wallis one-way analysis of variance is a nonparametric test that is analogous to a standard one-way analysis of variance but is distribution free and therefore appropriate for cases where results are not normally distributed (Howell, 1999). The test is based on the null hypothesis that a given number of independent samples being compared are from the same population, and when all the observations are ranked, (lowest value assigned rank 1, second lowest rank 2 and so on) the expectation is that the mean of the sum of ranks for each group will be equal for all groups (Ferguson and Takane, 1989).

When this statistical procedure was applied to the results of the flicker test there was found to be a statistically significant difference in rates of detection between the various experimental conditions at the 5% level of significance [chi-square = 6.445, p = 0.0399]. The mean score of the ranks of each experimental group showed that changes in hair were detected faster than changes in eyes and eyebrows, which were in turn detected faster than changes in nose and mouth. This is represented in Table 6-5. All the statistical results generated for this experiment can be found in Appendix G.

Table 6-5: Summary of Flicker Test Results in the Kruskal-Wallis Analysis

CONDITION	N	SUM OF SCORES	EXPECTED UNDER Ho	STD DEV UNDER Ho	MEAN SCORE
H	25	745.0	950.0	85.913874	29.800
E	25	988.0	950.0	85.913874	39.520
NM	25	1117.0	950.0	85.913874	44.680

Key: **H** = Hair, **E** = Eyes & Eyebrows, **NM** = Nose & Mouth

Mann-Whitney U tests carried out on all the possible pairings of conditions showed the difference between the H and NM conditions to be significant, while the difference between the H and E and the E and NM were not found to be significant. This is similar to the results which emerged from the pair-wise comparisons in experiment one.

7 DISCUSSION AND CONCLUSION

7.1 Question 1: The Qualitative Description of Change Detection Results

The rates of detection yielded by experiment one showed that, in the DP condition specifically, 70% of participants noticed the change. As mentioned in the literature review, this experiment can be said to belong to the narrative paradigm, since the images in the stimulus materials formed a story-like narrative sequence, and change detection referred to the number of participants who noticed the change. Although previous change blindness experiments did not utilize photo story stimuli like the one used in this study, the stimuli here were designed to mimic the scenarios used in change blindness experiments such as the door study (Simons and Levin, 1998, Levin et al., 2002) and the motion picture studies (Levin and Simons, 1997). The DP condition specifically mimics the whole object changes used in the abovementioned studies. As such, the results of the present study contribute to the existing body of work that demonstrates the change blindness phenomenon.

Results from the present study showed higher levels of detection than similar studies in the narrative paradigm. For example, in the door study (Simons and Levin, 1998), only 47% of participants noticed the change.

The DP condition is being considered here in isolation for the purpose of comparison, since the existing body of work on change blindness includes several examples of whole object changes where people were the object of change. The comparatively high rates of change detection yielded by this study could be explained in a number of ways.

The first explanation may relate to sample sizes. In the experiments summarized in Table 3-1, of those experiments which were categorised as belonging to the narrative paradigm and containing a person switch, sample sizes for each experiment ranged from 15 participants (Simons and Levin, 1998) to 39 participants (Levin et al., 2002) at most. In the present study, 90 participants took part in the DP condition of experiment one. It is suggested here that there may be a possibility that if all the (narrative paradigm) experiments detailed in Table 3-1 had sufficiently large samples, the levels of change detection observed for all of these might be more similar.

The second explanation may pertain to the stimulus used in this study as compared with the stimuli used in previous studies. Like the present study, experiments by Levin and Simons (1997), Simons and Levin (1998), and Levin et al. (2002) all used a person as the changing stimulus. In this study, however, the two people whose photographs were used in the experimental stimuli were selected for their marked dissimilarity in appearance.

In Kadosh's (2005) experiment, videos containing an actor switch across camera cuts were screened to participants, and the levels of change detection were determined using a questionnaire similar to the one used in the present study. Two videos were created, one containing black actors and one containing white actors. While the relationship between the race of the observer and race of the actors was found to have an effect on levels of change detection, it was noted that a surprisingly large number of observers failed to notice the change in the video containing a black actor switch, regardless of the observer's race. During the debriefing session, when participants had the opportunity to look at the videos again, several people noted that the pair of black actors was far more similar in appearance than the white actor pair. This led Kadosh (2005) to note that the relative similarity or dissimilarity in appearance of the people in the 'switch' may have had an impact on the rates of detection.

It is thus possible that the high rates of change detection in the present study resulted from the marked dissimilarity in appearance of the people in the switch, and that the levels of similarity (or lack thereof) in appearance of those actors used in similar change blindness experiments likewise affected the levels of change detection. In order to conclude that this was the case, however, further research would be required.

Furthermore, it may be argued that many previous studies involving people as the object of change involved observers undertaking some activity, for example giving directions, while exposed to the pre- and post- change stimuli, while observers in the present study adopted a more passive role, which may account for the higher rates of change detection. Levin et al. (2002) addressed such concerns in their experiments, and results did show higher rates of change detection when participants adopted a more passive role during pre- and post-change observation of the stimuli. This would seem to strengthen the argument for the important role of attention in the change blindness phenomenon, since fewer demands on an observer's attention appear to be associated with higher levels of change detection.

Experiment two, described as belonging to the flicker paradigm, yielded relatively fast rates of change detection as compared to other flicker studies. The average number of alternations required ranged from 1.84 in the H condition to 4 in the NM condition. As discussed in the literature review, the number of alternations required, on average, by participants in the Rensink et al. (1997) experiments ranged between 7.3 alternations and over 17 alternations.

The faster rates of detection could best be explained by the choice of stimulus images, since this was the main difference between the Rensink et al. (1997) experiments and this one, which replicated the Rensink et al. model in every other respect. Rensink et al. (1997) did address the question of the effect of the relative 'centrality' or 'marginality' of the changing object within the whole image, and showed that changes to objects of central interest were detected much faster. As noted in the literature review, objects of central interest were defined as those objects mentioned by three or more (out of five) observers in brief verbal descriptions of the images. Although this exercise was not carried out in the present study, it is possible that the facial features which underwent changes in experiment two would have been mentioned more frequently by naïve describers. That is, it is possible that these features are more salient than the objects of central interest as defined by the Rensink et al. (1997) study. Thus, the fast rates of detection observed here may be due to the use of a face as the changing image, the fact that the face was in the foreground and took up a large amount of space in the image, and the fact that the rest of the image was sparse and uncluttered. As such it could be said that the nature and composition of the experimental stimulus has an impact on change detection, and would make an interesting topic for future research.

In summary, across both experiments and all conditions, some level of failure of change detection (in experiment one) or delays in detection (in experiment two) were observed. This demonstrates the robustness of the change blindness phenomenon and contributes to the growing body of experimental work on the topic.

7.2 Questions 2 and 3: The Effect of Facial Feature Transformations on Change Detection and Implications

The pattern which emerged from the results of both experiments suggests that hair changes are more likely to facilitate change detection than changes to eyes and eyebrows, which in turn are more likely to facilitate detection than changes to noses and mouths. In the results section it was noted that while in experiment one the differences between the H and NM conditions were significant, the differences between the H and the EE conditions and the EE and NM conditions were not found to be statistically significant. However, the logistic regression analysis showed, based on the resulting odds ratios, that change detection did follow the pattern described above. This pattern can be interpreted in the context of change blindness and its possible underlying causes by turning back to the discussion in the literature review.

It was noted that in studies of the recognition and memory for individual facial features some features emerged as more salient than others. Hair emerged as the most salient of these features, with eyes ranking somewhat lower in salience and noses and mouths less salient still (Shepherd et al., 1981). These findings guided the construction of the experimental stimuli in the present study. The feature which was expected to be most salient, namely hair, was associated with higher levels of change detection that those features expected to be somewhat less salient, namely eyes, noses and mouths. When applied to the change blindness phenomenon this could have a number of implications.

Firstly, returning to Rensink's coherence theory of focused attention, it was suggested that focused attention is required to see change across interruptions, as such focused attention would act as a 'hand' that grasps a number of proto-objects, affording them the spatial and temporal coherence required to maintain their identity even across a masked transient. What the findings of the present study may begin to tell us is what objects or combination of objects this metaphorical 'hand' is most likely to 'grasp'.

It appears that the more salient features of an image are more likely to receive focused attention, and therefore less likely to produce change blindness than other, less salient features. It may be argued that data on feature salience is not unambiguous, which may allow for alternative explanations of the results of the present research. For example, it could be said that the relatively high rates of detection in the H condition were not the result of hair being the most salient feature per se, but that the position of the hair on the border or frame of the face may facilitate detection due to changes in figure-ground outline. A response to this would be that the framing effect of hair is an aspect that contributes to its relative salience, and may form part of the salience explanation as opposed to an alternative explanation. The question of what it is that renders some features more salient than others is not in question here, but rather which features, when altered, will produce higher levels of change detection. While it was not feasible in the present study, future research may employ strategies such as eye-tracking techniques in order to more accurately assess what parts of the face receive the most attention in facial recognition tasks and integrate these with change blindness tasks.

While it was already noted in the discussion of the differences between face and object recognition that the findings of this study may not be directly applicable to stimuli other than faces, this nevertheless serves as a starting point, and investigation of the role of feature salience of other objects in change blindness experiments is recommended for possible future research.

Secondly, the literature review included a discussion of the possibility that change blindness results from a failure to compare pre- and post-change objects, despite these being adequately retained in memory. The findings of the present study showed that some facial features when changed produced higher incidences of change blindness (in experiment one) or slower rates of change detection (as in experiment two). This goes some way towards challenging the comparison failure hypothesis.

The reasoning for this is that if the pre- and post-change images are adequately retained by the observer, and it is only a failure to compare the two that produces change blindness, then it would be expected that the different experimental conditions, such as the variations of the photo story used, would produce similar levels of change detection. This was not the case. As such it is suggested that comparison failure is not a sufficient explanation of the occurrence of change blindness.

This reasoning, however, does not necessarily rule out potential rival explanations of change blindness, such as processing failures at the level of encoding or retaining visual representations – in other words – memory. Although a pattern emerged where changes to various facial features produced different levels of change detection, and this weakens the comparison failure argument, it does not necessarily rule out the possibility that the relative salience of a feature will affect one's likelihood to visually represent and *remember* that feature over time, as opposed to simply affecting the likelihood to attend to that feature in the first instance.

However, research shows that good visual representations can be retained and remain accessible even under conditions that lead to change blindness (Varakin & Levin, 2006). This would suggest that poor memory for visual information cannot be solely responsible for the occurrence of change blindness. This puzzling finding led the researchers to stress the importance of distinguishing between what we are cognitively capable of doing and what we generally tend to do. Indeed, studies investigating the intentional control of visual memory show that individuals are able to intentionally forget some items from an array being studied in order to improve recall of the remaining items (Yotsumoto & Sekuler, 2006). Yotsumoto and Sekuler (2006) noted that the visual system selectively filters incoming information, in order to ensure that it remains manageable, but that exclusion occurs after the excluded items receive considerable visual processing.

This discussion is not dissimilar to the ideas put forward by Noe (2002), Clark (2002) and Mack (2002), in terms of suggesting that we have perceptual access to the visual world around us and stressing the important role of intention. It does suggest, however, that this intentional directing of visual experience occurs at a deeper level of processing, namely, memory, as opposed to attention.

Taken together, these various discussions and research findings point to possible links between attention, visual representations and memory in the context of change blindness, and so would call for more integrated research in the future. The focus of this research and the interpretation of its findings are, however, pitched at the level of attention. The reason for this is, firstly, that the role of attention in the change blindness phenomenon has been the focus of more theoretical discussion and research than the role of memory, and this is presumably because processing at the level of attention would seem to logically precede processing at the level of memory. In other words, it would be more sensible to explore possible failures of change detection at this level before attending to the questions surrounding deeper levels of processing. Secondly, and more importantly – as Varakin and Levin (2006) noted in their investigation of the relationship between visual memory and change blindness – although memory for visual representations can be good, people may still generally fail to detect changes to objects. The focus here is therefore on what visual experience in response to changing visual stimuli is generally like, as opposed to what it is potentially capable of, under special circumstances. From an operational standpoint alone, it is more efficient to first address those questions relating to more overt behaviour that can be more easily tested and measured (for example what parts of images we attend to and how this impacts on change detection).

An interesting result that emerged in experiment one is the large difference in change detection between the DP condition (70%) and the three other conditions (24%, 14%, and 13%) containing isolated facial feature transformations. The differences in the latter conditions were accounted for here by asserting that cue salience plays a role in change detection, but in order to interpret the relatively high levels of change detection for the different person condition, we return to the discussion of the various theories of visual experience presented in the literature review.

The claims of Noe (2002), Clark (2002) and Mack (2002), were compared and brought together to reveal an understanding of the nature of conscious visual experience which suggests that we do have a sense of a complete and rich visual world that is accessible to conscious awareness, and this sense is not an illusion. We have perceptual access to the world via our knowledge of sensorimotor contingencies (Noe, 2002) that allow us to attend, at will, to various aspects of a scene. In addition to this sensorimotor knowledge, our visual experience is mediated by semantic and memory-based knowledge which both informs our vision-related actions and our interpretation of the gist or meaning of scenes, which in turn directs our visual attention (Clark, 2002, Mack, 2002, and Noe, 2002).

These assertions can be used to interpret the high levels of detection recorded for the DP condition in experiment one. In this condition, the final frame of the photo story contained a photo of a person (let us say person B) not previously seen in the story, standing in the place that the narrative and one's expectation would dictate that the original occupant of the house (let us say person A) should be standing. This would have strongly affected the meaning, or gist of the story, and so, according to Mack (2002), would have attracted attention, in the form of conscious visual awareness of the change, that is change detection. The three other experimental conditions still contained person A in the final frame, only with some or other feature altered. The fact that person A was still there means that the gist or meaning of the story was less strongly affected by the change, and therefore attracted less attention, which resulted in much lower levels of change detection. To take this argument a step further, results of experiment two showed that when the gist or semantic understanding of the relationship between the images was of less importance (recall that this experiment fits into the flicker paradigm, so participants were actively looking for change between two still images, with no story or narrative dictating an expected relationship between them) the changes were detected extremely quickly. This demonstrates that while the change blindness phenomenon appears to be related to a failure at the level of purely perceptual mechanisms – there were delays in detection in experiment two

– the semantic content of the images also has an important effect on one's likelihood to notice, or fail to notice, change.

On the other hand, the fact that thirty per cent of participants failed to notice the change in the different person condition can be accounted for by the abovementioned discussion of visual experience as well. Our perceptual access to detailed visual information means that we are neither required to nor inclined to attend to all aspects of a scene all at once. While under normal circumstances this affords us the ability to process visual information efficiently and glean the gist of visual scenes quite accurately, in some cases this can result in errors, such as missing changes in change blindness experiments. As discussed above, however, this understanding may not provide a complete picture of all the variables involved in the change blindness phenomenon, and as mentioned, memory may be an additional variable affecting the observed levels of detection failure. In order to tease out all the variables involved in the process of change detection, more integrated research would be called for in the future.

Another finding of interest which emerged, although it was not explicit in the aims of the study, is the relationship between change detection and gender. In experiment one, a chi-square analysis of the results showed that females were better at detecting changes than were men. It should be noted, however, that the experimental stimuli contained men's faces. This could potentially be an influencing factor, and more extensive investigation of the role of gender in the change blindness phenomenon would be required before conclusions could be drawn.

Returning to the comments made in the literature review about the possible practical application of the findings of the present study, the above assertion that the more salient facial features are more likely to receive focused attention and therefore less likely to be susceptible to change blindness could be used to inform decisions relating to eye-witness testimony. For example, it was noted that feature retrieval techniques are used when presenting 'mug shots' to eye-witnesses for identification (Lee et al., 2004). That is, based on an initial subjective description of the suspect by the witness, photographs of offenders are presented to the witness in order of decreasing resemblance to their description as opposed to presenting witnesses with all available photographs. Knowledge of which facial features are less likely to lead to instances of change blindness could be helpful in directing the focus of such processes, constructing identification composites and so on, particularly in instances where the scene of the crime involved larger numbers of people or commotion, potentially leading to instances of change blindness.

7.3 Conclusion

A great deal of interest has been shown in the change blindness phenomenon, and its role within cognitive psychology. The surprising nature and robustness of the phenomenon have spurred a large volume of research that not only demonstrates the phenomenon itself but attempts to uncover its underlying causes, and in so doing create a deeper understanding of how visual information is processed under normal circumstances. This research attempted to contribute to this body of work by beginning to address the question of what draws attention to some scene elements and not others in a change detection task.

The primary aim of the study was to explore whether some facial features, when altered, were more likely to result in change blindness than others, in a change blindness experiment involving a person as the object of change. Findings from studies of the salience of facial features were used as a general guide, in order to determine if feature salience could be an influencing factor in drawing attention and therefore facilitating change detection. It was found that those features considered to be more salient, such as hair, indeed produced higher levels of change detection than those features considered less salient, such as noses and mouths. It is therefore concluded that the more salient aspects of an image are most likely to draw preferential attention in a change detection task, specifically in tasks involving images of faces or people.

A secondary aim of the study was to take the findings derived from the primary investigation and use these inferentially towards deepening an understanding of the possible mechanisms underlying the change blindness phenomenon. The results obtained appear to challenge, to some extent, the hypothesis that change blindness results from a failure to make between-view comparisons. This would suggest that the role of attention is more central to the phenomenon.

Finally, the findings of this study, when taken together with the theoretical discussion of the nature of visual experience addressed in the literature review led to the following conclusion: in order to see change, we need attention, and in order to attract attention, visual aspects (or changes to these) have to be both *salient* and *meaningful*. As such, the visual 'blunders' demonstrated by change blindness do not necessarily represent a flaw in our visual processing abilities, but rather the efficiency with which we normally interpret our visual world.

7.4 Limitations and Recommendations

7.4.1 Limitations of Experimental Design

Due to time constraints and the difficulty of accessing large samples, there was no random assignment of volunteers into the different experimental conditions. Random assignment is recommended based on the reasoning that if a group of subjects is recruited for a study and then randomly assigned to the various experimental conditions, the likelihood that the subjects in one experimental group will be fairly similar to the subjects in another experimental group with regards to all possible variables other than the independent variable being tested, increases (Babbie & Mouton, 2001). The advantage of this is that it strengthens the argument that the observed effects on the dependent variable are related to the independent variable, and not other variables. However, given that the phenomenon has been repeatedly demonstrated and replicated and has been shown to be robust, it is fair to assume that the possible effects of extraneous variables would not be of great concern. Furthermore, random assignment was sacrificed in an effort to gain access to larger samples, and indeed, the samples used in the present study were much larger than samples used in similar change blindness experiments (for example Simons and Levin, 1998, Levin et al., 2002).

Statistical analysis for experiment one showed that there was no relationship between age and change detection, and no gender-condition interaction, which suggests that these variables did not necessarily affect the change detection pattern observed in the results. However, as noted in Chapter 5, some participants viewed the photo story in large groups, and others in significantly smaller groups, or individually. It is possible that viewing it in a group affected the results in that participants may have gauged others' reactions, causing them to notice the change when they otherwise would not have done. It might be inferred, however, from the generally low levels of change detection in most of the experimental conditions that group membership as opposed to individual viewing did not have a large impact on the results. In addition, it is worth noting that the experimenter was present during all the trials (group and individual), and did not note the tendency of participants to communicate or discuss the experiment while it was taking place.

7.4.2 Limitations of Instruments

In the photographic images used in experiment one, the original image of the occupant of the house and the image used in the DP condition also differed slightly in background (see Figure 5-1). This may have had an effect on rates of change detection that were not accounted for by the person switch, and as such are noted as a limitation. However, given that the focus of this research was more specifically on those conditions involving individual facial feature transformations, and the backgrounds were identical for all the images used in these, this limitation should not greatly affect the conclusions drawn from the results of this study. Furthermore, of those participants who did notice the change in this condition, none mentioned the changes in the background of the image in their questionnaire responses. All noted the changing person in the foreground.

Additionally, as noted in the literature review, the use of a face as an experimental stimulus could be a limiting factor to the extent that the findings of this study may not be directly applicable to change blindness experiments containing non-face objects as the changing stimulus. Furthermore, in experiment one, the still photo images used may have been a limiting factor in a number of ways. Firstly, it is possible that the changes to isolated facial features in the final frame were interpreted as the same face presented at a different angle, making it look somewhat different, but not creating the impression that the said features had been replaced, especially in the eyes and eyebrows and nose and mouth conditions. Experiment two overcomes this potential problem to some extent, as participants were actively looking for an expected change.

Secondly, although the photo story stimuli can be considered more naturalistic than some experimental stimuli used in change blindness experiments in the past, photographic images preclude information that could facilitate change detection, such as voice quality and physical gestures and mannerisms. Although audio narration was used in the photo story experiment in place of written captions (which may have been a visual distraction for observers), a single narrator spoke for all the characters in the story. This is unlike change blindness experiments involving a person switch in real world interactions, where the voice quality of the individuals may act as another identifying feature which could facilitate change detection. The recommendations for future research that emerge from the general discussion above and the abovementioned limitations are discussed below.

7.4.3 Recommendations for Future Research

Given the differences between the mechanisms for face and object recognition, the findings of this study may not be directly applicable to the occurrence of change blindness involving non-face scenes or objects, and thus it is recommended that the role of cue salience for objects other than faces be investigated in the context of the change blindness phenomenon. Such investigations could mimic the experiments in the present study while using changes to objects or animals in place of people. While 'features' of objects do not necessarily correspond directly to the way we understand features of faces, such investigations could elaborate on aspects such as shape, colour and orientation of objects or parts of objects and the relative salience of these in the context of the facilitation of change detection.

Secondly, the present study's findings in the flicker test experiment showed faster rates of change detection than the Rensink et al. (1997) findings, and this was attributed to the nature of the images used. This could form the basis of future research into the relationship between the nature of visual scenes and images and change blindness. Investigations could highlight aspects of visual images which could potentially facilitate or hinder change detection, such as the number of items in each image, the locations of the changing item within an image (for example background or foreground), and the relative size of the changing item, and so on.

Thirdly, the findings that gender had an effect on the likelihood of noticing change in this experiment is also interesting and warrants future investigation. Fourthly, the limitations in the design and instruments used in this study discussed above could be considered when pursuing research of a similar nature. Random assignment, where possible, would be recommended to limit the effects of extraneous variables on the measure of interest. For change blindness experiments involving people as the objects of change, while the use of video clips were not technologically feasible in the present study (since isolated facial features were altered), it would be recommended, where possible, to use experimental stimuli that allow the inclusion of identifying cues such as voice and gesture, which may be important in facilitating change detection in real world scenarios.

Most importantly, however, would be the recommendation to address the differences between experimental paradigms in change blindness studies. As noted in the literature review and the discussion above, demonstration of the change blindness phenomenon within different experimental paradigms, that is narrative and flicker paradigms, could have important implications for the underlying mechanisms involved in change blindness and visual perception in general. This distinction has, to date, not been made explicit in the literature, and is strongly recommended for future investigation.

Careful consideration at the operational level would be required in order to tease out the potential differences between the role of narrative or semantic continuity/consistency (which, as suggested here, may hinder change detection if the change fails to disrupt these) and the role of purely visual, perceptual mechanisms (which, when unrelated to semantic consistency, are argued above to increase the likelihood of noticing change). Recent investigations have begun to touch on the role of the semantic versus the visual nature of changing objects in the context of change blindness, such as Hollingworth and Henderson (2002), who ensured that the change detection in their experiments relied on visual information by utilising ‘token-level’ changes to objects, in which objects were replaced by objects from the same semantic category. While Hollingworth and Henderson (2002) restricted their focus to the flicker paradigm, it would be recommended that future investigations extend such considerations into the broader context of different experimental paradigms.

REFERENCES

Babbie, E., & Mouton, J. (2001). *The Practice of Social Research.* Cape Town: Oxford University Press.

Ballard, D. H. (2002). Our perception of the world has to be an illusion. *Journal of Consciousness Studies, 9,* 54-71.

Bartlett, J. C., Searcy, J. H., & Abdi, H. (2003). What are the routes to face recognition? In Peterson, M. A. & Rhodes, G. (Eds.), *Perception of Faces, Objects and Scenes: Analytic and Holistic Processes* (pp. 21-52) New York: Oxford University Press.

Bruce, V., & Humphreys, G. W. (1994). Recognizing objects and faces. *Visual Cognition, 1,* 141-180.

Clark, A. (2002). Is seeing all it seems? Action, reason and the grand illusion. *Journal of Consciousness Studies, 9,* 181-202.

Davies, G., Ellis, H., & Shepherd, J. (1977). Cue saliency in faces as assessed by the "Photofit" technique. *Perception, 6,* 263-269.

Dennett, D. (1991). *Consciousness Explained.* New York: Little Brown.

Ferguson, G. A., & Takane, Y. (1989). *Statistical Analysis in Psychology and Education.* New York: McGraw-Hill.

Freire, A., Lee, K., & Symons, L. A. (2000). The face-inversion effect as a deficit in the encoding of configural information: Direct evidence. *Perception, 29*, 159-170.

Hollingworth, A., & Henderson, J. M. (2002). Accurate visual memory for previously attended objects in natural scenes. *Journal of Experimental Psychology: Human Perception and Performance, 28,* 113-136.

Howell, D. C. (1999). *Fundamental Statistics for the Behavioral Sciences.* Pacific Grove: Duxbury Press.

Kadosh, H. (2005). Change blindness in the South African context: Replicating the phenomenon. Unpublished Honours research essay.

Kanwisher, N. (2001). Faces and places: of central (and peripheral) interest. *Nature Neuroscience, 4,* 455-456.

Lee, E., Whalen, T., Sakalauskas, J., & Baigent, G. (2004). Suspect identification by facial features. *Ergonomics, 47,* 719.

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

Levin, D. T., Momen, N., Drivdahl, S. B., & Simons, D. J. (2000). Change blindness blindness: The metacognitive error of overestimating change-detection ability. *Visual Cognition, 7,* 397-412.

Levin, D. T. (2002). Change blindness blindness: A visual metacognition. *Journal of Consciousness Studies, 9,* 111-130.

Levin, D. T., Simons, D. J., Angelone, B. L., & Chabris, C. F. (2002). Memory for centrally attended changing objects in an incidental real-world change detection paradigm. *British Journal of Psychology, 93,* 289-302.

Mack, A. (2002). Is the visual world a grand illusion? A response. *Journal of Consciousness Studies, 9,* 102-110.

McKone, E., & Peh, Y. X. (2006). Memory conjunction errors for realistic faces are consistent with configural processing. *Psychonomic Bulletin & Review, 13,* 106-111.

McMullen, P. A., Fisk, J. D., Phillips, S. J., & Maloney, W. J. (2000). Apperceptive agnosia and face recognition. *Neurocase, 6,* 403-414.

Milner, A. D., & Goodale, M. A. (1998). The visual brain in action. *Psyche, 4.*

Mitroff, S. R., Simons, D. J., & Levin, D. T. (2004). Nothing compares 2 views: Change blindness can occur despite preserved access to the changed information. *Perception and Psychophysics, 66*, 1268-1281.

Noe, A. (2002). Is the visual world a grand illusion? *Journal of Consciousness Studies, 9,* 1-12.

O'Regan, J. K., & Noe, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences, 24.*

Pashler, H. (1988). Familiarity and visual change detection. *Perception and Psychophysics, 44,* 369-378.

Reinitz, M. P., Lammers, W. J., & Cochran, B. P. (1992). Memory-conjunction errors:Miscombination of stored stimulus features can produce illusions of memory. *Memory & Cognition, 20,* 1-11.

Rensink, R. A. (2000a). When good observers go bad: Change blindness, inattentional blindness, and visual experience. *Psyche, 6.*

Rensink, R. A. (2000b). Visual search for change: A probe into the nature of attentional processing. *Visual Cognition, 7,* 345-376.

Rensink, R. A. (2001). Change blindness: Implications for the nature of visual attention. In Jenkin, M. & Harris, L. (Eds.), *Vision & Attention* (pp.169-185) New York: Springer.

Rensink, R. A. (2002). Change detection. *Annual Review of Psychology, 53,* 245-277.

Rensink, R. A., O'Regan, J. K., & Clark, J. J. (1997). To see or not to see: The need for attention to perceive change in scenes. *Psychological Science, 8,* 368-373.

Searcy, J. H., & Bartlett, J. C. (1996). Inversion and processing of component and spatial- relational information in faces. *Journal of Experimental Psychology: Human Perception and Performance, 22,* 904-915.

Searcy, J. H., Bartlett, J. C., & Memon, A. (1999). Age differences in accuracy and choosing in eyewitness identification and face recognition. *Memory & Cognition, 27,* 538-553.

Shepherd, J., Davies, G., & Ellis, H. (1981). Studies of cue saliency. In Davies, G., Ellis, H., & Shepherd, J. (Eds.), *Perceiving and Remembering Faces* (pp. 105-131) London: Academic Press.

Simonoff, J. S. (1997). An unusual episode. *Journal of Statistics Education, 5.*

Simons, D. J. (1996). In sight, out of mind: when object representations fail. *Psychological Science, 7,* 301-305.

Simons, D. J., & Levin, D. T (1998). Failure to detect changes to people in a real-world interaction. *Psychonomic Bulletin and Review, 5,* 644-649.

Simons, D. J., & Chabris, C. F. (1999). Gorillas in our midst: sustained inattentional blindness for dynamic events. *Perception, 28,* 1059-1074.

Simons, D. J., & Ambinder, M. S. (2005). Change blindness: Theory and consequences. *American Psychological Society, 14,* 44-48.

Simons, D. J., & Rensink, R. A. (2005). Change blindness: Part, present, and future.*Trends in Cognitive Sciences, 9,* 16-20.

Stacey, P. C., Walker, S., & Underwood, J. D. M. (2005). Face processing and familiarity: Evidence from eye-movement data. *British Journal of Psychology, 96,* 407-423.

Tanaka, J. W., & Farah, M. J. (2003). The holistic representation of faces. In Peterson, M. A., & Rhodes, G. (Eds.), *Perception of Faces, Objects, and Scenes: Analytic and Holistic Processes* (pp. 53-74) New York: Oxford University Press.

Varakin, D. A., & Levin, D. T (2006). Change blindness and visual memory: Visual representations get rich and act poor. *British Journal of Psychology, 97,* 51-77.

Yotsumoto, Y., & Sekuler, R. (2006). Out of mind, but not out of sight: Intentional control of visual memory. *Memory and Cognition, 34,* 776-786.

APPENDIX A: PHOTO STORY

The Delivery Man

Slide 1

Slide 2

Slide 3

Slide 4

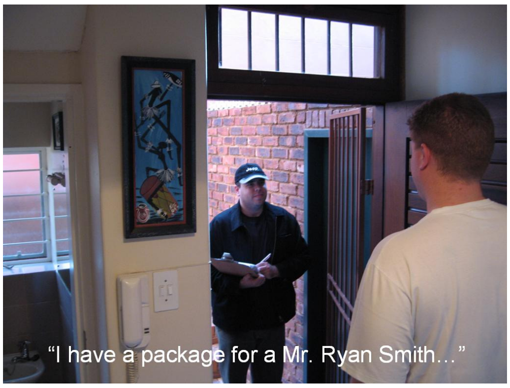

Slide 5

Slide 6

Slide 7

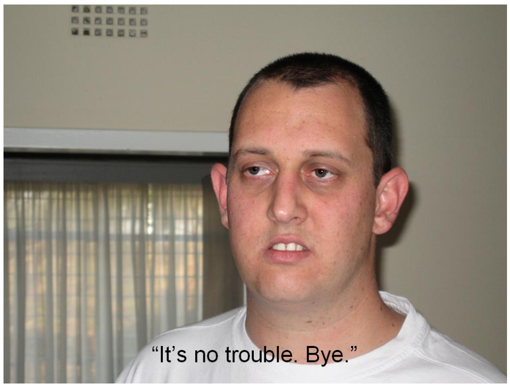

Slide 8 – Different Person

The End

Slide 9

Alternative Slide 8 – Different Hair

Alternative Slide 8 – Different Eyes

Alternative Slide 8 – Different Nose and Mouth

NB: Although subtitles appear on these slides, audio narration was used in the actual experiment.

APPENDIX B: QUESTIONNAIRE

Please complete the questionnaire below before turning the page

Age: ____________________

Gender: ☐ **F** ☐ **M**

Questionnaire Page 1

Please answer all the following questions as fully and honestly as possible

1. Briefly describe the events in the photo story as you remember them

____

Questionnaire Page 2

2. Did you notice anything unusual in the photo story?

☐ **Yes** ☐ **No**

If yes, please specify

____
____
____
____
____
____
____
____
____
____
____
____
____
____
____
____
____
____
____
____
____
____
____
____
____
____
____
____
____
____
____
____

Thank You

Questionnaire Page 3

APPENDIX C: INFORMED CONSENT FORM

Hadar Kadosh – Psychology Masters Research Project
Information Sheet

My name is Hadar Kadosh and I am a psychology Masters student at the University of the Witwatersrand. I would like you to take part in my research project.

Participation in this research will entail looking at a short photo story on a screen, and answering a brief questionnaire about the photo story. Both these tasks together should take no more than 15 minutes to complete. The photo story does not contain any disturbing images. Participation is voluntary, and no potential participant will be advantaged or disadvantaged in any way for choosing to participate or not to participate. Should you choose to participate, your anonymity will be guaranteed, as you will not be asked to provide any identifying information on the questionnaire, and there will be no way to link any questionnaire to a particular participant. The results of this study will be published in my research report. A short debriefing session will be carried out following the completion of the two abovementioned tasks, where I will discuss my research with you and answer any questions you may have.

If you choose to participate in this study **please complete the form that follows to indicate your consent**. Details of times and procedure will then be discussed with you. Should you want to receive feedback on the study once it is complete please send me an email with the subject line: "Study feedback". Feedback will then be sent to you. My email address is provided below.

Thank you,

Hadar Kadosh

Email: hadarkadosh@yahoo.com

Informed Consent Form – Page 1

Hadar Kadosh – Psychology Masters Research Project

I, ____, have read and understood the information sheet and am electing to participate in the study of my own free will.

Signature: ______________________________________

Date: ____________________

Informed Consent Form – Page 2

APPENDIX D: ETHICS CLEARANCE CERTIFICATE

UNIVERSITY OF THE WITWATERSRAND, JOHANNESBURG

Division of the Deputy Registrar (Research)

HUMAN RESEARCH ETHICS COMMITTEE (NON-MEDICAL)
R14/49 Kadosh

CLEARANCE CERTIFICATE **PROTOCOL NUMBER IH60908**

PROJECT Effects of Isolated Facial Feature Transformation in a Change Blindness

INVESTIGATORS Ms H Kadosh

DEPARTMENT Human & Comm Development/Psychology

DATE CONSIDERED

DECISION OF THE COMMITTEE* APPROVED UNCONDITIONALLY In-House Review

This ethical clearance is valid for 2 years and may be renewed upon application

DATE **CHAIRPERSON** ....
(Professor C Penn)

*Guidelines for written 'informed consent' attached where applicable

cc: Supervisor : M Pitman
Human & Comm Development

DECLARATION OF INVESTIGATOR(S)

To be completed in duplicate and **ONE COPY** returned to the Secretary at Room 10005, 10th Floor, Senate House, University.

I/We fully understand the conditions under which I am/we are authorized to carry out the abovementioned research and I/we guarantee to ensure compliance with these conditions. Should any departure to be contemplated from the research procedure as approved I/we undertake to resubmit the protocol to the Committee. **I agree to a completion of a yearly progress report.**

This ethical clearance will expire on 1 February 2008

Ethics Clearance Certificate

APPENDIX E: LOGISTIC REGRESSION

The following is a logistic regression analysis carried out on the results from experiment one, including pair-wise comparisons of all the conditions, and generated by SAS.

Key: **WF**: Different person condition, **H**: Different hair condition, **E**: Different eyes and eyebrows condition, **NM**: Different nose and mouth condition.

Logistic Regression Results

The LOGISTIC Procedure

Model Information		
Data Set	WORK.SORT3465	
Response Variable	NOTICED	NOTICED
Number of Response Levels	2	
Model	binary logit	
Optimization Technique	Fisher's scoring	

Number of Observations Read	299
Number of Observations Used	297

Response Profile		
Ordered Value	**NOTICED**	**Total Frequency**
1	Y	97
2	N	200

Probability modeled is NOTICED='Y'.

Note: 2 observations were deleted due to missing values for the response or explanatory variables.

Class Level Information				
Class	**Valu e**	**Design Variables**		
GENDER	**F**	1		
	M	-1		
CONDITION	**E**	1	0	0
	H	0	1	0
	NM	0	0	1
	WF	0	0	0

Model Convergence Status
Convergence criterion (GCONV=1E-8) satisfied.

Logistic Regression Results

The LOGISTIC Procedure

Model Fit Statistics		
Criterion	**Intercept Only**	**Intercept and Covariates**
AIC	377.256	302.706
SC	380.950	335.949
-2 Log L	375.256	284.706

Testing Global Null Hypothesis: BETA=0			
Test	**Chi-Square**	**DF**	**Pr > ChiSq**
Likelihood Ratio	90.5501	8	<.0001
Score	89.3411	8	<.0001
Wald	70.5433	8	<.0001

Type 3 Analysis of Effects			
Effect	**DF**	**Wald Chi-Square**	**Pr > ChiSq**
AGE	1	0.9746	0.3235
CONDITION	3	58.9340	<.0001
GENDER	1	4.7518	0.0293
GENDER*CONDITION	3	4.2753	0.2332

The LOGISTIC Procedure

Analysis of Maximum Likelihood Estimates									
Parameter			**DF**	**Estimate**	**Standard Error**	**Wald Chi-Square**	**Pr > ChiSq**	**Standardized Estimate**	**Exp(Est)**
Intercept			1	-0.8247	1.5947	0.2674	0.6050		0.438
AGE			1	0.0780	0.0790	0.9746	0.3235	0.0812	1.081
CONDITION	**E**		1	-2.5717	0.4400	34.1633	<.0001		0.076
CONDITION	**H**		1	-1.8404	0.3719	24.4891	<.0001		0.159
CONDITION	**NM**		1	-2.9151	0.4952	34.6480	<.0001		0.054
GENDER	**F**		1	0.5301	0.2432	4.7518	0.0293		1.699
GENDER*CONDITION	**F**	**E**	1	-0.8890	0.4449	3.9921	0.0457		0.411
GENDER*CONDITION	**F**	**H**	1	-0.0914	0.3722	0.0603	0.8060		0.913
GENDER*CONDITION	**F**	**NM**	1	-0.3109	0.4937	0.3965	0.5289		0.733

Odds Ratio Estimates			
Effect	**Point Estimate**	**95% Wald Confidence Limits**	
AGE	1.081	0.926	1.262

Association of Predicted Probabilities and Observed Responses			
Percent Concordant	80.1	**Somers' D**	0.622
Percent Discordant	17.8	**Gamma**	0.636
Percent Tied	2.1	**Tau-a**	0.275
Pairs	19400	**c**	0.811

Profile Likelihood Confidence Interval for Parameters					
Parameter			**Estimate**	**95% Confidence Limits**	
Intercept			-0.8247	-3.9430	2.3473
AGE			0.0780	-0.0787	0.2334
CONDITION	**E**		-2.5717	-3.5051	-1.7561

Profile Likelihood Confidence Interval for Parameters					
Parameter			**Estimate**	**95% Confidence Limits**	
CONDITION	**H**		-1.8404	-2.5922	-1.1288
CONDITION	**NM**		-2.9151	-4.0125	-2.0187
GENDER	**F**		0.5301	0.0561	1.0149
GENDER*CONDITION	**F**	**E**	-0.8890	-1.8184	-0.0428
GENDER*CONDITION	**F**	**H**	-0.0914	-0.8200	0.6470
GENDER*CONDITION	**F**	**NM**	-0.3109	-1.2335	0.7702

Wald Confidence Interval for Parameters					
Parameter			**Estimate**	**95% Confidence Limits**	
Intercept			-0.8247	-3.9502	2.3008
AGE			0.0780	-0.0769	0.2329
CONDITION	**E**		-2.5717	-3.4341	-1.7094
CONDITION	**H**		-1.8404	-2.5692	-1.1115
CONDITION	**NM**		-2.9151	-3.8858	-1.9445
GENDER	**F**		0.5301	0.0535	1.0066
GENDER*CONDITION	**F**	**E**	-0.8890	-1.7610	-0.0169
GENDER*CONDITION	**F**	**H**	-0.0914	-0.8210	0.6382
GENDER*CONDITION	**F**	**NM**	-0.3109	-1.2785	0.6568

Profile Likelihood Confidence Interval for Adjusted Odds Ratios				
Effect	**Unit**	**Estimate**	**95% Confidence Limits**	
AGE	1.0000	1.081	0.924	1.263

Wald Confidence Interval for Adjusted Odds Ratios				
Effect	**Unit**	**Estimate**	**95% Confidence Limits**	
AGE	1.0000	1.081	0.926	1.262

Contrast Test Results			
Contrast	**DF**	**Wald Chi-Square**	**Pr > ChiSq**
E vs H	1	2.4533	0.1173
E vs NM	1	0.3616	0.5476
H vs NM	1	4.2966	0.0382
E vs WF	1	56.7782	<.0001
H vs WF	1	53.8074	<.0001
NM vs WF	1	55.7503	<.0001

Contrast Rows Estimation and Testing Results									
Contrast	**Type**	**Row**	**Estimate**	**Standard Error**	**Alpha**	**Confidence Limits**		**Wald Chi-Square**	**Pr > ChiSq**
E vs H	PARM	1	-0.7314	0.4669	0.05	-1.6466	0.1838	2.4533	0.1173
E vs H	EXP	1	0.4812	0.2247	0.05	0.1927	1.2018	2.4533	0.1173
E vs NM	PARM	1	0.3434	0.5711	0.05	-0.7759	1.4627	0.3616	0.5476
E vs NM	EXP	1	1.4097	0.8051	0.05	0.4603	4.3176	0.3616	0.5476
H vs NM	PARM	1	1.0748	0.5185	0.05	0.0585	2.0910	4.2966	0.0382
H vs NM	EXP	1	2.9293	1.5189	0.05	1.0603	8.0932	4.2966	0.0382
E vs WF	PARM	1	-9.8989	1.3137	0.05	-12.4737	-7.3241	56.7782	<.0001
E vs WF	EXP	1	0.000050	0.000066	0.05	3.826E-6	0.000659	56.7782	<.0001
H vs WF	PARM	1	-9.1675	1.2498	0.05	-11.6171	-6.7180	53.8074	<.0001
H vs WF	EXP	1	0.000104	0.000130	0.05	9.011E-6	0.00121	53.8074	<.0001
NM vs WF	PARM	1	-10.2423	1.3717	0.05	-12.9309	-7.5537	55.7503	<.0001
NM vs WF	EXP	1	0.000036	0.000049	0.05	2.422E-6	0.000524	55.7503	<.0001

Classification Table									
	Correct		Incorrect		Percentages				
Prob Level	Event	Non-Event	Event	Non-Event	Correct	Sensi-tivity	Speci-ficity	False POS	False NEG
0.040	97	0	200	0	32.7	100.0	0.0	67.3	.
0.060	95	0	200	2	32.0	97.9	0.0	67.8	100.0
0.080	92	2	198	5	31.6	94.8	1.0	68.3	71.4
0.100	90	24	176	7	38.4	92.8	12.0	66.2	22.6
0.120	87	53	147	10	47.1	89.7	26.5	62.8	15.9
0.140	86	82	118	11	56.6	88.7	41.0	57.8	11.8
0.160	78	91	109	19	56.9	80.4	45.5	58.3	17.3
0.180	76	129	71	21	69.0	78.4	64.5	48.3	14.0
0.200	73	142	58	24	72.4	75.3	71.0	44.3	14.5
0.220	72	147	53	25	73.7	74.2	73.5	42.4	14.5
0.240	72	152	48	25	75.4	74.2	76.0	40.0	14.1
0.260	72	152	48	25	75.4	74.2	76.0	40.0	14.1
0.280	72	152	48	25	75.4	74.2	76.0	40.0	14.1
0.300	67	153	47	30	74.1	69.1	76.5	41.2	16.4
0.320	63	155	45	34	73.4	64.9	77.5	41.7	18.0
0.340	63	164	36	34	76.4	64.9	82.0	36.4	17.2
0.360	63	170	30	34	78.5	64.9	85.0	32.3	16.7
0.380	62	171	29	35	78.5	63.9	85.5	31.9	17.0
0.400	62	173	27	35	79.1	63.9	86.5	30.3	16.8
0.420	62	173	27	35	79.1	63.9	86.5	30.3	16.8
0.440	62	173	27	35	79.1	63.9	86.5	30.3	16.8
0.460	62	173	27	35	79.1	63.9	86.5	30.3	16.8
0.480	62	173	27	35	79.1	63.9	86.5	30.3	16.8
0.500	61	173	27	36	78.8	62.9	86.5	30.7	17.2
0.520	53	173	27	44	76.1	54.6	86.5	33.8	20.3
0.540	52	175	25	45	76.4	53.6	87.5	32.5	20.5
0.560	51	181	19	46	78.1	52.6	90.5	27.1	20.3

Classification Table

Prob Level	Correct		Incorrect		Percentages				
	Event	Non-Event	Event	Non-Event	Correct	Sensi-tivity	Speci-ficity	False POS	False NEG
0.580	48	184	16	49	78.1	49.5	92.0	25.0	21.0
0.600	47	184	16	50	77.8	48.5	92.0	25.4	21.4
0.620	46	185	15	51	77.8	47.4	92.5	24.6	21.6
0.640	45	185	15	52	77.4	46.4	92.5	25.0	21.9
0.660	45	186	14	52	77.8	46.4	93.0	23.7	21.8
0.680	44	186	14	53	77.4	45.4	93.0	24.1	22.2
0.700	44	186	14	53	77.4	45.4	93.0	24.1	22.2
0.720	44	186	14	53	77.4	45.4	93.0	24.1	22.2
0.740	43	187	13	54	77.4	44.3	93.5	23.2	22.4
0.760	34	188	12	63	74.7	35.1	94.0	26.1	25.1
0.780	6	191	9	91	66.3	6.2	95.5	60.0	32.3
0.800	2	196	4	95	66.7	2.1	98.0	66.7	32.6
0.820	1	198	2	96	67.0	1.0	99.0	66.7	32.7
0.840	0	200	0	97	67.3	0.0	100.0	.	32.7

APPENDIX F: HISTOGRAM FOR FLICKER TEST RESULTS

The following is a histogram generated using SAS showing the results obtained in experiment two.

The UNIVARIATE Procedure
Fitted Distribution for Alt

Moments			
N	75	**Sum Weights**	75
Mean	2.85333333	**Sum Observations**	214
Std Deviation	2.86966112	**Variance**	8.23495495
Skewness	3.38720734	**Kurtosis**	14.5403124
Uncorrected SS	1220	**Corrected SS**	609.386667
Coeff Variation	100.572236	**Std Error Mean**	0.33135992

Distribution analysis of: Alt

The UNIVARIATE Procedure
Fitted Distribution for Alt

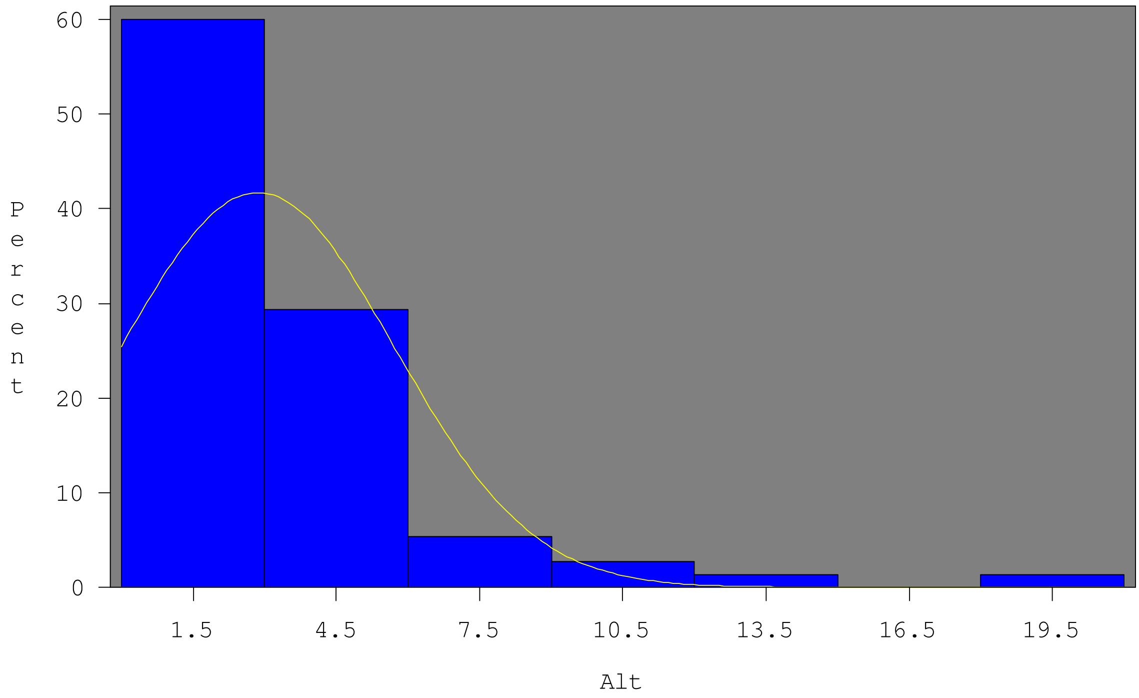

The UNIVARIATE Procedure
Fitted Distribution for Alt

Parameters for Normal Distribution		
Parameter	**Symbol**	**Estimate**
Mean	Mu	2.853333
Std Dev	Sigma	2.869661

Goodness-of-Fit Tests for Normal Distribution				
Test	**Statistic**		**p Value**	
Kolmogorov-Smirnov	**D**	0.30628584	**Pr > D**	<0.010
Cramer-von Mises	**W-Sq**	1.55324663	**Pr > W-Sq**	<0.005
Anderson-Darling	**A-Sq**	8.37767800	**Pr > A-Sq**	<0.005

Quantiles for Normal Distribution		
	Quantile	
Percent	**Observed**	**Estimated**
1.0	1.00000	-3.82250
5.0	1.00000	-1.86684
10.0	1.00000	-0.82429
25.0	1.00000	0.91778
50.0	2.00000	2.85333
75.0	3.00000	4.78889
90.0	6.00000	6.53095
95.0	9.00000	7.57351
99.0	19.00000	9.52916

APPENDIX G: KRUSKAL – WALLIS

The following is a Kruskal-Wallis analysis of the results obtained in experiment two, generated by SAS.

Key: **H**: Different hair condition, **E**: Different eyes and eyebrows condition, **NM**: Different nose and mouth condition.

Wilcoxon Scores (Rank Sums) for Variable Alt Classified by Variable Condition					
Condition	**N**	**Sum of Scores**	**Expected Under H0**	**Std Dev Under H0**	**Mean Score**
H	25	745.0	950.0	85.913874	29.800
E	25	988.0	950.0	85.913874	39.520
NM	25	1117.0	950.0	85.913874	44.680
Average scores were used for ties.					

Kruskal-Wallis Test	
Chi-Square	6.4450
DF	2
Pr > Chi-Square	0.0399

Kolmogorov-Smirnov Test for Variable Alt Classified by Variable Condition			
Condition	N	EDF at Maximum	Deviation from Mean at Maximum
H	25	0.800	1.00
E	25	0.480	-0.60
NM	25	0.520	-0.40
Total	75	0.600	
Maximum Deviation Occurred at Observation 31			
Value of Alt at Maximum = 2.0			

Kolmogorov-Smirnov Statistics (Asymptotic)			
KS	0.142361	KSa	1.232883

Cramer-von Mises Test for Variable Alt Classified by Variable Condition		
Condition	N	Summed Deviation from Mean
H	25	0.484741
E	25	0.112652
NM	25	0.307852

Cramer-von Mises Statistics (Asymptotic)			
CM	0.012070	CMa	0.905244