

Bangor University

DOCTOR OF PHILOSOPHY

Visual fame effects: the processing benefits of highly learnt images

Buttle, Heather

Award date:
2002

Awarding institution:
University of Wales, Bangor

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.


**University of Wales
Prifysgol Cymru**

**Visual Fame Effects: The Processing Benefits of Highly Learnt
Images.**

By

Heather Buttle

A thesis submitted to The School of Psychology, University of
Wales, Bangor, in partial fulfilment of the requirements for the
Degree of Doctor of Philosophy.



Summary

This thesis investigated the interaction of perception and memory when visually processing stimuli of varying familiarity. Specifically, it assessed whether advantages for processing famous (highly learned) versus non-famous (recently learnt) images were evident. A change-detection methodology was constructed, which required two briefly displayed images to be compared in memory. Each of the images contained two items, one of which changed into a different item. Experiments were conducted on a number of object classes: faces, landmarks, and consumer products, with the factor of main interest being whether the images contained famous or non-famous items. All categories of object benefited from the presentation of a famous item, while the exact pattern of effects differed depending on the object category. The divergence in the pattern of effects for certain object classes is explained by the degree of structural representation developed and maintained in conceptual short-term memory before transferral into traditional short-term memory (Potter, 1976, 1993, 1998). The advantage found for famous items, termed the *visual fame effect*, is explained in terms of efficient encoding mechanisms as described by robust representations (Tong & Nakayama, 1999) and the population-encoding hypothesis (Perret, Osram, & Ashbridge, 1998).

Keywords: fame, visual processing, attention, conceptual short-term memory, categorisation, face recognition, and object recognition.

Acknowledgements and Declaration

I wish to acknowledge and gratefully thank Prof. Jane Raymond for her rigorous supervision of my progress throughout the three-year process that led to the submission of this thesis. Without her ability to encourage my talents, while always adhering to a firm critique of my work, I would not have been as able or, as content in pursuing my research interests to this level of qualification.

I also owe a debt of gratitude to The School of Psychology at the University of Wales, Bangor. The department has supported me throughout the period of study leading to this submission. Furthermore, it was this department that enabled me to develop the skills required, through their undergraduate course, to undertake postgraduate study. The influence of the members of the department will never be forgotten.

I am most grateful to the ESRC (Economic and Social Research Council) for funding me through the PhD, by way of their Research Studentship.

This thesis is dedicated to my loving husband who I have intensely relied on for strength and energy from the start to the finish of this thesis. Without you by my side I would never have had the confidence or enthusiasm to pursue this research. Although this thesis may be complete, I will continue to rely on your understanding throughout the years to come, and I merely take this moment to thank you by dedicating the thesis to you.

Table of Contents

Declaration	ii
Summary	iii
Acknowledgements & Dedication	iv
Table of Contents	v
Chapter 1 - Introduction	1
Familiarity and Vision: The dynamic world around us	1
Face perception:	3
Cognitive Models	6
The Fusiform 'Face' Area	13
Configural Processing - the face inversion effect	16
Faces and the Right Hemisphere	20
Evidence of Familiarity Effects in Faces	24
Object recognition	30
Object Invariance	31
Categorisation and Expertise	36
Objects in Scenes	41
Visual Familiarity: what remains to be understood	49
Chapter 2 - General Methods	55

Chapter 3 - Fame Effects in Faces	59
Method	63
Experiment 1- upright	65
Experiment 2 - inverted	70
Experiment 3 - study and biography	74
General Discussion	79
Chapter 4 - Fame Effects in Landmarks	86
Experiment 4	92
Discussion	98
Chapter 5 - Fame Effects in Products	102
Method	104
Experiment 5 - Within-product type	107
Experiment 6 - Between-product type	113
General Discussion	117
Chapter 6 - Fame Measurements in Political Faces	124
Experiment 7	126
Discussion	131
Chapter 7 - General Discussion	134
Memory Differences in Visual Fame Effects	138
General Encoding Advantages	143
Market Research Applications	148
Summary	150
References	151
Appendix - Images	168

Chapter 1 - Introduction

Familiarity and Vision: The Dynamic World around Us.

The world is full of visually dynamic scenes and objects. Every time we move our heads we experience a slightly different view of the environment. When we turn corners or enter different rooms and buildings we are bombarded with a whole array of complex visual information. However, our experience of this world is one where we feel secure that the visual scene will be fully intact each and every time we open our eyes. In fact, the visual system and its associated brain regions are extremely successful in creating a visual experience that appears to be veridical. This is a great accomplishment if you consider that there are over twenty areas of the brain that extract different forms of visual information (e.g., shape, colour, depth) and then integrate them to create our actual visual experience.

Our visual system coherently and meaningfully interprets the visual input signals it receives into representations that become the perceived visual experience of the seeing organism. In order for this to occur, visual perception relies on encoding mechanisms that interact with an organism's previous knowledge and experience. Such interactions must be of a reciprocal nature. Information stored in memory should influence the visual processing of a stimulus. For example, when interacting with objects in our environment we require information from memory to guide attention to the salient attributes of an object (e.g., when talking to someone we look at their eyes and mouth more than their forehead). However, perceptual processing and encoding should also modulate stored memory representations in order to update and strengthen connections in memory networks. For example, if a friend radically alters their hairstyle, we need to be able to recognise that the friend is the same person before and after the change. Therefore, this new visual

information needs to be adequately incorporated into memory, in order to account for the person's change in appearance.

The examples cited above reflect the importance of visual learning and indicate that an individual's familiarity with coherent visual images (i.e., objects) influences both perceptual encoding and memory. Consequently it is anticipated that the greater the familiarity for any given object the more likely it is to benefit from encoding and memory advantages. Of central interest in this thesis is the issue of how encoding and perception are affected by the level of stimulus familiarity. Do highly learnt objects produce processing advantages when compared to recently learnt items? And how does the role of memory modulate any perceptual benefits of familiarity?

This thesis will examine perceptual advantages from viewing famous (highly learnt) over non-famous, but recently learnt objects. This investigation defines fame as describing any stimulus that is highly familiar and for which there is a well-established mental representation. Consider the act of picking up a favourite mug from a table that contains other mugs. A mug is an object that is familiar to us all, but the favourite mug is one that has a level of familiarity above that of most other mugs. We have developed a robust representation through repeated viewing and interaction that enables us to easily isolate it from the other competing options. It is important to emphasise that this description of fame pertains only to the individual's experience of an object, and not necessarily to a whole population's consensus as to what is or is not famous. For example, my husband's face is 'famous' to me, but is unlikely to be famous to you. Following on from this, the term *visual fame effect* is coined in this thesis in order to refer to any significant encoding advantages conferred on perception for highly learnt images.

The topic of visual familiarity is particularly interesting as it directly addresses the relationship between learning and vision. Here, behavioural techniques are used to investigate perceptual differences for distinct classes of famous and non-famous stimuli. Specifically, a change detection paradigm

was developed (i.e., participants had to detect an item that had changed in a visual display). A divergent pattern of fame effects will be demonstrated using faces, landmarks, and consumer products as different object classes. These findings are discussed in reference to previous theories of complex image perception that reflect a coherent explanation of encoding and memory processes.

The remainder of this chapter will discuss theories of face, object, and scene perception with the aim of establishing what is and is not already known about familiarity in visual stimuli.

Face Perception

Why discuss face perception before object recognition? Naturally, there is a tremendous amount of research and theory surrounding the issue of how we recognise objects. In fact, face perception could simply be considered a subsection of object recognition. However, there is little or no direct research for how visual object processing may be affected by high levels of familiarity. Certainly, there are no studies that I am aware of that compare 'famous' and 'non-famous' objects. Rather, the term 'fame' is restricted to the celebrity in human form, in respect to face, voice, and name recognition. Therefore, the literature involving objects and fame is sparse in comparison to the face perception literature.

The emphasis on face recognition arises from the fact that all humans are experts with this particular stimulus class. The prominent need for us to interact cohesively with people we encounter has obvious ramifications for how our visual system recognises faces. Faces need to be effectively processed in order for appropriate action to be instigated. Imagine (see Figure 1) you were required to oversee a small herd of sheep that included a checklist for each individual sheep. Without extensive experience with these types of

animal the task of identifying each sheep from only a brief introduction would be nearly impossible, with each animal appearing to look the same as the others. Overtime the task would become easier as isolated features could be learnt to aid recognition (e.g., 'Henrietta has a black patch at the bottom of her left hind leg'). Now imagine the problems we would encounter if each time we met a new group of people we couldn't tell them apart (it is one thing to fail to recall a recently introduced person's name, but to completely confuse them with other people would be highly embarrassing)¹.




Figure 1. Human faces are easier for us to recognise than other animals.

Thankfully this is not the situation most of us face. Rather, we have a visual system that is extremely proficient in recognising faces. This aptitude has led researchers to examine how the visual system recognises faces using perceptual tasks that involve famous and non-famous faces. In fact, Tong and Nakayama (1999) used faces in a visual search task that demonstrated that familiarity is a useful area to study. They developed a theory of *robust representations* through investigating the time-course of learning for a visual search task with both familiar/famous (i.e., the participant's own face) and unfamiliar/non-famous (i.e., stranger's faces). The term 'robust representation' is described as reflecting the endpoint of learning, which is

¹ Note that if this situation were reversed it would be equally true for sheep. Kendrick, Atkins, Hinton, Heavens, and Kerverne (1996) found that discrimination learning in sheep was better for familiar and unfamiliar facial stimuli of the same breed than of a different breed. This suggests that sheep face processing in this species is 'special' as has been claimed for human and non-human primates.

signified by asymptotic performance. Hence, with a highly learned face the robust representation that has developed will produce a flat performance function as it is already processed optimally, while a recently learned face will require learning, leading to a gradual improvement in performance with repeated exposure.

Using visual search tasks for self and stranger faces they discovered that the search slopes and intercepts were consistently faster for faces of the self rather than strangers' faces (both when self was target and when distractor). This occurred regardless of changes in view (e.g., profile or three-quarter angles) and number of stranger face presentations (the advantages persisted over 100s of presentations). Furthermore, while analysis of the trial positions produced a flat-function (no improvement with repeated exposure) for the face of the self, the analysis for the stranger faces revealed two components. First, there was a rapid reduction in response times to the unfamiliar face that occurred within a few presentations (e.g. 30). Second, the remaining trials produced a further reduction in reaction times, albeit much more gradual. Moreover, by the end of the experiment participants' reaction times still remained significantly faster for self over stranger faces.

From this behavioural data the authors suggest that robust representations may: (1) mediate rapid asymptotic visual processing (e.g., as demonstrated by the flat function for visual search with the face of the participant's self), (2) require extensive experience to develop (e.g., the stranger faces never produced as fast reaction times), (3) contain abstract or view invariant information (e.g., the advantage for familiar faces remained even in inverted conditions), (4) facilitate a variety of processes (e.g., can occur for tasks that do not explicitly require recognition of individual faces), and (5) demand less attentional resources (advantage occurred for images of the self as both target and distractor).

What remains unclear is whether these results would occur with other types of stimuli (e.g., the family pet versus other pets, the favourite mug

versus other mugs etc.), or whether this is an effect specific to face perception. This theory will be discussed in the ensuing chapters in order to provide support for the conclusions drawn from the experimental findings of this thesis. However, at this point in the thesis, it acts to show how familiarity can be studied, and demonstrates the advantages of exploiting the participants' expertise with the visual processing of faces.

Other theories have developed predictions of how familiar and unfamiliar faces are processed within models that seek to account for our abilities to recognise faces in general. These will be discussed in the following section.

Cognitive Models

One of the most influential approaches to face recognition is the functional model laid out by Bruce and Young (1986). In their model, recognition is considered to occur from any type of stored visual information extracted from faces that leads to an interaction with a number of functional components. The authors describe seven types of information (or codes) that can contribute to the face recognition process: pictorial, structural, visually derived semantic, identity-specific semantic, name, expression, and facial speech. The pictorial codes are merely the description of the picture that includes details such as lighting and grain. As far as the understanding of face recognition is concerned this aspect is perhaps the least interesting. Of far more relevance is the impact of the structural codes, where the essential details that differentiate one face from another are encoded. Accordingly these structural codes mediate everyday recognition of familiar faces and produce different codes to those of unfamiliar faces, as experience over time permits them to be elaborated and represented within recognition units. (Possibly such codes can account for the development of robust representations referred to earlier.) A familiar face is represented through an

inter-linked set of expression-independent structural codes, which contain information on head angles, global configurations (spatial relations of features), and distinctive features. For each known face there will be a face recognition unit that describes the critical aspects of the face, so that when a face is viewed the resemblance between the stored description and the input can be compared. When the correspondence is high there will be a strong signal from the unit to the cognitive system. Figure 2 depicts the model's layout, and as can be seen the visual information about a face enters the "structural encoding module". This process provides various visual descriptions at different levels of abstraction, where viewer-centred descriptions (visually-derived semantic) are analysed for expression, facial speech, sex, gender and race etc. This occurs for familiar and unfamiliar faces alike, as we are easily able to extract such salient information about faces that have just been encountered. Note that the Expression-independent descriptions, which access the face recognition units, are also derived from the structural encoding stage.

The face recognition units (FRUs) provide graded signals of resemblance that are then assessed by a decision process within the cognitive system. However, the basic level of activation can also be primed by recent use and also from contextual information that is fed back by the person identity nodes. The person identity nodes (PINs) allow recognition of the individual person, and can be accessed not just by a face, but also by information such as a name, a voice, clothing etc. When a PIN is accessed, the identity specific semantic codes become available, and it is only here that the name of a person can be obtained.

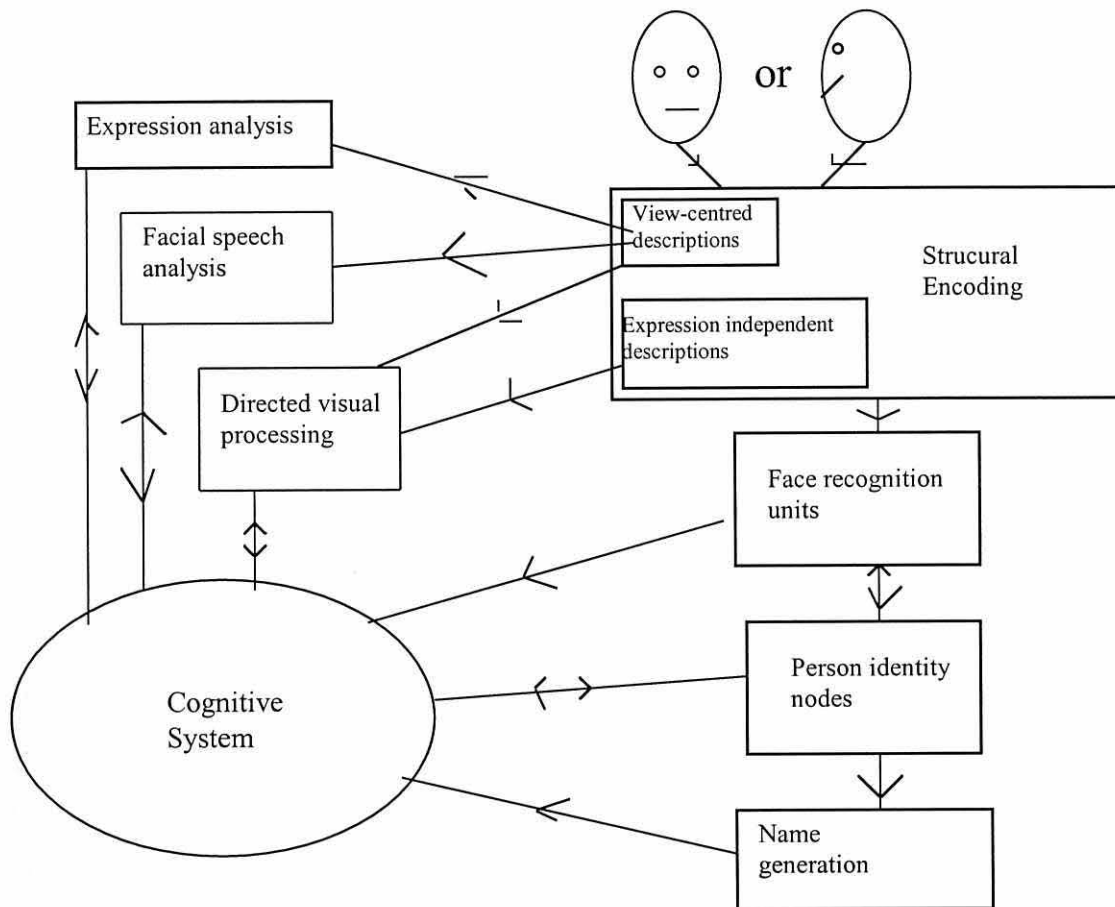


Figure 2. Bruce and Young's (1986) functional model of face recognition.

Intuitively it appears that naming should actually precede retrieval of semantic information. If we encounter a known acquaintance in the street, we are generally more anxious to recall the person's name than we are to remember his/her occupation. However, behavioural evidence argues otherwise. Johnston and Bruce (1990) required participant's to make match/mismatch judgements on pairs of faces from a set of eight faces. All of the faces were of celebrities; half were named John and half were named James. Within each of the name groups half were American and half were British, and within these subsets half were dead and half were alive. For example, of the four named John two were dead, one being American (John Wayne) and the other British (John Lennon). The match/mismatch decisions were based on whether the faces were the same name, same nationality, or

same life status (dead/alive). In all instances participants found the naming match decision a much harder task to make than the other semantic judgements. This suggests that there are separate semantic and naming codes, and that there is sequential processing from the former to the latter.

However, this issue is one of great controversy within the face recognition literature. While some studies support the distinction between the semantic information store and the name store others contradict this argument (e.g., Bredart, Valentine, Calder, & Gassi, 1995; Brennen, David, Fluchaire, & Pellat, 1996; 1999, Hanley, 1995; Hodges & Greene, 1998; Schweinberger, Burton, & Kelly, 2001). Evidence from face naming studies with Alzheimer dementia patients has produced some of the most vehement debates. Brennen's 1996 case study of such a patient revealed that an object/face's name could sometimes be recalled in the absence of appropriate semantic information ("naming without semantics"). This though contrasted with a report from Hodges and Greene (1998), who examined recognition, identification (the ability to produce accurate information) and naming abilities in 24 Alzheimer's patients (DAT type) along with 30 matched controls. Their data revealed that naming a famous face was possible only with semantic knowledge that could identify a person. In fact, 17% of the responses provided details about a person without being able to name them. While this particular debate over face naming abilities in Alzheimer patients' still continues, others have found evidence in normal participants that diverges from the Bruce and Young (1986) model. Schweinberger, Burton, and Kelly (2001) determined whether face names are accessed sequentially or in parallel to semantic information by using a speeded priming task with famous faces. The faces were primed by either partial semantic or partial name information. Significant priming occurred when a face naming task was preceded by a partial name prime, but not by a semantic prime. In converse there was significant priming for semantic tasks preceded by a semantic prime, but inconsistent results for name primes. The authors argue that this is

evidence for independence between access to personal semantics and names, and that it is possible that semantic and name information are processed in parallel as opposed to sequentially.

It is clear from Bruce and Young's model (1986) that regardless of whether we have previously encountered a face or not, there is full facial analysis at the stage of the view-centred descriptions. Clearly from our own experience we are almost instantly aware of the gender, race and age group associated with a new face. Furthermore, the visual processing that occurs at the point of the Face Recognition Units will depend on the familiarity of the individual face. If we have never encountered a particular face before, there is little possibility of recognition and therefore the pathway will not succeed in retrieving relevant information from the Person Identity Nodes.

This model of face recognition has been modified by Ellis and Young (1990) to incorporate another level that is relevant to the issue of face familiarity. They propose a two-route model that includes both a route similar to the Bruce and Young model (1986) that subserves the visual recognition of faces, and another separate route that subserves an affective component. Their model stems from research (Bauer, 1984, 1986; Tranel & Damasio, 1985) with prosopagnosic patients who, while failing to overtly recognise familiar faces, produce an autonomic response (i.e., SCR - skin conductance response) when the associated name is read aloud during viewing. The model assumes that in cases of this kind the primary visual route for face recognition is damaged, while an intact secondary, affective route produces the emotional response to the familiar faces.

Ellis and Young (1990) further support this account of visual processing in their reports of Capgras delusion (the belief that an impostor has replaced a close friend or relative). This is a case of having intact overt face recognition, but the affective response to the familiar face has been disrupted. Therefore, it appears that there is a double dissociation between prosopagnosia where face recognition is impaired, but the emotional associative response is intact,

and Capgras delusion where there is normal face recognition, but there is a failure to produce an autonomic response to familiar faces (Ellis, Young, Quayle, & de Pauw, 1997; Hirstein & Ramachandran, 1997).

Breen, Caine, and Colthart (2000) point out that these two-routes could be interpreted in ways not elaborated by Ellis and Young. One possibility could be a duplication of the FRU module, so that both the visual pathway and the affective pathway have an allotted FRU module. Alternatively, there could be just one pathway to the FRU that then branches into two pathways; one leads to the PINs and the other to the affective response. Breen et al. (2000) elucidates the ambiguity of this model in their account of face processing. Based on the original model of Bruce and Young (1986), the FRU module is seen to bifurcate into the PINs pathway and into the affective response pathway. This they believe is a more parsimonious account as duplication of the whole FRU would be far more demanding on resources. Of particular interest is the suggestion that the affective response will have increased activation for faces that have stronger emotional relationships with the viewer. This posits that the faces of close family members would produce efficacious affective responses when compared with less well-known faces (e.g., local taxi drivers). The modifications to the functional model of face processing offered by Breen et al. (2000) suggest that along with expected differences between known and unknown faces (at the level of the FRUs and PINs), there are likely to be processing differences for famous (highly learned/salient) faces when compared to faces that are recently learned or only occasionally encountered.

The issues surrounding these models are primarily concerned with the explicit act of face recognition and clearly aim to specify how retrieval of information relevant to viewed faces is achieved. However, a key concern of this thesis is how familiarity may influence visual processing efficiency regardless of whether a face is or is not recognised at the level of awareness.

For example, does viewing famous faces at a sub-threshold level confer processing benefits when compared to unknown faces?

Evidence from single-cell recordings in Macaque monkeys indicates that familiarity and experience do produce processing advantages in the form of supra-threshold responses (Oram & Perrett, 1992, Perrett, Oram, & Ashbridge, 1998, Perrett, Oram, Harries, Bevan, Hietanen, Benson, & Thomas, 1991). The neuronal input in response to a stimulus presentation should according to this view take longer to accumulate when the stimulus is unfamiliar or in an unfamiliar view than when it is familiar. Perrett, Oram and their colleagues used as stimuli faces and head angles to test this *population-encoding hypothesis*. Having selected 20 cells that are known to be selective for face views in temporal cortex, average responses were recorded to the presentation of head images. In order to compare the activity of the cells across the population, all firing rates were expressed as a percentage of the difference between the cell's maximum rate (set at 100%), and the cell's spontaneous activity (set at 0%). This normalisation allowed for each cell to make an equivalent contribution to the estimate of the population activity. Using this method the results showed that overall activity declined when the head was rotated away from the face view. Furthermore, the data for each view's presentation rose over time (cumulative response) and this rise was fastest for the face views. Hence, the rate of accumulation of response decreases in proportion to the angular rotation of the head from the face, or equally this can be expressed as a decreasing response for less familiar views. Similarly, it could be anticipated that familiar and unfamiliar exemplars of a given stimulus class (e.g., known and unknown faces) would produce similar patterns of accumulation.

These types of investigation with single cell recordings indicate that the level of familiarity with a face may impact the overall response time of face selective neurons. This, taken alongside Tong and Nakayama's theory of robust representations, strongly suggests that familiarity, and therefore

famous and non-famous visual stimuli, can lead to visual processing differences irrespective of any explicit demand for face recognition in the form of semantic information or naming. This is a critical point, as this thesis will aim to establish that visual fame effects (processing advantages from famous/ highly learned objects) are evident for faces, where no explicit recognition is required.

The Fusiform 'Face' Area

Single cell recording studies, coupled with cases of prosopagnosia (where brain lesions, typically located in the occipitotemporal region of the right hemisphere, produce selective deficits in recognising faces, while leaving recognition for other objects intact, for example see De Renzi, 1996), have been useful in that they suggest the existence of specific brain areas for the visual processing of faces. This has led researchers to investigate brain activation differences for faces and objects with brain imaging techniques, such as PET and fMRI (e.g., Haxby, Horwitz, Ungerleider, Maisog, Pietrini, & Grady, 1994; Puce, Allison, Gore, & McCarthy, 1995; Sergent, Ohta, & MacDonald, 1992). Such studies have revealed that certain brain regions are more active when viewing faces than other objects. In particular, an area of the fusiform gyrus has been associated with increased activation to faces. Kanwisher, McDermott, and Chun (1997) tested the selectivity of the fusiform gyrus by running multiple tests applied to the same cortical region within individual participants. Using a one-back task (i.e., judgements of consecutive repetitions of identical stimuli) with fMRI they searched for discrete regions of cortex (in particular regions of occipitotemporal cortex) that were specialised for face perception. Only one area, the fusiform gyrus, was consistently more active for face viewing. A number of further tests ruled out the possibility that the difference between the activation for faces and other objects was due to processes other than face selectivity. In one test,

participants viewed two-tone face images that were either intact or scrambled. Although mean luminance was constant between the two conditions, the intact faces produced greater fusiform gyrus activation. This ruled out the possibility that the low-level perceptual factor of luminance accounted for face and object differences in cortical activation. They also ruled out the possibility that the activation was due to viewing different exemplars of faces versus a variety of different objects, by testing the face exemplars alongside a set of house exemplars. Results again showed greater fusiform activation for faces over houses. Similarly, a test was made of whether the activation was due to faces being moving/living items as opposed to static objects, e.g., a chair. Even when body parts other than faces were tested, faces still produced more activation. Finally, a fourth alternative explanation was evaluated; this addressed the issue of sub-ordinate categorisation. If the increased activation of the fusiform gyrus was due to viewing objects that involved sub-ordinate categorisation, tasks involving other objects that required sub-ordinate categorisation would also lead to increased activation. This was tested by comparing activation from face stimuli with activation from hand stimuli (an equally difficult and attention demanding task), again activation was greater for face viewing. These imaging results provided the authors with evidence that the fusiform gyrus is selective for face stimuli, and has led to the cortical region being labelled the 'fusiform face area'.

However, other researchers have questioned whether the sub-ordinate categorisation account of increased fusiform activation has been reliably discounted (e.g., Gauthier, Skudlarski, Gore, & Anderson, 2000a; Gauthier, Tarr, Moylan, Anderson, Skudlarski, & Gore, 2000b). Gauthier et al (2000b) argue that although activation was stronger for face viewing than hand viewing, the hands and faces were not equivalent due to the level of expertise with the objects. If the hand stimuli were as expertly known they may have actually activated the fusiform-face area to a significant degree. The debate over whether the fusiform face area is truly face specific also extends to the

general issue of whether faces are a special case of object recognition. Although, this will be discussed in the section on object recognition later in the chapter, it is perhaps sufficient at this point to report here that this is an issue that continues to be debated extensively.

Another issue pertinent to the role of the fusiform gyrus in face processing is the relative activation for each area in the right and left fusiform gyri. A number of studies have revealed bilateral face-specific activity in the middle fusiform area (e.g., Halgren, Dale, Sereno, Tootell, Markinkovic, & Rosen, 1999; and Haxby, Horwitz, Ungerleider, Maisog, Pietrini, & Grady, 1994), while others report increased activation for faces in the fusiform gyrus of the right hemisphere versus the left hemisphere (Kanwisher, McDermott, & Chun, 1997; McCarthy, Puce, Belger, & Allison, 1999). A common view amongst these cognitive neuropsychologists is that the right hemisphere regions of the fusiform gyrus are specifically involved in face processing, while the fusiform regions of the left hemisphere are activated during general object recognition. Rossion, Dricot, Devolder, Bodart, Crommelinck, de Gelder, and Zoontjes (2000) tested the activation levels of the right and left fusiform gyrus on a face matching task that required either whole face judgements or part face judgements. The PET study revealed that the right fusiform gyrus (rFFA) was activated more by matching whole faces than by face parts. Furthermore, the left fusiform gyrus (lFFA) was activated more by the face parts matching task than by the whole face task. This was consistent with their prediction that as faces are processed faster in the right hemisphere (Hillger & Koenig, 1991; Leechey, Carey, Diamond, & Cahn, 1978; Levine, Banich, & Koch-Weser, 1988; Rhodes, 1993), a pattern consistent with faces being processed configurally (see following sub-section), the whole face task would lead to greater activation of the rFFA. Similarly, if the left hemisphere is advantaged in processing of features, the opposite pattern of activation would be expected for face parts.

Rossion et al's. (2000) study reveals differences in activation for right and left fusiform gyri depending on whether a face task requires whole versus feature matching. As stated, this links to research on the role of configural face processing and the respective properties of the two hemispheres. Both these factors need to be considered in some detail; hence the subsections that follow will elaborate these issues.

Configural processing – the face inversion effect.

The face inversion effect was originally investigated by Yin (1969). Recognition memory for faces and other objects that are normally presented in one orientation (e.g., houses) were compared. When the stimuli were presented and tested in their upright orientation, faces were better recognised than other items. However, when these stimuli were presented upside down, faces became the most difficult to recognise. Hence, Yin claimed that recognition memory for faces was disproportionately impaired by inversion. This effect has been successfully replicated in a variety of different conditions (e.g., Scapinello & Yarmey, 1970; Carey & Diamond, 1977; Young, Hellawell, & Hay, 1987).

The main theory for this face inversion effect is that face perception involves a greater reliance on configural (or holistic) information than on component featural information. When faces are inverted the familiar arrangement of facial features is disrupted, whereas the information extracted from other types of object tends to be based more on individual features. Furthermore, studies have shown that participants are more accurate at identifying parts of faces when they are presented within the whole object (face), than when they are presented in isolation (Davidoff and Donnelly, 1990; Tanaka and Farah, 1993). Interestingly, when the same paradigm was tested using scrambled faces, inverted faces, and houses, the previously

found superiority for identifying parts within the whole (as opposed to parts in isolation) was absent (Tanaka & Farah, 1993).




Figure 3. The Thatcher Illusion. The upright face on the right appears to be more grotesque than the inverted face on the left.

Similarly, the “Thatcher Illusion” demonstrates the effect of inversion on the configural processing of faces (e.g., Bartlett & Searcy, 1993; Thompson, 1980). If the spatial relations of a face are distorted by placing the mouth and eyes upside down within an upright face, the perception of that face is one of grotesque distortion (see Figure 3). However, if this face is inverted it is perceived as being almost normal. This unequal response to upright and inverted faces also occurs in other spatially distorted faces. However, when faces contain grotesque expressions (where feature information is odd, but the spatial relations are held constant) they appear to be distorted from the norm in both upright and inverted conditions. Again this suggests that inverting face stimuli disrupts configural information. Furthermore, Murray, Yong, and Rhodes (2000) suggest that there is a qualitative difference in upright and inverted face processing. They support this position through the

finding that unaltered and component-distorted faces are perceived to become more bizarre as they are oriented from 0 degrees to 180 degrees. However, Thatcher illusion faces showed a discontinuity in the function between 90 degrees and 120 degrees (i.e., the faces appear to be less bizarre).




Figure 4. The face on the left depicts the spatial relations (spatial configuration) in a face, while the right depicts the abstract holistic form of configuration.

While the evidence for configural processing in faces is reasonably established, the exact meaning behind this term is not so well determined. Leder and Bruce (2000) outline two possible interpretations for 'configural processing' in faces. One is that faces are processed holistically, in that faces are gestalts where by the features of a face are not explicitly represented (or at least to less of a degree). This view assumes that spatial relations between features will also not be represented (see Figure 4). An alternative interpretation termed 'relational' processing places emphasis on representations of spatial relationships between different local features (Diamond & Carey, 1986). Leder and Bruce (2000) have found experimental

support for the relational processing account. They tested the inversion effect on two types of facial stimuli; one consisted of a set of faces with the same local features, but different relational information, and the other set consisted of faces with the same spatial information, but different local details (the two sets did not differ in terms of holistic processing as the stimuli all shared the same arrangement of features). Using face stimuli created using the Mac-a-Mug program, the researchers found that it was only the first of these conditions that produced an inversion effect on a recognition task; the local differences failed to produce such an effect. This is entirely consistent with a relational processing account, as the faces that demonstrated an inversion effect only differed in terms of relational information.

It has also been proposed that configural disruption caused by inversion occurs at the encoding stage, as opposed to memory consolidation (Freire, Lee, & Symons, 2000). Comparisons between face encoding and memory (recognition) tasks indicate that once the face information has been encoded retention of both configural and featural information is equal. An encoding task requiring same/different discriminations to pairs of faces revealed that when differences were configural there was a decrease in accuracy from 81% for upright faces to 55% accuracy for inverted faces. However, on a memory task involving a match-to-sample task, where delays in stimulus presentation ranged from one to ten seconds there was no further decrease in accuracy for inverted and configurally changed trials (57% accuracy). This also held true for featurally changed trials, where neither the encoding or memory tasks led to significant decreases in accuracy for inverted stimuli. It appears that once configural or featural information is encoded, retention of these two types of information remains similar. Hence, Freire et al. (2000) suggest that perceptual processing of inverted faces creates an 'encoding bottleneck' that limits the input of configural information into memory.

Evidence from patients with neurological damage leading to visual agnosia for objects but normal face recognition for upright faces (Moscovitch,

M., & Moscovitch, D. A., 2000; Moscovitch, Wincour, & Behrmann, 1997), indicate that it is the integration of configural information from the internal details of a face (e.g., mouth, eyes, nose etc.) that leads to the inversion effect. CK is one such agnosic, who performs normally on upright recognition tasks of faces, but performs well below average with inverted face stimuli (the effect in CK is even more pronounced when only the internal facial features are viewed). Moscovitch and Moscovitch (2000) argue that the inversion of the faces makes configural information unavailable, making CK reliant on featural information. However, as with non-face objects CK is impaired in using featural information, and so performance falls markedly below controls on inverted face tasks, even though upright face tasks produce comparable performances.

The difference in processing of internal versus external facial details is particularly interesting as differences are also revealed for familiar and unfamiliar faces. A number of studies (Ellis, Shepherd, & Davies, 1979; Haig, 1986; Nachson, Moscovitch, & Umilita, 1995; Ross & Turkewitz, 1982; Young, Hay, McWeeeney, Flude, & Ellis, 1985) have observed that unfamiliar faces (recently introduced) are primarily recognised by external features (i.e., hair), while recognition of familiar faces relies on internal details, as much, or even more than external features. This implies that any potential difference in the processing of famous compared to non-famous faces may evolve from differences in configural processing. Famous faces would benefit the most from the global information contained in a face leading to a greater processing efficiency.

Faces and the Right Hemisphere

Differences in configural and featural perception of faces have been attributed to hemispheric specialisation for global and local processing. It has been commonly argued that the left hemisphere is specialised in analytic

processing (local/featural), while the right hemisphere is specialised in configural processing (Levy-Agresti & Sperry, 1968).

The research cited in the previous subsection (that indicates faces are processed configurally) fits this suggestion because, not only do face perception studies reveal biases toward the configural information contained in a face, there is also a right hemisphere (left visual field) advantage for face perception tasks (e.g., Anderson & Parkin, 1985; Rhodes & Wooding, 1989; Young, 1985). Anderson and Parkin (1985) successfully revealed a left visual field advantage when testing right and left visual field presentations of unfamiliar faces on a same/different face task. Young (1985), using a famous/non-famous judgement, also revealed a left visual field advantage for faces. Hence, the literature relating to visual field differences for face processing suggests a left visual field advantage (right hemisphere) that is consistent with both non-famous and famous faces being processed configurally.

Further support comes from the fact that the advantage shown by the right hemisphere (left visual field) for face recognition is actually eliminated (or greatly decreased) by inverting face stimuli (Hillger & Koenig, 1991; Leechey, Carey, Diamond, & Cahn, 1978). Hence, there seems to be a strong connection between the configural processing of faces and the right hemisphere. Furthermore, research relating to neurologically damaged patients also suggests that the right hemisphere preferentially processes faces. Individuals with right hemisphere damage tend to exhibit greater difficulty on face recognition tasks than those who sustain left hemisphere damage (Benton & Van Allen, 1968; De Renzi & Spinnler, 1966; Hecaen & Angelergues, 1962; Levy, Trevarthen, & Sperry, 1972; Milner, 1968; Newcombe, 1969; Warrington & James, 1967).

However, a series of studies by Sergent (1984, 1985) suggests that the two hemispheres may not strictly be involved in configural versus featural perceptual analysis. Rather, both may actually use the frequency of visual

information in different ways. When two faces were displayed one above the other and participants had to decide whether the faces were the same or different a right visual field (not left visual field) advantage was found (Sergent, 1984). In this case it was argued that participants performed the task by matching features of the face. This required local details, rather than global detail, and probably favoured the left hemisphere and the right visual field. Hence, it was possible to observe a right visual field advantage in some situations. Sergent (1985) explored a fuller account of the role of the hemispheres and spatial frequencies in face perception. She argued that the right hemisphere relied more on low-frequency information, while the left hemisphere relied on information from high spatial frequencies. Thus the right hemisphere processes configural information by using the low-spatial frequencies, while the left hemisphere processes fine detail, using the high-spatial frequencies. This allows encoding of the high resolution areas that tend to provide clear visual information for features.

The distinction between the high and low spatial frequency information and the role of the right and left hemispheres has led to the 'Double Filtering by Frequency' (DFF) theory (Ivry & Robertson, 1998). This account also states that the hemispheres are not specialised to process parts and wholes as such. Rather, the computations involving spatial frequency information by each hemisphere are structured differently and affect the part/whole analysis differentially. The theory postulates that an initial filtering stage selects task relevant frequency information from the visual image. This is then followed by a second filtering stage where the selected information is subject to asymmetric filtering by the two hemispheres. The left hemisphere will filter the relative high-spatial frequencies and the right hemisphere will filter the relative low-spatial frequencies.

Using this account Robertson and Ivry (1998) argue that the typical left visual field advantage for face perception can be explained in terms of their DFF theory. At the initial filtering stage the attentional process is proficient in

selecting configural information contained within the low-spatial frequencies. As the right hemisphere (left visual field) is biased toward these relatively low-spatial frequencies, a left visual field advantage for processing faces is witnessed. Furthermore, this account also makes provision for instances where face perception tasks reveal right visual field advantages. DFF predicts that if identification of a face requires attention to be allocated to distinctive features (as opposed to spatial relations between features or the overall Gestalt of a face) a high-spatial frequency bias should reveal a right visual field bias. In fact, Schyns and Oliva (1999) examined the spatial frequency biases involved in a variety of face judgment tasks. They used hybrid faces where two faces are superimposed on one another, with one containing visual information at low-spatial frequencies (LSF) and one containing information at high-spatial frequencies (HSF). After a participant made an initial facial judgement, the hybrid would be displayed to test how the first task biased subsequent perception. For instance, when required to decide whether an initial face was expressive or non-expressive, the hybrid task that followed revealed a bias toward the face displayed at HSF. In contrast, when participants had to make judgements as to what type of expression the first face exhibited the task transfer revealed a bias to the LSF face. Gender decisions led to an equal selection of the high/low spatial frequency face. Furthermore, the researchers were able to demonstrate that the initial learning of faces produced a bias toward LSF information, and that judgements about known/unknown faces also led to a LSF bias.

In respect to the role of familiarity in face processing the literature predicts generally that a left visual field advantage should occur and, this should be evident for both famous and non-famous faces alike. However, it may be anticipated that if faces are encoded configurally (using low spatial frequency information) and they are identified on this basis, any advantage for famous faces over non-famous faces may be greater in the left visual field. The comparison between famous and non-famous faces in the right visual

field may not lead to any discernable difference as the featural or high-spatial frequency information has not been optimised in the same way as the configural information for known faces. Any investigation of face processing involving multiple face displays and levels of familiarity should address this issue, if effects of famous versus non-famous faces are to be understood.

Evidence of familiarity effects in faces.

There have been a number of reports of double dissociations in prosopagnosia between familiar and unfamiliar face processing (e.g., Benton, & Van Allen, 1972; McNeil, & Warrington, 1991; Malone, Morris, Kay, & Levin, 1982; Warrington, & James, 1967), familiar and expression processing (Bowers, 1985), and unfamiliar and expression processing (e.g., Kurucz, 1979). These dissociations fit in well with the Bruce and Young model (1986) discussed earlier, as according to the model there are different functional pathways for unfamiliar, familiar, and expression processing of faces.

However, there are reasons to be sceptical of such reports. A major problem is that these dissociations are demonstrated across studies not within studies. Hence, one paper may report a patient who can recognise facial expressions but cannot identify familiar faces, and in another paper it will be reported that someone can recall familiar faces, but cannot recognise facial expressions. Researchers use different methods and tests to assess patients, making such cross-patient comparisons difficult. To help resolve this problem Young, Newcombe, de Han, Small, and Hay (1993) investigated possible dissociations in a group of ex-servicemen who had sustained unilateral brain injuries affecting posterior areas of the left or right cerebral hemisphere. Using this group it was possible to control for confounds which may have affected comparisons across previous studies. The authors used the term 'double dissociation' so that only selective impairments (a significantly reduced performance for one ability, with normal performance for all other abilities) were considered. The response latency data from the study revealed

that there was a selective deficit in the processing of facial expressions compared with familiar and unfamiliar recognition. However, impairments affecting familiar face recognition were not entirely independent from unfamiliar face recognition. It seems from the prosopagnosia literature that there is not necessarily a clear distinction between processing of familiar and unfamiliar faces.

However, other behavioural and neurological studies suggest that there is a difference in the processing of familiar and unfamiliar faces (e.g., Begleiter, Porjesz, & Wang, 1995; Dubois, Rossion, Schiltz, Bodart, Michel, Bruyer, & Crommelink, M. 1999; Hanley, Pearson, & Howard, 1990; Hanley, Smith, & Hadfield, 1998; Rhodes, & Tremewan, 1993). Begleiter, Porjesz, & Wang, (1993) established an event-related potential (ERP) correlate of a visual memory process. This was demonstrated by a reduction in the amplitude of the visual memory potential (VMP) to repeated pictures of unfamiliar faces compared to that obtained with novel faces (occurring between 170 and 240 msec). This led them to explore the differences between familiar (famous) and unfamiliar (non-famous) faces (Begleiter et al., 1995). Using a repetition-priming paradigm the authors reported no significant difference in response times between primed and unprimed unfamiliar faces, but there was a significant difference for primed and unprimed familiar face response times. The ERP recordings showed that the VMP was reduced for both the primed unfamiliar and familiar faces compared to the difference in the unprimed stimuli. However, the reduction was much greater for familiar primed faces than for unfamiliar primed faces.

Further work investigating ERPs and face stimuli has revealed a negative potential that peaks around 170 msec (N170) and responds preferentially to human faces (Bentin, Allison, Perez, Puce, and McCarthy, 1986). The N170 is distributed over posterior-inferior aspects of the temporal lobes, with greater activation in the right hemisphere. Bentin and Deouell (2000) suggest it is associated with a face-specific structural encoding

mechanism that is not involved in face identification. This suggestion is supported by the demonstration that there are no differences in N170 for tasks requiring explicit identification of familiar (famous) or unfamiliar (non-famous) faces and tasks where such faces are ignored. Instead they find familiarity differences for the N400 peak (familiar faces produce more negative potentials between 350-550msec, and more positive potentials between 550-800msec), which they term face-N400. N400 is associated with semantic activity (Kutas and Hillyard, 1980; McCallum, Farmer, and Peacock, 1984) and consequently Bentin and Deouell suggest that the negative component of face-N400 is associated with semantic activity of face identification. This they argue fits with the Bruce and Young model (1986) of the Person Identity Nodes (PINs).

Other studies (Dubois, Rossion, Schiltz, Bodart, Michel, Bruyer, & Crommelink, 1999; Wiser, Andreasen, O'Leary, Crespo-Facorro, Boles-Ponto, Watkins, & Hichwa, 2000) have investigated the processes underlying familiar and unfamiliar face recognition with positron emission tomography (PET). Dubois et al (1999) used PET to measure regional cerebral blood flow distribution to presentations of familiar and unfamiliar faces. In this case, familiar faces were those that had been trained during an experimental setting and were not familiar to the participants prior to enrolment on the study. The unfamiliar faces were unknown faces that had not been trained. The data on a number of tasks, including gender judgments, showed that for both types of face stimuli there were bilateral activations of the fusiform gyri, including what has been labelled the fusiform-face area (a region in the right fusiform gyrus), which is considered to be specifically devoted to face processing (e.g., Kanwisher, McDermott, & Chun, 1997). However, differences were found for familiar versus unfamiliar faces. For the unknown faces there was activation of the left amygdala a structure involved in implicit learning of visual representations (Bechara, Tranel, Damasio, Adolphs, Rockland, & Damasio, 1995). However, the known faces were noted to show a

relative decrease of activity in the early visual cortical areas (i.e., areas V1, V2, and V3) in comparison to the unknown faces.

While Dubois et al (1999) considered the effect of familiarity on how we process faces, Wiser et al (2000) looked at the effect in terms of memory. They showed that in the novel-memory task (recognition task for new faces) the frontal areas were activated almost exclusively, a finding that links novel faces with learning and short-term memory. They also noted that there was activation of the anterior cingulate, reflecting greater attentional demands of the novel compared to the familiar task (McIntosh, Grady, Haxby, & Horowitz, 1996). However, the familiar face recognition task produced a wide distribution of activation, including visual areas that seem to act as memory-storage sites. The greater activation reported for familiar faces when compared to brain activations associated with new faces have also been detected in fMRI studies (Leveroni, Seidenberg, Mayer, Mead, Binder, and Rao, 2000). Such studies reveal bilateral activations involving the prefrontal, lateral temporal, and mesial temporal (hippocampal and parahippocampal regions) areas. However, it is interesting to note that the behavioural data from the face recognition task used in these studies did not reveal any differences in performance between novel and familiar faces. It maybe that differences in behavioural data only become evident when tasks do not require explicit recognition, such as Tong and Nakayama's (1999) visual search task.

Clearly, the face perception literature indicates that advantages for visually processing famous faces versus non-famous faces are likely. However, the variety of methods used in assessing effects of familiarity present a mixed view of how familiarity affects visual processing. In some studies familiar faces are famous celebrities, while in others they are non-famous faces that are trained within an experimental setting. Furthermore, unfamiliar faces may be entirely novel (if viewed once only and for the first time) or develop some level of familiarity during the course of the

experimental session. As referred to earlier, Tong and Nakayama (1999) took these important differences into account when developing their theory of robust representations. They used faces that were extremely familiar (the participant's own face) and compared target detection with that of unfamiliar faces that were analysed over time for familiarity (e.g., the first 50 trials through to the last 50 trials). Using this method they revealed two different effects of repeatedly viewing an unknown face: the first was rapid learning within the first fifty trials. This was followed by a second effect: a gradual decrease in reaction times. Furthermore, the robust image (own face) maintained an advantage over the unknown but recently learned faces, and showed no improvement in reaction times with repeated viewing. Not only does this study contribute to our understanding of face perception, but it also suggests a good method for how future studies of familiarity and faces should be conducted.

An important point about the face perception literature concerns the controversy about whether faces are "special". Some authors (Farah, Wilson, Drain, & Tanaka, 1998; Gauthier, Behrmann, and Tarr, 1999; Gauthier, and Logothetis, 2000; Gauthier, Skudlarski, Gore, and Anderson, 2000; Gauthier, and Tarr, 1997; Gauthier, Tarr, Moylan, Anderson, Skudlarski, and Gore, 2000) have suggested that differences between the processing of faces and other objects could be explained in terms of expertise. For instance, Gauthier and Tarr (in press) have used 'Greeble' stimuli to investigate the configural/holistic advantage found for faces. The Greeble stimuli are 3D rendered objects that are meaningless to any one who is untrained on them. With training, participants can learn to identify certain categories of Greeble (e.g., the Greeble category 'radok' will have a set of features and configurations that separate it from other types of Greeble). Studies using behavioural techniques and fMRI have revealed that in experimental conditions, trained participants can rely on configural information in these stimuli, and such tasks will even activate the fusiform face area. Therefore, in

considering the effects of familiarity on processing faces, it is possible that the same applies to other object categories. Hence, face processing may differ only in degree and not in kind.

Object Recognition

In reviewing the literature pertinent to face perception and recognition a variety of studies (neurological damage, brain imaging, and behavioural) have assessed the importance of differences in familiarity. Specifically, comparisons have been made between famous and non-famous faces. In contrast, object recognition studies may compare novel objects (e.g., meaningless 2D or 3D shapes) with real world objects, but they do not directly consider famous or non-famous real world object comparisons. This though does not signify that objects cannot be famous or highly learnt. Imagine how impossible a job it would be for marketers and advertisers to promote their brand if familiarity was negligible. Likewise, the task of weekly shopping would become a more tiresome job if our familiarity with items failed to help us differentiate between things we want and don't want (see Figure 5: e.g., sauces, shampoos, bleaches, and drinks all come in different bottles but we are unlikely to process and react to them as if they were only one type of object). Clearly, familiarity and experience allows us to both categorise and interact correctly with objects. And as a consequence, some objects will be more famous than others.



Figure 5. Bottles have distinctive features in relation to the substance contained inside. Our familiarity with products allows us to correctly categorise edible from inedible products (e.g., sauce and shampoo).

Naturally, any theory of object recognition has to consider the role of experience in terms of how items can be identified and named. Such theories can give a useful insight into how processing of famous and non-famous objects may differ.

Object invariance (viewpoint-independent versus viewpoint-dependent recognition).

In the recognition-by-components model (RBC) of object recognition (Biederman, 1987), it is suggested that complex objects are depicted as spatial arrangements of basic component parts, and that these parts come from a set of limited shapes. These shapes are items such as cylinders and wedges, which Biederman labelled "geons". The type and organization of the geons can be matched against the structural models of objects (see Figure 6). Biederman (1987) also proposed that geons are defined by properties that are invariant over different views (viewpoint-independent), and that this allows robust object perception even when an object is presented from a novel viewpoint.




Figure 6. An illustration of how objects can be made up of 'geons' as described in Biederman's (1987) Recognition-by-components model of object recognition.

As evidence for this theory of object invariance, Biederman and Cooper (1992) used repetition-priming studies. In these studies, participants are required to view and name objects. When a previously seen object is repeated from one block of trials to another block of trials, naming responses are facilitated (become faster) compared with naming of non-repeated items. However, the priming effect is reduced if a different exemplar of the same category is presented (an upright piano instead of a grand piano). This affords a measure of visual priming at the level of the structural model rather than merely conceptual priming due to re-accessing the same object meaning.²

Using this technique, they were also able to demonstrate that visual priming was reduced by changing the object components (geons) from one presentation to the next. However, priming was unaffected if both presentations of the same object exemplar were viewed with the same object components visible. Biederman concluded from these and other experiments that priming is invariant over many types of viewpoint change as long as the image components are preserved.

In contrast, other studies have provided evidence for viewpoint-dependent recognition (Bulthoff & Edelman, 1992; Lawson & Humphreys, 1996; Lawson, Humphreys & Watson, 1994; Palmer, Rosch, & Chase, 1981). Typically these studies demonstrate an increase in naming latency for objects rotated away from the normal viewing orientation. Lawson and Humphreys (1996) ascertained that matching performance was influenced by the similarity between object views, with dissimilar views producing slower responses. However, the difference in performance on similar and dissimilar trials was reduced by familiarity (experience) with the stimulus. This is attributed to activation of a number of object representations that are

² Note that there is a rudimentary effect of familiarity in that recent exposures can produce superior recognition compared to never seen before items.

abstracted from some image characteristics, but still remain view specific. Under such conditions matching becomes less reliant on representations that are strongly related to the initial view.

However, evidence from neurobiology suggests there are both viewpoint dependent and independent mechanisms for object recognition (Turnbull, Carey, & McCarthy, 1997). This research supports the possibility that the processes underlying both mechanisms may be due to separate anatomical regions. Studies of brain lesions (Solms, Turnbull, Kaplan-Solms, & Miller, 1998; Turnbull, 1997) reveal that patients may be left with selective access to particular types of representation. Turnbull (1997) reported two patients with a double dissociation between object orientation and object knowledge. One patient with visual object agnosia had impaired performance in naming objects, but could provide the upright canonical orientation for the items. Another patient, with a visuo-spatial disorder could name objects correctly, but was unable to identify the correct upright canonical orientation. Turnbull (1997) suggests that the information about the canonical orientation of an object may be derived from stored object knowledge and also from visual information about object structure. Furthermore, Leek (1998) examined the effects of orientation with mono-oriented (two-dimensional) and poly-oriented (three-dimensional) objects and found an interaction between the two types. While mono-oriented items were dependent on viewpoint, the poly-oriented objects were not. The results led the author to conclude that objects are encoded in stored representations at familiar stimulus orientations.

Memory studies have also contributed to this issue (e.g., Srinivas, 1995). Short-term recognition is shown to be unaffected by rotation (when visible parts remain the same), but long term memory is sensitive to these rotations. Srinivas (1995) proposes that long-term recognition is mediated by representations that specify the viewpoint exactly, and that the invariance found in short term memory can be accounted for by descriptions being generated from multiple successive views.

The explanations for the longer naming latencies of rotated objects have also focused on the nature of the mental transformation (i.e., from the visual pattern information to the match with stored representations). One theory is that of the double-checking account (Corballis, 1988; & De Caro, S. A., & Reeves, A., 1995), where the longer naming latencies are explained by the participants' need to double check the identity of the object for these views. However, the results produced by Lawson and Jolicoeur (1998) indicate that accuracy on a nonspeeded task deteriorates the further an object is rotated from the normal canonical view. If double checking was a valid account there should be no significant differences between accuracy levels for different orientations as the system of double checking could not influence performance.

An alternative explanation is that of the dual route account (Jolicoeur, 1990; Lawson & Jolicoeur, 1998). In this account there are two independent routes to recognition: a transform-then-match route and an invariant features route. In the former, an object is identified by being transformed to a canonical view and then matched to an orientation sensitive representation (e.g., mental rotation). The latter route achieves recognition through extracting orientation invariant features from the object and then matching these features to stored feature representations. In accounting for their results Lawson and Jolicoeur (1998) argue that the effects observed in the unspeeded task occur because canonically oriented objects are matched more efficiently to stored representations.

Another alternative and compelling account, which has been discussed previously in reference to face perception, is the population-encoding hypothesis (Perret, Oram, & Ashbridge, 1998). Perret et al. (1998) argue that the idea of mental transformations are untenable, as previous studies (Farah and Hammond, 1988; Perret, Mistlin, Chitty, Harries, Newcombe, & De Haan, 1988) have shown that some patients who show no evidence of increased naming latencies with rotation (i.e., they take a long time on all items) are still

able to make orientation invariant object recognition judgments. Instead they propose that there is a physiological explanation. When an object's appearance is learned in a particular way (familiar orientation), there will be more cells tuned to these views than to less experienced (unusual) views of the object. Assuming that the neural process guiding the behavioural output acts on the basis of input neuronal activity, then the output will take longer the more unusual the view. This is because the activity amongst the population of cells selective for the object's appearance will accumulate more slowly when the object is seen in an unusual view, as there are fewer cells responding. There is experimental evidence for this cell population theory in the form of cumulative response (spikes) recordings. Wachsmuth, Oram, and Perret (1994) recorded cells in Macaque temporal cortex that respond to a whole body. Three conditions were recorded: where only a head was visible, where only the body was visible, and where both head and body were visible. The cell recording revealed that of the 35 cells recorded, 28% were responsive to the head alone, and 10% were responsive to the body alone. However, the response to the whole body accumulated fastest of all, and occurred for two reasons. Firstly, a sub-population of 19% of cells responded to only the head and body together, and secondly the sub-population for the head alone (28%) and that of the body alone (10%) will also respond to the combination. Hence, it can be argued that total brain activity will be greater in response to the visual image of an entire object than to one component. Likewise, those components that are more familiar or more often encountered will have more responsive cells (e.g., face versus body).

There is also growing support from other temporal cortex evidence (Logothetis, Pauls, and Poggio, 1995; Miyashita and Chang, 1988; Perret, Smith, Potter, Mistlin, Head, Milner, & Jeeves, 1984) that cell selectivity is biased towards images we experience as adults. Logothetis et al (1995) trained three monkeys to recognise novel objects from a specific viewpoint. When the monkeys were required to recognise the objects from viewpoints other than

those experienced in training, they found the task much more difficult to complete. Electrode recordings from inferotemporal cortex showed that a number of neurons were 'remarkably' selective for the trained views, while other neurons responded to the presentation of unfamiliar objects. Furthermore, Foldiak (1991) has suggested that learning mechanisms are coupled with visual experience. This might account for cell selectivity for objects that are frequently seen and also for cell selectivity for the viewing conditions in which objects are seen.

The role of familiarity thus far has been constrained to the effect of familiar versus unfamiliar viewing conditions, rather than the direct comparison of familiar and unfamiliar items. However, the various proposals of how we recognise objects are pertinent to the issue of perceptual differences in famous and non-famous items. In particular, the population-coding hypothesis clearly suggests that famous items should be processed faster and/or more efficiently than non-famous items, and that this may occur to a greater degree for objects that are 'over-learned' (i.e., objects that through environmental salience require an uncommon level of experience) such as face stimuli. What also appears to be of import to the overall question of familiarity is how learning and experience determine efficient recognition and interaction. One specific area that reveals learning effects is the study of categorisation.

Categorisation and expertise.

Categorisation is important to our understanding of how people class and group objects in their environment. Typically, we recognise objects from their basic level of categorisation (e.g., bird, table, car etc). This basic level is the point where classification is most efficient. There are also subordinate (e.g., a penguin is a type of bird), and superordinate (e.g., a bird is a type of animal) classes of categorisation that are accessed according to a person's level of expertise with the object or the task at hand. For example, a birdwatcher out

on a day's field trip will not be simply looking through his/her binoculars reporting 'there's a bird, there's another bird...'. Rather, the birdwatcher will process multiple levels of categorisation (e.g., 'oh there's a bird, it's a wader, and yes I believe it's a common sandpiper'). In contrast, the average person on investigating a strange movement outside will only consider the type of animal it is (e.g., 'oh it was only a bird'). This demonstrates that our ability to categorise objects will depend on our experience and familiarity with the items.

In turn, the way in which individuals conceptualise visual stimuli will affect the overall perception of the item. Lin and Murphy (1997) propose that background knowledge is required to explain (1) how features are associated to form a coherent concept, and (2) why some features are more relevant than others. In order to support their argument a series of experiments were conducted to establish an effect of conceptual background knowledge on both picture categorisation and perceptual identification. Two groups of participants were required to learn several categories of novel (artificial) objects. One group would learn one interpretation for each category (e.g., category *x* were 'animal catchers'), while the other group received completely different explanations (e.g., category *x* were 'pesticide sprays'). Both groups saw exactly the same pictures of the learning exemplars, with the same numbered parts being described. However, the key manipulation was that the features crucial to the function of one of the category interpretations were not crucial to the other. Participants were tested on a categorisation task, which involved judging whether an item belonged to a category they had learned. The items displayed either had the crucial features for one group without the critical features of the other group or vice versa. There was also a control condition where the item had no crucial features from either group descriptions. The results showed that when the item had consistent features (the crucial details) participants were significantly more likely to report it as being a category member (87% non-speeded task, 72% speeded task) than

those items with inconsistent features (19% non-speeded task, 26% speeded task). Furthermore, reaction times were 534 milliseconds faster in responding to consistent versus inconsistent items. There is a clear effect of background knowledge on the way in which these participants categorised the newly learnt items.

However, Lin and Murphy (1997) also wished to see if background knowledge would influence perceptual identification of items. Undergoing the same training methods as described previously, a different set of participants were required to complete a part-detection task. Here, either crucial or non-crucial parts could be missing, along with category prototypes that included all parts. Participants simply had to indicate whether the test items had all the parts. The results revealed that when an item was missing a less important part it was noticed less than if it had been crucial (74% versus 83% accuracy). This held true for speeded, non-speeded and briefly presented tasks, and led the authors to conclude that important features in terms of background knowledge are more perceptually salient than those features that are less important. Other researchers (e.g., Schyns and Rodet, 1997) have also speculated that the features used for representation in object recognition are flexible, in that they are adjusted to the perceptual experience and the categorisation history of the individual.

The topic of object categorisation has led to an interesting debate on just how unique face recognition is as a perceptual process. In particular recent research has questioned whether the fusiform gyri (which is often referred to as the fusiform face area) is truly specialised for faces, or whether it is used in subordinate category judgments (of which faces are more often categorised at the subordinate level). For instance, Gauthier, Anderson, Tarr, Skudlarski, and Gore (1997) reported that subordinate matching of objects activated the fusiform and inferior temporal gyri in a similar manner to that observed in face perception. However, this did not occur when participants were required to match the same objects at the basic-level of categorisation. Moreover, face

recognition can be considered to be a subordinate level categorisation where we differentiate between individual faces, whereas most other objects are usually considered at the basic level (e.g., when we recognise a chair we do not normally categorise the style of chair). Perhaps the distinction between categorising faces at the subordinate level and categorising non-face objects at the basic level can account for the difficulties prosopagnosics have with face recognition in comparison to general object recognition? It could be the case that the deficit in face recognition found in prosopagnosics is actually a deficit in subordinate level categorisation? Gauthier, Behrmann and Tarr (1999) found evidence for this hypothesis using a match-to-sample task. Recognition sensitivity was compared for face and non-face objects in two prosopagnosic participants. Using this demanding memory task (i.e., participants viewed an item then had to decide whether twelve other items were the same or different) the prosopagnosic participants were more affected by manipulations involving the level of categorisation than normal controls. Furthermore, in simultaneous matching tasks, (both where there was no time limit and where presentation times were set at 1500 msec) these prosopagnosics showed a deficit for both face and non-faces alike. Therefore, it seems that there is a relationship between subordinate categorisation and the deficits displayed in prosopagnosics. Controls show fusiform activation on both face tasks and subordinate level object tasks, and prosopagnosics can be shown to demonstrate wider deficits when subordinate level categorisations are required.

To take this argument a step forward and associate the role of expertise with these findings, it is important to demonstrate that experts with a given class of stimuli can be shown to process the stimuli in a manner similar to that of faces. In one of the earliest studies to imply this, Diamond and Carey (1986) demonstrated that when dog experts were tested with Yin's face inversion effect they produced impairments in recognising the faces of dogs just as if they had been human faces. This suggests that there is a specific relationship

between expertise and configural processing, in that spatial relations between object parts may aid expert object recognition. Other support for this position has come from functional imaging studies (fMRI), where bird and car experts (Gauthier, Skudlarski, Gore, & Anderson, 2000) have shown expertise effects involving the right and left fusiform face area as well as the occipital face area.

However, this debate is one that is still raging. Its importance here is to suggest that while the background literature on faces is far more advanced in the understanding of familiarity, there are indications that objects are likely to exhibit fame effects also. Theories of object recognition focus on how an object is usually viewed compared to different orientations or even different background information, suggesting familiarity is of immense importance to our recognition experience. Likewise, the object versus face debate indicates that any significant effects of famous faces may also be evident in famous objects, though possibly to a lesser degree.

Objects in Scenes.

A discussion of how we perceive familiar objects cannot be complete without considering the environment in which they are normally presented. Just as Lin and Murphy (1997) debated the importance of background knowledge on object recognition, it is necessary to investigate how the scene background or context might influence object perception. Here, the role of expertise and familiarity with stimuli will be highlighted because familiar objects are usually viewed within a familiar context. For example, if you went abroad on holiday and saw an unexpected but familiar face, recognition would probably take longer than normal. In fact, unexpected scene context can even lead to strange forms of misidentification for close family relations (Thompson, 1986). Thompson (1986) reports an occasion where one of his students' parents flew from Australia to London. Unknown to the parents the student was also due to fly out to London. It was arranged that the student would stand at a bus stop near her parents' lodgings with a companion positioned close by to make observations. When the parents emerged from their lodgings and saw their daughter they stopped abruptly. The father then approached the daughter and said hello to her. As instructed, the student turned towards her father, only to look straight through him (as if he was unknown to her). The father promptly broke off his greeting and apologised for the mistaken identity. Such forms of unusual identification behaviour for highly familiar stimuli in unfamiliar scenes, leads to a need to consider scene recognition effects alongside any familiarity effects that may occur in object recognition. While this thesis aims to assess fame effects for individual/isolated objects in order to inform at the theoretical level, it is always useful to try and consider how these effects will interact with more complex situations. Such considerations can help to evaluate the external validity of extrapolating laboratory behaviours to the real world.

Scene background can play an important part in how we categorise objects. As described earlier, people typically categorise objects at the basic

level, but with expertise they can categorise at the subordinate level. But, this leaves open the question of when do we use superordinate categorisations? Murphy and Wisniewski (1989) postulated that while basic level categorisations were both specific and distinctive (e.g., a fork is an item that has unique requirements such as prongs and handle), superordinate level categorisations were more abstract and not specific (e.g., cutlery refers to spoons, knives, forks, etc). However, a further difference is that superordinate categories tend to refer to plurals and not individual items. For instance, children understand superordinates as collection terms rather than class terms (Markham & Callanan, 1984). Murphy and Wisniewski (1989) propose that relational information is contained in superordinate concepts (e.g., relative locations of exemplars, functional relations between exemplars, and how concept exemplars interact with non-exemplars). Hence, if these relational forms of information are contained in such superordinate level concepts, then the basic level advantages for objects in isolation (e.g., Brownell, 1978; Jolicoeur, Gluck, & Kosslyn, 1984; Murphy & Brownell, 1985; Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976; Smith, Balzano, & Walker, 1978) should diminish when the objects are presented in scenes (i.e., they are presented amongst other exemplars from that category).

In order to test this hypothesis, Murphy and Wisniewski (1989) conducted an experiment whereby either a basic or superordinate level category name was displayed, followed by a 150-millisecond presentation of a scene or individual object. A cue indicated the object that should be compared to the category name (e.g., the category name 'furniture' followed by a cued chair would require a 'yes' response as the two matched). The reaction time data showed that when an object was displayed in isolation there was a 162msec response advantage for basic level decisions when compared to the superordinate categories (as predicted from previous studies). However, when the object was displayed in a scene the difference in reaction times was reduced to only 30msec (non-significant effect). The authors concluded that

the pattern of results was consistent with superordinate categorisation benefiting from configural information.

More recent work on the role of object categorisation in complex scenes has addressed the issue of expertise with items (Archambault, A., O'Donnell, C., & Scyhns, P.G., 1999; Werner, & Thies, 2000). Here the previous finding that how a person is familiarised with a set of items can influence how the items are perceived is extended to objects in naturalistic scenes. Archambault et al. (2000) used a change detection procedure to test whether features that were irrelevant for encoding an object were noticed significantly less than relevant features. The task required one group of participants to learn a set of computers at a specific level (e.g., 'This is John's computer') and a set of mugs at a general level (e.g., 'This is a mug'), while a second group learnt mugs at the specific level and computers at the general level. Once these stimuli had been successfully learnt, the participants viewed office scenes that were separated by a blank and then replaced with a change. The change could be a disappearance of a mug or computer, or a mug/computer could be replaced by another mug/computer. The key measurement was how many alterations between the two scenes it would take for participants to correctly locate and identify the change. The results revealed that when participants knew an object at the specific level of categorisation, they perceived a change almost immediately (1 repetition of the scenes), but it took much longer for objects learnt at the general level (7 repetitions of the scenes). This difference only occurred when the change involved a replacement; when the change was a disappearance, performance was equal between specific and general categories. This indicates that it was not a case of the specific items being scanned or attended more than the general items, but rather there was a difference in how the features of the items were perceived when both were equally scanned/attended.

It is interesting to see that an individual's experience and familiarity with a stimulus is an important factor in how an individual perceives the

stimulus both in isolation (a rather unnatural situation), and within naturalistic and complex scenes. This raises the question of what effect scene perception has on object identification. For instance, is object recognition independent of the context in which it occurs? Or does the expectancy raised by the scene influence the perception of the objects? Different theories of scene perception offer different answers to this complicated issue. According to the "functional isolation model" (Henderson & Hollingworth, 1999a), bottom-up visual analysis is enough to distinguish between entry-level categories of objects, and hence object identification is isolated from scene knowledge. This contrasts with the "perceptual schema model" (Biederman, Mezzanotte, & Rabinowitz, 1982) that proposes that expectations about a scene interact with the perception of the objects in it. Specifically, it suggests that the scene schema (memory representation) contains information about the objects and their spatial relationships, leading to facilitation in processing for objects consistent with the scene. This model is similar to the third main theory, the "priming model" (Palmer, 1975), which states that the scene is matched against long-term memory representations that prime representations of consistent objects. Therefore, less perceptual information is needed for consistent items compared to inconsistent items. The key difference between the priming model and the perceptual schema model is that the priming model argues that scene knowledge only influences the criterion used for recognising that an object type is present, whereas the perceptual schema model proposes that the scene directly influences the perceptual analysis of the object token.

So, what evidence is there for interactions between scene background and object recognition? A number of studies have attempted to determine how consistency between objects and scenes affects eye movements, processing time and recognition. An early study investigated how we perceive objects after we have viewed a scene (Palmer, 1975). Line drawings of objects were arranged into pairs, so that each pair consisted of objects with

a similar shape and appearance. An example is a mailbox and a bread bin; both are different objects expected in different scenes, but have similar characteristics. At the start of each trial a line drawing of a scene was presented for 2 seconds, followed by a delay of 1300 milliseconds. Then an object was presented for 20-120msec. Participants had to write down the name of the object they had just seen. The results showed that participants were more accurate in naming an object when it was consistent with the scene that preceded it, (e.g. a kitchen scene followed by the bread bin), than when the item was inconsistent (e.g. the kitchen scene followed by the mail box). If a kitchen scene was viewed, and then followed by a mailbox, participants tended to erroneously report the object as being a breadbin. Even though the target item was absent in the preceding scene, the context primed the perception of the singly displayed object.

More recent studies have concentrated on the role of single items within the scene itself. One such study that led to a great amount of interest used a method that employed "relational violations" (Biederman, Mezzanotte, & Rabinowitz, 1982). These violations fall into five categories: *Support*: when an object does not appear to rest on any surface in the scene (e.g., a sofa in the sky); *Interposition*: when the background passes through the object (i.e., the object looks transparent); *Probability*: when the object is unlikely to appear in the scene (i.e., inconsistency); *Position*: when an object is likely to occur in a scene, but not in the position depicted (e.g., a fire hydrant in a street scene appearing to be placed on top of a waste bin); and *Size*, when an object appears to be too large or too small in comparison to the scene it is in. Using constructed images, two experiments were conducted to assess violation detection and to investigate the role of these violations on object recognition. Participants were shown a target name at the beginning of each trial. When the image was presented (150msec) an object in the scene was cued. Following a mask the participants judged whether the target name was the same as the cued object. Three main results concerning object detection

were reported: first, objects with relational violations were detected more slowly and less accurately than normal objects; second, detection of normal objects was unaffected by another object in the scene undergoing a violation; and third, physical violations were not any more disruptive than semantic violations (e.g., support versus probability). In terms of violation detection, they found that semantic violations were actually detected more accurately than physical violations. The authors argue that this is evidence that an object's semantic relations are accessed along with its physical relations.

A number of methodological criticisms have been levied at this study and have been controlled in later studies (Hollingworth & Henderson, 1998a; Henderson & Hollingworth, 1999a; Hollingworth & Henderson, 1999). These will be discussed later in this section. However, one point that should be made now is that by presenting the target name prior to the scene image, participants were provided with a clue as to which part of a scene would be most likely to contain the specified object. Therefore, participants had an advantage in guessing where a consistent item would be placed. While it is clearly important to examine the effect of displaying the target name at the end of the trial, which was not done by Biederman et al. (1982), it is important to note that the presentation of the object name before the scene creates a more real-world task. It is very rare for us to choose to walk into a room without a plan of what we want to do or find in it.³ So, despite criticisms, Biederman et al. (1982) have provided a useful method of exploration of complex scenes.

A number of studies have been continuations of this work (Boyce & Pollastek, 1992; DeGraef, 1992; Hollingworth & Henderson, 1998a; Henderson & Hollingworth, 1999a; Hollingworth & Henderson, 1999), with a large number of researchers contributing to the debate of whether scene context

³ For example, if someone enters the kitchen in order to switch on the kettle, the task grabs their attention and they might not notice that the television remote control is also in the kitchen when it should be in the living room. It is only on rare occasions, like sitting in a doctor's waiting room that people have nothing in particular to attend to, and therefore are more likely to notice an "odd event".

affects object identification. DeGraef (1992) used line drawings and stimuli very similar to Biederman et al (1982). However, in these scenes a number of non-objects occurred along with real objects. Participants were asked to count the number of non-objects, whilst the experimenters recorded fixation times across the scenes. The fixation times for the real objects in violated scenes were compared to a base condition with no violations. Longer fixation times were found for objects undergoing the violations of "probability", "position", and "size". However, these longer fixations only occurred in the later stages of scene viewing, as at earlier stages violated and normal objects were fixated for the same amount of time.

Another method for measuring object recognition in complex scenes is the "wobble" paradigm used by Boyse and Pollastek (1992). Participants are asked to focus on a fixation spot that is then replaced by the image of a scene. After 75msec, an item in the scene wobbles (apparent motion) causing the participant to fixate it. They then have to name this item as quickly as possible. Response latency comparisons were made between scenes that had a consistent background (e.g., if the target was a bicycle it would be placed into a street scene background), and scenes that were inconsistent (e.g., if the target was a fire-hydrant it might be placed in a swimming pool scene). In addition, trials were included that had non-scene controls for each type. The main finding was that participants were 50msec faster to name target objects in the consistent background than in the inconsistent background. The researchers then went on to manipulate when the background was presented. Scene context information was made available either on the first, second, both, or neither fixations. The new set of results showed that each fixation had a context effect, with the first fixation in the centre of the display having a larger effect than the second fixation on the target. Boyse and Pollastek (1992) concluded that scene background affects the probability of identifying a target object in a single brief presentation and also the time taken to name a target object when it is fixated and clearly seen on a second fixation.

However, they point out that the use of context in decisions about objects is not mandatory. For example, Malcus, Klatsky, Bennet, Genarelli, and Biederman (1983) showed that when a target object is cued prior to the scenes appearance and participants are instructed to attend to just the object and ignore the background, then context is not used. Again there is strong evidence that the information extracted from complex scenes is task dependent, with focus of attention being a key issue.

Recent research by Hollingworth and Henderson (Hollingworth & Henderson, 1998a; Henderson & Hollingworth, 1999a; Hollingworth & Henderson, 1999), explored the use of context in complex scenes. They tried to control for problems with the original Biederman et al. (1982) study and investigated the role of eye movements and saccadic integration (Henderson, Weeks, & Hollingworth, 1999). Initially, they replicated Biederman et al's (1982) study and found similar effects, in that object detection was better for semantically consistent objects than for inconsistent objects (Hollingworth & Henderson, 1998a). Subsequently, they attempted to remove any response bias by manipulating the catch trials (i.e., when the target name was not present in the scene, but was reported by the participant). Unlike the original study, they controlled for participants responding "yes" more frequently for consistent trials (where it is possible that participants would have believed the object was likely to have been in the scene, whether it was or was not). When this response bias was removed from the experiment, the results showed no advantage for the detection of consistent objects. A further control was added so that instead of presenting the target name before the scene, it was displayed afterwards. This eliminated the participants' ability to predict the location of the consistent items. Interestingly, the results indicated that there was no effect of consistency.

These subsequent results contrast with the findings of Biederman et al (1982), and are used by Hollingworth and Henderson (1998a) to argue that object perception is not facilitated by consistent scene context. This reasoning

fits their “functional isolation model” and contrasts with the “perceptual schema model” and the “priming model” that were outlined previously.

Although this study indicates that scene context does not contribute to object recognition, it is also clear that scene information was processed to some level during the 200ms exposure. This is evidenced in the greater amount of false alarms (reporting an item as present, when it was absent) for consistent test items than inconsistent items. This indicates that on some trials the target item was not in the scene but the extracted scene knowledge lent itself to a guessing strategy. For example, participants may have developed the strategy of "if the target usually occurs in the scene, choose it". If this was indeed the case, it proved to be a useful tactic, leading to higher hit rates (correctly reporting an item as being present) in the consistent condition. This debate is clearly not fully resolved. However, it does provide evidence that in certain circumstances the familiarity an individual has with scenes and the objects contained within them can produce beneficial processing and/or retrieval.

Visual familiarity: what remains to be understood.

Two important points can be drawn from the preceding reviews. First, most research involving the role of familiarity in visual recognition has been restricted to the domain of face perception. Theories of face recognition make distinctions between familiar and unfamiliar faces, and these theories have been supported by both experimental and neuropsychological evidence. However, evidence for such distinctions within the general object recognition literature are scarce. Any effects of visual familiarity in objects appear to have been revealed as a by-product of other questions (e.g., such as recognition of objects at different orientations) that inevitably touch on issues relating to learning, knowledge, and expertise. Therefore, there seems to be an important gap in our understanding of human visual perception about general familiarity effects in object processing.

Second, visual familiarity should not be viewed as a simple dichotomy between familiar and unfamiliar items. Rather, there must be a progressive perceptual learning of any visual stimulus until it eventually reaches an encoding and retrieval stage that is optimal and cannot be processed any more efficiently (Tong and Nakayama, 1999). This optimal stage of experience with an object is what I shall term 'fame' throughout the remainder of the dissertation. This is when an object has been highly learnt with significant associations both perceptually and semantically. Whereas, we commonly associate human faces with being famous, we also have to consider that many forms of stimuli in our environment are also famous (e.g., the Millennium Dome).

This thesis aims to open up the discussion on famous versus non-famous effects on visual processing. Although it seems obvious that familiarity with an object should speed access to memory, it is less clear whether familiarity should similarly affect the speed of visual processing. By comparing highly learnt stimuli with stimuli that are just recently learnt I intend to demonstrate that (1) fame is an influential factor on how we process visual information, greater than mere familiarity; (2) fame is a term that should not be restricted to the domain of faces; (3), fame can be implicitly measured, offering a potential application for practical issues (e.g., consumer marketing evaluations).

In the experiments that are reported in the following chapters the test of famous versus non-famous visual stimuli used a change detection task that involved brief presentations. The purpose of these studies was to explore perceptual processing advantages bestowed by high familiarity, or 'fame'. To bypass the semantic and verbal memory demands of explicit recognition tasks, I used an implicit measure of image processing that involved locating a change in successive object presentations. This task is based loosely on the 'change blindness' tasks of Rensink, O'Regan, and Clark (1987). "Change blindness" refers to an inability to detect changes made to a scene after a brief

interruption. This interruption can be in the form of saccadic eye movement or through an external source of disruption, such as placing a blank interval between scenes. It is the latter option, in particular, that provides researchers with a useful tool for investigating the attentional and visual processing mechanisms that involve this failure to detect changes in scenes.




Figure 7. The hen changes from facing left to facing right. (Courtesy of Hollingworth and Henderson, 1999)

The main method of manipulation previously used has been the “flicker paradigm” (e.g., Pashler, 1988; Rensink et al., 1997). In this method a fixation point is presented and then an image of a scene is shown for 250 ms. It is then replaced by a blank inter stimulus interval of a 80 ms duration and the scene is then re-presented with an alteration made to it. The sequence is repeated until the participant makes a response to indicate that the change has been detected (see figure 7). Rensink, O’Regan, and Clark (1997) used this technique in order to assess whether changes to important object/areas in complex scenes were detected more rapidly than changes to incidental features. These areas were rated by a separate group of five participants who looked at the scenes and then reported the items in them. Central interests (important areas) were judged as items that were reported by three or more observers, and marginal interests (incidental areas) were items that were not mentioned by any of the group. Both types of change proved difficult to

detect, though the scene changes to marginal interests needed more than twice the amount of sequence alternations needed for central interest changes. On average it took 4.7 sec to detect a change to a central interest, and a long 10.9 sec for a marginal interest. These findings led the authors to propose that to detect an object change in a scene, that object must be given focused attention. Without focused attention it has been hypothesized that the visual memory trace is overwritten by subsequent stimuli (Enns & Di Lollo, 1997).

As previously discussed, change detection has been used to assess categorisation and expertise effects in object recognition (Archambault, O'Donnell, & Scyhns, 1999; Werner and Thies, 2000). Therefore, change detection offers a useful technique to adapt, and forms the basis of the experiments reported here. In these experiment a two-item display underwent a change to one of the items. Participants, had to make a categorisation decision and decide whether left or right items had changed.

Effects of familiarity on perceptual encoding and memory were of principal interest. If visual fame bestows an advantage on perception (as opposed to naming memory), change detection should be more accurate with famous than non-famous objects. This may occur through general speed of processing advantages and/or attentional differences, such as attentional capture by famous items or reduced attentional resource allocation for famous items. However, differences in how successfully information is encoded into memory may also interact with processing effects. If famous items are processed faster than non-famous items, encoding in Conceptual Short Term Memory (Potter, 1993, 1999) should allow representations of the perceptual details to be more structured for famous items (e.g., more information will have become structured in CSTM before the next conceptual/meaningful mask). Whereas, non-famous items may only permit encoding of semantic/categorical meaning, without details of the perceptual distinctions being available for accurate change detection. For example, a well-known landmark (e.g., Statue of Liberty) through rapid processing may

have its features encoded allowing comparisons to be made with other landmark items. While an unknown landmark with slower processing may only have semantic details available, such as "it's a building". If this building changed to another unknown building a failure in change detection would result. For instance, repetition blindness may occur because of a failure in perceiving the two events as being distinct, due to the semantic similarity of the items (Kanwisher, 1987; Kanwisher & Potter, 1989, 1990). Furthermore, expertise with different classes of object (e.g., faces versus landmarks) may also result in greater perceptual versus semantic advantages for detecting changes.

The following investigation addressed the issue of fame and its effects on visual processing and memory with a number of different object classes. Chapter 3 presents a series of studies using faces as famous and non-famous stimuli. Considering the literature reviewed at the beginning of this chapter, if fame has an effect on processing ability, it will most likely be evident in face stimuli. Chapter 4 aimed to extend the fame effects evident in faces to an inanimate class of object: landmarks. Landmarks provide a useful comparison with faces in that they are equally unique (e.g., there is only one 'Mel Gibson' and only one 'Eiffel Tower') and have clear examples that are famous. Chapter 5 explored the role of categorisation in object recognition and familiarity. Consumer products (e.g., shampoo, bottled drinks, etc) were used as famous and non-famous stimuli. These were different to objects such as the landmarks in several respects (1) they were not unique exemplars (e.g., Persil washing tablets are available in multiples unlike the Statue of Liberty), (2) they were likely to be more personally experienced (i.e., handled and used on a daily basis) potentially producing a greater affective response, (3) they were more easily categorised at different levels (e.g., subordinate = Persil, basic level = washing powder, superordinate = detergent), and (4) lexical information was contained within the object. Furthermore, Chapter 6 provided a demonstration of how the change detection technique can be

developed to reveal differences between individual famous images. The chapter describes a study that was conducted during the run-up to the British General Election in 2001, which indicated that not all political leaders have the same perceptual impact.

Finally, in Chapter 7 I discuss how effects of familiarity for highly learned stimuli informs current object and face recognition theories, along with what future directions need to be taken in this area of visual familiarity. The following chapter presents the general methodology adopted in the ensuing experimental studies.

Chapter 2 - General Methods

This chapter details the change detection method that was common to the experiments of Chapter 3 (faces), 4 (landmarks), and 5 (consumer products). Where details differ between experiments the relevant chapter will give clear details.

Participants

British undergraduates volunteered to participate in exchange for course credit. All reported normal or corrected to normal vision and were naive to the purpose of the experiment. Informed consent was obtained prior to participation.

Apparatus and Stimuli

The experiments were programmed in Psyscope (v.1.2.2) and run on PowerMacintosh computers. Responses were recorded via the computer keyboard. A chin rest was used to stabilize head position 97 cm from the computer monitor. Testing was conducted in a small room with low ambient illumination.

Each experimental session used a set of eight items: four non-famous (two from one category, two from a different category) and four famous (two from one category, two from a different category). For the face experiments (Chapter 3) the categories were male or female. The categories for the landmarks experiment (Chapter 4) were building or monument, and the categories for the product experiments (Chapter 5) related to the container type (e.g., box or bottle, can or bottle). All items were greyscale photographs. While low-level featural differences were inevitably going to occur (e.g., brightness, contrast, and position), items were chosen so that predictions based on low-level perceptual differences were reduced. For instance, where the 'Statue of Liberty' and the 'Eiffel Tower' differ in terms of contrast (light

and dark respectively), the two non-famous monuments were matched with these items on shape, ornateness, and contrast (again one light and one dark). However, this was a rather subjective measure and potential differences probably persisted. Interestingly, evidence from Ro, Russell, and Lavie (2001), indicates that when changes are made to arrays containing faces and non-face objects, faces are better detected than these other items, even when there are far greater perceptual differences between the changes of the non-face objects. This indicates that changes, albeit very similar, to items of interest or emotional significance can lead to better change detection than changes to less significant items with larger perceptual differences. Nevertheless, attempts were made here to minimise differences, but be aware they were not entirely eliminated.

Procedure

A typical trial is illustrated in Figure 8. Each trial commenced with a 1000 ms presentation of the fixation cross. Four 100 ms displays were then presented successively (without any inter-stimulus intervals). These consisted of (1) the first pair of items, (2) the pattern mask presented over each item, (3) a second pair of items, and (4) the pattern masks. The entire sequence lasted 400 ms. On every trial, one item in the second pair was different from the items in the first pair, but the location of the different image was random. An item from one category only changed to an item from the same category. The item that changed is called the target and the item that remained unchanged is called the distractor.

The participants' task was to report the change location (left or right) using a key press. They then reported the category of the items by pressing one of two keys for the left image and one from a different pair of keys for the right image. The category task prevented participants from attending to only half the display. Trials were self-paced with at least 1.5-second intervals. Only

trials for which both categories were correctly identified were included in the analysis of the change detection performance.

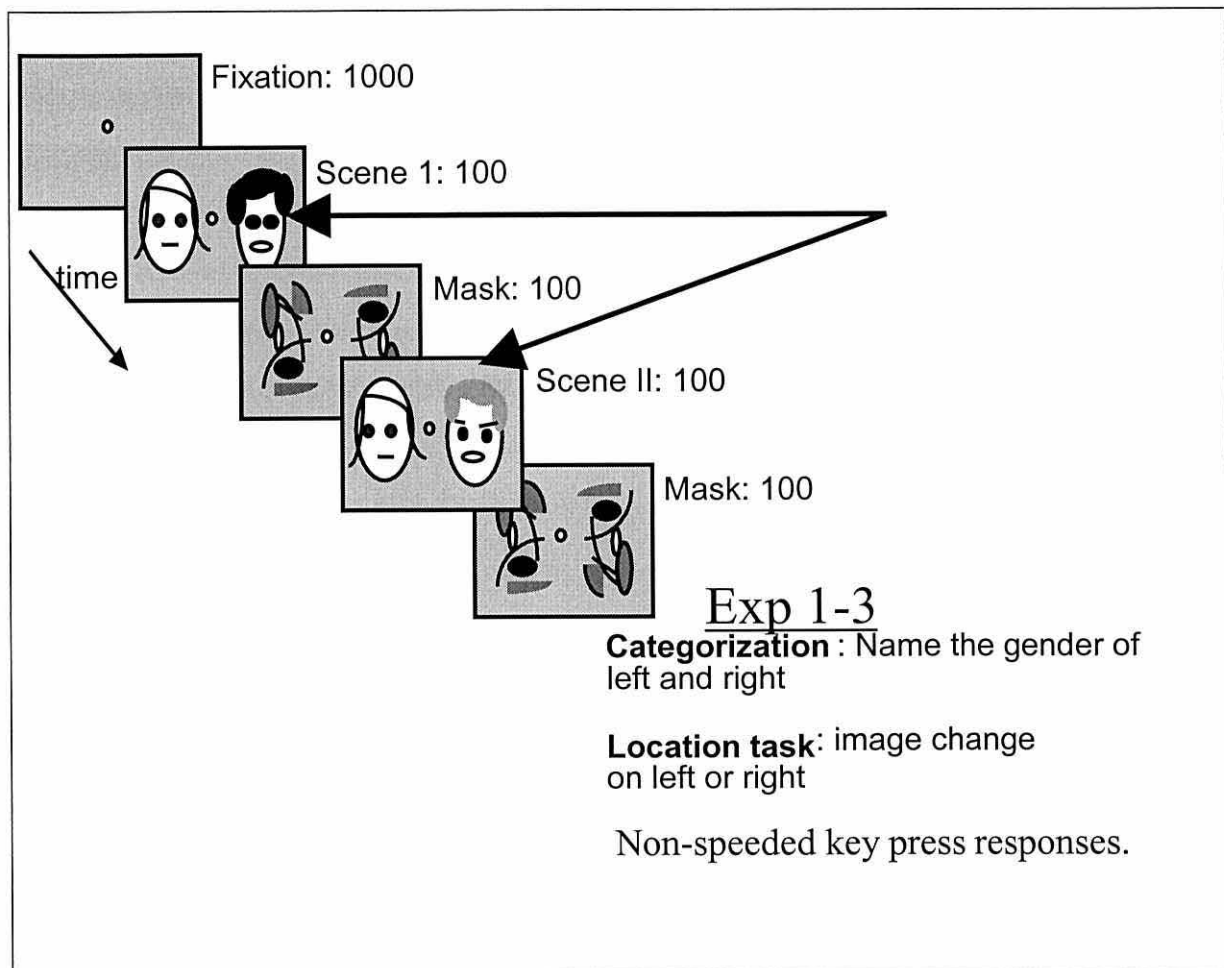


Figure 8. An illustration of a trial. Following a 1000ms fixation cross a pair of faces are briefly (100ms) presented followed by a scrambled image mask (100ms). A second pair of faces are presented (100ms) with mask (100ms). The second faces image includes one face as presented in the first image, and one, which is a completely different face.

A test session consisted of 384 trials presented in a pseudo-random order such that equal numbers of trials were presented for each condition and all necessary counterbalancing was maintained. Note, that while only eight items were used per participant in 384 trials, for full counterbalancing to occur, combinations of individual items and conditions were seen only once. For example, famous item A would only change to non-famous item C in the right visual field with famous item B as distractor on one occasion. Therefore,

only small sets of stimuli were tested. Within each set and category group, each item was changed into each of the other items an equal number of times. Moreover each item was presented an equal number of times and attributes (e.g., fame, category) of the items in the first presentation could not be used to predict the location of the change contained in the second presentation.

The experiment used four change conditions: a famous item changed into another famous item (FF), a non-famous item changed into another non-famous item (NN), a famous item became a non-famous item (FN), and a non-famous became a famous item (NF). On half of the trials, the distractor item was famous and on the remainder it was non-famous. Changes occurred on the left on half the trials and on the right on remaining trials. Additionally, changes occurred half the time to one of the two categories and half the time to the other category. Simultaneously presented item pairs were matched for category on half the trials and mismatched on remaining trials.

Participants completed 20 practice trials at the beginning of the session, using non-famous items that were not part of the test session. After the change detection task was completed participants were debriefed and thanked for their participation. The experiment took 45 minutes to complete.

Chapter 3 – Fame Effects in Faces

The social nature of humans makes face recognition an important functional aspect of visual perception. Face recognition involves two main components: (1) processing the visual stimuli comprising the face, and (2) linking the resultant perceptual representation with information (e.g., names, contexts, affect, etc.) in memory. According to Bruce and Young's (1986) face recognition model, the former component is separated into two distinct processes: view-centred descriptions (i.e., emotion, gender, and expression) and expression independent descriptions (i.e., face recognition). It is only the encoding of the expression independent descriptions that can lead to the retrieval of names (i.e., activation of the Person Identity Nodes)⁴. Therefore, an important aspect of this process is familiarity, i.e., the extent of previous exposure to an individual. Although it seems obvious that familiarity with a person should speed access to memory, it is less clear that familiarity should similarly speed visual processing. I sought to explore this question by assessing face perception using brief presentations and a change detection task that did not require identification. By comparing performance with famous versus just-familiar faces I was able to indirectly assess the effect of familiarity on the visual processing of faces.

Previous studies of face perception indicate that two different perceptual processes may mediate face recognition (Bartlett & Searcy, 1993; Moscovitch, Wincour & Behrmann, 1997; Rhodes, Brake & Atkinson, 1993). One is a *featural* mechanism that uses facial features, such as eyes, ears, etc., to identify individuals. The second is a *configural* mechanism that integrates information from the whole face and matches it to an internal template (e.g., Farah, Wilson, Drain and Tanaka, 1998). The latter process seems necessary because faces form a very homogeneous class of stimuli in which single

⁴ Refer to Chapter 1 for a more detailed overview of face recognition models.

features do not uniquely specify each member. Moreover, recognition is remarkably good in spite of variations in viewpoint, lighting, face expression, make-up, hairstyle, facial hair, etc. that alter the visual information in a face. Support for the notion that both processes are important in face recognition comes from neuropsychological research (Farah, Wilson, Drain & Tanaka, 1995; Moscovitch et al., 1997; Postma, Izendoorn, De Haan, & Edward, 1998) showing a double dissociation for featural and configural deficits. Additional support comes from studies in which face stimuli are manipulated in a way that disrupts configural processes but leaves local feature information intact. Such manipulations include inverting face stimuli (Tong & Nakayama, 1999; Valentine & Bruce, 1988; Yin, 1969), scrambling (Tanaka & Farah, 1993), “exploding” (Farah, Tanaka & Drain, 1995), and misaligning parts of the image in two dimensional (Moscovitch et al., 1997) or three dimensional space (Nakayama, Shimojo and Silverman, 1989). In general such manipulations disrupt face recognition performance but leave recognition of facial features intact (see Chapter 1). On the other hand, simple alterations to one or two facial features, leaving the configural image generally intact, can also impair face identification (Sinha & Poggio, 1996).

How might familiarity with a face affect either or both processes? Common sense suggests if visual learning under natural circumstances were to aid face processing it might preferentially speed configural, as opposed to featural processing, because configural, not featural, aspects of a face allow it to be distinguished from others over a wide range of viewing conditions. Support for this possibility can be found in a recent study by Tong and Nakayama (1999). Using a visual search procedure, they showed that search was more efficient for one’s own face, compared to that for a familiar stranger, among a heterogeneous field of other familiar strangers. Importantly they showed that although face inversions and view point variations slowed search in both cases, search for one’s own face was never rendered as slow as search for a familiar stranger’s face. They interpreted the

results to indicate that over-learning of one's own face produces a viewpoint invariant "robust representation" that speeds recognition. Their results favour the possibility that configural processes specifically become more efficient through over-learning.

The difference in processing of internal versus external facial details is particularly interesting because differences are apparent for familiar and unfamiliar faces. A number of studies (Ellis, Shepherd, & Davies, 1979; Haig, 1986; Nachson, Moscovitch, & Umilita, 1995; Ross & Turkewitz, 1982; Young, Hay, McWeeney, Flude, & Ellis, 1985) have observed that unfamiliar faces (recently introduced) are primarily recognised by external features (i.e., hair), while familiar faces use internal details as much or even more than external features. This implies that any potential difference in the processing of famous compared to non-famous faces may evolve from differences in configural processing. Famous faces would benefit the most from the global information contained in a face leading to a greater processing efficiency.

It has been suggested that the left hemisphere is specialised in analytic processing, while the right hemisphere is specialised in holistic processing (Levy-Agresti & Sperry, 1968). This is consistent with face perception research that has reported a right hemisphere (left visual field) advantage for face judgement tasks (Anderson and Parkin, 1985; Rhodes & Wooding, 1989; Young, 1985). Furthermore, just as inverted face stimuli can disrupt configural processing (e.g., Tanaka and Farah, 1993) the right hemisphere advantage can be eliminated (or greatly decreased) by inverting face stimuli (Leechey, Carey, Diamond, & Cahn, 1978). Hence, there seems to be a strong connection between the configural processing of faces and the right hemisphere.

Other studies have used imaging and electrophysiological techniques to reveal familiarity effects for faces. Begleiter, Porjesz, & Wang, (1993) established an event-related potential (ERP) correlate of a visual memory process. This was demonstrated by a reduction in the amplitude of the visual

memory potential (VMP) to repeated pictures of unfamiliar faces compared to that obtained with novel faces. This led them to explore the differences between familiar and unfamiliar faces (Begleiter, Pirjesz, & Wang, 1995). Using a repetition-priming paradigm the authors discovered that in terms of the behavioural data there was no significant difference between primed and unprimed unfamiliar faces, but there was a significant difference for primed and unprimed familiar faces. The ERP recordings showed that the VMP was reduced for both the primed unfamiliar and familiar faces compared to the difference in the unprimed stimuli. However, the reduction was much greater for familiar primed faces, as a significant difference was found between primed familiar and unfamiliar faces.

Using positron emission tomography (PET) to study brain activation sites while viewing familiar vs. unfamiliar faces, Dubois, Rossion, Schiltz, Bodart, Michel, Bruyer, & Crommelinck, (1999) reported bilateral activation of the fusiform gyri (including the right fusiform face area) for both types of faces. However, for familiar faces, a selective decrement in activity in the early visual cortical areas was reported. For novel faces, selective activation of the left amygdala was observed. These different patterns of brain activation support the idea that as familiarity for face stimuli develops, qualitative rather than simple quantitative changes in processing occur, i.e., different processing strategies recruit different brain areas.

In the present study I sought to explore perceptual processing advantages bestowed by high familiarity, or “fame”, in a face. To bypass the semantic and verbal memory demands of explicit recognition tasks, I used an implicit measure of image processing that involved detecting a change in successive face presentations. Two faces, on either side of a fixation point, were presented briefly and then masked. A similar pair of faces were again presented and then masked. One face in each frame stayed the same and one changed to a different person of the same gender. The task was to report the genders of the faces on the left and right and to report the location of the face

that changed. This required attention to be directed to both faces and representations of the first and second pair of images to be compared.

Because I was interested in the effects of familiarity, famous faces were used in all possible combinations with non-famous faces. If visual fame bestows an advantage on perception (as opposed to naming memory) then change detection should be more accurate with famous than non-famous faces. If, on the other hand, visual fame aids memory processes without enhancing visual processes, then simple detection of image changes should be unaffected by fame. Our procedure also allowed us to compare the effects of left and right hemifield presentations. In experiment 2, I assessed the relative contribution of configural vs. featural processes by repeating the procedure using inverted faces (since this manipulation is thought to selectively disrupt configural processes). The role of semantic memory was addressed in Experiment 3 by providing a pre-test study including fictional biographical sketches for all the faces used in the experiment. This allowed us to affirm two issues: (1) that the effect of 'fame' was not an advantage that could be easily induced in non-famous faces, and (2) that the advantage of a famous face was not due to access to semantic retrieval (e.g., at the P.I.N.s), but occurred at the structural encoding stage of face recognition.

Method

Please refer to Chapter 2 for general methods. Below are the method details specific to the face experiments of this chapter.

Participants

Forty-seven British undergraduates (41 females, 6 males), ranging in age from 18-49 years, volunteered to participate in exchange for course credit. Forty-one were right-handed and six were left-handed.

Apparatus and Stimuli

The experiment was run on an 8600/200 Power PC Macintosh computer. Stimuli were displayed on a 13-in. (33-cm) colour (75 Hz) monitor. Face stimuli were rectangular greyscale digital photographs subtending 5.6 deg by 4.1 deg. The centre of each face was positioned 3.2 deg along the horizontal meridian to the left and right of a small central black fixation cross. Two different sets of stimuli were presented in Experiment 1 (Set A and B) but only Set A was used in Experiment 2 and 3. Each set comprised eight faces: four non-famous (two males, two females) and four famous (two males, two females). Within each face set, contrast and clarity of each image and face size was adjusted to a roughly similar level. Set A faces were “natural” portraits that included neck, some shoulders and some variation in background grey level. In Set B the faces were isolated and presented without neck or shoulders on the same uniform grey background. These images had more contrast than those in Set A. All faces were frontal views, with both eyes and all other internal and external facial features clearly visible⁵. All faces, except one famous male and one non-famous male in each set with neutral expressions, were smiling.

Non-famous faces were randomly selected from a North American high school yearbook and from business web sites depicting employees. Famous faces were selected as highly familiar to British undergraduates at the time the study was conducted. Set A contained Diana, Princess of Wales (PD), (who had died 1.5 years prior to the study; Prince Charles (PC) whose image was often present in the news; Leonardo DiCaprio (LDC) who starred in the film “Titanic” that was at cinemas at the time of the study; and Jennifer Aniston (JA), who starred in the TV series “Friends” and in shampoo advertisements (both being aired at the time of the study).

⁵ Including both internal and external features should increase the likelihood of finding a difference in famous and non-famous faces, as both types of face rely on different details (as previously referred to).

Set B contained Tony Blair (TB), the then UK Prime Minister; William Hague (WH) the then UK Leader of the Opposition, (both appearing often in news clips); Carol Vorderman (CV), a UK television personality; and Cher (C), a pop star, both experiencing popularity at the time of the study.

A pattern mask was a scrambled collage of face parts similar in details to the faces used in the study. The same mask was used throughout the experiment.

Procedure

The procedure for this experiment is set out in Chapter 2. See figure 8 (pg 56) for a schematic illustration.

Experiment 1

Only upright faces were presented. Face Set A was used for 10 participants and Set B was used for 16 participants.

Results

Gender task. On average, participants correctly named the gender of the left and right faces on 83% of trials (s.e. = 6.8%). The difference in performance for Set A and Set B was non-significant. All participants scored 70% or better on this task.

An ANOVA on the percentage of correct gender identification responses for the gender task was conducted using three within factors (visual field: left or right; distractor: famous or non famous; and change condition: NN, FF, FN, and NF). No main effects or interaction effects were significant. A point to note about the data for the gender categorisation task is that the pattern of results does not account for the pattern of results seen in the change detection task. This indicates that the face analysis required for change detection was somewhat independent from the gender analysis

required by the task, supporting Bruce and Young's (1986) distinction of two processing routes.

Change detection. Only trials where both gender judgements were correct were analyzed. As the gender of the changing faces was a non-significant factor in all three experiments I collapsed this data, and will not discuss it further. An ANOVA on the percentage of correct location identification responses for the change detection task was conducted using one between (face set) and three within factors (visual field: left or right; distractor: famous or non famous; and change condition: NN, FF, FN, and NF). Since there was a non-significant effect of face set and all interactions involving face set were non-significant, further discussion of the data refers to the combined results of both groups.

Figure 9




Figure 9. Experiment 1: Group mean percent correct change location identification for each type of change condition. The first letter in the condition-label refers to the nature of the face before it changed (famous, f; or non-famous, n) and the second letter refers to its replacement. For, example, nf refers to the conditions wherein a non-famous face changed into a famous face. Vertical bars indicate ± 1 s.e.

The ANOVA indicated a non-significant effect of distractor type and non-significant interaction effects involving this factor.

Figure 9 shows the mean percent correct change location identification for each of the four change conditions. As can be seen in the Figure, visual fame clearly aids change detection. This is supported by a significant main effect for change condition $F(3, 75) = 7.227, p < .01$. The NN trials produced the lowest accuracy (mean = 80%), whereas in the NF condition participants were correct on average on 85% of trials. Bonferroni post-hoc comparisons indicate that the difference between the NN and NF conditions was significant ($p < .01$). Note that the performance levels for NF, FN and FF conditions were non-significantly different from one another.

Figure 10.




Figure 10. Experiment 1: Group mean percent correct change location identification for trials in which the face in the left (black bars) or right (white bars) visual field was replaced by a different face are shown for each of the four change conditions. Vertical bars indicate ± 1 s.e.

The significant main effect of change condition must be interpreted in light of a significant interaction effect of change condition and location of change $F(3, 75) = 4.654, p < .01$. As can be seen in Figure 10, group mean performance on the change detection task showed no effect of fame for faces presented to the right visual field. For faces presented to the left visual field, a clear fame effect is found that accounts for the pattern of results seen in Figure 10. For the NN condition, changes in the left visual field were detected significantly less often than the ones that occurred in the right visual field. Using a famous face in the changing image restored performance for left visual field presentations to a level equivalent to that for right visual field presentations. Bonferroni post-hoc comparisons indicate that the difference between the NN and NF conditions was significant when presented in the left visual field ($p < .01$).

Effect of session. A question that arises from this experiment concerns the effect of repeatedly viewing the same non-famous faces. Although the faces were novel on the first few trials, familiarity must have developed during the course of the session. To analyse for this effect, the trials were divided into quartiles and a three way repeated measures ANOVA (change condition X quartile X visual field) was conducted. The interaction of visual field and change condition was significant (as before) but the triple interaction with quartile was non-significant, meaning that the visual fame effect described above was present throughout the session and did not diminish. However, the analysis revealed a significant interaction between quartile and visual field, $F(3,75) = 3.863, p < .01$. As seen in Figure 11, in the first quartile, mean change detection for left hemifield stimuli (79%) was better than that for right hemifield changes (69%). After the first quarter of trials, these general hemifield effects disappeared. This supports previous studies indicating preferential processing of faces in the right hemisphere (Levy-Agresti & Sperry, 1968), while also suggesting that with practice on the task the left hemisphere can learn to use the visual information as efficiently as that of the

right. However, this explanation does not apply to the effect of a famous face, which maintained a left visual field advantage throughout the experiment. This pattern of effects suggests that the right hemisphere maintains superior access to stored representations through benefits in configural processing.

Figure 11.




Figure 11. Experiment 1: Group mean percent correct change location identification across quartiles in which the face in the left (black bars) or right (white bars) visual field was replaced by a different face are shown for each of the quartiles. Vertical bars indicate ± 1 s.e.

Discussion

When a brief, masked, bilateral display of two faces was followed by a similar display with one of the two faces being replaced by another, observers were more likely to detect the change in the display if a famous face was involved in the alteration. This effect was only found for changing stimuli presented to the left visual field and was not present for changes to faces presented to the right visual field. When left visual field changes only

involved non-famous faces, change detection performance was correct on merely 77% of trials. This is a surprisingly poor level of performance considering that substituting one face for another is a substantial alteration made to an object of central interest in the scene. When a non-famous face was substituted with a famous face in the left visual field, correct detection of the change location rose to 86%, a gain of nine percentage points. Hence, the 'visual fame effect' was limited to locating face changes in the left visual field.

An interesting feature of the quartile data indicated that while the effect of a famous face remained constant throughout the course of the experiment, there was an overall change in the pattern for the left and right hemifields. Initially, the change detection accuracy in the left visual field was superior to that of the right visual field (79% and 69% respectively). However, after the first quartile of trials, overall change detection accuracy in the right visual field increased to that of the left visual field. It appears that the faces benefited from the configural processing by the right hemisphere (left visual field) early on in the experiment. However, in the case of the left hemisphere (right visual field), participants required repeated experience with the task and stimuli in order for the featural and/or configural information to become beneficial.

Experiment 2

The advantages for detecting changes bestowed by famous faces observed for changes in the left visual field in Experiment 1, may have been due to superior configural processing of famous faces mediated by the right hemisphere. To test this possibility, Experiment 2 used the same procedure as before, this time with the face images inverted. Face inversion is thought to disrupt configural processing (e.g., Leder & Bruce, 2000). If configural processing accounted for the pattern of results in Experiment 1, then the advantage of famous faces would be diminished or eliminated in this task.

The stimuli of Face Set A (used in Experiment 1) were used in this experiment. Here, I inverted all the faces and masks. All other methods were the same as in Experiment 1. Eight females and three males ranging in age from 18 to 34 years participated.

Results

Gender task. On average, participants were correct on this task on 71% of trials (s.e. = 3%). An ANOVA on the percentage of correct gender identification responses for the gender task was conducted using three within factors (visual field: left or right; distractor: famous or non-famous; and change condition: NN, FF, FN, and NF). No main effect and no interaction effects were significant.

Figure 12.




Figure 12. Experiment 2: Group mean percent correct change location identification for trials in which the face in the left (black bars) or right (white bars) visual field was replaced by a different face are shown for each of the four change conditions. Vertical bars indicate ± 1 s.e.

Change detection. Only trials where both gender judgements were correct were included in the analysis. An ANOVA on the percentage of correct location identification responses for the change detection task was conducted using three within factors (visual field: left or right; distractor: famous or non-famous; and change condition: NN, FF, FN, and NF). No main effect and no interaction effects were significant. Although, there was a difference in overall performance in the left and right visual field (11%) this did not reach statistical significance, $F(1,10) = 1.37, p > .26$. Group mean performance for each change condition presented to each visual hemifield is shown in Figure 12. This null result was compared with the data of Experiment 1 (where both used Face set A). A mixed factorial ANOVA was calculated with one between-subject factor (face set: upright versus inverted) and two within-subject factors (side: left versus right; change: NN, NF, FN, and FF). This revealed a significant main effect of face set, with inverted face changes detected less than upright face changes. No other main effects or interactions were significant. Therefore, as predicted, the presentation of inverted faces in this experiment appeared to reduce the processing benefits conferred by famous faces seen in Experiment 1. However, there was no significant interaction between face set and fame. At present, it seems this non-significant interaction might simply reflect a lack of experimental power; future experiments, perhaps using within-subject designs, could address this issue.

Effect of Session

By dividing the trials into quartiles an analysis of the effect of repeatedly viewing the inverted face stimuli was made. A three way repeated measures ANOVA (change condition X quartile X visual field) was conducted. There were no significant main effects or interaction effects. Figure 13 displays the mean performance for each quartile presented to each hemifield. With the inverted face stimuli, performance in the right visual

field, only improved by 4% from first to last quartile. This suggests that the right visual field was unable to benefit from stimuli repetitions, as had been the case with upright presentations. This also accounts for the overall lower change location accuracy in the right visual field.

Figure 13




Figure 13. Experiment 2: Group mean percent correct change location identification across quartiles in which the face in the left (black bars) or right (white bars) visual field was replaced by a different face are shown for each of the quartiles. Vertical bars indicate ± 1 s.e.

Discussion

The inversion of faces resulted in a null effect of visual fame that supports the configural processing account offered for the visual fame effect of Experiment 1. The visual fame effect of Experiment 1 only occurred with changes presented to the right hemisphere (left visual field), a hemisphere known to preferentially process global detail. The inversion of the face stimuli in Experiment 2 appears to have disrupted the ability to process configural information as efficiently.

It is possible that fame did have an effect on visual processing, but this particular manipulation was insensitive. Hence, the visual fame effect was significantly reduced even if it was not entirely eliminated. This is consistent with the population coding hypothesis (Perret, Oram, & Ashbridge, 1998), which predicts that unusual (rarely encountered) viewpoints such as inverted faces would have fewer cells tuned to their orientation and therefore a slower accumulation rate than more common viewpoints. This logic can be extended as famous faces are more common than non-famous faces supporting the findings of Experiment 1. However, famous faces are invariably encountered in their upright orientation. When presented upside-down they do not have the same advantage of a large population of selectively tuned cells, and are therefore more on a par with the non-famous faces. This would clearly reduce the famous advantage as observed in Experiment 2.

Note, also that if low-level perceptual differences between items had led to the pattern of data in Experiment 1, the same pattern should have been evident in Experiment 2. This was not the case. Therefore, factors such as brightness and contrast could not have significantly contributed to the fame effect of Experiment 1.

Experiment 3

From examination of the quartile data for face set A in Experiment 1, it is clear that the effect of fame did not diminish over the course of the experiment. This indicates that the brief multiple presentations of the non-famous faces did not create familiarity equal to that of the famous faces. Here, I asked whether studying the faces before the experiment, would lead to similar accuracy for both famous and non-famous faces. This issue was addressed in Experiment 3 along with the possibility that semantic information (e.g. such as a label) may be required to reduce the advantage previously observed for famous faces. This would imply that access to the

PINs was required and that identity retrieval at the later stages of encoding was necessary to produce efficient change detection.

Eleven female British undergraduates participated in this experiment. The ages ranged from 18 to 30 years of age, with ten participants being right-handed and one left-handed. The experiment was identical to that of Experiment 1, except that only face set A was used, and participants went through a study period prior to the experimental test session.

In the study session participants were presented with eight A4 sheets of paper, each sheet contained the face, name, and biography for one of the faces. The biographical details were presented so that an accurate statement about the famous faces was made, along with a fictitious description of the non-famous faces. The non-famous descriptions were made to be as interesting as the famous faces. The statements were on average 80 words in total across a mean of five sentences per face. Participants were instructed to study the sheets at their own pace, and to stop when they believed they were able to identify each face (this process averaged eight minutes, with no participant taking more than ten minutes). The researcher then presented each face separately and asked for the name and another piece of information about the person. All participants were 100% correct and so proceeded to the experimental test session.

Results

Gender task. As one participant scored only 50% on this task their data was excluded from further analysis. On average, participants were correct on this task on 89% of trials (s.e. = 1.1%). An ANOVA on the percentage of correct gender identification responses for the gender task was conducted using three within factors (visual field: left or right; distractor: famous or non-famous; and change condition: NN, FF, FN, and NF). The main effects and interaction effects were non-significant. A between group ANOVA was

calculated for this data and the data from experiment 1 (Face Set A), but the effect of experiment group did not reach statistical significance.

Change detection. An ANOVA on the percentage of correct location identification responses for the change detection task was conducted using three within factors (visual field: left or right; distractor: famous or non-famous; and change condition: NN, FF, FN, and NF). Only trials where both gender judgements were correct were included in the analysis.

Figure 14.




Figure 14. Experiment 3: Group mean percent correct change location identification for each type of change condition. Vertical bars indicate ± 1 s.e.

Figure 14 shows the mean percent correct change location identification for each of the four change conditions. As can be seen in the Figure, fame clearly aids change detection in the same way as Experiment 1. This is supported by a significant main effect for change condition $F(3, 27) = 3.38, p < .05$. The NN trials produced the lowest accuracy (mean = 74%), whereas in the NF condition participants were correct on average on 80% of

trials. Bonferroni post-hoc comparisons indicate that the difference between the NN and NF conditions was significant ($p < .05$). Note that the performance level for NF, FN and FF conditions were non-significantly different from one another, a result that is consistent with the findings of Experiment 1.

Figure 15.




Figure 15. Experiment 3: Group mean percent correct change location identification for trials in which the face in the left (black bars) or right (white bars) visual field was replaced by a different face are shown for each of the four change conditions. Vertical bars indicate ± 1 s.e.

Visual Field Effects. The significant main effect of change condition must be interpreted in light of a significant interaction effect of change condition and location of change $F(3, 7) = 3.73, p < .05$. As can be seen in Figure 15, group mean performance on the change detection task showed no effect of fame for faces presented to the right visual field. For faces presented to the left visual field, a clear effect of fame was found that accounts for the pattern of results seen in Figure 14. For the NN condition, changes in the left visual field were detected significantly less often than those occurring in the right visual field. Using a famous face in the changing image restored performance for left

visual field presentations to a level equivalent to that for right visual field presentations. These results are entirely consistent with those of Experiment 1.

Figure 16

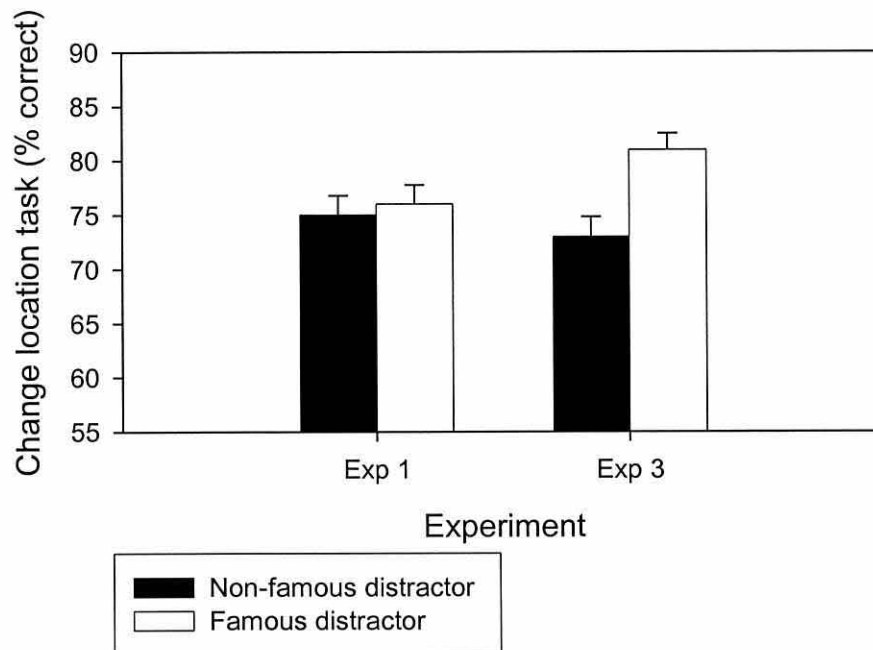


Figure 16. Experiment 1& 3: Group mean percent correct change location identification for type of distractor (non-changing image either famous or non-famous). Vertical bars indicate ± 1 s.e.

Distractor Effects. Unlike Experiment 1 there was a significant effect for distractor type, $F(3, 27) = 15.5, p < .01$. Inspection of the means reveals that the presence of a famous distractor aided detection of the change. Task performance when a famous distractor was present produced a mean of 81%, while a non-famous distractor reduced the mean to 73%. This new effect was compared with the data from face set A of the first experiment. A mixed factor ANOVA was conducted using three within factors (visual field: left or right; distractor: famous or non-famous; and change condition: NN, FF, FN, and NF), and one between factor (group: Experiment 1 or Experiment 3). This revealed that while the pattern of results for type of change and the

interaction of change and side were not different in the two experiments, there was a significant difference between groups when type of distractor was considered, $F(1, 18) = 5.8, p < .05$. Inspection of the means in Figure 16 shows that task performance remained the same between groups when a non-famous distractor was present, but a famous distractor increased mean accuracy in Experiment 3 (when there had been pre-exposure to the face). Note, that in this case the effect occurred regardless of visual field.

Discussion

Experiment 3 replicated the visual fame effect of Experiment 1, where an advantage for detecting changes involving famous faces in the left visual field was found. The training session did not appear to reduce the difference in change detection between famous and non-famous changes. This indicates that the speed of processing advantage for famous faces was acquired through long-term repeat exposure to a face (robust representations), and not merely through successful access to semantic information at the level of the PINS.

However, Experiment 3 revealed a significant effect of distractor, which was not evident in Experiment 1. The presence of a famous distractor led to significantly improved change detection performance when compared to non-famous distractors. The difference between experiments will be discussed in the following section.

General Discussion

Famous vs. non-famous faces were assessed using an indirect measure of face perception. The findings of Experiment 1 indicate that visual fame bestows an advantage to perceptual processing. When a brief, masked, bilateral display of two faces was followed by a similar display with one of the two faces being replaced by another, observers were more likely to detect

the change in the display if a famous face was involved in the alteration. This effect was only found for changing stimuli presented to the left visual field and was not present for changes to faces presented to the right visual field. Experiment 1 also showed that the advantage of a famous face did not diminish as the experiment progressed. This was evident in the quartile data, for although there was a general improvement in accuracy from quartile 1 through to quartile 4, the fame effect remained constant through out. As well as reducing overall performance, the inverted faces of Experiment 2 disrupted the effect of fame witnessed in Experiment 1, as no significant effects of famous face presentations were found. This suggests that a configural processing advantage for famous faces presented to the left visual field may account for the fame effect of Experiment 1, as inverting faces disrupts configural processing. Experiment 3 further established that 'fame' as identified by the fame effect in Experiment 1 could not be simply established through a study phase, and that the availability of semantic information about the famous faces (through activation of the PINs) did not account for the famous face advantage of Experiment 1. Even when participants studied the non-famous faces and learnt biographical details about each face, the effect of Experiment 1 was replicated.

Fame Effects

I now consider how the observed fame effect may be explained. There are a least two means by which fame in a face could enhance change detection. A first mechanism by which fame could aid change detection could be that famous, but not non-famous, faces are processed faster or more efficiently. In this speeded processing view, the images presented in each of the two displays are processed as quickly as possible, as they appear in the sequence. The sensory information in a famous face is matched more quickly or more efficiently to a stored representation (at the level of the FRUs) than that for a non-famous face, leading to a greater likelihood of establishing a

consolidated representation available to awareness before each post-stimulus mask is presented. Comparisons of the mental representations for the first and second images are made after the stimulus sequence is complete and change detection is reported. Therefore, performance depends on access to each representation.

The second possibility is that the appearance of a famous face in the first presentation of the display could attract or capture attention (Theeuwes, Kramer, Hahn, & Irwin, 1998; Yantis, 1993). Tong & Nakayama (1999) have shown that highly familiar faces are more easily detected in visual search tasks where distractors are familiar stranger's faces. Attentional capture by highly familiar words has been repeatedly demonstrated by the cocktail party effect (selective hearing of one's own name in a multi-voice conversation; Cherry, 1953). More recently, Shapiro, Caldwell & Sorensen (1997) demonstrated in a divided attention task (using the attentional blink paradigm, Raymond, Shapiro & Arnell, 1992) that selecting one's own name in a rapid successive series of words and names was better than selecting other names or other words. In the current task, attentional capture by a famous face might have facilitated consolidation of the face's representation before the disrupting visual mask was presented, thus permitting change detection. Engagement of attention at the location of a famous face may have also reduced the ability to process information in the competing face.

These two views would predict different patterns of data for the change location tasks reported in these three experiments. An attentional capture account predicts a directional effect in performance, i.e., fame should have had more impact when present in the first presentation (F-N trials) than when presented in the second presentation (N-F trials). Such directional effects were not observed. Although attention capture by famous faces may occur with longer stimulus presentations, the current data do not support the idea that fame had any special effect on attentional orienting *per se*. However, the data are not necessarily inconsistent with an effect of an attentional blink.

The term attentional blink (AB) refers to an inability in reporting the second of two targets presented within 200-400msec of each other in a RSVP (rapid serial visual presentation) stream (Raymond, Shapiro, & Arnell, 1992). It is suggested that this phenomena occurs as a result of limited attentional resources that do not permit the second item to be processed until after the first item has been processed (e.g., 400msec +). The presentation times of the change location task are entirely consistent with the crucial timing of the AB (i.e., 200msec SOA). Interestingly, the change detection data reported in experiment 1 and 3 revealed significant differences for the second presentation of a famous face (N-F trials) when compared to the non-famous baseline (N-N trials). However, there was a non-significant difference between the baseline and when a famous face was followed by a non-famous face (F-N trials). This suggests that when a salient known face was presented as the second item in a change it was more likely to survive the attentional blink than a presentation of a non-famous face in the second image. This would be consistent with the finding that a person's own name survives the attentional blink when detecting words in a RSVP stream when compared to other names or nouns (Shapiro, Caldwell, & Sorensen, 1997).

Consideration of the speed of processing account predicts a different outcome to that of the attentional capture account. It correctly predicts that there should be no statistical difference in N-F vs. F-N trials, and that performance on these trials should be similar to F-F trials. However, it also predicts that performance should be better when famous faces are used as the non-changing image than when non-famous faces are used. This result was not observed in Experiment 1, but was present in Experiment 3. This suggests that without revealing the famous faces (with name and biography) in the pre-test condition of experiment 3, the task was insensitive to this factor. It might be expected that as both famous and non-famous faces were exposed in the training session, there should be an equal priming effect for both types of face. Priming refers to the finding that exposure to a stimulus, e.g., a picture

of a car, will facilitate subsequent responses to that item, or a related item (e.g., Durso, & Johnson, 1979; Palmer, 1975). This would lead to no discernable differences between Experiment 1 and 3.

However, a number of studies previously cited would indicate that these results are consistent with a priming explanation. Consider Bruce and Young's (1986) model with a priming explanation, the activation level of a face recognition unit can be raised by the actual structural description of the viewed face, and also by the person identity node when a face is expected or has been recently seen. This suggests that pre-test exposure to the famous faces may have primed responses during the experiment.⁶ It is possible that the pre-test exposure primed the famous faces in the distractor condition to a greater extent than the non-famous faces, producing statistical power large enough to reveal a significant effect of distractor (Experiment 1 distractor power = 51%, Experiment 3 distractor power = 95%).

Visual Field Effect

The fame effect reported in Experiment 1 and 3 was localized to the left visual field. This finding indicates that the right cerebral hemisphere was accessing information that differentiates a famous from a non-famous face, even though this information was not required in the task. When changes occurred in the right visual field, the differences between famous and non-famous faces were non-significant, indicating that the left hemisphere detected changes using a somewhat different mechanism for which fame was irrelevant. Why might this occur?

One possibility is that the spatial frequency information (i.e., level of detail) in the display was analysed differently by each hemisphere. Previous studies have shown that in speeded tasks, the right hemisphere (accessing information presented to the left hemifield) preferentially processes low

⁶ In Experiment 1 participant's verbal recall of who had been presented in the experiment averaged 75% accuracy (3 out of 4 famous faces). However, in Experiment 3 all participants knew what faces to expect, due to the pre-test.

spatial frequency (coarse) information, whereas the left hemisphere (accessing information presented to the right hemifield) preferentially processes high spatial frequency (detailed) visual information (Ivry & Robertson, 1998; Kitterle, Christman & Hellige, 1990). Research has generally shown that, when bilateral stimuli are presented the right hemisphere processes information in a more global manner than the left hemisphere (see Ivry & Robertson, 1998 for a review). This suggests that part-based face processing might be more heavily favoured by the left hemifield whereas the right hemisphere might favour configural processing. Findings from face perception research are generally consistent with this lateralization of processing account (Gilbert, & Bakan, 1973; Levine, Banich, & Koch-Weser, 1988; Luh, 1998; Sergent, 1985). In the current study, the observation that the fame effect was only found with left hemifield presentations suggest that it is the configural face perception processes that cause the effect.

This configural processing interpretation was tested in Experiment 2 where performance on an inverted face task diminished the effect of fame. The inversion of faces is known to disrupt configural processing (Leder & Bruce, 2000). Hence, the advantage from having a famous face is likely due to superior configural processing by the right hemisphere at the stage of encoding. This is consistent with the proposal of Freire, Lee, and Symons (2000) that the perceptual processing of inverted faces creates an 'encoding bottleneck' that limits the input of configural information into memory. This account stemmed from their study of the effect of facial configuration (manipulating configural and featural information) on encoding and memory tasks. The encoding tasks produced a decrease in accuracy for inverted faces compared with upright faces, but a memory test (10 second delay since stimulus presentation) produced no further decrease in performance. Thus, suggesting that once configural or featural information is encoded, the retention of the information is similar.

One concern that needed addressing was whether the 'fame effect' would also occur for recently learned faces compared to novel faces. If true, this would imply that it was not fame (extensive learning) that produced the said effect, but mere priming of the previously seen faces. Our results reject this possibility; first the size of the effect in Experiment 1 (as measured by the means) did not decrease from the first quartile of trials through to the fourth quartile of trials. Furthermore, Experiment 3 provided a training session prior to the experimental test session. Even after learning the name of each non-famous face and associating biographical information, the advantage of a famous face remained consistent in size to that of Experiment 1. Therefore, the reported effects can be attributed to the famous (highly learned) status of the faces. This is consistent with the idea that perceptually encoding information of the famous faces allows contact with their particularly robust representations in memory (Tong & Nakayama, 1999).

The experiments presented in this chapter demonstrate that implicit measures of fame, such as the change location task, can reveal significant differences in the speed of processing of famous versus non-famous stimuli. The method adopted here has the potential to become a useful tool for evaluating fame that may be beneficial to market research companies who need to evaluate people's perceptions of fame, without always relying on "accurate" verbal reports.

Chapter 4 – Fame Effects in Landmarks

Introduction

A fame effect for faces was described in the change detection task of Chapter 3. The ability to detect a changing face from one of two faces was enhanced when the face involved in the change was famous (highly learnt) as opposed to new or recently learned. This effect occurred regardless of whether participants had any semantic knowledge (e.g., a name) for the non-famous faces. However, the effect was specific to changing presentations in the left visual field and could be eliminated by presenting inverted faces. This pattern of data was accounted for by an advantage for face processing in the right hemisphere that provided a configural (rather than part-based) representation that was disrupted by inversion. However, what role does 'fame' play in the visual processing of other classes of object? Will configural advantages be evident for non-face stimuli, or will such advantages be a face-specific effect?

There is a considerable amount of research suggesting that faces more than any other object type, rely on configural processing (e.g., Farah, Wilson, Drain and Tanaka, 1998). Furthermore, a face specific brain area, termed the fusiform face area, is greatly activated by the presentation of face stimuli rather than other object types. However, a debate continues over whether faces are truly a special case of object recognition, or whether these findings result from the level of expertise that we have with faces. In fact, expert recognition for a selection of objects including birds, cars, and dogs has been compared with that of face recognition (Diamond & Carey, 1986; Gauthier, Skudlarski, Gore, & Anderson, 2000; Tanaka & Curran, 2001). These studies have shown that when people have extensive knowledge and experience with a particular category of object, brain areas are activated that are similar to those activated by faces in the fusiform area (Gauthier, et al., 2000). Also, an

enhanced early negative component (N170, 164ms), associated with face perception, has been reported for objects in the field of expertise compared to objects that are not within the field of expertise (Tanaka & Curran, 2001). This suggests that high levels of familiarity with processing and encoding non-face images may share similarities with that of face stimuli. This raises the possibility that famous and non-famous differences in processing may be evident for other types of stimuli.

In attempting to study the effect of fame in non-face objects, there are two problematic issues in adopting the types of stimuli that are used to assess the differences between experts and non-experts (e.g., dogs or birds). First, the need to recruit participants who are expert in a particular field leads to a population that is less random, tends to be older than the usual pool of student participants, and does not represent how the majority of the population processes these objects. Alternatively, naïve participants could be trained on a category of object, but this is both time intensive and artificial. Rather, for the current purpose, it is more desirable to select a category of non-face objects that have established famous exemplars that are common to the majority.

The second problem with using these categories to compare object and face recognition is the issue of individuation. When we walk around our office building and encounter our colleagues we don't categorise them as 'human' or 'male/female', rather we instantly access the name of the person and other relevant information. For every face we encounter there is only one version (with the exception of course of identical twins), e.g., there is only one Nelson Mandela. However, whether a person is an expert or non-expert with birds, they are unlikely to encode them at such an individual level. A birdwatcher may be able to instantly differentiate between a swallow and a swift, but I doubt many can say 'There's Sam swallow, and over there is Suzy swift'! This means that where each face in a group of faces will be perceived at an individual level (each are distinct types), a group of gulls will be

perceived as multiple tokens of one type. Therefore, in order to generalise effects of fame to the normal population it is important to compare faces with a class of object that is distinctively individual, and where famous exemplars can be easily established.

For the reasons outlined above, the category of 'landmark' was chosen as a face comparison for the fame effect. Landmarks such as the 'Eiffel Tower' and the 'Statue of Liberty' are excellent examples of objects that would be considered famous by most Western people. Moreover, they are also unique in a similar way to faces; there is only one Millennium Dome (thankfully!) just as there is only one you. However, there is a further advantage in using landmarks as famous/non-famous objects; the literature on topographical representation (i.e., spatial knowledge) suggests strong similarities between face and place processing.

It has been proposed that the inferior temporal cortex is functionally involved in specialised face processing (Desimone, 1991; Farah, 1990; Kanwisher, McDermott, & Chun, 1997). In particular an area labelled the 'fusiform face area' has been implicated. However, a condition known as landmark agnosia often occurs with prosopagnosia, which suggests the fusiform area may also be involved in specialised landmark processing (Aguirre, Zarahn, & D'Esposito, 1998). In prosopagnosia, patients are unable to recognise individual faces, even though facial expression analysis remains intact. Similarly, landmark agnosia is primarily a deficit in the ability to use large environmental features for orientation (Hecaen, Tzortzis, & Rondot, 1980; Landis, Cummings, Benson, & Palmer, 1986; McCarthy, Evans, & Hodges, 1996; Pallis, 1955; Takahashi, Kawamura, Hirayama, & Tagawa, 1989; Whiteley & Warrington, 1978). Patients are able to describe routes and draw maps (i.e., intact spatial representation), but cannot recognise familiar landmarks and in particular, buildings.

Furthermore, there appears to be a dissociation between spatial relationship information and memory for landmarks. A condition known as

Egocentric Disorientation prevents patients from representing spatial relationships between objects, although they are able to identify landmarks. Aguirre, and D'Esposito (1997) proposed a divisional model whereby the ventral areas (which are specialised for representing environmental features: 'what') are more involved in the appearance of landmarks, than the location/position of landmarks.

Figure 17a




Figure 17b

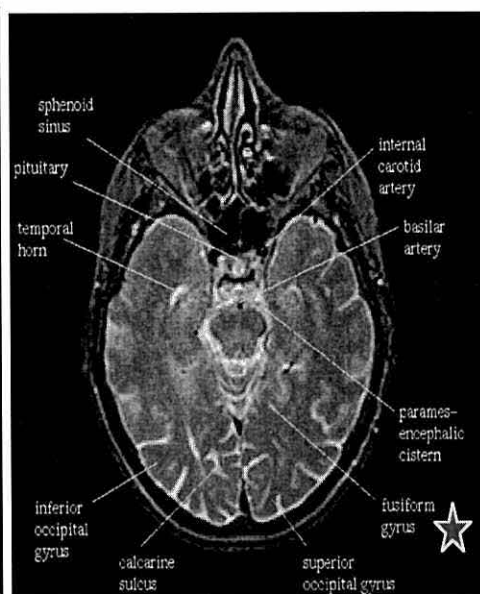


Figure 17a shows the parahippocampal gyri, which are implicated in spatial representation processes. Also shown is the lingual sulcus along with the fusiform gyri (Figure 17b) that are activated during landmark recognition. The images are taken from 'The Whole Brain Atlas' CD-ROM, Johnson and Becker.

Using a virtual reality town, Aguirre et al (1997) tested, in normal participants, the prediction that ventral rather than dorsal areas would be recruited for appearance versus position judgements. A consistent ventral/dorsal dissociation was found across participants. When appearance judgements (e.g., "how many windows"), were required there was greater activation of the lingual and inferior fusiform gyri. The position judgements revealed activation of an extensive area of posterior cortex, and the

parahippocampal gyrus. Figures 17a and 17b show the location of the relevant activations for the two tasks.

Because landmark appearance judgements recruited fusiform areas similar to those areas recruited by faces perception tasks, Aguirre, Zarahn, & D'Esposito (1998) used fMRI to measure the neural correlates of perception of exemplars from different stimulus categories, including buildings. They were able to identify voxels within ventral occipitotemporal cortex that responded significantly greater to buildings. These voxels were located medial and superior to the 'face' voxels. This suggests that there may be an area specialised for the perception and representation of buildings, that closely neighbours the area selective to face stimuli.

While the research reported by Aguirre and colleagues has focused on the visual processing of landmarks as individual buildings, other researchers have concentrated on place stimuli that involves general scene layouts rather than specific buildings. These studies revealed an area in the parahippocampal cortex named the Parahippocampal Place Area (PPA) that is associated with visual processing of places (Epstein, Harris, Stanley, & Kanwisher, 1999; Epstein & Kanwisher, 1998). This area has been shown in fMRI to respond selectively and automatically to passively viewed scenes. Moreover, responses in this brain area to objects and faces are weak, with the exception of landmark stimuli (individual buildings) that produce significantly greater responses than these other objects (though not as strong as scenes). Furthermore, the response to landmarks is significantly higher when they are familiar. Epstein et al. (1999) propose that this familiarity effect is because buildings help define the space around us. Hence, when a landmark is familiar it forms a partial scene, from which the whole scene can be imagined. This would lead to greater activation of the PPA, as the areas involved in both viewing and imagining are closely related (O'Craven & Kanwisher, 2000).

The evidence of specialised landmark processing in the brain suggests, that like face stimuli, there is a case for predicting that places should lead to an advantage in processing for famous over non-famous stimuli. There is even tentative evidence to suggest that, like face stimuli, landmarks may use configural rather than featural information when being identified. Hácaen, Tzortis, Rondot, and Loss (1980) reported that Landmark Agnosics relied on featural details to accurately perform a cathedral-matching task (e.g., a door or window, but not the whole building). This implies that under normal circumstances we use the whole of the landmark to identify it. If this is the case then landmark stimuli could yield an advantage for famous items that resembles the pattern of data found for faces in Chapter 3.

In the present study I sought to compare the perceptual processing advantages for famous faces evidenced in Chapter 3, with an effect of high familiarity for landmarks. This involved adopting the same change detection task as described in Chapter 2, a task that bypassed the semantic and verbal memory demands of explicit recognition tasks. The implicit measure involved detecting a change made to one of two images in successive presentations. Two landmarks, on either side of a fixation point, were presented briefly and then masked. A similar pair of landmarks were again presented, and then masked. One landmark in each frame stayed the same and one changed to a different landmark. The task was to report the location of the change and whether the landmarks on both the left and right were buildings or monuments. This required attention to be directed to both landmarks and representations of the first and second pair of images to be compared, as with the face experiments of Chapter 3. Because I was interested in the effects of familiarity, famous landmarks were used in all possible combinations with non-famous landmarks. If visual fame bestows an advantage on perceptual processing in a similar way to face stimuli, then a famous landmark should lead to an advantage when presented in the changing image. Furthermore, if landmark processing relies on configural rather than featural information

there should be a greater advantage for famous landmarks presented in the left visual field, as the right hemisphere is more specialised for global analysis (as opposed to local detail).

Experiment 4

Methods

Participants

Twenty-seven British undergraduates (17 female, 10 male), ranging in age from 18-44 years, volunteered to participate in exchange for course credit. Twenty-two were right-handed and five were left-handed.

Apparatus and Stimuli

The experiment was run on a 9600/200 Power PC Macintosh computer. Stimuli were displayed on a 13-in. (33-cm) colour (75Hz) monitor. Landmark stimuli were rectangular greyscale digital photographs subtending 5.6 deg by 4.1 deg. The centre of each landmark was positioned 3.2 deg along the horizontal meridian to the left and right of a small central black fixation cross. Two different sets of stimuli were presented (Set A and B) to two different groups of participant. Each set comprised eight landmarks: four non-famous (two buildings, two monuments) and four famous (two buildings and two monuments). Within each landmark set, contrast and clarity of each image was adjusted to a roughly similar level. Set A landmarks were isolated and presented on a uniform grey background, while Set B landmarks were central to a scene background (e.g., containing trees etc). This allowed for a comparison to be made between scenes and objects. All landmarks were central to the scene, with non-famous landmarks being equally as elaborate and detailed as the famous landmarks.

Non-famous landmarks were selected from a range of places on 'Art Explosion' CD-ROMs. Famous landmarks were selected as highly familiar to British undergraduates at the time the study was conducted. A pilot test was conducted to select only those famous landmarks that could be identified by 80% of participants (10 participants in total). This developed into the two sets of landmarks, Set A contained 'The Statue of Liberty', 'The Eiffel Tower', 'The Sydney Opera House and 'The Millennium Dome'. Set B contained 'The White House', 'The Houses of parliament (including Big Ben) ', 'The Empire State Building', and 'The Leaning Tower of Pisa'.

A pattern mask was a scrambled collage of landmark parts similar in detail to the landmarks used in the study. The same mask was used throughout the experiment.

Figure 18




Figure 18. An illustration of a trial. Following a 1000ms fixation cross a pair of landmarks are briefly (100ms) presented followed by a scrambled image mask (100ms). A second pair of landmarks are presented (100ms) with mask (100ms). The second landmarks image includes one landmark as presented in the first image, and one, which is a completely different landmark.

Procedure

The procedure was identical to the details in the General Method (Chapter 2) except for the choice of images and the categorisation task (see Figure 18 for schematic). In this experiment participants had to report the type of landmark viewed in both images (building or monument) by pressing one of two keys for the left image and one from a different pair of keys for the right image. As this categorisation task was somewhat unnatural (as opposed to judging the gender of faces) participants were instructed to class vertical structures (e.g., The Eiffel Tower) as monuments and horizontal structures (e.g., The White House) as buildings.

Results

Categorisation task.

On average, participants were correct on this task on 75% of trials (s.e. =3.5%). An ANOVA analysis of the categorisation data revealed that in judging whether a landmark was a building or a monument a famous item significantly increased accuracy in both the changing condition ($F(3, 26) = 7.224, p < .01$) and in the distractor condition ($F(1, 26) = 12.22, p < .01$). Therefore, it appears that access to pre-existing knowledge about a landmark enabled participants to decide which category to select. This occurred even though participants were guided to base their responses on the buildings structure (horizontal/vertical).

Change detection.

Only trials in which categorisation judgements (i.e., building or landmark) were correct for both left and right images were included in the analysis of change location data. As the category of the changing landmarks was a non-significant factor I collapsed the data, and do not discuss it further. An ANOVA on the percentage of correct location identification responses for the change detection task was conducted using one between (landmark set: a

or b) and three within factors (visual field: left or right; distractor: famous or non-famous; and change condition: NN, FF, FN, and NF). It is worth noting that these two stimuli sets used isolated landmarks (set a) and landmarks embedded in scenes (set b). However, there were no significant differences between landmark sets and all interactions involving landmark set were non-significant. Further discussion of the data refers to the combined results of both groups.

Figure 19

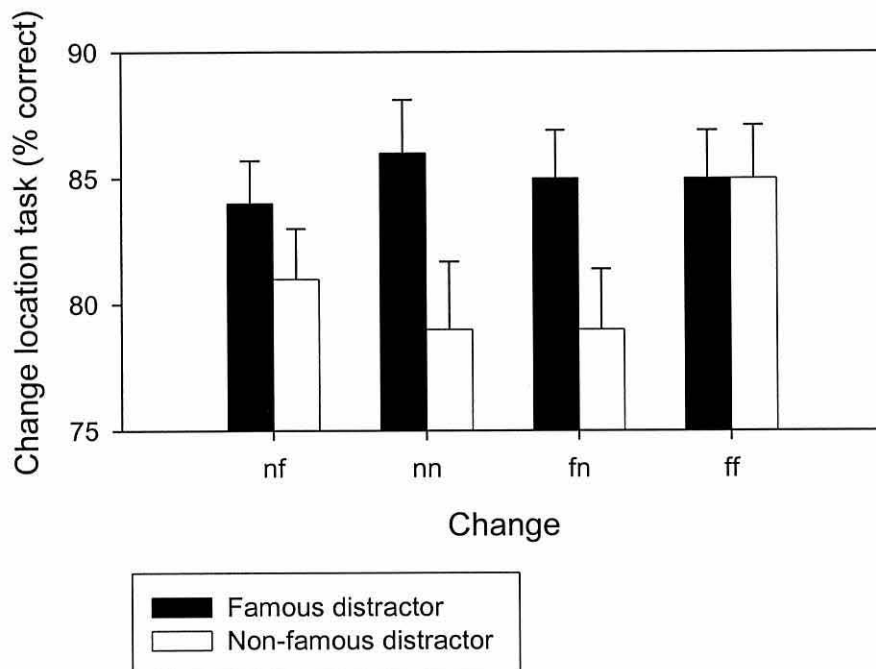


Figure 19. Group mean percent correct change location identification for trials in which the distractor (non-changing) landmark was famous (black bars) or non-famous (white bars) for each of the four change conditions. Vertical bars indicate ± 1 s.e.

Figure 19 shows the mean percent correct change location identification for each of the four change conditions. There was a significant main effect for distractor condition (non-changing image) $F(1,26) = 9.221, p < .01$. When a famous distractor was present in the display change detection performance was enhanced. The ANOVA also indicated non-significant main

effects of visual field and change condition, and non-significant interaction effects involving all three factors.

Figure 20.




Figure 20. Group mean percent correct change location identification for trials in which the distractor (non-changing) landmark was famous (black bars) or non-famous (white bars) for conditions when the second of the changing items was either famous or non-famous. Vertical bars indicate ± 1 s.e.

The effect of a famous distractor was largest in the NN change condition where the presence of a non-famous distractor led to change detection being on average 7% lower. However, in the condition FF, performance was the same for both famous and non-famous distractor conditions (mean = 85%). So on trials where the distractor was non-famous, but the change had two famous items the mean was 6% higher than when all items were non-famous (NN with non-famous distractor). This suggests that even though there were non-significant effects of change condition, a famous landmark in the changing image may have also aided task performance.

Further consideration of the non-famous means as depicted in Figure 19 indicates that when the changing image displayed a famous item as the second item (NF, FF) performance accuracy on non-famous distractor trials was enhanced relative to non-famous second images (NN, FN). A repeated measures ANOVA was calculated with two factors (distractor: famous, non-famous; and second image: famous, non-famous). This revealed a marginally significant interaction of distractor and second image ($F(1, 26) = 3.7, p < .06$). The means in Figure 20 show that when the distractor was non-famous a famous landmark presented as the second of the changing items improved accuracy scores (mean = 83%) when compared to non-famous second items (mean = 79%), $t = 2.12(26), < .02$.

Figure 21.




Figure 21. Group mean percent change location identification across quartiles. Vertical bars indicate ± 1 s.e.

Effect of session.

In order to assess the contribution of recent familiarity on the visual processing of the non-famous landmarks an ANOVA on the quartile effects was conducted. There were non-significant effects of change condition, distractor condition, and interactions. Although the ANOVA did not reveal any significant effects of quartile, the means indicate there was a small improvement in accuracy over time. Figure 21 shows that quartiles 1 and 2 both produced an overall accuracy of 80%, while quartiles 3 and 4 both produced accuracy scores of 82% and 83% respectively.

Discussion

Using an implicit measure of landmark processing that required participants to locate a change in a two-item display, the effect of famous versus non-famous landmarks was assessed. There were three main findings. First, the presence of a famous landmark in the display aided the detection of change. The clearest evidence for this was in the distractor (non-changing) condition, where a famous distractor produced greater accuracy than non-famous distractor trials. However, there was also some evidence that a famous item in the changing image also aided performance, though this was restricted to trials where the famous item was the second of the two changing images. A second finding was that there was no evidence of a visual field effect, unlike the face experiments of Chapter 3 (Experiments 1&3). A third finding was that there was only a slight improvement in change detection accuracy over the four quartile time course of the experiment.

The advantage for detecting changes when famous landmark distractors were present advocates that the visual fame effect for landmarks was due to speeded processing allowing access to stored representations. The advantage found with famous distractors indicates that participants relied on the representation formed by the longer presentation of the distractor (200ms

presentation compared to 100ms for each changing item). This may have led to participants inferring where the change had occurred (e.g., "If left image didn't change, the change must have been on the right").⁷ When the distractors, were famous landmarks the processing was faster and accuracy was greater.

However, participants were less able to benefit from a famous landmark in the changing condition, which suggests that the 100ms exposure (of the first or second landmark) did not permit adequate processing to occur, and therefore did not allow the different successive presentations to be easily detected. This does not mean that there were no benefits from the presence of a famous item in the changing image, but rather they were subtler than previously observed with face stimuli. For instance, the change condition FF (famous to famous) produced a high accuracy score regardless of distractor type. Furthermore, when there were only non-famous distractors those changes where a famous landmark was placed as the second of the changing items were more successfully detected than when a non-famous item was in the second image. Though the evidence for the effect of fame in the changing condition is more tentative (marginal significance) than the effect of distractor, it is important to consider what might underlie a famous advantage for the second of the two changing items. A likely explanation is that famous items are more resistant to attentional blink effects. This would allow comparisons to be made with the first image, leading to successful change detection. This explanation is further supported by the significant effect of fame in the changing conditions for faces, where greater accuracy was found for changes again involving the second of the two changing items (i.e., NF was significantly different to NN).

The lack of a visual field effect contrasts with the left visual field advantage found for faces. Face stimuli revealed an advantage for changes

⁷ The reports of participants during debriefing indicated that this was a feasible explanation. When asked how they felt they performed a number of participants reported that they had relied on detecting the image "that was on for longer" (distractor).

involving famous faces when presented to the left visual field, but not to the right visual field (the famous distractor advantage for faces in Exp. 3 occurred regardless of visual field). I suggested that this could be explained by preferential processing of configural information by the right hemisphere.

There are two possible accounts for why there was no evidence of a laterality effect: (1) faces really are special in comparison to other object categories, (2) specialised expertise is required such as personal experience with landmark stimuli. Support for the first possibility can be found in research, which directly compares faces with other classes of object. Ro, Russell and Lavie (2001) made a direct behavioural comparison of change detection performance between faces and objects when all were present in the display. Six items (one of which would be a face) were presented in a circle for 500ms. After a brief blank display, six items would be displayed again (500ms). The participants' task was to indicate which of the items had changed. Their results showed that participants were both significantly faster and more accurate when a face changed than when any other type of object (food, clothes, musical instruments, appliances, and plants) changed. The conclusion drawn by the authors is that faces play a special role in attention. Perhaps the difference in the ability to detect changes in the face and landmark tasks arose due to greater attention/arousal effects of faces that made the task more salient.

The other possibility stems from the fact that participants were unlikely to have personal experience with the famous landmarks and their surrounding areas. This meant that the participants did not have any background scene knowledge of where these landmarks occur, and how they should be navigated. The probable explanation for why landmarks should be preferentially processed is that they help define the space around us (Epstein, Harris, Stanley, & Kanwisher, 1999). In the case of these famous but rarely encountered landmarks there is no such connection with real experience and therefore the importance and saliency may be weaker compared to

objects/landmarks that are experienced in the physical sense. This relates to reports of configural processing with non-face stimuli in experts in whom there has been at least 10 years experience with an object category. For example dog experts have produced configural-processing effects when novice participants failed to do so (Diamond and Carey, 1986). In the landmark study here participants may not have been expert enough with these particular landmarks to benefit from configural processing. Moreover, landmarks unlike faces did not produce an improvement in task accuracy from the first through to the fourth quartile. While accuracy significantly increased for faces on trials where the change was presented to the right visual field (13% increase from quartile 1 to quartile 4), landmarks did not reveal any significant improvements for either visual field. This indicates that familiarity may take longer to establish in landmarks than in faces, and again suggests that experience with non-face stimuli may be crucial to how salient a stimulus is and to how reliant processing can be on configural information.

In summary, a visual fame effect was evident for landmark stimuli indicating a visual processing advantage for famous items. However, the landmark data did not reveal any evidence for preferential configural processing over featural processing, as would be indicated by effects of visual field. This disparity in visual fame effects between landmarks and faces could be due to the special nature of faces, or lack of personal involvement/experience with the famous landmarks.

Chapter 5 – Fame Effects in Consumer Products

Introduction

Differences in visual fame effects for two distinct classes of object (faces and landmarks) were demonstrated in Chapters 3 and 4. When face changes had to be located an advantage occurred when a famous face was involved in the changing image (target) of the left visual field. There was also some evidence that a famous landmark in the changing image could aid change detection. However, the strongest effect of fame in landmarks occurred when famous landmarks were presented as the non-changing images. These findings indicated that even though each famous landmark was just as unique as the famous faces, they were not processed in the same manner as face stimuli. With the landmark stimuli it appeared that participants relied on the stability of the non-changing landmark to infer where the change had taken place, rather than detecting the change itself. Moreover, effects of fame for faces and landmarks differed in the way that left visual field versus right visual field information were used. There was a left visual field advantage for faces, but no specific visual field advantage for landmarks. The difference in processing ability with these two classes of object may be due to the specialised nature of face processing.

In this Chapter, I investigate the effect of fame on a class of object that differs in a number of aspects from face and landmark stimuli. With everyday household goods (such as shampoo, laundry detergent, soft-drinks, tinned food etc), participants will have interacted in a number of ways with such objects. These interactions range from searching for and selecting items, to manipulating the items for use (e.g., opening, pouring etc). Naturally, on a daily basis we all manipulate a variety of products, from the shampoo we use to wash our hair, to the cereal we eat at breakfast. This provides an

opportunity for significant object familiarity, a feature that may have been limited with the landmark stimuli.

Another intriguing facet of product stimuli is the incorporation of words into objects. The combination of container and brand name leads us to easily categorise these objects at multiple levels. For example, if we viewed a can of Coca-Cola, we can easily recognise it as a drink (superordinate level categorisation), as a cola (basic level categorisation), and as Coca-Cola (subordinate level categorisation). Therefore, product stimuli offer a special opportunity to evaluate the relative contributions of lexical information, categorisation, and object recognition in an externally valid selection of items.

In this study I sought to explore the perceptual processing advantages conferred by high familiarity (fame) in branded products. Using the change detection task of Chapter 2, which did not require explicit semantic and verbal recognition, an implicit measure of image processing was acquired. Two products, on either side of a fixation point, were presented briefly and then masked. A similar pair of products were again presented and then masked. One product in each frame stayed the same and one changed to a different product. The task was to report the category of the products on the left and right and to report the location of the product that changed. This required attention to be directed to both products and representations of the first and second pair of images to be compared. Because I was interested in the effects of familiarity, famous products were used in all possible combinations with non-famous products. My procedure also allowed me to compare the effects of left and right hemi field presentations, as the right visual field may be expected to benefit from the lexical information contained on the product (e.g., Ivry and Robertson, 1998). In experiment 5 a within-product type change detection task (e.g., one brand of cola changed to another brand of cola) was used. Following the counter-intuitive finding that non-famous changes (NN) were detected more accurately than changes involving famous products, Experiment 6 was conducted in order to

eliminate the possibility that semantic repetition blindness (two-related items being coded as the same) could occur. This required between-category changes (e.g., a can of baked beans changed to a can of cola), where semantic associations between the changing items should not exist. If visual fame bestows an advantage on perception (as opposed to naming memory) change detection should be more accurate when a famous product is present in the display, than when only non-famous products are present. If, on the other hand, visual fame aids memory processes without enhancing visual processes, then simple detection of image changes should be unaffected by fame. These experiments allowed affirmation of two issues: (1) that a visual fame effect exists for product stimuli, and (2) that semantic blindness exists for real-world stimuli that contain both pictures and words.

General Methods

Participants

Thirty-two British undergraduates (25 females, 7 males), ranging in age from 18-32 years, volunteered to participate in exchange for course credit. Seventeen were right-handed and five were left-handed.

Apparatus and Stimuli

The experiment was run on a G4 Macintosh computer. Stimuli were displayed on a 19-in. (43-cm) colour (75Hz) monitor.

Product stimuli were rectangular greyscale digital photographs subtending 12 deg by 7.6 deg. The centre of each product was positioned 5.4 deg along the horizontal meridian to the left and right of a small central black fixation cross. Two different sets of stimuli were presented in Experiment 5 (Set A and B) but only one set was used in Experiment 6 (Set C). Each set comprised eight products: four non-famous and four famous. Within each

Table 1: Brand product and container type.

Set A:		
	<u>Famous</u>	<u>Non-famous</u>
<u>Washing-powder(box)</u>	Persil	Cheer
	Surf	Gain
<u>Shampoo (bottle)</u>	Organics	Abba
	Pantene	Nexus
Set B:		
<u>Cola (can)</u>	Coca-Cola	Tesco brand
	Pepsi	Safeway brand
<u>Lemonade (bottle)</u>	Sprite	Londis brand
	7up	Safeway brand
Set C:		
<u>Cola (can)</u>	Coca-Cola	Safeway brand
<u>Spaghetti/beans (can)</u>	Heinz	Tesco brand
<u>Lemonade (bottle)</u>	Sprite	Londis brand
<u>Shampoo (bottle)</u>	Organics	Nexus

product set, contrast and clarity of each image and product size was adjusted to a roughly similar level. Set A products were shampoo bottles (two famous and two non-famous) and washing-powder boxes (two famous and two non-famous). Set B products were lemonade bottles (two famous and two non-famous) and cola cans (two famous and two non-famous). Set C contained bottles (one famous and one non-famous shampoo and one famous and one non-famous lemonade) and cans (one famous and one non-famous cola and one famous and one non-famous spaghetti). Non-famous products were randomly selected from North American leading brands that are unknown to British consumers and unadvertised (low exposure) British brands. Famous products were selected from leading brands (frequently advertised – high

exposure) in Britain at the time the study was conducted (see Table 1). Each product was photographed against a plain white background using a digital camera (Sony, MiniDV Handycam vision, DCR-TRV 900E).

A pattern mask was a scrambled collage of product parts similar in details to the products used in the study. The same mask was used throughout the experiment.

Procedure

The procedure was identical to the details in the General Method (Chapter 2) except for the choice of images and the categorisation task (see Figure 22). In these experiments participants had to report the type of product container (Experiment 5 Set A: box or bottle; Set B: can or bottle, and Experiment 6 Set C: can or bottle).

Figure 22.




Figure 22. An illustration of a trial. Following a 1000ms fixation cross a pair of products are briefly (100ms) presented followed by a scrambled image mask (100ms). A second pair of products were presented (100ms) with mask (100ms). The second products image includes one product as presented in the first image, and one, which is a completely different product.

Experiment 5

Only within-category changes were made. For example, a shampoo would only change into another shampoo. Product Set A was used for 9 participants and Set B was used for 12 participants.

Results

Category task.

On average, participants correctly named the container type of the left and right products on 74% of trials (s.e. = 5.9%). The difference in performance for Set A and Set B was non-significant. All participants scored 70% or better on this task. An ANOVA revealed non-significant effects of change condition, distractor, and category. Therefore, there was no evidence that fame aided the categorisation task.

Change condition.

Only trials where both categorisation judgments were correct were analysed. An ANOVA on the percentage correct location identification responses for the change detection task was conducted using one between (product set) and three within factors (visual field: left or right; distractor: famous or non-famous; and change condition: NN, FF, FN and NF). Since there was a non-significant effect of product set and all interactions involving product set were non-significant, further discussion of the data refers to the combined results of both groups.

Figure 23 shows the mean percent correct change location identification for each of the four change conditions. In contrast to previous findings with other objects, visual fame can be seen to hinder change detection performance on this task. This is supported by a significant main effect for change condition, $F(3, 63) = 3.77, p < .05$. When a famous product was present

in the changing display, change detection performance was reduced on average by 4% when compared to a change that only involved non-famous products (NN = 69%).

Figure 23.

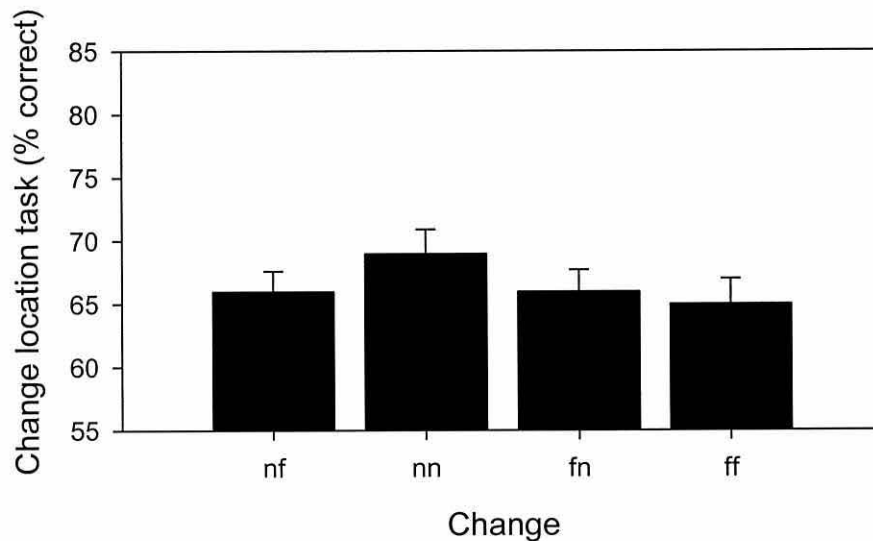


Figure 23. Group mean percent correct change location identification for each of the four change conditions. Vertical bars indicate ± 1 s.e.

The ANOVA also indicated non-significant main effects of visual field and distractor, and non-significant interaction effects involving all three factors.

Effect of session.

In order to assess the contribution of recent familiarity on the visual processing of the non-famous landmarks an ANOVA on the quartile effects were conducted. The main effect of quartile was non-significant, as were all the interaction effects involving this factor. This indicates that like landmark stimuli there was no benefit from repeatedly viewing the brief presentations of the non-famous products, and that the difference between famous and non-famous distractors was consistent from the first to the fourth quartile.

Discussion

The advantage for NN changes is in stark contrast to previous visual fame effects, where the presentation of a famous item led to perceptual benefits for both face and landmark stimuli. However, product stimuli differ from these other object categories in that products cannot be individuated to the same degree. While there is only one Statue of Liberty, and it is distinct from other landmarks such as the Eiffel Tower, products such as shampoo are displayed in multiples and do not differ greatly from one another (e.g., they are all contained in bottles). A significant visual difference between brands of shampoo is the names/logos that are displayed on the item (e.g., Organics or Pantene). However, these famous names are inextricably associated with the product type. Therefore, it is possible that the combination of physical similarity (e.g., shape of container) and product association may lead to a repetition blindness effect. When 'Pepsi' is followed by 'Coca-Cola' the activation of the type 'Cola' may not be instantiated into the two different occurrences. This would lead to a greater repetition blindness effect for famous items, than with non-famous products where the lexical information would be of less meaning. This explanation would account for the advantage of NN trials (69%) over FF trials (65%).

However, a debate exists in the literature surrounding the factors contributing to occurrences of repetition blindness. Before discussing these details it is important to discuss the background of repetition blindness (RB). RB as described earlier refers to a difficulty in detecting repetitions presented in serial visual presentations (Kanwisher, 1987; Kanwisher & Potter, 1989, 1990). This occurs at rapid presentations rates even when several intervening items separate the two incidences of the repeated item. The initial RB findings were restricted to the study of repeated words (Kanwisher, 1987) where sentences such as "It was work time so work had to get done," would be recalled minus the second occurrence (the word 'work'). This happened even

when the sentence became ungrammatical or if meaning was lost. Later studies established that RB could also arise from the repetition of pictures (Arnell & Jolicoeur, 1997; Bavelier, 1994; Kanwisher & Yin, 1993).

The RB phenomenon has been interpreted in terms of a distinction in type recognition (distinct visual categories) and token individuation (distinct spatiotemporally defined visual objects or episodes) in Kanwisher's (1987) 'token individuation' hypothesis. Naturally, the visual system is required to link these types and tokens together. RB is a case of recognising a repetition's type (e.g., the lexical/pictorial entry is activated for both occurrences), while failing to individuate (tokenise) the items as distinct events (e.g., Chun, 1997; Kanwisher, 1987, 1991). This failure results in an interpretation of the repetition of the type activation as being residual activation from the first instance. Therefore, the assimilation of the second occurrence into the first leaves only one token of the type to be registered consciously.

What kinds of repetition produce repetition blindness? Bavelier (1994) makes a number of assumptions about RB (see Table 2) which infer that the early stages of processing words involve linking visual and phonological types to a token, with semantic information only becoming incorporated later. This is supported by Kanwisher and Potter's (1990) failure to implicate semantics in repetition blindness for words. Using sentences containing noun synonyms as repeated items (e.g., "The company's new *toxin* might *poison* people accidentally.") no evidence to support the possibility that *semantic* blindness might occur was found. However, these assumptions indicate that the case of repetition blindness for pictures varies in its early processing stages from words. For pictures it may be expected that visual and semantic types are linked to a token early in processing, whereas phonological information is incorporated later. Kanwisher, Yin, and Wojciulik (1999) report findings that appear to support these propositions. In a series of experiments Kanwisher et al. (1999) consistently failed to produce RB for semantically related words. However, they did succeed in producing RB for semantically

related pictures. Using a rapid serial visual presentation that would include semantic repetitions (that were visually distinct), such as a picture of a helicopter followed later by a picture of an airplane, a deficit in reporting the second occurrence was found. However, there was no such effect for using the words 'helicopter' and 'airplane' in rapid presentations of words.

Table 2: Bavelier (1994): assumptions on the nature of repetition blindness.

1. There are multiple types available for integration into a token.
2. The rate of information accrual differs across types.
3. The order of information accrual is different for pictures and words.
4. Repetition of type information that is integrated into a token relatively early in processing is more likely to result in repetition blindness than repetition of information integrated later.

MacKay and Miller (1997) raise the concern that the use of noun synonyms in the same language failed to provide pairs of words with identical meanings. They argue that in the previous example using 'toxin' and 'poison' there are differences in both lexical and connotative meaning. Furthermore the authors suggest that within-language synonyms involve different lexical nodes, while only one lexical node underlies familiar words in bilingual speakers (e.g., *horses* and the Spanish equivalent *caballos*). When this hypothesis was tested using translation equivalents in RSVP (rapid serial visual presentation) streams a repetition blindness effect was evident. This supports MacKay and Miller's (1997) contention that two types of word blindness exist: orthographic-phonological and semantic. Supporting this position Parasuraman and Martin (2001) reported that semantic associations

between words could modulate RB. This led them to conclude that semantic and lower-level perceptual presentations can interact to affect performance.

Clearly the exact role of semantics in the early perceptual processing stages that give rise to the RB effect is yet to be fully established. However, the use of branded products in a briefly presented change detection task provides naturalistic stimuli from which to assess the potential for semantic blindness. Unlike the artificial nature of presentation in MacKay and Miller's (1997) study, which incorporated a mixture of two languages in their RSVP task (bilinguals tend not to mix two languages together in one sentence) most participants are experienced with associating a number of brands with a particular type of product (e.g., *Persil*, *Daz*, *Arial*, and *Bold* are all popular brands of detergent in Britain). Therefore, if famous brands are semantically associated with the product at an early stage of processing a change involving two related famous products could induce semantic blindness. This would lead to an advantage for changes involving non-famous brands that do not activate the semantic category.

To illustrate this point consider the detergent example. If the brands *Persil* and *Daz* (tokens) activate the category *washing-powder* (type) a briefly presented change from one to the other may not be detected (due to a failure in token individuation). However, if non-famous brands (e.g., The US brands *Cheer* and *Gain* are unknown in Britain) are involved in a change there are no available associations to the category washing-powder. This would decrease the likelihood of semantic blindness occurring and lead to a potential non-famous change detection advantage (the opposite result to face and landmark changes).

Experiment 6

In Experiment 5 an advantage for detecting non-famous product changes was found. This may have been due to a semantic repetition blindness effect of viewing famous products. If the association of two branded products to one category produced this proposed RB effect then the effect should be eliminated by detecting changes between-product types (e.g., a can of beans changes to a can of Cola) as opposed to within-product types (e.g., Coca-Cola changes to Pepsi). To test this possibility, Experiment 6 used the same procedure as before, this time using between category products (see Table 1).

The stimuli of Product Set C were used in this experiment. All other methods were the same as in Experiment 5. Nine females and one male ranging in age from 18 to 30 years participated.

Results

Category task.

On average, participants correctly named the container type of the left and right products on 76% of trials (s.e. = 4.3%). All participants scored 70% or better on this task. An ANOVA revealed non-significant effects of change condition, distractor, and category. A between subjects ANOVA revealed no significant differences between Experiment 5 and 6 on this categorisation task.

Change detection.

Only trials where both categorisation judgments were correct were included in the analysis. An ANOVA on the percentage of correct location identification responses for the change detection task was conducted

using three within factors (visual field: left or right; distractor: famous or non-famous; and change condition: NN, F, FN, NF).

Figure 24.




Figure 24. Group mean percent correct change location identification for each of the four change conditions. Vertical bars indicate ± 1 s.e.

Figure 24 shows the mean percent correct change location identification for each of the four change conditions. The graph indicates that there was no difference in change detection for famous or non-famous product changes. This is supported by a non-significant effect of change condition. Furthermore, the NN condition (mean = 74%) is no longer the more accurately detected change (FF mean = 76%).

Moreover, when the data from Set B (Experiment 5) and Set C (Experiment 6) were compared, as both shared 50% of the same stimuli, a significant difference was found for set used, $F(1, 20) = 5.6, p < .05$. Figure 25 displays the mean change detection accuracy for the two sets of data. It appears performance on the within-product type task (Set B) was on average 11% lower than the between-product type task (Set C). This was greatest for the FF condition (17%) where a semantic RB effect was most likely to have

occurred in Experiment 5 (two famous same product type items are more likely to be associated).

Figure 25.




Figure 25. Group mean percent correct change location identification for each of the four change conditions using Set B (Exp 5) and Set C (Exp 6). Vertical bars indicate \pm s.e.

Although there were no significant effects of change condition in Experiment 6, there was a significant effect of distractor, $F(1, 9) = 14.45, p < .01$. As Figure 26 indicates the presence of a famous distractor aided change detection performance in all change conditions. This resembles the pattern of results revealed with landmark stimuli in Chapter 4.

The ANOVA also indicated non-significant main effect of visual field, and non-significant interaction effects involving all three factors.

Figure 26.




Figure 26. Group mean percent correct change location identification for trials in which the distractor (non-changing) product was famous (black bars) or non-famous (white bars) for each of the four change conditions. Vertical bars indicate ± 1 s.e.

Effect of session.

In order to assess the contribution of recent familiarity on the visual processing of the non-famous landmarks an ANOVA on the quartile effects were conducted. There were non-significant effects of quartile condition, and interactions involving this factor. As in Experiment 5 it appears there was no benefit from repeatedly viewing the brief presentations of the non-famous products and no decrement in the accuracy difference between famous and non-famous stimuli.

General Discussion

Experiments 5 and 6 explored the effect of fame on the detection of changes to two simultaneously and briefly displayed product images. Experiment 5 used a within-product change task that revealed an unusual effect of fame. When one famous product changed to another famous product (both easily identified brand names) deterioration in change detection accuracy was evident when compared to non-famous to non-famous (NN) product changes. However, in Experiment 6 changes were made between-products categories. The performance on the change detection task produced no evidence of a NN change advantage, but instead revealed a detection advantage when famous products were presented as the distractor (non-changing image). Therefore Experiment 6 was consistent with the visual fame effect found for landmarks in Chapter 4.

The difficulty in detecting changes between two famous products of the same type is likely to be due to a repetition blindness effect for the display of the products brand names. Repetition blindness (RB) occurs when only one instance of a repetition is reported, as the second occurrence fails to be encoded as a separate token from the first. In the case reported here there is evidence of a special and controversial case of RB, termed by MacKay and Miller (1994) as semantic blindness. If we consider a change condition where an image of 'Pepsi' changes into 'Coca-Cola' clearly both are distinct examples of the product type 'COLA'. If semantic access to the type 'Cola' is achieved, it is possible that the two different brands were coded as one token of 'Cola' as opposed to two.

The RB literature reports that the effect occurs in rapidly presented stimuli where the two items follow in close succession (within 300msec, Chun, 1997). This is consistent with the presentation times involved in the change detection display (200msec SOA). Furthermore, according to Potter's (1993) Conceptual Short-Term Memory (CSTM) hypothesis, when a stimulus is identified, its meaning is rapidly activated and maintained briefly in CSTM.

Evidence for this hypothesis comes in three forms: 1) there is rapid access to semantic information about a stimulus; 2) the activated conceptual information can be used to build a structured representation of the information; and 3) there is rapid loss of information that does not become structured. Therefore, according to this account each of the changing product images could be identified along with its meaning. However, there was likely to be a loss of information that had not been structured. This raises the possibility that while both images were coded as 'COLA' the information distinguishing the two products may have not become stably encoded.

Potter (1999) clearly links the CSTM hypothesis with repetition blindness effects. The representation form of CSTM involves the activation of recognised items (familiar) to form new structures. Hence, existing types (and their associations) have to be activated in CSTM. However, the new structures will have to be represented by establishing tokens of the relevant types. In the case of RB, the structuring process that stabilises associated items in CSTM appears to collapse the two events into one. This CSTM hypothesis can be considered alongside Kanwisher's (1987, 1991) 'token individuation hypothesis', which argues that RB is a result of 'the failure to bring a fleeting type activation into awareness by stably representing the episode in which it was presented.' Hence the activation from the repeated item is interpreted as being residual activation from the first instance. Therefore, it seems plausible that the data using within-product changes is due to a semantic blindness effect. This pattern of data would not have been anticipated with the face or landmark stimuli of Chapters 3 and 4, as none of the famous items would activate the same associated meaning, even though all the items belonged to one large category of object (e.g., 'Prince Charles' does not have any major semantic connections with 'Madonna').

The RB interpretation of the data from Experiment 5 was further supported by the absence of a NN advantage in the between-product changes of Experiment 6. In this experiment, the famous changes would not permit the


same semantic activation (e.g., the shampoo 'Organics' and the soft drink 'Sprite' do not have a close semantic relationship, they are clearly two different types of product). In this case, participants benefited from the presence of a famous item, although only when it appeared as the distractor. Here, as in the case of landmarks it appears participants were unable to stabilise the representation encoding to the level required for detecting the actual changing stimuli. Instead, they relied on the extended viewing time (amount of neural activation) of the distractor (non-changing image) to infer where the change had occurred. Therefore, it seems that while visual fame effects did exist for landmark and product stimuli, these objects could not be encoded to the same level as faces in this briefly displayed change detection paradigm.

Tentative support for this position comes from an ERP (Event Related Potentials) study into the time course of the repetition blindness effect (Bavelier, 1999). Bavelier (1999) proposed that RB occurs at the bridge between perception and short-term memory processes. If so, RB should emerge after basic perceptual processes have been completed and higher-level visual processes are engaged. However, this effect will appear before short-term memory maintenance has been recruited. According to this view the effect of RB should occur between 100-200msec when high-level aspects of object recognition are engaged (e.g., Anllo-Vento, & Hillyard, 1996; Thorpe, Fize, & Marlot, 1996) and 400msec when short-term memory is unlikely to be consolidated (e.g., Ruchkin, Canoune, Johnson, & Ritter, 1995). Bavelier's (1999) findings are consistent with this description, where the main differences in ERPs to repeated and unrepeated stimuli occurred as early as 200msec after onset of the second instance. If RB occurs at the bridge between visual processing and encoding in short-term memory (as appears to be the case in Experiment 5), it is possible that the failure to find an advantage for changes involving famous (Experiment 6 with products and Experiment 4 with landmarks) items was due to a failure to compare the two items in STM.

Instead the longevity of activation from the non-changing image revealed the location of the distractor, and so the change could be inferred. This was enhanced when a famous item was displayed as the distractor, leading to greater activation and speeded processing. This raises the possibility that the advantage found for famous faces in the changing image may be due to successful comparisons in short-term memory. (This will be discussed in the General Discussion of Chapter 7.)

A controversial aspect of the studies presented in this chapter is the possibility that the lexical information (brand name) contained in the images contributed to the RB effect. A debate exists in the literature as to whether semantically related words can produce RB. While a repetition blindness effect has been established for semantically related pictures (e.g., Kanwisher, Yin, & Wojcicki, 1999), the evidence for a semantic RB for words is still contentious. Kanwisher, et al. (1999) failed to produce RB for semantically related items presented as words, even though when presented as pictures they did produce RB (e.g., 'mouth' and 'ear' stimuli are both body parts). In fact, Altarriba and Soltano (1996) revealed a facilitation effect from presenting two semantically related words (one concept in two different languages), suggesting token individuation was possible for translation equivalents. In contrast, MacKay and Miller (1994) demonstrated translation equivalents (Spanish/English) can produce repetition blindness.

Interestingly, Bavelier (1994) reported repetition blindness for 'rebus' sentences. A rebus sentence contains pictures as well as words, such as:

The duck quacked when the  was hungry.

Clearly, the visual image of a duck is vastly different to the lexical form. However, this finding does not necessarily implicate semantic/conceptual processing of the word instance. Bavelier's (1994)

proposal that word and picture stimuli follow different sequences of activation in the perceptual system, does not necessarily contradict the RB result. Instead she suggests that repetition blindness in this case is mainly mediated by the identity of their phonological representation. Using this conjecture to explain the repetition blindness produced by viewing two brands (exemplars) of one product (type), it is possible to argue that the product type was phonologically encoded and this was the factor that induced RB. However, in order to identify the product efficiently the brand name was required to be related to the product. (E.g., "Persil" and "Ariel" are not phonetically the same, but are semantically related.)

While the visual similarity of two within-product brands may lead to some RB, it was clear from the data of Experiment 5 that changes involving famous products (where the name can be associated to a product category) produced a greater RB effect than changes involving non-famous products (where the name is not associated to a product category). Furthermore, the between-product changes used visually similar items (e.g., can of beans, can of cola), which would be equally likely to produce RB as the within-product changes. However, performance on the between-product task was superior to that of the within-product task. Therefore, some of the RB effect for famous product changes is attributable to the semantic/conceptual activation of the brand name.

Two explanations can be offered for why repetition blindness occurred for famous product names while other tests of RB on words have failed to find an effect. One possibility is that, like the stimuli in MacKay and Miller's translation experiment, the pairs of items used in Experiment 1 have a more immediate association than the items used in other experiments. The word 'horse' has a direct relationship with the Spanish equivalent; similarly 'Persil' and 'Daz' have a direct relationship to the category washing powder. Hence, in terms of spreading activation these concepts are likely to be linked at an early stage of semantic analysis. However, take the example of 'helicopter'

and 'airplane', although instinctively the items appear to be closely related, it is not necessarily the case that when presented with the word 'helicopter' we activate the concept of 'mechanical air transport', we could relate the word to 'rescue' or 'military', just as 'airplane' may relate to 'holiday'. This explanation does not dispute Bavelier's (1994) assumptions that information from words and pictures are accrued in different orders. It suggests that while semantic access to picture information takes place early in processing (producing robust effects of RB), limited semantic access is available from briefly displayed words, although only through the immediate stages of semantic activation.

The second possibility is that the combination of brand name and visual container information allowed the visual image of a can, bottle etc to facilitate recognition of the brand and its product type. Perhaps the brand names alone would not be processed to the early stage of semantic/conceptual analysis described previously. Rather, the context of the word with a congruent object facilitated the recognition of product identity. It is even possible that there was some combination of these two accounts involved in the repetition blindness result. However, the first explanation is tempting as it provides a coherent account for the divergence in RB effects for words.

To summarise Experiments 5 and 6 appear to have found a semantic repetition blindness effect for consumer products that is mediated by the famous brand name displayed on the items. Experiment 5 demonstrated that when two famous brands of the same category are changed from one to the other, change detection accuracy is reduced compared to two non-famous brands being involved in the change. This is explained by activation of conceptual information that associates the product type of the two famous items together leading to a repetition blindness effect. However, when the potential for RB was eliminated (by making between-product changes in Experiment 6) an advantage was reported for change detection when the non-

changing image was famous. This replicates the visual fame effect found for landmark stimuli in Chapter 4.

Chapter 6 - Fame Measurements in Political Faces

Having established a visual fame effect for faces using a two-item change detection task I wanted to explore the possibility that this technique could be used to measure and compare individual famous faces. If this paradigm was sufficiently sensitive and could reveal differences between famous faces, it could provide a useful assessment tool for market researchers. For instance if an advertising company wishes to approach a football player to endorse a range of products, the quantitative assessment of which players are more famous could aid the selection. Similarly, market researchers could benefit from such a technique when addressing the issue of how well politicians are known.

While as a society we hope that voting decisions are made on sound political judgments, it is quite clear that any political party that has an unpopular or even an unattractive leader will suffer at the voters hands. Zebrowitz (1994) supported this view in her description of the political prospects of different politicians according to their type of face (e.g., mature, persuasive, powerful etc). Therefore, any technique for gaining an insight into people's perceptions of famous faces could benefit pollsters and market researchers.

An examination of the fame effect for the faces of the leaders of the three main political parties (Labour, Conservative, Liberal Democrats) would be an apt and useful way to assess whether the change detection paradigm would be sensitive to differences among famous faces. An opportunity arose to test the visual fame effect in July 2001. This was the month of the general election, where in the prior weeks the media presented a constant stream of information about the three main political parties and their leaders. The fortnight before the election date provided an occasion for assessing how

famous the three leaders were and whether the change detection technique could be adapted as a fame measurement tool.

The background to the 2001 election will now be summarised. A Labour government with Tony Blair as Prime Minister and Labour leader was elected in 1997 with a landslide victory usurping the previous Conservative government (under John Major). As a reasonably popular leader Tony Blair (sometimes referred in the media as 'Princess Tony' as he appeared to be as well liked as the late Princess Diana) naturally represented his party as leader in the 2001 election. However, due to the appalling defeat the Conservative party endured in 1997 John Major resigned as leader and was replaced by William Hague. William Hague was expected to turn the party's fortunes around and regain a large number of the parliamentary seats lost in 1997. However, due to his receding hair and rather round face he was often referred to as a "baby face" (e.g., Younge, 2001). This could be a problem for him as Zebrowitz (1997) reports that people with immature looking faces are often also regarded as being immature in terms of their personality. In addition, there was a further new leader entering the 2001 election campaign. Following the success of the Liberal Democrats in the elections of 1997, Paddy Ashdown who had led his party to their largest parliamentary presence with 48 seats, decided it was time for him to resign as leader. Charles Kennedy succeeded him, and while of a likeable demeanour it was considered to be a tough election for him and his party without the charismatic Ashdown and especially after the 1997 success (Compuserve news, 2001). Therefore, by timing the experimentation of the change detection task involving these political faces with the election campaign not only could the sensitivity of the measurement be tested but also a comparison of the leaders election outcomes could be made.

In this study the change detection task of Chapter 2 was adapted so that only those trial conditions that evidenced a strong and reliable visual fame effect were used. In Chapter 3 the visual fame effect for faces was

revealed by a processing advantage for detecting changes involving a famous face when compared to changes that only involved non-famous faces (NN trials). In particular the non-famous to famous changes (NF trials) produced the greatest effect. Therefore, in order to increase the likelihood of finding differences between the politicians' faces only NN and NF changes were made. However, though the fame effect had been limited to changes presented in the left visual field, changes were made to both visual fields to avoid participants only attending to one side of the display. The NN trials constituted the baseline from which to measure whether the changes involving political faces were significantly different.

Experiment 7

Method

Participants

Ten British volunteers (6 males, 4 females), ranging in age from 18-38 years, participated in exchange for £2.50 (for half an hour). All were right-handed, reported normal or corrected to normal vision, and were naïve to the purpose of the experiment. Informed consent was obtained prior to participation.

Apparatus and Stimuli

The experiment was programmed in Psyscope (v.1.2.2), and run on a PowerBook 1400cs Macintosh computer. Stimuli were displayed on an 11-in (29-cm) colour (75 Hz) monitor. Responses were recorded via the computer keyboard. Testing was conducted in a small room with low ambient illumination.

The presentation format and criteria for selecting the face stimuli was as reported in Chapter 2, with the following differences. Eight face images were randomly chosen from a selection of male business faces to be displayed

in the first of the two images in a trial (all non-famous), and eight different face images (including the three political faces) were randomly chosen to be displayed as the changed face in the second image. All of the non-famous faces were selected so that age, hairstyle, and, general appearance were similar to the politicians' faces. This was important in order to reduce perceptual differences between faces. For instance, William Hague had a receding hairline, while Tony Blair and Charles Kennedy both had a full head of hair. Therefore, both the non-famous first and second images also included both types of appearance. A point to remind ourselves is that problems of distinctiveness are difficult to fully resolve with such complex and naturalistic stimuli.

Note, that the politicians' faces were only ever presented as a changing second item. This meant that participants could have developed a strategy of reporting changes when a famous face was present. However, the purpose of this study was to test a brief version of the change detection test that would be appealing for market research use (e.g., time is a significant cost when needing to test thousands of participants). Therefore, not all conditions were tested. Furthermore, participants would have to view each of the faces as being familiar, if the strategy was to work. And of course, this is the very thing I wish to measure: Do participants view each of these faces as being famous?

Procedure

The procedure was the same as Chapter 2 with the following exceptions. The participants' task was to report whether a change had occurred and its location. Only trials for which both responses were correct were included in the analysis.

A test session consisted of 144 trials presented in a pseudo-random order. The faces to be displayed as the second item in the change (eight faces

in all) were displayed in eight trials on the left and eight trials on the right (128 trials). A further 16 trials contained no changes at all.

The experiment used only two change conditions: a non-famous face changed to another non-famous face (NN baseline) and a non-famous face changed to a famous political face (NF).

Once the experimental test session was complete the experimenter asked a series of questions. These questions asked, "Who of the three main party leaders has the most distinctive face?" "Who of the three do you like the most in terms of personality (not party politics)?" and "Who do you intend to vote for?"

The experiment took twenty minutes to complete.

Results

Change detection

On the catch trials, where there was no change, participants were on average 93% correct (7% errors).

Only trials where participants had correctly responded that there was a change and had successfully identified its location were included in the analysis of the accuracy data. A repeated measures ANOVA was conducted with two factors (location of change: left/right, and face used: Tony Blair, William Hague, Charles Kennedy, and the non-famous baseline). There was a significant effect of the location of a change ($F(1,9) = 15.9 < .01$). Inspection of the means revealed a 23% advantage for detecting changes in the left visual field (left = 84%, right = 61%).

There was also a main effect for the factor of face that changed ($F(3, 9) = 10.68 < .05$). Figure 27 shows that the changes that were least detected were the non-famous baseline changes (mean = 64%), while all three politicians faces produced higher accuracy (Charles Kennedy received the highest accuracy scores with a mean of 80% correct). Bonferroni post-hoc tests were calculated in order to establish where the significant difference occurred. Both

the faces of Charles Kennedy and Tony Blair were revealed to be significantly different than the non-famous baseline. However, the face of William Hague was not significantly different.

Figure 27

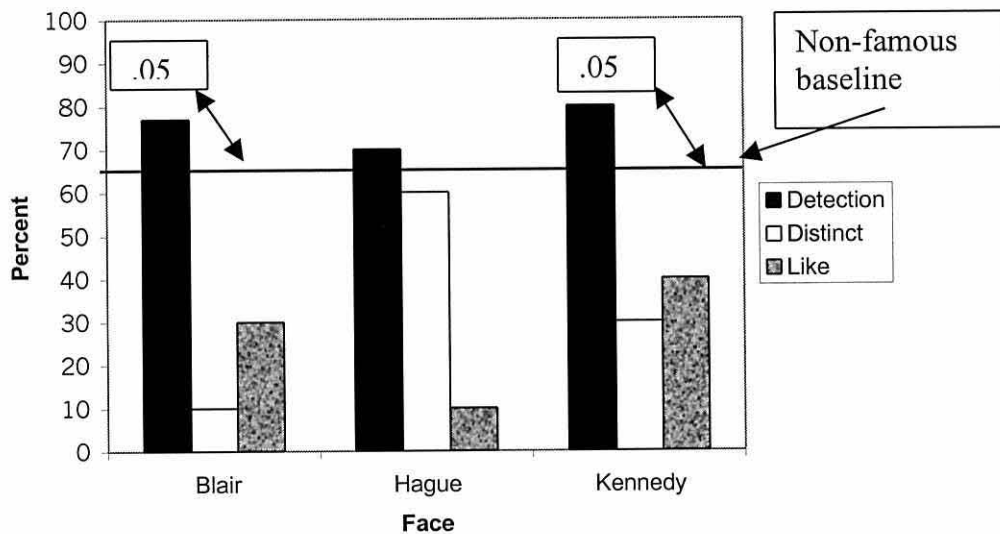


Figure 27 indicates the % correct performance on change detection/location task (black bars), the % of participants reporting the most distinctive face (white bars), and the % of participants reporting the most liked face (grey bars). The horizontal bar indicates the non-famous baseline on the change detection task, only Tony Blair and Charles Kennedy were significantly different.

In order to ensure that the lower accuracy in the non-famous baseline condition was not due to an outlier in the selection of the non-famous faces (e.g., one of the faces is so hard to detect it lowers the overall mean, when the other faces were actually similar in performance to the politicians faces) confidence limits were calculated. The mean scores of the five non-famous faces were all within the upper and lower confidence limits (see Figure 28). The mean score for the face of William Hague also fell within the two confidence limits, but the faces of Tony Blair and Charles Kennedy were above the upper confidence limit. This indicates that all five non-famous faces

led to a similar change detection performance, but the faces of Charles Kennedy and Tony Blair were better detected than this non-famous baseline.

There was a non-significant interaction of side of change and face.

Figure 28




Figure 28 show the upper 95% and lower 95% confidence limits (dashed lines). All of the means for the non-famous faces fall between the confidence limits (solid black squares), while two of the political faces have means above the upper confidence limit (black squares with white cross).

Distinctiveness, Liking, and Voting Intention.

The majority of the participants indicated that they did not intend to vote in the general election, and therefore responses to this question will not be considered further. When asked to indicate who had the most distinctive face out of the three leaders, 60% chose William Hague, 30% chose Charles Kennedy, and 10% chose Tony Blair. When asked who they liked the most, 40% indicated Charles Kennedy, 30% for Tony Blair, and 10% for William Hague. For a comparison of the responses, view Figure 27.

Discussion

An abridged version of the two-item change detection paradigm was used to assess whether individual faces could be compared to produce a measure of fame in certain target faces. Using only the critical change detection conditions of NF and NN (were the visual fame effect had been established as being largest) the three leaders of the main British political parties were evaluated against a non-famous baseline. Participants had to judge whether a change had occurred and if it had did it appear on the left or right of fixation. The results revealed that only two of the political leaders' faces were significantly better detected in a change than the non-famous faces.

There was a main effect of side of change with detection being superior when in the left visual field (23% advantage). This is consistent with the results from the change detection tasks of Chapter 3. Although there was no overall main effect of visual field in the original face study the analysis of the quartile data (the effect of experience with the task) revealed that in the first quartile (the first 96 trials) there was a significant advantage for detecting changes in the left visual field (13%). The study here only involved 144 trials, rather than 384 in the face experiments of Chapter 3. It is likely that the effect of visual field was due to the same superior processing by the right hemisphere, prior to the left hemisphere's improvement with repeated exposure to the faces.

The key finding from this experiment was that the different faces involved revealed significant differences. The ability to detect changes involving the face of William Hague (the leader of the Conservative party) did not differ significantly from that measured using the non-famous faces. However, changes involving both the face of Tony Blair (Labour party) and Charles Kennedy (Liberal Democrats) were detected significantly more than the baseline. Furthermore, confidence intervals were calculated and plotted against each of the famous and non-famous critical faces. This confirmed that

the faces used in the non-famous baseline were detected to a similar extent. Again the data for Hague's face fell within the two confidence limits along with the non-famous faces, while Blair and Kennedy's data points fell above the upper confidence limit. This data supports the potential application of the change detection paradigm as a measurement tool for market researchers.

Another interesting outcome of this study was the reports by the participants of their own perceptions about the three political leaders. When asked who they considered to have the most distinctive face the majority of respondents believed Hague to have the most distinctive face. This seems counterintuitive, as it would be expected that the face considered to be most distinctive would be the easiest to detect in a change. This was clearly not the case as the detection rates for Hague were lower than the other two political faces. Again, this helps to rule out the possibility that perceptual distinctness produced the 'fame' effect; as the most distinct famous face was the least 'famous' on this change detection task.

More in line with the change detection outcomes were the participants' reports of which politician they liked the most. Here the most liked politician (Kennedy) was also the face most detected in a change, and likewise the least liked politician (Hague) was the face least detected in a change. Although this pattern of results cannot attribute the outcome of one measure to the outcome of another measure, it does suggest that factors (e.g., affect) other than amount of experience (e.g. the face of Hague on television was as common as Kennedy's) and distinctiveness may lead to perceptual saliency and a visual fame effect. Perhaps participants are quicker to process faces that promote positive/negative affects than neutral affects. Stephens (1988) reported that although affect-eliciting faces did not produce improved memory when compared to neutral stimuli, they did increase the likelihood that a participant would claim a face to be familiar (induced familiarity). Whittlesea (1993) has argued that 'feelings of familiarity' where we have not actually encountered the stimuli before are due to fluent processing that is attributed

to a source in the past. It is possible that stimuli that elicit strong affective responses are processed rapidly and therefore produce a feeling of familiarity. If this is the case, then the political faces that elicit positive responses may be more accurately detected in the change detection paradigm than neutrally viewed faces as they benefit from speeded processing from two sources 1) the positive affect, and 2) genuine previous experience.

So, how did these politicians actually fair in the 2001 General Election? It was hard going for William Hague, the leader of the Conservative party. The party failed to recover from the huge 1997 losses to the Labour party, and only managed to gain one seat (+1%). The day after the election results William Hague announced resignation as leader of the party ⁷ The results were relatively satisfying for Tony Blair with only six losses from a previous total of 419 out of 659 seats (-1.5%). However, Charles Kennedy's first general election as leader of the Liberal Democrats saw a gain of six seats (+13%) taking the party to its highest number of parliamentary seats.

Although the study employed here used only a small sample of participants, the fates of the political leaders in the 2001 general election seemed to fit well with the fame measurement of the change detection task. With wider sampling of the general population this technique could be employed to give an implicit and objective measurement of people's perceptions of famous faces.

Further testing and refinements of the change detection paradigm may in future provide a useful evaluation tool for assessing the fame of a whole array of objects classes, which could benefit disciplines outside of cognitive psychology such as market research.

⁷ William Hague's replacement was Iain Duncan Smith. His fate could be a bleak as Hague's if appearances are anything to judge by. The 'Guardian' newspaper (August 18, 2001) reported that "After weeks of quietly submitting to unflattering analysis of his meagre scalp ...[he]... snapped back at comparisons with the virtually bald William Hague."

Chapter 7 - General Discussion

The previous chapters assessed the visual processing benefits of viewing famous items over the viewing of recently learned items using a change- detection task. The task required participants to detect which of two laterally displayed images changed from one presentation to the next. Three main categories of image were assessed: faces, landmarks, and consumer products. All produced processing benefits for famous items (see Table 3 for summary of findings).

Object Class	Manipulation	Outcome:	
		Fame in change advantage	Fame in distractor advantage
Faces:			
Exp. 1	2 item change detection (upright)	Yes	No
Exp. 2	2 item change detection (inverted)	No	No
Exp. 3	2 item change detection with biography study (upright)	Yes	Yes
Landmarks:			
Exp. 4	2 item change detection (upright)	Yes	Yes
Products:			
Exp. 5	2 item change detection (within-product type)	No (reverse effect)	No
Exp. 6	2 item change detection (between-product type)	No	Yes

Interestingly, the different categories of image produced variations in how famous images benefited task performance. Both face and landmark images produced an advantage for locating changes when the second of the changing items was famous. Landmark stimuli also produced an advantage for locating changes when the distractor (non-changing image) was famous. This also occurred for product stimuli, but faces only produced this effect when participants had undergone a study period prior to the experimental test.

Figure 29




Figure 29. Group mean percent correct change location identification for each of the object types (faces: Experiment 1, landmarks: Experiment 4, and products: Experiment 6) as a consequence of whether the second of the changing images was famous (black bars) or non-famous (white bars). Vertical bars indicate ± 1 s.e.

To elaborate on this distinction, I refer to the results showing an advantage in detecting changes involving a famous item in the change compared to the advantage for change detection when the non-changing image was famous (i.e., landmarks and products). The face experiments 1 and

3 (Chapter 3) revealed that fame aided change detection when involved in the actual change (albeit only when presented in the left visual field). This suggests that participants were abler to compare the first and second image in conceptual short-term memory. However, this may have arisen from an attentional effect, whereby a famous item, in the second of the changing images, was able to survive attentional blink effects more than non-famous images. There is even support from the landmark experiment (Exp 4) that famous second images aided change detection. Figure 29 shows the average effect of presenting a famous item (NF and FF) versus a non-famous item (NN and FN) as the second of the changing items, when the distractor (non-changing image) was non-famous. Here, all three object types exhibit an increase in performance when the second image was famous. (Note, however, that the product means have high standard errors). A mixed factorial ANOVA was calculated with one between-subject factor (Object type: faces, landmarks, and products) and two within-subject factors (Side of change: left, and right; second image in change: famous and non-famous). This revealed a significant main effect of object, $F(1, 61) = 3.32, p < .05$, with Bonferroni corrections indicating that changes involving faces were located significantly more than products. There was also a significant main effect of change, $F(1, 61) = 9.1, p < .01$, with changes involving a famous item as the second of the changing images, being detected better than non-famous second images. The main effect of side of change was non-significant, and all interactions were non-significant.

The fame advantage for landmark and product stimuli was also evident in the non-changing distractor condition. Figure 30 shows the effect of a famous versus non-famous distractor when the changing items were non-famous (NN). Here we see an advantage only for landmarks and products. One explanation might be that a speed of processing advantage for famous distractors allowed inference of where the change had occurred.

Figure 30

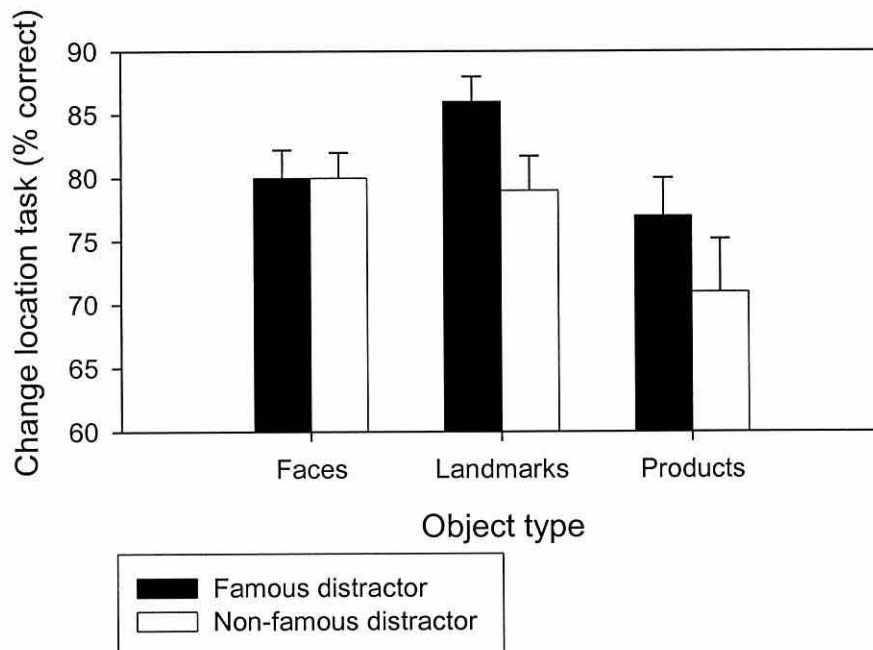


Figure 30. Group mean percent correct change location identification for each of the object types (faces, landmarks, and products) as a consequence of whether the distractor was famous (black bars) or non-famous (white bars) on NN change trials. Vertical bars indicate ± 1 s.e.

Four important ideas emerge from the study of these fame effects: 1) fame produces general encoding advantages; 2) more attentional resources are required to process non-famous items; 3) change detection ability is achieved through successful structural representation in conceptual short-term memory (CSTM); and 4) expert object recognition (i.e., faces) enhances these elements of fame.

Memory differences in visual fame effects.

Although there were processing advantages in the change detection task for all three types of object class, the pattern of these effects varied. The face experiments were the only object class to elicit a clear fame effect that required a famous image to be displayed as the changing item (NF trials versus NN trials), although landmarks produced a marginally significant advantage for famous items presented in the second of the changing items (NF and FF trials combined compared to NN and FN trials). In contrast, a fame effect for landmarks and between-category products was evident when a famous item was displayed as the non-changing item. While there appears to be general encoding advantages for all types of famous images, the variation in effects can be explained by differences in memory. In particular the ability to stabilise structured information in conceptual short-term memory (CSTM) may depend on expertise with objects that leads to differences in the ability to compare one image with another (as required in the change detection task).

Conceptual short-term memory was first described by Potter (1976) when reporting that a stimulus could be identified and meaning rapidly activated within 100ms of viewing the item. However, consolidation of this stimulus could easily be disrupted by another pictorial event that closely followed it. This indicated that although the meaning of a picture was understood (established by correct identification of target pictures in a sequence of tachiscopically presented images), awareness and memory for the item decayed (established by poor recognition memory for items). This gave rise to Potter's suggestion that stimulus identification and meaning is briefly maintained in CSTM, and that CSTM differs from sensory memory (e.g., iconic), conventional short-term memory, and long-term memory in three key details (see table 6).

Table 6

- | |
|---|
| <ol style="list-style-type: none">1. The rapidity with which stimuli reach a post categorical, meaningful level of representation.2. The rapid structuring of the representations.3. The lack of awareness (or immediate forgetting) of information that is not structured or otherwise consolidated. |
|---|

Table 6. Differences between CSTM and conventional memory according to Potter (1976, 1993, 1998).

Unlike short-term memory, which accounts for information that is processed and rehearsed after a number of seconds or minutes, and iconic memory that contains a momentary trace of the exact visual details, CSTM is able to rapidly extract semantic information (including associations) about a stimulus in at least 100ms of stimulus onset. Furthermore, the structured information can be used in a variety of ways on a range of tasks. However, any information that has not been incorporated into a structural description/representation, or likewise has not been selected as a target, is likely to be forgotten, sometimes even before it has reached awareness.

How this concerns the results of the change detection experiments can be seen in Potter's (1976, 1993, 1998) description of the effects of visual masking during a sequence of image presentations. A picture is identified in 100msec, but within that time it is vulnerable to visual masking by a new visual event. However, once identification is complete, the representation is maintained in the conceptual short-term store for a few milliseconds, where the information is then consolidated into memory. The representation is no longer susceptible to visual masking, but it can be disrupted by conceptual masking (a new meaningful visual event as opposed to a nonsense visual event like a scrambled mask). A new meaningful event warrants conceptual processing which then replaces the previous item in the conceptual short-term store. Only the part of the information that is immediately organised

into a meaningful structure will persist and be transferred into traditional short-term memory.

The change detection experiment required participants to match two images each presented for 100ms with scrambled masks displayed after each image (100msec). Each of the images displayed two items, one of which would differ from the first to the second presentation, and one that would remain the same. While the task of main interest was the ability to identify whether the change occurred on the left or right, participants also reported the category of each of the two items (faces: male or female; landmarks: building or monument; products: can or bottle). The categorisation task was conducted to ensure attention was divided between the two items, but it also revealed that the ability to extract the category of image from the brief presentation was generally resistant to the visual mask, as items were correctly categorised 78% of the time (faces =83%, landmarks =75%, and products = 75%). However, depending on the time needed to encode the structural details of the stimulus, conceptual masking may have disrupted memory consolidation, resulting in a failure to be able to compare the two images. Participants may have been able to extract the information needed to register that the first image contained a building (as opposed to a monument). However, the information that could distinguish one building from another may not have been incorporated into the structural representation, before the presentation of the second building. This would result in correct categorisation, but a failure to detect that a change had taken place.

This would explain why the fame effect was evident only when the non-changing image was famous for landmarks and between-category products. Unable to detect the change due to conceptual masking replacing the first item in the conceptual short-term store, the participants were able to recognise which item was the non-changing image, as activation in the CSTM was maintained from the first image to the second. This prolonged activation from the non-changing image allowed the participants to infer that the

change occurred in the other image. The 200msec exposure of the non-changing image provided extra time in which to develop a fuller structured description. In particular this benefits the famous item, as neural activation accumulates faster than for a non-famous item (although not fast enough to benefit from the 100msec presentation in the changing image).

In contrast the change detection task with within-product type changes revealed an advantage for changes involving non-famous products (NN) in the changing image. This divergent result further supports the role of CSTM in the performance difference in visual fame effects. The finding that the non-famous advantage arose from repetition blindness (RB) for changes with famous items heavily indicates that meaningful semantic information was accessed from the 100msec presentation. If 'Coca-Cola' changing to 'Pepsi' produced RB, it was due to the category label (type) 'Cola' being successfully maintained in CSTM, while the distinction (token) between the two products was rapidly forgotten, as it failed to be incorporated into the structured representation. This is in accordance with Kanwisher's Token Individuation account, where RB occurs due to a failure in establishing that two tokens or events occurred, even though the correct type had been established. Instead, the activation from the second event is interpreted as being residual activation from the first event and hence the two events are collapsed together, with only one being registered consciously. In terms of CSTM theory, the type recognition is structured, but the token individuation information fails to be structurally represented. Therefore, it appears that both landmark and product images benefit from efficient encoding of famous over non-famous items that allows rapid access to semantic information about the stimulus. However, CSTM for these objects is limited to the amount of information that can be structurally encoded, producing a pattern of fame effects that is restricted to change detection judgements that are dependent on the longer displays of the non-changing items.

In contrast, due to the expertise people have for recognising faces, the fame effect for faces allowed a successful comparison to be made between the two changing items. Although differences between one face and another are not as large as say differences in landmarks, people are expert at distinguishing a multitude of faces and retrieving salient information about them. As Chapter 1 outlined, there are numerous studies citing evidence for the special case of face recognition compared with object recognition in general. Although faces are a homogeneous class of stimuli, we tend to class individual faces at the basic level. This necessitates that an efficient process for encoding different faces must be developed if a person is to quickly identify and distinguish each face as it is seen. One theory is that our expertise with faces arises through an ability to take advantage of configural information. Furthermore, a processing and memorial advantage is likely to be evident for faces, as more cells will be tuned to the category of faces than most other types of object. Assuming that this is the case, it is possible that the 100msec presentation of the faces in the first image of the change detection task provided an adequate amount of time to stably structure a face representation in CSTM, that could then be maintained in conventional working memory allowing comparison with the second image. Thus participants were able to benefit from the superior processing of the famous face in the changing image, which then produced a fame effect for the changed item (as opposed to the non-changing image).

The difference in the pattern of visual fame effects for faces versus other objects can be accounted for in an explanation of conceptual short-term memory, where faces provide more structured representations in CSTM than other classes of object.

General Encoding Advantages

The advantage for detecting changes better when a famous item is present in a display versus when all stimuli are relatively unfamiliar, is entirely consistent with the population-encoding hypothesis (Perret, Oram, & Ashbridge, 1998). This theory stems from single cell recordings that reveal that a larger number of cells are tuned to familiar orientations (e.g., canonical views of objects) than to less often encountered views (e.g. inverted objects). It is suggested that neuronal activity drives the behavioural output and, therefore, any output will take longer the more unusual the view. This occurs because fewer cells respond to the unusual view, leading to a slower accumulation of activity for the population of cells responsive to an object's identity. A clear extension of this view is that familiar exemplars of an object class will have more responsive cells than unfamiliar or novel exemplars. This correctly predicts that speed of processing advantages would be evident for change detection tasks involving famous (extensively experienced) images when compared against changes where only novel or recently learned stimuli are presented. The data from all object classes (faces, landmarks, and products) revealed that famous items produced advantages that can be interpreted as improvements in speed of processing.

The face inversion experiment (Chapter 3, Experiment 2) further supports the notions of the population-encoding hypothesis. When the change detection task was conducted with all faces presented upside-down there was no advantage of having a famous item in the display. According to the population-encoding hypothesis, an upright famous face would benefit from neuronal activity that is specific to the identity of that particular face, as well, as the activity from cells responsive to (upright) faces in general. An upright non-famous face only activates the 'general face' cells. Thus, the total population of cells responding to a famous face would be greater and would lead to an encoding advantage.

Table 4

Stimulus	Percentage of the population of face cells active
A) Any upright face	80%
B) Any inverted face	60%
C) Highly learnt face (upright only)	20% (excluding general face cells)
D) Unknown face/ Inverted learnt face	0% (excluding general face cells)

Table 4 lists hypothetical percentages for a population of cells responsive to face stimuli under different viewing conditions, e.g., usual or unusual face orientation, and known or unknown face.

Table 5.

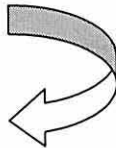

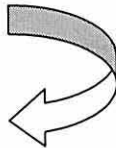

Famous face upright	$A+C=100\%$	
Non-famous face upright	$A+D=80\%$	
Famous face inverted	$B+D=60\%$	
Non-famous face	$B+D=60\%$	

Table 5. indicates the hypothetical totals for the inverted versus upright conditions with famous and non-famous faces. The white arrow indicates a significant difference between upright famous and non-famous faces, and the black arrow indicates a non-significant difference between inverted famous and non-famous faces.

However, in the case of the inverted faces, neither the famous nor non-famous face would have been sufficiently experienced to merit a population of cells that are responsive to a particular face. For example, as far as I can recall I have never viewed the ex-president of the United States, Bill Clinton from a completely inverted perspective, even though the face is well known

to me. Therefore, the population of cells that are active to both the presentation of a famous and non-famous face will be the 'inverted general face' cells (as it is an unusual view a small sub-population of cells will be active when compared to upright face orientation). In this case, there will be no advantage for viewing changes that contain famous as opposed to non-famous items. To help illustrate this point a simple arithmetic model is sketched out in Table 4 and 5.

This model shows that the culmination of cells responsive to an upright famous face is greater than those responsive to an upright non-famous face, and should produce a significant difference in performance on any task involving face perception. When these faces are presented upside down there are comparable amounts of cell activity (relative rather than absolute equality) for both famous and non-famous faces, leading to non-significant differences in processing. This was the case with the face experiments of Chapter 3. Change detection revealed a processing advantage for famous over non-famous faces in the upright conditions, but was not evident in the inverted presentation condition.

It appears that the visual processing advantages conferred by high familiarity in the change detection tasks aptly fit Perrett et al's (1998) population-encoding hypothesis. However, a more elaborate model of these encoding benefits has been described by Tong and Nakayama (1999). They used the term 'robust representation' to describe a form of visual representation that is acquired for highly over-learned stimuli (in particular faces) that are experienced in a variety of conditions and contexts. Moreover these robust representations are considered to potentially mediate optimal visual processing. Five defining properties of robust representations are suggested: 1) they mediate rapid asymptotic visual processing; 2) they require extensive visual experience; 3) they contain some abstract or view invariant information; 4) they facilitate a variety of visual and decisional processes; and 5) they demand less attentional resources.

These robust representations would reflect the endpoint of visual learning and the most extreme form of familiarity. Any stimulus that is considered to have formed a robust representation would be expected to produce an asymptotic performance on any given task involving it. This would be recognised by only a negligible improvement in task performance after further learning of the stimulus. In fact Tong and Nakayama (1999) illustrated with a visual search experiment for faces (including the participant's own face) that there were two components to the learning functions associated with newly experienced faces. At first there is a dramatic enhancement of recognition times for new faces between practice and trial (e.g., only 36 trials). This is followed by a weak (non-significant) trend of improved recognition performance across remaining test trials. This contrasts with the participants' performance when searching for their own faces, where a flat learning function occurs (i.e., the asymptotic performance of one of the defining properties of robust representations). Throughout the time course of the trials, the responses to the new faces never reach the speed of the familiar face. This pattern of results is consistent with the findings from the change detection experiments in this thesis. While, here, the responses during the practice trials were not recorded making it impossible to confirm whether rapid learning for the non-famous faces occurred, it is clear that the accuracy levels for the non-famous change trials were never as high as for famous trials. Analysis of the quartile data revealed that whereas there was a general improvement across all trials, the significant difference between N-F (non-famous to famous) and N-N (non-famous to non-famous) trials remained consistent throughout the experimental session. According to the concept of robust representations it appears that the recently learned but non-famous faces did not acquire the experience needed to produce asymptotic performance, and therefore failed to equal the performance of the robust famous faces.

The data from the change detection task involving landmarks and products produced a different pattern of visual fame effects. Although both types of stimulus revealed a visual fame effect this was mainly evident when the distractor (non-changing) image was famous. This can be contrasted with the face detection experiments where, in order to produce a fame effect, it was critical for the presence of a famous face to occur in the changing item. Furthermore, this advantage only occurred when changes were presented to the left visual field. This indicates that the processing of non-face object changes differed to that of face changes. However, across quartiles the time-course of responses did reveal that like the data from the face detection task the significant difference between famous and non-famous trials (albeit the distractor condition) remained consistent throughout the experiment. Again there is some indication that the famous items used in these change detection experiments may have formed robust representations.

Tong and Nakayama (1999) tentatively suggest that robust representations may arise from the development of an efficient visual code (a compact code). These efficient codes would reduce the redundancy of incoming visual input and thereby reduce the number of active neurones required to code a stimulus. They would also extract or explicitly code important information (features and spatial relations between features), as well as be influenced by an individual's visual experience. This is conceptualised by Tong and Nakayama as a 'reduction in the number of units or principal components needed to accurately describe a face'. Initially this appears to be in opposition to Perrett et al's (1998) population encoding hypothesis, as the compact visual code (robust representations) would require fewer rather than more cells to fire at the presentation of a famous face than to a presentation of a non-famous face. However, Perrett et al (1998) also argue that increased reaction times to unfamiliar views and objects is due to weak signals amongst a noisy background (i.e., low signal to noise ratio). On the other hand a strong selective signal from a familiar item with little noise

in the background would produce efficient coding and a fast response (i.e., high signal to noise ratio).

The visual fame effects presented within this thesis offer support for these visual encoding theories. Interestingly, previous studies of differences in familiarity and encoding have focused on reaction times (as with Perrett et al., 1998; and Tong and Nakayama, 1999). However, the change detection studies, I have presented, provide evidence that accuracy data can also be affected. The tasks involved detecting a change between the presentation of two images each displayed for only 100ms. To perform this task images would need to be compared within memory, but in the case of processing non-famous images there may have been an inadequate amount of time to encode the image to memory, unlike for famous items. Similarly, non-face objects may not have been as well stored in memory as faces, which we are all expert at recognising. In general this thesis has helped establish that famous items are processed and encoded by the visual system in a more effective way that is likely achieved through efficient neuronal coding leading to robust representations and the *visual fame effect*.

Market Research Applications

An important aspect of Brand Research is the assessment of consumers' awareness sets. The term "awareness set" refers to the products and brands a consumer can retrieve from long-term memory concerning a particular category. In general, a company will want to have its brand in a consumer's awareness set (Mowen and Minor, 2001). Traditional market research techniques adopt two approaches to brand awareness, both involve questionnaire based formats (Feldwick, 1998). The first approach is prompted awareness, where respondents simply have to recognise brand names (e.g., "Which of these brand names have you heard of?"). The second approach is spontaneous awareness, where the respondent has to recall brand names associated with a particular category (e.g., "What brands of shampoo can you

think of?”). These approaches are potentially problematic, in that, they measure explicit awareness, which can lead to response bias. For example, Feldwick (1998) reports that prompted awareness tasks often result in ceiling effects, where 90% or more of items are recognised. Clearly, an implicit test of brand awareness would be a useful addition to these traditional techniques.

Naturally, to be aware of a brand, the consumer must also have some familiarity with it. Therefore, an implicit measurement of fame could fill this gap in market research techniques. This thesis demonstrated that encoding advantages for famous objects when compared to non-famous objects in a two-item change detection paradigm led to a visual fame effect. Such advantages were present for a variety of object categories: faces, landmarks, and consumer products, indicating that this implicit technique could be tailored to solving market research questions.

Chapter 6 made a direct test of the sensitivity of the visual fame effect for individual items, by comparing the change detection accuracy for the leaders of the three main political parties with a non-famous baseline. The results from this albeit small sample revealed that two of the leaders significantly aided change detection, while one of the leaders was not significantly different to the baseline condition. This visual fame sensitivity to individual faces could provide a useful market research tool. For instance, if the three political faces had been three well-known football players, and advertisers wished to select one for a sportswear commercial the results from this technique would have eliminated at least one of the faces from the selection process.

However, caution needs to be taken in comparing famous images to non-famous images. The studies reported here relied on a small stimulus set that incorporated a certain amount of perceptual variation between images. Future studies should attempt to address these problems. For instance, increasing the number of items in a stimulus set would remove some concerns, as would the creation a database of non-famous images (e.g., faces)

that have been pre-ranked on qualities, such as brightness, contrast, size, features etc. Items from the database could then be matched with the to-be-tested famous items. In general, further work should address potential confounds and trade-offs between 'fame' and 'distinctiveness'. Providing caution is used in these respects, techniques such as change detection may prove to be a useful measure of visual fame effects.

Summary

The research contained within this thesis highlights the importance of learning and memory by investigating visual processing advantages for famous and non-famous objects. Evidence has been reported that famous stimuli are encoded more efficiently than non-famous (recently learned stimuli), supporting the contention that highly learned stimuli form robust representations that are developed over a prolonged period of time. Furthermore, it supports the population-encoding hypothesis that suggests familiar objects and views are processed faster than unfamiliar objects and views because of greater neural activation for known items. The differences that occurred in the pattern of the visual fame effect between the face stimuli and the other experiments, suggests that transfer of information into memory was more successful for faces than other objects. This is entirely consistent with the face literature showing expert face recognition in comparison to non-face object recognition.

Finally, the visual fame effect was tested in a mock commercial application that indicated the change detection technique could be employed to assist in market research that requires a measurement of the fame or familiarity of brand products. This concluded the research of this thesis, which aimed to illustrate that theories of learning and memory from experimental cognitive psychology can be married with consumer research issues to contribute across disciplines to inform applied problems.

REFERENCES

- Aguirre, G. K., & D'Esposito, M. (1997). Environmental knowledge is subserved by separable dorsal/ventral neural areas. *Journal of Neuroscience*, *17*, 2512-2518.
- Aguirre, G. K., Zarahn, E., & D'Esposito, M. (1998). Neural components of topographical representation. *Proceedings of the National Academy of Sciences*, *95*, 839-846.
- Allo-Vento, L., & Hillyard, S. A. (1996). Selective attention to the color and direction of moving stimuli: Electrophysiological correlates of hierarchical feature selection. *Perception & Psychophysics*, *58*, 191-206.
- Altarriba, J., & Soltano, E. G. (1996). Repetition blindness and bilingual memory: Token individuation for translation equivalents. *Memory & Cognition*, *24*, 700-711.
- Anderson, E., & Parkin, A. J. (1985). On the nature of the left visual field advantage for faces. *Cortex*, *21*(3), 453-459.
- Archambault, A., O'Donnell, C., & Schyns, P. G. (1999). Blind to object changes: When learning the same object at different levels of categorisation modifies its perception. *Psychological science*, *10*, 249-255.
- Arnell, K. M., & Jolicoeur, P. (1997). Repetition blindness for Pseudoobject pictures. *Journal of Experimental Psychology: Human perception and Performance*, *23*, 999-1013.
- Bartlett, J. C., & Searcy, J. (1993). Inversion and Configuration of faces. *Cognitive Psychology*, *25*, 281-316.
- Bauer, R. (1984). Autonomic recognition of names and faces: A neuropsychological application of the Guilty Knowledge Test. *Neuropsychologia*, *22*, 457-469.
- Bauer, R. (1986). The cognitive psychophysiology of prosopagnosia. In H. Ellis, M. Jeeves, F. Newcombe, A. Young, (Eds), *Aspects of face processing* (pp. 253-267). Dordrecht, The Netherlands: Martinus Nijhoff.
- Bavelier, D. (1994). Repetition blindness between visually different items: the case of pictures and words. *Cognition*, *24*, 700-711.

- Bavelier, D. (1994). Role and nature of object representations in perceiving and acting. In V. Coltheart (eds), *Fleeting Memories: Cognition of Brief visual Stimuli*, London: MIT Press.
- Bechara, A., Tranel, D., Damasio, H., & Adolphs, R. (1995). Double dissociation of conditioning and declarative knowledge relative to the amygdala and hippocampus in humans. *Science*, 269(5227), 1115-1118.
- Begleiter, H., Porjesz, B., & Wang, W. (1993). A neurophysiological correlate of visual short-term memory in humans. *Electroencephalography and clinical Neurophysiology*, 87, 46-53.
- Begleiter, H., Porjesz, B., & Wang, W. (1995). Event-related brain potentials differentiate priming and recognition to familiar and unfamiliar faces. *Electroencephalography and clinical Neurophysiology*, 94, 41-49.
- Bentin, S., Allison, T., Perez, E., Puce, A., & McCarthy, G. (1986). Electrophysiological studies of face perception in humans. *Journal of cognitive neuroscience*, 8, 551-565.
- Benton, A. L., & Van Allen, M. W. (1968). Impairment in facial recognition in patients with cerebral disease. *Cortex*, 4, 344-358.
- Benton, A. L., & Van Allen, M. W. (1972). Prosopagnosia and facial discrimination. *Journal of the Neurological Sciences*, 15, 167-172.
- Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, 94, 115-147.
- Biederman, I., & Cooper, E. E. (1992). Scale invariance in visual object priming. *Journal of Experimental Psychology: Human Perception and Performance*, 18, 121-133.
- Biederman, I., Mezzanotte, R. J., & Rabinowitz, J. C. (1982). Scene perception: Detecting and judging objects undergoing relational violations. *Cognitive Psychology*, 14, 143-177.
- Boyce, S. J., & Pollastek, A. (1992). An exploration of the effects of scene context on object identification. In K. Raynor's *Eye Movements and Visual Cognition: Scene Perception and Reading*, 227-242.
- Bredart, S., Valentine, T., Calder, A., & Gassi, L. (1995). An interactive activation model of face naming. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology*, 48(2), 466-486

Breen, N., Caine, D., & Coltheart, M. (2000). Models of face recognition and delusional misidentification: A critical review. *Cognitive Neuropsychology*, 17, 55-71.

Brennen, T., David, D., Fluchaire, I., & Pellat, J. (1996). Naming faces and objects without comprehension: A case study. *Cognitive Neuropsychology*, 13(1), 93-110

Brownell, H. H. (1978). *Picture perception and semantic memory*. Unpublished doctoral dissertation, John Hopkins University.

Bruce, V., & Young, A. (1986). Understanding face recognition. *British Journal of Psychology*, 77, 305-327.

Bulthoff, H. H., & Edelman, S. (1992). Psychophysical support for a two-dimensional view interpolation theory of object recognition. *Proceedings of the National Academy of Sciences of the USA*, 89, 60-64.

Carey, S., & Diamond, R. (1977). From piecemeal to configurational representation of faces. *Science*, 195, 312-314.

Cherry, E. C. (1953). Some experiments on the recognition of speech, with one and with two ears. *Journal of the Acoustical Society of America*, 25, 433-458.

Chun, M. (1997). Types and tokens in visual processing: A double dissociation between the attentional blink and repetition blindness. *Journal of Experimental Psychology: Human Perception and Performance*, 23, 738-755.

Compuserve news (2001). Kennedy faces first election as leader. Retrieved on 7th March 2002 from <http://www.copuserve.co.uk/channels/news/election/kennedyprofile.htm>.

Corballis, M. C. (1988). Recognition of disoriented shapes. *Psychological Review*, 95, 115-123.

Davidoff, J & Donnelly, N (1990) Object superiority: A comparison of complete and part probes. *Acta Psychologica*, 73, 225-243.

De Caro, S. A., & Reeves, A. (1995, November). *Object Recognition without mental rotation*. Paper presented at the annual meeting of the Psychonomic Society, Los Angeles.

DeGraef, P. (1992). Scene context effects and models of real-world perception. In K. Raynor's *Eye Movements and Visual Cognition: Scene Perception and Reading*, 243-259.

De Renzi, E. (1996). Prosopagnosia in two patients with CT-scan evidence of damage confined to the right hemisphere. *Neuropsychologia*, 24, 385-389.

De Renzi, & Spinnler, H. (1966). Impaired performance on color tasks in patients with hemispheric damage. *Cortex*, 3(2), 194-217.

Desimone, R. (1991). Face-selective cells in the temporal cortex of monkeys. *Journal of Cognitive Neuroscience*, 3, 1-8.

Diamond, R., & Carey, S. (1986). Why faces are not special: An effect of expertise. *Journal of experimental psychology: General*, 115, 107-117.

Dubois, S., Rossion, C., Schiltz, C., Bodart, J. M., Michel, C., Bruyer, R., & Crommelink, M. (1999). Effect of familiarity on the processing of human faces. *Neuroimage*, 9, 278-289.

Durso, F. T., & Johnson, M. K. (1979). Facilitation in naming and categorizing repeated pictures and words. *Journal of Experimental Psychology: Human Learning & Memory*, 5, 449-459.

Ellis, H. D., Sheperd, J. W., & Davies, G. M. (1979). Identification of familiar and unfamiliar faces from internal and external features: Source implications for theories of face recognition. *Perception*, 8, 431-439.

Ellis, H. D., & Young, A. (1990). Accounting for delusional misidentifications. *British Journal of Psychiatry*, 157, 239-248.

Ellis, H. D., Young, A. W., Quayle, A. H., & de Pauw, K. W. (1997). Reduced autonomic responses to faces in Capgras delusion. *Proceedings of the Royal Society, London(B)*, 264, 1085-1092.

Enns, J. T., & Di Lollo, V. (1997). Object substitution: A new form of masking in unattended visual locations. *Psychological Science*, 8, 135-139.

Epstein, R., Harris, A., Stanley, D., & Kanwisher, N. (1999). The parahippocampal place area: Recognition, Navigation, or encoding? *Neuron*, 23, 115-125.

Epstein, R., & Kanwisher, N. (1998). A cortical representation of the local visual environment. *Nature*, 392, 598-601.

- Farah, M. J. (1990). Selective impairments of object representation in prosopagnosia. *Journal of Clinical Neuropsychology*, 12, 32-32.
- Farah, M. J., Hammond, K. M. (1988). Mental rotation and orientation-invariant object recognition: dissociable processes. *Cognition*, 29, 29-46.
- Farah, M. J., Tanaka, J. R., & Drain, H. M. (1995). What causes the face inversion effect? *Journal of Experimental psychology: Human Perception and Performance*, 21, 628-634.
- Farah, M. J., Wilson, K. D., Drain, H. M., & Tanaka, J. R. (1995). The inverted face inversion effect in prosopagnosia: Evidence for mandatory, face specific perceptual mechanisms. *Vision Research*, 35, 2089-2093.
- Farah, M. J., Wilson, K. D., Drain, M., & Tanaka, J. N. (1998). What is "Special" about face perception? *Psychological Review*, 105, 482-498.
- Feldwick, P. (1998). Brand Research. In C. MacDonald and P. Vangelder (eds), *ESOMAR Handbook of market and Opinion Research* (pp 632-635). Amsterdam, ESOMAR.
- Foldiak, P. (1991). Learning invariance from transformational sequences. *Neural Computation*, 3, 194-200
- 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.
- Gauthier, I., Anderson, A. W., Tarr, M.J., Skudlarski, P., & Gore, J., C. (1997). Levels of categorisations in visual object studied with functional MRI. *Current Biology*, 7, 645-651.
- Gauthier, I., Behrmann, M., & Tarr, M. J. (1999). Can face recognition really be dissociated from object recognition? *Journal of cognitive neuroscience*, 11, 349-370.
- Gauthier, I., & Logothetis, N. K. (2000). Is face recognition not so unique after all? *Cognitive neuropsychology*, 17, 125-142.
- Gauthier, I., Skudlarski, P., Gore, J. C., & Anderson, A. W. (2000). Expertise for cars and birds recruits brain areas involved in face recognition. *Nature neuroscience*, 3, 191-197.

Gauthier, I., & Tarr, M. J. (1997). Becoming a "greeble" expert: Exploring mechanisms for face recognition. *Vision Research*, 37(12), 1673-1682.

Gauthier, I., & Tarr, M. J. (inprint). Unraveling neural mechanisms for expert object recognition: Bridging brain activity and behavior. *Journal of Experimental Psychology: Human performance and Perception*.

Gauthier, I., Tarr, M. J., Moylan, J., Anderson, J., Skudlarski, P., & Gore, J. C. (2000). Does visual subordinate-level categorisation engage the functionally-defined fusiform face area? *Cognitive neuropsychology*, 17, 143-163.

Gilbert, C., & Bakan, P. (1973). Visual asymmetry in the perception of faces. *Neuropsychologia*, 11, 355-362.

Hecaen, H., & Angelergues, R. (1962). Agnosia for faces (Prosopagnosia). *Archives of Neurology (Chicago)*, 7(2), pp. 92-100

Hecaen, H., Tzortzis, C., Rondot, P. (1980). Loss of topographic memory with learning deficits. *Cortex*, 16, 525-542.

Haig, N. D. (1986). Exploring recognition with interchanged facial features. *Perception*, 15, 25-247.

Halgren, E., Dale, A. M., Sereno, M. I., Tootell, R. B. H., Marinkovic, K., & Rosen, B. (1999). Location of human face-selective cortex with respect to retinotopic areas. *Human Brain Mapping*, 7, 29-37.

Hanley, J. R. (1995). Are names difficult to recall because they are unique? A case study of a patient with anomia. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology*. 48A(2), 487-506.

Hanley, J. R., Pearson, N., & Howard, L. A. (1990). The effects of different types of encoding task on memory for famous faces and names. *The Quarterly Journal of Experimental psychology*, 42a, 741-762.

Haxby, J. V., Horowitz, B., Ungerleider, L. G., Maisog, J. M., Pietrini, P., & Grady, C. L. (1994). The functional organization of human extrastriate cortex: A PET-rCBF study of selective attention to faces and locations. *Journal of Neuroscience*, 14, 6336-6353.

Henderson, J. M. (1997). Transaccadic memory and integration during real-world object perception. *Psychological Science*, 8,(1), 51-55.

Henderson, J. M., & Hollingworth, A. (1999a). High-level scene perception. *Annual Review of Psychology*, 50, 243-271.

Henderson, J. M., & Hollingworth, A. (1999b). The role of fixation position in detecting scene changes across saccades. *In press: Psychological Science*.

Henderson, J. M., Weeks, P. A., & Hollingworth, A. (1999). The effects of semantic consistency on eye movements during complex scene viewing. *Journal of Experimental Psychology: Human Perception and Performance*, 25, 210-228.

Hillger, L. A., & Koenig, O. (1991). Separable mechanisms in face processing: Evidence from hemispheric specialization. *Journal of Cognitive Neuroscience*, 3, 42-58.

Hirstein, W., & Ramachandran, V. S. (1997). Capgras syndrome: A novel probe for understanding the neural representation of the identity and familiarity of persons. *Proceedings of the Royal Society, London(B)*, 264, 437-444.

Hodges, J. R., & Greene, J. D. W. (1998). Knowing about people and naming them: Can Alzheimer's disease patients do one without the other? *Quarterly Journal of Experimental Psychology: Human Experimental Psychology*, 51A(1), 121-134.

Hollingworth, A., & Henderson, J. M. (1998a). Does consistent scene context facilitate object perception? *Journal of Experimental Psychology: General*, 127 (4), 398-415.

Hollingworth, A., & Henderson, J. M. (1998b). The role of eye movements in detecting changes to scenes in the flicker paradigm. *Poster presented at the Annual meeting of the Psychonomic Society, Dallas, November 1998*.

Hollingworth, A., & Henderson, J. M. (1999). Object identification is isolated from scene semantic constraint: Evidence from object type and token discrimination. *Acta Psychologica: Special Issue on Object perception and Memory*, in press.

Ivry, R. B. & Robertson, L. C. (1998). *The Two sides of Perception*. MIT Press: Cambridge, USA.

Johnson, K. A., & Becker, J. A. *The Whole Brain Atlas: Version 1*.

Johnston, R. A., & Bruce, V. (1990). Lost properties? Retrieval differences between name codes and semantic codes for famous people. *Psychological Research*, 52, 62-67.

Jolicoeur, P. (1990). Identification of disoriented objects: A dual systems theory. *Mind & Language*, 5, 387-410.

Jolicoeur, P., Gluck, M., & Kosslyn, S. M. (1984). Pictures and names: Making the connection. *Cognitive Psychology*, 16, 243-275.

Kanwisher, N. (1987). Repetition blindness: Type recognition without token individuation. *Cognition*, 27, 117-143.

Kanwisher, N. (1991). Repetition blindness and illusory conjunctions: Errors in binding visual types with visual tokens. . *Journal of Experimental Psychology: Human Perception and Performance*, 17, 404-421.

Kanwisher, N., McDermott, J., Chun, M., M. (1997). The fusiform face area: A module in human extrastriate cortex specialized for face perception. *Journal of Neuroscience*, 17, 4302-4311.

Kanwisher, N., & Potter, M., (1989). Repetition blindness: The effects of stimulus modality and spatial displacement. *Memory & Cognition*, 17, 117-124.

Kanwisher, N., & Potter, M., (1990). Repetition blindness: Levels of processing. *Journal of Experimental Psychology: Human Perception and Performance*, 16, 30-47.

Kanwisher, N., & Yin, C. (1993). Phonological and visual codes in repetition blindness for words and pictures. *Investigative Ophthalmology and Visual science*, 34, 1131.

Kanwisher, N., Yin, C., & Wojciulik, E. (1999). Repetition blindness for pictures: Evidence for the rapid computation of abstract pictorial descriptions. In V. Coltheart (eds), *Fleeting Memories: Cognition of Brief visual Stimuli*, London: MIT Press.

Kendrick, K. M., Atkins, K., Hinton, M. R., Heavens, P., & Keverne, B. (1996). Are faces special for sheep? Evidence from facial and object discrimination learning tests showing effects of inversion and social familiarity. *Behavioural Processes*, 38, 19-35.

Kitterle, F. L., Christman, S., & Hellige, J. B. (1990). Hemispheric differences are found in the identification, but not the detection of low versus high spatial frequencies. *Perception and Psychophysics*, 48, 297-306.

Kutus, M., & Hillyard, S. A. (1980). Reading senseless sentences: Brain potentials reflect semantic incongruity. *Science*, 207, 203-205.

Landis, T., Cummings, J. L., Benson, D. F., & Palmer, E. P. (1986). Loss of topographic familiarity – An environmental agnosia. *Journal of Clinical Experimental Neuropsychology*, 43, 132-136.

Lawson, R., & Humphreys, G. W. (1996). View-specificity in object processing: Evidence from picture matching. *Journal of Experimental Psychology: Human Perception and Performance*, 22, 395-416.

Lawson, R., & Humphreys, G. W., Watson, D. (1994). Object recognition under sequential viewing conditions: Evidence for viewpoint-specific recognition procedures. *Perception*, 23, 595-614.

Lawson, R., & Jolicouer, P. (1998). The effects of plane rotation on the recognition of brief masked pictures of familiar objects. *Memory & Cognition*, 26, 791-803.

Leder, H., & Bruce, V. (2000). When inverted faces are recognised: The role of configural information in face recognition. *The Quarterly journal of experimental psychology*, 53a, 513-536.

Leechey, S., Carey, S., Diamond, R., & Cahn, A. (1978). Upright and inverted faces: The right hemisphere knows the difference. *Cortex*, 14, 411-419.

Leek., E. C. (1998). Effects of stimulus orientation on the identification of common polyoriented objects. *Psychonomic Bulletin and Review*, 5, 650-658.

Leveroni, C. L., Seidenberg, M., Mayer, A. R., Mead, L. A., Binder, J. R., & Rao, S. M. (2000). Neural systems underlying the recognition of familiar and newly learned faces. *The Journal of Neuroscience*, 20, 878-886.

Levine, S. C., Banich, M. T., & Koch-Weser, M. (1988). Face recognition: A general of specific right hemisphere capacity? *Brain and Cognition*, 8, 303-325.

Levy, J., Trevarthen, C., & Sperry, R. W. (1972). Perception of bilateral chimeric figures following hemispheric deconnexion. *Brain*, 95(1), 61-78.

Levy-Agresti, J., & Sperry, R. W. (1968). Differential perceptual capacities in major and minor hemispheres. *Proceedings of the National Academy of Sciences*, 61, 1151.

Lin, E. L., & Murphy, G. L. (1997). Effects of background knowledge on object categorisation and part detection. *Journal of Experimental Psychology: Human Perception and Performance*, 23, 1153-1169.

Logothetis, N. K., Pauls, J., & Poggio, T. (1995). Shape representation in the inferior temporal cortex of monkeys. *Current Biology*, 5, 552-563

Luh, Karen E. (1998). Effect of inversion on perceptual biases for chimeric faces. *Brain and Cognition*, 37, 105-108.

MacKay, D. G., & Miller, M. D. (1994). Semantic blindness: Repeated concepts are difficult to encode and recall under time pressure. *Psychological Science*, 5, 52-55.

McCallum, W. C., Farmer, S. F., & Peacock, P. V. (1984). The effects of physical and semantic incongruities on auditory event-related potentials. *Electroencephalography and clinical Neurophysiology*, 59, 477-488.

McCarthy, R. A., Evans, J. J., Hodges, J. R. (1996). Topographic amnesia: Spatial memory disorder, perceptual dysfunction, or category specific semantic memory impairment? *Journal of Neurology and Neurosurgery*, 60, 318-325.

McCarthy, R. A., Puce, A., belger, A., & Allison, T. (1999). Electrophysiological studies of human face perception: Response properties of face-specific potentials generated in occipito-temporal cortex. *Cerebral Cortex*, 9, 431-444.

McIntosh, A. R., Grady, C. L., Haxby, J. V., & Ungerleider, L. G. (1996). Changes in limbic and prefrontal functional interactions in a working memory task for faces. *Cerebral Cortex*, 6(4), 571-584.

McNeil, J. E., & Warrington, E. K. (1991). Prosopagnosia: a reclassification. *Quarterly Journal of Experimental Psychology*, 43a, 267-287.

Malcus, L., Klatsky, G., Bennet, R. J., Genarelli, E., & Biederman, I. (1983). Focused attention in scene perception: Evidence against automatic processing of scene backgrounds. Paper presented at the 54th Annual meeting of the Eastern psychological Association, Philadelphia.

Malone, D. R., Morris, H. H., Kay, M. C., & Levin, H. S. (1982). Prosopagnosia: a double dissociation between the recognition of familiar and unfamiliar faces. *Journal of Neurology, Neurosurgery, and Psychiatry*, 45, 820-822.

Markham, E. M., & Callanan, M. A. (1984). An analysis of hierarchical classification. In R. Sternberg (Ed.), *Advances in the psychology of human intelligence* (Vol.2, pp. 325-365). Hillsdale, NJ: Erlbaum.

Milner, B. (1968). Visual recognition and recall after right temporal-lobe excision in man. *Neuropsychologia*, 6(3), 191-209.

Miyashita, Y., & Chang, H. S., 1988. Neuronal correlate of pictorial short-term memory in the primate temporal cortex. *Nature*, 331, 68-70.

Moscovitch, M., Wincour, G., & Behrmann, M. (1997). What is special about face recognition? Nineteen experiments on a person with visual object agnosia, and dyslexia but normal face recognition. *Journal of Cognitive Neuroscience*, 9, 555-604.

Moscovitch, M., & Moscovitch, D. A. (2000). Super face-inversion effects for isolated internal or external features, and for fractured faces. *Cognitive neuropsychology*, 17, 201-219.

Mowen, J. C., & Minor, M. S. (2001). *Consumer Behavior: A Framework*. New Jersey, Prentice Hall.

Murphy, G. L., & Brownell, H. H. (1985). Category differentiation in object recognition: Typicality constraints on the basic category advantage. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 11,70-84.

Murphy, G. L., & Wisniewski, E. J. (1989). Categorizing objects in isolation and in scenes: What a superordinate is good for. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 15, 572-586.

Murray, J. E., Yong, E., & Rhodes, G. (2000). Revisiting the perception of upside-down faces. *Psychological Science*, 11, 492-496.

Nachson, I., Moscovitch, M., & Umilita, C. (1995). The contribution of external and internal features to the matching of unfamiliar faces. *Psychological Research*, 58, 31-37.

Nakayama, K., Shimojo, S. & Silverman, G. H. (1989). Stereoscopic depth: Its relation to image segmentation, grouping, and the recognition of occluded objects. *Perception*, 18, 55-68.

Newcombe, F. (1969). *Missile wounds of the brain: A study of psychological deficits*. NY, Oxford U. Press. vi, pp.145.

O'Craven, K. M., & Kanwisher, N. (2000). Mental Imagery of faces and places activates corresponding stimulus-specific brain regions. *Journal of Cognitive Neuroscience*, 12, 1013-1023.

Oram, M. W., & Perrett, D. I. (1992). Time course of neural responses discriminating different views of the face and head. *Journal of Neurophysiology*, 68, 70-84.

Pallis, C. A. (1955). Impaired identification of faces and places with agnosia for colours. *Journal of Neurology, Neurosurgery, and Psychiatry*, 18, 218-224.

Palmer S. E. (1975). The effects of contextual scenes on the identification of objects. *Memory & Cognition*, 3, (5), 519-526.

Palmer, S.E., Rosch, E., & Chase, P. (1981). Canonical perspective and the perception of objects. In J. Long & A. D. Baddeley (Eds), *Attention and Performance*, IX. Hillsdale, N.J.: Lawrence Erlbaum Associates.

Parasuraman, R., & Martin, A. (2001). Interaction of semantic and perceptual processes in repetition blindness. *Visual Cognition*, 8, 103-118.

Pashler H. (1988). Familiarity and visual change detection. *Perception and Psychophysics*, 44(4), 369-378.

Perrett, D. I., Harries, M. H., Bevan, R., Hietanen, J. K., Benson, P. J., & Thomas, S. (1991). Viewer-centred and object-centred coding of heads in the macaque temporal cortex. *Experimental Brain Research*, 86, 159-173.

Perrett, D. I., Smith, P. A. J., Potter, D. D., Mistlin, A. J., Head, A. S., Milner, A., & D., Jeeves, M. A. (1984). Neurons responsive to faces in the temporal cortex: Studies of functional organization, sensitivity to identity and relation to perception. *Human Neurobiology*, 3, 197-208.

Perrett, D. I., Mistlin, A. J., Chitty, A. J., Harries, M., Newcombe, F., & De Haan, E. (1988). Neuronal mechanisms of face perception and their pathology. In C. Kennard, F. Clifford Rose (Eds), *Physiological Aspects of Clinical Neuro-ophthalmology*. Chapman and Hall, London, pp. 137-154.

Perrett, D. I., Oram, M. w., & Ashbridge, E. (1998). Evidence accumulation in cell populations responsive to faces: an account of generalisation of recognition without mental transformations. *Cognition*, 67, 111-145.

Postma, A., Izendoorn, R., De Haan, & Edward, H. F. (1998). Sex differences in object location memory. *Brain and Cognition*, 36, 334-345.

Potter, M. C. (1993). Very short-term conceptual memory. *Memory & Cognition*, 21, 156-161.

Potter, M. C. (1999). Understanding sentences and scenes: The role of conceptual short-term memory. In V. Coltheart (eds), *Fleeting Memories: Cognition of Brief visual Stimuli*, London: MIT Press.

Puce, A., Allison, T., Gore, J. C., & McCarthy, G. (1995). Face-sensitive regions in human extrastriate cortex studied by functional MRI. *Journal of Neurophysiology*, 74, 1192-1199.

Raymond, J. E., Shapiro, K. L., & Arnell, K. M. (1992) Temporary suppression of visual processing in an RSVP task: an attentional blink? *Journal of Experimental Psychology: Human Performance and Perception*, 18, 849-860.

Rensink, R. A., O'Regan, J. K., & Clark J. J. (1997). To see or not to see: The need for attention to perceive changes in scenes. *Psychological Science*, 1-6.

Rhodes, G. (1993). Configural coding, expertise, and the right hemisphere advantage for face recognition. *Brain and Cognition*, 22, 19-41.

Rhodes, G., Brake, S., & Atkinson, A. P. (1993). What's lost in inverted faces? *Cognition*, 47, 25-57.

Rhodes, G., & Tremewan, T. (1993). The Simon and Garfunkel effect: Semantic priming, sensitivity, and the modularity of face recognition. *Cognitive Psychology*, 25, 147-187.

Rhodes, G., & Wooding, R. (1989). Laterality effects in identification of caricatures and photographs of famous faces. *Brain & Cognition*, 9(2), 201-209

Ro, T., Russell, C., & Lavie, N. (2001). Changing faces: A detection advantage in the flicker paradigm. *Psychological Science*, 12, 94-99.

Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology*, 8, 382-439.

Ross, P., & Turkewitz, G. (1982). Changes in hemispheric advantage in processing facial information with increasing stimulus familiarisation. *Cortex*, 18, 489-499.

Rossion, B., Dricot, L., Devolder, A, Bodart, J., Crommelinck, M., de Gelder, B., & Zoontjes, R. (2000). Hemispheric asymmetries for whole-based and part-based face processing in the human fusiform gyrus. *Journal of Cognitive Neuroscience*. 12(5), 793-802

Ruchkin, D. S., Canoune, H. L., Johnson, R., & Ritter, W. (1995). Working memory and preparation elicit different patterns of slow wave event-related brain potentials. *Psychophysiology*, 32, 399-410.

Scapinello, K. F., & Yarmey, A. D. (1970). The role of familiarity and orientation in immediate and delayed recognition of pictorial stimuli. *Psychonomic Science*, 21, 329-330.

Schweinberger, S. R., Burton, A. M., & Kelly, S. W. (2001) Priming the access to names of famous faces. *British Journal of Psychology. Special Issue*, 92(2), 303-317

Schyns, P., & Oliva, A. (1999). Dr. Angry and Mr. Smile: When categorization flexibly modifies the perception of faces in rapid visual presentations. *Cognition*, 69(3), 243-265.

Schyns, P. G., & Rodet, L. (1997). Categorization creates functional features. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 23, 681-696.

Sergent, J. (1984). Configural processing of faces in the left and right cerebral hemispheres. *Journal of experimental psychology: Human perception and performance*, 10, 554-572.

Sergent, J. (1985). Influence of task and input factors on hemispheric involvement in face processing. *Journal of Experimental Psychology: Human Perception & Performance*, 11(6), 846-861.

Sergent, J. (1998). Influence of task and input factors on hemispheric involvement in face processing. *Journal of Experimental Psychology: Human Perception and Performance*, 11, 846-861.

Sergent, J., Ohta, S., & McDonald, B. (1992). Functional neuroanatomy of face and object processing: A positron emission tomography study. *Brain*, 115, 15-36.

Shapiro, K. L., Caldwell, J., & Sorensen, R. E. (1997). Personal names and the attentional blink: A visual "cocktail party" effect. *Journal of Experimental Psychology: Human Perception and Performance*, 23, 504-514.

Smith, E. E., Balzano, G. J., & Walker, J. (1978). Nominal, natural and semantic codes in picture categorisation. In R. L. Klatzky (Eds), *Semantic factors in cognition* (pp. 168). Hillsdale, NJ: Erlbaum.

Sinha, P. & Poggio, T. (1996). Role of learning in three-dimensional form perception. *Nature*, 384, 460-463.

Solms, M., Turnbull, O. H., Kaplan-Solms, K., & Miller P. (1998). Rotated Drawing: The Range of performance and anatomical correlates in a series of 16 patients. *Brain and Cognition*, 38, 358-368.

Srinivas K. (1995). Representation of rotated objects in explicit and implicit memory. *Journal of Experimental Psychology: Learning, Memory, and cognition*, 21, 1019-1036.

Takahashi, N., Kawamura, M., Hirayama, K., & Tagawa, K. (1989). Non-verbal facial and topographic visual object agnosia - a problem of familiarity in prosopagnosia and topographic disorientation. *No To Shinkei*, 41, 703-710.

Tanaka, J. W., & Curran, T. (2001). A neural basis for expert object recognition. *Psychological Science*, 12, 43-47.

Tanaka, J. W., & Farah, M. J. (1993). Parts and wholes in face recognition. *The Quarterly Journal of Experimental psychology*, 46a, 225-245.

Theeuwes Jan, Kramer, Arthur F., Hahn, Sowon, & Irwin, David E. (1998). Our eyes do not always go where we want them to go: Capture of the eyes by new objects. *Psychological Science*, 9, 379-385.

Thompson, P. (1980). Margaret Thatcher: A new illusion? *Perception*, 9, 483-4.

Thompson, D. M. (1986). Face recognition: More than a feeling of familiarity? In *Aspects of face processing*, (ed. H. D. Ellis, M. A. Jeeves, F. Newcombe, and A. Young), pp. 118-22. Martinus nijhoff Dordrecht.

Thorpe, S., Fize, D., & Marlot, C. (1996). Speed of processing in the human visual system. *Nature*, 381, 520-522.

Tong, F., & Nakayama, K. (1999). Robust representations for faces: Evidence from visual search. *Journal of experimental psychology: Human perception and performance*, 25, 1016-1035.

Tranel, D., & Damasio, A. R. (1985). Knowledge without awareness: an autonomic index of facial recognition by prosopagnosics. *Science*, 228, 1453-1454.

Turnbull, O. H. (1997). A double dissociation between knowledge of object identity and object orientation. *Neuropsychologia*, 35, 567-570.

Turnbull, O. H., Carey, D. P., & McCarthy, R. A. (1997). The neuropsychology of object constancy. *Journal of the International Neuropsychological Society*, 3, 288-198.

Valentine, T., & Bruce, V. (1988). Mental rotation of faces. *Memory and Cognition*, 16, 556-566.

Yantis Steven (1993). Stimulus-driven attentional capture. *Current Directions in Psychological Science*, 2, 156-161.

Yin, R. K. (1969). Looking at upside-down faces. *Journal of Experimental Psychology*, 81, 141-145.

Young, A. (1985). Familiarity decisions for faces presented to the left and right cerebral hemispheres. *Brain & Cognition*, 4(4), 439-450.

Young, A. W., Hay, D. C., McWeeney, K. H., Flude, B. M., & Ellis, A. W. (1985). Matching familiar and unfamiliar faces on internal and external features. *Perception*, 14, 737-759.

Young, A. W., Hellowell, D., & Hay, D. C. (1987). Configural information in face perception. *Perception*, 16, 747-759.

Young, A. W., Newcombe, F., de Haan, E. H. F., Small, M., & Hay, D. C. (1993). Face perception after brain injury: Selective impairments affecting identity and expression. *Brain*, 116, 941-959.

Younge, G. (2001). Little Billy aims for more altitude. *The Guardian*, may 17th. Retrieved on the 7th March 2002 from <http://www.guardian.co.uk/Columnists/Column/0,5673,492194,00.html>

Wachsmuth, E., Oram, M. W., & Perret, D. I. (1994). Recognition of objects and their component parts: Responses of single units in the temporal cortex of the macaque. *Cerebral Cortex*, 4, 509-522.

Wainwright M. (2001). Jokes wear thin for Duncan Smith. *Guardian Unlimited*, Aug 2001.

Warrington, E. K., & James, M. (1967). An experimental investigation of facial recognition in patients with unilateral cerebral lesions. *Cortex*, 3, 317-326.

Whiteley, A. M., & Warrington, E. K. (1978). Selective impairment of topographical memory: a single case study. *Journal of Neurology, Neurosurgery, and Psychiatry*, 41, 575-578.

Wiser, A. K., Andreasen, N. C., O'Leary, D.S., Crespo-Facorro, B., Boles-Ponto, L. L., Watkins, G. L., & Hichwa, R. D. (2000). Novel vs. well-learned memory for faces: A positron emission tomography study. *Journal of Cognitive Neuroscience*, 12, 255-266.

Werner, S., & Thies, B. (2000). Is "changeblindness" attenuated by domain-specific expertise? An expert-novices comparison of change detection in football images. *Visual cognition*, 7, 163-174.

Whittlesea, B. (1993). Illusions of familiarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 19, 1235-1253.

Zebrowitz, L. A. (1994). Facial maturity and political prospects: Persuasive, culpable, and powerful faces. In Schank, R. C., & Langer, E. (Eds) *Beliefs, reasoning and decision-making*. (pp 315-345). Hillsdale, NJ, US: Lawrence Erlbaum Associates.


Zebrowitz, L. A. (1997). *Reading Faces*. New York: The Perseus Books Group.

APPENDIX


Face stimuli used in Chapter 3.

A and B indicate the first and second set of faces used. Each set had four famous (top row of each set) and four non-famous faces (bottom row of each set).

A



B



Biography Information used in Experiment 3.

LEO

This is Leonardo DiCaprio. He is a professional actor who has appeared in many Hollywood films. He especially shot to fame when he starred alongside Kate Winslet in the film 'Titanic'. His latest release is 'The Beach', though the critics are still undecided about its success. This won't stop the popularity of the 'moody' star from being a popular choice in male pin-up.

CHARLES

This is Prince Charles. He is currently in line to inherit the British throne from his mother Queen Elizabeth II. Naturally, the media spotlight has shone on him to a considerable extent. He has been noted for his strong opinions on architecture, including his dislike for the Millennium Dome, and it has been reported he talks to his plants. However, readers of the 'Big Issue' recently voted him the people's voice.

JENNIFER

Jennifer Aniston is an actress who shot to fame in the extremely popular sitcom 'Friends'. This experience has led her into the world of Hollywood films and a large amount of media attention. Advertisers have taken advantage of her popularity and looks to boost Loreal's shampoo sales, with appearances and phrases such "...because I'm worth it!" Currently, she is known to be dating Hollywood hunk Brad Pitt.

DIANA

The late Princess Diana has been seen to be an icon for millions worldwide. She became a Princess through marriage and seemed to win the heart of the British nation. As well as stunning the public with designer dresses that only a Princess could wear, she became even adored through her charity work. It was a turning point for Aids charity workers when she was seen holding the hand of an Aids victim. Her untimely death left the nation in shock.

LEE

This is Lee Schofield. He is a children's TV presenter on one of the new digital networks. Lee's education set him up for the entertainment business, as he attended a drama and music school in London. Having finished his formal education he was deciding whether to continue into college, when this job opportunity arose. Considered a natural with pre-schoolers his future looks set to be very promising.

JOSH

Josh Matthews is a semi-professional musician. His band the 'dunes' of which he is lead guitarist are popular in clubs and bars around the country. However, their jazz-rock style has not yet earned them a record contract. To support his musical career Josh also writes review columns in NME, which provides his main source of income. His aims are set high, whilst at the same time being happy with having any opportunity to be involved in music one way or another.

JADE

American born Jade Kirby is a young country and western singer from Tennessee. Having seen the meteoric rise of singers such as Leanne Rimes and Shania Twain she is hoping to make into the mainstream of American charts. She is greatly encouraged by her parents who are also singers, and who act as managers to Jade. However, the current popularity of young singers like Brittany Spears means that the family will have to push hard to get Jade her big break. This does not detract from the fact that for a seventeen year old she has a really big powerful voice, which people will always want to hear.


HELEN

This is Helen Murphy whose profession is air stewardess on Virgin Airways. She appeared in one of the popular fly-on-the-wall documentaries, when she was shown passing the various stages required to become a stewardess. Her view of her television appearance was that she barely recalls the cameras. The stress involved in passing the course meant she had no time to dwell on the cameras. "Besides why would I need to worry about what I look like on film, when my job already demands an immaculate appearance?"



Landmark stimuli used in Chapter 4.

A and B indicate the first and second set of landmarks used. Each set had four famous (top row of each set) and four non-famous landmarks (bottom row of each set).

A



B



Product stimuli used in Chapter 5.

A and B indicate the first and second set of products used. Each set had four famous (top row of each set) and four non-famous products (bottom row of each set).

A



B



C

Included the cola cans and lemonade bottles of Set B, the shampoos of Set A (Organics and Nexus), plus:



Face Stimuli used in Experiment 7
1st Images:



2nd Images:

