University of Alberta

CHARACTERIZING GENERAL AND FACE SPECIFIC ERP CORRELATES OF FACE MEMORY

by

Reiko Graham ©

A thesis submitted to the Faculty of Graduate Studies and Research
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

Department of Psychology

Edmonton, Alberta

Fall, 2002

National Library of Canada
Acquisitions and Bibliographic Services
395 Wellington Street
Ottawa ON K1A 0N4
Canada

Bibliothèque nationale du Canada
Acquisitions et services bibliographiques
395, rue Wellington
Ottawa ON K1A 0N4
Canada

Your file Votre référence

Our file Notre référence

The author has granted a non-exclusive licence allowing the National Library of Canada to reproduce, loan, distribute or sell copies of this thesis in microform, paper or electronic formats.

L'auteur a accordé une licence non exclusive permettant à la Bibliothèque nationale du Canada de reproduire, prêter, distribuer ou vendre des copies de cette thèse sous la forme de microfiche/film, de reproduction sur papier ou sur format électronique.

The author retains ownership of the copyright in this thesis. Neither the thesis nor substantial extracts from it may be printed or otherwise reproduced without the author's permission.

L'auteur conserve la propriété du droit d'auteur qui protège cette thèse. Ni la thèse ni des extraits substantiels de celle-ci ne doivent être imprimés ou autrement reproduits sans son autorisation.

0-612-81195-6

Canada

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

UNIVERSITY OF ALBERTA

LIBRARY RELEASE FORM

Name of author: Reiko Graham

Title of thesis: Characterizing general and face specific ERP correlates of face memory

Degree: Doctor of Philosophy

Year this degree granted: 2002

Permission is hereby granted to the University of Alberta Library to reproduce single copies of this thesis and to lend or sell such copies for private, scholarly, or scientific research purposes only.

The author reserves all other publication or other rights in association with the copyright in the thesis, and except as herein before provided, neither the thesis nor any substantial portion thereof may be printed or otherwise reproduced in any material form whatever without the author's prior written permission.

10522 77th Avenue
Edmonton, Alberta, Canada T6E 1N1

Date: Oct. 02/2002

University of Alberta

Faculty of Graduate Studies and Research

The undersigned certify that they have read, and recommend to the Faculty of Graduate Studies and Research for acceptance, a thesis entitled CHARACTERIZING GENERAL AND FACE SPECIFIC ERP CORRELATES OF FACE MEMORY submitted by Reiko Graham in partial fulfillment of the requirements for the degree of Doctor of Philosophy.

Dr. Michael R.W. Dawson
(Supervisor)

FOR Dr. Roberto Cabeza
(Co-Supervisor)

Dr. Walter F. Bischof
(Committee Member)

Dr. Pierre Boulanger
(Committee Member)

FOR Dr. Hal Weinberg
(External Examiner)

Date: SEP. 30/02

For H.W.
(na^6)

Abstract

Despite inquiry, questions remain regarding the event-related potential (ERP) correlates of face recognition. ERP components that index the activity of cortical areas involved in general memory processes (transmodal areas) have yet to be firmly established. Similarly, the memory sensitivity of components thought to index the activity in face-specific processing areas (unimodal areas) is debatable. This series of four studies attempted to identify and characterise the nature of general and face specific ERP correlates of face memory. The first two studies demonstrate the generalizability of verbal ERP old/new effects to face recognition, suggesting that these effects index general memory processes. In the first study, episodic memory for faces with and without facial expression were compared. The second study describes a comparison of item and context memory for faces. In both studies, ERP effects were observed that were similar to those reported in studies of verbal memory. The laterality of these effects may be stimulus dependent. The memory sensitivity of ERP correlates specific to face memory are examined the next two studies. The third study reports the results of temporal analyses conducted on right temporal electrode channels with ANN's and ANOVA. Results suggest that early ERP's from unimodal areas do index memory related activity, but that this activity is in the form of higher-order relationships between voltage and time. Classification was achieved through coarse coding in the hidden units and subsets of adjacent timepoints seemed to be driving their activity. This provides support for the memory-sensitivity of ERP effects associated with face processing. The fourth and final study describes a series of topographic analyses of N170 conducted with conventional linear statistics and ANN's using two different procedures for determining peak

amplitude. Results suggest that early activity to faces, as indexed by ERP's, may indeed be memory-sensitive. The procedure used to determine peak amplitude appears to be a critical factor in the outcome of such analyses. As a whole, this series of studies provides a useful framework for understanding the ERP correlates of face memory. The experimental evidence revealed by these studies will hopefully extend our understanding of the timing and nature of the brain events underlying face processing.

Acknowledgments

I gratefully acknowledge the input of my committee in the conception and completion of this dissertation. In particular, I would like to thank my supervisors, Michael Dawson and Roberto Cabeza for their support and guidance over the years. I also extend my gratitude to Walter Bischof for his counsel and encouragement.

I am also indebted to Jamie Khau, Eileen Noel and Annette Colangelo for their help with data collection and analysis, and my former labmates Andrea Hughes and Florin Dolcos for their friendship and advocacy. I would also like to thank reviewers, both known and anonymous, for their useful comments on the different chapters of this dissertation.

This research was supported from grants from the National Sciences and Engineering Research Council of Canada (NSERC) and the Alberta Heritage Foundation for Medical Research to Roberto Cabeza, and grants from NSERC to Michael Dawson.

TABLE OF CONTENTS

CHAPTER 1: INTRODUCTION 1

Event-related potentials (ERP's) 3

Hemodynamic correlates of general visual memory 7

Electrophysiological correlates of general visual memory 8

Experimental evidence regarding ERP correlates of general visual memory 9

Unresolved issues 11

Hemodynamic and anatomical correlates of face specific memory 11

Electrophysiological correlates of face specific memory 14

Experimental evidence regarding ERP correlates of face specific memory 15

Unresolved issues 17

Rationale for Research 19

Thesis Overview 25

References 26

CHAPTER 2: EVOKED POTENTIALS OF RECOGNIZING HAPPY AND NEUTRAL FACES 35

Introduction 36

Materials and Methods 37

Results 39

Discussion 41

References 46

CHAPTER 3: DISSOCIATING THE NEURAL CORRELATES OF ITEM AND CONTEXT MEMORY: AN ERP STUDY OF FACE RECOGNITION 48
Introduction 49
Method 53
Results 55
Discussion 58
References 67
CHAPTER 4: ARTIFICIAL NEURAL NETWORKS AS ANALYTIC TOOLS IN AN ERP STUDY OF FACE MEMORY 71
Introduction 72
Methods 74
Results 76
Discussion 81
References 90
CHAPTER 5: THE IMPORTANCE OF PEAK ANALYSIS PROCEDURES IN THE EXAMINATION OF EARLY EVENT-RELATED CORRELATES OF FACE MEMORY........ 93

Introduction 94

Methods 100

Results 102

Discussion 107

References 114

CHAPTER 6: GENERAL DISCUSSION 118

APPENDIX A 127

APPENDIX B 132

LIST OF TABLES

Table 4-1: Summary of timepoints identified as significant predictors of net input via multiple regression for each hidden unit of each network. Numbers in the table represent the input number, as well as r and R^2 values for each regression; the timecourse is indicated along to bottom of the table to aid interpretation. 89

LIST OF FIGURES

Figure 2-1: Average voltage values (microvolts) obtained during recognition of happy and neutral faces at parietal and frontal electrodes on both hemispheres. 43

Figure 2-2: Mean amplitudes (microvolts) for hits and CR's (happy and neutral faces combined) at parietal electrodes from 375 ms to 750 ms. 44

Figure 2-3: Mean difference amplitudes (hits minus CR's, in microvolts) for happy and neutral faces at frontal electrodes from 750 ms to 1250 ms. 45

Figure 3-1: Average voltage values (μV) obtained during recognition of happy and neutral faces at parietal and frontal electrodes on both hemispheres. 65

Figure 3-2: Mean difference amplitudes for hit-hits (hit-hits minus CRs) and for hit-misses (hit-misses minus CRs) obtained at frontal and parietal electrodes (μV), collapsed across epochs and hemispheres. 66

Figure 4-1: Grand-averaged voltage values (μV) obtained from right temporal electrodes during the first 500 ms of face presentation. 85

Figure 4-2. Plot of cell means for the early epoch (100-200ms) for hits and CR's 86

Figure 4-3: The network architecture used in this study. Average standardized voltage values represent averaged standardized ERP's obtained from each subject for each ERP type. 87

Figure 4-4: Three-D scatterplots of hidden unit activity for each of the ANN's. Plots correspond to ANN's trained on data from right temporal electrode sites (FT8, T8 and TP8, respectively). 88

Figure 5-1. Grand averaged ERP waveforms corresponding to hits and correct rejections (CR's) 110

Figure 5-2: The network architecture adopted by the ANN used in this study. Average voltage values represent averaged topographic voltage values corresponding to the average latency of the N170 obtained from each subject for each ERP type (hits and CR's). Lines between the layers of the ANN represent weighted connections. 111

Figure 5-3: Two-dimensional scatterplots of A) net input to each hidden unit and B) hidden unit activity (net input after logarithmic transformation). 112

Figure 5-4: Topographic distribution of standard deviations: A) before weighting (raw), B) after weighting with connections to hidden unit 1, C) after weighting with connections to hidden unit 2. 113

CHAPTER 1

INTRODUCTION

Introduction

The human face is a complex and salient stimulus, conveying critical social information. Faces provide information about gender, age, intentions and health; telling us what a person is attending to and what they are feeling. Faces provide information about identity, information of substantial importance for interpersonal relationships. Face recognition, which has proven difficult to simulate with computer models, is an ability that normal humans can perform accurately and easily. Despite this facility, the character and timecourse of the neural processes underlying face recognition remain a matter of debate.

One distinction that may be useful in the description of face recognition is that memory can be viewed as the result of processing in transmodal and unimodal association areas (Mesulam, 1998). Transmodal areas are those that fulfill three criteria: they receive inputs from more than one sense or modality, cells do not show preferences for stimulus modality, and damage to these areas results in global behavioural deficits. In humans, transmodal areas include the prefrontal cortices, posterior parietal cortex, and parts of the lateral and medial temporal cortices. Unimodal areas also satisfy three conditions: they receive inputs from only one modality, neurons respond selectively to stimuli in a given modality, and damage to these areas gives rise to circumscribed deficits. In humans, unimodal visual areas include extrastriate areas and the fusiform, inferior, and middle temporal gyri. The interaction between transmodal and unimodal cortical areas ultimately gives rise to the phenomenon of conscious recollection.

Face recognition then can be conceptualized as the result of two types of processes; those associated with the activity of transmodal areas, and those associated with unimodal areas. Within the visual modality, memory for faces may also be conceptualized as the result of processing in areas that are common to memory for many types of visual stimuli, as well as areas that are specific to face processing. Memory-related activity in these cortical areas is amenable to examination with event-related potentials (ERP's). In the following sections, ERP's are introduced, as well as methodological concerns associated with their use in the study of general and face specific memory processes. This is followed by a review of the current literature regarding hemodynamic, anatomical and electrophysiological correlates of face memory as evidenced by neuroimaging and single cell studies, as well as from studies with high temporal resolution such as ERP or intracranial studies. For the purposes of this review, only studies employing unfamiliar faces as stimuli are considered. Studies employing famous faces are not reviewed because famous faces have associated semantic information, making their recognition more similar to that for words (Bruce & Young, 1986). Some of the unresolved issues regarding general and face specific ERP correlates of recognition will also be introduced, as well as the rationale for the research reported in subsequent chapters of this dissertation.

Event-related potentials (ERP's)

In this section, ERP's are introduced and briefly described. This is followed by a description of the experimental designs employed in the search for ERP correlates of general and face specific memory and the possible impact of these designs on experimental results. Various methods of analysing ERP data are also introduced, as well as the possible impact of these various methods on the description of general and face specific ERP correlates of face memory.

Although a more detailed description of ERP's is available in Appendix A, some discussion of ERP's is necessary here. ERP's refer to changes in the brain's electrical activity, time-locked to an event (i.e., stimulus presentation). These changes are very small in comparison to the background, random neural activity of the brain. Consequently, the signal is extracted by averaging across a number of trials, usually along experimental conditions, yielding an estimate of the neural activity elicited by different types of stimuli, conditions and/or disease states. ERP's are useful in the study of cognitive phenomena for two reasons. First, activity associated with the processing of various stimulus types can be recorded with a temporal resolution sufficient to detect changes in activity that last only for a few milliseconds, allowing for inferences regarding the time course of certain ERP components (Allan et al., 1998). Second, ERP's have utility in determining if different experimental conditions elicit functionally dissociable cognitive process (Ganis, Kutas & Sereno, 1996). Because of these properties, ERP's can be used to detect changes in brain electrical activity during memory. With regard to identifying the ERP correlates of general and face specific memory, correlates that can be evidenced with faces as well as other stimuli would qualify as correlates of general visual memory, while others will only be elicited expressly to faces.

One important issue to consider is what experimental design to use to study face recognition with ERP's. In studies of memory with unfamiliar faces, two different types of recognition paradigms are used. Old/new paradigms are typically used in the study of ERP correlates of verbal memory, and have been recently applied to the study of memory for unfamiliar faces (i.e. Graham & Cabeza, 2001a; Graham & Cabeza, 2001b). This paradigm consists of study blocks where subjects view unfamiliar faces, which are followed by test blocks that consist of studied and new faces. Participants are required to indicate which faces are studied and which are new. Study and test blocks are separated by several minutes, creating long lags between studied faces. Another paradigm used in ERP studies of face memory is the continuous recognition paradigm. Although there is a fair amount variation in the use of this paradigm, it usually consists of the presentation of unfamiliar faces, some which are repeated at various stimulus lags. Here, the subject is required to indicate which faces are repeated. In continuous recognition, the lag between repeated faces is typically quite short.

Choice of paradigm in the study of face memory could be important in the study of general and face specific memory effects. For example, ERP correlates of general memory may be insensitive to long stimulus lags because they are not modulated by short-term memory processes. However, face specific correlates could be modulated by short-term memory and memory effects may be attenuated or absent at longer lags. Attenuation would make it difficult to detect these effects with conventional linear statistics. Instead, more powerful pattern classification tools such as artificial neural networks may be necessary.

Another issue to consider is how data should be analysed when studying ERP correlates of face memory. This point is given considerable emphasis in Chapter 5, however some explanation should be given here. The method of analysis most commonly used in ERP studies is repeated-measures analysis of variance (RM-ANOVA). After averaging the raw EEG according to trial type for each subject, a number of different waveforms for each trial type are produced. These waveforms contain information about voltage changes over time. Researchers identify epochs of interest, which represent temporal windows in which trial types are most likely to differ. These then become levels of the independent variable in RM-ANOVA. The method by which epochs are chosen varies dramatically from study to study. Some researchers focus on only a few electrodes and epochs of interest, excluding those where differences between waveforms have not been documented or are not visibly apparent.

Other researchers choose an epoch window length and systematically analyse the entire ERP waveform at channels of interest (e.g. every 50 ms). Other researchers choose to conduct peak amplitude analyses on specific ERP components at specific channels of interest. As mentioned in Chapter 5, a number of different stragies for determining peak amplitude can be used for peak analyses. The choice of method for ERP analysis may not be as important in the study of general correlates of visual memory, where memory effects appear to be quite sustained (e.g., Allan et al., 1998). However, it could be of critical importance in the study of face specific correlates, where effects could be very transient.

It is clear that methodological issues regarding the collection and analysis of ERP's are a critical concern in the determination of general and face specific correlates of face memory. In the remainder of this chapter, I review the literature regarding hemodynamic and electrophysiological correlates of face memory with emphasis on suggesting putative general and face specific cortical areas. For the purposes of this review, general face memory areas are defined as those areas showing memory effects for other stimuli as well as faces, while face specific areas exhibit effects unique to faces. Given the flow of information through these areas, it is likely that face specific areas are recruited earlier than general visual memory areas (Mesulam, 1998). In the case where temporal information is available (electrophysiological correlates), face specific activity is also assumed to onset earlier than general activity. With regard to ERP correlates of general and face specific visual memory, emphasis is placed on issues that remain unresolved in the current literature. These issues are then be used to generate a framework for the studies that constitute the body of this dissertation.

Hemodynamic correlates of general visual memory

The results from a wealth of studies suggest that there is a common network of brain areas involved in the recognition of words and faces. Here, I summarize the results of neuroimaging studies that provide evidence that regardless of stimulus type, visual recognition is associated with activity in prefrontal, lateral parietal, anterior cingulate and cerebellar regions. These areas are activated during recognition of both words and faces, thereby fulfilling the criteria for general visual memory areas.

In a review of 257 neuroimaging studies, Cabeza and Nyberg (2000) found overwhelming evidence for the involvement of prefrontal cortex in verbal memory, which is also evident during face memory. This activity has been related to retrieval mode, retrieval effort, and retrieval success. Gur et al. (1997) used PET to compare verbal and face memory and did not find differences in rCBF in frontal regions during retrieval. Andreasen et al. (1996) compared rCBF changes during face perception and recognition and also found increases in prefrontal activation during memory. McDermott, Buckner, Peterson, Kelley & Sanders (1999) compared word and face retrieval with fMRI and found greater activity in the left inferior frontal regions during word retrieval and greater right inferior frontal activity during face retrieval. Both types of retrieval elicited activity in right frontal polar areas. This suggests that common brain areas seem to be involved in the retrieval of faces and words, but the laterality of activation may be sensitive to stimulus type. The finding of right prefrontal activations during retrieval has been reported in several other PET studies of verbal and face memory and strongly implicates prefrontal areas in general memory (e.g. Haxby et al., 1996; Henke, Buck, Weber, and Wieser, 1997; Nyberg et al., 1996; Tulving, Kapur, Craik, Moscovitch & Houle, 1994).

Other regions that are activated during both verbal and face memory include lateral parietal cortex, the cerebellum, and the anterior cingulate gyrus (Cabeza & Nyberg, 2000). In addition to frontal activity, Andreasen et al. (1996) also reported significant activations in right lateral parietal cortex, left cerebellum, and the anterior cingulate during face memory. Haxby et al. (1996) also reported a similar pattern of activations. Lateral parietal regions have been related to the perceptual component of recognition (Cabeza et al., 1997), cerebellar activation is thought to reflect self-initiated retrieval, while activation of the anterior cingulate is believed to index response selection and initiation (Cabeza & Nyberg, 2000). Therefore, it is likely that these areas are involved in general memory processes and are not specific to a certain stimulus type.

Electrophysiological correlates of general visual memory

Hemodynamic studies suggest that the recognition of any visual stimulus involves a network of cortical areas, including the frontal and parietal cortices. These findings support the idea that face recognition involves the recruitment of cortical areas supporting general visual memory. In the next two sections, I review ERP studies of verbal and face recognition, then discuss some enduring uncertainties regarding two possible ERP correlates of general visual memory. The results of studies of verbal recognition will be compared to analogous studies using unfamiliar faces as stimuli. Overall, these studies provide some support for common ERP correlates for word and face recognition. However, although results are suggestive of general processing mechanisms involved in face recognition, there is a significant lack of correspondence between the paradigms used in ERP studies of verbal and facial memory that makes direct comparison difficult. This shortcoming provides the motivation for Studies One and Two of this dissertation.

Experimental evidence regarding ERP correlates of general visual memory

ERP studies have associated recognition with two effects: the parietal effect and the frontal effect. During retrieval, ERP's to remembered items show a positivity occurring after approximately 400ms over left parietal regions relative to new items (e.g., Senkfor & Van Petten, 1998; Wilding, Doyle & Rugg, 1995; Wilding & Rugg, 1996; Wilding, 1999) This effect has been found consistently in ERP studies of verbal recognition memory and has been attributed to hippocampal-cortical interactions during retrieval (Allan et al., 1998). The second effect is a positivity over frontal regions, which onsets later than the parietal effect and has been strongly associated with the retrieval of contextual information (Wilding et al., 1995; Wilding & Rugg, 1996; Wilding, 1999). The frontal effect is more sustained, and has been found to be right lateralized with verbal stimuli (Wilding et al., 1995; Wilding & Rugg, 1996; Wilding, 1999). This effect has been attributed to the activity of the frontal cortex, and may be an index of post-retrieval monitoring (Allan et al., 1998) or retrieval success (Graham & Cabeza, 2001b). The majority of studies examining old/new effect have used verbal materials as stimuli (Allan et al., 1998), although a study has also reported the left parietal effect for correctly recognized pictures (Schloerscheidt & Rugg, 1997). However, these effects had not been previously reported with stimuli such as unfamiliar faces that require a minimum of verbal processing.

Endl et al. (1998) examined old/new effects with unfamiliar faces as stimuli, focusing specifically on activity at occipitoparietal sites. The time courses of the waveforms began to diverge after 200 ms, with maximum differences between them occurring during a phasic negative peak at approximately 350 ms. This differential negativity reappeared at 500 ms in a slow negative wave that was largest over left occipitoparietal sites. It is important to note that the effect described by Endl et al. (1998) is an negativity over occipitoparietal areas occurring at approximately 350 ms, which is larger for correctly identified target faces. This topography is different from the left parietal increases in positivity to targets reported by Allan et al. (1998).

Some evidence for ERP correlates of general visual memory can also be found in studies employing continuous recognition paradigms although admittedly, these types of paradigms do not map directly onto those used to elicit old/new effects with verbal stimuli. Potter and Parker (1989) examined memory for faces in an ERP repetition paradigm and found that ERP's to repeated faces were relatively more positive than ERP's to non-repeated faces. These effects were identical to those reported in other repetition studies using faces (Barrett, Rugg, & Perrett, 1988; Bentin & McCarthy, 1994; Itier & Taylor, 2002; Seecke et al., 1997) and words (Bentin & McCarthy, 1994). Muente et al. (1997) compared repetition effects produced by the correct recognition of faces to those produced by words. ERP's to words and faces had very different overall waveforms. However, when old/new difference waves were compared, the resulting waveforms were remarkably similar in both latency and topography. Therefore, this repetition effect probably reflects a later stage of retrieval that is not domain-specific. However, unlike the left parietal old/new effect reported by Allan and his colleagues, this old/new effect was maximal over right parietal sites, although the latency (~400 ms) was similar.

Unresolved issues

What is clear from this review of ERP studies of verbal and face memory is that inconsistencies in experimental results preclude drawing conclusions regarding the existence of ERP correlates of general visual memory. Whereas the ERP studies of verbal memory conducted by Rugg and colleagues employ an old/new recognition paradigm with lags between study and test sessions, ERP studies of face memory typically use continuous recognition paradigms which may be more appropriate for the study of short-term, rather than long-term memory processes. This methodological difference could give rise to inconsistencies across studies of verbal and facial memory. In addition, the ERP studies reviewed above describe old/new effects that appear to differ dramatically from study to study, making it impossible to establish whether old/new effects differ for faces and words. Because of this, ERP studies of recognition do not provide clear evidence for memory processes that are common to both types of stimuli. These issues can only be resolved with further research using old/new paradigms similar to those used in the study of verbal episodic memory. This will be the focus of Studies One and Two of this dissertation.

Hemodynamic and anatomical correlates of face specific memory

Unlike general memory processes, face specific processing areas are unique to faces. Various lines of evidence suggest that people and other primates have genetically predetermined mechanisms for the detection of faces and that specific brain areas subserve this function. In this section, I summarize the results of these studies, which range from single cell recordings in the macaque to clinical and hemodynamic neuroimaging studies (PET and fMRI) with humans. These studies suggest that areas of the temporal lobe, notably the inferior temporal gyrus and superior temporal sulcus in the macaque, and fusiform gyrus in humans fullfil the criteria of face specific processing areas.

Single cell recordings in the macaque shown that different populations of neurons in the temporal lobe respond selectively to biologically meaningful stimuli (Walsh & Perrett, 1994). These studies have established the existence of face-selective cells; cells in the lateral temporal lobe which fire to the presentation of faces (Desimone, 1991; Harries & Perrett, 1991; Rolls, 1984, Rolls, 1992). Two major clusters of cells have been identified: one in the ventral bank of the superior temporal sulcus (STS), another in the adjacent cortex on the inferior temporal gyrus (IT) (Desimone, 1991; Harries & Perrett, 1991; Rolls, 1984; Rolls, 1992). These cells had excitatory discharges with latencies between 80 and 180 ms (Perrett, Rolls, & Caan, 1982) and responded to faces in different sizes and orientations. Yamane, Kaji and Kawano (1988) recorded the responses of IT neurons in the macaque to human faces and found that responses were more similar when faces were physically similar. Activity was best predicted by the configuration of facial features, such as inter-eye distance, suggesting that these cells participate in the perception of faces through some form of population coding (Yamane et al., 1988).

Neurological studies have demonstrated that face perception and recognition can be selectively impaired (Farah, Levinson, & Klein, 1995) or spared (Moscovitch, Winocur & Behrmann, 1997) relative to the recognition of objects of equivalent visual complexity. Prosopagnosia, a deficit in the recognition of faces, is associated with bilateral damage to the inferior areas of the temporal cortex, in particular, the fusiform gyrus (Damasio, Damasio & Van Hoesen, 1982). These cases imply that different brain areas could be involved in the recognition of faces and the recognition of other objects and implicates the fusiform area in face processing.

Sergent et al. (1992) reported significant PET activations in widespread areas of the temporal cortex during face processing. Faces produced activations in the right cuneus, right occipital and occipitotemporal gyri, the right parahippocampal gyrus, as well as activations in the bilateral temporal poles and fusiform gyri. Similarly, PET studies by Haxby et al. (1991; 1992; 1994) reported selective activity in the fusiform gyrus during face processing. Bilateral fusiform activation to faces has since been replicated in other PET studies (Andreasen et al., 1996; Kapur, Friston, Young, Frith, & Frackowiak, 1995; Wicker, Michel, Henaff, & Decety, 1997).

MRI studies also implicate ventral temporal areas in unimodal face memory. An fMRI by Kanwisher, McDermott, and Chun (1997) identified an area of the fusiform gyrus that was significantly more active to faces than other objects. Puce, Allison, Gore and McCarthy (1995) reported significant activations in the midfusiform gyri to faces. At posterior locations, faces elicited more lateral activations in inferior temporal gyri. Puce, Allison, Asgari, Gore and McCarthy (1996) also reported bilateral fusiform activations to faces, as will as activations in the right occiptitotemporal and inferior occipital regions, and bilateral middle lateral temporal cortex. Another study by the same group of researchers showed that faces elicited activations that were specific only to the lateral portions of the right fusiform gyrus (McCarthy, Puce, Gore, & Allison, 1997). These results are congruent with the results of other fMRI studies which reported significant activations in ventral occipitotemporal areas specific to faces (Clark, Maisog, & Haxby, 1998; Clark et al., 1996; Leveroni et al., 2000).

In conclusion, PET and fMRI studies suggest that a specific neuroanatomical area, the fusiform gyrus, may subserve face perception. Evidence from Kanwisher, Tong and Nakamura (1998) suggests that this area may be involved only in early stages of face processing. The finding that blood oxygen level dependent (BOLD) differences in temporal lobe activation are not found between learned and novel faces with fMRI (Leveroni et al., 2000) could also support the involvement of this area only in early perceptual processing. Therefore, it is likely that face recognition involves other cortical areas that are in a more superior and lateral location, homologous to face selective areas in IT of the macaque. Furthermore, the identification of individuals most likely involves the involvement of general visual memory structures, and may therefore involve areas such as the frontal and parietal lobes, the anterior cingulate and the cerebellum. However, it is important to note that if face specific memory effects are manifested as changes in the timecourse of activations and not as increases in the magnitudes of the activations per se, then hemodynamic measures would be unable to detect these differences.

Electrophysiological correlates of face specific memory

Similar to hemodynamic studies with high spatial resolution, electrophysiological correlates of face memory obtained with methods with high temporal resolution also suggest that face recognition is subserved, at least in part, by specific brain areas. In the next two sections, I summarize the results of studies that stem from intracranial and ERP studies with humans and introduce some unresolved issues regarding ERP correlates of unimodal face memory. The review of intracranial and ERP studies suggests that face processing is strongly correlated with an early negative deflection at occipitotemporal electrodes and concurrent positive deflection over frontocentral areas that are consistent with activity of a face specific processing area and could reflect the activity of the fusiform gyrus. The issue of whether these potentials are, in themselves, memory sensitive will also be introduced and serves as the motivation for Studies Three and Four of this dissertation.

Experimental evidence regarding ERP correlates of face specific memory

In study using implanted depth electrodes, Allison et al. (1994) reported a surface-negative potential occurring approximately 200 ms after stimulus onset was evoked by faces but not other categories of objects, such as cars or scrambled faces (Allison et al., 1994). This potential was only observed in small areas of the left and right fusiform gyri, as well as the bilateral inferior temporal gyri. Electrical stimulation of the same areas often produced temporary anomia for familiar faces. Subsequent intracranial studies by the same group of researcher have replicated this finding of early face specific activity to faces from ventral temporal areas (Allison, Puce, Spencer & McCarthy, 1999; McCarthy, Puce, Belger & Allison, 1999; Puce, Allison, & McCarthy, 1999).

Bentin, Allison, Puce, Perez, & McCarthy (1996) compared the temporal and topographic characteristics of ERP's elicited to faces and other objects. Human faces evoked a negative potential at approximately 170 ms after face onset that was not evidenced in ERP's to non-face stimuli. This potential was maximal over the posterior temporal scalp areas and was larger on the right. It is likely that this potential is the ERP correlate of the fusiform activity recorded intracranially by Allison et al. (1994). The latency and topographic distribution of the N170 are also consistent with single cells studies that report excitatory face selective cells in IT with response latencies between 80 and 180 ms (i.e., Perrett et al., 1982). The N170 has since been studied extensively (e.g. Eimer, 2000a; Itier & Taylor, 2002).

George, Evans, Fiori, Davidoff and Renault (1996) compared ERP's elicited by intact and scrambled faces. They reported a positive potential that occurred after approximately 200 ms that was largest over the vertex (the centre of the top of the head). This face sensitivity of this potential, dubbed the vertex positive potential (VPP), has been observed in several studies (e.g. Jeffreys, 1993; Itier & Taylor, 2002). The topographic distribution of the VPP is consistent with activity in face processing areas such as the STS, and IT, as well as activity in the parahippocampal, and fusiform gyri (George et al., 1996). Due to their similar latencies, the N170 and the VPP are thought to reflect the activity of the same neural generator, although this remains a matter of debate.

In summary, the results of intracranial and ERP studies of face processing are consistent with studies describing hemodynamic correlates of face specific memory. Intracranial studies describe face-selective activity in the ventral temporal lobes that is consistent with fusiform activity reported by fMRI and PET studies. This activity occurs relatively early after face onset, and prior to the parietal and frontal effects described prior sections. The latency of this activity is similar to that of face-selective cells in the macaque temporal lobe. In addition, the latency and topography of the N170 and VPP map directly onto the results from intracranial studies. There are some discrepancies between the results of human and non-human studies. However, this may be due to a lack of homology between human and non-human primate face processing areas. These discrepancies may be due to displacement in the human brain, likely caused by the evolution of newer cortical areas (i.e., language areas) (Haxby et al., 1991).

Unresolved issues

One issue that remains unresolved with regard to electrophysiological correlates of face specific memory is the actual memory sensitivity of these effects. According to Mesulam (1998), the activity of unimodal (stimulus specific) areas alone is insufficient for recognition, which also involves transmodal (general) areas. However, what remains unclear is whether activity in ventral temporal areas is actually memory sensitive. Single cell recordings from the IT in the macaque show that face sensitive IT neurons alter their response rates as a face becomes familiar or when a novel face is added to a set of familiar faces (Rolls, Bayliss, Hasselmo, & Nalwa, 1989). Face sensitive IT neurons also show evidence of familiarity responses to known faces, even 24 hours after they were encountered, suggesting that these cells may be involved in long-term memory as well as perception (Fahy, Riches & Brown, 1993).

The evidence from humans regarding the memory sensitivity of early ERP correlates of face processing is more confusing. With regard to the N170, some researchers have reported reduced amplitude to repeated faces (Campanella et al, 2000; Endl et al., 1998), while Itier and Taylor (2002) show both decreased amplitude and latency of the N170 to repeated faces. These studies suggest that the N170 does index memory-related activity in face areas. Other researchers have not replicated these findings (Bentin & Deouell, 2000; Eimer, 2000a; Eimer, 2000b; Eimer, 2000c), leading to the interpretation that the N170 reflects the structural encoding of face information, in particular, the encoding global configural facial information. This view is also supported by evidence from fMRI (Kanwisher et al., 1998; Leveroni et al., 2000). The memory sensitivity of the VPP has been studied less extensively and also remains a matter of debate. Potter and Parker (1989) reported an early frontal positivity with a latency and topography similar to that of the VPP that was sensitive to face repetition. VPP memory effects have also been reported by Itier and Taylor (2002), but not by Rossion et al. (1999).

One possible explanation for the mixed results regarding the memory sensitivity of early ERP correlates of face processing is that they reflect repetition effects that can be attributed to perceptual priming. Indeed, most of the studies that reported early memory effects employed continuous recognition tasks where the lags between repeated faces were very short (e.g. Itier & Taylor, 2001), whereas studies employing longer lags between repeated faces typically did not show these effects (e.g., Eimer, 2000b). Therefore, it is possible that the memory-sensitive N170 and VPP effects will only be evidenced when repetition is relatively immediate and will not be seen at longer lags (Guillaume & Tiberghien, 2000; Itier & Taylor, 2002). However, even if these memory sensitive effects are due to perceptual priming, it is still not clear whether these effects disappear at longer lags that are typical of old/new recognition paradigms, or if they are merely attenuated. This question provides the motivation for Studies Three and Four of this dissertation.

Another, more methodological explanation could account for mixed results regarding the memory sensitivity of the N170 and VPP. Itier and Taylor (2002) report a small but consistent latency effect for both components, which may not have been detectable in other studies due to small sample sizes. Additionally, relatively few of the studies reviewed above directly examined peak latency and there is tremendous variability in the procedures used for peak amplitude analysis. For example, Eimer (2000c) did not find evidence of memory sensitivity, but did not directly examine latency. Additionally, the peak of the N170 for each subject was quantified as the mean amplitude between 140-190ms post-stimulus, regardless of individual variation. On the other hand, Itier and Taylor (2002) determined the peak latency and amplitude of the N170 and VPP individually for each subject and each condition and found evidence of repetition sensitivity in both latency and amplitude. Therefore, inconsistencies could be caused in part by the different methods that researchers employ to determine peak amplitude. The possibility that memory effects in face specific cortical areas are manifested as latency changes could provide an alternate explanation for fMRI results, which do not show BOLD differences in these areas during memory. The effects of different peak amplitude analyses on a large data sample will be examined in Study Four.

Rationale for Research

The goal of this chapter was to provide an introduction to current research directions and issues in the study of face recognition with regard to general and face specific correlates of face recognition. The evidence regarding hemodynamic correlates of general visual memory suggest that a network of cortical areas, primarily, the parietal and frontal lobes, are involved in the retrieval of various types of stimuli. Putative electrophysiological correlates of general visual memory, on the other hand, have yet to be firmly established. In particular, ERP studies of verbal and nonverbal memory for faces show so many inconsistencies in experimental results that they prohibit drawing conclusions regarding the nature of memory for faces. Because of this, ERP correlates of recognition have not yet provided clues to whether or not electrophysiological correlates of general visual memory can be firmly established. The first two studies in this dissertation address this concern by attempting to assess the generalizability of the parietal and frontal old/new ERP effects to face memory. In order to make the results comparable to those obtained with verbal materials, old/new paradigms will be employed that are similar to those typically used to elicit these effects with words.

Studies One and Two examine whether the parietal and frontal effects reflect general memory processing or activity specific to verbal memory. In Study One, our objective was to investigate these effects during face recognition. The left lateralization of parietal has been attributed to the verbal nature of the materials employed in most studies (Allan et al., 1998, Maratos, Allan & Rugg, 2000). ERP's were measured during the recognition of unfamiliar faces, which should involve a minimum of verbal mediation. Consistent with neuroimaging studies (Cabeza & Nyberg, 2000), the parietal effect should be elicited for faces and should be bilateral. In contrast with the parietal effect, the frontal effect is right lateralized, consistent with neuroimaging evidence of right frontal activity during retrieval (Nyberg et al., 1996). The laterality of the effect may also be affected by the emotional content of the stimuli. Moreover, it is unknown if stimulus valence influences ERP effects during memory. To investigate this, the frontal effect was examined during the recognition of happy and neutral faces. It was predicted that the frontal effect would be elicited during face recognition, but that its lateralization would be affected by stimulus valence.

In Study Two, the effect of context retrieval on the parietal and frontal old/new effects was investigated. Context memory refers to how, when, and where items are presented and includes source, temporal, and spatial memory. Study Two was designed to determine if facial expression acts as contextual information in the retrieval of identity. Specifically, I attempted to determine the sensitivity of the frontal and parietal effects to item and context memory during face recognition. There are two unresolved issues concerning ERP's for item and context memory. First, it is unknown if the context memory effect as manifested by the frontal effect can be generalized to nonverbal stimuli. Second, it is unclear if the parietal effect is sensitive to context retrieval. Consistent with ERP studies of verbal memory, the frontal effect should be elicited during face recognition and should be larger during context retrieval than during item retrieval. It was also predicted that there would be a dissociation between frontal and parietal effects: the frontal effect would be sensitive to context retrieval whereas the parietal effect would be sensitive to item retrieval.

With regard to face specific memory, evidence regarding hemodynamic and anatomical correlates suggests that ventral temporal cortex, namely the fusiform gyrus, may have a specialized role in face processing. Electrophysiological studies of face processing have identified two potentials, the N170 and the VPP, which appear to reflect face-selective activity in ventral temporal areas. However, the memory sensitivity of these components remains unclear. Specifically, it is unclear whether memory processes will affect the latency and/or amplitude of these components. There could be several explanations for inconsistencies in results.

First, temporal areas specific to face processing may not have a direct role in long-term memory processes; repetition-sensitive effects may be elicited under certain conditions that reflect transient perceptual priming effects. Second, the summation of these effects may be such that they cannot be detected from the scalp. Third, memory processes may recruit relatively small cortical areas whose activity may not always be detected by scalp recordings. Such effects may be further obscured by inter-subject differences and small sample sizes. A fourth possibility is that temporal areas for face processing are involved in recognition, but early ERP amplitude or latency effects are attenuated at longer study-test intervals such that they cannot always be detected with conventional linear methods. Studies Three and Four of this dissertation explored this possibility by comparing the ability of linear techniques to discriminate between ERP's elicited by remembered and novel faces with that of artificial neural networks (ANN's). Additionally, Study Four will discuss the interpretation of an ANN trained to discriminate between ERP types in the context of different strategies for peak analyses.

ANN's are appropriate for the analysis of ERP data for several reasons. ANN's are powerful analytic tools, particularly when they are employed for pattern recognition. Network architecture and function are discussed in detail in Chapters 4 and 5 and in Appendix B, however, for the sake of clarity, some explanation is given here. An ANN can be conceptualized as a multi-layered system for generating a desired response to a given input. It consists of interconnected layers of units that are analogous to neurons in their activation patterns and ability to operate in parallel (Bechtel & Abrahamsen, 1994). Because classification occurs by comparing the entire ERP against known types without a priori models, ANN's can learn non-formalizable knowledge for which there are no precise rules (neither brain regions nor thresholds for classification are defined) (Anderer, Saletu, Klöppel, Semlitsch & Werner, 1994). Therefore, they are able to detect complex, nonlinear interactions between input features that may remain undetected if linear analysis is employed. Finally, ANN's can tolerate noisy data, which is a hallmark of ERP data, and can therefore distinguish between groups with overlapping distributions and variances (Gupta, Molfese & Tammana, 1995). This ability to handle noise may be of particular importance when attempting to detect an attenuated signal, such as an early ERP memory effect. Therefore, ANN's are ideally suited for complex pattern recognition with ERP's, especially given that some interpretation of ANN's is possible.

ANN's have been employed EEG studies to discriminate between demented patients and normals with excellent accuracy relative to linear methods (e.g., Anderer et al., 1994; Klöppel, 1994a; Pritchard et al., 1994). ANN's have been also used to differentiate between ERP's produced by individuals with brain pathology and controls (e.g., Gupta et al., 1997; Klöppel, 1994b; Slater, Wu, Honig, Ramsay & Morgan, 1994). The major focus of these studies was on differentiating between disease states, which is at odds with the majority of ERP studies where the focus is to differentiate between cognitive states. In order to apply ANN's to a problem that is more in keeping with conventional research, Studies 3 and 4 were designed to investigate whether ANN's can differentiate between ERP's associated with different cognitive states.

Given that this is a novel research direction, I chose two fundamental ways to examine ERP's with ANN's: a temporal approach and a topographic approach. Study Three describes a temporal analysis of data from selected electrode channels. Because the goal was to index the activity of face specific processing areas which are thought predominate in right ventral temporal areas, right temporal electrode channels were chosen for analysis. The motivation for this study stemmed from the desire to characterize the timecourse of early ERP memory effects from these channels. The rationale for this stemmed from the assumption that if early latency activity was influencing recognition, it should be detectable. However, if signal attenuation is an issue, it may not be detectable with linear analysis but may be detectable with ANN's. This approach also allowed for the examination of the memory sensitivity of activity occurring prior to or after the N170. Furthermore, if an ANN could be trained to differentiate between peak topographic voltages values of the N170 to remembered and novel faces obtained during an old/new recognition task, their interpretation could aid in elucidating the timecourse of ERP effects recorded over face specific processing areas.

The second approach was a topographic analysis of the memory sensitivity of the N170 face effect. This study was motivated by the assumption that if N170 memory effects are attenuated by long lags between study and test, they may be detectable with an ANN but not with linear analysis of the same data. If such effects are contingent upon immediate repetition and disappear at longer lags, they should not be detectable with either type of analysis. As in Study Three, if ANN's were successful, I was interested in determining how the ANN was able achieve discrimination. This study also examined effects of different strategies for determining peak amplitude on the results of linear and ANN analyses. If different strategies yielded different results with either type of analysis, then methodological differences between studies could account for mixed results regarding the memory sensitivity of the N170 and VPP.

Thesis Overview

The goal of this research is to investigate the nature of general visual and face specific ERP correlates of face memory. A review of the literature regarding the electrophysiological correlates of face memory has revealed that questions still remain regarding the nature of these effects. The four experiments that comprise this dissertation were designed to address these uncertainties. The studies described in Chapters 2 (Study One) and 3 (Study Two) examine the generalizability of verbal ERP old/new effects to face recognition with the hope of identifying general ERP correlates of visual memory: Chapter 2 presents a comparison of memory for faces with and without facial expression, while Chapter 3 describes a comparison of item and context memory for faces. The studies reported in Chapters 4 (Study Three) and 5 (Study Four) examine the memory sensitivity of the electrophysiological correlates of face specific memory. In Chapter 4, I present the results of temporal analyses conducted on right temporal electrode channels with ANN's and the attempt to characterise the timecourse of early ERP face memory effects. Chapter 5 describes the results of topographic analyses of N170 conducted with conventional linear statistics and ANN's using two different procedures for determining peak amplitude. Finally, Chapter 6 summarizes the key findings of this stream of research and discusses them within the wider context of face recognition.

References

Allan, K., Wilding, E.L. & Rugg, M.D. (1998). Electrophysiological evidence for dissociable processes contributing to recollection. Acta Psychologica, 98, 231-252.

Allison, T., Puce, A., Spencer, D.D. & McCarthy, G. (1999). Electrophysiological studies of human face perception. I: Potentials generated in occipitotemporal cortex by face and non-face stimuli. Cerebral Cortex, 9(5), 415-430.

Anderer, P., Saletu, B., Klöppel, B., Semlitsch, H.V. & Werner, H. (1994).Discrimination between demented patients and normals based on topographic EEG slow wave activity: Comparison between z statistics, discriminant analysis and artificial neural network classifiers. Electroencephalography and Clinical Neurophysiology, 91, 108-117

Andreasen, N.C., O'Leary, D.S., Arndt, S., Cizadlo, T., Hurtig, R., Rezai, K., Watkins, G.L., Ponto, L.B., & Hichwa, R.D. (1996). Neural substrates of face recognition. Journal of Neuropsychiatry and Clinical Neurosciences, 8, 139-146.

Barrett, S.E., Rugg, M.D., & Perret, D.I. (1988). Event-related potentials and the matching of familiar and unfamiliar faces. Neuropsychologica, 26(1), 105-117.

Bechtel, W. & Abrahamsen, A. (1994). Connectionism and the mind: An introduction to parallel processing in networks. Cambridge, MA: Blackwell Publishers.

Bentin, S., Allison, T., Puce, A., Perez, E. & McCarthy, G. (1996). Electrophysiological studies of face perception in humans. Journal of Cognitive Neuroscience, 8(6), 551-565.

Bentin, S. & Deouell, L.Y. (2000). Structural encoding and identification in face processing: ERP evidence for separate mechanisms. Cognitive Neuropsychology, 17(1-3), 35-54.

Bentin, S. & McCarthy, G. (1994). The effects of immediate stimulus repetition on reaction time and event-related potentials in tasks of different complexity. Journal of Experimental Psychology: Learning, Memory and Cognition, 20(1), 130-149.

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

Cabeza, R., Kapur, S., Craik, F.I.M., McIntosh, A.R., Houle, S., & Tulving, E. (1997). Functional neuroanatomy of recall and recognition: A PET study of episodic memory. Journal of Cognitive Neuroscience, 9(2), 254-265.

Cabeza, R. & Nyberg, L. (2000). Imaging cognition: An empirical review of 275 PET and fMRI studies. Journal of Cognitive Neuroscience, 12(1), 1-47.

Campanella, S., Hanoteau, C., Depy, D., Rossion, B., Bruyer, R., Crommelinck, M.. & Guerit, J. M. (2000). Right N170 modulation in a face discrimination task: An account for categorical perception of familiar faces. Psychophysiology, 37(6), 796-806.

Clark, V.P., Maisog, J.M., & Haxby, J.V. (1998). FMRI study of face perception and memory using random stimulus sequences. Journal of Neurophysiology, 79(6), 3257-3265.

Damasio, A.R., Damasio, H., & Van Hoesen, G.W. (1982). Prosopagnosia: Anatomical basis and behavioral mechanisms. Neurology, 32, 331-341.

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

Eimer, M. (2000a). Effects of face inversion on the structural encoding and recognition of faces: Evidence from event-related brain potentials. Cognitive Brain Research, 10(1-2), 145-158.

Eimer, Martin (2000b). The face-specific N170 component reflects late stages in the structural encoding of faces. Neuroreport, 11(10), 2319-2324.

Eimer, M. (2000c). Event-related brain potentials distinguish processing stages involved in face perception and recognition. Clinical Neurophysiology, 111(4), 694-705.

Endl, W., Walla, P., Lindinger, G., Lalouschek, W., Barth, F.G., Deecke, L. & Wang, W. (1998). Early cortical activation indicates preparation for retrieval of memory for faces: An event-related potential study. Neuroscience Letters, 240, 58-60.

Fahy, F.L., Riches, I.P. & Brown, M.W. (1993). Neuronal activity related to visual recognition memory: Long-term memory and the encoding of recency and familiarity information in the primate anterior and medial inferior temporal and rhinal cortex. Experimental Brain Research, 96, 457-472.

Farah, M.J., Levinson, K.L. & Klein, K.L. (1995). Face perception and within-category discrimination in prosopagnosia. Neuropsychologia, 33, 661-674.

Ganis, G., Kutas, M. & Sereno, M.I. (1996). The search for "common sense": An electropysiological study of the comprehension of words and pictures in reading. Journal of Cognitive Neuroscience, 8(2), 89-106.

George, N., Evans, J., Fiori, N., Davidoff, J., & Renault, B. (1996). Brain events related to normal and moderately scrambled faces. Cognitive Brain Research, 4, 65-76.

Graham, R. & Cabeza, R. (2001a). Evoked potentials of recognizing happy and neutral faces. Neuroreport, 12(2), 245-248.

Graham, R. & Cabeza, R. (2001b). Dissociating the neural correlates of item and context

memory: An ERP study of face recognition. Canadian Journal of Experimental Psychology, 55, 154-161.

Guillaume, F. & Tiberghien, G. (2001). An event-related potential study of contextual modifications in a face recognition task. Neuroreport, 12, 1209-1216.

Gupta, L., Molfese, D.L. & Tammana, R. (1995). An artificial neural network approach to ERP classification. Brain and Cognition, 27, 311-330.

Gur, R.C., Ragland, D., Mozley, L.H., Mozley, D.P., Smith, R., Abass, A., Bilker, W., & Gur, R.E. (1997). Lateralized changes in regional cerebral blood flow during performance of verbal and facial recognition tasks: Correlations with performance and effort. Brain and Cognition, 33, 388-414.

Harries, M.H. & Perrett, D.I. (1991). Visual processing of faces in temporal cortex: Physiological evidence for a modular organization and possible anatomical correlates. Journal of Cognitive Neuroscience, 3(1), 9-24.

Haxby, J.V., Grady, C.L., Horwitz, B., Ungerleider, L.G., Mishkin, M., Carson, R.E., Herscovitch, P., Schapiro, M.B., & Rapoport, S.I. (1991). Dissociation of object and spatial visual processing pathways in human extrastriate cortex. Proceedings of the National Academy of Sciences USA, 88, 1621-1625.

Haxby, J.V., Ungerleider, L.G., Horwitz, B., Maisog, J.M., Rapoport, S.I., & Grady, C.L. (1996). Face encoding and recognition in the human brain. Proceedings of the National Academy of Sciences USA, 93, 922-927.

Henke, K., Buck, A., Weber, B., & Wieser, H.G. (1997). Human hippocampus establishes associations in memory. Hippocampus, 7, 249-256.

Itier, R.J. & Taylor, M.J. (2002). Inversion and contrast polarity reversal affect both encoding and recognition processes of unfamiliar faces: A repetition study using ERP's. Neuroimage, 15, 353-372.

Jeffreys, D.A. (1993). A face-responsive potential recorded from the human scalp. Experimental Brain Research, 96, 193-202.

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(11), 4302-4311.

Kanwisher, N. Tong, F., & Nakayama, K. (1998). The effect of face inversion on the human fusiform face area. Cognition, 68, B1-B11.

Kapur, N., Friston, K.J., Young, A., Frith, C.D., & Frackowiak, R.S.J. (1995). Activation of human hippocampal formation during memory for faces. Cortex, 31, 99-108.

Klöppel, B. (1994a). Application of neural networks for EEG analysis: Considerations and first results. Neuropsychobiolgy, 29, 39-46.

Klöppel, B. (1994b). Classification by neural networks of evoked potentials: A first case study. Neuropsychobiology, 29, 47-52.

Leveroni, C.L., Seidenberg, M., Mayer, A.L., Mead, L.A., Binder, J.R., & Rao, S.M. (2000). Neural systems underlying the recognition of familiar and newly learned faces. Journal of Neuroscience, 20(2), 878-886.

Maratos, E.J., Allan, K. & Rugg, M.D. (2000). Recognition memory for emotionally negative and neutral words: An ERP study. Neuropsychologia, 38, 1452-1465.

McCarthy, G., Puce, A., Belger, A. & Allison, T. (1999). Electrophysiological studies of human face perception. II: Response properties of face-specific potentials generated in occipitotemporal cortex. Cerebral Cortex, 9(5), 431-444.

McCarthy, G., Puce, A., Gore, J.S., & Allison, T. (1997). Face-specific processing in the human fusiform gyrus. Journal of Cognitive Neuroscience, 9(5), 605-610.

Mc Dermott, K.B., Buckner, R.L., Peterson, S.E., Kelley, W.M., & Sanders, A.L. (1999). Set-specific and code-specific activation in frontal cortex: An fMRI study of encoding and retrieval of faces and words. Journal of Cognitive Neuroscience, 11(6), 631-640.

Mesulam, M.-M. (1998). From sensation to cognition. Brain, 121, 1013-1052.

Moscovitch, M., Winocur, G., & Behrmann, M. (1997). What is so special about face recognition? Nineteen experiments on a person with visual object agnosia and dyslexia but normal face recognition. Journal of Cognitive Neuroscience, 9(5), 555-604.

Muente, T.F., Brack, M., Grootheer, O., Wieringa, B.M., Matzke, M. & Johannes, S. (1997). Event-related brain potentials to unfamiliar faces in explicit and implicit memory tasks. Neuroscience Research, 28, 223-233.

Nyberg, L., Cabeza, R., & Tulving, E. (1996). PET studies of encoding and retrieval: The HERA model. Psychonomic Bulletin and Review, 3(2), 135-148.

Perrett, D.I., Rolls, E.T. & Caan, W. (1982). Visual neurones responsive to faces in the monkey temporal cortex. Experimental Brain Research, 47, 329-342.

Potter, D.D. & Parker, D.M. (1989). Electrophysiological correlates of facial identity and expression processing. In J.R. Crawford & D.M. Parker (Eds.), Developments in clinical and experimental neuropsychology (pp. 137-150). New York: Plenum Press.

Pritchard, W.S., Duke, D.W., Coburn, K.L., Moore, N.C., Tucker, K.A., Jann, M.W. & Hostetler, R.M. (1994). EEG-based, neural net predictive classification of Alzheimer's disease versus control subjects is augmented by non-linear EEG measures. Electroencephalography and Clinical Neurophysiology, 91, 118-130.

Puce, A., Allison, T., Bentin, S., Gore, J.C., & McCarthy, G. (1998). Temporal cortex activation in humans viewing eye and mouth movements. The Journal of Neuroscience, 18(6), 2188-2199.

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(3), 1192-1199.

Puce, A., Allison, T. & McCarthy, G. (1999). Electrophysiological studies of human face perception. III: Effects of top-down processing on face-specific potentials. Cerebral Cortex, 9(5), 445-458.

Rolls, E.T. (1984). Neurons in the cortex of the temporal lobe and in the amygdala of the monkey with responses selective for faces. Human Neurobiology, 3, 209-222.

Rolls, E.T., Baylis, C.G., Hasselmo, M.E. & Nalwa, V. (1989). The effect of learning on the face selective responses of neurons in the cortex of the superior temporal sulcus of the monkey. Experimental Brain Research, 76, 153-164.

Rolls, (1990). A theory of emotion, and its application to understanding the neural basis of emotion. Cognition and Emotion, 4(3), 161-190.

Rossion, B., Campanella, S., Gomez, C. M., Delinte, A., Debatisse, D., Liard, L., Dubois, S., Bruyer, R., Crommelinck, M., & Guerit, J. -M. Task modulation of brain activity related to familiar and unfamiliar face processing: An ERP study. Clinical Neurophysiology, 110(3), 449-462.

Schloerscheidt, A.M. & Rugg, M.D. (1997). Recognition memory for words and pictures: an event-related potential study. Neuroreport, 8, 3281-3285.

Seeck, M., Michel, C.M., Mainwaring, N., Cosgrove, R., Blume, H., Ives, J., Landis, T. & Schomer, D.L. (1997). Evidence for rapid face recognition from human scalp and intracranial electrodes. Neuroreport, 8, 2749-2754.

Senkfor, A. J. & Van Petten, C. (1998). Who said what? An event-related potential investigation of source and item memory. Journal of Experimental Psychology: Learning, Memory, & Cognition, 24(4), 1005-1025.

Sergent, J.; Ohta, S. & MacDonald, B. (1992). Functional anatomy of face and object processing. Brain, 115, 15-36.

Slater, J.D., Wu, F.Y., Honig, L.S., Ramsay, R.E., & Morgan, R. (1994). Neural network analysis of the P300 event-related potential in multiple sclerosis. Electroencephalography and Clinical Neurophysiology, 90, 114-122.

Tulving, E., Kapur, S., Craik, F.I.M., Moscovitch, M. & Houle, S. (1994). Hemispheric encoding/retrieval asymmetry in episodic memory: Positron emission tomography findings. Proceedings of the National Academy of Sciences, 91, 2016-2020.

Walsh, V. & Perrett, D.I. (1994). Visual attention in the occipitotemporal processing stream of the macaque. Cognitive Neuropsychology, 11(2), 243-263.

Wilding, E. L. (1999). Separating retrieval strategies from retrieval success: an event-related potential study of source memory. Neuropsychologia, 37, 441-454.

Wilding, E. L., Doyle, M. C., & Rugg, M. D. (1995). Recognition memory with and without retrieval of context: An event-related potential study. Neuropsychologia, 33, 743-767.

Wilding, E. L., & Rugg, M. D. (1996). An event-related potential study of recognition memory with and without retrieval of source. Brain, 119, 889-905.

Yamane. S., Kaji, S. & Kawano, K. (1988). What facial features activate face neurons in the inferotemporal cortex of the monkey. Experimental Brain Research, 73, 209-214.

CHAPTER 2

EVOKED POTENTIALS OF RECOGNIZING HAPPY AND NEUTRAL FACES

A version of this chapter had been published in Neuroreport, Vol. 12, 2001

Event-related potential (ERP) studies have associated the recovery of information from personally experienced past events, or episodic retrieval, with two ERP effects: a parietal effect and a frontal effect. These effects have been repeatedly found using emotionally neutral words (Allan, Wilding & Rugg, 1998), emotionally neutral pictures of objects (Schloerschidt & Rugg, 1997), and more recently with emotionally negative words (Maratos, Allan & Rugg, 2000). In the present study, we investigated the lateralization of these effects during the recognition of happy and neutral faces.

The parietal effect refers to the finding that compared with ERP's for correctly rejected new words, ERP's for correctly recognized old words are more positive over parietal scalp regions (Allan et al., 1998). This effect has been attributed to hippocampal-cortical interactions during episodic recovery. The parietal effect is usually left lateralized, which has been attributed to the verbal nature of the materials employed in most studies (Allan et al., 1998; Maratos et al., 2000). However, the parietal effect was also left lateralized in a study using pictures of common objects (Schloerscheidt & Rugg, 1997). One possible explanation is that common objects have names, and hence, they involve some degree of verbal processing. To investigate the parietal effect for nonnameable stimuli, we measured ERP's during the recognition of unfamiliar faces. Since perception and memory of faces are associated with bilateral brain activity in functional neuroimaging studies (Cabeza & Nyberg, 2000), we predicted that the parietal effect for faces would be bilateral rather than left lateralized.

The frontal effect refers to the finding that compared with ERP's for correctly rejected new words, ERP's for correctly recognized old words are more positive over frontal scalp regions (Allan et al., 1998). In contrast with the parietal effect, the frontal effect is usually larger on the right hemisphere. This finding is consistent with functional neuroimaging evidence that frontal activity during episodic retrieval tends to be right lateralized (Nyberg, Cabeza & Tulving, 1996).

It is also possible that the lateralization of the frontal effect depends on emotional nature of the stimuli. For example, the valence hypothesis states that processing of pleasant stimuli is left lateralized whereas processing of unpleasant stimuli is right lateralized (Davidson, 1992), particularly in prefrontal regions (Davidson, 1995). This hypothesis has received considerable empirical support, (e.g., Ali & Cimino, 1997; Canli et al., 1999; Davidson, 1995; Reuter-Lorenz et al., 1983) but a study using ERP's and face stimuli yielded conflicting results (Laurian, Bader, Lanares & Oros, 1991). Moreover, it is unknown if the valence hypothesis applies to memory as well as to perception. To investigate whether the emotional nature of stimuli can modulate the frontal effect, we compared the frontal effect during the recognition of happy and neutral faces. According to the valence hypothesis, happy faces may show a left lateralized frontal effect.

Materials and Methods

Nineteen undergraduate students (9 male, 10 female) participated in the study for course credit. All subjects were right-handed and had normal or corrected-to-normal vision.

The stimuli consisted of 240 black-and-white photos of unfamiliar male and female faces (120 happy, 120 neutral), taken from the Purdue University, the University of Stirling, and the University of Northern British Columbia face databases. Stimulus size was 3.5 by 2.7 inches. The stimuli were coded using the Facial Action Coding System

(Ekman & Friesen, 1978) to ensure that neutral faces did not show any extraneous muscle displacement and that happy faces fulfilled objective criteria for the existence of a smile (Ekman & Friesen, 1975).

In each of 4 blocks, subjects studied 15 happy and 15 neutral faces in random order; they then performed a recognition test that included 30 studied faces (15 happy, 15 neutral) and 30 new faces (15 happy, 15 neutral). The assignment of faces to the four blocks and to the old and new conditions was counterbalanced across subjects. Study and test trials consisted of 5 events: a fixation for 500 ms, a photo of a face for 400 ms, a fixation for 1600 ms, a response selection screen until a response was made, and a blank screen for 1000 ms. The fixation was a 3.5 by 2.7 inches gray rectangle. The response selection screen prompted subjects to make a happy/neutral decision at study or an old/new decision at test. Subjects made their responses by pressing keys with different hands. Hand use was counterbalanced across subjects. At study, subjects were asked to try to remember the faces for a subsequent memory test.

EEG was collected from 32 silver/silver chloride electrodes embedded in a Quikcap (Neurosoft Inc., Sterling, Virginia) electrode cap. Recording locations were based on a variation of the 10/20 international electrode placement system (Jasper, 1958). Sites included frontopolar (FP1, FP2), frontal (F7, F3, FZ, F4, F8), frontocentral (FC3, FZ, FC4), central (C3, CZ, C4), centroparietal (CP3, CPZ, CP4), parietal (P3, PZ, P4), frontotemporal (FT7, FT8), temporal (T7, T8), temporoparietal (TP7, TP8), and occipital (O1, OZ, O2) electrodes. Linked mastoids served as a reference. Horizontal eye movements were monitored with bipolar electrodes on the outer canthus of each eye, and vertical movements were monitored from electrodes placed above and below the left eye.

EEG and EOG were recorded with a sampling rate of 500 samples/second for an epoch of 1700 ms starting 100 ms prior to the onset of a face. Electrode channels were amplified with a filter bandwidth of 0.03 - 50 Hz. Trials with values above 100 mV or below -100 mV, possibly due to blinks or eye movement artifacts, were excluded from the analyses. To ensure an adequate signal to noise ratio in the ERP's, all subjects had more than 16 artifact-free trials per item category.

Results

At study, the proportion of faces classified correctly, according to the classification determined by objective criteria, was similar ($F(1,18) = 1.8$, $p>.05$) for neutral (.92) and happy (.95) faces. At test, the proportion of hits and false alarms was also similar ($F < 1$) for neutral (.72 vs. .27) and happy (.74 vs. .28) faces.

Correctly recognized test faces were averaged according to 4 categories: (1) old neutral faces (neutral hits); (2) new neutral faces (neutral CR's); (3) old happy faces (happy hits); (4) new happy faces (happy CR's). The ERP results for electrodes of interest (parietal and frontal) are shown in Figure 2-1. The parietal effect is usually maximal around 400-600 ms whereas the frontal effect is usually maximal around 800-1200 ms. To investigate both effects using contiguous epochs, we used an early epoch from 375 ms to 750 ms and a late epoch from 750 ms to 1250 ms. Repeated measures analyses of variance (ANOVA's) on average voltage values were performed separately for parietal and frontal electrodes during both the early and late epochs. Each of the four ANOVA's included three within-subject factors: memory type (hits vs. CR's), emotion (happy vs. neutral), and laterality (right vs. left). The statistical measures reported correspond to raw ERP data, but all analyses were confirmed using normalized data (McCarthy & Wood, 1985).

---------- Insert Figure 2-1 about here ----------

We had hypothesized that the parietal effect during the recognition of unfamiliar faces would be bilateral. The results of the ANOVA on the early epoch at parietal sites indicated that the parietal effect was significant (main effect of memory: $F(1, 18) = 26.5$, $p < 0.0001$) and bilateral (laterality x memory interaction: $F(1, 18) < 1$, $p > .05$). These results are illustrated in Figure 2-2. Additionally, it indicated that ERP's for happy faces were more positive on the left than on the right hemisphere (emotion x laterality interaction: ($F(1, 18) = 5.0$, $p < 0.04$) and that this effect did not interact with memory (memory x emotion x laterality interaction: $F(1, 18) = 1.0$, $p > .05$). At frontal sites, there was also a significant main effect of memory type during the early epoch ($F(1, 18) = 26.8$, $p < 0.0001$). There was a tendency for frontal old/new effects to be larger over the left hemisphere for happy faces and over the right for neutral faces, but this result did not reach significance (memory x emotion x laterality interaction, $F(1, 18) = 4.0$, $p = .06$). No other effects were statistically reliable (F's < 1, $p > .05$).

---------- Insert Figure 2-2 about here ----------

We had also hypothesized that the lateralization of the frontal effect would be affected by emotion, in agreement with the valence hypothesis. The results of the ANOVA on the late epoch at frontal sites indicated that the frontal effect was significant (main effect of memory: $F(1, 18) = 6.1$, $p < .03$). Additionally, it demonstrated that the lateralization of the frontal memory effect was affected by emotion (memory x emotion x laterality interaction ($F(1, 18) = 5.3$, $p < .04$). As illustrated by Figure 2-3, this interaction occurred because the frontal effect (hits minus CR's) was larger on the left hemisphere for happy faces and larger on the right hemisphere for neutral faces. At parietal sites during the late epoch there was also a significant main effect of memory type ($\underline{F}(1, 18) = 7.7$, $\underline{p} < 0.02$). No other effects approached significance (all other $\underline{F}$'s < 2, $\underline{p} > .05$).

---------- Insert Figure 2-3 about here ----------

Discussion

The results confirmed our prediction that the parietal effect during the recognition of unfamiliar faces would be bilateral. They also confirmed our prediction that the lateralization of the frontal effect would be modulated by emotion in accordance with the valence hypothesis.

The bilateral parietal effect observed in this study could indicate that the left-lateralization of the parietal effect in previous ERP studies reflected the verbal nature of stimuli employed, i.e., words (Allan et al., 1998; Maratos et al., 2000) and nameable objects (Schloerscheidt & Rugg, 1997). If the ERP reflects cortico-hippocampal interactions (Allan et al., 1998), then a bilateral parietal effect could reflect bilateral hippocampal involvement. This idea is consistent with functional neuroimaging evidence of bilateral hippocampal activity during face processing (Cabeza & Nyberg, 2000). The parietal effect was not modulated by emotion. There was a significant emotion by laterality interaction during the same epoch; in accordance with the valence hypothesis, happy faces evoked larger positivities over the left hemisphere while neutral faces evoked bilateral positivities. However, the absence of a significant interaction between memory type, emotion, and laterality suggests that this effect was not sensitive to memory and may be perceptual.

The frontal effect was larger on the left hemisphere for happy faces but on the right hemisphere for neutral faces. The discovery of a left-lateralized memory effect for happy faces is congruent with the valence hypothesis. The finding of a right-lateralized memory effect for neutral faces is difficult to interpret. This result could signify that, like negative stimuli, the processing of neutral stimuli may also be lateralized to the right hemisphere. Alternatively, it may reflect the right-lateralization of prefrontal activity during episodic retrieval typically found in functional neuroimaging studies (Nyberg et al., 1996).

The present findings are the first to demonstrate that ERP's of face recognition can be modulated by facial expression in a manner that is predicted by the valence hypothesis of emotion processing. Whereas memory effects for happy faces were left lateralized over frontal areas, memory effects for neutral faces were right lateralized. Furthermore, this effect was dissociable from the parietal old/new effect in terms of both latency and topography. Whereas the parietal effect was not sensitive to emotion, the frontal effect was, possibly indicating that general recognition and emotion-specific recognition are dissociable neural processes. The fact that the emotion-sensitive memory effect over frontal regions occurred later than effect over parietal regions could signify that emotion-sensitive memory processes are engaged later than general recognition processes. The functional significance of these effects, including their cerebral generators, remains unclear. Nevertheless, the present results strongly suggest that facial emotion can modulate brain activity associated with face recognition.

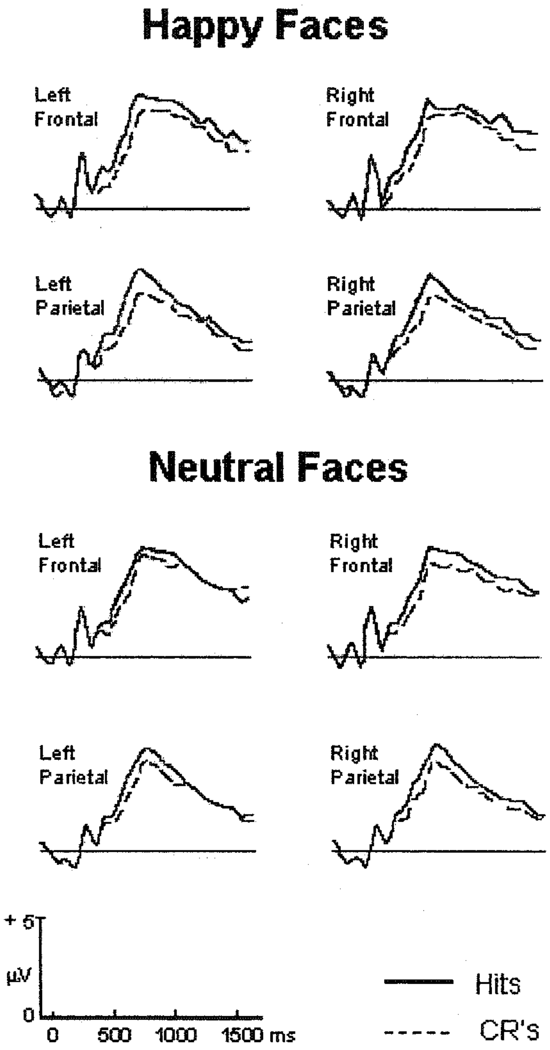

Figure 2-1. Average voltage values (microvolts) obtained during recognition of happy and neutral faces at parietal and frontal electrodes on both hemispheres

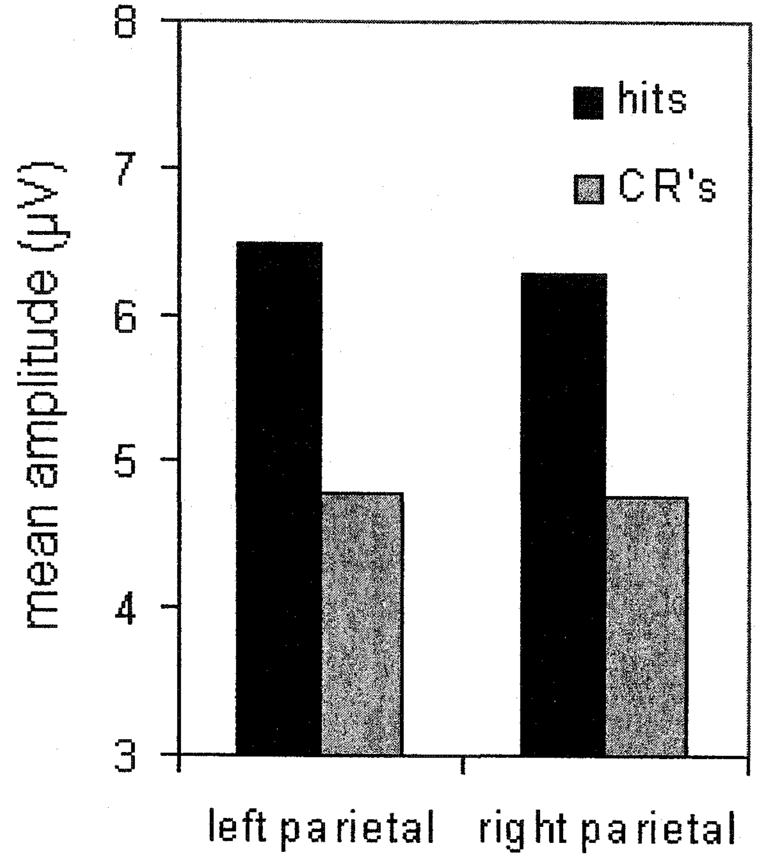

igure 2-2. Mean amplitudes (microvolts) for hits and CR's (happy and neutral faces combined) at parietal electrodes from 375 ms to 750 ms.

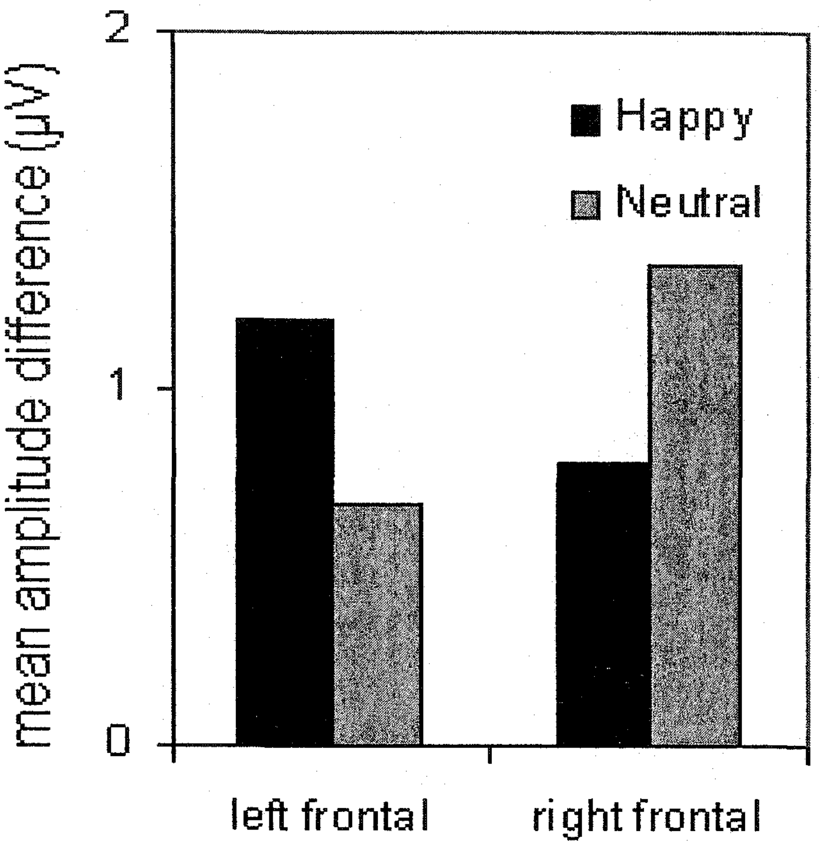

Figure 2-3. Mean difference amplitudes (hits minus CR's, in microvolts) for happy and neutral faces at frontal electrodes from 750 ms to 1250 ms.

References

Ali, N. & Cimino, C.R. (1997). Hemispheric lateralization for perception and memory for emotional verbal stimuli in normal individuals. Neuropsychology, 11(1), 114-125.

Allan, K., Wilding, E.L. & Rugg, M.D. (1998). Electrophysiological evidence for dissociable processes contributing to recollection. Acta Psychologica, 98, 231-252.

Cabeza, R. & Nyberg, L. (2000). Imaging cognition II: An empirical review of 275 PET and fMRI studies. Journal of Cognitive Neuroscience, 12(1), 1-47.

Canli, T., Desmond, J.E., Zhao, Z., Glover, G. & Gabrielli, J.D.E. (1999). Hemispheric asymmetry for emotional stimuli detected with fMRI. Neuroreport, 9, 3233-3239.

Davidson, R. J. (1992). Emotion and affective style: Hemispheric substrates. Psychological Science, 3(1), 39-43.

Davidson, R.J. (1995). Cerebral asymmetry, emotion, and affective style. In R.J. Davidson and K. Hugdahl (Eds.), Brain asymmetry (pp. 361-387). Cambridge, MA: MIT Press.

Ekman, P. & Friesen, W.V. (1978). Facial action coding system. Palo Alto: Consulting Psychologists Press.

Ekman, P. & Friesen, W.V. (1975).Unmasking the face. Palo Alto: Consulting Psychologists Press.

Jasper, H.H. (1958). The ten-twenty system of electrode system of the International Federation. Electroencephalography and Clinical Neurophysiology, 20, 371-375.

Laurian, S., Bader, M., Lanares, J., & Oros, L. (1991). Topography of event-related potentials elicited by visual emotional stimuli. International Journal of Psychophysiology, 10, 231-238.

Maratos, E.J., Allan, K. & Rugg, M.D. (2000). Recognition memory for emotionally negative and neutral words: An ERP study. Neuropsychologia, 38, 1452-1465.

McCarthy, G. & Wood, C.C. (1985). Scalp distributions of event-related potentials: An ambiguity associated with analysis of variance models. Electroencephalography and Clinical Neuropsychology, 62, 203-208.

Nyberg, L., Cabeza, R., & Tulving, E. (1996). PET studies of encoding and retrieval: The HERA model. Psychonomic Bulletin and Review, 3(2), 135-148.

Reuter-Lorenz, P.A., Givis, R.P., & Moscovitch, M. (1983). Hemispheric specialization and the perception of emotion: Evidence from right-handers and from inverted and non-inverted left-handers. Neuropsychologia, 21(6), 687-692.

Schloerscheidt, A.M. & Rugg, M.D. (1997). Recognition memory for words and pictures: an event-related potential study. Neuroreport, 8, 3281-3285.

CHAPTER 3

DISSOCIATING THE NEURAL CORRELATES OF ITEM AND CONTEXT MEMORY:

AN ERP STUDY OF FACE RECOGNITION

A version of this chapter had been published in the Canadian Journal of Experimental Psychology, Vol. 55, 2001

Remembering personally experienced past events, or episodic memory retrieval, involves both item and context retrieval. Item retrieval refers to remembering what events happened, while context retrieval refers to remembering when (temporal-order memory), where (spatial memory) or how (source memory) they happened. The distinction between item and context retrieval is important because lesion, functional neuroimaging, and electrophysiological evidence suggests that these two components of episodic retrieval involve different neural mechanisms. In the present study, we investigated the neural correlates of item and context retrieval using event-related potentials (ERP's).

Lesion evidence has associated item retrieval with the medial-temporal lobe (MTL), and context retrieval with the prefrontal cortex (PFC). In a classic study by Corsi and Milner (cited by Milner, 1971), patients with MTL lesions were impaired in item retrieval (recognition test) but not in context retrieval (temporal-order test) whereas patients with PFC lesions were impaired in context but not in item retrieval. Although this double dissociation may not be generalizable to other forms of item and context memory, numerous studies with human (e.g., Janowsky, Shimamura & Squire, 1989; Shimamura, Janowsky & Squire, 1990) and nonhuman primates (e.g., Petrides, 1991) have found that PFC lesions produce greater deficits in context retrieval than in item retrieval.

Functional neuroimaging evidence is also consistent with the idea that item retrieval is primarily associated with MTL function, and context retrieval, with PFC function. A positron emission tomography (PET) study found that MTL was more activated for item retrieval than for context retrieval, whereas PFC was more activated for context than for item retrieval (Cabeza et al., 1997). This double dissociation was significant when differences in task difficulty were statistically removed, and when item and context conditions with similar levels of performance were compared. Greater PFC activity for context than for item retrieval was also found in several functional magnetic resonance imaging (fMRI) studies (e.g., Eyler Zorilla, Aguirre, Zarahn, Cannon & D'Esposito, 1996; Nolde, Johnson & D'Esposito, 1998). Moreover, a recent PET study dissociated item and context memory as a function of aging. Compared to younger adults, older adults showed weaker PFC activity during context memory, but similar MTL activity during item memory (Cabeza et al., 2000). This result is consistent with evidence that age-related memory deficits are more pronounced for context than for item retrieval (Spencer & Raz, 1995) and that age-related atrophy is more pronounced in PFC than in MTL regions (Raz et al., 1997).

Electrophysiological evidence concerning the neural correlates of item and context memory has been provided by event-related potentials (ERP's) studies of word recognition (for a review, see Allan, Wilding, & Rugg, 1998) Compared with ERP's for correctly rejected new words (correct rejections or CR's), ERP's for correctly recognized studied words (hits) have been associated with two positive-going effects: the parietal effect, which is maximal over parietal electrodes between 400-800 ms, and the frontal effect, which is maximal over frontal electrodes between 700-1200 ms. The parietal effect is assumed to reflect MTL activity indirectly through its interactions with cortical areas, and the frontal effect is assumed to reflect PFC activity (Allan et al., 1998). Thus, on the basis of lesion (Milner, 1971) and functional neuroimaging (Cabeza et al., 1997) evidence, one would expect the frontal effect to be primarily sensitive to context retrieval, and the parietal effect, to item retrieval. ERP data tend to support these predictions.

In the ERP study by Wilding and Rugg (1996; see also Wilding Doyle & Rugg, 1995), subjects heard words in either a male or a female voice. At test, they read words and decided if the words were old or new, and in the case of words classified as old, if they were heard at study in a female or a male voice. This method allows for a distinction between trials associated with correct item and context decisions (hit-hits), and trials associated with correct item but incorrect context decisions (hit-misses). Although both parietal and frontal effects were larger for hit-hits than for hit-misses, the effect was more prominent in the frontal effect, particularly in the right hemisphere. Moreover, Senkfor and Van Petten (1998) found that the parietal effect was similar for a recognition test (old vs. new) and for a source test (same-voice vs. different-voice vs. new), whereas the frontal effect occurred only for the source task.

There are at least two unresolved issues concerning ERP's for item and context memory. First, it is unknown if the sensitivity of the frontal effect to context retrieval found in word recognition studies can be generalized to nonverbal stimuli. Although frontal effects have been found for nonverbal stimuli, such as pictures of common objects (Schloerscheidt & Rugg, 1997) and unfamiliar faces (Graham & Cabeza, in press), these studies did not include context retrieval conditions. Faces are suitable stimuli for the study of context retrieval because when humans remember faces, we extract and retrieve identity information (who the person is) and we also extract and retrieve contextual information, or episodic detail, such as facial expression. Therefore, the retrieval of facial expression is analogous to context retrieval seen with stimuli such as words. If the modulatory influence of context retrieval on the frontal effect reflects general context retrieval operations, then it should occur not only for verbal but also for nonverbal stimuli, such as faces.

Second, it is unclear if the parietal effect is actually insensitive to context retrieval. The assumption that the parietal effect reflects primarily MTL-mediated item retrieval implies that it should not be sensitive to variations in context retrieval, which supposedly depends on PFC function. Yet, the parietal effect has shown significant hit-hit vs. hit-miss differences in a few word recognition studies (e.g., Wilding & Rugg, 1996; Wilding, 1999). One possible explanation is that these differences reflect an enhancing effect of context retrieval on item retrieval. It is reasonable to expect that recovering context information (i.e., where, when, and how it happened) facilitates the recovery of item information (i.e., what happened). This effect may be more likely in the case of familiar stimuli, such as common words, because they are associated with many different contexts outside the experiment. Hence, deciding whether or not they occurred during the experiment can be facilitated by the recovery of specific details of the experimental context (e.g., speakers' voice). In contrast, stimuli that are first encountered during the experiment are more likely to become uniquely associated with the experimental context, and to be retrieved independently of context information recovery. If this idea is correct, then novel stimuli such as unfamiliar faces should yield a parietal effect that is insensitive to context retrieval.

The goal of the present study was to investigate the foregoing issues. Subjects studied unfamiliar faces in happy or neutral expressions, and at test, they were presented studied faces in the same or different expression intermixed with nonstudied faces, and for each face, they decided if it was studied in the same expression, studied in a different expression, or nonstudied. Thus, even if face identify and face expression were simultaneously encoded during study, the test allowed us to differentiate hit/hits, where both the face and the facial expression were remembered, from hit/misses, where the identity was correctly remembered but facial expression was not.

We expected three main results. First, we expected that, as in the case of the frontal effect for words (Wilding & Rugg, 1996), the frontal effect for faces would be larger for hit-hits than for hit-misses. We reasoned that if the modulatory influence of context retrieval on the frontal effect reflects general retrieval operations, then it should occur not only for verbal but also for nonverbal stimuli, such as faces. Second, we expected that the parietal effect for faces would be similar for hit-hits and hit-misses. This prediction was based on the assumption that if the contribution of context retrieval to item retrieval is reduced in the case of novel stimuli, then the parietal effect for unfamiliar faces should be relatively insensitive to context retrieval. Finally, we expected that the previous two results would yield a dissociation between frontal and parietal effects: the frontal effect would be sensitive to context retrieval whereas the parietal effect would be sensitive to item retrieval. Such a dissociation would support the notion that the parietal effect is an index item retrieval associated with MTL function, and the frontal effect is an index of context retrieval associated with PFC function, converging with lesion (Milner, 1971) and functional neuroimaging (Cabeza et al., 1997, 2000) evidence.

Method

Participants

Seventeen female undergraduate students participated in the study for course credit. All subjects were right-handed and had normal or corrected-to-normal vision.

Materials

The stimuli consisted of 240 black-and-white photos of unfamiliar male and female faces (120 happy, 120 neutral), taken from the Purdue University, the University of Stirling, and the University of Northern British Columbia face databases. Stimulus size was 3.5 by 2.7 inches. The stimuli were coded using the Facial Action Coding System (Ekman & Friesen, 1978) to ensure that neutral faces did not show any extraneous muscle displacement, and that happy faces fulfilled objective criteria for the existence of a smile (Ekman & Friesen, 1975).

Procedure

In each of four blocks, subjects studied 20 happy and 20 neutral faces in random order, and then they performed a recognition test that included: (1) 20 studied faces with the same expression as in the study phase (10 happy, 10 neutral); (2) 20 studied faces with a different facial expression than in the study phase (10 happy, 10 neutral); and (3) 20 new faces (10 happy, 10 neutral). The assignment of faces to the four blocks was counterbalanced across subjects. Study and test trials consisted of five events: a fixation for 500 ms, a photo of a face for 400 ms, a fixation for 1600 ms, a response selection screen until a response was made, and a blank screen for 1000 ms. The fixation was a 3.5 by 2.7 inches gray rectangle. The response selection screen prompted subjects to make a decision. At study, subjects made a happy/neutral decision by pressing keys with different hands, and also tried to remember the faces and their facial expressions for a subsequent memory test. At test, they made 3-choice recognition decisions: (a) studied with same expression, (b) studied with different expression, or (c) new. Responses (a) and (b) were assigned to one hand, and response (c) to the other hand. Left/right hand use was counterbalanced across subjects at both study and test.

ERP Methods

EEG was collected from 32 silver/silver chloride electrodes embedded in an electrode cap. Recording locations were based on a variation of the 10/20 international electrode placement system (Jasper, 1956). Sites included frontopolar (FPI, FP2), frontal (F7, F3, FZ, F4, F8), frontocentral (FC3, FZ, FC4), central (C3, CZ, C4), centroparietal (CP3, CPZ, CP4), parietal (P3, PZ, P4), frontotemporal (FT7, FT8), temporal (T7, T8), temporoparietal (TP7, TP8), and occipital (O1, OZ, O2) electrodes. Linked mastoids served as a reference. Horizontal eye movements were monitored with bipolar electrodes on the outer canthus of each eye, and vertical movements were monitored from electrodes placed above and below the left eye. EEG and EOG were recorded with a sampling rate of 500 samples/second for an epoch of 1700 ms starting 100 ms prior to the onset of a face. Electrode channels were amplified with a filter bandwidth of 0.03 - 50 Hz. Trials with values above 100 μV or below -100 μV, possibly due to blinks or eye movement artifacts, were excluded from the analyses. To ensure an adequate signal to noise ratio in the ERP's, subjects that had less than 16 artifact-free trials per item category were replaced.

Results

Behavioral data

At study, the proportion of faces correctly classified as happy or neutral (.97) was very high, suggesting that subjects could easily distinguish between the two expressions.

At test, the proportion of "studied" responses (studied-same plus studied-different) was significantly higher [t(16) = 12.6, $p < .001$] for studied faces (.76) than for new faces (.36). Among studied faces correctly recognized as studied, the proportion of faces whose expression was correctly classified as "same" or "different" (.69) was reliably greater [t(16) = 12.0, $p < .001$] than the proportion of faces whose expression was misclassified (.31). Thus, there was a good level of accuracy for both item and context memory retrieval.

ERP data

ERP's for correctly recognized test faces were averaged according to three categories. (1) Hits-hits: studied faces correctly recognized as studied and correctly classified as same or different expression; i.e., items associated with both successful item retrieval and successful context retrieval. (2) Hit-misses: studied faces correctly recognized as studied but incorrectly classified as same or different; i.e., items associated with successful item retrieval but unsuccessful context retrieval. (3) Correct Rejections (CR's): new faces correctly classified as new; i.e., items associated with neither item retrieval nor context retrieval.

The ERP results for electrodes of interest (parietal and frontal) are shown in Figure 3-1. Consistent with previous research, the parietal effect was maximal around 500-800 ms, and the frontal effect around 700-900 ms. To investigate both effects using contiguous epochs, we measured parietal ERP's from 550 ms to 750 ms, and frontal ERP's, from 750 ms to 900 ms. Differences in ERP amplitude were analyzed using repeated-measures ANOVA's corrected for inhomogeneity of covariance (Geisser-Greenhouse), and t-tests corrected for multiple comparisons (Bonferroni). The statistical measures reported correspond to raw ERP data, but all analyses were confirmed using normalized data (McCarthy & Wood, 1985).

---------- Insert Figure 3-1 about here ----------

Our first prediction was that the frontal effect for unfamiliar faces would be sensitive to successful contextual retrieval. To test this prediction, we conducted a 3 (item type: hit-hit vs. hit-miss vs. CR) x 2 (hemisphere: left vs. right) ANOVA on frontal ERP's. The analysis yielded a significant main effect of item type [$F(1.6, 26.1) = 11.0$, $p < .001$] but no reliable effect of hemisphere [$F(1, 16) = 1.5$, $p > .05$], or item type x hemisphere interactions ($F < 1$). Pairwise comparisons indicated that ERP's for hit-hits (4.9 μV) were reliably more positive than those for hit-misses [3.4 μV; $t(16) = 5.1$. $p < .001$] and CR's [4.2 μV; $t(16) = 2.8$. $p < .015$], with no difference between hit-misses and CR's [$t(16) = -1.9$. $p > .05$]. Thus, confirming our first prediction, the frontal effect was larger for hit-hits than for hit-misses

Our second prediction was that the parietal effect for unfamiliar faces would not be sensitive to contextual retrieval. To test this prediction, we conducted a 3 (item type: hit-hit vs. hit-miss vs. CR) x 2 (hemisphere: left vs. right) ANOVA on parietal ERP's. This analysis yielded a significant main effect of item type [$F(1.6, 25.7) = 9.5$, $p < 0.002$], but no reliable effect of laterality ($F < 1$), or item type x hemisphere interactions ($F < 1$). Pairwise comparisons indicated that ERP's for hit-hits were reliably more positive than ERP's for CR's [4.4 vs. 3.3 μV; $t(16) = 4.6$. $p < .001$] and hit-misses were also more positive than CR's [4.1 vs. 3.3 μV; $t(16) = 2.7$. $p < .015$], but there was no difference between hit-hits and hit-misses [$t(16) = 1.2$. $p > .05$]. Therefore, confirming our second prediction, the parietal effect was similar for hit-hits and hit-misses.

Finally, our third prediction was that the frontal effect would be sensitive to context retrieval but that the parietal effect would be sensitive only to item retrieval. To test this prediction, we subtracted CR's from hit-hits and hit-misses and conducted a 2 (item type: hit-hits vs. hit-misses) x 2 (region: frontal vs. parietal) ANOVA, collapsing over hemispheres. Confirming our third prediction, this analysis yielded a significant item type x region interaction ($F(1,16) = 11.0$, $p < .01$). This interaction can be seen in Figure 3-2 where mean amplitude differences for hit-hits (hit-hits minus CR's) are larger than those for hit-misses (hit-misses minus CR's) at frontal sites, whereas amplitude differences at parietal sites are equivalent in magnitude.

---------- Insert Figure 3-2 about here ----------

The two scalp distributions were also directly contrasted in a 2 (distribution: frontal vs. parietal) x 3 (item type: hit-hits vs. hit-misses vs. CR's) x 2 (hemisphere: right vs. left) ANOVA in order to demonstrate that the effects did come from different generators. This analysis gave rise to a significant item type x distribution interaction [$F(1.6, 24.9) = 9.4$, $p = .002$]. This result confirmed the results of earlier analyses, which found that the parietal effect was sensitive to item retrieval, being similar for hit-hits and hit-misses but smaller for CR's, while the frontal effect was sensitive only to context retrieval with hit-hits larger than hit-misses and CR's.

Discussion

The study provided three main results. First, there was a dissociation between the frontal effect, which was sensitive to context retrieval, and the parietal effect, which was sensitive only to item retrieval. This result is consistent with the results from lesion and functional neuroimaging research and supports the idea that context retrieval is associated with PFC activity, and item retrieval, with MTL function. Second, the frontal effect for unfamiliar faces was greater for hit-hits than hit-misses. This finding supports the generalizability of the modulatory influence of context retrieval on the frontal effect previously found in word recognition studies. Finally, the parietal effect for unfamiliar faces was similar for hit-hits and hit-misses. This finding supports the idea that the sensitivity of the parietal effect to context retrieval in word recognition studies reflected an influence of context retrieval on item retrieval that is attenuated when using novel stimuli. Below, we discuss the three results, and consider other related issues.

Dissociation between frontal and parietal effects

The main finding of the present study was a clear dissociation between the influences of context and item retrieval on frontal and parietal effects: the frontal effect was sensitive to context retrieval, whereas the parietal effect was sensitive to item but not to context retrieval. This dissociation is consistent with lesion (e.g., Milner, 1972) and functional neuroimaging evidence (Cabeza et al., 1997, 2000), and supports the idea that context memory is primarily dependent on PFC function, whereas item retrieval is primarily dependent on MTL function.

This idea is central to Moscovitch's (1992, 1994) component process model of memory. According to this model, episodic memory is mediated by two different systems: an associative MTL system and a strategic PFC system. The associative system automatically encodes information that has been consciously apprehended, and it automatically retrieves it whenever an appropriate cue is presented. This system cannot distinguish veridical from false memories, organize the retrieval output, or guide a retrieval search. These "intelligent" functions are provided by the strategic PFC system.

This view predicts that age effects should be more pronounced on episodic memory tests that are more dependent on the strategic system, such as context memory, than on tasks that rely primarily on the associative system, such as recognition. As mentioned previously, a recent PET study found age-related decreases in PFC activation during context retrieval together with age-invariant MTL activity during item retrieval (Cabeza et al., 2000).

Instead of a dissociation between item and context retrieval, the present results could also be described as a dissociation between familiarity and recollection (e.g., Atkinson & Juola, 1974; Mandler, 1980). Familiarity refers to the feeling—sometimes very strong—that an item or event occurred in the past, even though no other related contextual information can be retrieved. Recollection, in contrast, refers to the retrieval of specific information about the event, such as its spatiotemporal context or other associated information. Thus, familiarity would involve item retrieval without context retrieval, whereas recollection would involve both item and context retrieval. This does not imply that familiarity and recollection involve a similar level of item retrieval. As mentioned in the introduction, recovering contextual information may facilitate the recovery of item information, and hence, recollection is likely to involve a higher level of item retrieval than familiarity. This would explain why functional neuroimaging studies have found activation differences between recollection and familiarity (e.g., remember/know paradigm) not only in PFC (e.g., Henson et al., 1999), but also in MTL (Elridge et al., 2000).

However, it could be argued that the hit-hit/hit-miss differences found in the present study do not reflect differences between item and context retrieval, or between familiarity and recollection, but rather priming effects. Since at test half of the faces were presented in the same expression and half in a different expression, it is possible that the former generated more perceptual priming than the latter. Yet, it is unlikely such priming effects could account for hit-hit/hit-miss differences. First, hit-hit items included faces both in the same and in different expressions, and this applies also to hit-miss items. Thus, if priming effects occurred, they affected both hit-hit and hit-miss items. Second, in ERP studies, priming effects have not been associated with the parietal or the frontal effect but with a different electrophysiological marker. (Rugg et al., 1998). Finally, hit-hit/hit-miss ERP differences have been found in conditions in which perceptual priming was attenuated or eliminated by a study-test modality shift (Wilding & Rugg, 1996).

Frontal effect

The present study provided the first evidence that the modulatory influence of context retrieval on the frontal effect previously found for words (Wilding & Rugg, 1996; Senkfor & Van Petten, 1998) can be also found for faces. The fact that this modulatory influence can be found when stimuli are verbal and familiar as well as when they are nonverbal and novel stimuli, and when context information refers to voice quality as well as when it refers to facial expressions, suggest that this influence reflects general context retrieval operations.

The present results are also relevant to the issue of whether the frontal ERP effect reflects success in retrieving context information (Wilding and Rugg, 1996) or only the attempt to retrieve this information regardless of the outcome of the search (Senkfor and Van Petten, 1998). In Wilding and Rugg's (1996) study, the frontal effect was significantly larger for hit-hits than for hit-misses (Experiment 2), whereas in Senkfor and Van Petten's (1998) study, it was larger in a context retrieval task than in an item retrieval task, but did not show reliable hit-hit/hit-miss differences. Thus, Wilding and Rugg's results are consistent with the retrieval success hypothesis, whereas Senkfor and Van Petten's results are consistent with the retrieval attempt hypothesis. The present results favor the success hypothesis because the frontal effect showed clear hit-hit/hit-miss differences.

It is important to note some differences between the frontal effect in the present study and in word recognition studies. First, there is a difference in the duration of the frontal effect. Whereas in Wilding and Rugg's studies (Wilding et al., 1995; Wilding & Rugg, 1996), the frontal effect extended beyond 1400 ms, in the present study, it disappeared at about 900 ms. This inconsistency may reflect differences between processing words vs. faces, between retrieving voices vs. expressions, or it may reflect a difference between the double 2-choice response (old/new followed by male/female) used in Wilding and Rugg's studies versus the single 3-choice response (same/different/new) used in the present study. However, Senkfor and Van Petten (1998) used a single 3-choice response and found a frontal effect lasting at least 1200 ms.

Additionally, whereas in Wilding and Rugg's (1996) study hit-misses showed a frontal effect, i.e., they were more positive than CR's, in the present study they were nonsignificantly different from CR's, and even showed a trend to be *less* positive than CR's. This trend is similar to that found in Wilding (1999). The lack of a significant difference between hit-misses and CR's actually fits very well with the assumption that the frontal effects reflect successful context retrieval because both hit-miss and CR events occur in the absence of context retrieval. What is unclear is why other studies found a frontal effect for hit-misses (Wilding et al., 1995; Wilding & Rugg, 1996). We do not have a good explanation of the trend of hit-misses to be less positive than CR's in our study, but as noted above, this difference was not reliable.

Parietal effect

The parietal effect was similar for hit-hits and hit-misses. This implies that the parietal effect was insensitive to context retrieval and supports the idea that the parietal effect is an index of MTL-mediated item retrieval. This finding is counter to those found in word recognition studies, which have reported parietal effects that are sensitive to the retrieval of context (Wilding & Rugg, 1996, Wilding, 1999). As mentioned in the introduction, one explanation for the inconsistency of the parietal effect for context memory may be due to an enhancement of item retrieval, which is facilitated by successful context retrieval.

This enhancement may be sensitive to the nature of the stimuli used in an experiment, being more likely with highly familiar stimuli, such as words, which are encountered in many extra-experimental contexts. Therefore, the decision about whether or not a word occurred during an experiment (item memory) may be aided by the retrieval of contextual information. In contrast, the stimuli used in this experiment were faces that were unfamiliar to the subjects and therefore were not associated with extra-experimental contexts. Because the faces were first encountered during the experiment, item retrieval would be less likely to be improved by context retrieval and more likely to occur independently. However, it is possible that the finding of similar parietal effects for hit-hit and hit-miss items is not related to the use of novel stimuli but to the use of faces. Further research about this issue is warranted.

The finding that the parietal effect was similar for hit-hit and hit-miss items does not imply that MTL is not involved in context memory. First, the parietal effect is a very indirect measure of MTL function, and the lack of hit-hit/hit-miss differences on this effect may simply reflect a lack of sensitivity. Second, we investigated a particular form of context retrieval (i.e., retrieval of face expression during face recognition), and the results may not generalize to other forms of context memory. For example, there is a considerable amount of evidence that certain MTL components and their associated structures play an important role in spatial memory (for a review, see Aggleton & Brown, 1999). Finally, the present results apply only to *retrieval*, and do not provide information concerning the role of MTL during the *encoding* of context information.

It is important to note that the magnitude of the parietal effect was somewhat smaller than the one typically found in ERP studies of word recognition (e.g., Allan et al., 1998). Given the high false alarm rate obtained in the study, it is possible that guessing was a factor in the attenuated magnitude of the parietal effect. However, although small, the parietal effect was remarkably consistent across subjects.

Conclusions

The present study dissociated the neural correlates of item and context memory. The ERP parietal effect was sensitive to item retrieval, whereas the frontal effect was sensitive to context retrieval. These results are consistent with lesion and functional neuroimaging evidence associating item retrieval primarily with MTL function and context retrieval primarily with PFC function.

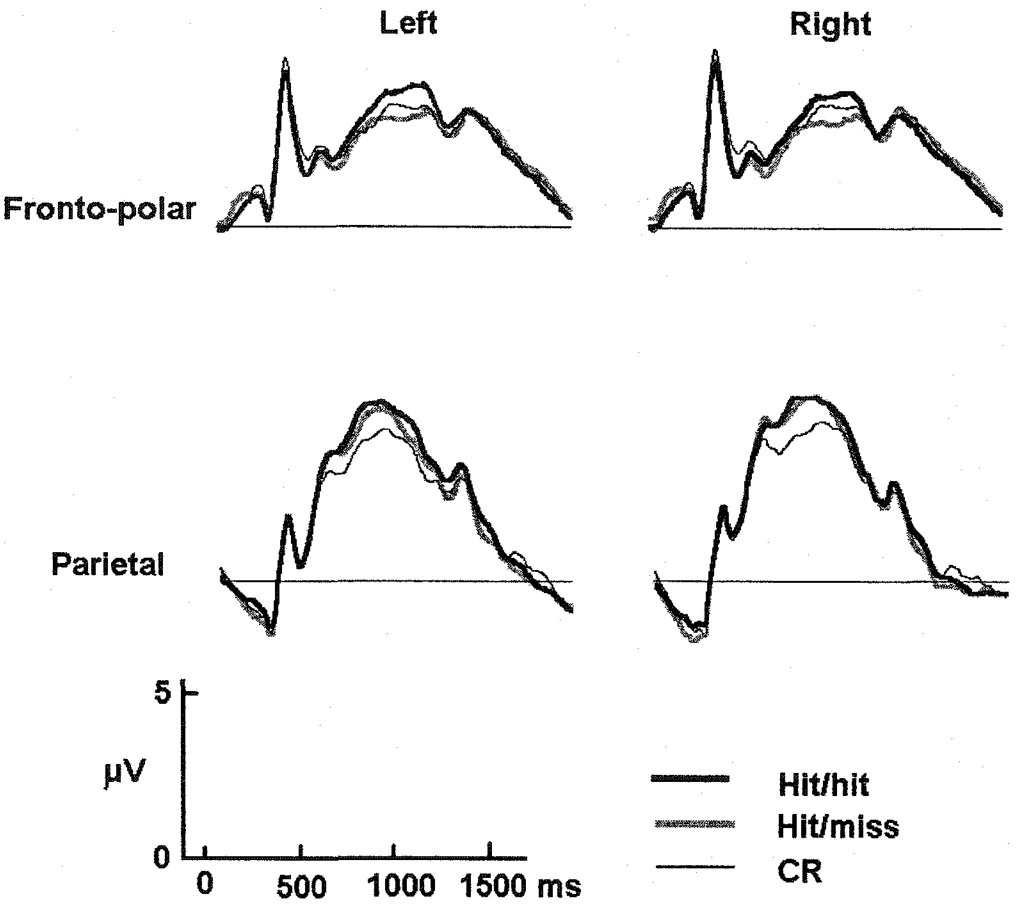

Figure 3-1. Average voltage values (µV) obtained during recognition of happy and neutral faces at parietal and frontal electrodes on both hemispheres.

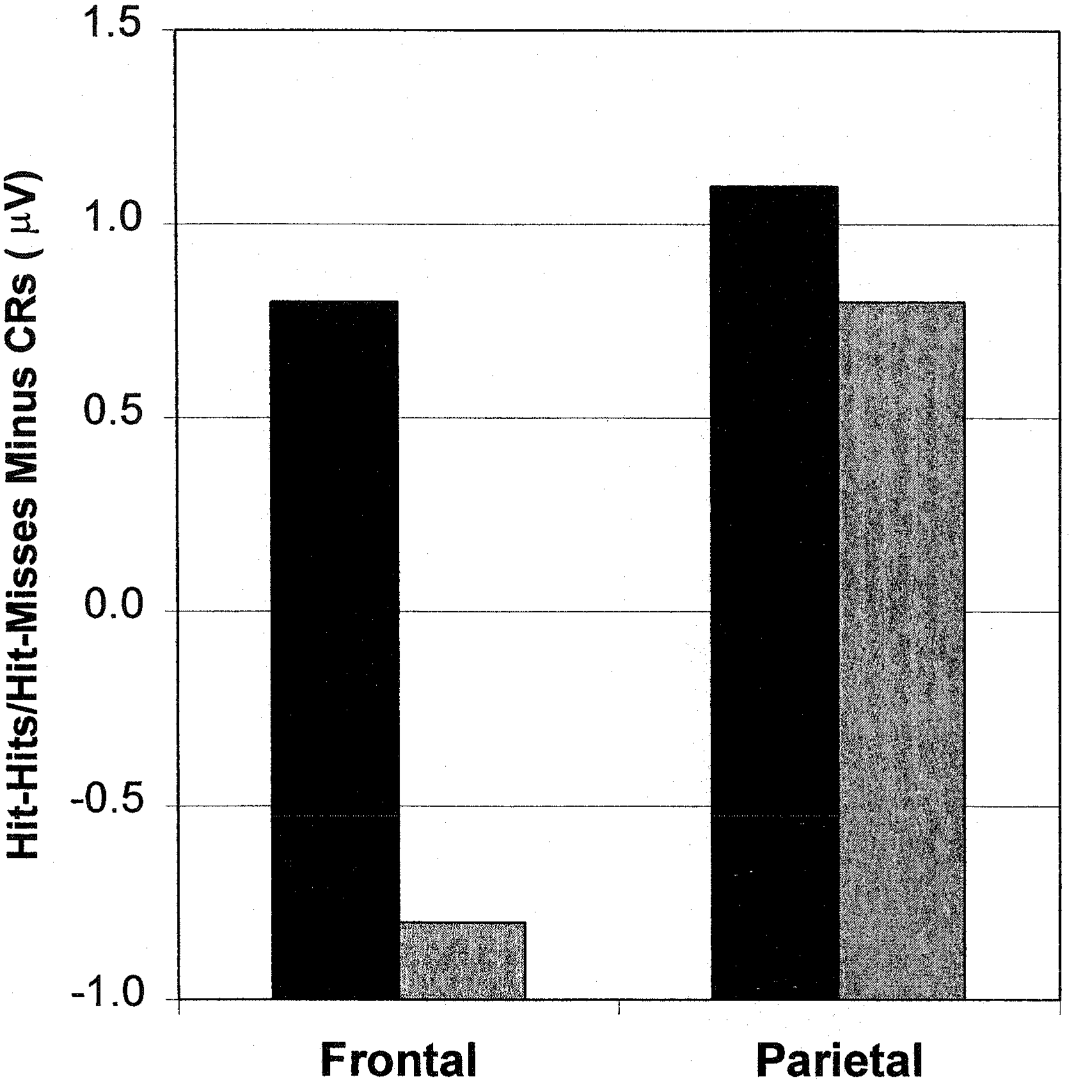

Figure 3-2. Mean difference amplitudes for hit-hits (hit-hits minus CR's) and for hit-misses (hit-misses minus CR's) obtained at frontal and parietal electrodes (µV), collapsed across epochs and hemispheres.

References

Aggleton, J. P., & Brown, M. W. (1999). Episodic memory, amnesia, and the hippocampal-anterior thalamic axis. Behavioral and Brain Sciences, 22, (425-289).

Allan, K., Wilding, E.L. & Rugg, M.D. (1998). Electrophysiological evidence for dissociable processes contributing to recollection. Acta Psychologica, 98, 231-252.

Atkinson, R. C. and J. F. Juola (1974). Search and decision processes in recognition memory. Contemporary developments in mathematical psychology. D. H. Krantz, R. C. Atkinson, R. D. Luce and P. Suppes. San Francisco, Freeman. 1: 243-293.

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

Cabeza, R., Anderson, N.D., Houle, S., Mangels, J.A., & Nyberg, L. (2000). Age-related differences in neural activity during item and temporal-order memory retrieval: A positron emission study. Journal of Cognitive Neuroscience, 12 (1), 197-206.

Cabeza, R., Mangels, J.A., Nyberg, L., Habib, R., Houle, S., McIntosh, A.R., & Tulving, E. (1997). Brain regions differentially involved in remembering what and when. Neuron, 19, 863-870.

Cabeza, R. & Nyberg, L. (2000). Imaging cognition II: An empirical review of 275 PET and fMRI studies. Journal of Cognitive Neuroscience, 12(1), 1-47.

Ekman, P. & Friesen, W.V. (1975). Unmasking the face. Palo Alto, CA: Consulting Psychologists Press.

Ekman, P. & Friesen, W.V. (1978). Facial action coding system. Palo Alto, CA: Consulting Psychologists Press.

Eldridge, L. L., B. J. Knowlton, et al. (2000). "Remembering episodes: a selective role for the hippocampus during retrieval." Nature Neuroscience, 3(11), 1149-1152.

Eyler Zorilla, L.T., Aguirre, G. K., Zarahn, E., Cannon, T.D., & D'Esposito, M. (1996). Activation of the prefrontal cortex during judgments of recency: A functional MRI study. Neuroreport, 7, 2803-2806.

Graham, R. & Cabeza, R. (2001). Event-related potentials of recognizing happy and neutral faces. Neuroreport, 12(2), 245-248.

Henson, R. N., M. D. Rugg, et al. (1999). "Recollection and familiarity in recognition memory: an event-related functional magnetic resonance imaging study." Journal of Neuroscience, 19(10): 3962-72.

Janowsky, J.S., Shimamura, A.P. & Squire, L.R. (1989). Source memory impairment in patients with frontal lobe lesions. Neuropsychologia, 27(8), 1043-1056.

Jasper, H.H. (1958). The ten-twenty system of electrode system of the International Federation. Electroencephalography and Clinical Neurophysiology, 20, 371-375.

McCarthy, G. & Wood, C.C. (1985). Scalp distributions of event-related potentials: An ambiguity associated with analysis of variance models. Electroencephalography and Clinical Neuropsychology, 62, 203-208.

Milner, B. (1971). Interhemispheric differences in the localization of psychological processes in man. British Medical Bulletin, 27, 272-277.

Moscovitch, M. (1992). Memory and working with memory: A component process model based on modules and central systems. Journal of Cognitive Neuroscience, 4, 257-267.

Moscovitch, M. (1994). Models of consciousness and memory. In M. Gazzaniga (Ed.), The cognitive neurosciences (pp. 1341-1356). Cambridge, MA: MIT Press.

Nolde, S. F., Johnson, M. K., & D'Esposito, M. (1998). Left prefrontal activation during episodic remembering: An event-related fMRI study. Neuroreport, 9, 3509-3514.

Petrides, M. (1991). Functional specialization within the dorsolateral frontal cortex for serial order memory. Proceedings of the Royal Society of London, Ser. B 246, 299-306.

Raz, N., Gunning, F. M., Head, D., Dupuis, J. H., McQuain, J., Briggs, S. D., Loken, W. J., Thornton, A. E., & Acker, J. D. (1997). Selective aging of the human cerebral cortex observed in vivo: differential vulnerability of the prefrontal gray matter. Cerebral Cortex, 7, 268-82.

Rugg, M. (1995). ERP studies of memory. In M. G. H. Coles and M. Rugg (Eds.), Electrophysiology of mind: Event-related brain potentials and cognition (pp. 132-170). Oxford, New York: Oxford University Press.

Rugg, M. D., Mark, R. E., Walla, P., Schloerscheidt, A. M., Birch, C. S., & Allan, K. (1998). Dissociation of the neural correlates of implicit and explicit memory. Nature, 392(6676), 595-598.

Schloerscheidt, A.M. & Rugg, M.D. (1997). Recognition memory for words and pictures: an event-related potential study. Neuroreport, 8, 3281-3285.

Senkfor, A.J., & Van Petten, C. (1998). Who said what? An event-related potential investigation of source and item memory. Journal of Experimental Psychology: Learning, Memory and Cognition, 24 (4), 1005-1025.

Shimamura, A.P., Janowsky, J.S., & Squire, L.R (1990). Memory for the temporal order of events in patients with frontal lobe lesions and amnesic patients. Neuropsychologia, 28(8), 803-813.

Spencer, W. D., & Raz, N. (1995). Differential effects of aging on memory for content and context: a meta-analysis. Psychology & Aging, 10, 527-39.

Squire, L.R. (1994). Declarative and nondeclarative memory: Multiple brain systems supporting learning and memory. In D. Schacter & E. Tulving (Eds.), Memory systems (pp. 203-231). Cambridge, MA: MIT Press.

Tovée, M.J. & Cohen-Tovée, E.M. (1993). Neural substrates of face processing models: A review. Cognitive Neuropsychology, 10(6), 505-528.

Wilding, E. L. (1999). Separating retrieval strategies from retrieval success: an event-related potential study of source memory. Neuropsychologia, 37, 441-454.

Wilding, E. L., & Rugg, M. D. (1996). An event-related potential study of recognition memory with and without retrieval of source. Brain, 119, 889-905.

Wilding, E. L., Doyle, M. C., & Rugg, M. D. (1995). Recognition memory with and without retrieval of context: An event-related potential study. Neuropsychologia, 33, 743-767.

CHAPTER 4

ARTIFICIAL NEURAL NETWORKS AS ANALYTIC TOOLS IN AN ERP STUDY OF FACE MEMORY

Memory can be conceptualized as the result of processing in two cortical areas: unimodal and transmodal association areas (Mesulam, 1998). Unimodal areas, such as the extrastriate areas and the fusiform, inferior, and middle temporal gyri, are modality specific and receive projections from primary sensory cortex. Transmodal areas receive inputs from more than one modality and include the prefrontal and posterior parietal cortices, as well as lateral and medial temporal areas. Within a modality, it is also possible to further conceptualize memory as being the result of general and stimulus specific processes. ERP studies have associated visual memory with two effects that likely reflect the activity of areas involved in general visual memory because they are elicited by different stimuli and experimental contexts. First, ERP's to remembered items show a positivity relative to new items that occurs after about 400ms over parietal regions. This has been found with both words (Senkfor & Van Petten, 1998; Wilding, 1999; Wilding & Rugg, 1996; Wilding, Doyle & Rugg, 1995) and faces (Graham & Cabeza, 2001a; Graham & Cabeza, 2001b). The second effect is a sustained positivity over frontal areas, which occurs later than the parietal effect and has also been found with words (Wilding, 1999; Wilding & Rugg, 1996; Wilding, Doyle & Rugg, 1995) and faces (Graham & Cabeza, 2001a; Graham & Cabeza, 2001b).

These results suggest that general visual memory effects occur relatively late after stimulus presentation. However, the timecourse of face specific memory activity remains uncertain. Neuroimaging (e.g., Kanwisher, Tong, & Nakamura, 1998) and intracranial (McCarthy, Puce, Gore & Allison, 1997) studies, with their superior structural resolution, have identified right ventral temporal lobe areas, particularly the right fusiform gyrus, as putative face sensitive areas. Although neurons in the macaque temporal lobe show early latency changes as a face becomes familiar (Rolls, Bayliss, Hasselmo & Nalwa, 1989), ERP evidence from humans is mixed. Electrophysiologically, fusiform activity is correlated with the N200, a negative deflection that occurs 200ms after face presentation and is maximal over right occipitotemporal areas (Bentin, Allison, Puce, Penez & McCarthy, 1996). Some studies have shown evidence of face memory effects over right temporal areas during the timeframe of the N170 (e.g., Campanella et al., 2000; Endl et al., 1998), while another study reports memory effects as early as 50ms (Seeck et al., 1997). However, other studies have not replicated these effects, and did not report memory effects until after 300 ms (e.g., Eimer, 2000a; Eimer, 2000b; Sommer, Heinz, Leuthold, Matt & Schweinberger, 1995).

There could be several explanations for the inconsistencies found in previous studies. We were interested in the possibility that memory-related activity could be represented in early latency ERP's from right temporal sites that are believed to exhibit face specific activity, but conventional methods of statistical analysis may not consistently detect them. The current study explored this possibility by comparing the ability of linear techniques and ANN's to discriminate between ERP's elicited by remembered and novel faces recorded from right temporal sites.

ANN's offer advantages over analyses based on the general linear model, particularly when they are employed for pattern recognition (Dawson et al., 1994). A more detailed description of ANN's is provided in Appendix B, but some introduction is warranted here. An ANN is a system consisting of interconnected layers of units, which generates a desired response to a given stimulus (Bechtel & Abrahamsen, 1994). A stimulus is encoded as a pattern of activity in a set of input units. The ANN's response is represented as a pattern of activity in output units. Intervening layers of units called hidden units detect higher-order stimulus regularities (Dawson et al., 1994). Rather than being programmed, an ANN is taught by providing it feedback about its performance. This is used by the ANN to adjust the weights of the connections between its layers of units in such a way that it gradually reduces its response error. ANN's can tolerate noisy data (Gupta, Molfese & Tammana,1995) and detect complex, nonlinear interactions between input features that are not normally detected by conventional techniques. Therefore, they have the potential to be powerful tools for ERP analysis.

The present study explored this potential by determining whether ANN's could distinguish between early latency ERP's recorded from right temporal sites to correctly recognized old and new faces, and comparing these results to those obtained using repeated measures ANOVA. Furthermore, if ANN's were successful, we were interested in determining if they could aid in elucidating the timecourse of ERP effects recorded over unimodal face areas.

Methods

Subjects

Fifty-four right-handed students with normal or corrected-to-normal vision participated in the study for course credit or paid participation. After obtaining written consent, participants were seated in a soundproof, shielded chamber and EEG recording equipment was applied. Data from 4 subjects were rejected because of equipment problems, while data from 3 were rejected due to excessive artifact. The final analysis included data from 47 subjects (19 male, 28 female).

Materials and Methods

The stimuli consisted of 240, 3.5 x 2.7 inch, black-and-white photos of unfamiliar faces taken from the Purdue University, the University of Stirling, and the University of Northern British Columbia face databases. In each of 4 blocks, participants studied 30 faces in random order, and then performed a recognition test that included 30 studied and 30 new faces. The assignment of faces to the blocks and to the old and new conditions was counterbalanced. Trials consisted of 5 events: a fixation (500 ms), a face (400 ms), a fixation (1600 ms), a response selection screen until a response was made, and a blank screen (1000 ms). The fixation was a 3.5 by 2.7 inch gray rectangle. At study, subjects were asked to remember the faces for a subsequent memory test. At test, a response screen prompted subjects to make an old/new decision. Subjects made their responses by pressing keys with different hands. Hand use was counterbalanced.

ERP methods

EEG was collected from 32 Ag/AgCl electrodes embedded in a Quikcap (Neurosoft Inc., Sterling, Virginia) electrode cap. Sites included frontopolar (FPI, FP2), frontal (F7, F3, FZ, F4, F8), frontocentral (FC3, FZ, FC4), central (C3, CZ, C4), centroparietal (CP3, CPZ, CP4), parietal (P3, PZ, P4), frontotemporal (FT7, FT8), temporal (T7, T8), temporoparietal (TP7, TP8), and occipital (O1, OZ, O2) electrodes with linked mastoids as a reference. Horizontal eye movements were monitored with bipolar electrodes on the outer canthus of each eye; vertical movements, from electrodes placed above and below the left eye. EEG was recorded with a sampling rate of 500 Hz for 1700 ms starting 100 ms prior to face onset. Channels were amplified with a filter bandwidth of 0.03 - 50 Hz. Trials with values above 100 or below -100 μV were excluded. ERP's were averaged according to 2 categories: correctly classified old faces (hits) and correctly classified new faces (CR's). The grand-averaged ERP's for the first 500 ms after face onset are shown in Figure 4-1 for the electrodes of interest (FT8, T8 and TP8).

---------- Insert Figure 4-1 about here ----------

Results

Two types of analyses were performed on the three electrode channels of interest: FT8, T8, and TP8. The first was a linear analysis of the data via repeated measures ANOVA; the second, analysis of the same data with ANN's.

Linear Analyses

Repeated measures ANOVA's were performed for each of the three channels on the first 500 ms of data after face onset from each of the temporal electrode channels. Data from each channel was parsed into 25 contiguous 20ms epochs that were also used as ANN inputs. Each electrode site was analysed with trial type (hit vs. CR's) and time (25 x 20ms epochs) as within-subjects factors. None of the ANOVA's was able to detect reliable differences between ERP's to hits and CR's during the first 500ms at any electrode site (main effect of type, type x time interaction, all F's < 1, p's>.05). In order to keep results as similar as possible to the ANN analysis, ANOVA's were repeated using within-subject standardized averaged voltage values for a 20ms epoch. Although it was not possible to assess main effects because of standardization, no time by type interactions were significant (all F's < 1, p's>.05).

Examination of Figure 4-1 revealed that waveforms for hits and CR's were virtually overlapping. However, some divergence was evidenced between 100 and

200ms, and from 340 to 500 ms. To rule out the possibility that differences existed within these epochs, repeated measures ANOVA's were performed for each of the three channels on two epochs: an early epoch from 100-200ms, and a later epoch from 340-500ms. Data in each epoch was parsed into contiguous 20ms epochs. For the early epoch, each electrode site was analysed with trial type (hit vs. CR) and epoch (5 x 20ms) as within-subjects factors. For the late epoch, each electrode site was analysed with trial type (hit vs. CR) and epoch (8 x 20ms) as within-subjects factors.

ANOVA's performed on the early epoch on electrode channels FT8 and TP8 were unable to detect reliable differences between ERP's to hits and CR's (main effect of type, type x time interaction, all $\underline{F}$'s < 2, $\underline{p}$'s$>.05$). The ANOVA performed on T8 did not detect any main effect of ERP type ($\underline{F}(1,46) = 0.51$, $\underline{p}>.05$). However, it did detect a small ERP type by time interaction ($\underline{F}(3,138) = 1.6$, $\underline{p}<.05$), which is depicted in the plot of the cell means in Figure 4-2. Post hoc paired t-test between corresponding timepoints for hits and CR's did not reveal any significant differences between the two ERP types. Examination of Figure 2 shows that none of the mean differences was larger than 1 microvolt, therefore, it is likely that the omnibus ANOVA is picking up a divergence from parallelism that cannot not be explained by any systematic differences between hits and CR's.

---------- Insert Figure 4-2 about here ----------

None of the ANOVA's performed on the later epoch were able to detect reliable differences between ERP's to hits and CR's during the epoch between 340 and 500ms at any of the three electrode sites (main effect of type, type x time interaction, all $\underline{F}$'s < 1, p's>.05). All ANOVA's showed a main effect of time, indicating that voltage values were changing over time.

ANN Analyses

Three ANN's (one for each channel) were trained to discriminate between the two ERP types. To promote training and reduce network complexity, an input representation scheme based on (Slater, Wu, Honig, Ramsay, & Morgan 1994) was used. Each ANN had 25 input units, each corresponding to the within-subject standardized averaged voltage values for a 20ms epoch, three hidden units and one output unit. For each network, both the hidden units and the output unit used a logistic activation function. This architecture is shown in Figure 4-3 and was adopted because pilot studies indicated that ANN's with these specifications could discriminate between hits and CR's with the fewest number of hidden units and could generalize well to new cases. The output unit was trained to generate a response of 1 to ERP's corresponding to hits, and a response of 0 to ERP's to CR's.

---------- Insert Figure 4-3 about here ----------

The data from 47 subjects was used to construct two data sets; a training and a test set. The training set consisted of data from 42 subjects (42 hits and 42 CR's), while the test set contained data from 5 randomly selected subjects (10 patterns, 5 of each type). Data from the test set was not used for network training. The ANN's were trained using the generalized delta rule (Rumelhart, Hinton & Williams 1986). Initially, all connection weights and processing unit biases were randomly assigned values between 1.0 and -1.0. The networks were trained with a learning rate of 0.075 and zero momentum. Weights and biases were updated after each training pattern presentation. Each pattern was presented once during a "sweep" of training; the order of pattern presentation was randomized before each sweep. The networks were trained until they converged (generated a hit for every pattern), where a hit was defined as response of .99 or higher when the desired response was 1, or .01 or lower when the desired response was 0. All ANN's were able to differentiate between ERP types. The first ANN (corresponding to the electrode FT8) converged after 3,640 sweeps; the second (corresponding to the electrode T8) after 2,263 sweeps, and the third (corresponding to the electrode TP8) after 2,828 sweeps. The prediction accuracy of each ANN was evaluated by exposing each network to the 10 test patterns. The first network (FT8) generalized to 80% of the new test patterns (3/5 CR's, 5/5 hits); the second (T8), to 90% (4/5 CR's, 5/5 hits); and the third (TP8), generalized to 80% (3/5 CR's, 5/5 hits).

Network Interpretation

To understand how the ANN's discriminated between hits and CR's, the activities of each hidden unit to each individual input pattern were examined. Alone, each unit was a fairly poor discriminator of ERP type. However, when the activities of all three units were examined simultaneously, it was possible to see how discrimination was achieved. As can be seen in the scatterplots of hidden unit activity in Figure 4-4, the ANN's were able to classify ERP's by representing input patterns in 3-D space that can be divided by a plane. Inputs representing hits fall on one side of this hypothetical plane; inputs corresponding to CR's fall on the other side.

---------- Insert Figure 4-4 about here ----------

Figure 4 demonstrates that hidden unit activity discriminates the two types of patterns. Hidden unit activity is a nonlinear function of net input, which in this study is the sum of the signals traveling from the input units through the weighted connections to the hidden unit. In order to determine if some inputs (timepoints) were more influential than others in classifying hits and CR's, we used multiple regression to determine if a subset of timepoints could account for most of the variance in the input to each hidden unit. A linear approach seemed appropriate since no nonlinear transformations occur at the level of net input. This more complex method of analysis was chosen over a straight interpretation of network weights because simple examination of connection weights revealed that inputs were not mapping specifically onto only one hidden unit. This indicated that the hidden units were not detecting local features in the data, and simple weight interpretation is only appropriate in the case where local representations were assumed (Hanson & Burr, 1990). Furthermore, due to large differences in voltage over time, the importance of a particular input in determining hidden unit activity is not only based on the magnitude of its connection weight to the hidden unit, but also on its magnitude before weighting. Regressions with net input as the predicted variable and weighted inputs as predictors permit the examination of both the connection weights and the initial values of the inputs.

Three regressions, one for each hidden unit, were performed on each of the three networks, with net input as the predicted variable. Timepoints entered as predictors were chosen on the basis of their standard deviations after weighting by the appropriate hidden unit, the rationale being that only timepoints with fairly large variance would contribute significantly to hidden unit activity. Inputs with the highest standard deviations after weighting were entered into each regression to yield the final regression equations.

A summary of these regressions is shown in Table 1, which shows the timepoints entered into the regression for each hidden unit and each channel, as well as measures of fit. Few inputs were necessary to account for variance in net input (one-third of the original inputs). All had high standard deviations after weighting, having the top third of all standard deviations for their respective hidden unit. There is also similarity between the timepoints entered, both within and across channels. These timepoints cover the entire range of the 500ms epoch; however, Table 4-1 shows some global characteristics of these inputs. Inputs between 100-200ms, as well as those between 400-500ms, appear to dominate each regression, with very few timepoints lying between 300-400ms. Early timepoints (<100ms) are also included in each regression. Additionally, timepoints between 200-300ms become more dominant in each solution as you move from the frontal channel (FT8) to the posterior channel (TP8).

---------- Insert Table 4-1 about here ----------

Discussion

While RM-ANOVA was unable to distinguish between ERP's, ANN's were able to differentiate between ERP's recorded over right temporal areas thought to be involved in the early stages of face processing. ANN's were able to discriminate between ERP's to correctly recognized old and new faces and generalized well to new test cases. This suggests that memory-related activity is represented in early latency ERP's from temporal areas and can be used to differentiate between hits and CR's. However, this information may be in the form of higher-order voltage/time relationships that cannot be consistently detected with linear methods. The success of the ANN's was not particularly surprising: ANN's have been used to differentiate between ERP's produced by individuals with brain pathology and controls (Gupta et al.,1995; Klöppel, 1994; Slater et al., 1994). However, a critical feature of these studies was their focus on identifying pathology. This study provides evidence that ANN's can also be useful in differentiating between cognitive states.

One reason for the inconsistencies between previous studies could be that early memory effects recruit relatively small cortical areas whose location and activity vary across individuals. This sample dependent variability could obstruct the detection of mean differences, particularly if long epochs are chosen for analysis. ANN's, on the other hand, can distinguish between groups with overlapping distributions and variances (Gupta et al., 1995) and classification occurs by comparing the entire timecourse of an ERP against those which belong to known types, without the aid of any a priori models. Effects in previous studies may also have been attenuated, either by virtue of their cortical generators or by the use of mastoid references. Since ANN's can tolerate noisy, incomplete data, they should be more robust to signal attenuation, as evidenced by the successful training of ANN's in this study despite the use of mastoid references.

Examination of hidden unit activity revealed that the ANN's were able to differentiate between hits and correct rejections by transforming inputs into three-dimensional pattern space. Further analysis revealed that timepoints that could account for a significant amount of variability in the net input to each hidden unit had high variability after weighting. This suggests that standard deviation after weighting could be useful for the interpretation of trained networks.

Timepoints identified as predictors of net input via multiple regression were surprisingly consistent between and within channels, but were not circumscribed to any particular epoch, suggesting that voltage values over the entire timecourse were necessary for correct classification. However, Table 4-1 shows an almost bimodal distribution of these timepoints, with inputs between 100-300ms and 400-500ms dominating each solution, as if the ANN's were detecting two features in the data. The first feature (100-300ms) may be a memory effect corresponding to the N170/VPP complex, while the second feature (400-500ms) may be the onset of the parietal effect, which was prominent after 500ms. Further examination of this first epoch is warranted; in particular, it would be interesting to see if an ANN can be trained to discriminate between hits and CR's when presented with data from the first 300ms after face presentation. While the second feature may be the onset of the parietal effect and could be interpreted negatively, these results are actually quite heartening at this preliminary stage, as they do provide some validation for network interpretation. The fact that the ANN's could detect this regularity before it became statistically significant is one reason to be optimistic about the use of ANN's for ERP analysis.

An interesting result of the regressions was the predominance of adjacent timepoints in their solutions. The fact that these timepoints were prominent in each solution was surprising given that adjacent timepoints should be highly correlated and therefore should not contribute much unique variance before weighting. The inclusion of adjacent timepoints as predictors in each regression, particularly between 100-300ms and 400-500ms, could signify that the ANN's are picking up differences in the onset latencies of both early and late ERP components. This suggestion can only be confirmed through topographic analysis of these epochs.

One surprising result of the regression analyses was the inclusion of timepoints before 100ms. This result does lend credence to the idea that there are rapid mechanisms for face recognition (Seeck et al., 1997); however, this result awaits replication. There also appears to be an anterior-posterior gradient with respect to timepoints between 200-300ms, with these timepoints appearing more prominently at the posterior than at the anterior temporal channel. The significance of this result is unknown and certainly warrants further investigation. Future inquiries will focus on the topographic distribution of specific ERP effects. However, the results of the present study do indicate that ANN's can be useful for the temporal analysis of ERP data.

The purpose of this study was to extend the ERP literature by re-examining ERP memory effects with a different analytic tool, ANN's. Much remains to be revealed about how ANN's classify ERP's and a framework for future analyses remains to be established. Nevertheless, preliminary results are encouraging and provide support for early unimodal memory effects and the continued use of ANN's in ERP analysis.

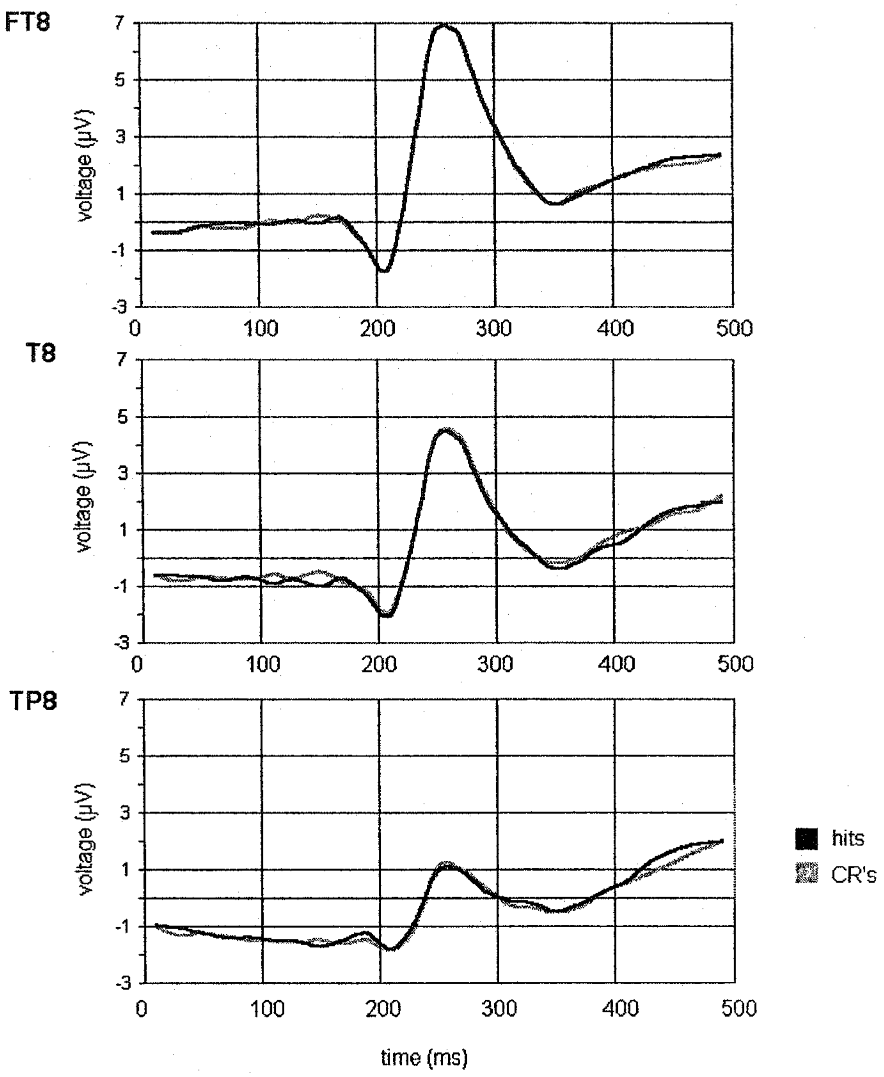

Figure 4-1. Grand-averaged voltage values (µV) obtained from right temporal electrodes during the first 500 ms of face presentation.

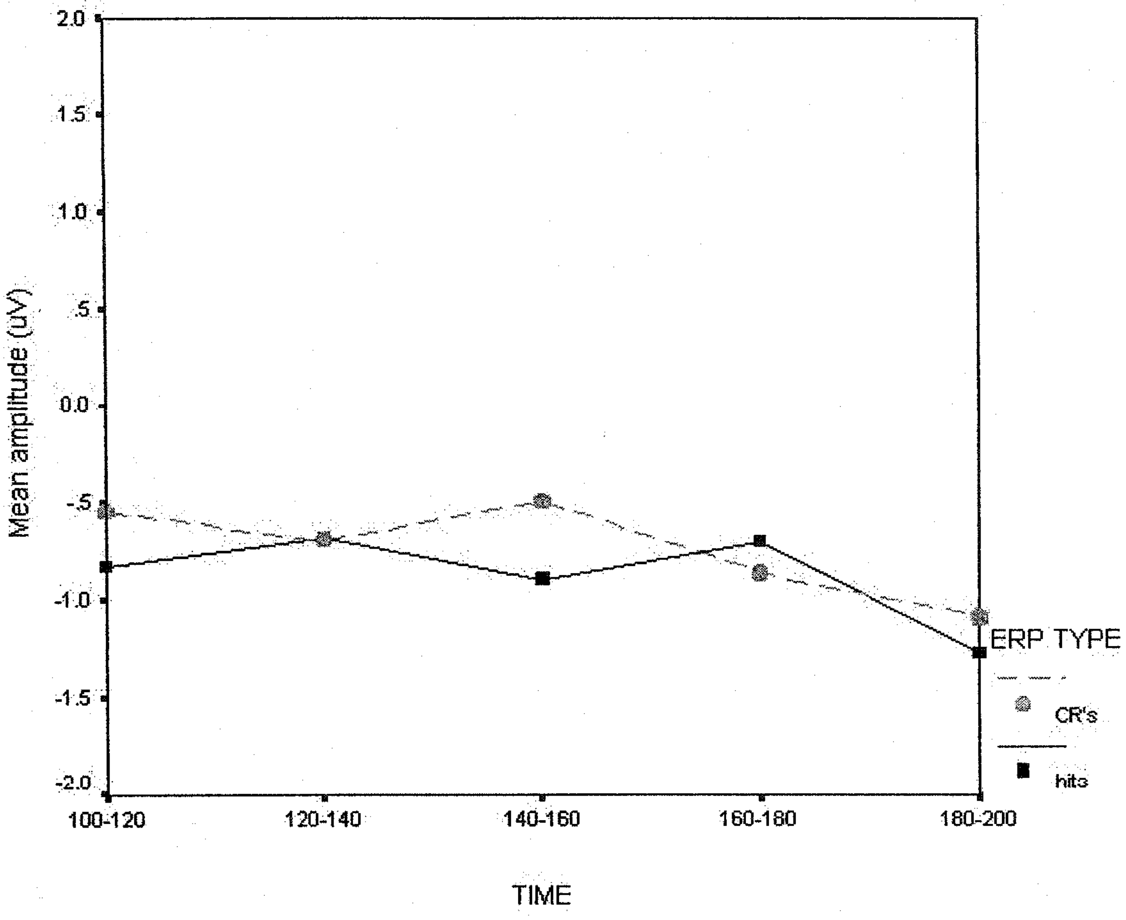

Figure 4-2. Plot of cell means for the early epoch (100-200ms) for hits and CR's

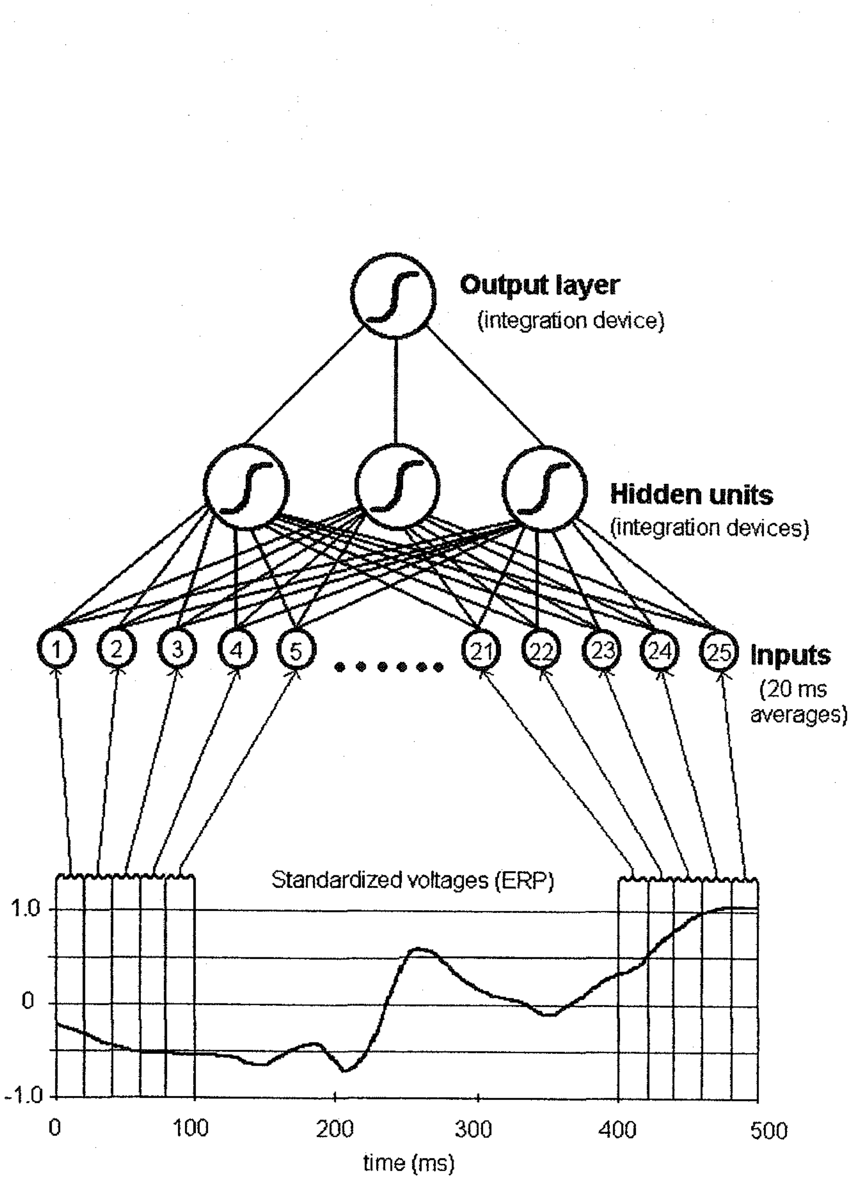

Figure 4-3. The network architecture used in this study. Average standardized voltage values represent averaged standardized ERP's obtained from each subject for each ERP type.

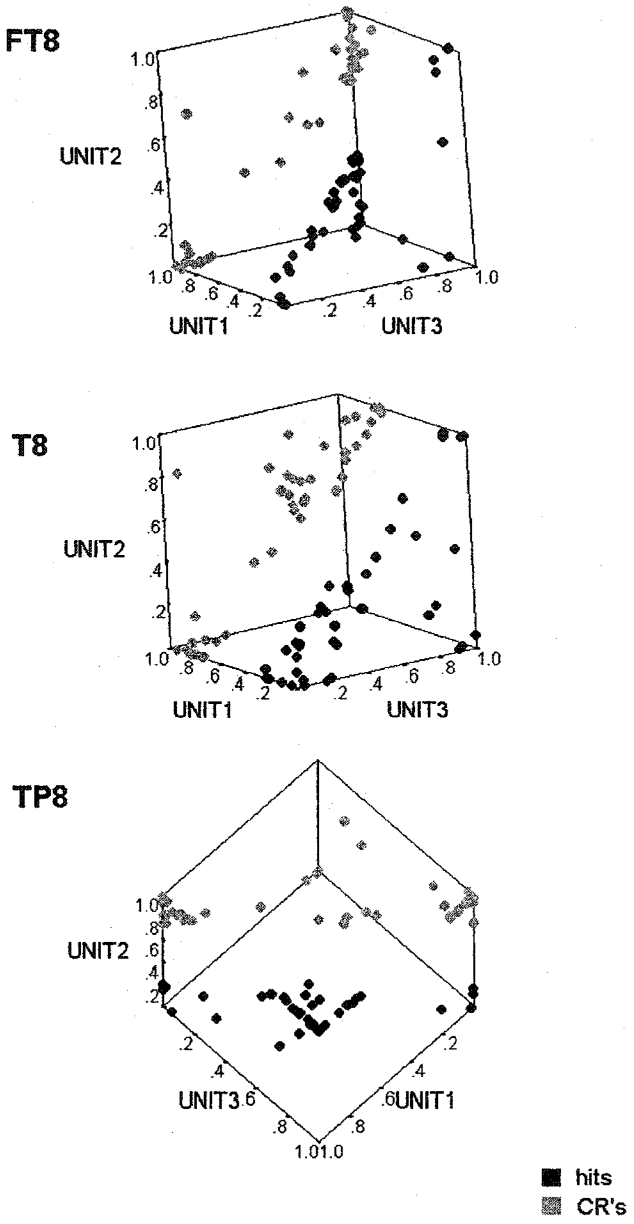

Figure 4-4. Three-D scatterplots of hidden unit activity for each of the ANN's. Plots correspond to ANN's trained on data from right temporal electrode sites (FT8, T8 and TP8, respectively).

Table 4-1. Summary of timepoints identified as significant predictors of net input via multiple regression for each hidden unit of each network. Numbers in the table represent the input number, as well as r and R^2 values for each regression; the timecourse is indicated along to bottom of the table to aid interpretation.

Inputs entered						r	R^2
FT8							
HU1:	4	7 8 9 10	15		22 24	.92	.85
HU2:	1 2	6 7 8		19 20	22	.83	.69
HU3:	1	7	11	16 19	23	.96	.92
T8							
HU1:		7 8	11 12 13		21 22 25	.91	.83
HU2:	1 2 3	6 9	12		22 25	.97	.93
HU3:	1	8	11 13	20	21 23 25	.98	.97
TP8							
HU1:	3	7 10	12 13 14		23 25	.94	.88
HU2:	1	7 8	12 13 14 15		25	.93	.86
HU3:	2	8 9	11 13 14	19	23	.96	.93

time (ms) 0 100 200 300 400 500

References

Bechtel, W. & Abrahamsen, A. (1994). Connectionism and the mind: An introduction to parallel processing in networks. Cambridge: Blackwell Publishers.

Bentin, S., Allison, T., Puce, A., Penez, E. & McCarthy, G. (1996). Electrophysiological studies of face perception in humans. Journal of Cognitive Neuroscience, 8, 551-565.

Campanella, S., Hanoteau, C., Depy, D., Rossion, B., Bruyer, R., Crommelinck, M.. & Guerit, J. M. (2000). Right N170 modulation in a face discrimination task: An account for categorical perception of familiar faces. Psychophysiology, 37(6), 796-806.

Dawson, M.R.W., Dobbs, A., Hooper, H.R., McEwan, A.J.B., Triscott, J. & Clooney, J. (1994). Artificial neural networks that use single-photon emission tomography to identify patients with probable Alzheimer's disease. European Journal of Nuclear Medicine, 21, 1303-1311.

Eimer, M. (2000a). Event-related brain potentials distinguish processing stages involved in face perception and recognition. Clinical Neurophysiology, 111(4), 694-705.

Eimer, M. (2000b). Effects of face inversion on the structural encoding and recognition of faces: Evidence from event-related brain potentials. Cognitive Brain Research, 10(1-2), 145-158.

Endl, W., Walla, P., Lindinger, G., Lalouschek, W., Barth, F.G., Deecke, L. & Wang, W. (1998). Early cortical activation indicates preparation for retrieval of memory for faces: An event-related potential study. Neuroscience Letters, 240, 58-60.

Graham, R. & Cabeza, R. (2001a). Event-related potentials of recognizing happy and neutral faces. Neuroreport, 12, 245-248.

Graham, R. & Cabeza, R. (2001b): Dissociating the neural correlates of item and context memory: An ERP study of face recognition. Canadian Journal of Experimental Psycholology, 55, 154-161.

Gupta, L., Molfese, D.L. & Tammana, R. (1995). An artificial neural network approach to ERP classification. Brain and Cognition, 27, 311-330.

Hanson, S.J. & Burr, D.J. (1990). What connectionist models learn: Learning and representation in connectionist networks. Behavioral and Brain Sciences, 13, 471-518.

Kanwisher, N., Tong, F. & Nakamura. K. (1998). The effect of face inversion on the human fusiform face area. Cognition, 68, B1-11.

Klöppel, B. (1994). Classification by neural networks of evoked potentials: A first case study. Neuropsychobiology, 29, 47-52.

McCarthy, G., Puce, A., Gore, J.S. & Allison, T. (1997). Face-specific processing in the human fusiform. Journal of Cognitive Neuroscience, 9, 605-610.

Mesulam, M-M. (1998). From sensation to cognition. Brain, 121, 1013-1052.

Rolls, E.T., Baylis, C.G., Hasselmo, M.E. & Nalwa, V. (1989). The effect of learning on the face selective responses of neurons in the cortex of the superior temporal sulcus of the monkey. Experimental Brain Research, 76, 153-164.

Rumelhart, D.E., Hinton, G.E. & Williams, R.J. (1986). Learning internal representations by error propagation; in D.E. Rumelhart and J. McLelland (Eds.): Parallel distributed processing: Explorations in the microstructure of cognition, Vol 1., 319-362. Cambridge: MIT Press.

Seeck, M., Michel, C.M., Mainwaring, N., Cosgrove, R., Blume, H., Ives, J., Landis, T. & Schomer, D.L. (1997). Evidence for rapid face recognition from human scalp and intracranial electrodes. Neuroreport, 8, 2749-2754.

Senkfor, A. J. & Van Petten, C. (1998). Who said what? An event-related potential investigation of source and item memory. Journal of Experimental Psychology: Learning, Memory, & Cognition, 24(4), 1005-1025.

Slater, J.D., Wu, F.Y., Honig, L.S., Ramsay, R.E., & Morgan, R. (1994). Neural network analysis of the P300 event-related potential in multiple sclerosis. Electroencephalography and Clinical Neurophysiology, 90, 114-122.

Sommer, W., Heinz, A., Leuthold, H., Matt, J. & Schweinberger, S.R. (1995). Metamemory, distinctiveness and event-related potentials in recognition memory for faces. Memory and Cognition, 23, 1-11.

Wilding, E. L. (1999). Separating retrieval strategies from retrieval success: an event-related potential study of source memory. Neuropsychologia, 37, 441-454.

Wilding, E. L., Doyle, M. C., & Rugg, M. D. (1995). Recognition memory with and without retrieval of context: An event-related potential study. Neuropsychologia, 33, 743-767.

Wilding, E. L., & Rugg, M. D. (1996). An event-related potential study of recognition memory with and without retrieval of source. Brain, 119, 889-905.

CHAPTER 5

THE IMPORTANCE OF PEAK ANALYSIS PROCEDURES IN THE EXAMINATION OF EARLY EVENT-RELATED CORRELATES OF FACE MEMORY

A version of this chapter is currently under review at Cognitive Brain Research

Face recognition seems rapid, however, early event-related potential (ERP) correlates of face memory have yet to be firmly established. Neuroimaging (e.g., Kanwisher, Tong & Nakayama, 1998) and intracranial (e.g., McCarthy, Puce, Gore & Allison, 1997) studies have identified the fusiform gyrus as an important cortical area for face processing. Activity in this area has been correlated with the N170 (Bentin, Allison, Puce, Perez & McCarthy, 1996) and with the vertex positive potential (VPP) (Jeffreys, 1993), deflections that occur approximately 200ms after face presentation. Is this activity memory-sensitive? Although neurons in the macaque temporal lobe alter their firing rates as a face becomes familiar (Rolls, Baylis, Hasselmo & Nalwa, 1989), ERP evidence with humans is inconsistent.

Some studies have reported reduced amplitude of the N170 to repeated faces (Campanella et al, 2000; Endl et al., 1998), while another showed both decreased amplitude and latency of the N170 to repeated faces (Itier & Taylor, 2002). These studies suggest that the N170 does index memory-related activity. Other researchers have not shown evidence of the memory sensitivity of this effect (Bentin & Deouell, 2000; Eimer, 2000a; Eimer, 2000b; Eimer, 2000c), leading to the interpretation that the N170 indexes the structural encoding of face information, in particular, the encoding global configural facial information. This view is also supported by evidence from fMRI (Kanwisher et al., 1998). With regard to the VPP, its memory sensitivity has been studied less extensively and also remains in question. Potter and Parker (1989) reported an early frontal positivity with a latency and topography similar to that of the VPP whose amplitude was sensitive to face repetition. VPP memory effects in the form of amplitude effects have also been reported by Itier and Taylor (2002), but not by Rossion et al. (1999). In the case of both components, it is unclear whether putative memory effects are in the form of amplitude and/or latency differences.

One possible explanation for inconsistencies regarding the memory sensitivity of the N170 and VPP is that they are modulated by perceptual priming effects due to repetition. Many of the studies describing memory effects in the N170 and VPP employed continuous recognition tasks where lags between repeated faces were short (e.g. Itier & Taylor, 2001), whereas studies employing longer lags between faces typically did not show these effects (e.g., Eimer, 2000b). Therefore, it is possible that N170 and VPP memory effects will not be observable at long lags and will only be evidenced when repetition is immediate (Guillaume & Tiberghien, 2000; Itier & Taylor, 2002). However, even if these effects are due to perceptual priming, it is still not clear whether they disappear at longer lags or if they are merely attenuated. One way to investigate this is to examine the memory sensitivity of the N170 and VPP in a paradigm that employs long lags between repeated faces. Given that the signal may attenuate at long lags, it is possible that it can only be detected with the power afforded by a large number of subjects.

Another, more methodological explanation could account for mixed results regarding the memory sensitivity of the N170 and VPP: inconsistencies could be due to variability in the procedures used for peak amplitude analysis. These can be conceptualized as belonging to one of three types. The first type involves choosing an epoch to represent peak amplitude whose latency is constant both within and across individual subjects. Procedures of this type would include choosing peak amplitude based on a prespecified epoch, or based on the grand averaged waveforms. A second method to determine peak amplitude involves choosing peak amplitude based on the averaged ERP's for each subject. With this method, peak latency is variable across individuals, while within-subject peak latency is constant. Although this procedure is rarely used, it may be more suitable than fixed epochs for the analysis of early phasic components such as the N170 and VPP, which show significant inter-subject variability in onset (e.g., Roisson et al., 1999). The third category of analysis procedures determine peak amplitude for each individual and condition separately, resulting in variable peak latencies both across and within individuals. The fundamental difference between the different types of procedures is whether peak latency remains constant across subjects and conditions.

The choice of peak analysis procedure could have an impact on the results of amplitude analysis because of interrelationships between amplitude and latency. Studies examining the N170 and VPP can also be grouped according to whether or not this relationship was directly examined. Researchers using fixed epochs for amplitude analysis typically do not examine the association between amplitude and latency effects. For example, Eimer (2000a) did not find evidence of memory effects when peak amplitude of the N170 was defined as the mean amplitude between 140-190ms post-stimulus, regardless of individual variation. Because the phasic nature of the N170, if memory effects are characterized by subtle changes in latency, then the choice of a wide epoch could preclude their detection. When latency was analysed, it was chosen individually for each subject and condition, such that latency and amplitude measures could not be directly compared. Other studies showing positive evidence regarding the memory sensitivity of the N170 (e.g. Endl et al., 1998) have used smaller fixed epochs and did not examine latency. Because of this, it is impossible to conclude that these amplitude differences are not actually due to differences in peak onset under different conditions.

Another group of studies do take the relationship between amplitude and latency into account in peak analysis. For example, Campanella et al. (2000) determined peak amplitude for each subject and condition individually and found amplitude memory effects for the N170 even after onset differences were taken into account. Although latency effects were not examined, visual inspection of the waveforms did not reveal obvious differences in peak latency. Roisson et al. (1999) also examined peak amplitude and latency of the N170 and VPP for each subject and condition and did not find evidence of memory effects in either amplitude or latency, a finding that could be attributed to task differences described above.

Another explanation for these findings could be that the small sample sizes that are typical of ERP experiments do not provide enough power to consistently detect these effects. Itier and Taylor (2002) also determined peak latency and amplitude of the N170 and VPP individually for a large number of subjects (N = 34) and reported a small but consistent latency effect for both components that may not have been detectable in other studies due to small sample sizes, even if it had been examined explicitly. Therefore, inconsistencies could be caused in part by the different methods that researchers employ to determine peak amplitude and further exacerbated by small sample sizes. We explored this by conducting our analyses on a large corpus of ERP data.

The purpose of this study was twofold. First, we were interested in the memory sensitivity of the N170 and VPP when elicited in a paradigm with long lags between studied faces. Old/new paradigms are typically used in the study of ERP correlates of verbal memory (e.g. Allan et al., 1998) and have also been applied to the study of face memory (e.g. Graham & Cabeza, 2001a; Graham & Cabeza, 2001b). This paradigm consists of study blocks where subjects view unfamiliar faces, which are followed by test blocks consisting of studied and new faces. Study and test blocks are separated by several minutes, creating long lags between studied faces. If early memory effects are contingent upon immediate repetition and disappear at long lags, they should not be detectable with this paradigm. However, if memory effects are attenuated by long lags between study and test, they may not be detectable with linear analysis, but may be detected with a more powerful pattern recognition tool, namely artificial neural networks (ANN's).

ANN's offer advantages over linear analyses, particularly when they are employed for pattern recognition (Dawson et al., 1994). This is because ANN's can detect complex interactions between input features that are not normally detected by conventional techniques. An ANN is a system consisting of interconnected layers of units that generates a desired response to a given stimulus, which is encoded as a pattern of activity in a set of input units (Bechtel & Abrahamsen, 1994). The ANN's response is represented as a pattern of activity in output units. Intervening layers of units called hidden units detect higher-order stimulus regularities (Dawson et al., 1994). Rather than being programmed, an ANN is trained by providing it feedback about its performance, which is used to adjust the weights of the connections between the layers of units. In this way, the ANN gradually reduces its response error until perfect discrimination is achieved. ANN's can tolerate noisy data and distinguish between groups with overlapping distributions and variances (Gupta, Molfese & Tammana, 1995). Therefore, they have the potential to be powerful tools for ERP analysis, particularly since network interpretation is possible.

A second objective of this study was to investigate the effects of different strategies for determining peak amplitude on the results of linear and ANN analyses. Because procedures differ in the way latency is represented across individuals and conditions, it is possible that the method used to determine peak amplitude could impact the interpretation of results. To examine this, we compared the results of two types of peak amplitude analyses, one that kept latency constant across conditions and another that sampled peak amplitude and latency separately for each subject and condition. This was done in order to examine the effect of separating latency and amplitude measures on the outcome of peak analyses. If different strategies yielded different results with either type of analysis, then methodological differences between studies could account for mixed results regarding the memory sensitivity of the N170 and VPP.

In summary, this study describes a series of topographic analyses conducted to examine the memory sensitivity of the N170 and VPP face effects. This study was motivated by the assumption that if early memory effects are attenuated by long lags between study and test, effects elicited in an old/new paradigm may be detectable with an ANN but not with linear analysis of the same data. Furthermore, if an ANN was successful, we were interested in determining how it was able achieve discrimination. We also examined effects of different strategies for determining peak amplitude on the results of these analyses, as well as the effect of employing a large sample. Ultimately, it is hoped that this research will help us to understand how methodological differences between studies could account for mixed results regarding the memory sensitivity of the N170 and VPP.

Methods

Subjects

Fifty-four right-handed students with normal or corrected-to-normal vision participated in the study for course credit or payment. After obtaining consent, participants were seated in a soundproof, shielded chamber and EEG recording equipment was applied. Data from 4 subjects were rejected because of equipment problems, while data from 3 were rejected due to excessive artifact. The final analysis included data from 47 subjects (19 male, 28 female).

Materials and Methods

The stimuli consisted of 240, 3.5 x 2.7 inch, greyscale photos of unfamiliar faces taken from the Purdue University, the University of Stirling, and the University of Northern British Columbia face databases. In each of 4 blocks, participants studied 30 faces in random order, and then performed a recognition test that included 30 studied and 30 new faces. The assignment of faces to the blocks and to the old and new conditions was counterbalanced. Trials consisted of 5 events: a fixation (500 ms), a face (400 ms), a fixation (1600 ms), a response selection screen until a response was made, and a blank screen (1000 ms). The fixation was a 3.5 by 2.7 inch gray rectangle. At study, subjects were asked to remember the faces for a subsequent memory test. At test, a response screen prompted subjects to make an old/new decision, which they made by pressing keys with different hands. Hand use was counterbalanced.

ERP methods

EEG was collected from 32 Ag/AgCl electrodes embedded in a Quikcap (Neurosoft Inc., Sterling, Virginia) electrode cap. Sites included frontopolar (FPI, FP2), frontal (F7, F3, FZ, F4, F8), frontocentral (FC3, FZ, FC4), central (C3, CZ, C4), centroparietal (CP3, CPZ, CP4), parietal (P3, PZ, P4), frontotemporal (FT7, FT8), temporal (T7, T8), temporoparietal (TP7, TP8), and occipital (O1, OZ, O2) electrodes, with linked mastoids as a reference. Horizontal eye movements were monitored with bipolar electrodes on the outer canthus of each eye; vertical movements, from electrodes placed above and below the left eye. EEG was recorded with a sampling rate of 500 Hz for 1700 ms starting 100 ms prior to face onset. Channels were amplified with a filter bandwidth of 0.03 - 50 Hz. Trials with values above 100 or below -100 μV were excluded. After collection, data was recalculated offline to obtain values relative to a common averaged reference for each subject. ERP's were averaged according to 2 categories: correctly classified old faces (hits) and correctly classified new faces (CR's).

Data Analyses

For both type of peak analyses, the N170 was defined as a negative potential peaking between 160 and 250 ms, maximal over occipitotemporal areas and the VPP, as a positive potential occurring during the same epoch, maximal over frontocentral areas. Both components were readily discernible in all participants and can be seen in the grand averaged ERP's in Figure 5-1. For the first peak analysis, averaged peak voltages of the N170 and VPP were determined for each subject by finding the peak latency of the each component for hits and CR's and averaging across conditions and components to obtain mean peak latency for each subject. Topographic voltage values corresponding to a 20ms window around that latency were then used as an estimate of peak amplitude. The decision to average across components was made because differences in the onset of the N170 and VPP were significantly shorter than the epoch window.

Peak values were then subjected to linear and ANN analyses. For the second peak analysis, peak voltages of the N170 and VPP were obtained by finding the peak latency and amplitude for each subject and condition individually. Topographic voltage values corresponding to a 20ms window around the peak latency for each condition were then used as estimates of peak amplitude, yielding separate measures for peak latency and amplitude for each component and condition.

---------- Insert Figure 5-1 about here ----------

Results

As mentioned above, two peak analyses were performed on the data. The first peak analysis was performed on peak voltages of the N170 and VPP for each subject , keeping peak latency constant within subjects. Topographic voltage values corresponding to a 20ms window around that latency were used as estimates of peak amplitude, which were then subjected to analysis with repeated measures analysis (RM-ANOVA) and ANN's. The second peak analysis was performed on peak voltage values of the N170 and VPP obtained by finding the peak latency and amplitude for each subject and condition individually, such that within-subject latency was variable. Topographic voltage values corresponding to a 20ms window around the peak latency for each condition were then used as estimates of peak amplitude, yielding separate measures for peak latency and amplitude for each component and condition. These values were then subjected to analysis with RM-ANOVA and ANN's.

Peak Analysis 1

Repeated measures ANOVA's corrected for inhomogeneity of covariance (Geisser-Greenhouse) and t tests corrected for multiple comparisons (Bonferroni) were performed on the peak topographic voltage values for the N170 and VPP. Voltage over the entire head during the peak of the N170 was analysed with trial type (hits vs. CR) and electrode location as within-subjects factors. The ANOVA was not able to detect reliable differences between topographic voltages to hits and CR's (main effect of type, type x location interaction, all F's< 1). There was a significant effect of electrode location, F (2.4,109.5) = 8.7, $p<0.001$, indicating that voltage values varied across the scalp.

An ANN was also trained to discriminate between peak topographic voltage values corresponding the two trial types. It had 30 input units, each corresponding to the averaged voltage values for one electrode channel, two hidden units and one output unit. The output unit was trained to generate a response of 1 to ERP's corresponding to CR's, and a response of 0 to ERP's to hits. The ANN was a hybrid of the type described by (Dawson & Schopflocher, 1992), in which the hidden units used a logistic activation function, and the output unit used a Gaussian activation function. This architecture is shown in Figure 5-2 and was adopted because pilot studies indicated that ANN's with these specifications could discriminate between hits and CR's with the fewest number of hidden units.

---------- Insert Figure 5-2 about here ----------

The ANN was trained using a modification of the generalized delta rule used to train ANN's with Gaussian activation functions (Dawson & Schopflocher, 1992). Initially, all connection weights and processing unit biases were randomly assigned values between 1.0 and -1.0. The network was trained with a learning rate of 0.001 and zero momentum. Weights and biases were updated after the presentation of each training pattern. Each input pattern (topographic voltages for an individual subject) was presented once during a "sweep" of training; the order of pattern presentation was randomized before each sweep. The ANN was trained until it generated a hit for every pattern, where a hit was defined as response of .99 or higher when the desired response was 1, or .01 or lower when the desired response was 0. The ANN was able to differentiate between hits and CR's, converging after 30,180 sweeps.

In order to understand how the ANN discriminated between hits and CR's, the activities of each hidden unit to each individual input pattern were examined. Alone, each unit was a poor discriminator of ERP type, however, when the activities both were examined simultaneously, it was possible to see how discrimination was achieved. As can be seen in the scatterplots of net input and hidden unit activity in Figure 5-3, the ANN was able to classify ERP's by representing input patterns in two-dimensional (2-D) space. Inputs representing CR's fall into a circumscribed area within this space, whereas those corresponding to hits fall outside of this area.

---------- Insert Figure 5-3 about here ----------

To determine if some electrode locations were more influential than others in determining hidden unit activity, we used multiple regression to determine if a subset of inputs could account for most of the variance in the activities of each hidden unit. In order to keep inputs and activity in the same metric, regressions were performed with net input (hidden unit activity before non-linear transformation) as the predicted variable. Two regressions were performed, one for each hidden unit. Electrode locations entered as predictor variables were chosen on the basis of their standard deviations after weighting by the connections to the appropriate hidden unit. The rationale was that only locations with large variances would contribute significantly to net input and thus to hidden unit activity; therefore, locations with the 10 highest standard deviations after weighting were included as predictors. This is shown in Figure 5-4, which depicts the standard deviations for each electrode location before weighting, and after weighting with connections to hidden units 1 and 2, respectively. Channels used are predictors are shown as dark grey and black areas in Figure 5-4b and c.

---------- Insert Figure 5-4 about here ----------

The regression performed for hidden unit 1 activity included 10 channels, 8 from frontal and central areas; F3, FT7, FC3, FCZ, FC4, CZ, C4, and CPZ, and 2 from occipitoparietal areas; P8 and OZ. This regression equation was statistically significant ($\underline{r}$ = .89, $\underline{F}$(10,83) = 31.7; p<.001), accounting for 79% of the variance in the net input to hidden unit 1. The regression for unit 2 activity included 3 frontocentral channels; FC3, C3 and C4, and 7 channels from occipitoparietal areas; CPZ, CP4, PZ, P4, P8, OZ, and O2. This regression equation was also statistically significant ($\underline{r}$ = .92, $\underline{F}$(10,83) = 44.0; p<.001), accounting for 84% of the variance in net input to unit 2. Thus, it would appear that the neural network was able to discriminate between hits and CR's by detecting amplitude differences over frontocentral and occipitoparietal areas.

Peak analysis 2

Two separate RM-ANOVA's were performed, one with peak latency of the N170 and VPP as the dependent variable, the other with peak amplitude as the dependent variable and trial type (hit vs. CR) as a within subject variable. In order to discover any interactions between the N170 and VPP, component type (N170 vs. VPP) was also included as a within-subjects factor in each ANOVA. Peak latency corresponded the peak of the N170 and VPP for each subject and condition. Peak amplitudes corresponded to a 20ms window around the peak of the N170 and VPP for each condition. Based on grand averaged peak maxima, values from P8 were chosen to represent the N170 and values from FC3 were chosen to represent the VPP. All RM-ANOVA's were corrected for inhomogeneity of covariance (Geisser-Greenhouse) and t-tests corrected for multiple comparisons (Bonferroni).

For latency, there was a significant main effect of component type ($\underline{F}(1,46) = 14.6$, $p > .001$), reflecting the tendency for the N170 to onset 2.4ms earlier than the VPP. There was also a significant main effect of trial type ($\underline{F}(1, 46) = 4.5$, $\underline{p} > .05$) indicating that hits onset 2.4ms earlier than CR's. There was no trial by component type interaction ($\underline{F} < 1$). The RM-ANOVA conducted with amplitude as the dependent variable did not reveal any trial type or trial type by component interactions ($\underline{F}$'s < 2, $\underline{p} >.1$), reflecting but did reveal a main effect of component type ($\underline{F}(1,46) = 147.8$, $\underline{p} < .001$), reflecting the reverse polarity of the N170 and the VPP.

To confirm that amplitude differences were no longer present, an ANN identical to the one used previously was employed to discriminate between peak topographic voltage values corresponding the two trial types after the removal of latency information. The ANN was unable to train on this task, even with the addition of 13 extra hidden units. Therefore, amplitude differences were not detectable after peak latency was taken into account.

Discussion

This study provided two main results. First, the N170 and VPP face effects do show evidence of memory sensitivity when elicited in a paradigm with long lags between studied faces. However, these memory effects appear to be attenuated by long lags between study and test. Second, the interpretation of these effects also differed depending on the method of peak analysis employed. Below, we discuss these results in more detail and consider their relevance as an explanation for conflicting reports of the memory sensitivity of early face-sensitive ERP components.

One explanation that has been ventured regarding the N170 and VPP is that they are modulated by perceptual priming effects due to repetition (Guillaume & Tiberghien, 2000; Itier & Taylor, 2002). However, rather than disappearing at longer lags, in this study they appear to be attenuated. In the first analysis, amplitude effects were not detectable with linear analysis, but were detected with a more powerful pattern recognition tool (ANN's). The advantage of the ANN could have been due to the fact that it does not use any a priori models whereas standard statistical analysis focuses on specific relationships between electrode channels and ERP types (Slater et al., 1992). In other words, an ANN can take voltage values over the entire head into account when discriminating between hits and CR's. This could be an asset if amplitude differences are small, but consistent across several channels. Indeed, the clustering of channels identified as significant predictors of net input tends to predominate over scalp areas where either the N170 or VPP are evident in the grand averages (see Figure 1). It is also interesting to note that both P8 and FC3, channels where the N170 and VPP were prominent in most subjects, were significant predictors of net input for both hidden units. The success of the

ANN's was not very surprising: ANN's have been used to classify ERP's produced by individuals with brain pathology and controls (e.g., Wu, Slater, Honig & Ramsay, 1993). This study provides evidence that ANN's can also be useful in differentiating between cognitive states.

In the second set of analyses, the effect of memory on the N170 and VPP was manifested as a small reduction in latency of both components to remembered faces that was detectable via ANOVA. With latency differences taken into account, amplitude differences were no longer detectable with either ANOVA or (viably) with ANN's. The latency effect observed in this study was very small, accounting for approximately 9% of the total sum of squares and is estimated to be at least half the size of the latency effect reported by Itier and Taylor (2002). These results could be interpreted to indicate that amplitude effects disappear at longer intervals between repeated faces, while latency effects are more robust but are attenuated at longer lags. Attentuation may occur to such a degree that only a large sample size would be sufficient to detect them. However, this speculation awaits confirmation through more detailed examination of how these effects change over varying lags. Nevertheless, sample size may be a critical consideration in the endeavour to establish the memory sensitivity early ERP components, especially if long intervals intervene between repeated faces.

Choice of peak analysis procedure appears to be a critical factor in the interpretation of the nature of early ERP face memory effects. When latency was held constant across individuals allowing only amplitude to vary, the ANN was able to detect amplitude differences at channels where the N170 and VPP were prominent. However, when peak latency was taken into account, amplitude differences were no longer evident.

Instead, ANOVA was able to detect a significant decrease in the latency of the N170 and VPP to remembered faces, suggesting that the amplitude effect detected by the ANN in the first analysis was due to subtle differences in the onset latencies of hits and CR's.

These results could have several implications for interpretation. If latency is held constant within a narrow epoch, relatively small differences in latency could mistakenly be interpreted as amplitude differences. If latency is held constant within a large epoch, these differences may be undetectable. Therefore, choice of epoch length could also be an important factor in the assessment of early phasic ERP components. Nevertheless, it is clear from this study that in order to fully understand how memory modulates early face sensitive potentials both amplitude and latency must be taken into account. It is possible that the contradictory reports regarding the memory sensitivity of these components are at least in part attributable to methodological differences between studies.

In this paper, a series of topographic analyses examining the memory sensitivity of early ERP face effects suggests that the N170 and VPP are modulated by memory. These early memory effects could be attenuated by long lags between study and test such that they can only be detected with large data samples or powerful pattern recognition devices such as ANN's. However, the attribution of effects to amplitude or latency differences depends highly upon the method of peak analysis that is used. The results of these analyses could help to reconcile the mixed results regarding the memory sensitivity of the N170 and VPP.

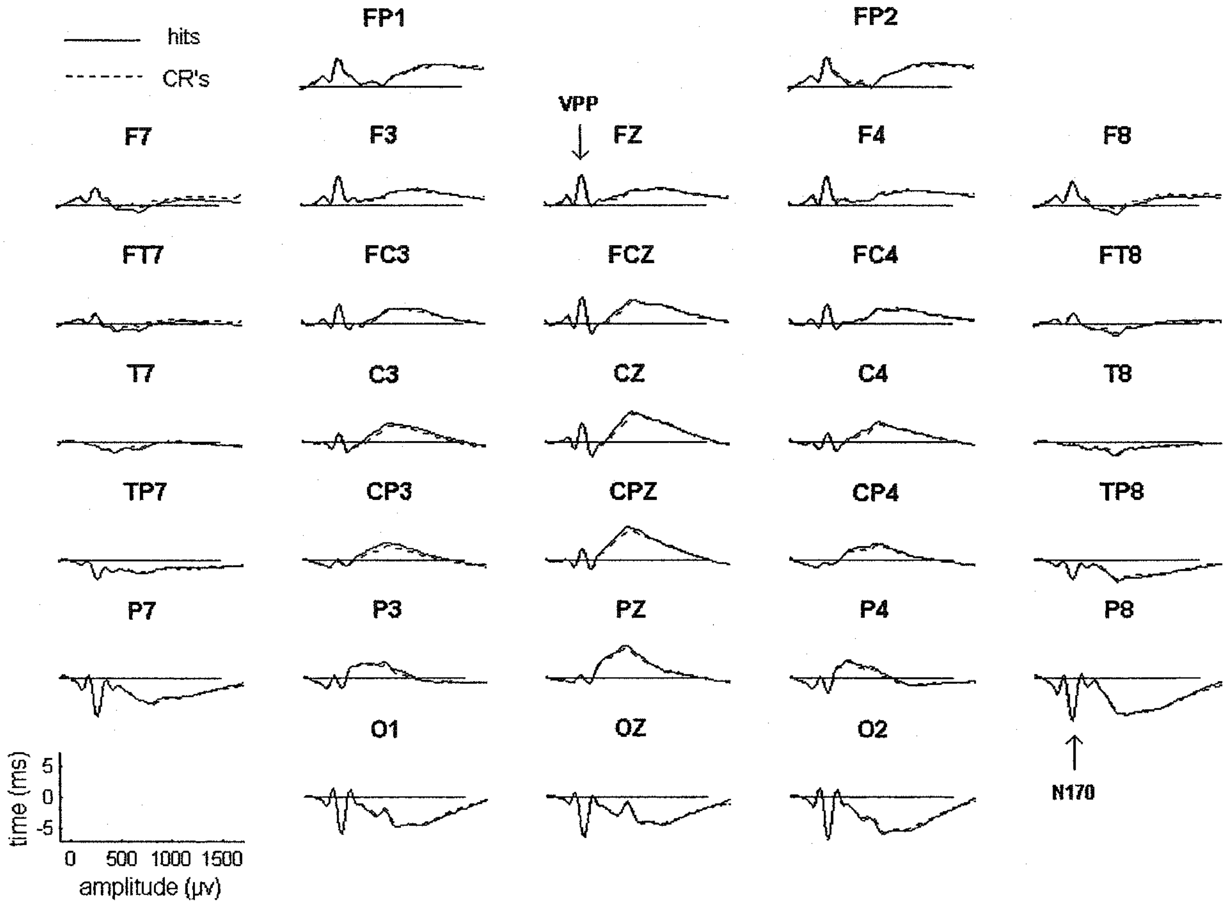

Figure 5-1. Grand averaged ERP waveforms corresponding to hits and correct rejections (CR's)

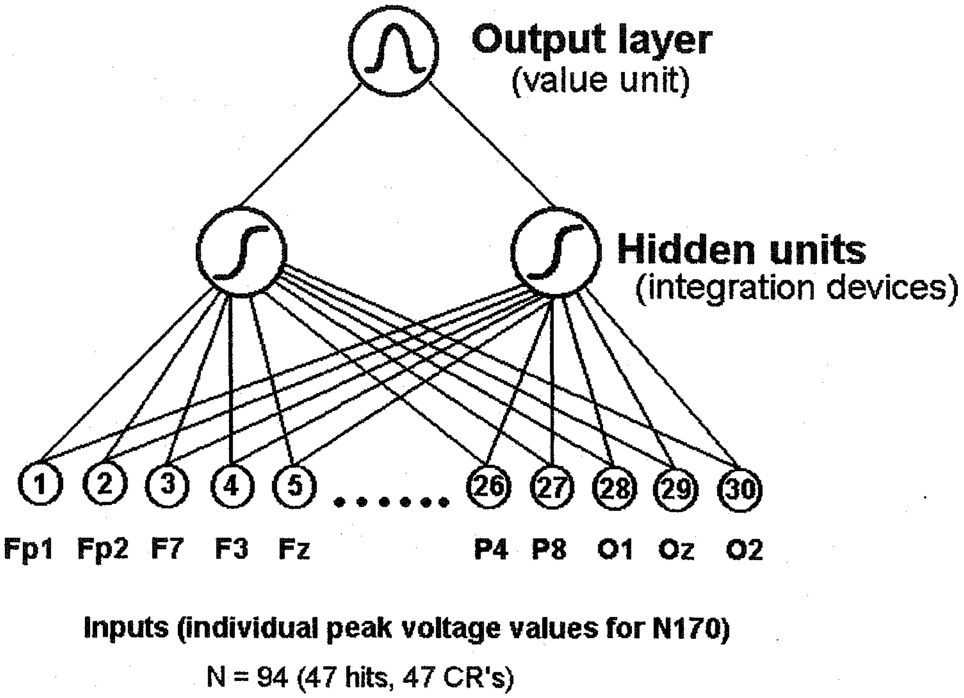

Figure 5-2. The network architecture adopted by the ANN's used in this study. Average voltage values represent averaged topographic voltage values corresponding to the average latency of the N170 obtained from each subject for each ERP type (hits and CR's). Lines between the layers of the ANN represent weighted connections.

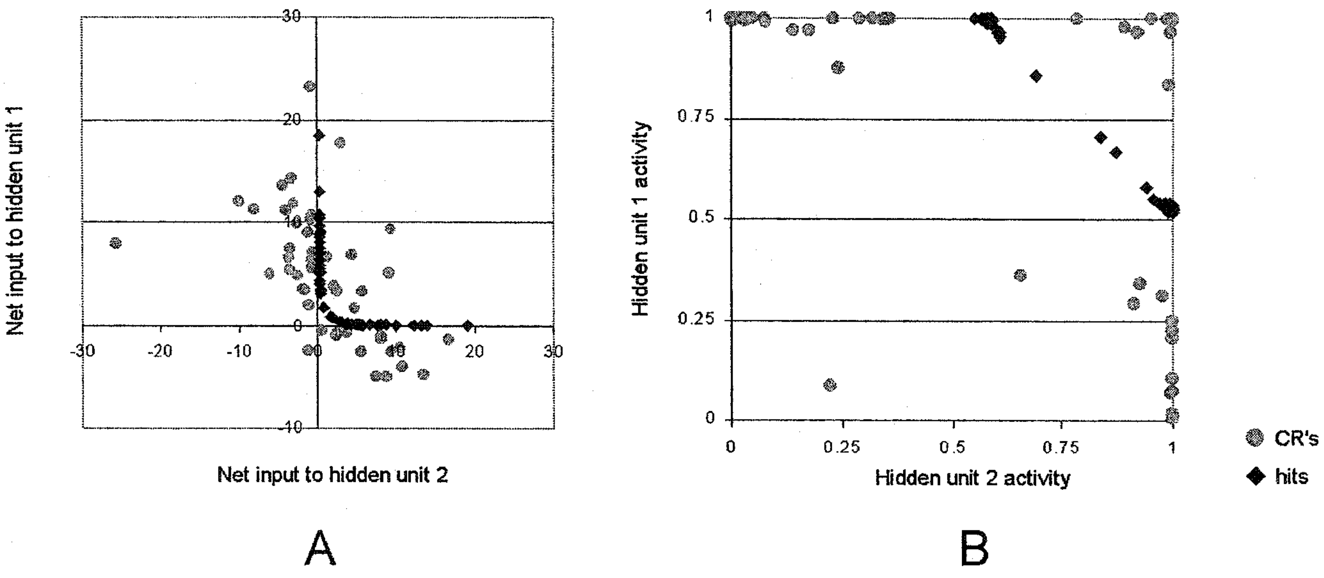

Figure 5-3. Two-dimensional scatterplots of A) net input to each hidden unit and B) hidden unit activity (net input after logarithmic transformation).

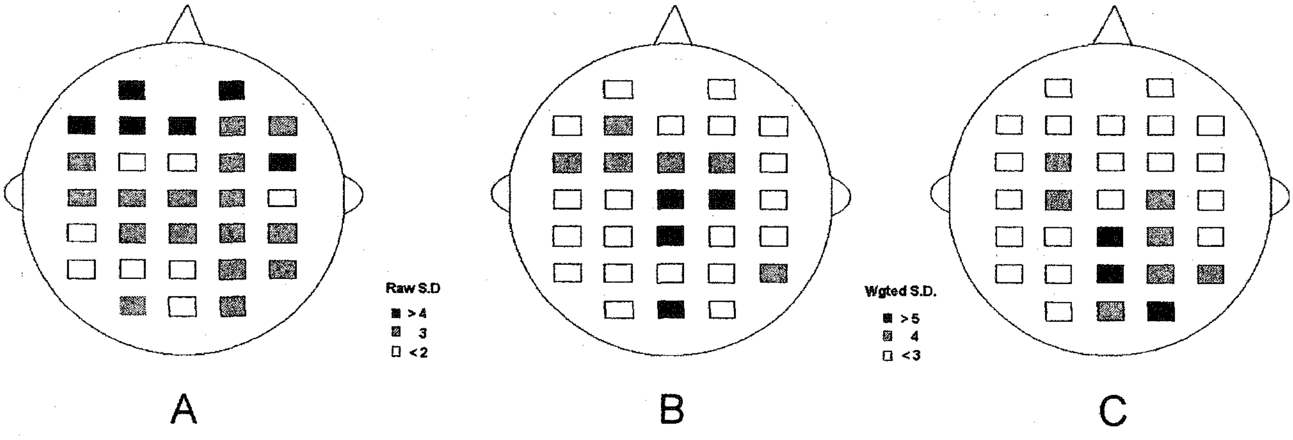

Figure 5-4. Topographic distribution of standard deviations A) before weighting (raw S.D.) B) after weighting with connections to hidden unit 1 C) after weighting with connections to hidden unit 2.

References

Allan, K., Wilding, E.L. & Rugg, M.D. (1998). Electrophysiological evidence for dissociable processes contributing to recollection. Acta Psychologica, 98, 231-252.

Bechtel, W. & Abrahamsen, A. (1994). Connectionism and the mind: An introduction to parallel processing in networks. Cambridge, MA: Blackwell Publishers.

Bentin, S. & Deouell, L.Y. (2000). Structural encoding and identification in face processing: ERP evidence for separate mechanisms. Cognitive Neuropsychology, 17(1-3), 35-54.

Bentin, S., Allison, T., Puce, A., Perez, E. & McCarthy, G. (1996). Electrophysiological studies of face perception in humans. Journal of Cognitive Neuroscience, 8(6), 551-565.

Campanella, S., Hanoteau, C., Depy, D., Rossion, B., Bruyer, R., Crommelinck, M.. & Guerit, J. M. (2000). Right N170 modulation in a face discrimination task: An account for categorical perception of familiar faces. Psychophysiology, 37(6), 796-806.

Dawson, M. R.W., Dobbs, A., Hooper, H.R., McEwan, A.J.B., Triscott, J. & Cooney, J. (1994). Artificial neural networks that use single-photon emission tomography to identify patients with probable Alzheimer's disease. European Journal of Nuclear Medicine, 21(12), 1303-1311.

Dawson, M.R.W. & Schopflocher, D. P. (1992). Modifying the generalized delta rule to train networks of nonmonotonic processors for pattern classification. Connection Science, 4, 19-31.

Eimer, M. (2000a). Effects of face inversion on the structural encoding and recognition of faces: Evidence from event-related brain potentials. Cognitive Brain Research, 10(1-2), 145-158.

Eimer, M. (2000b). Event-related brain potentials distinguish processing stages involved in face perception and recognition. Clinical Neurophysiology, 111(4), 694-705.

Eimer, Martin (2000c). The face-specific N170 component reflects late stages in the structural encoding of faces. Neuroreport, 11(10), 2319-2324.

Endl, W., Walla, P., Lindinger, G., Lalouschek, W., Barth, F.G., Deecke, L. & Wang, W. (1998). Early cortical activation indicates preparation for retrieval of memory for faces: An event-related potential study. Neuroscience Letters, 240, 58-60.

Graham, R. & Cabeza, R. (2001a). Evoked potentials of recognizing happy and neutral faces. Neuroreport, 12(2), 245-248.

Graham, R. & Cabeza, R. (2001b). Dissociating the neural correlates of item and context

memory: An ERP study of face recognition. Canadian Journal of Experimental Psychology, 55, 154-161.

Guillaume, F. & Tiberghien, G. (2001). An event-related potential study of contextual modifications in a face recognition task. Neuroreport, 12, 1209-1216.

Gupta, L., Molfese, D.L. & Tammana, R. (1995). An artificial neural network approach to ERP classification. Brain and Cognition, 27, 311-330.

Itier, R.J. & Taylor, M.J. (2002). Inversion and contrast polarity reversal affect both encoding and recognition processes of unfamiliar faces: A repetition study using ERP's. Neuroimage, 15, 353-372.

Jeffreys, D.A. (1993). A face-responsive potential recorded from the human scalp. Experimental Brain Research, 96, 193-202.

Kanwisher, N., Tong, F., & Nakayama, K. (1998). The effect of face inversion on the human fusiform face area. Cognition, 68, B1-B11.

McCarthy, G., Puce, A., Gore, J.S., & Allison, T. (1997). Face-specific processing in the human fusiform gyrus. Journal of Cognitive Neuroscience, 9(5), 605-610.

Potter, D.D. & Parker, D.M. (1989). Electrophysiological correlates of facial identity and expression processing. In J.R. Crawford & D.M. Parker (Eds.), Developments in clinical and experimental neuropsychology (pp. 137-150). New York: Plenum Press.

Rolls, E.T., Baylis, C.G., Hasselmo, M.E. & Nalwa, V. (1989). The effect of learning on the face selective responses of neurons in the cortex of the superior temporal sulcus of the monkey. Experimental Brain Research, 76, 153-164.

Rossion, B., Campanella, S., Gomez, C. M., Delinte, A., Debatisse, D., Liard, L., Dubois, S., Bruyer, R., Crommelinck, M., & Guerit, J. -M.(1999). Task modulation of brain activity related to familiar and unfamiliar face processing: An ERP study. Clinical Neurophysiology, 110(3), 449-462.

Seeck, M., Michel, C.M., Mainwaring, N., Cosgrove, R., Blume, H., Ives, J., Landis, T. & Schomer, D.L. (1997). Evidence for rapid face recognition from human scalp and intracranial electrodes. Neuroreport, 8, 2749-2754.

Slater, J.D., Wu, F.Y., Honig, L.S., Ramsay, R.E., & Morgan, R. (1994). Neural network analysis of the P300 event-related potential in multiple sclerosis. Electroencephalography and Clinical Neurophysiology, 90, 114-122.

Wu, F.Y., Slater, J.D., Honig, L.S. & Ramsay, R.E. (1993). A neural network design for event-related potential diagnosis. Computers in Biology & Medicine, 23(3), 251-264.

CHAPTER 6

GENERAL DISCUSSION

The goal of this dissertation was to examine face recognition within the framework of general and face specific memory processes with the hope of clarifying electrophysiological or ERP correlates of face memory. With regard to general visual memory processes, the first two studies provided evidence that the frontal and parietal ERP effects can also be elicited during face recognition using two different memory tasks, thereby fulfilling the criteria for general correlates of visual memory. The first study described the memory sensitivity of these effects to faces elicited in an old/new paradigm similar to that used with verbal stimuli. The second study reported the memory sensitivity of the parietal and frontal ERP effects to faces elicited in a context retrieval paradigm. The second two studies examined the memory sensitivity of putative face specific ERP correlates of face processing. The first of these studies was a temporal analysis of data from right temporal channels, which is believed to index the activity of unimodal face areas. The second study presented the results of two different topographic analyses of peak voltage values corresponding to the N170 and VPP, two potentials believed to be correlated with unimodal face memory. Together, these four studies help to elucidate the roles of general visual memory areas and face specific processing areas in face recognition.

General visual memory correlates

A central issue that arose from a review of the current literature regarded the electrophysiological correlates of general visual memory processes. Neuroimaging studies suggest that a network of cortical areas, primarily the parietal and frontal lobes, are involved in the retrieval of various types of stimuli. Putative ERP correlates of transmodal memory, on the other hand, have yet to be firmly established due to inconsistencies in experimental methods. The goal of the first two studies of this dissertation was to examine the generalizability of two ERP correlates of verbal memory to face memory.

In ERP studies, verbal recognition memory is associated with two positivities: one over parietal regions, and one over frontal regions. With nameable neutral stimuli, such as words or common objects, the parietal effect is usually left lateralized, and the frontal effect is usually right lateralized. We investigated the lateralization of these effects for nonnameable emotional stimuli: unfamiliar faces with happy and neutral expressions. The parietal effect was bilateral, suggesting that the left lateralization of this effect in studies using nameable stimuli reflected verbal processing. The frontal effect was left lateralized for happy faces, but right lateralized for neutral faces. This finding was consistent with the valence hypothesis, which posits that processing of pleasant emotions is lateralized to the left hemisphere and suggests that the frontal ERP effect may be modulated by the emotional valence of a stimulus.

This study suggested that the frontal and parietal effects index general memory processes. Their elicitation with faces, as well as words, provided evidence that these components are ERP correlates of transmodal memory. However, the frontal and parietal can also be elicited with different memory paradigms, such as source retrieval and temporal order paradigms. To ensure that these effects would also be elicited with a context retrieval task employing faces as stimuli, ERP's were recorded during a source memory paradigm. The results of this study are reported in Chapter 3.

In this study, the neural correlates of item and context retrieval were examined using ERP's. Subjects studied unfamiliar faces with happy or neutral expressions, and at test, they decided whether test faces were studied in the same or in a different expression, or were new. Both the frontal and parietal effects were elicited during this task, providing further evidence that these ERP effects index the activity of areas involved in general visual memory. The parietal ERP effect, which is hypothesized to indirectly reflect medial-temporal lobe (MTL) function, was sensitive to item retrieval, whereas the frontal ERP effect, which is thought to reflect prefrontal cortex (PFC) function, was sensitive to context retrieval. Converging with lesion, functional neuroimaging (PET and fMRI), and ERP evidence, these results support the notion that item retrieval is primarily associated with MTL function whereas context retrieval is primarily associated by PFC function.

Taken together, the results of these studies provide support for the notion that the parietal and frontal ERP effects index general recognition processes. These processes are engaged during a variety of tasks employing various stimulus types, converging with evidence from hemodynamic studies that implicate the involvement of prefrontal and parietal areas in general visual memory.

Face specific correlates

The impetus for Chapters 4 and 5 of this dissertation was to attempt to identify and characterize ERP correlates of face specific memory processes using conventional linear analyses and ANN's. Because this was a novel research direction, two fundamental ways to examine ERP's with ANN's were chosen: a temporal approach and a topographic approach. Chapter 4 describes the results of temporal analyses conducted on right temporal electrode channels with ANN's and an attempt to characterise the timecourse of early ERP face memory effects. The results of topographic analyses of N170 conducted with conventional linear statistics and ANN's using two different procedures for determining peak amplitude are presented in Chapter 5. The goal of this research was to determine the memory sensitivity of ERP effects that are believed to index the activity of face specific processing areas.

Despite inquiry, the timecourse of event-related potential (ERP) correlates of face memory has yet to be confirmed. In particular, the memory sensitivity of early ERP components that may be specific to faces has remained elusive. In Chapter 4, the possibility that early latency differences are modulated by memory was investigated with linear methods and artificial neural networks (ANN's). The rationale for this study was that early ERP memory effects exist but may not be reliably detected by linear analysis. Instead, a more powerful pattern recognition tool such as an ANN may be necessary to detect differences between ERP's from right temporal areas elicited to recognized and novel faces. ANOVA's were unable to distinguish between ERP types; however, ANN's were. Results suggest that early latency ERP's from right temporal areas do index memory related activity, but that this activity is in the form of higher-order relationships between voltage and time. Wiretapping revealed that classification was achieved through coarse coding in the hidden units and that a subset of timepoints seemed to be driving their activity.

Timepoints that were identified as predictors of net input by multiple regression were surprisingly consistent between and within channels. These timepoints had an almost bimodal distribution, with inputs between 100-300ms and 400-500ms dominating each solution. The first time window may correspond to the N170/VPP complex, while the second may correspond to the onset of the parietal effect. In addition, the predominance of adjacent timepoints in the solutions of the regression suggests that the ANN's were picking up differences in the onset latencies of both early and late ERP components. The possibility that early ERP correlates of face processing could be memory sensitive was examined further in a topographic analysis of the data, which is described in Chapter 5.

The purpose of the final study was to examine the memory sensitivity of two face sensitive ERP deflections, the N170 and VPP. We investigated the possibility that memory-related voltage changes are represented in ERP's but methodological considerations could be a factor in the outcome of results. We examined the effects of two peak analysis procedures in assessing the memory-sensitivity of the N170 and VPP elicited in an old/new recognition paradigm using analysis of variance (ANOVA) and artificial neural networks (ANN's). When latency was kept constant within subjects, ANOVA was unable to detect differences between ERP's to remembered and new faces; however, an ANN was. Network interpretation indicated that the ANN was detecting amplitude differences at occipitotemporal and frontocentral sites. When peak latency was taken into account, ANOVA was able to detect a significant decrease in onset latency of the N170 and VPP to remembered faces and amplitude differences were not detectable, even with an ANN.

Results suggest that the N170 and VPP do index memory-sensitive activity in face sensitive areas. This effect may be a repetition priming phenomenon that is attenuated at long lags between faces and can only be detected with sufficient power. Furthermore, choice of peak analysis procedures appears to be critical to the interpretation of phasic memory effects obtained from ERP data. These results provide evidence that methodological differences could account for mixed results regarding the memory sensitivity of the N170 and VPP evidenced in previous studies.

Summary

The experiments described in this thesis are an attempt to understand the neural mechanisms underlying face recognition as indexed with ERP's. The distinction between general and face specific memory provides a valuable framework for understanding the phenomenon of face memory, allowing us to determine which correlates are common to all memory processes and those which may be specific to face processing.

The results of the first two studies in this dissertation (Chapters 2 and 3) converge with ERP studies of verbal memory and together with neuroimaging studies, provide compelling evidence that the frontal and parietal effects are functional correlates of transmodal memory. The lateralization of these components does appear to have some stimulus specificity. However, these differences could be attributed to higher order stimulus properties (i.e. verbal vs. nonverbal stimuli, valenced vs. neutral stimuli) that cross stimulus domains and thus, are not stimulus specific in nature. In addition, the relatively late onsets of these effects are suggestive of ERP correlates of general visual memory. To what extent these effects also generalize to other modalities such as auditory memory presents interesting avenues for future inquiry.

The second two studies (Chapters 4 and 5) represent the attempt to determine the memory sensitivity of ERP's that have been uniquely associated to faces and attest to the utility of artificial neural networks in the analysis of ERP data. The results of these studies imply that early activity recorded over areas known to be involved in face processing is sensitive to the recognition status of a face and that this sensitivity could be manifested in the form of a reduction in latency for remembered faces. Whether this effect is due to priming, and the nature of this effect with varying lags between repeated faces awaits further scrutiny. However, preliminary results suggest that this effect is attenuated at longer lags. These studies also embody an attempt to reconcile contradictory results regarding early face-sensitive components that have been reported in the literature. Methodological inconsistencies that result from different recognition tasks and peak analysis procedures could account for mixed results with reference to the memory sensitivity of the VPP and N170.

By incorporating face recognition into a more general model of recognition, it is possible to understand those processes which make face memory special and those which it shares with other forms of retrieval. With regard to ERP correlates of general visual memory, future research should focus on range of experimental conditions under which the frontal and parietal effects can be elicited to determine if these effects are specific to visual memory or if they generalize to other modalities. Another fruitful line of inquiry would be to attempt to determine the generators of the parietal and frontal effects using source analysis and/or hemodynamic neuroimaging techniques, in order to provide information regarding the cortical substrates of these effects. In addition, the lateralization of these effects with different stimulus categories (e.g. words vs. faces, neutral vs. happy) warrants further scrutiny. Characterization of these processes in healthy young adults can then be used to aid in understanding how memory processes change with normal aging.

With regard to face specific ERP correlates of memory, several lines of research could also be pursued. Notably, the sensitivity of early ERP correlates of face processing to various cognitive manipulations remains unknown. For example, whether the N170 and VPP are sensitive to cognitive factors such as recognition status, selective attention and working memory load are enduring questions. In addition, the cortical generators of the N170 and the VPP in particular still remain matter of debate. Future research should attempt to determine the sources of the N170 and VPP, as well as identify other putative face specific ERP correlates, while behavioural manipulations could be devised to dissociate these potentials.

This corpus of research represents an attempt to describe and characterize event-related correlates of face recognition in terms of general visual and face specific memory and introduces intriguing possibilities for future research. In particular, technological advances which now allow ERP's to be recorded simultaneously during fMRI scanning will now allow for direct comparison of structure and function within subjects. Additionally, continued scrutiny of current analysis methods and the development of new and powerful analytic tools for data analysis will continue to increase our insight into the phenomenon of face recognition and how it is actualized in the human brain.

APPENDIX A

Electroencephalography has allowed us to index the electrical activity of the brain via scalp recordings, providing valuable insight into the neural events associated with different disease and cognitive states. One way of representing electrical activity is with event-related potentials (ERP's); event-locked potentials that correspond to the average change in the time course of an electroencephalogram (EEG). In order to understand the significance of ERP's in the study of cognitive phenomena, it is important to be familiar with how they are obtained, quantified and interpreted. The following provides an introduction to ERP's and discusses their particular strengths and weaknesses[1].

ERP's are recorded from various electrode sites on the scalp and refer to the changes in the brain's electrical activity that are time-locked to some distinct event (i.e., the presentation of a stimulus). These changes are very small in comparison to the background EEG; the continuous random neural activity of the brain. In order for ERP's to be discernible, raw EEG associated with a given trial type is averaged, allowing for small, consistent changes to be detected from a background of EEG noise. Random unrelated activity is assumed to average out over trials, leaving waveforms that are estimates of the time-course of neural activity elicited by different classes of stimuli or cognitive states.

The resulting ERP waveform can be parsed into various components. Various methods exist for determining individual components; however, usually these are identified through visual inspection of the grand averaged waveforms. Positive and negative peaks, both phasic and sustained, can then be selected for subsequent

---

[1] For a more complete description of ERP methodology see Picton, Lins & Scherg, 1995.

examination. The nomenclature of these components varies, but can be one of three basic types. Components can be named in terms of their polarity and peak latency (e.g. N170, P300), their peak distribution on the scalp (e.g. vertex positive potential, parietal effect), or the conditions under which they are typically elicited (e.g. readiness potential, mismatch negativity).

ERP's can be one of two types. Exogenous ERP's occur at early latencies (10-100 msec) and are dependent on physical stimulus parameters, such as stimulus intensity, rather than on the psychological characteristics of the subject or the testing situation. These ERP's are assumed to reflect the activity of primary sensory areas. Longer latency deflections (>100 msec) are referred to as endogenous ERP's. Endogenous ERP's differ from exogenous ERP's in that their scalp distribution is often independent of the modality of the eliciting stimulus, and they are dependent on behavioural and psychological processes related to an event. ERP's, in particular, endogenous ERP's are useful in the study of memory and other cognitive functions for two reasons.

First, scalp-recorded neural activity associated with the processing of various stimuli can be recorded with a temporal resolution sufficient to detect transient changes in activity that last only for a few milliseconds. Therefore, sensitive estimates regarding the time course of certain ERP components can be made, allowing the researcher to determine the time required by the nervous system to discriminate between different types of stimuli (i.e., remembered vs. novel faces; Allan et al., 1998). ERP's also have tremendous utility in the investigation of whether different experimental conditions elicit functionally dissociable cognitive processes. In demonstrating a dissociation of an ERP component from behaviour, it is possible to make inferences about the processes that give rise to various ERP components.

The following rationale, adapted from Ganis, Kutas and Sereno (1996) serve as useful criteria for the dissociation of ERP components elicited under different experimental conditions: 1) if components are identical in their onset and/or peak latency, amplitude and distribution, strong evidence exists that common brain mechanisms are involved in the processing of the two stimuli; 2) ERP components that differ only in terms of latency (i.e., differ only in onset and/or peak latency, but not amplitude or distribution) are compatible with the notion of common processes; however, the stimulus class showing shorter latencies may have privileged access to processing resources; 3) ERP components that differ only in amplitude, with no differences in latency or topographic distribution also favour a common system; however, the stimulus class showing larger amplitudes may require more processing resources; 4) the absence of an ERP component for one class of stimuli, as compared to the other provides very weak evidence that the stimuli engage functionally dissociable cognitive processes, with one stimulus class eliciting the involvement of a different population of neural generators; 5) if ERP's to the two classes of stimuli are similar in morphology, but differ in their distribution, both stimuli are processed in similar ways, but by different brain areas; 6) if ERP effects differ in both morphology and distribution, strong evidence exists for functionally dissociable cognitive processes.

ERP studies offer an advantage over behavioural studies of memory because they allow the researcher to examine the temporal relationship between a stimulus and subsequent neural activity in participants while they perform the actual task. Therefore, ERP's are an index of processing that may show independence from behaviour (i.e., reaction time or accuracy). The temporal resolution of ERP's is a primary strength of this methodology. However, the spatial resolution of ERP's leaves much to be desired. For the most part, it is assumed that ERP's index extracellular potentials associated with the electrical activity of groups of cortical columns; however, volume conduction associated with scalp recordings makes it difficult to determine the exact sources of ERP activity. Although improved source localization algorithms are continually under development, convergent evidence from studies with superior spatial resolution (i.e. hemodynamic neuroimaging) should also be considered.

References

Ganis, G., Kutas, M. & Sereno, M.I. (1996). The search for "common sense": An electropysiological study of the comprehension of words and pictures in reading. Journal of Cognitive Neuroscience, 8(2), 89-106.

Picton, T.W., Lins, O.G., & Scherg, M. (1995). The recording and analysis of event-related potentials. In F. Boller and J. Grafman (Eds.), Handbook of neuropsychology (Vol.10), (pp.3-73).

APPENDIX B

Artificial neural networks (ANN's) are computer simulations of networks consisting of simple, interconnected processing devices. Despite the simplicity of their components, ANN's have powerful pattern recognition capabilities. The following serves as an introduction to ANN's, describing their elements and operation. Emphasis is placed on the describing properties of ANN's relevant to those used in this dissertation. Nevertheless, some general statements regarding ANN's can be made. In general, an ANN can be conceptualized as a multi-layered system that produces a specific response to an input stimulus. ANN's consist of three basic components: processing units, modifiable connections between units, and a learning procedure.

Processing units are the basic elements of ANN's and perform the computations in the network. The activity of any unit depends on three factors: the net input function, the activation function, and the output function. The net input function determines the total signal to a given unit, typically the sum of all signals that a unit receives from the other units to which it is connected, also known as the sum of the weighted inputs. Therefore, the net input to any unit *i* that receives input from *j* other units can be described by the following equation:

$$\text{Net input}_i = \sum_j w_{ij} \, \text{output}_j$$

The activation function of a unit determines the internal activity of the unit, which is in turn dependent upon the net input to the unit. Several different activation functions exist that are be used to describe unit type. The fundamental difference between the types of units characterized by different activation functions is how they transform net input into unit activity. For example, one common processing unit type is the integration device, which transforms net input into unit activity with a sigmoidal or logarithmic function:

$$\text{Unit activity}_i = \frac{1}{1 + e^{-(\text{net input}_i + \text{bias})}}$$

Another unit type, the value unit, transforms net input with a Gaussian function (Dawson & Schopflocher, 1992):

$$\text{Unit activity}_i = e^{-\pi(\text{net input}_i - \text{bias}_i)^2}$$

Finally, the output function determines the magnitude of the signal output by the unit, a value that is usually equivalent to unit activity.

Connections are essential for communication between units. In ANN's, processing units are connected to other processing units, usually in parallel. These connections are usually weighted, that is they alter the nature and strength of the signal sent along the connection as function of its current weight. Connection weights can be either positive (excitatory) or negative (inhibitory), and can be modified by the network as it learns. The strength of the connection is determined by the size or absolute value of the weight. Networks of units can be interconnected in many different ways to form different network architectures. This pattern of connectivity essentially defines the causal relations between the processing units. However, attention will be limited to the simple three-layered architecture used in this dissertation. The bottom layer of the network or input layer provides external input to the network. The top layer or output layer send signals outside the network itself. An intervening layer of units, the hidden layer, do not have external access and only communicate directly with the input and output layers.

Learning rules are algorithms that are used to make changes to the modifiable connections between units. Although they are not a necessary component of all ANN's, they are critical to networks that learn because they are used to optimize the performance of the network. Networks employing learning rules must be trained, which normally consists of the repeated presentation of input patterns. In this manner, randomly determined connection weights can be modified to that the network is better able to produce the appropriate response to any given set of input patterns. One of the most popular learning rules is the generalized delta rule (see Rumelhart, Hinton & McClelland, 1986).[1] This rule has since been modified to accommodate networks with value units (Dawson & Schopflocher, 1992).[2]

With this rule, learning proceeds by presenting the ANN with input patterns from the training set. For any given pattern, this results in signals sent along weighted connections from input layer to the hidden units, which then sends signals to the output layer. At this level, the actual activation of the output units is compared to their desired outputs (specified by the investigator) and an measure of error is obtained. In this manner, a measure of total error is obtained for the presentation of the entire training set. This error term is used to modify the weights of the connections between the output layer and the hidden layer. Error is then sent through these newly modified connections to the hidden layer. Units in the hidden layer use this information to modify connection weights between the hidden and input layers, such that total error is reduced. Changes in

---

[1] For a more detailed mathematical description of the generalized delta rule, see Rumelhart, Hinton & McClelland, 1986.

[2] For a more detailed mathematical description of the modified generalized delta rule, see Dawson & Schopflocher, 1992.

connection weights to the presentation of the training set are typically quite small and training often requires numerous presentations of the training set, also known as sweeps or epochs. However, through application of the learning rule and the repeated presentation of input patterns, networks gradually learn to produce correct responses to the input patterns in the training set. When this occurs, the ANN has converged and training ends. By the end of this procedure, the network will have a specific pattern of connection weights relative to the initial random connections, which can then be used for network interpretation.

References

Dawson, M.R.W. & Schopflocher, D. P. (1992). Modifying the generalized delta rule to train networks of nonmonotonic processors for pattern classification. Connection Science, 4, 19-31.

Rumelhart, D.E., Hinton, G.E.. & McClelland, J.L. (1986). Learning internal representations by error propagation. In J.L. McClelland, D.E. Rumelhart & the PDP Research Group (Eds.), Parallel Distributed Processing (Vol.1, pp. 45-76). Cambridge, MA: MIT Press.