

Effects of Priming on Subsequent Associative Memory: Testing Prediction Error and Attentional Accounts



Alexander J. Kaula

MRC Cognition and Brain Sciences Unit
University of Cambridge

This dissertation is submitted for the degree of
Doctor of Philosophy

Darwin College

June 2018

Abstract

Effects of Priming on Subsequent Associative Memory: Testing Prediction Error and Attentional Accounts

Alexander J. Kaula

The Predictive Coding framework is a general theory of brain function which proposes that the brain creates a hierarchical model of the world, with higher levels predicting, based on previous experience, inputs from lower levels (and ultimately the sensory input). This framework entails feedback connections carrying predictions and feedforward connections carrying error signals. Divergences of inputs from those expected are termed prediction errors (PE), and indicate the possibility of updating the model to improve future performance. Thus, learning should be driven by PE. Feedforward and feedback signalling have been widely studied in the fields of reward learning and perception, but although there are strong reasons to expect related processes in memory, less work has been done to investigate this. One difficulty addressing this question concerns the role of attention in memory formation; although the roles of PE and attention are theoretically distinct, when events are surprising we are likely to attend more to them, and attending to events makes them more likely to be remembered. The aim of this research is therefore to de-confound effects of PE and attention on memory, in order to test the explanatory power of the predictive coding framework applied to memory processes, both at the behavioural and, using neuroimaging techniques, at the neural level.

Preface

This project was carried out at the Medical Research Council Cognition and Brain Sciences Unit (MRC-CBU) under the supervision of Professor Rik Henson, and the secondary supervisor was Dr Mike Anderson. My research was supported by an MRC Studentship.

This dissertation is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared in the Preface and specified in the text. It is not substantially the same as any that I have submitted, or, is being concurrently submitted for a degree or diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Preface and specified in the text. I further state that no substantial part of my dissertation has already been submitted, or, is being concurrently submitted for any such degree, diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Preface and specified in the text. It does not exceed the prescribed word limit for the relevant Degree Committee.

Data and results reported in Greve, Cooper, Kaula, Anderson, & Henson, (2017) are not presented here, but came from an experimental paradigm similar to one that is used extensively in this thesis.

Acknowledgements

First and foremost, I thank my supervisor, Rik Henson, for whose supervision I am forever grateful. Rik has been an inspiring and considerate guide. He has been careful and patient in his guidance, pointing out potential footholds and hazards while being sure of reminding me that the direction of the work remained my own.

We have been grappling with a slippery beast of an idea. It has been an adventure with its share of blind alleys, secret trapdoors and bewildering wide-open spaces, but it has been a rewarding one. Rik has led by example, showing disinterest with regard to the results of this work and its consequences (or lack thereof) for 'his' theory, content to let the results be whatever they are. Remaining good-humoured and supportive throughout, I have never heard him curse my many lapses. Rik has treated me with respect, kindness, and decency, encouraging me at every turn to be the best scientist I can be.

I am grateful to Jiří Čevora, friend, housemate, labmate – for illuminating discussions at all times of the day and night – and for knowing when it has been time to continue (or stop) over a beer. I am grateful to Andrea Greve for her insightful conversation and support throughout this project, and to Mike Anderson for his wise words and useful pointers. And, to my dear friend Taylor Schmitz, who helped me through proof reads and conceptual tangles alike. My thanks also to Michael Ewbank, a friend and a mentor.

The remaining people to whom I owe thanks at the CBU are too numerous to name: the CBU is especially collegiate, and there are many – my cohort, colleagues, friends, and acquaintances here who have generously given up their time to help me, not least those who work hard to keep the administrative paperwork flowing where it must.

Finally, to my sister, Miranda, who has stood by me throughout this project, and walked with me through the most difficult of times.

This thesis is dedicated to the memory of Mum.

Table of Contents

Abstract	1
Preface	2
Acknowledgements	3
Chapter 1 General Introduction	9
Priming Effects on Subsequent Associative Memory	10
Initial PSM Experiment of Gagnepain et al.	11
PIMMS and PSM	12
Other Related Studies	16
Item-Context Interactions	20
Perceptual Attention and Working Memory	22
Resource Accounts of PSM	24
Project Overview	24
Chapter 2 Testing Resource Accounts I: Temporal and Perceptual Resources.	27
Introduction	28
Experiment 1: Temporal Resource Account	30
Method	31
Participants and Materials:	31
Procedure:	31
Statistical Analysis	34
Results	35
Discussion	36
Experiment 2: Perceptual Resources	37
Method	38
Results	39
Discussion	41
Chapter Summary	42
Chapter 3 Testing Attentional Accounts II: Central Attention.	43
Introduction	44
Method	44
Results	45
Discussion	46

Experiment 4(a)	48
Method	48
Participants and materials	48
Procedure	49
Statistical analysis	49
Results	50
Discussion	50
Experiment 4(b)	50
Method	51
Participants	51
Procedure	51
Statistical analysis	52
Results	53
Primary Face Task	53
Secondary Duration Report Task	54
Discussion	55
Meta-Analysis Across Experiments 1-4	56
Chapter Summary	58
Chapter 4 Testing Attentional Accounts III: Computational Model	61
Introduction	62
Relating resources to performance	62
Modelling PRT, PSM, and load task performance	64
1. Participant Resources and Resource Allocation	64
2. Effect of Priming	65
3. Correlations Between PRT and PSM	66
4. Effects of Load	70
Experiments 3 & 4b	73
Chapter Summary and Discussion	76
Chapter 5 Using fMRI to Examine the Neural Basis of PSM	81
Introduction	82
Predictions of Resources Account	82
Predictions of PIMMS Account	85
Summary	86
Method	87
Participants and Materials	87
Adaptation of Behavioural Paradigm for fMRI	87

fMRI Acquisition and Preprocessing	88
Analysis	89
First-level	89
Second level	90
Region of Interest (ROI) analysis	90
Whole brain analysis	91
Representational Similarity Analysis (RSA)	91
Beta-series Connectivity ROI analysis	92
Results	93
Behavioural Results	93
Univariate ROI Results	93
Whole-Brain Analysis	98
Representational Similarity Analysis	98
Beta-series Connectivity Analysis	101
Discussion	102
Summary of Results	102
Univariate Results and Resource Accounts	104
Multivariate Results and The PIMMS Account	106
Chapter Summary and Potential Caveats	106
Chapter 6 General Discussion	109
Overview	110
Summary	110
Main Results	111
Limitations and Possible Future Directions	113
Behavioural	113
Correlations Between Outcome Measures and Presumed Shared Resources	114
Neuroimaging	114
Prediction Error and Attention	115
Reconciling or Deciding Between PIMMS/PE and Resources Accounts	116
Concluding Remarks	118
Bibliography	119

Chapter 1

General Introduction

Priming Effects on Subsequent Associative Memory

This thesis concerns the effect of prior exposure to a stimulus on the creation of new episodic memory, specifically the ability to associate that stimulus with a background context (associative memory). Effects of priming are typically seen in indirect measures such as in speeded RTs, or biased decisions as a result of prior exposure. However, several studies have observed that priming a stimulus can also improve its subsequent memory, not only for the item itself, but for the priming episode (e.g., Gagnepain et al., 2011; Gagnepain, Lebreton, Desgranges, & Eustache, 2008). This will be called “the effect of Priming on Subsequent Memory” (PSM). PSM is interesting because priming is traditionally thought to affect implicit memory, such as facilitation of perceptual and semantic processes (Henson, 2003), but not explicit (episodic) memory per se (Cohen & Squire, 1980; Schacter & Tulving, 1994). Henson and Gagnepain (2010) introduced a novel memory framework (‘Predictive, Interactive, Multiple Memory Signals’, PIMMS) which provides a possible explanation for the effect of priming on associative memory: priming is assumed to sharpen the sensory evidence for an item, which in the presence of an uninformative context (prior), causes a larger prediction error (PE), and PE is assumed to drive learning of new associations, such as that between the context and the item. This Introduction will present the design and results of the original Gagnepain et al. studies, before explaining the PIMMS-based account in more detail. Other studies examining the effect of priming on memory will then be considered, and an alternative explanation raised in terms of attentional resources, whereby priming ‘frees-up’ of resources which can then improve encoding of item-context associations. The purpose of this thesis was to compare the PIMMS account and the resource accounts, in particular by testing predictions of the resource account. The different types of resources (e.g., temporal, perceptual and central) will be briefly considered in the rest of the Introduction, before previewing the four behavioural experiments, computational modelling, and functional neuroimaging experiment that were conducted to address the two accounts of PSM.

Initial PSM Experiment of Gagnepain et al.

Gagnepain et al. (2008) showed that priming a set of words by presenting them during a simple task on one day led them to be better-associated with the context in which they were encountered the following day during performance of a different task (compared to unprimed words). The experiment had 3 phases: a priming phase, a critical study phase, and a (surprise) test phase (see Figure 1). In the priming phase, participants listened to single words, followed after a short interval by a spoken phoneme, and were required to indicate the presence or absence of that phoneme in the word they had just heard. The task therefore encouraged reflection on the perceptual qualities of the words, with the intent being to enhance priming effects. Words were pseudo-randomly ordered and split into 3 groups: high-primed, low-primed, and distractors ('filler' words). Over the course of the task, high-primed words were encountered 3 times, low-primed words once, and distractors once. The next day, the high- and low-primed words were encountered during an auditory lexical decision task, in which they were presented together with new words and pseudowords. Importantly, all the words were presented against a 'background context' of either the sound of applause, or the sound of a motorcycle race. A surprise test followed this critical study phase, in which participants were presented with the high-primed, low-primed and unprimed words from the critical study phase, plus never-before heard foils. The task was to indicate for each word whether it had been presented in the previous lexical decision task, and if participants thought it had, then to indicate whether this memory was accompanied by recollection of episodic details ('Remember', 'R' response), or if not, whether they were confident they had heard the word but could not recall episodic details ('Know', 'K' response), or if they were guessing ('G' response), based on the Remember/Know procedure of Tulving (1985). Finally, in case of an R response, participants were asked to indicate in which context sound (motorcycle race or applause) the word had been heard, and if correct, this was called an $R_{\text{Justified}}$ response.

The number of correct K responses was significantly increased by each level of the priming condition (i.e. most ‘ K ’ responses for high-primed words). This is not particularly surprising, because across the two days, high-primed words had been seen more times in total than low-primed words, and low-primed words more times in total than unprimed words. More surprising was the finding that priming also increased the number of R and $R_{\text{Justified}}$ responses, which indicates subjectively and objectively better memory for the particular occurrence of a word during the Study phase (if anything, prior presentation in the priming phase might have been expected to cause confusion of the episode in which a primed word was presented, i.e., impair memory through proactive interference). Moreover, for the high-primed words, the amount of RT priming was correlated with the number of R and $R_{\text{Justified}}$ responses. This priming effect on subsequent memory (PSM) is the subject of this thesis. Note that Gagnepain et al. (2011) used a very similar paradigm and replicated this result, while also addressing the underlying neural basis of these effects, as will be discussed in Chapter 5.

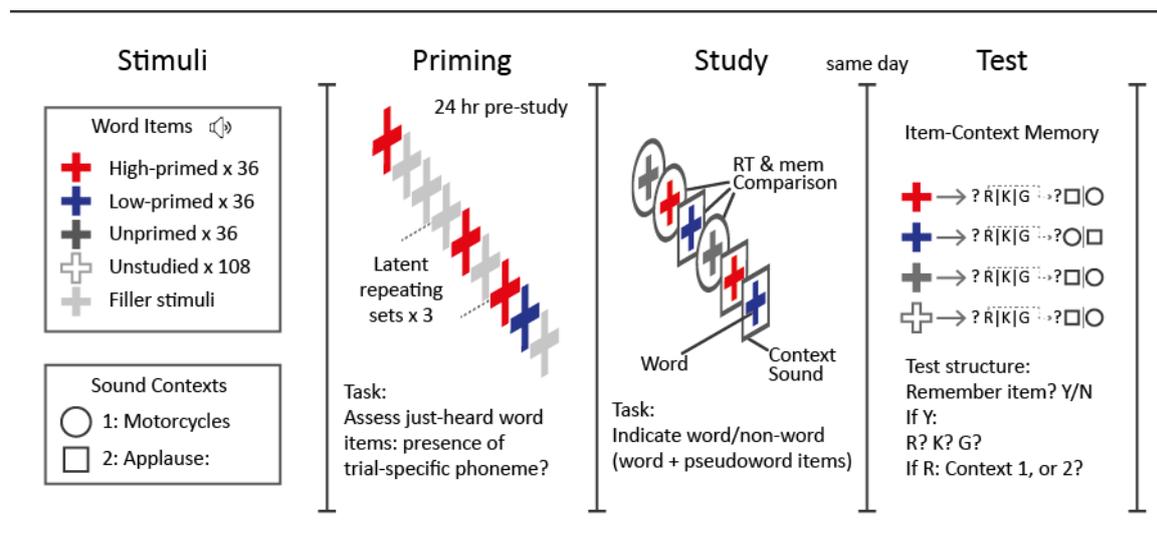


Figure 1: Simplified design of Gagnepain et al. (2008)

In a study using a related design, Poppenk, Köhler, and Moscovitch (2010) showed a commensurable effect to that described in Gagnepain et al. (2008), using visually-presented word stimuli. Poppenk et al. found that within-session prior exposure to proverbs, as well as existing prior knowledge, each enhanced subsequent recall memory for the study-phase task context in which those proverbs had been encountered, where this varied over study trials between rating either the vividness, or the valence, of the proverbs.

PIMMS and PSM

While the initial account by Gagnepain et al. (2008) for their results focused on the presence of reinstated perceptual traces within the recalled critical study episode, Henson and Gagnepain (2010) subsequently proposed a reinterpretation of PSM in terms of the PIMMS framework. PIMMS is based on the Predictive Coding (PC) hypothesis of brain function (e.g. Feldman & Friston, 2010; Knill & Pouget, 2004; Lee & Mumford, 2003; Rao & Ballard, 1999, for reviews) which argues that hierarchical message-passing in the neocortex solves the problem of inferring the causal structure of the environment. Optimal statistical inference given new data about the world requires combining beliefs with new evidence. Beliefs are conveyed by top-down feedback, which has been extensively studied in perception (e.g. Rao, 2005), action (e.g. Fletcher & Frith, 2009), and reward learning (e.g. Schultz, 1997). According to PC, efficient processing of information by the brain relies on minimizing the disparity, or prediction error (PE), between these top-down feedback (prediction) signals and bottom-up (evidence) signals. According to a Bayesian framework, the top-down signals encode priors, i.e., the expected inputs given the context, or a probabilistic representation attempting account for the likely cause of the inputs (Bastos, Usrey, Adams, & Mangun, 2012). In the absence of any priors, inputs from our senses would simply propagate up through the system as 'sensory evidence'. However, in the presence of top-down predictions, the signals propagating up through the system become error signals conveying the aspect of current input that is not predicted from higher levels, and through recurrent interactions, these errors alter the predictions themselves. This interactive process of perception is assumed to be rapid: it has been argued that the whole hierarchy may settle on a coherent model which minimises PE across multiple cortical levels within 60-80 ms of interactions (Lee & Mumford, 2003).

The PIMMS framework extends the idea of PC into the study of memory. Within this framework, the function of memory is considered as the long-term reduction of PE; the slower, synaptic-based complement to the rapid, activity-based processes of perception. PIMMS predicts that memory formation is driven by residual PE (after perceptual

processing has stabilised), which indicates unexplained input and therefore the necessity of updating future predictions.

PE at any level of granularity can be viewed as the divergence (lack of overlap) between two probability density functions: the prior (expected input) and the likelihood (sensory evidence). These functions are combined under Bayesian principles to yield a posterior, which is the updated model. Both of these probability densities may be more or less sharply defined, so as to cover a greater or lesser range of possible values. This may be thought of as the precision of the prediction: a precise prediction covers only a small range of values and is therefore sharply peaked, whereas the more uncertain a prediction, the flatter will be its prior. Greve et al. (2017) demonstrate a variety of ways to vary PE in experimental contexts (see Figure 2, below).

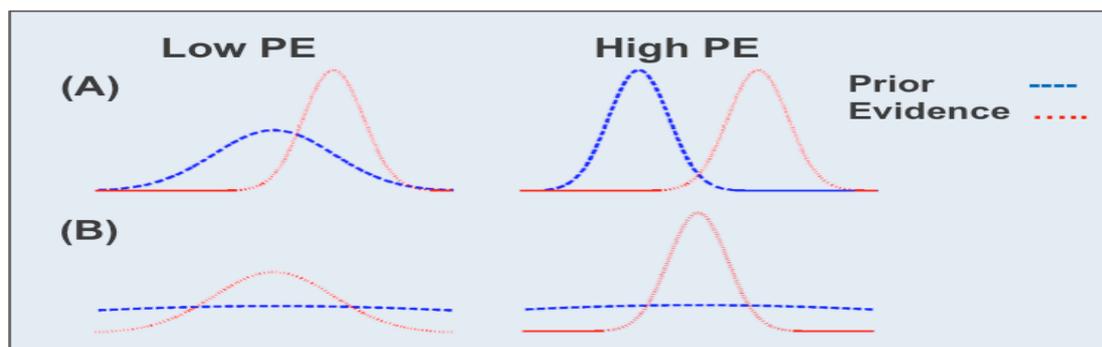


Figure 2: Panels A & B show 2 different approaches to manipulating PE in a PIMMS framework. Y-axes represent probability; x-axes represent a hypothetical perceptual dimension. In their Experiment 2, Greve et al. (2017) varied the precision of the prior (blue) while the sensory evidence (red) was constant (Panel A). In Experiment 3 of Greve et al., and the experiments of Gagenpain et al. (2008, 2011), and those in the present thesis, the PIMMS hypothesis is that sensory evidence is varied (via repetition priming), while keeping the (weak) prior constant (Panel B).

Manipulation of the precision of either of the prior or likelihood functions yields variation in PE and hence should lead to greater or lesser learning. Greve et al. (2017) showed just such effects. In Experiment 2, they showed that increasing the precision of a prior (Panel A, Figure 2), in the presence of divergent likelihood, improves memory performance. More specifically, the precision with which a scene predicted a particular face was shown to modulate later memory for the pairing of that scene with a new face. In a training phase of the experiment, ‘Low Precision’ scenes were presented with a new face on each occasion, whereas ‘High Precision’ scenes were always shown with the same face. In a subsequent Study phase, these scenes were paired with a novel face, and

memory for this novel scene-face pairing was tested in a final Test phase (using a forced-choice procedure in which all the choices were from the study phase). When the expectations that were violated in the Study phase were more precise in the 'High Precision' condition, and hence PE greater, the association between the scene and the (violating) face was indeed better remembered.

Conversely, in Experiment 3 of Greve et al. (2017), rather than manipulating the precision of the prior belief, the precision of the sensory evidence was manipulated while the prior remained consistent (Panel B, Figure 2). This involved presenting some faces (without any scenes) several times in the training phase (priming faces), and then pairing them with new scenes in the study phase, together with pairings of new (unprimed) faces and scenes. Because the scenes were new, and the pairing with the face was arbitrary, the prior expectation of a face given a scene was flat (uninformative). As expected, memory for the face-scene pairing from the study phase was better for primed than unprimed faces, consistent with a greater PE between the flat prior is accompanied by a sharper likelihood function.

Experiment 3 of Greve et al., (2017) also used an orthogonal manipulation of image quality of the face, which was also assumed to affect the precision of its likelihood function, and again the PIMMS prediction of better scene-face memory for less degraded (sharper) faces was confirmed. This manipulation is less relevant to the current discussion about priming and Gagnepain et al.'s (2008, 2011) results, though is relevant to Experiment 2 in Chapter 2 of this thesis.

The same PIMMS account, that priming sharpens the likelihood (of auditory words presented against background sounds), resulting in greater PE in the context of an uninformative prior (the two sounds) and hence better word-sound association (and also leading to increased R responses because of this association) can therefore apply to Gagnepain et al.'s results.

Other Related Studies

Although dissociations between implicit and explicit memory have been much studied, there are fewer studies looking at commonalities between these two types of memory, and only a handful that have specifically examined the effects of priming on subsequent memory. The two most relevant are a study by Wagner, Maril, and Schacter, (2000), which found that increased behavioural and neural priming correlated with worse subsequent memory for studied items, and a study by Turk-Browne, Yi, and Chun, (2006) which found the apparently converse effect, though using a different paradigm (Panels A & C, Figure 3). These studies are presented in more detail below, and their explanations are compared to that of PIMMS above, but it is worth noting up front that one important difference is that these studies focused on subsequent memory for the items themselves, rather than associations between items and their context.

Wagner et al. (2000) combined priming and subsequent memory for incidentally-encoded words in an fMRI study which manipulated the lag between repeated word presentations: some words had been seen on the previous day, while the others were seen 20 mins before. The scanning was then done in a subsequent study phase, in which words from both lags were intermixed with filler words for a lexical decision task, and then a final recognition memory test was performed outside the scanner, in which participants had to decide whether a test word had been presented in the scanner.

Prior studies have shown that short lags lead to increased priming, whereas longer lags between repetitions tend to improve subsequent memory (the “massed versus spaced” effect, e.g. Ebbinghaus, 1885/1913). Priming of (semantic processing of) words has been associated with reduced activation – so-called ‘Repetition suppression’ (RS; Henson, 2003) – in left inferior prefrontal cortex (LIPC), whereas subsequent memory for words has been associated with increased activation of LIPC (e.g., Wagner, 1998). Wagner et al. (2000) therefore hypothesised that words showing greater RS in LIPC would be remembered less well. Indeed, behavioural priming, as measured by RT difference between unprimed and primed words, was found to correlate negatively with the probability of subsequent recognition. This was also reflected in the fMRI data: participants who showed greater subsequent memory advantages for primed items

showed least evidence of RS when comparing LIPC activation between unprimed and primed trials in the long-lag condition. The authors propose that this finding supports an ‘encoding variability’ account (Martin, 1968), whereby spacing repetitions (long lag) increases the variability with which they are processed (e.g., increases the degree of semantic elaboration), which increases the chance of the item subsequently being remembered. Conversely, (semantic) priming is higher when an item is processed more similarly on each presentation, which is more likely at short lags.

The explanation of Wagner et al. (2000) for the inverse relationship between priming and subsequent memory seems reasonable. However, some alternative hypotheses are worthy of consideration. One simple possibility is that subsequent memory is related to the time spent processing an item (i.e., a “time on task” account). Because larger priming is associated with shorter RTs on second presentation, this means there was less time spent processing a primed item in the critical study phase, and so less time to encode it fully.

A second possibility is that, because Wagner et al. (2000) only measured recognition memory, they may have missed a correlation between greater priming and better subsequent context memory. This is because recognition memory can be supported by an acontextual feeling of familiarity (separate from explicit recollection of the study phase, Yonelinas, 2002), and so it is possible that subsequent memory performance was influenced mainly by familiarity (for which the PIMMS account of PE driving encoding of context-item associations does not apply). Admittedly, there are counter-arguments to this familiarity-based account: 1) familiarity would be expected to be higher in the short-lag condition, since the initial presentations were closer to the time of test than the initial presentations in the long-lag condition, yet overall memory performance was worse in that condition; and 2) the high confident recognition decisions that Wagner et al., analysed were more likely to be driven by recollection (though high levels of familiarity can produce high confidence). Nonetheless, the design did not allow a clear separation of memory for the item from memory for its context.

Turk-Browne et al. (2006) also examined interactions between priming and subsequent memory, this time using scene stimuli, with participants making an indoor/outdoor judgement on scenes presented twice during a single fMRI scanning session. The scan

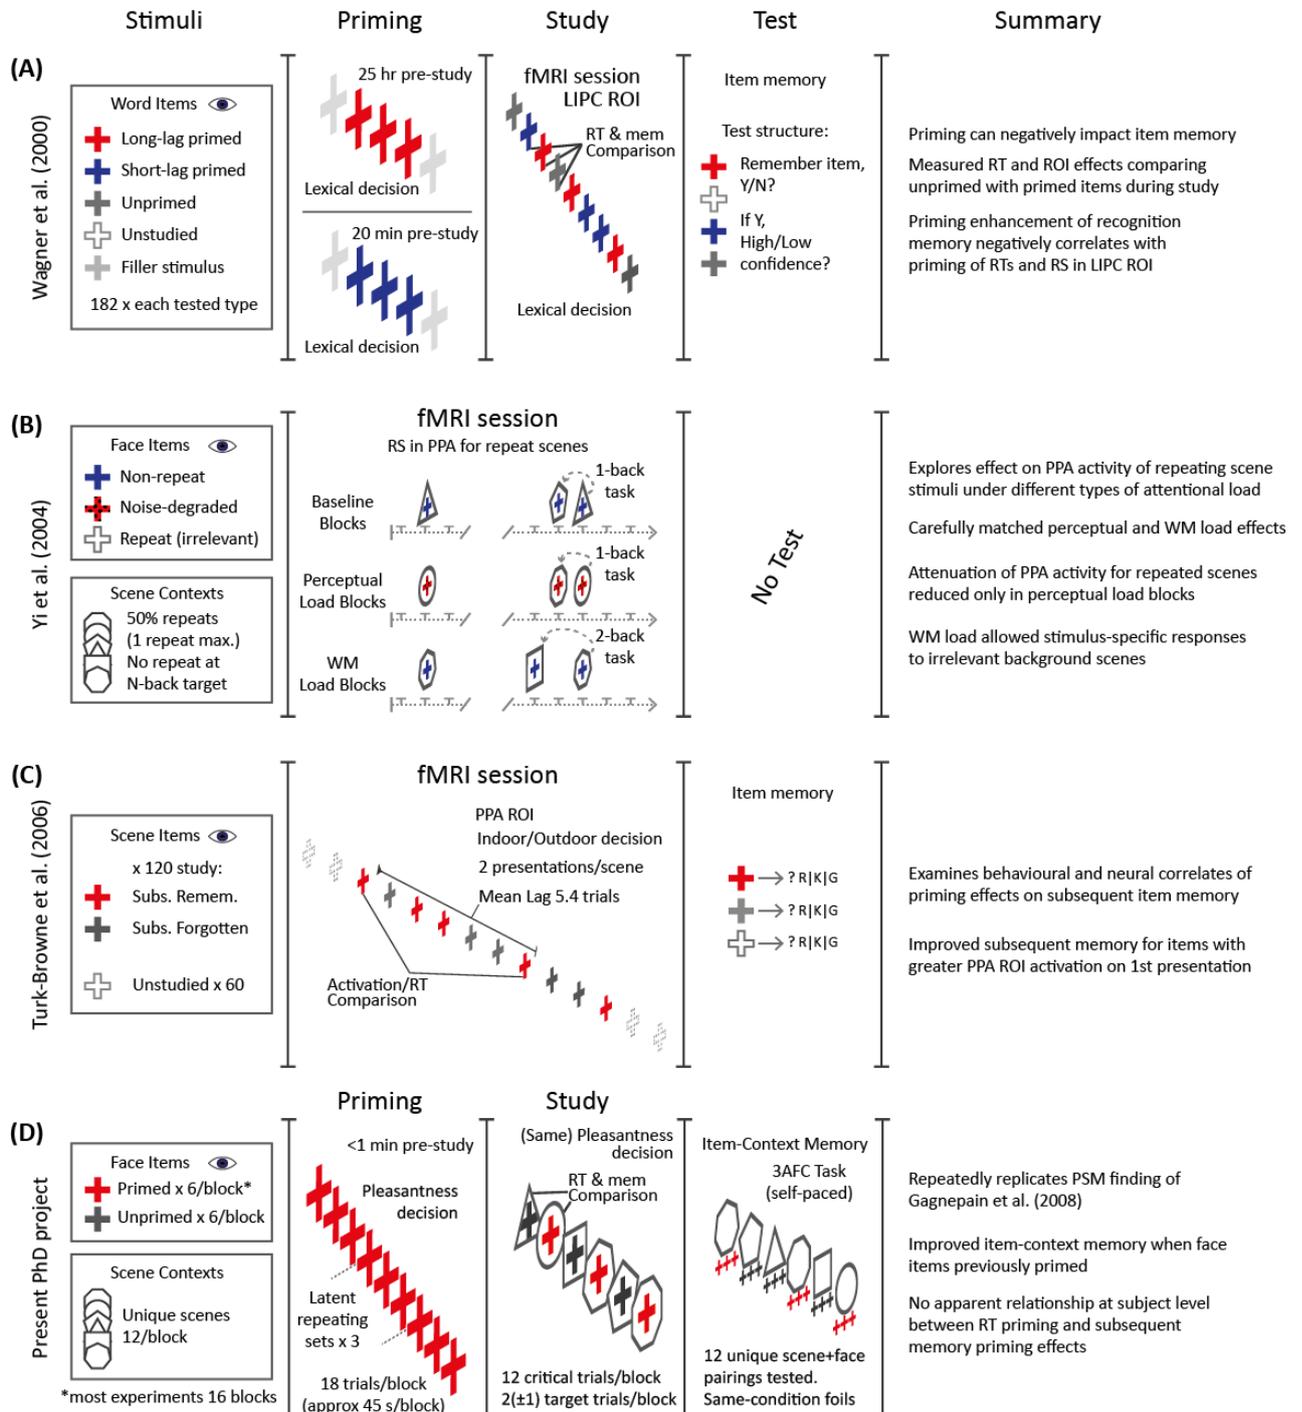


Figure 3: (chronological by row) detail 3 relevant studies (A-C), and (D) basic paradigm used for the present thesis. The studies concern either priming effects on subsequent item memory (A&C), the effect of attention on processing of irrelevant peripheral stimuli (B), or priming effects on subsequent associative memory (D).

was followed by a surprise memory test, and trials were divided into subsequently remembered and forgotten trials (see Panel C, Figure 3). Unlike Wagner et al. (2000), Turk-Browne et al. were able to analyse brain activity during the first priming encounter with each item (as well as second encounter), and relate this neural activity to subsequent memory. Their main aim was to examine the relationship between RT priming effects (difference between 1st and 2nd presentations) and neural priming effects (focusing on the parahippocampal place area, PPA), as a function of whether the scene was later remembered (unfortunately they did not directly correlate behavioural or neural priming with subsequent memory, as did Wagner et al.). Most relevant to the present discussion however, their data show that the greater behavioural priming and neural RS for remembered than forgotten scenes was driven by a greater response to the 1st presentation of a scene when it was remembered.

Although at first glance, the results of the Turk-Browne et al. (2006) and Wagner et al. (2000) studies appear contradictory, in that Wagner et al. generally found a negative relationship between priming and subsequent memory, whereas Turk-Browne et al. generally found a positive relationship. However, the discrepancy may be explained by the different ways in which the two studies measured priming. Wagner et al. defined priming as the difference between items presented previously and unprimed items not seen previously. As noted above, primed items had shorter RTs on average, so the reduced memory associated with greater priming could reflect reduced time-on-task (less processing/encoding time) for the 2nd presentation. Turk-Browne et al. defined priming as the difference between 1st and 2nd presentation of the same items. In this case, items with longer RTs on the 1st presentation will tend to show a larger priming effect. Given that Turk-Browne et al. also showed that neural activation for the 1st presentation was greater for remembered than forgotten items, the total processing time across both presentations could be greater for remembered than forgotten items (this could increase the contribution of familiarity to the recognition decision for example). Thus a time-on-task account would appear to be able to explain both findings.

Both Wagner et al. (2000) and Turk-Browne et al. (2006) discuss the effects of priming and subsequent memory in terms of selection of task-relevant features for processing, and in terms of attentional effects on the perceptual system's representation of the stimuli, proposing that a change in the perceptual response to the stimulus could itself have facilitatory effects for the decision processes, as well as making it less effortful to encode an item for subsequent recall. However, as noted above, these studies measured subsequent memory at the item level, but are unable to conclusively support claims about context memory. Evidence of context memory requires for example the use of R/K judgments or tests of 'source memory', as assessed by Gagnepain et al (2008; 2011); or more generally, according to the PIMMS hypothesis, any memory that requires the formation of new associations, such as that between an item and its context (or between two items). The experiments used in this thesis therefore measured associative memory (using a forced choice recognition task, in which the choices are equally familiar; see Chapter 2). These experiments also employed a three-phase design like those of Wagner et al. and Gagnepain et al. (2008, 2011), in which priming effects were measured as the difference between items seen previously and new (unprimed) items, thereby controlling for one-shot item-context encoding across both item types (Panel D, Figure 3).

Item-Context Interactions

The relationship between item memory and context memory was considered previously by Jurica and Shimamura (1999), who argued that they are negatively related. These authors showed that circumstances that improve item memory tend to impair context memory. Participants either generated answers to questions presented on a screen, or read answers to those questions. A face (which functioned as the context here) was also present on the screen. When participants generated their own answers, memory was enhanced relative to when they read the answers (the well-known "generation effect", Slamecka & Graf, 1978). However, memory for the face (context memory) was worse after generation than reading. The authors argued that this reflects a form of "item-context" trade-off: when item memory is enhanced by elaboration of the item, context memory is impaired. This can be explained by any type of resource account: if the total processing resource is fixed, then the more of that resource given to processing the item

(e.g., its perceptual or conceptual details), the less resource available to process other information around that item, i.e., context.

An fMRI study by Yi, Woodman, Widders, Marois, and Chun (2004, Panel B, Figure 3) investigated neural effects of this 'trade off' between processing an attended item and its unattended background context, by varying processing demands for centrally-presented, task-relevant, opaque face photos, overlaid on larger images of scenes, which were thus less foveal than the faces. In baseline-condition blocks, participants performed a simple repeat-detection (1-back) task on faces that appeared in the centre of a composite scene-face image. Unbeknown to participants, scene images were also periodically repeated. Attenuation of PPA activity for repeated compared with 1st presentation of scenes (RS) suggested that despite being irrelevant to the task, stimulus-specific processing did take place. Two experimental manipulations were then compared to this baseline. In 'perceptual load' blocks, the central face images were degraded with salt and pepper random noise to make them less distinguishable, while participants performed the same 1-back task on them. As evidence of the effectiveness of the manipulation, 1-back performance fell significantly. During these perceptual load blocks, RS in the PPA to repeated scenes was abolished. In a third condition, a 'working memory (WM)' load, rather than perceptual load, was introduced. In WM load blocks, the visual conditions were identical to the baseline condition, but participants performed a more taxing 2-back task on the face stimuli. As expected, performance was worse than for the 1-back version. However, in contrast to perceptual load blocks, PPA activity was again sensitive to scene repetition in the WM load blocks. Thus, where perceptual load has the effect of narrowing the focus of perceptual processes, producing an early-selection effect, attending to central representations does not interfere with stimulus-specific processing of task-irrelevant material, consistent with late-selection. Yi et al. propose that this shows primarily that task demands determine the level of neural processing that unattended stimuli will undergo: perceptual attentional processes may be distinguished from central attentional ones (also see Burgess, Dumontheil, & Gilbert, 2007), and that attention should not be thought of as a unitary phenomenon. The term 'working memory' is not conceptually distinct from that of

‘central attention/resources’, and throughout the present work, except where referring to others’ terms, such as in the present discussion of Yi et al., ‘central’ load/attention/resources will be the preferred terminology.

Returning to Jurica and Shimamura’s item-context trade-off theory, it is yet unknown whether the trade-off reflects limited perceptual resources (early selection) or limited central resources (late selection). This question is addressed by various manipulations of attentional load in the present thesis, and the Yi et al. (2004) data are also relevant to the fMRI study in Chapter 5.

Perceptual Attention and Working Memory

More generally, the idea of attentional “trade-offs”, and the study by Yi et al. (2004), are consistent with the load hypothesis of Lavie (1995), which proposes that for all kinds of goal-oriented task performance, top-down control is required so that attention does not ‘spill over’ into the processing of irrelevant stimuli. Such attentional control may broadly be said to originate from a fronto-parietal network (see e.g. Lavie, Beck, & Konstantinou, 2014; Seitz & Watanabe, 2009; Watanabe & Funahashi, 2014). Whereas Yi et al. showed neural consequences of perceptual vs central load, at a behavioural level the contrast between perceptual attentional processes and central attentional processes is clearly illustrated in two classic studies.

In the first, Lavie and Cox (1997) demonstrated the effects of perceptual load on the processing of distracting, task irrelevant stimuli. In a visual search task with a factorial design, participants looked for a target ‘X’ or an ‘N’ among a circular array of letters with similar orthographic features (‘H’, ‘M’, ‘W’, ‘K’, ‘Z’, high-load condition), or among an array consisting solely of ‘O’s (low-load condition). Adjacent to the search array appeared a task-irrelevant ‘flanker’, or distractor letter, either congruent or incongruent with the target letter (‘X’ or ‘N’ for ‘X’ search, ‘N’ or ‘X’, for ‘N’ search). A congruency effect was found in low-load trials, with faster RTs when the flanker was congruent with the search target, but under high load no such effect was seen. Thus, Lavie and Cox conclude that when the perceptual system does not have to ‘work hard’ for task performance, its capacity will ‘spill over’ into the processing of irrelevant stimuli.

De Fockert, (2001) then showed a reversal of this perceptual load effect: the congruency effect, which was abolished under high perceptual load, could instead be *enhanced* by an increase in central load. De Fockert hypothesised that central attention is necessary for the maintenance of task demands – i.e. direction of perceptual selection, such that increasing central demands leads to diminished control over perceptual processing. While in an MRI scanner, participants carried out 2 unrelated tasks, a central attentional task, and a simultaneous selective attention task. In the central attention task, which had high and low load conditions, participants were required to retain a 5-digit sequence of numbers in their head during several trials of a second, selective attention task. When probed with a digit after several trials of the selective attention task, participants were required to indicate the next number in the sequence. In the low load condition, the sequence was always “0 1 2 3 4”, and in the high load condition, the sequence could be any ordering of those numbers, changing on each presentation. In the concurrent attention task, the requirement was to indicate whether the name appearing as text onscreen was of a pop star or a politician, while faces behind the text were either congruent with the name, incongruent (i.e. highly distracting), or unfamiliar.

At the behavioural level, the results showed that the congruency effect (i.e., slowed RTs when faces and text were incongruent) was exaggerated for trials occurring during times of high versus low central load. De Fockert (2001) argued that directing attention toward appropriate (i.e. task-relevant) stimuli, requires the maintenance by central attention of ‘stimulus priorities’, such that placing additional demands on that central system would lead to a reduced ability of the perceptual system to select only high-priority stimuli, and thus central load increases distractor processing. At the neural level, de Fockert showed that, along with increased activation of a fronto-parietal network previously implicated in tasks where central demand varied, activation in regions of ventral temporal cortex associated with processing of faces was more widespread under high central load, congruent with the increased stimulus-relevant processing of scene stimuli found by Yi et al. (2004).

Resource Accounts of PSM

The purpose of reviewing the studies of perceptual and central load above, and of the item-context trade-off theory, is to introduce an alternative account of the original PSM effect reported by Gagnepain et al. (2008). This resource account is based on the idea of a limited resource that can be allocated to either processing an item or processing its context. When items are primed, they require less resource to process, thereby “freeing-up” resource for processing surrounding context, and hence improving subsequent associative memory (PSM). The main question concerns the nature of this resource.

The first, simplest possibility is the ‘temporal resource’ account. According to this account, when the response to an item is faster because of priming, there is more ‘leisure-time’ available before the onset of the next trial, during which the context can be processed – in other words, PSM arises purely from more time spent processing context information. This relates to the time-on-task hypothesis used to reconcile the effect of priming on item memory in the Wagner et al. (2000) and Turk-Browne et al. (2006) studies above (and would predict that these studies would have observed better context memory after priming, had that been measured).

The second possibility is the ‘perceptual resource’ account, in which priming reduces the perceptual load of Lavie (1995). In the Gagnepain et al study, for example, priming could have rendered the word stimuli more discriminable against the background sound context, resulting in (as in Yi et al.’s 2004 study) greater incidental processing of the background sound, and hence improving source memory.

The third possibility is the ‘central resource’ account, in which priming facilitates semantic processing of items, reducing demands for a central, generalised executive control process, which can then be recruited for better association of an item and its context, i.e., for encoding of new associations.

Project Overview

This thesis explores these different ‘attentional resource’ accounts, in an attempt to distinguish them from the PE-based account of PIMMS. Across four behavioural

experiments in Chapters 2 and 3, Gagnepain et al.'s (2008) basic PSM effect is replicated many times, but using stimuli in the visual rather than auditory domain. The main reason for changing the stimuli (in addition to generalising the effect) was to use faces and scenes as the items and contexts, given the evidence that these stimulus classes are processed in different brain regions, and therefore allow potential further insight by using fMRI (as in the Yi et al., 2004, study; separating the effects of priming in auditory cortex on auditory words and sounds would be more difficult). On initial consideration, these four behavioural experiments appeared to rule out the temporal resource and perceptual resource accounts, and provide only limited evidence for a possible central resource account. However, computational modelling in Chapter 4 demonstrated the situation to be more complicated, such that either a perceptual resource or central resource account could explain the results of the behavioural experiments. In brief, the failure to find an interaction between priming and concurrent load (perceptual or central) on subsequent memory does not rule out a resource account, when the relative amounts of resource are unknown, and nor does the failure to find a correlation between priming and subsequent memory, when the function relating resource to performance is unknown and can be nonlinear. As a response to the limitations of behavioural data demonstrated by the modelling work, Chapter 5 describes an fMRI experiment that attempted to more directly test the neural predictions of the resource versus PIMMS accounts of PSM. Some limited evidence is found that is consistent with both accounts. Chapter 6 then summarises the main findings and attempts to bring together the various theoretical accounts of this puzzling phenomenon.

Chapter 2
Testing Resource Accounts I:
Temporal and Perceptual Resources.

Introduction

The results of Gagnepain et al., (2008, 2011) showed a surprising effect of priming on subsequent memory for the association of word items and the background sounds with which they appeared. In this chapter and the next, PSM is shown in the visual domain using face ('item') and scene ('context') stimuli, and explored in relation to 3 possible resource-based accounts. These resource accounts of PSM assume that when more of a resource is needed to encode an item, less remains for encoding the spatiotemporal context in which that item occurred. Thus, any effect that priming has on associative memory is due to priming freeing up resources so that they can be used in a way that leads to better encoding of the item-context pairing.

The first experiment addresses a simple freeing of time: a 'temporal resource' account. In this account, the PSM seen in Gagnepain et al. (2008, 2011), is explained simply by the fact that primed items were responded-to faster in trials where the time between the start of one trial and the start of the next (stimulus onset asynchrony, SOA) is fixed. This would provide an explanation for the significant relationship between the Priming Effects on RT (PRT) and PSM: the more time is freed by priming, the more 'leisure time' will be available before the next trial onset to process the item-context pairing, leading to better encoding of the context of primed items than unprimed items. For instance, in the Study Phase, if participants have 1000 ms per trial to make a decision about a face (the primary task), and if initial perception of an unprimed face takes 500 ms, then 500 ms is left over to process the background scene and face-scene pairing. However, if prior exposure enables a face to be processed in only 400 ms, then there is clearly more time (600 ms) left over to encode the face-scene pairing. This is the account tested in Experiment 1. In this account, resource-allocation is a simple, serial process.

In contrast to the above focus on the serial allocation of temporal resources, a more sophisticated account refers to *attentional* resource allocation, whether serial or parallel. According to this account, the advantage for primed item-context pairs arises due to reduced processing demands for primed items freeing up attentional resources which serve item-context binding. According to the experimental results of Yi et al., (2004) as discussed in the Introduction and in line with the 'item-context' trade-off of

Jurica and Shimamura (1999), reduced perceptual load for the processing of primed items should lead to greater processing of the background scene image, which may result in improved memory for the association of face and scene for primed trials.

The precise paradigm we used is shown in Figure 4. Faces were first primed in an initial Training Phase. Following this, in the Study Phase, scenes were presented for a short period after which a superimposed face (either primed or unprimed), was presented in the centre of the scene. After a short period of distraction, face-scene recollection was then measured in a final Test Phase, during which participants performed a 3-alternative forced-choice test (3AFC) in which they had to select the one of three faces that had been paired with a given scene. The critical question was whether people's memory for the pairing of a face and scene (as measured in the Test phase) was better when the face that was paired with that scene (in the prior Study Phase) had also previously been presented (primed) in the Training Phase. It is important to note that the two foils in the final 3AFC Test Phase were other faces from the Study phase, and from the same condition (i.e. primed or unprimed) as the target face, so that correct performance could not be achieved simply by item memory (i.e., relative familiarity of each face).

To minimise the influence of intentional memorisation strategies, we used an incidental task during Training and Study phases, in which participants simply decided as quickly as possible whether each face was pleasant or unpleasant (see Methods for more details). The RTs to make this decision provided an independent measure of perceptual priming (i.e., priming effects on RT, PRT). To ensure that participants also paid some attention to the scenes (which could otherwise be ignored), we added a secondary task in the Study Phase in which participants also had to press a different key to a pre-specified target scene (any scene that contained the moon). These targets occurred infrequently (and without a face) and played no role in the main analyses. Note that piloting showed that only 12 unique scene-face pairings (6 primed and 6 unprimed) were possible in order for memory performance to be sufficiently above chance (i.e. off floor, but below ceiling). Therefore, in order to obtain enough trials in total, the Priming-Study-Test cycle was repeated across 16 blocks (with unique stimuli in each block).

The first experiment explored the temporal resources account, comparing performance in blocks where the ‘leisure time’ after a response was determined by the response latency, with trials in which such leisure time was fixed. If freeing of temporal resources is a sufficient explanation for the priming effect found by Gagnepain *et al.*, then such effects should be abolished in the Fixed Leisure condition, i.e. when faster responses for primed items do not free-up extra leisure time before the next trial.

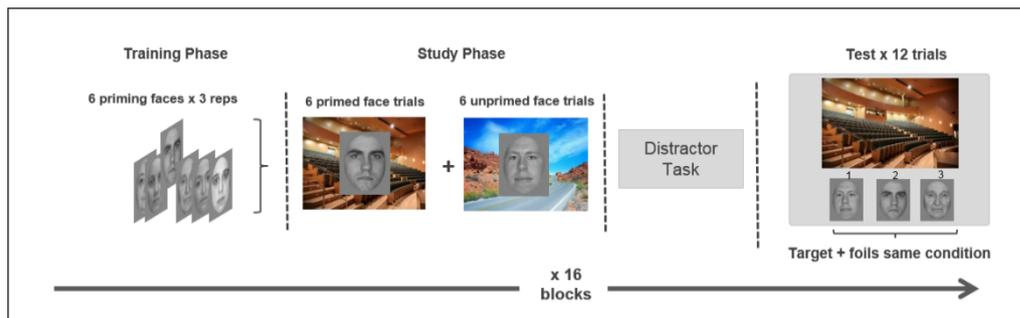


Figure 4: The three main phases of Experiments 1-3. 12 scene-face pairs were shown per block, in 16 blocks. In both Priming and Study phases, the primary task was to judge pleasantness of the faces. During Study Phase an occasional target task was to spot scenes containing the moon (not shown). Distractor task (10 s odd/even number classifying). Test was 3AFC, self-paced. Blocks ~ 4 mins/block. The main manipulations described here involve conditions during study phase, in alternating blocks.

Experiment 1: Temporal Resource Account

To test the “temporal” resource account, we created two conditions that differed in the timing of the Study trials, namely whether the trials had a fixed SOA (regardless of the speed of the participant’s response) or a variable SOA that depended on the response (akin to a self-paced design). Thus, for one half of the Priming-Study-Test blocks, trial length was fixed, meaning that an earlier response would be followed by a longer period before the following trial would begin (“Fixed-SOA” condition). In the other half of the blocks, this ‘leisure period’ after response was fixed, i.e. SOA was variable (“Fixed-Leisure” condition). According to the temporal resources account, the amount of leisure time after a response is the critical determinant of face-scene encoding, and no priming effect on associative memory should be found in the Fixed-Leisure condition. Furthermore, if PSM is a function of leisure time, then PRT should predict PSM in the Fixed-SOA condition: participants who gained greater RT advantage for primed trials will have had greater leisure time for the encoding of the scene-face association, and therefore greater subsequent memory advantage for primed trials.

Method

Participants and Materials

16 participants (6 male) were recruited aged 18-35 ($M = 25$, $SD = 4.1$), and paid £6 for their time, according to ethics protocol CPREC 2005.08. Two participants were excluded because they did not perform significantly above chance at test. All were right handed and had normal or corrected-to-normal vision.

A set of 192 colour 1024 x 768 resolution images of indoor and outdoor scenes (without any obvious foreground objects), plus a further set of 40 scenes (same resolution) in which the moon featured clearly, were created from the set obtained from the internet by Greve et al., 2017. A set of 192 grayscale 255 x 321 resolution face stimuli was created by taking neutral expression face images from the Stirling Database (pics.stir.ac.uk), in which the model was face-on and gazed directly at the camera. An oval mask was then applied to minimise differences in hairstyle. Eye height and centrality within the oval mask was kept constant between images. Images which included facial piercings, heavy make-up or facial hair were excluded.

Face images were divided into 4 sets for counterbalancing, taking care to balance the sets so that subjectively similar faces were evenly distributed, i.e. with roughly equal numbers of male and female, more and less aged, and plumper and thinner faces. Each face could then appear in one of four conditions in the 2x2 study design: Fixed-SOA Primed, Fixed-SOA Unprimed, Fixed-Leisure Primed, or Fixed-Leisure Unprimed. The 4 sets were then allocated to conditions across participants via a Latin square.

Procedure

The experiment consisted of 4 phases, repeated across 16 blocks. Participants sat at a computer and used a keyboard to respond. After a brief summary of the tasks given by the experimenter, instructions were presented onscreen. A short (approx. 3 minute) practice session consisting of shortened versions of each task was completed prior to the beginning of the experiment proper.

Task 1: Training Phase

The Training Phase consisted of 18 evenly-spaced trials in which 6 faces were presented three times in a pseudorandom order (with all faces being presented, in a random order, before being repeated). A trial began with a white fixation cross in the centre of the screen for 500ms, after which a face image appeared, with the fixation cross remaining visible over the face, always sitting approximately halfway down the nose. Participants indicated with a speeded response whether they found the face to be a 'pleasant' or 'less pleasant' (than some subjective average) face. When a response was made, feedback was given by the white crosshair changing to black. Each face image was presented for 1750ms, after which the next trial began, giving an SOA of 2250ms. After 18 trials which took ~45s, the software presented a brief reminder of the instructions for the next task, with participants indicating with a key press when they were ready to continue.

Task 2: Study Phase

The Study Phase consisted of 12 study trials plus 2 (+/- 1) target trials interspersed randomly. Study stimuli in each trial were pairings of scenes with either primed or unprimed faces. Trial timing parameters alternated in each block, with blocks of either Fixed SOA or Fixed Leisure (see Figure 5, below), counterbalanced with respect to whether participants started with a Fixed SOA or Fixed Leisure block. Trial sequences were pseudo-randomly generated such that neither primed nor unprimed trials occurred for more than 3 successive trials. Each study trial began with a scene presented in the centre of the screen for 1000 ms. Target trials were scenes that included a moon, to which participants responded with a speeded key-press (spacebar). All other non-target scenes were followed by a face image, overlaid centrally on the scene image, with a white fixation cross appearing in the centre of the face image, as in the Training Phase. After a further 400ms, the images disappeared from the screen, leaving only the fixation cross. As in the Training Phase, participants indicated their speeded judgement about whether the face was 'more pleasant' or 'less pleasant', with feedback given in the form of the fixation cross turning to black. The response window was either fixed at 800 ms (Fixed-SOA blocks) or else the trial ended 485 ms after response (Fixed-Leisure blocks),

immediately after which the next trial began. At the end of the task, reminder instructions were presented for the next task.

Task 3: Distraction Phase

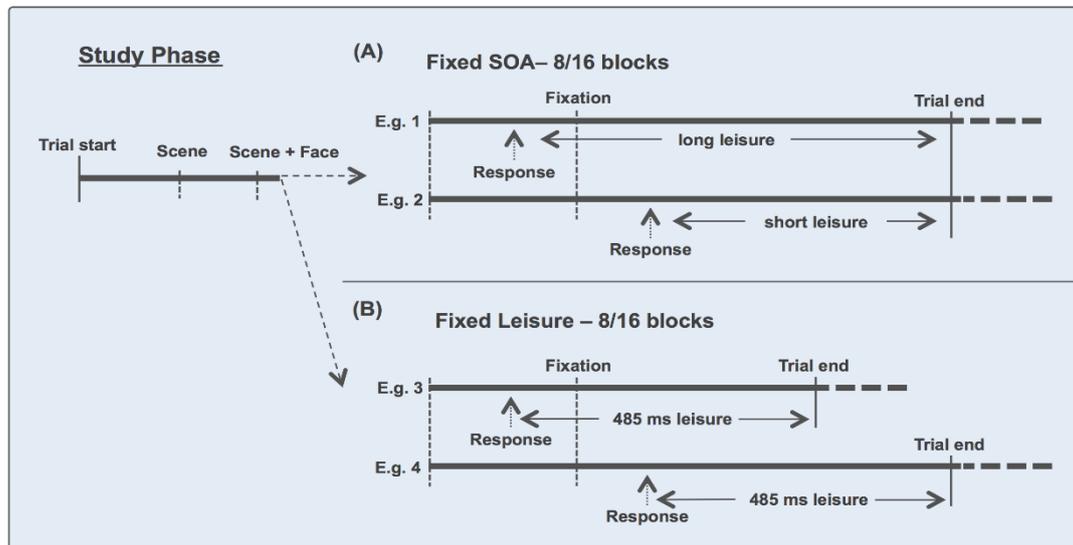


Figure 5: Study Phase Structure in alternating blocks. Panel A shows how Fixed SOA trials produce leisure time variation depending on response speed, whereas in Fixed Leisure blocks (B) this leisure time is constant regardless of response speed.

In order to prevent contributions from working memory, and to minimise recency effects, whereby the later-presented face-scene pairings would be better remembered, a short distractor task followed the Study Phase. This consisted of five trials of an odd/even number-categorisation task with a fixed SOA of 2000ms. At the start of each trial, a white fixation cross appeared in the centre of the screen, replaced after 250 ms with a randomly selected number between 1 and 100. When the participant responded, feedback was given via replacement of the number with either a green (correct response) or red (incorrect) fixation cross, which remained onscreen until the end of the fixed trial time. At the end of the task, reminder instructions were presented for the next task.

Task 4: Test Phase

The Test Phase consisted of a 3-alternative-forced-choice (3AFC) task, where all pairings encountered in the study phase were tested, yielding 12 trials. Trials ended when a response was given. In each trial, a scene was presented from the study phase, together

with 3 face images presented at 75% scale and arrayed below the scene image, numbered 1-3. The position of the target and the 2 foils was pseudo-randomised such that the target could not appear in the same position for more than 3 consecutive trials, and appeared in each position an equal number of times. Target and foils were always chosen from the same condition (i.e. primed/unprimed) and were always of the same sex. At the end of this task, participants were informed that the block had ended and given reminder instructions for Task 1 of the following block, or informed at the end that the experiment had ended.

Statistical Analysis

Given the predicted monotonic decrease in RTs with repetition in the Priming Phase, priming was tested by a one-tailed linear trend analysis across the three presentations. RTs from the Study phase were analysed in a two-way ANOVA, with factors of Priming and Study-SOA, as were the memory accuracy scores from the Test phase. In both Training and Study phases, RTs at the participant level which were outside 2SDs from the median RT were defined as outliers, and these trials excluded. In cases of an outlying Study trial, its corresponding Test trial was also excluded from analysis.

A positive correlation was predicted between PRT and PSM in the Fixed-SOA condition, given the hypothesised the (priming) increase in leisure time depends directly on (priming) reduction of RTs.

Results

The linear trend analysis showed a significant decrease of RTs with repetition during the Priming phase, as expected, $T(31) = 2.07$, $p = .023$, one-tailed (Figure 6A).

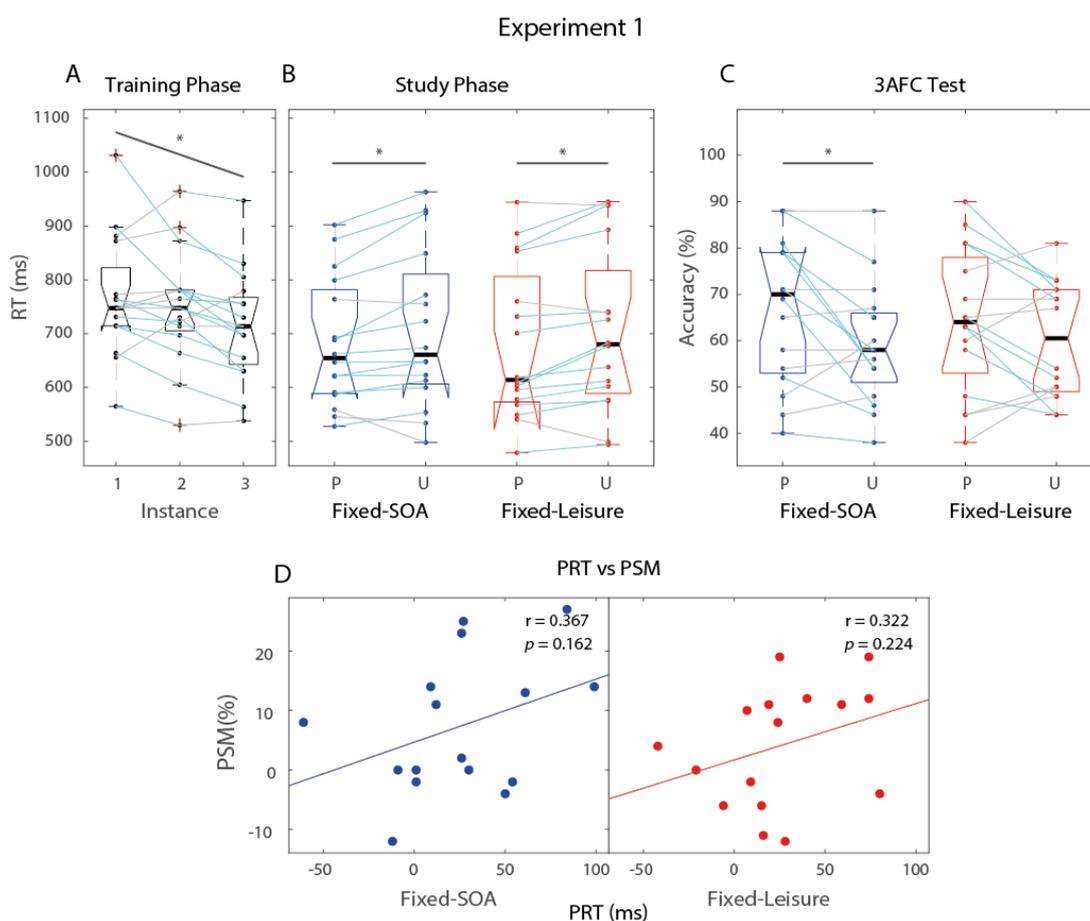


Figure 6: Panels A-C show mean scores ($N = 16$) during 3 phases of Experiment 1. Effect of repetition during Priming (A) and Study (B) on RTs for face pleasantness judgments. Panel C shows that the effect of priming on memory accuracy during final 3AFC Test phase occurs regardless of leisure time. Asterisks show significant linear trend (A) or pairwise differences (B+C) at $p < .05$ one-tailed. Panel D shows no significant relationship was found between PRT and PSM in either condition.

The two-factor ANOVA on Study phase RTs showed a main effect of priming, $F(1,15) = 12.8$, $p = .003$. As expected, responses to primed faces were faster than to unprimed faces (Figure 6B). The main effect of Study-SOA was not significant, nor was any interaction between Study-SOA and priming, $F_s < 1$.

At the Test phase, the 2x2 ANOVA revealed a main effect on memory of priming the face images, $F(1,15) = 5.68$, $p = .031$. Fixing the leisure time after a response did not produce a significant change to the priming effect, with no main effect of leisure time, $F(1,15) < 1$, nor any interaction, $F(1,15) = 1.71$, $p = .21$ (Figure 6C). Simple effects showed that priming significantly improved memory in Fixed SOA ($M = 7.31\%$, $SE = 2.85$), $T(15) = 2.57$, $p = .011$, one-tailed), and produced a trend effect in the Fixed Leisure ($M = 4.1\%$, $SE = 2.5\%$), $T(15) = 1.61$, $p = .064$, one-tailed) conditions.

Discussion

The experiment had two main aims: firstly to replicate in a visual modality the findings of Gagnepain et al. (2008) that primed stimuli are better associated with a particular context than are unprimed stimuli, and second to check that this priming effect remained when the time between the response and the next trial was equated (Fixed-Leisure condition). Both of these aims were confirmed. This suggests that the effect of priming is not simply to “free-up” more time to encode the face-scene pairing. Note that, although the interaction between priming and leisure condition was not significant ($F < 1$), there was a numerical pattern of a greater effect of priming on memory in the Fixed-SOA than Fixed-Leisure conditions. This may be significant if more participants were tested. Nonetheless, we do not see this null result as a problem for our experiment, because our main question was whether priming was significant in the Fixed-Leisure condition, for which we obtained a significant positive result. Thus, the main claim from this experiment is not that time-on-task does not affect memory performance, but rather that it is not sufficient to explain the effect of priming on associative memory originally found by Gagnepain et al., and now replicated here. In further support of this claim, the lack of any observed correlation between PRT and PSM in the fixed-SOA condition suggests that priming advantage in this condition was independent of time saved by priming facilitation of RT.

However, while this experiment rules out one of the simplest types of “resource” account – time-on-task – there are other definitions of resource that might apply. For example, primed faces might require fewer attentional resources in the Study phase, even while the face is being processed, regardless of the amount of leisure time following the face

(i.e. a “parallel” processing account, rather than the “serial” time-on-task account). The problem with such attentional accounts is to operationalise the specific type of processing resource. In Experiment 2, we tested an alternative, perceptual resource account.

Experiment 2: Perceptual Resources

A second resource-based explanation for the recollection advantage for previously encountered (primed) items compared with unprimed ones found in Gagnepain et al. (2008) and Experiment 1 (above) could be that, in trials when an item has been processed previously, perceptual resources are freed for the task of binding items with their contexts. Experiment 2 examines this explanation by comparing trials in which perceptual load is low (therefore more perceptual attentional resources are available) with trials in which perceptual load is high (perceptual resources therefore less available).

As discussed in Chapter 1, manipulation of perceptual load has been widely shown to affect processing of surrounding stimuli (see Lavie, 2005 for review). According to Burgess et al., (2007), stressing perceptual attention requires: i) the immediate availability of information to be processed, ii) that the processing requires target perceptual features of the present stimulus, and iii) that the responses and rules governing them are relatively well-learned. Fulfilling these criteria, for the current experiment, we increased perceptual load for half of the face items by adding pixel noise to the images, as was also used successfully by the Yi et al. (2004) study considered in Chapter 1. We used the Fixed-Leisure procedure from Experiment 1, to ensure that time-on-task remained controlled, but added a new manipulation of perceptual load by degrading half of the face images with pixel noise when they appeared during Study Phase. Note that all faces were presented clearly during Training Phase exposure, in order to maximise priming effects. If degraded faces require greater attentional resources to process (i.e. impose a higher perceptual load), then the effect of priming on memory should be reduced relative to that for clear faces (i.e. there should be an interaction between priming and degradation).

To show that our perceptual load manipulation affected face processing, it is important to demonstrate a main effect of degradation on memory (even if this does not interact with priming). However, if we only used degraded faces during the Study phase (and presented only clear faces at Test), then any overall detrimental effect of perceptual degradation at study could simply reflect a reduced perceptual overlap between Study and Test stimuli. In other words, degrading faces at Study would induce a Study-Test mismatch in the face images (and in the extreme case, participants might not recognise the same face at Test when the noise is removed). Therefore, we added a second (within-subjects) factorial manipulation of perceptual degradation at Test, as well as Study. If Study-Test match is an important determinant of memory, then there would be a main effect of matching versus mismatching conditions. If perceptual load is an important determinant of memory, then there should be an interaction between priming and image-clarity, regardless of study-test match.

Method

32 participants (10 male), age 18-35 ($M = 25$, $SD = 4.6$) years were recruited. Three participants were excluded because they did not perform significantly above chance during test, as shown by a permutation test. Due to the addition of another factor (degradation of images during Test Phase), and the comparatively weak effect seen in Experiment 1, the sample size was doubled to maintain the sensitivity of the experiment. The degradation of faces during the Study phase alternated between 'Clear' blocks and 'Degraded' blocks. Degraded face images were produced using a MATLAB function which took a random set comprising a given proportion (in this case .57, based on piloting) of the pixels in the image and replaced them with grey pixels (see Figure 7 for examples of images). Degradation was judged enough to make the face harder to identify, but not to the extent that the degraded version could not be identified as the same person as the clear version. To improve overall Test Performance, face images remained onscreen during Study Phase trials for 800ms rather than the 400ms of Experiment 1. In all other regards, the procedure for Experiment 2 remained identical

to that of Experiment 1, and statistical analyses likewise were conducted in the same fashion.



Figure 7: Example of clear face image (left) as used throughout present study, and a degraded version (right) of the same face, as used in Experiment 2.

Results

The omnibus 2 (test degradation, within-participants) x 2 (study degradation, within-participants) x 2 (priming, within-participants) ANOVA on Test phase performance showed no significant main effect of test degradation, $F < 1$, and no evidence of any effect of whether study and test degradation matched, i.e., no evidence of a two-way interaction between study and test degradation, $F < 1$, nor of a three-way interaction between study, test and priming, $F(1,31) = 3.2$, $p = .082$. We therefore collapsed across the Test degradation factor in Figure 8 and the subsequent analyses below.

Linear trend analysis showed significant decrease of RTs with repetition during the Training Phase, as expected, $T(63) = 6.24$, $p < .001$ (Figure 8A).

A two-factor ANOVA on Study Phase RTs showed a main effect of priming, $F(1,31) = 37.3$, $p < .001$. As expected, responses to primed faces were faster than to unprimed faces (Figure 8B). The ANOVA also showed a main effect of image quality, with clearly presented faces responded to more quickly than degraded ones, $F(1,31) = 8.46$, $p = .007$, as also expected. However, there was also a significant interaction between degradation

and priming, $F(1,31) = 8.97$, $p = .005$. Further analysis revealed that the priming effect for clear faces ($M = 56$ ms, $SE = 7.6$ ms) was greater than that for degraded faces ($M = 31$ ms, $SE = 8.9$ ms), $T(31) = 2.99$, $p = .005$.

A $2 \times 2 \times 2$ ANOVA revealed a main effect on test performance of priming the face images, $F(1,31) = 21.8$, $p < .001$. Degradation during Study Phase also produced a significant effect, with degraded faces remembered less well than clearly presented faces, $F(1,31) = 29.3$, $p < .001$. However, the interaction did not reach significance, $F(1,31) < 1$, and one-tailed simple effects showed that priming significantly improved memory in both Clear ($M = 7.25\%$, $SE = 2.09$), $T(31) = 3.47$, $p = .002$, and Degraded ($M = 6.62$, $SE = 1.72$), $T(31) = 3.85$,

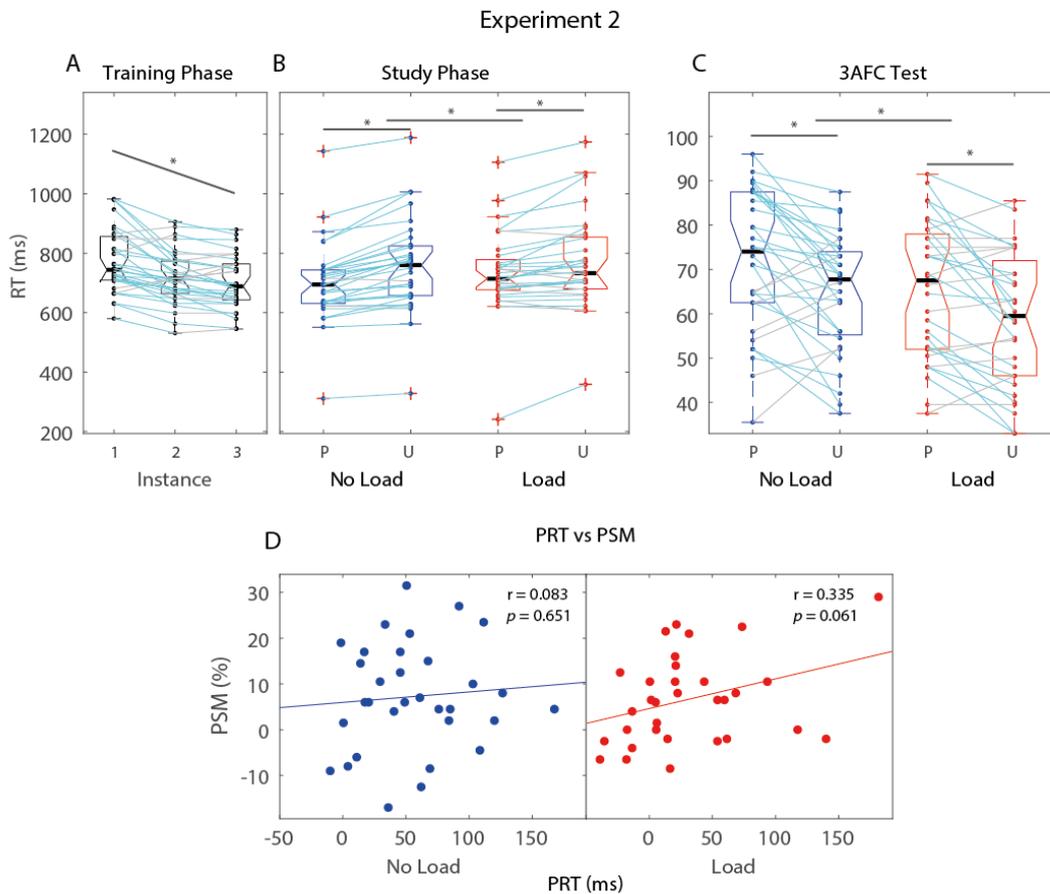


Figure 8: Panels A-C show mean scores ($N = 32$) during 3 phases of Experiment 2. Priming of face pleasantness RTs was found during Training (A) and Study (B), in addition to a main effect (longer bar) of degradation during Study. Panel C shows effect of priming on memory accuracy during final 3AFC Test phase and effect (longer bar) of degrading face images during Study Phase, but no interaction. Panel D shows relationship between priming effects at study (PRT) and test (PSM) with a trend positive correlation for the Perceptual Load condition.

$p < .001$, conditions (Figure 8C). No effect was found of perceptual match between Study and Test Phase, $F(1,31) < 1$.

Discussion

The aim of Experiment 2 was to test the idea that the priming effect seen in Experiment 1 could be explained by a freeing up of perceptual attentional resources for the primed faces during the Study Phase. This possibility was tested by factorially manipulating the perceptual degradation of the face images, at study and/or test. There was no effect of whether the degradation level of a face image was the same at study and test, suggesting that there was no concern about test performance being influenced by the matching of study-test format, so analyses were collapsed across the test degradation factor. Though degradation at study did reduce test performance overall (which is important to show that the manipulation was strong enough to have some effect), it did not abolish PSM (which remained significant even for degraded faces); nor was there any evidence that it reduced the magnitude of PSM (i.e., any interaction between degradation and PSM did not reach significance). Thus Experiment 2 provided no support for a perceptual attention resource account.

Interestingly, as in Experiment 1, there was no significant positive correlation between PRT and PSM (though there was a trend in the high perceptual load condition). While this again questions the temporal resource account, it is also puzzling according to the PIMMS account: if a greater speed-up of RTs comes from greater sharpening of the likelihood, then PE and hence PSM should also increase with PRT. One possibility is that PRT does not reflect purely perceptual speed-up/sharpening; it could be dominated for example by conceptual processing of the faces, or more likely by retrieval of S-R bindings from the priming phase (Henson, Eckstein, Waszak, Frings, & Horner, 2014). In any case, the question of PRT-PSM correlations is again addressed in Chapter 4.

Although significant priming was found even for high perceptual load, it is possible that priming would be abolished if the perceptual load had been higher still, given that maximum load was unknown. (Note this is unlike Experiment 1, where the Fixed-Leisure

condition does completely match the time-on-task, so the presence of a significant priming effect in this condition is sufficient to reject the time-on-task account). Thus, the question remains of whether the lack of significant interaction between priming and perceptual load is a type II error.

Even if there is no true interaction in the data, it is possible that an attentional resource account would not predict such an interaction for the particular loads we used. This is because it is possible that all four conditions do differ in attentional resources, but the degree of resource is ordered in a linear fashion (greatest for degraded, unprimed faces and least for clear, primed faces), such that there would be main effects of both degradation and priming, but not an interaction (i.e. the pattern found here). While theoretically possible, this linear ordering cannot exist for all levels of resource, because the definition of resource implies that at some point the total resource is exceeded, at which point an interaction would be expected (such that priming is abolished completely at high enough perceptual load). Indeed, in general one might expect a nonlinear effect of perceptual load as the resource limit approaches (asymptotes). These issues are considered formally in Chapter 4. Moreover, the issue of where exactly conditions lie on a resource continuum applies to all such manipulations and is difficult to refute. So, while keeping this caveat in mind, it is concluded that there is currently no empirical support for a perceptual load account of the effect of priming on subsequent associative memory.

Chapter Summary

Experiments 1 and 2 found no support for either a temporal or a perceptual resources account. The temporal resource account is eliminated by simply controlling for 'leisure time' so that priming does not affect this. As discussed here and at greater length in the context of the modelling work in Chapter 4, the absence of an interaction pattern should not be seen as evidence that perceptual resources are not freed by priming processes. Nevertheless, it remains possible that priming frees up central, rather than perceptual resources, and this is addressed in Experiments 3 and 4 of Chapter 3.

Chapter 3

Testing Attentional Accounts II: Central Attention.

Introduction

Experiments 1 and 2 provided no evidence that the effect of priming on subsequent memory (PSM) is affected by temporal resources or by perceptual attentional resources respectively. A further possibility is that the attentional resources that are freed by priming are “central”, rather than “perceptual”, resources. Processes engaging central attentional resources include maintenance, refreshment, rehearsal, and manipulation of offline, internal representations (Baddeley, 2003; Chun & Johnson, 2011). Experiment 3 addressed this possibility by adding, in alternate blocks, a secondary task during Study Phase that requires central attentional resources for successful performance, reducing their availability for face-scene processing, with PSM being compared between low and high attentional load item-context study phases. In order to place load on central attention without visually altering the Study Phase, an additional auditory task was introduced in High Load blocks, in which participants were required to maintain a 2-digit number, and to update it each time a tone was heard. This satisfies the criteria of Burgess et al., (2007) for central, stimulus-independent processing, in that the information being attended is not present in the environment, coming instead from a previously-witnessed episode, or being self-generated, with the responses made referring to these internal representations. In baseline (Low Load) blocks, participants were asked to ignore tones. If priming frees central attentional resources, then adding a task for which such resources are necessary should reduce their availability and therefore reduce the associative memory advantage for primed items.

Method

In order to maintain power to detect what was hypothesised to be a weak effect, the larger sample size from Experiment 2 was kept. Hence 32 participants (12 males) aged 18-35 ($M = 23$, $SD = 3.2$) years were recruited. 1 participant was replaced due to extremely slow Study Phase responses, while 3 did not perform significantly above chance overall in the Test Phase, as shown by a permutation test. 1 further participant was replaced due to an extremely outlying pattern of Test performance. As with Experiment 2, the design maintained the same overall structure as Experiment 1, with all blocks employing Fixed-Leisure trial time during the Study phase. Instead of blockwise changes in face-

image quality, Study phases alternated blockwise between ‘Attend Tones’ and ‘Ignore Tones’. = 500 ms, σ

During Study phases, participants wore headphones through which they heard occasional tones. Tones were timed to appear with a probability of .6 during the period immediately before onset of a scene stimulus, with their precise occurrence during this 500 ms period sampled from a normal distribution with $\mu = 250$ ms, $\sigma = 125$ ms. Prior to the beginning of the Study phase, instructions presented onscreen indicated whether tones should be attended to (High Load blocks), or ignored (Low Load blocks). High Load block instructions showed a ‘starting number’, between 10 and 87, and participants were instructed to begin maintaining this number in their head while they carried out the rating task, adding 1 to the number each time they heard a tone. At the end of the task, participants were prompted onscreen to report the final tally, which they input via the keyboard. During ‘Ignore’ blocks, participants were instructed to ignore the tones.

Results

The linear trend analysis showed a significant decrease of RTs with repetition during the Training phase, as expected, $T(63) = 6.21$, $p < .001$ (Figure 9A).

The two-factor ANOVA on Study phase RTs showed a main effect of priming, $F(1,31) = 35.8$, $p < .001$. As expected, responses to primed faces were faster than to unprimed faces (Figure 9B). Interestingly, the presence of a secondary task did not seem to affect Study task performance, with no significant main effect ($F < 1$) nor interaction, $F(1,31) = 2.13$, $p = .155$. Secondary task performance was close to ceiling, with a median 7 correct out of 8 responses in total.

At the Test phase, the 2x2 ANOVA revealed a main effect of priming on subsequent memory, $F(1,31) = 12.4$, $p = .001$, consistent with Experiments 1-2. A main effect of secondary task showed that attending to tones in the Study phase resulted in worse test performance, $F(1,31) = 49.0$, $p < .001$, as expected (Figure 9C). However, though the PSM was numerically smaller in the High Load ($M = 2.8\%$, $SE = 1.4\%$), than No Load ($M = 4.8\%$, $SE = 1.3\%$) condition, the interaction did not reach significance, $F(1,31) = 1.33$, $p =$

.258. Analysis of the simple effects showed that priming improved memory significantly in both the Ignore Tones condition, $t(31) = 3.58$, $p < .001$, one-tailed, and the Attend Tones condition $t(31) = 1.93$, $p = .031$, one-tailed).

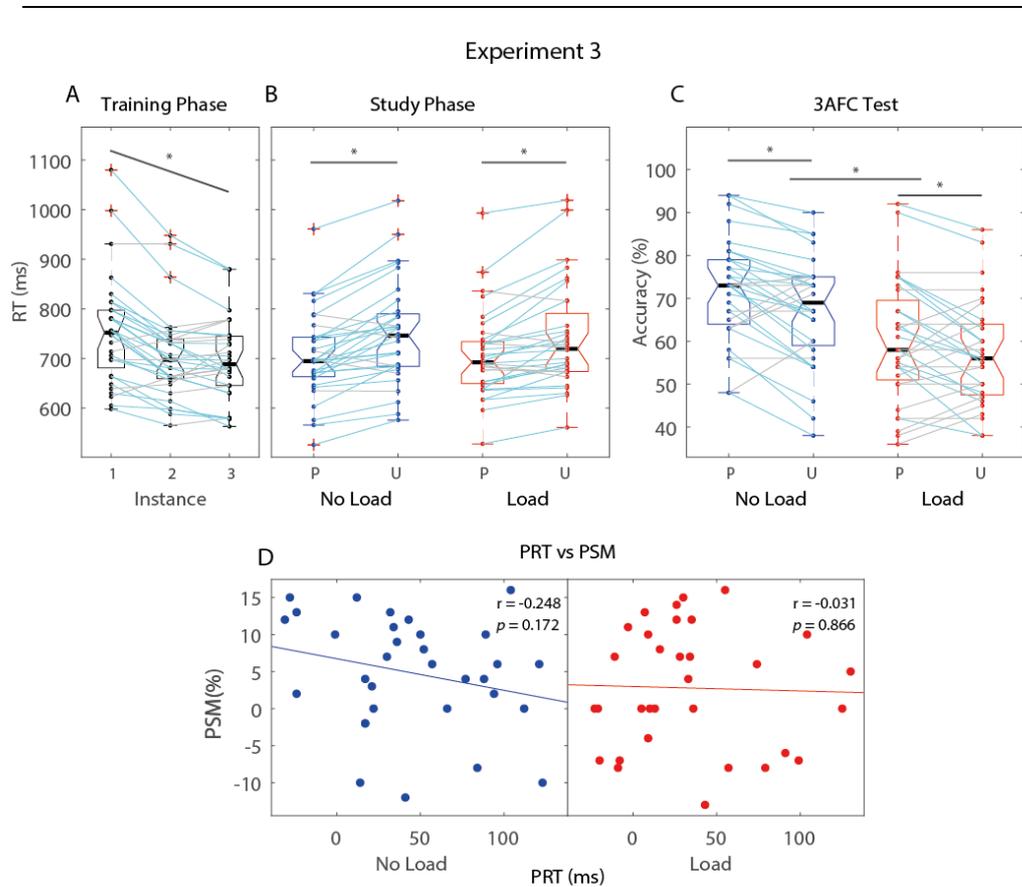


Figure 9: Panels A-C show mean scores ($N = 32$) during 3 phases. Priming of face pleasantness RTs was found during Training (A) and Study (B), but not affected by secondary task during Study. Panel C shows effect of priming on memory accuracy during final 3AFC Test phase and main effect of secondary task during Study Phase, but no interaction. See Figure 6 legend for more details. Panel D shows no observed relation between priming effects on RT and subsequent memory.

Discussion

The results of Experiment 3, which varied, in alternating blocks, the load on central attentional resources, were somewhat unclear. We predicted that the PSM would be smaller under high load, and this was true numerically, but the interaction between PSM and Load was not significant. Moreover, PSM was not only significant in the Low Load (ignore tones) condition, but also in the High Load (attend tones) condition, albeit

smaller in magnitude. One possibility is that the difference in central attentional load between the two conditions was not large enough to detect an interaction.

For example, the tones in the Low Load condition may still have made some demand on central resources, even though they were supposed to be ignored. This is because the Low Load condition alternated with the High Load condition, and used the same tones, so the cues (tones) that were irrelevant in one block had nonetheless been relevant in the previous block. This may have made them difficult to ignore completely, e.g., participants may have momentarily performed the High Load task during Low Load blocks.

Another possible reason that the difference in central attentional load between the two conditions was not large enough may be because the High Load condition was not sufficiently resource-demanding. Central demands could be increased by, for example, presenting tones with different frequencies and making the amount to increment the running total depend on the pitch of the tone. Given the mean Test accuracy scores in High Load blocks (primed $M = 58.2\%$, unprimed $M = 55.9\%$) were well above chance (33%), it seems that such a more extreme manipulation of central load could be tolerated.

A final problem with Experiment 3 is that the measure of secondary task performance (tone counting) was too coarse to be sensitive to effects of priming, given that the tone total is only given at the end of a block, such that secondary task performance cannot be split by primed vs. unprimed trials. This is important because it is possible that the effect of priming may be seen on performance of the secondary task rather than priming task (depending which task the participant prioritizes). In other words, participants may have a fixed amount of central resources, but decide to always allocate enough resources to the primary (face) task to ensure a certain performance level, such that it is the amount of resources left for the secondary task performance that is affected by the resources freed up by priming.

Experiment 4 addressed all of the above concerns by using: 1) no tones in the lower load task (i.e. now a No Load condition), 2) a more difficult task on the tones in the High Load condition, and 3) a high load secondary task that provides a more continuous, trial-by-trial measure of performance.

Experiment 4(a)

The purpose of Experiments 4a and 4b were to provide more compelling evidence for an interaction between central load and PSM, by using more extreme levels of load than used in Experiment 3.

One obvious way to increase load in the High Load condition would be to increase the number of tones presented each trial, and require an answer at the end of each trial (so that the effect of priming can be measured on the secondary task too). For example, participants could be asked to maintain and update a two-digit number based on a sequence of high or low tones presented throughout the trial, and report the total at the end of the trial. Initial piloting however showed that 3AFC memory accuracy fell to chance levels in this case. Moreover, the binary outcome of secondary task performance for each trial (i.e., correct or incorrect) may not have sufficient resolution to detect an effect of priming (e.g., given that a proportion of correct trials could be guesses).

Therefore, we sought a continuous measure of performance for each trial of the secondary task, based on a single tone presented at the start of the trial. In Experiment 4a, we piloted a duration judgment task for this purpose, and ran 6 participants just to check that we could detect an effect of central attentional load on such duration judgments (without employing any face priming task).

Method

Participants and Materials

6 participants were recruited internally (3 males, aged $M = 25$, $SD = 3.9$). All were right-handed and had normal or corrected-to-normal vision. The experiment was carried out using the Psychtoolbox package running in MATLAB on a desktop PC using a 1040 x 1280 pixel 15" LCD screen, with responses entered through a USB keyboard. Auditory

cues were played through headphones using Asio sound drivers as recommended for use in psychophysics experiments by the producers of Psychtoolbox software.

Procedure

Participants sat at a desk wearing headphones. Instructions were presented on the screen, and clarified verbally where participants expressed any doubt. The experiment was divided into 16 blocks of 12 trials. Blocks alternated between High and No Load, with participants' starting condition counterbalanced. In the High Load condition, participants were shown a 3 digit 'starting number' on the instruction screen (akin to Experiment 3 High Load condition, but with an extra digit to increase load in an otherwise low-demand task). In both High and No Load conditions, trial structure was as follows: a tone of either low (500 Hz) or high (1000 Hz) was played, whose duration was sampled from a uniform distribution between 300 ms and 800 ms. 200ms after tone onset, a response cue in the form of the word 'Now!' appeared onscreen, and a 100ms response-onset window began. On seeing the cue, participants were required to press and hold the spacebar for a duration that replicated as accurately as possible the duration of the tone from the beginning of the trial.

For High Load blocks, participants were instructed to increment the number in their heads by 1 after a low tone and by 2 after a high tone. At the end of the block, participants were prompted to enter the number they had reached. In No Load blocks, there was no number to maintain and update.

Statistical Analysis

Participants were expected to have a worse representation of tones in the High compared with the No Load condition. Task performance was assessed as the Pearson correlation between the tone and keypress duration, with the expectation that under high load this correlation should be weakened.

Results

All participants showed significant correlations ($p < .05$) between their keypresses and the durations of the tones which they heard, in both High and No Load blocks. These correlations were Fisher-transformed for comparison of High and Low Load conditions. Tone and keypress durations were significantly less correlated in the High Load ($M = .593$, $SD = .139$) than the No Load condition ($M = .714$, $SD = .032$), $t(5) = 2.37$, $p = .032$. Participants were worse at repeating a tone's duration under conditions of working memory load.

Discussion

Experiment 4a showed that an increase in central attentional load can be detected by a decrease in the accuracy with which the duration of a tone could be maintained, as measured by the correlation across trials between actual and estimated durations. The results suggest that a duration report measure should be sensitive to differences in central attentional resources, and therefore in Experiment 4b, we combined this secondary task with the face-priming task of Experiments 1-3.

Experiment 4(b)

To improve on the sensitivity of Experiment 3 to the hypothesised dependence of PSM on central attentional resources, we firstly used the duration report task piloted in Experiment 4a, which would allow us, for the first time, to detect a consequence of the priming effect on the secondary, rather than primary, task. Secondly, we removed tones altogether in the No Load blocks, so they could not cause any interference (load) by virtue of their relevance in the alternating High Load blocks. Thirdly, the secondary task was made more taxing than in Experiment 3, and adapted in such a way that it would dovetail with the duration report task detailed above: As with Experiment 3, participants were shown a two-digit 'starting number' at the beginning of High Load blocks, but this time were instructed to increment this number by 1 for each low (500 Hz) tone, and by 2 for each high (1000 Hz) tone.

One possible outcome was that under the No Load condition, the PSM on the primary task would be replicated, but that it would be significantly attenuated in the High Load condition. Another possible outcome was that the effect of priming would not be detected on the primary task, but rather on the secondary task, where the central resources freed up by priming would be visible as improved performance on the secondary task in the High Load condition for primed relative to unprimed trials (there was no measure of secondary task performance in the No Load condition).

Moreover, the extent to which the priming effect is detected on the primary versus secondary task might even vary across participants, owing to differences in their perceived importance of the two concurrent tasks. In this case, one might also expect to see a negative correlation (in the High Load condition) between the size of the priming effect on the primary task (RT speed-up) and the size of the priming effect on the secondary task (higher correlation between true and estimated tone duration). On the other hand, participants might differ in their overall amount of central resource, such that a positive correlation is seen between the priming effects on primary and secondary tasks (i.e., those with fewer resources might have greater potential to show bigger priming effects on both tasks).

Method

Participants

32 participants (9 males) aged 18-35 ($M = 23$, $SD = 4.4$) years were included in the analysis. There was a high rate of exclusion from this experiment, with participants finding it extremely challenging and failing either to produce a significant correlation between tones heard and the duration of their keypresses ($N = 16$), or to perform significantly above chance at test ($N = 5$), or both ($N = 4$).

Procedure

As with Experiments 2 and 3, the design maintained the same overall structure as Experiment 1, with identical Training, Distraction, and 3AFC Test phases, and with all blocks employing Fixed-Leisure trial time during the Study phase. As with Experiment

3, Study phases alternated blockwise between High and No Load conditions, with starting condition counterbalanced.

During Study phase in No Load blocks, participants carried out the categorisation task as it had appeared in Experiments 1 & 2, i.e. with 500 ms fixation before scene onset, a further 1 s until face onset, and 800 ms scene-face presentation, and fixed 485 ms leisure time after response and before subsequent trial onset.

The stimuli and procedure were identical to experiment 3 except for the following changes: during Study Phase, in addition to the categorisation task (pleasant or unpleasant judgement), participants carried out two concurrent tasks associated with tones which they heard at the beginning of each trial. Tones were equiprobably high (1000 Hz) or low (500 Hz) pitched, with tone onset at either 200 ms or 700 ms after trial onset, and tone durations sampled from a normal distribution, $\mu = 500$ ms, $\sigma = 175$ ms, matched within blocks so that tone pitches, onsets, and durations were equated between primed and unprimed trials. Prior to the beginning of the Study Phase, instructions presented on the screen indicated whether or not the coming task would involve hearing and responding to tones (High Load blocks only). For High Load blocks, the instructions showed a 'starting number', between 10 and 87, and participants were instructed to begin maintaining this number in their head while they carried out the categorisation task, adding 1 to the number each time they heard a low tone and 2 each time they heard a high tone. Due to the complicated nature of the task participants were performing, leisure time was extended by 1250 ms, after which the text 'Now!' appeared onscreen, which cued participants to repeat as accurately as possible, with a keypress, the duration of the tone they had heard at the beginning of the trial. At the end of the block, participants were prompted to enter the number that their internal tally had reached.

Statistical analysis

Analyses of Experiment 4 Training, Study, and Test data remained identical to that employed in Exps. 1-3, with the additional exclusion criterion of trials where no duration report was recorded. In addition to the ANOVA of Study Phase RT data, duration report data were analysed separately, assessing correlations between tone and report duration.

Results

The linear trend analysis showed a significant decrease of RTs with repetition during the Training phase, as expected $T(63) = 6.38, p = < .001$ (Figure 10A).

Primary Face Task

The two-factor ANOVA on Study phase RTs showed a main effect of priming, $F(1,31) = 24.7, p < .001$. As expected, responses to primed faces were faster than to unprimed faces (Figure 10B). Concurrent task load also made Study Phase categorisation RTs slower, $F(1,31) = 35.8, p < .001$ (unlike in Experiment 3), but did not show any interaction with the priming effect on RTs, $F(1,31) = 2.70, p = .11$.

At the Test phase, the 2x2 ANOVA revealed a main effect of PSM, $F(1,31) = 4.26, p = .047$, consistent with Experiments 1-3. A main effect of secondary task showed that performing a concurrent task in the Study phase resulted in worse test performance, $F(1,31) = 62.9, p < .001$, as expected (Figure 10C). The priming effect was largely driven by significant PSM in the no load condition ($M = 3.3\%$, $SE = 1.3\%$), $t(31) = 2.52, p = .017$, with no significant effect seen in the high load condition ($M = 1.4\%$, $SE = 1.8\%$), $t(31) = .782, p = .44$. However, the interaction in this experiment, as in Experiment 3, again failed to reach significance, even when adding back the 16 participants who were excluded because they performed at chance on the secondary task, $F(1,47) = .120, p = .279$.

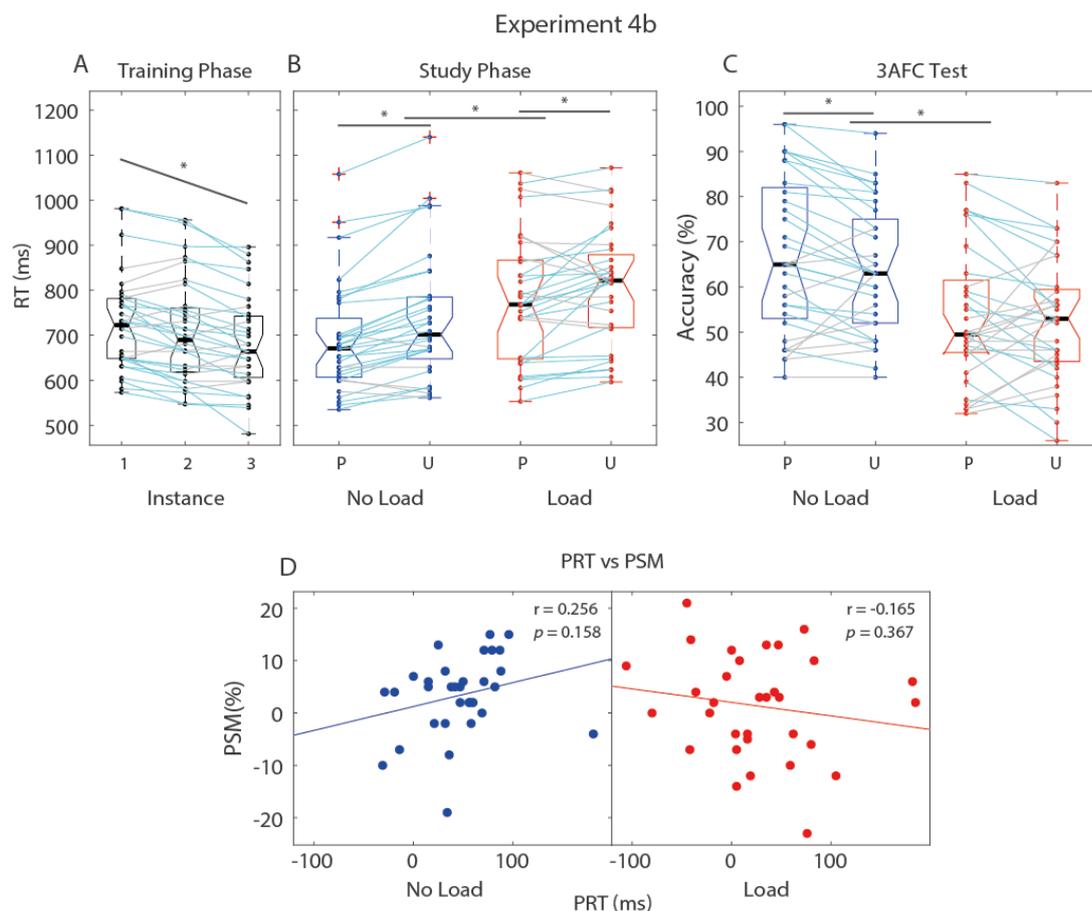


Figure 10: Panels A-C show mean scores ($N = 32$) during 3 phases. Priming of face pleasantness RTs was found during Training (A) and Study (B), but not affected by secondary task during Study. Panel C shows effect of priming on memory accuracy during final 3AFC Test phase and main effect of secondary task during Study Phase, but no interaction. See Figure 6 legend for more details.

Secondary Duration Report Task

Correlations between tone duration and duration report were Fisher-transformed and Primed and Unprimed conditions were compared. The correlation coefficients for the secondary task did not differ between Primed ($M = .437$, $SD = .130$) and Unprimed ($M = .438$, $SD = .202$) trials, $t(31) = 0.027$, $p = .510$, one-tailed. (This was true even when adding back the 5 participants who were excluded because they performed at chance on the primary task, $t(36) = 0.41$, $p = .344$, 1-tailed). In other words, participants did not differ in their ability to maintain tone durations depending on whether faces were primed or unprimed (Figure 11).

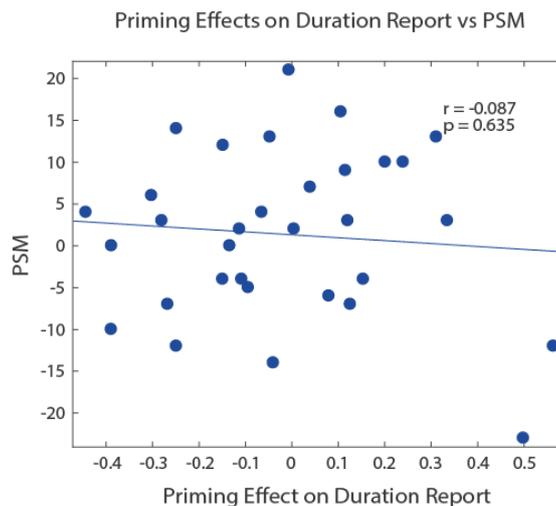


Figure 11: scatter plot shows no relationship between PSM and priming effects on secondary (tone duration report) task.

Discussion

The aim of Experiment 4 was to test whether PSM found in Experiments 1-3 can be explained by priming 'freeing up' central attentional resources during the study phase, and these additional resources then being available for binding the face and scene, leading to improved performance in the later Test Phase. Despite increasing the difference in High and No loads relative to Experiment 3, and now measuring secondary task performance too, Experiment 4 still failed to find conclusive evidence for this central resource hypothesis. Though the PSM effect was numerically smaller under High than No load, and indeed only significant in the No Load condition (i.e., Experiment 4 is the first experiment not to find that PSM was significant under High load), positive evidence from the critical interaction between Load and PSM failed to reach significance.

Meta-Analysis Across Experiments 1-4

We can order conditions in Experiments 2-4 along a continuum of central load, from minimal load in the Intact condition of Experiment 2, intermediate load in the High Load condition of Experiment 3 (where there was essentially only one counting task to perform on tones), to highest load in the High Load condition of Experiment 4 (where there were essentially two tone tasks: to count them and judge their duration). The No Load condition of Experiment 4 may also have minimal load (and participants may have nonetheless been fatigued because of the interleaved, highly demanding High Load blocks), while the Degraded condition of Experiment 2 and Low Load condition of Experiment 3 might have slightly higher loads, though the predictions are less clear. Therefore, while we plot the PSM for all 6 conditions in Figure 12, we did not include these latter 3 conditions in the analysis because of the uncertainty of their load and because this also enabled us to conduct a single between-participant (rather than mixed) analysis of the linear trend across the three remaining conditions.

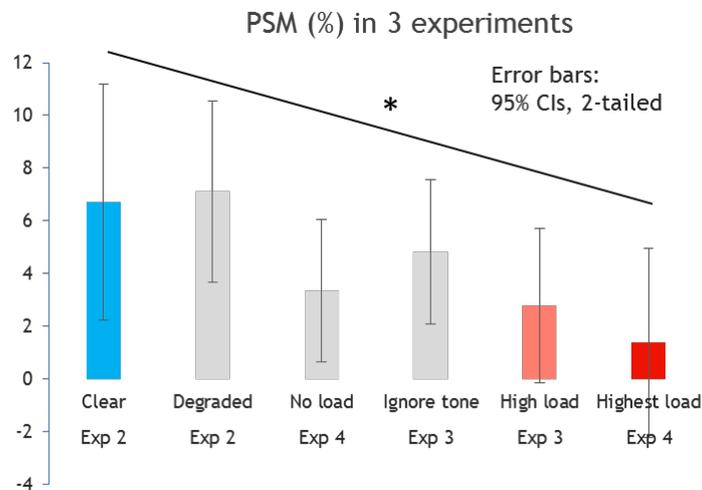


Figure 12: Bars show significant linear decrease of PSM as a function of increasing Central Attention Load across the 3 between-participants (each N = 32) conditions.

A simple linear regression analysis showed that increasing Central Attention Load significantly reduced PSM, $F(1,94) = 4.32$, $p = .040$, with an R^2 of .044. In other words, by combining across Experiments 2-4, we have significant positive evidence that PSM depends on central resources (though effect size is small, i.e. <5%). This result supports

the hypothesis that priming of faces leads to the greater availability (or ‘freeing up’) of central attention, supporting better encoding of item-context pairings later tested in 3AFC Test Phase.

Nonetheless, alternative explanations of these findings remain. One uninteresting explanation is that, as performance on the 3AFC test approaches chance levels (of 33%), there is simply less “room” to detect an effect of priming (i.e., a “floor” effect). However, we note that even under the highest central load condition above (the High Load condition of Experiment 4), performance was significantly greater than chance ($M=53.8\%$, was significantly greater than 33%, $t(31)=7.59$, $p < .001$). Moreover, when we plot PSM (difference between 3AFC performance for primed and unprimed) against baseline performance (unprimed, no load) across all four PSM experiments (hence $N = 112$), the positive correlation that would be predicted by a floor effect is not seen (Figure 13).

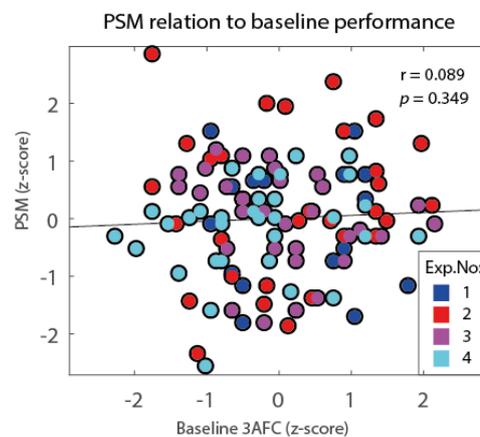


Figure 13: Scatter plot shows no correlation ($N = 112$) between baseline performance and PSM across the four PSM experiments reported in chapters 2 and 3.

No evidence for relationship between PRT and PSM

Assuming the above patterns of means does not reflect a floor effect (and if all participants have the same amount of resources; see Chapter 4), then a simple resource account would seem to predict that the priming effect on subsequent memory (PSM) should correlate positively with the priming effect on RTs at Study (i.e., participants for

whom priming “frees up” more resources should show a greater benefit on subsequent memory). However, across all four experiments, no significant correlation was seen between RT and subsequent memory effect within participants (Figure 14). This finding is explored in detail in Chapter 4, where modelling reveals that only very specific (and unlikely, experimentally) conditions would be likely to produce correlations, even if the resources account is correct.

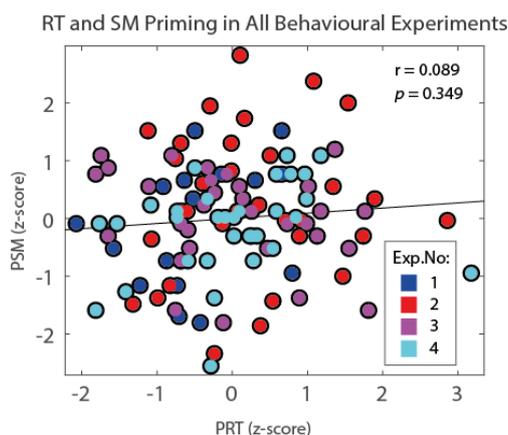


Figure 14: No apparent relationship between priming effects at study on RTs and at test on 3AFC performance across 4 experiments.

Chapter Summary

When combined for analysis with the baseline condition from Experiment 2, the experiments presented in this chapter give qualified support to a central resources account of PSM. The addition of a concurrent central attentional load was shown to interfere with the priming memory advantage, consistent with the idea that priming reduces load on central attentional resources, which are redirected to processing of the scene-face association, yielding improved memory performance.

Although supportive of a resources account, it is important to note it was not the intention to rule out the PIMMS account. In the case that priming does have the effect of increasing scene-face PE, it may be that a central attentional or control process is necessary for the processing of PEs. This idea will be returned to in Chapter 6, the General Discussion.

In the next chapter, a formal model of the resources account is developed. Simulations using this model are shown able to reproduce basic PSM and also qualitatively the results of the addition of perceptual and central load manipulations. The model answers the question of why no reliable relationship was seen in these experiments between PRT and PSM, and shows we should not be confident in rejecting the perceptual resources account based on the null result of no interaction between load and priming.

Chapter 4

Testing Attentional Accounts III: Computational Model

Introduction

Experiments 1-4 attempted to address experimentally the idea that a particular type of resource is “freed up” by prior exposure to stimuli (in primed trials), and that this freeing of resources accounts for the PSM effect. No significant interaction of Load and PSM was found in Experiment 2, providing no evidence of perceptual resources being freed by priming. However, by combining data from three experiments with monotonically increasing central load, significant evidence was found that reducing central attentional resources does reduce the PSM effect (although the effect size was small). The present chapter addresses the question from a computational perspective. By formalising our hypotheses, it is possible to gain further insight into possible explanations for the results of Experiments 1-4. In particular, the modelling will make apparent the difficulty of inferring the underlying resource model simply from interaction patterns and correlations in data.¹

Relating resources to performance

Norman and Bobrow (1975) describe the performance-resource function in its most general terms as monotonically non-decreasing. In theory, a wide variety of such functions are plausible, but a reasonable case would be a sigmoid: at low levels of resource allocation, task performance is impossible, and after this threshold there is an improvement in performance as resources are applied, followed by a tailing-off as additional resources no longer improve performance due to sensory or response limits:

$$P = \frac{1}{\left(1 + e^{-(r-d) \dot{E} s}\right)}$$

where P is performance, r is resources, s is the sharpness of the sigmoid and d is the “difficulty” of the task. (In general, we would expect $P = 0$ if $r = 0$, which is true if r ranges from $-\infty$ to $+\infty$, but in reality r must be finite, so P is allowed to be slightly above zero when $r = 0$). If r ranges between 0 and 1, and $s = 0.1$, then Figure 15 shows two different difficulties: $d = 0.35$ (harder) and $d = 0.20$ (easier). As can be seen from

¹ The MATLAB code with which this model was implemented and the simulations generated, is freely available online at the following URL: www.github.com/AlexJKaula/PhDmodelCode

comparing Line B with Line A, a lower difficulty (for example, because of priming) means a higher performance (P) results from the same level resources (r).

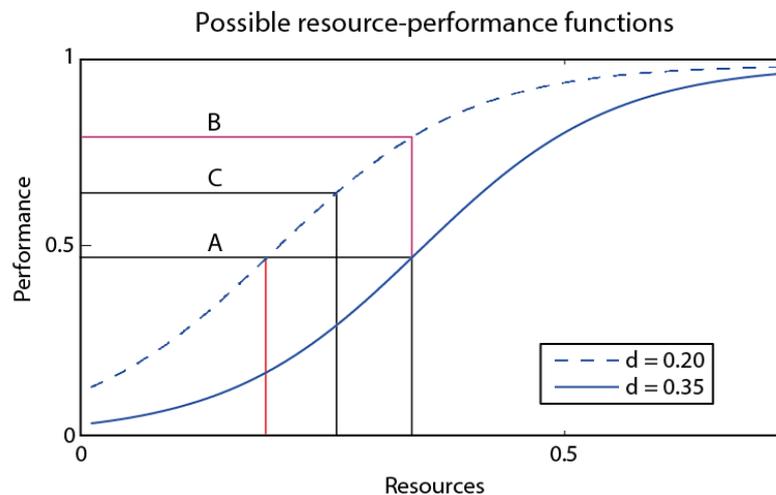


Figure 15: Schematic illustration shows how changing difficulty parameter, d , can allow improved performance with reduced resources. Line A shows normal performance. Line B shows that when difficulty parameter d is reduced, performance for the same amount of resources is increased. Line C shows increased performance with reduced resources. Red vertical line between x-axis and Line A indicates resources required for maintenance of performance when change in d makes the task easier.

Under these basic assumptions, an important feature is clear even before any further modelling takes place: given we do not know the form of the function, and cannot directly measure the hypothesised resources, it is possible to produce several qualitatively different data patterns from a single underlying resource. Figure 16 shows two such patterns. Panel B shows an interaction pattern (as across Experiments 3-4), which is often used to argue for a single underlying resource. However, Panel A shows that the same single resource can also produce two main effects in the absence of an interaction, which is often used to argue for separate resources (as in Experiment 2). Whether or not an interaction pattern is found depends only on whether there is a difference in the gradient of the resource-performance function between primed and unprimed data points under high and low load, and thus in the absence of further information, one cannot infer the number of resources from basic performance levels in a 2x2 design. (Only by finding a reversed association in a 2x3 design can more than one resource be inferred, Dunn & Kirsner, 1988.)

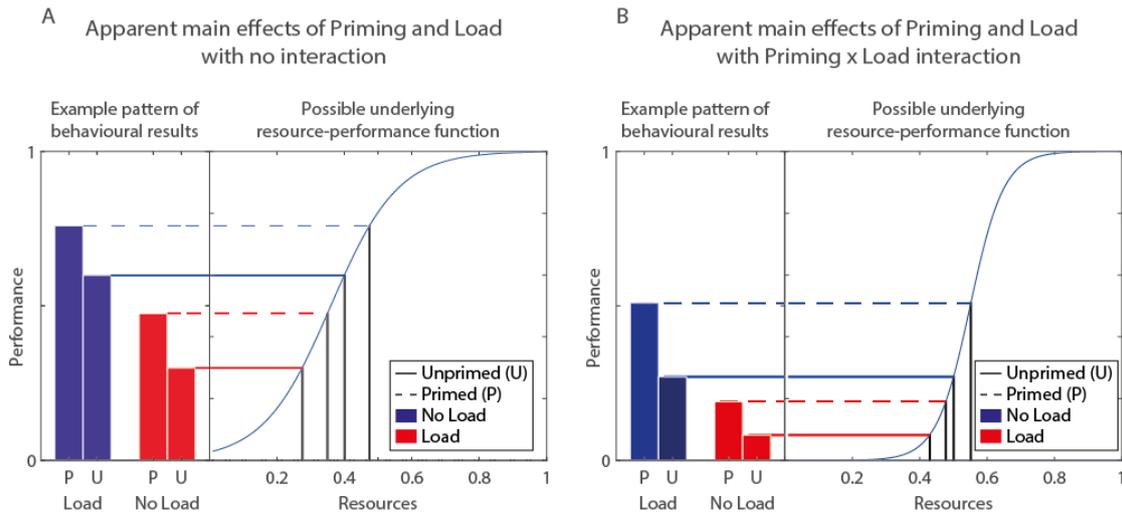


Figure 16: Schematic shows how distinct data patterns can arise from sampling different points on a single (nonlinear) performance-resource function. Panel A shows main effects of both priming and load, but no interaction, while Panel B shows a priming-by-load interaction.

However, performance on the final associative memory task is not the only measure in our experiments: there are also measures of performance (PRT) during the study task (which is when resources are assumed to impact) and also performance on the secondary task (in Experiments 3-4). To explore how these measures relate to each other, according to a single resource account, it helps to formalize a model of these tasks. It turns out that this model can easily simulate the patterns of performance on all measures, under any of the various assumptions concerning resource allocation described below. Furthermore, depending on how resource allocation under load is simulated, the model may produce highly correlated, anti-correlated, or non-correlated measures of PRT and PSM.

Modelling PRT, PSM, and load task performance

1. Participant Resources and Resource Allocation

In the model, each participant, s , is assumed to possess some total amount of resource R_s , sampled from a normal distribution, with the proportion allocated to each of $i = 1 \dots M$ simultaneous tasks being defined as $0 < a_i < 1$, such that $\sum_{i=1}^M a_i = 1$ and $r_{is} = a_i R_s$. For the basic PSM paradigm without any concurrent load (as in Experiment 1), $M = 2$, such that $r_{1s} = a_1 R_s$ represents the resources allocated (by participant s) to the

primary Study task of making face pleasantness judgments (hereafter, “Face Task”) and $r_{2s} = a_2R_s$ or $r_{2s} = (1 - a_1)R_s$ represents the resources remaining for the “second task” of encoding the scene-face associations into memory (hereafter, “Memory Task”). Finally, to introduce some random variability, measurement noise was added, $u \sim N(0, \sigma)$, to the final performance values.

2. Effect of Priming

As alluded to earlier, priming can be modelled as decreasing the difficulty, d , of the face task, such that performance increases for a given level of resource, r . Thus $d_{1p} < d_1$, where p indicates a primed trial. Changing parameter d has the effect of translating the resource-performance function in question along the x -axis, as seen already in Figure 15. This means that the participant can reduce the proportion of resources allocated to this primary task by as much as allowed by the change in difficulty, $d_1 - d_{1p}$ produced by priming, and still maintain performance of that task. This in turn releases more resources for the “second task” of memory encoding, i.e. $r_{2sp} = a_{2p}R_s > (1 - a_1)R_s = r_{2s}$. More resources released means that memory encoding is better for primed trials, explaining basic PSM.

In reality, we know that performance on the primary task also improves, suggesting that participants (on average) do not release all the resources that are saved by priming. Instead, it is assumed that participants maintain a proportion, $0 < x_p < 1$, of the resources freed by priming, i.e., $r_{1ps} = a_1R_s - x_p(d_1 - d_{1p})$ and hence $r_{2ps} = (1 - a_1)R_s + x_p(d_1 - d_{1p})$.

Figure 14 shows results of simulating the same number of participants ($N=32$) as in Experiment 2. Performance on the Face Task improves following priming (Panel A), concurrent with improved performance on the Memory Task (Panel B). Note that, because accuracy is close to ceiling for the Face Task, higher performance is expressed as shorter RTs, so we used a simple linear equation, $RT = 1000(1 - P) + 300$, to map better performance onto shorter RTs (i.e., PRT; Panel C). Interestingly however, with this minimal noise level, no significant correlation is produced between PRT and PSM (Panel D). The reason for this lack of correlation is explored in the next section.

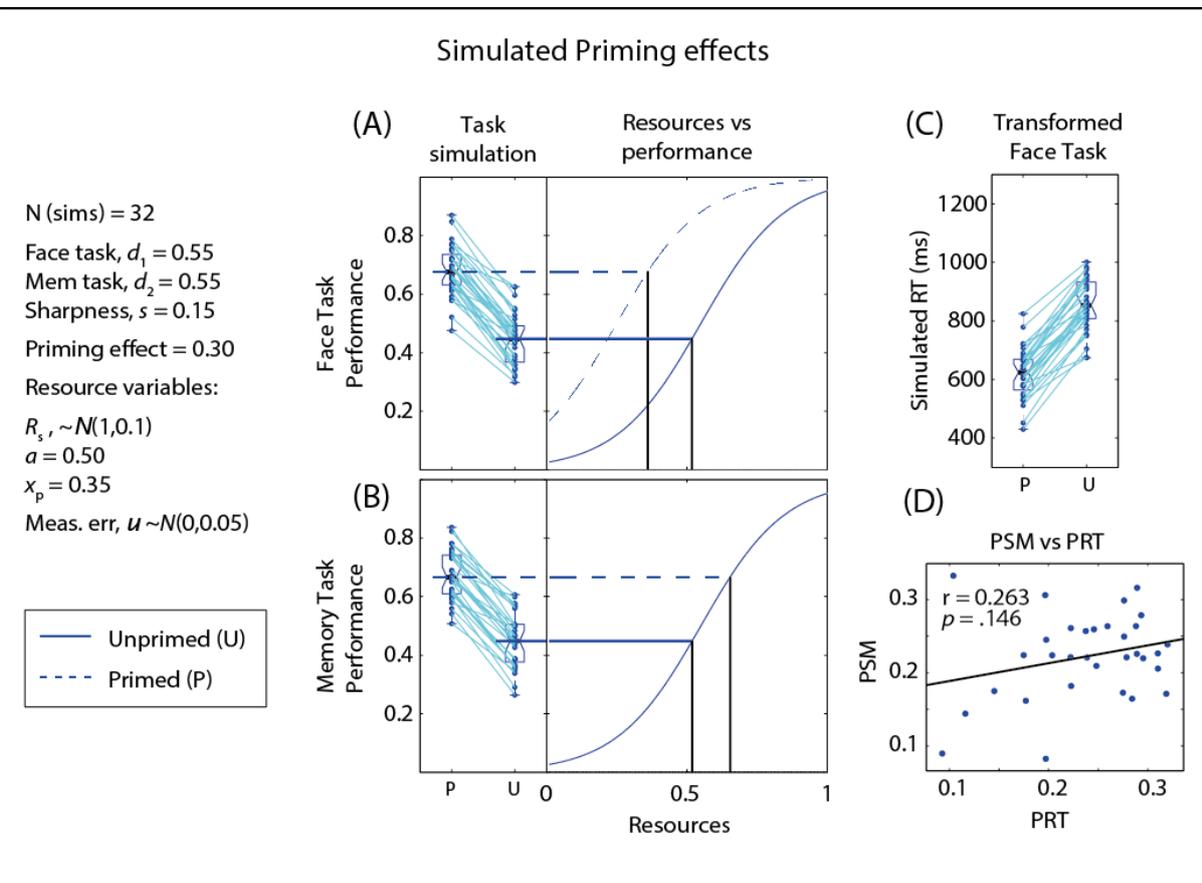


Figure 17: Basic simulation ($N=32$) of PRT and PSM, illustrating how priming can be modelled as freeing up resources from Face Task for Memory Task. Lines between boxplot markers show individual simulations. N.B. this simulation shows an exaggerated effect for illustration purposes. Variation in simulated performance comes from randomly sampling values of simulated-subject overall resources, R_s , and by addition of measurement error for each data point.

3. Correlations Between PRT and PSM

In Experiments 1-4, we did not find correlations between PRT and PSM (the exception was the positive correlation in the perceptual load condition of Experiment 2), and this

seemed puzzling if they depend on the same resource, for which one might expect a negative correlation (i.e., performance on one task might be expected to decrease while performance on another increases, as more resources are diverted from former to latter). However, simulation of the simplest version of the experiment (i.e. Face Task and Memory Task, with no perceptual or central load) shows that simply by varying only the allocation parameter, a_1 , a range of positive, zero or negative correlations can be produced (see Figure 18).

To understand the reasons for this, first, measurement error, u , is removed from the simulation. Performance of the two primary (Face and Memory) tasks is then simulated, randomly sampling R_s while keeping the s and d sigmoid parameters equated for both tasks, and keeping a_1 and x_p both fixed at .5, simulating participants allocating resources evenly between the tasks (Panel A of Figure 18). In this case, PRT and PSM are almost perfectly positively correlated. As can be seen from the illustration, this positive correlation arises because, under these assumptions about allocation, priming effects on both RT and Memory are produced from the same regions of the sigmoids underpinning task performance.

However, as a_1 is allowed to vary towards either 0 or 1 (entailing asymmetries of R_s allocation), correlations become negative (Panel B). Again, the illustration helps understand this feature of the model: for higher-than-average-value samples of R_s , higher performance in the task allocated more of the resources places a participant at the lowest-gradient part of a decreasing-gradient part of the sigmoid, thus producing limited priming effect. Meanwhile, in the other task, task performance for the same participant will be produced from the highest-gradient part of an increasing-gradient part of the sigmoid, producing a relatively large priming effect, and thus we see the negative correlation between PRT and PSM.

If a_1 is set between values producing strong positive and negative correlations, then sampling R_s produces unusual patterning in the relationship between PSM and PRT (Panel C). Under this situation, the pattern (which could be described as divergent values of PSM produced for similar values of PRT) may be explained by the fact that

both PRT and PSM are difference scores, so higher and lower R_s values may place simulated participants in similar-gradient regions of one sigmoid and differing-gradient regions of the other.

The relationships shown in the simulations are delicate: small differences in a_1 tip the balance from positive, to 'strange', to negative. However, it seems unrealistic to assume that every participant uses the same re-allocation of their resources between the two tasks. If random sampling of a_1 from a uniform distribution $U(0,1)$ is now added, then, as shown in Panel D, any apparent relationship between PRT and PSM is no longer seen. Note that in all other simulations shown and discussed in this chapter, the parameter a_1 is kept fixed at .5, and measurement error, $u \sim N(0,0.05)$ is sufficient to obscure what would otherwise be positive correlations between PSM and PRT.

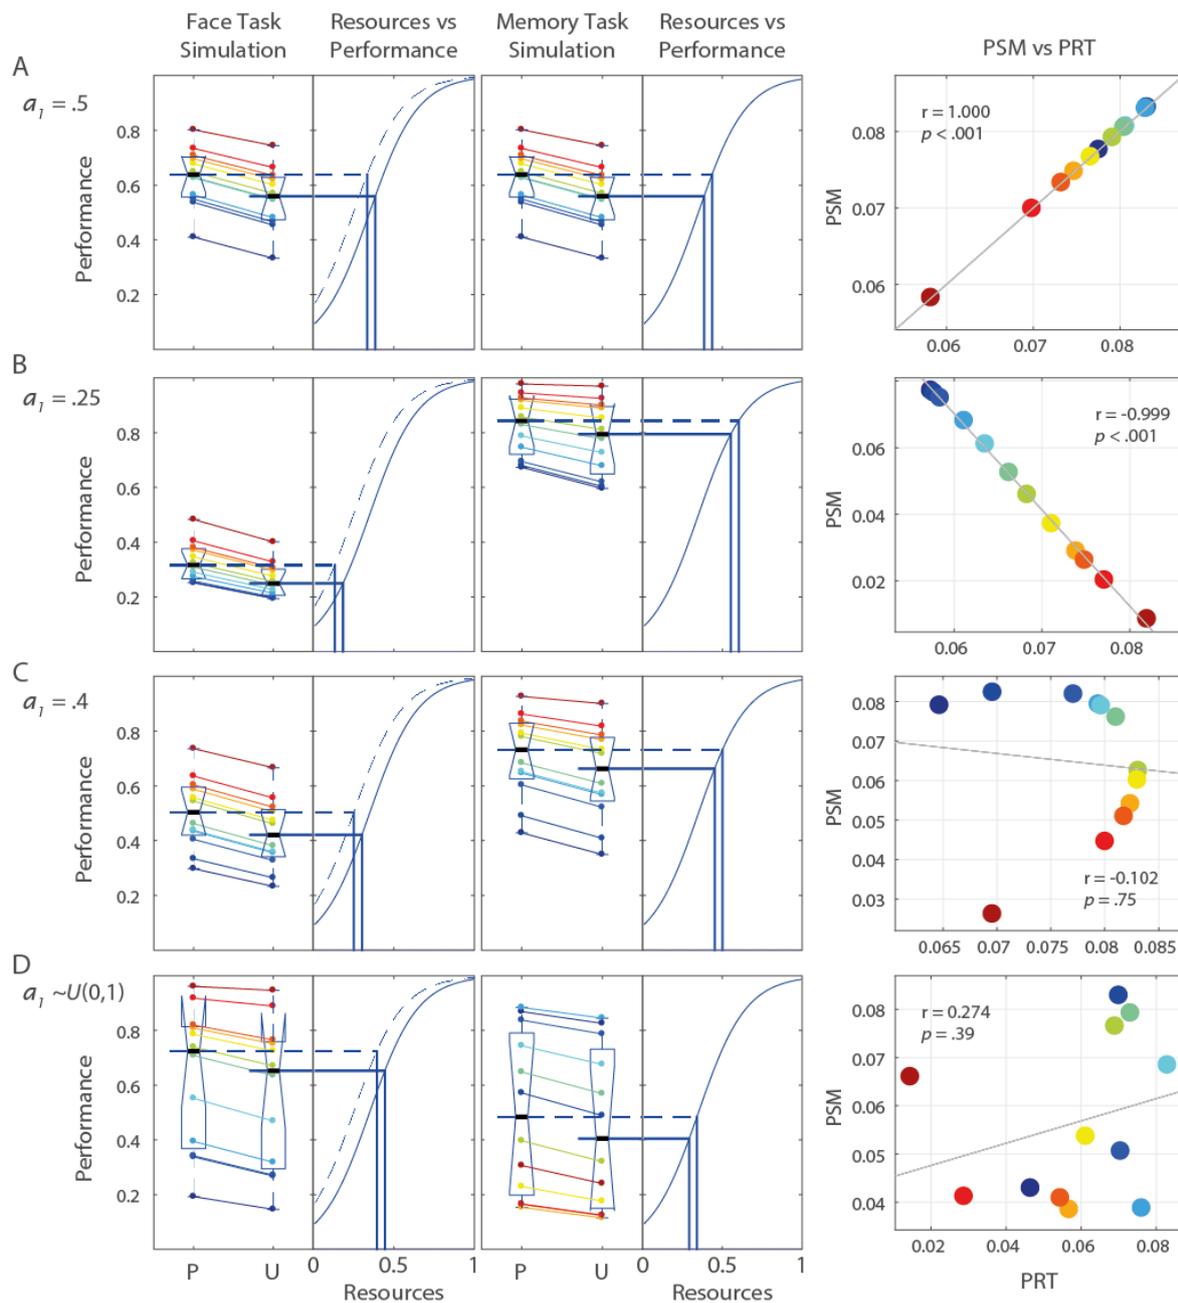


Figure 18: Panels A-D show 4 ($N=12$) simulations of primed (P) and unprimed (U) performance in Face and Memory tasks. Overall resources, R_s , were randomly sampled from $N(0.75, 0.15)$ each with a different value for allocation parameter, a_1 , with all other settings equated and measurement error, u , set to zero. Colours used are consistent for individual simulations across panels, for ease of comparing performance across tasks and with scatter plots.

4. Effects of Load

4.1 Perceptual Load

A simple way to model the effect of adding perceptual load is to assume that it increases the difficulty of the face task, meaning more resources are required to maintain performance, i.e. $d_{1l} > d_1$, where l stands for load (the low load condition is identical to the default condition in the previous simulations). In the load condition, the participant must either redirect resources from the memory task, or perform the face task less well, or both. Since in the experiment we observed that on average participants did perform both tasks less well, we can assume that some proportion, x_l , of the additional resources that would have been required to maintain performance in the face task (equal to the change in the parameter d_1 due to the addition of perceptual load), are redirected from the memory task, i.e., for unprimed trials:

$$r_{1ls} = a_1 R_s + x_l (d_{1l} - d_1)$$

$$r_{2ls} = (1 - a_1) R_s - x_l (d_{1l} - d_1)$$

whereas for primed trials:

$$r_{1pls} = a_1 R_s - x_p (d_1 - d_{1p}) + x_l (d_{1l} - d_1)$$

$$r_{2pls} = (1 - a_1) R_s + x_p (d_1 - d_{1p}) - x_l (d_{1l} - d_1)$$

This model can easily produce the qualitative pattern of results in Experiment 2, with main effects on RT and subsequent memory of both priming and load, as shown in Figure 19.

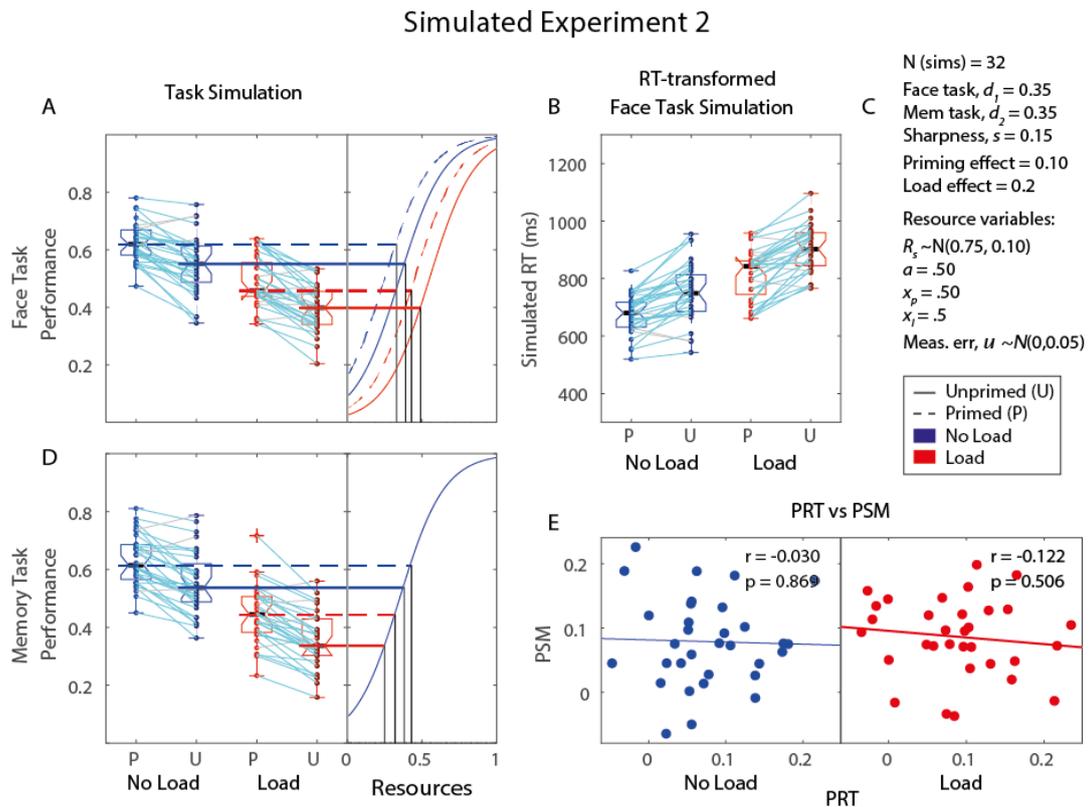


Figure 19: Qualitative reproduction of the results of Experiment 2. Panels A and B show PRT in the Face Task, and also an effect of perceptual load. Panel D shows effects of both priming and load on subsequent memory performance, with no apparent interaction. Lines plotted between points panels A, B, and D, connect single subject performance in Primed and Unprimed conditions of No Load and Perceptual Load. Cyan lines indicate an apparent priming effect, grey lines ‘negative’ priming. Heavier lines linking boxplots with sigmoid curves show median performance in each of the 4 conditions, and how that performance maps to the sigmoid functions that underlie performance.

4.2 Central Load

In Experiment 3, a secondary central load task was added to the study phase, and in Experiment 4b this task was replaced with a more challenging one. It is possible to model both of these experiments by introducing a third task with its own (independent) resource-performance curve, and with some amount, r_{3s} , of a participant’s overall resources being directed to Task 3 performance, as determined by the task’s allocation parameter, a_3 . The most obvious difference between how Experiment 2 and Experiments 3 & 4b are modelled is that, instead of resources being transferred from Memory Task to Face Task, they are instead transferred away from both tasks in some proportion (see Figure 20 for illustration).

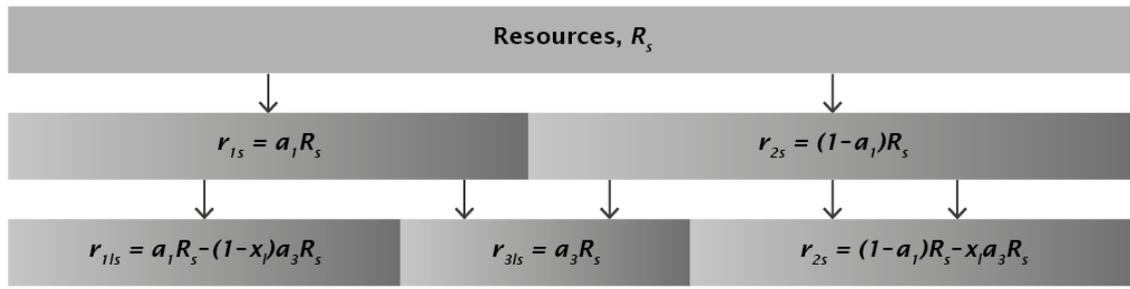


Figure 20: initial resource allocation (1st row), resource allocation in No Load condition (2nd row), with resources split between Face and Memory tasks, and in (central) Load condition (3rd row), where a secondary task is added to Face and Memory tasks.

x_l can continue to be used, but in this case it will determine the proportion of the resources required for the central load task to be taken from the Memory Task, $x_l a_3 R_s$, and from the Face Task, $(1 - x_l) a_3 R_s$. Thus in unprimed trials under load, resources for the three tasks will be:

$$r_{1ls} = a_1 R_s - (1 - x_l) a_3 R_s$$

$$r_{2ls} = (1 - a_1) R_s - x_l a_3 R_s$$

$$r_{3ls} = a_3 R_s$$

In Experiment 3, there was not a trial-by-trial measure of performance, so it was not possible to address whether performance was improved in primed trials. In Experiment 4a, a measure was found that was sensitive to changes in central attentional resources (The Duration Task), and Experiment 4b addressed whether priming might affect performance in such a task. This can be modelled by adding a further parameter, x_{p2} , which is the proportion of the resources freed from the Face Task by priming which go toward Memory Task performance, the remainder of which (i.e. $1 - x_{p2}$) will be added to secondary (Duration Task) resources, so in primed trials under load, resources are distributed thus:

$$r_{1pls} = a_1 R_s - (1 - x_l) a_3 R_s - x_p (d_1 - d_{1p})$$

$$r_{2pls} = (1 - a_1) R_s - (1 - x_l) a_3 R_s + x_{p2} x_p (d_1 - d_{1p})$$

$$r_{3ps} = a_3 R_s + (1 - x_{p2}) x_p (d_1 - d_{1p})$$

In practice, there was no advantage of priming for Duration Task trials, so $x_{p2} = 1$, and all resources freed from the Face Task by priming go towards performance of the Memory Task.

Experiments 3 & 4b

In Experiment 3, performance of the secondary load task was at ceiling. Furthermore we observed a main effect of load on memory performance, but no such effect on RTs in the study phase, so we can assume that performance of that task was accomplished using resources redirected away from the Memory, not Face Task, and thus set the proportion of load demand to be met by Memory Task resources, x_l , to 1. Figure 21 shows the results of simulating Experiment 3. As apparent in Panels A and B, with x_l set to 1, no load effect is seen on performance of the Face Task: the demands of the secondary task are met wholly by resources redirected from the memory task. Panel D therefore shows a large main effect of load on subsequent memory. In line with my experimental data, in which performance of the secondary load task was at ceiling in Experiment 3 and in which Memory Task performance was not reduced to the same extent as in Experiment 4b, it is simply assumed that resource allocation to the load task, a_3 , is smaller than in Experiment 4b. Performance of the secondary task is not simulated, but it is assumed that r_{3ls} is always sufficient for maximum performance (the underlying difficulty parameter, d_3 , simply being low enough to accommodate this).

Simulated Experiment 3

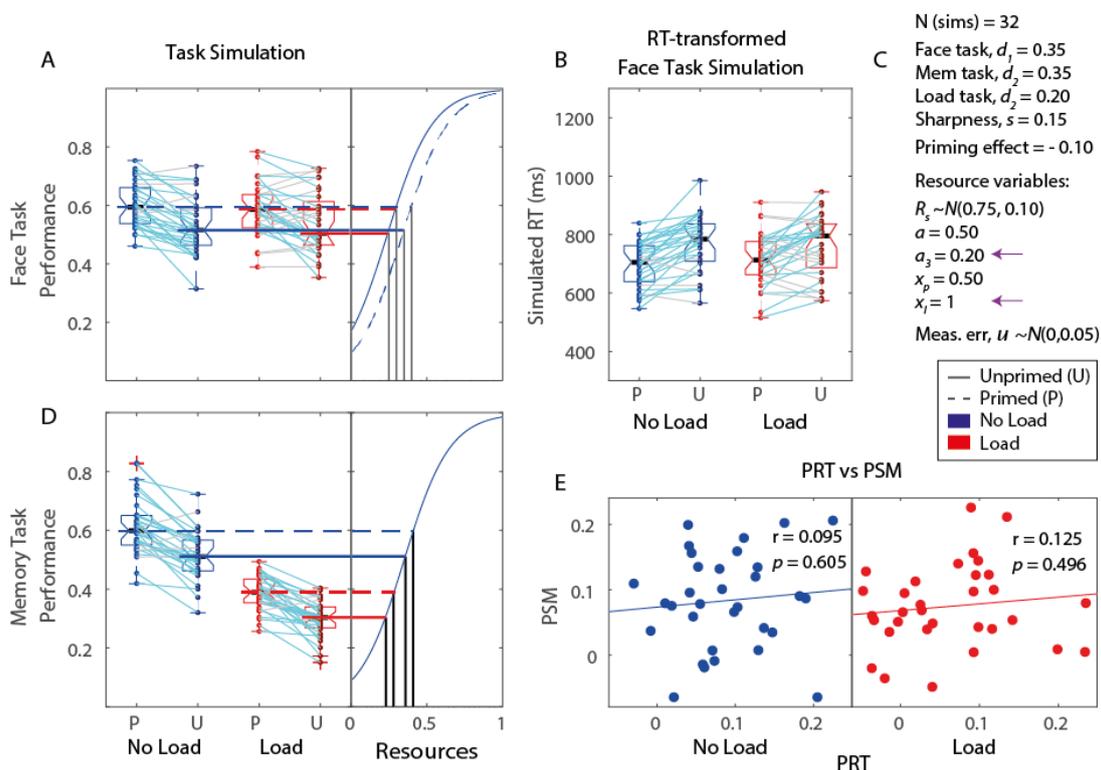


Figure 21 Qualitative reproduction of the results of Experiment 3, in which we added Central Load by including a dual task during the study phase in half of the blocks. Panels A and B show PRT with no effect of load in the Face Task. Panel D shows the effects of both priming and load on subsequent memory performance. Consistent with the behavioural results, scatterplots in panel E do not show any clear relationship between PRT and PSM. Panel C shows the settings used in this simulation, with small arrows highlighting the two parameters which were changed between simulations of Experiments 3 and 4b.

In contrast to the results of Experiment 3, Experiment 4b produced a main effect of load on RTs and also on memory performance. In addition, the design afforded a secondary-task performance measure during each trial, in which it was hypothesised an effect of priming may be seen, although none was found. Experiment 4b was therefore simulated (see Figure 22), by adjusting 2 parameters (see small arrows, Panel C) of the Experiment 3 model and adding another parameter. The secondary task load allocation parameter, a_3 , was increased from 0.20 to 0.25, reflecting the increased load of the secondary task, and the load distribution parameter, x_l , was reduced from 1 (all resources redirected from Memory Task) to 0.65, reflecting load's effect on Face Task performance in Experiment 4b. As with previous simulations, and consistent with behavioural results, no obvious relationship was found between PRT and PSM (Panel E). A new parameter, x_{p2} , was included, which was the proportion of resources freed from the Face Task by

priming which would go towards Memory Task performance, and the remainder of which would go towards secondary (Duration) task performance. Since no significant effects of priming on Duration Task performance were seen, this new parameter, x_{p2} was set to 1, and the simulation does not produce any priming effect on Duration Task performance (Panel F).

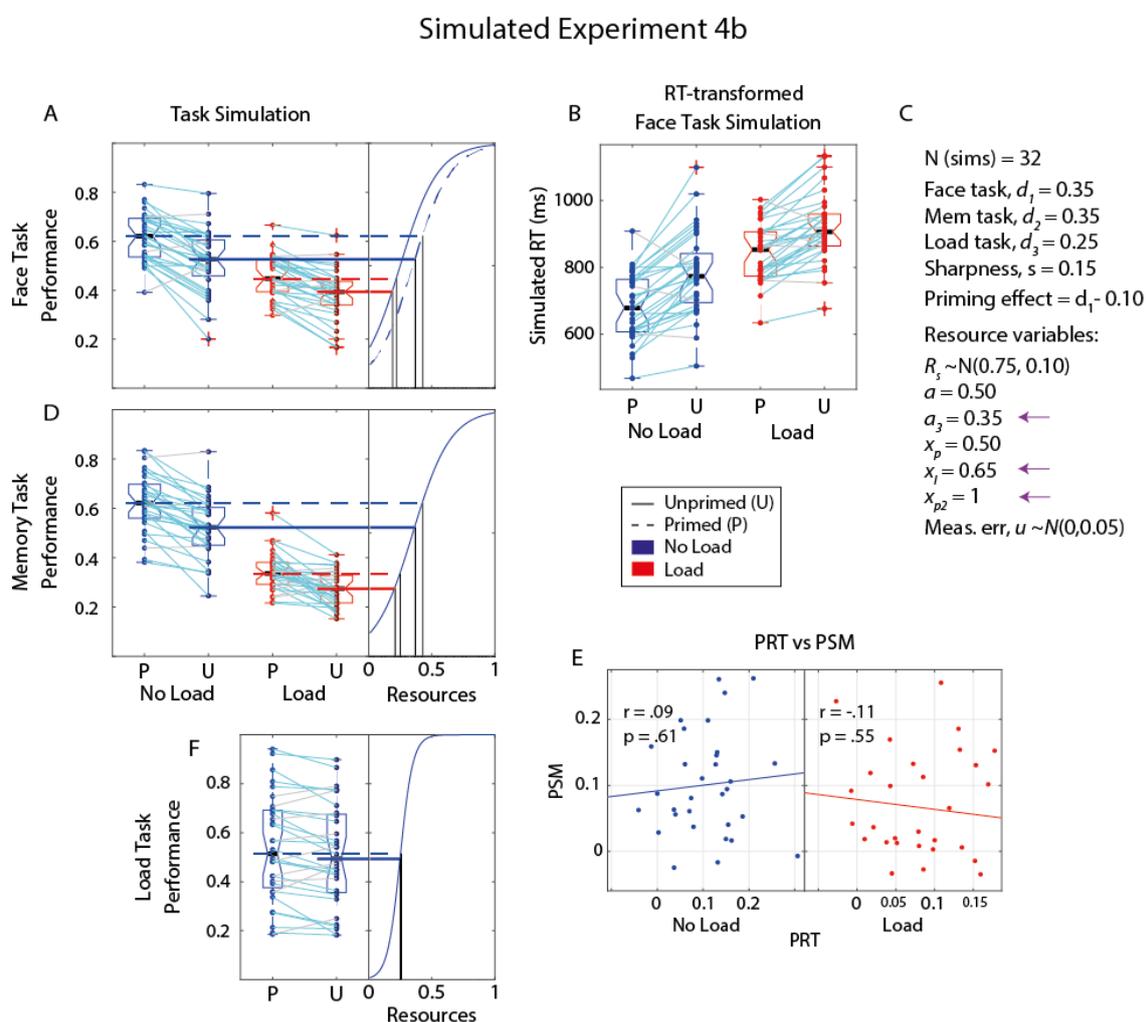


Figure 22: Qualitative reproduction of the results of Experiment 4b, in which we further increased Central Load by including a more difficult dual task during the study phase in half of the blocks. Panels A and B show PRT and effect of load in the Face Task. Panel D shows effects of both priming and load on subsequent memory performance, with the suggestion of an interaction between load and PSM. Consistent with behavioural results, panel E does not show any clear relationship between PRT and PSM. Panel F shows simulated performance on the secondary load task, with no effect of priming. Panel C shows the settings used in this simulation, and small arrows highlight the two parameters which were changed from the simulation of Experiment 3, and the addition of parameter x_{p2} (see text).

Chapter Summary and Discussion

This chapter began by showing that, with only minimal assumptions about how resources may relate to performance (that monotonic, but nonlinear, e.g., sigmoidal), and the simple idea that priming affects the difficulty of the Face Task, one can reproduce the basic PSM of Experiment 1. Indeed, when adding an additional attentional load of any kind, it is easy to see how different points on a nonlinear resource-performance function can produce a range of interaction patterns between load and priming on memory performance: from interaction to no interaction. However, it is less obvious that this pattern of performance on the Memory Task can be fit simultaneously with performance on the Face Task, and with the general lack of correlation between Face and Memory task performance (since *prima facie*, one would assume some dependency between these measures if they share the same total resource).

Therefore a relatively simple computational model was implemented to fit the data from Experiments 2-4b. The effects of attentional load were modelled in two structurally-different, but mathematically-equivalent ways, reflecting the difference in the nature of the perceptual load and the central load manipulations: it was assumed that perceptual load affects the difficulty of the task (like priming), whereas central load was assumed equivalent to adding a third task. In both cases, the result was to change the resources available for the Face Task and Memory Task. In other words, regardless of whether load is considered as making the Face Task more difficult (perceptual load, as in Experiment 2), or as having its own independent resource-performance function (central load, as in Experiments 3 and 4b), the model can straightforwardly generate the patterns of results observed across all four experiments. Thus one cannot refute this resources model account of PSM based on these behavioural data.

Despite their mathematical equivalence, there is a conceptual difference between perceptual and central load. For perceptual load, the change to the stimulus itself was assumed to affect the difficulty of the task, rather than add a new task *per se*. Perception of the stimulus logically precedes making a decision about it, and adding noise to the image did indeed increase RTs in the Study Phase. Thus while perceptual load could be

modelled instead as reducing resources available, it seemed more sensible to model in the converse way to priming, with absorption rather than release of resources.

For the central load however, a third task was assumed to be imposed that drew its own resources from the other tasks. Adding a new task entails a further resource-performance curve for that task, which can potentially be used to explain any behavioural measures from the central load (duration) task (even though in Experiment 4b, no effect of priming was seen for the duration task). By doing so, it was assumed that such a task would not directly interfere with a perceptual decision, and neuroimaging studies suggest somewhat separable systems subserving perceptual- and central-attentional tasks (e.g. see Burgess et al., 2007). Despite this, the model lays bare the fact that whether one system or another is preferentially recruited would be indistinguishable from a behavioural point of view. It might be tempting to conclude that, since there was no effect of central load on Face Task RTs in Experiment 3, the secondary task did not interfere with the Face Task. However, the model illustrates we do not have evidence for this: whether or not there is an effect of load on Face Task RTs depends only on the value of the x_l parameter. If $x_l = 1$, then regardless load demands on resources, they are entirely met by Memory Task resources, so no RT effect will be observed.

Since, *prima facie*, it seems reasonable to assume that a secondary task should have its own function relating resources to performance, that is how secondary tasks were implemented in the model. The model's design was constrained by the behavioural design as well as the findings: in Experiment 3, performance was at ceiling, and did not offer a trial-by-trial measure of performance – thus in the model its resource-performance curve is implicit and we simply assume the difficulty parameter to be small enough that performance will always be at ceiling. In Experiment 4b however, there was a measure of performance during each trial, the Duration Task, so to maintain congruence with this aspect of the design, performance of this additional task was simulated. The proportion of the simulated participants' overall resources, R_s , that would be allocated to the Duration Task was fixed with parameter a_3 , and the new

parameter, x_{p2} , was set to 1 so that resources freed from the Face Task by priming would all go toward improved performance of the Memory Task – and no difference would be seen in Duration Task performance between unprimed and primed trials.

As well as capturing mean performance levels in the various tasks, the model proved capable of producing a range of correlations across participants in their performance of those tasks, from negative to positive correlations. This illustrates the danger of intuitions that, for example, drawing from the same fixed resource means that the correlations should be negative. Again, the key aspect of the model that enabled this flexibility is the nonlinear sigmoid function, such that just by changing one parameter (the allocation of resource between tasks, a), performance is determined by different regions on the sigmoid that can have sharper or shallower gradients for one task relative to the other. (The further addition of minimal measurement noise can also easily disguise any true but weak correlations, resulting in nonsignificant correlations given the number of participants tested here; see also (Berry, Shanks, & Henson, 2008).

Despite the fact that the model was the simplest that could be devised which satisfied the basic conception of resources, it contains several free parameters, and it was necessary to constrain some of them in order to understand the model. For example, the resource-performance curves for the Face Task and Memory Task were kept identical, which in practice seems unlikely. Furthermore, the concept of priority was realised in the combination of several variables: the initial allocation between tasks, a , and the proportion of demanded or freed resources for the task concerned, x_l and x_p . These values were fixed once, and for all participants, whereas in practice the priority given to one task or another is likely to have varied considerably between participants.

These potential extra degrees of freedom in the model make it even harder to reject on the basis of behavioural data alone. Thus, in order to learn more about how resources may be freed, it may be helpful to examine the neural underpinning of PSM. One of the questions this work sought to address was whether either freeing of perceptual resources, or of central attentional resources could explain PSM. However, although this is an interesting question to ask, this project sets out to address a different question: whether a ‘pure resources’, or a prediction-error-based (PIMMS) account is a better

explanation for PSM. Here, again, the two explanations we wish to compare make different predictions about what to expect if we look at overall brain activity and also patterns of activity, and it is these predictions which I address in the next chapter.

The model of PSM proposed here, and shown not to be falsifiable on the basis of behavioural data, is conceptually simple: resources are freed by priming and redirected to mnemonic processes. The PIMMS account is somewhat different: improved memory performance when faces are primed arises because the perceptual system responds to such a face by producing clearer evidence (a more precise likelihood) for primed faces. Another way to put this is that the brain response to the perceptual evidence must be more distinctive, better separating the particular face that has been seen from the average face. This notion of representational ‘sharpening’ does not make a specific prediction regarding univariate activation, but it does predict that patterns of activity associated with seeing individual primed faces will be more distinctive (and thus less correlated with one another), than in the case of an unprimed face. In Chapter 5, these predictions will be considered in more detail. It is shown, using representational similarity analysis (RSA; Kriegeskorte, Mur, & Bandettini, 2008), that priming faces does lead to more distinctive patterns of brain response in fusiform face area (FFA), a key face-processing region, and the question is then considered of whether this may support a more nuanced account of PSM than the simple resource-based one modelled in this chapter.

Chapter 5

Using fMRI to Examine the Neural Basis of PSM

Introduction

In Chapter 4, the modelling of the earlier experiments showed that conclusions based on behavioural experiments regarding resource-based accounts of PSM based should be treated with caution, since the patterns of results found in those experiments can be reproduced while making only minimal assumptions about relationships between resource allocation and performance. One fruitful way to take forward this line of enquiry may be to address what happens at the neural level during training and study periods in a modified version of the paradigm used to generate PSM in Experiments 1-4.

Analysis of activity in brain regions that are selectively activated by our stimuli will allow us to address the critical question of whether PSM can be explained by the simple resources account examined in earlier behavioural experiments and modelled in Chapter 3, or whether PIMMS can offer a more parsimonious explanation. Pictures of faces and places have been repeatedly shown to elicit stimulus-specific responses in the brain, each with respective eponymous regions, the fusiform face area (FFA, Kanwisher et al., 1997) and parahippocampal place area (PPA, Epstein & Kanwisher, 1998). Indeed, this was one of the original reasons for using faces and scenes (rather than the words and sounds of Gagnepain et al., 2008). Furthermore, the hippocampus is well-established as important for encoding of associative memories, as confirmed by the PSM fMRI study of Gagnepain et al. (2011). Thus, the paradigm used in Experiments 1-4 is well suited to a region-of-interest (ROI) analysis focusing on FFA, PPA, and hippocampus ROIs. Examining univariate activation in each ROI can test the resources account, assuming that activation is related to amount of resources devoted to a particular task (Shallice, 2003), while examining multivariate patterns in each ROI can test the PIMMS prediction of sharpened sensory evidence after priming. Furthermore, by examining connectivity between ROIs, we can test the PIMMS explanation offered for the findings of Gagnepain et al. (2011), as we expand below.

Predictions of Resources Account

Under the simple formulation of the resource account modelled in Chapter 4, PSM is explained by the reduced demand for performance of the Face Task freeing-up resources

for performance of the Memory Task. Assuming that greater use of resources entails greater metabolic demand and hence higher fMRI signal in regions using those demands, the basic resource account predicts reduced activation in face-processing regions for primed trials. These freed resources might either be directed to increased processing of scenes and/or of the scene-face association. In other words, priming-related decreases in FFA activation should be accompanied by *increases* in PPA and/or hippocampus, and the size of these neural priming effects should be positively correlated between regions (across trials or participants). Furthermore, the size of these neural priming effects should correlate, across participants, with the size of behavioural PSM. Note that although it would be possible to use logistic regression to analyse individual trial data, this is not done in the present experiments, since item-level differences are controlled only between, rather than within, participants.

There is some ambiguity concerning the subsequent memory effects in FFA according to the simple resources account. Freeing of resources may happen by virtue of priming, as forms the basis of the resources account of priming; or due to variability in the stimuli (i.e. some faces requiring less processing effort than others). In these cases, improved memory would tend to coincide with reduced activity in FFA. In other words, as well as reduced FFA activation for primed trials, we might expect reduced FFA activity for subsequently remembered trials (i.e. main effects of priming and memory). However, it could be that greater univariate activation of FFA helps memory in other ways, e.g., owing to some trials receiving greater overall attention for some reason, or involving more difficult stimuli to process, which are associated with better subsequent memory (and which potentially overcomes any reduction in resources freed by priming).

For the PPA ROI, the simple resource account predictions seem relatively straightforward. As mentioned above, an 'item-context trade-off' should see resources freed from the processing of items (faces) in FFA, redirected for processing of contexts (scenes) in PPA, so reduced activation for primed trials in FFA should be accompanied by increases in PPA.

Like the FFA, the resources account prediction for the hippocampal ROI is also somewhat ambiguous. In typical studies of memory, hippocampal activity has been shown to predict successful subsequent recollection (e.g. see Uncapher & Rugg, 2005). However, Gagnepain et al. (2011) found that priming reversed this relationship: for unprimed items, activation in bilateral hippocampal ROIs was greater for subsequently remembered compared with familiar and missed trials, whereas for primed items, activation was reduced for subsequently remembered compared with familiar or missed trials. Gagnepain et al. hypothesise a freeing-up of resources, from auditory item processing, to auditory context processing, with increased context-processing making binding of item and context less effortful for the hippocampus. Whereas Gagnepain et al. used auditory stimuli for both item and context, making it difficult to infer processing of item from processing of context, in the present experiment, we can separate processing of item and context information by the stimulus-specific ROIs, i.e., FFA and PPA. Thus, it may be possible to revisit this interpretation of Gagnepain et al., asking whether priming-related decreases in FFA activation produce corresponding increases in PPA, and decreases in hippocampal activation.

Note that priming may free up more central resources (as suggested in Chapter 3), rather than perceptual resources. In Chapter 4, it was shown that while the results of the behavioural experiments could not rule out perceptual resource freeing as a mechanism of PSM, the interaction pattern found when analysing PSM with increasing central load was at least suggestive of a role of central attention in generating the effects. While the intention was not to distinguish between perceptual and central accounts, an intuitive notion of reduced Face Task difficulty relates this to some sort of perceptual processing, even if this is itself the result of reduced higher-level attentional demands. It is possible that reduced central demands still cause reduced activation in perceptual regions: in this case PSM effects should be found in higher-order regions (e.g. fronto-parietal executive regions, perhaps part of the Multiple Demand Network, Duncan, 2010). To test this idea a whole-brain voxel-wise search for PSM effects was also performed.

Predictions of PIMMS Account

The PIMMS account of PSM states that, when there is an uninformative prior regarding the probability of a specific face given a specific scene, then more precise sensory evidence (owing to priming) results in a greater PE (between scene and face) and hence better encoding of the scene-face association (Gagnepain & Henson, 2010; Greve et al, 2017). Thus, if FFA activation reflects this type of prediction error (PE), then priming should increase FFA activation.

However, FFA activation is normally decreased, rather than increased, by repetition (RS), and the conventional predictive coding account of this is reduced PE owing to synaptic changes that improve predictions between layers of a perceptual processing hierarchy (Henson, 2003). In other words, improved perceptual predictions about face features from higher-level visual regions (such as anterior temporal lobe) would reduce the PE in FFA, in addition to possible reductions in input (PE) from lower-level visual regions. Thus FFA is likely to be influenced not only by associative predictions from other stimuli (such as scenes) but also perceptual predictions about the face stimuli themselves. This means that the univariate predictions of PIMMS for FFA are unclear.

A clearer prediction of PIMMS is that priming should sharpen the representation of a face in FFA (i.e., sharpen sensory evidence), which should appear as increased dissimilarity between faces when they have been primed, relative to unprimed. This can be tested by correlating the pattern across voxels for individual trials (i.e., RSA) and testing whether the mean similarity is lower for primed vs unprimed trials. Most importantly, PIMMS predicts that the difference in similarity (degree of sharpening) should correlate, across participants, with PSM.

A second prediction of PIMMS is that improved memory should be associated with increased interactions between brain regions, such as PPA and FFA for learning a scene-face pairing (and/or possibly between hippocampus and PPA/FFA). These interactions would correspond to increased functional connectivity between fMRI data. As mentioned above, Gagnepain et al (2011) used Dynamic Causal Modelling (DCM; Friston, Harrison, & Penny, 2003) and found greater effective connectivity between

sensory cortex (in this case superior temporal gyrus, STG) and hippocampus for subsequently remembered items, regardless of univariate activation. The authors discuss this in general terms as supporting the idea that connectivity between medial temporal lobe (MTL) regions and neocortex may be more important for subsequent memory than simple activation. Given uncertainty about DCM for event-related designs (Henson et al., 2012), a simpler form of connectivity was used, based on single-trial estimates (Beta-series regression, Rissman et al., 2004) to test the PIMMS prediction that connectivity between FFA, PPA and hippocampus is greater for subsequently remembered versus subsequently forgotten trials.

Summary

As reviewed above, the resources account makes the specific univariate prediction that FFA activation for primed trials will be reduced compared with unprimed trials, and that this reduced activation will correspond to increases in activation in PPA. Such increased activation in PPA (reflecting more robust scene processing) may reduce the workload of hippocampus for binding face and scene stimuli in memory, thereby predicting a corresponding reduction in hippocampal activation. These effects could be expected to correlate with PSM across participants. Predictions of memory effects in FFA and of priming effects in hippocampus are unclear.

Whereas the resources account makes univariate predictions, PIMMS predictions are mainly at the multivariate level. The main prediction made by PIMMS is that previous exposure to faces should render the neural response (particularly in FFA) more distinctive, reducing within-condition correlations among primed trials, compared with unprimed trials. This effect would be expected to correlate with PSM. A further prediction of PIMMS is that subsequent memory improvement for primed trials will be brought about by increased connectivity between the ROIs, as increased prediction error is resolved.

Method

Participants and Materials

24 healthy volunteers, (12 females), aged 19-36 ($M = 26.0$, $SD = 4.5$) years, were recruited from the local population, and paid £10/hr for their time, as agreed by Cambridge Psychological Research Committee (PRE.2016.055). A reduced set of face and scene stimuli was created, comprising 144 of the original 192 unique scene and face stimuli used in the previous priming experiments, with the scenes removed from the main experiment used instead for the (PPA) functional localiser. The faces used for the FFA localiser came from 3 sets: NimStim (Tottenham et al., 2009), Karolinska Directed Emotional Faces (KDEF) (Lundqvist and Litton, 1998) and FERET (Phillips et al., 2000).

Adaptation of Behavioural Paradigm for fMRI

The basic design most closely resembled the no load condition of Experiment 2, and was adapted for fMRI, and behaviourally piloted on a small sample ($N = 3$), all of whom showed the expected direction of PSM. The main change made to the experimental paradigm affected the study phase: in order to be able to analyse single-trial data, it was necessary to lengthen Study trials to approx. 8 s, and we therefore inserted a 5 s period immediately after the face judgement. To prevent participants from mind-wandering or going to sleep during this period, they were required to make odd/even judgments to 2-digit numbers as they appeared. This digit task in each study phase trial was self-paced, with a new number appearing onscreen as soon as participants pressed a key (note that this differs from the 16 s digit task which was employed as a distractor between study and test phases in all PSM experiments reported in this thesis, and modelled separately in the fMRI analysis). This additional task also served as an active baseline, since previous studies have shown that hippocampus tends to be more active during rest than during a demanding, non-mnemonic task (Stark & Squire, 2001). The insertion of this additional task into Study Phase trials lengthened blocks by ~1 min, and therefore to keep the whole experiment to a manageable length for participants, the overall number of training-study-test blocks was reduced from 16 to 12. This meant that participants would be exposed to fewer stimuli overall, although the absence of a load manipulation

meant that in this experiment there would be an increase in the number of no-load trials relative to previous experiments. The 12 blocks were split into 4 runs of ~15 min each (i.e., 3 blocks per run).

The Training phase was kept identical to Experiments 2, 3 and 4b. Beyond the changes discussed above to the Study phase, there were two minor changes: first, the Test Phase became time-limited to 5 s rather than self-paced, where the 5s was chosen to capture 95% of 3AFC response times based on the prior behavioural experiments. Secondly, all onscreen instructions, which appeared before the start of each task in every block, were set to appear for a fixed period of 16 s.

To enable demarcation of face- and scene-relevant ROIs, an 8 min 'functional localiser' scanning run was appended at the end of the 4 'main experiment' runs. Participants were instructed to perform a '1-back' task, and were shown 'place', 'face', and 'scrambled' images in 16 same-category blocks comprised of 8 trials, pressing a button whenever a stimulus was repeated (repeats occurred 2 ± 1 times each block). Stimuli were onscreen for 1600 ms, followed by 400 ms ISI with a fixation cross. Between blocks, an 8 s filler screen presented the text 'please wait'.

fMRI Acquisition and Preprocessing

The MRI data were collected using a Siemens 3 T TIM TRIO system (Siemens, Erlangen, Germany). MRI data preprocessing and univariate analysis used the SPM12 software (Wellcome Department of Imaging Neuroscience, London, UK, www.fil.ion.ucl.ac.uk/spm), implemented in the AA 5.0 batching software (<https://github.com/rhodricusack/automaticanalysis>).

The functional images were acquired using T2*-weighted data from a Gradient-Echo Echo-Planar Imaging (EPI) sequence with a multi-band (factor = 4) acquisition. A mean of 736 (SD = 4.3) volumes were acquired in each of 4 runs of the main experiment, plus $M = 499$, $SD = 4.0$ volumes from the fifth localiser run. Each volume contained 64 slices (acquired in interleaved order within each excitation band), with a slice thickness of 2 mm with no interslice gap (for whole brain coverage; TR = 1.19 s; TE = 30 ms; flip angle = 74 degrees; FOV = 192 mm × 192 mm; voxel-size = 2 mm × 2 mm × 2 mm).

A structural image was also acquired with a T₁-weighted 3D Magnetization Prepared Rapid Gradient Echo (MPRAGE) sequence (repetition time (TR) 2250 ms, echo time (TE) 3.02 ms, inversion time (TI) 900 ms, 230 Hz per pixel; flip angle 9 deg; field of view (FOV) 256 x 240 x 192 mm; GRAPPA acceleration factor 2).

The structural images were rigid-body registered with an MNI template brain, bias-corrected, segmented and warped to match a sample template using DARTEL (Ashburner, 2007). This template was subsequently affine-transformed to standard Montreal Neurological Institute (MNI) space. The functional images were then spatially realigned across and within the 5 runs, interpolated in time to correct for the different slice acquisition times, rigid-body coregistered to the structural image and then transformed to MNI space using the warps and affine transforms from the structural image, and resliced to 2x2x2mm voxels. For the whole-brain analysis and functional localisation of ROIs, these normalised images were then smoothed by 6mm FWHM. For the subsequent ROI analysis, this smoothing was omitted.

Analysis

First-level

A General Linear Model (GLM) was constructed for each participant for the main experimental runs, consisting of 18 neural components per run: one for each trial-type. Trial types were: Training (instances 1, 2, and 3), Study (Unprimed Subsequently Remembered, Unprimed Subsequently Forgotten, Primed Subsequently Remembered and Primed Subsequently Forgotten), and Test (Unprimed Remembered, Unprimed Forgotten, Primed Remembered and Primed Forgotten). These trials were modelled as events, and the digit task and instructions were each modelled as epochs of 16 s. Each neural component was convolved with a canonical haemodynamic response function (HRF) to create a regressor in the GLM. Six additional regressors representing the 3 rigid body translations and rotations estimated in the realignment stage were included to capture residual movement-related artefacts. Finally, the data were scaled to a grand mean of 100 over all voxels and scans within a session.

The GLM was fit to the data in each voxel. The autocorrelation of the error was estimated using an AR(1)-plus-white-noise model, together with a set of cosines that functioned to highpass the model and data to 1/128 Hz, fit using Restricted Maximum Likelihood (ReML) as described in Friston et al. (2002). The estimated error autocorrelation was then used to “prewhiten” the model and data, and ordinary least squares used to estimate the model parameters.

Second level

To compute subsequent memory and priming effects, the parameter estimates for the 11 trial-types of interest were averaged across the 4 runs, and then subjected to the two types of analysis below.

Region of Interest (ROI) analysis

We defined the FFA and PPA functionally, based on group-level results from the localiser run. The first-level models consisted of 3 regressors for ‘faces’, ‘places’ and ‘scrambled’ blocks, modelled as 16 s epochs convolved with the canonical HRF. Contrast images for faces versus scrambled, and of places versus scrambled, were taken into one-sample T-test second-level models across participants. Functional localiser data from one participant was accidentally lost, so group-level ($N = 23$) ROIs were established and used, rather than discarding otherwise-informative experimental data from that participant.

In order to obtain suitable ROIs for subsequent analysis, different methods were used for FFA and PPA: for FFA, alpha was first reduced to family-wise-error (FWE) corrected $p < .001$. Bilateral clusters were then chosen on the lateral fusiform gyrus, ignoring clusters in hippocampus and occipital cortex, yielding a mask consisting of 258 voxels. For the functional definition of PPA, given widespread activity surviving FWE correction alpha $p < .001$, instead a t -statistic threshold of $T(46) > 10$ was used, choosing the surviving bilateral clusters on the parahippocampal gyrus, and yielding a mask comprising 590 voxels. For the hippocampal ROI, an anatomically-defined mask was used, obtained from The LONI Probabilistic Brain Atlas (Shattuck et al., 2008), and containing 2365 voxels. Figure 23 below shows locations and extents of these 3 ROIs.

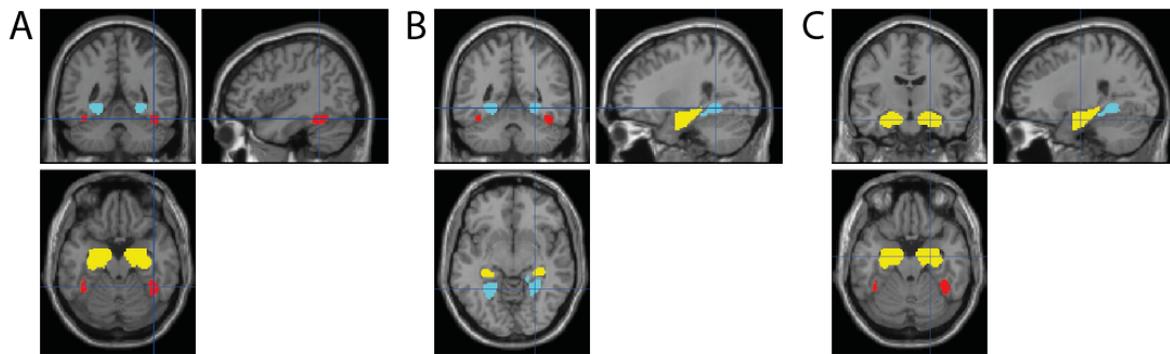


Figure 23 ROI masks used in univariate and multivariate analysis. FFA (red blobs, highlighted in Panel A), PPA (cyan blobs, highlighted in Panel B), and hippocampus (yellow blobs, highlighted in Panel C).

For each ROI, a single fMRI timecourse was defined by taking the first temporal component of a singular-value decomposition of the ROI data in that participant. These data were then fit, yielding betas for each subject for our conditions of interest (the 3 Training Phase conditions and the 4 Study and 4 Test conditions). These values were then compared in 3 separate repeated-measures ANOVAs, one for each of the 3 phases of the experiment.

Whole brain analysis

To check for any effects outside the ROIs above, contrast images for each trial-type and participant were then entered into a second GLM corresponding to a repeated-measures ANOVA, which treated participants as a random effect. Within this model, Statistical Parametric Maps (SPMs) were created of the T-statistic for the various effects of interest, using a single pooled error estimate for all contrasts, whose nonsphericity was estimated using ReML. The SPMs were thresholded at $p < .001$ uncorrected, and clusters of >50 voxels whose peak or extent survived $p_{fwe} < .05$ corrected using Random Field Theory (RFT) were reported.

Representational Similarity Analysis (RSA)

RSA of each ROI was done for the 4 Study Phase conditions (Unprimed Subsequently Remembered, Unprimed Subsequently Forgotten, Primed Subsequently Remembered and Primed Subsequently Forgotten). The timecourses for each voxel were fit using

Least Squares Single (LSS, see e.g. Abdulrahman & Henson, 2016), in which a new GLM was estimated for each trial, with all other trials modelled by one regressor per condition.

Patterns of activity across voxels in each ROI were correlated between every pair of trials, and mean correlation coefficients (Pearson's r) across trials within each condition were Fisher-transformed and entered into a repeated measures ANOVA with 2 factors, Priming (P,U), and Subsequent Memory (SR, SF). Since such correlations would be expected to decrease as the signal-to-noise ratio (SNR) decreases (assuming reduced signal but constant, additive noise), the mean activation for each condition and participant was included as a covariate in a second ANCOVA, to control for such SNR differences.

Beta-series Connectivity ROI analysis

Finally, a Beta-series connectivity analysis was conducted to address whether the correlation of individual trial estimates between the ROIs during Study phase differed as a function of condition. For this analysis, LSS was again used to estimate single-trial responses, but now on the SVD timecourse for each ROI (rather than each voxel separately). Fisher-transformed Pearson's correlation coefficients across the Beta-series for each pair of ROIs were entered into the same 2x2 ANOVA as above.

Note that Dynamic Causal Modelling (DCM) of the raw timecourses to estimate effective connectivity between ROIs, as was done by Gagnepain et al (2011), was not performed here. This is because recent research in my group (by Tibon and Henson together with Zeidman in London), as well as prior work (Henson et al., 2012) suggests that DCM for event-related responses is very unstable.

Results

Behavioural Results

Participant 15 was excluded from behavioural and imaging analyses because of outlying number of missed study phase responses (69 of 144 trials) ($M = 19.5$, $SD = 16.9$).

The linear trend analysis for the Training Phase data did not show a significant decrease of RTs with repetition during the Training phase, $F(1.42, 31.20) = 2.02$, $p = .080$, 1-tailed. However, given previous results in Experiments 1-4, a post-hoc T-test was carried out, which showed a significant decrease in RTs from 2nd instance ($M = 791$ ms, $SD = 110$ ms) to 3rd instance ($M = 761$ ms, $SD = 110$ ms), $T(22) = 2.81$, $p = .010$.

A T-test of Study phase RTs showed a significant effect of priming, $T(22) = 5.76$, $p < .001$. As expected, responses to primed faces ($M = 800$ ms, $SD = 76.2$ ms) were faster than to unprimed faces ($M = 843$ ms, $SD = 94.7$ ms).

At the Test phase, the T-test showed no significant effect on memory of priming the face images, $T(22) = 0.33$, $p = .747$. Contrary to expectations (and all four previous behavioural experiments), memory for scene-face pairings with primed faces ($M = 62$ %, $SD = 14.5$ %) was not significantly different than with unprimed faces ($M = 62$ %, $SD = 13.7$ %). The correlation between PRT and PSM was not significant.

Univariate ROI Results

The results for all three phases and all three ROIs are shown in Figure 24, below. The main interest was in the Study phase, but results from the Training and Test phases are shown for completeness. A number of participants showed outlying univariate activations in one or more ROIs in one or more conditions, based on their data falling outside the whiskers ($1 \times IQR$) of the respective boxplots (see Figure 24). None of these participants behaved unusually, and therefore where relevant, results of analyses are first reported including these participants, then excluding them.

Training Phase

For Training Phase analysis, decreasing responses to repeating presentations were hypothesised in the FFA, and there were not strong hypotheses regarding PPA and hippocampus. The FFA did indeed show a significant linear decrease, $T(45) = 2.81$, $p = .002$, 1-tailed. With outlier participant data excluded, this result remained, $T(43) = 3.58$, $p < .001$. With all data included, the PPA showed a nonsignificant linear increase, $T(45) = 1.61$, $p = .115$, but with the 3 outlying participants' data removed, this effect became significant, $T(39) = 2.49$, $p = .016$. The hippocampus showed no significant effect, $T(45) = 1.43$, $p = .159$, even with outliers excluded, $T(45) = 0.88$, $p = .380$.

Study Phase

FFA

The 2x2 ANOVA for the FFA during Study showed, the main effect of Priming was in the expected decreased response for primed relative to unprimed trials, though only trending towards significant according to a one-tailed test ($F(1,22) = 2.35$, $p = .070$). With outlier values removed, this became more significant ($F(1,22) = 2.92$, $p = .050$, 1-tailed). There was also a main effect of subsequent memory, $F(1,22) = 10.33$, $p = .004$, with greater activity when a pairing was later remembered than later forgotten. Outlier removal made little difference to this result ($F(1,16) = 11.80$, $p = .003$). There was no interaction, neither with all data included, $F(1,22) = 0.22$, $p = .647$, nor with outliers excluded, ($F(1,16) = 1.125$, $p = .305$).

PPA

For the PPA, there was a significant effect of priming, $F(1, 22) = 10.847$, $p = .002$, 1-tailed, in the opposite direction to FFA, whereby activation was greater when a face had been primed than when the face was unprimed. There was also a main effect of subsequent memory, $F(1, 22) = 16.05$, $p = .001$, this time in the same direction as the FFA. There was no interaction, $F(1, 22) = 0.13$, $p = .719$.

Hippocampus

For the hippocampus, there were no significant effects of priming, $F(1, 22) = 0.224, p = .321$, and nor did one appear when the 2 outliers were excluded, ($F(1,20) = 0.024, p = .878$), but like FFA and PPA, activation was greater when pairings were subsequently remembered than forgotten, $F(1, 22) = 12.41, p = .002$, with this effect remaining reliable when data were reanalysed without outlying values ($F(1,20) = 10.41, p = .004$). Again, there was no interaction, $F(1, 22) = 0.13, p = .717$, and nor did one appear when outliers were removed ($F(1,20) = 0.214, p = .649$).

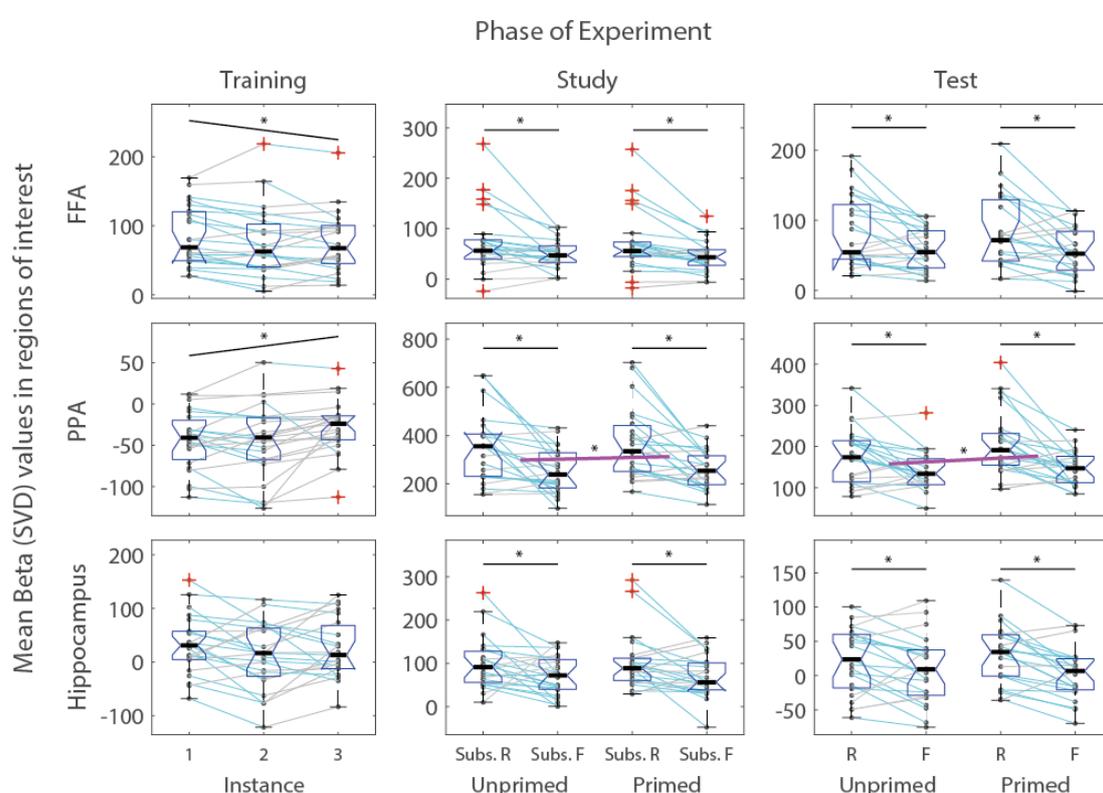


Figure 24: Results of ROI-based analysis. In Study and Test phases, black lines with asterisks indicate significant main effects of Subsequent Memory; magenta lines with asterisks indicate significant main effects of Priming.

Correlation analysis

According to the resources account, the size of the priming-related decreases in FFA (reflecting resources that are freed-up) should be related to the size of the priming-related increases in brain regions associated with subsequent memory, e.g. PPA and/or hippocampus (i.e. negative relationship between priming effects in FFA and priming

effects in other ROIs). Activation difference was therefore correlated between primed and unprimed trials across participants for FFA and PPA, and for FFA and hippocampus (Figure 25). Contrary to expectations, in both cases priming effects were strongly positively correlated: a positive correlation was found between FFA and PPA priming effects, Pearson's $r = .704$, $p < .001$, and a positive correlation was also found between FFA and hippocampal univariate priming effects, Pearson's $r = .650$, $p = .001$.

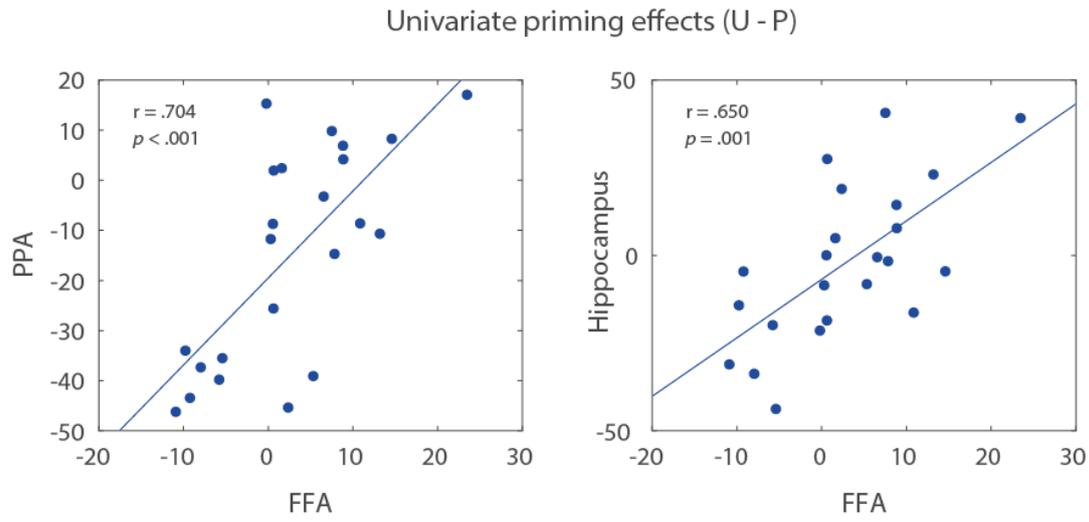


Figure 25: Scatter plots show relationship between univariate suppression effects of priming in N=23 participants in FFA and PPA ROIs, and between FFA and hippocampus ROIs.

A further prediction of the resources account was that effects of priming on univariate activation should correlate with subsequent memory advantage for primed trials (PSM), most particularly in the FFA. The correlations were therefore analysed between effects of priming on univariate activation, and PSM (Figure 26, below) for each of the 3 ROIs. Visual inspection of the scatterplot showed 1 obvious outlier. These data were therefore analysed with and without this participant. When including this participant, no correlations were significant (see red text in figure). However, when this participant was excluded from the correlation analysis, in line with the prediction of the simple resources account, it was found that priming-related reduction in FFA activation was positively correlated with PSM, Pearson's $r = .42$, $p = .027$, 1-tailed. However, rather than the predicted *increased* activity in PPA correlating with PSM, reflecting freed up resources from FFA, the correlation analysis showed that *reduced* PPA activation for

primed trials correlated with PSM, Pearson's $r = .59$, $p = .004$. No relationship was found between priming effects in the hippocampus and PSM.

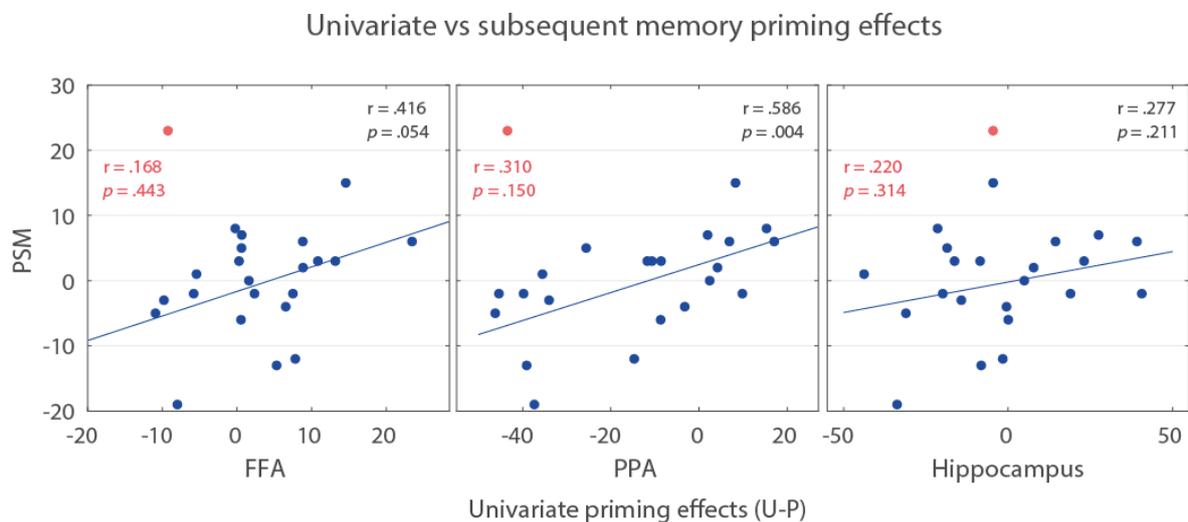


Figure 26: Scatter plots show relationship between behavioural (memory) and neural (univariate) effects of priming in the ROIs. Text in red shows results including PSM behavioural outlier (red marker). Black text and least squares fit line do not include the behavioural outlier.

Test Phase

In the FFA, there was no significant main effect of priming during Test phase, $F(1, 22) < 0.01$, $p = .952$. There was a significant memory effect, $F(1, 22) = 12.00$, $p = .002$, 2-tailed, with greater activity for remembered trials. There was no interaction, $F(1, 22) = 0.26$, $p = .614$.

In the PPA, there was a significant main effect of priming, $F(1, 22) = 22.92$, $p < .001$, with greater activity for primed items than for unprimed items, and this was unaffected by the removal of the 2 outliers, ($F(1,20) = 19.05$, $p < .001$). There was also a main effect of memory, $F(1, 22) = 14.28$, $p = .001$, with greater activity for remembered items, with outlier removal producing little change in the result ($F(1,20) = 11.81$, $p = .003$). This effect was qualified by an interaction, $F(1, 22) = 4.74$, $p = .040$, whereby the memory activation was greater for primed than unprimed trials, but this interaction was no longer significant when outlying data were excluded from analysis ($F(1,20) = 2.61$, $p = .122$).

In the hippocampus, the main effect of priming was not significant, $F(1, 22) = 1.70$, $p = .206$, but the main effect of memory was significant, $F(1, 22) = 9.53$, $p = .005$, in the same direction as FFA and PPA, i.e., greater activity for remembered trials. There was no interaction, $F(1, 22) = 4.17$, $p = .053$.

Whole-Brain Analysis

Study Phase neural responses were explored across the whole brain. There were no clusters showing effects of priming, whether enhancement or suppression. There was 1 cluster ($k = 255$) which showed greater activation in Middle Occipital Gyrus for subsequently remembered trials, $p_{fwe} = .017$, peak-level $T = 4.13$. There were no clusters showing interactions between either priming enhancement and subsequent memory, nor priming suppression and subsequent memory.

Representational Similarity Analysis

Next, the PIMMS prediction that priming of faces reduces the similarity between their representations in FFA (by sharpening them) was tested. This analysis was restricted to the main Study phase of interest. As with the univariate analyses above, results are presented both with and without outlying data (Figure 27).

FFA

In the FFA, there was a significant main effect of Priming, $F(1,22) = 13.53$, $p < .001$, 1-tailed, with reduced similarity for Primed trials ($M = .245$, $SD = .155$) relative to Unprimed trials ($M = .260$, $SD = .149$). This priming-related decrease remained after covarying out univariate effects for each participant (i.e., the mean across voxels), $T(65) = 2.58$, $p = .003$, 1-tailed, suggesting that the decrease in Pearson correlation did not owe simply to reduced SNR with priming. Exclusion of 3 outlying participants' data had only a small effect on the result, $F(1,19) = 15.23$, $p < .001$, 1-tailed $.356$, which again remained reliable when covarying out univariate effects, $T(56) = 2.71$, $p = .004$, 1-tailed.

There was no main effect of memory, $F(1,22) = 0.89$, $p = .356$, on pattern similarities, and this was unchanged by the exclusion of the 3 outlying participants $F(1,19) = 1.28$, $p = .102$. There was no interaction, $F(1,22) = 0.231$, $p = .635$, and exclusion of outliers did not change this, $F(1,19) = 0.076$, $p = .785$.

PPA

In the PPA, there was a significant main effect of Priming, $F(1,22) = 4.35$, $p = .049$, with greater similarity for primed than unprimed trials, though this effect did not remain after correcting for the increased univariate activation for primed trials, $T(65) = 1.41$, $p = .081$, 1-tailed. After excluding data from 1 participant with outlying results, this effect became less reliable ($F(1,21) = 3.770$, $p = .066$), and covarying out univariate effects reduced reliability further ($T(62) = 1.29$, $p = .100$, 1-tailed).

There was also a main effect of subsequent memory on pattern similarity, $F(1,22) = 10.05$, $p = .004$, with greater similarity for subsequently remembered than subsequently forgotten trials. However, as with the priming effect, this effect was no longer reliable when adjusting for the fact that the mean signal was also higher for remembered trials, $T(65) = 0.10$, $p = .460$, 1-tailed. When 1 participant's data were excluded because outlying, this general result remained the same, $F(1,21) = 8.332$, $p = .009$, with covarying out univariate activity removing this apparent main effect, $T(62) = 0.209$, $p = .418$, 1-tailed.

Although the interaction was a non-significant trend without correction for univariate activation, $F(1,22) = 3.21$, $p = .087$, when correcting for univariate activation, this became a significant interaction, $T(65) = 2.44$, $p = .009$. Post-hoc t-tests showed that pattern similarity was higher only in the unprimed condition for subsequently remembered trials ($M = .57$, $SD = .196$) compared with subsequently forgotten ($M = .52$, $SD = .192$), $T(22) = 2.83$, $p = .010$. With the outlying participant removed, the interaction was not significant ($F(1,21) = 2.93$, $p = .102$), until the univariate activations were covaried out, when the interaction became significant, $T(62) = 2.358$, $p = .011$, 1-tailed. Post-hoc t-tests showed that pattern similarity, in the unprimed condition only, was higher for remembered ($M = .55$, $SD = .175$) than for forgotten ($M = .50$, $SD = .176$) trials, $T(21) = 2.60$, $p = .017$.

Hippocampus

With all data included, the hippocampus showed no significant effect of Priming, $F(1,22) = 1.12$, $p = .301$, nor of subsequent memory, $F(1,22) = 1.62$, $p = .217$. However, there was a significant interaction between these factors, $F(1,22) = 4.44$, $p = .047$, and this was

reinforced when correcting for univariate activation. Post-hoc t-tests showed, similar to the PPA, increased pattern similarity for unprimed remembered ($M = .13$, $SD = .070$), compared with unprimed forgotten ($M = .12$, $SD = .060$) trials, $T(22) = 2.21$, $p = .038$, but no difference between primed remembered and primed forgotten trials, $T(22) = 0.51$, $p = .616$.

Removal of 2 participants' data affected the RSA results in the hippocampus. Priming significantly reduced pattern similarity in the hippocampus, $F(1,19) = 6.71$, $p = .018$, and when univariate activity was covaried out, this effect remained, $T(56) = 2.07$, $p = .021$, 1-tailed. There remained no effect on subsequent memory of pattern similarity, $F(1,19) = .597$, $p = .032$, and the interaction ceased to be significant, $F(1,20) = 2.31$, $p = .144$, with this remaining true when univariate activation was covaried out, $T(59) = 1.52$, $p = .067$, 1-tailed.

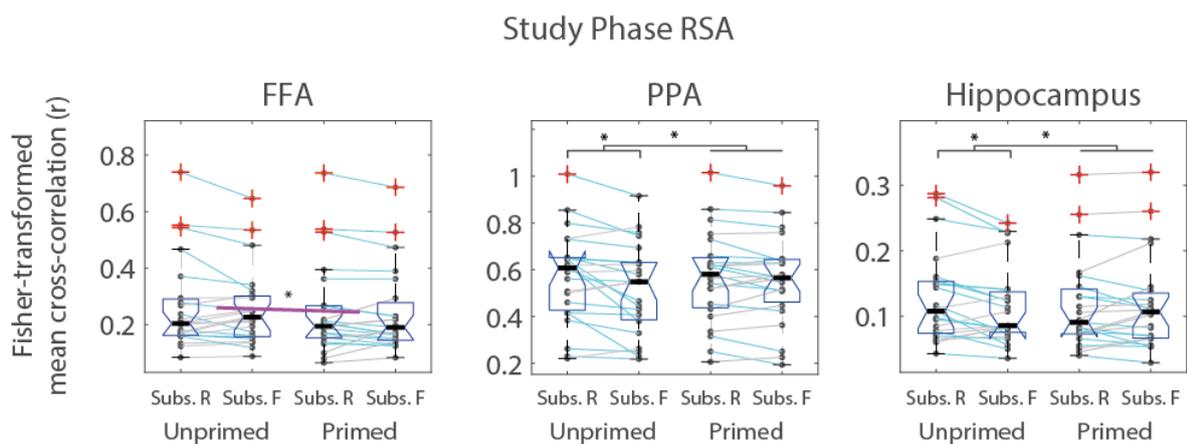


Figure 27: results of RSA analysis, showing reduced correlation between patterns in FFA when faces had been primed. (Note that Fisher-transformed Pearson's r may take on values exceeding 1.)

Correlations between RSA measures and behaviour

PIMMS predicts that the behavioural advantage of priming comes from more distinctive neural responses to primed faces, giving rise to increased PE. Therefore, correlation was tested between the significant effect of priming on pattern similarity, and PSM (Figure 28). There was no evidence of any relationship between these measures, Pearson's $r = -.238$, $p = .274$.

Priming effects on behaviour (PSM) and on trial pattern similarity

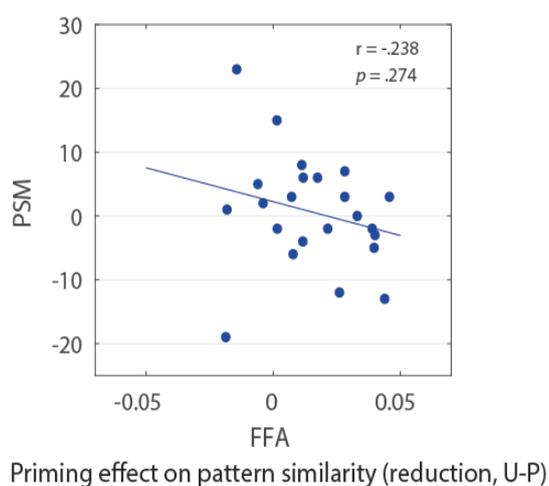


Figure 28: Scatter plot of priming effects on pattern similarity and subsequent memory across (N=23) participants in FFA ROI.

Beta-series Connectivity Analysis

Finally, trial-based connectivity between the 3 ROIs was compared across participants, addressing whether connectivity during Study Phase between FFA and PPA, FFA and hippocampus, and PPA and hippocampus, differed as a function of Priming and Subsequent Memory.

FFA – PPA

There was no evidence of significant functional connectivity differences between FFA and PPA depending on Priming, $F(1,22) = 0.43$, $p = .520$, nor on whether a scene-face pairing would be later recalled, $F(1,22) = 1.37$, $p = .254$, and nor was there any interaction, $F(1,22) = 0.56$, $p = .464$.

FFA – Hippocampus

There was also no evidence of any significant functional connectivity differences between FFA and hippocampus depending on Priming, $F(1,22) = 0.021$, $p = .886$, nor on whether a pairing would be later recalled, $F(1,22) = 0.04$, $p = .846$, nor was there any interaction, $F(1,22) = 1.82$, $p = .191$.

PPA – Hippocampus

There was no evidence of any significant functional connectivity differences between PPA and the hippocampus, depending on Priming, $F(1,22) = 0.06$, $p = .805$, nor on whether a pairing would be later recalled, $F(1,22) = 1.37$, $p = .254$, and nor was there any interaction, $F(1,22) = 0.35$, $p = .562$.

Discussion

fMRI was used with a modified version of the PSM experiments described and modelled in Chapters 2 – 4 to address the two main accounts of PSM at a neural level. Analyses were concentrated on Study Phase data. The simple resources account, modelled in Ch. 4, made univariate predictions, but no clear multivariate predictions. Conversely, the PIMMS account made no clear univariate predictions, but predictions for representational similarity analysis (RSA) and functional connectivity.

Summary of Results

Behavioural measures (RTs) during Training and Study Phase were as expected, with speeding of responses over successive presentations during Training (apart from a blip for second presentations for some unknown reason), and a very reliable effect of priming on RTs at Study. However, at test there was no significant PSM, which was surprising (given that it was demonstrated in four experiments in Chapters 2-3, plus a fifth experiment in Greve et al, 2017). The main univariate result during the Study Phase, with outlying data removed, were: 1) a reduction in activation for primed trials in FFA, though this was weak, 2) the opposite pattern of increased activation for primed trials in PPA, which was very robust (the hippocampus showed no significant effect of priming). The FFA and PPA results are in line with the resources prediction of a trade-off between processing demands in item- and context-processing regions. However,

contrary to this prediction of the resources account of a negative correlation between priming effects in FFA and the other ROIs, priming effects in the FFA were actually *positively* correlated with those in both PPA and the hippocampus.

There were significant subsequent memory effects in all three ROIs, with each showing greater univariate activation during subsequently remembered trials. This is contrary to what might be expected if fewer resources are needed for remembered trials, though would be consistent with random fluctuations across trials in other factors that increase FFA activity and improve memory, such as attention. Correlation analysis across individuals (with a visually identified outlier removed), indicated, in line with the resources hypothesis, that participants whose FFA activity was most affected (reduced) for primed trials, also showed greatest priming advantage for subsequent memory. Again, however, this relationship also held for activation in PPA, where the resources account predicts the converse: If resources were being 'freed' from the FFA to the PPA to give a priming advantage for subsequent memory, then participants with greater 'freeing' (i.e. priming reduction) in FFA ought to have shown correspondingly greater *enhancement* for primed items in PPA or hippocampus.

Whole brain analysis did not reveal any results of relevance to the theoretical accounts this thesis attempts to address: for example, neither priming nor subsequent memory effects were seen in other regions which might be associated with central attention (such as the Multiple Demand Network, Duncan, 2010).

The RSA results showed that responses to faces in FFA were made less similar to one another by priming. This supports the PIMMS prediction that priming faces leads to more distinctive responses in a face-processing area, and as a result producing greater PE and improved subsequent memory. However, it was further predicted that this effect of priming on distinctiveness of neural responses should correlate across participants with PSM, yet no apparent relationship between effects of priming on FFA pattern similarity and subsequent memory was found.

RSA in PPA showed that in the unprimed condition, pattern similarity was higher among subsequently remembered than subsequently forgotten trials. A comparable pattern was seen in the hippocampal ROI, with subsequently remembered trials more similar to each other than subsequently forgotten trials. However, when 2 participants with outlying data were removed from the hippocampus RSA, a different pattern emerged, and, when averaging across memory conditions, pattern similarity was decreased by priming. This result is again compatible with the PIMMS explanation of priming producing more distinctive neural responses, in line with a more detailed model-update in a region encoding the scene-face encounter. However, in exploratory correlation analysis, as with FFA, priming pattern similarity effects at the participant level were not found to be related to memory effects.

Increased iterative processing between regions associated with a bigger model update predicts greater connectivity would be found for subsequently remembered trials potentially between PPA and FFA, or potentially either or both of these regions and the hippocampus. The DCM results of Gagnepain et al. showed increased connectivity between hippocampus and STG for subsequently remembered trials, congruent with the PIMMS account. However, across all combinations of the ROIs, beta-series connectivity analysis did not show any connectivity differences produced by priming, nor reflected in subsequent memory. This could simply reflect low power for beta-series regression in this design.

Univariate Results and Resource Accounts

The univariate results do not strongly support the resources account, at least in its simple form. There was a weak reduction in univariate activity which was predicted by the idea of less-effortful processing of faces during primed trials. However, the negative correlation of priming effects across individuals between ROIs, which would be predicted if processes in PPA or the hippocampus had been competing for resources with FFA, was not found. Instead, results showed that priming effects on neural activity, where present in individual participants, were positively correlated between FFA and PPA, as well as between FFA and the hippocampus. This correlation should be treated with caution: overall activation across priming conditions in FFA and PPA might be

highly correlated, owing to global changes across the brain, so the difference, primed-unprimed, may also tend to be correlated for that reason.

An intriguing result was that rather than finding a reversal of the FFA effect, the relationship between PSM and univariate effects in PPA were in the same direction: reduced activation during primed trials in PPA predicted a subsequent memory advantage for primed trials. This provides some evidence against the idea proposed by Gagnepain et al. 2011 that increased context processing activity for primed items as a result of item-context trade-off meant that the hippocampus had to work less hard to encode primed-item trials.

Across participants, correlations between ROIs were different between remembered and forgotten conditions, and a regression analysis was therefore carried out, comparing the slopes of the FFA-PPA relationship for remembered trials, with that relationship for forgotten trials. The same analysis was also carried out for FFA-hippocampus. In both cases, slopes were reliably more positive across remembered trials, FFA – PPA, $t(22) = 5.43$, $p < .001$, and FFA – Hippocampus, $t(22) = 5.97$, $p < .001$, plausibly suggesting some differences in co-activation of these regions which related to subsequent memory. A simple explanation for this could again be related to global attention: during trials that are subsequently remembered, activity might be increased across all three ROIs, but during trials which are later forgotten, arbitrary metabolic fluctuations across the ROIs would tend to make activation less similar.

In sum, the pattern of the univariate results as a whole do not provide clear evidence to support the resources account. Although some of the pattern is congruent with that account, other results run contrary to what would be expected: no evidence was found for a trade-off between regions, and although the subsequent memory advantage for primed items at the participant level was found to be weakly related to reduction in univariate activity for primed items in FFA, the priming reduction in PPA activity was also found to predict subsequent memory advantage.

Multivariate Results and The PIMMS Account

Returning to the RSA results, again these are not conclusive, though some support is shown for an important aspect of the PIMMS account: even when adjusting for univariate differences in FFA (e.g., reduced responses for primed items, which may reduce SNR and hence mean correlation), correlations between trials were reduced by previous exposure to the faces, suggesting that neural responses to faces were indeed more distinctive during primed trials. In the PIMMS account, this should on the one hand increase activation both in this region and in PPA, because increased PE is entailed. In line with this, univariate effects of priming were indeed positively correlated across participants. Considered alongside the univariate reduction of activation in FFA, the reduction in pattern similarity for primed trials is a potential mechanism of resource freeing. However, follow-up analysis did not reveal any relationship between these univariate and RSA priming effects.

Contrary to what would be expected, no relationship was observed between effects of priming on pattern similarity and on subsequent memory, so it is not yet known whether pattern distinctiveness as measured by cross-correlations between trials of the same type was relevant to memory performance.

The subsequent memory effect shown strongly in PPA, and weakly in the hippocampus, of increased similarity between voxel patterns for subsequently remembered trials, runs counter to the idea of more distinctive representations being better remembered. A simple explanation for this (trial-level) effect is similar to that given above for the higher correlations at the participant level between ROI activation for remembered than forgotten trials: in trials where PPA or hippocampal activity was engaged, then patterns may be somewhat similar, whereas in trials when either region was not particularly task-engaged, random fluctuations would produce reduced correlations between trials.

Chapter Summary and Potential Caveats

The present fMRI study did not provide clear support for the PIMMS account of PSM, or for the simple resources account as modelled in Chapter 4. Several factors contribute to the ambiguity of the results: First and most obviously, there was no group-level PSM

(though correlation analysis of participant-level effects did reveal some behaviourally-relevant neural effects of priming). Second, the study was somewhat exploratory, and contains many analyses without correction for multiple comparisons. Confidence in the validity of its findings would be increased by running an experiment with a similar design, but limiting analyses to those from this experiment in which results were significant, and correcting α to reflect the (much reduced) number of comparisons.

The fact that no significant effects were revealed by the connectivity analysis might reflect lack of power, and while DCM may be unstable in rapid event-related designs like the present experiment (Henson et al., 2012), it could be that a psycho-physiological interactions (PPI, O'Reilly et al., 2012) analysis would be more sensitive to such effects. Sensitivity could potentially be slightly increased by defining ROIs separately for each participant: unfortunately, functional localiser data were accidentally lost for one participant, and the loss of power which would have resulted from excluding this participant's experimental data was not thought to be worth exchanging for a potential benefit of defining ROIs individually.

Most importantly, a future PSM fMRI experiment must establish subsequent memory effects of priming in the scanner. It was unexpected to find no overall PSM once the experiment was adapted for MRI. There is some irony in this: having searched for ways to abolish the effect in Chapters 2 and 3, the adjustment needed for measuring brain correlates of PSM may have abolished the effect! Note that the absence of PSM would not be explained by 'scanner attentional load', since mean performance at test does not differ from previous no-load conditions (whereas in all previous experiments, attentional load has diminished memory performance). However, an interesting possibility is that the immediate post-trial period (during which time participants performed an attentionally-engaging odd-or-even-number distractor task, introduced to give a better baseline for hippocampal activation) may have been important for some kind of encoding for which there was a specific priming advantage. Note that although Experiment 1 did control for time on task, as did all subsequent experiments, this typically only made a small difference (approx. 100 ms) to trial times. There was still

around 700 ms before next trial onset in the fMRI experiment, with no load, during which time such effects could conceivably have arisen. The possibility of such an immediate post-trial effect described above was not apparent in the experiments used in Chapters 2 and 3, in which Study Phase inter-trial intervals were much lower, but the design used in this fMRI experiment could potentially be adapted to examine whether inter-trial-delay attentional load has an effect on PSM. A simple way to do this would be to change the Study Phase inter-trial distractor task slightly. The present inter-trial distractor was relatively demanding, with the onscreen number replaced immediately on response, and participants working hard to get through as many 'mini-trials' as possible. If the interspersed number task was run at a more relaxed pace, with, for example, a 1.5 s SOA, then this may have reduced attentional demand. This could be tested in a factorial manner, within-subjects, in a pilot study. Regardless of any factorial manipulation, more-thorough behavioural piloting than used here will be necessary for a subsequent experiment, and may give a clue as to why no PSM was found in this study.

From the perspective of the resources account modelled in Chapter 4, the fMRI study presented here may lend some qualified support: reduction in FFA activity for primed faces was weakly correlated with PSM, suggestive of resource-freeing. However, one of the most salient findings of the modelling was that non-linearities in resource-performance functions make it hard to draw strong conclusions without a formal model. Thus, even though PSM correlates with priming-related reductions in both FFA and PPA, without knowing the underlying functions relating these activations to performance, Chapter 4 showed that one cannot be confident about the direction of these correlations. The second main finding was that distinctiveness of neural responses in both FFA and hippocampus ROIs was increased for primed items. This is in line with the predictions of PIMMS but does not contradict the simple resources account. The possibility of further neuroimaging work based on the current study is discussed in the next chapter, the General Discussion.

Chapter 6

General Discussion

Overview

This thesis has addressed two accounts of priming effects on subsequent (associative) memory (PSM): the increased-PE-driven memory update of item and context according to PIMMS, and the attentional resources account of a trade-off between processing of item and context.

Summary

Much of the thesis was concerned with testing the attentional resources account of PSM, based on the idea that resources could be freed up by priming (because processing of the primed item becomes less effortful). Through behavioural experiments examining the effect of priming faces on subsequent memory for face-scene associations, an attempt was made to adjudicate between several ways in which priming could free-up resources and hence facilitate memory encoding: i) increasing the time (temporal resource) available after a response has been made to encode the face-scene association, ii) reducing the perceptual load needed for processing faces and hence benefiting perceptual processing of the scene, and iii) reducing higher-order “central” load and hence benefiting (e.g. semantic) encoding of face-scene associations. The first, “time-on-task” account was not supported because PSM was not significantly affected when this “leisure” time was equated across primed (P) and unprimed (U) conditions. The perceptual and central load accounts were addressed by examining whether PSM was reduced by a simultaneous load, but any evidence for such a priming-by-load interaction was weak, only being observed when pooling conditions across two central load experiments. Importantly however, computational modelling showed that the absence of a priming-by-load interaction does not rule out a resources account: indeed, once a nonlinear (but monotonic) relationship between performance and resources is allowed, a resources account can predict any pattern of interaction with such load manipulations. Thus the behavioural data appear insufficient to rule out a perceptual or central attentional load account of the PSM, therefore failing to distinguish it from the PIMMS account. Given the insufficiency of the behavioural experiments, an fMRI experiment was conducted to investigate whether neural data were more consistent with the resources account or the PIMMS account. Unfortunately, aspects of both

accounts were supported: the mean activation in PPA increased with priming, supporting increased resources for perceptual processing of the scene, but priming also produced more distinct multivoxel patterns in FFA, supporting the PIMMS prediction of sharper sensory evidence. Thus overall, the results do not strongly favour one account for PSM over the other. Below, other possible experimental approaches are considered that might provide evidence to favour one account in future, as well as possible reconciliation of these theoretical viewpoints. First, though, the main results of the experiments and modelling are considered in more detail.

Main Results

Across Chapters 2 & 3, four behavioural experiments were used to address whether a resources-based account was sufficient to explain PSM. First, the simple temporal resources explanation for the results of Gagnepain et al. (2008), where faster responses for primed items would leave longer 'leisure time' before the next trial, was eliminated by simply fixing the time available. The following experiment investigated whether priming might cause a reduction in demands on the perceptual system, freeing up resources for encoding association. Reducing availability of such perceptual resources in the same way as Yi et al. (2004) did not have an impact on PSM. Therefore, work turned from addressing perceptual processes, to asking whether priming might free central processes: combining central load conditions from 2 experiments showed that central attentional tasks interfered with the benefits of priming, suggesting a common resource. However, even with a large sample size across experiments, no correlation was found between priming effects on RTs (PRT) and PSM across participants, which appeared difficult to reconcile with a common resource, which *prima facie* would appear to predict a positive relationship (the more resources freed by priming, the more available to encode associative memories). However, the computational model in Chapter 4 showed that neither the interaction pattern on mean performance nor the correlation across participants is uniquely predicted by a resources account.

The formal model of the resources account in Chapter 4 made clear that, given simple and plausible assumptions concerning the sigmoidal relationship between a

hypothesised resource and performance, a variety of interaction patterns, including no interaction, are all possible. Simulations using the model could reproduce the basic PSM, and the qualitative patterns of results seen in Experiments 2, 3, and 4. It was surprising to find that correlations between simulated PRT and PSM could range from strongly positive, through uncorrelated, to strongly negative, depending only on slight adjustments to model parameters. These all arose because of the simple nonlinearity in the resource-performance curve. A slight change to the allocation/priority parameter for a simulation run was sufficient to change the relationship between PRT and PSM dramatically because of the way that the gradients of the respective curves correspond to one another when resources are 'traded off'. In practice, allocation is likely to have differed considerably between participants, and there is no reason to expect the latent curves of the different tasks to be similar, whereas in the model they are identical, further complicating attempts to predict correlations between outcome measures.

The final experiment compared PIMMS and resource trade-off accounts at a neural level using fMRI. For reasons that are unclear, the behavioural PSM was not found in this experiment, despite the design being very similar to that used in Experiments 1-4. Nevertheless, there were effects of priming observed at the neural level. A simple resources account would appear to predict decreased FFA activity, with commensurable increase in PPA and/or Hippocampal activity for primed items (i.e., a negative correlation between effect of priming on FFA versus PPA/Hippocampus). However, although there was a trend for reduced FFA activation in primed trials, correlations between priming effects in FFA and PPA were positive, which is the reverse of the trade-off predicted by a resources account. A trend for a positive correlation was also found between effects of priming (U-P) on neural activity in FFA and effects on memory (PSM), perhaps suggestive of freed resources. However, the same positive relationship was also found between neural and behavioural effects of priming in PPA, which is the reverse of the resources prediction (i.e. that associative memory improvement is due to greater recruitment of PPA in primed trials, so PPA U-P would be expected to be more negative for participants who showed more PSM, the reverse of what was found).

The PIMMS prediction was that priming should make the neural response in the FFA more distinctive. RSA bore out this prediction, with voxel patterns for primed trials

significantly less similar to one another than for unprimed trials (and this remained after covarying out mean activity, suggesting it was not just a consequence of reduced SNR after priming). This is generally consistent with a sharpening of the likelihood function.

Limitations and Possible Future Directions

Behavioural

As well as laying bare some of the interpretational issues with the behavioural work, the modelling highlights another, perhaps deeper problem. The hypothetical resource-performance functions are a very general problem that pervades many types of psychological experiments – many experiments manipulate load and conclude that two tasks share a resource or not, depending on whether or not there is an interaction. In other words, resource-based accounts are “slippery”: without knowing independently the nature of the resource-performance function and the points on those functions that relate to a particular experimental condition (and particular participant), just about any pattern of means or correlations seems possible.

One approach which might appear promising at first glance would be trial-level analysis of the behavioural PSM experiments, to address whether, when processing appears facilitated (e.g. where decision task RT in Study Phase was low), subsequent successful recall is more likely. However, the model again shows that interpretation of this would be problematic: facilitation effects might appear as these reduced RTs, but equally could be transferred elsewhere, and there would be no principled way for the commensurate changes in performance to either a load task, or the associative memory task, to be interpreted. One possible avenue to mitigate this issue would be experimental to render observable the underlying resource-performance curves of the respective tasks. This would require systematic manipulation of resource availability (without imposing a dual task), such as testing people at different levels of sleep deprivation for example.

Correlations Between Outcome Measures and Presumed Shared Resources

As already mentioned, one unexpected finding of the modelling work in Ch. 4 was that, when considering trade-offs between tasks with assumed non-linearities relating the hypothesised resource to performance, whether or not correlations in performance changes correlate, and the direction of any such correlation, was found to be rather arbitrary. By making very small changes either to difficulty parameters of modelled tasks, or to the presumed priority given to one or another by a participant, correlation structure changed dramatically. There is no reason to think that this concern should not apply very broadly across cognitive science: claims of modularity often appeal to correlations (or their absence) between performance in tasks said to rely on the same processing module (e.g. Craik, Govoni, Naveh-Benjamin, & Anderson, 1996, and for discussion see Henson, 2006; Poldrack, 2006; White & Poldrack, 2013). The simulations explored in Chapter 4 raise the question of whether such correlations are merely artefactual, in which case they will have fortuitously bolstered some good claims, and erroneously encouraged type I error in other cases. Examining the literature for potential cases such as this could contribute to ongoing efforts to enhance the reproducibility of cognitive science such as The Reproducibility Project (Alexander et al., 2012). In addition, further exploration of this model and its simulations could lead to development of a principled taxonomy of such patterns of correlations, aiding analysis and characterisation of relationships between processes.

Neuroimaging

Problems similar to those discussed regarding behavioural data also pervade interpretation of fMRI data, since the relationship between resources and BOLD signal could also be nonlinear (Henson, 2011). In this light, the univariate results of the fMRI experiment, already somewhat mixed in their support for the trade-off hypothesis as formulated for the experiment, become yet more equivocal.

If one favoured the central resources account of PSM, then regardless of whether the freeing-up of resources is the result of an unspecified process, as in the current formulation of the resource account, or whether facilitation is due to a sharpening of

the likelihood function, further neuroimaging work could address the level at which priming produces changes. Such work could proceed using “reverse inference” (see Poldrack, 2006): we have competing hypotheses, and there is strong support in the existing literature for functional separation of mechanisms pertaining to perceptual, vs central, attention. Thus, one way to approach the question of whether priming acts on more perceptual or central aspects of processing would be to compare activity in perceptual attentional regions, with activity in regions more associated with central attentional processes.

The PIMMS account of PSM did not make clear predictions at the univariate level, and resources account did not make clear predictions at the multivariate level, and thus direct comparison of these resources and PIMMS hypotheses is rendered difficult. Although the PIMMS account does suggest more distinctive representations at the neural level, and indeed the finding of reduced correlations between primed trials does suggest that neural responses became more distinctive, this does not distinguish the PIMMS account from the resources account, which is not inconsistent with such a change to neural response as a result of priming. Note that sharpening of stimulus representation over voxels suggests a more distinctive neural response, although the converse of this is not true: a more distinctive neural response could produce no effect at the voxel level, since relevant features of a neural response making it a more distinctive could simply be impossible to detect in BOLD signal.

Prediction Error and Attention

More broadly, one explanation for the effect of central load on priming from a PIMMS perspective is that PE occurs and is then responsible for a shift of attention. Under this conception, attention would moderate the effect of PE on memory. Feldman and Friston (2010) and Kok et al. (2012) explore models in which attention enhances the precision of PE by increasing the postsynaptic gain of neurons encoding PE. This idea reconciles the well-documented signal-enhancement effects of attention to a predicted stimulus (see e.g. Desimone & Duncan, 1995), with the attenuation of signal that might be expected under predictive coding when a stimulus is correctly predicted. PE may be

seen as a signal indicating (potential) salience, which guides attention. Behavioural Experiments 2,3 and 4 may be seen as involving factorial manipulation of attention and PE (i.e., load affects attention, priming affects PE), but, again, the modelling highlights the problems that would need to be overcome for the absence of a load effect in a behavioural experiment to lend support to separable processes of PE and attention in PSM.

Such attention-based explanations of purported PE effects on subsequent memory are also considered by Greve et al., 2017, and have shown themselves difficult to refute in the behavioural paradigms. The central feature of the PIMMS account of PSM is that primed items generate more PE. This PE should have a short latency (though this also could depend on whether PE occurs at an early or later stage of processing), suggesting timing as a way to distinguish PE and attentional effects on subsequent memory. While the haemodynamic response reflected in BOLD integrates neural activity over timescales in the order of seconds, and thus may not have the temporal resolution given current methods to reflect differences between first-pass feedforward neural responses that occur within 100ms post-stimulus-onset, compared with top-down control processes of around 300ms, nevertheless EEG and MEG are well able to address such millisecond timescales. This could be addressed by testing whether M/EEG priming effects relevant to subsequent memory are seen predominantly early, around 100 ms, or later, around 300 or even 600 ms, as with well-established top-down attentional (e.g., Li, Gratton, Yao, & Knight, 2010; Polich, 2007) and subsequent memory (for review see Cohen, Pell, Edelson, Ben-Yakov, Pine, & Dudai, 2015) effects.

Reconciling or Deciding Between PIMMS/PE and Resources Accounts

Although the PIMMS claim is that there is something special about PE, and a resources account eschews PE, a problematic aspect is the nevertheless similar predictions that both accounts make when verbally described: the resources account states that resources are freed because priming has facilitated face processing, and the PIMMS account states that priming makes the sensory evidence more precise, which also facilitates processing.

One can be more specific about the accounts and still show apparent isomorphisms, or at least close relation between concepts relevant to PIMMS, and those relevant to an attentional account. In biased competition theory (Desimone & Duncan, 1995), and contemporary models associated with it which do not make use of PE (such as divisive normalisation, Carandini & Heeger, 2012; Reynolds & Heeger, 2009), bottom-up, or stimulus-driven, processes are those in which some aspect of a stimulus leads it to win the competition for representation (such as a bright light suddenly appearing in a previously-unattended location). In such a case, representations of other items in a neuronal RF will be subdued as a result of lateral inhibition, and the integrative aspect of biased competition sees to it that the stimulus wins the competition for representation throughout the processing hierarchy. Similarly, an unpredicted stimulus, or aspect of a stimulus, produces PE, and thus should entail processing for the resolution of the PE and update of the model. On the one hand this relatedness is not surprising, since both are attempts to explain observed phenomena: on the other hand, in order to distinguish between such accounts it will be necessary to find where they differ in their predictions.

On this basis, promising additional neuroimaging work would involve developing more detailed computational models so as to make precise quantitative predictions concerning either resource account, or prediction-error driven update. This work would include formal model comparison taking complexity into account, and would have the benefit of showing where predictions are identical, and where or in which versions of models they may differ. Spratling (2008, 2013), for example, shows that some common variants of predictive coding and biased competition models are formally identical. A literature review would reveal the current best candidates in terms of computational and algorithmic parsimony, and detailed simulations could be compared with the existing fMRI data. New fMRI data could then be acquired and only the best models kept for comparison.

It is possible that neuroimaging using somewhat indirect measures of brain activity such as fMRI and M/EEG may not have sufficient resolution to address different predictions

such models may make, and perhaps multi-site single-unit recordings will be necessary to disambiguate models relying on PE from those that do not. It may be that there is no difference to be seen, and that the difference is really a difference of language.

Concluding Remarks

This thesis has been an attempt to adjudicate between two accounts of the phenomenon of The Effect of Priming on Subsequent Memory (PSM). In the process of the work, PSM has been repeatedly replicated, and in a different modality from its original form in the work of Gagnepain et al. (2008). The simple temporal resources account was straightforwardly ruled out. However, although some evidence was shown in support of the central attentional resource account, modelling then showed that the results of the perceptual resources experiment were not sufficient to rule out that account. Furthermore, neither of those accounts were shown to be incompatible with the PIMMS account of PSM. The fMRI experiment then showed limited support for both accounts but did not disambiguate them. Further work has been suggested which might help further to distinguish the accounts, but it is not certain that they will be found to be separable.

Bibliography

- Alexander, A., Barnett-Cowan, M., Bartmess, E., Bosco, F. A., Brandt, M., Carp, J., ... Vianello, M. (2012). An open, large-scale, collaborative effort to estimate the reproducibility of psychological science. *Perspectives on Psychological Science*, 7(6), 657–660. <http://doi.org/10.1177/1745691612462588>
- Baddeley, A. (2003). Working memory: looking back and looking forward. *Nature Reviews. Neuroscience*, 4(10), 829–39. <http://doi.org/10.1038/nrn1201>
- Bastos, A., Usrey, W., Adams, R., & Mangun, G. (2012). Canonical Microcircuits for Predictive Coding. *Neuron*, 76(4), 695–711. <http://doi.org/10.1016/j.neuron.2012.10.038>
- Berry, C. J., Shanks, D. R., & Henson, R. N. A. (2008). A Single-System Account of the Relationship Between Priming, Recognition, and Fluency. *Journal of Experimental Psychology: Learning Memory and Cognition*, 34(1), 97–111. <http://doi.org/10.1037/0278-7393.34.1.97>
- Burgess, P. W., Dumontheil, I., & Gilbert, S. J. (2007). The gateway hypothesis of rostral prefrontal cortex (area 10) function. *Trends in Cognitive Sciences*, 11(7), 290–8. <http://doi.org/10.1016/j.tics.2007.05.004>
- Burgess, P. W., Gilbert, S. J., & Dumontheil, I. (2007). Function and localization within rostral prefrontal cortex (area 10). *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 362(April), 887–899. <http://doi.org/10.1098/rstb.2007.2095>
- Carandini, M., & Heeger, D. J. (2012). Normalization as a canonical neural computation. *Nature Reviews Neuroscience*, 13(1), 51–62. <http://doi.org/10.1038/nrn3136>
- Chun, M. M., & Johnson, M. K. (2011). Memory: enduring traces of perceptual and reflective attention. *Neuron*, 72(4), 520–35. <http://doi.org/10.1016/j.neuron.2011.10.026>
- Cohen, N., Pell, L., Edelson, M. G., Ben-Yakov, A., Pine, A., & Dudai, Y. (2015). Peri-encoding predictors of memory encoding and consolidation. *Neuroscience and Biobehavioral Reviews*, 50, 128–142. <http://doi.org/10.1016/j.neubiorev.2014.11.002>
- Cohen, N., & Squire, L. (1980). Preserved learning and retention of pattern-analyzing skill in amnesia: dissociation of knowing how and knowing that. *Science*, 210(4466), 207–210. <http://doi.org/10.1126/science.7414331>
- Craik, F. I. M., Govoni, R., Naveh-Benjamin, M., & Anderson, N. D. (1996). The effects of divided attention on encoding and retrieval processes in human memory. *Journal of Experimental Psychology: General*, 125(2), 159–180. <http://doi.org/10.1037/0096-3445.125.2.159>
- De Fockert, J. W., Rees, G., Frith, C. D. and Lavie, N. (2001). The Role of Working Memory in Visual Selective Attention. *Science*, 291(5509), 1803–1806.

<http://doi.org/10.1126/science.1056496>

- Duncan, J. (2010). The multiple-demand (MD) system of the primate brain: mental programs for intelligent behaviour. *Trends in Cognitive Sciences*, 14(4), 172–179. <http://doi.org/10.1016/j.tics.2010.01.004>
- Dunn, J. C., & Kirsner, K. (1988). Discovering Functionally Independent Mental Processes: The Principle of Reversed Association. *Psychological Review*, 95(1), 91–101. <http://doi.org/10.1037/0033-295X.95.1.91>
- Ebbinghaus, H. (2013). Memory: A Contribution to Experimental Psychology. *Annals of Neurosciences*, 20(4), 155–156. <http://doi.org/10.5214/ans.0972.7531.200408>
- Epstein, R., & Kanwisher, N. (1998). A cortical representation of the local visual environment. *Nature*, 392, 598. <http://dx.doi.org/10.1038/33402>
- Feldman, H., & Friston, K. J. (2010). Attention, uncertainty, and free-energy. *Frontiers in Human Neuroscience*, 4(December), 215. <http://doi.org/10.3389/fnhum.2010.00215>
- Fletcher, P. C., & Frith, C. D. (2009). Perceiving is believing: a Bayesian approach to explaining the positive symptoms of schizophrenia. *Nature Reviews. Neuroscience*, 10(1), 48–58. <http://doi.org/10.1038/nrn2536>
- Friston, K. J., Harrison, L., & Penny, W. (2003). Dynamic causal modelling. *NeuroImage*, 19(4), 1273–1302. [http://doi.org/10.1016/S1053-8119\(03\)00202-7](http://doi.org/10.1016/S1053-8119(03)00202-7)
- Gagnepain, P., Henson, R. N. A., Chételat, G., Desgranges, B., Lebreton, K., & Eustache, F. (2011). Is neocortical-hippocampal connectivity a better predictor of subsequent recollection than local increases in hippocampal activity? New insights on the role of priming. *Journal of Cognitive Neuroscience*, 23(2), 391–403. <http://doi.org/10.1162/jocn.2010.21454>
- Gagnepain, P., Lebreton, K., Desgranges, B., & Eustache, F. (2008). Perceptual priming enhances the creation of new episodic memories. *Consciousness and Cognition*, 17(1), 276–87. <http://doi.org/10.1016/j.concog.2007.03.006>
- Greve, A., Cooper, E., Anderson, M., & Henson, R. N. A. (2014). Does Prediction Error drive one-shot declarative learning? *Quarterly Journal of Experimental Psychology*.
- Greve, A., Cooper, E., Kaula, A., Anderson, M. C., & Henson, R. N. A. (2017). Does prediction error drive one-shot declarative learning? *Journal of Memory and Language*, 94(2017), 149–165. <http://doi.org/10.1016/j.jml.2016.11.001>
- Henson, R. N. A. (2003). Neuroimaging studies of priming. *Progress in Neurobiology*, 70(1), 53–81. [http://doi.org/10.1016/S0301-0082\(03\)00086-8](http://doi.org/10.1016/S0301-0082(03)00086-8)

- Henson, R. N. A. (2006). Forward inference using functional neuroimaging: Dissociations versus associations. *Trends in Cognitive Sciences*, 10(2), 64–69. <http://doi.org/10.1016/j.tics.2005.12.005>
- Henson, R. N. A. (2011). How to discover modules in mind and brain: The curse of nonlinearity, and blessing of neuroimaging. A comment on Sternberg (2011). *Cognitive Neuropsychology*, 28(June 2015), 209–223. <http://doi.org/10.1080/02643294.2011.561305>
- Henson, R. N. A., Eckstein, D., Waszak, F., Frings, C., & Horner, A. J. (2014). Stimulus – response bindings in priming. *Trends in Cognitive Sciences*, 18(7), 376–384.
- Henson, R. N. A., & Gagnepain, P. (2010). Predictive, interactive multiple memory systems. *Hippocampus*, 20(11), 1315–26. <http://doi.org/10.1002/hipo.20857>
- Henson, R. N. A., & Rugg, M. D. (2003). Neural response suppression, haemodynamic repetition effects, and behavioural priming. *Neuropsychologia*, 41(3), 263–270. [http://doi.org/10.1016/S0028-3932\(02\)00159-8](http://doi.org/10.1016/S0028-3932(02)00159-8)
- Henson, R. N. A., Wakeman, D., Phillips, C. & Rowe, J. (2012). Effective Connectivity between OFA and FFA during face perception: DCM of evoked MEG, EEG and fMRI responses. Poster presented at the 18th Annual Meeting of the Organization for Human Brain Mapping.
- Jurica, P. J., & Shimamura, A. P. (1999). Monitoring item and source information: Evidence for a negative generation effect in source memory. *Memory & Cognition*, 27(4), 648–656. <http://doi.org/10.3758/BF03211558>
- Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The fusiform face area: a module in human extrastriate cortex specialized for face perception. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 17(11), 4302–11. <http://doi.org/10.1098/Rstb.2006.1934>
- Knill, D. C., & Pouget, A. (2004). The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends in Neurosciences*, 27(12), 712–9. <http://doi.org/10.1016/j.tins.2004.10.007>
- Kok, P., Rahnev, D., Jehee, J. F. M., Lau, H. C., & De Lange, F. P. (2012). Attention reverses the effect of prediction in silencing sensory signals. *Cerebral Cortex*, 22(9), 2197–2206. <http://doi.org/10.1093/cercor/bhr310>
- Kriegeskorte, N., Mur, M., & Bandettini, P. (2008). Representational similarity analysis - connecting the branches of systems neuroscience . *Frontiers in Systems Neuroscience* . <https://www.frontiersin.org/article/10.3389/neuro.06.004.2008>
- Lavie, N. (1995). Perceptual load as a necessary condition for selective attention. *Journal of Experimental Psychology: Human Perception and Performance*, 21(3), 451–468. <http://doi.org/10.1037/0096-1523.21.3.451>

- Lavie, N. (2005). Distracted and confused?: selective attention under load. *Trends in Cognitive Sciences*, 9(2), 75–82. <http://doi.org/10.1016/j.tics.2004.12.004>
- Lavie, N., Beck, D. M., & Konstantinou, N. (2014). Blinded by the load: attention, awareness and the role of perceptual load. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369(1641). <http://rstb.royalsocietypublishing.org/content/369/1641/20130205.abstract>
- Lavie, N., Hirst, A., de Fockert, J. W., & Viding, E. (2004). Load theory of selective attention and cognitive control. *Journal of Experimental Psychology. General*, 133(3), 339–54. <http://doi.org/10.1037/0096-3445.133.3.339>
- Lee, T. S., & Mumford, D. (2003). Hierarchical Bayesian inference in the visual cortex. *Journal of the Optical Society of America A*, 20(7), 1434. <http://doi.org/10.1364/JOSAA.20.001434>
- Li, L., Gratton, C., Yao, D., & Knight, R. T. (2010). Role of frontal and parietal cortices in the control of bottom-up and top-down attention in humans. *Brain Research*, 1344, 173–184. <http://doi.org/10.1016/j.brainres.2010.05.016>
- Lundqvist D., Litton J-E. (1998). The Averaged Karolinska Directed Emotional Faces-AKDEF. AKDEF CD ROM. Psychology section, Karolinska Institutet, Stockholm.
- Martin, E. (1968). Stimulus meaningfulness and paired-associate transfer: An encoding variability hypothesis. *Psychological Review*, 75(5).
- Naveh-Benjamin, M., Guez, J., & Marom, M. (2003). The effects of divided attention at encoding on item and associative memory. *Memory & Cognition*, 31(7), 1021–1035. <http://doi.org/10.3758/BF03196123>
- Norman, D. A., & Bobrow, D. G. (1975). On data-limited and resource-limited processes. *Cognitive Psychology*, 7(1), 44–64. [http://doi.org/10.1016/0010-0285\(75\)90004-3](http://doi.org/10.1016/0010-0285(75)90004-3)
- O'Reilly, J. X., Woolrich, M. W., Behrens, T. E. J., Smith, S. M., & Johansen-Berg, H. (2012). Tools of the trade: Psychophysiological interactions and functional connectivity. *Social Cognitive and Affective Neuroscience*, 7(5), 604–609. <http://doi.org/10.1093/scan/nss055>
- Phillips P.J., Moon H., Rizvim S.A., Rauss P.J. (2000). The FERET evaluation methodology for face-recognition algorithms. *IEEE Transactions on pattern analysis and machine intelligence* 22, 1090-1104.
- Poldrack, R. A. (2006). Can cognitive processes be inferred from neuroimaging data? *Trends in Cognitive Sciences*, 10(2), 59–63. <http://doi.org/10.1016/j.tics.2005.12.004>
- Poppenk, J., Köhler, S., & Moscovitch, M. (2010). Revisiting the novelty effect: When

- familiarity, not novelty, enhances memory. *Journal of Experimental Psychology: Learning Memory and Cognition*, 36(5), 1321–1330. <http://doi.org/10.1037/a0019900>
- Polich, J. (2007). Updating P300: An integrative theory of P3a and P3b. *Clinical Neurophysiology*, 118(10), 2128–2148. <http://doi.org/10.1016/j.clinph.2007.04.019>
- Rao, R. P. (2004). Hierarchical Bayesian Inference in Networks of Spiking Neurons. *Advances in Neural Information Processing*, 1113–1120.
- Rao, R. P. N., & Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1), 79–87. <http://doi.org/10.1038/4580>
- Reynolds, J. H., & Heeger, D. J. (2009). The Normalization Model of Attention. *Neuron*, 61(2), 168–185. <http://doi.org/10.1016/j.neuron.2009.01.002>
- Rissman, J., Gazzaley, A., & D'Esposito, M. (2004). Measuring functional connectivity during distinct stages of a cognitive task. *NeuroImage*, 23(2), 752–763. <http://doi.org/10.1016/j.neuroimage.2004.06.035>
- Schacter, D. L., & Tulving, E. (1994). What are the memory systems of 1994? In *Memory systems 1994*. (pp. 1–38).
- Schultz, W. (1997). A Neural Substrate of Prediction and Reward. *Science*, 275(5306), 1593–1599. <http://doi.org/10.1126/science.275.5306.1593>
- Seitz, A. R., & Watanabe, T. (2009). The phenomenon of task-irrelevant perceptual learning. *Vision Research*, 49(21), 2604–2610. <http://doi.org/10.1016/j.visres.2009.08.003>
- Shallice, T. (2003). Functional imaging and neuropsychology findings: How can they be linked? *NeuroImage*, 20(SUPPL. 1), 146–154. <http://doi.org/10.1016/j.neuroimage.2003.09.023>
- Shattuck, D. W., Mirza, M., Adisetiyo, V., Hojatkashani, C., Salamon, G., Narr, K. L., ... Toga, A. W. (2008). Construction of a 3D probabilistic atlas of human cortical structures. *NeuroImage*, 39(3), 1064–1080. <http://doi.org/10.1016/j.neuroimage.2007.09.031>
- Slamecka, N. J., & Graf, P. (1978). The Generation Effect: Delineation of a Phenomenon. *Journal of Experimental Psychology: Human Learning and Memory*, 4(6), 592–604.
- Spratling, M. W. (2008). Reconciling predictive coding and biased competition models of cortical function. *Frontiers in Computational Neuroscience*, 2(October), 4. <http://doi.org/10.3389/neuro.10.004.2008>
- Spratling, M. W. (2013). Distinguishing theory from implementation in predictive coding accounts of brain function. *Behavioral and Brain Sciences*, 36(3), 231–232.

<http://doi.org/10.1017/S0140525X12002178>

- Stark, C. E. L., & Squire, L. R. (2001). When zero is not zero: The problem of ambiguous baseline conditions in fMRI. *Proceedings of the National Academy of Sciences*, 98(22), 12760–12766. <http://doi.org/10.1073/pnas.221462998>
- Tottenham N., Tanaka J., Leon A. C., McCarry T., Nurse M., Hare T.A., Marcus D.J., Westerlund A., Casey B.J., Nelson C.A. (2009). The NimStim set of facial expressions: judgments from untrained research participants. *Psychiatry Research*, 168, 242-249.
- Tulving, E. (1985). Memory and Consciousness. *Canadian Psychology*, 26(1), 1–12. <http://doi.org/10.1037/h0080017>
- Turk-Browne, N. B., Yi, D. J., & Chun, M. M. (2006). Linking implicit and explicit memory: Common encoding factors and shared representations. *Neuron*, 49, 917–927. <http://doi.org/10.1016/j.neuron.2006.01.030>
- Uncapher, M. R., & Rugg, M. D. (2005). Encoding and the durability of episodic memory: A functional magnetic resonance imaging study. *Journal of Neuroscience*, 25(31), 7260–7. <http://doi.org/10.1523/JNEUROSCI.1641-05.2005>
- Wagner, A. D. (1998). Building Memories: Remembering and Forgetting of Verbal Experiences as Predicted by Brain Activity. *Science*, 281(5380), 1188–1191. <http://doi.org/10.1126/science.281.5380.1188>
- Wagner, A. D., Maril, A., & Schacter, D. L. (2000). Interactions Between Forms of Memory: When Priming Hinders New Episodic Learning. *Journal of Cognitive Neuroscience*, 12(supplement 2), 52–60. <http://doi.org/10.1162/089892900564064>
- Watanabe, K., & Funahashi, S. (2014). Neural mechanisms of dual-task interference and cognitive capacity limitation in the prefrontal cortex. *Nature Neuroscience*, 17(4), 601–11. <http://doi.org/10.1038/nn.3667>
- White, C. N., & Poldrack, R. A. (2013). Using fMRI to Constrain Theories of Cognition. *Perspectives on Psychological Science*, 8(1), 79–83. <http://doi.org/10.1177/1745691612469029>
- Yi, D. J., Woodman, G. F., Widders, D., Marois, R., & Chun, M. M. (2004). Neural fate of ignored stimuli: dissociable effects of perceptual and working memory load. *Nature Neuroscience*, 7(9), 992–6. <http://doi.org/10.1038/nn1294>
- Yonelinas, A. P. (2002). The Nature of Recollection and Familiarity: A Review of 30 Years of Research. *Journal of Memory and Language*, 46(3), 441–517. <http://doi.org/10.1006/jmla.2002.2864>