

Minerva Access is the Institutional Repository of The University of Melbourne

Author/s:

Lilburn, Simon David

Title:

Information limits within visual short-term memory

Date:

2016

Persistent Link:

<https://hdl.handle.net/11343/116417>

Terms and Conditions:

Terms and Conditions: Copyright in works deposited in Minerva Access is retained by the copyright owner. The work may not be altered without permission from the copyright owner. Readers may only download, print and save electronic copies of whole works for their own personal non-commercial use. Any use that exceeds these limits requires permission from the copyright owner. Attribution is essential when quoting or paraphrasing from these works.

Information limits within visual short-term memory

Simon David Lilburn

<http://orcid.org/0000-0001-8820-8188>

Doctor of Philosophy

June 2016

Melbourne School of Psychological Sciences
The University of Melbourne

Submitted in total fulfilment of the requirements of the degree of Doctor of Philosophy.

Produced on archival quality paper.

Abstract

This thesis examined separable information constraints upon visual short-term memory through application of signal detection theory and diffusion modelling (R. Ratcliff, 1978) applied to a series of related tasks. Following D. K. Sewell, S. D. Lilburn, and P. L. Smith (2014), I employed simple psychophysical paradigms using the near-threshold presentation of stimulus information and controlling the complexity of decision-making through a post-stimulus probe to obtain both accuracy and response time data.

The first part of this thesis examines the relationship between two common visual short-term memory procedures—orientation discrimination and change detection—to characterise the effect of the task type on responding. Modelling of observer sensitivity for Experiment 1 showed that a good account of the data can be obtained with the use of the sample-size relation between performance and memory load, with the addition of an item in change detection trials to account for the effect of probe array on performance. This was supported in Experiment 2 by the higher stimulus contrasts for change detection trials required to offset the decrement in accuracy. Response time modelling using the diffusion model further supported this account, with the addition of a constant time for encoding and comparing the probe array in change detection decisions and the inclusion of an intrusion process to model the entry of non-target information into the decision.

The second part of this thesis expanded the modelling of the first part to examine the constraints upon orientation information in two fine orientation discrimination experiments. These experiments required observers to judge the direction of small angular offsets from a known referent. The change in observer sensitivity across different target angular offset conditions was well captured by a Gaussian-shaped tuning function, centred on the orientation of the referent, and weighting squared sensitivity directly. This information constraint was found to be independent of the sample-size relation, leading to the conclusion that the division of memory resources may be separate from the quality of stimulus information. Stimulus exposure duration was modelled as a linear increase in observer sensitivity, also independent of the sample-size or tuning function constraints. Diffusion modelling with these three constraints placed on the drift rate provided a parsimonious description of both the response proportions and response time distributions of a large memory experiment with a small number of parameters.

In all, the progression of both sensitivity and response time modelling through this thesis aims to demonstrate three separable constraints on the fundamental capacity of visual short-term memory—the sample-size relation, the tuning channel constraint on orientation information, and the linear constraint on information growth—and the importance of considering the role of decision processes explicitly.

Declaration

This is to certify that—

1. The thesis comprises only my original work towards the degree of Doctor of Philosophy;
2. Due acknowledgement has been made in the text to all other material used; and,
3. The thesis is fewer than 100,000 words, exclusive of tables, figures, footnotes, references, and appendices.

Simon Lilburn

Acknowledgements

My supervisors, Prof. Philip Smith and Dr David Sewell, are exemplary scientists and mentors: rigorous and methodical, thoughtful and supportive, passionate and compassionate. I do not know of any individuals I would prefer to have as supervisors. Although I own any mistakes within this thesis, I share any of the praise with them. I cannot thank them enough for their guidance, kindness, and insight.

I would also like to give a special thank you to Dr Elaine Corbett and Saam Saber. The stimulating discussions and excellent company in the Vision and Attention Lab have made for a wonderful place to work.

I must also thank Daniel Bennett, Christina van Heer, Katharina Voigt, Bowen Fung, Daniel Rosenblatt, Dr Kelly Trezise, Jacob Paul, Jian Chen, Shi Xian Liew, Annie Blunden, Robert de Lisle, Gabriel Ong, Christine Podwysocki, Dr Jason Forte, Dr Simon Cropper, Dr Stefan Bode, Dr Daniel Little, Dr Meredith McKague, Dr Chris Groot, and Prof. John Trinder. Their support and enthusiasm have made this period an immense pleasure.

Thank you to my family and friends: Mum, Dad, Gop, Talki, Chris, Michael, Elizabeth, Sam, Rebecca, Neha, Jialing, and the many, many others. Too many to name. I am so lucky to have such good people in my life.

Thank you to the observers for their time and good humour.

This work used Free (as in Freedom) Software extensively. I owe a debt of gratitude to the developers of Linux, gcc, ssh, tmux, LaTeX, emacs, Python (particularly Numpy and Scipy), and R.

Contents

1	Overview of the present thesis	1
2	General literature review	3
2.1	Visual and informational persistence	4
2.2	Functionally distinct visual memory stores	6
2.3	The storage capacity limits of visual short-term memory	9
2.4	The temporal dynamics of visual short-term memory	20
2.5	An information limit on short-term memory	27
3	Change detection and orientation discrimination: sensitivity	33
3.1	Signal detection theory	41
3.2	Experiment 1	47
3.3	Experiment 2	63
3.4	General discussion	69
4	Change detection and orientation discrimination: response time	77
4.1	Experiment 1	81
4.2	Experiment 2	108
4.3	General discussion	126
5	Fine orientation discrimination: sensitivity	143
5.1	Experiment 3	147
5.2	Experiment 4	161
5.3	General discussion	170
6	Fine orientation discrimination: response time	181
6.1	Experiment 3	183
6.2	Experiment 4	195
6.3	General discussion	203
7	A summary and discussion of the current work	217
7.1	The effect of decision type on observer performance	218
7.2	The effect of stimulus discriminability on observer performance	223
7.3	Methodological implications	226
7.4	Theoretical implications	229
7.5	Future research	233

List of Figures

3.1	Predicted proportion of correct responses for different SDT models. . . .	38
3.2	The relationship between detection and discrimination in the theory of recognition due to Tanner (1956).	42
3.3	The differencing model of Sorkin (1962).	45
3.4	Overview of Experiment 1 presentation order.	49
3.5	Group average accuracy data for Experiment 1.	51
3.6	Accuracy data for each observer in Experiment 1.	52
3.7	Predictions for a differencing model with sample-size increment against group average data in Experiment 1.	58
3.8	Predictions for a differencing model without sample-size increment against group average data in Experiment 1.	59
3.9	Predictions for an independent observations model against group average data in Experiment 1.	60
3.10	Predicted proportion of correct responses for a differencing model against an independent observations and a discrimination model.	62
3.11	Group average accuracy data for Experiment 2.	67
3.12	Accuracy data for each observer in Experiment 2.	68
4.1	Group average mean response time data for Experiment 1.	82
4.2	Mean response time data for each observer in Experiment 1.	83
4.3	Response time kernel density estimates for two observers in Experiment 1.	85
4.4	Different theoretical response time component distributions.	88
4.5	Results of the convolution analysis for observer BF in Experiment 1.	91
4.6	Results of the convolution analysis for observer DB in Experiment 1.	92
4.7	Results of the convolution analysis for observer KT in Experiment 1.	93
4.8	Results of the convolution analysis for observer SL in Experiment 1.	94
4.9	Results of the convolution analysis for observer SS in Experiment 1.	95
4.10	A schematic of the diffusion model of Ratcliff (1978).	97
4.11	The construction of the quantile–probability plot (QPP).	99
4.12	The effect of the intrusion process on the leading edge of response times.	102
4.13	QPP for the best fitting diffusion model against the group average data in Experiment 1.	106
4.14	Group average mean response time data for Experiment 2.	109
4.15	Mean response time data for each observer in Experiment 2.	110
4.16	Results of the convolution analysis for observer BF in Experiment 2.	114
4.17	Results of the convolution analysis for observer DB in Experiment 2.	115
4.18	Results of the convolution analysis for observer KT in Experiment 2.	116
4.19	Results of the convolution analysis for observer SL in Experiment 2.	117

List of Figures

4.20	Results of the convolution analysis for observer SS in Experiment 2.	118
4.21	QPP for the diffusion model with independent boundary separation values across decision types against the group average data in Experiment 2.	122
4.22	QPP for the diffusion model with independent intrusion rate values across decision types against the group average data in Experiment 2.	123
4.23	QPP for the diffusion model with independent boundary separation and intrusion rate values across decision types against the group average data in Experiment 2.	124
5.1	Overview of Experiment 3 presentation order.	148
5.2	Group average accuracy data for Experiment 3.	150
5.3	Accuracy data for each observer in Experiment 3.	151
5.4	The connection between oriented stimuli and the underlying psychophysical detectors.	153
5.5	Three different forms of detector tuning functions.	154
5.6	The relationship between target and non-target detectors in the tuning channel model.	155
5.7	Predictions of the extended signal detection theory model against the group average data of Experiment 3.	160
5.8	Overview of Experiment 4 presentation order.	163
5.9	Group average accuracy data for Experiment 4.	164
5.10	Predictions of the extended signal detection theory model against the group average data of Experiment 4.	168
5.11	Accuracy data for observer HA in Experiment 4.	176
5.12	Accuracy data for observer LA in Experiment 4.	176
5.13	Accuracy data for observer ME in Experiment 4.	177
5.14	Accuracy data for observer SN in Experiment 4.	178
5.15	Accuracy data for observer TS in Experiment 4.	178
6.1	Group average mean response time data for Experiment 3.	185
6.2	Mean response time data for each observer in Experiment 3.	186
6.3	QPP for the best fitting diffusion model against the group average data in Experiment 3.	192
6.4	Nondecision time parameter estimates for each observer from the best fitting diffusion model in Experiment 3.	193
6.5	QPP for a diffusion model configuration without flexible trial-to-trial drift variability against the group average data in Experiment 3.	194
6.6	Group average mean response time data for Experiment 4.	196
6.7	QPP for the diffusion model configuration that best fit the group average data Experiment 4.	201
6.8	Mean response time data for observer HA in Experiment 4.	211
6.9	Mean response time data for observer LA in Experiment 4.	211
6.10	Mean response time data for observer ME in Experiment 4.	212
6.11	Mean response time data for observer SN in Experiment 4.	213

6.12 Mean response time data for observer TS in Experiment 4. 213

List of Tables

3.1	Different SDT model configurations for Experiment 1.	54
3.2	The best fitting SDT models for the group average data in Experiment 1. . .	57
3.3	Mean stimulus contrasts for Experiment 2.	66
3.4	The best fitting SDT models for observer BF in Experiment 1.	74
3.5	The best fitting SDT models for observer DB in Experiment 1.	74
3.6	The best fitting SDT models for observer KT in Experiment 1.	74
3.7	The best fitting SDT models for observer SL in Experiment 1.	74
3.8	The best fitting SDT models for observer SS in Experiment 1.	75
3.9	The best fitting parameter estimates for the SDT models in Experiment 1. .	76
4.1	Goodness-of-fit and information criteria statistics for convolution analysis of Experiment 1.	89
4.2	Best fitting parameters for the convolution analysis of Experiment 1. . . .	90
4.3	Parameters controlling the extended diffusion model.	103
4.4	The best fitting diffusion models for the group average data in Experiment 1.	105
4.5	Goodness-of-fit and information criteria statistics for convolution analysis of Experiment 2.	111
4.6	Best fitting parameters for the convolution analysis of Experiment 2. . . .	112
4.7	The best fitting SDT models for the group average data in Experiment 2. .	120
4.8	The best fitting diffusion models for observer BF in Experiment 1.	134
4.9	The best fitting diffusion models for observer DB in Experiment 1.	134
4.10	The best fitting diffusion models for observer KT in Experiment 1.	134
4.11	The best fitting diffusion models for observer SL in Experiment 1.	135
4.12	The best fitting diffusion models for observer SS in Experiment 1.	135
4.13	The best fitting parameter estimates for the diffusion model with a sample-size increment in Experiment 1.	136
4.14	The best fitting parameter estimates for the diffusion model without a sample-size increment in Experiment 1.	137
4.15	The best fitting diffusion models for observer BF in Experiment 2.	138
4.16	The best fitting diffusion models for observer DB in Experiment 2.	138
4.17	The best fitting diffusion models for observer KT in Experiment 2.	138
4.18	The best fitting diffusion models for observer SL in Experiment 2.	139
4.19	The best fitting diffusion models for observer SS in Experiment 2.	139
4.20	The best fitting parameter estimates for the diffusion model with independent boundary separation values across decision types in Experiment 2.	140
4.21	The best fitting parameter estimates for the diffusion model with independent intrusion rate values across decision types in Experiment 2. . . .	141

List of Tables

4.22	The best fitting parameter estimates for the diffusion model with independent boundary separation and intrusion rate values across decision types in Experiment 2.	142
5.1	Different SDT model configurations for Experiment 3.	157
5.2	The best fitting SDT models for the group average data in Experiment 3. .	158
5.3	Different SDT model configurations for Experiment 4.	167
5.4	The best fitting SDT models for the group average data in Experiment 4. .	168
5.5	The best fitting signal detection theory models for observer AB in Experiment 3.	174
5.6	The best fitting signal detection theory models for observer CVH in Experiment 3.	174
5.7	The best fitting signal detection theory models for observer SIS in Experiment 3.	174
5.8	The best fitting signal detection theory models for observer SL in Experiment 3.	175
5.9	The best fitting signal detection theory models for observer XL in Experiment 3.	175
5.10	The best fitting SDT models for observer HA in Experiment 4.	179
5.11	The best fitting SDT models for observer LA in Experiment 4.	179
5.12	The best fitting SDT models for observer ME in Experiment 4.	179
5.13	The best fitting SDT models for observer SN in Experiment 4.	180
5.14	The best fitting SDT models for observer TS in Experiment 4.	180
6.1	Different diffusion model configurations for Experiment 3.	190
6.2	The best fitting diffusion models for the group average data in Experiment 3. .	191
6.3	Different intrusion rate configurations for Experiment 4.	199
6.4	The best fitting diffusion models for the group average data in Experiment 4. .	199
6.5	The best fitting diffusion models for observer AB in Experiment 3.	208
6.6	The best fitting diffusion models for observer CVH in Experiment 3.	208
6.7	The best fitting diffusion models for observer SIS in Experiment 3.	208
6.8	The best fitting diffusion models for observer SL in Experiment 3.	209
6.9	The best fitting diffusion models for observer XL in Experiment 3.	209
6.10	The best fitting parameter estimates for the diffusion model that best fit the group average data in Experiment 3.	210
6.11	The best fitting diffusion models for observer HA in Experiment 4.	214
6.12	The best fitting diffusion models for observer LA in Experiment 4.	214
6.13	The best fitting diffusion models for observer ME in Experiment 4.	214
6.14	The best fitting diffusion models for observer SN in Experiment 4.	215
6.15	The best fitting diffusion models for observer TS in Experiment 4.	215
6.16	The best fitting parameter estimates for the diffusion model configuration that best fit the group average data in Experiment 4.	216

Chapter 1

Overview of the present thesis

Visual short-term memory, or visual working memory, is the phenomenon of perceptual information persisting both beyond the cessation of sensory input, and beyond the transient and labile retention of activation in early visual systems (Phillips, 1974). This memory system appears to be fundamental within visual perceptual processing more generally: the characteristics of how information is stored in visual short-term memory are evident even in very basic visual tasks, such as cued detection (Ratcliff & Rouder, 2000; Smith, Ratcliff, & Wolfgang, 2004). As such, the study of representations within the visual short-term memory system—their formation, their maintenance, the constraints and structure upon their storage, and the decision-making processes that make some demand on these representations to—has been a central line of inquiry in the study of visual perception over the last fifty years.

This thesis will examine the properties of these memory representations through a series of psychophysical experiments with a consistent methodological design. In particular, the experiments presented in this thesis will use small-N designs with stimuli presented near the sensory threshold, unlike much of the work conducted in the visual memory and cognition literature, but strongly related to the work conducted in low-level visual psychophysics (e.g., Carrasco, Penpeci-Talgar, & Eckstein, 2000; Downing, 1988; Duncan, 1980; Lee, Koch, & Braun, 1997; Palmer, 1990; Smith, 1995; Smith & Ratcliff, 2004; Thomas & Gille, 1979). These designs allow for an extensive consideration of the decision processes that relate representation to response—a topic somewhat neglected in the broader visual short-term memory literature—and to examine the distributions of response times under various experimental manipulations. This specificity of the inferences provided by these designs can, I hope, allow for a powerful set of constraints on quantitative models of visual short-term memory.

Four experiments comprise the main body of the thesis. These investigations are divided into two larger themes: the structure and temporal dynamics of the memory repre-

sentations; and the structure and temporal dynamics of two means of assaying memory representations. The two themes are, of course, fundamentally inseparable and do not represent some fine taxonomic distinction, but are used here as an expository convenience. Each part will contain one chapter on modelling observer sensitivity and one chapter on modelling response time distributions.

The first theme, the structure and temporal dynamics of the experimental paradigms used to investigate memory representations, will be covered in Chapters 3 and 4. These chapters will cover the two experiments examining the sensitivity and response time characteristics of change detection and (orthogonal) orientation discrimination paradigms. The first of the two chapters in this part will consider the changes in observer sensitivity with changes in the size and exposure duration of the memory array and duration; the second of these chapters will consider changes in response time—in addition to accuracy—using the diffusion model of Ratcliff (1978).

The second theme, the structure and temporal dynamics of the memory representations, will be covered in Chapters 5 and 6. These chapters will cover the two experiments examining the sensitivity and response time characteristics of fine orientation discrimination from memory representations. Like the preceding part, the first of the two chapters will consider sensitivity and the second of the two chapters will focus on the response time data.

The next chapter will provide a general review of the literature concerning visual short-term memory and, in particular, the themes of storage capacity and the temporal properties of representation formation.

Chapter 2

General literature review

The visual world for healthy adults is furnished with perceptual objects that maintain consistent form over time and are perceived complete without conscious effort. The experience of visual perception—the richness and fluidity present—provides very little indication of the nature of the initial sensory input: the raw sensory material projected from the retina is characterised by sudden disruptions from saccades (Irwin & Andrews, 1996) and corruption from neural noise (Pelli, 1991).

This conscious experience of visual information is the product of incoming sensory information integrated across time to form coherent perceptual objects. The memory processes that serve to integrate information provide a consistent basis from which decisions about the world can be made and meaning can be extracted: processes which may far exceed (by an order of magnitude) the duration for which the visual stimulus is available (Ratcliff & Rouder, 2000). The role of such memory processes is central in visual processing, seated at the nexus between sensory representation and broader decision-making and interceding in even simple visual tasks such as detection (Smith & Ratcliff, 2004) and correcting gaze trajectories (Hollingworth & Luck, 2009).

In general, integrative processes in vision fall into two categories¹: the low-level persistence of sensory signals after their offset unable to survive subsequent input and which may be used for decision-making, labelled *iconic memory*, after Neisser (1967); and, the higher-level, adaptive maintenance of information across time that is more selective, able to survive changes to the visual input, and more behaviourally bound. The second of these phenomena is termed *visual short-term memory*, or *visual working memory*. It is this latter memory phenomenon that shall be the focus of this thesis.

This chapter will provide a general overview of some of the core themes of empirical

¹Hollingworth (2004) also distinguishes between labile forms of “informational” and “visible” persistence, following Coltheart (1980), from low-level (“iconic”) persistence. We will consider these forms of persistence to be identical to the more generally recognised phenomenon of the “iconic” sensory buffer.

research pertaining to visual short-term memory. The discussion will proceed as follows. I will first examine the early historical work investigating memory processes in the visual perceptual system. I will then examine the more recent empirical work and debate surrounding the structure of the visual short-term memory store and its storage capacity. The temporal dynamics of the memory system will then be reviewed, leading to a discussion of work that myself and colleagues have conducted that examines both storage capacity limitations and the temporal dynamics of visual short-term memory (the direct precursor to the current thesis).

2.1 Visual and informational persistence

Modern work examining short-term memory phenomena in visual perception originates with the *partial report* paradigm of G. Sperling (1960). Over a series of experiments, Sperling examined the ability of observers to report from memory a specified subset of a larger array of letters presented for a very brief interval. Observers were first shown an array of letters, split into rows. After the offset of the letter array, and a specified interval had elapsed, observers were then cued to report a single row of the letter array with a tone (the pitch of the tone indicating the row to report). The principal finding was that observers could accurately report any specified subset of the memory array when cued with a brief period after the offset of the stimulus; in one instance, an average of seven letters from an array of nine could be accurately recalled if the observer was cued within 150 ms. The longer the period—known as the *probe delay*—between the offset of the memory array and the presentation of the report tone, the fewer letters could be accurately identified. At probe delays of approximately 500 ms in conditions presenting nine and twelve letters, performance plateaued at a level of accuracy which indicated that only about four or five elements in total were available.

The inference drawn by Sperling (G. Sperling, 1960, 1963) from these data was that all visual information was available for decision-making for a period shortly after the stimulus offset, indicating the presence of some detailed but short-lived memory representation (or *memory trace*). Although there had been previous indications that visual information could persist in some representation for a brief period after presentation (Cattell, 1885; Erdmann & Dodge, 1898), and the phenomenological persistence of vision is self-evident, the partial report was a precise solution to the problem of quantifying the extent of and change in a quickly decaying memory trace. The partial report task also resembles the post-stimulus probe paradigm (Downing, 1988), popular in attentional re-

search and a precursor to the research presented in this thesis.

G. Sperling (1963), using the data from an earlier study reported by Baxt (1871), further demonstrated the sensitivity of this short-lived visual information to interference from subsequent visual stimulation. Letter recognition subsequent to an immediate flash of light after presentation was essentially nil; recognition after some period of delay before the presentation of the interfering flash increased linearly, indicating a recovery of approximately one letter for every 10 ms that the interfering light was delayed.

The memory phenomenon identified by Sperling was later called *iconic memory* by Neisser (1967), and became the prototype of a sensory buffer in the information processing *modal memory model* of Atkinson and Shiffrin (1968). The memory model divided memory phenomena into three modes: an initial sensory buffer, like the iconic memory store or the analogous “echoic” memory store in the auditory domain, which could hold modality-specific information in a high fidelity representation, but could not be rehearsed and decayed quickly; a modality-independent “short-term memory” store, holding information for active mental processing and constrained by a limited capacity (at the time presumed to be around seven items, following Miller, 1956), and which requires active rehearsal; and a “long-term memory” store, with a virtually limitless storage capacity and an indefinite lifespan for representations stored, but which requires intensive and slow elaborative processing (Craik & Lockhart, 1972; Craik & Watkins, 1973) for trace formation to occur.

Critically, the short-term memory store capable of sustaining a memory representation was argued to be auditorially or lexically mediated, with only the iconic memory store operating directly upon visual representations. The decline in performance seen in partial report tasks was imputed to be the rate of decay of the memory store, a fixed property not under behavioural control. The limited baseline performance seen in partial report tasks at longer probe delays was argued to be the result of a recoding (or “scanning”) of the modality-bound visual information into a more general, rehearsable, short-term memory representation (Atkinson & Shiffrin, 1968; G. Sperling, 1963)². The distinction between the two apparent rates of memory retention—limitless immediate recall and a truncated amount of information that persisted beyond some delay in instructions—anticipated the later work of Phillips (1974) on visual short-term memory. As the stimuli used by Sperling and, later, by Averbach and Sperling (1961) were alphanumeric, the

²G. Sperling (1963) calls the short-term memory store, the “Auditory Information Store” system, in contradistinction to his nomenclature for the iconic memory store, the “Visual Information Storage” system, but there is no reason to think that this system is functionally distinct from the short-term memory system of Atkinson and Shiffrin (1968). The identification between the two is only possible when verbalisable stimuli are used.

durable representations required could easily be argued to be the product of verbal recoding, rather than being intrinsically visual. Subsequent studies were required to indicate the presence of more durable representations that were intrinsically visual in constitution.

2.2 Functionally distinct visual memory stores

The modal model of memory, and the sketch of a memory hierarchy given by G. Sperling (1963), only defined a single form of memory specific to the visual domain. The short-term memory store, capable of active maintenance of a small number of representations, was originally described in terms that would be general across cognitive and perceptual domains, largely reflecting research—like that of L. Peterson and Peterson (1959) and Miller (1956)—that dealt with digit span and alphabetic trigrams task: discrete and declarative units of information. The concept of memory provided by this model is one that naturally lends itself to discussion of rehearsal and “control processes” that involve well-defined symbolic operations on stored information rather than purely visual information.

Shortly after the proposal of the Atkinson–Shiffrin modal model, however, investigations by Posner and colleagues (Posner, 1967; Posner, Boies, Eichelman, & Taylor, 1969; Posner & Keele, 1967) demonstrated some retention of visual information beyond the short timespan of iconic memory traces seen in partial report paradigms. In a study conducted by Posner and Keele, observers were asked to indicate whether the identity of two pairs of letters, successively presented in the same location, matched. In one set of trials, identity was maintained by the case of the letters was switched (that is, an “a” would be matched with an “A” or vice versa). In another set of trials, the identity and case were maintained between the two pairs, meaning that the visual information was identical. Posner and Keele reported an advantage in mean response time when the visual aspect of the stimulus was maintained across consecutive arrays when compared to simply maintaining the identity of the letters. This advantage persisted even when the interval between the two letter sets (the inter-stimulus interval) was 1,500 ms, indicating a purely visual benefit for comparison, a benefit that could not be accounted for by verbal recoding, beyond the accepted range of iconic memory. Kroll, Parks, Parkinson, Bieber, and Johnson (1970) further demonstrated that alphabetic stimuli were less prone to retroactive interference by verbal shadowing (i.e., reciting a list of aurally presented letters) when compared to stimuli that were presented aurally (prior to the shadowing

task). The different rates of correct retrieval between aurally presented and visually presented stimuli over a range of retention intervals—up to 25 seconds—was taken as evidence for the existence of a separate visual short-term memory store.

Additional evidence for the existence of a short-term memory store specific to the visual domain was found in the ability to compare stimuli unlikely to be verbally encoded across retention intervals longer than the persistence of iconic memory. Cermak (1971) showed that complex radial figures, created by the addition of Fourier components³, were able to be retained and compared across timescales far exceeding the rate of decay for iconic memory traces. The nature of the stimuli allowed a strong *a priori* argument to be made that iconic representations were not being verbally recoded. Likewise, Phillips and Baddeley (1971) asked observers to compare two non-declarative random block patterns—matrices of 5×5 black and white squares in random arrangements—separated by a retention interval and an interceding separate random interrupting mask. Phillips and Baddeley showed that, while change detection performance was well above chance after a retention interval of three seconds, performance decreases substantially after this.

The design employed by Phillips and Baddeley was also employed and extended upon by Phillips (1974). These studies form the basis for all modern visual short-term memory investigations employing *change detection*: a series of stimuli, forming the *memory array*, are presented to the observer for a fixed duration; the observer must retain these stimuli within memory over the presentation of a backwards interruption mask (Di Lollo, Lowe, & Scott, 1974; Kahneman, 1968) and a retention interval (sometimes called the *stimulus onset asynchrony*, or SOA); finally, the observer must indicate whether a second series of stimuli, the *probe array*, matches the first array or not.

Phillips (1974), in using manipulations to the random block pattern change detection task and expanding the earlier change detection work conducted with Baddeley, provided the first comprehensive account of the visual short-term memory store. In particular, Phillips described four properties of visual short-term memory that form the *sine qua non* of any definition for the store. First, the store appeared highly limited in its ability to maintain complex information: increases in the complexity of the random block patterns led to decreases in observer accuracy, with change detection performance under ceiling but high for 4×4 matrices (around 75% accuracy at a retention interval of nine seconds) to scarcely above chance for 8×8 matrices at a retention interval of three seconds. Second, the information that was maintained in visual short-term memory ap-

³These stimuli are now known as “radial frequency” stimuli (Habak, Wilkinson, Zakher, & Wilson, 2004; Shepard & Cermak, 1973; Smith, Lee, Wolfgang, & Ratcliff, 2009; Wilkinson, Wilson, & Habak, 1998).

peared to be stable for much longer periods of time than would be attributable to iconic memory. Performance without masks was very high for comparisons of any complexity across very small (less than a second) retention intervals, with a sharp decrease in accuracy proportional to the complexity of the display at a one second retention interval, and a gradual decline in performance thereafter. Third, performance in change detection did not seem to be affected by presentation of a checkerboard mask immediately after the presentation of the memory array. Last, performance at longer retention intervals was not affected by movement in the location of the probe array relative to the memory array, indicating that the operations on the visual information stored within the short-term memory store were not contingent on maintaining retinotopic spatial arrangements.

These findings provided evidence for the capacity of the visual perceptual system to retain a reduced amount of information (what Phillips called a “schematic” representation of the visual scene) with some independence of spatial positioning from the original stimulus. Representations within this memory system were found to be robust to subsequently presented visual stimulation, and could be retained on the order of seconds after the cessation of the stimulus. Pashler (1988) demonstrated that these same properties were obtained, unchanged, in the case of familiar letters rather than non-declarative stimuli: after the involvement of iconic memory was precluded by using longer retention intervals, the properties of memory were consistent across masking conditions, different stimulus exposure durations, or whether the letterforms were reflected horizontally or not.

As these studies were being carried out, the characterisation of short-term memory of Atkinson and Shiffrin elaborated upon by Baddeley and Hitch (1974), who argued for a new view of short-term memory that was composed of multiple interacting systems operating on different modalities and controlled by a central “executive control” system, known as *working memory*⁴. Their idea of a working memory system—including a proposed system specifically for visual storage—has largely replaced the more general concept of short-term memory, especially in clinical fields, and informs many of the current debates regarding visual short-term memory, particularly that of storage capacity.

⁴For a more thorough review and discussion of the concepts, see recent reviews by Baddeley and colleagues (Baddeley, 1992, 2012; Baddeley & Sala, 1996).

2.3 The storage capacity limits of visual short-term memory

A central line of inquiry in the subsequent examination of visual short-term memory has been the capacity limitations of the system to maintain multiple representations simultaneously. The question of storage capacity has taken a primary position in the visual short-term memory literature as a route to examine the fundamental architecture of the memory system.

A great deal of attention on the question of visual memory storage capacity derives from a series of influential empirical studies reported by Luck and Vogel (1997). Luck and Vogel adapted the change detection paradigm of Phillips (1974) to examine how the composition of items in the memory array affect change detection performance, as measured in terms of proportion of trials correctly identified. Rather than manipulating the configuration of patterns within a random block matrix, Luck and Vogel used suprathreshold and spatially separated objects designed to be highly distinct. The principal finding of the experiments reported was that change detection performance was affected by the number of objects in the memory array, rather than any of the visual characteristics that comprised those objects—the *features* of the objects⁵. This finding was obtained regardless of whether the constitutive features of the object were relevant to the decision at hand, relevant to change detection performance, or not. In more concrete terms, observer performance did not significantly differ in cases where detection of an item in the memory array was contingent upon a single feature value (for instance, the shape of the item, orientation, colour, size, etc.), against trials where monitoring multiple feature values was crucial for accurate performance. Changes in observer accuracy were only seen with manipulations in the total number of distinct items in the memory array, the total number of items to be remembered. In particular, a decrement in the proportion of correct responses was only seen, on average, in the case where more than four items were required to be maintained simultaneously to make a change detection judgement. This basic result was extended, in one of their experiments, to a case where one or more of four different visual characteristics (colour, orientation, size, etc.) could change in memory arrays of two, four, or six distinct items, with no apparent difference between observer change detection sensitivity seen. This result also extended to items composed

⁵The term feature in the realm of visual perception refers broadly to the constitutive perceptual components of a visual entity, usually defined by reference to either the distributed representation of visual items by neural assemblies that operate as selective filters, following Hubel and Wiesel (1959, 1963), or by reference to the components used for forming stimulus judgements, following Garner and Felfoldy (1970).

of two values of the same feature—specifically, a coloured square inset within a differently coloured square—where observer accuracy was not significantly different between conditions where either or both of the colours could change. Subsequent studies conducted by Wheeler and Treisman (2002) and Olson and Jiang (2002) have demonstrated, however, a significant decrement in observer performance when using squares composed of two colours, contrary to the findings of Luck and Vogel⁶.

In all experiments, any in the majority of change detection research that followed, the use of high-contrast, highly discriminable stimuli provided a strong *a priori* case to argue that any performance decrements observed were the product of memory capacity limitations along, rather than limitations in attending to or encoding stimuli into memory. This assumption was further supported by absence of significant change in observer performance with changes to the stimulus exposure duration (from 100 ms to 500 ms); this point will be discussed in further detail in a following section (see §2.4.2).

The selection of visual information into memory and its subsequent retention at the level of objects, rather than some other unit of visual information, has a precedent in the visual perception literature. Vogel, Woodman, and Luck (2001) noted a strong affinity between the result that distinct visual objects are the apparent quantum of visual short-term memory capacity and the earlier work of Duncan (1984) which demonstrated that observer performance for reporting multiple visual features in a display is improved when the visual features belonged a single distinct visual entity⁷. The selection benefits for visual features when grouped into objects lends support to the broader claim that visual representations are organised, at least partially, at the level of the object. This construction of visual short-term memory was also presaged by Kahneman and Treisman (1984) and expanded later by Kahneman, Treisman, and Gibbs (1992) in their argument for *object files*: “episodic” groupings of feature information, bound together in memory as objects to maintain identity and continuity over time.

The invocation of object files as a product of separate visual features bound together is analogous, as discussed by Wheeler and Treisman (2002), to the process of “chunking” (Miller, 1956) from domain-general working memory. More generally, the conceptualisation of visual short-term memory objects as unitary but discrete representations further strengthens the link between the change detection tasks presented above and

⁶Some consensus has been formed following subsequent experimental investigation that there may be smaller feature-level effects on performance in addition to the established large object-level memory array size effects on participant performance (Fougnie, Asplund, & Marois, 2010; Woodman & Vecera, 2011).

⁷A similar result was reported by Treisman, Kahneman, and Burkell (1983) who showed that object groupings lessen the cost of response time incurred by distractor information.

2.3 The storage capacity limits of visual short-term memory

domain-general short-term memory tasks where, such as the digit span (J. Jacobs, 1887) or *n*-back task (Kirchner, 1958), where the elements to be remembered are discrete articulatory symbols. In a notable review on capacity limitations in various modalities of working memory, Cowan (2001) argued for a unified conception of working memory, based on a discrete object-based limit of around four item “slots”, basing the argument on the striking consistency in estimates of memory capacity between modalities. This view of working memory as possessing consistent storage capacity constraints was formed not only on the change detection work of Phillips (1974), Pashler (1988), and of Luck and Vogel (1997), but also on the object-tracking work of Pylyshyn and Storm (1988) showing an object-based limit to the number of objects that could be tracked concurrently and the subsequent visual object enumeration work of Trick and Pylyshyn (1994) which suggested a qualitative change in object enumeration beyond four items. Adapting the work of Pashler (1988), Cowan proposed a relationship between the observed proportions of hits and false alarms in change detection experiments to estimate the underlying item capacity of visual short-term memory, making the assumption that performance is determined in the first instance by whether an item has been stored in the memory system and, then, by guessing if the item is absent from memory.

Although this view of memory unifies the visual short-term memory literature with the domain-general working memory literature, it was contested in subsequent empirical work. Wheeler and Treisman (2002), as mentioned above, did not replicate a key result found by both Luck and Vogel (1997) and Vogel and colleagues (2001), that displays of perceptual objects that are composed from two values of the same feature—in this case, two coloured squares, one inset another, called *colour-colour conjunctions*—show equivalent retention capacity limitations to displays of single-feature objects. Wheeler and Treisman found, instead, that observer accuracy on displays for many configurations of colour-colour conjunctions was at the level seen for displays twice as large, implying that each colour is treated in the memory system as a different item. Wheeler and Treisman further demonstrated that, although observers might be sensitive to changes between individual feature values, observer accuracy substantially diminishes when feature values are swapped between items in the memory array. This finding indicates that, while all of the features of an object might be stored within the memory system simultaneously, this does not imply that there exists a single representation that binds these features together into a unitary object.

The utility of a discrete item-based account of memory capacity decreases substantially without an operational definition for itemhood, if it does not encompass the binding of constituent components into a representational whole (as argued by Wheeler &

Treisman, 2002) or cannot be intuited from the figure itself (as argued by Sakai & Inui, 2002). Alternative accounts of visual short-term memory capacity limitation have emphasised the role of information in visual stimuli at a level below whole objects, arguing for the sensitivity of the system to feature content. Alvarez and Cavanagh (2004) attempted to examine the relationship between the complexity of visual objects and storage capacity, the rate at which observer accuracy declined as the memory array size increases. Multiple classes of stimuli were used in a change detection paradigm—including coloured squares and letters (following earlier change detection work), shaded cubes, randomly generated closed polygons, Chinese characters, greyscale isometric cubes, and line drawings of common household objects. The estimated storage capacity varied substantially between different stimulus types, with estimates of a memory capacity of over four items for coloured squares and around one and a half items for the greyscale shaded isometric cubes. These storage capacity estimates from the change detection tasks, using the Pashler–Cowan formula, were then compared to the corresponding observer performance in a visual search task. Specifically, visual search performance was examined as a function of the observer response time against the number of distractors in the search display; a measure known as the *search slope*. A search slope is a measure of the efficiency of visual search, the ability to quickly identify the target among distractors of the same stimulus class, and is often used to make inferences about the nature of stimulus representation⁸ (Treisman & Gelade, 1980). Alvarez and Cavanagh found a striking correspondence between the estimated storage capacity from the change detection task and the search slope from the visual search task, which they imputed as relationship based on the visual complexity of each of the stimulus types.

The conclusions drawn from this relationship have been controversial, however. Eng, Chen, and Jiang (2005), in a replication of the study, found that changing the stimulus exposure duration minimised the correspondence between memory capacity and visual search slopes, suggesting that the encoding of stimulus information and formation of a memory representation might be the limiting factor. Eng and colleagues also showed that this relationship further declined with subsequent practice sessions, again implicating an attentional or encoding stage rather than the memory maintenance stage itself. A similar conclusion was reached by Awh, Barton, and Vogel (2007), and later Barton, Ester, and Awh (2009), who reported that, when the inter-item (class) similarity was controlled, complexity had no effect on visual short-term memory capacity. Awh

⁸Somewhat against better advice (Townsend, 1972, 1990), researchers have used an invariance of response time over search display size to make inferences about the structure of the visual search procedure.

2.3 The storage capacity limits of visual short-term memory

and colleagues conducted change detection experiments using simple stimuli (coloured squares) and two classes of complex stimuli (shaded cubes and Chinese characters). Following Alvarez and Cavanagh (2004), they found considerable decrements in observer accuracy in detecting changes between memory–probe array pairs using complex stimuli when compared to those using simple stimuli. Awh and colleagues extended this result, however, to examine change detection using complex stimuli in cases where the class of stimuli may differ between the memory array and probe array; that is, where a stimulus location may have a change in the class of the stimulus as well as the identity. In these trials where the stimulus class changed between the memory and probe arrays (cross-class trials), observer accuracy was comparable to the detection of changes in simple stimuli. Indeed, in examining observer performance, a strong correlation ($r = 0.88$) was observed between cross-class capacity estimates, using the Pashler–Cowan formula, and simple colour memory capacity estimates. Awh and colleagues argued that if memory storage was the locus of degraded performance with higher item counts, then differing levels of inter-item similarity would only be a limiting factor if an item was stored in memory. They concluded that the substantial decrements in observer performance seen in change detection tasks employing complex stimuli were due to the similarity of items within the stimulus category—that is, due to limitations in the ability to compare two stimuli in the perceptual task—rather than due to any fundamental storage limit of the memory system itself.

The disagreement in the locus of observer performance constraints in a change detection task highlight a limiting factor with drawing conclusions about the characteristics of memory maintenance from the paradigm: behavioural data may reflect constraints in many different mental processes beyond simply the storage capacity⁹ limitation of the memory system in maintaining multiple representations simultaneously. In particular, observer performance in change detection tasks is reliant on the encoding and maintenance of the initial presentation, the comparison of the active memory representations to the probe array, and the decision process to map the output of a comparison to a response. Questions relating to the encoding stage are often examined through manipulations to the stimulus exposure duration (as done by, for instance, Vogel, Woodman, & Luck, 2006). The nature of the comparison procedure, however, is often not subjected to isolated theoretical analysis¹⁰. Likewise, any decision stage after the comparison is also

⁹It is more correct to specify the capacity constraints under discussion here are *asymptotic* limits on storage capacity. The constraints on storage capacity also show dynamic variation over the course of the retention interval, as will be discussed in §2.4.1 and §2.5.

¹⁰There are exceptions to this, such as the study conducted by Hyun, Woodman, Vogel, Hollingworth, and Luck (2009) and work in transsaccadic memory, which will be reviewed in the next chapter.

rarely considered in isolation.

In total, work in change detection has not reached a consensus on the essential structure of the memory system. Partially, this lack of resolution may come from a lack of consideration regarding the psychological processes that map the internal representations to patterns of responding. Wilken and Ma (2004) investigated this relationship in the context of change detection. In doing so, they provided two important contributions to the literature: first, they considered the predictions of the Pashler–Cowan formulation of object-based memory limitations on observable ROC curves; second, they introduced another experimental paradigm into the visual short-term memory literature, *continuous report*. These two themes will be examined in more depth in the next two subsections.

2.3.1 ROC based analyses

Wilken and Ma (2004) noted that the discrete object models of visual short-term memory capacity (those argued for by, e.g., Awh et al., 2007; Barton et al., 2009; Cowan, 2001; Luck & Vogel, 1997; Vogel et al., 2001), are analogous to *high threshold* models (e.g., Luce, 1963), once popular in psychophysics to describe observer sensitivity in simple detection tasks, but have been shown to be difficult to reconcile with observer data, as discussed below. The Pashler–Cowan formulation of item capacity implies an “all-or-none” storage of items: either stimulus material is encoded into the memory system, if spare capacity is available, leading to a correct response; otherwise, the observer must guess. This conjecture predicts that observed hit rate should be proportional to the false alarm rate over the entire range of response biases for each given memory array size. When the hit rate is plotted as a function of the false alarm rate, a graphical display known as the *receiver operating characteristic* (or ROC) curve, storage systems of these types would predict a characteristic linear function.

High threshold models in sensory detection, with the same “all-or-none” representation of the encoding process, also predict these linear functions. These predictions are not observed in sensory detection tasks, however (e.g., Swets, 1961; Swets, Tanner, & Birdsall, 1961; Tanner & Swets, 1954b, 1953, 1954a)¹¹. Instead, ROC data in detection tasks are curved, consistent with a theory of representation and decision that posits—in the case of simple detection—responses are made on the basis of an internal criterion value providing a response threshold over perceptual signals suffused with Gaussian distributed noise. These assumptions form the basis of signal detection theory (Green & Swets, 1966; W. W. Peterson, Birdsall, & Fox, 1954), and will be discussed in greater detail in §3.1.

¹¹See also the excellent review by Luce and Green (1974).

2.3 The storage capacity limits of visual short-term memory

Consistent with the curved ROC functions found in simple detection, as well as other psychophysical procedures, Wilken and Ma found that observer data in a visual short-term memory change detection task did not conform to the linear pattern expected by the Pashler–Cowan discrete object model. They instead found that ROC data was curved when participant confidence was used as a means of examining the trade off between hits and false alarms. They showed that this pattern of results conformed to predictions of two species of “differencing” model taken from the signal detection theory literature (Noreen, 1981; Sorkin, 1962; to be introduced and discussed in §3.1.2); in particular, Wilken and Ma found support in a majority of cases for a maximum absolute difference (MAD) model. From these model fits, they concluded that the curved ROC data obtained from the change detection paradigm supported a continuous resource account of visual memory storage, in agreement with the conclusions of Alvarez and Cavanagh (2004) and Eng et al. (2005).

The finding that ROC functions were curved over responses conditioned on levels of confidence for a single set size has been subsequently challenged. Rouder et al. (2008) and, later, Rouder, Morey, Morey, and Cowan (2011) demonstrated striking empirical support for a discrete item limit, with the ROC curves showing a linear tradeoff between the hit rate and the false alarm rate with different levels of change probability, rather than confidence ratings, in conditions where the memory array was composed of two, five, and eight items. A model comparison between the discrete item Pashler–Cowan model against a signal detection theory model unanimously supported the fixed-capacity storage model. Although these results could be the result of systematic guessing behaviour, they mean that an “all-or-none” account cannot be dismissed on the basis of ROC data.

The lack of theoretical resolution from work employing the change detection paradigm—either when examining accuracy data or ROC functions—has meant that other experimental procedures have been employed to further constrain the debate; principal among these additional paradigms is the continuous report task, also following from Wilken and Ma and reviewed in the next section.

2.3.2 Empirical studies of storage capacity using continuous report

The support of a continuous resource model of memory storage based on change detection ROC data was the central theoretical contribution of the work by Wilken and Ma (2004). This conclusion, as shown in the last section, has been challenged. A methodological contribution of their work, however, has been adopted in subsequent visual memory experiments: in addition to using a standard change detection paradigm, Wilken

and Ma also employed a *continuous report* task (Prinzmetal, Amiri, Allen, & Edwards, 1998; Prinzmetal, Nwachuku, Bodanski, Blumenfeld, & Shimizu, 1997). Continuous report tasks bear a strong resemblance to experimental paradigms previously used in visual and auditory psychophysics, *methods of adjustment*, where an observer must replicate one characteristic of a presented stimulus standard by manipulating a second stimulus using a continuous means of input.

Wilken and Ma (2004) conducted three experiments using a continuous report paradigm: one requiring the reproduction of the colour value of a presented square; one requiring the reproduction of the orientation of a presented Gabor patch; and one requiring the reproduction of the spatial frequency of a presented Gabor patch. In each, the observer was first shown a memory array of stimuli, with the number of stimuli differing between two and eight items on each trial. After a retention interval of 1,500 ms, the observer was then shown a report cue over the location of the stimulus to reproduce. A series of candidate values of the focal stimulus attribute (colour, orientation, spatial frequency) were presented, and the observed instructed to select the value closest to their memory of the target stimulus.

The data produced from these experiments were examined in the form of distributions of the selected stimulus values, computed as the difference from the actual stimulus value. The data obtained by Wilken and Ma showed that the dispersion of these reproduction error distributions, measured in terms of the standard deviation of the observed responses, increased smoothly in all three conditions. They argued that this supported their primary conclusion regarding the nature of decision-making in change detection tasks, that storage could best be conceptualised as a continuous model of perceptual encoding suffused with noise.

Later studies have followed the methodological example of Wilken and Ma, usually utilising colour as the stimulus attribute to reproduce and requiring responses on a circular domain (a colour wheel of pure hue values). Two key advantages are claimed with the use of a continuous report paradigm over existing change detection tasks: first, the continuous nature of the response entry and the continuous nature of the stimulus presents a putative link between the nature of the underlying representation and the nature of the observed response; and, second, there is an increase in the resolution of the response data, allowing a more detailed specification of theoretical predictions. Although Wilken and Ma argued that their continuous report data supported a continuous resource model of memory storage, and that it was a more direct measure than the change detection ROC data they did examine against specific model predictions, they did not attempt any quantitative model fitting on their reproduction error data. Subsequent studies have employed

direct model fitting to examine model predictions against the observed data.

Bays and Husain (2008) extended the work of Wilken and Ma to examine colour, orientation, and location (displacement) reproduction errors using Gaussian distributions fitted to response distributions, although using a two-alternative task. From the variance of a Gaussian distribution fitted to observer data over different memory array sizes conditions, Bays and Husain concluded that memory representations shared a common resource both at the level of objects and at the level of features¹². On the basis of an eye movement manipulation, they further concluded that the allocation of this resource to different items and features within the display was mediated through the allocation of spatial attention.

The conclusions of Wilken and Ma and of Bays and Husain that the monotonically increasing dispersion of reproduction errors with increases in memory array size has been contested with subsequent model fits. The additional resolution of data from the continuous response task has meant that model specifications of greater sophistication can be more readily distinguished between than might otherwise be the case in change detection tasks. Zhang and Luck (2008) examined data from a colour reproduction continuous report task over differing memory array size conditions against model predictions from a simple mixture model. Their mixture model included a “guess” component, modelled as a uniform distribution across the entire stimulus domain, as well as a “memory” component distributed as a von Mises distribution centred on the actual stimulus value¹³. Examining the maximum likelihood estimates of the mixture models, Zhang and Luck showed a discontinuity in the dispersion of the von Mises component with an increase in the memory array size: where there was an increase in the dispersion in the response distribution between one and three items, there was no increase between a memory array of three items and one of six items¹⁴. The uniform, “guess”, component of the mixture model, however, increased linearly.

Zhang and Luck concluded that their data supported a “slots + averaging” account: a hybrid model combining a discrete item limit with a flexibly allocated resource. In this model, visual short-term memory is characterised as able to simultaneously maintain a

¹²Bays and Husain indicated that the bilinear function of accuracy used to posit the existence of a discrete items memory limit, following Luck and Vogel (1997), may be a function of the suprathreshold nature of the change detection decision to be made. This conclusion is consistent with the findings of Sewell et al. (2014), discussed in §2.5.

¹³The von Mises distribution is similar, but not identical, to a Gaussian distribution on the circle. It is described by two parameters: one location parameter and one concentration parameter, operating in an analogous way to the mean and standard deviation of a Gaussian distribution, respectively.

¹⁴Some caution about these results are warranted, given that only one memory array size condition was tested that exceeded the putative item limit given by previous papers.

fixed number of visual objects. When the number of items to be maintained exceeds this fundamental limit, the number of “slots”, then performance is contingent on whether the item to be accessed was stored, as per the slot account of Luck and Vogel: performance is based on memory if the particular item required was stored, otherwise a guess is made. When the number of items to be maintained falls below this upper limit, however, it is predicted that a single item can be held within multiple unoccupied slots, effectively increasing the fidelity with which an item can be reproduced. Zhang and Luck argued that this account predicts small changes in the fidelity with which items can be reproduced below some (small) item limit but predicts no change in reproduction error beyond a certain rate, only an increase in the rate of “guess” based errors. These were consistent with the mixture model parameters obtained through maximum likelihood estimation.

This account was also supported by Cowan and Rouder (2009), who directly disputed the conclusions of Bays and Husain (2008). Cowan and Rouder noted that the discrete object-based model proposed by Bays and Husain was not exactly the “slot + averaging” model proposed by Zhang and Luck: specifically, it lacked the “averaging” mechanism seen in memory array sizes under the upper item limit, which allowed a flexibility in the precision when dealing with small displays. With this modification, the model accounts for the data obtained by Bays and Husain, as well as the data obtained by Rouder et al. (2008), as discussed above. In response, Bays and Husain (2009) extended their study to use a larger range of stimulus (displacement) values to demonstrate that a “slot + averaging” model could not account for data from larger memory array sizes¹⁵ They also noted that while flexibly allocated resources had some claim to a plausible neural realisation, through population coding (Pouget, Dayan, & Zemel, 2000), the same was not evident in the case of the “slots + averaging” model. A subsequent model analysis by Bays, Catalao, and Husain (2009) demonstrated that, when non-target information is also taken into account in modelling continuous report data, the uniform “guess” component of Zhang and Luck appears to be substantially diminished and the von Mises “memory” component again smoothly increases in dispersion—following a power law function—as the memory array size increases.

The continuing debate over the best characterisation of visual short-term memory organisation, and the lack of a simple resolution, has led to consideration of more complex models. Van den Berg, Shin, Chou, George, and Ma (2012) expanded the class of models previously applied by Zhang and Luck in examining behaviour in the continuous report

¹⁵The treatment of displacement as a stimulus value comparable to colour and orientation is somewhat interesting, given that spatial context is critical in defining the limits of objecthood and that visual short-term memory appears to be fundamentally mediated by some sort of spatial encoding (see, for instance, Jiang, Olson, & Chun, 2000).

2.3 The storage capacity limits of visual short-term memory

paradigm. Van den Berg and colleagues discarded the constraint that memory could be described as a fixed, scalar quantity describing either the item limit—where, at most, a small number of items possess the same fidelity in memory—or fixed resource pool—where there is an inverse relationship between the number of items in memory and the fidelity at which any item is stored. Instead, they proposed a doubly-stochastic model: the response observed is drawn from a von Mises distribution with a dispersion parameter drawn from another distribution (selected to be a gamma distribution for flexibility). In Bayesian model fitting and comparison using root-mean-square error¹⁶, the doubly-stochastic model was able to capture the maximum likelihood estimates of a Zhang and Luck-style mixture model (viz. the circular standard deviation and the mixture component). These conclusions were also supported by studies conducted by Keshvari and colleagues in examining change detection (Keshvari, van den Berg, & Ma, 2012, 2013). In a more recent paper, van den Berg, Awh, and Ma (2014a) expanded consideration of a variable precision doubly-stochastic model by way of a large-scale factorial comparison of Bayesian models of continuous report over ten continuous report datasets (e.g., Bays et al., 2009; van den Berg et al., 2012; Wilken & Ma, 2004; Zhang & Luck, 2008)¹⁷. Combinations of both the form of item limits (e.g., a set number of items remembered, a flexible number following a Poisson distribution, etc.) and the form of the precision of the item representations (e.g., a set precision shared amongst items, a “slots + averaging” mechanism of precision, variable precision, etc.) were tested, as well as models with and without non-target responding (i.e., Bays et al., 2009). Using improvements in AIC as the model selection criterion, they concluded that a model where a variable number of items (following a Poisson distribution) are encoded, each with a (gamma distributed) variable precision and which might lead to non-target responding on some trials, best fit the data.

The recent work from van den Berg and colleagues, although computationally intensive, is open to some amount of criticism. First, although ameliorating one problem of the earlier model comparison reported by van den Berg et al. (2012) by using an objective function that had some sensitivity to the complexity of the model under consideration, both AIC and BIC may be inappropriate for use as a model selection criterion when comparing models with wholly different functional forms¹⁸. Work by Myung, Pitt, and col-

¹⁶A goodness-of-fit measure without any correction for model complexity.

¹⁷Three datasets from the University of Oregon under Anderson, Awh, and colleagues (Anderson & Awh, 2012; Anderson, Vogel, & Awh, 2011) have been subsequently retracted due to inappropriate data handling procedures. The main conclusions, and limitations, regarding the work of van den Berg and colleagues are unlikely to be compromised on this basis.

¹⁸The work of the current thesis also engages in model comparisons using the BIC, although changes in parameters in the current work usually correspond to allowing flexibility in model fitting across different experimental manipulations rather than the addition of a hyperprior across the entire experiment.

leagues (Myung & Pitt, 1997; Navarro, Pitt, & Myung, 2004; Pitt, Kim, & Myung, 2003; Pitt & Myung, 2002; Pitt, Myung, & Zhang, 2002) highlights the complexity of conducting competitive model fitting exercises using Bayesian models with different levels of functional flexibility. It may be unsurprising that the most flexible model was ranked above others, and with the more flexible models ranked generally more highly than those with less flexibility. Second, the models are defined on the measurement level: a variable precision, variable object-number model is not given in terms of a proposed mechanism, but simply as an output. Although some justification for the variable precision nature was given in terms of populations of Poisson spiking neurons under divisive renormalisation (see Ma, 2010), the correspondence between model parameters and any proposed interpretation is somewhat loose without subsequent direct testing. Outside of experiments conducted with the express purpose of direct manipulation and characterisation of each of the parameters, it is difficult to know whether flexibility in the parameters is required due to otherwise unaccounted for, but theoretically unrelated, variability in the observed responses or is related to the proposed mechanisms.

2.4 The temporal dynamics of visual short-term memory

The work reviewed in the last section provided a short overview of the most intensive line of inquiry among the various topics examined in the broader visual memory literature, the capacity of the visual short-term memory system to hold multiple representations active simultaneously. All accounts of these storage capacity limitations—slot-based accounts, resource-based accounts, or some hybrid of the two—cast these limitations as being stable properties of the system: the maximum capacity of the system to hold information beyond any initial encoding constraints. This follows the original visual short-term memory results of Phillips (1974), where memory representations can be kept active and available for decisions made after an extended interval, on the scale of seconds after the presentation of the stimulus. Such an account is held in contrast to the rapid decay of information seen in iconic memory tasks and, like the item-based accounts of the memory system reviewed in the last section, is seen as analogous to the rehearsal of verbal information (like lists of digits) in general cognitive tasks.

Yet, even as investigations regarding the storage capacity inform theoretical accounts of the structure of visual short-term memory, examination of the how observable performance characteristics of the memory system might change over time has been underemphasised. Of the research examining the time course of visual short-term memory

representations, two themes have emerged: the stability of representations over different retention intervals; and, the initial entrance of sensory information into the memory system and its encoding into a robust representation. We shall examine these in the following sections.

2.4.1 Maintenance over the retention interval

The invariance of observer performance in change detection tasks over different retention intervals indicates that some representation of visual information, accurate enough to be used in detecting changes, persists seconds after the removal of the original sensory input. Indications that some visual information persisted beyond the decay of iconic memory was first demonstrated by Posner and colleagues (Posner, 1967; Posner et al., 1969), as well as by Kroll et al. (1970), Cermak (1971), and Phillips and Baddeley (1971). Phillips (1974) provided the first comprehensive demonstration of the extent of this memory system: indicating that a small number of items could be retained without apparent loss of overall accuracy for retention intervals exceeding five seconds. This extends to discrimination of the constituent features of simple visual objects over long retention intervals, such as the spatial frequency of luminance gratings (Magnussen, Greenlee, Asplund, & Dyrnes, 1990; Regan, 1985) or the velocity of drifting luminance gratings (Magnussen & Greenlee, 1992).

The ability for the memory system to maintain a visual memory representation may not be without some dependency on the content of the visual memory representation itself. Using an identical change detection task to Phillips (1974), Paivio and Bleasdale (1974) showed that some information could be retained for longer timescales—performance in a condition with large-scale changes was undiminished with a thirty second retention interval—but showed that greater similarity between the memory and probe arrays led to a faster decrease in accuracy as the retention interval increased.

Paivio and Bleasdale also found that mean response times continued to increase with increasing retention intervals, even after a plateau in the error rate; they interpreted this result as indicating a loss of accessibility of the memory representation, rather than the complete cessation of maintenance.

Like Paivio and Bleasdale (1974), Cornelissen and Greenlee (2000) also found a relation between change detection performance using displays of 10×10 random block patterns (both colour and luminance), and the complexity of the stimuli being stored. In an experiment using near-threshold stimuli, a linear relationship between observer change detection sensitivity and memory–probe stimulus similarity was found: performance

increased as similarity between the memory stimulus and probe stimulus decreased. Cornelissen and Greenlee showed a shift in observer change detection sensitivity, for all levels of complexity, when comparing between retention intervals of one and three seconds. This uniform downward shift in sensitivity was used to conclude that a slow decay process was occurring.

The conclusion that steadily decreasing accuracy with increases in the retention interval represented a slow decay of representations from visual short-term memory representation was challenged by Zhang and Luck (2009). Using both a colour and radial frequency continuous task, Zhang and Luck attempted to fit the mixture model of memory introduced in their earlier paper discussing a “slots + resource” account of memory (described in §2.3.2). Rather than showing a slow decrease in either the ability to access memory traces as argued by Paivio and Bleasdale, or a slow increase in the dispersion of responses indicating a loss of memory fidelity, the maximum likelihood parameters showed almost no increase in the variance of the von Mises distributed “memory” component of the model, but an increase in the rate of guessing (a uniform distribution across the entire response domain) from a retention interval of four seconds to one of ten seconds. Zhang and Luck interpreted this as the “sudden death” of memory traces, arguing that models of memory represented as dynamical systems could display this “all-or-none” behaviour if the system state was displaced from a stable attractor state of maintenance—a prediction of many recurrent neural network models of active memory (Barak & Tsodyks, 2014; Camperi & Wang, 1998; Wang, 2001; Wei, Wang, & Wang, 2012; Zipser, Kehoe, Littlewort, & Fuster, 1993).

One potential confound of such a finding, acknowledged by Zhang and Luck and foreshadowed by Paivio and Bleasdale, was that the apparent increase in guessing may have been due to a decrement in the accessibility of the memory representations—for instance, a loss of the location–item binding—rather than a destruction of the memory information itself. Donkin, Nosofsky, Gold, and Shiffrin (2014) attempted to address this ambiguity by conducting a colour-based continuous report study using only single-item memory arrays, which obviate the need to use location–item mappings. Donkin and colleagues also anticipated that employing single item displays might confound any conclusions drawn regarding the memory system if observers verbally recode the stimulus, a process that might be expected to occur less readily with larger memory array sizes. In order to account for any effect of verbal labelling, pairs of “labelling” trials were conducted, where observers were asked to label each memory item and then, some trials later, replicate the colour patch—using the colour wheel reporting method of the memory trials. Extending the Zhang and Luck two-part mixture model, Donkin and col-

leagues added a third Gaussian component representing the contribution of verbal labelling to the obtained responses and fixed to data from the verbal labelling conditions. They concluded that, on the basis of model fits with complexity taken into account (via AIC), a model with stable precision estimates leading to changes in the guess rate—the “sudden death” model following Zhang and Luck—better accounted for the data than a model with gradual changes to the variance of the memory component, but only when verbal labelling was taken into account.

These studies indicate that memory representations appear to be maintained with considerable stability until either a response is made or the memory representation fails. The process of maintenance is not entirely without cognitive control, however. Griffin and Nobre (2003) demonstrated that presenting a symbolic cue with 80% validity during the retention interval of a change detection task led to a significant and substantial increase in accuracy for valid trials and decrease for invalid trials compared to trials where no cue was presented. The valid “retrocue” condition also produced substantially faster mean response times when compared to trials without a cuing condition. Both the accuracy and response time shifts were comparable to the effect of an analogous cue presented prior to the memory array to shift spatial attention. An analysis of ERP waveforms suggested large similarities in the neural signature of both pre-cuing and retrocuing, consistent with the overall pattern of orienting spatial attention. Differences in the ERP waveform between pre-cuing and retrocuing were found early in retention interval in prefrontal cortex, with the retrocuing condition having enhanced positive activation, and later in the retention interval in the posterior parietal and occipital cortices. Griffin and Nobre suggested these represented an increase in the control processes of the memory store and in the maintenance of the visual information itself, respectively. These conclusions largely agree with functional neuroimaging results (Curtis & D’Esposito, 2003; Pessoa, Gutierrez, Bandettini, & Ungerleider, 2002; Todd & Marois, 2004; Xu & Chun, 2006) and later studies (Nobre et al., 2004).

The retrocuing effects reported by Griffin and Nobre have prompted additional investigations into the benefits of such cues. Matsukura, Luck, and Vecera (2007) found that, while change detection performance decreases as the duration of the retention interval increases in trials without any retrocue presented, trials in which a perfectly informative retrocue is displayed to the observer did not lead to any difference in performance between a short retention interval (500 ms) and a longer retention interval (2000 ms). They concluded that retrocuing did not prioritise retrieval of information from the memory store, but rather protected the representation from decay. This conclusion is consistent with the finding reported by Makovski and Pertzov (2015) that retrocues not only protect

memory representations from deletion, in the style of Zhang and Luck, but also prevent location–object binding errors, of the type reported by Bays and Husain (2008)¹⁹.

2.4.2 Encoding

The other stream of visual short-term memory research examining the temporal dynamics of memory representations has focused on their initial formation from sensory data. Evidence from the partial report tasks conducted by Sperling and colleagues (Averbach & Sperling, 1961; G. Sperling, 1960, 1963) indicates that durable visual short-term memory representations are not formed instantaneously upon the exposure of a stimulus, but take some time to form. This is indicated by the change in observer performance when adjusting the length of the stimulus exposure duration prior to presenting an interruption backwards mask. Observer performance when a mask is displayed within milliseconds of the stimulus onset is at chance levels. Performance increases substantially, but subject to the visual short-term memory storage capacity constraint, when masks are presented hundreds of milliseconds or more after the presentation of the stimulus. The mask itself precludes the involvement of iconic memory representations, indicating that information has been transferred from the sensory and iconic representations into durable visual short-term memory representations.

Gegenfurtner and Sperling (1993) provided a first examination of the rate of transfer from iconic memory representations to visual short-term memory. Using a letter-based partial report task, Gegenfurtner and Sperling manipulated both the delay times before the report cue was displayed to the observer—the audio signal to the observer about which subset of letters would require recall—and, in one experiment, the time after stimulus offset prior to the presentation of a backwards mask. In trials where the report cue was presented immediately after stimulus offset, maximum performance could be obtained by transferring only elements from the cued subset into visual short-term memory as quickly as possible. In cases where the cue was only shown after a longer retention interval, the accuracy functions were consistent with observers attempting to encode, non-selectively, any available information in the display. This was a consistent strategy displayed by observers given sufficient practice on the task: to encode any information available until the report cue was presented, and then focusing only on cued information. The strategy employed by observers was not only due to the limited overall storage capacity of the memory system, but also the limited rate at which this information could be encoded. Gegenfurtner and Sperling demonstrated that the data obtained were well fit by

¹⁹See §2.3.2 for further discussion.

2.4 The temporal dynamics of visual short-term memory

a model of attention, following the spatial–temporal model described by G. Sperling and Weichselgartner (1995) which predicted that the rate of encoding was well described as the product of both spatial proximity to cued location and the stimulus discriminability.

Further evidence for the limited capacity for the process of visual short-term memory encoding has also been indicated by work in rapid serial visual presentation (RSVP) tasks. In the RSVP paradigm, a series of items—usually alphanumeric characters—are displayed in a single location in rapid succession, usually at a rate of around ten items per second, with each new element replacing the last. In some cases, the observer is asked to report the identity of the target stimulus—denoted by a display of a cue or change in colour (e.g., Weichselgartner & Sperling, 1987) or by the display of a concurrent stimulus (e.g., Reeves & Sperling, 1986). A consistent finding of RSVP experiments in which participants were asked to report only a single target item from the stream of stimuli is the asymmetrical pattern of errors: observers often report the identity of the stimulus presented shortly after the display of the target item but rarely report items presented before the presentation of the target (Broadbent & Broadbent, 1987; N. G. Kanwisher, 1987; Reeves & Sperling, 1986; Weichselgartner & Sperling, 1987). The prevalence of these post-target intrusion errors was attributed to the temporal dynamics of an attentional gating process: the speed of presentation was faster than the ability of an attentional process to filter out subsequent irrelevant information, allowing it to reach the visual short-term memory store (Chun & Potter, 1995; Reeves & Sperling, 1986; G. Sperling & Weichselgartner, 1995).

Subsequent work using the RSVP paradigm required observers to report multiple stimuli presented in the stimulus stream. A consistent finding from a number of studies (Chun & Potter, 1995; Raymond, Shapiro, & Arnell, 1992; Shapiro, Raymond, & Arnell, 1994; Weichselgartner & Sperling, 1987) is that, in trials where the observer can and must report the identity of the target stimulus, there is a large deficit in the detection of a probe item presented afterward in the stimulus stream—a phenomenon known as the *attentional blink*. This deficit follows a regular pattern in which the item presented immediately subsequent to the target is detected at the rate that the probe would be detected at without the presence of the target, but probes in positions subsequent to that suffer a large decrement in detection performance until roughly 500–600 ms later. Chun and Potter (1995) proposed a two-stage model to account for these data where, subsequent to a stage of rapid detection where distractors may be filtered out on the basis of surface features, a process of consolidation is undertaken to convert sensory representations into durable visual short-term memory representations and the process of identification and responding begins. Chun and Potter argued that it is this second stage that is limited in

capacity, causing the attentional blink phenomenon as items subsequent to the target that are able to pass the initial filtering cannot be consolidated and identified until the end of processing for the target item.

The constraint in the processing resources available to consolidate memory representations were also identified in a dual-task paradigm by Jollicœur and Dell'Acqua (1998). In their task, a visual array of either one or three characters—either consonants or common punctuation—was presented to the observer, to be retained for the duration of the trial. After the presentation of the characters, and the presentation of a backwards mask, a tone was presented which required an immediate speeded response. Subsequent to the response for the tone, the observer then recalled the letters retained, and typed their response on a computer keyboard. Jollicœur and Dell'Acqua found that, although accuracy for the tone identification and letter report were very high, the response times for the speeded tone response were much slower when the interval between the presentation of the letters and the presentation of the tone was short (around 300 ms), with responses slowly speeding to an asymptote over a second afterwards. These delays were also heightened when displays of three items were to be retained compared to displays of a single item. Jollicœur and Dell'Acqua interpreted this initial delay in responding to the presentation of the tone as the product of amodal capacity limitations caused by the consolidation of the visual representations. These capacity limitations show an overall constraint—impeding the processing of other information—that is, critically, sensitive to the amount of information being consolidated.

The dual-task and RSVP studies presented above demonstrate some limitation in the processing of visual information for identification. The exact contribution of consolidation, however, is unclear as both mix the processes of consolidation with “decision” processes that map this durable representation to a response: the identification of the stimulus and the selection of a response on the basis of this identification. In an attempt to isolate the contribution of consolidation alone, Vogel et al. (2006) used a change detection task—like that used by Luck and Vogel (1997) and Vogel et al. (2001)—where the decision components of the task are separated from the memory consolidation and maintenance components: although consolidation can happen upon the presentation of the memory array, the decision of which response to enter can only begin upon the presentation of the probe array. Vogel and colleagues adjusted both the size of the memory array, between one and four coloured squares, as well as the duration of the stimulus onset asynchrony (SOA) period between the presentation of the memory array (shown for 100 ms) and the mask array (200 ms). Detection of changes to displays of a single item were detected with almost perfect accuracy even at the shortest SOA period,

117 ms; changes in displays of two items were detected with below ceiling performance, roughly 90% accuracy, at this short SOA period, but reached ceiling performance with the longer SOA period of 234 ms. Changes to the larger displays of three and four items were detected with performance far below these smaller displays at shorter SOA periods, with performance gradually increasing as the SOA period was extended. When examining just change detection performance in displays with four items, Vogel and colleagues used a series of very short SOA periods (between 17 and 83 ms), as well as longer delays (167 ms and 217 ms), to demonstrate not only a monotonic increase in the performance as the period in which stimulus information was available increased, but to characterise the rate of this increase. Using the Pashler–Cowan capacity formula, they characterised the rate of performance increases as constituting a linear increase in capacity, of around 50 ms per an item.

A tenfold decrease in the previous estimate of consolidation time, Vogel and colleagues provided a clearer estimate of the capacity limitations aside from processes of identification and response selection. Their interpretation of this limitation, estimated at 50 ms/item, was framed in terms of the “all-or-none” object account of their previous work on overall storage capacity limitations. More recent work by Bays, Gorgoraptis, Wee, Marshall, and Husain (2011) used a continuous report paradigm—following Bays and Husain (2009)—to jointly estimate both the storage capacity limitation and the initial rate of encoding through analysing the variance of a von Mises distribution fit to responses and centred on target stimulus values (reviewed in §2.3.2). Looking at memory array sizes of six items, exceeding the putative “slot” account limits of four items, Bays and colleagues found no hard limit to overall capacity constraints. Instead, they found an interaction between the rate of encoding and the overall memory load, with the precision at which observers could estimate the colour of a target stimulus increasing with the memory array exposure duration and decreasing with an increase in the number of items to retain, both well described by a power law relationship: that is, $P \sim \chi^{-\lambda}$, where P is the precision of the response, χ is either the exposure duration of the stimulus array or the size of the stimulus array, and λ is an estimated scaling parameter.

2.5 An information limit on short-term memory

The previous two sections characterised two major themes in visual short-term memory research: capacity limitations in the ability to maintain representations simultaneously; and, the temporal dynamics of visual short-term memory representations with particu-

lar attention on the constraints seen in the rate of forming representations from sensory data. As evinced by the conclusions drawn by Vogel et al. (2006) or Bays, Gorgoraptis, et al. (2011), these two strands of the literature have been largely seen as reflecting independent properties of the memory system: visual information is first subject to a constraint on the rate of memory formation and then, within memory, subject to overall structural constraints—either at level of visual items or visual information—of the overall holding capacity of the memory system.

Recent work, reported by Sewell et al. (2014), has suggested that the initial constraint on the formation and consolidation of memory representations may also indicate an overall constraint on the information that can be concurrently sustained within the system. Extending the work of Vogel et al. (2006), Sewell and colleagues manipulated both the exposure duration of a memory array and the number of items within that array. Rather than using a change detection paradigm, following a large portion of the visual short-term memory literature, they used a two-choice orthogonal orientation task: observers were presented with a memory array composed of between one and four horizontally or vertically oriented Gabor patches for a duration of between 50 ms and 200 ms. After the presentation of a high-contrast checkerboard mask, the observer was cued to a location on the display corresponding to one of the stimuli in the memory array and required to report on the orientation of that stimulus.

Several design choices distinguished this study from previous work conducted in the visual short-term memory literature. The use of a two-choice orthogonal orientation task and Gabor patches as stimuli, rather than coloured squares, followed a precedent set in the attentional literature (e.g., Carrasco et al., 2000; Lee et al., 1997; Smith & Ratcliff, 2004) due to the low-level properties of Gabor patches²⁰. This allows for a precise characterisation of observer sensitivity without introducing the systematic bias seen in “yes–no” detection tasks. The information in the display available for the formation of a memory representation was also controlled. The contrast of the memory array was presented near the detection threshold, interleaved with dynamic Gaussian noise, and adjusted such that discrimination performance for the smallest memory array size—one item in the display—was near chance for the shortest memory array exposure duration (50 ms) and near ceiling for the longest memory array exposure duration (200 ms). The usage of a post-stimulus probe to indicate the location of the stimulus to be retrieved and identified, rather than allowing a decision to be made about all locations in the dis-

²⁰Thomas and Gille (1979) showed that observer performance on orientation discrimination between Gabor patches with greater in angular displacement of about 20° angular difference was comparable in sensitivity to detection of Gabor patches.

play as done in some change detection tasks, further meant that observer performance would reflect constraints within the memory system without the contribution of additional variability as the uncertainty regarding the location of the stimulus with changed with manipulations of the size of the memory array (Downing, 1988; Shaw, 1982). The precise control of stimulus contrast, the duration for which the stimulus display was available, and the control on decision-level uncertainty using the post-stimulus probe allowed overall inferences about changes in observer performance to be attributed to differences in the amount of stimulus information relevant to decision-making specifically²¹.

The main finding of Sewell and colleagues was that observer performance as the memory array size increased could be characterised in terms of a simple relationship:

$$d'_m = d'_1 / \sqrt{m},$$

where d' represents the signal detection theory measure of sensitivity (the subscript denoting the display size of the sensitivity measure) and m representing the size of the memory array. That is, the observer sensitivity in discrimination for a display where a number of items must be maintained is related to the sensitivity when retaining a single item²². This finding, known as a *sample-size* relationship (Bonnell & Hafter, 1998; Bonnell & Miller, 1994; Lindsay, Taylor, & Forbes, 1968; Palmer, 1990; Shaw, 1980; Taylor, Lindsay, & Forbes, 1967), was unaffected by whether the stimuli in the memory array were presented simultaneously—as in a standard memory task—or sequentially, where each item was shown for a time equivalent to the entire array display time in the simultaneous presentation condition.

The interpretation offered by Sewell and colleagues was that, rather than being a constraint on memory formation, as it had been previously interpreted, this performance constraint was one on the overall rate of information that can be sustained by the visual short-term memory system. An encoding limitation would predict a decline in performance in simultaneous condition, but not in the sequential condition where each item was available for at least 50 ms (and for as long as 200 ms). This interpretation was fur-

²¹Smith (1995) and Ratcliff and Rouder (2000) provided empirical demonstrations from response time modelling that the effective information available for decision-making in simple visual tasks—detection and discrimination between two letters, respectively—can be described as the integration of stimulus contrast with respect to noise over the stimulus presentation time. This finding follows from Bloch's law, in which the observer sensitivity for the detection of a stimulus is the product of the intensity of the stimulus and the duration of stimulus availability (H. G. Sperling & Jolliffe, 1965).

²²An equivalent way of expressing this relationship, $\sum_{i \in 1 \dots m} (d'_i)^2 = c$, where c is a constant and the index i runs across items to the memory display size m , shows that the sum of squared sensitivities for a display equals a constant.

ther demonstrated with a quantitative comparison between a flexible interpretation of an “all-or-none” account of item formation, where the formation of a memory representation occurs instantly as a Poisson event (Bundesen, 1987, 1990; Bundesen, Pedersen, & Larsen, 1984; Townsend, 1972). In almost all cases, quantitative models implementing the sample size relationship were able to account for the data when compared in terms of the BIC (Schwarz, 1978). (A sample-size model without any estimated parameters was also fit and was competitive to a remarkable degree, closely approximating the data.)

This sample-size relationship offers a process interpretation: it is the relationship that is obtained when normally-distributed information is sampled at a fixed rate, leading to a square-root improvement in the signal-to-noise ratio. The relationship is, however, highly restrictive in the behaviour it can predict, and greatly constrains quantitative models of visual short-term memory dynamics. Smith and Sewell (2013) provided a comprehensive model of observer decision-making characteristics, extending the work of Smith and Ratcliff (2009), in a range of visual tasks such as simple detection, visual search, and the two-choice orthogonal orientation task of Sewell and colleagues. The model was based on competitive interactions between representations, modelled using interacting systems of shunting equations (Elias & Grossberg, 1975; Grossberg, 1982, 1987, 1988), to allow a full expression of memory dynamics through time. Recently, Smith, Sewell, and Lilburn (2015) has demonstrated that these relationships can be obtained from plausible neural mechanisms, the integration and normalisation of the difference of pairs of Poisson shot-noise processes, extending a previous result (Smith, 2010). Taken together, these models provide both an account of response time behaviour in visual short-term memory tasks, with neurally realisable basis.

The sample-size constraint on information capacity within visual short-term memory, as well as the findings and methodology of Sewell and colleagues, form the basis of this thesis. The increased precision in the control of the rate of information available for the formation of memory representations allowed by a simple two-choice near-threshold orthogonal orientation discrimination task with post-stimulus probing affords greater specification between the relationship between observer behaviour the properties of both the memory system and decisions made on the basis of information taken from the memory system when compared to high-contrast change detection, where the quantity of information available is not as tightly controlled. The main subject of the first part of this thesis will be a more complete characterisation of the differences seen between change detection tasks and two-choice orthogonal orientation task used by Sewell and colleagues. The results of this first part will then be used to examine an extension to the two-choice near-threshold methodology where the observer must discriminate

2.5 An information limit on short-term memory

between small differences in orientation, allowing characterisation of the precision at which this information is stored.

Chapter 3

Change detection and orientation discrimination: sensitivity

This chapter will examine the relationship between a two-alternative orientation discrimination paradigm and a change detection paradigm. Although these tasks are intrinsically related in terms of stimulus presentation and memory retention, they differ in the decision required. Further defining the requirements of the tasks allows a greater understanding of how the observed responses relate to the underlying memory representations.

The primary focus of the visual short-term memory literature has been on properties of the memory system, rather than the nature of the tasks used. As outlined in the preceding review, a large proportion of the literature has focused on the effects of storage capacity limitations in maintaining multiple representations simultaneously. The principal question arising within this research has been on the question of the architecture of visual short-term memory; specifically, whether visual short-term memory is best characterised as storing visual information at the level of the discrete object, at the information below the level of the object (for instance, at the level of individual features), or some hybrid or hierarchical account (e.g., Brady & Alvarez, 2011; Zhang & Luck, 2008). This question, despite a large amount of empirical investigation, remains unresolved. Part of the resolution of this theoretical point will be found in the quantitative specification of the link between observed data and proposed models of storage.

A common distinction in the division of mental processes is the division between processes encoding, maintaining, and operating upon mental representations and those processes which map internal representations to response categories based on the task demands. The latter category are known collectively as decision processes, those pro-

⁰Portions of this chapter have been previously presented at the 2013 Experimental Psychology Conference held in Adelaide.

cesses that identify psychological (in our case, perceptual) representations—usually considered to be a procedure of statistical inference working on representations that follow a known distribution (Luce, 1977). The distinction was foreshadowed in the mental chronometry work of Donders (1868/1969), and was more fully detailed by Thurstone (1927) in considering experimental tasks where a simple comparison on some specified perceptual dimension was required.

Although many theorists consider the processes that might govern the storage and maintenance of visual short-term memory representations, usually in terms of metaphors relating to the storage architecture such as the discrete object-based “slot” and the flexibly allocated “resource”, a detailed evaluation of the decisions these tasks require is often overlooked. (A notable counterexample, however, is given by Wilken & Ma, 2004.) These considerations may not compromise the overall finding that memory is capacity limited, but might affect the exact degree to which the apparent changes in observer performance are due to manipulations in display size are attributed to constraints on the memory system. Work in visual attention supplies a clear example: although orienting spatial attention improves performance in target detection across many possible locations (Posner, Snyder, & Davidson, 1980), the full extent to which any enhancement in the perceptual representation is due to attention can only be examined when the decisional effects of changing the number of relevant locations—and, thus, the number of potential independent decisions—is taken into account, either experimentally or through explicit quantitative modelling (Baldassi & Burr, 2004; Gould, Wolfgang, & Smith, 2007; Palmer, Verghese, & Pavel, 2000; Shaw, 1982, 1984). In the case of visual short-term memory, the Pashler–Cowan formulation and the mixture model of continuous report performance due to Zhang and Luck (2008), for instance, assume a direct equivalence between response and representation. Yet, any bias in estimate might be particularly pronounced given the small number of experimental procedures used in visual short-term memory work: the change detection paradigm (following Phillips, 1974) and the continuous report paradigm (following Wilken & Ma, 2004) constituting the vast majority of major work in the field.

A large share of the very recent debate regarding storage capacity has focused on data obtained using a continuous report procedure. The continuous report task has grown substantially in popularity, particularly since the work of Zhang and Luck (2008) and Bays and Husain (2008), due to the conceptual similarity between the stimulus attributes being examined—usually the orientation or colour of an item—and the means of responding. Descriptions of the connection between the memory representation and responding have, like change detection, not featured prominently in theoretical discus-

sion. Two descriptions for the possible relationship between observed patterns of responding in a continuous report task to an underlying neurally plausible representation have been provided by Ma (2010) and Bays (2014). These accounts represent perceptual information stored within visual short-term memory using a common model of neural action: populations of orientation-specific neurons modelled as point processes emitting Poisson-distributed spike trains, where the entire rate of population action is controlled by divisive normalisation (Pouget et al., 2000)¹ and, although they have a theoretical parsimony in identifying neural codings directly with probabilistic memory representations, the relationship between observable data and the parameters controlling these representations have not been specified. By way of example, a result by Smith (2016) showed that, whilst continuous report tasks can be modelled within the context of sequential sampling models by using a natural extension of the popular diffusion model (Ratcliff, 1978) to the circular domain, the exact effect of an experimental procedure on observed data cannot be fully identified without response time data².

Although this observation does not preclude a more detailed consideration of the decision properties of the continuous response task, I shall focus largely on the change detection task in this chapter. The benefit of this focus is that, whilst the behavioural characteristics of the continuous report task are still a point of theoretical exploration, there are several clear decision models for the change detection task (which will be reviewed in the next section). The analysis of observer behaviour in this task is also aided by the close correspondence of (near-threshold) change detection task to the experimental paradigm employed by Sewell et al. (2014), two-choice orthogonal orientation discrimination. Thus, the aim of the current chapter—and the main focus of this first section of the thesis—shall be to characterise the decisional properties of change detection with relation to those of orientation discrimination. To this end, the experiments presented are designed to differ only in the decision asked of the observer, either to report the identity of a previously presented stimulus or to report the presence of a change between the memory and probe arrays; the presentation of the memory array and the timing of the retention interval prior to this decision stage are identical.

The distinguishing feature between orientation discrimination and change detection is the need to detect a difference between a stored memory representation and current sensory input. From an intuitive position, it might be expected that observer sensitivity would be decreased for comparing memory representations to current sensory input

¹These models are broadly consistent, albeit differently framed, with work relating to the sample-size information limit outlined by Smith (2015).

²Specifically, Smith (2016) demonstrated that the effect of response bias has an identical effect on functions of response variability as the effect changing the quality of the memory representation.

when compared to simply reporting a feature of an array item from memory—as additional sensory input may compete for representation in the highly limited memory store and any additional external sensory input would provide an additional source of noise in the decision-making process. There is reason to believe, however, that perceptual comparison may occupy a privileged position in the functioning of visual short-term memory. Hyun et al. (2009) showed an asymmetrical sensitivity for detecting changes: differences between two sequentially presented arrays were identified with a low error rate regardless of the number of items to be compared, both when all items changed or a single item changed, yet the error rate increased substantially over different array sizes when observers were required to identify if the arrays were the same when no item changed. On the basis of electrophysiological correlates of change detection performance, Hyun and colleagues further argued that covert attention may be automatically drawn to locations where a discrepancy between a memory representation and the current sensory input is detected. Specifically, they showed that memory load manipulations in the change detection experiment did not lead to changes in the onset of the N2pc component³ in an averaged ERP waveform, previously associated with shifts in visual attention (Luck, Girelli, McDermott, & Ford, 1997; Luck & Hillyard, 1994a, 1994b). Changes in the onset of the N2pc component have been closely associated with changes in the timing of attentional orienting (Woodman & Luck, 1999, 2003); the lack of a change indicating that orienting occurs automatically when changes are present regardless of the size of the display. A change in the latency of the P3 component⁴, usually associated with stimulus identification, across different array size conditions was also found, argued by Hyun and colleagues to be a signature of a subsequent limited-capacity verification process. In total, they argued that the detection of changes between memory representations and sensory input is supported an automatic orienting process, unlimited in processing capacity, followed by a limited-capacity identification and verification process.

This advantage for detecting changes derive from the close connection between visual short-term memory characteristics and the preservation of perceptual continuity during saccadic eye movements. Evidence has suggested that visual short-term memory is involved with maintaining information over saccadic eye movements: Irwin and Andrews (1996) found that individuals could retain three to four items—including colour, identity, and position—during saccades. Subsequent authors (Currie, McConkie, Carlson-

³A negative deflection in the waveform found around 200 ms after probe array presentation in posterior cortical regions (over the medial and lateral occipital cortex, and the posterior temporal cortex) contralateral to the presentation of the change.

⁴A positive deflection in the waveform found around 300 ms after the probe array presentation.

Radvansky, & Irwin, 2000; McConkie & Currie, 1996) have argued that visual short-term memory provides the continuity of visual experience over saccades, by retaining information about the saccadic target during the eye movements. Henderson and Hollingworth (2003) demonstrated that whilst change detection performance during saccadic eye movements was found to be quite poor when the exact identity of a stimulus was manipulated, the rates of detection when a change deleted the target of a saccade were very high. Further, Hollingworth, Hyun, and Zhang (2005) showed that deletions between two subsequently presented arrays of points presented in a grid could be detected and localised with a high level of reliability, particularly when simple local configurations could be utilised. Hollingworth, Richard, and Luck (2008) also demonstrated that these memory representations, beyond ensuring phenomenological continuity, may also mediate the guidance to eye movements to specific targets by correcting any errors in trajectory by comparing visual information after transit of the eyes.

A favoured sensitivity for detecting discrepancies between stored representation and the current sensory input has also been shown experimentally. Makovski, Sussman, and Jiang (2008), in a rare example of a direct comparison between experimental procedures in the visual short-term memory literature, and a close analogue of the current study, examined observer performance in a colour-based change detection task and a two-alternative forced-choice (2AFC) task. The structure of the 2AFC task was designed to closely match the structure of the change detection task: a memory array of three colour patches—identical to that used in the change detection condition—was displayed for 500 ms, after which a 1000 ms blank retention interval was shown. In both conditions, a single location was selected as the target: in the 2AFC condition, the observer was presented with two patches, one corresponding to the colour of the target patch and a foil alternative; in the change detection condition, all elements other than the target were removed, and the observer was asked to indicate whether the remaining patch was identical to that displayed in the memory array⁵. Using d' values computed from the table provided by Hacker and Ratcliff (1979), Makovski and colleagues showed that observer sensitivity was consistently lower in the 2AFC condition compared with the change detection condition, regardless of the similarity of the alternatives the presence of retro-cuing (see §2.4.1), or the probe positioning.

This result was, as noted by Makovski and colleagues, unexpected. The conversion table provided by Hacker and Ratcliff equates performance from *m*AFC tasks to perfor-

⁵Displaying only a single target patch in the probe array during a change detection has been shown to decrease performance when compared to displaying the entire array (Jiang et al., 2000), although this does not compromise any of the findings of Makovski et al. (2008).




Figure 3.1: A comparison between the observed proportion correct obtained for a given sensitivity parameter (d'_a) assuming three different detection models: A) 2AFC discrimination; B) independent observations same-different model; C) χ^2 differencing model. The discrimination and differencing models show accuracy predictions with different levels of bias (set with the likelihood ratio, denoted β). The independent observations model shows model predictions under different assumptions of sensitivity during the second covert discrimination (d'_b); the solid line represents discrimination when the sensitivity of the first and second covert discriminations are equal.

mance on same–different tasks (a distinction introduced in the next section) in terms of the bias-free sensitivity and correcting for the difference in the underlying experimental procedure; as such, it would be expected that, given each decision is operating on the same memory representation, the sensitivity between decision types would be equal. Observer performance in terms of the raw proportion of correct trials was reported in an appendix of the paper to be higher in the 2AFC task in some instances than in the change detection task, but the direction was not consistent. Such a result is also not compatible with standard signal detection theory predictions: according to differencing models (the optimal model of responding in this task; see Dai, Versfeld, & Green, 1996), the proportion of correct trials would be higher in a 2AFC task than a same–different task when equating for similarity (see Figure 3.1).

One potential resolution for this result may be found with distinguishing between 2AFC in the general psychophysical literature and the task employed by Makovski and colleagues which shares structural similarities with these tasks. The classical 2AFC task in the psychophysical literature, following the work regarding comparative judgement conducted by Thurstone (1927), requires a known stimulus or known stimulus property (such as colour or pitch) to be identified or compared amongst two (or m in the case of m AFC) alternatives. (Stimuli might also be presented sequentially, particularly in the case of aurally presented stimuli, in which case the task would be designated as a two-interval forced-choice procedure or 2IFC.) In these cases, the stimuli elicit some internal representation that is then compared against an internal standard, or compared against each other on a known stimulus dimension.

In the task conducted by Makovski and colleagues, the comparison does not involve an internal stimulus standard, or comparison on a stimulus dimension, but the comparison between an internal memory representation against two externally presented stimulus standards. This task is thus structurally similar to both a 2AFC task in the classical sense and a change detection (or same–different) task: two alternatives are compared against an internal representation, but there is no fixed standard or singular dimension that these external elements are compared to. The decrement in observer sensitivity in the 2AFC task conducted by Makovski and colleagues may, therefore, be an increase in the noise of the decision process—much like the role of location uncertainty in detection tasks. As the number of comparisons to be made increases, the overall error rate also increases as each independent comparison has some rate of error.

Rather than isolating the process by which discrepancies between the external sensory input and an internal memory representation are detected, the task used by Makovski and colleagues doubled the number of change comparisons to be made on each trial,

making a direct contrast between the tasks difficult. The current study follows from this previous research, but requires participants to—in the two-choice orientation discrimination task—judge the identity of the memory representation to an internal stimulus standard, following Sewell, Lilburn, and Smith (2014). The aim of the experiments presented is to, in the first instance, examine the difference in observer performance characteristics between a near-threshold change detection task and near-threshold orthogonal orientation discrimination task, where the trials of the two decision types have been designed to share the same timing and display prior to the presentation of the probe array. These tasks will be examined, in this chapter, with respect to observer sensitivity and accuracy. This analysis will continue in the next chapter with a consideration of the response time data.

As detailed in the preceding review, the experiment conducted by Sewell, Lilburn, and Smith (2014) was theoretically significant as it showed the existence of a limit on the quantity of information; this limit presenting as a “sample-size” relationship between the observer sensitivity and the memory array size. The ability to identify this relationship was due, in part to the sensitivity of the experimental design: the use of a task with a simple, well-specified decision model, the use of simple visual stimuli, and the use of near-threshold stimuli embedded in dynamic noise to control the rate of available stimulus information. Although near-threshold change detection using Gabor patches has been previously conducted (for a review, see Magnussen & Greenlee, 1999), it is unclear whether the sample-size relationship would also appear when a change detection decision rather than an orientation judgement is required. If the change detection is privileged then the amount of information required to reliably detect changes may be much less than that required to identify stimuli in orientation discrimination, in which case the sample-size relationship may not be obtained, indicating that the information limit on orientation discrimination may not be intrinsic to the memory system—as was proposed by Sewell and colleagues—but a consequence of the task structure.

The first experiment examines the effect of the memory array size on performance between differing decision types. The next section will discuss the basic signal detection theory framework that will be used to analyse the results of the first experiment, after which the first experiment will be presented.

3.1 Signal detection theory

Signal detection theory (Green & Swets, 1966; W. W. Peterson et al., 1954) is one of the most successful theoretical frameworks for examining decision processes in psychology and builds upon the distinction, given earlier, of distinguishing mental processes that encode and maintain the representation of a perceptual (or, more generally, psychological) signal from those that define the decision procedure made utilising that signal. The two decision models deriving from signal detection theory that are relevant to the examination of the current experiments are those for 2AFC discrimination judgements and those for same–different judgements. The next section reviews these two models prior to the presentation of experimental results.

3.1.1 Discrimination judgements

An analysis of the orientation discrimination task in terms of signal detection theory relies on a conceptualisation of the visual system as a set of parallel detectors (or channels): selective filters which transduce sensory information into information regarding the presence of specific feature values (such as a specific orientation) within a region of the visual field known as the receptive field of the detector (Graham, 1989; Marr, 1982/2010). The use of the Gabor patch as the principal stimulus type in this thesis, and in the broader visual psychophysical literature, is due to the close correspondence between the luminance profile of the stimulus and the receptive field profile of orientation-specific filters located in visual cortex (Daugman, 1980, 1984). This close correspondence is exhibited in the unusual sensitivity of larger (over 20°) discrimination judgements made, comparable to that of simple detection of the gratings (Thomas & Gille, 1979).

In the orientation discrimination task, it is assumed that each of the stimulus locations in the memory array is within the receptive fields of a set of orientation-specific detectors. These detectors are generally modelled as recoding orientation-specific information in the display by computing the inner products between a set of oriented windowing functions and luminance information from the retina (Graham, 1989). The role of the decision stage is to map the output of these detectors (or, *encoded stimulus intensity*) to a response category (Thomas, 1985).

The model of the decision stage follows from the theory of recognition given by Tanner (1956). Each decision-relevant orientation is represented as a single detector. The response of two detectors to a presented orientation in an orientation discrimination judgement might be entirely independent; that is, the difference between the preferred




Figure 3.2: The geometric relationship between detection and discrimination according to the theory of recognition due to Tanner (1956). Normal distributions are represented by concentric circles, representing lines of isoprobability. The distribution centred at μ_N represents background noise on two detectors: one for the stimulus attribute A and one for the attribute B. The distance between μ_N and μ_A , d'_A , represents the ability to detect A in noise (likewise for B). The distance between A and B, $d'_{A,B}$, represents the ability to discriminate between A and B.

orientation of each detector is sufficiently large that presentation of one elicits no change in response in the other detector. In some cases, when using smaller orientation differences, the responses of detectors may be correlated such that the presentation of a single orientation changes the response of multiple detectors to varying degrees. Discussion of these models will be kept for §5.1.2, where data from an experiment employing stimuli with small angular differences.

For Experiments 1 and 2, the stimuli are horizontally and vertically oriented Gabor patches. These stimuli are assumed to fall within the receptive field of two independent detectors that respond with equal sensitivity to their respective Gabor patches. This special case provides a simple relationship between observer sensitivity in discriminating the presence of a stimulus from background noise (detection sensitivity) and the sensitivity in discriminating between the two orientations:

$$d'_{A,B} = \sqrt{2}d'_A,$$

where d'_A is the distance between the mean intensity of noise and the mean intensity of the stimulus signal (this is denoted d'_A , but is equal to d'_B); $d'_{A,B}$ is the distance be-

tween the mean response for each of the two detectors. $\sqrt{2}$ advantage of two-alternative discrimination over simple detection of a stimulus is a consequence of a geometric interpretation of detector responses: if two detectors are of equal sensitivity and statistically independent of one another, then they can be represented as orthogonal vectors in Euclidean space, scaled in terms of units of standard deviation of the background neural noise; by Pythagoras' theorem, the distance between the end of the vectors is then equal to $\sqrt{2}$ multiplied by the distance of either of the vectors. Figure 3.2 demonstrates this relationship. The overall sensitivity parameter d' for two-alternative discrimination is computed from the observed data using the formula

$$d' = 2z [P(C)],$$

where $z[\cdot]$ is the standard normal quantile function (a z-score) and $P(C)$ is the overall proportion of correct responses.

The use of stimuli presented near the detection threshold means that the intensity of detector responses to stimuli is not completely separable from the intensity of detector responses to background noise. As such, there will always be a trade-off between response categories must be made. These trade-offs are modelled in signal detection theory more generally as response biases. Although there is no reason, *a priori*, for observers to favour one response option over the other in a two-alternative forced choice procedure, with bias not conventionally modelled in two-alternative forced-choice tasks, this does not preclude observers from preferring one response option over another. Any response bias on an individual observer level, however, will likely not be seen when examining averaged data. Bias is conventionally modelled as the observer setting a criterion value based on the likelihood ratio of detector activation given a stimulus presentation. Optimal responding occurs when the observer selects a likelihood ratio of one; that is,

$$\frac{\phi_A(c)}{\phi_B(c)} = 1,$$

where c is a given criterion value (in terms of detector activation), ϕ is the Gaussian pdf for the detector corresponding to either stimulus category A or B, identified by the subscript.

3.1.2 Same-different judgements

Unlike two-choice orientation discrimination, there is no single accepted model for change detection tasks in the signal detection theory literature. Change detection can be mod-

elled as a “same–different” judgement, where two classes of model are applicable to the analysis of observer responses: the first is a *differencing* model, due to Sorkin (1962) following Lamphiear and Birdsall (1960); the second is an *independent observations* model, due to Noreen (1981).

Differencing models, when applied to a change detection task, assume that the intensity of the target stimulus in the memory array and the intensity of target stimulus in the probe array can be combined, and the difference of the two outputs can be compared to a criterion level to determine a response. The original published model of Sorkin (1962) presented the differencing model stated in terms of multiple normal distributions representing the response of a psychophysical detector for stimuli above and below a target stimulus standard on a given dimension (Sorkin was originally concerned with the auditory psychophysics of frequency discrimination). Two criteria are set below and above the stimulus standard to represent the differences from the stimulus standard on specified feature dimension (see Figure 3.3). When these criteria are equal (or when the decision does not take into account the sign of the difference), the model given by Sorkin—as he noted and as earlier noted by Lamphiear and Birdsall (1960)—resolves to a discrimination in terms of stimulus energy (squared stimulus intensity) using a non-central χ^2 distribution. The historical difficulty of numerically approximating the distribution function of the non-central χ^2 distribution with arbitrary degrees of freedom has meant that the form using normal distributions has been the more commonplace⁶.

The logic of the χ^2 version of model follows simply from the version with normal distributions. A χ^2 distribution with k degrees of freedom is the sum of k normally distributed random variables squared. In the case of the conventional central χ^2 distribution, these normal random variables have mean of zero and a standard deviation of one (i.e., standard normal random variables); in the case of the non-central χ^2 distribution, the mean of the random variables is equal to the non-centrality parameter. The noise distribution is represented as a central χ^2 distribution with a single degree of freedom: the distribution of the squared difference of a single draw from a standard normal distribution (equivalent to a squared draw from a standard normal distribution), representing background noise. The signal (change) distribution is represented as a non-central χ^2 distribution with a single degree of freedom and a non-centrality parameter equal to the squared sensitivity, $(d')^2$: the distribution of the squared difference of a normally distributed random variable from a fixed signal of magnitude d' (equivalent to the sum of

⁶It is worth noting due to the popularity of the form of the differencing model given by Macmillan and Creelman (2004), using three normal distributions to determine the predicted hit rate and false alarm rate, that an apparent error in the formula means that probability mass is lost in the computation, where the tails of the probability distributions representing stimuli other than the standard cross both criteria.




Figure 3.3: The differencing model of Sorkin (1962). Three normal distributions represent intensities elicited for stimuli shown to the observer: a stimulus standard with a mean of μ_C and two stimuli with an offset of δ_S on a feature dimension. The two solid lines either side of the standard are criteria, c_- and c_+ , which represent the difference in feature dimension from μ_C before the observer responds that a stimulus is “different”. The hatched area between the criteria represents the probability of responding “same”. As the criteria are equally distant from the mean value of the standard, this decision can be modelled using χ^2 distributions.

the signal plus the random variable). The hit rate HR of this version of the differencing model is, for a given criterion value c and sensitivity value d' ,

$$HR = 1 - F\left(c; 1, (d')^2\right),$$

where the two parameters of the non-central χ^2 distribution are the degrees of freedom and non-centrality parameter, respectively. The false alarm rate FA of this version of the differencing model is

$$FA = 1 - F(c; 1),$$

with only one parameter for the central χ^2 distribution, the degrees of freedom. Different from standard signal detection theory models, c is scaled in terms of energy (squared stimulus intensity) rather than simply stimulus intensity. This formulation simplifies the differencing model in the case where signed differences do not matter.

The other class of model of same–different judgement—known as the independent observations models—does not combine the intensity of the stimuli themselves, but simply combines the results of two separate decisions made about the stimuli. In effect, the independent observations model is the product of two covert discrimination judgements to form a single decision: one discrimination judgement made on the probed representation encoded from the memory array and one discrimination judgement made on the probe array; if the result of both decisions is equal, then a response of “same” is elicited, otherwise a response of “different” is elicited.

Dai, Versfeld, and Green (1996) showed that each of the two classes of model could

be considered as optimal for different regimes of stimulus selection. When stimuli are sampled uniformly from a fixed set of potential values that can be known exactly by the observer, there is no correlation between the value of the memory stimulus and the probe stimulus, and an independent observations model is optimal. In the case where stimuli are sampled along some dimension and the foil stimuli are defined by some offset of a stimulus standard, there is a correlation between the value of the memory stimulus and the probe stimulus, and a differencing model is optimal. Although the former instance applies to the tasks presented in this thesis, there are reasons why the differencing model may provide a better account of performance in the specific change detection tasks presented here over the independent observations model. First, although the set of possible stimuli is fixed across the entire experiment, any covert decision regarding the identity of the first set of stimuli in the memory array would have to either be made for each item in the probe array—implicating some verbal recoding process (contra Pashler, 1988)—or be delayed until the display of the post-stimulus probe. As there is no evidence of verbal recoding in visual short-term memory experiments with similar duration of stimulus exposure and even longer retention intervals, verbal recoding is somewhat unlikely. Second, the probe array was presented for an indefinite time period at high contrast, meaning that—if the observer uses all information available to them—the representation of the probe should be essentially noiseless. In this case, an independent observations account would predict exactly the same performance in terms of the raw proportion of trials correct in the change detection case as the orientation discrimination case: no additional noise is being added to the decision. Any increment or decrement in performance seen in change detection when compared to orientation discrimination means that some additional source of noise is being added.

In the first experiment presented, the presentation of the memory array is consistent with respect to timing and stimulus contrast across trials where different types of decisions are required. This allows some basis of comparison between different classes of signal detection theory model for trials where a change detection decision is required. As there may be differences in the proportion of correct trials obtained for change detection and orientation discrimination trials, against a strong interpretation of the independent observations model, both classes of models will be fit to the change detection data—along with the standard discrimination model fit to the orientation discrimination data.

3.2 Experiment 1

Experiment 1 examined observer performance in change detection when compared to orientation discrimination across different memory array sizes, and equating the contrast at which the memory array is presented.

3.2.1 Method

Participants

Five observers participated in this study: myself (SL), a member of the Vision and Attention Laboratory (SS), and three paid naïve observers (KT, BF, and DB) from the University of Melbourne. Each observer, with the exception of myself, was briefed about the general nature of the study (with the specific predictions regarding the outcome omitted from this briefing, however) and signed a consent form prior to participation. Each observer from outside of the laboratory was remunerated AUD \$12 for each session completed.

Each observer completed a variable number of practice and calibration sessions, followed by five experimental sessions. Practice sessions were undertaken to familiarise the observers with the task and responding in a timely manner; calibration sessions were undertaken to control for individual differences in performance over stimulus contrast settings. All sessions lasted approximately thirty-five minutes each, with regular periods in between blocks of trials for breaks.

Stimuli

The stimuli used in the memory and probe arrays were oriented Gabor patches: Gaussian vignettted 3.5 cpd sinusoidal luminance gratings subtending 0.97° of visual angle at half-height. The form of the Gabor patches was as given by Graham (1989, p. 53). The patches were oriented either horizontally or vertically and placed on a mean luminance field of 30 cd/m^2 . Patches could be presented at four locations, diagonally located at a distance of 2.3° from a central fixation cross subtending 0.29° of visual angle.

Previous research has indicated that, under normal viewing conditions, observers can encode near-threshold Gabor patches in around 50–60 ms (Liu, Wolfgang, & Smith, 2009; Smith, 2000b, 2010; Smith & Ratcliff, 2004) and that coarse orientation discrimination sensitivity of Gabor patches is equivalent to that of detection sensitivity (Thomas & Gille, 1979). To maximise the effect of stimulus duration on memory formation, the period required to encode the stimuli was elongated by embedding the stimuli in dynamic

noise (Ratcliff & Smith, 2010; Sewell et al., 2014; Smith, Ratcliff, & Sewell, 2014). The noise patches were composed of 4×4 pixel blocks of luminance sampled from a truncated Gaussian distribution with a mean of the background luminance and a variance scaled to fit within 20% of the entire luminance range of the display. Noise patches were displayed on alternating frames to the stimulus display during the stimulus display period, meaning that each 10 ms of stimulus display was immediately followed by 10 ms of noise. This approach has been previously reported in the visual short-term memory literature by Sewell et al. (2014). The backwards (interrupting) masks used to truncate and disrupt any iconic memory trace were high-contrast checkerboards sized to ensure that the location of the original stimulus was completely occluded.

In this experiment, the contrast of stimuli was fixed to be equal across different decision types, and adjusted for each observer during the practice sessions to span as much of the range of accuracy as possible between both decision types.

Apparatus

Stimuli were generated on a Cambridge Research System ViSaGe framestore and presented on a gamma-corrected 21" Sony Trinitron Multiscan G520 monitor, running at a resolution of 1024×768 pixels and driven at 100 Hz (giving a frame duration of 10 ms). Custom C++ software was used to generate the stimuli, to control trial presentation, and to receive and record responses. Observers performed the task in a dimly lit observation booth at a viewing distance of 100 cm. Viewing position was stabilised with a chinrest.

Procedure

A 3×2 within-subjects design was used, composed of three memory set sizes (1, 2, and 4 items) and two decision types (orthogonal orientation discrimination and change detection). Each session of the experiment consisted of 384 trials, yielding a total of 1,920 trials per observer. Trial presentation order of differing memory set size conditions was randomised within alternating blocks of a single decision type. The decision type of the first block was also randomly selected prior to each session.

Each trial began with a 1,000 ms uniform field, followed by the presentation of the central fixation cross for 1,500 ms. A stimulus array of between one and four orthogonally oriented Gabor patches was then presented for 150 ms, followed by the presentation of a high-contrast checkerboard mask in all stimulus locations for 200 ms. A post-mask gap showing only the fixation cross was then presented for a total of 200 ms.




Figure 3.4: A schematic overview of the presentation order in Experiment 1. In both change detection (CD) and orientation discrimination (OD) trials, a fixation cross, memory array of between one and four items, and backwards mask are presented. Following the presentation of the backwards mask, a report cue is presented, overlaid on the checkerboard, for orientation discrimination trials; in the change detection trials, a probe array, with a probed stimulus, is displayed. Note: this diagram is not to scale; Gabor patches and checkerboard masks have been replaced with symbols for clarity.

In the orientation discrimination decision condition, the post-mask gap was followed by the presentation of a report cue, indicating the location in the memory array corresponding to the item the observer was to report the orientation. Responses were entered via button press and response times were recorded. Observers were instructed to enter their response as quickly as possible without compromising their accuracy. The report cue remained on screen until a response was entered, and auditory feedback was presented to the observer after the response was entered.

In the change detection decision condition, the presentation of the memory array was identical to the orientation discrimination condition. Following the post-mask gap, the probe array was presented instead of a report cue. The probe array consisted of oriented Gabor patches, corresponding to each item in the stimulus array, with one item marked by a report cue. Each of the items in the probe array not marked by the report cue shared the orientation with their respective items in the stimulus array; the item marked by the report cue was either unchanged from the stimulus array, or the orientation changed (from horizontal to vertical or vice versa). Observers in change detection trials were to indicate whether any change was detected at the probed location. Like the report cue in the orientation discrimination condition, the probe array remained on screen until the observer entered a response. Auditory feedback followed the entry of a response.

3.2.2 Results

For the analysis of sensitivity and accuracy, no data were filtered from the overall dataset. (Measures of response time obtained in this experiment will be examined in the next chapter.)

Figure 3.5 shows observer averaged accuracy (the proportion of correct responses) as a function of the memory array size and the decision type; Figure 3.6 shows these data for each observer. In all cases, the proportion of correct responses decreases as the memory size increases, and trials requiring an orientation discrimination decision have a higher proportion of correct responses than trials requiring a change detection decision.

A mixed-effects logistic regression was used to provide a preliminary analysis of the data, prior to conducting the signal theory analysis. The logistic regression was conducted on the proportion of correct trials, with the memory array size and decision type as fixed effects regressors and the observer was treated as a random effect on the intercept. A significant main effect of memory array size on observer accuracy was seen, $\beta = -0.164$, $SE = 0.028$, $p < 0.001$, with accuracy decreasing as the memory array size increased. A significant effect of the decision type on observer accuracy was also ob-




Figure 3.5: Group average accuracy data from Experiment 1 across different memory array sizes and conditioned on the decision type. Error bars represent one standard error of the mean. CD = change detection; OD = orientation discrimination.




Figure 3.6: Accuracy data from Experiment 1, displayed for each observer, across different memory array sizes and conditioned on the decision type. Error bars represent one standard error of the mean. CD = change detection; OD = orientation discrimination.

served, $\beta = 0.270$, $SE = 0.112$, $p = 0.016$, with the orientation discrimination decision type being higher in accuracy than change detection. The interaction between memory array size and decision type on observer accuracy was not significant, $\beta = -0.033$, $SE = 0.040$, $p = 0.408$.

The significant difference in the proportion of correct responses between decision types means that a strong version of the independent observations model is not supported: as the presentation of the memory array and the timing of the retention interval were held constant between decision types, the requirement to encode the probe array in order to make a decision is detrimental to overall performance, despite the fact that the probe array is available to the observer until a decision is made.

To examine the data in terms of signal detection theory introduced in the previous section, a set of signal detection theory models was created and fit both to each observer individually and to group average data created using the mean response proportions of each individual observer. Each model was fit to both the orientation discrimination trials and change detection trials jointly, using the two-alternative discrimination model for the former and either a differencing model or an independent observations model for the latter. The models fit the response proportions (e.g., the proportion of “Change” response given a “No Change” trial; the proportion of a “Horizontal” response given a “Horizontal” target; etc.), rather than the overall proportion correct to provide a more precise account of the data. A list of each of the different model configuration factors is listed in Table 3.1.

A set of 72 candidate model configurations using a differencing rule were constructed for each observer. In models where the sample-size constraint was enforced, a single d' was estimated for all memory array sizes; where the sample-size constraint was not enforced, each memory array size had an estimated d' term. Response bias could either be estimated for the orientation discrimination trials (i.e., bias towards responding one orientation over another), change detection trials (i.e., bias either toward responding that a change was detected or that no change was detected), both types of trial, or neither type of trial. In the cases where a bias parameter was not estimated, the observer was either assumed to behave optimally, selecting a criterion level that maximised the proportion of correct trials for a given decision type (the *optimal* response strategy); or, the observer was assumed to set the probability of responding “change” or “no change” to be equal (the *equal* response strategy). Last, an optional parameter allowed an increment to the denominator of the sample-size relation for change detection trials only: either fixed at zero, one, or freely estimated. This parameter represents an additional load in the sample-size relationship, due to comparison with an external sensory referent. When

Table 3.1: Different model configuration factors in the construction of the signal detection theory models for Experiment 1.

Model factor	Description	Factor levels
CD rule	The type of decision rule employed in change detection trials.	<i>Differencing</i> (the decision of the observer is based on the difference between the sensory intensity of the first array and the sensory intensity of the second array); <i>Independent observations</i> (the decision of the observer is based on two covert discrimination decisions).
SS constraint	Whether the sample-size relationship is enforced, leading to a constraint on sensitivity between different memory array sizes.	<i>Yes</i> (sensitivity is constrained to the sample size relation, $\sum_m (d'_m)^2 = \text{const.}$); <i>No</i> (the effect of memory array size on sensitivity is freely estimated).
Denominator	Whether an additional item or items are added to the denominator of the sample size relationship.	<i>0</i> (the sample-size relationship is the same for orientation discrimination and change detection trials); <i>1</i> (an additional item is added to the sample-size relationship in change detection trials, $d'_m = d'_1 / \sqrt{m+1}$); <i>Estimated</i> (an additional estimated constant is added to the sample-size relationship in change detection trials, $d'_m = d'_1 / \sqrt{m+c}$).
Bias	Whether the response bias of the observer is freely estimated (either for both conditions, or for a single condition) or is set to maximise accuracy.	<i>Both</i> (response bias is freely estimated for both conditions); <i>Orientation</i> (response bias is freely estimated for the orientation discrimination trials); <i>Change</i> (response bias is freely estimated for the change detection trials); <i>Equal</i> (the response bias means that the probability of responding “change” or “no change” is equal); <i>Optimal</i> (no response bias estimated; observer maximises accuracy).
d' weight	Freely estimated weighting parameter added to all sensitivity terms in change detection (in the case of the sample-size constraint, this affects the d'_1 value; without the sample-size constraint, it weights all estimated terms).	<i>Yes</i> (weight estimated and multiplied with change detection sensitivity value); <i>No</i> (no weight estimated).

fixed at one, this comparison load indicates that the item marked by the post-stimulus probe is also subject to the competitive interactions governing items within memory (Sewell et al., 2014; Smith, 2015; Smith et al., 2015); a fixed value of zero indicates no additional load; and estimated values greater than one indicate that processing the probe array places a greater load on memory. All included, the most flexible model with a differencing rule tested had eight estimated parameters: four d' parameters, one for each memory array size; one d' weighting for change detection trials; two bias parameters, one for orientation discrimination trials and one for change detection trials; and an estimated additional load parameter. This model was not saturated, given that thirty-two data points were fit (one for each response alternative, giving a total of 16 degrees of freedom). The most constrained model had only a single parameter: a single d' parameter.

For each memory array size, a d' value was shared between the orientation discrimination and change detection trials. In some model configurations, however, an additional weighting constant was estimated for change detection trials. In these configurations, a single value α was estimated for all memory array conditions, and multiplied with the sensitivity of the orientation discrimination trials to allow conversion from the sensitivity in the discrimination model to the differencing model. When the sample-size relationship was enforced, this weighting affected the numerator of the sample-size relationship,

$$d'_{m,CD} = \frac{\alpha \cdot d'_{1,OD}}{\sqrt{m}},$$

where the subscript provides the memory array size and the decision type. When the sample-size relationship was not enforced, the same constant was applied to the estimated sensitivity value for each memory array size,

$$d'_{m,CD} = \alpha \cdot d'_{m,OD}.$$

Although both change detection and orientation discrimination judgements are based on the same memory representations, their demand on those representations may differ: it is possible that a same–different performance requires less overall effective information from memory representations when compared to orientation discrimination, where an identification is made. In these cases, the sensitivity of change detection would be higher than otherwise expected.

A more restricted set of 16 candidate models using an independent observations rule was created for each observer. In these models, the sensitivity of the first covert decision was shared with the estimated sensitivity of the orientation discrimination trials. The

second covert decision was also fixed to be at the same level of sensitivity—despite the high contrast presentation and indefinite presentation time—but could be changed by an optional d' weighting parameter, like in models using the differencing rule. In the independent observations rule, this weighting parameter only changes the sensitivity of the second covert decision, as it is assumed that the first covert decision is identical to the decision made in the orientation discrimination condition. Like the models using the differencing rule, the sensitivity of both covert decisions and the orientation discrimination condition could be forced to conform to the sample-size relation (in which case, only a single d' parameter was estimated) or not (meaning all four array sizes had an estimated d' value). Response bias was also implemented in the same way as model configurations using the differencing rules: response bias could either be estimated for the orientation discrimination trials (which also estimates bias in the first covert decision in change detection trials); bias could be estimated for the second covert decision in change detection trials; bias could be estimated for orientation discrimination and change detection completely; or not estimated at all (in which case, like in the differencing model, the optimal response strategy is modelled). All included, the most flexible model with an independent observations model tested had seven estimated parameters: four d' parameters, a d' weighting parameter for change detection trials, and two bias parameters.

For each candidate model, the best fitting set of parameters were those that minimised a chi-squared objective function,

$$\chi^2 = \sum_{i \in C} N_i \sum_{j \in R} \frac{(p_{ij} - \pi_{ij})^2}{\pi_{ij}},$$

where the index pair i and j run over the set R of each response type (e.g., responding “change” given no change) and the set C of each experimental condition (memory array and decision type pair), N_i is the number of observations in a given experimental condition, p_{ij} is the proportion of observed responses in this cell, π_{ij} is the predicted proportion of responses in this cell. The chi-squared objective function was minimised using the modification provided by Byrd, Lu, Nocedal, and Zhu (1995) of the limited-memory quasi-Newton method of described by Broyden, Fletcher, Goldfarb, and Shanno independently in 1970 (Broyden, 1970; Fletcher, 1970; Goldfarb, 1970; Shanno, 1970), known as the L-BFGS-B algorithm.

All candidate models were ranked in terms of the Bayesian Information Criterion (BIC;

Table 3.2: The five best fitting signal detection theory models, in terms of BIC, for the group average data from Experiment 1. The denominator represents the extra element in the denominator of the sample-size relation (for differencing models only).

#	Model type	SS constraint	Denominator	Bias	d' weight	N. pars	χ^2	BIC
1	Diff.	Yes	0	Equal	No	1	12.991	20.542
2	Diff.	Yes	1	Equal	Yes	2	5.516	20.618
3	Ind.Obs.	Yes	—	Optimal	Yes	2	7.806	22.917
4	Diff.	Yes	0	Equal	Yes	2	9.128	24.238
5	Diff.	Yes	0	Orientation	No	2	10.692	25.826

Notes: SS constraint = Sample size constraint; IO = Independent observations rules; Diff. = Differencing rule; BIC = Bayesian Information Criterion; N. pars = Number of free parameters.

Schwarz, 1978),

$$\text{BIC} = G^2 + k \cdot \ln(N),$$

where G^2 is the maximum likelihood test, $G^2 = 2 \sum_{i \in C} N_i \sum_{j \in R} p_{ij} \cdot \ln(p_{ij}/\pi_{ij})$; k is the number of free parameters; N_i is the number of observations in condition i ; and N is the total number of observations. (The χ^2 function was minimised in this case, rather than G^2 , as it provided more stable minimisation.) The five best fitting models, in terms of the lowest BIC values, for the group average data are presented in Table 3.2 and shown in Tables 3.4–3.8 at the end of the chapter for clarity.

The parameters for each of the best fitting models employing a differencing rule and the best fitting models employing an independent observations rule are shown for each observer, in Table 3.9. Note that the best fitting model configuration for the group average does not have an item increment in the denominator of the sample-size relationship for change detection trials. This model configuration is marginally better in BIC terms than a model configuration with the denominator in the sample-size relation, which also requires the estimation of a d' weight, but a relatively poorer fit in terms of the goodness-of-fit statistic. Given the large qualitative difference in the fit, visible in Figures 3.7 and 3.8, and the small difference in the BIC values, the exact rank ordering of the model configurations is spurious. Each model provided predictions for each response outcome, rather than the overall proportion correct, meaning each observer had 32 data points rather than 8, giving a total of 16 degrees of freedom. The group average data, averaging over individual observer data, is underdispersed as indicated by the fact that χ^2 goodness-of-fit values are smaller than the residual degrees of freedom (the data degrees of freedom minus the model degrees of freedom) leading to an overly conservative




Figure 3.7: Predictions for the second best fitting model, in terms of BIC, to the group average data. This model uses a differencing rule, adhering to the sample-size constraint, with the addition of a single item to the denominator of the sample-size relation in change detection. The dashed lines represent observed data from the experiment, the solid lines represent the predicted data from the model.

penalisation of free parameters in the computation of the BIC⁷.

Figure 3.7 shows the predictions of the differencing model where a single item is added to the denominator of the sample-size relationship in the change detection trials; Figure 3.8 displays the predictions for the model with no denominator in the sample-size relation. Finally, Figure 3.9 displays the predictions of the best fitting model with an independent observations rule against the group average. The parameters for the best fitting model configurations for each observer are presented at the end of the chapter in Table 3.9.

⁷Overdispersed data, the more common occurrence, requires a correction to the information criterion value to appropriately scale the penalty incurred by additional parameters, leading to the QAIC (see Burnham & Anderson, 2003).




Figure 3.8: Predictions for the best fitting model, in terms of BIC, to the group average data. This model uses a differencing rule, adhering to the sample-size constraint, with no modification to the denominator of the sample-size relation in change detection. The dashed lines represent observed data from the experiment, the solid lines represent the predicted data from the model.




Figure 3.9: Predictions for the third best fitting model, in terms of BIC, to the group average data. This model uses an independent observations rule, adhering to the sample-size constraint. The dashed lines represent observed data from the experiment, the solid lines represent the predicted data from the model.

3.2.3 Discussion

The aim of the first experiment was to examine whether observer performance was comparable across two different decision types used in the visual short-term memory literature: two-choice orientation discrimination and change detection. The result obtained—most clearly seen in the group average data—was that trials requiring a change detection decision had lower performance in terms of the proportion of correct responses than the trials requiring an orientation discrimination decision. The similarity in form of the accuracy functions was apparent, although sometimes obscured by noise in the data.

The signal detection theory models using either change detection rule fit well on a qualitative level, but only when a weighting parameter could be introduced to convert the sensitivity predicted for the orientation discrimination trials into a sensitivity appropriate for the change detection data. Almost all of the best fitting models were those that employed a differencing rule. Of those models, models of the change detection task where a single additional item was added to the denominator of the sample-size relation tended to fare better than those where no additional item was added to the sample-size relationship. There were two exceptions to this finding: the group average data and observer BF. In the case of the group average data, the model where no additional item was added to the denominator was only slightly better in terms of BIC than a model where an additional item was added to the denominator (and the goodness-of-fit statistic without any penalisation for complexity was better in the latter case). The difference in BIC, however, was insubstantial: compare a BIC value of 20.542 for the model with no additional load in the sample size denominator to a BIC value of 20.618 with an additional element in the denominator (see Table 3.2). In the case of observer BF, differencing models with no change to the sample-size relationship accounted for four of the five best fitting models in terms of BIC.

Each of the best fitting models—regardless of the rule used in the change detection trials—was constrained by the sample-size relationship, consistent with the result of Sewell et al. (2014). The success of the sample-size relationship in describing the change in sensitivity with changes in the memory array size across both types of decision indicates that the information limit described by Sewell and colleagues is not due simply to a specific task strategy or a byproduct of the decision stage in orientation discrimination, but fundamental to the memory system itself. This result gives strong support to the argument that both orientation discrimination and change detect operate on the same memory representations, but that change detection tasks lead to a higher rate of errors due to additional complexity in decision procedure.




Figure 3.10: A comparison between the proportion correct values obtained by the differencing model ($P_D(C)$) when compared to other models ($P_x(C)$) given the same d' value. The solid line represents a comparison of the (χ^2) differencing model against an independent observations model, where the d' for the first covert discrimination is equal to the d' of the second covert discrimination. The dashed line represents a comparison in proportion correct between the differencing model and a 2AFC discrimination model. The dotted line on the diagonal shows the identity relation between proportion correct values. All models were assumed have optimal criteria values for a given sensitivity value.

This finding is in contradiction to the assertion of Hyun et al. (2009) that the comparison of sensory input to memory representations is unlimited in capacity. There is some evidence, however, to suggest that the process of detecting discrepancies between the active memory representation and the sensory input is somewhat privileged when compared to orientation discrimination, almost all of the better fitting models requiring the estimated sensitivity for the orientation discrimination condition to be weighted in the change detection condition such that change detection performance was better than would otherwise be expected. The size of this weight was more stable between participants in models employing the differencing rule, indicating that the change detection sensitivity is around one and a half times the sensitivity estimate of the orientation discrimination task. The locus of such an advantage is unclear. Although the weighting on d' improves change detection performance more than would otherwise be expected given the level of change detection performance, this was better fit to a multiplicative weight in the numerator of the sample size relationship (to the d') rather than an additive constant to the denominator (to the division of d' , via an estimated increment to the sample size relation). This may be due to the fact that less effective information is required from a visual short-term memory representation to make a same–different judgement, when compared to a discrimination (or identification) judgement—even though both are subject to the same overall information limits. Partly, this difference in d' may also be contingent on the difference in the mapping of d' to observed proportion correct between the three models. Figure 3.10 shows the relationship between proportion correct estimated by a differencing model against proportion correct estimated from an independent observations model and a 2AFC discrimination model given the same d' value. (Recall that, in this model fitting procedure, the proportion correct was a byproduct of the d' , and that conditional response proportions were estimated directly.)

Further exploration of the relationship between these two different tasks, for different periods of stimulus availability, will be explored in the next section.

3.3 Experiment 2

The similarity in the functional form of observer accuracy data over different memory array sizes between decision types in the first experiment was evident both at the group average level and from the lack of a statistically significant interaction between memory array size and the type of decision required. At the level of individual participants, however, the relationship between the two tasks was not as evident.

One key result of the previous experiment was that each type of decision—orientation discrimination or change detection—appeared to be well described overall by the sample-size information limit reported by Sewell, Lilburn, and Smith (2014), but that this relationship also required the sensitivity of the change detection task to be higher than would otherwise be expected on the basis of the sensitivity of the orientation discrimination task. Further, in many instances the sample-size relationship required the addition of an extra item to the denominator of the sample-size relationship to account for change detection results.

The implications of this result in orientation discrimination over different stimulus exposure conditions for change detection were previously unexplored but, on the basis of the results of the previous experiment, it was expected that the pattern of results in change detection would be consistent with those in orientation discrimination. Experiment 2 attempted to characterise the relationship between the two tasks over different levels of stimulus exposure duration under a larger visual short-term memory load of four items, where competition would be maximal. Furthermore, as competition between memory representations appears to be proportional to strength of memory traces, uniformly increasing the available information for a change detection decision by increasing the stimulus (memory array) contrast should offset the decrease in overall correct responses due to the difference in decision procedure. The calibration of stimulus contrast in Experiment 2 was, therefore, conducted separately for differing decision types in attempting to—as far as is possible—equate observer performance.

Sewell, Lilburn, and Smith (2014) argued that the sample-size constraint on information was also obtained for different lengths of stimulus exposure. They showed a linear growth in squared sensitivity with increasing exposure durations, even with sequentially presented stimuli, and demonstrated that this finding was consistent with the effect of representations in memory competing for processing or maintenance resources proportional to the strength of the representation even when, at shorter stimulus exposure durations, relatively little information about the stimulus had been encoded (Smith et al., 2015).

Data from this experiment will also be used for response time analysis, to be reported in the next chapter.

3.3.1 Method

Participants

The observers in Experiment 2 were the same as those in Experiment 1. Each observer completed a variable number of practice and calibration sessions, generally more than those completed for Experiment 1, followed by 15 experimental sessions.

Stimuli and apparatus

The stimuli and apparatus were the same as used in Experiment 1, with the key difference that the stimulus contrast of the memory array used in the orientation discrimination trials and the contrast of the memory array used in the change detection trials were not constrained to be equal. The stimulus contrast was designed to maximise the total range of observer performance in terms of accuracy as much as possible given the larger memory array size and the greater range memory array exposure durations.

The stimulus (memory array) contrast was adjusted on an observer-by-observer basis to attempt to equate performance between the two types of decision required.

Procedure

A 3×2 within-subjects design was used, composed of three memory array exposure durations (100, 150, and 200 ms) and two decision types (orthogonal orientation discrimination and change detection). Each session of the experiment consisted of 432 trials, yielding a total of 6,480 trials per observer.

The stimulus presentation regime was identical to that of the first experiment, with the exception that the memory array was always four items (presented diagonally from the fixation cross) presented for one of the three exposure durations. Additionally, there were some slight differences in the overall presentation timing when compared to the first experiment: the fixation cross was presented for 500 ms and the high-contrast backwards checkerboard mask was presented for 100 ms. The rest of the timing parameters were identical to the first experiment.

3.3.2 Results and discussion

As in Experiment 1, no data was filtered from the overall dataset for the analysis of observer accuracy in Experiment 2. Figure 3.11 shows the observer averaged accuracy in

Table 3.3: The mean stimulus (memory array) contrasts for the orientation discrimination (OD) and change detection (CD) conditions for each observer as a proportion of the total luminance range (between 0.0 and 1.0).

Observer	OD contrast	CD contrast
BF	0.274	0.439
DB	0.295	0.304
KT	0.266	0.452
SL	0.197	0.249
SS	0.274	0.397

terms of the proportion of correct responses; Figure 3.12 shows these data for each observer.

Each observer was calibrated on the different decision types independently to equate the overall proportion of correct responses in the two different types of decision required. Table 3.3 shows the mean display contrast of the memory array for each observer over the course of the majority of experimental sessions (small adjustments were made to observer contrast to offset learning effects during the course of the experiment). In each case, the contrast of the change detection trials was higher than those of the orientation discrimination trials to offset the difference in accuracy due to the decision type. This difference in contrast supports the modelling results of Experiment 1 in showing a lower sensitivity for change detection decisions when compared to orientation discrimination decisions. Further, this difference in sensitivity appears to be uniform across different exposure durations.

To examine whether the change in the proportion of correct responses over stimulus (memory array) exposure durations, a mixed-effects logistic regression was used. Like the regression of Experiment 1, the logistic regression here was conducted on the proportion of correct responses with the stimulus exposure duration and decision type treated as fixed effects and the individual differences of the observer on the intercept treated as a random effect. The stimulus exposure duration was scaled to be in units of seconds, rather than milliseconds, to stabilise estimation of the model parameters. A significant main effect of stimulus exposure duration on observer accuracy was observed, $\beta = 6.181$, $SE = 0.370$, $p < 0.001$, with accuracy increasing with increases to duration that the memory array was visible to the observer. No significant effects were seen for the effect of decision type on observer accuracy, $\beta = -0.088$, $SE = 0.081$, $p = 0.248$, or in the interaction between stimulus exposure duration and the decision required, $\beta = 0.399$, $SE = 0.529$, $p = 0.450$.




Figure 3.11: Group average accuracy data from Experiment 2 across different stimulus (memory array) exposure durations and conditioned on the decision type. The contrast of the memory array was adjusted within each decision type condition for each observer to equate performance. Error bars represent one standard error of the mean. CD = change detection; OD = orientation discrimination.




Figure 3.12: Accuracy data from Experiment 2, displayed for each observer across different stimulus (memory array) exposure durations and conditioned on the decision type. Error bars represent one standard error of the mean. CD = change detection; OD = orientation discrimination.

These results are consistent with what would be expected given a single underlying memory constraint in both orientation discrimination and change detection decisions: with the addition of greater stimulus information in the change detection task offsetting the decrease in the overall correct response rate due to the additional noise caused by the decision procedure, the profile of change in the proportion of correct responses is identical between tasks.

3.4 General discussion

The aim of the current experiments was to examine the relationship between change detection and two-alternative orthogonal orientation discrimination in terms of observer accuracy and sensitivity. In both experiments, a clear relationship was found between the two types of task: the proportion of correct responses for change detection decisions was lower than that of orientation discrimination decisions when the amount of stimulus information was equated (i.e., the contrast of the memory array was equal across decision type conditions), and the overall profile of accuracy was identical across stimulus exposure durations when the memory array contrast was used to offset the decrement in accuracy due to the the decision process in change detection.

Modelling of the data obtained in Experiment 1 using different decision rules taken from signal detection theory indicated that the relationship between memory array size and observer accuracy was accounted for by the sample size relationship, following the work of Sewell, Lilburn, and Smith (2014). The best fitting models employed a differencing rule, rather than the optimal independent observations rule, for change detection tasks and often included a small increment to the denominator to the sample-size relation in the change detection indicating that the probe element may have been encoded to the memory system, competing with representations from the memory array.

Overall, these results contradict the conclusions of Makovski and Jiang (2009) that observer sensitivity in a two-alternative forced choice task is less than that of a change detection task. As indicated in the introduction, their result may have been due to the design of their task, where the two response options were presented after the retention interval. Each trial, therefore, required two comparisons of the memory array. On the basis of the current findings, it would be expected that the decrement in sensitivity seen in the two-alternative forced choice task may have occurred due to this additional comparison. This, along with the change to the sample-size denominator, also indicates that the probe array may compete for representational resources in the memory system dur-

ing a change detection decision.

These results have clear implications for work in visual short-term memory attempting to estimate the storage capacity of the memory system from observer accuracy data obtained in a change detection experiment. The Pashler–Cowan formula relating observed proportions of hits and false alarms directly to the probability that an item has been stored in memory does not take into account these interactions between the memory system and the probe stimulus, nor does it account for the noise of the comparison process itself, an argument also echoed by Wilken and Ma (2004). As such, capacity estimates in the change detection literature—based on the Pashler–Cowan formula—systematically underestimate the capacity of the visual short-term memory system to represent items simultaneously. Further, the decrement due to the encoding of the probe array may be contingent on the complexity of the item. Awh, Barton, and Vogel (2007) argued that an earlier study conducted by Alvarez and Cavanagh (2004), where the complexity of classes of visual items was associated with the number of items of that class that could be stored in memory, was the result of comparison errors between the memory representation and probe item. The current results support this argument, indicating that items in the probe array are subject to the same memory constraints as representations formed from the memory array. In addition to this support, the current results suggest a direct means for testing claims regarding errors due to a comparison process would be to compare mAFC discrimination performance to an equated post-stimulus probe change detection task, and examining the magnitude of the difference in sensitivity found between the two tasks.

The current results also potentially implicate contemporary visual short-term memory work employing a continuous report paradigm. Previous studies have reported that filled displays during the retention interval can affect observer performance in a change detection experiment (Makovski & Pertzov, 2015; Makovski, Shim, & Jiang, 2006), indicating that visual short-term memory representations—whilst robust to interruption masking—may not be impervious to all subsequently presented information. The decrement in performance noted by Makovski and colleagues (2009) for their 2AFC task further reinforces this relationship. The more complex responding required by continuous report tasks, where a stimulus standard is manipulated to recreate a stored perceptual representation, is difficult to appraise in terms of its effect on the perceptual representation itself. Further research where joint fitting of continuous report data and data obtained in tasks which require comparison to an internal stimulus standard or a fixed comparator may be beneficial in further specifying the directionality of the relationship between representation and response.

3.4 General discussion

The accuracy data obtained in these experiments represents an important, but incomplete, account of responding in these tasks. The response time data obtained in the two experiments presented in this chapter will be examined in the following chapter, allowing further specification of the decision processes that distinguish orientation discrimination from change detection from the memory processes that unite these tasks.

Supplementary tables and figures

Best fitting signal detection theory model fits for individual observers in Experiment 1.	Tables 3.4–3.8
---	----------------

Best fitting parameters for the top signal detection theory models for each individual	Table 3.9
--	-----------

Table 3.4: The five best fitting signal detection theory models, in terms of BIC, for the observer BF in Experiment 1.

#	Model type	SS constraint	Denominator	Bias	d' weight	N. pars	χ^2	BIC
1	Diff.	Yes	0	Equal	No	1	26.124	33.484
2	Diff.	Yes	0	Orientation	No	2	21.388	36.552
3	Diff.	Yes	0	Change	No	2	23.081	38.106
4	Diff.	Yes	0	Equal	Yes	2	24.522	39.086
5	Ind.Obs.	Yes	—	Optimal	Yes	2	24.621	39.212

Notes: Following the format of Table 3.2.

Table 3.5: The five best fitting signal detection theory models, in terms of BIC, for the observer DB in Experiment 1.

#	Model type	SS constraint	Denominator	Bias	d' weight	N. pars	χ^2	BIC
1	Diff.	Yes	1	Both	Yes	4	16.572	46.103
2	Diff.	Yes	1	Orientation	Yes	3	30.671	52.635
3	Diff.	Yes	Estimated	Both	Yes	5	15.692	52.816
4	Diff.	Yes	0	Both	No	3	30.552	53.452
5	Diff.	Yes	0	Both	Yes	4	24.089	53.873

Notes: Following the format of Table 3.2.

Table 3.6: The five best fitting signal detection theory models, in terms of BIC, for the observer KT in Experiment 1.

#	Model type	SS constraint	Denominator	Bias	d' weight	N. pars	χ^2	BIC
1	Diff.	Yes	1	Orientation	Yes	3	35.705	58.956
2	Diff.	Yes	0	Orientation	No	2	44.133	59.833
3	Diff.	Yes	Estimated	Orientation	Yes	4	34.162	64.988
4	Diff.	Yes	0	Orientation	Yes	3	42.914	66.131
5	Diff.	Yes	1	Orientation	No	2	49.94	66.621

Notes: Following the format of Table 3.2.

Table 3.7: The five best fitting signal detection theory models, in terms of BIC, for the observer SL in Experiment 1.

#	Model type	SS constraint	Denominator	Bias	d' weight	N. pars	χ^2	BIC
1	Diff.	Yes	1	Equal	Yes	2	28.173	43.397
2	Ind.Obs.	Yes	—	Optimal	Yes	2	30.28	45.021
3	Diff.	Yes	0	Equal	No	1	38.66	47.6
4	Diff.	Yes	0	Equal	Yes	2	33.142	48.468
5	Diff.	Yes	1	Orientation	Yes	3	25.804	49.346

Notes: Following the format of Table 3.2.

3.4 General discussion

Table 3.8: The five best fitting signal detection theory models, in terms of BIC, for the observer SS in Experiment 1.

#	Model type	SS constraint	Denominator	Bias	d' weight	N. pars	χ^2	BIC
1	Diff.	Yes	1	Both	Yes	4	11.028	41.271
2	Diff.	Yes	Estimated	Both	Yes	5	9.736	47.506
3	Diff.	Yes	0	Both	No	3	26.423	48.678
4	Diff.	Yes	0	Both	Yes	4	19.019	49.255
5	Diff.	Yes	Estimated	Both	No	4	26.423	56.237

Notes: Following the format of Table 3.2.

Table 3.9: The minimum χ^2 parameter estimates for the differencing model and independent observations models that best fit each individual observer.

Differencing rule							
Observer	d'	d' weight	Orientation β	Change β	Denominator	χ^2	BIC
Avg	1.79	1	1	—	0	12.991	20.542
BF	1.728	1	1	—	0	26.124	33.484
SS	1.47	1.667	1.208	1.484	1	11.028	41.271
SL	1.754	1.569	1	—	1	28.173	43.397
DB	1.875	1.555	0.79	1.505	1	16.572	46.103
KT	1.69	1.425	0.794	—	1	35.705	58.956
Independent observations rule							
Observer	d'	d' weight	Orientation β	Change β		χ^2	BIC
Avg	1.544	2.799	1	1	—	7.806	22.917
BF	1.621	1.79	1	1	—	24.621	39.212
SS	1.519	2.531	1.29	0.091	—	31.535	60.138
SL	1.666	3.537	1	1	—	30.28	45.021
DB	1.774	4.44	0.798	1	—	33.028	54.727
KT	1.694	2.1	0.733	0.135	—	65.235	90.213

Notes: These models are both constrained by the sample-size relation. “Denominator” refers to the addition to the denominator of the sample-size relation in change detection trials using a differencing rule when compared to orientation discrimination trials.

Chapter 4

Change detection and orientation discrimination: response time

Theoretical treatments of response time have largely been missing in the visual short-term memory literature. The majority of empirical work examining visual short-term memory has employed accuracy as the principal dependent variable, usually either in the dichotomous sense of correct and incorrect responses in a change detection paradigm or in the continuous sense with reproduction error in continuous report tasks. As outlined in the previous chapter, and in the preceding literature review, the main models of the architecture of memory storage have been inferred, for the most part, from patterns of observer performance over experimental manipulations of memory load. Any utility of considering response proportions and response time jointly is often, therefore, overlooked as accuracy-only models relate changes in observer performance to constraints on the memory representations directly, without considering the contribution of any intervening decision process.

In the last chapter, I reported the results of two experiments examining the role of the task structure of observer performance, demonstrating a systematic decrease in observer accuracy in change detection tasks when compared to a two-alternative orthogonal orientation discrimination task. An analysis of these data in terms of models drawn from signal detection theory indicated that the decrement in performance in the change detection task was consistent with the cost of an additional item—presumably the probe item—being stored in memory. The explanation given for these data provided by the models affirm the idea that decision processes must be considered explicitly in order to interpret behavioural data. The consideration of response time further specifies this link

⁰Portions of this chapter have been previously presented at the 2015 Australasian Mathematical Psychology Conference held in Shoal Bay.

between representation and response. This chapter extends the analysis of the experiments presented in the last chapter to investigate the distributions of response time.

The use of response time analysis in examining short-term memory characteristics has pedigree in experimental psychology: Sternberg (1969) generalised the method of Donders (1868/1969) in describing the retrieval of representations from short-term memory through the use of mean response times. By observing changes in mean response time with manipulations in the experimental procedure—for instance, by changing the discriminability of the stimuli or the nature of the response—Sternberg was able to outline a detailed mechanistic model of short-term memory “scanning”: the cost of memory access times. The subsequent investigation of the memory scanning paradigm by Ratcliff (1978), beyond the level of the mean response time to consider the entire distribution of response times, allowed even greater detail in providing a theoretical treatment of memory retrieval and decision-making.

Sequential sampling theories, and the diffusion model in particular, are the basis for the theoretical treatment of response time distributions in the visual short-term memory literature and more generally. These models are analogous to a dynamic version of the models from signal detection theory outlined in the previous chapter (§3.1): the models assume intrinsic variability (usually attributed to neural noise) in underlying perceptual representations, and cast the determination of a response from these representations as a process of statistical inference. In the case of signal detection theory, the perceptual representations within a single trial are random variables drawn from stationary distributions, with the distance between the noise distribution and signal distributions (as well as a criterion value) determining the response probabilities. In the case of sequential sampling models, the response is determined through a sampling procedure—either modelled in discrete time or continuous time—where a threshold is met through the integration of many samples drawn from a distribution. Depending on the assumptions placed upon the dynamic time course of decision-making in these models, this distribution can either be stationary (representing a fixed distribution of stimulus information given a particular experimental condition) or can vary as a function of time (in the case of models where stimulus information changes dynamically) or as a function of the state of the decision process (as in the case of modelling a decision as an Ornstein–Uhlenbeck process; see Smith 2000a). In the current work, I shall be focusing solely on the simplest case, where the distribution of stimulus information is stationary across the course of the decision process, differing only between experimental manipulations.

The assumption of sampling stimulus information over time in a sequential sampling decision procedure supplies both the response probabilities—as signal detection theory

does—and also a distribution over response times. The number of samples taken to reach a response threshold in the discrete case, or the time taken to reach a threshold if sampling at a continuous rate, is added to a non-decision time component to obtain an overall response time.

The joint estimation of response proportions and response times imposes more stringent constraints on the specification of the decision process and on the representations supporting the decision process, at a level of precision that might otherwise be underdetermined or non-identifiable at the level of response proportions or mean response times considered separately. Ratcliff and Rouder (2000) used this additional specificity to differentiate between competing accounts of visual short-term dynamics in an examination of backwards masked two-alternative letter identification. The experimental paradigm examined was similar to the single item, orientation discrimination conditions of the first experiment presented in the previous chapter. As only a single item was required for a response, rather than the post-stimulus probe paradigm I employed in the preceding chapter, any decision about the identity of the stimulus could proceed from the first moment where stimulus information was physically available. Ratcliff and Rouder considered two claims about the stability of the perceptual representation underlying the decision process: the first claim was that the stimulus could be integrated during the exposure duration of the stimulus, but would be lost upon presentation of the backwards mask (consistent with the findings of early iconic memory and masking research; see, §2.1); the second claim was that information from the entire stimulus presentation period could be integrated, and the amount of integrated information could be retained after the presentation of the backwards mask. Their application of a modified diffusion model supported the second account, indicating that the identification was based on drawing perceptual evidence from a stable representation and implying a role for visual short-term memory even in a very simple identification task.

Smith et al. (2004) refined this account with a more elaborated form of the diffusion model. Examining observer sensitivity in cued near-threshold detection, Smith and colleagues extended the diffusion model to include both the integrated neural response from stimulus presentation—as Ratcliff and Rouder had—but also a dynamic attentional gate on the rate at which information could be extracted from different locations on the stimulus display. Their finding that observer sensitivity was heightened for cued stimuli only when stimuli were backwards masked—known as *mask-dependent cuing*, was consistent with predictions of the elaborated diffusion model. Like Ratcliff and Rouder, Smith and colleagues found evidence for a stable, short-term memory representation consisting of sensory information—modulated by attentional dynamics—used as evi-

dence in perceptual decision making.

Both of the studies reported by Ratcliff and Rouder (2000) and Smith, Ratcliff, and Wolfgang (2004) provide a compelling demonstration with response time modelling that visual short-term memory representations intercede in quite simple visual tasks, but do not attempt to specify the storage or architecture characteristics of that memory store. Subsequent work by Donkin, Nosofsky, Gold, and Shiffrin (2013) used a response time model related to the diffusion model, known as the linear ballistic accumulator (LBA; Brown & Heathcote, 2008), to examine performance in a change detection task replicating that of Rouder et al. (2008). Using two families of response time models, Donkin and colleagues attempted to distinguish between the predictions of a discrete item or slot model and the predictions of a resource model. The two families of linear ballistic accumulator model differed in the manner in which guessing was handled: in one case, all responses were made on the basis of the accumulation of information from a stored representation, with the quality of the information stored dependent on the number of items simultaneously represented within the memory system; in the other case, responses were either based on a precise memory-based accumulation procedure—indicating the proper storage of an item—or, if the stimulus probed was not stored in memory, based on an information-free “guessing” process. The former model was proposed to be analogous to the flexibly divisible resource model of memory, extended to predict response time distributions; the latter model was fixed such that only a small number of items could be stored, making the model analogous to a discrete-item slot-based model. Donkin and colleagues found strong support for the mixture memory–guess model over a model with continuous degradation of the memory quality, indicating support for a slot-based model, with some evidence that there may be small changes in memory quality when comparing displays smaller than the discrete item limit.

The studies examined demonstrate some precedent for looking at response time as a dependent variable in visual tasks involving memory. In particular, the work of Donkin and colleagues indicates that a sequential sampling response time model can be applied to a change detection task—a proposition that is not trivial. Response time models in the visual domain often deal with speeded responses in perceptual tasks that require very little in the way of higher order elaboration which might break the nexus between observed response times and experimental manipulations; examples include response time modelling of visual detection (Smith, 1995; Smith et al., 2009), brightness discrimination (Ratcliff, 2002; Ratcliff & McKoon, 2007; Ratcliff & Rouder, 1998; Ratcliff, Thapar, & McKoon, 2001; Ratcliff, Thapar, & McKoon, 2003), letter identification (Ratcliff & Rouder, 2000; Ratcliff & Smith, 2010), and motion discrimination (Mazurek, Roit-

man, Ditterich, & Shadlen, 2003; Palmer, Huk, & Shadlen, 2005; Ratcliff & McKoon, 2008). Although change detection is a straightforward task, it relies on the construction, maintenance, and retrieval of a memory representation which may necessitate further considerations about the application of a response time model. The demonstration by Donkin and colleagues is that sequential sampling techniques can be applied to change detection tasks to obtain meaningful conclusions about the structure of the memory representation, even at a broad level.

This chapter continues the response time analysis of visual short-term memory tasks with the consideration of two-choice orientation discrimination and its relationship to change detection. These response time data are taken from the experiments presented in the preceding chapter. As was argued in the last chapter, a comparison between the different types of memory tasks may be illuminating in defining the decision process, more clearly specifying the way memory representations are involved in the decision process, and characterising the role of the task demands in the variation observed in the data. Each experiment will be discussed and analysed in turn in the next two sections.

4.1 Experiment 1

The method and accuracy results for the first experiment are discussed in §3.2.1 and §3.2.2, respectively. Recall that the intent of the first experiment was to examine the relationship between two different visual memory tasks—orientation discrimination and change detection—and where the memory array size was manipulated across trials and the stimulus contrast was fixed to be equal across the two tasks.

4.1.1 Results

For the descriptive statistics shown in this section, no response time data was filtered from overall dataset. Figure 4.1 shows the mean response time for each decision type and memory array size condition, averaged across observers; Figure 4.2 shows these data for each observer separately.

A mixed-effects linear model was conducted to provide a preliminary analysis of the response time data at the mean level, prior to further analysis at the level of the entire distribution. The regression was conducted on mean response time as the dependent variable and, like the analysis reported in §3.2.2, the regression treated the memory array size and the decision type as fixed effects (change detection treated as a baseline for the decision type dichotomous variable), with the individual differences of the observer




Figure 4.1: Group average data from Experiment 1 across different memory array sizes and conditioned on the decision type. Error bars represent one standard error of the mean. CD = change detection; OD = orientation discrimination.

4.1 Experiment 1




Figure 4.2: Mean response time data from Experiment 1, displayed for each observer across different memory array sizes and conditioned on the decision type. Error bars represent one standard error of the mean. CD = change detection; OD = orientation discrimination.

treated as a random effect on the intercept. A significant main effect of memory array size on response time (reported in this section in milliseconds) was significant, $\beta = 88.607$, $SE = 3.1869$, $p < 0.001$. A significant main effect of the decision type was also seen, $\beta = -165.727$, $SE = 14.986$, $p < 0.001$, with the change detection condition slower than the orientation discrimination condition. No significant interaction was found in the effect of memory array size and the effect of decision type (with change detection as a baseline) on response time, $\beta = 9.694$, $SE = 5.472$, $p = 0.076$.

As is evident from both from Figures 4.1 and 4.2, and from the mixed-effects modelling, there is a systematic and consistent effect of memory array size and decision type on mean response time. Most strikingly, change detection decisions seem to add a constant cost in response time at the mean level. This effect can be seen clearly in the plots for the kernel density estimate of the probability density function (pdf) for response times, shown in Figure 4.3. In order to examine this effect more closely, prior to examining the predictions of a full response time model (the diffusion model), the next section will examine convolution as a method for providing some relationship between the response time distributions obtained in either task.

4.1.2 Convolution analysis

A qualitative similarity between the mean response times in the orientation discrimination task and change detection task over all levels of memory array size can be clearly seen in the data for all observers. In each case, the mean response time for the change detection task appears to be a constant time slower—estimated to be approximately 165 ms slower—than the mean response time in the orientation discrimination. A parsimonious explanation for this delay is that an additional stage of processing involving the encoding and comparison of the probe item in the target location is required prior to the decision stage when performing in a change detection task, but not in an orientation discrimination task. This explanation is consistent with the difference in the overall observer sensitivity between the two tasks, where change detection performance seems to incur the storage cost of an additional item—presumably the probe item—when compared to performance in the orientation discrimination conditions.

A constant additive relationship is apparent quantitatively through the use of the mixed-effects linear modelling reported in the last section. A relationship is qualitatively apparent in Figure 4.3 for the entire response time distribution. A straightforward quantitative means of demonstrating this relationship at the level of the distribution, under the assumption that the encoding and comparison of the probe array is an independent pro-

4.1 Experiment 1




Figure 4.3: Kernel density estimates for two observer response time distributions in Experiment 1, across different memory array sizes and conditioned on the decision type. The small lines at the bottom of the chart show the mean response time for each condition. The thicker lines show the orientation discrimination distribution and mean, the thinner lines show the change detection distribution and mean.

cess, is through the use of convolution.

By convention, response times are thought of in terms of random processes, and as the final completion time of multiple mental subprocesses, each of which can itself be represented as a random process (Luce, 1986; Townsend & Ashby, 1983). When these subprocesses are mutually exhaustive and independent, the overall response time can be decomposed as the summation of a number of random variates, each representing the duration of each subprocess. The assumption of independence simplifies this decomposition. In particular, the pdf of the sum of independent random variables is equal to the convolution of each pdf for the constitutive processes; that is, assuming A and B are random variables—representing the finishing time of two psychological subprocesses—and p_A and p_B are their respective density functions, the pdf for the overall response time (p_R) will be equal to

$$p_R = p_A * p_B,$$

where $*$ denotes the convolution operation. The process of deconvolution allows the extraction of one component density function using the density functions of the summed process and the remaining constitutive processes¹. For simplicity, we shall call the distribution that is assumed to be the product of the entire set of psychological processes under examination the *composite distribution* and we shall call the distribution assumed to contain a subset of the psychological processes under examination the *component distribution*.

The convolution and deconvolution of two signals—probability distributions or otherwise—are computed numerically using the discrete Fourier transform. The convolution theorem states that given two functions over time² t , $f(t)$ and $g(t)$, and their respective Fourier transforms, $\mathcal{F}(f)$ and $\mathcal{F}(g)$, the convolution of the two functions can be found by the pointwise product of their Fourier transforms,

$$(f * g) \propto \mathcal{F}^{-1} [\mathcal{F}(f) \cdot \mathcal{F}(g)],$$

denoting the inverse Fourier transform as $\mathcal{F}^{-1}$ (Kreyszig, 2011, pp. 574–575). The constant of proportionality depends on the normalisation constant of the Fourier transform and, in the case of probability distributions, simply requires the convolved distribution, $f * g$, to be renormalised to unity.

¹ A more modern technique, often used in image and audio signal analysis, known as “blind deconvolution” does not require any of the original constitutive components, but employs the maximum likelihood or Bayesian estimation of theoretically plausible selected kernel functions specified *a priori* to recover the original signal, somewhat similar to the technique given below.

² The argument t is omitted subsequently for clarity.

In practice, the deconvolution from empirical response time data is hampered by noise in both the composite and component density functions used, and by the truncation caused by floating point representations (Smith, 1990). These problems can be ameliorated by use of a lowpass filter, using a standard window function, or by convolving a theoretically plausible response time distribution specified *a priori* to the empirical component distribution to obtain a prediction of the composite distribution. The remainder of this chapter will focus on the latter approach. I refer to the plausible distribution as the *model distribution*. By adjusting parameters such as the location (mean) of the model distribution, convolving the model distribution with the component distribution, and comparing the resulting convolved distribution against the composite distribution, a close approximation of deconvolution can be achieved without the problems of artifacts due to noise and a loss of precision in the measurements. This does not obviate the problem of model mimicry, raised by Sheu and Ratcliff (1995), but does provide preliminary evidence of a simple relationship between the experimental conditions prior to a more comprehensive analysis.

The method of deconvolution is as follows. First, after incorrect responses were filtered from the dataset, a kernel density estimate of the component distribution is obtained. A Gaussian kernel was used in the current analysis, as the derivative of the kernel was smoother approaching zero than other common kernel selections such as a rectangular function or the Epanechnikov kernel. The bandwidth of the kernel was scaled such that the standard deviation of the Gaussian function was fifty milliseconds, providing a good trade-off between finer detail and a smooth overall distribution. The kernel density estimate of the component distribution was then convolved with a model distribution. Different functional forms of the model distribution were used: a delta function, a Gaussian distribution, an exponential distribution, and a uniform distribution. Apart from the delta function, these distributions had two parameters: a location parameter (translation across the time axis) and a scale parameter (a dispersion or variance parameter); the delta function only possesses a non-negative location parameter.

These particular functional forms were selected both due to their use in the response time modelling literature (for an extensive review, see Luce, 1986) and due to their distinct functional forms. The Gaussian distribution is obtained as the distribution for a sum of independent and (nearly) identically distributed random processes; the exponential distribution represents a psychological process in which the chance of a still-active process finishing in the next moment is always a constant (that is, an exponential distribution has a constant hazard function); the uniform distribution indicates equal probabilities of a psychological process finishing between two times; and a delta function




Figure 4.4: The different forms of the model distributions used in the convolutional analysis: a) a Gaussian distribution; b) an exponential distribution; c) a uniform distribution; and, d) a delta function.

indicates that the process linking the two distributions is invariant across all trials (that is, the process completes in constant time).

After the component distribution was convolved with the model distribution, the resulting distribution was compared with the empirical composite distribution by first constructing bins from twenty-five quantiles equally spaced between probabilities of 0.01 and 0.99. The number of observed response times within these bins was compared with the number of predicted response times according to the convolved distribution using bins with the same bin edges (that is, the same response times). A goodness-of-fit statistic, G^2 was computed between the number of predicted observations within each bin against the number of empirical observations within each bin:

$$G^2 = 2 \sum_i \sum_j N_{ij} \sum_k p_{ijk} \cdot \ln \frac{p_{ijk}}{\pi_{ijk}},$$

where i indexes the decision type condition, j indexes the memory array size condition, k indexes the bin, N_{ij} is the total number of observations in the specified condition, p_{ijk} is the proportion of observed change detection observations in a given bin, π_{ijk} is the proportion of predicted change detection observations within a given bin. The location and scale parameters of the model distribution were then iteratively adjusted by minimizing the objective function using the Nelder–Mead simplex method (Nelder & Mead, 1965) to obtain an estimate of the maximum likelihood parameter set. The information criterion BIC was then computed for the best fitting parameter set $BIC = G^2 + k \ln(N)$, where k is number of parameters to construct the model distribution (one in the case of the delta function, two in the other cases) and N is the number of total observations for the observer.

Table 4.1 shows the goodness-of-fit statistics in terms of G^2 and BIC for each distribution type and each observer. In examining the G^2 and BIC results across different

Table 4.1: Goodness-of-fit and information criteria statistics for the convolution analysis of Experiment 1, conditioned on observer and on distribution type.

Observer	Gaussian		Delta		Uniform		Exponential	
	G ²	BIC	G ²	BIC	G ²	BIC	G ²	BIC
BF	229.773	244.098	230.087	237.249	229.71	244.035	230.005	244.33
DB	194.745	209.19	204.08	211.303	193.271	207.717	148.091	162.536
KT	145.058	159.428	198.284	205.469	150.887	165.258	141.682	156.053
SL	203.552	218.335	209.853	217.245	203.774	218.557	201.166	215.949
SS	190.149	204.47	190.07	197.23	190.15	204.47	190.15	204.47

distributions within the results of each observer, the difference between the goodness-of-fit statistics is rarely large. For observers DB, KT, and SL, an exponential distribution convolved with the orientation discrimination distribution to estimate the change detection response time distribution produces the lowest overall BIC. For observers BF and SS, the delta function produces the lowest overall BIC; that is, the best fitting solution when corrected for complexity is simply shifting the orientation discrimination function to be slower in time. As convolution with the delta function and convolution with any of the other distributions only differs by a single parameter in terms of the model complexity, the penalty imposed by the BIC is not severe. Given that the two winning convolving functions were the exponential distribution, where shorter times for the completion of the comparison process are more probable than longer completion times, and the delta function, where all comparison processes take the same length of time, it is reasonable to conclude that the speed of encoding and comparing the probe array with the memory array is relatively invariant and other factors, including the decision time from that compared representation, form the majority of the variance seen in the response time distribution.

Figures 4.5–4.9 show the density estimates of both the observed orientation discrimination and change detection response time distributions and the result of the convolution with the delta function and with the best-fitting exponential distribution. Table 4.2 shows the maximum likelihood parameters for the both the delta function and exponential distribution models by observer. Note that, in the case of BF and SS, the dispersion parameter of the exponential distribution is small, meaning that the best fitting exponential distribution approximates the delta function (the delta function is obtained as the limit of the exponential distribution as the dispersion parameter approaches zero). This is consistent with the BIC rankings and the minor contribution of the comparison process to the overall variance in the response time distributions. I will elaborate upon

Table 4.2: The maximum likelihood parameters for the convolution analysis of Experiment 1 for the exponential and delta convolving functions, displayed by observer. The exponential function has both a dispersion parameter (also called a variance or scale parameter) and an offset parameter (also called a mean or location parameter). The delta function only has an offset parameter.

Observer	Delta	Exponential	
	Offset	Dispersion	Offset
BF	161.868	13.63	147.157
DB	186.094	97.676	107.159
KT	122.16	109.332	39.412
SL	188.146	44.531	148.759
SS	129.032	0.014	128.394

these using diffusion modelling in the next section.

4.1 Experiment 1




Figure 4.5: The results of convolving a delta function (left panels) and an exponential distribution (right panels) with kernel density estimate of the correct trials in orientation discrimination conditions—shown by the thin unbroken line—in order to obtain a fit to the density estimate of the correct trials in the change detection condition—shown by the dashed line. The result of the convolution is shown with the thick unbroken line. The data shown above are for observer BF in Experiment 1.




Figure 4.6: The results of the convolution analysis using a delta function (left panels) and an exponential distribution (right panels). The data shown above are for observer DB in Experiment 1, following the format of Figure 4.5.

4.1 Experiment 1




Figure 4.7: The results of the convolution analysis using a delta function (left panels) and an exponential distribution (right panels). The data shown above are for observer KT in Experiment 1, following the format of Figure 4.5.




Figure 4.8: The results of the convolution analysis using a delta function (left panels) and an exponential distribution (right panels). The data shown above are for observer SL in Experiment 1, following the format of Figure 4.5.

4.1 Experiment 1




Figure 4.9: The results of the convolution analysis using a delta function (left panels) and an exponential distribution (right panels). The data shown above are for observer SS in Experiment 1, following the format of Figure 4.5.

4.1.3 Diffusion model

The deconvolution results presented above provide a preliminary analysis of the entire response time distributions obtained in Experiment 1. Although these results show clear evidence that the distributions obtained in the change detection conditions can be thought of as being the product of all of the processes present in the orientation discrimination condition plus an additional comparison or probe array processing component, the analysis itself does not provide any insight into the nature of the processes present in the orientation discrimination distribution or the additional processes present in the change detection condition, other than their existence. In order to provide a more extensive theoretical treatment, additional assumptions are required.

One of the most successful and widely applied approaches to the joint analysis of response time distributions with response probabilities is the diffusion model due to Ratcliff (1978). The diffusion model is a member of a larger class of sequential-sampling decision models. These models assume that a decision process draws samples over time from perceptual (or other mental) representations, using each sample drawn to incrementally move towards a final decision, each sample taken building upon the accumulated evidence to that point. Evidence is sampled until there is sufficient evidence to execute a response with the determination of what constitutes sufficient evidence differing between classes of models: in some (accumulator) models, such as the LBA mentioned in the introduction, evidence is accumulated independently for different response options until the absolute level of evidence for one option reaches a threshold; in other models, such as the diffusion model, a response is executed when the balance of evidence favours one alternative over the others³. The threshold which defines the amount of information, either in absolute terms or relative terms, required to complete the decision and execute the response is known as the decision or absorbing boundary.

In the case of a judgement between two alternatives, as in the standard diffusion model, the decision can be thought of as point progressing leftward on the half-plane at a constant rate, buffeted up and down by random perturbations, until reaching a decision boundary running parallel to the direction of time⁴ (see Figure 4.10). The distance between the point at which the decision starts and a decision boundary represents the amount of evidence for a response alternative required before a response is initiated,

³There are some models that attempt a balance of the two approaches, as in the case of the leaky competing accumulator model of Usher and McClelland (2001), where negative interaction between the two accumulators means that, although independently modelled, the balance of information also affects the time taken to reach a decision.

⁴Some versions of the diffusion model allow non-parallel decision boundaries, allowing time-dependent changes in the amount of information required to make a decision.




Figure 4.10: A schematic of the parameters controlling the diffusion model. Evidence for an alternative (in this case, the correct response option) is accumulated through time, with an average level of accumulation defined by a drift rate drawn for each trial from a normal distribution with mean ν and standard deviation η . This accumulation is subject to moment-by-moment variability which is normally distributed, as seen by the perturbations in the sample path. This parameter is not independently identifiable and is usually assigned an arbitrary value which is treated as the fundamental scaling parameter for the model. In this thesis, I have set it to unity. The decision process is terminated when the amount of evidence reaches one of the two absorbing barriers, separated by a distance denoted a . The component of the response time due to processes outside of the main decision process is accounted for by a uniform distribution with mean time T and range s_T , which is added to the time of the completion of the decision process.

allowing a sequential sampling model to account for the speed–accuracy trade-off: the larger the distance, the more information required to initiate a response, but the slower the decision will take on average. The position of the starting point of the evidence accumulation process relative to the boundaries defines the response bias: if the start point is closer to one boundary, this response alternative will be favoured over the other than would otherwise be the case, representing a bias in the initial balance of evidence; if the start point is equally distant from both boundaries, no response alternative will be preferred, the initial balance of evidence favours no response alternative. In some versions of sequential sampling models, random variation in the exact starting point of the decision is also modelled, allowing for “fast errors”, although this modification will be omitted in this thesis.

Different sequential sampling models of decision-making also vary on the presence and characteristics of variability in the sampling process. In the case of the diffusion model, this incremental sampling of evidence occurs continuously through time, rather than in discrete steps, and the information sampled from the perceptual representation is assumed to be perturbed by Gaussian distributed noise. In this way, it is analogous to a continuous time version of the signal detection theory models discussed in the last chapter, and the application of the diffusion model in this chapter building upon the signal detection theory models of the last chapter will further demonstrate that connection. The magnitude of evidence against the magnitude of the background noise is analogous to the sensitivity parameter, d' , in representing overall stimulus discriminability. Unlike the signal detection theory model, however, the magnitude of the evidence can differ on a trial-by-trial basis, drawn from a Gaussian distribution itself.

For this thesis, I consider a two-alternative version of the diffusion model which does not model response bias. In this model, the decision process begins between two boundaries of equal distance away from the start point. This decision process is modelled using five parameters. The first parameter, the *drift rate* denoted ν , is the mean level of evidence (scaled in terms of standard deviations of the within-trial noise) favouring the correct response alternative: a drift rate of zero indicates no useful stimulus information and will result in equal response probabilities for correct responses and errors; larger values indicate a higher level of stimulus information, leading to greater accuracy predicted. The second parameter, *drift variability* denoted η , is the standard deviation of the distribution from which the individual drift rates for each trial are drawn from, with a drift variability of zero indicating the same drift rate is used in each trial. Note that this is different from the diffusion coefficient, denoted σ , which is the within-trial variability in the accumulation of evidence. For identifiability, this coefficient is fixed to one. The third




Figure 4.11: The construction of a quantile-probability plot (QPP). Given two responses, labelled A and B, the response time density functions (shown on the left hand side) are divided into prespecified quantiles (in this thesis, following convention, these quantiles are the 10th, 30th, 50th, 70th, and 90th percentiles). These quantiles are then plotted on the y-axis of a scatterplot, shown on the right hand side of the above figure, using the overall proportion of the response as the value on the x-axis. This shows both accuracy and response time distributional data clearly and compactly: both the moments of a response time distribution, the relative position of that distribution to other distributions in time, and the corresponding response proportion can easily be read from the diagram.

parameter indicates the distance between the decision boundaries, denoted a , defining the amount of information required to complete the decision and initiate a response. The fourth and fifth parameters are the mean nondecision time and the nondecision time range, denoted T (often T_{er} in the literature) and s_T respectively. These parameters control the additional time added to the completion time of the decision process to account for nondecision time in the final response time, such as the time taken to encode the stimulus and the time taken to execute a response.

By constraining different sets of these parameters to vary with the experimental manipulations, different interpretations can be given to the relationship between the experiment and the modelled decision process. For instance, if the drift rate, v , is constrained to be equal between the two decision type conditions in estimating the current experiment, this indicates that—whilst the response proportions or response times might be different between the decision type conditions—the amount of effective information available to

the decision process was assumed to be equal. In this chapter, a series of models will be fit to the experimental data—by observer, with the inclusion of group average data—and their goodness-of-fit will be compared using the BIC to account for differences in model complexity. To fit the diffusion model, I will employ the chi-squared method (e.g., Ratcliff & Smith, 2004) of comparing model predictions of binned data to actual quantile data (at the 10th, 30th, 50th, 70th, and 90th quantiles, following convention), similar to the method used to fit the convolution results in the section above. Rather than minimizing the chi-squared objective function, the G^2 function was used—like that used in the convolution analysis. Additionally, also like the preceding analysis, the Nelder–Mead simplex procedure was used to iteratively minimize the objective function. Unlike the mixed-effects model and convolution analysis, however, response time outliers were removed: trials with responses entered under 100 ms or over 2,750 ms were filtered out, less than 0.3% of the total dataset.

In fitting the current experiment, several modifications to the standard diffusion model had to be made. One notable aspect of each experiment presented in this thesis—as well as the experiment previously reported by Sewell, Lilburn, and Smith (2014)—is that incorrect responses appear to be substantially slower than correct responses even at the smallest quantiles, seen as a “bowing” of the 0.1 quantile (the 10th percentile) known as “leading edge” of the quantile–probability plot (or QPP; see Figure 4.11). Although bowing in the response time quantiles for incorrect responses is often seen in later quantiles, the leading edge of a QPP is often flat, indicating a consistency in the speed of the fastest responses regardless of many different types of experimental manipulation, such as stimulus contrast or stimulus exposure duration (see, however, Smith et al., 2014). Across visual short-term memory conditions, however, a difference in the leading edge can be seen between correct and incorrect responses. Although Ratcliff and Smith (2010) reported a bowing in the leading edge across experimental conditions when simple stimuli (letters) were presented in dynamic noise—like the stimuli used in the current experiments—the bowing occurred between discriminability conditions, rather than between the response outcomes (i.e., correct and incorrect responses). It is unlikely that such a difference in early responding could be caused by a difference in response bias between correct and incorrect responses, as this would indicate at least some insight by an observer into the accuracy of a decision prior to response entry.

One alternative explanation is that, in addition to the relatively fast and accurate memory retrieval process modelled by the main diffusion process, a secondary *intrusion* process models the possibility of sampling from an uncued item (i.e., stimulus information that does not correspond to the location of the probe). This is modelled as a separate dif-

fusion process with a non-positive (zero or negative) drift rate, denoted v_α ; an additional time component, denoted T_α , added to the main nondecision time component; and is combined with the main diffusion process using an intrusion parameter, α . This intrusion parameter α differs from the actual mixture probability in that the probability of sampling from an uncued location changes with the size of the memory array. Assuming that the chance of an unprobed item intruding into the decision process is uniform and that only items which differ from the target identity are sampled from, then the probability of sampling an uncued item u in an m item display is

$$P(u_m) = \sum_{i < m} \alpha \cdot \left(\frac{i}{m}\right) \cdot \binom{m-1}{i} p^i (1-p)^{m-i-1},$$

where p is the probability of a stimulus with a different identity than that of the target (for instance, the probability of a horizontal item when the target is vertical or vice versa). In the current experiment, each orientation is equally probable, so p is set to $\frac{1}{2}$. This means that the intrusion process represents the delaying sampling of stimuli with conflicting identities to the target. A similar mixture process was used by Sewell, Lilburn, and Smith (in press) to allow for “delayed guessing”: in that instance, the intrusion rate was identical to the mixture proportion of the guess process and the drift rate of the guess process was fixed at zero. In the current experiment, when the intrusion drift rate is estimated to be zero, the model predicts a proportion of delayed guessing like that used by Sewell and colleagues, except the mixture probability changes as a function of the memory array size.

The introduction of an intrusion diffusion process allows the prediction of asymmetrical response time distributions across response outcomes: where incorrect responses are slower than correct responses even at the earliest quantiles, consistent with the current data (see Figure 4.12).

In addition to the modification of the diffusion process to include the possibility of unprobed stimulus information intruding into a decision, additional modifications of the diffusion process can be made on the basis of the signal detection theory modelling presented in the previous chapter. The change in performance across different memory array size conditions in both the orientation discrimination and the change detection trials adhered closely to the predictions of the sample size relationship, where observer sensitivity given a display of m items decreased by the square root of the number of items $d'_m = \frac{d'_1}{\sqrt{m}}$. This relationship—as a limit on the effective information available to the decision process—corresponds directly to a constraint on the drift rate v for an m




Figure 4.12: Differences in the leading edge of the response time distributions between standard diffusion model predictions (in the left hand panel denoted "A") and the diffusion model with a negative intrusion process (in the right hand panel denoted "B"). Note the exaggerated asymmetry in the leading edge between correct and incorrect responses (the left hand and right hand side of the quantile closest to the bottom of the figure) for the diffusion model with the intrusion process.

Table 4.3: A list of the parameters controlling the extended form of the diffusion model, with an additional intrusion process, presented in this chapter.

Symbol	Name	Description
v	Drift rate	The average quality of stimulus information used to drive the decision process, where larger positive values indicate stronger evidence accumulated toward a correct response outcome and negative values indicate evidence accumulated towards an incorrect response outcome.
η	Drift variability	The standard deviation in the distribution describing trial-by-trial variability in the drift rate.
a	Boundary separation	The distance between the two response boundaries, representing the amount of evidence required to initiate a response.
T	Nondecision time	The minimum time added to the decision process to account for the contribution of nondecision processes (encoding and responding) in the overall response time.
s_T	Nondecision time range	The range of times for the completion of nondecision processes.
α	Intrusion rate	A weight to determine the probability of information being sampled from non-probed stimulus information.
T_α	Intrusion time	The additional nondecision time incurred by the accessing of non-probed stimulus information.
v_α	Intrusion drift	The quality of evidence derived from non-probed stimuli. This parameter is constrained to be non-positive, meaning that information must be either contributing towards incorrect decisions or represent (zero drift) noise.

item display,

$$v_m = \frac{v_1}{\sqrt{m}}.$$

The signal detection theory model comparison in the last chapter also indicated that the decrement in observer sensitivity in change detection trials when compared to orientation discrimination trials is well captured by an increment in the denominator of the sample size relationship to reflect the additional load of the probe item. This last assumption will be one factor in the many configurations tested.

In all, fifty-two different model configurations were examined for each of the five observers and for the group average (denoted “Avg” in tables and charts). These configurations differed in the number and types of constraints placed upon parameters, and in whether an increment was added to the denominator of the sample size relationship in change detection trials. The large number of the diffusion model configurations tested necessitates a naming system to distinguish between different parameterisations of the

model. In each model, each of the main parameters (see Table 4.3) is estimated for all experimental conditions (that is, for every trial regardless of the memory array size or the decision type) unless otherwise specified. Parameters which are estimated for each condition are not listed in the model name. Parameters which are listed in the model name are adjusted independently depending on the condition. Parameters with a superscript s are estimated separately for each memory array size condition, parameters with a superscript t are estimated for each decision type condition. By way of example, a model with η^t in the name indicates that the value for the trial-to-trial drift variability in the orientation discrimination condition is separate to the drift variability value for change detection trials, but within a decision type all of the drift variability values are the same. Parameters with the superscript $*$ are estimated for each experimental condition (for each level of the memory array size and each decision type), giving a total of eight parameters. The exception to this is the drift rate which may be estimated across trial types, but cannot be separately estimated across memory array conditions, given the sample size constraint.

In addition to the parameter constraints in the model name, the presence of the string “-CD+” in the name indicates that an extra item has been added to the denominator of the sample size relation in change detection trials, following the findings of the last chapter. For clarity, the name of each model configuration is prefixed with the string “MOD-” to indicate that it is a model name.

Of the fifty-two model configurations tested, the top five for each observer, as ranked by BIC, are presented in Table 4.4 for the group average data and in Tables 4.8–4.12 for each individual observer. The tables for individual observers are presented at the end of the chapter for clarity. For three of the five observers, and for the group average data, the best fitting model in terms of BIC has the configuration MOD-CD+-T*: where drift, drift variability, boundary separation, nondecision time variability, the intrusion rate, the intrusion time, and the intrusion drift are all kept equal across experimental conditions, and a nondecision time value is estimated for each experimental condition. This model configuration also includes an additional item within the denominator of the sample size relationship for change detection trials, consistent with the findings of the previous chapter. The number of nondecision time values estimated (eight values in total with one for each experimental condition), although high as a proportion of the total number of estimated values, is also justifiable: the convolutional analysis presented above provides clear evidence that the response time distributions of the change detection and orientation discrimination trials are related by a constant, consistent with an additional nondecision time component. Sewell, Lilburn, and Smith (in press) also reported that, when

Table 4.4: The top five best fitting diffusion model configurations, in terms of BIC, for the group average data in Experiment 1.

#	Model	Free parameters	G^2	BIC
1	MOD-CD+-T*	15	77.797	191.752
2	MOD-CD+- v^t -T*	16	81.074	202.626
3	MOD-T*	15	91.274	205.229
4	MOD-CD+- α^t -T*	16	86.405	207.957
5	MOD-CD+- α^t -T* $-\alpha^t$	17	79.746	208.895

fitting the diffusion model to a near-threshold orientation discrimination task identical to the orientation discrimination trials examined here, that the nondecision time values estimated varied principally as a function of the memory array size, rather than exposure duration of the memory array, which they took as evidence for a limited capacity memory retrieval process. It is also worth noting that the best fitting model includes only a single estimated drift rate: the memory array effects on response time and accuracy are well accounted for by the sample size constraint. A single drift rate is also consistent with the fact that the memory array presentation was identical, including the stimulus contrast, in trials for both the decision types. Figure 4.13 show the fit of the MOD-CD+-T* configuration to the group average data. One of the observers, KT, not best fit by the model configuration MOD-CD+-T* required additional flexibility in the drift rate between trials of different decision types and additional flexibility in the estimation of the rate of the intrusion process. All of the best fitting models for the observer did, however, still require the decrement in the sample-size denominator for change detection trials. The other observer for whom the MOD-CD+-T* configuration was not the preferred model—observer SS—had substantially higher goodness-of-fit values for all diffusion configurations examined, indicating a poorer overall fit.

The parameter estimates for the MOD-CD+-T* model configuration, as well as for MOD-T*, the second best fitting model for many observers, are presented at the end of this chapter in Tables 4.13 and 4.14.




Figure 4.13: QPP for the diffusion model configuration that best fits the group average data in Experiment 1 in terms of BIC: the MOD-CD+-T* model. In this model, a single parameter was estimated for each of the drift rate, drift variability, boundary separation, nondecision time variability, intrusion rate, the drift of the intrusion process, and the additional time of the intrusion process across all trial types. Nondecision time was estimated separately for each trial type, totalling eight free T parameters (one value for each of the four array size levels within each of the two decision type levels). The observed data, the group average response, is displayed as empty circles joined with dashed lines; the model predictions are presented as filled squares connected by solid lines.

4.1.4 Discussion

Both the convolution analysis and the diffusion model fits to data obtained in Experiment 1 were consistent with the findings from the previous chapter that observer sensitivity is decreased when a change detection decision is required, as compared to an orientation discrimination decision, and that this decrement is well accounted for by adding an additional item to the denominator of the sample-size relation. In addition, the response time data show a clear relationship between the orientation discrimination and change detection response time distributions, which were related either by the addition of a constant or a random variable with small variance. A natural interpretation for this additional process is that it represents the encoding and storage of the probed stimulus, prior to the commencement of a decision. This is also supportive of the proposition that the accuracy decrement seen from orientation discrimination performance to change detection performance is due to the additional load of the probe item within memory over the course of the entire decision.

In addition to the change in nondecision time observed between the two tasks, the best fitting diffusion model also required the nondecision parameters to increase with the size of the memory array, like the model fits reported by Sewell, Lilburn, and Smith (in press). The interpretation of these values by Sewell and colleagues was that, as mentioned earlier, the time taken to find and retrieve an active representation within the memory system increased as a function of the number of concurrently maintained items. The increasing nondecision times estimated for increasing memory array sizes support this conclusion. As mentioned above, Sewell and colleagues concluded that this pattern of nondecision time values reflected a limited capacity to hold elements within visual short-term memory in a highly activated state, ready for retrieval upon cueing. This conclusion is consistent with the concentric activation framework proposed by Oberauer and colleagues Oberauer (2002, 2006); Oberauer and Bialkova (2009); Oberauer and Hein (2012), where activation and retrieval of items stored within working memory—more broadly defined than visual short-term memory—relies on a serial process that cannot be divided when multiple items are required.

In addition to the observable effects ascribed to the retrieval of probed stimulus information from memory, a potential signature of unprobed information entering into the decision process was found. The asymmetry in the earliest quantiles between distributions taken from the correct and incorrect responses does not generally appear in simple detection data (e.g., Ratcliff & Rouder, 2000; Smith, 1995; Smith et al., 2014), but appears in visual short-term memory tasks such as the ones reported in this thesis and

the experiment reported by Sewell and colleagues (2016). To account for this asymmetry, a separate intrusion process was assumed, allowing for incorrect information to be accumulated on occasion with an additional time component. This additional process allows the model additional flexibility in accounting for response time biases at the earliest quantiles.

Each of these results highlight, along with the results of the last chapter, the non-triviality of the retrieval of a representation from memory. The role of representation access and retrieval is an integral component of theories of working memory generally, but is overlooked within current visual short-term memory theories. This is despite, as shown here, these processes might contribute substantially and systematically to both the accuracy and response time profiles.

More generally, the success of the modified diffusion model—with an intrusion process and modified sample size constraints—to capture the effect of memory load on accuracy and response time across qualitatively different types of decisions provides a unification of previous response time analyses in the literature, which had previously treated change detection (Donkin et al., 2013) and orientation discrimination (Sewell et al., in press) separately. The next section applies both the convolution analysis and the diffusion modelling to the second experiment, where accuracy was equated.

4.2 Experiment 2

Recall that the intent of the second experiment introduced in the last chapter (§3.3) was to examine the relationship between orientation discrimination and change detection over different stimulus exposure durations, where the stimulus contrast for each task was manipulated to equate performance in terms of accuracy. An increase in the contrast of memory array, allowing observer performance to be equated between the two tasks in terms of accuracy, also provides a means to dissociating the performance decrement seen in change detection from the additional nondecision time in change detection trials. The previous chapter contains details of the the method and accuracy results for this experiment, discussed in §3.3.1 and §3.3.2 respectively. Following the presentation of the last section, this section will first examine the response time data at the mean-level, using mixed-effects linear modelling, prior to a convolution analysis and, finally, a analysis of diffusion model fits.




Figure 4.14: Group average data from Experiment 2 across different stimulus exposure durations and conditioned on the decision type. Error bars represent one standard error of the mean. CD = change detection; OD = orientation discrimination.

4.2.1 Results

Like the preceding experiment, no response time data was filtered from the overall dataset for the descriptive and inferential statistics presented in this section. Figure 4.14 shows the mean response time for each decision type and stimulus exposure duration, averaged across observers; Figure 4.15 shows these data for each observer separately.

A mixed-effects linear model was conducted, with response time as the dependent variable, the decision type (treating change detection as a baseline) and stimulus exposure duration treated as fixed effects, and the individual differences in response times treated as a random effect on the intercept. The stimulus exposure duration was scaled to be in seconds, rather than milliseconds, to increase the stability of the model fitting routine. A significant main effect of exposure duration was found on response time, $\beta = -621.359$, $SE = 59.666$, $p < 0.001$, with an increasing stimulus exposure duration leading to a decreasing estimated response time. A significant main effect of decision type was also found, $\beta = -128.014$, $SE = 13.117$, $p < 0.001$, with the predicted response




Figure 4.15: Mean response time data from Experiment 2, displayed for each observer across different stimulus exposure durations and conditioned on the decision type. Error bars represent one standard error of the mean. CD = change detection; OD = orientation discrimination.

Table 4.5: Goodness-of-fit and information criteria statistics for the deconvolution analysis of Experiment 2, conditioned on observer and on distribution type.

Observer	Gaussian		Delta		Uniform		Exponential	
	G ²	BIC	G ²	BIC	G ²	BIC	G ²	BIC
BF	372.44	389.147	372.435	380.788	372.457	389.164	372.457	389.164
DB	219.213	236.031	219.988	228.397	219.16	235.977	219.08	235.898
KT	137.924	154.383	254.278	262.507	202.977	219.436	127.685	144.144
SL	700.327	717.107	789.321	797.71	752.847	769.626	598.972	615.751
SS	154.738	171.436	154.738	163.087	154.765	171.463	154.747	171.445

times for the change detection task slower than those for the orientation discrimination task. Last, a significant two-way interaction between stimulus exposure duration and decision type was found, $\beta = 293.397$, $SE = 84.380$, $p < 0.001$. This two-way interaction can be seen in the flatter curve between exposure duration conditions in change detection trial types when compared to orientation discrimination trial types.

In comparison to the mean response time results of Experiment 1, this reduced effect of exposure duration on mean response times for change detection trials compared to orientation discrimination trials may reflect the increased amount of stimulus information presented in the change detection trials: recall that, in order to equate performance between change detection and orientation discrimination whilst keeping the same trial timing, the change detection memory array was presented at a higher contrast, meaning that the effect of limiting visual information through truncating stimulus presentation may have been lessened.

4.2.2 Convolution analysis

A convolution analysis was also conducted for Experiment 2. The analysis followed the same procedure as that reported for Experiment 1 above (§4.1.2): first, trials with incorrect responses were filtered from the dataset; second, a kernel density estimate using a Gaussian kernel function was generated from the orientation discrimination trials to form the component distribution; third, the convolving distribution—either an exponential, Gaussian, or uniform distribution, or a delta function—was generated and convolved with the component distribution; last, the resulting composite distribution was segmented into twenty-six bins from twenty-five quantiles and compared to the empirical change detection response time distribution. Table 4.5 display the goodness-of-fit statistics in terms of G^2 and BIC for each distribution type and each observer.

Table 4.6: The maximum likelihood parameters for the convolutional analysis of Experiment 2 for the exponential and delta convolving functions, displayed by observer. The exponential function has both a dispersion parameter (also called a variance or scale parameter) and an offset parameter (also called a mean or location parameter). The delta function only has an offset parameter.

Observer	Delta	Exponential	
	Offset	Dispersion	Offset
BF	198.848	0.002	199.1
DB	181.642	26.432	156.754
KT	101.308	152.705	0.0
SL	220.936	55.657	165.924
SS	156.079	1.445	154.914

As the overall fit statistics in Table 4.5 show, convolution with the delta function was the best fitting relationship between the orientation discrimination and change detection response time distributions for three of the five observers: observers BF, DB, and SS. For observers KT and SL, convolution with an exponential distribution provided the best overall fit. The set of best fitting parameters for convolution with the delta function and convolution with an exponential distribution can be found in Table 4.6.

The advantage of the delta function in this experiment over the exponential function in BIC terms may reflect the difference in the stimulus contrast—and, therefore, the amount of stimulus information—between the orientation discrimination and change detection conditions. In particular, when examining the density functions of observer BF or observer SL—who both have poor overall fits—it is clear that the composite distribution formed by convolution do not fit the empirical change detection distribution well, with the empirical distribution showing a greater peakedness for faster response times. Convolving additional distributions with the component distribution can only add dispersion to the resulting composite distribution (or, in the case of the delta function, maintain the same level of dispersion), meaning that the higher level of skewness of the change detection distribution cannot be fit by assuming additional processing stages.

Although it appears that there is an additional component of encoding and comparing the probe array to the memory array, like that seen in Experiment 1 and captured by the delta function (or the offset parameter of the exponential distribution), it is difficult with the convolutional analysis alone to separate the effects of the decision stage from the additional processing. In the parameter estimates for the exponential distribution for the data of observer KT, for instance, the offset parameter is estimated at zero, indicating

that increased dispersion in the exponential distribution is partially accounting for the offset in response time. Unlike the first experiment, where the stimulus contrast meant that it was plausible a major source of variance in the decision was identical between task types, the difference in the decision dynamics in the present experiment indicates that, whilst a convolutional approach is useful, a diffusion model approach may provide additional clarity.

4.2.3 Diffusion model

The approach in fitting the diffusion model to Experiment 2 was largely the same as the approach described for Experiment 1. One notable difference, however, is the fact that the more restrictive constraints placed upon drift rates in the previous experiment relied on the fact that the memory load was directly manipulated and that the memory array was invariant between the two decision types: the drift rates were determined by the imposition of both the sample size constraint and the change to the denominator of the sample size constraint between decision types. In the current experiment, the exposure duration of the memory array, rather than the size of the memory array, was manipulated within different decision types. The change in the contrast of the memory array between the two decision types provided additional discriminable stimulus information in the change detection task, to offset the decrement induced by the encoding of the probe array. This means that the modification to the sample size denominator may not be able to be used to predict both types of decision in the task from a single drift rate, but could also lead to differences beyond a change to the drift rate alone.

Building upon the diffusion modelling for Experiment 1, one hundred different model configurations were constructed from different constraints placed upon the relationships between model parameters and the experimental manipulations (in this case, the stimulus exposure duration of the memory array and the type of decision the observer was instructed to make). The large number of additional model configurations, when compared to Experiment 1, were due to the fact that the drift rate was allowed to vary across stimulus exposure duration and decision type in this experiment, where the drift rate in the first experiment was constrained by the sample size relationship. These model configurations follow the same naming conventions as those given for Experiment 1 (a description of the parameters is given in Table 4.3). Following the modelling for the previous experiment, each of the model configurations was fit to the data of each observer independently using an iterative minimisation of a G^2 statistic by the Nelder–Mead simplex procedure. The results of the top five model fits by BIC are given for the group




Figure 4.16: The results of convolving a delta function (left panels) and an exponential distribution (right panels) with kernel density estimate of the correct trials in orientation discrimination conditions—shown by the thin unbroken line—in order to obtain a fit to the density estimate of the correct trials in the change detection condition—shown by the dashed line. The result of the convolution is shown with the thick unbroken line. The data shown above are for observer BF in Experiment 2.




Figure 4.17: The results of the convolution analysis using a delta function (left panels) and an exponential distribution (right panels). The data shown above are for observer DB in Experiment 2, following the format of Figure 4.16.




Figure 4.18: The results of the convolution analysis using a delta function (left panels) and an exponential distribution (right panels). The data shown above are for observer KT in Experiment 2, following the format of Figure 4.16.




Figure 4.19: The results of the convolution analysis using a delta function (left panels) and an exponential distribution (right panels). The data shown above are for observer SL in Experiment 2, following the format of Figure 4.16.




Figure 4.20: The results of the convolution analysis using a delta function (left panels) and an exponential distribution (right panels). The data shown above are for observer SS in Experiment 2, following the format of Figure 4.16.

average in Table 4.7, and for each observer in Tables 4.15–4.19 at the end of the chapter. The parameters for each observer for model configurations $\text{MOD-}\nu^d\text{-}\alpha^t\text{-}T^t$ (a model configuration that fit well for observer KT and the group average data), $\text{MOD-}\nu^d\text{-}T^t\text{-}\alpha^t$ (a model configuration that fit well for observers BF and SS), and $\text{MOD-}\nu^d\text{-}\alpha^t\text{-}T^t\text{-}\alpha^t$ (a model that fit well for observers BF, KT, SS, and the group average) are presented in Tables 4.20, 4.21, and 4.22, at the end of the chapter.

When compared to the model fits obtained from Experiment 1, the rankings of models for Experiment 2 are somewhat more inconsistent, with no clearly preferred models shared across observers. Each of the winning configurations requires additional flexibility in relating the drift rate parameter, ν , to data across conditions. This additional flexibility is both due to the fact that the exposure duration of the memory array directly affects the quality of stimulus information in the decision process (the effect of increasing stimulus exposure duration should be a linear increase in the squared drift rate; see, §5.2.3), and the fact that, unlike in Experiment 1, the display contrast of the memory array was adjusted for each observer to equate performance—in terms of accuracy—between the different decision type conditions. Note also that, unlike Experiment 1, there is no clear preference for models with the increment of an item within the sample size constraint in accounting for change detection performance, likely due to the fact that any performance decrement that would be observed due to the encoding and comparison of elements within the probe array was offset by the increased stimulus contrast. In cases where the drift rate was not constrained to be equal across orientation discrimination and change detection trials—that is, in model configurations with either ν^t or ν^* in their designations—the effect of an increment in the denominator of the sample size relationship for change detection trials would be completely offset by the estimated drift rate parameter in the numerator, rendering the CD+ and non-CD+ model types identical in their estimates. (In practice, these models can differ slightly in the values of their objective functions due to inefficiencies in the optimization procedure in finding the maximum likelihood parameter set.)

Like the best fitting model configurations from Experiment 1, each of the best fitting models in Experiment 2 across observers requires some relaxation in the constraints on the nondecision time parameter, T . In observers other than DB and SL, and for fitting the group average data, the best fitting models favoured constraining the nondecision time parameter within decision type conditions. This finding lends additional support to the argument made by Sewell, Lilburn, and Smith (in press) that systematic changes in nondecision time during an orientation discrimination task reflected the cost of accessing stored items in memory, changing across but not within trials of different memory

Table 4.7: The top five best fitting diffusion model configurations, in terms of BIC, for the group average data in Experiment 2.

#	Model	Free parameters	G^2	BIC
1	MOD- ν^d - α^t - T^t	12	102.274	207.58
2	MOD- ν^d - α^t - T^t - α^t	13	95.391	209.472
3	MOD-CD+- ν^d - α^t - T^t - α^t	13	96.543	210.624
4	MOD-CD+- ν^d - T^t	11	115.513	212.043
5	MOD-CD+- ν^d - α^t - T^t	12	107.417	212.723

array sizes—as was also found in Experiment 1. The difference in nondecision time between decision types, where change detection response times are slower than orientation discrimination response when comparing equivalent experimental conditions, also supports the finding of both the previous experiment and the convolutional analysis that an extra cost of response time is induced by the encoding and comparison of the probe array.

Beyond the consistency in each of the top ranking models requiring some flexibility in fitting the drift rate and in the nondecision time parameters, the best fitting models also required additional flexibility. This flexibility was obtained by allowing either the boundary separation parameter, α , to be estimated separately within each decision type condition, or by allowing the intrusion rate parameter, α , to be estimated separately across trials of different decision types. Observer KT and the group average data were best fit by the model configuration MOD- ν^d - α^t - T^t , a model configuration with 12 free parameters where the drift rate varied as a function of stimulus duration, and the boundary separation and nondecision time varied as a function of the decision type. The need to use an additional boundary separation parameter in this way is not unexpected: as the different decision types were grouped into blocks, and the tasks themselves require different decisions to be made, the possibility of a response bias between tasks may reflect a strategy on the part of the observer.

Alternatively, observer BF was best fit by the model configuration MOD- ν^d - T^t - α^t , a model configuration with 12 free parameters where drift rate varied as a function of stimulus duration, and the nondecision time and intrusion rate parameter varied as a function of the decision type. Given the estimated value of ν_α is zero (see Table 4.20), these results indicate a substantially higher rate of delayed guessing on the part of these observers in the change detection trials. Observer DB required full drift and nondecision time flexibility in fitting. Observer SS required only drift rate changes across duration and an increment in the nondecision time for change detection. Allowing the combination

of both nondecision time and the intrusion rate parameters to be flexible over different decision type conditions also fit the individual and group average data well. The model configuration $\text{MOD-}\nu^d\text{-}\alpha^t\text{-}\tau^t\text{-}\alpha^t$, a model configuration with 13 free parameters, fit best for observer SL, ranked highly in fitting the data for observers BF, KT, and SS, as well as in fitting the group average. Although not as highly ranked as models with either the nondecision time or the intrusion rate freely estimated across trials of different decision types (and the other constrained to be equal across decision types), these results indicate that, overall, there appears to be significant effects of the type of response strategy and guessing required to fully characterise the results in the current experiment.

Part of the inconsistency in the rank order of model configurations and the requirement to consider response strategies and guessing rates directly may be a result of the difficulty of the experiment, with the group average accuracy ranging from just below 60% of trials correct in the 100 ms condition to just above 70% of trials correct in the 200 ms (see Figure 3.11). The requirement to equate performance in each condition also necessitated an extended calibration procedure, preceding the experimental sessions for each observer. The conjunction of both a difficult experiment, using four-item memory arrays, and an extended calibration procedure, may have led to observers engaging in a response strategy to offset the overall difficulty of the task.

A display of model predictions against group average data is provided in Figures 4.21, 4.22, and 4.23 for the model configurations $\text{MOD-}\nu^d\text{-}\alpha^t\text{-}\tau^t$, $\text{MOD-}\nu^d\text{-}\tau^t\text{-}\alpha^t$, and $\text{MOD-}\nu^d\text{-}\alpha^t\text{-}\tau^t\text{-}\alpha^t$, respectively.




Figure 4.21: A QPP of the group average data against the predictions of the model, $MOD-\nu^d-\alpha^t-T^t$: a model in which—in addition to the drift rate being estimated for each exposure duration condition and the nondecision time estimated for each decision type (change detection or orientation discrimination)—the boundary separation parameter was also separately estimated for each decision type condition. The model predictions are represented by filled squared connected by solid lines; the observed data is shown by unfilled circles connected by broken lines.

4.2 Experiment 2




Figure 4.22: A QPP of the group average data against the predictions of the model, $MOD-\nu^d-T^t-\alpha^t$: a model in which—in addition to the drift rate being estimated for each exposure duration condition and the nondecision time estimated for each decision type (change detection or orientation discrimination)—the intrusion rate parameter, α , was also separately estimated for each decision type condition. The model predictions are represented by filled squares connected by solid lines; the observed data is shown by unfilled circles connected by broken lines.




Figure 4.23: A QPP of the group average data against the predictions of the model, $\text{MOD-}\nu^d-\alpha^t-T^t-\alpha^t$: a model in which—in addition to the drift rate being estimated for each exposure duration condition and the nondecision time estimated for each decision type (change detection or orientation discrimination)—both the boundary separation parameter and the intrusion rate parameter were also separately estimated for each decision type condition. The model predictions are represented by filled squared connected by solid lines; the observed data is shown by unfilled circles connected by broken lines.

4.2.4 Discussion

The aim of Experiment 2 was to examine whether the same relationship between response times for orientation discrimination and change detection trials in Experiment 1 would hold when exposure duration was varied. The performance of observers was equated by adjusting the memory array display contrast between different decision types for each observer individually. A striking pattern of response time differences was obtained, where the mean response times of observers was uniformly slower in the change detection condition when compared to the orientation discrimination condition, despite the near-equal level of accuracy between the two types of tasks. This difference in response times at the mean level was supported by the application of a convolution analysis where, in addition to extra dispersion in the change detection condition, there was a consistent and large translation of the entire response time distribution in the change detection condition later in time when compared to the distribution of response times in the orientation discrimination condition.

In examining this distribution-level shift in response time between decision type conditions with a diffusion model, the best fitting model configurations consistently required the nondecision time parameter of the diffusion model to be estimated separately for each decision type. The effect of the stimulus exposure duration was reflected in the drift rate estimates with the best fitting model configurations for observers BF, KT, and SS, and the group average data requiring separate estimates of drift rate for each exposure duration condition. The other observers, DB and SL, also had free drift parameters across different exposure duration conditions, but required further flexibility in freely estimating the drift rate parameters across different decision type conditions as well.⁵

The diffusion model results for Experiment 2 were, however, somewhat less consistent than those seen in Experiment 1. All observers, in addition to the nondecision time estimated as being longer for change detection trials and the drift rate parameters estimated as being larger with longer memory array exposure durations, required some flexibility in either the boundary separation or intrusion rate parameters⁶. As discussed in the preceding section, these may represent some strategy on the part of observers to ameliorate some of the difficulty of the task: either to attempt to shift in response caution—the

⁵Some of the top ranking models for observers BF, DB, KT, and the group average data, also included models where the drift rate of the intrusion process was also freely estimated across different exposure duration conditions, indicating that the intrusion process might be similarly affected by the overall level of stimulus information available to the memory system. This evidence is, however, relatively weak and would require further elaboration in future studies.

⁶As mentioned, given the v_{α} rate estimated in these model configurations is zero, this difference in the intrusion rate α can be seen as differences in delayed guessing.

behavioural analogue of the diffusion model's boundary separation parameter—or delayed guessing behaviour, perhaps as a result of some failed memory accessing process (see Sewell et al., in press). There is increasing recent evidence that strategic concerns play an integral role in observer behaviour during difficult short-term memory tasks. Bengson and Luck (2015) found that capacity estimates for a suprathreshold change detection task changed with differing instruction sets, emphasising either the need to remember all items within a display or stressing the need to maximise performance. More recently, Donkin, Kary, Tahir, and Taylor (2016) demonstrated that observer performance in change detection tasks is well captured by models that assume a mixture of discrete slot style memory representations and flexible resource style allocation. One systematic difference in observer responding noted by Donkin and colleagues was between the blocked design of an experiment and strategy use. In an experiment where memory array size conditions were blocked together, a much higher proportion of resource-like encoding was noted; when the memory array size changed on a trial-to-trial basis, a pure slot-like strategy was engaged by observers. The ability for observers to predict the number of elements on the next trial was shown to have a large effect on the encoding strategy used by observers. Given the limitations upon visual short-term memory and the difficulty of the current experiment, it is possible that observers attempted to maximise performance between different decision types—particularly as trials of different decision types were within a predictable block structure.

Both changes in boundary separation and intrusion rate between decision types represent a significant, if not substantial, contribution of the task dynamics to the overall behavioural profile but which may not be captured in simple models that do not explicate both how the characteristics of the memory representation influence the observed behaviour and how the nature of the decision, deriving from the circumstances of the task, can be seen in responding. This topic shall be elaborated upon in the next section.

4.3 General discussion

In this chapter, the response time data obtained in the two experiments described in the preceding chapter was examined. For each experiment, a linear mixed-effects analysis was conducted on the mean response times, followed by a convolutional analysis operating on the entire response time distribution for each experimental condition, and concluding with the application of a version of the diffusion model extended using the signal detection theory models of the last chapter as the basis. In both experiments, a

clear response time advantage for orientation discrimination trials were found over the corresponding change detection trials, regardless of the size of the memory array, the exposure duration of the memory array, or whether the contrast of the memory array was adjusted to equate performance. Following from the results of the last chapter, this increase in the time to make change detection responses is likely to be the product of encoding the probe array, and then comparing the probe array to the stored representations within memory.

The diffusion analysis provided additional elaboration on this process. In Experiment 1, it was found that simple extensions of the diffusion model, following the results of the signal detection theory analysis presented in the last chapter, provided a very good overall fit to observer data. In the best fitting model configuration overall, the drift rates were constrained to fit the sample size relation previously discussed by Sewell, Lilburn, and Smith (2014), and change detection performance and orientation discrimination performance was predicted from a single drift rate by incremented the denominator of the sample size relationship in the change detection case by one to reflect the increased memory load due to the processing of the probe array. In addition to the fixed relationship between the drift rate and performance across all experimental conditions, the best fitting model configurations also required a separate nondecision value to be estimated for each experimental condition: the increase in nondecision time as the memory array size increased was also found by Sewell, Lilburn, and Smith (in press), who argued that it represented a limited-capacity retrieval process; the increase in nondecision time across the type of decision following the results of both the linear mixed-effects analysis and the convolutional analysis. Last, like Sewell, Lilburn, and Smith (in press), an additional mixture process was included to account for the asymmetry observed between the correct and incorrect responses at the earliest quantiles (known as the leading edge). This process was given a more general form, where the drift rate of the process could be freely estimated by the model, to indicate the presence of unprobed “intruding” information into the decision process, but was estimated to be zero in all cases.

In Experiment 2, the results of the diffusion analysis were less consistent, partially owing to the nature of the experimental design in which the memory array was presented at different contrasts in the orientation discrimination and change detection conditions. Additional drift rate flexibility was required in all observers, either to allow for the increase of stimulus information into the decision process as the exposure duration of the memory array increased, or to account for both the effect of stimulus exposure duration and the effect of the increased stimulus contrast in the change detection condition. The top model configurations also required additional flexibility either by allowing sep-

arate boundary separation or delayed guessing rates (i.e., the intrusion rate parameter) estimates across the different two decision types, perhaps reflecting a difference in the response strategy employed by different observers.

Simply put, for both experiments, the best fitting model configuration showed that a single diffusion model, with modifications, can predict observer responses between two different types of decision-making using a single set of parameters. For the best fitting model configurations obtained in Experiment 1, orientation discrimination performance could be characterised using a single free drift rate parameter (estimating performance in the single-item memory array condition). The reduction in performance in conditions with larger memory array sizes was predicted by assuming that the drift rates scale following the sample-size relation. Change detection performance could be characterised by use of the same drift rate as orientation discrimination performance, with the addition of an item in the denominator of the sample-size relation. This additional item may be interpreted as the probe array competing for limited memory resources with the memory array representations. In both Experiments 1 and 2, the encoding and comparison of the probe array in the change detection decisions also appeared to add a constant (or near-constant) time to the response times of the corresponding orientation discrimination conditions. Last, in both experiments, the detailed patterns of response time and accuracy are better captured by a model that assumes there is a failure to retrieve the relevant item from memory on a proportion of trials (like Sewell et al., *in press*) or that the decision is influenced by the contents of memory outside the probed location. (The parameterisation of this intrusion process will be examined in further detail in Chapter 6.)

The role of the retrieval of information from memory and the decision process has figured prominently in the interpretation of the results—particularly in examining the diffusion model results. The process of retrieving information from visual short-term memory is often not explicitly modelled⁷, unlike work in working memory more generally (as reviewed in Sewell et al., *in press*), with behavioural results interpreted almost exclusively in terms of differences in the memory representation across different experimental manipulations. The results from this study show that the assumption is too simplistic. The evidence presented in this chapter, and the last, suggests that experimental procedures which require the comparison of stored memory representations to subsequently presented standards—either required to detect a change or, in the case of continuous report, to replicate the stored memory representation—are placing an extra

⁷Although Bays et al. (2009) did include the presence of retrieval errors in their model of a continuous report task.

cost on the limited resources of memory. Further, as shown both here and in the work reported by Sewell, Lilburn, and Smith (in press), there appears to be a consistent cost in the time taken to access representations within memory that grows as a function of the number of items stored in memory.

Both the cost of the probe array in change detection and the increasing item access times as the number of concurrent elements within memory grows reflect the behavioural signature of task and decision constraints. The asymmetry in the leading edge of the response time distributions—present in the data presented in this chapter, as well as in Chapter 6, and in the work conducted by Sewell and colleagues—also requires some further consideration of the role of memory retrieval. One distinction that might be instructive in examining the role of retrieval from memory is the distinction between the overall visual information stored and representations required for indexing that aggregate information. This distinction distinguishes between the processes that maintain visual information within memory and the processes that associate the contents of memory with locations in the visual field, allowing information to be retrieved when probing particular locations.

The delineation between memory content and memory indexing has a basis in the functional neuroimaging literature. Xu and Chun (2006) conducted a change detection study measuring observer performance using simple and complex stimuli while examining patterns of BOLD activation via fMRI. They found three sites of cortical activity related to change detection performance: activation of the inferior intraparietal sulcus (IPS) correlated with the number of items presented in the memory array, but was invariant between stimulus classes of different complexity, and reached a ceiling level of activation with the presentation of four items (consistent with the four item limit conjectured in slot model accounts of memory); activation in the lateral occipital cortex (LOC) and the superior IPS, however, was correlated with both the number of items to be simultaneously maintained and the complexity of those items, indicating that these regions are sensitive to the total informational load of representations maintained in memory. Xu and Chun argued that the role of the inferior IPS, previously found to be associated with the control of spatial attention (Corbetta & Shulman, 2002; N. Kanwisher & Wojciulik, 2000), allows the maintenance of multiple attentional loci in the visual field, corresponding to the location of items stored in memory, and further supported by differences in activation in the inferior IPS observed when the target location was manipulated. The role of the superior IPS and the LOC, then, could be more naturally associated with constraints modelled using the sample size information limit: areas limited by the total amount of visual material to keep active simultaneously.

The inclusion of an indexing process as distinct from processes that store visual information, and the consideration of the process of retrieving items from memory, does not sit easily with simple resource models or, more generally, models that do not clearly specify how the correspondence between stored information and spatial location may be maintained during memory retention intervals and resolved when retrieval is required. As previously argued by Sewell et al. (2014), the “slots + averaging” model of Zhang and Luck (2009) does not readily admit the existence of an indexing process. This model implies a level of intricate bookkeeping as the association between visual items and memory representations is not isomorphic, in that one (physical) stimulus can be represented by multiple representations, which may be redistributed upon presentation of additional task-relevant information. Sewell and colleagues noted that such model predictions relating to multiple copies of the same visual element to increase response precision is incompatible with the finding of Vogel and Machizawa (2004), where contralateral delay activation (CDA), an electrophysiological correlates of visual short-term memory storage, correlates the number of total stimuli stored, without any indication of multiple representations for the same item. More recent ERP evidence from Gao et al. (2011) indicates that the magnitude of CDA may track the number of distinct stimulus identities in memory, rather than the total number of locations in memory, however there is no indication that subcapacity representations are duplicated: the magnitude of CDA for retention of a single item (in a single location) appears to be the same as the magnitude of CDA for multiple items of the same identity.

The most parsimonious account of an index would be one that operates completely spatially, where stored representations are retrieved on the basis of their location in the visual field. Such an index would be consistent with the role of the inferior IPS argued for by Xu and Chun, and would reflect the spatiotopic structure of visual cortex. The existence of a spatially mediated retrocuing effect (Griffin & Nobre, 2003) suggests that, at least in part, memory representations within the visual short-term memory system are intimately tied to some representation of the original location of the memory element. The study by Jiang et al. (2000), where change detection performance was preserved under a limited number of transformations to the spatial relationship between the memory array and the probe array further indicates the intrinsic spatial characteristics of visual short-term memory representations. This account of a spatially mediated index would be consistent with an interpretation of the “intrusion” process as reflecting an inefficiency in the post-stimulus probe to confine memory retrieval purely to the specified location, allowing information indexed by other locations in the visual display to contribute to some responses.

A completely spatially mediated index cannot, however, account for above chance performance in the case of stimuli presented sequentially in the same location. The activation of the superior IPS and LOC areas was differentiated by Xu and Chun when examining trials where multiple stimuli were presented in the same location sequentially: the superior IPS decreased in activity and the LOC increased in activity in trials where information was presented sequentially in the same location when compared to trials where information was presented simultaneously in distinct locations. An extension of the current study could examine the characteristics of intruding information in instances where the distractor information might either closely surround the probed target—testing the probe efficiency in selecting the correct information—or, via sequential presentation, is presented in the same location as the probed target. In such an experimental design, the current modelling results—where both nondecision time and the intrusion process capture retrieval and decision variability—would provide a means of examining the effect of stimulus configuration directly, allowing greater leverage upon questions of fundamental memory architecture.

As such, the utility of the current findings, beyond the results regarding the relationship between near-threshold change detection and orientation discrimination observer performance, extends to future studies where the models and analyses may be adapted to further define the boundary between the observable effects of constraints on memory representations, and the observable characteristics of the decisions based on those representations. The next part of this thesis, starting with the next chapter, will examine the nature of the representations within the memory system by using a variant of the near-threshold two-alternative orientation discrimination paradigm used in this chapter, using the models developed in this section as the basis of the analysis.

Supplementary tables and figures

Best fitting diffusion model fits for individual observers in Experiment 1. Tables 4.8–4.12

Best fitting parameters for the model configuration MOD–CD+– T^* in Experiment 1. Table 4.13

Best fitting parameters for the model configuration MOD– T^* in Experiment 1. Table 4.14

Best fitting diffusion model fits for individual observers in Experiment 2. Tables 4.15–4.19

Best fitting parameters for the model configuration MOD– v^d – T^t – α^t in Experiment 2. Table 4.20

Best fitting parameters for the model configuration MOD– v^d – α^t – T^t in Experiment 2. Table 4.21

Best fitting parameters for the model configuration MOD– v^d – α^t – T^t – α^t in Experiment 2. Table 4.22

Table 4.8: The top five best fitting diffusion model configurations, in terms of BIC, for the data of observer BF in Experiment 1.

#	Model	Free parameters	G^2	BIC
1	MOD-CD+-T*	15	129.47	242.871
2	MOD- v^t -T*	16	126.592	247.554
3	MOD-T*	15	134.479	247.88
4	MOD-CD+-T*- α^t	16	128.703	249.665
5	MOD-CD+- a^t -T*	16	129.895	250.857

Table 4.9: The top five best fitting diffusion model configurations, in terms of BIC, for the data of observer DB in Experiment 1.

#	Model	Free parameters	G^2	BIC
1	MOD-CD+-T*	15	103.853	217.239
2	MOD-CD+- a^t -T*	16	99.949	220.894
3	MOD-CD+-T*- α^t	16	102.407	223.351
4	MOD-CD+- v^t -T*	16	102.828	223.773
5	MOD-CD+- v^t - a^t -T*	17	95.76	224.264

Table 4.10: The top five best fitting diffusion model configurations, in terms of BIC, for the data of observer KT in Experiment 1.

#	Model	Free parameters	G^2	BIC
1	MOD-CD+- v^t -T*- α^t	17	186.025	314.547
2	MOD-CD+- v^t - a^t -T*	17	190.458	318.98
3	MOD-CD+- v^t -T*- α^t - T_α^t	18	183.453	319.535
4	MOD-CD+- v^t - a^t -T*- α^t - T_α^t	19	176.046	319.687
5	MOD-CD+- v^t - a^t -T*- α^t	18	187.845	323.927

Table 4.11: The top five best fitting diffusion model configurations, in terms of BIC, for the data of observer SL in Experiment 1.

#	Model	Free parameters	G^2	BIC
1	MOD-CD+-T*	15	175.274	291.299
2	MOD-CD+-v ^t -T*	16	173.983	297.743
3	MOD-CD+-T*-α ^t	16	175.22	298.98
4	MOD-CD+-a ^t -T*	16	175.5	299.26
5	MOD-T*	15	185.554	301.579

Table 4.12: The top five best fitting diffusion model configurations, in terms of BIC, for the data of observer SS in Experiment 1.

#	Model	Free parameters	G^2	BIC
1	MOD-a ^t -T*	16	258.756	379.684
2	MOD-a ^t -T*-α ^t	17	258.376	386.862
3	MOD-CD+-v ^t -a ^t -T*	17	262.05	390.535
4	MOD-a ^t -T*-α ^t -T _α ^t	18	258.362	394.406
5	MOD-CD+-v ^t -a ^t -T*-α ^t	18	261.983	398.027

Table 4.13: The best fitting parameter estimates for the model configuration MOD-CD+-T* across all participants in Experiment 1. When multiple independent values are estimated for a single parameter, the values are presented in columns for different decision type conditions (change detection on the left, orientation discrimination on the right) and in rows for increasing memory array sizes.

Model	Observer	ν	s_ν	α	T		s_T	α	T_α	ν_α
MOD-CD+-T*	Avg	0.232	0.173	0.121	0.256	0.088	0.195	0.652	0.263	0
					0.293	0.156				
					0.337	0.212				
					0.369	0.256				
BF	0.141	0	0.115	0.421	0.365	0.238	0.156	0.625	0.216	-0.004
					0.388	0.219				
					0.421	0.303				
					0.467	0.318				
DB	0.253	0.161	0.146	0.376	0.294	0.098	0.15	0.574	0.262	0
					0.322	0.175				
					0.376	0.194				
					0.369	0.329				
KT	0.192	0.004	0.108	0.359	0.25	0.123	0.117	0.313	0.217	-0.152
					0.322	0.203				
					0.359	0.261				
					0.399	0.282				
SL	0.319	0.23	0.118	0.335	0.241	0.045	0.127	0.541	0.151	0
					0.284	0.151				
					0.335	0.187				
					0.344	0.186				
SS	0.245	0.189	0.183	0.265	0.13	0	0.145	0.311	0.495	-0.282
					0.211	0.082				
					0.265	0.17				
					0.384	0.198				

Table 4.14: The best fitting parameter estimates for the model configuration MOD-T* across all participants in Experiment 1. When multiple independent values are estimated for a single parameter, the values are presented in columns for different decision type conditions (change detection on the left, orientation discrimination on the right) and in rows for increasing memory array sizes.

Model	Observer	ν	s_ν	α	T	s_T	T_α	ν_α
MOD-T*	Avg	0.192	0.16	0.122	0.283	0.119		
					0.321	0.182		
					0.364	0.243	0.124	0.614
					0.422	0.308	0.27	0
BF	BF	0.121	0	0.113	0.344	0.206		
					0.379	0.215		
					0.409	0.296	0.245	0.462
					0.456	0.31	0.242	0
DB	DB	0.226	0.16	0.147	0.298	0.093		
					0.318	0.171		
					0.375	0.19	0.14	0.567
					0.367	0.325	0.277	-0.001
KT	KT	0.18	0.001	0.104	0.265	0.134		
					0.309	0.205		
					0.36	0.254	0.15	0.567
					0.4	0.278	0.151	-0.032
SL	SL	0.277	0.226	0.12	0.244	0.038		
					0.289	0.15		
					0.335	0.195	0.142	0.478
					0.345	0.185	0.142	-0.002
SS	SS	0.222	0.172	0.126	0.212	0.054		
					0.267	0.128		
					0.319	0.222	0.12	0.817
					0.427	0.303	0.452	0

Table 4.15: The top five best fitting diffusion model configurations, in terms of BIC, for the data of observer BF in Experiment 2.

#	Model	Free parameters	G^2	BIC
1	MOD- γ^d - T^t - α^t	12	143.743	249.053
2	MOD-CD+- γ^d - T^t - α^t	12	151.501	256.811
3	MOD- γ^d - T^t - α^t - γ_α^d	14	143.897	266.76
4	MOD- γ^* - T^* - α^t - T_α^t	20	92.753	268.271
5	MOD-CD+- γ^d - T^t - α^t - γ_α^d	14	150.509	273.371

Table 4.16: The top five best fitting diffusion model configurations, in terms of BIC, for the data of observer DB in Experiment 2.

#	Model	Free parameters	G^2	BIC
1	MOD- γ^* - T^*	18	261.867	419.829
2	MOD-CD+- γ^* - T^*	18	264.105	422.068
3	MOD- γ^d - T^t - α^t - T_α^t - γ_α^d	15	290.544	422.18
4	MOD-CD+- γ^d - T^t - α^t - T_α^t - γ_α^d	15	294.263	425.899
5	MOD- γ^d - T^t	11	332.194	428.727

Table 4.17: The top five best fitting diffusion model configurations, in terms of BIC, for the data of observer KT in Experiment 2.

#	Model	Free parameters	G^2	BIC
1	MOD- γ^d - α^t - T^t	12	111.938	217.25
2	MOD-CD+- γ^d - α^t - T^t	12	113.581	218.893
3	MOD- γ^d - α^t - T^t - α^t	13	109.048	223.136
4	MOD-CD+- γ^d - α^t - T^t - α^t	13	109.584	223.672
5	MOD- γ^d - α^t - T^t - α^t - γ_α^d	15	108.755	240.396

Table 4.18: The top five best fitting diffusion model configurations, in terms of BIC, for the data of observer SL in Experiment 2.

#	Model	Free parameters	G^2	BIC
1	MOD- $\gamma^* - a^t - T^* - \alpha^t - T_\alpha^t$	21	190.519	374.76
2	MOD-CD+- $\gamma^* - a^t - T^* - \alpha^t$	20	206.52	381.988
3	MOD- $\gamma^* - a^t - T^* - \alpha^t - T_\alpha^t - \gamma_\alpha^t$	22	189.336	382.351
4	MOD-CD+- $\gamma^* - a^t - T^* - \alpha^t - T_\alpha^t$	21	200.354	384.595
5	MOD-CD+- $\gamma^* - a^t - T^* - \alpha^t - T_\alpha^t - \gamma_\alpha^t$	22	200.367	393.381

Table 4.19: The top five best fitting diffusion model configurations, in terms of BIC, for the data of observer SS in Experiment 2.

#	Model	Free parameters	G^2	BIC
1	MOD- $\gamma^d - T^t$	11	320.368	416.909
2	MOD-CD+- $\gamma^d - T^t - \alpha^t$	12	312.372	417.69
3	MOD- $\gamma^d - T^t - \alpha^t$	12	313.264	418.582
4	MOD- $\gamma^* - T^*$	18	262.698	420.674
5	MOD-CD+- $\gamma^d - T^t$	11	325.998	422.539

Table 4.20: The best fitting parameter estimates for the model configuration MOD- v^d - T^t - α^t across all participants in Experiment 2. When multiple independent values are estimated for a single parameter, the values are presented in columns for different decision type conditions (change detection on the left, orientation discrimination on the right) and in rows for increasing stimulus exposure duration conditions.

Model	Observer	v	s_v	α	T	s_T	α	T_α	v_α
MOD- v^d - T^t - α^t	Avg	0.135							
		0.255	0	0.136	0.279	0.122	0.213	0.388	0.443
		0.315						0.243	-0.16
BF		0.153							
		0.276	0.133	0.106	0.339	0.133	0.151	0.36	0.763
		0.35						0.181	-0.002
DB		0.174							
		0.35	0	0.119	0.318	0.174	0.197	0.697	0.641
		0.35						0.191	-0.018
KT		0.081							
		0.18	0	0.127	0.263	0.145	0.23	0.932	0.99
		0.244						0.182	-0.015
SL		0.112							
		0.277	0	0.121	0.305	0.108	0.147	0.23	0.348
		0.35						0.216	-0.234
SS		0.124							
		0.301	0	0.127	0.305	0.155	0.146	0.796	0.923
		0.35						0.44	-0.001

Table 4.21: The best fitting parameter estimates for the model configuration $\text{MOD-}\gamma^d\text{-}\alpha^t\text{-}\tau^t$ across all participants in Experiment 2. When multiple independent values are estimated for a single parameter, the values are presented in columns for different decision type conditions (change detection on the left, orientation discrimination on the right) and in rows for increasing stimulus exposure duration conditions.

Model	Observer	γ	s_γ	α	τ	s_τ	α	τ_α	γ_α
$\text{MOD-}\gamma^d\text{-}\alpha^t\text{-}\tau^t$	Avg	0.125							
		0.292	0	0.121	0.114	0.293	0.157	0.153	0.237
		0.35						0.769	-0.013
BF		0.152							
		0.242	0.001	0.109	0.121	0.324	0.116	0.148	0.184
		0.323						0.363	-0.153
DB		0.173							
		0.35	0	0.119	0.12	0.322	0.173	0.186	0.184
		0.35						0.658	-0.02
KT		0.071							
		0.166	0	0.139	0.115	0.223	0.181	0.217	0.179
		0.234						0.969	-0.01
SL		0.101							
		0.321	0.001	0.114	0.101	0.294	0.133	0.15	0.168
		0.35						0.667	-0.008
SS		0.124							
		0.302	0	0.125	0.129	0.304	0.155	0.149	0.442
		0.35						0.859	-0.001

Table 4.22: The best fitting parameter estimates for the model configuration $\text{MOD-}\gamma^d-\alpha^t-\text{T}^t-\alpha^t$ across all participants in Experiment 2. When multiple independent values are estimated for a single parameter, the values are presented in columns for different decision type conditions (change detection on the left, orientation discrimination on the right) and in rows for increasing stimulus exposure duration conditions.

Model	Observer	γ	S_γ	α	T	S_T	α	T_α	γ_α
$\text{MOD-}\gamma^d-\alpha^t-\text{T}^t-\alpha^t$	Avg	0.123	0	0.121	0.113	0.298	0.155	0.151	0.704
		0.29		0.121	0.113	0.298	0.155	0.151	0.821
		0.35		0.121	0.113	0.298	0.155	0.151	0.821
BF		0.153	0.001	0.108	0.121	0.329	0.113	0.144	0.293
		0.243		0.108	0.121	0.329	0.113	0.144	0.44
		0.327		0.108	0.121	0.329	0.113	0.144	0.183
DB		0.173	0	0.118	0.12	0.321	0.173	0.199	0.694
		0.35		0.118	0.12	0.321	0.173	0.199	0.631
		0.35		0.118	0.12	0.321	0.173	0.199	0.21
KT		0.083	0	0.139	0.116	0.227	0.178	0.239	0.92
		0.18		0.139	0.116	0.227	0.178	0.239	0.99
		0.237		0.139	0.116	0.227	0.178	0.239	0.18
SL		0.098	0.002	0.113	0.098	0.303	0.131	0.146	0.537
		0.316		0.113	0.098	0.303	0.131	0.146	0.829
		0.35		0.113	0.098	0.303	0.131	0.146	0.184
SS		0.124	0	0.127	0.127	0.306	0.155	0.15	0.798
		0.302		0.127	0.127	0.306	0.155	0.15	0.922
		0.35		0.127	0.127	0.306	0.155	0.15	0.451

Chapter 5

Fine orientation discrimination: sensitivity

The first part of this thesis investigated the role of the experimental procedure on the pattern of observer performance seen. It was found that the sample-size relationship, plus an increment in the denominator of this relationship for change detection trials, provided a parsimonious account of both observer accuracy and response time across different memory paradigms. The second part of the thesis will extend this result to examine the relationship between constraints on feature information and the sample-size relationship. In this part, I will use a fine orientation discrimination paradigm, generalising the orthogonal orientation discrimination paradigm employed in the last section, to demonstrate that a simple weighting of the sensitivity function can account for the effect of stimulus discriminability on performance. In particular, the overall sample-size information limit on memory appears to have a separate effect on observer performance from the effect of stimulus discriminability due to feature confusability.

One of the key findings of the experiments reported in the first part of this thesis was that observer accuracy decreased uniformly in trials requiring a change detection decision compared to those requiring an orientation discrimination decision. In the analysis of the data obtained in Experiment 1, a set of differencing models provided the best overall account of the pattern of observer performance, with some of the better fitting models indicating that performance was affected by encoding of the probe item competing for memory resources during the change detection decision.

This decrease in performance indicates that a comparison against an internal stimulus standard, rather than an external sensory reference, increases the overall sensitivity of the testing procedure: higher observer performance is obtained with less stimulus infor-

⁰Portions of this chapter were presented at the 2015 Australasian Experimental Psychology Conference in Sydney.

mation (either a shorter stimulus exposure duration or, as in the case of Experiment 2, a lower stimulus display contrast). These findings do not invalidate previous claims based on change detection (or continuous report) data about the presence or severity of the capacity limitations in visual short-term memory, but may mean that estimates of the storage capacity limit may be biased due to the need to account for the storage requirements of the probe array. The effect of the probe array may be even more pronounced in the case of experiments employing near-threshold stimuli, where an additional level of control can be placed on observer performance. Although not often used in visual short-term memory experiments, near-threshold stimuli allow a finer control over the rate of available stimulus information to be encoded into memory. Sewell, Lilburn, and Smith (2014) used this additional control in an orientation discrimination experiment to describe a fixed limit on the rate at which information can be maintained and consolidated in memory.

Sewell and colleagues argued that the sample-size relationship existed at a level below that of the more regularly observed discrete item limit, described by Luck and Vogel (1997). The quantitative visual short-term memory model of Smith and Sewell (2013), for instance, implements both a fixed overall memory limit on discrete items and a shared information limit between those items, based on the divisive normalisation of neural populations that code features (Smith, 2015; Smith et al., 2015). The idea of multiple levels of constraints on visual memory has been canvassed before in the literature, with some evidence indicating that feature information is shared between object representations in hierarchical feature “bundles” (Brady & Alvarez, 2011; Brady, Konkle, & Alvarez, 2009, 2011) or clusters (R. Jacobs & Orhan, 2013; Orhan & Jacobs, 2013) which allow observers to more efficiently store information between separate items in a memory array. This decoupling between discrete items and their constitutive features once again raises the importance of considering not only the overall form of memory arrays, in terms of the number of items to be remembered, but also the content of those items.

In contemporary visual short-term memory research, the role of features within memory representations has been predominantly explored through the use of the continuous report paradigm (see, §2.3.2). As indicated in previous chapters, these tasks are designed to provide a clear relationship between the continuous (and, usually, circular) structure of the feature being probed—usually the hue of a coloured patch or the orientation of an item on a wheel—and the continuous (and, usually, circular) means of observer response. As such, the dispersion of responses in these tasks is often interpreted as directly reflecting the structure of a representation held in memory. As shown from the first part of this thesis, this relationship might be, at best, somewhat compromised by the need to

compare an internal memory representation to an externally presented stimulus referent and, at worst, neglect entirely the decision procedure for mapping an internal representation to a response option. The alternative is to use a task which requires comparison to an internal stimulus standard—like the orientation discrimination task used by Sewell and colleagues—where any variation in the responses may reflect more accurately constraints on the internal representation, yet these tasks have rarely been employed to examine the constraints on feature information.

Sewell and colleagues used only two stimulus alternatives—horizontally and vertically oriented Gabor patches—in examining memory performance for near-threshold stimuli. These orientations were selected due to the result by Thomas and Gille (1979) which indicates that the performance characteristics of Gabor patch orientation discrimination for differences in response alternatives greater than 20° are identical to performance characteristics in simple visual detection. For differences in orientation between response options that are under approximately 20° , Thomas and Gille argued, both stimulus alternatives fall within the orientation selectivity bandwidth of a single psychophysical detector leading to decreased performance in orientation discrimination as the activity elicited by a single orientation can no longer be attributed to a single response alternative. This observed behavioural difference between orthogonal orientation discrimination and fine orientation may arise even in the early cortical representation of orientation information. Discriminability between smaller angular differences may rely on the output of many orientation-selective detectors with different preferred orientations (Ben-Yishai, Bar-Or, & Sompolinsky, 1995), rather than a single psychophysical channel. The temporal properties of sensitivity to the higher spatial frequencies required to differentiate between small angular differences also differ from the sensitivity over time to low spatial frequency information: the psychophysical channels that are active for transient visual stimulation demonstrate a select passband for larger spatial differences when compared to the enhanced sensitivity for higher spatial frequencies shown by psychophysical channels which respond to sustained stimulation (Enroth-Cugell & Robson, 1966; Kulikowski & Tolhurst, 1973; Smith, 1995). In the context of constructing a visual short-term memory representation, this means that the information available for fine orientation discrimination may not be available as rapidly as seen in experiments using large stimulus differences, and may not show the linear growth in squared sensitivity with stimulus exposure duration seen in previous experiments (Sewell, Lilburn, & Smith, 2014).

Beyond the time course of information required to make a fine orientation discrimination decision, the previous use of only orthogonal orientation values also means that

the relationship in terms of storage constraints between the sample-size relation and feature information—in this case, orientation information—is somewhat unclear. Continuous report work would indicate that an increase in the number of items to be stored would lead to a decrease in the overall fidelity of feature information, making discrimination more difficult (Bays et al., 2009; Zhang & Luck, 2008). Bays, Gorgoraptis, Wee, Marshall, and Husain (2011) further argued that increasing the amount of stimulus information, through longer stimulus exposure durations, would also increase the precision of the memory representation. Critically, they showed an interaction between these two factors: both decreases in the memory array size and increases in the exposure duration of the memory array led to less response variability in reproducing the colour of a cued memory item. They argued that, as more information is available to store, the precision of the memory representation itself increases—with the asymptotic limit on precision being given by the number of memory representations maintained simultaneously, reflecting an overall resource limit. The effect of stimulus exposure duration on the precision of memory reproduction would suggest, in line with the sample-size limit, that the rate at which information is encoded into the memory system is limited (or fixed), and that the precision of the feature information should be correlated with the sample size limit. Further specification of the relationship between these data and the properties of the memory system is difficult without an explicit model for a decision stage, and without experimental constraints imposed on the amount of stimulus information available through the use of near-threshold stimuli.

As indicated, the current chapter will explore the interaction between the sample-size information limit and stored information of a visual feature by using a two-choice fine orientation discrimination paradigm. This paradigm aims to generalise the orthogonal orientation discrimination task introduced in the previous section by requiring observers to report on the difference in orientation of a presented and probed stimulus with reference to a known stimulus standard: in Experiment 3, observers reported whether the probed Gabor patch is offset clockwise or counterclockwise from the diagonal position; in Experiment 4, observers reported whether the probed Gabor patch is offset clockwise or counterclockwise from a vertical orientation. This chapter will consider the accuracy data and the sensitivity parameters derived from those data using a model from signal detection theory; the following chapter will continue this analysis using response time modelling. After a description of the method and overall results for the first experiments, a general description of the modelling framework based on signal detection theory will be given.

5.1 Experiment 3

The first experiment presented in this chapter, Experiment 3, is a two-choice fine orientation discrimination experiment. In this experiment, observers were to encode a memory array and report whether a subsequently probed item was rotated either in a clockwise or counterclockwise direction from a diagonal position. The diagonal referent was chosen to avoid observers using the rectangular outline of the monitor to facilitate their discrimination judgement. The discriminability of the target item was determined by the amount of rotation from the diagonal position, or the *angular offset*. The exact range of angular offset was consistent across observers and selected to give a range of performance from near-chance discrimination accuracy at the smallest angular offset, and near-ceiling discrimination accuracy at the largest angular offset (for a memory array of a single item). In this experiment, the size of the memory array was also manipulated—along with the angular offset of the target item—to examine the interaction between the sample-size constraint and orientation feature values.

5.1.1 Method

Participants

Five observers participated in this study: myself (SL) and four paid observers naïve to the aims of the study (AB, CVH, SIS, and XL) drawn from the University of Melbourne. Each observer, with the exception of myself, was briefed about the general nature of the study (with the specific predictions regarding the outcome omitted from this briefing, however) and signed a consent form prior to participation. Each observer other than myself was remunerated AUD \$12 for each session completed.

Each observer completed a variable number of practice and calibration sessions. Practice sessions were undertaken to familiarise the observers with the task and responding in a timely manner; calibration sessions were undertaken to control for individual differences in performance over stimulus contrast settings. All sessions lasted approximately thirty-five minutes each, with regular periods in between blocks of trials for breaks.

Stimuli and apparatus

Following the experiments presented in the first part, the memory array was composed of oriented Gabor patch stimuli. Unlike the experiments presented in the previous section, however, these Gabor patches were presented at angles around the diagonal ($\pi/4$




Figure 5.1: A schematic overview of the timing and structure of each trial in Experiment 3.

rad or $3\pi/4$ rad) axis. Eight different levels of angular offset of the target and distractor stimuli, the difference in rotation from the diagonal stimulus, were used. The minimum angular offset of 0.04 radians from the diagonal position and a maximum angular offset of 0.39 radians from the diagonal position—a position of “zero discriminability”, indicating that discrimination for stimuli presented with that orientation would necessarily be at the level of chance. The range of angular offset positions was selected after a pilot study to ensure almost chance levels of performance at the smallest offset from the zero discriminability point and near-ceiling performance for the largest offset when presenting a single item in the memory array. An equal number of each of the angular offset levels were used in each experiment; the angular offset for each of the distractor items was selected with uniform probability from the eight levels.

Stimulus contrast was also manipulated between observers to approximate a consistent level of performance at the maximum levels of angular offset between observers.

Procedure

A 4×8 within-subjects design was used, composed of four memory array sizes (1, 2, 3, and 4 items) and eight levels of target angular offset from a diagonal position. Each session of

the experiment consisted of 384 trials, yielding a total of 3,840 trials per observer after the full ten experimental sessions. Trial presentation order of differing memory array sizes and differing target angular offsets was randomised.

Each trial began with a 1,000 ms uniform field, followed by the presentation of a central fixation cross for 1,500 ms. A stimulus array of between one and four Gabor patches at the four diagonal positions around the central fixation cross, oriented with offsets from either a northwest–southeast ($3\pi/4$ rad) or northeast–southwest ($\pi/4$ rad) diagonal position, was presented for 200 ms. The presentation of the memory array was followed by the presentation of a high-contrast radial (bullseye) mask for 200 ms. The backwards mask was then followed by the presentation of the report cue at one location corresponding to an item in the memory array.

Observers were asked to judge whether the probed memory item was rotated clockwise or counterclockwise from a diagonal position. Responses were entered via button press and response time were recorded. Observers were instructed to enter their response as quickly as possible without compromising accuracy.

5.1.2 Results

For the analysis of sensitivity and accuracy, no data was filtered from the overall dataset. Stimulus orientation will be described in terms of “angular offset”: the difference between the orientation of the stimulus and a stimulus of zero discriminability (i.e., a diagonal stimulus with orientation $\pi/4$ rad or $3\pi/4$ rad).

Figure 5.2 shows the average observer accuracy for different memory array sizes as a function of the target angular offset; Figure 5.3 shows these data on an observer-by-observer basis.

A mixed-effects logistic regression was used to provide a preliminary analysis of the data. The logistic regression was conducted on the proportion of correct responses, with the memory array size and angular offset treated as fixed effects and the individual differences of the observer on the intercept treated as a random effect. No significant main effect of memory array size was found, $\beta = -0.165$, $SE = 0.038$, $p = 0.305$. A significant main effect of angular offset was, however, found, $\beta = 6.524$, $SE = 0.464$, $p < 0.001$, with larger angular offsets from the zero-discriminability (diagonal) position leading to higher levels of accuracy. A significant interaction between memory array size and angular offset was also found, $\beta = -0.882$, $SE = 0.164$, $p < 0.001$. This interaction indicates that the decrement in accuracy seen as the memory array size increased was seen more prominently in higher angular offsets.




Figure 5.2: Group average accuracy data from Experiment 3 across different memory array sizes as a function of different angular offset levels. Error bars represent one standard error of the mean. SS = Memory array (set) size.

5.1 Experiment 3




Figure 5.3: Accuracy data as a function of the target angular offset level taken from Experiment 3. The data presented is displayed for each observer, with the accuracy shown across different memory array sizes.

Signal detection theory modelling for fine orientation discrimination

In modelling observer accuracy for the first experiment, the signal detection theory framework previously presented in §3.1 was expanded to deal with tasks requiring the observer to discriminate between small angular differences. The simplest assumption is that sensitivity will decrease monotonically with smaller angular differences from the zero-discriminability orientation and that this effect will be independent of the sample-size relation. Two related model types were tested, one with a weighting function—called a “tuning function” for reasons given below—multiplied with squared sensitivity and one where the weighting function was multiplied directly with sensitivity. In the former case, the relationship between sensitivity in a specific angular offset and memory array size condition was defined as

$$(d'_{i,m})^2 = \frac{(1 - f(b(\theta_i - a))) \cdot (d'_1)^2}{m},$$

where i and m are the angular offset index and memory array size, respectively; d'_1 is the (estimated) sensitivity for an orthogonal discrimination from a single item memory array; $f(\cdot)$ is the tuning function, either a modified Gaussian, triangle, or cosine function (detailed below); and a and b are the estimated offset and bandwidth parameters. In the case where the tuning function was applied directly to d' , rather than $(d')^2$, this relationship was

$$d'_{i,m} = \frac{(1 - f(b \cdot (\theta_i - a))) \cdot d'_1}{\sqrt{m}}.$$

These models generalise the orthogonal discrimination models presented in the first part of this thesis. The model for orthogonal orientation discrimination assumed that the orientation values of the stimuli presented corresponded to a single psychophysical detector. As outlined in that section, these detectors are assumed to operate as orientation-based selective filters: maximally responding when an oriented stimulus with their “preferred orientation” value is displayed (Graham, 1989; Hubel & Wiesel, 1963; Thomas & Gille, 1979). In the case of orthogonal orientations, it was assumed that the detectors would not respond to presentation of their non-preferred stimulus, meaning that the response of all detectors would be uncorrelated, and all false alarms would be due to background neural noise.

In modelling discriminations between small orientation differences, the relationship between stimulus orientation and the response of the detectors needs to be expanded. Rather than treating each orientation value as being associated with a separate detector,




Figure 5.4: A schematic diagram of the connection between an oriented external stimulus and the underlying psychophysical channels. A Gabor patch stimulus, represented using the $\textcircled{\diagup}$ symbol, is presented to the observer. The inner product, denoted using the $\langle \cdot, \cdot \rangle$ operators, is taken between the presented stimuli and the receptive fields, represented using the $\textcircled{\diagup}$ to denote a centre-surround receptive field, to obtain different levels of activation based on the level of correspondence between the orientation of the stimulus and that of the filter.

it is assumed that a range of orientations elicit a response from a given detector. The level of this response is assumed to be determined by the distance of the presented orientation from the preferred orientation of the receptor (see Figure 5.4). The function which relates the angular difference of the stimulus to the response of the detector is known as the detector's sensitivity (or tuning) function.

In the current models, the effect of the orientation-selectivity of the zero-discrimination detectors or detectors corresponding to the incorrect response category may either be on sensitivity directly (corresponding to a change in the sensitivity measure d') or on squared sensitivity, or $(d')^2$. The possible relationship between the orientation-selective tuning (or sensitivity) function and $(d')^2$ directly arises from predictions of observer behaviour based on neural models of memory. Smith (2015) has previously shown, extending work previously given by Smith (2010), that populations of Poisson neuron models predict a sample-size information constraint arising from the divisive inhibition of the firing rates.

An equal distribution in the rate of recruiting and maintaining neural resources between represented items, giving rise to the general sample size case, can be manipulated in tasks where the representation or attentional weighting of stimuli is not equal




Figure 5.5: The three different forms of tuning function examined: a Gaussian function, a cosine function, and a triangular function. Each function has a maximum at the preferred orientation of the detector (θ_P) and monotonically decreasing response for other orientations. The arrow marker $\uparrow$ shows the angular offset of zero detector response. (Note: this is not true in the case of the Gaussian function, where the detector response approaches zero asymptotically with increasing angular offset.)

(Smith, Lilburn, Corbett, Sewell, & Kyllingsbæk, under review). The proposed mechanism of attentional allocation biases the pool of neurons representing information held within memory to preferentially code the attentionally selected item over all other items. A similar idea, in terms of the broader dynamics of the system, is described by Smith et al. (2009) and Smith and Sewell (2013). The effect of this attentional selection is a weighting directly on squared sensitivity, $(d')^2$.

A similar argument can be given for stimuli based on the orientation selectivity of the neural populations: although the population remains of a fixed size, and information from a presented display is sampled with a fixed rate, the rate of neural firing may be changed due to the alignment of the orientation profile of a decision-relevant receptor—represented as a pool of Poisson neurons—and the orientation of the presented stimulus. If λ_M represents the maximum firing rate of those neurons, then the mean and variance of the Poisson process will be equal to λ_M and stimulus discriminability, the signal-to-noise ratio, will be equal to $\lambda_M / \sqrt{\lambda_M} = \sqrt{\lambda_M}$ (Smith, 2015). Given a tuning function, $f(\theta)$, which maps the stimulus orientation to a weight, $f : (-\pi, \pi) \rightarrow [0, 1]$, then $d' = \sqrt{f(\theta) \lambda_M}$. The tuning functions are modified in two ways to limit the range of firing rates: first, the output of the function is fixed to unity if the angle θ is negative; second, the output of the function is fixed so that the output remains non-negative (i.e., $\max(0, f(\theta))$). A bandwidth and offset parameter scales and translates the stimulus orientation, respectively, within the tuning function. The bandwidth parameter modifies the range of orientations which result in a non-zero firing rate. The offset parameter changes the distance between the orientation of the detector at the angle of the internal standard and the orientation of the presented stimulus: that is, although the tuning function is assumed to be centred on the orientation of the referent, in practice the psychological representation of the referent may not be veridical, necessitating an offset




Figure 5.6: Assuming that orientation information is coded using a densely packed array of orientation selective detectors, any orientation close to the point of zero-discriminability (θ_0) will elicit a response from both target and non-target detectors. The proportion of the response accounted for by target detectors (the white area) determines the effective stimulus information available for encoding to the memory system. The response of non-target detectors (the grey area) contributes to incorrect decisions near the zero-discriminability orientation.

parameter. In practice, the exact form of the tuning function rarely matters greatly, so long as it is a function which predicts a monotonic decrease in the output of a detector with increasing angular difference. The functional forms examined presently, used in the wider literature (see, e.g., Seung & Sompolinsky, 1993; Thomas & Gille, 1979), are a Gaussian function with non-negative domain, a cosine function with a domain between zero and $\pi/2$, and a triangular function which linearly interpolates maximum responding with zero angular difference to zero detector response at a larger angular difference (see Figure 5.5).

In the current models, I assume that orientation-specific detectors are densely packed across the entire stimulus domain. Two sets of detectors are relevant in the discrimination decision: one set of detectors with preferred orientations clockwise to the standard and one set of detectors with preferred orientations that are counter-clockwise to the standard. When decisions are made regarding targets that have a large angular difference to the point of zero discriminability, observer performance will be largely dependent on the sensitivity of the correct set of orientation detectors to respond to a near-threshold stimulus; that is, incorrect responses will be principally due to noise and accuracy will approach the asymptotic level of performance seen when making orthogonal orientation discriminations given the stimulus contrast. When target stimuli are close in orientation to the standard orientation, the perceptual representation underlying the decision is assumed to involve both the set of detectors that correspond to the correct rotation away from the standard, and some responding by the incorrect set. Thus, the tuning functions of the detectors corresponding to the incorrect response category must be

taken into account (see Figure 5.6). For simplicity, in the models presented, false alarms due to activation of the set of detectors corresponding to the incorrect response category will be modelled with a single tuning function, centred on the zero discriminability orientation. This is entered into the model through a weighting of the sensitivity, $(1 - f(\cdot))$, where f is the tuning function, indicating that as the output of the tuning function increases, overall sensitivity decreases. Importantly, the effect of the tuning function on predicted sensitivity is wholly independent from the effect of the sample-size relation, indicating separable constraints.

Different model configurations varying in the number of estimated parameters, the functional form of the tuning function, the type of sensitivity weighted, and the constraints between parameters were individually generated and fit to data, giving a total of 72 model configurations for each observer; see Table 5.1 for the factors and their levels.

Table 5.1: Differing model configuration factors in the construction of the extended signal detection theory models for Experiment 3.

Model factor	Description	Factor levels
SS constraint	Whether the sample-size relationship is enforced, leading to a constraint on sensitivity between different memory array sizes.	Yes (sensitivity is constrained to the sample size relation, $\sum_m (d'_m)^2 = \text{const.}$); No (the effect of memory array size on sensitivity is freely estimated).
Curve	The functional form of the relationship between sensitivity and the target angular offset as defined by the tuning function.	<i>Gaussian</i> (half-Gaussian function); <i>Cosine</i> (quarter-cycle cosine); <i>Triangle</i> (decreasing linear function from one to zero over the bandwidth).
Bandwidth	Whether the width of the tuning function (that is the selectivity of the detector over different orientations) is estimated for each memory array size condition, or is constrained to be equal across all conditions.	<i>Free</i> (the bandwidth parameters are estimated for each memory array size condition); <i>Equal</i> (the bandwidth parameters are constrained to be equal across conditions).
Squared	Whether the tuning function is weighting squared sensitivity or sensitivity directly.	Yes (tuning function is applied to squared-sensitivity); No (tuning function is applied to sensitivity).
Offset	The centre of the tuning function—either fixed at zero (i.e., centred on the point of zero discriminability) or estimated.	<i>Free</i> (the offset is estimated for each memory array size condition); <i>Equal</i> (the offset is estimated but constrained to be equal across all memory array size conditions); <i>Zero</i> (the offset is fixed to zero).

Table 5.2: Five signal detection theory models extended to account for fine orientation discrimination that best fit, in terms of BIC, the group average data in Experiment 3.

Curve	SS constraint	Bandwidth	Offset	Squared	G^2	BIC
Gaussian	Yes	Equal	Zero	Yes	9.292	25.799
Cosine	Yes	Equal	Zero	Yes	11.035	27.541
Triangular	Yes	Equal	Zero	No	11.975	28.482
Gaussian	Yes	Equal	Zero	No	15.017	31.524
Gaussian	Yes	Equal	Equal	Yes	9.292	34.052

Refer Table 5.1 for nomenclature.

For each specific model tested, the best fitting parameter was found by minimising the objective function,

$$G^2 = 2 \sum_{i \in A} \sum_{j \in B} N_{i,j} \sum_{k \in R} p_{i,j,k} \cdot \ln \frac{p_{i,j,k}}{\pi_{i,j,k}},$$

where the index i runs over the set of the eight angular offset conditions, A ; the index j runs over the set of memory array conditions (from a single item presented to four items presented), B ; the index k runs over the set of response types (clockwise or counter-clockwise), R ; $N_{i,j}$ is the number of trials in each experimental condition of a specified angular offset and memory array level; and, $p_{i,j,k}$ and $\pi_{i,j,k}$ are the observed and predicted number of response proportions of the type indicated by the triple $\langle i, j, k \rangle$. The five best-fitting models for the group average data is presented in Table 5.2, selected on the basis of BIC values. The five best-fitting models for each observer, in BIC terms, are presented in tables at the end of this chapter for clarity (in Tables 5.5–5.9).

In general, the extended signal detection theory model with the inclusion of an orientation-specific tuning function, appears to capture the main qualitative variance observed in the data: the systematic effect of angular offset on performance can most clearly be seen in the change of accuracy across offset conditions in single-item displays. Like the model results reported in Chapter 3, the presence of a sample-size constraint between observer accuracy and the size of the memory array was supported by the data, with all of the top five models in terms of BIC for each observer (other than observer XL) employing this constraint. The lack of support for the sample-size constraint from the performance of observer XL appears to be due to the overall reduction in the magnitude of any effect of memory array size.

The best fitting model configuration in all cases—except for the observer XL—conformed to the sample-size relation and used a Gaussian-shaped tuning function, operating on the squared-sensitivity, where the preferred orientation of the tuning function was cen-

tred on the point of zero-discriminability. No predicted interaction between the bandwidth of the tuning function and the size of memory array was found: only a single bandwidth was used to predict the change in accuracy for each memory array size. Figure 5.7 displays the predictions of this model configuration presented against the group average observer data.

5.1.3 Discussion

The aim of Experiment 3 was to examine the relationship between observer performance over different levels of stimulus discriminability, and to investigate a generalisation of the sample-size relationship previously reported by Sewell et al. (2014). The data obtained showed a systematic effect of both stimulus discriminability and memory array size on observer accuracy, clearly seen in the group average data, and well accounted for by a simple model which used a Gaussian-shaped tuning function centred on the zero-discriminability orientation to modify observer sensitivity.

The finding that the best fitting model predicted a relationship between the angular offset of the target stimulus and squared-sensitivity, rather than simple sensitivity, intimates that the effect of the angular offset condition affected the variability of information entering the visual short-term memory system, rather than simply affecting stimulus confusability at the decision stage. Although this is consistent with the interpretation of response variability in continuous response tasks as reflecting a change in memory “precision” (Bays et al., 2009; Zhang & Luck, 2008), that interpretation is inconsistent with the lack of any interaction between the memory array size and the estimated bandwidth parameter. If it was assumed that an underlying resource constraint in the memory system affected the precision at which observers could discriminate between fine orientation, then one might expect to see an increasing bandwidth as the memory array size increased. The current data suggest, however, no structural change to the overall precision at which orientation can be represented in memory system, but a change in the signal-to-noise ratio.

This conclusion is somewhat difficult to appraise immediately given the variability in the data, particularly when examining individual observer data. This variability in the individual observer data—despite the large number of observations per cell—indicates that comparisons to an oblique internal standard may be difficult. (As will be shown in the next chapter, however, the response time data for this experiment is also consistent in showing the effects of memory array size.)

Further exploration of the relationship between the stimulus discriminability and mem-




Figure 5.7: Predictions of the extended signal detection theory model against the group average data. This model possessed a Gaussian-shaped tuning function, with the bandwidths of those functions constrained to be equal between both exposure durations and memory array sizes and the centre of the function at the zero-discriminability orientation. The effect of memory array size on observer sensitivity is constrained to follow the sample size relation. This model was the best fitting model (in terms of BIC) for the group average data. In the figure, dashed lines represent the observed data (cf. Figure 5.2) and the solid lines represent the prediction of the model.

ory array size, with the addition of an experimental manipulation of the stimulus exposure duration, will be examined in the next experiment.

5.2 Experiment 4

The previous experiment used a two-alternative fine-orientation discrimination to demonstrate a systematic relationship between increasing stimulus discriminability and increasing observer performance, independently of the information constraint on observer performance described by the sample-size relation. The second experiment expanded this methodology to examine the relationship between stimulus discriminability, memory array size, and observer performance with different levels of stimulus exposure duration. Rather than using an oblique stimulus as the standard to compare against, and potentially increasing the variability in the observed data due to uncertainty regarding the standard, this experiment used a vertical stimulus as the standard.

The conclusions drawn from previous studies examining visual short-term memory (Gegenfurtner & Sperling, 1993; Ratcliff & Rouder, 2000; Sewell et al., 2014; Smith et al., 2004; Vogel et al., 2006), as well as the first part of this thesis, is that increasing the stimulus exposure duration increases the effective amount of information for perceptual decision-making held within the memory system. This conclusion is compatible, as stated in the introduction of this chapter, with the findings of Bays, Gorgoraptis, Wee, Marshall, and Husain (2011). Bays and colleagues found that increases in the stimulus exposure duration were related to decreases in the variability of responding in a continuous report task. Their conclusion was that the effect of increasing stimulus exposure duration, like decreasing the number of items to be maintained simultaneously, could increase the precision at which memory representations are stored. However, the apparent increase in precision may be due to additional information available to the decision process, rather than a change in the structure of the representation itself. As shown from the equal bandwidth estimated between memory array size conditions in the first experiment, the effect of the increasing the amount of available information for the probed item may not lead to a higher quality of information—changing the profile of stimulus discriminability—but simply increasing the discriminability across all orientations uniformly as the amount of total information in the memory system increases.

These arguments will be explored in the following experiment.

5.2.1 Method

Participants

Five observers participated in this study, all paid observers who were naïve to the aims of the study selected from the University of Melbourne (HA, LA, ME, SN, and TS). Each observer was briefed about the general nature of the study and signed a consent form prior to participation. Each observer other than myself was remunerated AUD \$12 for each session completed.

As with the first experiment, each observer completed a variable number of practice and calibration sessions. All sessions lasted approximately thirty-five minutes each, with regular breaks between blocks of trials for observers to rest.

Stimuli and apparatus

The experiment largely followed the design presented for the previous experiment, using arrays of one to four Gabor patches presented around a central fixation cross. Unlike the first experiment, however, the Gabor patches were oriented clockwise or counterclockwise from a vertical—rather than diagonal—position (the point of zero discriminability). Four different levels of angular offset were used: 0.1, 0.3, 0.5, or 0.7 radians away from the point of zero discriminability which was consistent between observers.

Stimulus contrast was also manipulated between observers in attempting to provide the maximum range between the most difficult condition (four items presented for 100 ms, with the target very close to a vertical orientation) above chance, and the least difficult condition (a single item presented for 200 ms, with the target item oriented 0.7 radians away from a vertical position) below ceiling performance.

Procedure

A $4 \times 3 \times 4$ within-subjects design was used: four memory array sizes (1, 2, 3, or 4 items displayed); three stimulus exposure durations (100, 150, or 200 ms); and four levels of target angular offset from a vertical position. Each session of the experiment consisted of 384 trials, yielding a total of 5,760 trials per observer after the full fifteen experimental sessions. The timing for presentation (see Figure 5.8) was exactly the same as the first experiment, except that the stimulus exposure duration was variable, changing on a trial-to-trial basis.




Figure 5.8: A schematic overview of the timing and structure of each trial in Experiment 4.

5.2.2 Results

Figure 5.9 shows the average accuracy across observers for each memory array size and stimulus exposure duration as a function of the angular offset of the target; Figures 5.11–5.15 show the accuracy for each individual observer, and are presented at the end of the chapter.

As with the first experiment, prior to the discussion of the sensitivity model based signal detection theory, a mixed-effects logistic regression model was used to provide a preliminary analysis of the data. The logistic regression was conducted on the proportion of correct responses, with memory array size, stimulus exposure duration, and target angular offset treated as fixed effects. Individual differences of each observer on the intercept were treated as a random effect. Stimulus exposure duration was scaled to seconds to increase the stability of the model estimation routines. A significant main effect of the memory array size on observer accuracy was found, $\beta = -0.177$, $SE = 0.083$, $p = 0.032$, with accuracy predicted to decrease as memory array size increased. The main effect of stimulus exposure duration on accuracy was not significant, $\beta = -2.435$, $SE = 1.381$, $p = 0.077$; the main effect of target offset magnitude on accuracy was also not significant, $\beta = -0.487$, $SE = 0.355$, $p = 0.170$. The two-way interaction be-




Figure 5.9: Group average accuracy data from Experiment 4, showing the proportion of correct responses as a function of the target angular offset (in radians), conditioned by memory array size (different lines on the plot) and by the memory array exposure duration (different panels in the plot).

tween memory array size and stimulus exposure duration was significant, $\beta = 1.212$, $SE = 0.515$, $p = 0.019$, with the difference between accuracy levels between memory array size conditions predicted to increase as the exposure duration increased. The two-way interaction between memory array size and angular offset of the target was also significant, $\beta = 0.321$, $SE = 0.132$, $p = 0.015$, with the increase in accuracy with larger target offset magnitudes predicted to rise more slowly in larger memory array size conditions. The final two-way interaction between stimulus exposure duration and the angular offset of the target was also significant, $\beta = 10.061$, $SE = 2.253$, $p < 0.001$, with the increase in accuracy with larger target offset magnitudes predicted to rise more quickly with larger exposure durations. Last, the three-way interaction between the memory array size, the stimulus exposure duration, and the angular offset of the target was significant, $\beta = -3.979$, $SE = 2.253$, $p < 0.001$, with the increase in accuracy levels with larger target offset magnitudes predicted to increase more quickly for smaller memory array sizes as exposure duration increased than for larger memory array sizes.

5.2.3 Sensitivity modelling

The models earlier presented in §5.1.2 expanded the two-alternative orientation discrimination model presented earlier in §3.1 to account for data where the possible target orientations were not orthogonal. The prediction of that model was that the angular offset of the target item would directly affect sensitivity—or squared sensitivity—by an appropriately chosen tuning function, independent of the effect of memory array size described by the sample size relation. In the current experiment, the exposure time of the stimulus was also manipulated between trials, in addition to the discriminability of the target stimulus.

In order to account for the effect of exposure duration, different relationships between sensitivity and stimulus exposure duration, in terms of the milliseconds of memory array presentation, were examined. A strong interpretation of a sample size account would predict that the rate of information processing in the visual short-term memory system is fixed, giving a linear growth in the squared sensitivity function with the stimulus exposure duration (Sewell et al., 2014, in press). This account, although naturally arising from the constant rate sampling properties of the sample size model, assumes that observed sensitivity is identical with the rate of information formation, and that the information being sampled from remains consistent over the sampling period. The use of high spatial frequency information in fine orientation discrimination decisions suggests that perceptual representations may take longer to form and consolidate than perceptual representations required to make orthogonal orientation discrimination judgements, as reviewed earlier, and that non-linearities may arise due to this formation process. In fitting the current experiment, preliminary model fitting indicated that substantially better fits between the model and observed data were obtained by a linear relationship between the stimulus exposure duration and observed sensitivity—rather than squared sensitivity. The relationship between sensitivity and exposure duration used for the models where the tuning function directly weights squared sensitivity was, therefore,

$$d'_{i,m} = \frac{\sqrt{f(b \cdot (\theta_i - a))} \cdot d'_1}{\sqrt{m}} \alpha \tau,$$

where the indices i and m run over the different target angular offset and memory array size conditions, respectively; α is an estimated rate of sensitivity growth; and, τ is the stimulus exposure duration in milliseconds. In the model configurations where the

tuning functions weighted sensitivity, the models were of the form

$$d'_{i,m} = \frac{f(b \cdot (\theta_i - a)) \cdot d'_1}{\sqrt{m}} \alpha \tau.$$

In all, 720 model configurations were tested for each observer and for the composite (averaged) observer. In addition to the factors manipulated in the model fitting for the previous experiment, additional factors and factor levels were added to characterise the effect of exposure duration on sensitivity (see, Table 5.3). The five best fitting models, in terms of BIC, for the group average data are shown in Table 5.4. Tables 5.10–5.14 display the best fitting models for each observer and are presented at the end of this chapter for clarity.

In examining the best fitting models for each of the observers, the effect of exposure duration on sensitivity was seen as a linear multiplicative factor on sensitivity, rather than squared sensitivity, in all observers except ME. The sample-size constraint relating sensitivity to memory array size was also supported in each of the observers, consistent with the model results presented in Chapter 3 and the earlier results of Experiment 3. Although there is support overall for the dominant model configuration seen in Experiment 5—with a Gaussian-shaped tuning curve centred on the orientation of zero-discriminability, and with equal bandwidths between conditions—this model configuration is not the top model according to BIC for every observer. For the data of observer HA, a cosine curve fits the accuracy data slightly better (the configuration with the Gaussian-shaped tuning function is the second best fitting). For the data of observer LA, there are differences predicted in the bandwidth of the tuning function between exposure duration conditions, but the model with equal bandwidths across all conditions is ranked fourth. For the data of observer ME, the dominant model does not rank within the top five (it ranked 93rd of 720 model configurations) and the top configurations do not seem to be consistent with the model configurations seen for other observers. Last, for the data of observer SN, a triangular tuning function is preferred slightly over the Gaussian-shaped tuning function, which is ranked second.

Largely, however, the dominant model shown in Experiment 3 seems to be consistent with the data of Experiment 4, with the bandwidths largely remaining equal across experimental conditions and the growth in sensitivity apparently linear with increasing exposure duration. The deviations in rankings of model configurations appears to be due to the variability in the individual data, not seen in the composite observer data. The fit between this model configuration and the observed group average data can be seen in Figure 5.10.

Table 5.3: Differing model configuration factors for manipulated in constructing the extended signal detection theory models for Experiment 4.

Model factor	Description	Factor levels
SS constraint	Whether the sample size relationship is enforced, leading to a constraint on sensitivity between different memory array sizes.	<i>Yes</i> (sensitivity is constrained to the sample size relation, $\sum_m (d'_m)^2 = \text{const.}$); <i>No</i> (the effect of memory array size on sensitivity is freely estimated).
Curve	The functional form of the relationship between sensitivity and the target angular offset as defined by the tuning function.	<i>Gaussian</i> (half-Gaussian function); <i>Cosine</i> (quarter-cycle cosine); <i>Triangle</i> (decreasing linear function from one to zero over the bandwidth).
Bandwidth	Whether the width of the tuning function (that is the selectivity of the detector over different orientations) is estimated for each memory array size condition, is estimated for each exposure duration condition, is estimated freely in each condition, or is constrained to be equal across all conditions.	<i>Free</i> (the bandwidth parameters are estimated freely for each memory array size and exposure duration condition); <i>Free-ExpDur</i> (the bandwidth parameters are estimated for each exposure duration condition); <i>Free-Size</i> (the bandwidth parameters are estimated for each memory array size condition); <i>Equal</i> (the bandwidth parameters are constrained to be equal across conditions).
Squared	Whether the tuning function is weighting squared sensitivity or sensitivity directly.	<i>Yes</i> (tuning function is applied to squared-sensitivity); <i>No</i> (tuning function is applied to sensitivity).
Offset	The centre of the tuning function—either fixed at zero (i.e., centred on the point of zero discriminability) or estimated.	<i>Free</i> (the offset is estimated for each exposure duration and memory array size condition); <i>Free-ExpDur</i> (the offset is estimated for each exposure duration level); <i>Free-Size</i> (the offset is estimated for each memory array size); <i>Equal</i> (the offset is estimated but constrained to be equal across all memory array size and exposure duration condition); <i>Zero</i> (the offset is fixed to zero).
Times	The effect of exposure duration on sensitivity.	<i>Free</i> (each exposure duration has a freely estimated marginal effect on sensitivity); <i>Linear</i> (sensitivity is linearly related to exposure duration); <i>None</i> (sensitivity is unrelated to exposure duration).

Table 5.4: Five signal detection theory models extended to account for fine orientation discrimination that best fit, in terms of BIC, the group average data in Experiment 4.

Curve	SS constraint	Bandwidth	Offset	Times	Squared	G^2	BIC
Gaussian	Yes	Equal	Zero	Linear	Yes	19.854	45.83
Cosine	Yes	Equal	Zero	Linear	Yes	24.551	50.527
Triangular	Yes	Equal	Equal	Linear	Yes	19.461	54.096
Triangular	Yes	Equal	Zero	Linear	No	28.122	54.098
Gaussian	Yes	Equal	Equal	Linear	Yes	19.707	54.341

Refer Table 5.3 for nomenclature.




Figure 5.10: Predictions of the extended signal detection theory model against the data of the composite observer. This model possessed a Gaussian-shaped tuning function, with the bandwidths of those functions constrained to be equal between both exposure durations and memory array sizes and the centre of the function at the zero-discriminability orientation. The effect of memory array size on observer sensitivity is constrained to follow the sample size relation and the effect of exposure duration is linear (on sensitivity, not squared sensitivity). This model was the best fitting model (in terms of BIC) for the composite observer. In the figure, dashed lines represent the observed data (cf. Figure 5.9) and the solid lines represent the prediction of the model.

5.2.4 Discussion

The aim of Experiment 4 was to further investigate the relationship between the content of stimuli stored in visual short-term memory, in this case the orientation value for each element in the memory store, and overall performance in a simple two-alternative discrimination task. Building upon Experiment 3, Experiment 4 also included the experimental manipulation of the stimulus (memory array) exposure duration—from 100 ms to 200 ms. The results of the extended signal detection theory model fitting exercise was consistent with the findings of the previous experiments reported in the thesis: the effect of stimulus discriminability on performance was well captured by a tuning function directly weighting the squared-sensitivity parameter; moreover, the compatibility of the observed performance with the sample-size relationship was supported by the consistently better results obtained by models with the sample-size restriction between sensitivity and memory array size imposed over models where the relationship between sensitivity and memory array size is freely estimated.

The effect of stimulus exposure duration on performance was captured by a linear relationship, but this relationship was multiplied with the sensitivity parameter directly, rather than the sensitivity parameter squared as with the reciprocal of the memory array size and the tuning function. The relationship between sensitivity and the stimulus exposure duration is contrary to a strong, mechanistic interpretation of the sample-size model which predicts a constant rate of information accumulation across time, leading to a linear change in squared sensitivity. When examined in a preliminary model fitting, however, the models with a squared time coefficient fit substantially better than those with a linear time coefficient.

The lack of linear growth in the squared sensitivity does not threaten the overall validity of the sample size model as a description of resource division as, although constant rate of accumulation is a consequence of a strict interpretation of the sample-size model—the invariance of the sum of squared d' arises from samples being drawn steadily from an underlying perceptual representation—the division of resources and the rate of resource accumulation are not mutually inclusive. Instead, it may be the case that the discriminative information required for the decision in this task, provided by high spatial frequency detectors which have a slower response to stimulation (Watt, 1987), may mean that the underlying perceptual representation is not fully stabilised or stationary over the period of sampling, leading to an acceleration in the rate of squared sensitivity increasing. This faster-than-linear increase in squared sensitivity appears to be fit by a quadratic function in this case. By contrast, the rate of squared sensitivity growth using lower spa-

tial frequency detectors in the orthogonal discrimination case reported by Sewell, Lilburn, and Smith (2014) appears to stable over the entire decision period, consistent with the faster response of detectors tuned to lower spatial frequencies. Attentional factors have also been previously noted as changing the apparent rate of sensitivity increases with increasing stimulus information (Smith et al., under review), captured by weighted sample-size models (Bonnell & Hafter, 1998; Bonnell & Miller, 1994).

5.3 General discussion

The experiments presented in this chapter aimed to extend the previous literature showing a sample-size capacity constraint in visual short-term memory by using stimuli where similarity was directly manipulated. By using a two-alternative fine orientation discrimination task, it was argued that these data provided an insight into the structure and capacity of the visual memory system without increasing the complexity of the decision process required of the observer. In turn, this allowed the application and extension of simple signal detection theory models introduced earlier in this thesis.

The results obtained from these experiments show a clear and systematic effect of the difference in orientation between the target and an orientation of zero discriminability on observer performance. This systematic effect was well captured by an extension to the signal detection theory presented in Chapter 3 to account for the decrement in performance caused by detectors with preferred orientations that correspond to an incorrect response option (that is, coding for orientations counterclockwise to the zero discriminability orientation when the presented stimulus is, in fact, clockwise). By accounting for the false alarms caused by these detectors, weighting squared sensitivity directly using a tuning function over orientation values, the change in observed performance seen over orientation levels was predicted by the model.

Like the models presented in Chapter 3, the sample size relation between observer sensitivity and the memory array size was almost uniformly supported in both experiments presented in this chapter, with the exception of observer XL in Experiment 3. This suggests that the information capacity placed upon representations maintained actively within visual short-term memory does not appear to interact with the content of those items.

Furthermore, model configurations with the constraint that the equal bandwidths of the tuning functions between experimental conditions tended to outperform models where those bandwidths were freely estimated for different experimental conditions.

This result is inconsistent with the general finding of continuous report paradigms that memory array size or exposure duration change the precision at which a representation is maintained (Bays et al., 2009; Bays, Gorgoraptis, et al., 2011; Bays & Husain, 2008; Zhang & Luck, 2008), with higher stimulus exposure durations and lower memory array sizes leading to more “precise” representations, inferred from less response dispersion. The discrepancy between no apparent change in the resolution of orientation discrimination, as indicated by the invariance of the bandwidth between experimental conditions, and changes in response dispersion in continuous report tasks can be explained by a consideration of the way that the decision process differs between the two, following the comments made in the introduction to this chapter. The examination of continuous report data in terms of the variability of responses alone can not differentiate between systematic changes in the quality of stimulus information held within the visual memory system and systematic changes at the decision level, such as changes in response bias (Smith, 2016). As such, noise in the decision stage is indistinguishable with noise in the representation itself when using a continuous report paradigm without some consideration of the response time. By contrast, the extended signal detection theory model proposed in this chapter indicates that changes in the apparent response variability are due to the signal-to-noise ratio of the psychophysical detectors to background noise; the structure of the memory system does not change with additional information, nor does the representation become more precise in the sense that the tuning functions of detectors change. Indeed, according to this model, an increase in the amount of stimulus information with a longer stimulus exposure duration or fewer items to maintain simultaneously leads to both an increase in the hit rate due to activation of detectors corresponding to the correct response category, but also an increase in false alarms due to increasing activation of detectors corresponding to the incorrect response category. Further specification of the relationship between the angular offset of the target and the information available to the decision process will be given in the next chapter, where response time modelling using the diffusion model will be undertaken.

The second part of this chapter also examined the relationship between stimulus exposure duration and sensitivity, with evidence to suggest that sensitivity—rather than squared sensitivity—is linearly related to exposure time. As indicated in the preceding discussion of Experiment 4, this is not the natural prediction of a mechanistic interpretation of the sample size model or the results reported by Sewell, Lilburn, and Smith (2014), however. The faster-than-linear growth in squared sensitivity—fit with a linear relation between exposure duration and sensitivity—appeared to predict the data better than a linear relationship between exposure duration and squared sensitivity. This

may be due to the slower response to stimulation exhibited by the high spatial frequency psychophysical detectors required to make a fine orientation discrimination decision, compared with the faster response of low spatial frequency detectors required for orthogonal orientation discrimination. Further research comparing directly both the fine orientation discrimination procedure of the kind reported here and orthogonal orientation discrimination procedure of the kind reported by Sewell, Lilburn, and Smith (2014) is required to verify if the different temporal profiles of sensitivity increasing with exposure duration change with stimulus discriminability.

The next chapter extends the analysis of this chapter by considering the response time data of Experiments 3 and 4, and attempting to fit response time distributions using an analogous version of the signal detection theory models expressed in terms of the diffusion model.

Supplementary tables and figures

Best fitting signal detection theory model fits for individual observers in Experiment 3.	Tables 5.5–5.9
Accuracy data for individual observers in Experiment 4.	Figures 5.11–5.15
Best fitting signal detection theory model fits for individual observers in Experiment 4.	Tables 5.10–5.14

Table 5.5: Five signal detection theory models extended to account for fine orientation discrimination that best fit, in terms of BIC, observer AB in Experiment 3.

Curve	SS constraint	Bandwidth	Offset	Squared	G^2	BIC
Gaussian	Yes	Equal	Zero	Yes	51.781	68.288
Cosine	Yes	Equal	Zero	Yes	55.184	71.69
Gaussian	Yes	Equal	Zero	No	56.624	73.131
Triangular	Yes	Equal	Equal	Yes	49.987	74.747
Triangular	Yes	Equal	Zero	No	58.948	75.455

Refer Table 5.1 for nomenclature.

Table 5.6: Five signal detection theory models extended to account for fine orientation discrimination that best fit, in terms of BIC, observer CVH in Experiment 3.

Curve	SS constraint	Bandwidth	Offset	Squared	G^2	BIC
Gaussian	Yes	Equal	Zero	Yes	42.061	58.567
Cosine	Yes	Equal	Zero	Yes	42.873	59.38
Triangular	Yes	Equal	Zero	Yes	45.579	62.086
Triangular	Yes	Equal	Zero	No	46.453	62.959
Gaussian	Yes	Free	Zero	Yes	24.485	65.751

Refer Table 5.1 for nomenclature.

Table 5.7: Five signal detection theory models extended to account for fine orientation discrimination that best fit, in terms of BIC, observer SIS in Experiment 3.

Curve	SS constraint	Bandwidth	Offset	Squared	G^2	BIC
Cosine	Yes	Equal	Zero	Yes	34.574	51.08
Gaussian	Yes	Equal	Zero	Yes	34.863	51.369
Gaussian	Yes	Equal	Zero	No	37.855	54.362
Gaussian	Yes	Equal	Equal	No	34.107	58.867
Cosine	Yes	Equal	Equal	Yes	34.221	58.98

Refer Table 5.1 for nomenclature.

Table 5.8: Five signal detection theory models extended to account for fine orientation discrimination that best fit, in terms of BIC, observer SL in Experiment 3.

Curve	SS constraint	Bandwidth	Offset	Squared	G^2	BIC
Gaussian	Yes	Equal	Zero	Yes	38.142	54.649
Cosine	Yes	Equal	Zero	Yes	40.008	56.514
Triangular	Yes	Equal	Zero	No	40.814	57.321
Gaussian	Yes	Equal	Zero	No	41.085	57.591
Gaussian	Yes	Equal	Equal	No	37.736	62.496

Refer Table 5.1 for nomenclature.

Table 5.9: Five signal detection theory models extended to account for fine orientation discrimination that best fit, in terms of BIC, observer XL in Experiment 3.

Curve	SS constraint	Bandwidth	Offset	Squared	G^2	BIC
Triangular	No	Equal	Equal	No	21.308	70.828
Gaussian	No	Equal	Equal	No	22.184	71.704
Gaussian	No	Equal	Equal	Yes	22.381	71.901
Cosine	No	Equal	Equal	No	22.479	71.999
Triangular	No	Equal	Equal	Yes	23.115	72.635

Refer Table 5.1 for nomenclature.




Figure 5.11: Accuracy data from Experiment 4 for observer HA, showing the proportion of correct responses as a function of the target angular offset (in radians), conditioned by memory array size (different lines on the plot) and by the memory array exposure duration (different panels in the plot).




Figure 5.12: Accuracy data from Experiment 4 for observer LA, showing the proportion of correct responses as a function of the target angular offset (in radians), conditioned by memory array size (different lines on the plot) and by the memory array exposure duration (different panels in the plot).




Figure 5.13: Accuracy data from Experiment 4 for observer ME, showing the proportion of correct responses as a function of the target angular offset (in radians), conditioned by memory array size (different lines on the plot) and by the memory array exposure duration (different panels in the plot).




Figure 5.14: Accuracy data from Experiment 4 for observer SN, showing the proportion of correct responses as a function of the target angular offset (in radians), conditioned by memory array size (different lines on the plot) and by the memory array exposure duration (different panels in the plot).




Figure 5.15: Accuracy data from Experiment 4 for observer TS, showing the proportion of correct responses as a function of the target angular offset (in radians), conditioned by memory array size (different lines on the plot) and by the memory array exposure duration (different panels in the plot).

Table 5.10: Five signal detection theory models extended to account for fine orientation discrimination that best fit, in terms of BIC, observer HA in Experiment 4.

Curve	SS constraint	Bandwidth	Offset	Times	Squared	G ²	BIC
Cosine	Yes	Equal	Zero	Linear	Yes	75.361	101.337
Gaussian	Yes	Equal	Zero	Linear	No	75.599	101.575
Cosine	Yes	Equal	Zero	Linear	No	76.311	102.287
Gaussian	Yes	Equal	Zero	Linear	Yes	76.868	102.845
Gaussian	Yes	Equal	Equal	Linear	Yes	75.246	109.881

Refer Table 5.3 for nomenclature.

Table 5.11: Five signal detection theory models extended to account for fine orientation discrimination that best fit, in terms of BIC, observer LA in Experiment 4.

Curve	SS constraint	Bandwidth	Offset	Times	Squared	G ²	BIC
Gaussian	Yes	Free-ExpDur	Zero	None	Yes	55.231	89.866
Cosine	Yes	Free-ExpDur	Zero	None	Yes	58.227	92.862
Triangular	Yes	Free-ExpDur	Zero	None	No	59.227	93.862
Gaussian	Yes	Equal	Zero	Linear	Yes	68.9	94.876
Gaussian	Yes	Free-ExpDur	Zero	None	No	61.234	95.869

Refer Table 5.3 for nomenclature.

Table 5.12: Five signal detection theory models extended to account for fine orientation discrimination that best fit, in terms of BIC, observer ME in Experiment 4.

Curve	SS constraint	Bandwidth	Offset	Times	Squared	G ²	BIC
Gaussian	Yes	Free-Size	Zero	Linear	No	105.407	157.359
Gaussian	Yes	Equal	Free-Size	Linear	No	101.062	161.673
Gaussian	Yes	Free-Size	Equal	Linear	No	102.229	162.84
Cosine	Yes	Equal	Free-Size	Linear	Yes	102.635	163.246
Gaussian	Yes	Equal	Free-Size	Linear	Yes	102.689	163.299

Refer Table 5.3 for nomenclature.

Table 5.13: Five signal detection theory models extended to account for fine orientation discrimination that best fit, in terms of BIC, observer SN in Experiment 4.

Curve	SS constraint	Bandwidth	Offset	Times	Squared	G^2	BIC
Triangular	Yes	Equal	Zero	Linear	Yes	77.221	103.197
Gaussian	Yes	Equal	Zero	Linear	Yes	82.053	108.03
Gaussian	Yes	Equal	Equal	Linear	Yes	77.246	111.881
Gaussian	Yes	Equal	Equal	Linear	No	77.519	112.154
Gaussian	Yes	Equal	Zero	Linear	No	86.52	112.497

Refer Table 5.3 for nomenclature.

Table 5.14: Five signal detection theory models extended to account for fine orientation discrimination that best fit, in terms of BIC, observer TS in Experiment 4.

Curve	SS constraint	Bandwidth	Offset	Times	Squared	G^2	BIC
Gaussian	Yes	Equal	Zero	Linear	Yes	94.726	120.702
Cosine	Yes	Equal	Zero	Linear	Yes	98.794	124.77
Triangular	Yes	Equal	Zero	Linear	No	99.992	125.968
Triangular	Yes	Equal	Zero	Linear	Yes	100.938	126.914
Gaussian	Yes	Equal	Equal	Linear	No	93.074	127.709

Refer Table 5.3 for nomenclature.

Chapter 6

Fine orientation discrimination: response time

In the previous chapters of this thesis, I attempted to provide a model of the multiple constraints on the visual short-term memory system by placing an emphasis on the structure of the task used to examine memory performance and the decision process that arises from that task structure. The previous chapter attempted to model the effect of memory array size, the exposure duration of the memory array, and stimulus discriminability simultaneously through the use of a modified version of the two-alternative orientation discrimination paradigm introduced in Chapter 3, and the application of an extended signal detection theory model also outlined in that chapter. Chapter 4 continued the analysis of the experiments in Chapter 3 by examining response time distributions, and where the response time models introduced in Chapter 4 had constraints placed upon them following the models presented in Chapter 3. Similarly, in this chapter I examine the response time data for the experiments presented in the last chapter and extend the modelling presented into the response time domain.

In the introduction of Chapter 4, some of the existing work using response time distributions within the visual short-term memory literature was reviewed. Two critical studies, the first conducted by Ratcliff and Rouder (2000) and the second conducted by Donkin et al. (2013), demonstrated the utility of considering both accuracy and response time jointly by providing analyses of competing accounts of the memory system. It would not be possible to differentiate between these these competing accounts if only accuracy was considered alone, and these studies required the extra constraints provided by response time data. Both studies used simple, well described experimental paradigms as the basis for analysis: two-alternative forced choice responses and change detection.

Although response time methods allow a precision in model specification that cannot

be obtained by examining accuracy alone, the requirements for large datasets, clearly defined decision processes in designing experimental procedures, and an emphasis on speeded responses may help explain the paucity of response time modelling in the visual short-term memory literature to date. In particular, there has been a lack of response time modelling employing data obtained using the popular continuous report paradigm. Continuous report has been used to study the structure and fidelity of representations held within visual short-term memory, but the manner in which responses are made—the difficulty in providing speeded responses and the lack of consensus regarding the underlying decision process—makes response time modelling difficult.

The difficulty of obtaining response time data in continuous report tasks has meant that theoretical questions about the representation of feature information in visual short-term memory have not been examined in depth using response time data. Although, as previously discussed, the value of the continuous report paradigm is the direct analogy between the response procedure with the assumed form of the internal representation, and the precision of the response data, it is possible to examine information at the feature level whilst also using an experimental paradigm that controls the complexity of the decision required to make a response.

One study to have attempted response time distribution modelling from a procedure with a focus of visual information at the feature level was reported by Pearson, Raškevičius, Bays, Pertzov, and Husain (2014). Pearson and colleagues used a change detection procedure where the probed item was displaced either in terms of angle or lateral position from the corresponding item in the memory array. Observers were asked to report whether the probed item was displaced towards or away from the fixation point. Response times were, in the first two experiments, then conditioned by the magnitude of the displacement and the size of the memory array. Model fitting, using a very simple response time model known as LATER (Carpenter, 1981; Carpenter & Williams, 1995; Reddi & Carpenter, 2000), across this data was conducted to examine whether the effect of the size of the offset and the number of items in the memory array changed the overall quality of stimulus information used in the decision process or the response bias. Pearson and colleagues found that the effect of both memory array size and the magnitude of the offset systematically changed the quality of information, with the recovered parameter estimates increasing linearly with the slope of the psychometric function for accuracy across displacement conditions.

Although the relationship described by Pearson and colleagues is suggestive, their analysis lacked many key details common to response time modelling in perceptual and memory research (e.g., Donkin et al., 2014; Ratcliff & Rouder, 2000; Sewell & Smith, 2012;

Smith et al., 2009): the fit between the model predictions and the observed response time distributions was difficult to ascertain; only a simple test between the effect of the quality of stimulus information and response bias was examined; and the model used was unnecessarily restricted, unable to model response proportions and response times simultaneously, and with little flexibility for discriminating between effects of nondecision time, variability in the stimulus information, guessing processes. As a result of this lack of detail and further elaboration, the conclusions of the study reported by Pearson and colleagues were somewhat broad: namely, that the retrieval of memory representations required a decision process operating on information both subject to item-level and a feature-level capacity constraints.

The decision and response time models discussed in the preceding chapters provides a much more elaborated set of assumptions and constraints with which to examine the current data. In addition, the two experiments introduced in the previous chapter, and further examined in this chapter, used a simpler two-alternative fine orientation discrimination paradigm, where the rate of the information available for memory encoding is controlled through the use of a near-threshold stimulus contrast and the complexities of the change detection paradigm were avoided. As discussed in Chapters 3 and 4, the use of a two-alternative orientation discrimination task, rather than a change detection task, also increases the sensitivity with which capacity constraints on the memory representations within visual short-term memory can be measured by removing the additional variability added by encoding and comparison of the probe array.

The structure of this chapter is as follows. I will examine each experiment in turn. For each experiment, I first provide a general characterisation of the mean response times across experimental conditions and apply a simple mixed-effects linear model of mean response time. Then I present a diffusion model analysis for experiment. I then provide an analysis of the implications of the result. The chapter will close with a general discussion, providing a summary of the findings from this chapter and the entire thesis.

6.1 Experiment 3

The aim of Experiment 3 was to examine the relationship in observer performance between constraints upon information at the individual feature level and constraints upon whole items stored within visual short-term memory. Although related to work that has previously been conducted using a continuous report paradigm (e.g., Bays, Gorgoraptis, et al., 2011), the response requirements of the task make precise response time modelling

difficult. In contrast, the near-threshold two-alternative fine orientation discrimination paradigm of Experiment 3 is amenable to response time modelling, and the results from Chapters 3 and 4 provide a clear basis for analysis.

Recall that in this experiment, between one and four oriented Gabor patches were presented to the observer. The observer was to indicate whether a single probed item held within memory was rotated either counterclockwise or clockwise from a diagonal position, the magnitude of this rotation is the “target angular offset”. Eight angular offset values were selected to range from near-chance performance for angles close to a diagonal orientation to near-ceiling performance. The full method and accuracy results for Experiment 3 are discussed in §5.1.1 and §5.1.2, respectively.

6.1.1 Results

For the descriptive statistics shown in this section, no response time data was filtered from overall dataset. Figure 6.1 shows the mean response time for each decision type and memory array size condition, averaged across observers; Figure 6.2 shows these data for each observer separately.

In both the group average data and across the individual observers, a large difference can be seen between the mean response times of single item memory arrays and the mean response times of larger array sizes. This may indicate that decisions may begin prior to the presentation of the post-stimulus probe when the probe is irrelevant in distinguishing the target from the distractors.

A mixed-effects linear model was conducted as a preliminary analysis of the response time data at the mean level. This regression on mean response time treated the memory array size and the offset magnitude (in radians) as fixed effects, with the individual differences of the observer treated as a random effect on the intercept. A significant main effect of memory array size on response time (in milliseconds) was found, $\beta = 53.809$, $SE = 4.781$, $p < 0.001$, indicating an increase in mean response time as memory array size increased. A significant main effect of offset magnitude on response time was also found, $\beta = -506.685$, $SE = 62.842$, $p < 0.001$, with response times predicted to decrease as offset magnitude increased. Last, a significant two-way interaction was found between memory array size and offset magnitude, $\beta = 82.074$, $SE = 22.946$, $p < 0.001$, with the differences in mean response time for different memory array sizes lessened with smaller offset magnitudes.




Figure 6.1: Group average mean response time data as a function of offset magnitude from Experiment 3, displayed for each memory array size. Error bars represent one standard error of the mean.




Figure 6.2: Mean response time data as a function of offset magnitude from Experiment 3, displayed for each observer across different memory array sizes. Error bars represent one standard error of the mean.

6.1.2 Diffusion model results

In the previous chapter, I provided a model for observer performance over different levels of angular offset and argued that observer performance is well characterised by a “tuned channel” model, where a Gaussian-shaped tuning function centred on the point of zero discriminability weights the sensitivity measure d' directly. The constraints on sensitivity imposed by angular offset values in these tuned channel models can also be applied to the drift rate in the diffusion model, allowing the predictions of the signal detection theory models to be extended into the response time domain—like the extension of the models in Chapter 3 to the diffusion models of Chapter 4.

In contrast to the many model variants examined for both experiments in Chapter 4, in this chapter the scope of model fitting is a restricted subset of the total number of potential model configurations relating the experimental manipulations to different constraints placed upon the diffusion model. Instead, the focus of these models will be—in the first instance—to examine whether a large number of experimental conditions can be fit simultaneously using a (relatively) small number of assumptions and parameters and, beyond this, to compare different parameterisations of nondecision time, drift variability, and the intrusion process (introduced in Chapter 4).

In large part, the effect of angular offset and memory array size is confined to functions on the drift rate of the diffusion model, the average rate of evidence accumulation. Recall that the best fitting models of the last chapter indicated that the squared sensitivity parameter $(d'_{i,m})^2$ for a fine orientation discrimination trial with a memory array size of m elements and a target angular offset indexed by i was independently constrained by both a sample size relationship and, as mentioned, a Gaussian-shaped tuning function centred on the point of zero discriminability producing the simplified expression,

$$(d'_{i,m})^2 = \frac{(1 - \phi(b\theta_i)) \cdot (d'_1)^2}{m},$$

where θ_i is the angular offset of the target in radians for the i th offset condition, ϕ is the modified Gaussian pdf (normalised and with constrained values for negative angular offsets), b is the bandwidth of the tuning function, and d'_1 is the (estimated) sensitivity for an orthogonal orientation discrimination trial with a single item. In the diffusion model, sensitivity is proportional to the drift rate parameter v and, thus, the drift rate for the same experimental condition can be expressed in the same manner,

$$v_{i,m} = \frac{\sqrt{(1 - \phi(b\theta_i))} \cdot v_1}{\sqrt{m}}.$$

This weighting function on the drift rate in the diffusion model provides a restricted way by which experimental manipulations of target angular offset and memory array size can have a specific influence over both response proportions and response time. The response time models of Chapter 4 demonstrated that close fits could be achieved for an orthogonal orientation discrimination task (as well as for change detection task) by also allowing nondecision time to vary as a function of memory array size. This result was consistent with the findings of Sewell, Lilburn, and Smith (in press), who posited that the increases seen in nondecision time across different sizes of the memory array represented changes in the cost of accessing items within memory. The strongest restriction to place on this pattern of increasing nondecision time values for increasing memory array sizes is to assume nondecision time increases as a linear function of memory array size, reducing the number of parameters that need to be estimated in the model from the number of memory array size conditions (four in the current experiment) to two parameters, an intercept and slope parameter. Models configurations with this restriction will contain the label T^l in their name, with the superscript l representing a linear increase over memory array size. Models configurations with the substring T^* will have no constraints on the relationship between nondecision time and memory array size, with a separate value estimated for each array size condition, using the convention introduced in Chapter 4.

Different drift rates for the intrusion process were also examined. In model fits reported by Sewell, Lilburn, and Smith (in press), substantially better model fits were found when allowing for a “delayed guessing” process—a diffusion process with a drift rate of zero—mixed with some freely estimated probability into the main evidence-based process, as reviewed in Chapter 4. In the model configurations reported in Chapter 4, this mixture process was generalised, with the drift rate for this mixture process—termed an intrusion process—estimated freely to allow for any non-positive drift rate. This was interpreted as allowing for the possibility of distractor stimulus information, with the opposite response category to the target, being accessed mistakenly and being used in some proportion of trials. For the model configurations presented in this chapter, both the zero-drift “guess” process and the freely estimated “intrusion” process were tested. In addition, a third class of model configurations was examined, where the intrusion process was constrained to have a fixed negative drift rate, with the value of that drift rate corresponding to sampling a distractor item of the average angular offset (that is, of average discriminability) from the opposite response category, which I will call a “distractor average” intrusion process. The intrusion process drift rate in this model is not, therefore, directly estimated but is dependent on both the drift rate and bandwidth pa-

rameters. (Note that, in all three cases, the mixture probability of the intrusion process increases as a function of the distractor information of the opposite response category to the target, consistent with the model provided in Chapter 4.) The sampling of distractor information modelled using a negative drift rate in the intrusion process is similar to the use of the “location error” parameter used by Bays et al. (2009) in describing trials where observers would replicate a distractor item, rather than the target item, in a continuous report task.

Last, flexibility in the trial-to-trial variability of the drift rate in the main decision process, denoted η , was also free to vary between memory array size conditions in some of the model configurations tested. Trial-to-trial drift variability reflects variability in the quality of the stimulus information available to the decision process from trial to trial. Allowing variability in the amount of trial-to-trial variability of the drift rate across memory array size conditions allows additional flexibility in the prediction of both the fastest response time quantiles, and in the tails. Trial-to-trial variability is often fixed across the entire experiment, reflecting a consistent set of factors influencing the quality of a representation used for the decision-making across the entire experiment. In the current experiment, however, the use of an intrusion process which changes across memory array size conditions may change the effect of trial-to-trial variability across memory array size conditions. As larger values of η can mean that there is a non-negligible negative drift component in the main decision process, it is expected that these parameters are not fully decorrelated: part of the effect of the intrusion process in modelling the entrance of distractor information into the main decision process may be handled by the main process.

In total, twelve different model configurations were examined; Table 6.1 provides a summary of each of the model configuration factors manipulated. Each of the model configurations was fit to the response time and response proportion data of each observer and the group average by iterative minimisation of the G^2 goodness-of-fit statistic using the Nelder–Mead algorithm as described in Chapter 4. Following the format of Chapter 4, Table 6.2 displays the five model configurations that best fit the group average data. The top five model configurations in fitting the data for each of the individual observers are displayed in Tables 6.5–6.9, presented at the end of this chapter for clarity.

Model configurations with the “distractor average” intrusion drift rate provided the best fit for two of the five individual observers, and for the group average data. The three remaining observers (CVH, SIS, and SL) were best fit by model configurations involving delayed guessing zero-drift intrusion drift process. Model configurations where the estimation of trial-by-trial variability was free to vary between memory array size conditions

Table 6.1: Differing model configuration factors manipulated in examining Experiment 3.

Model factor	Description	Factor levels
Nondecision time (T)	Constraints imposed on estimating the nondecision time value across different memory array size conditions.	<i>Linear</i> (T^L ; nondecision time increases as a linear function of array size); <i>Free</i> (T^S ; nondecision time is freely estimated for each array size condition).
Drift variability (η)	Constraints imposed on estimating the value of the drift variability parameter over different memory array size conditions.	<i>Free</i> (η^S ; drift variability is freely estimated for each array size condition); <i>Fixed</i> (only a single drift variability value is used for all experimental conditions).
Intrusion drift rate (α_v)	Constraints imposed on estimating the value of the drift rate of the intrusion process over different memory array size conditions.	<i>Free</i> (α_v^S ; an intrusion drift rate is estimated freely for each memory array size condition); <i>Zero-drift</i> (α_v^0 ; the intrusion drift rate is fixed at zero, leading to delayed guessing); <i>Distractor average</i> (α_v^- ; the intrusion drift rate is fixed to the drift rate associated with the mean distractor stimulus angular offset).

Table 6.2: The top five best fitting diffusion model configurations, in terms of BIC, for the group average data in Experiment 3.

#	Model	Free parameters	G^2	BIC
1	MOD- η^s - T^s - v_{α}^-	14	182.267	297.734
2	MOD- η^s - T^s - v_{α}^o	14	202.965	318.432
3	MOD- T^s - v_{α}^o	11	229.184	319.908
4	MOD- T^s - v_{α}^-	11	229.262	319.986
5	MOD- η^s - T^l - v_{α}^-	12	222.063	321.035

provided the best fit for three out of the five observers and for the group average. Last, the constraint for linearly increasing nondecision time estimates with increasing memory array sizes was only obtained for a single observer, observer CVH. The other best fitting models required nondecision time to be freely estimated across memory array size conditions.

Predictions for the model configuration that best fit the group average data, MOD- η^s - T^s - v_{α}^- , are displayed against the observed data in Figure 6.3. This model configuration requires a total of 14 free parameters: four drift variability parameters, four nondecision time parameters, a single drift rate (as the tuning function and the sample-size relation provide a value for the drift rate for each of the experimental conditions), a single boundary separation parameter, a single bandwidth parameter, a single nondecision time range parameter, a single intrusion rate, and a single intrusion process time parameter. The model captures the qualitative profile of accuracy and response time across different memory array size and target angular offset conditions, with the most pronounced difference between the model prediction and observed data seen in the smaller angular offsets in memory arrays consisting of three items. The accuracy in trials with larger angular offsets and arrays of four items is slightly underestimated. Overall, however, the model captures both response times at both the leading edge of the distributions (the 0.1 quantile) and the tails of the distributions (the 0.9 quantile), including the asymmetry between correct and incorrect responses. Parameters for this model can be found in Table 6.10, presented at the end of this chapter.

6.1.3 Discussion

The close fit between the model predictions and the data obtained in Experiment 3 show that the diffusion model with both sample size and “tuned channel” constraints on the drift rate alone can provide a compelling account of the two-choice fine orientation dis-




Figure 6.3: A QPP displaying the predictions of the diffusion model configuration $\text{MOD-}\eta^s\text{-T}^s\text{-v}_\alpha^-$ against the group average data of Experiment 3, styled with solid lines and filled squares, against the observed group average data, rendered in open circles and dashed lines. The configuration $\text{MOD-}\eta^s\text{-T}^s\text{-v}_\alpha^-$ includes different drift variability and nondecision time values for each memory array size condition, with the intrusion drift rate set to the average distractor value. SS = Memory array size.




Figure 6.4: Nondecision time parameter estimates for the each individual observer and for the group average data as a function of memory array size using the model configuration $\text{MOD-}\eta^s\text{-T}^s\text{-v}_{\alpha}^s$.

crimination paradigm. Some measure of the predictive power of the model is shown by the relatively parsimonious treatment of the dataset: with the best fitting model configuration providing a good overall account of response time distributions for both correct and incorrect responses for thirty-two experimental conditions using only 14 parameters (the second best model requiring only 11 parameters for a fit of comparable accuracy).

No single consistent model configuration was recovered across both the group average data and the individual observers, although there were commonalities. The clearest commonality between the best fitting model configurations was the rejection of a linear constraint on nondecision times across memory array size conditions, with the model fits to all but a single observer preferencing greater flexibility in the estimation of non-decision time at the expense of a greater number of parameters and, therefore, a greater penalisation in BIC terms. From Figure 6.4, it can be seen that the nondecision time for single item memory arrays is, in some cases, an order of magnitude less than the non-decision time for arrays of a greater number of stimuli. This result reflects a confound in the post-stimulus probe task which would be difficult to account for without modelling




Figure 6.5: A QPP displaying the predictions of the diffusion model configuration $\text{MOD-T}^s - \sqrt{v_\alpha}$ against the group average data of Experiment 3, styled with solid lines and filled squares, against the observed group average data, rendered in open circles and dashed lines. This model configuration is identical to that of Figure 6.3, except only a single drift variability value is used in the predictions of all memory array size conditions. SS = Memory array size.

the full response time distribution and indicates that, in the case of discrimination decisions made for memory arrays consisting of a single item, the effective decision time begins prior to the presentation of the post-stimulus probe. In such trials, the probe is unnecessary as a signal to commence the decision process and the observer need not wait for its presentation. This difference in the nonddecision time of single item memory arrays against memory arrays with multiple items was not found by Sewell et al. (in press), potentially reflecting the difficulty of this task¹.

Last, a slight majority of the individual observers, the group average, had model fits

¹ The observer CVH appears not to follow this pattern, having low nonddecision time estimates for all memory array size conditions; this appears to be the result of the smaller drift rate and larger boundary separation estimates.

that preferred additional flexibility in the estimation of trial-by-trial drift variability. Figure 6.5 shows the predictions of the model configuration $\text{MOD}-T^s-v_{\alpha}^-$, the second-best fitting model configuration to the group average data—identical to the best fitting model configuration, $\text{MOD}-\eta^s-T^s-v_{\alpha}^-$ (predictions against group average data shown in Figure 6.3), with the exception that only a single drift variability parameter was estimated. As is evident from Figure 6.5, the additional flexibility in drift variability allowed the diffusion model to capture bowing in the tails of the response time distributions, as well as moderating the extent of the sample size relationship on accuracy at the larger angular offset conditions in conditions where one or four items were presented in the memory array. The requirement for this level of flexibility could be due to the variability in the observed data, however, rather than a reflection of theoretically meaningful processes and constraints. Despite the large number of trials within each of the datasets, and although a systematic profile of change in response time and accuracy are clear over memory array size and angular offset conditions, variability in the data may be due in part to the difficulty of the task. Unlike change detection or orthogonal orientation discrimination, fine orientation discrimination requires observers to reach a decision outcome by comparison with an internal standard of comparison, which may be associated with some level of uncertainty. The problem of an uncertain internal standard for comparison may have been especially acute given that an oblique angle was selected as the stimulus standard to compare against. Consistent psychophysical evidence has indicated that the discrimination and detection for obliquely oriented stimuli is consistently poorer than that of horizontal or vertical oriented stimuli, known as the “oblique effect” (Appelle, 1972; Furmanski & Engel, 2000; Heeley, Buchanan-Smith, Cromwell, & Wright, 1997; Regan & Beverley, 1985). The difficulty associated with an oblique internal standard may have, therefore, contributed to the difficulty of the task and motivated the shift to a vertical standard for comparison in the next experiment.

6.2 Experiment 4

The aim of Experiment 4 was to extend the fine orientation procedure of Experiment 3 in characterising the effect of both the size of the memory array as well as the stimulus exposure duration on observer performance over different levels of stimulus discriminability. Rather than using angular offsets from an oblique stimulus as the standard of comparison, as in Experiment 3, observers were required to indicate whether the probed stimulus was rotated clockwise or counterclockwise from a vertical position. This was




Figure 6.6: Group average mean response time data as a function of offset magnitude from Experiment 4, displayed for each memory array size and stimulus exposure duration. Error bars represent one standard error of the mean.

intended to mitigate some of the difficulty of the previous experiment and reduce observer variability. Furthermore, although fewer angular offset conditions were used, the range of angular offsets examined was larger, to increase the stability of observer performance across conditions.

The full method and accuracy results for Experiment 4 can be found in §5.2.1 and §5.2.2, respectively.

6.2.1 Results

Figure 6.6 shows the mean response time for each decision type and memory array size condition, averaged across observers. The mean response times for each of the individual observers are shown at the end of this chapter, in Figures 6.8–6.12.

A mixed-effects linear model was conducted to provide a preliminary analysis of the response times (in milliseconds) at the mean level. The regression treated the memory array size, stimulus exposure duration (in milliseconds), and the offset magnitude of the target (in radians) as fixed effects, with the individual differences of the observer treated as a random effect on the intercept. The only effect to be found significant was the change in mean response times over different memory array size conditions, $\beta = 55.930$, $SE = 22.728$, $p = 0.014$, with increases in the number of items in the memory array leading to an increase in the mean response time.

6.2.2 Diffusion model results

The diffusion model applied to the data in Experiment 3 shown above added a Gaussian-shaped tuning function to relate target angular offset to drift rate, following the signal detection theory models presented in the last chapter. The weighting of (squared) drift rate by a tuning function was combined with the sample-size relation, also used in Chapter 4. In Experiment 4, the introduction of stimulus exposure duration provides another experimental manipulation of the amount of overall stimulus information available to the observer. The signal detection theory models for Experiment 4 presented in the last chapter demonstrated that a good fit to data was obtained by model configurations that related stimulus exposure duration to observer sensitivity directly, via a linear relation, rather than weighting squared sensitivity. The reason given was that, over the exposure durations examined, the use of orientation information with a high spatial frequency may lead to a slower rate of information accumulation or non-linearities when compared to representations based on orthogonal orientations. In the diffusion models applied to data obtained in Experiment 4, the drift rate was constrained using the relation

$$v_{i,m} = \frac{\sqrt{(1 - \phi(b\theta_i))} \cdot v_1}{\sqrt{m}} \gamma \tau$$

where the indices i and m run across different target angular offset conditions and memory array sizes, respectively; γ is the (estimated) growth of drift rate as stimulus exposure duration increases; ϕ is the tuning function, a modified Gaussian pdf described in Chapter 5; and, τ is the stimulus exposure duration in milliseconds.

An increase in the drift rate of the diffusion process with exposure duration indicates that, as the exposure duration of the memory array increases, the amount of stimulus information available to the decision process increases. The inclusion of an intrusion process, however, means that changes in the exposure duration of the memory array may also cause a change in the drift rate of the intrusion process. The model configurations for Experiment 3 included intrusion drift rates that were estimated freely across memory array size conditions, as well as fixed to zero (indicating delayed guessing) and set to the drift rate associated with the average angular offset of the distractors. In this experiment, the number of intrusion drift rate constraints that could be used in model configurations was expanded. Configurations using a “distractor average” parameterisation (denoted v_{α}^-) predict that the drift rate of the intrusion process should increase with stimulus exposure duration. This interpretation follows the logic that, as a greater quality of stimulus information is obtained for the target with a longer memory array

duration, so too does the quality of the distractor information increase. In that scenario, the efficiency of the post-stimulus probe in selecting only the relevant information is unchanged by the presentation characteristics of the memory array, and distractor information intruding into the decision process is fixed. This strong account is plausible, but unlikely: one may expect that, as the quality of the memory representation increases with a greater quantity of stimulus information obtained over a longer interval, so too does the ability to distinguish target and distractor information improve—or, at least, remain unchanged. A number of alternative parameterisations were tested, each listed in Table 6.3, including a number of alternatives that did not require the addition of a free parameter—such as the zero-drift intrusion process (i.e., delayed guessing) introduced in the diffusion model fit to the previous experiment, as well as parameterisations where the intrusion drift rate is constant for each of stimulus exposure duration conditions. The inclusion of different constant intrusion drift rate configurations do not imply that the intrusion process is tied to a specific level of information growth at a certain stimulus exposure duration, but represent different magnitudes of a constant negative drift rate for the intrusion process.

Given the small difference in goodness-of-fit for different intrusion drift rates parameterisations, it is unlikely that—without a directed experiment—that the data are sufficient to adequately discriminate between closely related parameterisations, although the ranking of the model fits should be informative with respect to the general magnitude of the intrusion drift rate and the functional form of the intrusion drift rate over stimulus exposure duration conditions (linearly increasing, linearly decreasing, constant, or freely estimated).

Finally, some model configurations allowed flexibility in the estimation of trial-to-trial variability. In fitting the diffusion model to Experiment 3, trial-to-trial drift variability was also free to vary between memory array size conditions. This was found to be a necessary component in fitting the last experiment, capturing additional variability in the leading edge and tails of the response time distribution. Part of this variability may have been due to variability in the observer's internal standard of comparison or affected by relative judgement, where the observer used the orientation of the distractor or distractors to determine the response category of the target. In the last experiment, this problem may have been heightened, given the use of oblique orientations as the stimulus standard, although the difficulty may not be entirely ameliorated by the use of a vertical orientation as the standard for comparison. As such, model configurations where trial-to-trial variability are freely estimated for each memory array size condition will be included.

Table 6.3: Different types of intrusion drift rate constraints that could be used in model configurations for Experiment 4.

Constraint name	Symbol	Description
Zero-drift	v_{α}°	Predicts a (memory array size dependent) proportion of trials will be delayed guessing.
Distractor average	v_{α}^{-}	Intrusion drift rate is set by the mean angle of distractors, with the absolute value of the intrusion drift rate increasing with stimulus exposure duration.
Minimum distractor drift	$v_{\alpha}^{\min}$	Intrusion drift rate is set by the mean angle of distractors at the shortest exposure duration (100 ms), and held constant across stimulus exposure duration conditions.
Average distractor drift	v_{α}^{avg}	Intrusion drift rate is set by the mean angle of distractors at the average exposure duration (150 ms), and held constant across stimulus exposure duration conditions.
Maximum distractor drift	$v_{\alpha}^{\max}$	Intrusion drift rate is set by the mean angle of distractors at the longest exposure duration (200 ms), and held constant across stimulus exposure duration conditions.
Linear drift	v_{α}^l	Intrusion drift rate varies over stimulus exposure duration conditions as a linear function.
Free estimation	v_{α}^t	The intrusion drift rate in each stimulus exposure duration condition is freely estimated.
Free constant	v_{α}	A single value is estimated for the intrusion drift rate and used as a constant across each stimulus exposure duration condition.

Table 6.4: The top five best fitting diffusion model configurations, in terms of BIC, for the group average data in Experiment 4.

#	Model	Free parameters	G ²	BIC
1	MOD- η^s - T^s - $v_{\alpha}^{\min}$	15	237.932	367.74
2	MOD- η^s - T^s - v_{α}°	15	240.409	370.217
3	MOD- η^s - T^s - $v_{\alpha}^{\max}$	15	241.806	371.614
4	MOD- η^s - T^s - v_{α}^t	18	232.016	387.785
5	MOD- η^s - T^l - v_{α}°	13	284.571	397.071

In total, 48 alternative models were fit to each observer, as well as to the group average data. The top ranking models for the group average data in terms of BIC are displayed in 6.4; the top ranking models for each of the individual observers are shown in Tables 6.11–6.15, displayed at the end of the chapter for clarity. For two out of the five individual observers, as well as for the group average, the best-fitting model used a constant non-zero intrusion drift rate determined using the linear increase of stimulus information with exposure duration: one observer was best fit by the drift rate of the mean angle of the distractors given 150 ms of stimulus exposure (the average stimulus exposure duration), with one observer and the group average data best fit by a model configuration that had the intrusion drift rate set by the mean angle of the distractor elements given 100 ms stimulus exposure (the minimum stimulus exposure duration). In three of the five observers, a model configuration with an intrusion drift rate of zero (delayed guessing) provided the best account of the data.

The best-fitting model for four out of the five individual observers, and for the group average data, required flexibility in the estimation of the nondecision time parameter, with a separate value estimated for each memory array size condition. Like the best-fitting models from the last experiment, the speed of observer responding in the trials where only one item was presented for retention was far quicker than would be predicted for a linear increase in nondecision time.

Last, the best fitting model configuration for a majority of the individual observers, as well as for the group average data, also required additional flexibility in the estimation of trial-to-trial variability. Again, the additional flexibility in the model allowed the correct estimation of quantiles in response time distributions both at the leading edge and in the tails across different memory array size conditions.

Figure 6.7 displays the prediction of the model $\text{MOD-}\nu^s\text{-T}^s\text{-}\nu_{\alpha}^{\min}$, the best-fitting model in BIC terms for the group average data, with the observed group average data. This model configuration requires 15 free parameters: four boundary separation parameters, four drift variability parameters, a single drift rate, a single bandwidth parameter, a single parameter for the linear growth of the drift rate with stimulus exposure duration (denoted γ), a single nondecision time range parameter, a single intrusion rate, and a single intrusion process time parameter (the intrusion drift rate is not estimated). Table 6.16, at the end of the chapter, shows the best-fitting parameters for this model configuration across observers.




Figure 6.7: A QPP displaying the predictions of the diffusion model configuration $MOD-\gamma^s - T^s - \gamma_{\alpha}^{min}$ against the group average data of Experiment 4, styled with solid lines and filled squares, against the observed group average data, rendered in open circles and dashed lines. SS = Memory array size; ED = Exposure duration of memory array (in milliseconds).

6.2.3 Discussion

Like the close fits between the model predictions and the observed data for the last experiment, the overall pattern of observed data in Experiment 4 is very well accounted for by the modified diffusion model, particularly given the large number of data points (both correct and incorrect response times and response proportions in 48 different conditions) given the small number of parameters (the best fitting model had 15 parameters). The parsimony of this model suggests that influence of the modifications to the diffusion model map on to the experimental manipulations well.

In addition to examining the ability of sample-size constraint, the tuned channel model, and a linear increase in sensitivity to predict performance in a large fine orientation discrimination task, the diffusion model fits to Experiment 4 also examined the overall changes in the intrusion process across stimulus exposure duration conditions. Rather than finding an increasing linear relationship between intrusion drift rate and stimulus exposure duration, as was found for the drift rate of the main decision process, the best-fitting models for all observers and for the group average kept the intrusion drift rate fixed at a specified value for each stimulus exposure duration: either at zero or a small negative value. In examining this functional form, a flat intrusion drift rate across time is somewhat problematic. If it is assumed that the intrusion process represents inefficiency in the post-stimulus probe to localise retrieval from stored memory representations (that is, to filter out distractor information), then it would be natural to expect either an increasing or decreasing intrusion drift rate with an increasing stimulus duration. If the efficiency of the retrieval process to extract target information increased over time, it would be expected that this would lead to better localisation of the decision process to target, better filtering of non-target information, and a decrease in the drift rate of the intrusion process over time. If the efficiency of the probe was constant, and did not change with stimulus exposure duration, then it would be expected that quality of distractor representations would increase with time, leading to greater negative drift rates for the intrusion process as distractor representations accumulated a greater quality of information.

Given the small changes in the goodness-of-fit statistic for different intrusion drift rate values and parameterisations, it is perhaps unwise to draw strong conclusions from the current model fits without recourse to a directed experiment to test any claims made. A future experiment, where the efficiency of retrieval could be examined directly through further control of the distracting information, would be an informative complement to the current study and provide a natural avenue for the application of the extended diffu-

sion model.

One additional complicating factor for the stable estimation of intrusion process properties, and inferences relating to the functional form of these properties, is the inclusion of trial-to-trial variability. This will be discussed, along with the general findings of this chapter, in the next section.

6.3 General discussion

The aim of Experiments 3 and 4 was to examine the effect of different levels of stimulus discriminability upon observer performance, and to investigate the relationship between stimulus discriminability and other information constraints upon memory due to memory array size manipulations and memory array exposure duration manipulations.

The analyses shown in this chapter extended the signal detection theory models presented in the previous chapter into the response time domain. Both the sample-size relation and the tuned channel bandwidth model appear to model the change in observer performance over manipulations in memory array size and stimulus discriminability well via a specific influence on the rate of information available to the decision process, the drift rate. This also applies to the effect of stimulus exposure duration which, in the best-fitting model configurations, was well captured by a linear change in drift rate. The success of the model in fitting joint response time and response proportion data, as seen with the group average data in Figure 6.3 for Experiment 3 and Figure 6.7 for Experiment 4, further strengthens the claims of the previous chapter that the tuned channel model represents a fundamental constraint on the memory system itself, changing the quality of information available to a decision process, rather than confusability in the decision stage. This point is particularly salient when comparing the current results to the Bays, Gorgoraptis, Wee, Marshall, and Husain (2011), where changing variability in the response distribution of a continuous report task was argued to reflect a change of the “precision” of a memory representation, which might be interpreted as a change to the structure of the representation itself. This change in precision, however, appears from the current data to be simply the increase in the stimulus signal against background noise. In fitting the sensitivity data in previous chapter and the joint response time and response proportion data in the current, the parameterisation of the bandwidth function of the tuned channel constraint remained the same over both memory array size and exposure duration manipulations, suggesting that the structure of the memory representation is stable but the amount of effective stimulus information available to the

decision process increases.

Further, there are no indications that these different information constraints on the representation held within visual short-term memory interact: their influence on observer performance is independent of one another. The neural interpretation of the sample-size limit, provided by Smith (2015), is that the stimulus discriminability when coded by a population of fixed size comprising Poisson neurons under normalisation will asymptotically approach $d'_1/\sqrt{m}$: the discriminability of a single stimulus, d'_1 divided by the square root of the number of items in the display, m . That is, as discussed in Chapter 2, the sample-size relation represents the equal division of fixed quantity of resources between maintained representations—a fundamental constraint based upon the size of memory. The tuned channel constraint represents the output of a population of orientation-selective detectors with fixed response characteristics for given stimulus orientations. The signal-to-noise ratio of this population of orientation-selective detectors appears to remain constant for a given stimulus orientation regardless of other experimental factors, such as the memory array size or the exposure duration of the memory array. The locus of this tuned channel constraint is not clear from these data, however: the constraint may exist prior to the entrance of information into the memory system, perhaps due to the fixed resolution for orientation information imposed by the architecture of orientation-selective neurons in the primary visual cortex, or due to the memory system itself. The Neural Theory of Visual Attention (NTVA) of Bundesen, Habekost, and Kyllingsbæk (2005, 2011) suggests one plausible neural instantiation of visual short-term memory is the positive feedback of activation in the primary visual cortex from projections originating at a memory control system in the thalamic reticular nucleus. That is, the substrate of visual short-term memory representation is the same as the substrate of visual coding more generally (what NTVA terms “unselective” visual processing). This account would be compatible with the current findings in suggesting that the orientation discrimination limits seen in the current experiments arise from the limits on receptor discrimination at the level of primary visual cortex. Further specification of the feature representation within visual short-term memory may be achieved through studies examining other feature types (such as colour), and their interaction.

As mentioned in the previous section, open questions remain about the nature of the intrusion process used in the current response time modelling. One complicating factor in considering questions about the changes in intrusion drift rate with memory array exposure duration is the inclusion of trial-to-trial variability for the drift rate of the main decision process in the best-fitting model configurations. As larger values of η can mean that the main process also includes the non-zero probability of a negative drift rate, part

of the function of an intrusion process may be served by the main process. As intrusion drift rate and trial-to-trial variability may be correlated, easy interpretation of the functional form or magnitude of the drift rate of the intrusion process might be very difficult, particularly without a directed experiment. An ideal approach would be estimate both the intrusion process and the trial-to-trial variability in the main process from a model where all stimulus locations are represented within the model. In the current models, the interaction between the target information and distractor information maintained within memory are not explicitly handled (beyond the interpretation placed on the intrusion process), but models like that of Smith and Sewell (2013) allow some purchase on this question by modelling the dynamics of all extant representations maintained in memory and can be used. With additional assumptions about the connection between the strength of the individual representations, and the relationship between multiple representations and the final decision process, a model like that presented by Smith and Sewell would be able to impose constraints upon both trial-to-trial variability as well as the intrusion process simultaneously.

The next chapter—the final chapter of this thesis—will cover some of these future modelling directions, as well as providing a brief overview of the findings in this and previous chapters, and some of the proposed implications of these findings.

Supplementary tables and figures

Best fitting diffusion model fits for individual observers in Experiment 3. Tables 6.5–6.9

Best fitting parameters for the model configuration MOD– $\eta^s - T^s - v_{\alpha}^-$ in Experiment 3. Table 6.10

Mean response times for individual observers in Experiment 4. Figures 6.8–6.12

Best fitting diffusion model fits for individual observers in Experiment 4. Tables 6.11–6.15

Best fitting parameters for the model configuration MOD– $v^s - T^s - v_{\alpha}^{\min}$ in Experiment 4. Table 6.16

Table 6.5: The top five best fitting diffusion model configurations, in terms of BIC, for observer AB in Experiment 3.

#	Model	Free parameters	G^2	BIC
1	MOD- $T^s-v_{\alpha}^-$	11	514.463	605.197
2	MOD- $T^s-v_{\alpha}^o$	11	516.691	607.424
3	MOD- T^s-v_{α}	12	515.545	614.527
4	MOD- $\eta^s-T^s-v_{\alpha}^-$	14	509.367	624.846
5	MOD- $\eta^s-T^s-v_{\alpha}^o$	14	510.967	626.447

Table 6.6: The top five best fitting diffusion model configurations, in terms of BIC, for observer CVH in Experiment 4.

#	Model	Free parameters	G^2	BIC
1	MOD- $T^l-v_{\alpha}^o$	9	507.444	581.657
2	MOD- T^l-v_{α}	10	507.251	589.71
3	MOD- $T^s-v_{\alpha}^o$	11	502.61	593.315
4	MOD- T^s-v_{α}	12	500.321	599.272
5	MOD- $\eta^s-T^l-v_{\alpha}^o$	12	501.088	600.039

Table 6.7: The top five best fitting diffusion model configurations, in terms of BIC, for observer SIS in Experiment 4.

#	Model	Free parameters	G^2	BIC
1	MOD- $\eta^s-T^s-v_{\alpha}^o$	14	490.062	605.578
2	MOD- $T^s-v_{\alpha}^o$	11	515.473	606.236
3	MOD- $T^s-v_{\alpha}^-$	11	519.396	610.158
4	MOD- T^s-v_{α}	12	525.429	624.443
5	MOD- $\eta^s-T^s-v_{\alpha}$	15	502.077	625.844

Table 6.8: The top five best fitting diffusion model configurations, in terms of BIC, for observer SL in Experiment 4.

#	Model	Free parameters	G^2	BIC
1	MOD- $\eta^s - T^s - v_\alpha^o$	14	454.457	569.896
2	MOD- $\eta^s - T^s - v_\alpha$	15	447.63	571.314
3	MOD- $\eta^s - T^s - v_\alpha^-$	14	465.51	580.949
4	MOD- $\eta^s - T^l - v_\alpha^o$	12	599.069	698.017
5	MOD- $\eta^s - T^l - v_\alpha^-$	12	602.675	701.623

Table 6.9: The top five best fitting diffusion model configurations, in terms of BIC, for observer XL in Experiment 4.

#	Model	Free parameters	G^2	BIC
1	MOD- $\eta^s - T^s - v_\alpha^-$	14	434.628	550.085
2	MOD- $\eta^s - T^l - v_\alpha^-$	12	473.841	572.805
3	MOD- $\eta^s - T^s - v_\alpha^o$	14	464.239	579.696
4	MOD- $\eta^s - T^s - v_\alpha$	15	465.222	588.926
5	MOD- $\eta^s - T^l - v_\alpha^o$	12	495.454	594.417

Table 6.10: The best fitting parameter estimates for the model configuration $MOD-\eta^s-T^s-v_\alpha$ across all participants in Experiment 3. When multiple independent values are estimated for a single parameter, the values are separated by commas for increasing memory array sizes (1, 2, 3, and 4 items respectively).

Observer	v	b	η	α	T	s_T	α	T_α	v_α
Avg	0.326	0.054	0.23, 0.129, 0.065, 0.058	0.135	0.021, 0.096, 0.116, 0.131	0.316	0.443	0.17	-0.164
AB	0.567	0.127	0.262, 0.312, 0.268, 0.225	0.145	0.013, 0.154, 0.215, 0.207	0.357	0.001	0.007	-0.12
CVH	0.267	0.114	0.139, 0.119, 0.096, 0.084	0.198	0, 0.036, 0, 0.024	0.244	0.467	0.27	-0.125
SIS	0.356	0.08	0.228, 0.204, 0.203, 0.209	0.115	0, 0.176, 0.222, 0.248	0.38	0.03	0.123	-0.144
SL	0.366	0.055	0.271, 0.048, 0, 0	0.123	0.006, 0.103, 0.132, 0.155	0.183	0.497	0.146	-0.163
XL	0.652	0.037	0.577, 0.276, 0.122, 0.08	0.105	0.08, 0.109, 0.14, 0.135	0.334	0.33	0.095	-0.183




Figure 6.8: The mean response time data of observer HA as a function of offset magnitude from Experiment 4, across different memory array sizes and stimulus exposure duration. Error bars represent one standard error of the mean.




Figure 6.9: The mean response time data of observer LA as a function of offset magnitude from Experiment 4, across different memory array sizes and stimulus exposure duration. Error bars represent one standard error of the mean.




Figure 6.10: The mean response time data of observer ME as a function of offset magnitude from Experiment 4, across different memory array sizes and stimulus exposure duration. Error bars represent one standard error of the mean.




Figure 6.11: The mean response time data of observer SN as a function of offset magnitude from Experiment 4, across different memory array sizes and stimulus exposure duration. Error bars represent one standard error of the mean.




Figure 6.12: The mean response time data of observer TS as a function of offset magnitude from Experiment 4, across different memory array sizes and stimulus exposure duration. Error bars represent one standard error of the mean.

Table 6.11: The top five best fitting diffusion model configurations, in terms of BIC, for observer HA in Experiment 4.

#	Model	Free parameters	G^2	BIC
1	MOD- η^s - T^s - $v_{\alpha}^{\min}$	15	669.454	799.157
2	MOD- η^s - T^s - v_{α}^o	15	675.132	804.834
3	MOD- η^s - T^s - $v_{\alpha}^{\max}$	15	679.436	809.138
4	MOD- η^s - T^s - v_{α}	16	673.301	811.65
5	MOD- η^s - T^s - v_{α}^{avg}	15	713.482	843.184

Table 6.12: The top five best fitting diffusion model configurations, in terms of BIC, for observer LA in Experiment 4.

#	Model	Free parameters	G^2	BIC
1	MOD- η^s - T^s - v_{α}^o	15	745.406	875.011
2	MOD- η^s - T^l - v_{α}^o	13	764.436	876.76
3	MOD- η^s - T^s - v_{α}^l	16	747.637	885.882
4	MOD- η^s - T^s - v_{α}	16	748.761	887.005
5	MOD- η^s - T^l - v_{α}	14	768.834	889.798

Table 6.13: The top five best fitting diffusion model configurations, in terms of BIC, for observer ME in Experiment 4.

#	Model	Free parameters	G^2	BIC
1	MOD- η^s - T^s - v_{α}^o	15	1079.826	1209.568
2	MOD- η^s - T^s - v_{α}^i	18	1054.22	1209.91
3	MOD- η^s - T^s - $v_{\alpha}^{\min}$	15	1090.758	1220.5
4	MOD- η^s - T^s - v_{α}^{avg}	15	1095.44	1225.182
5	MOD- η^s - T^s - $v_{\alpha}^{\max}$	15	1114.317	1244.058

Table 6.14: The top five best fitting diffusion model configurations, in terms of BIC, for observer SN in Experiment 4.

#	Model	Free parameters	G^2	BIC
1	MOD- η^s - T^s - v_α^o	15	722.469	851.868
2	MOD- η^s - T^l - v_α^o	13	741.735	853.88
3	MOD- η^s - T^l - $v_\alpha^{\min}$	13	747.645	859.79
4	MOD- η^s - T^s - v_α	16	723.158	861.183
5	MOD- η^s - T^l - v_α	14	740.93	861.702

Table 6.15: The top five best fitting diffusion model configurations, in terms of BIC, for observer TS in Experiment 4.

#	Model	Free parameters	G^2	BIC
1	MOD- η^s - T^s - v_α^{avg}	15	771.13	901.695
2	MOD- η^s - T^s - v_α^-	15	781.81	912.375
3	MOD- η^s - T^s - v_α^{max}	15	814.658	945.223
4	MOD- T^s - $v_\alpha^{\min}$	12	900.336	1004.788
5	MOD- T^s - v_α^{avg}	12	901.567	1006.019

Table 6.16: The best fitting parameter estimates for the model configuration $\text{MOD}-\eta^s-T^s-\gamma_{\alpha}^{\text{min}}$ across all participants in Experiment 4. When multiple independent values are estimated for a single parameter, the values are separated by commas for increasing memory array sizes (1, 2, 3, and 4 items respectively).

Observer	ν	b	γ	η	α	T	S_T	α	T_{α}	ν_{α}
Avg	0.74	0.166	0.002	0.2, 0.095, 0.057, 0.004	0.152	0, 0.056, 0.081, 0.1	0.217	0.516	0.207	-0.072
HA	1.145	0.097	0.001	0.204, 0.112, 0.077, 0	0.146	0, 0.05, 0.092, 0.094	0.205	0.586	0.298	-0.043
LA	1.2	0.171	0.001	0.156, 0.087, 0.013, 0.007	0.164	0.025, 0.07, 0.046, 0.071	0.159	0.823	0.228	-0.04
ME	0.438	1.628	0.015	0.227, 0.186, 0.155, 0.104	0.196	0.004, 0.096, 0.107, 0.125	0.193	0.144	0.231	-0.187
SN	0.469	0.064	0.003	0.276, 0.121, 0.074, 0.076	0.126	0, 0.002, 0.052, 0.085	0.287	0.312	0.193	-0.128
TS	0.059	0.074	0.024	0.131, 0.082, 0.119, 0.155	0.162	0.024, 0.143, 0.234, 0.269	0.022	0.028	0.205	-0.912

Chapter 7

A summary and discussion of the current work

The aim of this thesis was to examine observer performance in a set of visual short-term memory tasks, designed to be consistent in their methodology and controlled in terms of the complexity of decision process required. The consistency and control in the design of the experiments allowed a level of continuity in the development of a theory for describing observer performance. Instead of following the focus of the broader visual short-term memory field on questions of storage capacity specifically, I have attempted to build a quantitative model of observer performance under different visual short-term memory incrementally for different visual short-term memory manipulations, by first considering the relationship between an observed response in a memory task and the underlying representation; that is, I attempted to proceed in the development of a model by considering the decision process directly. In this way, the work of this thesis distinguished itself from investigations in the broader visual short-term memory literature both in methodological approach (using small-N, near-threshold designs taken from visual psychophysics) and in theoretical approach (using response times as a central component in testing predictions).

The thesis was divided into two main parts, each comprising two experiments: the first section examined the relationship between two established psychophysical paradigms to examine the role of decision-making in visual short-term memory tasks; the second section developed upon the models provided in the first section to examine observer performance over different levels of stimulus discriminability. Here I summarise and review the major conclusions and themes from each of these two sections in turn.

7.1 The effect of decision type on observer performance

The aim of the first section was to distinguish the role of the task-specific decision process from central constraints on memory by comparing observer performance across different experimental paradigms equated to be identical in stimulus presentation and retention. In the two experiments contained within this section, two psychophysical paradigms used in the visual short-term memory literature were compared: (orthogonal) orientation discrimination and change detection. Each differed only in the presentation of the probe and the type of decision required: in the orientation discrimination condition, a post-stimulus report cue was presented at a location corresponding to an item presented in the memory array and observers were to indicate the orientation of the probed item; in the change detection condition, observers were presented with a report cue and a probe array with stimuli corresponding to those presented in the memory array—with the target item either identical or changed in orientation, and the non-target items all identical in orientation—and observers were to indicate whether they detected a change in the probed location.

Both experiments, as well as the experiments presented in the second section of the thesis, used oriented near-threshold Gabor patches as stimuli. Although these stimuli, and the use of stimuli near the detection threshold generally, are somewhat uncommon in the visual short-term memory literature, these stimuli are common in low-level psychophysical tasks, owing to their close correspondence to the profile of receptive fields in early visual cortex (Daugman, 1984; Graham, 1989) and overall control over constraints on performance afforded by careful manipulation of timing and display contrast. The use of near-threshold stimuli, as well as the interleaved dynamic noise, was used to control the rate of information that was available for the memory system to accumulate stimulus information and form a durable representation. In addition to this control on the rate of representation formation, the complexity of the decision in all tasks was controlled through the use of a post-stimulus probe (Downing, 1988). By probing only a single stimulus location, even in the change detection tasks, any variability in observer performance due to the increasing number of independent decisions to be made over manipulations of the array size was controlled for. These controls on stimulus generation and presentation followed those used by Sewell et al. (2014) to identify the sample-size relationship between observer performance and memory array size. This relationship seems to be a fundamental information limit upon memory that had largely escaped comment in the visual short-term memory literature. The careful control of presentation and decision complexity allowed, in this thesis, a detailed consideration of both the constraints

7.1 The effect of decision type on observer performance

upon memory retention and processing, as well as the characteristics of the tasks used to assess memory.

A cursory examination of the accuracy data across different memory array sizes in Experiment 1 showed that, in addition to a systematic decrease in accuracy as the size of the memory array increased, there was a consistent decrement in performance when observers were required to make a change detection decision when compared to orientation discrimination trials. This result was also supported by the outcome of the calibration procedure in Experiment 2, where observer accuracy was equated through the manipulation of stimulus contrast between the different decision types. In Experiment 2, there was a consistent increase in accuracy with an increasing exposure duration of the memory array and the display contrast of the memory array in change detection trials was much higher than that of the memory array in orientation discrimination trials in order to equate performance. This indicates that the amount of information in change detection tasks must be substantially increased to equal performance in the corresponding orientation discrimination trials or, equivalently, there is a decrease in the amount of effective information available to a decision process in change detection when compared to orientation discrimination under the same presentation conditions.

One strong conclusion to draw from this finding is that all visual information used in subsequent decision-making must be encoded to visual short-term memory prior to further processing, even if that information is still visible. The probe array in both of the experiments presented in the first part of the thesis was displayed at high contrast and present until a response was entered. As it is unlikely that any degradation in performance could be due to sensory coding factors in the encoding of the probe array—and, in the first experiment, the presentation of the memory array was identical—it follows that the observed decrement in performance is due to either to a constraint upon the decision process or the effect of the probe array on memory. To characterise this accuracy relationship systematically, these accuracy data were subsequently examined using a set of models taken from signal detection theory.

A large set of models based on signal detection theory were constructed, differing in the constraints placed upon their parameterisation and the type of decision rule applied, and fit to the observed data. The model configurations that ranked highest in terms of BIC were largely consistent across observers, with strong support for the sample-size relationship across observers and across tasks. Critically, the relationship between the different decision tasks was well captured by models that included an additional item in the denominator of the sample-size relationship in the change detection task to account for the additional load of the probe array. (The best fitting model in terms of BIC was,

as mentioned in the chapter, a model which placed no increment in the denominator of the sample-size relationship for change detection tasks, but this was likely due to the fact that the penalty applied to the goodness-of-fit statistic was too great given the underdispersion of the group data; as such, I placed greater emphasis on the second-best fit, which was substantially better fitting in unpenalised goodness-of-fit terms.)

The process of encoding the probe array and the comparison of the probe array to the memory array was also reflected in the response time data, with the mean response time pattern across both exposure duration and memory array size conditions apparently identical in between orientation discrimination and change detection trials, with the exception that mean response times in the change detection condition were longer (on the order of about 150–200 ms). This qualitative relationship was further examined using both a form of deconvolution with theoretically plausible response time component distributions, and with the application of the diffusion model of Ratcliff (1978). The best fitting deconvolution model configurations in the first experiment indicated that the relationship between the orientation discrimination and change detection conditions was best accounted for by the convolution of the orientation discrimination response time distribution within each memory array size condition with an exponential distribution with a delay of roughly 100–150 ms and decay rates (dispersion parameters) of between approximately 50–100 ms, for observers who were well fit by the exponential distribution (the remaining two observers were well fit by a model with just a delay parameter, with no variability in the convolving distribution). Likewise, the application of a convolutional analysis to the second experiment in the first part provided similar results: the data of three of the observers were best fit (in BIC terms) by simply shifting the orientation discrimination response time distribution backwards in time to match the delay of the change detection distribution; the data of the remaining two observers were best fit by an exponential distribution with a decay rate in the range of 50–150 ms. (In one observer, convolution with an exponential distribution fit best, and the offset parameter was estimated at a little above zero, with all of the delay in the change detection response time distribution handled by additional dispersion in the convolving distribution.)

Taken together with the accuracy data, the results from these analyses suggested that decision-making in the change detection task can be thought of as largely the same as the decision-making process within the orientation discrimination task, with the addition of a process for encoding the probe array and comparing the probe array element to maintained representations within memory. The time taken to complete this encoding and comparison process would account for the additional time taken by responses in the change detection trials compared to the orientation discrimination trials and if, as

mentioned, the memory store itself is the substrate for probe array representation during the comparison process, then the sample-size information constraint upon the memory system would provide an explanation for the uniform decrement in accuracy observed in change detection conditions when compared to their corresponding orientation discrimination conditions.

These dual signatures were further characterised for each experiment with the joint analysis of both response time distributions and accuracy data using a modified version of the diffusion model (Ratcliff, 1978). The critical modification made to the diffusion model was to include an “intrusion” process, a separate diffusion process with an additional non-decision time component and an additional drift rate parameter, which provided a way to capture the asymmetry in the leading edge (visible in the QPP figures as the 10th percentile of the response time distribution) between correct and incorrect responses. This asymmetry in the leading edge was previously reported by Sewell et al. (in press), who demonstrated that a substantially improved fit could be obtained by the inclusion of a “delayed guessing” process: a zero-drift diffusion process with an additional delay mixed with the main diffusion process at a constant mixture probability ¹. Like the fitting process for the signal detection theory and convolution models, a set of alternative models were generated by imposing different constraints between the experimental conditions and the estimated parameters, these model configurations were then ranked by BIC. In fitting the first experiment, a single model configuration consistently ranked as the best fitting model by BIC: a diffusion model with the drift rate constrained across memory array size conditions by the sample-size relationship, with additional item added to the denominator of the sample-size relationship to relate performance in the orientation discrimination trials to performance in the change detection trials, and the mean non-decision time varying between memory array size conditions and between decision types.

Both the qualitative pattern of accuracy and of the response time distribution was captured in this model across both memory array size conditions and decision types with a single set of fifteen parameters for each observer. The sample-size constraint upon the drift rates supports previous findings—both in modelling sensitivity (Sewell et al., 2014) and in diffusion modelling through a constraint on drift rate (Sewell et al., in press)—as well as the signal detection theory models reported in the first part of this thesis. The consistent requirement of the best fitting model configurations to estimate the non-decision

¹The current modelling work was completed prior to the inclusion of a “delayed guess” process in the paper by Sewell et al. (in press), although it followed work by Donkin et al. (2013) who modelled memory retrieval failures using a zero evidence accumulation process.

time values freely across memory array sizes as well as across trials of different decision types was also expected given both the previous findings in the literature and the current convolution analysis: Sewell, Lilburn, and Smith (in press) reported that non-decision time increased as a function of memory array size, potentially as a result of an increase in the time taken to access memory representations with a growing memory load; the results of the convolution analysis are consistent with the additional time required to store and compare the probe array to an existing memory representation.

In fitting the second experiment with the diffusion model, the difference between decision type conditions in the presentation of the memory array—coupled with an extended training period for observers—meant that no single model configuration was shared by all observers. Similar to the best fitting model configuration obtained for the first experiment, all of the highest ranking model configurations for the second experiment required some flexibility in the estimation of drift rates—either allowing drift rate to increase as a function of the exposure duration of the memory array or allowing the drift rate to be estimated for each exposure duration condition as well as both of the decision types. Additionally, in fitting data for the majority of the observers, as well as the group average data, the best fitting model only required mean non-decision time to differ between trials of different decision types, reflecting the additional time required for change detection. (In the remaining two observers, mean non-decision time was estimated across both exposure duration conditions and decision type conditions.)

Beyond relaxing the constraints upon drift rate and mean non-decision time to fit the second experiment, the best fitting models also required additional flexibility in either allowing the boundary separation or intrusion rate to vary across decision type conditions. In both cases, as discussed in Chapter 4, this may reflect the difficulty of the task, leading to strategic behaviour by observers learned over the long training period. A difference in boundary separation would indicate that individual observers require different amounts of stimulus information to initiate a response in orientation discrimination trials compared to change detection trials. A difference in the intrusion rate is more difficult to interpret without a clear model of the mapping between the intrusion process and memory retrieval, potentially indicating a difference in the efficiency of the post-stimulus report cue in localising the retrieval of information from memory (discussed below, see §7.4.2).

The analysis of accuracy data and response time data from both experiments, by means of signal detection theory models and convolution analyses in the first instance, and through the application of the diffusion model to analyse response proportions and response times jointly, provide a clear relationship between the two different types of vi-

sual short-term memory tasks examined. Changes in observer performance across the first experiment, in which memory array size was manipulated between trials, were well accounted for by the sample-size relation. In both experiments, the addition of an intrusion process into the main diffusion process was able to capture the asymmetry between correct and incorrect responses in the leading edge of the response time distribution. Beyond these findings, models of sensitivity based on signal detection theory and the diffusion model configurations provided the basis for analysis in the second section of the thesis.

7.2 The effect of stimulus discriminability on observer performance

The aim of the second section of the thesis was to develop the orientation discrimination experimental procedure presented in the first section of the thesis, as well as the signal detection theory and response time models in the previous chapters, to examine the relationship between the storage of feature information within memory (specifically, orientation information) and observer performance. The two experiments presented in the latter half of this thesis required discrimination of a post-stimulus probed near-threshold stimulus within a larger memory array, based upon a known standard (either an oblique stimulus in Experiment 3 or a vertical stimulus in Experiment 4). The angular offsets of the target stimulus were selected to provide a range of different observer performance levels, from near chance to a level of performance approximating orthogonal orientation discrimination, in addition to the performance constraint imposed by the display contrast of the stimulus. In Experiment 3, the memory array size was manipulated with a large range of target angular offset manipulations were used; in Experiment 4, a small range of target angular offset conditions were examined, but the exposure duration of the memory array was manipulated in addition to the memory array size.

A clear systematic and non-linear relationship between target angular offset and observer accuracy was evident upon examination of the accuracy data from Experiment 3, as well as a relationship between memory array size and accuracy. In order to account for the relationship of target angular offset to observer accuracy, a modification was made to the 2AFC signal detection theory model presented in the first section of the thesis to relate observer sensitivity to target angular offset through the assumption that discrimination performance changes around the zero-discriminability orientation (diagonal orientations in Experiment 3) as a function of the response character-

istics of a orientation-selective detector centred on that orientation, with monotonically increasing performance for the discrimination of orientations further away from the zero-discriminability orientation. A large number of model configurations, differing in the functional form of the tuning function and the constraints upon parameter estimation, were fit to the data of each individual observer, as well as the group average data, and compared in terms of BIC. One model configuration provided the best fit for each observer (apart from a single observer, XL) as well as for the group average data: a model where observer sensitivity was determined both by the sample-size relation and by a Gaussian-shaped tuning function weighting squared sensitivity independently. The bandwidth of the tuning function—the selectivity of the detector centred on the zero-discriminability orientation—was unchanged across the different memory array size conditions, indicating that the response characteristics of the tuned channel appear to be unrelated to memory load.

Similar non-linear relationships between target angular offset, memory array size, and observer accuracy were seen in Experiment 4, with the addition of increasing observer accuracy seen for an increasing exposure duration of the memory array across all conditions. Like the extended signal detection models applied to Experiment 3, a large number of model configurations were individually fit to each observer's data as well as to the group average data. There was slightly more inconsistency in the best fitting model for each observer, however the results largely agreed with the model results from Experiment 3, with the Gaussian-shaped weighting on squared sensitivity combined with the sample-size constraint imposed between memory array size conditions performing best for the group average data, and within the top overall models for all but one observer. The effect of stimulus exposure duration on observer sensitivity was best captured by a linear growth in observer sensitivity with increasing exposure duration. This finding was somewhat contrary to a strong mechanistic interpretation of the sample-size relation, which would predict a constant rate of increase in squared sensitivity rather than sensitivity directly. This finding may reflect the difference between the temporal response of the high spatial frequency detectors required in fine orientation discrimination, compared to the apparent linear increase in squared sensitivity for reported by Sewell, Lilburn, and Smith (2014) using an orthogonal orientation discrimination paradigm.

A large and systematic effect of increasing mean response times with increasing memory array sizes was apparent in Experiment 3, with a smaller (but significant) effect of increasing mean response times with smaller target differences from the zero-discriminability orientation. Only the effect of memory array size was apparent on the mean response times in Experiment 4, with no significant effect of target angular offset or stimulus du-

ration found using a linear mixed-effects model.

The less prominent effects of the experimental manipulations in Experiments 3 and 4 on mean response times further warranted investigation of the entire response time distribution, as well as the response proportions, using the diffusion model. The diffusion modelling for both experiments followed the modelling in the first part of the thesis, with the inclusion of an intrusion process. The constraints between experiment manipulations imposed upon sensitivity in the signal detection theory models were applied to the drift rate of the diffusion models. Rather than fitting each of the different tuned channel configurations using the more computationally intensive diffusion model, the Gaussian-shaped tuning function centred on the zero-discriminability orientation and weighting the squared drift rate was applied to all model configurations. In both Experiments 3 and 4, the sample-size relation and the tuned channel models were applied directly to the squared drift rate, independently of one another; in Experiment 4, drift rate was constrained to be a linear function of exposure duration.

An examination of the fits of the diffusion model for both the experiments in the second part of the thesis demonstrated a close overall fit to a large number of experimental conditions with a relatively small number of parameters. The accuracy and response time quantiles were well accounted for by the sample-size constraint and the tuning function constraint across memory array size and target angular offset manipulations in both Experiments 3 and 4, and the stimulus exposure duration manipulation was well accounted for by the linear constraint on drift in Experiment 4. In addition to the constraints placed upon the drift rate, different parameterisations of the drift rate for the intrusion process as well as the trial-to-trial variability in the drift rate of the main process were also examined. The best fitting model configuration for a slight majority of observers in Experiment 3, as well as for the group average data, fixed the drift rate of the intrusion process to be the drift rate associated with the average orientation of the distractor items. For Experiment 4, a constant drift rate for the intrusion process provided the best overall fit across stimulus exposure duration conditions, contrary to easy interpretation. In both experiments, the fit of the model configurations was also sufficiently improved by additional flexibility in the estimation of trial-to-trial variability to offset the penalty imposed by BIC for the additional parameters.

The diffusion models, as well as the signal detection theory models, provided a precise account of the information constraints upon memory representations in both of experiments of the second half of the thesis. In every experiment in the thesis, the performance decrement seen as the size of the memory array increased was well accounted for by the sample-size model, demonstrated by Sewell et al. (2014). The changes in per-

formance seen in the fine orientation discrimination tasks followed the predictions of a tuned channel model. Notably, the bandwidth and shape of the tuning function used to relate measures of stimulus information, either squared sensitivity or squared drift rate, to the angular offset of the target from the zero-discriminability orientation was not affected by either the memory array size or exposure duration manipulations, indicating that this constraint is not due to the fixed information limit reflected in the sample-size relationship, but may be an architectural constraint on the fidelity with which orientation may be represented. These two constraints were also independent of the effect of stimulus exposure duration.

Simply, taken together, these results suggest that stimulus exposure duration affects the amount of overall stimulus information available to be encoded in to memory, the tuned channel relation represents a fundamental constraint in the resolution of orientation information that can be encoded in information, and the sample-size relation represents the division of limited resources between the information that can be maintained in memory simultaneously. These results, as well as the results from the first part of the thesis, have implications for the broader memory literature, both in terms of methodological considerations in the design of memory experiments and in current theories of memory structure. Some of the methodological implications will be examined in the next section, followed by some of the larger theoretical implications.

7.3 Methodological implications

One important theme that I hope to have highlighted in the motivation for and discussion of the experiments within this thesis is the intimate relationship between the design of visual short-term memory experiments and the specificity and the structure of the quantitative models that can be applied to observed data. The principal experimental paradigms used to examine visual short-term memory, and visual short-term memory storage capacity in particular, have been the change detection and continuous report paradigms, almost exclusively using suprathreshold stimuli. Although the results from both paradigms allow some constraints to be placed on models of visual short-term memory and, therefore, upon the inferences drawn from visual short-term memory work, more detailed modelling is enabled when the correspondence between the properties of the stimulus, the task demands, and the observed responses is well understood—by, for instance, controlling the number of stimulus alternatives or through control of display contrast and presentation length of the memory array.

7.3.1 Change detection

As reviewed in Chapter 2, change detection procedures have provided some of the most important results in the visual short-term memory literature, such as the results reported given by Phillips (1974), and the work of Luck, Vogel, and colleagues (Luck & Vogel, 1997; Vogel & Machizawa, 2004; Vogel et al., 2001). These findings have been used to argue for an account of visual short-term memory where the storage capacity of the system is limited to a small, fixed number of objects, usually estimated around four or five items. Although each experiment in this thesis employed memory arrays below this limit, memory capacity constraints were still observed, both in the form of the sample-size limit and as a decrement in accuracy observed in change detection over orientation discrimination.

In comparing the results of an orientation discrimination task and a change detection task with the stimulus presentation properties fixed in Experiment 1 using a signal detection theory model (a two-alternative discrimination model for the orientation discrimination task and a differencing model for the change detection task), the pattern of accuracy and response time results was remarkably similar in form. Both Experiments 1 and 2 demonstrated that the sample-size relation was obtained in a near-threshold change detection procedure, in addition to obtaining the sample-size relation over orientation discrimination data, as reported by Sewell et al. (2014). These constraints on the memory system were found below any putative item limit, and their relationship to that item limit is unclear. Recent results by Donkin et al. (2016) indicated that the appearance of a slot or discrete item limit may reflect a shift in the encoding strategy for observers, and that resource-like responding can be seen in cases where observers have clear expectations about the number of elements that will be presented on a trial-to-trial basis. Donkin and colleagues noted that the common constraint between resource-like and slot-like responding was of the form of the sample-size relation examined here. Prior to this, the sample-size representation was only seen in experiments where the display contrast was limited, restricting the amount of available stimulus information, and where the memory array sizes were small, such as in the work of Sewell and colleagues. These findings reinforce the point that detailed modelling and controlled experimentation, of the kind presented here and of the kind engaged in by Donkin and colleagues, where the encoding and maintenance of stimuli is explicitly considered as well as the decision process relating a memory representation to an observed response is required to resolve—in some part—the relationship between slot-based accounts and resource-based accounts.

In addition to demonstrating the sample-size relation in a near-threshold change de-

tection paradigm, one of the clearest findings from the current thesis is that the four to five item storage capacity estimate may be systematically biased due to the obligatory encoding and maintenance of the probe array elements in memory. The magnitude of the difference in accuracy between the orientation discrimination and change detection tasks was predicted as the addition of an item into the sample-size relation in the change detection task, indicating the additional load on memory imposed by the probe array. This obligatory encoding of probe array material, regardless of the fact that the information is still visible to the observer, suggests that storage capacity estimates drawn from change detection performance alone may require an explicit correction for the additional memory load incurred in encoding the probe array. In addition, given that the current studies used a post-stimulus probe paradigm to control decision complexity, the only decision-relevant item in the probe item was known to the observer. Many studies, however, do not specify the target location, requiring observers to make a separate change detection decision for each item presented, potentially increasing both the errors due to noise in the decision stage and potentially increasing the load of the probe array on memory. A further study in which different requirements for making a comparison to the probe array are compared directly to orientation discrimination performance would be beneficial, particularly if conducted using a near-threshold paradigm.

The role of the probe array and response complexity will also be examined in the next section.

7.3.2 Continuous report

Many recent visual short-term memory studies have employed a continuous report paradigm, attempting to examine how memory constraints interact with the reproduction of a stimulus representation. This thesis did not investigate a continuous report task as there is no decision model for response time and accuracy comparable to those for orientation discrimination and change detection (Smith, 2016, has recently proposed such a model, but its empirical properties are yet to be investigated).

The lack of a detailed model of the decision process in continuous report means that “precision” accounts of memory (Bays et al., 2009; Zhang & Luck, 2008) identify response variability directly with memory variability. This direct identification may be misleading, as the results of Experiments 3 and 4 showed. Performance in these experiments was well described by a model in which the effects of memory load reduced the signal-to-noise ratio of memory representations without changing the channel tuning. Neurally, this is consistent with a system in which memory load affects the number of

neurons that are recruited to code a stimulus but not their tuning. This conclusion is strengthened by the demonstration by Smith (2016) that changes in response variability due to increased stimulus information cannot be distinguished in a continuous report task from changes in the response criterion without using response time data.

The method of stimulus reproduction may also have an effect on the observer performance in the continuous report paradigm. Given that the probe array places an apparent additional load on the memory system, it is possible that the process of reproducing the target stimulus from memory may also impose a cost in terms of storage capacity. The mixture modelling of Zhang and Luck (2009) indicates that this cost would be fixed over the retention interval, but it is yet to be determined whether precision remains stable over the interval of response selection and execution, when comparisons are required. Any cost of comparison may be heightened by the total number of object-level decisions, as was found by Woodman and Vecera (2011). Further experimentation, and modelling constrained by the sample-size relationship, would help understand the decremented imposed by memory retrieval comparison in terms of the overall storage capacity limitations of visual short-term memory.

7.4 Theoretical implications

Each of the major findings of this thesis place strong restrictions on theories of visual short-term memory: the existence and independence of the information constraints identified in the second part of the thesis, the existence of an intrusion process in the diffusion modelling, and the obligatory encoding of the probe array in change detection. Several of these restrictions exist at the level of the encoding and maintenance of representations: the tuned channel relationship between orientation and stimulus information may exist at the level of encoding or memory architecture; the sample-size relation is a restriction on the total information that can be simultaneously maintained, further affected by the performance decrement in change detection; and the linear increase in stimulus information is a restriction on encoding.

The intrusion process appears to be a phenomenon arising at the retrieval and decision stage as retrieval fails to access target information and non-target information enters into the decision process, like the non-target responding identified by Bays, Wu, and Husain (2011). At the very least, this suggests that the decision process is a non-trivial aspect of task performance and should be a consideration in models of visual short-term memory.

These restrictions on modelling will be examined in the next two sections looking at the

“slots + resources” model of Zhang and Luck (2008) and the multielement integrated system model of Smith and Sewell (2013).

7.4.1 Slots + resources

The hybrid “slots + resources” model of Zhang and Luck (2008) sought to unify conflicting results from change detection and continuous report experiments which demonstrated both a fixed upper limit on the number of items that could be maintained simultaneously in memory as well as changes in response variability for memory array size smaller than the item limit. The coexistence of a discrete item limit with some flexibility in the allocation of resources appears, *prima facie*, to be consistent with the findings of Donkin et al. (2016), albeit in the inverse direction: where Zhang and Luck contended that the fundamental limit was the discrete item limit, and resource flexibility was permitted when the number of memory array items was smaller than the item limit, Donkin and colleagues demonstrated that strategic slot-like behaviour could exist, constrained in size by a fundamental resource relationship.

The results from the current thesis demonstrate that this relationship may be little more complex. As discussed in the general discussion of Chapter 6 (§6.3), the sample-size relation is consistent with the overall resource limit also identified by Donkin and colleagues. The sample-size relation, however, appears to be independent of the fundamental precision of the memory system to represent orientation. The dynamic redistribution of resources under the item limit, argued by Zhang and Luck, where one memory item may be represented multiple times to increase the precision of the representation itself is not evident from the current data. The changes in response variability observed by Zhang and Luck would be explained in terms of the current results by the higher signal-to-noise ratio for target detectors over distractor detectors and background noise.

Last, the relationship between the memory representation and the observed data is further specified by the current results. Zhang and Luck used a mixture model comprising a von Mises memory distribution and a uniform guessing distribution, where the variability of the von Mises distribution and the mixture probability were freely estimated. The current data places hard constraints on both the response distributions and the response time distributions that are expected with different experimental manipulations, in both the change in the memory distribution and in non-memory variability. The sample-size relation provides a strong constraint on the change in response variability across memory array size manipulations. This is consistent with the findings of Bays and Husain (2008), who demonstrated that response variability changes as a power function

of the number of items in the memory array. The change in stimulus exposure duration also predicts a decrease in response variability as additional information is available for encoding and storage in the memory system. The work by Bays, Gorgoraptis, et al. (2011) showed an interaction between the size of the memory array and exposure duration of the memory array on response variability, but the current data would suggest that this interaction is a product of both manipulations increasing the overall quality of stimulus information in memory. In the current modelling, the two factors enter into the sensitivity model and the diffusion model independent of one another.

The intrusion process predicts that, on some proportion of trials related to the proportion of items contrasting with the target identity shown in the memory array, non-target information may intrude into the memory process and guide observed behaviour. Non-target reproduction has been noted by Bays, Wu, and Husain (2011), who showed that there was structure in the reproduction responses that were distant from the mean of the target distribution, and appeared to be clustered around distractor feature values. Further research is required to examine if the asymmetrical quantiles between correct and incorrect responses observed in the response time data obtained in this thesis is related to non-target reproduction directly.

Specific quantitative predictions are difficult to make for the mixture model formulation provided by Zhang and Luck without a more complete model relating the target-related stimulus information retrieved from memory to response variability directly. As discussed, the recent work of Smith (2016) provides a productive avenue to unify the model constraints raised in this thesis with the data obtained in a continuous report task.

7.4.2 Integrated system model

The integrated system model (Smith et al., 2009) and its multielement generalisation (Smith & Sewell, 2013) present a theoretically parsimonious and comprehensive model of observer performance in simple visual tasks. Unlike measurement models—such as those employed by Zhang and Luck, by Bays, Gorgoraptis, et al. (2011), or by van den Berg, Awh, and Ma (2014b)—the integrated system model and the multielement integrated system model allow a dynamic process representation of the encoding and maintenance of information in visual short-term memory, and both the response proportions and response time distributions that are obtained in decisions made from those representations. Although the predictions that result from these models offer a greater level of specificity than mixture models, the complexity in integrating additional constraints

is potentially also much greater.

The relationship between memory array size and observer performance is predicted as a consequence of the central dynamics of the model. The sample-size relation is an integral component of the multielement integrated system model. Competitive inhibition between maintained memory representations leads to a steady division of resources, with the strength of stimulus information resulting predicted by the sample-size relation. Each experiment presented in the current thesis have providing support for this prediction.

The integration of the tuned channel constraint into the multielement integrated system model is an open question. Although the multielement integrated system model represents interstimulus similarity, this parameter acts upon the selection of stimuli into memory, in order to predict the distractor homogeneity effect (Duncan & Humphreys, 1989; Mazyar, Berg, & Ma, 2012) where visual search performance is facilitated by the similarity of distractors to one another. A more natural position for the tuned channel constraint is within the interface between the representation of maintained memory traces and the decision process. During the retention interval, and prior to the decision, the delay of the post-stimulus probe means that a decision process is not likely to begin, unless in the case of single item memory arrays where the non-decision time estimates indicate the onset of the decision either coincides with the probe or slightly precedes it. Once the probe has been presented, and the target has been localised, retrieval of the representation can begin (entailing some response time cost itself), and the decision can commence. The angular difference between the retrieved memory representation and the internal standard (such as an oblique or vertical stimulus) may then modulate the drift rate in the decision process.

The interface in the multielement integrated system model between visual short-term memory and the final decision process is a minimal relationship: the drift rate of the diffusion process should be proportional to the strength of the visual short-term memory trace. Extending this relation to include the relationship between memory representation and the standard to compare against is one way of increasing the predictive power of the model to account for the current results. Another extension arising from the current work is the inclusion of probe array encoding and comparison. An additional constant to the non-decision time of the decision process, as well as a decrement to the trace strength of the maintained representations as the probe array elements are stored in memory would allow the prediction of change detection results using the multielement integrated system model.

This interface could also be modified to allow for the full prediction of intruding infor-

mation, by allowing some proportion of samples to be drawn from non-target memory traces. Unlike the diffusion modelling presented in preceding chapters, the multielement integrated system model provides a dynamic representation of the state of each item in memory, meaning that the non-target information is explicitly confined by the model rather than an estimated intrusion drift rate parameter. The specification of the probability for non-target sample draws could then be investigated in directed experiments. One potential route for further investigation is that, in reflecting the inefficiency of the post-stimulus probe to localise retrieval to the target location, the probability of non-target information could be related to the distance of stimuli locations from the probe. A spatially mediated retrieval process is compatible with both the sensitivity of change detection tasks to some, but not all, transformations of the correspondence between the memory and probe arrays (Jiang et al., 2000) and to retrocuing work, where the strength of memory traces is manipulated through an orienting cue (Griffin & Nobre, 2003). Additional proposals for future experimentation is presented in the next section.

7.5 Future research

Although the experimental procedures and modelling presented in the current thesis provide some open questions for future research they also, in part, supply methods for resolving those questions. By providing clear models for describing patterns of observed data, and providing experiments that closely correspond to elements in those models, some purchase can be gained on these issues.

One of the most prominent open questions in the visual short-term memory literature is in the adjudication between discrete slot and flexible resource accounts of memory storage. Donkin et al. (2016) distinguished between a fundamental resource limit and response strategies that produced slot-like responding. The current thesis supported the idea of a fundamental resource limit in the sample-size model, below any putative object limit. In Experiments 1, 3, and 4 the memory array size was varied on a trial-by-trial basis, which Donkin and colleagues suggested may induce slot-like responding. Although no discrete item limit was observed in these experiments, this was likely due to the fact that the maximum memory array size in each of these experiments—four items—is below that of many change detection experiments.

In Experiment 2, potential changes in strategic responding were observed, but not between different memory array size manipulations (only a single array size of four items was used for each trial). Instead, the differences in strategy—manifesting as differences in

the boundary separation and intrusion probability recovered in the diffusion analysis—were seen between different decision type conditions, but not seen in Experiment 1, which employed the same design but had a shorter training period and an identical display characteristics between the decision type trials for the presentation of the memory array. The decision type manipulation, between trials of orientation discrimination and change detection, varied between prespecified blocks rather than randomly between each trial. As the overall difficulty of the task—shown by the lower levels of accuracy—and the predictability of the decision requirements of each trial would suggest that observers could place an incentive on observers to maximise performance through subtle changes in deciding, a future experiment could further examine this strategic difference by examining a design where the decision type of each trial varies on a trial-by-trial basis.

Future experiments with memory arrays consisting of more than four items could also be examined in experimental procedures using response time information and near-threshold stimulus presentation to further define the strategic changes obtained when observers must accurately respond given larger memory array sizes. The diffusion model decision parameters controlling the intrusion process and the boundary separation parameters could be used to characterise changes under observer control. This could also be supplemented by an explicit examination of the order or speed of stimulus encoding, as was conducted using the attention-weighted sample-size model reported by Smith et al. (under review), where memory resources were divided unequally, showing a bias in the encoding of stimulus locations. These analyses are contingent on careful examination of the response time distributions obtained in these tasks, and require careful experimental design.

Further specification of the intrusion process also presents an opportunity for future research. In the current thesis, and in the analyses reported by Sewell et al. (in press), the addition of a delayed diffusion process mixed with the main diffusion process greatly improved diffusion model fits to tasks reliant on retrieval from visual short-term memory. These additional processes were able to account for an asymmetry in the leading edge of the response time distribution between correct and incorrect responses, indicating some delay in responding in even the earliest incorrect responses. Although Sewell and colleagues constrained the mixture process to have a drift rate of zero, indicating zero stimulus information and delayed guessing, the best fitting parameters and model configurations for each of the diffusion analyses used negative drift rates, implying that distractor information that contrasted with the target identity was contributing the decision process.

The inefficiency of the post-stimulus probe to localise the retrieval of stimulus infor-

mation in memory to the target information alone or to fail entirely, leading to guessing, provides an important avenue for further theoretical development. Examining how stimulus and probe arrangements may affect the properties of the intrusion process experimentally would be a logical further step, by changing the probability of distractors to contain contrasting identity information and changing the spatial arrangement of the probe to the distractor locations to test for a spatially mediated retrieval process. As indicated in the last section, the integration of a retrieval failure account into the integrated system model, where a dynamic representation of each stimulus location is explicitly accounted for within the model, would provide further leverage on this question.

Last, Experiments 3 and 4 provided an examination of the relationship between the sample-size constraint on representations within memory and the underlying feature information within those representations. The tuned channel constraint provides a description of the change in observer performance with a change in the angular difference between an item held within memory and the standard for comparison. A productive avenue for future investigation would be to examine whether the same constraint exists in simple discrimination tasks for other features, such as colour, and whether independence is seen in constraints upon representations which comprise multiple features. Although colour and orientation have been examined extensively in the visual short-term memory literature, I hope to have demonstrated in the current thesis that additional traction can be gained through simple relationships derived from behaviour in simple, well understood tasks.

Visual short-term memory sits at the intersection between low-level vision and higher-order cognitive processes, between the representation of visual information in the brain and the control of decision-making and memory. The characterisation of the constraints that exist upon visual memory, and the nature of decisions made from information held within visual short-term memory, provide further detail about how simple sensory information becomes the basis for complex behaviour. The centrality of visual short-term memory means that its properties are primal in the perception of the visual world—they are, in some modest way, the prism through which we see the world.

References

- Alvarez, G. A., & Cavanagh, P. (2004). The Capacity of Visual Short-Term Memory is Set Both by Visual Information Load and by Number of Objects. *Psychological Science*, 15(2), 106–111. doi: 10.1111/j.0963-7214.2004.01502006.x
- Anderson, D. E., & Awh, E. (2012). The plateau in mnemonic resolution across large set sizes indicates discrete resource limits in visual working memory. *Attention, Perception, & Psychophysics*, 74(5), 891–910. doi: 10.3758/s13414-012-0292-1
- Anderson, D. E., Vogel, E. K., & Awh, E. (2011). Precision in Visual Working Memory Reaches a Stable Plateau When Individual Item Limits Are Exceeded. *The Journal of Neuroscience*, 31(3), 1128–1138. doi: 10.1523/JNEUROSCI.4125-10.2011
- Appelle, S. (1972). Perception and discrimination as a function of stimulus orientation: the “oblique effect” in man and animals. *Psychological bulletin*, 78(4), 266. Retrieved from <http://psycnet.apa.org/journals/bul/78/4/266/>
- Atkinson, R. C., & Shiffrin, R. M. (1968). Human Memory: A Proposed System and its Control Processes. In K. W. Spence & J. Taylor (Eds.), *Psychology of Learning and Motivation* (Vol. 2, pp. 89–195). Academic Press.
- Averbach, E., & Sperling, G. (1961). Short-term storage in Vision. In C. Cherry (Ed.), *Symposium on information theory* (pp. 196–211). London: Butterworth.
- Awh, E., Barton, B., & Vogel, E. K. (2007). Visual Working Memory Represents a Fixed Number of Items Regardless of Complexity. *Psychological Science*, 18(7), 622–628. doi: 10.1111/j.1467-9280.2007.01949.x
- Baddeley, A. (1992). Working memory. *Science*, 255(5044), 556–559. doi: 10.1126/science.1736359
- Baddeley, A. (2012). Working Memory: Theories, Models, and Controversies. *Annual Review of Psychology*, 63(1), 1–29. doi: 10.1146/annurev-psych-120710-100422
- Baddeley, A., & Hitch, G. (1974). Working Memory. In G. H. Bower (Ed.), *Psychology of Learning and Motivation* (Vol. 8, pp. 47–89). Elsevier. doi: 10.1016/s0079-7421(08)60452-1
- Baddeley, A., & Sala, S. D. (1996). Working Memory and Executive Control [and Discussion]. *Philosophical Transactions of the Royal Society of London B: Biological*

References

- Sciences*, 351(1346), 1397–1404. doi: 10.1098/rstb.1996.0123
- Baldassi, S., & Burr, D. C. (2004). “Pop-out” of targets modulated in luminance or colour: the effect of intrinsic and extrinsic uncertainty. *Vision Research*, 44(12), 1227–1233. doi: 10.1016/j.visres.2003.12.018
- Barak, O., & Tsodyks, M. (2014). Working models of working memory. *Current Opinion in Neurobiology*, 25, 20–24. doi: 10.1016/j.conb.2013.10.008
- Barton, B., Ester, E. F., & Awh, E. (2009). Discrete resource allocation in visual working memory. *Journal of Experimental Psychology: Human Perception and Performance*, 35(5), 1359–1367. doi: 10.1037/a0015792
- Baxt, N. (1871). Ueber die Zeit, welche nöthig ist, damit ein Gesichtseindruck zum Bewusstsein kommt und über die Grösse (Extension) der bewussten Wahrnehmung bei einem Gesichtseindrucke von gegebener Dauer. *Archiv für die gesamte Physiologie des Menschen und der Tiere*, 4(1), 325–336. doi: 10.1007/BF01612494
- Bays, P. M. (2014). Noise in Neural Populations Accounts for Errors in Working Memory. *The Journal of Neuroscience*, 34(10), 3632–3645. doi: 10.1523/JNEUROSCI.3204-13.2014
- Bays, P. M., Catalao, R. F. G., & Husain, M. (2009). The precision of visual working memory is set by allocation of a shared resource. *Journal of Vision*, 9(10). doi: 10.1167/9.10.7
- Bays, P. M., Gorgoraptis, N., Wee, N., Marshall, L., & Husain, M. (2011). Temporal dynamics of encoding, storage, and reallocation of visual working memory. *Journal of Vision*, 11(10), 6. doi: 10.1167/11.10.6
- Bays, P. M., & Husain, M. (2008). Dynamic Shifts of Limited Working Memory Resources in Human Vision. *Science*, 321(5890), 851–854. doi: 10.1126/science.1158023
- Bays, P. M., & Husain, M. (2009). Response to Comment on “Dynamic Shifts of Limited Working Memory Resources in Human Vision”. *Science*, 323, 877. Retrieved from <http://www.paulbays.com/pdf/BayHus09.pdf>
- Bays, P. M., Wu, E. Y., & Husain, M. (2011). Storage and binding of object features in visual working memory. *Neuropsychologia*, 49(6), 1622–1631. doi: 10.1016/j.neuropsychologia.2010.12.023
- Bengson, J. J., & Luck, S. J. (2015). Effects of strategy on visual working memory capacity. *Psychonomic Bulletin & Review*, 23(1), 265–270. doi: 10.3758/s13423-015-0891-7
- Ben-Yishai, R., Bar-Or, R. L., & Sompolinsky, H. (1995). Theory of orientation tuning in visual cortex. *Proceedings of the National Academy of Sciences*, 92(9), 3844–3848. doi: 10.1073/pnas.92.9.3844
- Bonnel, A.-M., & Hafter, E. R. (1998). Divided attention between simultaneous au-

- ditory and visual signals. *Perception & Psychophysics*, 60(2), 179–190. doi: 10.3758/BF03206027
- Bonnel, A.-M., & Miller, J. (1994). Attentional effects on concurrent psychophysical discriminations: Investigations of a sample-size model. *Perception & Psychophysics*, 55(2), 162–179. doi: 10.3758/BF03211664
- Brady, T. F., & Alvarez, G. A. (2011). Hierarchical Encoding in Visual Working Memory Ensemble Statistics Bias Memory for Individual Items. *Psychological Science*, 22(3), 384–392. doi: 10.1177/0956797610397956
- Brady, T. F., Konkle, T., & Alvarez, G. A. (2009). Compression in visual working memory: Using statistical regularities to form more efficient memory representations. *Journal of Experimental Psychology: General*, 138(4), 487–502. doi: 10.1037/a0016797
- Brady, T. F., Konkle, T., & Alvarez, G. A. (2011). A review of visual memory capacity: Beyond individual items and toward structured representations. *Journal of Vision*, 11(5). doi: 10.1167/11.5.4
- Broadbent, D. E., & Broadbent, M. H. P. (1987). From detection to identification: Response to multiple targets in rapid serial visual presentation. *Perception & Psychophysics*, 42(2), 105–113. doi: 10.3758/BF03210498
- Brown, S. D., & Heathcote, A. (2008). The simplest complete model of choice response time: Linear ballistic accumulation. *Cognitive Psychology*, 57(3), 153–178. doi: 10.1016/j.cogpsych.2007.12.002
- Broyden, C. G. (1970). The convergence of single-rank quasi-Newton methods. *Mathematics of Computation*, 24(110), 365–382. doi: 10.1090/S0025-5718-1970-0279993-0
- Bundesen, C. (1987). Visual attention: Race models for selection from multielement displays. *Psychological Research*, 49(2-3), 113–121. doi: 10.1007/BF00308676
- Bundesen, C. (1990). A theory of visual attention. *Psychological Review*, 97(4), 523–547. doi: 10.1037/0033-295X.97.4.523
- Bundesen, C., Habekost, T., & Kyllingsbæk, S. (2005). A Neural Theory of Visual Attention: Bridging Cognition and Neurophysiology. *Psychological Review*, 112(2), 291–328. doi: 10.1037/0033-295X.112.2.291
- Bundesen, C., Habekost, T., & Kyllingsbæk, S. (2011). A neural theory of visual attention and short-term memory (NTVA). *Neuropsychologia*, 49(6), 1446–1457. doi: 10.1016/j.neuropsychologia.2010.12.006
- Bundesen, C., Pedersen, L. F., & Larsen, A. (1984). Measuring efficiency of selection from briefly exposed visual displays: A model for partial report. *Journal of Experimental Psychology*, 10(3), 329–339.

References

- Burnham, K. P., & Anderson, D. R. (2003). *Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach* (2nd edition ed.). New York: Springer.
- Byrd, R., Lu, P., Nocedal, J., & Zhu, C. (1995). A Limited Memory Algorithm for Bound Constrained Optimization. *SIAM Journal on Scientific Computing*, 16(5), 1190–1208. doi: 10.1137/0916069
- Camperi, M., & Wang, X.-J. (1998). A Model of Visuospatial Working Memory in Pre-frontal Cortex: Recurrent Network and Cellular Bistability. *Journal of Computational Neuroscience*, 5(4), 383–405. doi: 10.1023/A:1008837311948
- Carpenter, R. H. S. (1981). Oculomotor Procrastination. In D. Fisher, R. Monty, & J. Senders (Eds.), *Eye movements: cognition and visual perception*. (pp. 237–246). Hillsdale, NJ: Lawrence Erlbaum.
- Carpenter, R. H. S., & Williams, M. L. L. (1995). Neural computation of log likelihood in control of saccadic eye movements. *Nature*, 377(6544), 59–62. doi: 10.1038/377059a0
- Carrasco, M., Penpeci-Talgar, C., & Eckstein, M. (2000). Spatial covert attention increases contrast sensitivity across the CSF: support for signal enhancement. *Vision Research*, 40(10–12), 1203–1215. doi: 10.1016/S0042-6989(00)00024-9
- Cattell, J. M. (1885). The Inertia of the Eye and Brain. *Brain*, 8(3), 295–312.
- Cermak, G. W. (1971). Short-term recognition memory for complex free-form figures. *Psychonomic Science*, 25(4), 209–211. doi: 10.3758/BF03329095
- Chun, M. M., & Potter, M. C. (1995). A Two-Stage Model for Multiple Target Detection in Rapid Serial Visual Presentation. *Journal of Experimental Psychology*, 21(1), 109–127. doi: 10.1037/0096-1523.21.1.109
- Coltheart, M. (1980). Iconic memory and visible persistence. *Perception & Psychophysics*, 27(3), 183–228. doi: 10.3758/BF03204258
- Corbetta, M., & Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*, 3(3), 201–215. doi: 10.1038/nrn755
- Cornelissen, F. W., & Greenlee, M. W. (2000). Visual memory for random block patterns defined by luminance and color contrast. *Vision Research*, 40(3), 287–299. doi: 10.1016/S0042-6989(99)00137-6
- Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, 24(01), 87–114. doi: 10.1017/S0140525X01003922
- Cowan, N., & Rouders, J. N. (2009). Comment on "Dynamic Shifts of Limited Working Memory Resources in Human Vision". *Science*, 323(5916), 877–877. doi:

- 10.1126/science.1166478
- Craik, F. I. M., & Lockhart, R. S. (1972). Levels of processing: A framework for memory research. *Journal of Verbal Learning and Verbal Behavior*, 11(6), 671–684. doi: 10.1016/S0022-5371(72)80001-X
- Craik, F. I. M., & Watkins, M. J. (1973). The role of rehearsal in short-term memory. *Journal of Verbal Learning and Verbal Behavior*, 12(6), 599–607. doi: 10.1016/S0022-5371(73)80039-8
- Currie, C. B., McConkie, G. W., Carlson-Radvansky, L. A., & Irwin, D. E. (2000). The role of the saccade target object in the perception of a visually stable world. *Perception & Psychophysics*, 62(4), 673–683. doi: 10.3758/BF03206914
- Curtis, C. E., & D'Esposito, M. (2003). Persistent activity in the prefrontal cortex during working memory. *Trends in Cognitive Sciences*, 7(9), 415–423. doi: 10.1016/S1364-6613(03)00197-9
- Dai, H., Versfeld, N. J., & Green, D. M. (1996). The optimum decision rules in the same-different paradigm. *Perception & Psychophysics*, 58(1), 1–9. doi: 10.3758/BF03205469
- Daugman, J. G. (1980). Two-dimensional spectral analysis of cortical receptive field profiles. *Vision Research*, 20(10), 847–856. doi: 10.1016/0042-6989(80)90065-6
- Daugman, J. G. (1984). Spatial visual channels in the Fourier plane. *Vision Research*, 24(9), 891–910. doi: 10.1016/0042-6989(84)90065-8
- Di Lollo, V., Lowe, D. G., & Scott, J. P. (1974). Backward masking and interference with the processing of brief visual displays. *Journal of Experimental Psychology*, 103(5), 934–940.
- Donders, F. C. (1868/1969). On the speed of mental processes. *Acta Psychologica*, 30, 412–431. doi: 10.1016/0001-6918(69)90065-1
- Donkin, C., Kary, A., Tahir, F., & Taylor, R. (2016). Resources masquerading as slots: Flexible allocation of visual working memory. *Cognitive Psychology*, 85, 30–42. doi: 10.1016/j.cogpsych.2016.01.002
- Donkin, C., Nosofsky, R., Gold, J., & Shiffrin, R. (2014). Verbal labeling, gradual decay, and sudden death in visual short-term memory. *Psychonomic Bulletin & Review*, 1–9. doi: 10.3758/s13423-014-0675-5
- Donkin, C., Nosofsky, R. M., Gold, J. M., & Shiffrin, R. M. (2013). Discrete-slots models of visual working-memory response times. *Psychological Review*, 120(4), 873–902. doi: 10.1037/a0034247
- Downing, C. J. (1988). Expectancy and visual-spatial attention: Effects on perceptual quality. *Journal of Experimental Psychology: Human Perception and Performance*,

References

- 14(2), 188–202. doi: 10.1037/0096-1523.14.2.188
- Duncan, J. (1980). The locus of interference in the perception of simultaneous stimuli. *Psychological Review*, 87(3), 272–300. doi: 10.1037/0033-295X.87.3.272
- Duncan, J. (1984). Selective attention and the organization of visual information. *Journal of Experimental Psychology: General*, 113(4), 501–517. doi: 10.1037/0096-3445.113.4.501
- Duncan, J., & Humphreys, G. W. (1989). Visual search and stimulus similarity. *Psychological Review*, 96(3), 433–458. doi: 10.1037/0033-295X.96.3.433
- Ellias, S. A., & Grossberg, S. (1975). Pattern formation, contrast control, and oscillations in the short term memory of shunting on-center off-surround networks. *Biological Cybernetics*, 20(2), 69–98. doi: 10.1007/BF00327046
- Eng, H. Y., Chen, D., & Jiang, Y. (2005). Visual working memory for simple and complex visual stimuli. *Psychonomic Bulletin & Review*, 12(6), 1127–1133. doi: 10.3758/BF03206454
- Enroth-Cugell, C., & Robson, J. G. (1966). The contrast sensitivity of retinal ganglion cells of the cat. *The Journal of Physiology*, 187(3), 517–552. doi: 10.1113/jphysiol.1966.sp008107
- Erdmann, B., & Dodge, R. (1898). *Psychologische Untersuchungen über das Lesen auf experimenteller Grundlage*. Halle: Niemeyer.
- Fletcher, R. (1970). A new approach to variable metric algorithms. *The Computer Journal*, 13(3), 317–322. doi: 10.1093/comjnl/13.3.317
- Fougnie, D., Asplund, C. L., & Marois, R. (2010). What are the units of storage in visual working memory? *Journal of Vision*, 10(12), 27–27. doi: 10.1167/10.12.27
- Furmanski, C. S., & Engel, S. A. (2000). An oblique effect in human primary visual cortex. *Nature Neuroscience*, 3(6), 535–536. doi: 10.1038/75702
- Gao, Z., Xu, X., Chen, Z., Yin, J., Shen, M., & Shui, R. (2011). Contralateral delay activity tracks object identity information in visual short term memory. *Brain Research*, 1406, 30–42. doi: 10.1016/j.brainres.2011.06.049
- Garner, W. R., & Felfoldy, G. L. (1970). Integrality of stimulus dimensions in various types of information processing. *Cognitive Psychology*, 1(3), 225–241. doi: 10.1016/0010-0285(70)90016-2
- Gegenfurtner, K. R., & Sperling, G. (1993). Information transfer in iconic memory experiments. *Journal of Experimental Psychology: Human Perception and Performance*, 19(4), 845–866. doi: 10.1037/0096-1523.19.4.845
- Goldfarb, D. (1970). A family of variable-metric methods derived by variational means. *Mathematics of Computation*, 24(109), 23–26. doi:

- 10.1090/S0025-5718-1970-0258249-6
- Gould, I. C., Wolfgang, B. J., & Smith, P. L. (2007). Spatial uncertainty explains exogenous and endogenous attentional cuing effects in visual signal detection. *Journal of Vision*, 7(13). doi: 10.1167/7.13.4
- Graham, N. V. S. (1989). *Visual pattern analyzers* (No. no. 16). New York: Oxford University Press.
- Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics*. New York: Wiley.
- Griffin, I. C., & Nobre, A. C. (2003). Orienting Attention to Locations in Internal Representations. *Journal of Cognitive Neuroscience*, 15(8), 1176–1194. doi: 10.1162/089892903322598139
- Grossberg, S. (1982). Contour Enhancement, Short Term Memory, and Constancies in Reverberating Neural Networks. In *Studies of Mind and Brain* (pp. 332–378). Springer Netherlands.
- Grossberg, S. (1987). Competitive learning: From interactive activation to adaptive resonance. *Cognitive Science*, 11(1), 23–63. doi: 10.1016/S0364-0213(87)80025-3
- Grossberg, S. (1988). Nonlinear neural networks: Principles, mechanisms, and architectures. *Neural Networks*, 1(1), 17–61. doi: 10.1016/0893-6080(88)90021-4
- Habak, C., Wilkinson, F., Zakher, B., & Wilson, H. R. (2004). Curvature population coding for complex shapes in human vision. *Vision Research*, 44(24), 2815–2823. doi: 10.1016/j.visres.2004.06.019
- Hacker, M. J., & Ratcliff, R. (1979). A revised table of d' for M-alternative forced choice. *Perception & Psychophysics*, 26(2), 168–170. doi: 10.3758/BF03208311
- Heeley, D. W., Buchanan-Smith, H. M., Cromwell, J. A., & Wright, J. S. (1997). The oblique effect in orientation acuity. *Vision Research*, 37(2), 235–242. doi: 10.1016/S0042-6989(96)00097-1
- Henderson, J. M., & Hollingworth, A. (2003). Eye movements and visual memory: Detecting changes to saccade targets in scenes. *Perception & Psychophysics*, 65(1), 58–71. Retrieved from <http://link.springer.com/article/10.3758/BF03194783>
- Hollingworth, A. (2004). Constructing Visual Representations of Natural Scenes: The Roles of Short- and Long-Term Visual Memory. *Journal of Experimental Psychology: Human Perception and Performance*, 30(3), 519–537. doi: 10.1037/0096-1523.30.3.519
- Hollingworth, A., Hyun, J.-S., & Zhang, W. (2005). The role of visual short-term memory in empty cell localization. *Perception & Psychophysics*, 67(8), 1332–1343. Retrieved from <http://link.springer.com/article/10.3758/BF03193638>

References

- Hollingworth, A., & Luck, S. J. (2009). The role of visual working memory (VWM) in the control of gaze during visual search. *Attention, Perception, & Psychophysics*, 71(4), 936–949. doi: 10.3758/APP.71.4.936
- Hollingworth, A., Richard, A. M., & Luck, S. J. (2008). Understanding the function of visual short-term memory: Transsaccadic memory, object correspondence, and gaze correction. *Journal of Experimental Psychology: General*, 137(1), 163–181. doi: 10.1037/0096-3445.137.1.163
- Hubel, D. H., & Wiesel, T. N. (1959). Receptive fields of single neurones in the cat's striate cortex. *The Journal of Physiology*, 148(3), 574–591. Retrieved from <http://jp.physoc.org/content/148/3/574.full.pdf>
- Hubel, D. H., & Wiesel, T. N. (1963). Shape and arrangement of columns in cat's striate cortex. *The Journal of Physiology*, 165(3), 559–568. Retrieved from <http://jp.physoc.org/content/165/3/559.full.pdf>
- Hyun, J.-s., Woodman, G. F., Vogel, E. K., Hollingworth, A., & Luck, S. J. (2009). The comparison of visual working memory representations with perceptual inputs. *Journal of Experimental Psychology: Human Perception and Performance*, 35(4), 1140–1160. doi: 10.1037/a0015019
- Irwin, D. E., & Andrews, R. V. (1996). Integration and accumulation of information across saccadic eye movements. In T. Inui & J. L. McClelland (Eds.), *Attention and Performance XVI: Information Integration in Perception and Communication*. Cambridge, MA: MIT Press.
- Jacobs, J. (1887). Experiments on "Prehension". *Mind*, 12(45), 75–79. Retrieved from <http://www.jstor.org/stable/2246990>
- Jacobs, R., & Orhan, A. E. (2013). A Probabilistic Clustering Theory of the Organization of Visual Short-Term Memory. *Journal of Vision*, 13(9), 1368–1368. doi: 10.1167/13.9.1368
- Jiang, Y., Olson, I. R., & Chun, M. M. (2000). Organization of visual short-term memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26(3), 683–702. doi: 10.1037/0278-7393.26.3.683
- Jolicœur, P., & Dell'Acqua, R. (1998). The Demonstration of Short-Term Consolidation. *Cognitive Psychology*, 36(2), 138–202. doi: 10.1006/cogp.1998.0684
- Kahneman, D. (1968). Method, findings, and theory in studies of visual masking. *Psychological Bulletin*, 404–425.
- Kahneman, D., Treisman, A., & Gibbs, B. J. (1992). The reviewing of object files: Object-specific integration of information. *Cognitive Psychology*, 24(2), 175–219. doi: 10.1016/0010-0285(92)90007-O

- Kahneman, D., & Treisman, A. M. (1984). Changing views of attention and automaticity. In R. Parasuraman & D. R. Davies (Eds.), *Varieties of Attention*. New York, NY: Academic Press.
- Kanwisher, N., & Wojciulik, E. (2000). Visual attention: Insights from brain imaging. *Nature Reviews Neuroscience*, 1(2), 91–100. doi: 10.1038/35039043
- Kanwisher, N. G. (1987). Repetition blindness: Type recognition without token individuation. *Cognition*, 27(2), 117–143. doi: 10.1016/0010-0277(87)90016-3
- Keshvari, S., van den Berg, R., & Ma, W. J. (2012). Probabilistic Computation in Human Perception under Variability in Encoding Precision. *PLoS ONE*, 7(6), e40216. doi: 10.1371/journal.pone.0040216
- Keshvari, S., van den Berg, R., & Ma, W. J. (2013). No Evidence for an Item Limit in Change Detection. *PLoS Comput Biol*, 9(2), e1002927. doi: 10.1371/journal.pcbi.1002927
- Kirchner, W. K. (1958). Age differences in short-term retention of rapidly changing information. *Journal of Experimental Psychology*, 55(4), 352–358.
- Kreyszig, E. (2011). *Advanced Engineering Mathematics* (10th ed.). Hoboken, NJ: Wiley.
- Kroll, N. E., Parks, T., Parkinson, S. R., Bieber, S. L., & Johnson, A. L. (1970). Short-term memory while shadowing: recall of visually and of aurally presented letters. *Journal of Experimental Psychology*, 85(2), 220.
- Kulikowski, J. J., & Tolhurst, D. J. (1973). Psychophysical evidence for sustained and transient detectors in human vision. *The Journal of Physiology*, 232(1), 149–162. doi: 10.1113/jphysiol.1973.sp010261
- Lamphiear, D. E., & Birdsall, T. G. (1960). *Approximations to the chi-square distributions with applications to signal detection models* (Tech. Rep. No. 101). The University of Michigan: Electronic Defense Group.
- Lee, D. K., Koch, C., & Braun, J. (1997). Spatial vision thresholds in the near absence of attention. *Vision Research*, 37(17), 2409–2418. doi: 10.1016/S0042-6989(97)00055-2
- Lindsay, P. H., Taylor, M. M., & Forbes, S. M. (1968). Attention and multidimensional discrimination. *Perception & Psychophysics*, 4(2), 113–117. doi: 10.3758/BF03209520
- Liu, C. C., Wolfgang, B. J., & Smith, P. L. (2009). Attentional mechanisms in simple visual detection: A speed–accuracy trade-off analysis. *Journal of Experimental Psychology: Human Perception and Performance*, 35(5), 1329–1345. doi: 10.1037/a0014255
- Luce, R. D. (1963). A Threshold Theory for Simple Detection Experiments. *Psychological Review*, 70(1), 61–79. doi: 10.1037/h0039723
- Luce, R. D. (1977). Thurstone's discriminial processes fifty years later. *Psychometrika*, 42(4), 461–489. doi: 10.1007/BF02295975
- Luce, R. D. (1986). *Response Times: Their Role in Inferring Elementary Mental Organization*

References

- (1st ed.). New York, NY: Oxford University Press.
- Luce, R. D., & Green, D. M. (1974). Detection, discrimination, and recognition. *Handbook of perception*, 2, 299–342.
- Luck, S. J., Girelli, M., McDermott, M. T., & Ford, M. A. (1997). Bridging the Gap between Monkey Neurophysiology and Human Perception: An Ambiguity Resolution Theory of Visual Selective Attention. *Cognitive Psychology*, 33(1), 64–87. doi: 10.1006/cogp.1997.0660
- Luck, S. J., & Hillyard, S. A. (1994a). Electrophysiological correlates of feature analysis during visual search. *Psychophysiology*, 31(3), 291–308. doi: 10.1111/j.1469-8986.1994.tb02218.x
- Luck, S. J., & Hillyard, S. A. (1994b). Spatial Filtering During Visual Search: Evidence From Human Electrophysiology. *Journal of Experimental Psychology*, 20(5), 1000–1014.
- Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, 390(6657), 279–281. doi: 10.1038/36846
- Ma, W. J. (2010). Signal detection theory, uncertainty, and Poisson-like population codes. *Vision Research*, 50(22), 2308–2319. doi: 10.1016/j.visres.2010.08.035
- Macmillan, N. A., & Creelman, C. D. (2004). *Detection Theory: A User's Guide* (2nd ed.). Mahwah, N.J: Psychology Press.
- Magnussen, S., & Greenlee, M. W. (1992). Retention and Disruption of Motion Information in Visual Short-Term Memory. *Journal of Experimental Psychology*, 18(1), 151–156.
- Magnussen, S., & Greenlee, M. W. (1999). The psychophysics of perceptual memory. *Psychological Research*, 62(2-3), 81–92. doi: 10.1007/s004260050043
- Magnussen, S., Greenlee, M. W., Asplund, R., & Dyrnes, S. (1990). Perfect visual short-term memory for periodic patterns. *European Journal of Cognitive Psychology*, 2(4), 345–362. doi: 10.1080/09541449008406212
- Makovski, T., & Jiang, Y. V. (2009). The role of visual working memory in attentive tracking of unique objects. *Journal of Experimental Psychology: Human Perception and Performance*, 35(6), 1687–1697. doi: 10.1037/a0016453
- Makovski, T., & Pertzov, Y. (2015). Attention and memory protection: Interactions between retrospective attention cueing and interference. *The Quarterly Journal of Experimental Psychology*, 68(9), 1735–1743. doi: 10.1080/17470218.2015.1049623
- Makovski, T., Shim, W. M., & Jiang, Y. V. (2006). Interference from filled delays on visual change detection. *Journal of Vision*, 6(12), 11–11. doi: 10.1167/6.12.11
- Makovski, T., Sussman, R., & Jiang, Y. V. (2008). Orienting attention in visual working memory reduces interference from memory probes. *Journal of Exper-*

- imental Psychology: Learning, Memory, and Cognition*, 34(2), 369–380. doi: 10.1037/0278-7393.34.2.369
- Marr, D. (1982/2010). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. Cambridge, Mass: The MIT Press.
- Matsukura, M., Luck, S. J., & Vecera, S. P. (2007). Attention effects during visual short-term memory maintenance: Protection or prioritization? *Perception & Psychophysics*, 69(8), 1422–1434. doi: 10.3758/BF03192957
- Mazurek, M. E., Roitman, J. D., Ditterich, J., & Shadlen, M. N. (2003). A Role for Neural Integrators in Perceptual Decision Making. *Cerebral Cortex*, 13(11), 1257–1269. doi: 10.1093/cercor/bhg097
- Mazyar, H., Berg, R. v. d., & Ma, W. J. (2012). Does precision decrease with set size? *Journal of Vision*, 12(6), 10. doi: 10.1167/12.6.10
- McConkie, G. W., & Currie, C. B. (1996). Visual stability across saccades while viewing complex pictures. *Journal of Experimental Psychology*, 22(3), 563–581. doi: 10.1037/0096-1523.22.3.563
- Miller, G. A. (1956). The magical number seven, plus or minus two: some limits on our capacity for processing information. *Psychological Review*, 63(2), 81–97.
- Myung, I. J., & Pitt, M. A. (1997). Applying Occam's razor in modeling cognition: A Bayesian approach. *Psychonomic Bulletin & Review*, 4(1), 79–95. doi: 10.3758/BF03210778
- Navarro, D. J., Pitt, M. A., & Myung, I. J. (2004). Assessing the distinguishability of models and the informativeness of data. *Cognitive Psychology*, 49(1), 47–84. doi: 10.1016/j.cogpsych.2003.11.001
- Neisser, U. (1967). *Cognitive Psychology* (1edition ed.). New York, NY: Appleton-Century-Crofts.
- Nelder, J. A., & Mead, R. (1965). A Simplex Method for Function Minimization. *The Computer Journal*, 7(4), 308–313. doi: 10.1093/comjnl/7.4.308
- Nobre, A. C., Coull, J. T., Maquet, P., Frith, C. D., Vandenberghe, R., & Mesulam, M. M. (2004). Orienting Attention to Locations in Perceptual Versus Mental Representations. *Journal of Cognitive Neuroscience*, 16(3), 363–373. doi: 10.1162/089892904322926700
- Noreen, D. L. (1981). Optimal decision rules for some common psychophysical paradigms. In S. Grossberg (Ed.), *Mathematical Psychology and Psychophysiology* (pp. 237–280). Psychology Press.
- Oberauer, K. (2002). Access to Information in Working Memory: Exploring the Focus of Attention. *Journal of Experimental Psychology*, 28(3), 411–421.

References

- Oberauer, K. (2006). Is the focus of attention in working memory expanded through practice? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 32(2), 197–214. doi: 10.1037/0278-7393.32.2.197
- Oberauer, K., & Bialkova, S. (2009). Accessing information in working memory: Can the focus of attention grasp two elements at the same time? *Journal of Experimental Psychology: General*, 138(1), 64–87. doi: 10.1037/a0014738
- Oberauer, K., & Hein, L. (2012). Attention to Information in Working Memory. *Current Directions in Psychological Science*, 21(3), 164–169. doi: 10.1177/0963721412444727
- Olson, I. R., & Jiang, Y. (2002). Is visual short-term memory object based? Rejection of the “strong-object” hypothesis. *Perception & Psychophysics*, 64(7), 1055–1067. doi: 10.3758/BF03194756
- Orhan, A. E., & Jacobs, R. A. (2013). A probabilistic clustering theory of the organization of visual short-term memory. *Psychological Review*, 120(2), 297–328. doi: 10.1037/a0031541
- Paivio, A., & Bleasdale, F. (1974). Visual short-term memory: A methodological caveat. *Canadian Journal of Psychology*, 28(1), 24–31. doi: 10.1037/h0081973
- Palmer, J. (1990). Attentional limits on the perception and memory of visual information. *Journal of Experimental Psychology: Human Perception and Performance*, 16(2), 332. Retrieved from <http://psycnet.apa.org/journals/xhp/16/2/332/>
- Palmer, J., Huk, A. C., & Shadlen, M. N. (2005). The effect of stimulus strength on the speed and accuracy of a perceptual decision. *Journal of Vision*, 5(5), 1–1. doi: 10.1167/5.5.1
- Palmer, J., Verghese, P., & Pavel, M. (2000). The psychophysics of visual search. *Vision research*, 40(10), 1227–1268.
- Pashler, H. (1988). Familiarity and visual change detection. *Perception & Psychophysics*, 44(4), 369–378. doi: 10.3758/BF03210419
- Pearson, B., Raškevičius, J., Bays, P. M., Pertzov, Y., & Husain, M. (2014). Working memory retrieval as a decision process. *Journal of Vision*, 14(2), 2. doi: 10.1167/14.2.2
- Pelli, D. G. (1991). Noise in the visual system may be early. *Computational models of visual processing*, 147–151.
- Pessoa, L., Gutierrez, E., Bandettini, P. A., & Ungerleider, L. G. (2002). Neural Correlates of Visual Working Memory: fMRI Amplitude Predicts Task Performance. *Neuron*, 35(5), 975–987. doi: 10.1016/S0896-6273(02)00817-6
- Peterson, L., & Peterson, M. J. (1959). Short-term retention of individual verbal items. *Journal of Experimental Psychology*, 58(3), 193–198. doi: 10.1037/h0049234
- Peterson, W. W., Birdsall, T., & Fox, W. (1954). The theory of signal detectability. *Trans-*

- actions of the IRE Professional Group on Information Theory, 4(4), 171–212. doi: 10.1109/TIT.1954.1057460
- Phillips, W. A. (1974). On the distinction between sensory storage and short-term visual memory. *Perception & Psychophysics*, 16(2), 283–290. doi: 10.3758/BF03203943
- Phillips, W. A., & Baddeley, A. D. (1971). Reaction time and short-term visual memory. *Psychonomic Science*, 22(2), 73–74. doi: 10.3758/BF03332500
- Pitt, M. A., Kim, W., & Myung, I. J. (2003). Flexibility versus generalizability in model selection. *Psychonomic Bulletin & Review*, 10(1), 29–44. doi: 10.3758/BF03196467
- Pitt, M. A., & Myung, I. J. (2002). When a good fit can be bad. *Trends in Cognitive Sciences*, 6(10), 421–425. doi: 10.1016/S1364-6613(02)01964-2
- Pitt, M. A., Myung, I. J., & Zhang, S. (2002). Toward a method of selecting among computational models of cognition. *Psychological Review*, 109(3), 472–491. doi: 10.1037/0033-295X.109.3.472
- Posner, M. I. (1967). Characteristics of visual and kinesthetic memory codes. *Journal of Experimental Psychology*, 75(1), 103. Retrieved from <http://psycnet.apa.org/journals/xge/75/1/103/>
- Posner, M. I., Boies, S. J., Eichelman, W. H., & Taylor, R. L. (1969). Retention of visual and name codes of single letters. *Journal of experimental psychology*, 79(1p2), 1. Retrieved from <http://psycnet.apa.org/journals/xge/79/1p2/1/>
- Posner, M. I., & Keele, S. W. (1967). Decay of Visual Information from a Single Letter. *Science*, 158(3797), 137–139. doi: 10.1126/science.158.3797.137
- Posner, M. I., Snyder, C. R., & Davidson, B. J. (1980). Attention and the detection of signals. *Journal of Experimental Psychology: General*, 109(2), 160–174. doi: 10.1037/0096-3445.109.2.160
- Pouget, A., Dayan, P., & Zemel, R. (2000). Information processing with population codes. *Nature Reviews Neuroscience*, 1(2), 125–132. doi: 10.1038/35039062
- Prinzmetal, W., Amiri, H., Allen, K., & Edwards, T. (1998). Phenomenology of Attention: 1. Color, Location, Orientation, and Spatial Frequency. *Journal of Experimental Psychology*, 24(1), 261–282.
- Prinzmetal, W., Nwachuku, I., Bodanski, L., Blumenfeld, L., & Shimizu, N. (1997). The Phenomenology of Attention. *Consciousness and Cognition*, 6(2–3), 372–412. doi: 10.1006/ccog.1997.0313
- Pylyshyn, Z. W., & Storm, R. W. (1988). Tracking multiple independent targets: Evidence for a parallel tracking mechanism. *Spatial Vision*, 3(3), 179–197. doi: 10.1163/156856888X00122
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85(2), 59–108. doi:

References

- 10.1037/0033-295X.85.2.59
- Ratcliff, R. (2002). A diffusion model account of response time and accuracy in a brightness discrimination task: Fitting real data and failing to fit fake but plausible data. *Psychonomic Bulletin & Review*, 9(2), 278–291. doi: 10.3758/BF03196283
- Ratcliff, R., & McKoon, G. (2007). The Diffusion Decision Model: Theory and Data for Two-Choice Decision Tasks. *Neural Computation*, 20(4), 873–922. doi: 10.1162/neco.2008.12-06-420
- Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. *Neural Computation*, 20(4), 873–922.
- Ratcliff, R., & Rouder, J. N. (1998). Modeling Response Times for Two-Choice Decisions. *Psychological Science*, 9(5), 347–356. doi: 10.1111/1467-9280.00067
- Ratcliff, R., & Rouder, J. N. (2000). A diffusion model account of masking in two-choice letter identification. *Journal of Experimental Psychology: Human Perception and Performance*, 26(1), 127–140. doi: 10.1037/0096-1523.26.1.127
- Ratcliff, R., & Smith, P. L. (2004). A Comparison of Sequential Sampling Models for Two-Choice Reaction Time. *Psychological Review*, 111(2), 333–367. doi: 10.1037/0033-295X.111.2.333
- Ratcliff, R., & Smith, P. L. (2010). Perceptual Discrimination in Static and Dynamic Noise: The Temporal Relation Between Perceptual Encoding and Decision Making. *Journal of Experimental Psychology*, 139(1), 70–94.
- Ratcliff, R., Thapar, A., & McKoon, G. (2001). The Effects of Aging on Reaction Time in a Signal Detection Task. *Psychology & Aging*, 16(2), 323–341.
- Ratcliff, R., Thapar, A., & McKoon, G. (2003). A diffusion model analysis of the effects of aging on brightness discrimination. *Perception & Psychophysics*, 65(4), 523–535. doi: 10.3758/BF03194580
- Raymond, J. E., Shapiro, K. L., & Arnell, K. M. (1992). Temporary suppression of visual processing in an RSVP task: An attentional blink? *Journal of Experimental Psychology: Human Perception and Performance*, 18(3), 849–860. doi: 10.1037/0096-1523.18.3.849
- Reddi, B. a. J., & Carpenter, R. H. S. (2000). The influence of urgency on decision time. *Nature Neuroscience*, 3(8), 827–830. doi: 10.1038/77739
- Reeves, A., & Sperling, G. (1986). Attention Gating in Short-Term Visual Memory. *Psychological Review*, 93(2), 180–206.
- Regan, D. (1985). Storage of spatial-frequency information and spatial-frequency discrimination. *Journal of the Optical Society of America A*, 2(4), 619. doi: 10.1364/JOSAA.2.000619

- Regan, D., & Beverley, K. I. (1985). Postadaptation orientation discrimination. *Journal of the Optical Society of America A*, 2(2), 147–155. doi: 10.1364/JOSAA.2.000147
- Rouder, J. N., Morey, R. D., Cowan, N., Zwilling, C. E., Morey, C. C., & Pratte, M. S. (2008). An assessment of fixed-capacity models of visual working memory. *Proceedings of the National Academy of Sciences*, 105(16), 5975–5979. doi: 10.1073/pnas.0711295105
- Rouder, J. N., Morey, R. D., Morey, C. C., & Cowan, N. (2011). How to measure working memory capacity in the change detection paradigm. *Psychonomic Bulletin & Review*, 18(2), 324–330. doi: 10.3758/s13423-011-0055-3
- Sakai, K., & Inui, T. (2002). A feature-segmentation model of short-term visual memory. *Perception*, 31(5), 579–589. doi: 10.1068/p3320
- Schwarz, G. (1978). Estimating the Dimension of a Model. *The Annals of Statistics*, 6(2), 461–464. Retrieved from <http://www.jstor.org/stable/2958889>
- Seung, H. S., & Sompolinsky, H. (1993). Simple models for reading neuronal population codes. *Proceedings of the National Academy of Sciences*, 90(22), 10749–10753. Retrieved from <http://www.pnas.org/content/90/22/10749>
- Sewell, D. K., Lilburn, S. D., & Smith, P. L. (2014). An Information Capacity Limitation of Visual Short-Term Memory. *Journal of Experimental Psychology: Human Perception and Performance*. doi: 10.1037/a0037744
- Sewell, D. K., Lilburn, S. D., & Smith, P. L. (in press). Object selection costs in visual working memory: A diffusion model analysis of the focus of attention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*.
- Sewell, D. K., & Smith, P. L. (2012). Attentional control in visual signal detection: Effects of abrupt-onset and no-onset stimuli. *Journal of Experimental Psychology: Human Perception and Performance*, 38(4), 1043–1068. doi: 10.1037/a0026591
- Shanno, D. F. (1970). Conditioning of quasi-Newton methods for function minimization. *Mathematics of Computation*, 24(111), 647–656. doi: 10.1090/S0025-5718-1970-0274029-X
- Shapiro, K. L., Raymond, J. E., & Arnell, K. M. (1994). Attention to Visual Pattern Information Produces the Attentional Blink in Rapid Serial Visual Presentation. *Journal of Experimental Psychology*, 20(2), 357–371.
- Shaw, M. L. (1980). Identifying attentional and decision-making components in information processing. In R. S. Nickerson (Ed.), *Attention and Performance VIII* (pp. 277–296). Hillsdale, NJ: Erlbaum.
- Shaw, M. L. (1982). Attending to multiple sources of information: I. The integration of information in decision making. *Cognitive Psychology*, 14(3), 353–409. doi:

References

- 10.1016/0010-0285(82)90014-7
- Shaw, M. L. (1984). Division of attention among spatial locations: A fundamental differences between detection of letters and the detection of luminance increments. In H. Bouma & D. G. Bouwhuis (Eds.), *Attention and performance X* (pp. 109–121). Hillsdale, NJ: L. Erlbaum Associates.
- Shepard, R. N., & Cermak, G. W. (1973). Perceptual-cognitive explorations of a toroidal set of free-form stimuli. *Cognitive Psychology*, 4(3), 351–377. doi: 10.1016/0010-0285(73)90018-2
- Sheu, C.-F., & Ratcliff, R. (1995). The Application of Fourier Deconvolution to Reaction Time Data: A Cautionary Note. *Psychological Bulletin*, 118(2), 285–299.
- Smith, P. L. (1990). A note on the distribution of response times for a random walk with Gaussian increments. *Journal of Mathematical Psychology*, 34(4), 445–459. doi: 10.1016/0022-2496(90)90023-3
- Smith, P. L. (1995). Psychophysically principled models of visual simple reaction time. *Psychological Review*, 102(3), 567–593. doi: 10.1037/0033-295X.102.3.567
- Smith, P. L. (2000a). Attention and luminance detection: Effects of cues, masks, and pedestals. *Journal of Experimental Psychology: Human Perception and Performance*, 26(4), 1401–1420. doi: 10.1037/0096-1523.26.4.1401
- Smith, P. L. (2000b). Stochastic Dynamic Models of Response Time and Accuracy: A Foundational Primer. *Journal of Mathematical Psychology*, 44(3), 408–463. doi: 10.1006/jmps.1999.1260
- Smith, P. L. (2010). From Poisson shot noise to the integrated Ornstein–Uhlenbeck process: Neurally principled models of information accumulation in decision-making and response time. *Journal of Mathematical Psychology*, 54(2), 266–283. doi: 10.1016/j.jmp.2009.12.002
- Smith, P. L. (2015). The Poisson shot noise model of visual short-term memory and choice response time: Normalized coding by neural population size. *Journal of Mathematical Psychology*, 66, 41–52. doi: 10.1016/j.jmp.2015.03.007
- Smith, P. L. (2016). Diffusion theory of decision making in continuous report. *Psychological Review*.
- Smith, P. L., Lee, Y.-E., Wolfgang, B. J., & Ratcliff, R. (2009). Attention and the detection of masked radial frequency patterns: Data and model. *Vision Research*, 49(10), 1363–1377. doi: 10.1016/j.visres.2008.04.024
- Smith, P. L., Lilburn, S. D., Corbett, E., Sewell, D. K., & Kyllingsbæk, S. (under review). The attention weighted sample size model of visual short-term memory: Attention capture predicts resource allocation and memory load. *Cognitive Psychology*.

- Smith, P. L., & Ratcliff, R. (2004). Psychology and neurobiology of simple decisions. *Trends in Neurosciences*, 27(3), 161–168. doi: 10.1016/j.tins.2004.01.006
- Smith, P. L., & Ratcliff, R. (2009). An integrated theory of attention and decision making in visual signal detection. *Psychological Review*, 116(2), 283–317. doi: 10.1037/a0015156
- Smith, P. L., Ratcliff, R., & Sewell, D. K. (2014). Modeling perceptual discrimination in dynamic noise: Time-changed diffusion and release from inhibition. *Journal of Mathematical Psychology*, 59, 95–113. doi: 10.1016/j.jmp.2013.05.007
- Smith, P. L., Ratcliff, R., & Wolfgang, B. J. (2004). Attention orienting and the time course of perceptual decisions: response time distributions with masked and unmasked displays. *Vision Research*, 44(12), 1297–1320. doi: 10.1016/j.visres.2004.01.002
- Smith, P. L., & Sewell, D. K. (2013). A competitive interaction theory of attentional selection and decision making in brief, multielement displays. *Psychological Review*, 120(3), 589–627. doi: 10.1037/a0033140
- Smith, P. L., Sewell, D. K., & Lilburn, S. D. (2015). From shunting inhibition to dynamic normalization: Attentional selection and decision-making in brief visual displays. *Vision Research*, 116, 219–240. doi: 10.1016/j.visres.2014.11.001
- Sorkin, R. D. (1962). Extension of the Theory of Signal Detectability to Matching Procedures in Psychoacoustics. *The Journal of the Acoustical Society of America*, 34(11), 1745–1751. doi: 10.1121/1.1909115
- Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs: General and Applied*, 74(11), 1–29. doi: 10.1037/h0093759
- Sperling, G. (1963). A Model for Visual Memory Tasks. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 5(1), 19–31. doi: 10.1177/001872086300500103
- Sperling, G., & Weichselgartner, E. (1995). Episodic theory of the dynamics of spatial attention. *Psychological Review*, 102(3), 503–532. doi: 10.1037/0033-295X.102.3.503
- Sperling, H. G., & Jolliffe, C. L. (1965). Intensity-Time Relationship at Threshold for Spectral Stimuli in Human Vision. *Journal of the Optical Society of America*, 55(2), 191–198. doi: 10.1364/JOSA.55.000191
- Sternberg, S. (1969). The discovery of processing stages: Extensions of Donders' method. *Acta Psychologica*, 30, 276–315. doi: 10.1016/0001-6918(69)90055-9
- Swets, J. A. (1961). Is There a Sensory Threshold? *Science*, 134(3473), 168–177. Retrieved from <http://www.jstor.org/stable/1708294>
- Swets, J. A., Tanner, W. P., & Birdsall, T. G. (1961). Decision Processes In Perception. *Psychological Review*, 68(5), 301–340. doi: 10.1037/h0040547

References

- Tanner, W. P. (1956). Theory of Recognition. *The Journal of the Acoustical Society of America*, 28(5), 882–888. doi: 10.1121/1.1908504
- Tanner, W. P., & Swets, J. (1954b). The human use of information—I: Signal detection for the case of the signal known exactly. *Transactions of the IRE Professional Group on Information Theory*, 4(4), 213–221. doi: 10.1109/TIT.1954.1057461
- Tanner, W. P., & Swets, J. A. (1953). *A new theory of visual detection* (Tech. Rep.). The University of Michigan: The University of Michigan.
- Tanner, W. P., & Swets, J. A. (1954a). A decision-making theory of visual detection. *Psychological Review*, 61(6), 401–409. doi: 10.1037/h0058700
- Taylor, M. M., Lindsay, P. H., & Forbes, S. M. (1967). Quantification of shared capacity processing in auditory and visual discrimination. *Acta Psychologica*, 27, 223–229. doi: 10.1016/0001-6918(67)99000-2
- Thomas, J. P. (1985). Detection and identification: how are they related? *Journal of the Optical Society of America A*, 2(9), 1457–1467. doi: 10.1364/JOSAA.2.001457
- Thomas, J. P., & Gille, J. (1979). Bandwidths of orientation channels in human vision. *Journal of the Optical Society of America*, 69(5), 652–660. doi: 10.1364/JOSA.69.000652
- Thurstone, L. L. (1927). Psychophysical Analysis. *The American Journal of Psychology*, 38(3), 368–389. doi: 10.2307/1415006
- Todd, J. J., & Marois, R. (2004). Capacity limit of visual short-term memory in human posterior parietal cortex. *Nature*, 428(6984), 751–754. doi: 10.1038/nature02466
- Townsend, J. T. (1972). Some Results Concerning the Identifiability of Parallel and Serial Processes. *British Journal of Mathematical and Statistical Psychology*, 25(2), 168–199. doi: 10.1111/j.2044-8317.1972.tb00490.x
- Townsend, J. T. (1990). Serial vs. Parallel Processing: Sometimes They Look like Tweedle-dum and Tweedledee but They Can (And Should) be Distinguished. *Psychological Science*, 1(1), 46–54. Retrieved from <http://www.jstor.org/stable/40062391>
- Townsend, J. T., & Ashby, F. G. (1983). *The Stochastic Modeling of Elementary Psychological Processes*. Cambridge: Cambridge University Press.
- Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, 12(1), 97–136. doi: 10.1016/0010-0285(80)90005-5
- Treisman, A. M., Kahneman, D., & Burkell, J. (1983). Perceptual objects and the cost of filtering. *Perception & Psychophysics*, 33(6), 527–532. doi: 10.3758/BF03202934
- Trick, L. M., & Pylyshyn, Z. W. (1994). Why Are Small and Large Numbers Enumerated Differently? A Limited-Capacity Preattentive Stage in Vision. *Psychological Review*, 101(1), 80–102.
- Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: The

- leaky, competing accumulator model. *Psychological Review*, 108(3), 550–592. doi: 10.1037/0033-295X.108.3.550
- van den Berg, R., Awh, E., & Ma, W. J. (2014a). Factorial comparison of working memory models. *Psychological Review*, 121(1), 124–149. doi: 10.1037/a0035234
- van den Berg, R., Awh, E., & Ma, W. J. (2014b). Factorial comparison of working memory models. *Psychological Review*, 121(1), 124–149. doi: 10.1037/a0035234
- van den Berg, R., Shin, H., Chou, W.-C., George, R., & Ma, W. J. (2012). Variability in encoding precision accounts for visual short-term memory limitations. *Proceedings of the National Academy of Sciences*, 109(22), 8780–8785. doi: 10.1073/pnas.1117465109
- Vogel, E. K., & Machizawa, M. G. (2004). Neural activity predicts individual differences in visual working memory capacity. *Nature*, 428(6984), 748–751. doi: 10.1038/nature02447
- Vogel, E. K., Woodman, G. F., & Luck, S. J. (2001). Storage of features, conjunctions, and objects in visual working memory. *Journal of Experimental Psychology: Human Perception and Performance*, 27(1), 92–114. doi: 10.1037/0096-1523.27.1.92
- Vogel, E. K., Woodman, G. F., & Luck, S. J. (2006). The time course of consolidation in visual working memory. *Journal of Experimental Psychology: Human Perception and Performance*, 32(6), 1436–1451. doi: 10.1037/0096-1523.32.6.1436
- Wang, X.-J. (2001). Synaptic reverberation underlying mnemonic persistent activity. *Trends in Neurosciences*, 24(8), 455–463. doi: 10.1016/S0166-2236(00)01868-3
- Watt, R. J. (1987). Scanning from coarse to fine spatial scales in the human visual system after the onset of a stimulus. *Journal of the Optical Society of America A*, 4(10), 2006. doi: 10.1364/JOSAA.4.002006
- Wei, Z., Wang, X.-J., & Wang, D.-H. (2012). From Distributed Resources to Limited Slots in Multiple-Item Working Memory: A Spiking Network Model with Normalization. *The Journal of Neuroscience*, 32(33), 11228–11240. doi: 10.1523/JNEUROSCI.0735-12.2012
- Weichselgartner, E., & Sperling, G. (1987). Dynamics of automatic and controlled visual attention. *Science*, 238(4828), 778–780. doi: 10.1126/science.3672124
- Wheeler, M. E., & Treisman, A. M. (2002). Binding in short-term visual memory. *Journal of Experimental Psychology: General*, 131(1), 48–64. doi: 10.1037/0096-3445.131.1.48
- Wilken, P., & Ma, W. J. (2004). A detection theory account of change detection. *Journal of Vision*, 4(12). doi: 10.1167/4.12.11
- Wilkinson, F., Wilson, H. R., & Habak, C. (1998). Detection and recognition of radial frequency patterns. *Vision Research*, 38(22), 3555–3568. doi:

References

- 10.1016/S0042-6989(98)00039-X
- Woodman, G. F., & Luck, S. J. (1999). Electrophysiological measurement of rapid shifts of attention during visual search. *Nature*, 400(6747), 867–869. doi: 10.1038/23698
- Woodman, G. F., & Luck, S. J. (2003). Serial Deployment of Attention During Visual Search. *Journal of Experimental Psychology*, 29(1), 121–138.
- Woodman, G. F., & Vecera, S. P. (2011). The cost of accessing an object's feature stored in visual working memory. *Visual Cognition*, 19(1), 1–12. doi: 10.1080/13506285.2010.521140
- Xu, Y., & Chun, M. M. (2006). Dissociable neural mechanisms supporting visual short-term memory for objects. *Nature*, 440(7080), 91–95. doi: 10.1038/nature04262
- Zhang, W., & Luck, S. J. (2008). Discrete fixed-resolution representations in visual working memory. *Nature*, 453(7192), 233–235. doi: 10.1038/nature06860
- Zhang, W., & Luck, S. J. (2009). Sudden Death and Gradual Decay in Visual Working Memory. *Psychological Science*, 20(4), 423–428. doi: 10.1111/j.1467-9280.2009.02322.x
- Zipser, D., Kehoe, B., Littlewort, G., & Fuster, J. (1993). A spiking network model of short-term active memory. *The Journal of Neuroscience*, 13(8), 3406–3420. Retrieved from <http://www.jneurosci.org/content/13/8/3406>