The Role of Attention in Working Memory

Edwin S. Dalmaijer

Dissertation for the degree of Doctor of Philosophy

Experimental Psychology

University of Oxford

Supervision:

Prof. Masud Husain

Dr. Mark Stokes

1. Table of contents

The Role of Attention in Working Memory

2. Abstract

Attention is usually defined as the ability to select a limited amount of information from the abundance that we are continuously faced with, and short-term memory is often described as a temporary storage for this selected information. However, this description does not do justice to the complexities of the interplay between attention and short-term memory. This thesis explores how attention acts in and on short-term memory. It describes how attentionally selected information flows into short-term memory, and finds that encoding of several items can occur in parallel. It also demonstrates that encoding can occur in a serial fashion too, with behaviourally relevant items being selected and encoded first, and other items second. This thesis not only investigates the initial allocation of these resources during encoding, but also the re-allocation of short-term memory resources during eye movements. Specifically, it fails to replicate the finding that resources are re-allocated to saccade targets, but it does find that the location of items is stored at a higher precision than items' other features are. Finally, this thesis describes how cancellation tasks (multi-target visual search tasks in which participants have to find and cross out targets among distractors) can be used to measure spatial attention and short-term memory in healthy adults and patient groups. Data is presented to show that cancellation performance can potentially be improved in subtle ways in stroke patients with hemispatial neglect syndrome, using the noradrenergic agonist guanfacine. In sum, attention acts as a gateway to and (re-)allocator of short-term memory resources, and is an important target for pharmacological interventions in sensitive groups.

2.1. A note on the course of this DPhil

The work contained in this thesis is perhaps more diverse than one would normally expect. Although all the work pertains to attention and short-term memory, the employed methods and populations differ in particular between Chapters 4-6 on the one hand, and Chapters 7-8 on the other. Specifically, the former present fundamental work on the nature of attention and short-term memory, whereas the latter describe work on neuropsychological assessment and a pharmacological intervention. The EU Marie Curie FP7 grant titled INDIREA (nr. 606901) that funded my DPhil was focussed on finding new tools and interventions to diagnose and rehabilitate inattention in different patient groups. The grant was headed by the late Prof. Glyn Humphreys, and the work presented in Chapters 7-8 were very much a part of its intended direction. After the unfortunate passing of Prof. Humphreys, my work drifted away from patients, as is evidenced by the fundamental work presented in Chapters 4-6. The non-chronological ordering of the chapters is for the benefit of the reader, who will hopefully agree that reading on the architecture of attention and short-term memory first improves their appreciation of the later chapters on assessment and rehabilitation of attention and short-term memory in stroke patients.

3. Introduction

As every cognitive psychologist can tell you, William James once famously said that “everybody knows what attention is” (James, 1890). This is an appealing quote, because attention is intuitively easy to understand. Despite the abundance of information in the world, healthy adult humans are clearly able to focus on only a subset of that information. As James himself remarked, attention can be focussed on information from the outside world, but can also be focussed on information that exists only in the mind. That is: attention can be pointed inward, to interact with information that is stored in short-term memory. Although we all know this from subjective experience, it is important to go beyond an anecdotal understanding of the matter. In this thesis, I will be exploring the exact mechanism of how attention selects and (re-)encodes information into short-term memory, how these processes can be measured with a patient-friendly test, and how stroke patients who suffer from hemispatial neglect syndrome (which is characterised by deficits in spatial attention and short-term memory) can potentially be helped by pharmacological intervention.

3.1. The attentional spotlight

One influential metaphor for how attention works, is the *attentional spotlight*. Much like an actual spotlight, attention is said to highlight one particular location in the world at any one time (Crick, 1984; Duncan, 1984; Posner, 1980; Posner, Snyder, &

Davidson, 1980), and it can even zoom in or out to cover a larger or smaller area (Eriksen & St. James, 1986). The point here is straightforward: An individual has a limited attentional resource that acts much like an actual spotlight. It can be directed to light up a single location, and when the attended area around this location is extended the light becomes more diffuse. Although this theory of attention is of a rather descriptive level, it does make a strong prediction: Somewhere during the process of perception, the signal from one particular location is boosted.

That attention can in fact boost the processing of visual information in humans has been demonstrated with positron emission tomography in the visual cortex in general (Corbetta, Miezin, Dobmeyer, Shulman, & Petersen, 1991; Corbetta, Miezin, Shulman, & Petersen, 1993); in area V5/MT that is primarily sensitive to movement, using single-cell recordings in monkeys (Motter, 1994), and using magnetic resonance imaging (MRI) in humans (O'Craven, Rosen, Kwong, Treisman, & Savoy, 1997; Treue & Maunsell, 1996); and even in areas as early in the visual hierarchy as V1 (Brefczynski & DeYoe, 1999). In the last study, attention was modulated by directing participant's attention to central vision at the start of a trial, and at increasing eccentricity during the course of a trial. Participants kept their gaze fixated on the display centre, as the attentional cue moved along from central to peripheral vision. This attentional manipulation, during which a large stimulus was visible, was compared to a condition in which only a part of the stimulus was presented in a similar central-to-peripheral progression. The blood-oxygenation level dependent (BOLD) response as measured through functional magnetic resonance imaging (fMRI) followed the presented part of the stimulus in the visual condition, but importantly also followed the cue in the attentional condition. This illustrated the modulation of activity in a brain

area as early as V1, thereby effectively demonstrating a core prediction of the attentional spotlight hypothesis: Attention boosts perception at the attended location.

3.2. Moving beyond the spotlight metaphor

One prediction of the attentional spotlight hypothesis is that attention only boosts perception around a single location. If more than one stimulus is displayed, only one can be attentionally selected, or the attentional focus will have to 'zoom out' to incorporate all stimuli.

In a clever design, Kramer and Hahn (1995) tested exactly this. They showed two target stimuli in pre-cue boxes on an imaginary circle around the display centre. Participants were required to report the identity of these stimuli. The clever aspect of the task was that distractor stimuli were presented in between the target stimuli. In half of all trials, the pre-cue boxes appeared first, followed by the targets and distractors 150 milliseconds later, followed by masks 60 ms later. In the other half, placeholders (mask stimuli) appeared together with the cue boxes, followed by the targets and distractors 150 milliseconds later, and again followed by the masks 60 ms after stimulus onset. After mask onset, participants were asked to indicate whether the two target stimuli were the same or different. Distractor stimuli could either be the same or different from the target stimuli. Crucially, in the trials where placeholders appeared first, whether the distractors matched the targets or not did not affect participant's response time or accuracy in their judgement of whether the targets were the same or different. In addition, when the targets and distractors had a sudden onset (i.e. without the initial

placeholders), the distractors still didn't affect participants' response times or accuracy when the targets matched. The only condition in which participants were affected by the distractors, was when they appeared with an onset (i.e. without the initial placeholders), and when the targets did not match: Here, participants were slower and less accurate to indicate a mismatch if the distractors were different compared to when they were the same. This suggests that participants could attend to both locations without being distracted by what happened in between these locations, unless what happened in between the locations happened with a sudden onset. (Kramer & Hahn, 1995)

If the spotlight of attention is indivisible, it would have stretched out to encompass both target stimuli in the previous paragraph, and therefore also the distractors in between. However, the authors report less distraction from the in-between stimuli than one would expect in this scenario, and therefore Kramer and Hahn argued that the 'attentional beam had split in two'. One alternative explanation is that participants shifted their attention rapidly between the two stimuli, but at a stimulus exposure duration of 60 ms, this seems unlikely. In fact, presentation times under 200 milliseconds are generally considered to be too short to attend to one item, process it, and then attend to another item (Bundesen, 1990; Duncan, Ward, & Shapiro, 1994; Reeves & Sperling, 1986; Weichselgartner & Sperling, 1987).

In a follow-up study, the same authors demonstrated that indeed attention could be split between locations, and sustained there as long as distractors did not appear between the attended locations (Hahn & Kramer, 1998). In addition, they demonstrated that the two non-contiguous attended locations could lie within the same hemifield, thereby ruling out that their results were caused by each hemifield having its own

attentional spotlight.

Several years later, Kramer and Hahn's psychophysical work was followed up with an fMRI study of even greater elegance. Here, McMains and Somers presented rapid serial visual processing (RSVP) streams on five locations in the visual field: centrally (within the fovea), and in each quadrant. An RSVP stream consists of rapidly alternating (every 173 milliseconds) letters and numbers, save from the central RSVP stream, which was digits only. In one condition, participants only had to monitor one RSVP stream (bottom left), and to respond when a digit appeared and matched or did not match a target digit. In the other condition, participants had to simultaneously attend to two RSVP streams (top left and bottom right), and to respond when digits in both streams matched or not. Crucially, if there is only one attentional spotlight, BOLD activity in the attend-2 condition would span across neural populations that have their receptive fields in the fovea. In contrast, if attention can be divided, then attention-modulated BOLD activity is expected to occur only in those neural populations that have their receptive fields in the attended quadrants. McMains and Somers report modulation only in parts of the cortex that had been identified using retinotopic mapping as being sensitive for the quadrants where the attended RSVP streams occurred, but not in areas receptive for the fovea. This was not only true when RSVP streams were attended in opposite hemifields, but also when they were presented within one hemifield. (McMains & Somers, 2004)

Notably, a highly similar result was obtained using steady-state visual evoked potentials (SSVEP) which occur in electroencephalographic (EEG) signals as a response to several stimuli, each being presented at a unique frequency. SSVEPs for two items in

non-contiguous locations were enhanced by participants' attending to them, without intermediately presented stimuli being affected (Müller, Malinowski, Gruber, & Hillyard, 2003).

In sum, the attentional spotlight in its strictest interpretation is a bad metaphor for attention. While attention definitely boosts the processing of information from the attended location, there can be more than one non-contiguous attended location. In other words: Attention can be directed to multiple sources of information, in parallel.

3.3. Feature-based attention and the 'item'

A different way of thinking about attentional boosting of perceptual information is not a boost per location, but rather per item. One item can be defined as a grouping of features that correspond to the same source. For example, a red square can be an item, and so can a green bar. Aside from its shape, the bar has three defining characteristics that set it apart from other bars: its location, its colour, and its orientation.

Kastner and colleagues presented four items (complex figures of 2x2 degrees of visual angle) either sequentially or simultaneously, and argued that if items are competing for attention, their overall BOLD activity should be lower during simultaneous presentation compared to sequential presentation due to mutual inhibition. This was indeed found, and in addition it was found that attending to one of the stimuli reduced the suppressive effects of other stimuli. In sum, attending to an item reduces the suppressive effects of other items on the attended item. (Kastner, De Weerd, Desimone, & Ungerleider, 1998)

The previous paragraph treats ‘attention to items’ in much the same way as earlier paragraphs described ‘attention to locations’. In a practical sense, the distinction between locations and items seems almost arbitrary. However, it appears attention can be directed to non-location based features such as colour and orientation. When participants are asked to view two centrally presented Gabor stimuli, and asked to judge their orientation (clockwise or counter-clockwise compared to some reference), they were more sensitive (more likely to notice and less likely to make a false alarm) to gratings that could randomly appear in the periphery (Rossi & Paradiso, 1995). This demonstrates that when a participant is attending to an item with a particular grating, there is an attentional bias for that grating regardless of where it occurs in the visual field. Similar studies have replicated this effect for different features (orientation, motion, and colour) and task types (A. Cohen & Magen, 1999; Kumada, 2001; Sàenz, Buraĉas, & Boynton, 2003). For a review, see (Maunsell & Treue, 2006). It should be noted that the theoretical notion of feature-based attention was already incorporated in a mathematical model of visual attention that preceded these studies by five years (Bundesen, 1990), on which more will follow later.

Another interesting note on feature-based attention comes from a recent study in which participants held an item of a particular colour in short-term memory. In between seeing the item and being presented a memory probe, participants were presented with a binocular stimulus. In one eye a series of alternating ‘Mondrian’ stimuli (an array of brightly coloured rectangles) were presented, and in the other eye an item in a similar or dissimilar colour to the memorised item. The result was that participants initially saw only the frequently changing Mondrians, but slowly noticed the item breaking through. This method is called ‘breaking continuous flash suppression’, and is reviewed by

(Gayet, Van der Stigchel, & Paffen, 2014). Crucially, when the suppressed item's colour matched that of the memorised colour, the item would break through earlier compared to when the suppressed item's colour was seen before but not memorised or when it was unrelated to the current trial (Gayet, Paffen, & Van der Stigchel, 2013).

3.4. Items in short-term memory

The previously described prioritised access to conscious awareness for new information that matches the content of short-term memory (Gayet et al., 2013) already hinted at an intimate interplay between the content of short-term memory and attentional selection. In addition, the aforementioned notion of items is not only useful in conceptualising attention, but it is also a useful construct to describe how attentionally selected information ends up in short-term memory.

As described before, items can be viewed as connected stimulus features. According to the highly influential *feature integration theory* (Treisman, 1982; Treisman & Gelade, 1980), binding different features together in one item requires attentional processing. Specifically, the theory proposes that the integration of conjunctions that consist of more than one visual feature requires sequential attention. For example in a visual search task with such stimuli, attention would be deployed serially, binding the features of each of the stimuli (e.g. orientation and colour), until the one stimulus that stands out in either orientation or colour is found.

It is important to note that connected stimulus features are not only conceptualised in the context of items, but also in the context of hierarchical

connections between memorised stimulus features (Brady, Konkle, & Alvarez, 2011). The distinction is not quite as subtle as it might seem at first. Items are bound objects with for example a location, an orientation, and a colour; but their features are bound together, and not to features in other items. However, the work I reviewed previously is clear in its suggestion that attending to and memorising an item with a particular feature has consequences that reach beyond that item, and boost performance for the detection and memory of items with the same feature. Thus it can be conceived that memorised features can interact, and that items are temporary bonds across different feature layers. In other words: Items can, and perhaps should, be conceptualised as hierarchical feature bundles (Brady et al., 2011).

3.5. A brief history of visual short-term memory

Before discussing how items or ‘hierarchical feature bundles’ make their way into short-term memory after being attentionally selected, it is important to explore what short-term memory is. As per usual in cognitive science, in the beginning there was William James. This chapter quoted him on the concept of attention, but the same resource also outlines a distinction between what he referred to as ‘primary’ and ‘secondary’ memory (James, 1890). In James’ view, primary memory is for an immediate sort of information such as the topic of a current conversation. Secondary memory is for information that might not currently be in the mind, but is available to be called upon when needed; for example the name of the song that topped Billboard’s ‘Hot Country Songs’ list on 26 November 2011 (*Sparks Fly* by Taylor Swift) and other

potentially useful facts about the world. These concepts survived in contemporary scientific literature, in remarkably similar conceptualisations, as ‘short-term memory’ and ‘long-term memory’.

In their seminal work, Baddeley and Hitch describe short-term memory as three boxes that are connected by arrows: The visuospatial sketchpad for remembering non-verbal information, the phonological loop for continuously repeating memorised verbal information, and the central executive for moving information into and out of either store (Baddeley & Hitch, 1974).

Although Baddeley’s conceptualisation of short-term memory is entirely descriptive and lacks predictive power, it remains a popular tool to shape thinking about short-term memory. Others have since produced more testable hypotheses on how information is stored in short-term memory. Perhaps the most influential model in the current literature is the slot model which postulates that short-term memory consists of a limited and natural number of ‘slots’ that can each hold one item (Luck & Vogel, 1997, 2013a). At the core of this model is the idea that short-term memory is a discrete and quantised resource. In other words, short-term memory can only hold so many items, and cannot store any more items once it is filled up. How many items short-term memory can hold, is said to differ between people, with an average of around four items of visual information (Luck & Vogel, 2013a; Vogel, Woodman, & Luck, 2001). (A distinction is again made between verbal and non-verbal information.)

An important prediction of this model is that participants should be as accurate when remembering four items with two defining features (e.g. location and colour) as when remembering four items with three defining features (e.g. location, colour, and

orientation), as they remember on a per-item basis. This has indeed been documented by slot model founders (Awh, Barton, & Vogel, 2007; Luck & Vogel, 2013a; Vogel et al., 2001). It should be noted that this does not necessarily hold for more complex items that are defined by a greater number of features (Alvarez & Cavanagh, 2004; Eng, Chen, & Jiang, 2005), although this is often explained away as a consequence of the higher similarity between complex stimuli (Fukuda, Awh, & Vogel, 2010).

Another important prediction of the slot model is that at capacity, all memorised items should have roughly the same precision, regardless of how many additional stimuli were presented in the memory array. For example, a participant with a short-term memory capacity of four items will remember four items from a 4-item memory array, but also four items from an 8-item memory array. If one could isolate which four items were in the participant's short-term memory, they should be able to demonstrate whether the precision of item representations is the same for the 4-item and 8-item arrays. Two pieces of information support this prediction: Recall precision for memorised items, and evidence from an event-related potential (ERP) in electroencephalographic (EEG) signal names the 'contra-lateral delay activity' (CDA). Both are further explored in the next paragraphs.

The first comes from a highly influential study by Zhang and Luck, in which participants were presented with 1, 2, 3, or 6 items (squares defined by their location and colour) in each trial. Using a clever new method in which participants could indicate on a continuous colour wheel what colour they thought one randomly probed item had, Zhang & Luck obtained a continuous distribution of errors. This distribution could be characterised by the weighted sum of a Von Mises distribution (analogous to a

normal distribution but with connected extremes) and a uniform distribution. When a probed item is in memory, the probability that the associated response is of a low error is high, and adheres quite well to a normal distribution. When a probed item is not in memory, every response is equally likely, which is a pattern that adheres to a uniform distribution. Zhang & Luck estimated the weight of each of the distributions, and were thus able to estimate separately the representations of items that were in memory and those that were not. Their measure of representational precision was the standard deviation of the Von Mises curve, and they found that this metric did not differ between the 3 and 6-item conditions, thus confirming a crucial slot model prediction. (Zhang & Luck, 2008)

Another interesting paper by the same authors tested the same hypothesis from a different angle. Here, participants were encouraged to trade off precision for individual items and the total number of remembered items. Specifically, participants were encouraged to remember four items with higher or lower precision by being probed with the full colour space or with only a few discrete options when providing their response. In different experiments, participants were encouraged to increase the number of items they remembered by being offered rewards for different ranges of errors (short range for high precision, wide range for low precision). None of the manipulations, save from an immediate reward, changed the estimated number of items that participants remembered, nor their estimated representational precision. (Zhang & Luck, 2011)

When individuals retain an item in short-term memory, this leaves a trace in the electroencephalography (EEG) around the posterior parietal, lateral occipital, and posterior temporal parts of their scalp contra-lateral to where the memorised item was

presented. This signal can be measured my measuring the EEG activity from the onset of the memory array and through a retention interval. To optimise signal, memory arrays can be presented left or right of fixation (with another array on the other side that can be ignored). The EEG signal is averaged across all trials to compute the event-related potential (ERP) that is named contra-lateral delay activity (CDA). Interestingly, the amplitude of CDA increases as a function of the number of items in memory, but reaches an asymptote roughly around the average memory capacity. This is again in line with a crucial slot model prediction. (Vogel & Machizawa, 2004)

Curiously, the asymptotes described above only appear when capacity is reached, but not before. The precision of recall decreases from 1 to 4 items (Zhang & Luck, 2008); and the CDA amplitude increases from 1 to 3 items, and this increase correlates with individuals' short-term memory capacity (Vogel & Machizawa, 2004). This indicates that when not all slots are filled, participants memorise fewer items at higher resolutions. If slots were truly discrete, this should not happen: One item should be stored in one slot, and therefore be remembered at the precision that that slot allows for. That items from under-capacity-sized arrays are remembered at a higher precision conflicts with the core of the slot model: independent and quantised slots.

To circumvent their own logic, slot model theorists have proposed that slots can 'double up' so that the same item can be stored in multiple slots (Luck & Vogel, 2013a; Zhang & Luck, 2008). An alternative explanation is that information in slots competes with each other, and therefore the fewer items are present, the less interference occurs, and thus the higher representational precision is. Regardless of which theoretical patch you prefer, it hurts the original premise of independent slots, and greatly infringes on

the parsimony of the original theory.

A different way of conceptualising short-term memory is as a limited but fluid resource that can be flexibly divided over several items (Bays, Catalao, & Husain, 2009; Bays & Husain, 2008; Ma, Husain, & Bays, 2014; Wilken & Ma, 2004). At its core, the prediction that this resource model makes is simple: Short-term memory resources are limited and finite, and thus there should be a trade-off between the number of items participants can remember and the representational precision of each memorised item.

One crucial prediction by this model is already illustrated in aforementioned research: With low numbers of items, the number of remembered items negatively correlates with the precision of recall (Bays et al., 2009; Bays & Husain, 2008; Fukuda et al., 2010; Luck & Vogel, 1997, 2013a; Vogel & Machizawa, 2004; Vogel et al., 2001; Zhang & Luck, 2008).

A point of crucial debate, however, is recall precision at higher numbers of recalled items. A previous paragraph outlines how recall precision is the same for memorised items from arrays consisting of 3 or 6 items (Zhang & Luck, 2008). However, in a highly similar experiment that employs highly similar modelling (weighted sum of a cumulative normal distribution and uniform distribution fitted to binary change-detection data), this asymptote is not found (Bays & Husain, 2008).

The discrepancy between both findings is addressed in a replication of the influential Zhang and Luck study (Zhang & Luck, 2008) that employed a subtly different analysis method (Bays et al., 2009). Where Zhang and Luck only considered two options, remembering and not remembering, Bays and colleagues pointed out a third option: The possibility that individuals confused items. When probed with the

location of one item, perhaps sometimes participants actually respond with the colour of another item. Considering the previously described non-independent nature of items in short-term memory (Brady et al., 2011), these swapping errors seem a natural consequence of the system's architecture. Indeed, when another component was incorporated into Zhang and Luck's model to account for potential swaps of location and colour, Bays and colleagues found the predicted negative relationship between the number of remembered items and the precision of each item's representation (Bays et al., 2009).

In sum, when accounting for swap errors, there is an inverse relationship between the number of items in short-term memory and their precision, regardless of the number of presented items. This is in fundamental disagreement with predictions derived from the slot model, but does agree with resource model predictions.

3.6. A brief future of visual short-term memory

As always in science, understanding moves forward. Both the slot and resource model are most likely highly simplified versions of how the brain actually operates, and one would be hard-pressed to find serious researchers who actually believe the brain has discrete neuronal slots (but see: (Deco & Rolls, 2008; Lisman & Idiart, 1995; Luck & Vogel, 2013a; Wei, Wang, & Wang, 2012)), or a set of all-purpose neurons to flexibly distribute over to-be-memorised items. At the scale of behaviour these levels of explanation might be enough, but more detail is needed to more accurately describe the underlying mechanism of short-term memory maintenance.

Two interesting findings illustrate the fleeting nature of short-term memory. The first is that memorised items are inherently unstable: There is evidence of their gradual decline from memory (Reitman, 1974), for their abrupt and complete disappearance (Zhang & Luck, 2009), and for the idea that maintaining items in short-term memory requires attentional effort (Zokaei, Heider, & Husain, 2014). Another interesting finding is that short-term memory and long-term memory seem to share a precision threshold below which items cannot be remembered at all, which is a potential indicator of a lower-bound on representational precision (Brady, Konkle, Gill, Oliva, & Alvarez, 2013). Taken together, these findings suggest that representations are subject to noise, and that they are lost when the signal-noise ratio becomes to low.

One model that allows for such a dynamic is a bump-attractor model in which a feature (in this case the location of a saccade target) in short-term memory is represented as the summed activity of neural populations that code for that specific feature, and in which (through inter-neural dynamics) peak activity travels across populations (Wimmer, Nykamp, Constantinidis, & Compte, 2014). A different way of describing the process of maintenance is through the encoding strength of neural populations coding for memorised items' features: Increase in the number of memorised items decreases the relative signal strength of associated neuronal populations, and thus reduces the signal-noise ratio, thereby increasing the likelihood of errors and forgetting (Bays, 2014, 2015).

The most promising current models assume conjunctions of features can be encoded by populations of neurons that are sensitive to a particular range within a feature (i.e. they have a preferred orientation, colour, or location), combined with

neurons that can code for combinations of features (Manohar, Zokaei, Fallon, Vogels, & Husain, 2017; Schneegans & Bays, 2017a). This approach is confronted by the fundamental issue that the brain has a limited amount of neurons, and that the number of conjunction neurons required to represent a feature space exponentially increases with the number of features. However, Schneegans and Bays propose a clever way to account for this, by proposing (and elegantly demonstrating) that an item's location can act as a special feature to which all its other features are bound (Schneegans & Bays, 2017a). This elegant way of binding information significantly reduces the required neuronal infrastructure, and fits with the notion of items as hierarchical feature bundles (Brady et al., 2011).

Within this thesis, the notion of location as a special feature is touched upon in chapter 6: *Dynamic re-allocation of visual short-term memory resources to saccade targets?* In short, the results described in that chapter support the view that memory for item location is of a far greater precision than memory for other features (colour and orientation).

3.7. Short-term memory versus working memory

A popular term in contemporary writing is 'working memory'. The concept of working memory greatly overlaps with that of short-term memory, but by design there is a subtle difference in how active each system is thought to be. That is, the 'working' in working memory reflects that information is not only stored, but also manipulated. By contrast, in this framework short-term memory is just a passive storage.

In an impressive effort, Engle and colleagues tested participants on 11 tasks, each employed in the literature to test either short-term or working memory. In addition, they subjected participants to tests of fluid intelligence. Structural equation modelling of results on all memory tasks suggested that a solution with two underlying variables matched the data best, although the two were highly correlated. The authors also report that while working memory correlated with fluid intelligence, short-term memory did not. (Engle, Tuholski, Laughlin, & Conway, 1999)

A similar investigation from a partly-overlapping research team employed highly similar methodology and replicated these findings (Conway, Cowan, Bunting, Therriault, & Minkoff, 2002).

It should be noted that working memory tasks in the above studies were almost exclusively ones that required participants to remember information while potentially competing information was presented. For example, Conway and colleagues employ a reading span task where participants remember words while having to read sentences out loud (including comprehension tests), an operation span task during which participants were required to remember words while doing arithmetic, and a counting span task during which participants were required to remember digits while counting aloud (Conway et al., 2002). By contrast, the short-term memory tasks required participants to remember a series of words in four variations (words drawn from an unlimited or a fixed pool of words, and with or without articulatory suppression). Clearly, the primary distinction between these two sets of tasks is how resistant participants' short-term memory should be to interference. (In fact, a good strategy in the working memory tasks is to commit information to long-term memory, a store that is

notably less resistant to interference, and recall it from there when required. Perhaps an ability to do this is what underlies fluid intelligence?)

In my view, the latent factor that Engle and Conway and their respective colleagues labelled ‘working memory’ could just as well have been labelled ‘resisting interference’ or ‘efficiently processing information’. These are undeniably important cognitive abilities, and it seems more than reasonable to assume they support fluid intelligence. However, the conclusion that working memory is a different system from short-term memory is unwarranted. A description that fits the results just as well is that of a single short-term memory store, with a mechanism that allows for transfer of information in and out of this store. When the capacity of the store and the ease of transfer are independent, this system would be equivalent to having independent working and short-term memory stores. (Note that this equivalence only refers to the observable behaviour; obviously the outlined systems themselves differ in their architecture!)

The arguments outlined above are not particularly novel, and seem implicit in contemporary research. Some authors systematically refer to both short-term and working memory, sometimes for the same tasks. In a large body of the literature, the terms seem interchangeable. In honour of this tradition of confusion, the title of this thesis refers to working memory, but the actual text will systematically refer to short-term memory. Here, short-term memory is defined as the entire system of temporary storage of information, including a limited capacity, as well as an independent information processing rate.

3.8. The architecture of visual short-term memory

Perhaps it was implicitly clear through the previous paragraphs, but it should also be made explicit: This thesis takes a systems approach to cognition. Specifically, it operates in the layer between ‘box-and-arrow’ overviews of cognitive architectures (Baddeley & Hitch, 1974), and the implementation of cognitive faculties in biologically plausible neural networks (Schneegans & Bays, 2017a).

The specific system I have in mind when writing is that of a largely signal-driven perceptual system, from which information needs to be selected, and then temporarily stored in visual short-term memory. Attention is defined as every point where information is boosted (selected) and transferred to a state of higher permanence (encoding), and thus quite a catch-all term, which is not inconsistent with the literature, and there is reason to belief the processes of selection and maintenance are in fact related. This is supported by e.g. the work of Zokaei and colleagues, which is discussed in greater detail belowmore on that later. In addition, others have written on the idea that items maintained in visual short-term can be thought of as attentionally-activated internal representations (Cowan, 2001; Kiyonaga & Egner, 2013). Although interesting, these ideas are just beyond the scope of this thesis, and will not be discussed further.

Attentional selection is discussed in earlier paragraphs, but not the nature of short-term storage. In the traditional view, short-term memory is the sustained activity of neural populations in occipito-temporal, parietal, or frontal cortex that are sensitive to the memorised item (J. D. Cohen et al., 1997; Constantinidis, Franowicz, & Goldman-Rakic, 2001; Constantinidis & Steinmetz, 1996; Courtney, Ungerleider, Keil, & Haxby,

1997; Curtis & D'Esposito, 2003; Funahashi, Bruce, & Goldman-Rakic, 1989; Petit, Courtney, Ungerleider, & Haxby, 1998).

As an addition to sustained activity, some have suggested that information can be stored in an "activity-silent" way, encoded within the connections between neurons (Lewis-Peacock, Drysdale, Oberauer, & Postle, 2012; Sprague, Ester, & Serences, 2016; Sreenivasan, Curtis, & D'Esposito, 2014; Stokes, 2015; Wolff, Ding, Myers, & Stokes, 2015; Wolff, Jochim, Akyürek, & Stokes, 2017). This theory is highly compatible with ideas on how long-term memory is stored, and more importantly an elegant idea that can be implemented in biologically plausible modelling of neural populations. The preferred methodology of most activity-silent studies is impressive. Specifically, because of a lack of sustained activity, the 'hidden' or 'latent' state in which information was encoded had to be probed in a different way. For example, Wolff and colleagues present an unrelated visual stimulus to 'ping' the visual system while it was maintaining an item in short-term memory, and employ multivariate statistical classifiers to test whether the ping induced distortions in the EEG signal that were systematically related to the item in memory (Wolff et al., 2015, 2017). Other authors have used similar techniques in fMRI (Sprague et al., 2016).

Although there is little question that these results indicate that item-related information is present in the brain signal, it was recently argued that this might not necessarily mean that information was stored in neuronal connection weights. In fact, simulations show that it might well be sustained activity that produced the multivariate decoding accuracy (Schneegans & Bays, 2017b).

Regardless of in what form information is stored in visual short-term memory, it

requires active maintenance: When individuals engage in a task that taxes their attention, the load of that task is negatively correlated with their recall performance (Zokaei, Heider, et al., 2014). On the basis of these findings, Zokaei and colleagues suggested that maintaining the binding between memorised items' features is supported by the same system that is involved in the attentional selection of perceptual information.

Another fundamental property of the architecture of visual short-term memory, is that not all items are maintained equally. In fact, there seems to be one item in a 'privileged state' (Zokaei, Manohar, Husain, & Feredoes, 2014) or in the 'focus of attention' (Souza & Oberauer, 2016). A popular method to investigate this state is a regular short-term memory task where participants are required to maintain a certain number of items in short-term memory, followed by a *retro-cue*, which tells participants which of the items is most likely to be probed. Participants then shift their internal attention to the cued item, which thus moves into the focus of attention or privileged state, without the other items disappearing from short-term memory. There are other tasks that achieve the same effect, some perhaps even more convincingly so; for an impressive systematic investigation, see (Zokaei, Ning, Manohar, Feredoes, & Husain, 2014).

Perhaps the most convincing demonstration of the existence of a privileged state in short-term memory comes from Zokaei and colleagues. They presented participants with random dot kinematograms that were defined by two features: their colour and their direction of movement. While participants maintained these items in short-term memory, one of the items was retro-cued to bring it into the focus of attention. After the

retro-cue, transcranial magnetic stimulation (TMS) was applied to the vertex (sham stimulation), or to area V5 (also know as MT), a middle temporal visual area that is involved in the processing of visual motion. When sham TMS was applied, the retro-cued item was recalled with a higher precision that the non-cued items, which were also remembered but at a lower precision. Remarkably, when TMS was applied to V5, the retro-cued item seemed to have been forgotten, but the non-cued items were still remembered as precise as in the control condition. In sum, it seems as though the 'privileged state' in short-term memory is a re-activation of primary sensory cortical areas, and this is where item's representations are most vulnerable to being changed. (Zokaei, Manohar, et al., 2014)

3.9. From attention to short-term memory

Now that the preceding paragraphs have introduced attention and short-term memory, and even some of their important overlaps and similarities, the following paragraphs will introduce the first two empirical chapters in this thesis. These are concerned with how attentionally selected information is encoded into visual short-term memory, with a particular focus on the encoding mechanism.

According to the feature-integration theory of visual attention, items that are conjunctions of features need to be attended serially so their features can be bound veridically (Treisman, 1982; Treisman & Gelade, 1980). It has since been extended to argue that this is a necessary pre-requisite for items to be encoded into visual short-term memory (Wheeler & Treisman, 2002).

The popularity of this idea can hardly be overstated. In the field of visual search, there exist a multitude of models that are all the same in their core ideas: Visual search is characterised by a massively parallel and unconscious sweep across the visual field that picks up on simple features such as colours that stand out, which is followed by a serial deployment of attention on each individual item in the visual field (Cave & Wolfe, 1990; Duncan, 1984; Hoffman, 1979; Wolfe, 1994). This line of thinking can explain pop-out effects in visual search, which occur when one item is different from all other items (for example a red square among green circles), and response time is independent from the number of items in the search array. In addition, it explains why response time in visual search for complex items is predicted by the number of items in a search array: Those items are apparently too complex to be differentiated in the initial parallel stage, and require being picked up by the attentional spotlight that is deployed serially.

There exist notable exceptions to the idea that binding features into a single item requires focussed attention. For example, Houck and Hoffman demonstrated that the McCollough effect is equally present in items that are and that are not in the locus of spatial attention (Houck & Hoffman, 1986). The McCollough effect is an after-effect of post-retinal origin: during the adaptation phase, participants are presented with alternating horizontal and vertical gratings of different colours, and the after-effect is that non-coloured alternating gratings of the same orientation appear in the complementary colour to the adapted (McCollough, 1965). Houck and Hoffman argued that this effect is dependent on bound features within each item, and thus that its occurrence in a pre-attentive item is contrary to what feature integration theory would predict.

Some have partly accounted for these exceptions by suggesting that features in the first stage are “bundled” rather than bound (Wolfe & Cave, 1999), although it is unclear what the difference between bundling and binding is. Others have fully embraced the notion of the parallel encoding of items into short-term memory, and made it a central element of their model. In his ‘Theory of Visual Attention’ (TVA), Bundesen (1990) provides a mathematical framework for the attentional selection of information, and its consecutive encoding into short-term memory. (Apparently encoding is so integral to the theory of attention that it isn’t even mentioned in the title.)

In the TVA framework, each item in the visual scene is assigned an attentional weight. This weight is the sum across all the item’s features of the product of the strength of the sensory evidence that this item belongs to a specific category (e.g. red) and the pertinence value of that category. Pertinence can be set centrally to prioritise certain features over others, for example to attend to certain locations, or to attend specifically to red items. (Note that this allows for feature-based attention!) These attentional weights are then used to compute the hazard rate of each item, which is essentially the product of the strength of sensory evidence for an item belonging to a particular category (TVA is essentially a theory of pigeon-holing), the bias for that particular category (based on goals, task sets, priming, etc.), and the relative attentional weight. (Bundesen, 1990)

Because TVA concerns a rate function per item, it describes a race between each of the items from the first moments of visual perception until they are allocated a slot in short-term memory. Thus one of TVA’s central assumptions is that short-term memory consists of quantised slots, but it could be redefined to fit with a resource model,

without damaging any of its core predictions. From the sum of all rate functions, one can compute what the total processing capacity is (Bundesen, 1990). This determines how many items can be encoded per unit of time.

TVA can be fitted to data from whole report tasks, in which participants are presented with a number of letters, all of which they are required to remember and reproduce. The crucial manipulation is the exposure duration of stimuli before they are masked. Bundesen showed that with increased exposure duration a higher number of items are correctly remembered, until an asymptote is reached. This process is best fitted with the above described model with three constraints: There is a minimal exposure duration before which no information is encoded, the encoding process is rate limited (a finite number for the processing capacity), and the short-term memory store is capacity limited (a finite number of slots) (Bundesen, 1990). Examples for these values given by Bundesen's analysis are 25 ms for the minimally effective exposure duration, an encoding rate of about 17 items per second, and a short-term memory capacity of just under 3 elements.

More recently, Vogel and colleagues presented a number of experiments. The first employed a change detection task in which 1, 2, 3, or 4 items were presented with various encoding durations. The items were coloured squares (visible for 100 ms), which were followed by a grey screen (117 – 584 ms), and then followed by masks that were squares with four coloured tiles (visible for 200 ms). The probe array was identical to the stimulus array (half of all trials), or had one item with a changed colour (half of all trials). Participants' responses on whether a change had occurred were at ceiling for all encoding durations in the 1 item condition, but were affected by encoding duration

for all other array sizes. In addition, array size was negatively correlated with performance. In sum, performance was worse for larger arrays, and performance rose as a function of exposure duration. (Vogel, Woodman, & Luck, 2006)

In their second experiment of the same study, Vogel and colleagues used a single memory array size of 4 items, and shorter encoding durations: 100 ms of stimulus exposure, followed by 17, 33, 50, 67, 83, 167, or 217 ms of grey screen, and then followed by the same mask as for the first experiment. Vogel and colleagues computed the number of stimuli in memory at each exposure duration via Cowan's equation K = correct rejections * (hit rate + array size – 1) (Cowan, 2001). This value increased with encoding duration, at a rate of about 1 stimulus per 50 ms, and with an asymptote at 2.5 items that was reached just before 200 ms. (Vogel et al., 2006)

This work confirms Bundesen's assertion that short-term memory encoding is both rate and capacity limited, even if the values don't necessarily overlap. This could be due to the difference in stimulus material, i.e. letters in Bundesen's data, and coloured squares in Vogel and colleagues'.

Taking a step forward, Bays and colleagues introduced slightly more complex items, as well as responses on a continuous scale. In their experiment, they presented participants with four items: bars that were characterised by their orientation (180 degree range) and colour (one out of six options). The bars were presented briefly, at exposure durations of 25-1000 milliseconds, after which they were masked by a very large array of masks than spanned across nearly the entire screen. The mask disappeared after 100 milliseconds, leaving a grey screen for 1000 milliseconds, which was then replaced by the probe. The probe was a central bar in one of the presented items'

colours, and participants had to adjust the bar to the orientation of the probed item. (Bays, Gorgoraptis, Wee, Marshall, & Husain, 2011)

Importantly, Bays and colleagues used stimuli that were characterised by three features: location, orientation, and colour; and they probed by colour when asking for an item's orientation, without giving participants information on items' original location. This required participants to have bound each item's colour to its orientation. It is important to realise this, as feature-integration theory would predict that participants would only be able to accurately recall an item after attending to it to bind its features together.

Bays and colleagues computed the standard deviation of the error distribution at each exposure duration. The lower this standard deviation is, the tighter its grouping around 0 error, and thus the more accurate its recall. Precision was quantified as the reciprocal of each error distribution's standard deviation. This precision value increased with exposure duration, thus showing a positive encoding rate; and it reached an asymptote after about 120 milliseconds (regardless of set size!), thus showing a short-term memory capacity limit. In line with the work described in previous paragraphs (Bundesen, 1990; Vogel et al., 2006), Bays and colleagues show that the memory array size affects both the rate of encoding and the asymptote per stimulus, thereby providing more evidence that short-term memory encoding is both rate and capacity limited. (Bays et al., 2011)

Although there is consensus on short-term memory encoding being a rate-limited process, and that short-term memory has a limited capacity, the actual encoding mechanism remains unclear. Although violations of feature-integration theory have been

reported (Houck & Hoffman, 1986), and the results described in the previous paragraph (Bays et al., 2011) would not have been possible under strict feature-integration theory, it is still the de-facto theory of how items are encoded.

At the heart of feature-integration theory lies its insistence on the requirement of serial processing of items. An illustration of how popular this idea still is, can be found in recent theoretical work on a computational account of short-term memory, which assumes serial encoding (Manohar et al., 2017). This is not without grounds, as recent empirical work has been framed in support of a serial encoding mechanism: The aforementioned study by Vogel and colleagues (Vogel et al., 2006), and more recent work by Liu and Becker that is described in the next paragraph. (As an aside, it is interesting to note how a model proposed by the founders of the resource model is best supported by studies reported by slot model founders and supporters.)

In a recent study, Liu and Becker (2013) tested 12 undergraduate students, using a task in which participants were presented with two gradients of a particular orientation. These could be presented simultaneously for 150 milliseconds, followed by a mask that stayed on for 200 milliseconds, a blank screen for 500 milliseconds, and then followed by a probe that asked participants to report the orientation of one of the stimuli. In another condition, stimuli were presented sequentially, with one stimulus appearing for 150 milliseconds, then a mask appearing for 200 milliseconds, which was followed by a grey screen for 500 milliseconds, after which the second stimulus appeared for 150 milliseconds, again followed by a mask for 200 milliseconds, and a grey screen for 500 milliseconds; the probe followed in the same fashion, asking participants about the orientation of one of the presented gratings. Participants

responded on a continuous scale, and thus the resulting error distributions could be fit to a weighted sum of a Von Mises distribution with a standard deviation and a uniform distribution (Zhang & Luck, 2008). Potential swap errors were not accounted for. Crucially, Liu and Becker showed that the recall precision was the same for the simultaneous and serial presentations, but the guessing rate was higher in the simultaneous condition than in the serial condition. Liu and Becker argue that if items were consolidated in short-term memory in parallel, there should not be a difference between the conditions, and thus argue that their results indicate a serial bottleneck. (Liu & Becker, 2013)

Unfortunately, Liu and Becker did not account for the rate limit in encoding: Their sequential position allows 150 milliseconds of exposure duration for each stimulus, and their simultaneous condition allows 150 milliseconds for both stimuli (Liu & Becker, 2013). This means that the effective encoding duration was 300 milliseconds in the sequential condition, and only 150 milliseconds in the simultaneous condition, rendering any difference between the two likely due to the difference in encoding duration. It should also be noted that aforementioned theoretical (Bundesen, 1990) and empirical (Bays et al., 2011) work has demonstrated that the encoding rate per stimulus is directly related to the amount of stimuli that are processed: Memorising more items allows for fewer encoding resources per item. Hence, the difference between the sequential and simultaneous condition could as easily have occurred if encoding happened in parallel.

It would be unfair to dismiss Liu and Becker without also mentioning their additional analysis, which was a comparison between two models. The first model was

labelled a ‘serial model’, and employed one free parameter for the standard deviation of the two Von Mises distributions fitted to the error distributions from the sequential and parallel conditions, but two free guessing parameters (i.e. the proportion of trials in the uniform distribution that complements the Von Mises distribution), one for each condition. In what they label the ‘parallel model’, there are free parameters for the standard deviations of the Von Mises curves fitted to error distributions in both, and there is only a single free guessing parameter for both conditions. The ‘serial’ model fitted the data better, and thus Liu and Beck argued that encoding must occur in series rather than in parallel, with one item being committed to short-term memory at a time. (Liu & Becker, 2013)

Interestingly, it is entirely unclear why these models were chosen, and how they are constructed. The supposed serial model allows for a single encoding rate (quantified by the fitted standard deviation of the error distribution), which is in line with a serial model, as the items would be processed one-by-one by the same channel with a limited encoding resource. The supposed parallel model allows for two encoding rates, which is in line with a parallel model, as the two items would be encoded by two separate channels, each with half the limited amount of encoding resources. However, the supposed parallel model keeps the guessing rate fixed between the two unmatched sequential and simultaneous conditions, which is definitely not what a parallel model would predict, as the guessing rate is not independent from the encoding rate. Hence, a truly parallel model would have also allowed for two guessing rates.

One potential explanation for Liu and Becker’s invalid choice of models is that they intended to keep the amount of parameters equal to facilitate an easier model

comparison. Another is that they considered their choice to be valid, because they did not see the flaw in their unmatched design. Regardless, the only valid conclusion one can draw from Liu and Becker's data is that longer exposure durations allow for better recall performance, which is in line with previous research (Bays et al., 2011; Bundesen, 1990; Vogel et al., 2006).

In sum, there is currently little (valid) empirical evidence in favour of serial encoding of items into visual short-term memory, but there is also no evidence for parallel encoding of items into short-term memory. *Chapter 4: Encoding of information into visual short-term memory is a parallel process* presents the first evidence in favour of parallel encoding. In it, I present the results of four experiments with whole-report tasks and responses on a continuous scale, modelled using a computational framework that can distinguish between serial and parallel encoding of items.

In *Chapter 5: Serialisation of visual short-term memory encoding in the presence of reward*, I explore whether humans can adjust their encoding process to work serially instead by rewarding their recall of only one item. Although the sample is small (15 participants), computational modelling demonstrates that encoding did indeed occur in a serial fashion, seemingly with one item being committed to short-term memory within 40 milliseconds, a value that is very close to earlier estimates of item-based serial encoding rates (Vogel et al., 2006).

3.10. Re-allocation of short-term memory resources

Earlier in this chapter, a study was cited to illustrate the idea that the maintenance of information in short-term memory requires attentional efforts (Zokaei, Heider, et al., 2014). It was also illustrated that the same attentional process is involved in encoding information into visual short-term memory in the first place (Bays et al., 2011; Bundesen, 1990; Vogel et al., 2006). What has not been covered, is whether the limited amount of short-term memory resources can be redistributed between items in short-term memory. Re-allocation of resources in this case does not refer to bringing an item into the focus of attention (Souza & Oberauer, 2016; Zokaei, Ning, et al., 2014), but rather to 'free up' some short-term memory resources for a new item, by taking them away from other items.

In one of the experiments presented by Bays and Husain, participants were required to remember 5 items in short-term memory. However, they were required to fixate every item in the memory array. Crucially, during the saccade to the last item, the screen was blanked, which meant participants never actually fixated the final item. Despite this, Bays and Husain report that recall precision for the last (and unfixated) item was far better than for the other items. Thus, they argue, short-term memory resources must have been re-allocated to final item – the target of the last saccade – even before the saccade was made. In their view, saccade targets must be represented at a greater resolution in order for the eye to land sufficiently accurately. (Bays & Husain, 2008)

It is unclear whether this effect is truly a re-allocation of short-term memory

resources, or rather caused by a pre-saccadic shift in attention (as presented by e.g. (Rolfs, Jonikaitis, Deubel, & Cavanagh, 2011)). To potentially answer this question, the aforementioned experiment from Bays and Husain will be replicated and expanded upon in Chapter 6: *Dynamic re-allocation of visual short-term memory resources to saccade targets?* Unfortunately, the question will remain largely unanswered, as I could not replicate the effect reported by Bays and Husain, albeit with a different design. Although this suggests the effect might not be quite as general or robust as initially thought, the exact nature of the effect if it does occur remains elusive.

3.11. Stroke-induced deficits of attention and short-term memory

Moving beyond the fundamental work presented in the first three empirical chapters of this thesis, the other two empirical chapters will focus on clinical topics concerning deficits in attention and short-term memory that can be caused by stroke. This is not a trivial matter: Almost half of all stroke patients initially suffer from impaired attention (Lesniak, Bak, Czepiel, Seniow, & Czlonkowska, 2008). In 25-50% of stroke victims this is due to hemispatial neglect syndrome (Appelros, Karlsson, Seiger, & Nydevik, 2002; Buxbaum et al., 2004; Nijboer, Kollen, & Kwakkel, 2013), which is characterised by an attentional bias away from the contra-lesional hemifield (Bays, Singh-Curry, Gorgoraptis, Driver, & Husain, 2010). This syndrome is caused primarily due to lesions in the right hemisphere (Ringman, Saver, Woolson, Clarke, & Adams, 2004), but can also occur after lesions in the left hemisphere where it is more

rare but can also be more severe (Ten Brink, Verwer, Biesbroek, Visser-Meily, & Nijboer, 2017).

In addition to the characteristic bias in spatial attention, some patients with neglect syndrome may also display deficits in additional domains. One of these deficits is impaired suistained attention (Robertson, Manly, Beschin, et al., 1997; Robertson, Manly, Andrade, Baddeley, & Yiend, 1997; Robertson, Mattingley, Rorden, & Driver, 1998), and another is impaired short-term memory (Husain & Rorden, 2003; Malhotra et al., 2005; Parton et al., 2006).

3.12. Patient-friendly testing of cognitive functions

One popular task to measure spatial attention in stroke patients is the cancellation task – a multi-target visual search paradigm. In this task, participants are required to find targets among distractors. When they find a target, they cross it out ("cancel" it). Cancellation tasks have long been done with pen and paper, but have made an appearance as a computerised test in the late 1990s (Donnelly et al., 1999).

Pen-and-paper cancellation tests offer very limited information, as they can only tell a clinician which items were cancelled and which were not. This is not for a lack of trying to collect additional information: A variety of methods has been employed to circumvent the issue, including frame-by-frame video analysis, instructing participants to verbally cancel targets, an experimenter monitoring search behaviour, and regularly swapping pencil colours (Mark, Woods, Ball, Roth, & Mennenmeier, 2004; Samuelsson, Hjelmquist, Jensen, & Blomstrand, 2002; Warren, Moore, & Vogtle, 2008; Weintraub &

Mesulam, 1988).

Despite some of these creative solutions, computerised cancellation tests are much more informative, because they register what target was clicked when. This opens up a whole host of measures, including the possibility to record perseverations (consecutive cancellations of the same item), the speed of cancellations, and most importantly the cancellation path. From this path, one can derive the amount of revisited items, but also the structure of the participants search behaviour.

An additional option that computerised tests offer is to not visibly mark cancelled items. In this 'invisible' version, participants are required to remember which items they have cancelled, and thus the test is sensitive to deficits in short-term memory (Parton et al., 2006).

What measures can be derived from computerised cancellation tests, and how to compute them, is described in much greater detail in *Chapter 7: Using cancellation tasks to assess spatial attention, short-term memory, and executive functioning*. In this chapter, I not only provide a review of all measures derived from cancellation tasks in the current literature, but also introduce additional metrics, and illustrate their typical ranges in a small sample of control participants and neglect patients.

In addition, I present a study of several hundred healthy controls who participated in a brief (under 2 minutes) internet-based cancellation test. By applying machine-learning algorithms on the cancellation measures computed from their behaviour, I illustrate the existence of several cognitive profiles within the healthy cancellation-test-performing population.

3.13. Rehabilitation of inattention after stroke

Spontaneous recovery from neglect occurs in 30-40% of affected patients (Nijboer, Kollen, et al., 2013). Unfortunately, those who do not recover from neglect are hospitalised longer and face more problems in daily life than stroke patients without neglect (Nijboer, Van de Port, Schepers, Post, & Visser-Meily, 2013). Because of the severe and chronic nature of neglect syndrome, it is important that effective treatments are sought.

Many treatments have been tested in the past. One promising and non-invasive intervention is prism adaptation. In this treatment, neglect patients were made to wear prism glasses that shifted their vision in an ipsi-lateral direction; so that when they then remove the glasses after being adapted, attention shifts towards the contra-lesional direction (Rossetti et al., 1998). Some reports concluded that on the whole, prism adaptation modulates performance on pen-and-paper tasks in some reported studies (De Wit, Ten Brink, Visser-Meily, & Nijboer, 2016). However, the effects are extremely small and inconsistent, and not apparent on tests of neglect's effects on daily living. This is powerfully illustrated in a recent randomised, sham-controlled trial with a sufficiently large sample (70 patients), intensive treatment (2 weeks), and no less than seven measurements (baseline, 1, 2, 3, 4, 6, and 14 weeks post-treatment) showed no difference between patients who underwent prism adaptation and those who received sham adaptation (Ten Brink, Visser-Meily, et al., 2017).

An additional avenue of exploration is pharmacological treatment. A variety of treatments has been more-or-less successfully implemented, including cholinergic

(generally positive), dopaminergic (mixed results), and noradrenergic interventions (only improved sustained attention) (van der Kemp, Dorresteijn, Ten Brink, Nijboer, & Visser-Meily, 2017). The primary focus in the literature has been on dopaminergic interventions (Barrett, Crucian, Schwartz, & Heilman, 1999; Fleet, Valenstein, Watson, & Heilman, 1987; Geminiani, Bottini, & Sterzi, 1998; Grujic et al., 1998; Hurford, Stringer, & Jann, 1998; Mukand et al., 2001; Pierce & Buxbaum, 2002), which have unfortunately often lacked double-blinding and placebo controls.

Perhaps the most promising dopaminergic treatment results have been presented by Gorgoraptis and colleagues, who used rotigotine on 16 patients with left-sided neglect (right-hemispheric damage) in a double-blind and placebo-controlled design wherein patients were tested before, during, and after treatment. Their results demonstrate a drug-induced improvement of pen-and-paper cancellation performance (number of targets found and spatial bias), but no drug effects on independent tests of sustained attention and short-term memory. (Gorgoraptis et al., 2012)

Targetting a different system, Malhotra and colleagues used the noradrenergic agonist guanfacine in a double-blind and placebo-controlled trial of three patients. Two patients improved on computerised invisible cancellation tests, specifically in the number of targets that they found, and in the time they spend on the task. (Malhotra, Parton, Greenwood, & Husain, 2006)

These results are promising, but the study was only a pilot with a mere sample size of three. In addition, the improvements can be attributed to improved spatial attention (more exploration and thus more targets found), improved sustained attention (in line with extended time-on-task, allowing for more time to find more targets),

improved short-term memory (reducing the likelihood of previously cancelled targets to be cancelled), or a combination of the three. To answer this question, a randomised, placebo-controlled, cross-over study with 13 neglect patients was conducted, in which a time-limited invisible cancellation test, a sustained attention test, and a short-term memory test were employed. This study is discussed in chapter 8: *The effects of guanfacine on visual search performance, sustained attention, and short-term memory after stroke*.

In short, an increase in the number of targets was found on guanfacine compared to placebo, but sustained attention, short-term memory, and search organisation were unaffected. In fact, the latter were statistically reliable null-effects, as demonstrated using Bayesian statistics.

3.14. Summary

In sum, attention and short-term memory can be viewed as a continuum among which information from the outside world travels, is filtered to what is currently important for the individual, and temporarily retained. However, the relationship between attention and short-term memory is more complex. Attention does not only act as a gatekeeper to short-term memory, but can also act on internally stored information. Attention can bring information into focus within short-term memory, as well as suppress and remove information from short-term memory. Both modalities can be independently affected in stroke, as measured by simple tools such as the cancellation task. Fortunately, recent studies have provided a promising avenue for pharmacological

intervention to ameliorate inattention in neglect syndrome.

4. Encoding of information into visual short-term memory is a parallel process

Parts of this chapter are in preparation as:

Dalmaijer, E.S., Manohar, S.H., Lowcock, G., Poullias, C., Stokes, M.G., & Husain, M. (in prep.). Parallel encoding of information into visual short-term memory.

4.1. Abstract

Visual information is abundantly present in the outside world. Selecting and retaining the elements that are relevant to behaviour is crucial to survival. It is generally accepted that individuals can only retain a limited amount of information, as visual short-term memory has a highly limited capacity. However, it remains hotly debated whether this information is encoded into visual short-term memory one element at a time, or whether encoding of several elements can occur in parallel. In order to answer this question, I employed a novel adaptation of the classic whole report paradigm. In four experiments, participants were briefly (0-580 ms) presented with one or two stimuli that were immediately followed by a mask, and asked to remember its/their colour (experiment 1 and 3) or orientation (experiment 2). Shortly after mask onset (500 ms in experiments 1-3, 1100 ms in experiment 4), participants had to respond to both stimuli using a continuous report scale, but were free to choose which to respond to first. Their performance improved as a function of stimulus exposure duration, and was better for the first response than for the second in the two-item condition. When fitted to a mixture model of parallel encoding, the results show close adherence, with no indication of serial processing. This is true for colour (experiment 1), replicated with orientation (experiment 2), and even when stimuli are presented within the same hemifield so that they are first presented to one cerebral hemisphere (experiment 3). In addition, the electroencephalographic (ERP and oscillatory power over time) response is identical for both stimuli (experiment 4). These results suggest that information from different stimuli can be encoded into visual working memory at the same time, although not necessarily at the same rate.

4.2. Introduction

When interacting with the external world, individuals selectively attend to information that is relevant to their goals and behaviour. Because not all of this information is continuously available, it is important that some of it can be temporarily stored in short-term memory. There is a general consensus in the literature that short-term memory is of a highly limited capacity, although the precise architecture is still a topic of debate with some arguing short-term memory is organised in discrete 'slots' for single items (Luck & Vogel, 1997, 2013b; Zhang & Luck, 2008), and others proposing that it is in fact a fluid resource that can be flexibly allocated across items (Bays et al., 2009; Bays & Husain, 2008; Ma et al., 2014; Wilken & Ma, 2004).

A fundamental question remains how information is encoded into short-term memory.

4.2.1. The short-term memory encoding mechanism

Several lines of evidence have revealed that encoding is not an instantaneous process. When stimuli are presented simultaneously for a brief amount of time before being masked, the amount of items remembered by participants increases with stimulus exposure duration (Bundesen, 1990; Vogel et al., 2006). More recent work suggests that with increased encoding time, the precision of recall for items increases (Bays et al., 2011), suggesting that during encoding the representation of an item in short-term memory becomes progressively less coarse.

A more contested topic is *how* items are encoded, with roughly three lines of contemporary thinking. The first proposes that information is encoded into short-term memory in a serial fashion: One object is encoded first, and the next object is encoded only after the first has finished (Wheeler & Treisman, 2002). This is an influential idea that has recent support in the literature (Liu & Becker, 2013), and is a fundamental assumption for computational models of short-term memory that employ feature-based codes (for example (Manohar et al., 2017)).

In these models, units that code for the single features of one item (e.g. colour, orientation, and location) are simultaneously activated, thereby providing a representation of one item. When two (or more) items would activate this system at once, there would be a fundamental confusion between the unique features of the item. For example, when a green and a red bar would be processed at the same time, the colour units that are sensitive to red and green would be active at the same time, as would the populations that are sensitive to the locations of both items. Such a situation would mean that it is unclear which colour occupies which location. Thus, feature-based models of short-term memory require sequential processing (and binding) of item features.

A slightly different proposal is for a two-stage model. In the first stage a coarse representation is obtained from items in view (parallel), while in the second stage these items are individually processed in a serial manner (Cave & Wolfe, 1990; Duncan, 1984; Hoffman, 1979; Wolfe, 1994). It is important to note that the first stage is considered to be a pre-attentive and automatic process, and that the encoding of information into short-term memory is part of the second stage, because it is assumed

that encoding requires attention to a single object. In this sense, two-stage models actually support the serial encoding hypothesis.

Contrary to the serial hypothesis, is the proposal that information can be encoded into short-term memory in parallel. In one computational instance of this view, all (attended) items are considered to 'race' against each other towards short-term memory, where the first few will fill up the (limited) available slots (Bundesen, 1990). The psychophysical evidence for parallel encoding into short-term memory is scare. As discussed in Chapter 3: *Introduction*, there is no direct empirical evidence for it, although there are studies that demonstrate that attention can be allocated to more than one non-contiguous location at the same time (Hahn & Kramer, 1998; Kramer & Hahn, 1995; McMains & Somers, 2004; Müller et al., 2003), which is a key prerequisite for parallel encoding. In addition, some studies have interpreted their results in line with parallel encoding, although they were not direct evidence of it (Wilken & Ma, 2004).

In terms of biology, one could hypothesise that information is encoded and maintained in different populations of parietal and prefrontal neurons, through sustained activity (Constantinidis et al., 2001; Constantinidis & Steinmetz, 1996; Fuster & Alexander, 1971) or through updating of synaptic weights between them ((Stokes, 2015) but see (Schneegans & Bays, 2017b)). This view would be consistent with the resource model of short-term memory (Bays et al., 2009; Bays & Husain, 2008; Ma et al., 2014; Wilken & Ma, 2004), in the sense that one population of neurons can be flexibly subdivided and allocated to the retention of specific items.

Although the direct translation of a theoretical resource to a neuronal population seems simplistic, it is supported by emerging evidence: In simulations of bump-attractor

networks (often used to model short-term memory) an increase in population size is associated with an increased peak amplitude of the representation and a decrease of its variance (Curtis, Mackey, Ding, Wang, & Winawer, 2016). Moreover, analysis of population receptive fields (pRF; see (Wandell, Dumoulin, & Brewer, 2007)) computed from human functional magnetic resonance imaging (fMRI) studies in which participants perform a short-term memory task, demonstrates that the size of multiple visual field maps in frontoparietal cortex correlates with recall precision (Curtis et al., 2016).

In addition, recent computational work demonstrates that a single population of units can in fact act as a temporary information storage with characteristics that are remarkably similar to human short-term memory (Matthey, Bays, & Dayan, 2015). In this new model, conjunctive units that code for the combination of several object features were used in addition to the aforementioned feature-based units that code for single features. Thus items essentially leave their individual imprint in the neuronal population, and can do so at the same time.

In sum, the parallel hypothesis is biologically plausible, compatible with the major theories of short-term memory architecture, and is supported by recent computational accounts of short-term memory. However, as discussed in Chapter 3: *Introduction*, empirical work has mostly aligned with the serial hypothesis, or at least has been interpreted as such. The parallel encoding hypothesis, while supported by studies that demonstrate that attention can operate on multiple locations in parallel (Hahn & Kramer, 1998; Kramer & Hahn, 1995; McMains & Somers, 2004; Müller et al., 2003), has (to my knowledge) not been directly tested.

4.2.2. Modelling behaviour to investigate encoding

In this chapter, I directly contrast the serial and parallel encoding hypotheses by using a novel adaptation of the classic whole report task (Bundesen, 1990). In this type of experiment, participants are presented with one or more stimuli, which are masked after a variable time. After a fixed maintenance period, participants are asked to recall all the stimuli they remember. In my implementation, participants are presented with colours or orientations, and respond to all presented stimuli by using a continuous report scale. The resulting error distributions can be fitted to computational models of short-term memory functioning (outlined below in section 4.3.8: *Mixture model of short-term memory encoding*), for which the parallel and serial encoding hypotheses predict different outcomes.

More specifically, when two briefly presented stimuli have to be stored in short-term memory, the serial encoding hypothesis predicts that the probability of one item being encoded should be dramatically higher than the probability of two stimuli being encoded, when exposure durations are very short. In contrast, the parallel encoding hypothesis predicts that both stimuli should be encoded at the same time, and therefore result in a relatively high probability of both stimuli being encoded.

It has been shown that in monkeys that perform dual tasks, the frontal lobes essentially divide the goals at hand between the hemispheres (Charron & Koechlin, 2010). In addition, it has been demonstrated that pRFs are biased towards contra-lateral space even when participants were attending centrally, suggesting that there is a degree of hemispheric specialisation with regards to processing stimuli in the contra-lateral

hemifield (Sheremata & Silver, 2015). However, it should be noted that "parallel encoding" here does not refer to the processing of one stimulus per hemisphere. Instead, it is proposed that the parallel encoding of multiple items into short-term memory is a universal phenomenon even within a single hemisphere.

In experiments 1 – 3, I employ a variable stimulus exposure duration to probe behaviour at several points during the encoding process. The resulting error distributions can be fitted to a novel model that describes the short-term memory encoding process, and yields estimates of how many items are encoded into short-term memory at a given time. Under the serial encoding hypothesis, one item should be encoded first, followed by the second item, and thus a sharp peak is expected for the probability of only one item being encoded at short exposure durations. By contrast, under the parallel hypothesis, all items should be encoded at the same time, thus resulting in a relatively higher chance of all items being encoded at lower exposure durations.

4.2.3. Using EEG to investigate encoding

Attentional selection and the following encoding of selected information into short-term memory can be probed with electroencephalography (EEG). For example, an event-related potential that has been associated with attentional selection, is the P1 or P100 component. This is a positive peak in the EEG over parietal cortex, which is thought to reflect the processing of a visual stimulus, and it is attenuated when said stimulus was unattended (Luck et al., 1994; Mangun & Hillyard, 1991), or even

enhanced when said stimulus was attended to (Sauseng et al., 2005), in particular when attention is transient rather than sustained (Eimer & Forster, 2003). Interestingly, the same P100 enhancement occurs when an unrelated visual stimulus appears on the location of a memorised item (Awh, Anllo-Vento, & Hillyard, 2000). This could suggest the attentional rehearsal of information that is stored in short-term memory. In other words: It could be that individuals serially deploy internal attention to each of the items in short-term memory.

An additional hallmark of spatial attention and short-term memory, is oscillatory power in the 8-12 Hz range (from here on referred to as *alpha*, as is convention in the literature). Specifically, one can observe a decrease in alpha power contra-lateral to an attended location (Robert M. Mok, Myers, Wallis, & Nobre, 2016; Sauseng et al., 2005; van Ede, Szebényi, & Maris, 2014), and an increase in alpha that scales with memory load during retention of contra-laterally presented information in short-term memory (Jensen, Gelfand, Kounios, & Lisman, 2002). Either way, an increased task-related lateralisation of alpha power is associated with shifts of attention towards locations that are expected to contain (pre-cued) or have contained (retro-cued) memorised information (Myers, Walther, Wallis, Stokes, & Nobre, 2015). In fact, task-unrelated alpha lateralisation correlates with the precision of recall in a short-term memory task (Myers, Stokes, Walther, & Nobre, 2014).

The increase in alpha during maintenance has been described in modelling work, which proposes that serially and internally attending to items in short-term memory is reflected in oscillations in the 45-60 Hz range (referred to as *gamma*) that are superimposed on alpha oscillations (Jensen & Lisman, 1998; Lisman & Idiart, 1995). It

should be noted that although this work makes a strong biological neural argument, its implementation is on the level of response-times, and thus several steps removed from actual data.

The aforementioned findings appear to conflict, as they associate both an attenuation and an increase of alpha power. To ameliorate this to some extend, it's important to address what causes reductions in alpha. In recent work, fluctuations in alpha were associated with internal attention to one out of two items stored in short-term memory. Specifically, a decrease of alpha power was observed contra-lateral to the stimulus that was most likely to be probed at a given time (van Ede, Niklaus, & Nobre, 2017). This is in line with the consensus on alpha attenuation contra-lateral to attended locations in the outside world. In sum, decreases in alpha power over parietal cortex seem to indicate shifts of attention towards items in the environment or represented in short-term memory.

Another method of assessing information in short-term memory is to apply a multivariate classification algorithm to quantify how well EEG signals predict the contents of short-term memory. This is analogous to short-term memory decoding methods in functional magnetic resonance imaging (Harrison & Tong, 2009). In general, this approach requires raw EEG data that was collected while participants were viewing or keeping in mind some information (such as a grating with a particular orientation), which is used to train a multivariate model or classification algorithm. The fitted parameters for said model or classification algorithm are then used on different set of raw EEG data, where its classification accuracy is determined by comparing the model's prediction with the actual stimulus. In effect, this approach offers a quantification of

how much stimulus information is present in the EEG signal. For an overview of the method and its applications, please refer to (Stokes, Wolff, & Spaak, 2015).

Experiment 4 will be the same as experiments 1 and 2 in that one or two oriented gratings are visible to the left and/or right of fixation. The staggered exposure durations (0 – 580 ms) that are intended to probe recollection after various encoding times can be reduced to a single exposure duration (200 ms), because EEG reflects the continuous process of encoding during the stimulus exposure. In addition, the mask duration (500 ms during which the items are maintained in short-term memory) will be extended to allow for the tracking of short-term memory information in the EEG signal.

In the single-stimulus condition, one grating will be shown either left or right of fixation. This should result in a contra-lateral P1 in the ERP from parietal electrodes during encoding, as visual attention will be guided towards the stimulus. In addition, it should produce an increased lateralisation of alpha power during the stimulus and mask presentation, reflecting the encoding and maintenance of the item in short-term memory. This should be driven by an attenuation of alpha power contra-lateral to the stimulus.

More important are the predictions for the double-stimulus condition. Here, two stimuli are presented simultaneously: One left and one right of fixation. Under the serial encoding hypothesis, encoding a stimulus requires attending to it (Cave & Wolfe, 1990; Duncan, 1984; Wheeler & Treisman, 2002; Wolfe, 1994). In line with this hypothesis, one would expect to see a difference in P1 amplitude (or less likely: delay between the onsets of the P1) contra-lateral to each stimulus, as well as in the onsets of the alpha attenuations, as each follows an attentional shift.

Conversely, under the parallel encoding hypothesis, one would expect the

encoding of both items to occur at the same time. Thus, the onsets and amplitudes of the P1 in response to both stimuli, and the onsets alpha attenuation should occur at the same time.

4.3. Methods

4.3.1. Participants

Participants were recruited via the University of Oxford's Experimental Psychology recruiting website for healthy volunteers, with permission from the local ethics committee, and in accordance with the declaration of Helsinki. In total, 89 people took part in the experiments reported here (11 in experiment 1, 16 in experiment 2, 22 in experiment 3, and 40 in experiment 4).

4.3.2. Procedure

Participants came into the Department of Experimental Psychology for a single session. They were briefed, and provided written informed consent before the experiment was started.

In experiments 1 – 3, before starting the task an EyeLink 1000 was calibrated to each individual, which allowed its use to measure eye movements throughout the task. In experiment 4, before the participant could start the task, electrodes were placed on

the face for electrooculography (EOG), on the mastoids to act as references, and in a cap on the head participant's head for EEG.

Participants then performed the task, which consisted of 900 trials. The eye-tracking drift was tested every 20 trials, and the tracker was recalibrated when the drift exceeded 2 degrees of visual angle. Participants were allowed a break every 100 trials.

In experiments 1 – 3, the session lasted for no longer than two hours, and participants were compensated for their time with 16 pounds. In experiment 4, the session lasted for no longer than three hours, and participants were compensated for their time at a rate of 10 pounds per hour.

4.3.3. Setup

Stimuli were presented on a 21 inch (53.3 cm) ViewSonic P227f CRT monitor (screen dimensions: 40.5 x 30.5 cm), with a resolution of 1024 x 768 pixels, and a refresh rate of 100 Hz. The monitor was positioned at a distance of 62 cm from the participant's eyes.

The EyeLink 1000 was of the tower mount variety, and recorded the right eye at 1000 Hz. It was calibrated with a 9-point grid, with a 5-point grid or a 3-point line used as fall-backs if a 9-point calibration failed.

4.3.4. Task

The task in experiments 1, 2, and 4 consisted of two types of randomly intermixed trials: single and double stimulus presentation. There were two stimulus locations placed on the horizontal midline of the screen: one to the left, and one to the right of the screen centre. In experiment 3, there were four different stimulus locations: top-left, top-right, bottom-left, and bottom-right. These were equidistant from the screen's horizontal and vertical meridians, and were used in four combinations: top and bottom row for bi-lateral presentation, and left or right column for uni-lateral presentation. Stimuli had a diameter of 2.7 degrees of visual angle, and were presented at an eccentricity of 7.5 degrees of visual angle (measured from the screen centre).

Each trial started with the presentation of a central fixation cross. This was followed by the brief presentation of one or two stimuli (see section 4.3.5 *Stimuli* for more stimulus details). In experiments 1 – 3, exposure durations could be 10, 20, 40, 70, 120, 200, 360, and 580 ms. In addition, a 0 ms condition was used, in which the stimulus was not shown (instead, the mask was visible for an additional 10 ms). In experiment 4, the exposure duration was set at 200 ms. After the stimulus exposure, the stimuli were replaced by a mask (see section 4.3.5 *Stimuli* for more mask details). The mask consisted of coloured (experiments 1 and 3) or black-and-white (experiments 2 and 4) static elements, and was visible for 500 ms (experiments 1 – 3) or 1100 ms (experiment 4). After this, the central fixation cross disappeared, while the mouse cursor became visible, and two colour wheels appeared around the still-visible masks (only in experiment 1 and 3). The trial would end only if participants had made a response to both stimuli, *in the order of their own choosing*, and then pressed the space bar. See

Figure 4.1 for a visual representation of an example trial.

In order to recall the memoranda, participants were asked to use the mouse cursor to click on the stimulus or the colour wheel, which oriented itself towards the clicked location. By holding down the left mouse button, and dragging the mouse cursor around, participants could fine-tune their response. This allowed participants to choose their response from all possible colours or orientations, allowing them to precisely indicate what the memoranda looked like.

The benefit of using this continuous report (also referred to as the “method of adjustment”) is that it results in a quantifiable error: The difference between the presented stimulus and the participants report is a direct measure of the precision of participants’ recall. For example, when a stimulus was a grating with an orientation of 90 degrees, a participant’s response of 93 degrees is more accurate than a response of 120 degrees. This is in contrast with binary judgements where participants are asked whether a probe display is different from memoranda or not, where a response is either correct or not.

To prevent the stimulus to interfere with memorised features, masks remained visible until they were clicked. At the first click, the mask would turn into a colour (experiments 1 and 3) or oriented grating (experiment 2) that was of the colour or orientation that corresponded with the mouse click.

4.3.5. Stimuli

Colours in experiment 1 and 3 were sampled from CIE L*a*b*, a perceptually uniform colour space. Specifically, the colours were sampled from the L* = 50 plane, at an eccentricity of 22 around the whitepoint (where a* = 0 and b* = 0). They were presented as a uniformely coloured disk with a diameter of 2.7 degrees of visual angle. Within each condition cell, stimulus colours were chosen from a total of 360 different possible colours at uniform intervals, making it so that there was never a bias in probability density across the colour space. Mask colours were randomly sampled from the same 360 colours. The mask was a grid of 64 by 64 randomly coloured squares, masked by a disk with a diameter of 2.7 degrees of visual angle.

Orientations in experiments 2 and 4 were presented as a sinusoidal grating of 5 cycles per stimulus, masked by disks with a diameter of 2.7 degrees of visual angle and hard edges (essentially a Gabor with hard edges instead of a Gaussian mask). The mask stimulus consisted of a 32 by 32 grid of randomly chosen grey tones, masked by a disk with a diameter of 2.7 degrees of visual angle.

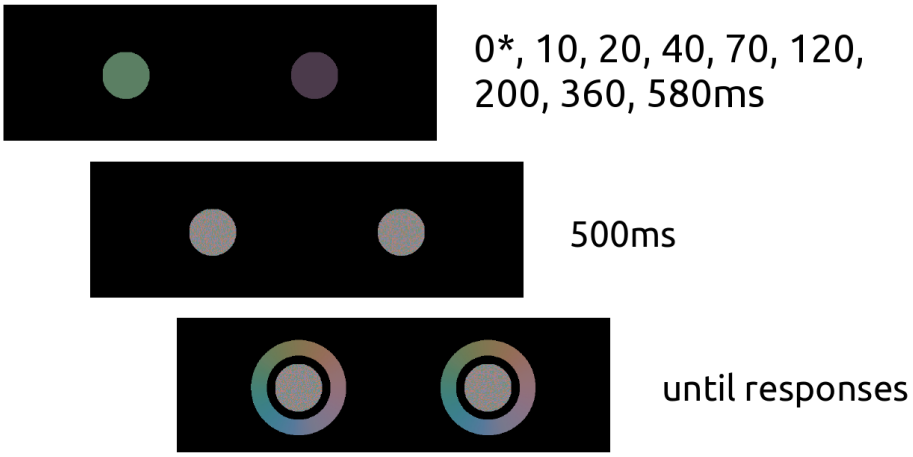

Figure 4.1 *– Trial from experiment 1 to illustrate the general structure of trials in all experiments. Each trial starts with the presentation of stimuli. After 10, 20, 40, 70, 120, 200, 360, or 580 milliseconds, each stimulus is replaced by stimulus-appropriate mask. 500 milliseconds after mask onset, the mouse cursor (not shown) becomes visible. In experiments that use coloured stimuli, cursor appearance is accompanied by the presentation of colour wheels (a cut-out from CIE L*a*b* space) around the stimulus.*

4.3.6. EEG acquisition

In experiment 4, a total of 67 Ag/AgCl sintered electrodes (EasyCap, Herrsching, Germany) were placed on each participant's head. Two were used for horizontal (one electrode on the non-nasal side of each eye), and another two for vertical EOG (one electrode over and one under the right eye). Two electrodes were used as reference electrodes (one placed on each mastoid). Sixty electrodes were placed in a cap according to the 10-20 system (**Figure 4.2**), and one electrode (AFz) was used as the ground. Data was recorded using a NeuroScan SynAmps RT amplifier and Curry 7 (Compumedics NeuroScan, Charlotte, NC, USA), at a rate of 1000 Hz. During

recording, electrode impedances were kept below 5 k Ohm.

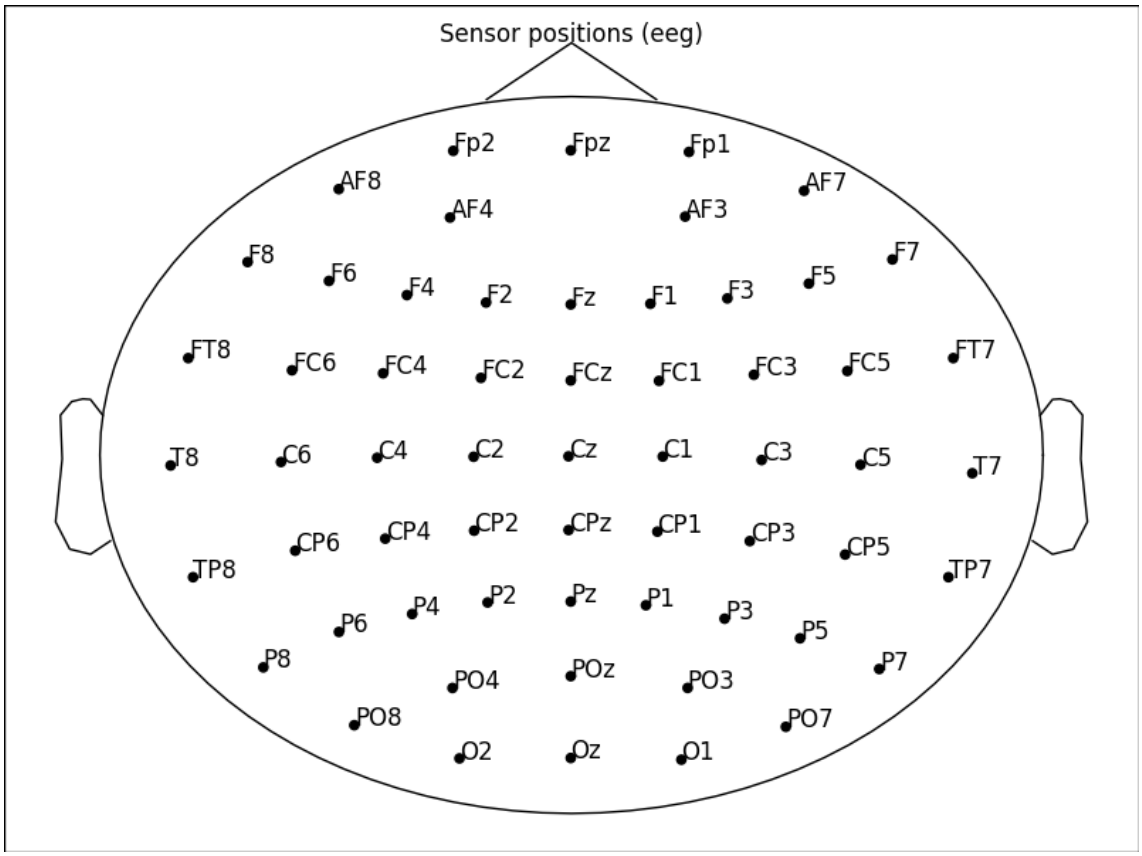

Figure 4.2 *– Placement of electrodes on a flattened representation of the scalp. Not shown is Afz, which was used as ground.*

4.3.7. EEG handling

During recording, an online notch filer for 50 Hz was used to filter out potential interference from the light net. In addition, a low-pass filter of 200 Hz was applied. Offline, an infinite-impulse response (IIR) filter was used to retain the frequencies between 1 and 45 Hz. After filtering, an independent-component analysis (ICA) was

used to identify the noise components in the EEG signal, which were automatically detected and zeroed out, after which the EEG signal was reconstructed from the components. Data epochs for each trial were selected from 100 ms before to 1200 ms after target onset. A logratio baseline correction was applied, using the pre-target onset window (-100 to 0 ms). Event-related potentials (ERP) were computed as the average signal within each condition. Time-frequency responses were computed using Morlet wavelets for frequencies of 1 to 45 Hz (in steps of 1 Hz), with F/2 cycles per frequency (where F is the frequency in Herz). The analysis was performed using MNE Python (Gramfort et al., 2013).

4.3.8. EEG Analysis

Electrodes of interest for ERP and time-frequency analyses were P7 and P8, both parietal electrodes, each located over a different hemisphere. In the single-stimulus condition, stimuli could appear on the left and the right of the central fixation. For each participant, the ERPs from electrodes contra-lateral to the stimulus were averaged, as were those from the ipsi-lateral electrode. These averaged ERPs were treated as time-series, and analysed on a per-sample basis. For each sample, a paired-samples t-test was performed to test the difference between the contra-lateral and the ipsi-lateral response. Naturally, the contra-lateral electrode was expected to show a response to the stimulus.

In the double-stimulus condition, the process was slightly more complex. Stimuli were shown left and right of the central fixation, and thus both electrodes of interest were expected to show stimulus-related activity. Because participant was

allowed to choose which item to respond to first (after the stimulus presentation of 200 ms, and the maintenance period of 1100 ms), it was known which stimulus the participant would respond to first. Therefore, it was possible to determine which electrode was contra-lateral to what would become the first-chosen stimulus. Within each participant, ERPs were averaged for this electrode, and for the electrode that was ipsi-lateral to the first response (and thus contra-lateral to the second response).

For the time-frequency analysis, the same combinations of contra- and ispi-lateral electrodes were used to compare against each other. Comparisons were made between each point in the time * frequency matrix.

4.3.9. Mixture model of short-term memory encoding

Because participants respond on a continuous scale, their responses range from 0 to 360 degrees in colour space (experiments 1 and 3), or 0-180 degrees in orientation space (experiments 2 and 4). I computed the error by taking the circular difference between participants' responses and the colour or orientation of the presented stimuli, and then transformed the result so that it would be in the range -pi to pi. The result is an error distribution where 0 corresponds with no error (perfect recall of the stimulus colour or orientation).

The error distribution for a single stimulus can be described as a combination of trials in which the stimulus was remembered (resulting in an error close to 0, with a certain deviation), trials where the stimulus was forgotten (and thus a guess had to be made), and when the stimulus was confused with a different stimulus (resulting in a

response that was close to the other stimulus).

Mathematically, the error distribution can be described as a combination of a Von Mises distribution centred around 0 to represent the trials in which the target was remembered, a uniform distribution to represent the trials in which a participant guessed (because each response is equally likely then) (Zhang & Luck, 2008), and a Von Mises distribution centred around the other stimulus to represent the trials in which a participant made a swap error (Bays et al., 2009).

This approach works when participants are responding to one item, usually randomly selected from an array of several items. However, in the current case participants respond to two items at the same time. A computational account for this has to consider two error distributions at the same time, one for each item.

In each trial, there are three possibilities: 1) A participant remembers both items, and the error distributions for the first and second response can be described by two Von Mises distributions; 2) remembers only one of the items and guesses the other, in which case the error distributions for the first and second response can be represented by one Von Mises and one uniform distribution; or 3) remembers neither item and has to guess both, and the error distributions can be described by two uniform distributions.

In this model, there are five parameters: The probability of two items being remembered (P_{both}), the probability of one item being remembered and not the other (P_{one}), and the probability of neither item being remembered (P_{none}); as well as the spreading parameters of the Von Mises distributions that describe the first (SD_{first}) and the second response (SD_{second}). Because P_{both}, P_{one}, and P_{none} have to add up to 1, only two are free parameters, and the third is dictated by the other two ($P_{none} = 1 - P_{both} - P_{one}$).

This is summarised in equation 1.

(1)

$$P_{(both)}\dot{\Phi}(0,\kappa_{first})\dot{\Phi}(0,\kappa_{second})+w_{first}(P_{(one)}\dot{\Phi}(0,\kappa_{first})\frac{\dot{1}}{2\pi})+w_{second}(\frac{1}{2\pi}\dot{P_{(one)}}\dot{\Phi}(0,\kappa_{second}))+P_{(non)}\frac{\dot{1}}{2\pi}\frac{\dot{1}}{2\pi}$$

Where:

- $\Phi(0, \kappa_{first})$ is a Von Mises distribution with centre 0 and spreading parameter κ_{first}.
- κ_{first} is the spreading parameter for the error distribution associated with the first response, which can be converted to standard deviation SD_{first}.
- κ_{second} is the spreading parameter for the error distribution associated with the second response, which can be converted to standard deviation SD_{second}.
- $1 / 2\pi$ is the uniform distribution associated with random guessing.
- w_{first} is the proportion of P_{one} trials in which the first stimulus is remembered and the second forgotten.
- w_{second} is the proportion of P_{one} trials in which the first stimulus is forgotten and the second is remembered. It is defined as $1 - w_{first}$.

While equation 1 describes both error distributions at the same time, its parameters can be used to describe the individual error distributions. Equation 2 describes the error distribution for the first response, and equation 3 describes the error distribution for the second response.

$$(2) \quad p(\theta_{first}) = P_{(both)}\dot{\Phi}(0, \kappa_{first}) + P_{(one)}w_{first}\dot{\Phi}(0, \kappa_{first}) + P_{(one)}w_{second}\frac{\dot{1}}{2\pi} + P_{(non)}\frac{\dot{1}}{2\pi}$$

$$(3) \quad p(\theta_{second}) = P_{(both)}\dot{\Phi}(0, \kappa_{second}) + P_{(one)}w_{first}\frac{\dot{1}}{2\pi} + P_{(one)}w_{second}\dot{\Phi}(0, \kappa_{second}) + P_{(non)}\frac{\dot{1}}{2\pi}$$

Crucially, according to the serial encoding hypothesis one item is encoded first, and the other item is encoded second. Thus, it predicts that the probability of remembering only one item P_{one} would increase before the probability of remembering two items P_{both}. On the other hand, according to the parallel encoding hypothesis both items will be encoded at the same time, and thus the probability of both items being remembered P_{both} should rise without a preceding spike in the probability of one item being remembered P_{one}.

It should be noted that even under the parallel encoding hypothesis P_{one} will be 0, due to the stochastic nature of the model. If we assume two channels are encoding one item each with a probability of P_{both} that both items are encoded with enough quality to be recalled at the time of the mask onset (which blocks further encoding of the items), then the probability that only one channel has finished encoding an item is given by

equation 4.

(4) $$P_{(one|both)} = 2\sqrt{\dot{P}_{(both)}}\left(1-\sqrt{\dot{P}_{(both)}}\right)$$

4.3.10. Model simulations

The model outlined in equations 1-4 is new, and thus it is yet unclear whether it is capable of accurately describing error distributions, or whether it can distinguish between serial and parallel encoding. To illustrate that the model can accurately capture both properties, I fitted the described model to simulated data. Simulations are powerful tools, as they allow the construction of a dataset of which the ground truth is know. How well the fitted parameters compare to the ground truth quantifies how good the model can capture the data. In addition, both serial and parallel encoding can be simulated by using different simulation parameters. If the model shows differentiation between the two encoding simulations, than one can be confident that the model would also be able to differentiate between serial and parallel encoding in empirical data.

Two simulations will be run, one for serial encoding and one for parallel encoding. In each, eight exposure durations will be simulated: 0, 10, 20, 40, 70, 120, 200, and 360 milliseconds. For each exposure duration, N trials will be simulated by randomly drawing two samples from a different underlying ‘true’ error distribution for each. These will reflect the first response and the second response that would be given by real participants.

Non-guessed first responses are drawn from a Von Mises distribution centred around 0 radians, with a standard deviation of 0.4 radians. Non-guessed second responses will be drawn from a Von Mises distribution centred around 0 radians, with a standard deviation of 0.6 radians. The standard deviation of the second response if higher than that of the first response, because the second response is expected to be of lower precision. In addition, guessed first and second responses will be drawn from uniform distributions between negative pi and pi radians.

The relative number of trials guessed and non-guessed trials will also be simulated. Specifically, it will be determined by the simulated number of trials in which both stimuli are remembered (neither guessed), the amount of trials that the first stimulus is remembered (second response is a guess), and the amount of trials in which the second item is remembered (first response is a guess). The simulated proportions of trials in which both or one stimulus is remembered, is varied according to two factors: Exposure duration, and encoding mechanism. The precise values are reported in Table 4.1, and in general abide by two rules: 1) The higher the exposure duration, the higher proportion of trials will have two non-guessed responses, and 2) During serial encoding the proportion of only one item being remembered peaks before the proportion of trials in which both trials are remembered can rise, whereas in parallel encoding the proportion of both items being remembered rises gradually and the chance of only one item being remembered follows equation 4.

Each simulation will be run with 1000 trials per cell to establish whether the employed methods can accurately capture the underlying ground truth. In addition, the same simulations will be run with 50 trials per cell, which corresponds to the amount of

trials per cell for each individual participant. If the fitting is still accurate at 50 trials, it means one can be confident in the parameter fits for individual participants.

Table 4.1 *– Parameters per simulation. Values rounded to two decimals. Simulated standard deviations (SD) are the same for serial and parallel simulations.*

Exposure duration (ms)	**Serial P(both)**	**Serial P(one)**	**Parallel P(both)**	**Parallel P(one)**	**SD_{first}**	**SD_{second}**
0	0.00	0.00	0.00	0.00	0.40	0.60
10	0.00	0.00	0.00	0.00	0.40	0.60
20	0.00	0.00	0.00	0.00	0.40	0.60
40	0.00	0.40	0.10	0.43	0.40	0.60
70	0.10	0.80	0.20	0.49	0.40	0.60
120	0.60	0.30	0.60	0.35	0.40	0.60
200	0.80	0.10	0.80	0.19	0.40	0.60
360	0.80	0.10	0.80	0.19	0.40	0.60

At 1000 trials per cell, the model fitting was able to highly accurately describe the simulated error distributions (**Figure 4.3**). Even at 50 trials, the simulated error distributions were well described by the parameter estimates from the employed model fitting procedure (**Figure 4.4**). Note that the simulated distributions were fitted to equation 1. The error distributions were then reconstructed from the fitted parameters using equations 2 and 3.

0 ms

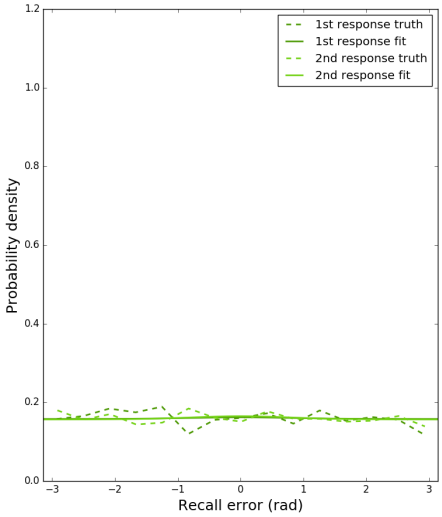
1st response truth
1st response fit
2nd response truth
2nd response fit
Probability density
Recall error (rad)

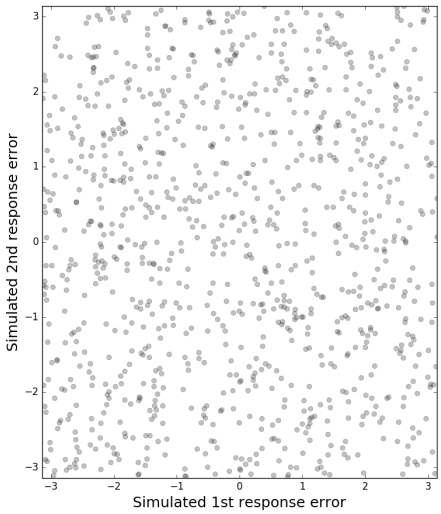
Simulated 2nd response error
Simulated 1st response error

10 ms

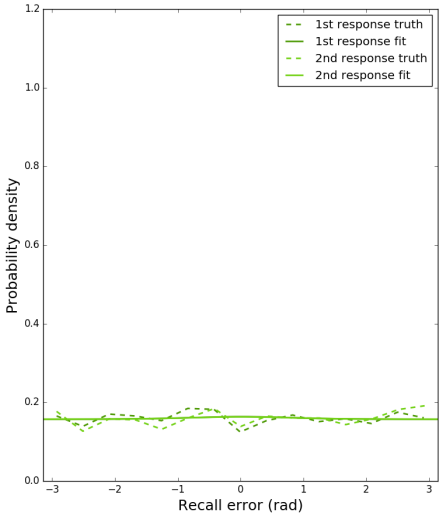
1st response truth
1st response fit
2nd response truth
2nd response fit
Probability density
Recall error (rad)

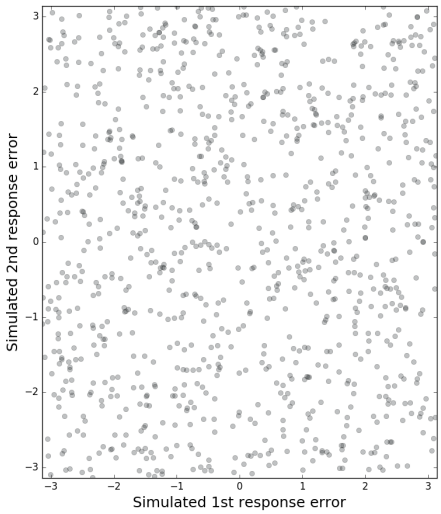
Simulated 2nd response error
Simulated 1st response error

20 ms

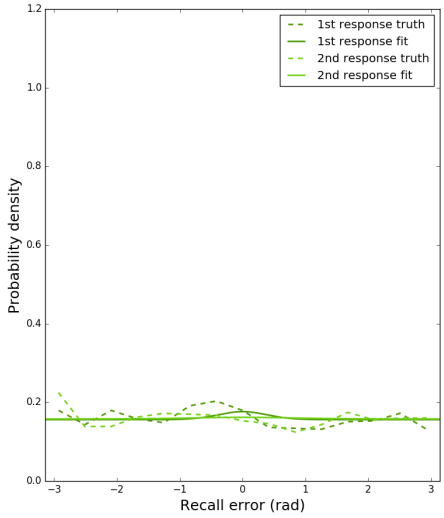
1st response truth
1st response fit
2nd response truth
2nd response fit
Probability density
Recall error (rad)

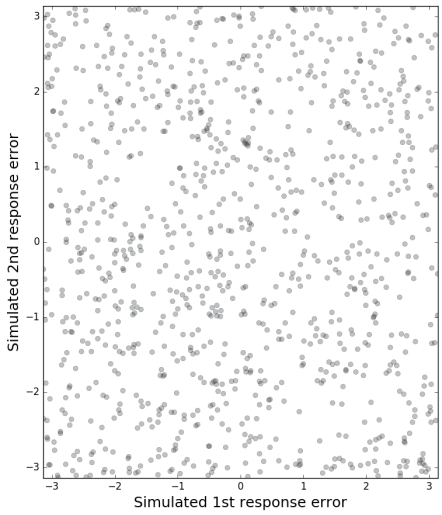
Simulated 2nd response error
Simulated 1st response error

40 ms

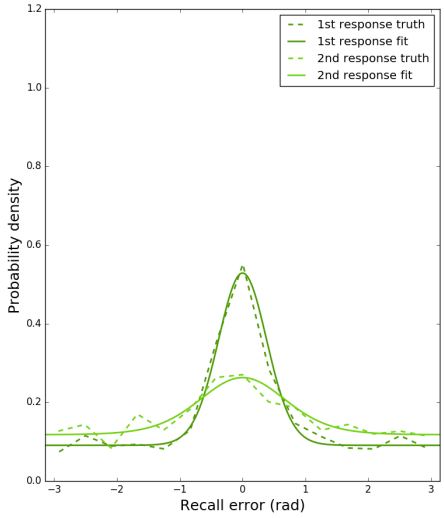
1st response truth
1st response fit
2nd response truth
2nd response fit
Probability density
Recall error (rad)

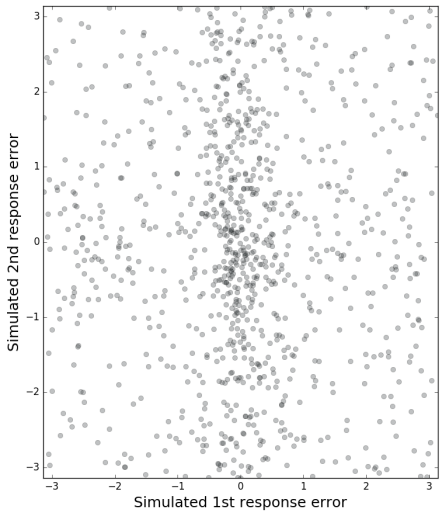
Simulated 2nd response error
Simulated 1st response error

70 ms

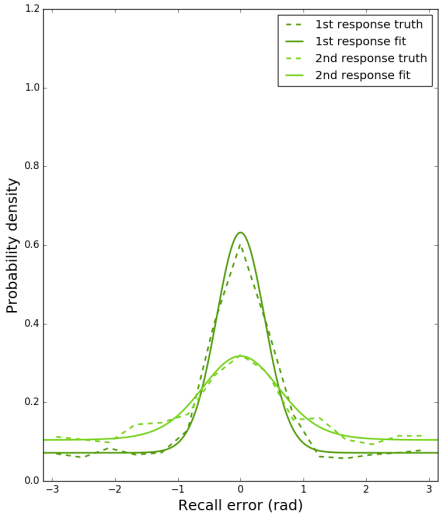
1st response truth
1st response fit
2nd response truth
2nd response fit
Probability density
Recall error (rad)

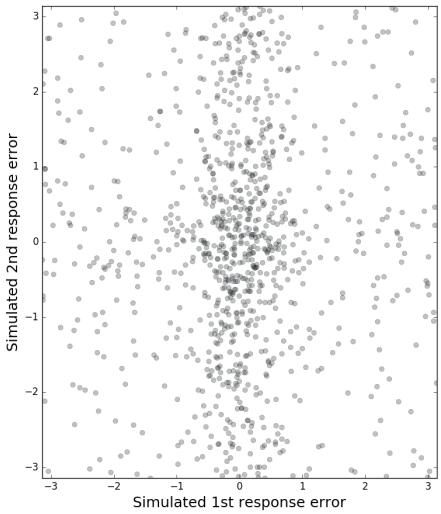
Simulated 2nd response error
Simulated 1st response error

120 ms

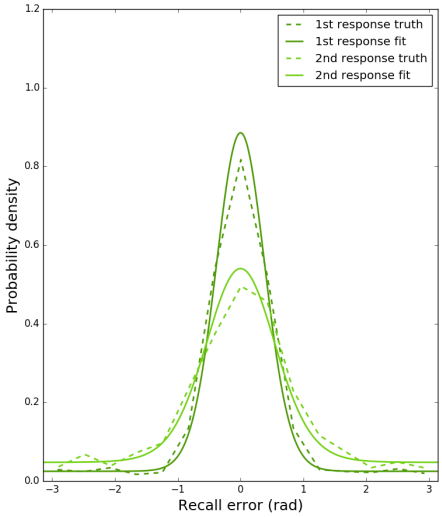
1st response truth
1st response fit
2nd response truth
2nd response fit
Probability density
Recall error (rad)

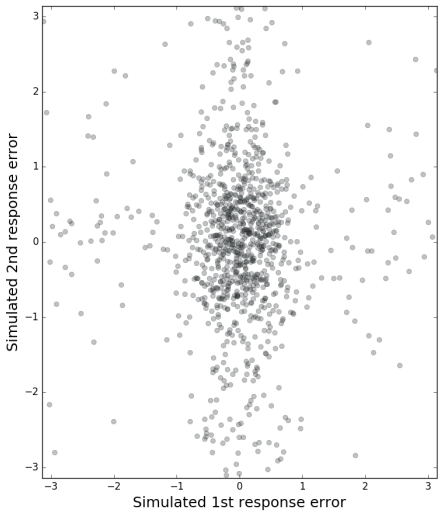
Simulated 2nd response error
Simulated 1st response error

200 ms

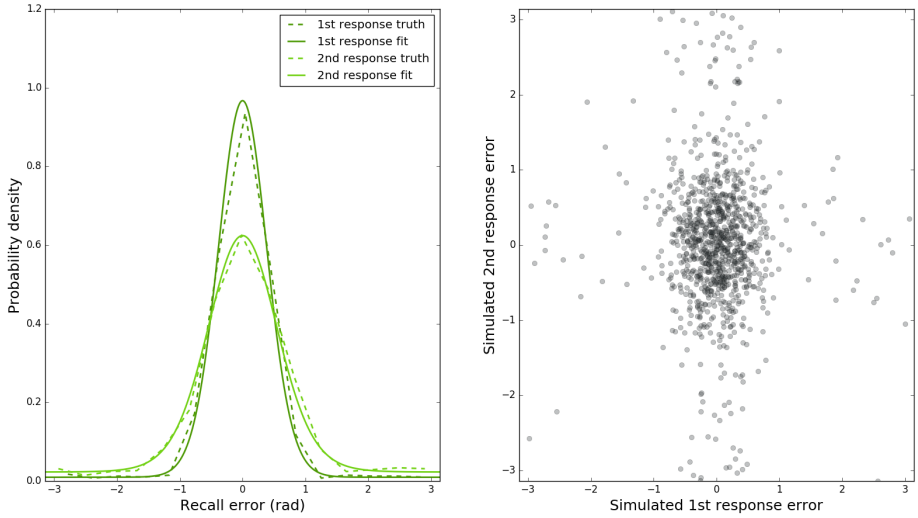

360 ms

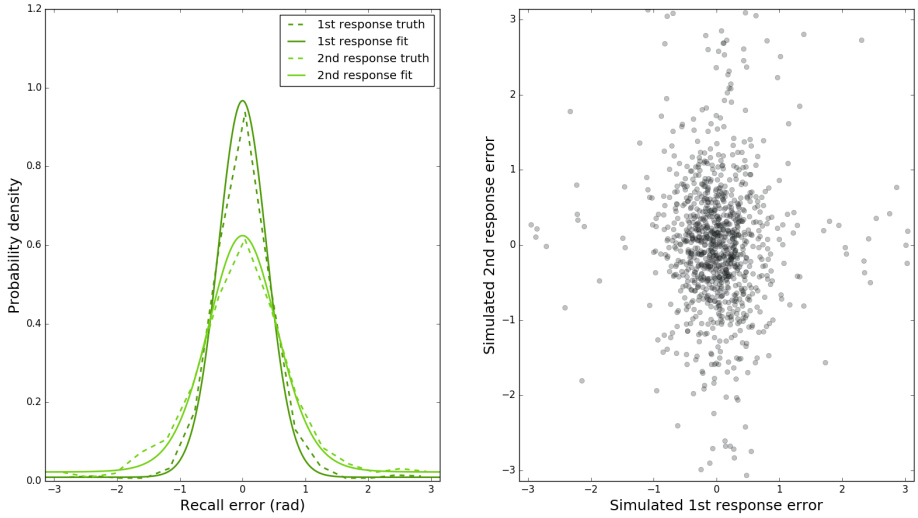

Figure 4.3 *– Simulated first and second responses in a whole-report task with two stimuli. Each row reflects an exposure duration. First and second samples (N=1000) are plotted in the scatter plots in the right column, with the first response error on the x-axis, and the second response error on the y-axis. Each dot represents a simulated sample. In the left column, histograms of the simulated errors for the first and second response are drawn in dashed lines. The solid lines are the associated model fits. Dark lines are for the first response; lighter lines for the second response.*

0 ms

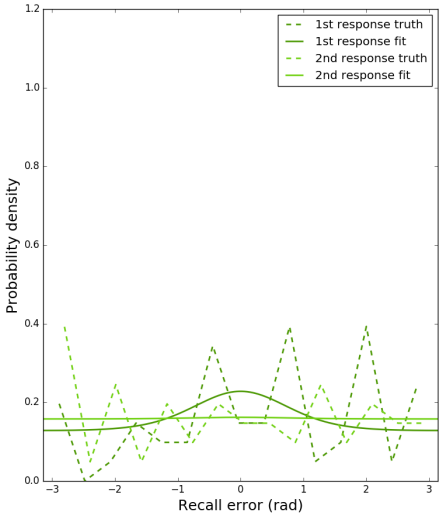
1st response truth
1st response fit
2nd response truth
2nd response fit
Probability density
Recall error (rad)

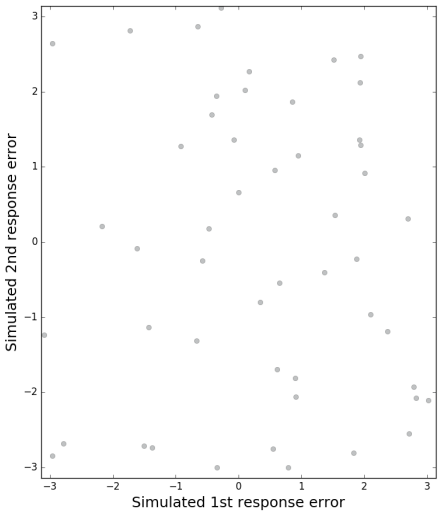
Simulated 2nd response error
Simulated 1st response error

10 ms

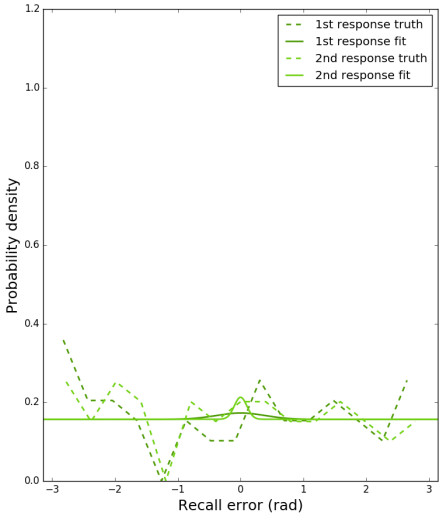
1st response truth
1st response fit
2nd response truth
2nd response fit
Probability density
Recall error (rad)

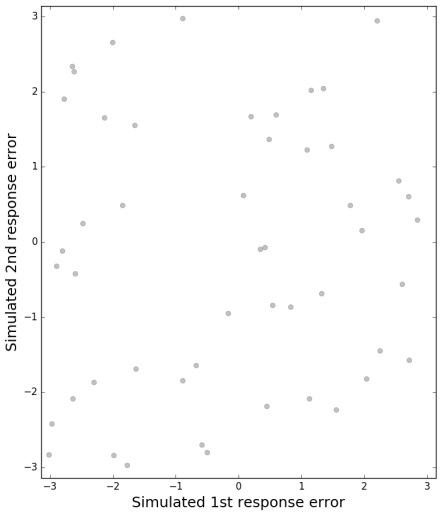
Simulated 2nd response error
Simulated 1st response error

20 ms

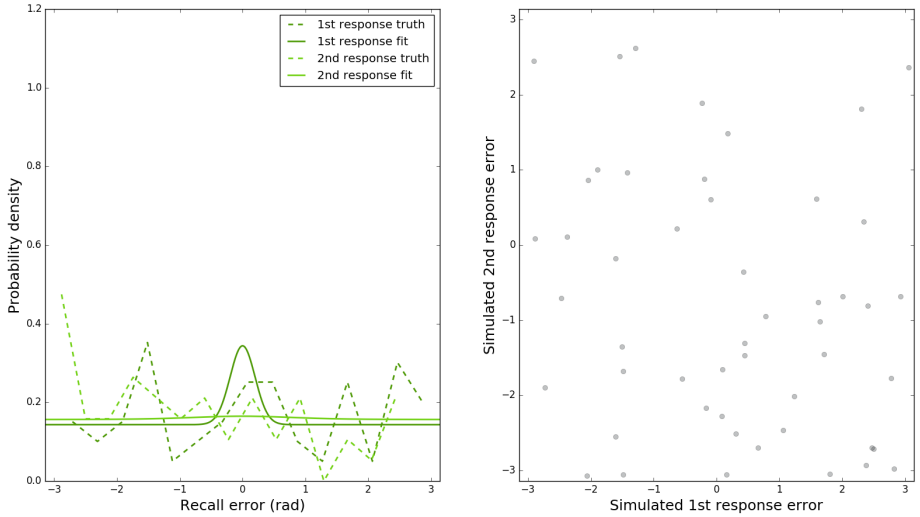
1st response truth
1st response fit
2nd response truth
2nd response fit
Probability density
Recall error (rad)
Simulated 2nd response error
Simulated 1st response error

40 ms

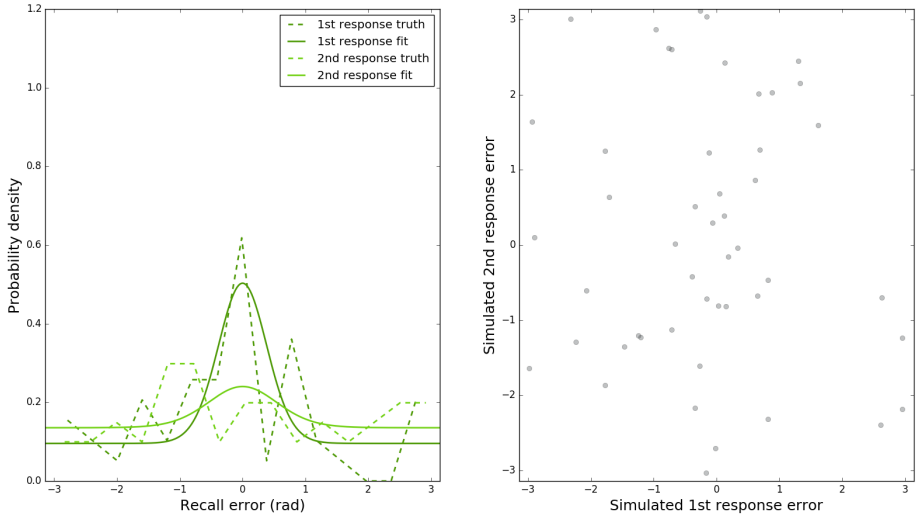
1st response truth
1st response fit
2nd response truth
2nd response fit
Probability density
Recall error (rad)
Simulated 2nd response error
Simulated 1st response error

70 ms

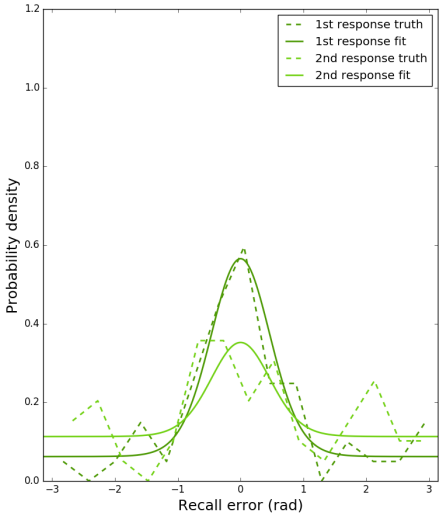
1st response truth
1st response fit
2nd response truth
2nd response fit
Probability density
Recall error (rad)

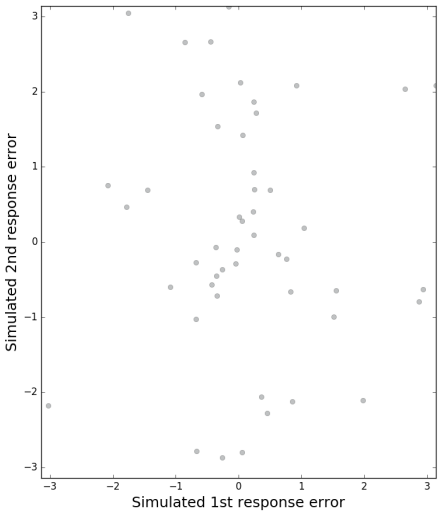
Simulated 2nd response error
Simulated 1st response error

120 ms

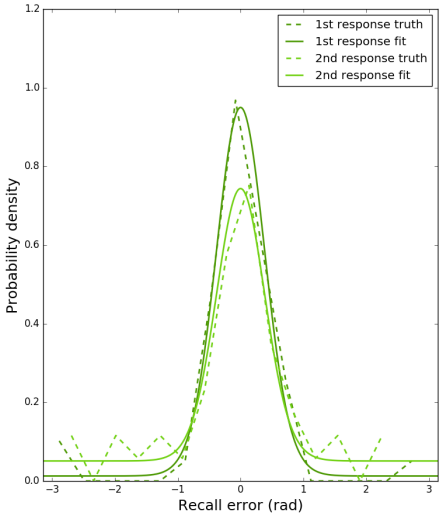
1st response truth
1st response fit
2nd response truth
2nd response fit
Probability density
Recall error (rad)

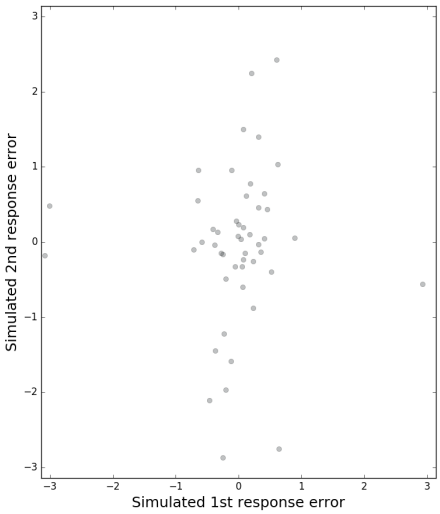
Simulated 2nd response error
Simulated 1st response error

200 ms

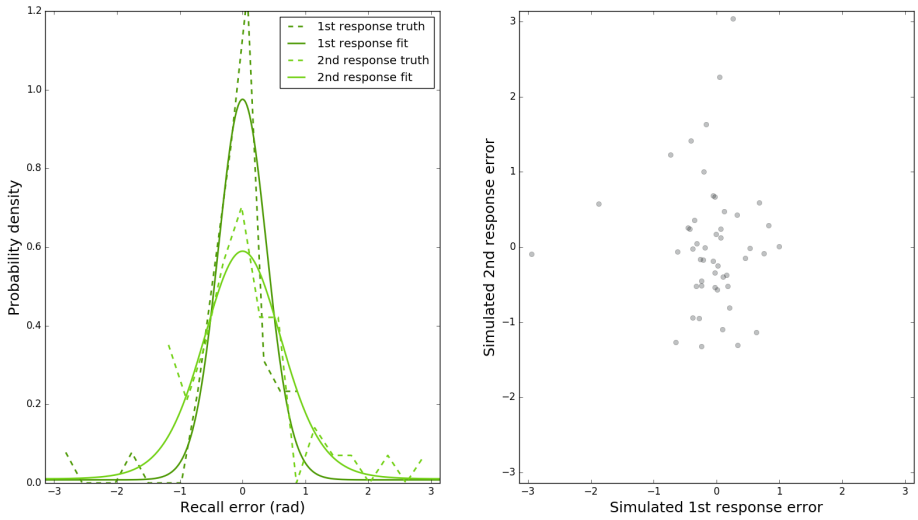

360 ms

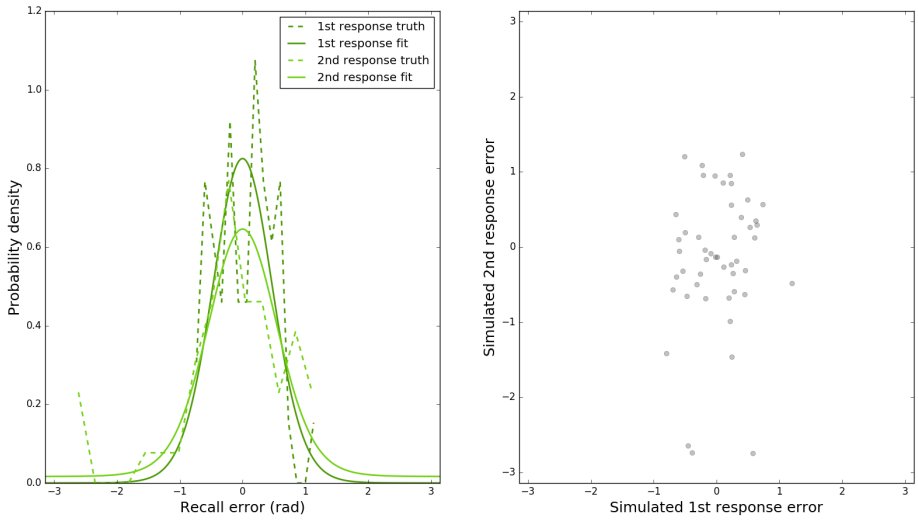

Figure 4.4 *– Simulated first and second responses in a whole-report task with two stimuli. Each row reflects an exposure duration. First and second samples (N=50) are plotted in the scatter plots in the right column, with the first response error on the x-axis, and the second response error on the y-axis. Each dot represents a simulated sample. In the left column, histograms of the simulated errors for the first and second response are drawn in dashed lines. The solid lines are the associated model fits. Dark lines are for the first response; lighter lines for the second response.*

The crucial comparison between the serial and parallel encoding hypotheses is whether P(one) (the proportion of trials in which one item is remembered and one is guessed) abides by equation 4. This equation describes the chance level P(one) in a parallel encoding framework. This chance level P(one) is referred to as P(one|both), because it depends on P(both). If the model can accurately distinguish between serial and parallel encoding, the serial simulation should present an estimated P(one) that peaks above the predicted value from equation 4. By contrast, the parallel simulation should show a P(one) that closely resembles P(one|both).

In the 1000 trials simulation, the expected pattern is exactly what can be observed after model fitting (**Figure 4.5**). Even the 50 trials simulation abides quite well to the expected pattern (**Figure 4.6**).

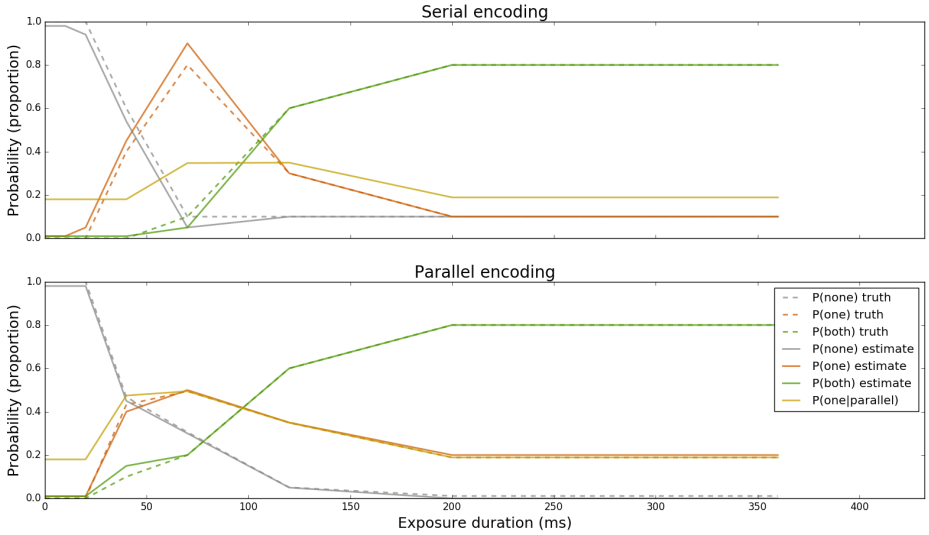

Figure 4.5 *– Fitted parameters to a simulation (1000 trials per cell). In green is the probability that both stimuli were remembered P(both), i.e. in which the first and second response errors were randomly drawn from two Von Mises distributions. In orange is the probability that one item was remembered but not the other P(one), i.e. that one response error was drawn from a Von Mises distribution and the other response error from a uniform distribution. In grey is the chance that neither stimulus is remembered P(none), which is 1 minus the probabilities that both stimuli or one stimulus were remembered. The yellow line describes 'chance level' P(one) when encoding occurs in parallel P(one|both), computed through equation 4. Solid lines are fitted parameters, dotted lines are the actual values used in the simulations. The top panel shows a serial encoding simulation, which is accurately picked up by the fitted model parameters that clearly show P(one) to peak over P(one|both). In the parallel simulation, the empirical P(one) and chance-level P(one|both) overlap, demonstrating that the model can also detect parallel encoding.*

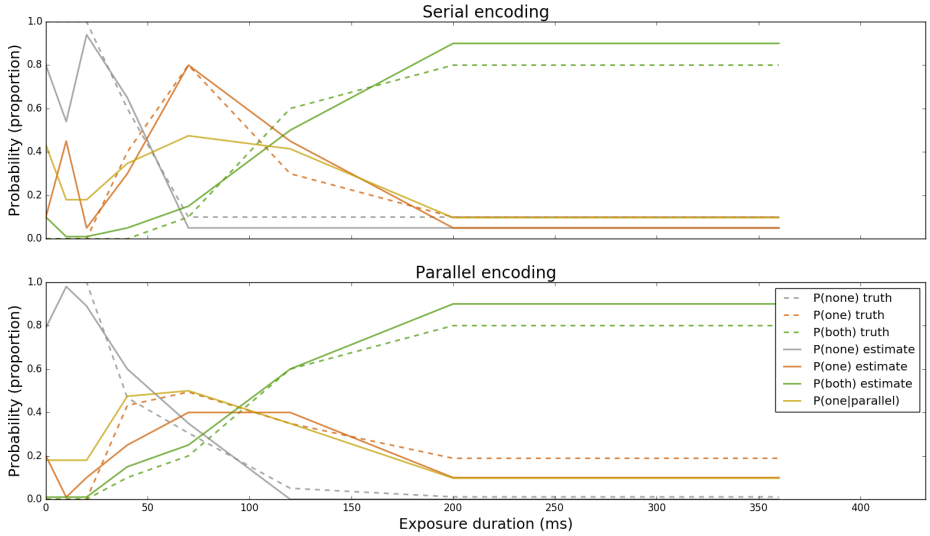

Figure 4.6 *– Fitted parameters to a simulation (50 trials per cell). In green is the probability that both stimuli were remembered P(both), i.e. in which the first and second response errors were randomly drawn from two Von Mises distributions. In orange is the probability that one item was remembered but not the other P(one), i.e. that one response error was drawn from a Von Mises distribution and the other response error from a uniform distribution. In grey is the chance that neither stimulus is remembered P(none), which is 1 minus the probabilities that both stimuli or one stimulus were remembered. The yellow line describes 'chance level' P(one) when encoding occurs in parallel P(one|both), computed through equation 4. Solid lines are fitted parameters, dotted lines are the actual values used in the simulations. The top panel shows a serial encoding simulation, which is accurately picked up by the fitted model parameters that clearly show P(one) to peak over P(one|both). In the parallel simulation, the empirical P(one) and chance-level P(one|both) overlap, demonstrating that the model can also detect parallel encoding.*

In sum, the simulations show that my fitting procedure for equation 1 works: Individual error distributions can be reconstructed from the estimated parameters using equations 2 and 3, and comparing the parameter estimate for P(one) can be compared to the result of equation 4 (using the estimated P(both)) to distinguish between serial and parallel encoding.

4.3.11. Analysis

In experiments 1 and 2, the design was within-participants, and had two factors: exposure duration (nine levels: 0, 10, 20, 40, 70, 120, 200, 360, and 580 ms), stimulus (three levels: single response in the one-item condition, first response in the two-item condition, second response in the two-item condition), and number of stimuli (two levels: one and two items). In experiment 3, there were three factors: exposure duration (same nine levels), stimulus (two levels: first and second response), and visual field of presentation (two levels: uni-lateral and bi-lateral).

The absolute angular error is computed as the circular difference between participants' responses and the target stimuli's orientation of colour. It was analysed using one repeated-measures ANOVA per experiment, with the factors and levels outlined above. Post-hoc tests were performed when appropriate (i.e. when a main or interaction effect was statistically significant), using related-samples t-tests.

The aforementioned model was fitted to the error distributions for the first and second response in the two-item condition, obtained for each exposure duration, using log-likelihood maximisation (cf. (Bays et al., 2009)). This resulted in estimates of the

parameters P_{both}, P_{one}, P_{none}, SD_{first}, and SD_{second} for each exposure duration. In addition, using equation 2, the chance-level P_{one} could be obtained from P_{both}. The crucial comparison was between the empirical P_{one} and the chance-level P_{one}. This comparison was performed using a repeated-measures ANOVA with two factors: exposure duration (nine levels, as above), and P_{one} type (two levels: empirical and chance).

Related-samples t-tests were performed to test the difference between the empirical and chance P_{one} for every exposure duration, because a difference is expected only at low exposure durations (under the serial encoding hypothesis).

To account for multiple comparisons, a Holm-Bonferroni correction (Holm, 1979) was performed for each experiment separately. This correction aims to reduce the family-wise error rate, and is performed by sorting all p-values within a family of tests from lowest to highest. With an overall alpha set at 0.05, the alpha level for each test increases with rank order. That is, for *m* tests, the corrected alpha level for the p-value at index *k* would be 0.05 divided by $m + 1 – k$. Here, we do 9 t-tests each time (one for each exposure duration), so the most stringent alpha would be $0.05 / (9 + 1 – 1) = 0.0056$, and the least stringent would be $0.05 / (9 + 1 – 9) = 0.05$. If the first p-value is below the first corrected alpha, all p-values until (but not including) the first p-value that exceeds its corrected alpha-level are considered significant.

4.4. Results

4.4.1. Experiment 1 – Whole-report of coloured disks

4.4.1.1. Absolute angular error

To assess the effects of stimulus type (single stimulus, double-stimulus first response, and double-stimulus second response) and exposure duration on the average absolute angular error (**Figure 4.7**), a repeated-measures ANOVA was conducted. This revealed a main effect of exposure duration, $F(8, 80) = 178.24$, $p < 0.001$, $\eta_{partial}^2 = 0.95$, a main effect of stimulus, $F_{Greenhouse\text{-}Geisser}(1.10, 10.95) = 39.78$, $p < 0.000$, $\eta_{partial}^2 = 0.80$, and an interaction effect between stimulus and exposure duration, $F(16, 160) = 19.24$, $p < 0.000$, $\eta_{partial}^2 = 0.66$.

Post-hoc related-samples t-tests revealed no statistically significant difference between the single item and the first response in the double-item condition (all p were over a Holm-Bonferroni corrected alpha level of 0.05). There was also no statistically significant difference between the first and the second response in the double-item condition for exposure durations 0 – 40 ms (all p were over a Holm-Bonferroni corrected alpha level of 0.05), but a difference became apparent from 70 – 580 ms (all $p < 0.002$).

These results indicate that recall accuracy was better in the single-item condition, and better for the first response in the two-item condition, compared to the second response in the two-item condition. In addition, recall error improved with exposure time.

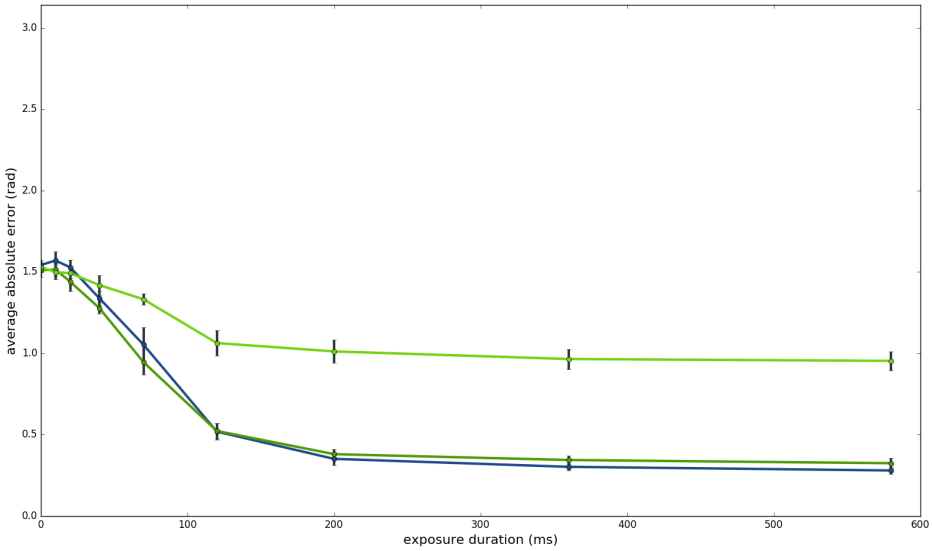

Figure 4.7 *– Average absolute error (circular distance between the target and the response; y-axis) as a function of exposure duration (x-axis). Average errors in the single-stimulus condition is plotted in dark blue. Average error in the two-stimulus condition is plotted in green for the first response, and light green for the second response. Error bars indicate standard error of the mean.*

4.4.1.2. Encoding model parameters

Recall that under the serial hypothesis a significant difference between the empirical and chance P_{one} is expected only at low exposure durations. The difference between the empirically obtained and chance P_{one} (**Figure 4.8**) was investigated using a repeated-measures ANOVA. This revealed a statistically significant main effect of exposure duration, $F_{Greenhouse\text{-}Geisser}(2.7, 27.0) = 4.79$, $p = 0.010$, $\eta_{partial}^2 = 0.32$, and an interaction effect between P_{one} type and exposure duration, $F_{Greenhouse\text{-}Geisser}(3.86, 38.62) =$

3.44, $p = 0.018$, $\eta_{partial}^2 = 0.26$. There was no main effect of P_{one} type, $F(1, 10) = 1.54$, $p = 0.243$. Planned related-samples t-tests revealed no statistically significant differences (**Table 4.2**) between the empirical and chance P_{one} at any exposure duration.

Table 4.2*. Results from related-samples t-tests between the empirical and chance P_{one} at each exposure duration. The alpha level is Holm-Bonferroni corrected. Statistically significant differences are marked with an asterisk.*

Exposure duration (ms)	*t* value	*p* value	alpha
0	-2.10	0.062	0.00833
10	-1.71	0.119	0.0125
20	-0.69	0.508	0.0500
40	1.04	0.325	0.0250
70	2.56	0.028	0.00556
120	2.27	0.047	0.00714
200	1.96	0.078	0.0100
360	1.58	0.145	0.0167
580	2.34	0.041	0.00625

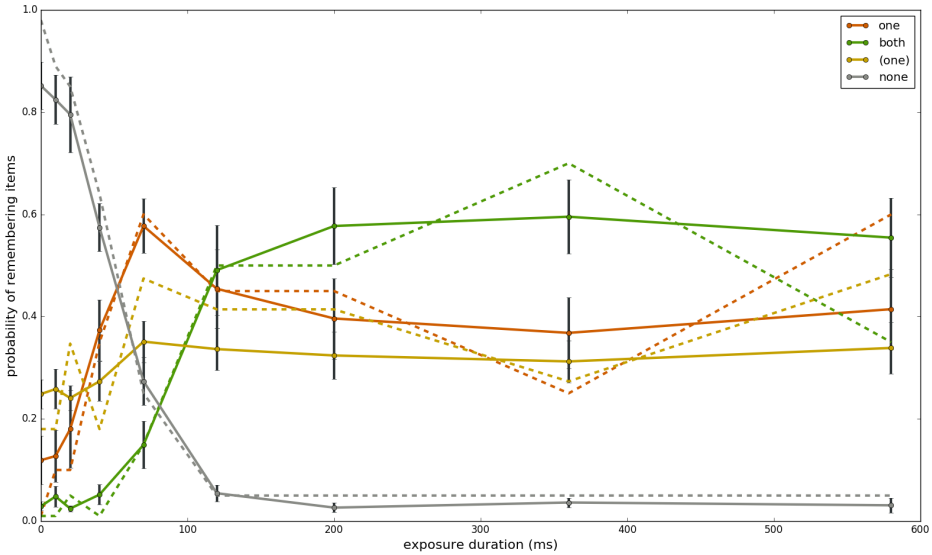

Figure 4.8 *– Fitted probabilities of both items being remembered (green), only one item being remembered (orange line), and neither item being remembered (grey line). The yellow line is the chance-level P(one) when encoding would occur in parallel, computed using equation 4 and the estimated P(both). Solid lines indicate averaged participant-level fits, and dashed lines indicate group-level fits (all participant data taken together within the same model). Error bars indicate standard error of the mean.*

4.4.1.3. Summary

These results indicate that participants' recall improved with increasing exposure duration. The lack of a difference between the empirical and chance P_{one} is in line with the parallel encoding hypothesis.

These data were collected with coloured circles sampled from CIE L*a*b* space. To replicate the findings, I conducted the same whole-report experiment with oriented gratings. These were briefly presented for the same range of exposure

durations. Again, in each trial one or two stimuli could be presented on the left and/or the right of the display. Participants were again free to choose the order in which they responded in the two-stimulus trials.

4.4.2. Experiment 2 – Whole-report of oriented gratings

4.4.2.1. Absolute angular error

To assess the effects of stimulus type (single stimulus, double-stimulus first response, and double-stimulus second response) and exposure duration on the average absolute angular error (**Figure 4.9**), a repeated-measures ANOVA was conducted. This revealed a main effect of exposure duration, $F_{Greenhouse\text{-}Geisser}(19.52, 44.04) = 147.77$, $p < 0.001$, $\eta_{partial}^2 = 0.91$, a main effect of stimulus, $F_{Greenhouse\text{-}Geisser}(1.30, 19.52) = 68.59$, $p < 0.000$, $\eta_{partial}^2 = 0.82$, and an interaction effect between stimulus and exposure duration, $F(16, 240) = 17.89$, $p < 0.000$, $\eta_{partial}^2 = 0.54$.

Post-hoc related-samples t-tests revealed no statistically significant difference between the single item and the first response in the double-item condition (all p were over a Holm-Bonferroni corrected alpha level of 0.05), save from at an exposure duration of 120 ms ($p < 0.001$). There was also no statistically significant difference between the first and the second response in the double-item condition for exposure durations 0 – 10 ms (both p were over a Holm-Bonferroni corrected alpha level of 0.05), but a difference became apparent from 20 – 580 ms ($p = 0.023$ at 20 ms, and all $p < 0.001$ from 40 ms).

As in experiment 1, these results indicate that recall accuracy was better in the single-item condition, and better for the first response in the two-item condition, compared to the second response in the two-item condition. In addition, recall error improved with exposure time.

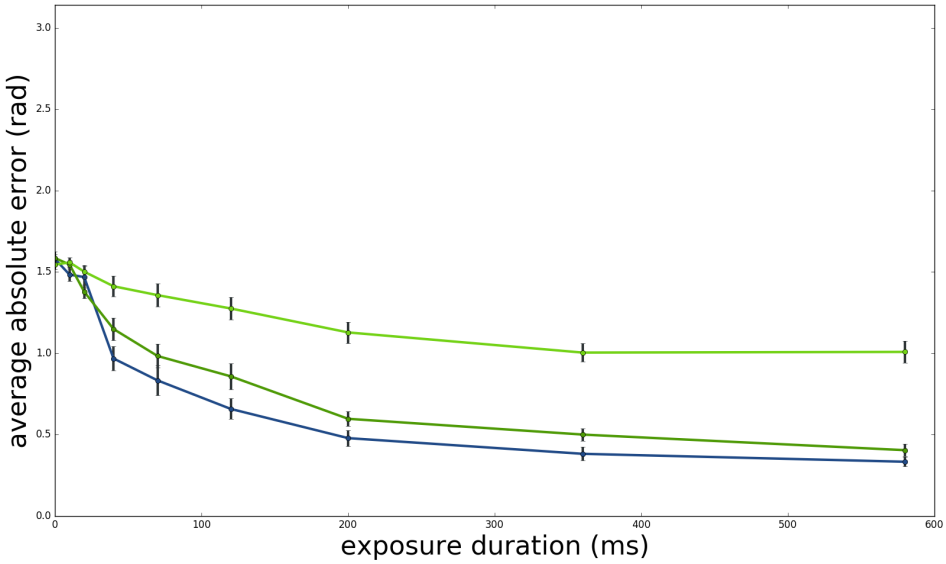

Figure 4.9 *– Average absolute error (circular distance between the target and the response; y-axis) as a function of exposure duration (x-axis). Average errors in the single-stimulus condition is plotted in dark blue. Average error in the two-stimulus condition is plotted in green for the first response, and light green for the second response. Error bars indicate standard error of the mean.*

4.4.2.2. Encoding model parameters

The difference between the empirically obtained and chance P_{one} (**Figure 4.10**) was investigated using a repeated-measures ANOVA. This revealed a statistically

significant main effect of exposure duration, $F_{Greenhouse\text{-}Geisser}(4.06, 60.93) = 18.63$, $p < 0.001$, $\eta_{partial}^2 = 0.55$, and an interaction effect between P_{one} type and exposure duration, $F(8, 120) = 7.05$, $p < 0.001$, $\eta_{partial}^2 = 0.32$. There was no main effect of P_{one} type, $F(1, 15) = 2.58$, $p = 0.129$. Planned related-samples t-tests revealed that at 0 and 10 ms the empirical P_{one} was statistically significantly lower than chance (potentially reflecting a 'better-than-parallel' encoding; see below). At 200 ms, the empirical P_{one} exceeded chance, which is somewhat in line with the serial encoding hypothesis. All other differences between the empirical and chance P_{one} at any exposure duration were not statistically significant. For an overview, see **Table 4.3**.

Table 4.3*. Results from related-samples t-tests between the empirical and chance P_{one} at each exposure duration. The alpha level is Holm-Bonferroni corrected. Statistically significant differences are marked with an asterisk.*

Exposure duration (ms)	*t* value	*p* value	alpha
0	-3.97	0.001*	0.00625
10	-4.10	< 0.001*	0.00556
20	-1.44	0.170	0.0250
40	2.13	0.050	0.0125
70	2.65	0.018	0.0100
120	0.97	0.346	0.0500
200	3.68	0.002*	0.00714
360	1.72	0.106	0.0167
580	2.75	0.015	0.00833

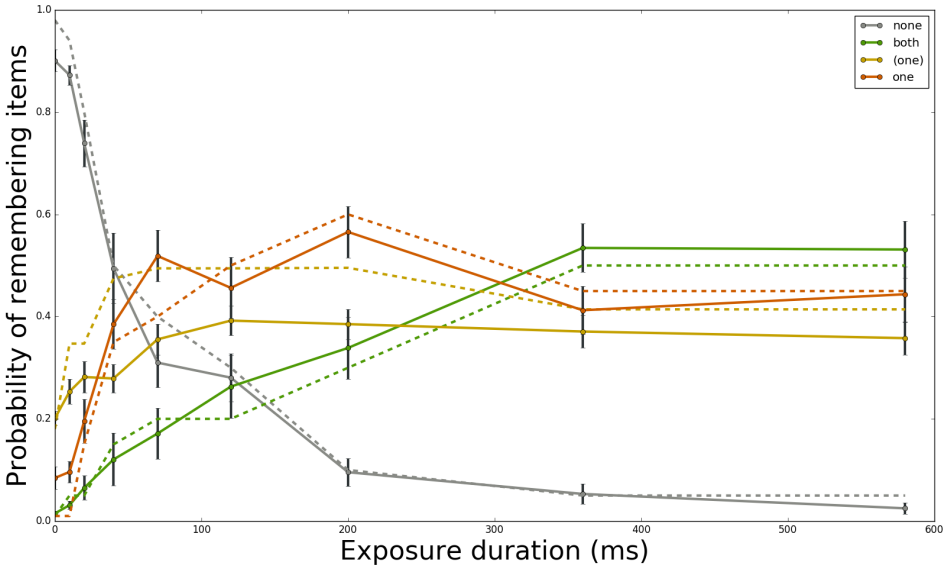

Figure 4.10 *– Fitted probabilities of both items being remembered (green), only one item being remembered (orange line), and neither item being remembered (grey line). The yellow line is the chance-level P(one) when encoding would occur in parallel, computed using equation 4 and the estimated P(both). Solid lines indicate averaged participant-level fits, and dashed lines indicate group-level fits (all participant data taken together within the same model). Error bars indicate standard error of the mean.*

4.4.2.3. Summary

These results indicate that participants' recall improved with increasing exposure duration. The general lack of a difference between the empirical and chance P_{one} (reflected by a lack of a main effect in the repeated-measures ANOVA) is in line with the parallel encoding hypothesis. These results replicated those obtained in experiment 1 where coloured disks were used as stimuli.

At 0 and 10 ms exposure durations, the chance P_{one} exceeded the empirical P_{one}.

In general, this alludes to 'super-parallel' processing, where the processing of one stimulus helps the processing of another stimulus, thereby allowing both to be encoded quicker than one would expect when stimuli were processed in parallel but completely independently. Given that this difference occurs so early, and even in the 0 ms condition (during which no stimulus was presented), indicates that they are probably best accounted for by chance rather than having a biological basis.

At the 200 ms exposure duration the empirical P_{one} exceeded chance. This result might be in line with the serial encoding hypothesis, but should be interpreted with care, given the lack of a main effect of P_{one} type in the repeated-measures ANOVA. In addition, the difference occurs quite late: It suggests that processing of the second stimulus only started after 200 ms. Furthermore, there is no such effect at longer durations (360 or 580 ms).

In experiments 1 and 2, stimuli were presented bi-laterally. It might be possible to argue that each hemisphere could therefore process its own stimulus independently. Hence, within each hemisphere there could still occur serial encoding. To account for this, I conducted experiment 3, which presented two coloured stimuli on each trial (like the double-item condition from experiment 1), either in one or in both hemifields.

4.4.3. Experiment 3 – Whole-report of uni- and bilateral coloured disks

4.4.3.1. Absolute angular error

To assess the effects of presentation (uni-lateral or bilateral), stimulus type (first or second response), and exposure duration on the average absolute angular error (**Figures 4.11 and 4.12**), a repeated-measures ANOVA was conducted. This revealed a main effect of exposure duration, $F_{Greenhouse\text{-}Geisser}(2.59, 54.32) = 332.88$, $p < 0.001$, $\eta_{partial}^2 = 0.94$, a main effect of stimulus, $F(1, 21) = 72.06$, $p < 0.000$, $\eta_{partial}^2 = 0.77$, and a main effect of presentation, $F(1, 21) = 7.87$, $p = 0.011$, $\eta_{partial}^2 = 0.27$. There were no statistically significant interaction effects.

In both the uni-lateral and the bi-lateral presentation conditions, there were no statistically significant differences between the first and the second response for exposure durations 0 – 40 ms (all *p* were over a Holm-Bonferroni corrected alpha level of 0.05), but a difference became apparent from 70 – 580 ms (all $p < 0.001$).

For the first response, there were no statistically significant differences between the uni-lateral and bi-lateral presentation (all *p* were over a Holm-Bonferroni corrected alpha level of 0.05). For the second response, there were no statistically significant differences between the uni-lateral and bi-lateral presentation (all *p* were over a Holm-Bonferroni corrected alpha level of 0.05), save from at exposure durations of 40 ms ($p = 0.002$), where the average absolute angular error was lower for the bi-lateral presentation condition.

The ANOVA results suggested that the presentation of bilateral stimuli resulted in subtly better recall than the presentation of unilateral stimuli. However, the lack of

such differences in post-hoc t-tests indicates that this is a very subtle effect.

As in experiments 1 and 2, recall error was lower for higher exposure durations, and recall error was lower for the first than for the second response.

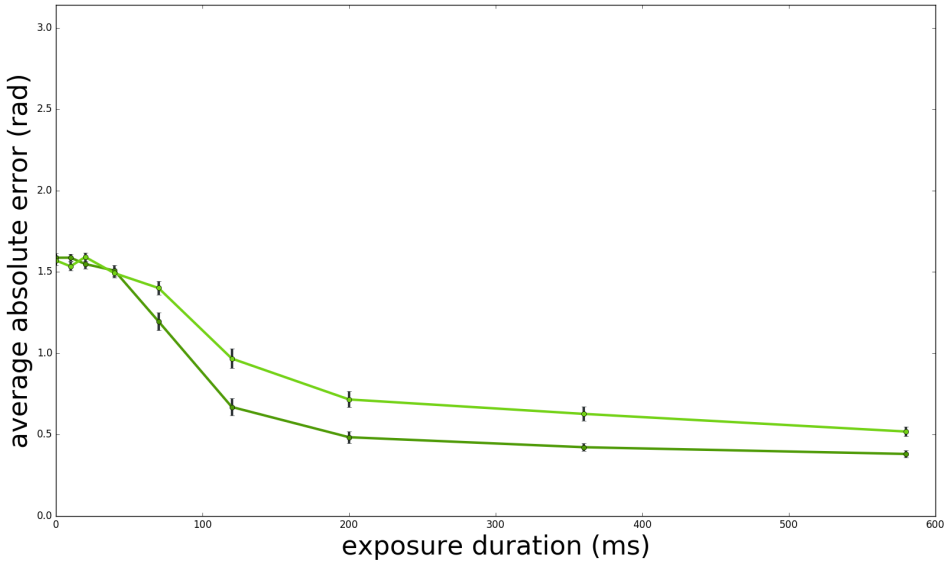

Figure 4.11 *– Presentation in two hemifields. Average absolute error (circular distance between the target and the response; y-axis) as a function of exposure duration (x-axis). Average error in is plotted in green for the first response, and light green for the second response. Error bars indicate standard error of the mean.*

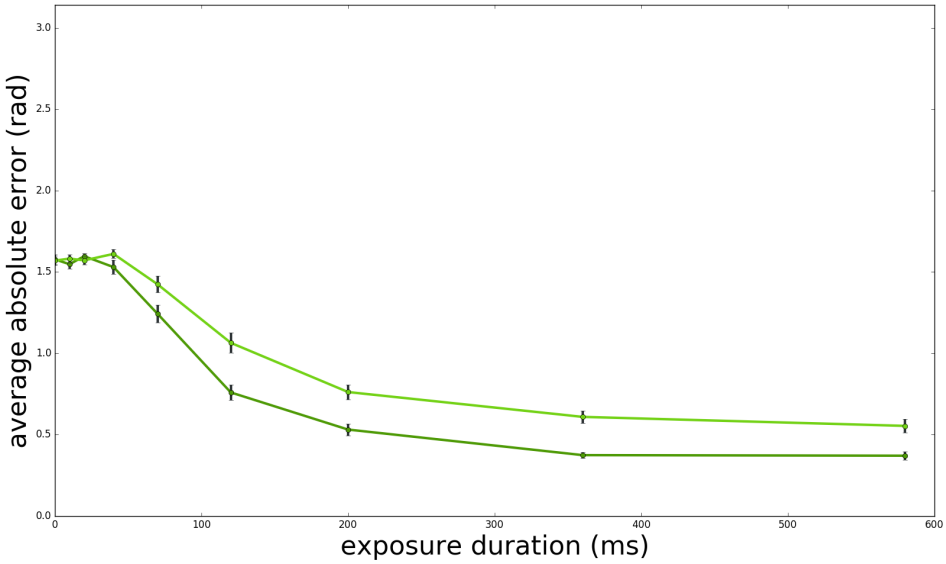

Figure 4.12 *– Presentation in one hemifield. Average absolute error (circular distance between the target and the response; y-axis) as a function of exposure duration (x-axis). Average error is plotted in green for the first response, and light green for the second response. Error bars indicate standard error of the mean.*

4.4.3.2. Encoding model parameters

The difference between the empirically obtained and chance P_{one} (**Figures 4.13 and 4.14**) was investigated using a repeated-measures ANOVA. This revealed a statistically significant main effect of exposure duration, $F_{Greenhouse\text{-}Geisser}(3.26, 68.43) = 13.25$, $p < 0.001$, $\eta_{partial}^2 = 0.39$, a main effect of P_{one} type, $F(1, 21) = 139.20$, $p < 0.001$, $\eta_{partial}^2 = 0.87$, and an interaction effect between P_{one} type and exposure duration, $F(8, 168) = 6.25$, $p < 0.001$, $\eta_{partial}^2 = 0.23$. There was no main effect of presentation (uni- or bi-lateral), $F(1, 21) = 1.39$, $p = 0.252$, and no statistically significant interaction effects.

Planned related-samples t-tests revealed statistically significant differences

between the empirical and chance P_{one} at any exposure duration at exposure durations 0, 10, 20, 70, and 200 ms in the uni-lateral presentation condition, and at exposure durations 0, 10, 20, 40, 120, and 200 ms. In all of these cases, the chance P_{one} exceeded the empirical P_{one}. For an overview, see **Table 4.4**.

Table 4.4. Results from related-samples t-tests between the empirical and chance P_{one} at each exposure duration. The alpha level is Holm-Bonferroni corrected. Statistically significant differences are marked with an asterisk.

Presentation	Exposure duration (ms)	*t* value	*p* value	alpha	rank
uni-lateral	0	-5.88	< 0.001*	0.00278	1
uni-lateral	10	-3.74	0.001*	0.00500	9
uni-lateral	20	-5.28	< 0.001*	0.00294	2
uni-lateral	40	-2.41	0.025	0.0167	16
uni-lateral	70	-4.05	< 0.001*	0.00455	8
uni-lateral	120	-2.47	0.022	0.0125	15
uni-lateral	200	-3.41	0.003*	0.00556	10
uni-lateral	360	-0.82	0.423	0.0500	18
uni-lateral	580	-2.66	0.015	0.00833	13
bi-lateral	0	-4.20	< 0.001*	0.00385	6
bi-lateral	10	-4.57	< 0.001*	0.00357	5
bi-lateral	20	-3.38	0.003*	0.00625	11
bi-lateral	40	-4.67	< 0.001*	0.00333	4
bi-lateral	70	-2.16	0.042	0.0250	17
bi-lateral	120	-4.89	< 0.001*	0.00313	3
bi-lateral	200	-4.15	< 0.001*	0.00417	7
bi-lateral	360	-2.70	0.013	0.00714	12
bi-lateral	580	-2.61	0.016	0.0100	14

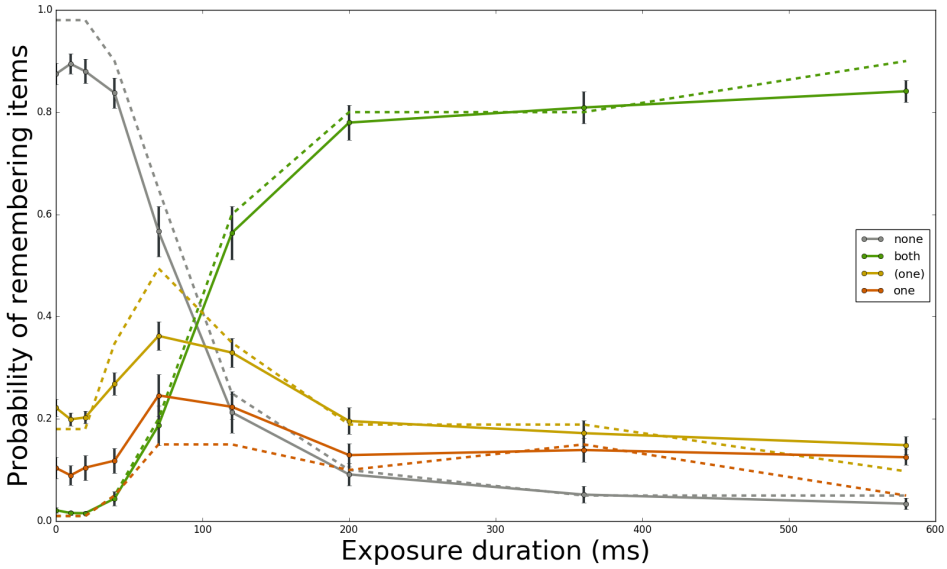

Figure 4.13 *– Fitted probabilities of both items being remembered (green), only one item being remembered (orange line), and neither item being remembered (grey line). The yellow line is the chance-level P(one) when encoding would occur in parallel, computed using equation 4 and the estimated P(both). Solid lines indicate averaged participant-level fits, and dashed lines indicate group-level fits (all participant data taken together within the same model). Error bars indicate standard error of the mean.*

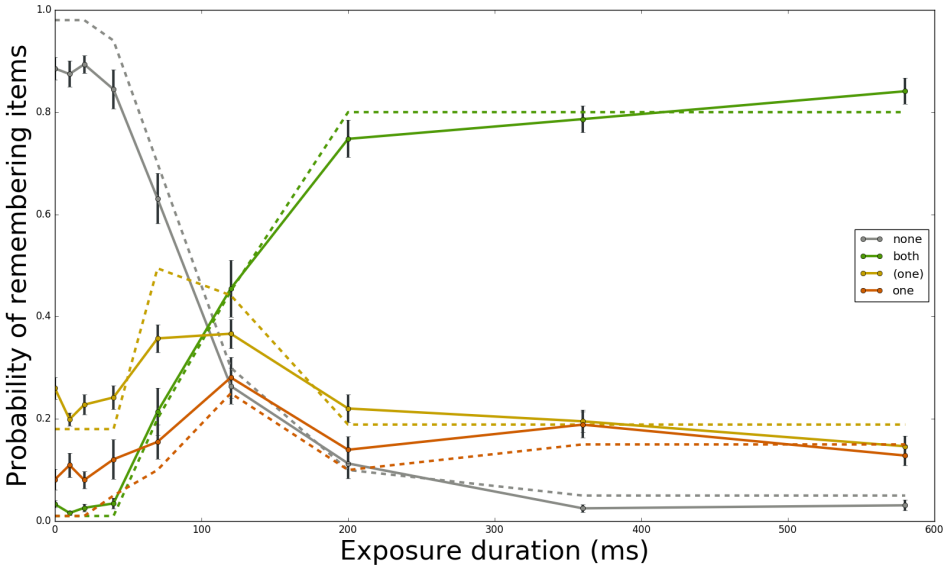

Figure 4.14 *– Fitted probabilities of both items being remembered (green), only one item being remembered (orange line), and neither item being remembered (grey line). The yellow line is the chance-level P(one) when encoding would occur in parallel, computed using equation 4 and the estimated P(both). Solid lines indicate averaged participant-level fits, and dashed lines indicate group-level fits (all participant data taken together within the same model). Error bars indicate standard error of the mean.*

4.4.3.3. Summary

These results indicate that participants' recall improved with increasing exposure duration. In addition, there seemed to be some evidence that processing one stimulus in each hemisphere improved performance in comparison to processing two stimuli in the same hemisphere (main effect of presentation in the repeated-measures ANOVA of average absolute angular distance), but this was a subtle effect (evidenced by the lack of statistically significant post-hoc tests).

Crucially, that the empirical P_{one} never exceeded chance is in line with the parallel encoding hypothesis. If anything, that the chance P_{one} quite consistently exceeded the empirical P_{one} can point to 'super-parallel' processing, where the processing of one stimulus improves the concurrent processing of another.

4.4.4. Experiment 4 – Whole-report of oriented gratings in EEG recording

4.4.4.1. Behaviour

Responses in the single-stimulus condition were the most accurate (**Figure 4.15**, left panel). In the double-stimulus condition, the first response (**Figure 4.15**, central panel) was more accurate than the second response (**Figure 4.15**, right panel). These results replicate those in the previous experiments.

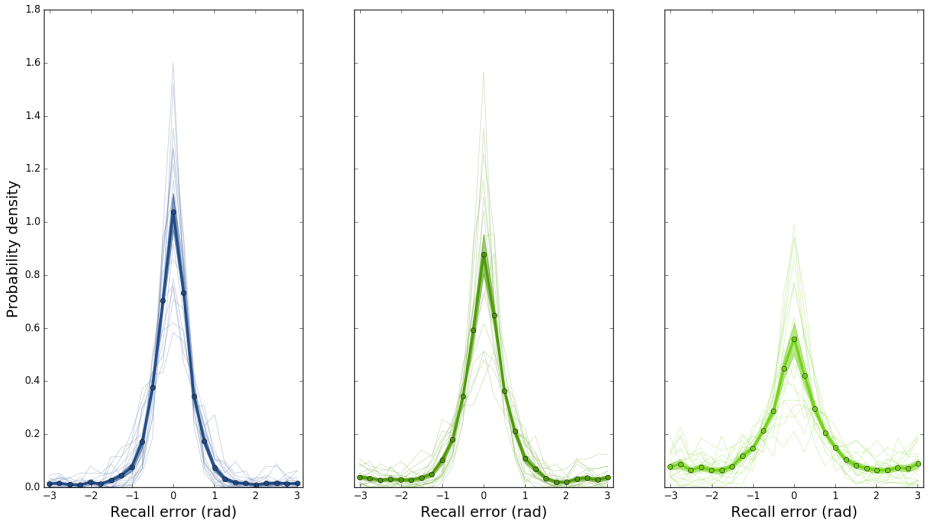

Figure 4.15 *– Histograms of the error distributions in experiment 4. Error in the single-item condition is presented in the left panel (dark blue). Errors in the two-item condition is presented in the central panel (green) for the first response, and in the right panel (light green) for the second response. Thick lines represent the group averages, with the standard error of the mean represented by the shaded area around the thick line. Individual participants are drawn as thin and transparent lines.*

4.4.4.2. Event-related potentials

In the single-item condition (**Figure 4.16A**), the parietal electrode contra-lateral to the stimulus shows a positive deflection around 100 ms after stimulus onset, and a pattern of further positive deflections after mask onset. The ipsi-lateral electrode shows a positive deflection around 200 ms, and seems to generally be anti-correlated with the contra-lateral electrode. The topography of these deflections is depicted in **Figure 4.16C and 4.16D**, and indeed shows contra-lateral positive deflection followed by ipsi-lateral positive deflection over parietal cortex.

In the double-stimulus condition (**Figure 4.16B**), the same positive deflection occurs around 100 ms, contra-lateral to both stimuli. A similar deflection occurs 100-150 ms after mask onset. Crucially, the ERP contra-lateral to the stimulus that will be the first response does not differ from the ERP contra-lateral to the second response (neither in onset nor amplitude). The topography (**Figure 4.16E**) of the effect indeed shows bilateral positive deflection over parietal cortex 100 ms after stimulus onset.

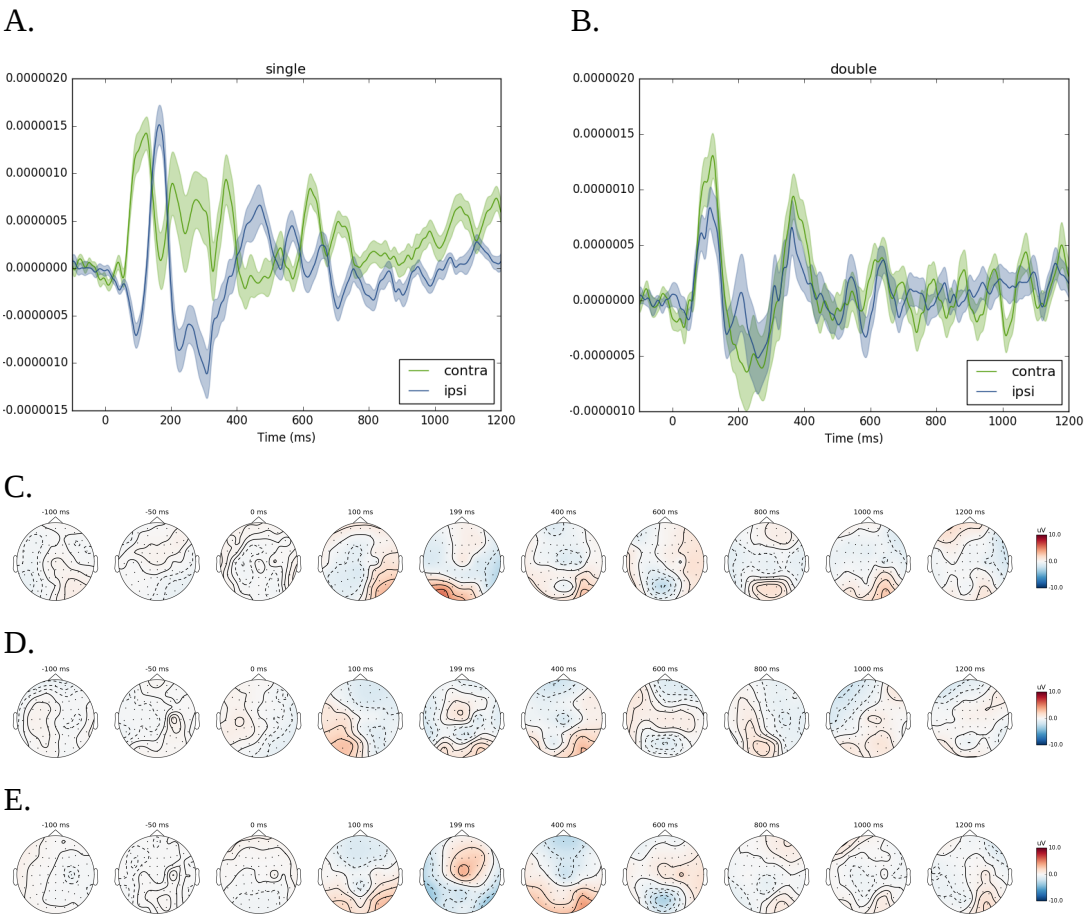

Figure 4.16 *– **A)** Event-related potential (Volts) in the single-stimulus condition from electrode P7/P8 contra-lateral (green) or ipsi-lateral (dark blue) to the presented item. Lines indicate group averages, and shaded areas indicate the standard error of the mean. **B)** Event-related potential (Volts) in the double-stimulus condition from electrode P7/P8 contra-lateral (green) or ipsi-lateral (dark blue) to the first respons to (note that the ipsi-lateral item is thus contra-lateral to the second response). Lines indicate group averages, and shaded areas indicate the standard error of the mean. **C)** Topography of the potential in the single-stimulus condition when it was shown on the left side of the display. **D)** Topography of the potential in the single-stimulus condition when it was shown on the right side of the display. **E)** Topography of the potential in the double-stimulus condition, where one item was presented on the left and the other on the right side of the screen. **A**-**E)** Time was set to 0 at stimulus onset. Presented timepoints are -100, -50, 0, 100, 199, 400, 600, 800, 1000, 1200 milliseconds. Microvolt scale ranges from -10 (blue) to 10 (red).*

4.4.4.3. Time-frequency

In the single-stimulus condition, a statistically robust decrease in power in the contra-lateral parietal electrode compared to the ipsi-lateral electrode is observed in the alpha frequency band from 400 – 600 ms after stimulus onset (**Figure 4.17A**).

In the double-stimulus condition, a brief decrease in power in the electrode contra-lateral to the stimulus that will be responded to first compared to the second response stimulus is visible 500-600 ms after stimulus onset (Figure **4.17B**).

A.

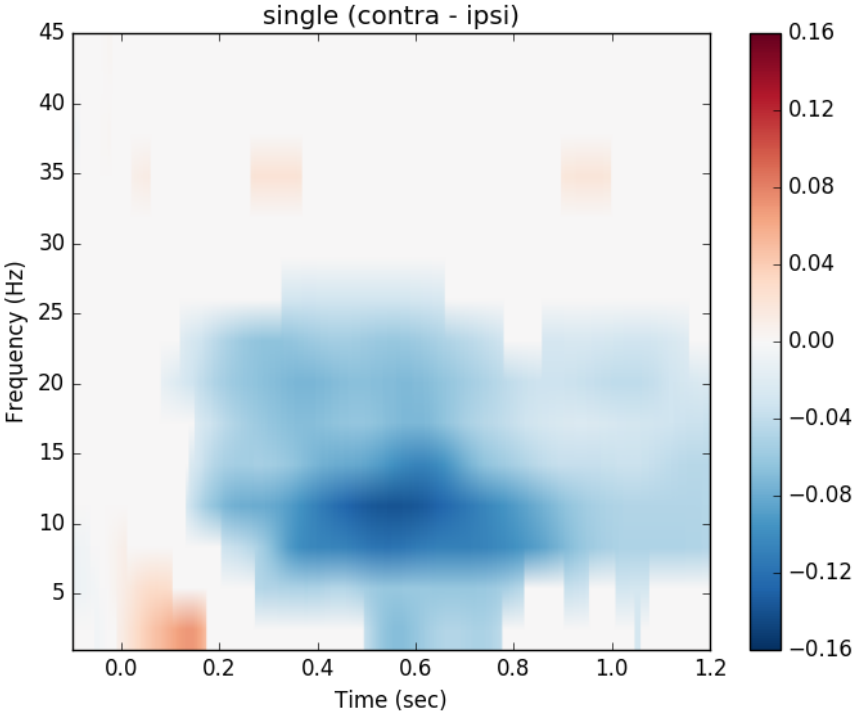
single (contra - ipsi)
Frequency (Hz)
Time (sec)
45
40
35
30
25
20
15
10
5
0.0
0.2
0.4
0.6
0.8
1.0
1.2
0.16
0.12
0.08
0.04
0.00
−0.04
−0.08
−0.12
−0.16

B.

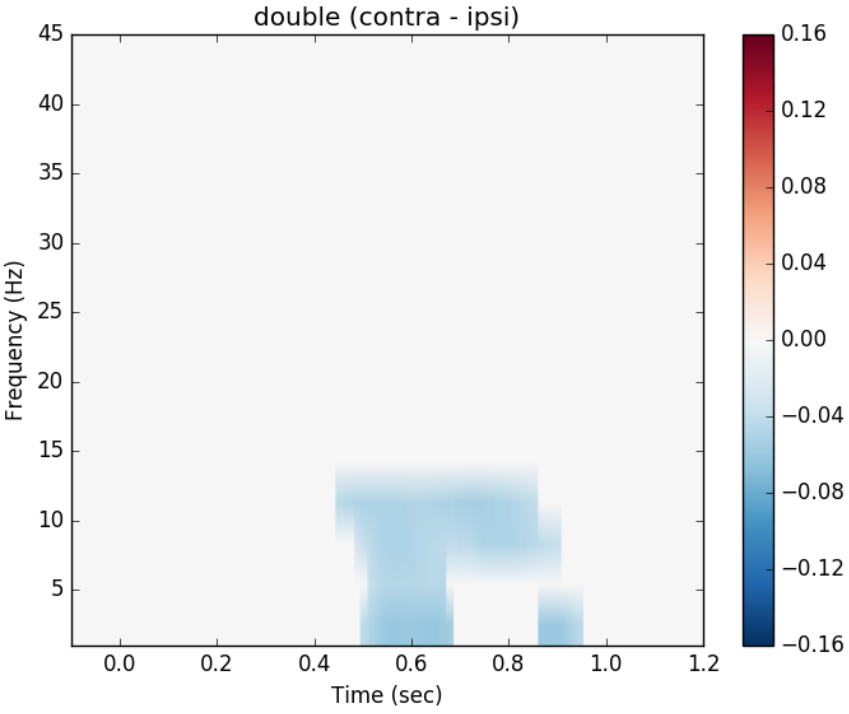
double (contra - ipsi)
Frequency (Hz)
Time (sec)
45
40
35
30
25
20
15
10
5
0.0
0.2
0.4
0.6
0.8
1.0
1.2
0.16
0.12
0.08
0.04
0.00
−0.04
−0.08
−0.12
−0.16

***Figure 4.17** – **A)** Oscillatory power difference between contra- and ipsi-lateral (with regards to stimulus presentation) electrode P7/P8 in the single-item condition. **B)** Oscillatory power difference between contra- and ipsi-lateral (with regards to the first response) electrode P7/P8 in the double-item condition.*

4.5. Discussion

In three experiments, participants were presented with one or two stimuli with a unique colour in CIE L*a*b* space or orientation, uni- or bi-laterally, and with a variable exposure duration (0 – 580 ms) before the onset of a mask. A reduction in all conditions of the average absolute angular error with exposure duration indicates that longer encoding time improved stimulus encoding, as shown previously using different stimuli (Bays et al., 2011; Bundesen, 1990; Vogel et al., 2006). More importantly, computational modelling of the error distributions demonstrated that the probability of participants remembering both stimuli increased with exposure duration, accompanied by a decrease in the probability of participants to remember none of the stimuli.

Crucially, the probability of participants remembering only one stimulus did not exceed what one would expect when two channels are independently processing one stimulus each. This finding runs counter to the prediction of the serial encoding hypothesis, which states that stimuli are encoded in sequence, and thus the probability of only one stimulus should peak at lower exposure durations. These results suggest instead that stimuli were encoded in parallel, even when presented in the same

hemifield.

In a fourth experiment (another whole-report task), participants were presented with one or two (concurrent) stimuli for 200 milliseconds, and then with one or two masks for 1100 milliseconds. Hereafter, they were asked to reproduce all stimuli with the method of adjustment, in an order of their own choosing. Parietal ERPs were not statistically different between the hemisphere contralateral to the stimulus of the first response and that of the second. The time-frequency spectrum did also not differ between the two stimuli, save for a brief window in the alpha interval (8-12 Hz) during the maintenance period (about 500 ms after mask onset). This reflects a stronger alpha desynchronisation over the hemisphere contralateral to the first response, compared to the second response. These findings are in line with existing literature on alpha power, and attention and visual short-term memory (Jensen et al., 2002; R.M. Mok, Myers, Wallis, & Nobre, 2016; Myers et al., 2014, 2015; Sauseng et al., 2005; van Ede et al., 2014). More importantly, these findings are consistent with a framework in which stimuli can be encoded into visual short-term memory in parallel.

4.5.1. Behavioural results

That recall precision improves gradually as a function of exposure duration is not a surprising finding, nor is it new. It has been demonstrated and modelled in seminal work (Bundesen, 1990), as well as in more contemporary research (Bays et al., 2011; Vogel et al., 2006).

On the other hand, direct empirical evidence of *parallel* encoding of information

into visual short-term memory is highly novel, and an important addition to contemporary literature. It has long been assumed that stimuli are encoded in a serial fashion, and this has been a foundation of many models of visual search (Cave & Wolfe, 1990; Duncan, 1984; Wheeler & Treisman, 2002; Wolfe, 1994) and short-term memory functioning (Manohar et al., 2017).

Contrary to the serial hypothesis, it has been proposed by some researchers that information might be encoded into short-term memory in parallel (Bundesen, 1990; Wilken & Ma, 2004). This view is supported by the finding that attention can be directed at multiple non-contiguous locations (Hahn & Kramer, 1998; Kramer & Hahn, 1995; McMains & Somers, 2004; Müller et al., 2003). Recent theoretical advances in short-term memory research describe the system as a limited resource that can be flexibly allocated (Bays et al., 2009; Bays & Husain, 2008; Ma et al., 2014; Wilken & Ma, 2004), and computational work that supports this view is compatible with the simultaneous imprinting of information from different items in the same neural population (Matthey et al., 2015). In addition, there is emerging evidence to support the biological plausibility of this paradigm (Curtis et al., 2016).

The findings presented in this chapter demonstrate that encoding is most likely to occur in parallel, under circumstances of the tasks deployed here.

4.5.2. EEG results

In a serial encoding framework, one could predict a difference in the onset of ERPs or changes in oscillatory power between the electrodes over hemispheres

contralateral to the first-encoded item, and those contralateral to the second. In a parallel framework, both stimuli would be encoded at the same time, and thus the ERP and oscillatory responses contra-lateral to both stimuli are expected to have the same onsets. It could be argued that the amplitudes of the ERP and oscillatory responses to both stimuli do not necessarily have to be the same, because stimuli can be encoded at different rates (as evidenced in experiments 1 – 3).

A strong assumption in experiment 4 is that the first response is to the first-encoded stimulus. If this does not hold, the averaging of ERPs and time-frequency spectra to the stimuli contra-lateral to the first and second responses would not overlap with encoding order, and would thus dramatically undermine the validity of the presented comparisons. One argument against this, is that the onsets of particularly the ERPs show very little variance, which suggests that regardless of how the data is ordered, no difference in onsets would exist. That said, the data presented for experiments 1 – 3 provide stronger evidence than those for experiment 4, thus the neural responses presented here should merely be considered to not be against parallel encoding.

In a future experiment, it would be very interesting to employ multivariate decoding methods. In a design like the current, these would offer a trial-by-trial estimates of which stimulus was encoded first (when decodability onsets were in fact different, which would only be the case during serial encoding). This would have addressed the potential shortcoming of the current study outlined above (i.e. the assumption that first-encoded equals first-responded to), as it would reflect the actual encoding order of both stimuli.

One could argue that in order to decode the orientation of stimuli, they need to be in the focus of attention (Cowan, 1988; Oberauer, 2002), which would be in line with the current literature (LaRocque, Lewis-Peacock, Drysdale, Oberauer, & Postle, 2013). However, studies that are similar to the one presented here have demonstrated decodability during the presentation of multiple stimuli (Wolff et al., 2015, 2017). It is important to ensure that the stimuli are not too small to elicit a potent enough response in (primary) visual cortices for it to reliably affect the EEG signal (Stokes et al., 2015).

My ERP findings corroborate earlier work that demonstrates a positive deflection around 100 ms after stimulus onset (Eimer & Forster, 2003; Sauseng et al., 2005). In addition, it replicates studies that link alpha desynchronisation to the maintenance of information in visual short-term memory (Jensen et al., 2002; Myers et al., 2014, 2015). Interestingly, the apparent reflection of preparation to respond to one stimulus before the other in the difference between the alpha desynchronisation contra- and ipsi-lateral to the first-response is very much in line with a recent study that demonstrated that alpha power is attenuated in the hemisphere contra-lateral to a stimulus that is brought into the focus of attention in visual short-term memory (van Ede et al., 2017).

In sum, the EEG results presented in this chapter do not speak against the parallel encoding of information into visual short-term memory, and confirm earlier work that demonstrates that alpha power is attenuated over parietal cortex contra-lateral to stimuli that are attended to and encoded into visual short-term memory. However, a much stronger illustration of parallel encoding could come from multivariate statistics that could ‘decode’ items from EEG signal. The presence or absence of offsets in the

decodability of two items during encoding would be a more direct test of whether they are encoded in series or in parallel.

5. Serialisation of visual short-term memory encoding in the presence of reward

5.1. Abstract

Human behaviour and cognition are adaptive, and can be optimised to the environment. In a neutral context, encoding information from several items into visual short-term memory happens in parallel. However, it is unclear how encoding proceeds in a context where the retention of information from particular items is rewarded. Here, I present data from a whole-report experiment in which two items are briefly presented (0-360 ms), then masked for 500 ms, and then probed. Participants are to respond to both items with the method of adjustment, but are free to choose in which order they do so. One item was pre-cued, which indicated the presence of a monetary reward in one condition, and neutral feedback in the other. Reward and neutral trials were intermixed, and participants were signalled the condition at the start of each trial. Results show that in both conditions, participants are more likely to choose the pre-cued item, but slightly more so in the reward condition compared to the neutral feedback condition. Recall performance was better for the first response compared to the second response, with no difference between the conditions. Crucially, computational modelling demonstrated that the probability of participants remembering one item peaked at low exposure durations, followed by a rise in the chance of both items being remembered at higher exposure durations. These results suggest that participants encoded items in a serial fashion, i.e. the pre-cued item first, and the other item second. This was true in both conditions, which I attribute to participants perceiving the neutral feedback on their performance as inherently rewarding. In sum, reward biased participants' visual short-term memory encoding to operate in a serial rather than a parallel fashion.

5.2. Introduction

The visual world is an information-rich environment, and human interests are continuously changing. For example, an oncoming car deserves a pedestrian's immediate attention, and a high-resolution representation in visual short-term memory, to allow for a precise eye movement towards the car, and to assess whether it presents a danger or not. However, when the car indicates that it will be turning before the pedestrian crossing, the pedestrian's attention is better employed to look for potential dangers elsewhere – the car's representation is now only filling up highly limited and precious space in short-term memory. Even for something as pedestrian as crossing a street, humans need flexible cognition, with a highly dynamic interplay between attentional selection and short-term memory storage.

In the previous chapter, I proposed that information from multiple items can be encoded into visual short-term memory at the same time, in parallel. However, items in all four experiments described in the previous chapter were of equal behavioural relevance, and thus probed the encoding mechanism in an unbiased state. In this chapter, I will explore how the encoding mechanism is affected when an item is made behaviourally relevant.

Previous work has shown that short-term memory resources can be rapidly re-allocated under changing goals (Bays & Husain, 2008). Specifically, Bays and Husain presented an array of several items, and monitored participant's eye movements during their visual exploration (and encoding) of the array. Crucially, during participant's last eye movement, from the penultimate to the final item, the display was blanked. The

result was that participant's never actually foveated the final item. Regardless, recall precision was dramatically better for the final item, which suggests that participants rapidly (during a single fixation) shifted their short-term memory resources from previously fixated items to the to-be-foveated item.

This mechanism is adaptive: A saccade towards the periphery requires a high-acuity representation of the target, to prevent over- or under-shooting. Other items, however, have already been fixated, and can therefore be retained with lower acuity. The result is a gist representation of the environment, with only behaviourally targets being retained at a high resolution. The next chapter will focus in-depth on the re-allocation of short-term memory resources during eye movements.

There are other methods, however, for inducing behavioural reward that do not necessarily require manipulating the availability of information across eye movements. For example, previous research (among which my own) has demonstrated that items that are associated with negative consequences can automatically attract attention (Mulckhuyse, Crombez, & Van der Stigchel, 2013; Mulckhuyse & Dalmaijer, 2016; Schmidt, Belopolsky, & Theeuwes, 2015, 2017). In these designs, a task-irrelevant feature is typically associated with an aversive stimulus (a shock or a sound) using Pavlovian conditioning: A participant repeatedly sees two types of stimuli (e.g. red and green), and the averse stimulus is applied after the presentation of one of the stimuli, but not after the other. After participants learn the association between the visual stimulus and aversive stimulus, and also the safe associated with the non-threatening visual stimulus, both visual stimuli are used as a task-irrelevant distractors in a task that measures attention. This could be the additional singleton paradigm, for example, in

which a search array of neutral stimuli and one target is presented, with a distractor stimulus appearing shortly after the onset of the search array (Theeuwes, Kramer, Hahn, & Irwin, 1998). It has been demonstrated that the additional distractor stimulus captures more eye movements when it was previously associated with an aversive stimulus compared to when it is not. In other words: Fear guides attention.

The precise mechanism of fear-guided attentional capture is debated: Some researchers argue for a 'direct route' from the retina, through the superior colliculus, via the pulvinar, to the amygdala. This theory of a 'direct route' for visual information via the amygdala to quickly filter out threatening information rests primarily on the notion that the superior colliculus is an important structure for establishing attentional priority maps (Munoz, 2002), and also reciprocally connects to the amygdala via the pulvinar (LeDoux, 2000, 2003; Vuilleumier, 2005). Empirical evidence exists of threatening information being processed more quickly than non-threatening information (Bannerman, Milders, de Gelder, & Sahraie, 2009; Bannerman, Milders, & Sahraie, 2009; LoBue, Matthews, Harvey, & Stark, 2014; Reynolds, Eastwood, Partanen, Frischen, & Smilek, 2009), but in these experiments the threatening information was always task-relevant up to some degree. When distractors were task-irrelevant, no difference exists between the time-course of attentional capture by threatening and non-threatening stimuli, suggesting that while threatening stimuli were more potent in capturing attention, they do not do so via a quicker mechanism (Mulckhuyse & Dalmaijer, 2016). Regardless of the exact route, most researchers agree that fear guides attention in a bottom-up fashion (i.e. stimulus-driven, and without too much cognitive control).

Another potent ‘grasper’ of attention is reward. In animal experiments, this is often food (e.g. fruit juice for macaque monkeys), but in humans the preferred experimental reward is money. Early studies of the effects of reward on attention studied task-relevant rewards overwhelmingly demonstrated that reward-associated stimuli capture attention, and that the absence of an expected reward can suppress attentional inhibition of the associated stimulus (Della Libera & Chelazzi, 2006; Hickey, Chelazzi, & Theeuwes, 2010a, 2010b; Peck, Jangraw, Suzuki, Efem, & Gottlieb, 2009; Serences, 2008). Perhaps more interestingly is that when introduced as a stimulus quality, rewards still capture attention. In a task where participants are required to identify the cardinally oriented line (horizontal or vertical) among diagonal lines, coloured shapes (circles or rectangles) around the line stimuli can be used as task-irrelevant stimulus information (Theeuwes, 1992).

When a particular colour is consistently paired with a high (80%) or low (20%) reward probability, those stimuli will capture attention, as evidence by higher response times and lower accuracy rates to stimuli in trials where any reward-associated colour was present compared to when none was present (Anderson, Laurent, & Yantis, 2011). Interestingly, the effect scales with reward probability, and is correlated with working memory capacity: higher capacity is associated with lower reward-driven attentional capture (Anderson et al., 2011). For contemporary conceptual replications, see (Bucker & Theeuwes, 2017; Donohue et al., 2016; Failing & Theeuwes, 2017). In sum, the attentional effects of reward mirror those of threat, but the underlying mechanism is less hotly debated.

Important for the current chapter is whether the attentional bias towards reward-

associated stimuli is translated into better (or different) short-term memory encoding. Fortunately, a recent study has indeed demonstrated that task-unrelated rewards can improve short-term memory, and that this effect is specific to encoding (Klink, Jeurissen, Theeuwes, Denys, & Roelfsema, 2017). In some of their experiments, Klink et al. present three stimuli (gratings with an orientation, much like the stimuli in the previous and this chapter) for 300 or 3000 ms, each positioned within a coloured circle. The stimuli were followed by a screen with placeholders that was visible for 2000 ms (the maintenance phase), and then a screen with placeholders and a single probe that had to be adjusted to the memorised orientation. The colour of the circles around stimuli or placeholders indicated how much reward would be available upon correct recollection of the stimulus orientation: None, a little (5 points), or a lot (50 points). Crucially, these colours could be presented either at the same time as the stimuli (during encoding), during the maintenance phase, or during the probe (during recall). When presented during the maintenance phase, the colouring had no effect, which suggests participants did not discard non-rewarded items, nor did they re-allocate short-term memory resources from non-rewarded stimuli to reward-associated stimuli during the maintenance phase. When presented during recall, the reward information also did not influence recall accuracy.

The reward information only had an effect when presented together with the memory array: Recall accuracy scaled with reward. Specifically, there was poorer recollection for no reward, average recollection (as in all other described conditions where reward information had no effect) for low reward, and better recollection for high rewards. This suggests that reward information is factored in only during attentional selection and the consecutive encoding of information into short-term memory.

A note of caution: The reward-effect described in the previous paragraph only appeared when the encoding time was 3000 ms, but not when it was 300 ms (Klink et al., 2017). This is curious, as my results in the previous chapter suggest that the precision of the encoded information plateaus around 350-400 ms. My experiments differ from those of Klink et al., in the sense that I used one or two stimuli, where Klink et al. used three. Perhaps participants felt overwhelmed by three stimuli at short exposure durations. An alternative, and perhaps more likely explanation is that the coloured circles around the stimuli had to be parsed before the stimuli themselves could be allocated attention and encoding resources. In this case, 300 ms might simply not have been enough to fist parse the colours, and then encode the orientations, forcing participants to encode all stimuli equally instead. Perhaps Klink et al. would have been better off using a pre-cue instead, so that participants could pre-allocate attention and/or short-term memory resources to the locations where reward-associated stimuli would appear.

Although the discussed study by Klink et al. is highly elegant in demonstrating a bias towards reward-associated stimuli during short-term memory encoding, it leaves open how exactly this bias operated on the mechanism of encoding. To remind the reader: In the previous chapter, I suggested that during short-term memory encoding, memory resources are divided up into as many channels as there are attentionally selected items. These channels then each encode their assigned item into short-term memory, and they operate in parallel. In this framework, reward can operate in two distinct ways. The first option is that reward biases the division of resources over items. For example, if two stimuli are presented, and only one is reward-associated, then a higher proportion of the available resources might be allocated to the reward-associated

item. An alternative option is that only reward-associated stimuli are assigned channels, and that other stimuli are only encoded after all reward-associated stimuli are. In our example, this would mean one channel is assigned to the reward-associated stimulus, taking up all encoding resources. Only after encoding of the reward-associated stimulus will the channel open up and be assigned to the other stimulus.

In this chapter, I aim to distinguish between these two options by employing a similar experiment to those presented in the previous chapter. Two stimuli will be presented on each trial, with a variable exposure duration. This will allow me to employ the same computational modelling that is employed in the previous chapter, which can distinguish between serial and parallel encoding of items into short-term memory. Crucially, a pre-cue will be presented on each trial, indicating which stimulus will be reward-associated and which will not be. This resolves the aforementioned potential issue with Klink et al.'s study, during which reward information and the stimuli had to be parsed during the presentation of the stimuli, making it unclear what information was being processed first. The reward participants receive will be directly related to their recall error, to emphasise the importance of remembering the rewarded stimulus.

In addition to monetary reward, I will also employ a condition during which one stimulus will be pre-cued, but will not be rewarded. Instead, participants will simply receive feedback on how well they recalled the pre-cued stimulus. This is to distinguish between the effects of monetary reward and feedback, as the monetary reward condition will also indirectly provide feedback on participants' accuracy (when they receive a high reward, they know their performance was good).

5.3. Methods

5.3.1. Participants

Participants were recruited via the University of Oxford's experimental psychology recruiting website for healthy volunteers, with permission from the local ethics committee, and in accordance with the declaration of Helsinki. 15 participants took part in the experiment presented here. They had a mean age of 24.2 years (SD = 3.9, range = 18-31), 7 were male, and 8 were female.

5.3.2. Procedure

Participants visited a building associated with the Department of Experimental Psychology on a single occasion, for two hours. They were briefed, and provided written informed consent before starting the task.

At task onset, an EyeLink 1000 was calibrated using a 9-point calibration (or a 3-point fallback calibration if the participant’s eyes could not be calibrated in the vertical direction). The EyeLink 1000 was of the tower mount variety, and recorded the right eye at 1000 Hz.

The task was presented on a 21 inch (53.3 cm) ViewSonic P227f CRT monitor (screen dimensions: 40.5 x 30.5 cm), with a resolution of 1024 x 768 pixels, and a refresh rate of 100 Hz. The monitor was positioned at a distance of 62 cm from the participant's eyes.

The task consisted of 900 trials. A drift check for the eye tracker was performed every 20 trials, during which participants pressed the Space key while looking at a

central fixation target. At measurement errors of over 2 degrees of visual angle, the eye tracker was recalibrated. Participants were allowed a break every 100 trials.

Participants were compensated for their time at a base rate of 7 pounds per hour. They received additional compensation between 0 and 2 pounds per hour, depending on their individual performance during the task. In sum, participants were compensated with 14 pounds, plus up to 4 pounds.

5.3.3. Task

Each trial started with a central white fixation cross (visible for 1000 milliseconds), which briefly (100 milliseconds) turned bold and coloured green in half of the experiment, and red in the other. A green cross indicated that the following trial would be performed for a monetary reward, and red indicated that no monetary reward would be provided.

After the fixation cross returned to white, a recording was played of a voice saying “right” or “left”. This pre-cued which item would be provided feedback for. In neutral trials, this feedback was simply how accurate a participant was on a range from 0 to 100. In reward trials, this feedback indicated how many credits the participant had won, also on a range from 0 to 100.

Following the pre-cue, two items appeared for a variable exposure duration, immediately followed by a masking stimulus. This could be 20, 40, 80, 160, or 320 ms, and a condition of 0 ms was also included in which no gratings were shown, but instead a single frame (10 ms) of the mask was shown.

The items were sinusoidal gratings with a hard edge. They had a spatial

frequency of 5 cycles per stimulus, a diameter of 2.7 degrees of visual angle, at an eccentricity of 7.5 degrees of visual angle. The difference in orientation between the two gratings was at least 10 degrees, but otherwise independent between the stimuli. Orientations were sampled from a uniform distribution ranging between 0 and 180 degrees. Masks were two identical 32x32 grids of greyscale static, and were regenerated for each trial.

500 ms after mask onset, the mouse cursor became visible, which probed participants to respond to both stimuli. A response could be made by clicking on either stimulus, at which point the mask was replaced by a grating that was oriented towards the mouse cursor. Participants could hold down the mouse button while moving the cursor to rotate the stimulus to the orientation in their memory. Each trial progressed only after both stimuli were responded to, and the Space button was pressed.

After responding, participants were presented with feedback on their recall of the pre-cued item. The feedback indicated the error, and was scaled to 100 points/credits at a 0 degree recall error (best possible performance), and 0 points/credits at a 90 degree recall error (worst possible performance). In neutral trials, feedback would be "Accuracy is X points", and in reward trials it would be "Won X credits". This procedure made the procedure in both conditions identical, with the crucial exception of credits won in reward trials being translated into a monetary reward.

5.3.4. Analysis

The absolute error was computed as the angular difference between each presented item's orientation and the associated response. This resulted in two error

distributions per condition: one for the first response, and one for the second. These distributions were modelled according to the methods presented in the previous chapter. The resulting parameters include the proportion of guessed responses (Bays et al., 2009); and the proportions of trials in which neither items, only one item (but not the other), or both items were remembered. In addition, it can be computed what the chance is of one item being remembered given the chance that both items are remembered, when assuming parallel encoding into visual short-term memory. A comparison between this computed parameter and the parameter estimated from the empirically obtained data can address whether encoding happened serially or in parallel. For further details including equations for each model, see the previous chapter's Methods section.

Models were run within each reward condition, and within each exposure duration. After parameters were estimated for each exposure duration, they were compared between conditions using repeated-measures ANOVAs run in JASP version 0.8.2.0 (JASP Team, 2016).

To assess performance, a repeated-measures ANOVA with factors reward condition (2 levels: reward and neutral), choice order (2 levels: first and second response), and exposure duration (6 levels: 0, 20, 40, 80, 160, 320 ms) was run on the absolute error, and also one on the estimated proportion of guessing.

Crucially, to assess visual short-term memory encoding, a repeated-measures ANOVA was run with factors reward condition (2 levels: reward and neutral), proportion of trials in which only one item is remembered (2 levels: empirical and estimated under the assumption of parallel encoding), and exposure duration (6 levels: 0, 20, 40, 80, 160, 320 ms).

Post-hoc related-samples t-tests were performed on the appropriate cells following a statistically significant main effect.

5.4. Results

5.4.1. Reward manipulation check

The pre-cue manipulation failed if behaviour was not affected, in which case participants would choose the pre-cued item half of the time. This is not the case, with the pre-cued item being chosen in about 90% of the trials on average, across all exposure durations.

The reward manipulation would have failed if the proportion of choices for the pre-cued stimulus was the same between the reward and neutral conditions. This was barely the case, with a repeated-measures ANOVA with factors reward condition (2 levels: monetary reward and neutral feedback) and exposure duration (6 levels: 0, 20, 40, 80, 160, and 320 ms) on the proportion of trials in which the pre-cued stimulus was not chosen resulted in a main effect of reward, $F(1, 14) = 4.65$, $p = 0.049$, $\eta^2 = 0.25$, and a main effect of exposure duration, $F(5, 70) = 2.54$, $p = 0.036$, $\eta^2 = 0.15$. Post-hoc related-samples t-tests demonstrated that participants were more likely to choose the pre-cued stimulus in the reward condition at exposure durations 0 ($t(14) = 2.42$, $p = 0.030$), 160 ($t(14) = 2.48$, $p = 0.027$), and 320 ms ($t(14) = 2.24$, $p = 0.042$), but not at other exposure durations (all $p > 0.120$).

In sum, these results demonstrate that the proportion of trials in which the pre-cued item was chosen was slightly higher in the reward condition than in the neutral

condition. However, the most important finding is that participants were highly likely to choose the pre-cued item in both conditions, suggesting that they were motivated to choose the item for which they would be given feedback on their recall error.

5.4.2. Recall performance under reward

A repeated-measures ANOVA on the average absolute angular recall error (**Figure 5.1**) with factors reward condition (2 levels: monetary reward and neutral feedback), choice order (2 levels: first and second response), and exposure duration (6 levels: 0, 20, 40, 80, 160, and 320 ms) revealed a main effect of exposure duration, $F(5, 70) = 74.03$, $p < 0.001$, $\eta^2 = 0.84$, and a main effect of choice order, $F(1, 14) = 107.50$, $p < 0.001$, $\eta^2 = 0.89$, but not of reward condition, $F(1, 14) = 2.28$, $p = 0.153$.

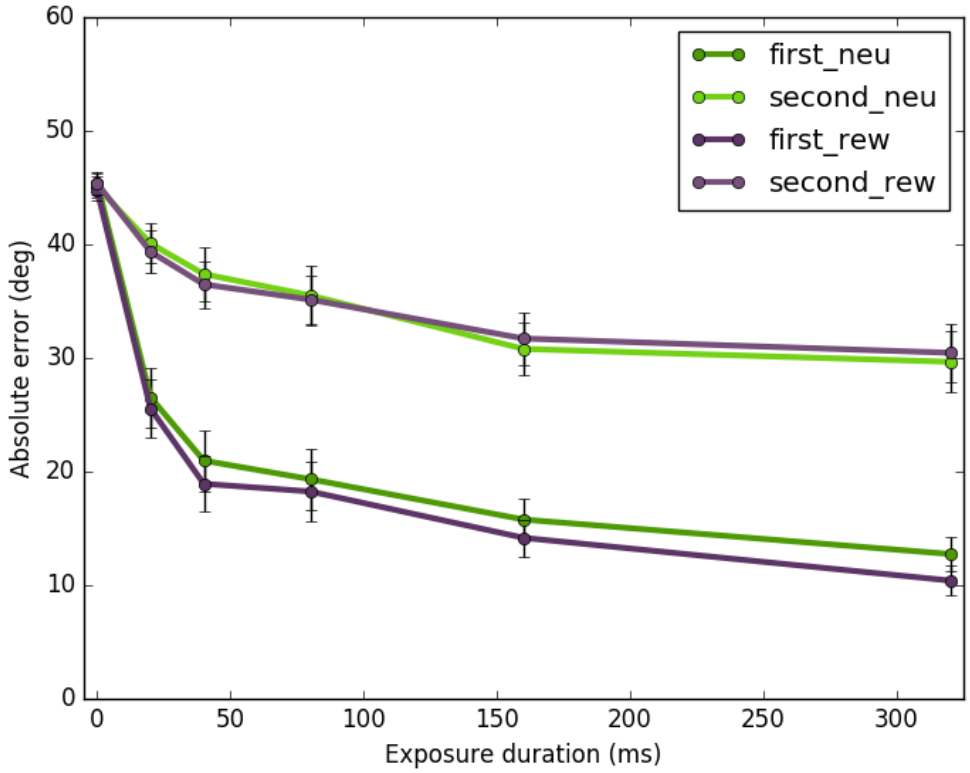

Figure 5.1 *– Average absolute error as a function of exposure duration. Green lines represent the neutral feedback condition, purple lines the monetary reward condition. Darker lines represent the first response, lighter lines the second. Lines indicate group averages, error bars indicate the standard error of the mean.*

A repeated-measures ANOVA on the average proportion of guessing (i.e. forgetting the stimulus orientation; **Figure 5.2**) with factors reward condition (2 levels: monetary reward and neutral feedback), choice order (2 levels: first and second response), and exposure duration (6 levels: 0, 20, 40, 80, 160, and 320 ms) revealed a main effect of exposure duration, $F(5, 70) = 69.89$, $p < 0.001$, $\eta^2 = 0.83$, and a main effect of choice order, $F(1, 14) = 102.37$, $p < 0.001$, $\eta^2 = 0.88$, but not of reward

condition, F(1, 14) = 3.48, p = 0.083.

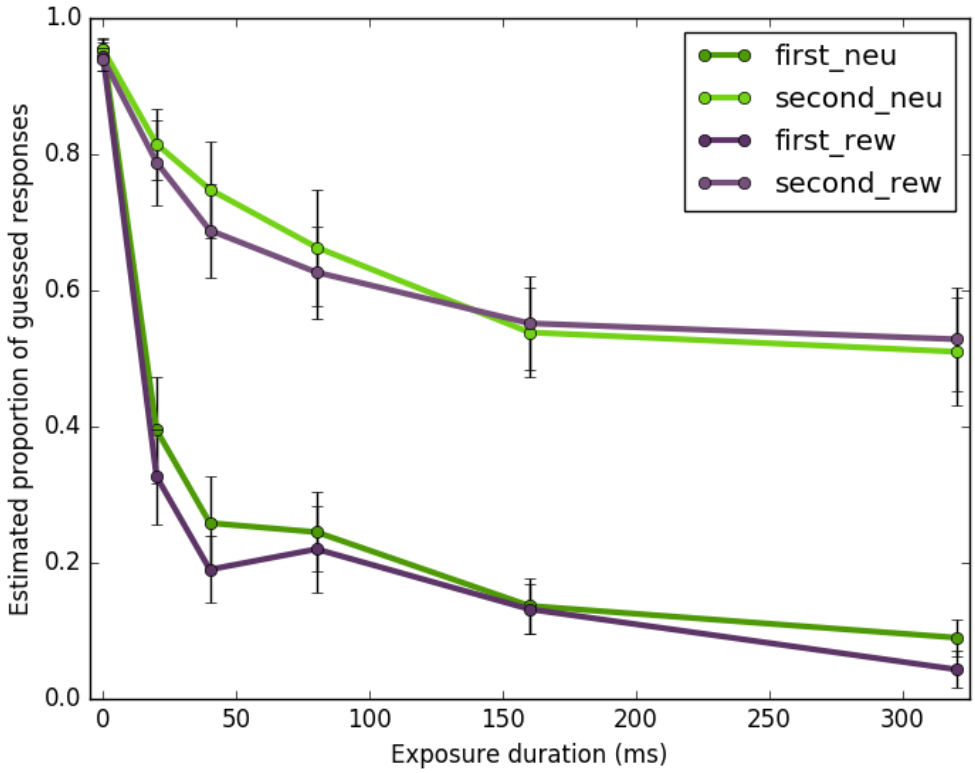

Figure 5.2 *– Estimated proportion of guessed responses as a function of exposure duration. Green lines represent the neutral feedback condition, purple lines the monetary reward condition. Darker lines represent the first response, lighter lines the second. Lines indicate group averages, error bars indicate the standard error of the mean.*

These results indicate that performance was better for longer exposure durations, and for the first response compared to the second response, but that performance did not differ between the monetary reward and neutral feedback conditions. Histograms of the

error distributions for each condition are plotted in **Figure 5.3**.

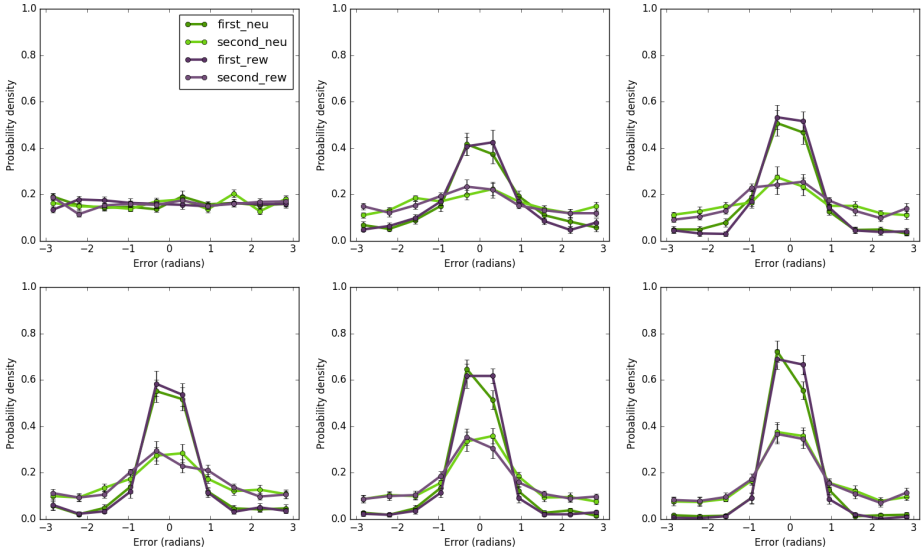

Figure 5.3 *– Histograms of error distributions. The top row is for exposure durations 0, 20, and 40 milliseconds, and the bottom row for exposure durations 80, 160, and 320 milliseconds. Green lines represent the neutral feedback condition, purple lines the monetary reward condition. Darker lines represent the first response, lighter lines the second. Lines indicate group averages, error bars indicate the standard error of the mean.*

5.4.3. Encoding strategy under reward

The proportion of trials in which only one item (but not the other) was remembered was estimated using a computational model described in the previous chapter (**Chapter 4, equation 1**). In addition, the chance level given a parallel short-

term memory encoding process for that parameter was also computed (**Chapter 4, equation 4**). If encoding happened in a serial process, one item would be encoded at a time. This means that the chance of only one item being remember should peak at short exposure durations. However, if encoding would occur in parallel, both items would be encoded at the same time, and such a peak in would not be observed; at least not above chance level, which is computed using the proportion of trials in which both items are remembered using an equation given in the Methods section of the previous chapter.

In sum, the crucial comparison is between the empirically obtained proportion of trials in which only one item is remembered, and the estimated proportion of trials in which only one item would be remembered under the assumption of parallel encoding (**Figure 5.4**). A repeated-measures ANOVA with factors parameter (2 levels: empirical and estimated), reward condition (2 levels: monetary reward and neutral feedback), and exposure duration (6 levels: 0, 20, 40, 80, 160, and 320 ms) revealed a main effect of reward, $F(1, 14) = 6.87$, $p = 0.020$, $\eta^2 = 0.329$, and a main effect of exposure duration, $F(5, 70) = 8.08$, $p < 0.001$, $\eta^2 = 0.365$, but not of parameter.

Given the fact that the difference between parameters is expected to only be evident at lower exposure durations, a main effect might be drowned out by the lack of a difference between the parameters at higher exposure durations. To accommodate for this, related-samples t-tests were performed within each reward-condition, between the empirically obtained and the estimated proportion of trials in which only one stimulus was remembered. The results are summarised in **Table 5.1**.

Table 5.1 – *The outcomes of related-samples t-tests between the empirically obtained and the estimated proportion of trials in which only one stimulus was remembered. Statistically significant t values are indicated in bold, and are accompanied by the associated uncorrected p value.*

Exposure duration	**Monetary reward**	**Neutral feedback**
0 ms	-1.14	-1.25
20 ms	0.86	-0.50
40 ms	**2.38** (p = 0.032)	**3.04** (p = 0.009)
80 ms	1.37	0.21
160 ms	1.21	-1.32
320 ms	1.30	0.15

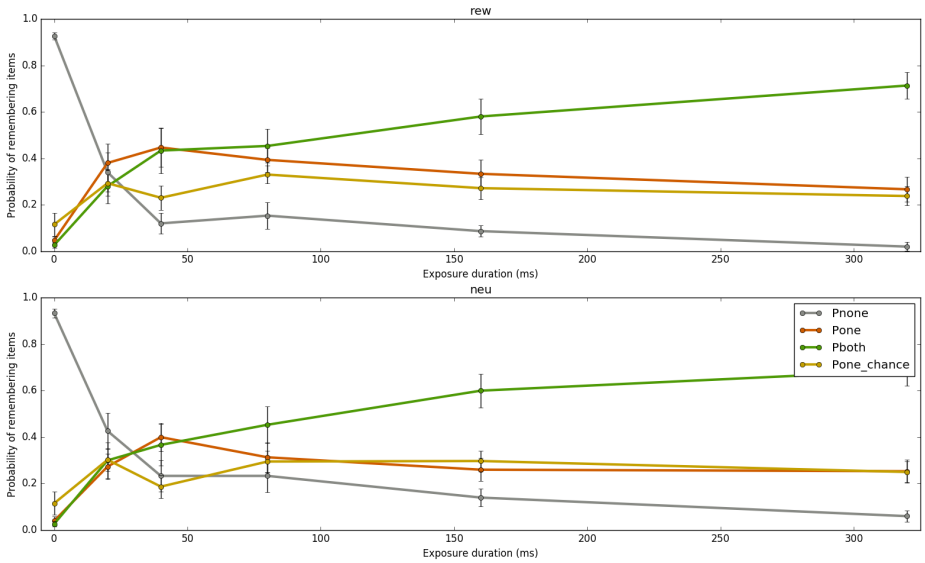

Figure 5.4 *– Estimated parameter estimates for the probability that both items are remembered P(both) (green), that only one item is remembered P(one) (orange), or that none of the items are remembered P(none) (grey). In addition, there is the chance-level P(one) which indicates where P(one) would be if encoding happened in parallel. The top panel represents the monetary reward condition, and the bottom panel the neutral feedback condition. Lines indicate group averages. Error bars indicate standard error of the mean. The differences observed between P(one) and chance-level P(one) are statistically significant only at 40 milliseconds (p=0.032 for the reward condition and p=0.009 for the neutral feedback condition; only the latter would survive multiple comparisons correction).*

These results indicate that the empirically obtained proportion of trials in which only one item was remembered was higher than predicted from a parallel encoding

framework. This suggests items were encoded into visual short-term memory in a serial (one-by-one) fashion.

5.5. Discussion

In this chapter I present the results of a whole-report task in which participants are briefly (0-360 ms) presented with two items, followed by a mask, and then a response phase in which participants are required to recall both items. Crucially, participants get feedback on only one of the items, and they know which one from a pre-cue before stimulus onset. In half of the trials, participants know that this feedback is also translated into a monetary reward. Participants chose to respond first to the pre-cued item (and then to the other) in about 90 percent of all trials, indicating that they were motivated by the feedback. Participants were marginally more likely to respond first to the pre-cued item in the monetary reward condition compared to the neutral feedback condition, which indicated that they were only slightly more motivated by monetary reward than by only feedback on their own performance. This preference did not result in recall performance differences: Recall error and guessing proportion were lower for the first response than for the second, and with increasing exposure duration, but there was no effect of reward.

Crucially, at exposure durations of 40 ms, participants were more likely to remember only one item (compared to what was predicted by a parallel visual short-term memory encoding framework), suggesting that they were encoding the pre-cued stimulus first, and the uncued stimulus second. An encoding duration of around 40

milliseconds is very much in line with previous work that suggests the encoding of a single coloured item occurs in roughly 50 milliseconds (Vogel et al., 2006).

5.5.1. Rewarding feedback

One failure of this experiment is to employ a truly neutral condition: In what was dubbed the 'neutral' condition, participants still had a strong tendency to respond to the pre-cued item first. In a truly neutral condition, participants would have opted to respond to the pre-cued item in 50 percent of all trials, rather than the 90 percent reported here. This shows that the feedback they received in the neutral condition was in fact not quite so neutral at all, but rather carried an inherent value to participants.

This presents a problem: The conditions here were not just 'neutral' and 'rewarding', but rather potentially 'inherently rewarding', and 'inherently and also monetarily rewarding'. Fortunately, this is not a stand-alone experiment: the previous chapter's experiments 1-3 present a comparable experimental design without any feedback manipulation, and can thus serve as a comparison of how information is encoded into short-term memory when all stimuli are equal.

5.5.2. Encoding information into visual short-term memory

The results in this chapter partly replicate those presented in the previous chapter and by others (Bays et al., 2011): More time to encode a stimulus means better recall of that stimulus.

What does differ between this chapter and the previous is the encoding sequence. In the previous chapter, I presented the results of four experiments, and

argued that information from two items can be encoded into visual short-term memory in parallel. The results presented in this chapter suggest that information can also be encoded in series, with one item being committed to short-term memory at a time. This follows directly from the manipulations employed: In the previous chapter all stimuli were of equal importance, whereas in this chapter it was made more appealing to participants to recall one stimulus in particular.

That participants employ serial encoding when one stimulus is of more importance to them, indicates that short-term memory resources can already be biased at the level of encoding. Previously, it has been argued that short-term memory resources can be flexibly distributed between stimuli (Ma et al., 2014). It has also been argued that short-term memory resources can be dynamically redistributed during behaviour, to bias recall accuracy towards stimuli of more behavioural relevance, for example to better memorise the location of stimuli that will be the target of a saccade (Bays & Husain, 2008). In addition, there exists a large literature on how attentional selection is biased towards behaviourally relevant information, for example because stimuli are associated with reward (Anderson et al., 2011; Bucker & Theeuwes, 2017; Failing & Theeuwes, 2017) or with threat (Mulckhuyse & Dalmaijer, 2016; Schmidt et al., 2017). In fact, it has recently been shown that (provided ample presentation time) reward-associated stimuli are prioritised during short-term memory encoding (Klink et al., 2017).

What was yet unclear, is *how* behavioural relevance such as reward information changes the encoding mechanism. The current experiment addresses this question, by demonstrating that behaviourally relevant information not only biases the allocated proportion of attention and short-term memory resources, but that it can actually

fundamentally change the cognitive operations during short-term memory encoding. More specifically, reward-associated information is encoded first, and non-rewarded information is processed second. This suggests that human short-term memory encoding does not always operate in parallel, but can also operate in a serial manner if that is better suited to the environment. Thus the process appears to be quite flexible in its deployment, depending upon behavioural goals.

6. Dynamic re-allocation of visual short-term memory resources to saccade targets?

6.1. Abstract

It has been proposed that to be able to accurately make a saccade towards a target, one needs a high-resolution representation of the target location in visual short-term memory. This means that short-term memory resources have to be re-allocated towards saccadic targets before each saccade. A strong prediction of this theory is that (at saccade onset or during a saccade) *recall performance* regarding target features is better for the current saccadic target than for previously fixated items. Empirical evidence for this is provided by Bays and Husain (2008, Science). This chapter aims to replicate and extend their findings. Three experiments were run that required participants to fixate 4 or 5 items, each defined by 2 or 3 features (location, orientation, and colour), and memorise all of them. After a brief maintenance period (3 seconds), one of the items was probed, and participants had to recall one of its features. Crucially, the display was blanked during the saccade towards the final item, so that participants never foveated it. Results indicated that all items were memorised with roughly equal precision, which directly contrasts with the findings of Bays and Husain. Whether items were recalled by the method of adjustment or by change detection did not affect the results. Finally, location was recalled with greater precision than colour or orientation, and no difference in recall error existed between colour and orientation. In sum, the current chapter presents a failed replication of Experiment 2 of Bays and Husain (2008). Potential reasons are discussed.

6.2. Introduction

Even when the world around us is stable, our interests and goals can change with every eye movement. For example, when searching for a particular paper on their messy desk, a graduate student needs to select all items that look like papers (attentional selection), foveate each of the selected items to read the title, and then commit the foveated location to short-term memory to avoid re-fixating it. During the search process, the student will need a high-resolution representation of a to-be-foveated paper, in order for the planned saccade to land in the intended location without over- or under-shooting. In terms of short-term memory, the described process requires that a gist representation is maintained of where all articles are. In addition, each previously fixated item requires a low-resolution representation of its location and identity (“not what I’m looking for”). Finally, a high-resolution representation needs to be generated before each saccade to the next paper, so that it will land on the paper rather than somewhere in the general vicinity. This requires a highly flexible reshuffling of short-term memory resources between each saccade.

Evidence for this dynamic re-allocation of short-term memory resources to saccade targets has been demonstrated by (Bays & Husain, 2008). In each trial of the second experiment reported by Bays and Husain, five stimuli were presented: One in the centre of the display, and four in random positions on an imaginary circle around the centre. Participants were instructed to look at each of the stimuli on the circle, and then finally make an eye movement towards the central stimulus. Crucially, during this final eye movement, the display was blanked. This meant that participants never actually

foveated the central stimulus. After the display was blanked for 250 milliseconds, a probe display appeared for 250 ms. The probe display contained one randomly selected item form the memory array, which could be displaced (location task) or rotated (orientation task) in either direction (three change magnitudes were used for each direction). The location task was performed by 8 participants, and the orientation task by 8 different participants. Participants reported the direction of displacement or rotation. Crucially, Bays and Husain reported a large benefit for the recall precision for the last item in both tasks, which was recalled with a precision of 3 to 5 times larger than all other items. They conclude that short-term memory resources must there have been shifted towards the saccade target before saccade onset.

This finding of short-term resource allocation towards saccadic targets has since been conceptually replicated twice by studies showing increased precision of recall for items that were saccade targets (Oostwoud Wijdenes, Marshall, & Bays, 2015; Shao et al., 2010), but has never been replicated directly.

The first experiment presented in this chapter was intended to be a direct replication of Bays & Husain. It employs five Gabor patches, each presented on an imaginary circle around the centre of the screen, and each with a unique orientation. As in Bays & Husain, participants were instructed to foveate and remember every stimulus, and the screen was blanked during the final saccade. Slightly different from Bays & Husain was the response method: I employed the method of adjustment rather than change detection. Also different was the positioning of stimuli: Bays and Husain put four stimuli on an imaginary circle, and a fifth in the centre; whereas in the first experiment presented here all five stimuli were on an imaginary circle (with the same

radius as in Bays & Husain).

The aforementioned results are in line with the idea that a gist representation is held in visual short-term memory for items that are not directly behaviourally relevant (e.g. saccadic targets), but that are still of importance to the individual (e.g. items that can be probed later). One prediction that one could make, is that the gist-representations are more likely to fall back to stereotypical representations. For example, a participant might not remember the exact colour of an item, but they are likely to recall that it was a shade of red.

Exactly this hypothesis was tested in a rather clever way by (Joseph, Iverson, et al., 2015). They employed an auditory short-term memory task. Participants were presented with 1, 2, or 4 speech sounds. Specifically, the stimuli were sampled from a continuous 'vowel stimulus space' see (Iverson, Smith, & Evans, 2006), which has the advantage of resembling commonly used sounds in regular speech (British English vowel sounds). When only remembering one item, participants were quite good at recalling the exact sound using the method of adjustment, which in this case was a continuous response scale with continuous vocalisation rather than a visual representation. Crucially, when participants had to recall two or even four stimuli, their recollection was progressively more biased towards the real-life vowel sounds (such as they occur in British English). From these data, one could conclude that when representations in short-term memory were less precise (because more stimuli had to be recalled), participants tended towards the stereotypical representation. A colour analogy would be that participants did not recall the exact shade, but still recollected the semantic label associated with a colour category, e.g. 'red'.

One could question whether auditory short-term memory is organised in a similar way to visual short-term memory, and thus whether Joseph, Iverson et al.'s findings can be generalised to vision. According to published work from the same group, auditory short-term memory behaves in much the same way as its visual counterpart. It shows set-size effects with more memorised items resulting in poorer performance, and also the typical primacy (first presented item recalled better than rest of the sequence) and recency (most recently presented item recalled better than rest of the sequence) effects (Kumar et al., 2013). Features of auditory stimuli that originate in the same location even seem to be bound together into objects (Joseph, Kumar, Husain, & Griffiths, 2015). In sum, there does not appear to be a clear reason to disregard Joseph, Iverson et al.'s findings in the context of visual short-term memory.

The only comparable evidence of similar 'stereotyped' representations in visual short-term memory is a study that demonstrated recall precision was not uniform across colour space (Bae, Olkkonen, Allred, Wilson, & Flombaum, 2014). Recall precision was best, on average, for regions in CIE L*a*b* space that corresponded with red, yellow, and green; and lower for 'in-between' colours (this pattern was quite universal, although individual differences were present). This could suggest that participants recalled a semantic label rather than a specific colour, but could just as easily mean that participants are better equipped for the fine-grained perception of aforementioned colours. The latter idea is supported by Bae et al., who demonstrate a similar pattern for an experiment in which recall followed immediately upon stimulus presentation, and thus required on perception more than short-term memory.

The second experiment in this chapter aims to more precisely assess whether

visual short-term memory regresses to a more gist-like (stereotyped) representation when stimuli cease to be behaviourally relevant. It will use four stimuli in the same paradigm as experiment 1: Stimuli are presented on an imaginary circle, participants foveate three of them, and the screen is blanked during participants' saccade towards the last stimulus. The stimuli were butterflies presented in a certain orientation, and with a particular two-part colouring (the colour combination varied in a continuous way). Crucially, specific combinations of colours along the continuum have previously been associated with semantic labels. The hypothesis is that recall precision would be best for the final item (as per (Bays & Husain, 2008)), and more biased towards the associated semantic labels that act as anchors for canonical, gist-like representations (as per (Joseph, Iverson, et al., 2015)).

To provide the reader with a preview of the results: Bays and Husain's results were not replicated in the first two experiments, and thus I decided to run a closer replication. This third and final experiment consisted of two parts: One with a change detection response (like their original study), and one with a response according to the method of adjustment (cf. experiments 1 and 2). Stimuli were triangles with two equal sides that varied in location along an imaginary circle centred around the screen centre, colour (in CIE L*a*b*), and orientation. In each trial, participants were told what feature would be probed, then presented with four triangle stimuli, required to fixate all of them, and again faced with a blanked screen during their saccade towards the final item. One stimulus would be chosen at random, and response would be given via change detection, or on an analogue scale.

The hypothesis was that the experiment with the analogue scale would likely not

replicate Bays & Husain, in line with experiments 1 and 2. However, the change-detection experiment employed the same response method as Bays & Husain, and was thus expected to have a higher chance of replicating the original study.

6.3. Methods and Results

6.3.1. Experiment 1

6.3.1.1. Task

A total of 9 participants were tested for experiment 1. One participant was excluded for not finishing the task: They quit after 44 trials, when it became apparent that the eye tracker could not reliably calibrate. They were recruited through the Department of Experimental Psychology's participant website, and compensated for their time at a rate of 8 pounds per hour. All participants were briefed before participating, and provided written informed consent. The experiment lasted for up to two hours.

At task onset, an EyeLink 1000 was calibrated using a 9-point calibration. The EyeLink 1000 was of the tower mount variety, and recorded the right eye at 1000 Hz.

The task was presented on a 21 inch (53.3 cm) ViewSonic P227f CRT monitor (screen dimensions: 40.5 x 30.5 cm), with a resolution of 1024 x 768 pixels, and a refresh rate of 100 Hz. The monitor was positioned at a distance of 62 cm from the participant's eyes.

The task lasted for an intended 400 trials with a maximum of 2 hours, after

which the experiment was terminated. A new trial was generated on every re-fixation of an earlier fixated stimulus, resulting in an average of 386 trials (SD=66, min=344, max=538). During each trial, participants were presented with a central fixation for 800-1100 ms. At stimulus onset, this fixation would be replaced by 5 Gabor stimuli with a diameter of 1.25 degrees of visual angle, and a spatial frequency of 5 cycles per stimulus. The stimuli were positioned on an imaginary circle around the display centre with a radius of 8 degrees of visual angle, and a minimal distance between each of the stimuli of 44 degrees (corresponding to an inter-stimulus distance of at least 6 degrees of visual angle). Each of the stimuli had a unique orientation that was sampled from a flat distribution between 0 and 180 degrees, with a minimal angular separation from the other stimuli of 20 degrees.

Participants were instructed to look at each stimulus. When they foveated at a yet unfixated stimulus for 150 ms or over, a beep was played to signal that their fixation was registered (cf. (Bays & Husain, 2008)). During the final saccade from the penultimate stimulus to the final stimulus, the screen was replaced by a grey screen, which was visible for 250 ms (cf. (Bays & Husain, 2008)).

After the grey screen, one randomly selected stimulus was re-displayed, and participants were required to recall its orientation. They did so by using the left joystick on an Xbox 360 controller, which they could point in any direction to manipulate the orientation of the stimulus. Participants could then press the A button on the controller to confirm their response, which would advance the experiment to the next trial.

6.3.1.2. Analysis

Recall error was defined as the angular difference between the actual stimulus orientation and participant's responses. Responses were organised according to the order of fixations of the stimuli, i.e. n-3, n-2, n-1, n, and n+1, where n is the current fixation (n+1 is the target of the final saccade, during which the screen was blanked). This resulted in one error distribution per fixation-ordered stimulus.

Error distributions were fitted a mixture model by (Bays et al., 2009), presented in greater detail in the first empirical chapter of this thesis (on short-term memory encoding). In short, the fitting procedure uses a likelihood minimisation procedure to estimate the parameters of a model with four parameters. This model describes recall error distributions as summations of weighted distributions. The first of these is a uniform distribution that represents guessing behaviour, because when participants guess, every response is equally likely. The next is a Von Mises distribution centred around 0 with a free spreading parameter *k* (can be recomputed into the standard deviation of an equivalent normal distribution), which represents recall when an item is remembered, because here responses around the actual stimulus orientation are more likely. The final distribution is a Von Mises distribution centred around one of the non-targets (stimuli that were presented but not probed), with the same spreading parameter as the target-centred distribution. This last distribution reflects misbinding behaviour: When a participant correctly recalls the orientation of a non-target stimulus that they confuse with the target's location, responses around that non-target orientation are most likely. Each of these distributions is weighted by a parameter so that the sum of these weights add up to 1. Essentially, the weight is the proportion of trials in which a participant forgot (or did not properly encode) the probed target's orientation, in which

the participant remembered the probed target's orientation, and in which the participant swapped a non-probed target's orientation with the target's orientation. The spreading parameter reflects the precision at which stimuli are encoded. Specifically, precision is defined as 1 over the standard deviation (computed from k) of the fitted Von Mises distributions.

Of note are the average absolute error (as a 'raw' metric), and the precision of the target's fitted Von Mises distribution. These will be analysed in a repeated-measures ANOVA with one factor fixation order with 5 levels (n-3, n-2, n-1, n, and n+1). Related-samples t-test will determine whether the final stimulus (n+1) was indeed recalled at a different absolute recall error or precision than other stimuli.

6.3.1.3. ***Results and Discussion***

Absolute error

A repeated-measures ANOVA on the average absolute error (**Figure 6.1**) with factor fixation (5 levels: n-3, n-2, n-1, n, and n+1) revealed a significant main effect of fixation, $F_{Greenhouse\text{-}Geiser}(1.38, 9.676) = 5.47$, $p = 0.034$, $\eta^2 = 0.44$. Post-hoc paired-samples t-tests revealed no difference between any of the conditions (all Bonferroni-corrected $p > 0.05$), with the exception of stimulus fixations n-2 and n, $t(7) = 4.82$, $p_{Bonferroni} = 0.019$, and stimulus fixations n-1 and n, $t(7) = 4.74$, $p_{Bonferroni} = 0.021$.

These results suggest that the absolute error was not equal across all fixated stimuli. Specifically, the last-fixated stimulus (n) was recalled with lower error than the second (n-2) and third (n-1) fixated stimuli. The hypothesis that the to-be-fixated

stimulus (n+1) would be recalled at a lower error than other stimuli was not confirmed.

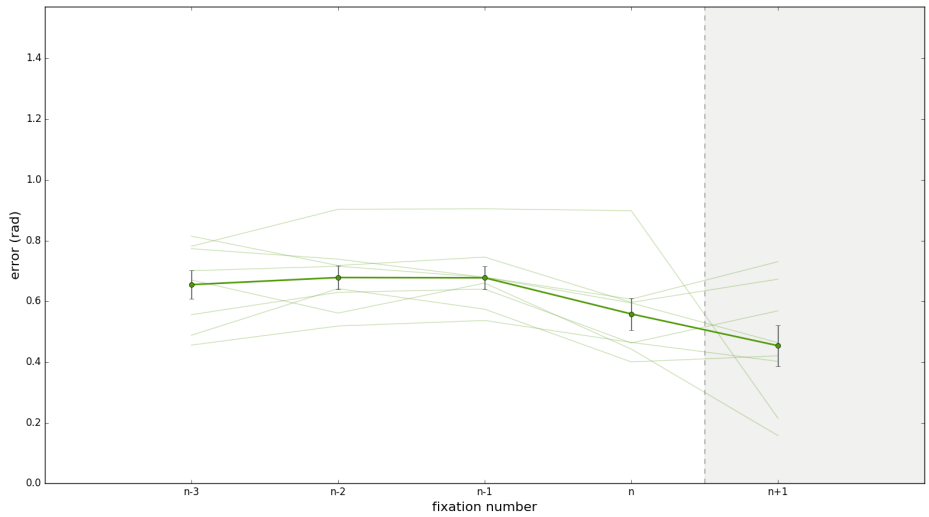

Figure 6.1 *– Average absolute error per fixated stimulus. The solid line represents a group average, the error bars the standard error of the mean. Thinner transparent lines represent individual participants. The dotted line represents the screen blanking, and the area to its right the stimulus that was never fixated.*

Mixture model parameters

A repeated-measures ANOVA on the estimated precision (**Figure 6.2**) with factor fixation (5 levels: n-3, n-2, n-1, n, and n+1) revealed no main effect of fixation, $F_{Greenhouse\text{-}Geiser}(1.39, 10.43) = 1.85$, $p = 0.207$. The to-be-fixated stimulus did not differ from any of the fixated stimuli in related-samples t-test (all Bonferroni-corrected $p = 1$).

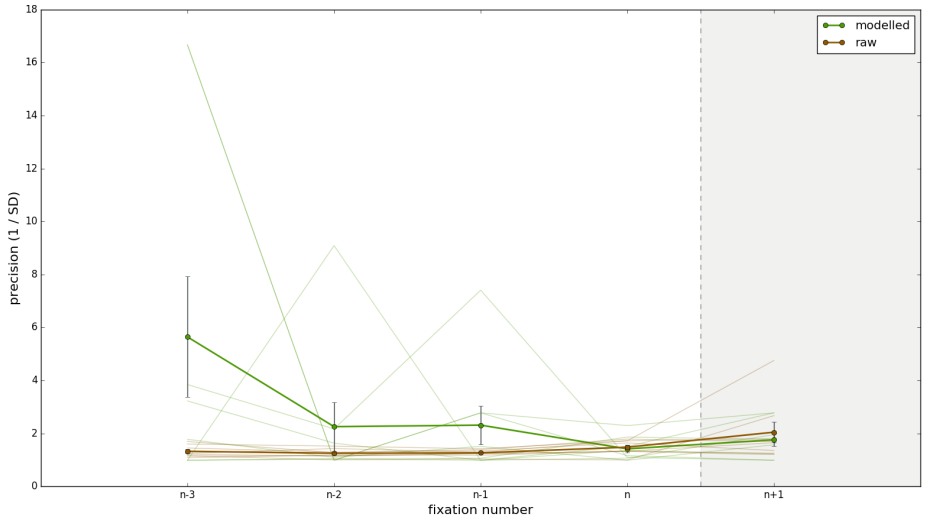

Figure 6.2 *– Estimated precision (green) and 'raw' precision (reciprocal of the standard deviation of the error; orange) as a function of fixation order. Lines represent group averages, and error bars standard error of the mean. Thinner transparent lines represent individual participants. The dotted line represents the screen blanking, and the area to its right the stimulus that was never fixated.*

A repeated-measures ANOVA on the estimated probability of guessing (**Figure 6.3**) with factor fixation (5 levels: n-3, n-2, n-1, n, and n+1) revealed no main effect of fixation, $F(4, 28) = 2.05$, $p = 0.115$. The to-be-fixated stimulus did not differ from any of the fixated stimuli in related-samples t-test (all Bonferroni-corrected $p > 0.598$).

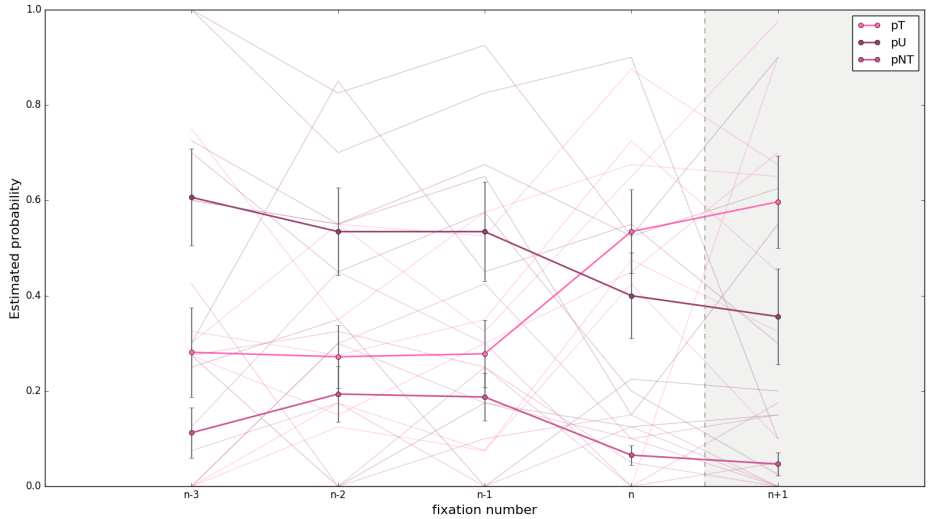

Figure 6.3 *– Estimated proportion of guessing (purple), of swapping (pink), and of remembering a stimulus (hot pink) as a function of fixated order. Solid lines represent group averages, error bars the standard error of the mean. Thinner transparent lines represent individual participants. The dotted line represents the screen blanking, and the area to its right the stimulus that was never fixated.*

These results suggest that a stimulus' place in the fixation order had no effect on the probability of it being forgotten, nor on the precision of its encoding. This directly contrasts with the hypothesis that the to-be-fixated stimulus would be recalled at a higher precision than stimuli that were fixated.

6.3.2. Experiment 2

6.3.2.1. Task

A total of 7 participants were tested for experiment 2. They were recruited through the Department of Experimental Psychology's participant website, and compensated for their time at a rate of 8 pounds per hour. All participants were briefed before participating, and provided written informed consent. The experiment lasted for up to two hours.

At task onset, an EyeLink 1000 was calibrated using a 9-point calibration. The EyeLink 1000 was of the tower mount variety, and recorded the right eye at 1000 Hz. The task was presented on the same equipment as experiment 1.

The general task outline in experiment 2 was highly similar to experiment 1, with a four differences: In experiment 2 butterfly stimuli were used that had an additional feature (location and orientation as in experiment 1, plus colour), there were only 4 instead of 5 stimuli, either the orientation or the colour of the target could be probed (indicated by a pre-cue at the start of a trial), and the experiment was preceded by a learning phase during which participants learned to associate 8 points in butterfly colour space with non-words.

The butterflies had a black body and wing outlines, and could thus be oriented in 360 degrees. The top and bottom of their wings were coloured in different shades from the same colour space. The top colour would be sampled from a circle with radius 22 around the whitepoint in CIE L*a*b* space, at an L of 50. The bottom colour would be sampled from a circle with radius 22 around the whitepoint in L*a*b* space at an L of

85, rotated 180 degrees. During the response phase, when participants rotated the joystick on their game controller, both colours changed according to the orientation of the joystick.

The training phase that preceded the experiment consisted of participants seeing 3 presentations of 8 equally interspaced points along the colour space: at 0, 45, 90, 135, 180, 225, 270, 315 degrees. During each presentation, one non-word was presented along with the coloured butterfly. Employed non-words were ‘ags’, ‘daf’, ‘dox’, ‘fud’, ‘kos’, ‘ims’, ‘vof’, and ‘zim’, which were selected based on their roughly equal bigram and trigram frequencies in English (Duyck, Desmet, Verbeke, & Brysbaert, 2004; Jones & Mewhort, 2004; Solso, Barbuto, & Juel, 1979). After each combination of colours and word was presented thrice in randomised order, participants’ recall for the combinations was tested. Each combination was presented with all 8 non-words, and participants were to click on the associated non-word. When participants correctly identified a combination for three times in a row, it was considered learned. On an error, the count was reset, and participants were provided with feedback on the correct association. Participants only started the main task after successfully completing the training phase.

6.3.2.2. Analysis

As for experiment 1, both the absolute angular error and the precision of the target’s fitted Von Mises distribution were subjected to a repeated-measures ANOVA, but this time with two factors: Feature (two levels: orientation and colour), and place in stimulus order (n-2, n-1, n, n+1). Related-samples t-tests were again used to determine

whether the final stimulus n+1 was recalled with a different error or precision compared to the stimuli n-2 to n within a feature condition.

6.3.2.3. Results and Discussion

Absolute error

A repeated-measures ANOVA on the average absolute error (**Figure 6.4**) with factors stimulus feature (2 levels: orientation, and colour) and fixation (4 levels: n-2, n-1, n, and n+1) revealed no effect of feature, $F(1, 6) = 2.74$, $p = 0.149$, and a significant main effect of fixation, $F(3, 18) = 4.82$, $p = 0.012$, $\eta^2 = 0.45$. There was also a significant interaction between fixation and feature, $F(3, 18) = 11.08$, $p < 0.001$, $\eta^2 = 0.65$.

Post-hoc paired-samples t-tests revealed a significant difference between the colour and orientation conditions, $t=2.11$, $p_{\text{Bonferroni}} = 1.0$. Further post-hoc paired-samples t-tests revealed no significant differences between any of the places in the fixation order (all Bonferroni-corrected $p > 0.05$).

These results indicate that recall error was higher in the orientation condition; an effect that seems to be exclusively driven by a somewhat higher error for the orientation of the to-be-fixated stimulus. The expected lower error for the to-be-fixated stimulus compared to fixated stimuli could not be found.

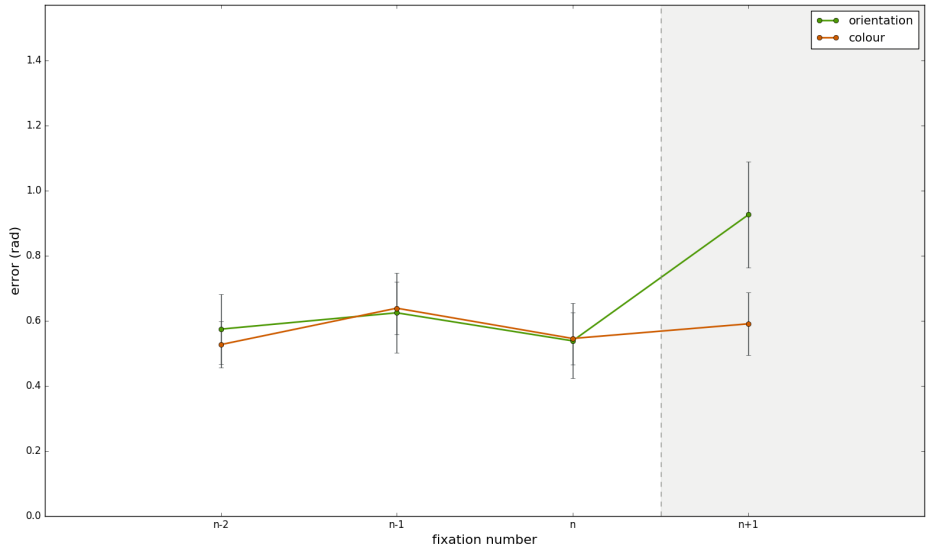

Figure 6.4 *– Average absolute error for orientation (green) and colour (orange) as a function of fixation order. Solid lines represent a group averages, error bars the standard error of the mean. The dotted line represents the screen blanking, and the area to its right the stimulus that was never fixated.*

Mixture model parameters

A repeated-measures ANOVA on the estimated recall precision (**Figure 6.5**) with factors stimulus feature (2 levels: orientation, and colour) and fixation (4 levels: n-2, n-1, n, and n+1) revealed no effect of feature, $F(1, 6) = 1.38$, $p = 0.285$, and no main effect of fixation, $F(3, 18) = 1.64$, $p = 0.216$. There was no interaction between fixation and feature, $F(3, 18) = 0.58$, $p = 0.639$.

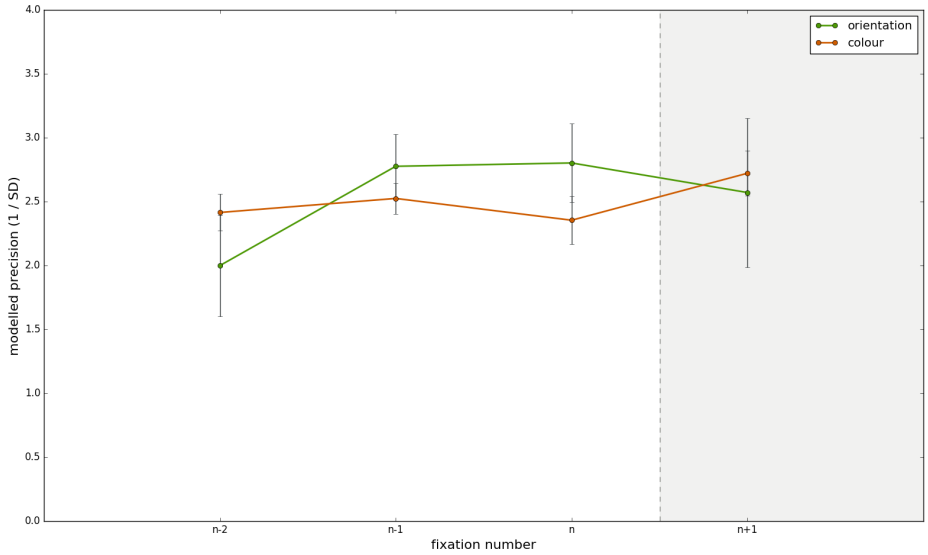

Figure 6.5 *– Estimated precision for orientation (green) and colour (orange) as a function of fixation order. Solid lines represent a group averages, error bars the standard error of the mean. The dotted line represents the screen blanking, and the area to its right the stimulus that was never fixated.*

A repeated-measures ANOVA on the estimated guessing probability (**Figure 6.6**) with factors stimulus feature (2 levels: orientation, and colour) and fixation (4 levels: n-2, n-1, n, and n+1) revealed no effect of feature, $F(1, 6) = 1.21$, $p = 0.313$, and a significant main effect of fixation, $F(3, 18) = 4.28$, $p = 0.019$, $\eta^2 = 0.42$. There was no interaction between fixation and feature, $F(3, 18) = 0.52$, $p = 0.672$.

Post-hoc paired-samples t-tests revealed no difference between the colour and orientation conditions ($p_{Bonferroni} = 0.121$). Further post-hoc paired-samples t-tests revealed no significant differences between any of the places in the fixation order (all

Bonferroni-corrected $p > 0.05$).

These results suggest that the expected difference between the to-be-fixated stimulus and all previously fixated stimuli could not be found.

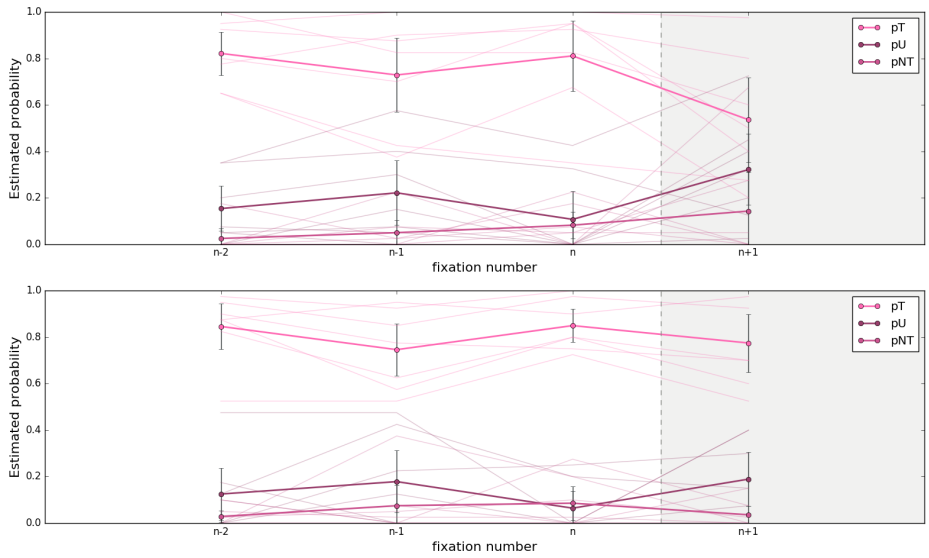

Figure 6.6 *– Estimated probabilities for guessing (purple), swapping (pink), and guessing (hot pink) for orientation (top panel) and colour (bottom panel) as a function of fixation order. Solid lines represent a group averages, error bars the standard error of the mean. Thinner transparent lines represent individual participants. The dotted line represents the screen blanking, and the area to its right the stimulus that was never fixated.*

Granularity

Although no formal analysis was undertaken, it is clear form the stimulus * response scatterplot (**Figure 6.7**) that no grouping around anchoring points occurred. More gist-like representations would look like clustered responses around the semantic labels (non-words) associated with particular points in colour space. The continuous relationship between stimulus and response orientation indicates that colours were not remembered in this way. This could well be due to the novel nature of the labels in colour space. It is more than likely that participants instead used the colour labels they have learned and used throughout their lives, e.g. ‘green’ and ‘red’. Although the double-coloured stimuli were designed to discourage using this familiar scheme, it is quite possible that participants simply ignored one half of the butterfly to encode only the other half with a semantic colour label.

It should be noted that despite this possible alternative explanation, recall in general was rather good for all stimuli’s colours, and roughly equal to their recall of stimulus orientations. This suggests participants were in fact able to remember all stimuli at a sufficiently high resolution to not require a semantic gist representation.

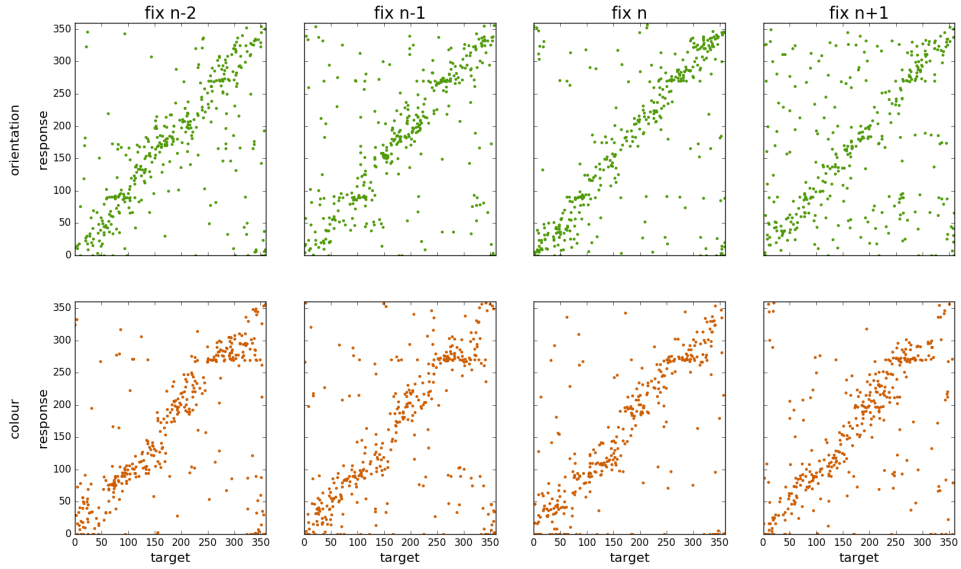

Figure 6.7 *– Target orientation (x-axis) and participants' responses (y-axis) for orientation (top row, green) and colour (bottom row, orange) as a function of fixation order (columns from left to right: n-2, n-1, n, and n+1). Each dot represents a single trial from a single participant.*

6.3.3. Experiment 3

6.3.3.1. Task

Experiment 3 was conducted in two parts. In the first part, the method of adjustment was used as a response option. In the second part, a change detection response was employed (cf. (Bays & Husain, 2008)). Experiment 3 was conducted in much the same vain as experiments 1 and 2, with the exception that stimuli were now triangles with two equal sides. These triangles had an orientation in a 360 degree space,

a colour in a 360 degree space (a circle with a radius of 22 in CIE L*a*b* space with L=50), and a location in 360 degree space (on an imaginary circle around the display centre). All of these features could be probed, and a 100% valid pre cue indicated in advance which stimulus feature would be probed in that trial.

Participants came in for two sessions of two hours each, on different days. One of the sessions would be using the change detection method, and the other the method of adjustment (continuous input); the order differed between participants. A total of 11 participants took part in the study, all doing the continuous input task variety, but only 7 also doing the change detection task. 4 participants dropped out before their second session. That all of them stopped after a continuous input session is not a coincidence: This study was run close to Oxford's summer break, during which students do not tend to stay in the city, and most drop-outs quoted this as their reason for not being able to attend the second session. The second half of the sample were those planned in for doing continuous input first, and change detection second.

The continuous input trials were presented in the same way as those in experiments 1 and 2. The change detection trials were different. Instead of a probe that they could orient with a joystick, participants were presented with a probe stimulus that was rotated along one of its feature axes: Rotated in orientation, its colour shifted along colour space, or its location shifted along the imaginary circle around the display centre. These rotations could be clockwise or counter-clockwise by 5, 20, or 45 degrees (cf. (Bays & Husain, 2008)).

6.3.3.2. Analysis

For the continuous input, as for experiments 1 and 2, both the absolute angular error and the precision of the target's fitted Von Mises distribution were subjected to a repeated-measures ANOVA, but this time with two factors: Feature (three levels: location, orientation, and colour), and place in stimulus order (four levels: n-2, n-1, n, n+1). Related-samples t-tests were again used determine whether the final stimulus n+1 was recalled with a different error or precision compared to the stimuli n-2 to n within a feature condition.

The change detection data did not provide a continuous error measure, but rather a proportion correct. This is due to every trial being a dichotomous choice between "clockwise" or "counter-clockwise" rotation of the probe compared to the target stimulus. Because participants are more likely to detect a change at larger probe rotations, their responses form a cumulative normal distribution: A low rate of "clockwise" responses for probes that were rotated counter-clockwise by 45 degrees, a roughly equal rate of "clockwise" responses at rotations of 5 degrees, and a high rate of "clockwise" responses for probes that were rotated clockwise by 45 degrees.

Cf. (Bays & Husain, 2008), I fitted the proportion of clockwise responses to a cumulative normal distribution with a uniform component. The derived parameters are conceptually equal to those obtained from the mixture modelling of continuous error distributions, and reflect the precision of recall, as well as the probability of forgetting. These were subjected to the same 3x4 repeated-measures ANOVA as the continuous input parameter estimates.

6.3.3.3. Results and Discussion

Continuous input – absolute error

A repeated-measures ANOVA on the average absolute error (**Figure 6.8**) with factors stimulus feature (3 levels: orientation, colour, and location) and fixation (4 levels: n-2, n-1, n, and n+1) revealed a significant main effect of feature, $F(2, 20) = 25.80$, $p < 0.001$, $\eta^2 = 0.72$, and a significant main effect of fixation, $F(3, 30) = 3.65$, $p = 0.024$, $\eta^2 = 0.27$. There was no interaction between fixation and feature, $F(6, 60) = 1.79$, $p = 0.116$.

Post-hoc paired-samples t-tests revealed no difference between the colour and orientation conditions ($p_{Bonferroni} = 1.0$), but a significant difference between the location and orientation and between the location and colour conditions (both $p_{Bonferroni} < 0.001$).

Further post-hoc paired-samples t-tests revealed no significant differences between any of the places in the fixation order (all Bonferroni-corrected $p > 0.05$), with the exception of the first fixation (n-2) and the last fixation (n), $t=3.42$, $p_{Bonferroni} = 0.010$, and the second fixation (n-1) and the last fixation (n), $t=3.98$, $p_{Bonferroni} = 0.002$.

These results suggest that the absolute error was not equal across all fixated stimuli and conditions. Specifically, the average absolute error was lower in the location condition than in the colour and orientation conditions (but did not differ between colour and orientation). More importantly, the last-fixated stimulus (n) was recalled with lower error than the first (n-2) and second (n-1) fixated stimuli. The hypothesis that the to-be-fixated stimulus (n+1) would be recalled at a lower error than other stimuli was not confirmed.

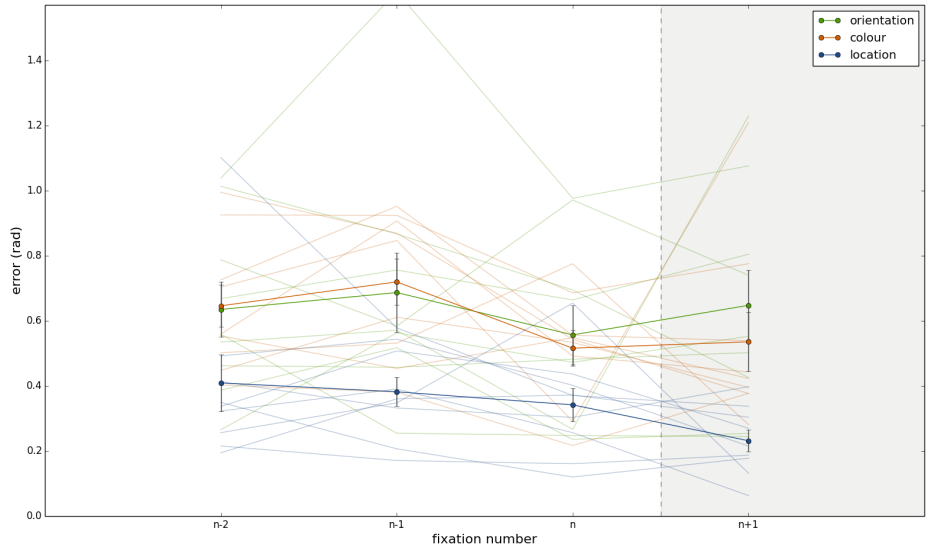

Figure 6.8 *– Average absolute error for orientation (green), colour (orange), and location (blue) as a function of fixation order. Solid lines represent a group averages, error bars the standard error of the mean. Thinner transparent lines represent individual participants. The dotted line represents the screen blanking, and the area to its right the stimulus that was never fixated.*

Continuous input – mixture model parameters

A repeated-measures ANOVA on the estimated **precision** (**Figure 6.9**) with factors stimulus feature (3 levels: orientation, colour, and location) and fixation (4 levels: n-2, n-1, n, and n+1) revealed a significant main effect of feature, $F(2, 20) = 28.45$, $p < 0.001$, $\eta^2 = 0.74$, and no effect of fixation, $F(3, 30) = 0.43$, $p = 0.733$. There was no interaction between fixation and feature, $F_{\text{Greenhouse-Geisser}}(1.84, 18.40) = 0.91$, $p = 0.413$.

Post-hoc paired-samples t-tests revealed a significant difference between the orientation and colour conditions, t = 3.36, $p_{\text{Bonferroni}} = 0.005$, a significant difference between the location and orientation conditions, t = 3.44, $p_{\text{Bonferroni}} = 0.004$, and between the location and colour conditions, t = 10.60, $p_{\text{Bonferroni}} < 0.001$.

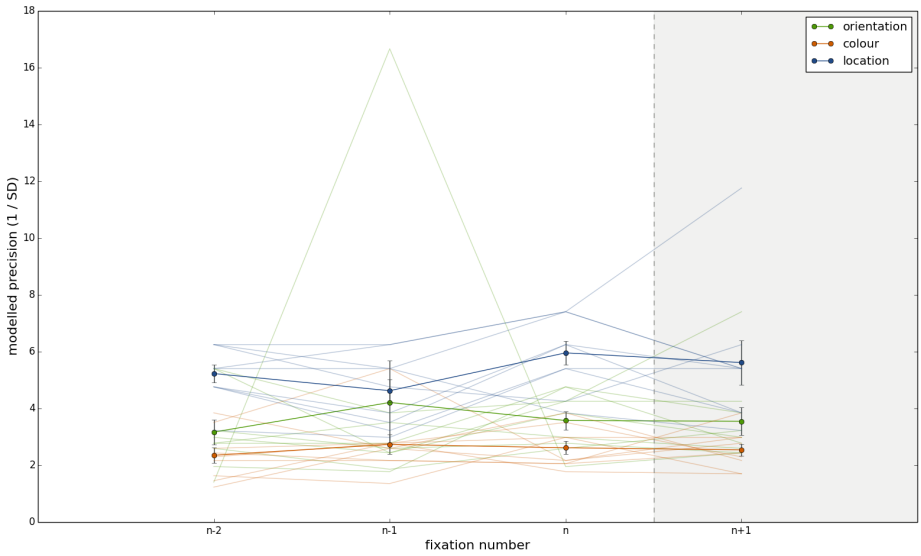

Figure 6.9 *– Estimated precision for orientation (green), colour (orange), and location (blue) as a function of fixation order. Solid lines represent a group averages, error bars the standard error of the mean. Thinner transparent lines represent individual participants. The dotted line represents the screen blanking, and the area to its right the stimulus that was never fixated.*

A repeated-measures ANOVA on the estimated **probability of guessing** (**Figure 6.10**) with factors stimulus feature (3 levels: orientation, colour, and location) and fixation (4 levels: n-2, n-1, n, and n+1) revealed a significant main effect of feature, F(2,

20) = 10.77, $p < 0.001$, $\eta^2 = 0.52$, no main effect of fixation, $F(3, 30) = 2.10$, $p = 0.121$. There was no interaction between fixation and feature, $F_{greenhouse\text{-}Geisser}(3.00, 30.00) = 0.90$, $p = 0.452$.

Post-hoc paired-samples t-tests revealed no difference between the orientation and colour conditions ($p_{Bonferroni} = 0.212$), a significant difference between the location and orientation conditions, $t = -4.85$, $p_{Bonferroni} < 0.001$, and no difference between location and colour conditions ($p_{Bonferroni} = 0.104$).

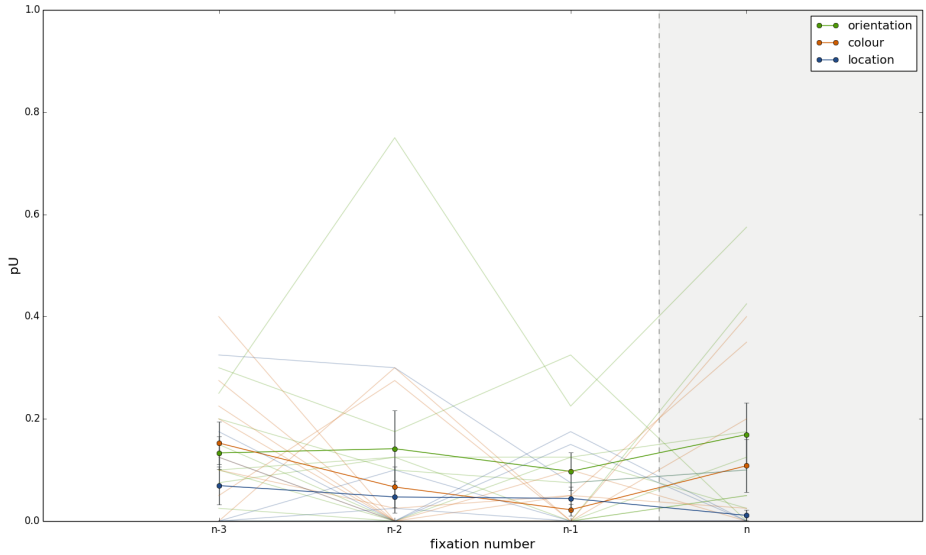

Figure 6.10 *– Estimated probability of guessing for orientation (green), colour (orange), and location (blue) as a function of fixation order. Solid lines represent a group averages, error bars the standard error of the mean. Thinner transparent lines represent individual participants. The dotted line represents the screen blanking, and the area to its right the stimulus that was never fixated.*

These results indicate that recall performance was better for location than for colour or orientation (higher precision and lower guessing probabilities). In addition, these results indicate that the expected difference between the to-be-fixated stimulus and all other stimuli could again not be observed.

Change detection – model parameters

A repeated-measures ANOVA on the estimated **probability of guessing** (**Figure 6.11, bottom panel**) with factors stimulus feature (3 levels: orientation, colour, and location) and fixation (4 levels: n-2, n-1, n, and n+1) revealed no effect of feature, $F(2, 12) = 0.25$, $p = 0.783$, and no effect of fixation, $F(3, 18) = 1.43$, $p = 0.268$. There was no interaction between fixation and feature, $F_{Greenhouse\text{-}Geisser}(2.15, 12.90) = 0.32$, $p = 0.748$.

A repeated-measures ANOVA on the estimated **precision** (**Figure 6.11, top panel**) with factors stimulus feature (3 levels: orientation, colour, and location) and fixation (4 levels: n-2, n-1, n, and n+1) revealed no effect of feature, $F(2, 12) = 2.67$, $p = 0.110$, and no effect of fixation, $F_{Greenhouse\text{-}Geisser}(1.34, 8.06) = 3.39$, $p = 0.096$. There was no interaction between fixation and feature, $F(6, 36) = 0.89$, $p = 0.129$.

These results indicate that the hypothesised difference between the to-be-fixated stimulus and fixated stimuli was not present in the data.

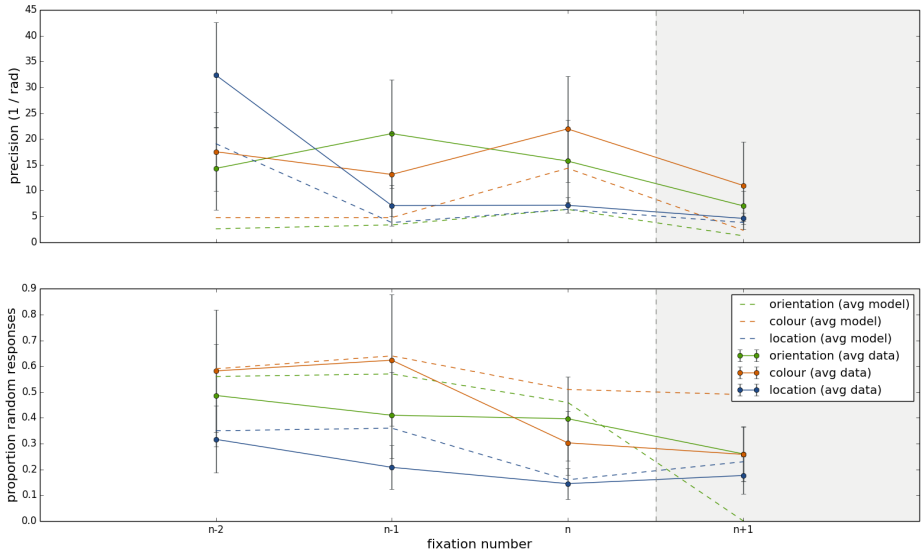

Figure 6.11 *– Estimated precision (top panel) and probability of guessing (bottom panel) for orientation (green), colour (orange), and location (blue) as a function of fixation order. Solid lines represent a group averages, error bars the standard error of the mean. Dotted lines represent group-level fits (all trials of all participants in one model fit). The dotted line represents the screen blanking, and the area to its right the stimulus that was never fixated.*

6.4. General Discussion

In three experiments, I investigated whether visual short-term memory resources were re-allocated to saccade targets before saccade onset. Despite using different stimuli (sinusoidal gratings, butterflies, and triangles), stimulus features (orientation, colour, and location), and different response methods (continuous input, and change detection),

each of the three experiments produced similar results: Average error, recall precision, and guessing probability do not differ between a stimulus that is about to be fixated, and stimuli that were fixated. These results do not support the idea that short-term memory resources are re-allocated towards fixation targets, as that would result in improved recall performance for the to-be-fixated stimulus.

6.4.1. What about Bays & Husain?

An interesting question is why I failed to replicate (Bays & Husain, 2008), and a potential answer can be found in the differences between the original study and the results presented here.

One difference is the number of stimuli. Bays & Husain employed 5 stimuli, whereas experiments 2 and 3 use four. However, experiment 1 did employ 5 stimuli, and also did not replicate Bays & Husain. Moreover, the theory is that visual short-term memory resources are re-allocated to saccadic targets, and there is little reason to expect this not to happen for 4 stimuli instead of 5.

Another difference is stimulus placement. Bays & Husain placed four stimuli on an imaginary circle around the display centre, where the fifth was presented. This stimulus thus was always present in the periphery, and participants were forced to always fixate this central stimulus last. In sum, the final stimulus was not only the last in the fixation sequence, but also special in other ways: It was the only *predictably* placed stimulus as it was always in the same location, it had the same distance to all other stimuli (which in turn had different inter-spacing), and participants were aware throughout all their fixations that the final stimulus would be the central one. In the

experiments presented here, participants were free to foveate the stimuli in the order of their own choice. This means that the last stimulus was just like all other stimuli: Unpredictably placed, and whether it was the last stimulus was entirely dependent on participants' behaviour. I would argue that this is more akin to natural vision.

Perhaps the most striking difference is the response modality. Experiments 1, 2, and 3 allow participants to respond using the method of adjustment, in which participants can provide continuous input on a circular scale. Bays & Husain employ change detection, in which participants were required to make a clockwise / counter-clockwise decision regarding a rotated probe stimulus. When I substituted continuous input for change detection in experiment 3, two results stood out: Participants were worse when asked how a probe stimulus was rotated compared to the original stimulus (compared to when they had to freely recall the stimulus orientation), and the to-be-fixated stimulus benefit that Bays & Husain reported was still not present. It is curious to see that the change detection task was harder for participants, as it seems to depend on recognition rather than recall, which is often thought to require a less strong memory trace (at least in long-term memory research, see e.g. the two-step model of recall (Hogan & Kintsch, 1971)). Perhaps the rotated probe interfered with participants' memory, and perhaps it was a very unnatural task. In daily life one would recall something from memory more often than judging whether a situation has changed.

In sum, the experiments presented here are more naturalistic than the experiments presented by Bays & Husain. Seemingly insignificant details might have helped participants in the original study to recall the final item better than all other items. At the very least, this suggests that the effect reported by Bays & Husain is not

quite as robust as originally assumed. Maybe short-term memory resources are not always redistributed to behaviourally relevant items, and perhaps they are only redistributed up to a certain extent that depends on context.

6.4.2. Superior recall for location

One striking result in experiment 3 is that the location of stimuli is recalled with much larger precision than either the colour or the orientation of stimuli. These three features were only in the same space due to arbitrarily chosen reasons (choice for locations on a circle, choice for colours in a circular space, and choice of stimuli that could point in one direction), so this finding could be trivial. However, it does speak to a particular set of theories of short-term memory organisation that treat location as a special feature of items. Indeed, these theories suggest that the binding of information associated with a single item is to the location of this item. Intuitively, this arrangement makes sense: Items occupy space, and therefore the feature that several items are least likely to share is their exact location.

A contemporary model in this class predicts that space is the feature to which other features are bound, as though the location is the central information in a look-up table (Schneegans & Bays, 2017a). In this model, when presented with a feature of one of several items stored in short-term memory, this feature is compared to all stored items'. Then the location of the closest matching item is found (accessible due to binding of features to locations). Using this location, the other features of this item can then be looked up. An example of a context in which this could occur, is when a participant is presented with one remembered item's colour and asked "What was this

item's orientation?".

In this model, it is crucially important that participants recall location with a high resolution, for it would hamper recall of all items and their features if their locations become confusable. In this light, it does not seem trivial that participants' recall was better for items' location than for other features in experiment 3.

7. Using cancellation tasks to assess spatial attention, short-term memory, and executive functioning

Parts of this chapter have been published as:

Dalmaijer, E.S., Van der Stigchel, S., Nijboer, T.C.W., Cornelissen, T.H.W., & Husain, M. (2015). CancellationTools: All-in-one software for administration and analysis of cancellation tasks. *Behavior Research Methods*, *47*(4), p. 1065-1075. doi: 10.3758/s13428-014-0522-7

7.1. Abstract

In a cancellation task, a participant is required to search for and cross out (“cancel”) targets, which are usually embedded among distractor stimuli. The number of cancelled targets and their location can be used to diagnose the neglect syndrome after stroke. In addition, the organisation of search provides a potentially useful way to measure executive control over multi-target search. Although many useful cancellation measures have been introduced, most fail to make their way into research studies and clinical practice due to the practical difficulty of acquiring such parameters from traditional pen-and-paper measures. Here I present new, open-source software that is freely available to all. It allows researchers and clinicians to flexibly administer computerized cancellation tasks using stimuli of their choice, and to directly analyse the data in a convenient manner. The automated analysis suite provides output that includes almost all of the currently existing measures, as well as several new ones introduced here. In addition, I present the data collected from 523 healthy participants over the internet to provide a frame of reference to interpret cancellation scores in. Furthermore, I employ the introduced toolbox and machine-learning techniques to find clusters of participants with similar scores. This is a data-driven way of establishing independent sub-groups with similar cognitive profiles in a large sample, and resulted in four major clusters. Two of these were small clusters with relatively many omissions (perhaps due to low task motivation), and relatively many revisits (perhaps due to low short-term memory). In addition, there were two larger groups, one with relatively poor and one with good search organisation. The presented tools and results demonstrate that cancellation tasks can provide rich data on several domains of cognition.

7.2. Introduction

Almost half of all stroke patients initially suffer from impaired attention (Lesniak et al., 2008). One of the most severe stroke-induced attention deficits is hemispatial neglect, a syndrome where patients disregard what happens towards contralesional space. It occurs in 25-50% of stroke victims (Appelros et al., 2002; Buxbaum et al., 2004; Nijboer, Kollen, et al., 2013), predominantly after damage to the right hemisphere (Ringman et al., 2004). Stroke patients suffering from neglect are hospitalized longer and face profound problems in daily life (Nijboer, Van de Port, et al., 2013; Nys et al., 2005). Although spontaneous recovery occurs, about 30-40% of individuals with neglect still suffer from the syndrome after a year (Nijboer, Kollen, et al., 2013). Neglect is associated with many negative factors, for example it appears to have a suppressive effect on upper-limb motor recovery (both synergism and strength) especially over the first 10 weeks post-stroke (Nijboer, Kollen, & Kwakkel, 2014). Importantly, neglect does not only impair spatial attention, but has also been associated with reduced sustained attention (Robertson, 2001; Robertson, Manly, Beschin, et al., 1997), and short-term memory (Husain & Rorden, 2003; Malhotra et al., 2005; Parton et al., 2006).

Because of its severity, it is important that good tools are available to diagnose neglect syndrome, and to support research on potential rehabilitation methods. One type of test that is widely used for assessment measures multi-target visual search. Such cancellation tasks require participants to cross out (“cancel”) all stimuli of a certain type, often while ignoring stimuli of all other types (distractors). These search tasks

have gained immense popularity in cognitive neuropsychology, and have proven their worth both in clinical and research environments.

Cancellation performance is not only a measure of interest in patient groups, but in other sets of participants as well. For example, a recent study on a wide age range of healthy adults described search patterns on cancellation tasks in a qualitative manner (e.g. "*horizontal left-to-right*"), and concluded that no significant differences exists between different age groups (Warren et al., 2008). However, this investigation lacked more sensitive measures of search organisation that have been shown to improve with age in children (Woods et al., 2013). Healthy elderly people tested two years before dementia require significantly more time to complete a cancellation task than elderly individuals who did not develop dementia (Fabrigoule et al., 1998). Differences in performance within demented patients became apparent when tests of a higher attentional load were deployed: patients with Alzheimer's disease performed as accurately as patients with multi-infarct dementia on a low-load cancellation task, but were both less accurate and faster on a cancellation task that required more selective and divided attention (Gainotti, Marra, & Villa, 2001). Principal component analysis of a range of neuropsychological tests, including cancellation, indicates there might be a common factor underlying performance deterioration for in the pre-clinical stage of Alzheimer's disease, perhaps associated with a general ability to control cognitive processes (Fabrigoule et al., 1998).

All of the findings summarized above could profitably be extended with more sensitive measures of cancellation performance and search organisation. When diagnosing neglect, the primary measures of cancellation tasks are usually the amount

and spatial spread of omissions (non-cancelled targets). However, there is emerging evidence that neglect syndrome constitutes more than just lateralized deficits (Husain & Rorden, 2003), and deficits of spatial working memory or sustained attention might contribute, for which additional indices of cancellation performance might be helpful.

Numerous measures of general performance, timing, and search strategy that can be derived from cancellation tasks have been suggested in the literature (for an overview, see section *3. Supported Measures*). However, data collection for these measures is often performed using labour intensive and perhaps sub-optimal procedures, e.g. frame-by-frame video analysis (Mark et al., 2004; Woods & Mark, 2007), monitoring of "verbal cancellation" (Samuelsson et al., 2002), "observing and recording the predominant search pattern" during a task by a human observer (Warren et al., 2008), or asking patients to change the colour of their pencil every 10-15 cancellations (Weintraub & Mesulam, 1988). A more efficient way of analysing search patterns would be to use a computerized cancellation task, with which cancellation positions and order can be recorded without the risk of human error.

Although the first reports of computerized cancellation software date back 15 years (Donnelly et al., 1999), the currently available packages are very limited in either the number of supported tasks (Donnelly et al., 1999; Wang, Huang, & Huang, 2006), or the supported measures (Rorden & Karnath, 2010; Wang et al., 2006), and none of them provide both task presentation and data analysis (CACTS by (Wang et al., 2006) is reported to be able to do both, but is not available for download). Therefore, most labs use custom software and most clinicians still prefer pen-and-paper tests.

Due to the lack of practically useful software, the field is currently in a situation

in which ample theoretically valid measures exists (Donnelly et al., 1999; Hills & Geldmacher, 1998; Malhotra et al., 2006; Mark et al., 2004; Rorden & Karnath, 2010; Samuelsson et al., 2002; Warren et al., 2008; Weintraub & Mesulam, 1988), of which most are validated on a small scale in research studies, but very few can be applied on a large scale in clinical practice or research studies due to the aforementioned practical issues.

In this chapter, I present a potential solution: CancellationTools, a package that combines the administration and the analysis of cancellation tasks, supporting almost all types of cancellation tests, and outputting almost all of the currently available research measures. The software is designed to be as user-friendly as possible, by using a very straightforward interface, and the option to import a scanned task that allows users to use their preferred cancellation task type. Additionally, CancellationTools supports touchscreen input, which is very comparable to pen-and-paper cancellation, for example in the sense that it allows bedside testing. The package is open source, and is available to download for free.

In addition, I use the presented software on a sample of over 523 healthy participants who did an online cancellation task. This sample was stratified for age and level of education, in order to allow for statistically meaningful analyses of these demographic factors on cancellation measures. In neuropsychology, it is generally assumed that cancellation performance is not affected by such factors, but this assumption has never been tested in a healthy sample.

Furthermore, and perhaps more interestingly, the online sample is large enough to allow for the employment of ‘Big Data’ tools to find statistical regularities in large

datasets. Specifically, cluster analysis is a data-driven way of identifying participants with similar feature profiles (each measure computed by CancellationTools can be seen as a feature of a participant), and who are thus likely to share common underlying (latent) factors. In other words, I am interested in identifying the different *cognitive profiles* that are present in the healthy population.

7.3. Methods

7.3.1. CancellationTools and the Landolt C task

CancellationTools has been written completely in Python (Dalmaijer, 2017; Van Rossum & Drake, 2011), using as few dependencies as possible. The graphical user interface (GUI) has been written from scratch using the PyGame toolbox, and the software to analyze and visualize data has been written using the NumPy (Oliphant, 2007) and Matplotlib (Hunter, 2007) packages. All of these are open source projects, that are maintained by a large community of volunteers.

The software can be downloaded for free from www.cancellationtools.org. It is released under the GNU General Public License version 3 (Free Software Foundation, 2007), which ensures that it can be used, shared and modified by anyone. The source code is publicly available and managed via GitHub, which stimulates programming with frequent feedback, version control, and collaboration on a large scale – all according to the *best practices for scientific computing* as formulated by Wilson et al. (2014).

A simplified version of the application can be used online. Due to copyright

issues, I cannot allow users to upload their own tasks to our website. I do provide different versions of the Landolt C cancellation task. After online completion of a task, a raw data file can be downloaded, which can later be analyzed via the offline version. No data will be permanently stored or accessed by the authors of CancellationTools, or any third party. An advantage of the online runtime, is that it can be accessed from computers that do not allow installation of new software (e.g. in most hospitals), or via tablet devices (e.g. Apple's iPad) that are gaining increasing popularity in neuropsychological testing.

Currently, the standalone version of CancellationTools is only available on Windows. Users of other operating systems can choose between running the application from source via Python, or using the online runtime to test participants and a PC for data analysis.

I have aimed to keep the software as user-friendly as possible, without constraining functionality. The graphical user interface (GUI) is tailored to be operated smoothly via touch screen devices and traditional PCs, and is both visually appealing and intuitive. Tasks can be set up and started within a minute. Analysing data can be done with a minimum of two mouse clicks.

CancellationTools' default cancellation task is a Landolt C cancellation task, as described by (Parton et al., 2006). The stimuli are circles with or without a gap, displayed in rows and columns with a random spatial jitter for each stimulus (**Figure 7.1**). A user is free to choose the types and number of targets and distractors, the foreground and background colour, the input type (mouse or touchscreen), and whether cancellation marks should be visible or not. The optimal placement of the stimuli (i.e.

the number of rows and columns) is automatically calculated based on the display resolution. The placement of targets is pseudo-random, as they are placed evenly over the width of the screen. In the example task depicted in **Figure 7.1**, this means that four targets are present in every column.

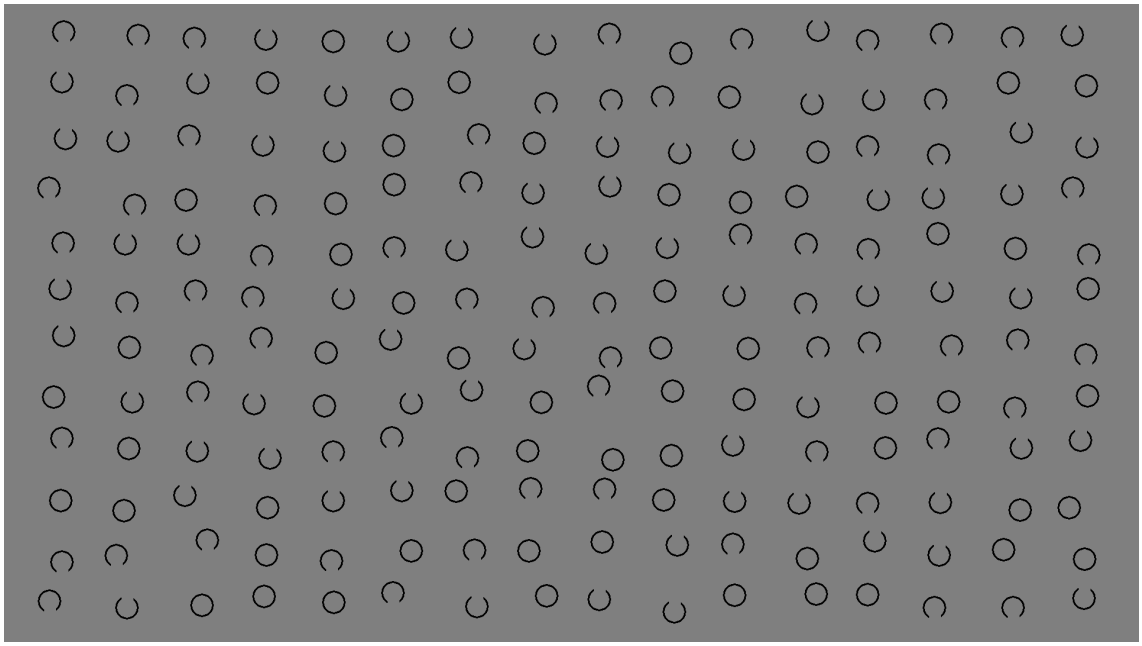

Figure 7.1 *– A Landolt C cancellation task. Here the circles with top openings are the targets (64 in total), and the circles without or with bottom openings are distractors (128 in total).*

For researchers and clinicians who prefer to work with a different cancellation task, CancellationTools has an option to import scanned tasks. If users select this option, they are asked to provide an image file. The image is automatically scaled to the display resolution, and a user can proceed to manually indicate where the targets are. The task is then saved, and is available for future use in task administration and analysis.

7.3.2. Supported Measures

I have attempted to include all of the currently existing measures that can be derived from cancellation tasks, which can be broadly divided into four categories: measures of *biases* in spatial attention, of *short-term/working memory*, of *search organisation*, and of *general performance*. Furthermore, to complement or improve on existing measures, I have devised some new metrics too (e.g. the standardised angle, see below). I have not included qualitative descriptions of cancellation path structure (Samuelsson et al., 2002; Warren et al., 2008; Weintraub & Mesulam, 1988), or an algorithm to categorise search organisation with a semantic label (Huang & Wang, 2008). In my view, these do not provide much further insight into cancellation performance than the included qualitative measures and visualizations.

7.3.2.1. Omissions

CancellationTools reports the total number of omissions and the omissions per half of the search array, which have traditionally been used to diagnose neglect. These values are to be interpreted using standardised scores, depending on what task is employed. Traditionally, a relatively large number of omissions has been used as one index of neglect, but the left-right omissions ratio is potentially more informative and has been used widely. For example, a recent study on a large sample (55 neglect patients, 138 non-neglect patients, and 119 controls) by Rabuffetti et al. (2012) reported that neglect patients show a large directional (left vs. right) imbalance in omissions, compared to healthy controls and patients with left or right lesions without neglect.

7.3.2.2. Revisits

A revisit is a cancellation of a previously cancelled target. Some authors refer to this kind of response in the cancellation literature as '*perseveration*'. However, perseverations are often used as a term associated with a (frontal) lack of ability to inhibit. In neglect research, there is evidence that while some patients might have a problem with the ability to inhibit re-cancelling a previously visited item, others re-cancel because of a deficit in spatial working memory (Mannan et al., 2005). Therefore, it could be preferred to use the empirically descriptive term 'revisit'.

Revisits can occur immediately, when a participant cancels the same target twice in a row – analogous perhaps to perseveration. A delayed revisit occurs when a participant goes back to a previously cancelled target, after cancelling other targets (Mannan et al., 2005). The number of revisits correlates with measures of disorganized search, as the best R (see below), the inter-cancellation distance, and the number of cancellation path intersections (Mark et al., 2004). (Parton et al., 2006) reported that neglect patients demonstrated a higher number of revisits than non-neglect patients, an effect that was especially apparent when no cancellation marks were visible, i.e. when patients had to remember which targets they had previously visited. In this touch screen study, the median number of intervening targets was 8. The authors argued that a possible underlying mechanism for such revisiting behaviour might therefore be a deficit in spatial working memory.

My software provides the option of using an invisible cancellation condition, should users wish to use this type of search display which can provide a more sensitive measure of left-right biases in neglect, and allows investigation of the role of spatial working memory in cancellation tasks (Wojciulik, Rorden, Clarke, Husain, & Driver,

2004).

7.3.2.3. Standardised inter-cancellation distance

Inter-cancellation distance refers to the Euclidean distance between two consecutively cancelled targets (sometimes divided by the number of targets) and has been used to asses search behaviour (Huang & Wang, 2008; Mark et al., 2004; Wang et al., 2006; Woods & Mark, 2007). I introduce a new measure that originates from the inter-cancellation distance, but is comparable across different tasks: the *standardised* inter-cancellation distance (**Figure 7.2**). This is the mean inter-cancellation distance, divided by the mean distance between each target and its nearest neighbouring target. A low standardised inter-cancellation distance originates from cancelling targets that are in close proximity of each other, and reflects an organized search pattern. Both the average and standardised inter-cancellation distance are calculated and reported by CancellationTools.

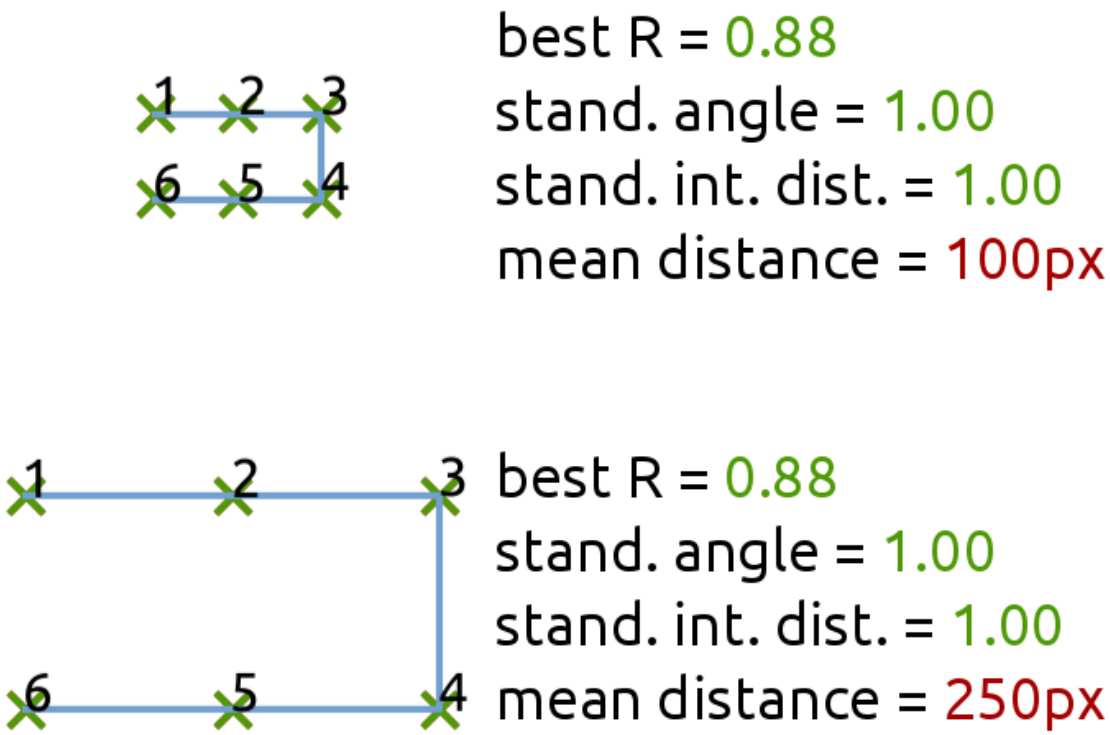

Figure 7.2 *– Two examples of a cancellation paths. The top path was obtained from a target grid with a 100 pixel inter-spacing, the bottom path from a task with 250 pixel inter-spacing. The search organisation is identical for both paths, yet the mean inter-cancellation distances are not. Other measures of search organisation (best R, standardised angle and standardised inter-distance) are.*

7.3.2.4. Centre of Cancellation

The centre of cancellation (CoC), introduced by (Binder, Marshall, Lazar, Benjamin, & Mohr, 1992) and popularised by (Rorden & Karnath, 2010), is the average horizontal position of all cancelled targets, standardised so that a value of -1 corresponds with the leftmost, and 1 with the rightmost target. The CoC is a very elegant measure of neglect severity, as it captures an attentional gradient rather than a bimodal decision (i.e. left field is or is not impaired). Additional to the horizontal CoC,

CancellationTools provides the vertical CoC, where -1 corresponds with the topmost target, and 1 with the target that is closest to the bottom of the task.

7.3.2.5. Timing

The total amount of time a participant spends on a cancellation task might be an indication of the participant's sustained attention for the task. The average inter-cancellation time (sometimes dubbed *latency index*) differs between healthy controls and brain-damaged patients, but also between neglect and non-neglect patients (Rabuffetti et al., 2012a). It could hypothetically serve as a measure of executive functioning, as it reflects how much processing time a participant needs to find and cancel a new target.

7.3.2.6. Search speed

The search speed is the average of all inter-cancellation distances divided by all inter-cancellation times (Equation 5), and has been introduced and validated by Rabuffetti et al. (2012), who show that controls are slightly faster than brain-damaged patients. This is not surprising, as the same study reports lower inter-cancellation times for patients than for controls.

(5) $$v_{search} = \frac{\sum_{i=1}^{n-1} \frac{s_i}{t_i}}{n-1}$$

Where:

n is the number of cancellations

s is the distance between two consecutive cancellations

t is the time between two consecutive cancellations

7.3.2.7. Quality of search (Q) score

A measure of the quality of search, is the Q score introduced by (Hills & Geldmacher, 1998). The Q score combines speed and accuracy in a single measure, and is calculated using Equation 6. A high Q score reflects a combination of a high number of cancelled targets, and a high cancellation speed. This index does not seem to be task independent: (Huang & Wang, 2008) found that Q scores in healthy undergraduates were higher for unstructured arrays compared to structured arrays. The number of correct responses for both task types did not differ, meaning that the difference in Q scores was driven by a higher time-on-task for the structured array. However, one should be careful when interpreting these results, as the terms 'structured' and 'unstructured' only applied to the distractors in this study: The targets locations were the same for both tasks, and only the distractors (the noise) were distributed either with or without equal spacing.

(6) $$Q = \frac{N_{cor}^{2}}{N_{tar} \cdot t_{tot}}$$

Where:

N_{cor} is the number of cancelled targets (correct responses)

N_{tot} is the total number of targets

t_{tot} is the total time spent on the task

7.3.2.8. Intersections rate

(Donnelly et al., 1999) counted the total number of cancellation path intersections: the number of times a cancellation path crosses itself (**Figure 7.3**). Mark et al. (2004) and Rabuffetti et al. (2012) divided the intersections total by the amount of produced markings to correct for search path length, resulting in the *intersections rate.* Rabuffetti et al. use the term *crossing index*, which differs from the intersections rate in one aspect: the total amount of intersections is divided by the total amount of markings, whereas the intersections rate of Mark et al. is calculated by dividing the amount of intersections by the total amount of markings excluding immediate revisits. As the cancellation path is only determined by cancelled targets, I define the intersections rate as the total amount of path intersections divided by the amount of cancellations that are not immediate revisits (Equations 7-12).

An efficient search pattern includes as few intersections as possible. In other words, a high rate of intersections would be indicative for unsystematic exploration. Rabuffetti et al. (2012) have shown that this measure can differentiate between different

groups of participants: controls < non-neglect right brain damage < non-neglect left brain damage < neglect right brain damage.

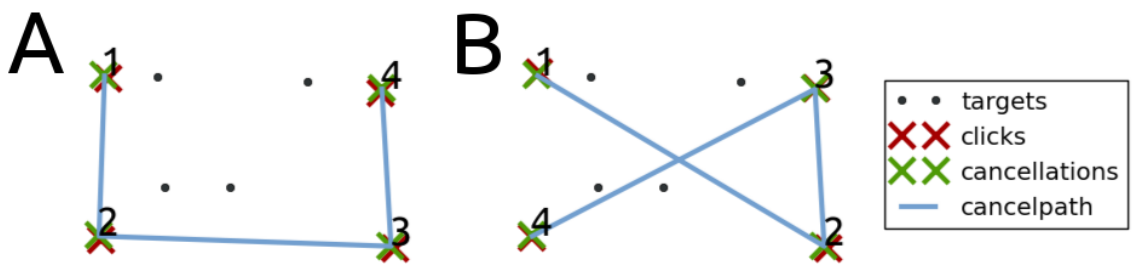

Figure 7.3 *– Examples of a cancellation paths. A) A path that does not cross itself; resulting in no intersections and an intersection rate of 0. B) A path that does cross itself; resulting in 1 intersection and an intersection rate of 0.25.*

(7) $$D_x=(x_{1,i}\cdot y_{2,i}-y_{1,i}\cdot x_{2,i})\cdot(x_{1,j}-x_{2,j})-(x_{1,i}-x_{2,i})\cdot(x_{1,j}\cdot y_{2,j}-y_{1,j}\cdot x_{2,j})$$

(8) $$D_y=(x_{1,i}\cdot y_{2,i}-y_{1,i}\cdot x_{2,i})\cdot(y_{1,j}-y_{2,j})-(y_{1,i}-y_{2,i})\cdot(x_{1,j}\cdot y_{2,j}-y_{1,j}\cdot x_{2,j})$$

(9) $$D=(x_{1,i}-x_{2,i})\cdot(y_{1,j}-y_{2,j})-(y_{1,i}-y_{2,i})\cdot(x_{1,j}-x_{2,j})$$

(10) $$(P_x,P_y)=\left(\frac{D_x}{D},\frac{D_y}{D}\right)$$

(11)

$$N_{intersect}=\sum_{i=1}^{n-1}\sum_{j=i+1}^{n-1}(D>0)\wedge((x_{1,i}<P_x<x_{2,i})\wedge(x_{1,j}<P_x<x_{2,j}))\wedge((y_{1,i}<P_y<y_{2,i})\wedge(y_{1,j}<P_y<y_{2,j}))$$

(12) $$r_{intersect}=\frac{N_{intersect}}{N_{cancellation}-N_{imm.revisit}}$$

Where:

$(X_1$, $y_1)$	is the starting coordinate of the line between two consecutive cancellations (cancellation X)
$(X_2$, $y_2)$	is the ending coordinate of the line between two consecutive cancellations (cancellation $X+1$)
$(P_x$, $P_y)$	is the coordinate of the intersection between two inter-cancellation lines
n	is the number of inter-cancellation lines (not to be confused with the number of cancellations)

7.3.2.9. Best R

(Mark et al., 2004) coined a quantitative measure for assessing cancellation strategy, which can be viewed as a formalisation of the qualitative ways in which some researchers have tried to describe cancellation paths (Samuelsson et al., 2002; Warren et al., 2008; Weintraub & Mesulam, 1988). The best R is defined as the highest absolute value of the Pearson correlation between cancellation rank number and either horizontal or vertical cancellation position (Equation 13, **Figure 7.4**), and should increase with search efficiency. The most efficient way of performing a cancellation task, is to start searching at an extremity (e.g. the left), and proceed the search in one general direction (e.g. rightward, or downward), alternating moving up and down on the perpendicular direction (e.g. upward and downward, or leftward and rightward), as is depicted in **Figure 7.4A**.

(13) $R_{best} = max(|R_{hor}|, |R_{ver}|)$

Where:

R_{hor} is the Pearson correlation coefficient of the horizontal position of all cancellations and their rank numbers

R_{ver} is the Pearson correlation coefficient of the vertical position of all cancellations and their rank numbers

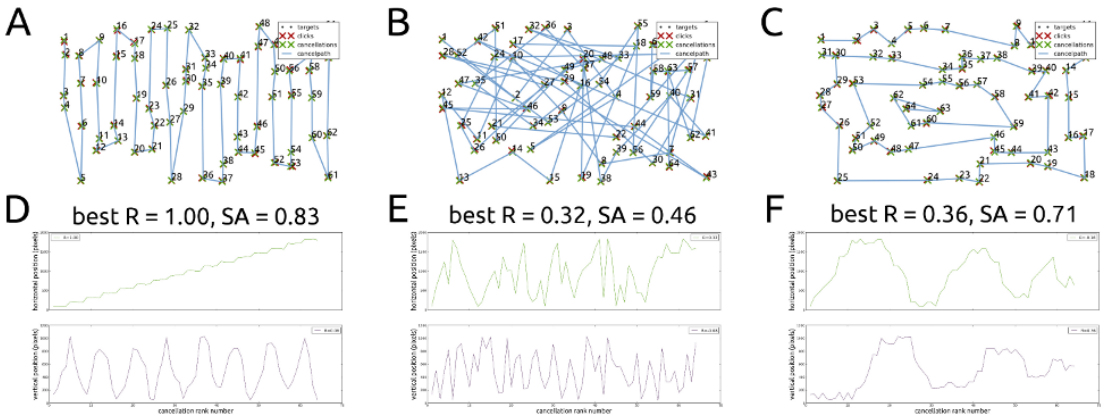

Figure 7.4 *– Three search paths to illustrate the best R and the standardised angle. A) An efficient search, starting on the left of the field and proceeding in a general rightward direction, results in a high correlation between the cancellation rank number and the horizontal cancellation position. B) An inefficient search has no general direction, and leads to low correlations between the cancellation rank number and both the horizontal and vertical cancellation position. C) Another efficient search path, resulting in a lower best R, but a high standardised angle. D-F) Best R plots of cancellation path A-C.*

7.3.2.10. Standardised angle

One of the possible cancellation paths that is efficient, but will nonetheless result in a relatively low best R, is a circular path that starts in the extremes of the cancellation task, and gradually moves inward, or spirals (**Figure 7.4C**). What characterizes this kind of path and the paths that do result in a high best R (e.g. **Figure 7.4A**), is the occurrence of predominantly horizontal and vertical lines between cancellation locations. I introduce a measure that can differentiate between horizontal and vertical paths (associated with an optimal search strategy) on the one hand, and diagonal lines (associated with a suboptimal strategy) on the other (Equations 14 and 15). As the inter-cancellation angle approaches 45º, the standardised angle approaches 0. In contrast, inter-cancellation angles approaching either 90º or 0º will result in a standardised angle that approaches 1 (**Figure 7.5**). Therefore, a high standardised angle is potentially an indication of an efficient cancellation process.

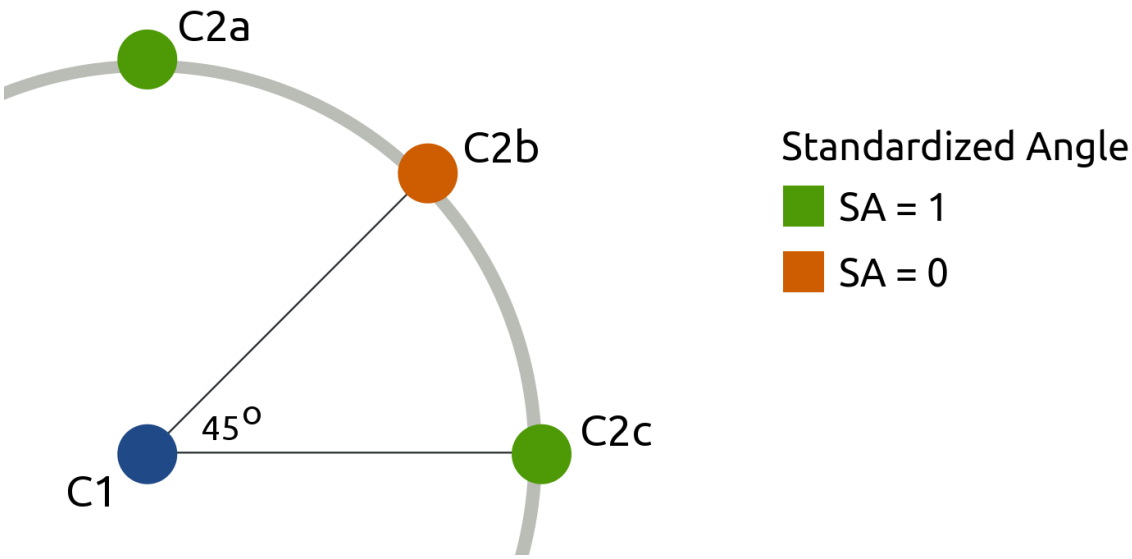

Figure 7.5 *– Illustration of the standardised angle. C1 is a cancelled target, C2a-c are potential consecutive cancellations. The standardised angle is 1 for vertical (between C1 and C2a) or horizontal (between C1 and C2c) inter-cancellation angles, and nears 0 for diagonal angles (between C1 and C2b). Paths containing predominantly horizontal and vertical lines are considered to be more efficient, therefore a high standardised angle is an indication of an efficient search pattern.*

(14) $$\gamma = \arcsin\left(\frac{\Delta y}{d}\right)$$

(15) $$\gamma_{standardized} = \frac{\sum_{i=1}^{n} \left| \frac{2 \cdot [\gamma_i]}{90} - 1 \right|}{n}$$

Where:

- γ is the angle between two consecutive cancellations
- Δy is the vertical distance between two consecutive cancellations
- d is the Euclidean distance between two consecutive cancellations
- n is the total amount of inter-cancellation angles between consecutive cancellations that are not immediate revisits

7.3.2.11. First marking

Age has a significant influence on measures of search organisation. Specifically, the mean inter-cancellation distance and the amount of intersections decrease as age increases, while the best R increases, demonstrating an improvement in search organisation over time (Woods et al., 2013). Another index that increases with age, is the likelihood of the first cancellation to be in the top-left quadrant of the search array. CancellationTools provides the location of the first marking in standardised space, so that the top left of the search array is (0,0) and the bottom right (1,1). These standardised locations are comparable between different task types and sizes. A qualitative description (e.g. “top-left”) of the quadrant in which the first cancellation happened is also available.

7.3.3. Overview

To give a preliminary indication of the ranges of the summarized cancellation measures, I tested small samples of healthy adults (N = 10) and right-hemisphere patients with leftward neglect (N = 10). They were tested on Landolt C cancellation tasks that consisted of 64 targets (opening on top) and 128 distractors (50% without opening, and 50% with an opening on the bottom), on which cancellation markings were invisible, and the time limit was 2 minutes. The averages, standard deviations, and 95% confidence intervals of all CancellationTools' quantitative measures are listed in **Table 7.1**. These values should not be regarded as norm scores. More elaborate studies on larger samples include Rabuffetti et al. (2012) (omissions, revisits, inter-cancellation distance and time, cancellation speed, and amount of path intersections in healthy controls, stroke patients with and without neglect), (Woods & Mark, 2007) (inter-cancellation distance, intersection rate, and best R in a healthy and a non-neglect stroke patient sample), (Parton et al., 2006) (immediate and delayed revisits in stroke patients with and without neglect), and (Rorden & Karnath, 2010) (centre of cancellation in neglect and non-neglect patients with right hemisphere damage).

Theoretically task-independent measures (provided there is a relatively equal spread of targets over the search array) are left-right omission ratio, standardised inter-cancellation distance, centre of cancellation, average inter-cancellation speed, intersections rate, and location of the first cancellation in standardised space. Whether this theoretical task-independency holds up in practice, might be determined in future research.

Table 7.1 *– Averages, standard deviations (between round brackets), and 95% confidence intervals (between square brackets) of a healthy sample and a neglect patient sample, collected using 1280x1024 pixels Landolt C cancellation tasks with 64 targets and 128 distractors, invisible cancellation markings, and a time limit of 2 minutes.*

	Healthy sample (N=10)	**Neglect sample (N=10)**
Omissions	2.5 (2.0)	41.6 (12.6)
(total)	[1.3, 3.7]	[33.8, 49.4]
Omissions in left half	1.4 (1.6)	27 (6.6)
	[0.4, 2.4]	[22.9, 31.1]
Omissions in right half	1.1 (1.2)	14.6 (7.3)
	[0.4, 1.8]	[10.1, 19.1]]
Revisits	2.3 (3.7)	27.8 (18.3)
(total)	[0.0, 4.6]	[16.5, 39.1]
Immediate revisits	0.6 (1.1)	12 (20.7)
	[-0.1, 1.3]	[-0.8, 24.8]
Delayed revisits	1.7 (2.7)	15.8 (13.3)
	[0.0, 3.4]	[7.5, 24.1]
Horizontal centre of	0.04 (0.02)	0.55 (0.31)
cancellation	[0.03, 0.06]	[0.36, 0.75]

Vertical centre of cancellation	0.13 (0.03)	0.01 (0.11)
	[0.11, 0.14]	[-0.06, 0.07]
Task duration	76.2 (14.3)	116.1 (6.9)
(sec.)	[67.3, 85.1]	[111.8, 120.4]
Mean inter-cancellation time	1.17 (0.19)	2.48 (0.87)
(sec.)	[1.06, 1.29]	[1.95, 3.02]
Q score	0.80 (0.14)	0.09 (0.08)
	[0.71, 0.89]	[0.04, 0.14]
Mean inter-cancellation	210.6 (32.7)	235.0 (40.6)
distance	[190.3, 230.9]	[209.9, 260.2]
(pixels)		
Standardised inter-cancellation	2.30 (0.40)	2.57 (0.48)
distance	[2.05, 2.55]	[2.28, 2.87]
Speed	0.18 (0.03)	0.11 (0.04)
(pixels per second)	[0.16, 0.20]	[0.08, 0.13]
Mean inter-cancellation angle	41.1 (29.8)	57.1 (13.0)
	[22.6, 59.5]	[49.0, 65.2]
Standardised inter-cancellation	0.74 (0.07)	0.64 (0.07)
angle	[0.70, 0.78]	[0.59, 0.68]
Best R	0.99 (0.01)	0.50 (0.29)
	[0.99, 0.99]	[0.32, 0.68]

Path intersections	1.2 (2.1)	19.3 (20.5)
	[-0.1, 2.5]	[6.6, 32.0]
Intersection rate	0.02 (0.04)	0.44 (0.43)
	[0.00, 0.04]	[0.18, 0.71]
First cancellation x-coordinate	0.03 (0.00)	0.77 (0.20)
(standardised space)	[0.03, 0.03]	[0.64, 0.89]
First cancellation y-coordinate	0.06 (0.01)	0.32 (0.17)
(standardised space)	[0.06, 0.06]	[0.22, 0.43]

7.3.4. Summarized measures

In CancellationTools, two kinds of summarized results are produced. The first is a single A4-sized, high-quality PDF document that contains an overview of all outcome measures, as well as a plot of the cancellation path (**Figure 7.6A-B**) and a heatmap of the cancelled targets (**Figure 7.6C-D**). This kind of output is potentially useful in a clinical setting, where a medical professional can consecutively run a task and an analysis, and add a print the results for a patient's file. Furthermore, a simple text file is created, which can be opened with spreadsheet software (e.g. OpenOffice Calc, or Microsoft Excel) and statistics packages (e.g. SPSS), and can also be easily processed using custom analysis scripts. Using the text files, researchers can extract data from a large group of participants for further analysis. Also available is an image of the

cancellation task with all cancellation markings that a participant made (i.e. as the participant saw the task upon finishing), and a text file of all raw click (or touch) times and coordinates.

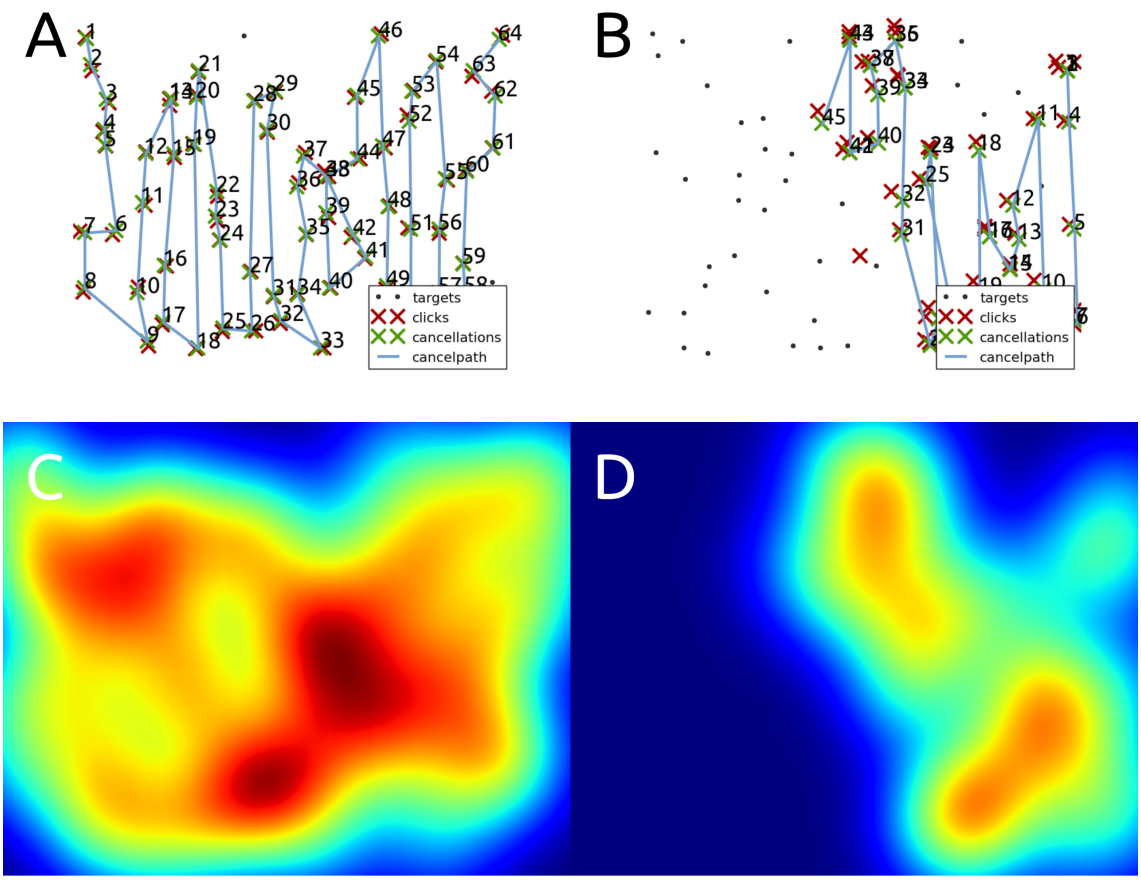

Figure 7.6 *– Output files of two example data sets, collected on a Landolt C cancellation task with invisible cancellation markings. A) A cancellation path of a typical healthy adult. B) The cancellation path of a patient suffering from severe leftward neglect. C and D) Heatmaps of the cancellation path presented in A and B.*

7.3.5. Data visualization

Several plots are created by each CancellationTools analysis. These give further insight into the performance of single participants, and can be used in addition to the measures described above. These plots include the aforementioned cancellation path and heatmap. The cancellation path (**Figure 7.6A-B**) gives a clear view of a participant's cancellation behaviour, e.g. to help with the interpretation of measures of disorganized search. A plot of the relation between the cancellation rank number and either the horizontal or vertical position of the cancelled target (**Figure 7.5D-F**) gives an indication of how organized a participants search was (Mark et al., 2004; Woods & Mark, 2007).

Heatmaps of fixation locations illustrate the deployment of attention in 2D-space, as is demonstrated by (Bays et al., 2010). With the current cancellation heatmaps we aim to create a similar visualization of spatial attention. The presented pilot testing is promising on both an individual level (**Figure 7.6C-D**), and on a group level (**Figure 7.7**). Additional heatmaps are provided based on the locations of omissions, and on the locations of path intersections, to give an indication of the spatial properties of these measures.

For the cancellation and omission a Gaussian kernel is added to the location of each cancelled or missed target. The resulting field is then scaled to the heatmap that would result from an optimal performance on the cancellation task in question, which means that heatmaps are comparable between individuals and tasks. Heatmaps for individual data from a healthy individual and a neglect patient are displayed in **Figure**

7.6C-D. Averaged heatmaps of a healthy and a neglect sample are shown in Figure 8, and show an even spread of cancellations across the search array in healthy people (Figure 8A), whereas neglect patients show a rightward bias (**Figure 7.7B**). Neglect patients also display a leftward bias of omissions (**Figure 7.7D**), whereas our healthy sample shows a lack of omissions (**Figure 7.7C**).

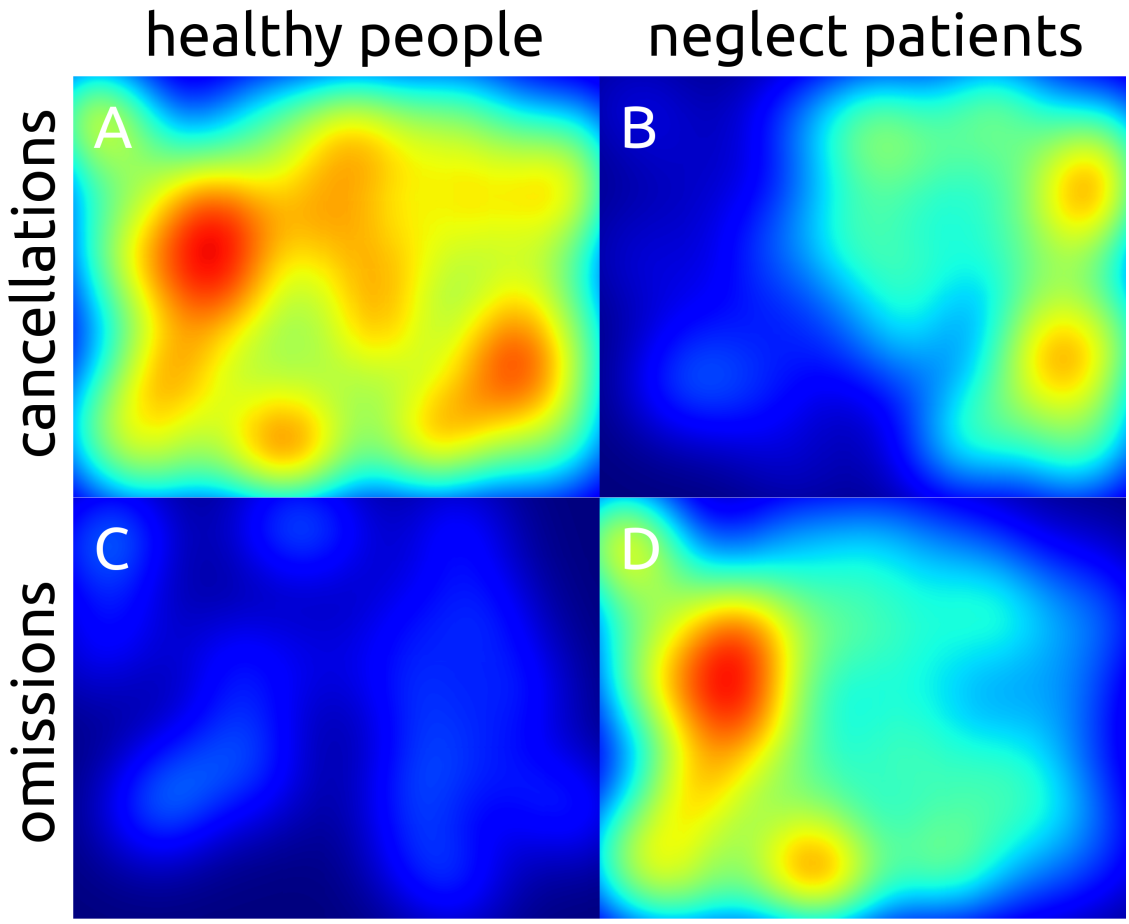

Figure 7.7 *– Cancellation heatmaps (A-B) and omission heatmaps (C-D) for a healthy sample (N=10, A and C), and a sample of right-hemisphere, leftward neglect patients (N=10, B and D). The data was collected using 1280x1024 pixels Landolt C cancellation tasks with 64 targets and 128 distractors, invisible cancellation markings, and a time limit of 2 minutes.*

7.3.6. Online Data Collection

A Landolt C cancellation task like the one described in this chapter (**Figure 7.1**) was used via a JavaScript implementation on Prolific Academic (a website that allows for online data collection). The task had a time limit of two minutes, and markings were visible. Participants were awarded one euro for their participation. Ethical permission for the study was obtained according to local legislation.

The sample was stratified by allowing for the inclusion of only particular age and education bins for each run (runs were collected in parallel). When the study went offline, a total of 535 participants had taken part. Of these, 266 were female and 269 were male. The average age was 43, with a standard deviation of 14, a minimum of 19, and a maximum of 84. In total, 28 participants had no formal qualifications, 116 had secondary school (GCSE) diploma, 101 had a College (A levels) diploma, 124 had an undergraduate degree (BA/BSc/other), 99 had a graduate degree (MA/Msc/MPhil/other), and 67 had a doctorate (PhD/MD/other). Twelve participants were excluded for not clicking on any stimulus in the task, leaving 523 participants for analysis.

7.3.7. Online Data Analysis

Data from 523 participants was first processed using CancellationTools, which computed the measures described in this Methods section. Typically, how well participants perform on a cancellation task depends on four cognitive modules: Spatial

bias, short-term memory, processing speed, and executive functioning (search organisation). To quantify these, I opted for using four measures: the total number of target omissions (a popular and simple index of neglect), the total number of revisits (correlates with short-term memory), the task duration (a proxy for processing speed), and the best R (quantifies search organisation). I opted for not including more variables to avoid the 'curse of dimensionality': The tendency that clustering analyses have to perform increasingly poorly when including more features.

After computation of the metrics, they are still in non-unified spaces. For example, best R scores exist in a range between 0 and 1, but the duration is measured in seconds (and can thus range from under ten to several hundreds). These ranges are relatively arbitrary, and will throw off most machine-learning algorithms. To account for this, values are transformed within each measure to a range between 0 and 1.

These transformed values were then fed into a series of k-means cluster analyses, as a matrix of 523 observations by 4 features. K-means clustering requires a user to set the number of clusters k, and then organises k cluster centroids in a random order. On each iteration, the algorithm tries to minimise the distance between a cluster centroid and the points that lie closer to that centroid than to any other centroid, which are considered part of that cluster. Once the centroids have reached their optimal position (with minimal inter-cluster position variability, and maximal between-cluster variability), the algorithm stops.

At this point in the analysis, a cluster-coefficient can be computed for each data point. This coefficient is scaled between -1 and 1, and is based on how close a data point is to its assigned cluster centroid. A value of 1 means that it's perfectly aligned with the

cluster centroid, a value of 0 means that it's not closer to its own centroid than to another, and a value of -1 means that the point is aligned with a the centroid of a cluster that it was not assigned to (and thus that the point might be better assigned to a different cluster). The quality of a clustering solution can be determined by averaging the cluster-coefficients of all samples. Negative values indicate many data points were likely assigned to the wrong cluster, values around zero indicate poor clustering, and positive values indicate good clustering. The general rule of thumb is that values between 0 and 0.25 indicate that the data shows no cluster structure, values between 0.25 and 0.5 indicate potentially arbitrary clustering, values between 0.5 and 0.75 indicate reasonable clustering, and values between 0.75 and 1 indicate good clustering.

Although the k parameter is user-defined, it is not necessarily an arbitrary choice. It is a common practice to run several k-means clustering analyses, and choose the solution with the highest cluster-coefficient. This is the method I will apply here too. A cluster analysis will be run for the range k=1 to k=10, and the best solution will be chosen based on the cluster-coefficient. If that coefficient is over 0.5, it will be considered as evidence for a cluster structure being present in the data.

For visualisation purposes, the data is subjected to a multi-dimensional scaling procedure. More specifically, the data is processed using a t-distributed stochastic neighbour embedding, or t-SNE (Van der Maaten & Hinton, 2008). This technique aims to reduce the number of dimensions in which the data is defined to (typically) two, making it easier to plot in a single graph. It does so by preserving local structure but not necessarily global structure, making distances in the resulting plot relatively arbitrary. It is important to note that this is a stochastic process, and will thus not always produce

the same result. It is also important to note that t-SNE is only used to produce illustrations; not for the purpose of making statistical comparisons.

7.4. Results

7.4.1. Population averages

Age is not a significant predictor of any measure, save from the duration (time-on-task) in which it accounted for 10 percent of the variance in the whole sample (**Table 7.2**). On average, older participants required more time to complete the task, at a rate of 0.59 seconds of time-on-task per year-of-age.

Table 7.2 *– Explained variance (R^2) for cancellation measures when predicted with age in a linear regression, either in the entire sample ("All" column), or within a group with the same educational background. Bold R^2 values indicate the associated p value was below 0.05 (exact p values are reported within the same cell).*

	All	**None**	**Secondary**	**College**	**Undergrad**	**Grad**	**Doctorate**
	(N=523)	**(N=27)**	**(N=112)**	**(N=100)**	**(N=124)**	**(N=94)**	**(N=66)**
Omissions	0.0000011	0.054	0.00039	0.033	0.00074	0.0016	0.0040
Revisits	0.00039	0.0081	0.020	0.00049	0.0069	0.023	0.049
Best R	0.0026	0.047	**0.037** *p*=0.0418	0.022	0.0029	0.022	0.0027
Duration	**0.10** *p*<0.00001	0.064	**0.29** *p*<0.00001	**0.080** *p*=0.00459	**0.13** *p*<0.0001	0.036	**0.12** *p*=0.00421

An ANCOVA with age as co-variate and education as fixed factor (6 levels: none, secondary, college, undergrad, grad, doctorate) revealed no effect of level of education on the total number of omissions, $F(5, 516) = 1.31$, $p = 0.258$; and no effect of age on the total number of omissions, $F(1, 516) = 0.000060$, $p = 0.994$.

An ANCOVA with age as co-variate and education as fixed factor (6 levels: none, secondary, college, undergrad, grad, doctorate) revealed no effect of level of education on the total number of revisits, $F(5, 516) = 1.43$, $p = 0.211$; and no effect of age on the total number of revisits, $F(1, 516) = 0.227$, $p = 0.634$.

An ANCOVA with age as co-variate and education as fixed factor (6 levels:

none, secondary, college, undergrad, grad, doctorate) revealed no effect of level of education on best R, $F(5, 516) = 1.24$, $p = 0.290$; and no effect of age on best R, $F(1, 516) = 1.80$, $p = 0.180$.

An ANCOVA with age as co-variate and education as fixed factor (6 levels: none, secondary, college, undergrad, grad, doctorate) revealed no effect of level of education on task duration, $F(5, 516) = 1.46$, $p = 0.203$; and an effect of age on task duration, $F(1, 516) = 60.71$, $p < 0.001$, $\eta^2 = 0.10$. The latter effect mirrors the aforementioned correlation between age and time-on-task.

7.4.2. K-Means Clustering

The solution of a 4-means cluster analysis proved to produce the highest cluster-coefficient (**Figure 7.8**), and should thus be taken as the best solution. The value is over 0.5, and should thus be treated as evidence for a clustered (multi-modal) structure being present in the data (for a full silhouette plot, see **Figure 7.9D**).

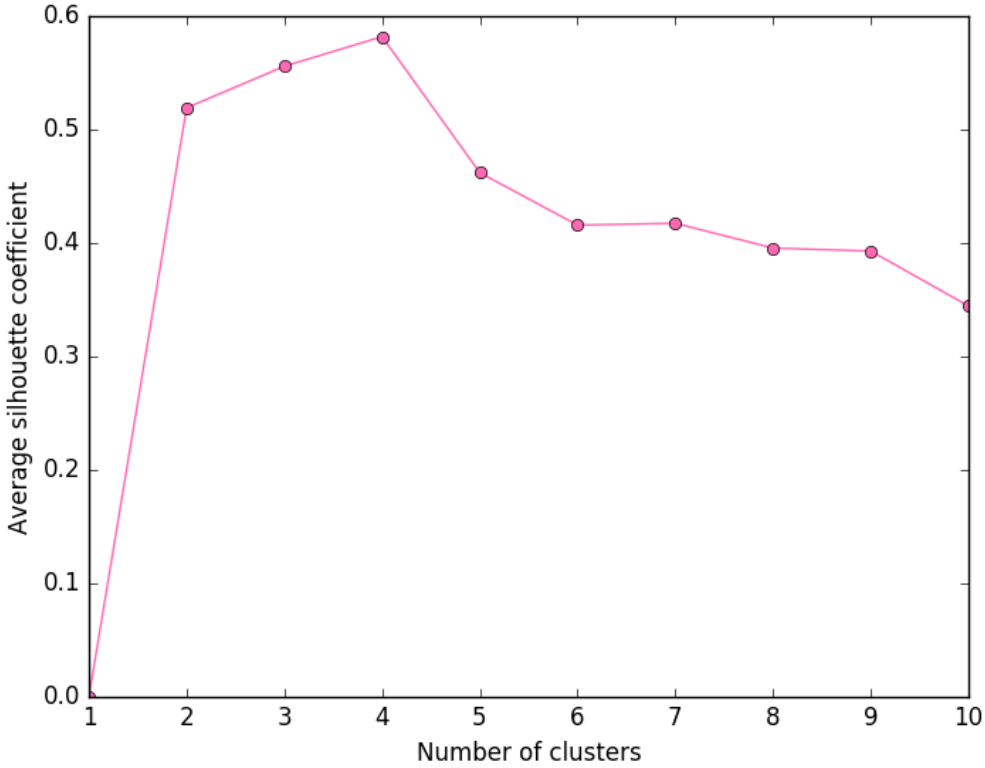

Figure 7.8 *– Cluster coefficient (y-axis) as a function of cluster size k (x-axis). A higher cluster coefficient indicates better clustering. The best solution here is for k=4 clusters.*

For the best k=4 solution, the produced clusters are one with high omissions of 18 participants, one with high revisits and 18 participants, one with a low best R of 125 participants, and one with a high best R of 362 participants. Task duration was slightly lower in the high omission group, but not by much (**Figure 7.9A**). Scatterplots of feature space are presented in **Figure 7.9A and 7.9B**. A scatterplot of the clusters in reduced space (using t-SNE) is plotted in **Figure 7.9C**. Cluster coefficients for each sample are plotted, per cluster, in **Figure 7.9D**. A plot of each cluster's average scores on the four features (best R, task duration, number of revisits, and number of omissions) in min-max scaled space is plotted in **Figure 7.10A**.

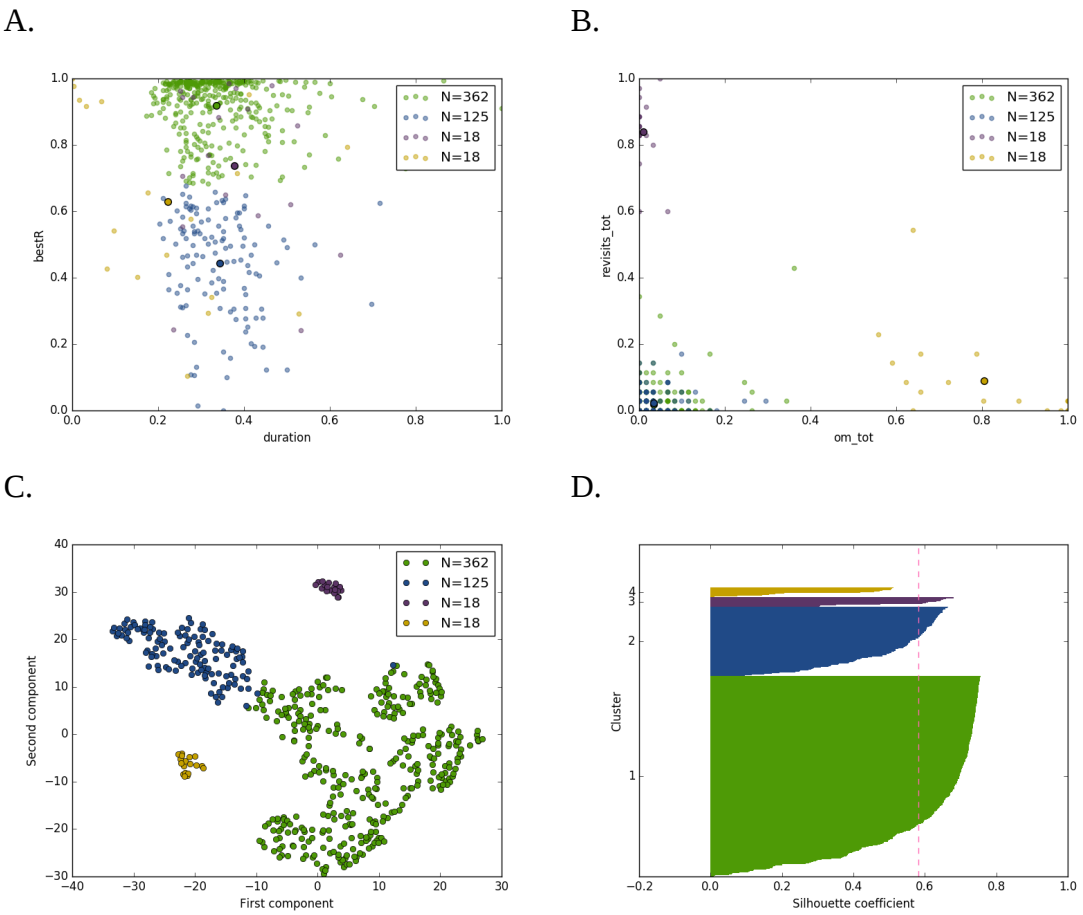

Figure 7.9 *– Descriptive plots of the four clusters, one with structured search organisation (N=362, green), one with unstructured search organisation (N=125, dark blue), one with a high number of revisits (N=18, purple), and one with a high number of omissions (N=18, yellow). Each dot in the scatterplots represents an individual participant.* ***A)*** *Scatterplot of best R (y-axis) and duration (x-axis), clearly dissociating the organised (N=362, green) and disorganised search (N=125, dark blue) clusters. Values are min-max scaled to fit between 0 and 1.* ***B)*** *Scatterplot of the number of revisits (y-axis) and the number of omissions (x-axis), clearly dissociating between the high revisits cluster (N=18, purple), the high omissions cluster (N=18, yellow), and the other two clusters.* ***C)*** *Scatterplot in along two components. Space was reduced from four features to two components using t-distributed stochastic neighbour embedding (t-SNE). In this reduced dimensional view, it is easier to see the four separate clusters.* ***D)*** *Silhouette plot of the k=4 cluster solution. Each sample is organised along the y-axis, sorted from high to low cluster coefficient (x-axis) within each cluster. The dotted line represents the average cluster coefficient of the k=4 solution.*

None of the clusters differ significantly from the whole sample in their proportion of males (all $p > 0.052$ in one chi-square test per cluster), or in their proportion of different levels of education (all $p > 0.17$ in chi-square test per cluster); see **Figure 7.10B**. Average age also does not differ between the clusters (**Figure 7.10B**). These results indicate that cancellation performance does not predict demographic values, which mirrors the aforementioned results that showed no evidence that demographics predict cancellation performance (save from a minor effect of age on duration).

A.

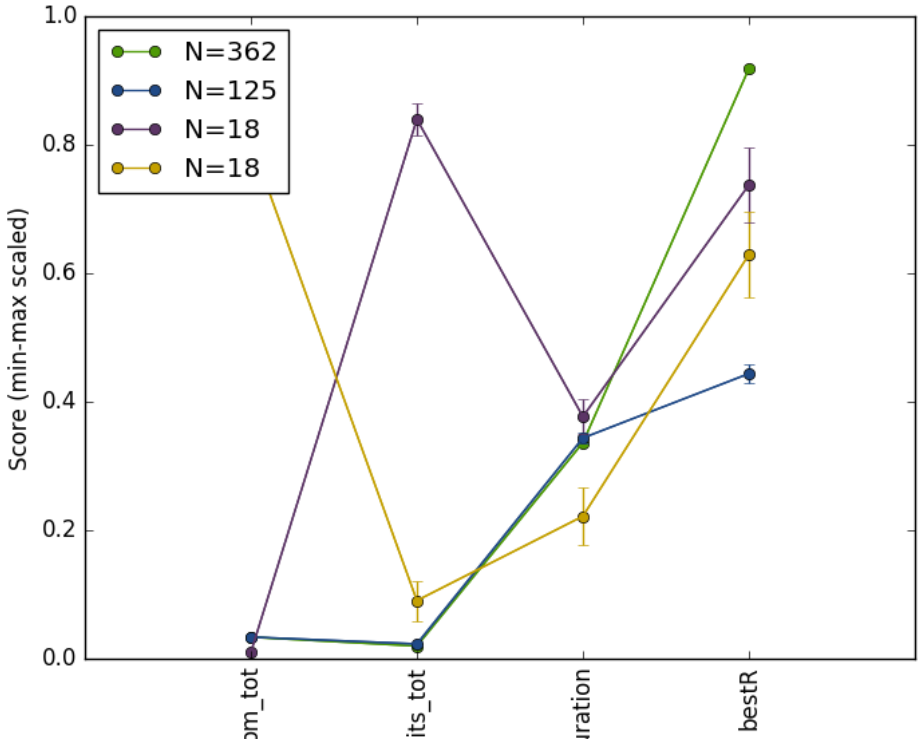
N=362
N=125
N=18
N=18
Score (min-max scaled)
1.0
0.8
0.6
0.4
0.2
0.0
its_tot
bestR

B.

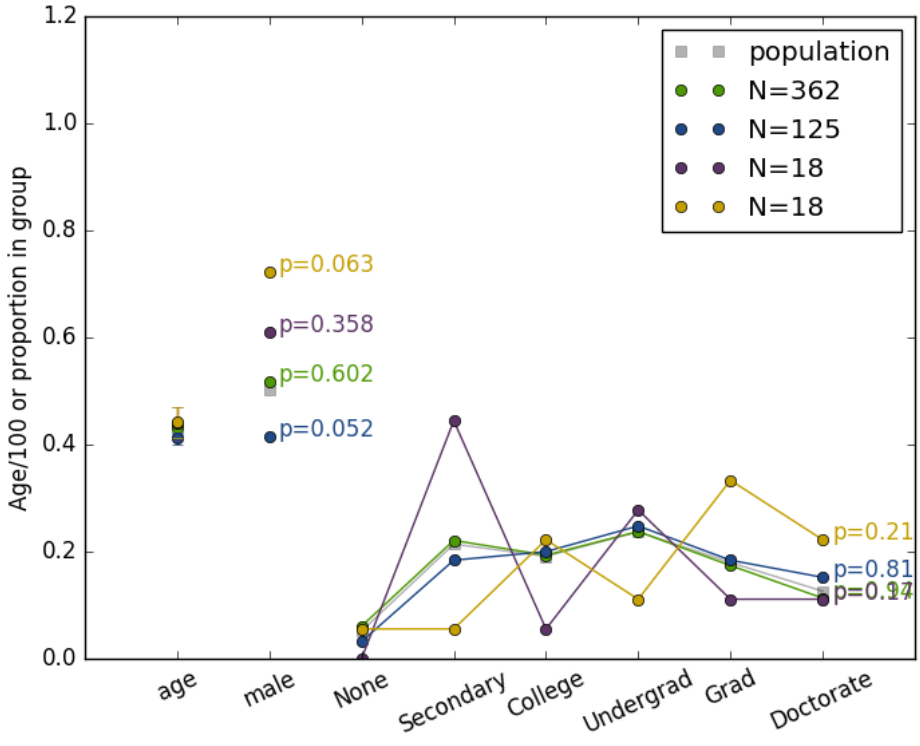
population
N=362
N=125
N=18
N=18
Age/100 or proportion in group
1.2
1.0
0.8
0.6
0.4
0.2
0.0
p=0.063
p=0.358
p=0.602
p=0.052
p=0.21
p=0.81
age
male
None
Secondary
College
Undergrad
Grad
Doctorate

Figure 7.10 – *Descriptive plots of the four clusters, one with structured search organisation (N=362, green), one with unstructured search organisation (N=125, dark blue), one with a high number of revisits (N=18, purple), and one with a high number of omissions (N=18, yellow). Solid lines represent within-cluster averages, and error bars the standard error of the mean.* ***A)*** *Min-max scaled averages (y-axis) for four features (number of omissions, number of revisits, task duration, and best R. The clusters differ from each other on these features per definition, as these are the values used to identify them.* ***B)*** *Average age, proportion of males, and proportions of highest education (None, secondary school, college, university undergraduate, university graduate, and university or medical doctorate) in all four clusters. Annotated p-values report significance of chi-square tests that tested whether a cluster's make-up significantly differed from the whole sample (plotted in grey squares).*

7.5. Discussion

There is a need to quantify multi-target visual search performance on cancellation tasks. I made an effort to summarise all of the currently available measures that can be derived from cancellation task data. In the new software package CancellationTools, I included all relevant measures from the currently available literature in an application that can be used to administer a computerised cancellation task, and to analyse the resulting data with the click of a button. I have aimed to make

this software as flexible as possible, e.g. by allowing users to incorporate their own scanned tasks into the software, whilst keeping an eye on simplicity. The result is a user-friendly interface that can be employed both in clinical and research settings. CancellationTools is open source, and free to download and use by anyone.

Furthermore, a sample of 523 healthy participants was collected using an online cancellation task. In this sample, age and level of education did not affect cancellation performance. A cluster analysis identified four cognitive profiles: Participants who make many omissions (N=18), those who make many revisits (N=18), participants with relatively bad search organisation (N=125), and individuals with relatively good search organisation (N=362). Within both larger groups, participants still show intra-individual variability in the time they take to complete the cancellation task, and in their search organisation scores.

7.5.1. CancellationTools and its measures

I have introduced two new measures of search organisation: the standardised inter-cancellation distance and angle. The former is an improvement of the existing mean inter-cancellation distance, which takes into account the distances between targets within a search array, therefore allowing comparisons of cancellation performance on different tasks. The standardised inter-cancellation angle can be viewed as complimentary to the best R, as it is robust to situations where the best R does not reflect search organisation optimally (**Figure 7.4C**). Even though the best R and standardised inter-cancellation angle seem to differentiate between our small test

groups, a much larger difference between healthy people and leftward neglect patients is observed in the intersections rate, suggesting that this might be the clearest measure of search organisation.

CancellationTools is already useful to clinicians, as it provides quantitative data on established measures of neglect (e.g. number of omissions), as well as qualitative data that provides better insight in patient behaviour than pen-and-paper cancellation tests (e.g. cancellation path plots). However, for the majority of the measures summarised above, there are currently no norm scores to compare individual test results to. The value ranges that we provide based on our pilot testing (**Table 7.1**) serve as a preliminary indication of how neglect patients and healthy controls differ on different measures, and should not be treated as a clinical directive.

Apart from our newly introduced standardised angle measure, all of the indices we report have been validated on a small scale in the articles in which they were coined. A few have been validated on a larger scale in the study of Rabuffetti et al. (2012), but it is arguable whether this provides enough data to base norm scores on.

7.5.2. Cancellation measures in the healthy population

The results presented here show that neither age nor level of education affect cancellation performance, with the exception of age accounting for a small part (ten percent) of the variance in the time it takes for healthy participants to complete the task (not regularly used as an index of performance). This is an important finding, as it

supports the long-standing (but untested) assumption that cancellation performance is not affected by demographical factors. Stroke and its associated syndromes can occur at any point in life, and thus it is important to test stroke patients with metrics that are not sensitive to potential confounds such as level of education or age.

In addition to reassuring clinicians, the current results on healthy participants demonstrate that cancellation tasks can be used to assess cognition. Four different groups emerged from a k-means cluster analysis that was run on 523 observations with 4 features each. The chosen features were the number of omitted targets, the number of revisits (delayed and immediate), the best R to index search organisation, and the task duration to index processing speed. These were chosen to be the most representative of the cognitive domains that a cancellation task aims to test: spatial attention (omissions), short-term memory (revisits), search organisation / executive functioning (best R), and general performance (time-on-task).

The resulting clusters included two small clusters, with 18 members each. The first of these was a cluster with many omissions. This cluster did not show a significantly different spatial bias from other clusters, and thus might instead just reflect a group of participants that were not motivated to perform the task in the correct way. This is supported by the clusters comparatively low time-on-task. The second small cluster showed many revisits. This is surprising behaviour, as the markings were visible during the task. Hence the problem in these participants might not necessarily be with their short-term memory, but rather with impulse control.

A larger cluster with 125 members was characterised by relatively poor search organisation. Within this group, there was still a reasonable variability of best R scores,

as was there for task duration (= time-on-task). This demonstrates that although these participants share a cognitive profile, there are still individual differences.

The final cluster identified here consisted of 362 members, and was characterised by well-structured search, but did not differ from the 125-member group on task duration. In a large proportion of the sample, best R scores plateaued, although some variability remained. Perhaps more interesting is that there was still considerable variability in the time-on-task even among participants with a plateau-level search organisation. This, and the fact that the poor and good search clusters did not differ in average task duration, suggests that processing speed is an individual characteristic that is independent from the ability to search in an organised way.

One could debate the choice of four representative measures over using all available measures as features in the cluster analysis. When ran with more features, roughly the same clustering results occurred, at the cost of lower clustering coefficients. This is likely a by-product of 'the curse of multi-dimensionality': the more features are analysed in a cluster analysis, the less clear results become. To reduce dimensionality, I argue that it is useful to focus on representative measures, especially when some correlate very strongly with each other, and when results are not affected in meaningful ways after omitting excess features.

7.5.3. Interpretation of cancellation measures in stroke

The results from the healthy sample presented here outline what one could learn from a stroke patient's test scores on a cancellation task. In particular, it illustrates that very few healthy participants miss many targets, and thus that a reasonably low number of omissions (say 5) would already be indicative of potential neglect. On the other hand, a low score on a search organisation measure does not necessarily mean that the patient's stroke is to blame. As demonstrated here, a large proportion of the healthy population will actually score relatively poorly, and thus a poor score after stroke is not necessarily a product of that stroke. Unless a pre-stroke measurement exist, low search organisation should be interpreted with care.

7.5.4. Conclusion

By making CancellationTools publicly available, and by demonstrating the utility of large datasets and machine-learning tools, I hope to inspire large-scale international collaborations to pool data, from healthy people and patient groups, on all of the measures summarised here. By removing practical boundaries that previously prevented large-scale testing, CancellationTool opens up exciting new research possibilities. One of those possibilities is highlighted here: 'Big Data' statistical tools can identify independent clusters of participant with similar cognitive profiles.

8. The effects of guanfacine on visual search performance, sustained attention, and short-term memory after stroke

Parts of this chapter have been published as:

Dalmaijer, E.S.* & Li, K.*, Gorgoraptis, N., Leff, A., Cohen, D., Parton, A., Husain, M., & Malhotra, P. (in press). A randomised double-blind, placebo-controlled crossover study of single-dose guanfacine in unilateral neglect following stroke. *Journal of Neurosurgery, Neurology, and Psychiatry*.

8.1. Abstract

After stroke, some patients develop hemispatial neglect, a syndrome in which their spatial attention is biased away from the contra-lesional visual field. In addition, stroke patients can suffer from a reduced ability to sustain attention, and from deficits of executive function that can result in for examle disorganised search. Here, I present the results of a randomised, double-blind, placebo-controlled, crossover study of guanfacine, a noradrenergic agonist that is thought to work primarily through frontal alpha-2 receptors. It has previously been shown to improve attention and working memory in monkeys and humans, and promising results in neglect syndrome have been reported in a pilot study. Thirteen stroke patients with hemispatial neglect were recruited, and tested on an invisible cancellation task (see Chapter 7), a sustained attention task, and a spatial working memory task. There was a small but significant increase in the number of targets that patients found on a cancellation task when given 2 mg of guanfacine. However, there was no convincing evidence for or against beneficial effects of guanfacine on patients' spatial attention bias. Analysis of the data using Bayesian statistics revealed reliable evidence against any effects of guanfacine on search organisation in the cancellation task. In addition, there was a reliable lack of evidence for any effects of guanfacine on sustained attention in a continuous performance task, nor on a spatial working memory task. After a single 2 mg dose, guanfacine did not improve attentional bias, sustained attention, or working memory in neglect patients. However, it did increase the amount of correctly marked targets in a cancellation task compared to placebo.

8.2. Introduction

Almost half of all stroke patients initially suffer from impaired attention (Lesniak et al., 2008). One of the most severe stroke-induced deficits is hemispatial neglect, a condition which occurs predominantly after damage to the right hemisphere (Ringman et al., 2004) in which patients' attention is biased away from contra-lesional space (Bays et al., 2010). Stroke victims suffering from neglect are hospitalized longer and have profound problems in daily life (Nys et al., 2005). Despite its serious impact, no widely accepted and effective therapies exist.

It has been proposed that neglect is not only a disorder of spatial attention, but comprises non-spatial attentional deficits as well (Husain & Rorden, 2003; Robertson, 2001). Reduced vigilance is one of these deficits, and is associated with more severe neglect (Robertson, Manly, Andrade, et al., 1997). Functional imaging studies in healthy individuals indeed shows that the right frontal and parietal lobes, which are often damaged in neglect, are activated during maintained vigilance (Lawrence, Ross, Hoffmann, Garavan, & Stein, 2003; Pardo, Fox, & Raichle, 1991; W. Sturm et al., 1999; Walter Sturm & Willmes, 2001). Furthermore, neglect patients are impaired in sustaining attention, and the degree of this impairment is correlated with reduced recovery (Robertson, Manly, Andrade, et al., 1997). In sum, impaired vigilance is an important contributor to neglect severity, and neuropharmacological improvement of this deficit could potentially decrease the clinical severity of the syndrome.

Most research on drug treatment of neglect has used dopaminergic therapies, providing conflicting results (Barrett et al., 1999; Fleet et al., 1987; Geminiani et al.,

1998; Grujic et al., 1998; Hurford et al., 1998; Mukand et al., 2001; Pierce & Buxbaum, 2002), and often lacking double-blinding and placebo controls. Despite the dopaminergic focus in the existing literature, there is substantial evidence that links vigilance to central noradrenergic pathways (Aston-Jones, Rajkowski, & Alexinsky, 1994; Berridge & Waterhouse, 2003; Foote, Aston-Jones, & Bloom, 1980; Smith & Nutt, 1996; Witte & Marrocco, 1997). Indeed, a small but double-blind, randomized and placebo-controlled pilot study involving only three patients has shown that guanfacine (an α_{2A} noradrenergic agonist) improved leftward space exploration in two right-hemisphere neglect patients (Malhotra et al., 2006).

Guanfacine has been used in the effective treatment of Attention-Deficit Hyperactivity Disorder, ADHD (Biederman et al., 2008; Scahill et al., 2001). Remarkably, even a single dose of guanfacine improves working memory in healthy humans (Jäkälä et al., 1999), as well as non-human primates (Arnsten, Cai, & Goldman-Rakic, 1988; Avery, Franowicz, Studholme, Dyck, & Arnsten, 2000), and improves performance on attentional tasks in non-human primates (Arnsten & Contant, 1992), presumably via alpha-2A receptors in prefrontal cortex (Ji, Ji, Zhang, & Li, 2008).

Guanfacine is an agonist of the alpha-2A adrenoreceptor that was originally developed as an antihypertensive, but has been found to have effects on attention. Guanfacine differs from alpha-2 agonist clonidine in that it is a phenylacetylguanidine derivative and it is more selective in its alpha-2 agonism (Bream, Lauener, Picard, Scholtysik, & White, 1975; Westfall & Westfall, 2006).

Guanfacine has been shown to improve performance in a delayed response task in healthy aged monkeys, and decrease distractibility during such a task (Arnsten &

Contant, 1992). In healthy humans, guanfacine has been shown to improve paired associates learning as well as planning and spatial working memory (Arnsten et al., 1988; Jäkälä et al., 1999), whilst in ADHD it leads to improvements in sustained attention (Scahill et al., 2001). Guanfacine is thought to exert its positive effects via the alpha-2A adrenergic receptor subtype in pre-frontal cortex (Avery et al., 2000; Ji et al., 2008), a region that has been implicated in the regulation of attention (Robertson & Garavan, 2004).

The aforementioned study of guanfacine in neglect patients (Malhotra et al., 2006) employed both pen-and-paper and two computerized cancellation tasks including one that did not allow visible marking of previously cancelled targets, which required patients to keep track of the locations of previous cancellations (Parton et al., 2006). When Malhotra et al. employed this type of cancellation task, two out of three patients who were treated with guanfacine demonstrated increased time spent on the task, which suggests guanfacine improved the ability to sustain attention. These patients also demonstrated greater exploration of the left side of space, although it is unclear whether this was driven by the extended time-on-task, or even by improved working memory performance. One patient with extensive prefrontal damage did not show any improvement, which was in line with the idea that guanfacine works via the alpha-2A receptors in frontal cortex (Ji et al., 2008).

To replicate and to further explore the underlying mechanism of reported beneficial effects of guanfacine in visual neglect, 13 stroke patients with and without extensive frontal damage were tested. To decompose the improved time-on-task and spatial exploration reported by Malhotra et al. (2006), an invisible cancellation task *with*

a fixed time limit, a test of sustained attention, and a working memory task were employed. It was hypothesised that guanfacine would improve sustained attention (as in Malhotra et al., 2006) and potentially working memory (as in Jäkälä et al., 1999), but not necessarily cancellation performance when there is a fixed time limit.

My role in the project was analysis of the data obtained using the software I have developed (see Chapter 7) and a combined frequentist and Bayesian statistical approach.

8.3. Methods

8.3.1. Patient sample

Stroke patients were screened for neglect using the Mesulam shape (Mesulam, 1985) and BIT star (B. A. Wilson, Cockburn, & Halligan, 1987) cancellation tasks. Patients who demonstrated evidence of robust visual neglect when tested twice with these cancellation tests (specifically an overall score on one or both tests less than 75% total, with five or more omissions on the left than on the right) were considered for inclusion. Further inclusion criteria were an age of over 18 years, stroke onset of at least two weeks prior to testing, and the ability to give informed consent.

Exclusion criteria included concomitant illnesses that may affect the interpretation of any findings, a labile blood pressure following stroke, a systolic blood pressure less than 100 mmHg or diastolic blood pressure less than 70 mmHg, having

recently started using new antihypertensive medication within three weeks before testing, hepatic or renal dysfunction, neuroleptic medication, a brain tumour diagnosis, a weight of less than 55 kg, pregnancy or current breast feeding, severe coronary insufficiency or myocardial infarction in the six months prior to testing, and dysphasia, dementia, or any other cognitive or physical impairment that would prevent a patient from providing consent or performing standard clinical tests for neglect.

A total of 13 patients was included, ten of whom had frontal brain damage of different severities, and three of whom had no frontal damage (**Figure 8.1**). **Table 8.1** provides an overview of their ages and basic test scores at the time of inclusion.

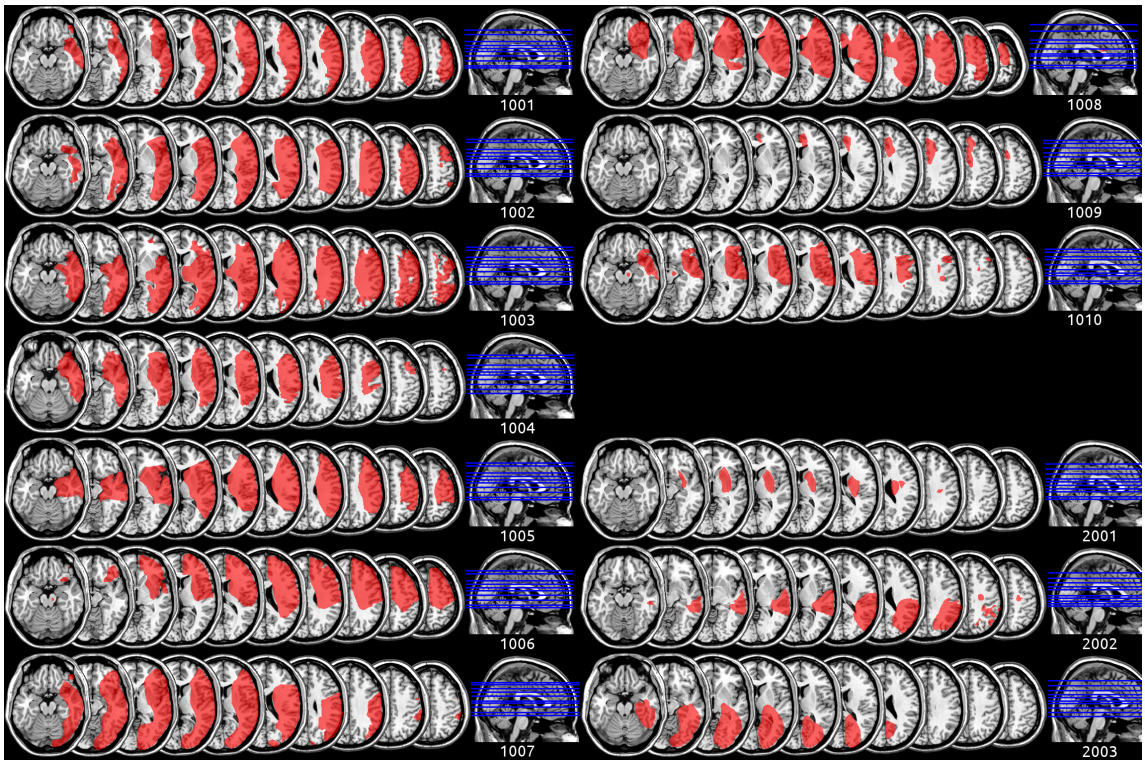

Figure 8.1 *– Hand-drawn lesions of patients in an structural MRI scan in standardised space. Patients 1001-1010 had frontal cortical involvement, patients 2001-2003 did not.*

Table 8.1. *Patient demographics.*

Patient	Age (years)	Time Since Stroke (months)	Cancellation Test Scores at Screening						Cancellation Test Scores at Time of Testing					
			BIT Star			Mesulam Shape			BIT Star			Mesulam Shape		
			Left	Right	Total	Left	Right	Total	Left	Right	Total	Left	Right	Total
1001	42	27	21	21	42	21	26	47	19	26	45	19	29	48
1002	66	3.25	0	12	12	0	11	11	0	17	17	0	10	10
1003	45	49	15	27	42	5	28	33	11	27	38	2	28	30
1004	58	14	0	14	14	0	13	13	0	18	18	0	14	14
1005	61	33.5	16	26	42	19	30	49	10	24	34	1	23	24
1006	63	2.75	14	27	41	16	15	31	26	27	53	2	29	31
1007	74	6	0	8	8	0	4	4	0	2	2	0	1	1
1008	64	1.25	8	21	29	1	26	27	25	26	51	11	24	35
1009	72	3	27	27	54	24	30	54	19	18	37	1	24	25
1010	74	6	27	27	54	16	28	44	27	24	51	7	28	35
2001	63	7	24	27	51	13	27	40	26	26	52	19	26	45
2002	75	6.5	20	26	46	8	18	26	21	14	35	8	12	20
2003	64	3.75	0	9	9	0	8	8	7	20	27	0	15	15

8.3.2. Guanfacine

Single 2 mg oral doses of guanfacine (brand name Estulic) were encapsulated by Nova Pharmaceuticals Ltd, who also provided the matching placebo preparation.

8.3.3. Procedure

Patients were tested on five consecutive days. On days one, three, and five, they were tested on a task battery that consisted of a computerised blind cancellation task (Parton et al., 2006), a continuous performance test (Riccio, Reynolds, & Lowe, 2001),

and a spatial working memory task. More details on individual tasks are provided below.

On days two and four, patients were tested on the same battery twice, once before oral administration of 2 mg of guanfacine or placebo, and once after. Whether placebo was administered on day two and guanfacine on day four or vice versa, was counter-balanced between patients.

8.3.4. Sustained attention task

Tests of sustained attention usually require participants to respond to a subset of sequentially presented stimuli, while ignoring others (Riccio et al., 2001; Robertson, Manly, Andrade, et al., 1997). In this study, a task was employed in which the targets were downward-pointing triangles (either red or green), and non-targets were upward-pointing red triangles. This task was taken from (Barceló, Suwazono, & Knight, 2000). Stimuli were presented for 1000 to 1500 ms, and were interleaved with inter-stimulus intervals of 1000-1500 ms. Patients had to press a button when a target was presented, and withhold a button press when no target was presented. In total, 320 stimuli were shown, of which 40 were green targets, 40 were red targets, and 240 were non-targets. The task lasted approximately 10 minutes.

Patients' accuracy and response times (only for correct responses) were computed, together with the proportion of hits, misses, false alarms, and correct rejections; as well as patients' response sensitivity (d') and bias (criterion, c) in terms of

signal-detection theory (Stanislaw & Todorov, 1999). In addition, patients' response time variability, which is commonly used as an index of sustained attention (with higher variability indicating poorer attention), was also computed. In order to track patients' sustained attention over the course of the test, their correct responses were examined in five bins. Each bin contained a minimum of five trials, and reaction time variability was calculated as the standard deviation of all response times within a bin.

8.3.5. Spatial working memory task

This test was taken from (Malhotra et al., 2005). In each trial, patients were shown a sequence of purple dots that could be presented at one of 10 different locations along the vertical midline of the computer screen (five above and five below a central fixation cross). After seeing a sequence, patients were presented with a probe display that contained nine black placeholders and a single purple dot, after which they were required to indicate whether the probed location was part of the sequence. Dot sequences varied in length from one to five stimuli, and became progressively longer during the course of the experiment, with an increase of one per ten trials. (Trials 1 – 10 were of length 1, 11-20 of length 2, …, and 41-50 of length 5.)

The average accuracy (proportion of correct responses) was computed to quantify performance.

8.3.6. Invisible cancellation task

A cancellation task requires participants to find and mark targets among distractors, usually leaving a visible 'cancellation' of each marked target. A computerised version of the task was employed that allowed patients to touch targets without visibly marking them as 'cancelled' (Dalmaijer, Van der Stigchel, Nijboer, Cornelissen, & Husain, 2015; Malhotra et al., 2006; Parton et al., 2006). This required patients to remember which targets they had already found.

Cancellation performance can be quantified to provide indices of patients' spatial attention and the quality of their search organisation. To examine patients' spatial attention, we computed the difference between the number of correctly cancelled targets on each half of the search array (right minus left), and the centre of cancellation (Binder et al., 1992). To examine patients' search organisation, we computed the best R (Mark et al., 2004), the intersections rate (Donnelly et al., 1999; Mark et al., 2004), the absolute and standardised inter-cancellation distance (Dalmaijer et al., 2015; Mark et al., 2004), and the standardised inter-cancellation angle (Dalmaijer et al., 2015). In addition, we computed the number of targets that patients revisited after the initial cancellation, and metrics of general performance such as the target processing speed (Rabuffetti et al., 2012b) and the total number of cancellations. These metrics are explained in greater detail in the previous chapter, and in (Dalmaijer et al., 2015).

8.3.6.1. Software

The aforementioned metrics were extracted using CancellationTools (Dalmaijer et al., 2015), which is standardised software that provides nearly all cancellation indices that are reported in the literature, and has been used in contemporary neuropsychological research (Smit et al., 2015; Ten Brink et al., 2016; Ten Brink, Van der Stigchel, Visser-Meily, & Nijboer, 2015).

Three different task versions were used that were matched in the number and spread of targets and distractors, to prevent learning of the target locations in each test. Each patient saw each version only twice, and never in direct succession. A strict time limit of two minutes was enforced on each task.

8.3.7. Bayesian data analysis

Baseline performance was determined for each patient by averaging the scores on all days, with the exception of the post-administration tests on day 2 and 4, which instead informed individual patients' performance on placebo and guanfacine. Group averages and differences were computed between treatment type (baseline, guanfacine, and placebo), across individuals.

To test whether there was an effect of treatment type, I employed repeated-measures ANOVAs. Drug was a factor in all analyses, with three levels: baseline, guanfacine, and placebo. For the sustained attention task, time bins were included as an additional factor, with five levels: one for each time bin. This allowed me to asses patients' performance over the course of the experiment. For the working memory task,

sequence length was included as an additional factor, with five levels: one for each sequence length (one to five stimuli).

Traditional (frequentist) repeated-measures ANOVAs produce p-values, which can inform one whether the null hypothesis should be rejected or not, but not how well it is supported by the data. To address this, I performed Bayesian repeated-measures ANOVAs, which produce a Bayes Factor (BF_{10}). This is the probability of the alternative hypothesis ("*guanfacine changes patients' performance*") divided by the probability of the null hypothesis ("*guanfacine does not change patients' performance*"). In essence, the Bayes Factor is a quantification of how much confidence one can have in either hypothesis. I interpret the results following the guidelines of (Jeffreys, 1961), which treat a Bayes Factor of 3 as roughly equivalent to a p-value of 0.05 in support of the alternative hypothesis. Conversely, a Bayes Factor of 1/3 would give equivalent support for the null hypothesis.

Data was handled in custom Python (Dalmaijer, 2017; Van Rossum & Drake, 2011) software, using the NumPy and SciPy libraries (Oliphant, 2007) for computations, and the Matplotlib library (Hunter, 2007) for plotting. All statistical analyses were performed in JASP, version 0.7.1.12 (JASP Team, 2016).

8.4. Results

8.4.1. Touchscreen Cancellation test – general performance

There was a significant main effect of drug on the total number of targets found, with a mean of five more targets cancelled on guanfacine compared to placebo, $F(2, 24) = 5.66$, $p = 0.010$, $\omega^2 = 0.26$, $BF_{10} = 4.926$. Post-hoc paired-sampled t-tests revealed a significant improvement in the total number of targets found between baseline (M = 28.4, SD = 13.91) and guanfacine (M = 31.15, SD = 15.09), $t(12) = -2.21$, $p = 0.047$, Cohen's d = -0.613, $BF_{10} = 1.687$; and between the placebo (M = 26.15, SD = 14.29) and guanfacine conditions, $t(12) = -2.93$, $p = 0.013$, Cohen's d = -0.813, $BF_{10} = 4.806$. Importantly, by contrast, there was no significant difference between baseline and placebo conditions, $t(12) = 1.52$, $p = 0.154$, $BF_{10} = 0.704$). These results provide moderate evidence of a significant effect of drug on search performance with the total number of targets found increasing by five on average on guanfacine compared to placebo (**Figure 8.2A**).

Cancellation speed (in pixels per millisecond) was computed as the average of the point-wise division of inter-cancellation distance (in pixels) by the average inter-cancellation time (in milliseconds), with higher numbers being better. A repeated-measures ANOVA revealed no main effect of drug on cancellation speed, $F(2, 24) = 0.06$, $p = 0.945$, $BF_{10} = 0.188$; and a direct comparison of guanfacine (M = 0.10, SD = 0.03) with placebo (M = 0.10, SD = 0.3) revealed no difference in cancellation speed, $t(12) = -0.262$, $p = 0.798$, $BF_{10} = 0.287$.

These results provide moderate evidence for there being no effect of drug on cancellation speed, and no difference between guanfacine and placebo treatments.

8.4.2. Touchscreen Cancellation test – directional bias

Independent repeated-measures ANOVAs revealed no main effect of drug on the number of cancellations on the left side of the cancellation task, $F_{greenhouse\text{-}Geisser}(1.27, 15.20) = 2.02$, $p = 0.175$, $BF_{10} = 0.654$; no main effect of drug on the number of cancellations on the right side of the cancellation task, $F(2, 24) = 2.44$, $p = 0.109$, $BF_{10} = 0.844$; no main effect of drug on the difference between the cancellations on the right and left sides of the cancellation task (**Figure 8.2B**), $F(2, 24) = 0.39$, $p = 0.683$, $BF_{10} = 0.231$, and no main effect of drug on the centre of cancellation (**Figure 8.2C**), $F(2, 24) = 2.45$, $p = 0.108$, $BF_{10} = 0.848$. A direct comparison of the difference in right and left cancellations in the placebo (M = 12.0, SD = 7.99) and guanfacine (M = 11.92, SD = 11.36) conditions revealed no difference between the two, $t(12) = 0.03$, $p = 0.976$, $BF_{10} = 0.278$.

In sum, these results provide no evidence of an effect of drug on the centre of cancellation, or on the amount of cancellations on the left or the right side independently. However, they also do not provide evidence for the absence of an effect of drug. There does seem to be moderate evidence for there not being an effect of drug on spatial bias defined by the difference in cancellations on each side of the cancellation task (right minus left), specifically there is moderate evidence for guanfacine and placebo having the same effect.

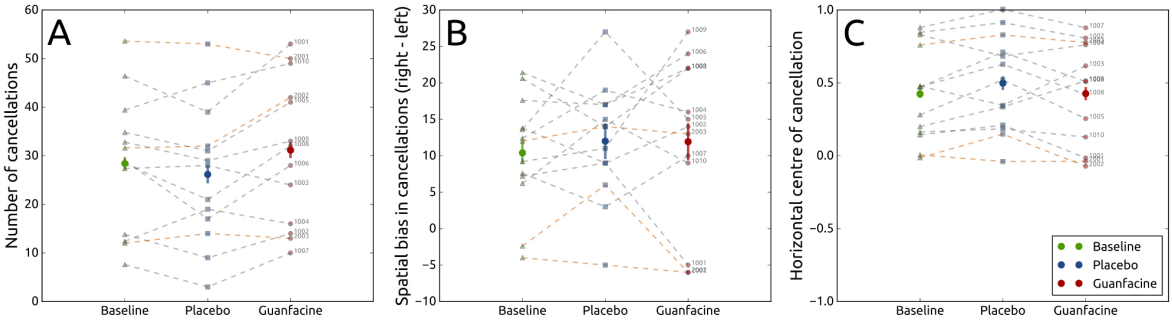

Figure 8.2 *– Patients' test scores during baseline (green) measurements, on placebo (blue), and on guanfacine (red). Large filled circles represent group averages, with error bars representing the within-subjects standard error of the mean. Dashed lines represent individual participants. Grey dashed lines represent patients with an extent of frontal cortical damage, and orange dashed lines without. Patient numbers that correspond with the lesion maps in Figure 8.1 are annotated.* ***A)*** *Total number of cancelled targets (maximum is 64).* ***B)*** *Difference in the number of cancelled targets on the right and on the left halves of the task.* ***C)*** *Average position on the horizontal axis of the cancellation test, 'centre of cancellation' (-1 is leftmost target, 1 is rightmost target).*

8.4.3. Touchscreen Cancellation test – revisits

Independent repeated-measures ANOVAs revealed no main effect of drug on the total number of revisits in this 'invisible' cancellation task where items already cancelled are not left marked on the screen (**Figure 8.3A**), $F(2, 24) = 1.43$, $p = 0.259$, $BF_{10} = 0.451$; no main effect of drug on the number of immediate revisits (perseverations; **Figure 8.3B**), $F(2, 24) = 1.15$, $p = 0.334$, $BF_{10} = 0.387$; and no main

effect of drug on the number of delayed revisits (**Figure 8.3C**), $F(2, 24) = 2.37$, $p = 0.115$, $BF_{10} = 0.805$. These results provide no conclusive evidence on whether drug had an effect on revisits.

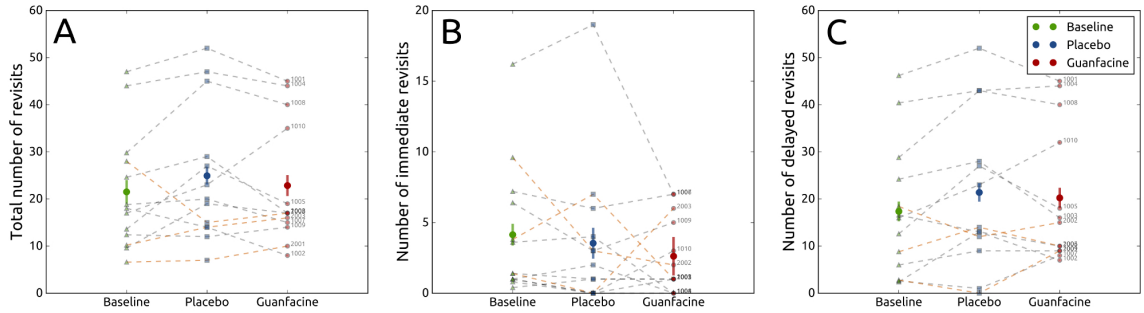

Figure 8.3 *– Patients' test scores during baseline (green) measurements, on placebo (blue), and on guanfacine (red). Large filled circles represent group averages, with error bars representing the within-subjects standard error of the mean. Dashed lines represent individual participants. Grey dashed lines represent patients with an extent of frontal cortical damage, and orange dashed lines without. Patient numbers that correspond with the lesion maps in Figure 8.1 are annotated.* ***A)*** *Total number of revisits.* ***B)*** *Number of immediate revisits, or 'perseverations'.* ***C)*** *Number of delayed revisits.*

8.4.4. Touchscreen Cancellation test – search organisation

The best R was computed as the highest correlation between cancellation rank number, and either the corresponding vertical or the horizontal cancellation coordinate. Potential values range between 0 (not organised) and 1 (organised). A repeated-measures ANOVA found no main effect of drug on the best R, $F(2, 24) = 0.37$, $p = 0.694$, $BF_{10} = 0.235$; and a paired-samples t-test found no difference in the best R between guanfacine ($M = 0.51$, $SD = 0.29$) and placebo ($M = 0.49$, $SD = 0.26$), $t(12) = 0.31$, $p = 0.760$, $BF_{10} = 0.290$. These results provide moderate evidence that drug had no effect on the best R, and that there was no difference between guanfacine and placebo (**Figure 8.4A**).

The intersection rate was computed as the number of times a participant's search path intersected with itself, divided by the total number of cancellations that are not immediate revisits. Independent repeated-measures ANOVAs revealed no main effect of drug on the total number of search path intersections, $F_{Greenhouse\text{-}Geisser}(1.41, 16.86) = 0.67$, $p = 0.475$, $BF_{10} = 0.280$; and no main effect of drug on the intersection rate, $F_{Greenhouse\text{-}Geisser}(1.18, 14.10) = 0.50$, $p = 0.521$, $BF_{10} = 0.253$. Direct comparisons (paired-samples t-tests) between guanfacine and placebo revealed no difference in the total number of intersections (guanfacine: $M = 18.77$, $SD = 24.67$; placebo: $M = 17.08$, $SD = 16.71$), $t(12) = 0.35$, $p = 0.733$, $BF_{10} = 0.294$; nor in the intersections rate (guanfacine: $M = 0.37$, $SD = 0.51$; placebo: $M = 0.30$, $SD = 0.30$), $t(12) = 0.55$, $p = 0.596$, $BF_{10} = 0.317$. These results provide moderate evidence that drug had no effect on the intersections total (**Figure 8.4B**) and rate (**Figure 8.4C**), and that there was no difference between guanfacine and placebo.

The inter-cancellation distance was computed as the average distance (in pixels) between each cancellation and the next. The standardised inter-cancellation distance was defined as the average inter-cancellation distance in pixels, divided by the average distance between each target its nearest neighbouring target (which makes this metric comparable between different cancellation tasks). Independent repeated-measures ANOVAs revealed no main effect of drug on the average inter-cancellation distance, $F_r(2, 24) = 0.86$, $p = 0.435$, $BF_{10} = 0.316$; and no main effect of drug on the standardised inter-cancellation distance, $F(2, 24) = 0.48$, $p = 0.624$, $BF_{10} = 0.253$. Direct comparisons (paired-samples t-tests) between guanfacine and placebo revealed no difference in the average inter-cancellation distance (guanfacine: M = 218.4, SD = 43.58; placebo: M = 224.6, SD = 28.79), $t(12) = -0.71$, $p = 0.493$, $BF_{10} = 0.345$; nor in the standardised inter-cancellation distance (guanfacine: M = 2.45, SD = 0.47; placebo: M = 2.52, SD = 0.30), $t(12) = -0.74$, $p = 0.472$, $BF_{10} = 0.353$. These results provide moderate evidence that drug had no effect on the average (**Figure 8.4D**) and standardised (**Figure 8.4E**) inter-cancellation distances, and that there was no difference between guanfacine and placebo.

The standardised angle is defined as the angle between consecutive cancellations, standardised so that a value of 1 corresponds with a completely cardinal angle, and a value of 0 with a completely diagonal angle. A repeated-measures ANOVA found no main effect of drug on the standardised angle, $F(2, 24) = 0.07$, $p = 0.937$, $BF_{10} = 0.187$; and a paired-samples t-test found no difference in the standardised angle between guanfacine (M = 0.66, SD = 0.075) and placebo (M = 0.65, SD = 0.11), $t(12) = 0.05$, $p = 0.959$, $BF_{10} = 0.279$. These result provides moderate evidence that drug had no

effect of the standardised angle, and that there was no difference between guanfacine and placebo (**Figure 8.4F**).

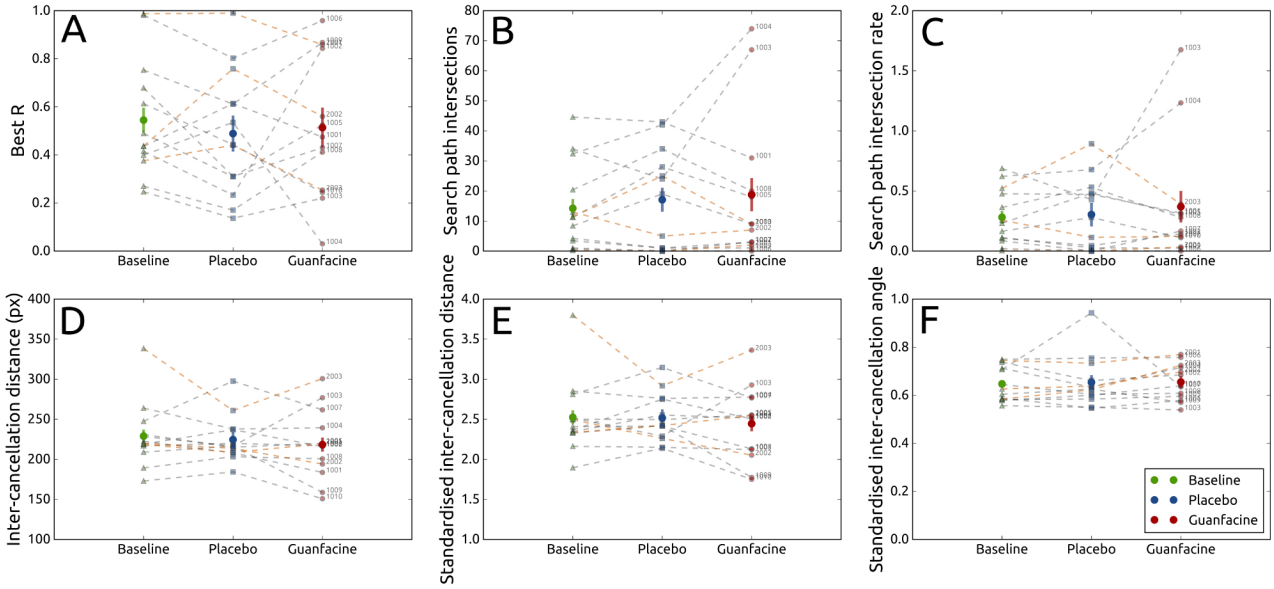

Figure 8.4 *– Patients' test scores during baseline (green) measurements, on placebo (blue), and on guanfacine (red). Large filled circles represent group averages, with error bars representing the within-subjects standard error of the mean. Dashed lines represent individual participants. Grey dashed lines represent patients with an extent of frontal cortical damage, and orange dashed lines without. Patient numbers that correspond with the lesion maps in Figure 8.1 are annotated.* ***A)*** *Correlation between cancellation rank order and either horizontal or vertical coordinate, best R.* ***B)*** *Number of search path intersections.* ***C)*** *Search path intersection rate.* ***D)*** *Average inter-cancellation distance.* ***E)*** *Standardised inter-cancellation distance.* ***F)*** *Standardised inter-cancellation angle.*

8.4.5. Sustained attention test – signal detection

Independent repeated-measures ANOVAs revealed no main effect of drug on hit rate (**Figure 8.5A**), $F(2, 24) = 2.84$, $p = 0.078$, $BF_{10} = 1.061$; no main effect of drug on false alarm rate (**Figure 8.5B**), $F_{Greenhouse\text{-}Geisser}(1.12, 13.39) = 1.51$, $p = 0.243$, $BF_{10} = 0.474$; no main effect of drug on sensitivity (d'; **Figure 8.5C**), $F(2, 24) = 1.05$, $p = 0.367$, $BF_{10} = 0.359$; and a main effect of drug on response bias (c; **Figure 8.5D**), $F(2, 24) = 5.31$, $p = 0.012$, $\omega^2 = 0.24$, $BF_{10} = 4.199$. Post-hoc paired-samples t-tests indicated that there was no difference in response bias (c) between the baseline (M = 0.69, SD = 0.44) and guanfacine (M = 0.72, SD = 0.42) condition, $t(12) = -0.47$, $p = 0.650$, $BF_{10} = 0.306$; but that there was a difference in response bias (c) between the baseline and placebo (M = 0.51, SD = 0.45) conditions, $t(12) = 2.73$, $p = 0.018$, Cohen's d = 0.76, $BF_{10} = 3.544$; as well as between the guanfacine and placebo conditions, $t(12) = 3.51$, $p = 0.004$, Cohen's d = 0.973, $BF_{10} = 11.555$.

These results provide no conclusive evidence of a presence or absence of an effect of drug on hit rate, false alarm rate and sensitivity (d'). The results do provide moderate evidence that placebo reduced the response bias (c) to a less conservative value compared to baseline (values of c between 0 and 1 indicate a bias to not responding), and that guanfacine had no effect on response bias (compared to baseline).

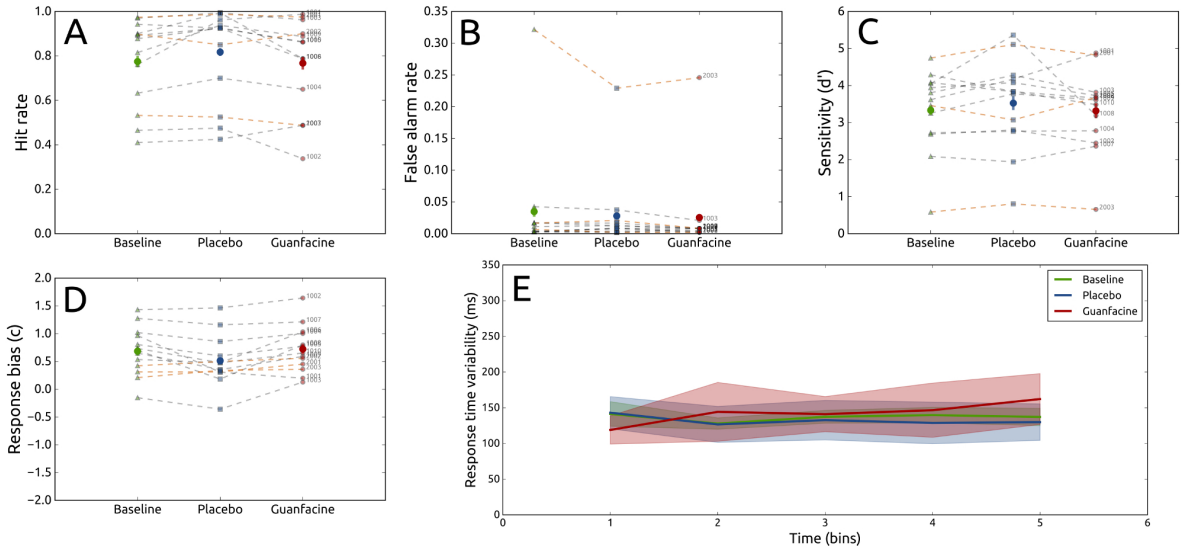

Figure 8.5 *– Patients' test scores during baseline (green) measurements, on placebo (blue), and on guanfacine (red). Large filled circles represent group averages, with error bars representing the within-subjects standard error of the mean. Dashed lines represent individual participants. Grey dashed lines represent patients with an extent of frontal cortical damage, and orange dashed lines without. Patient numbers that correspond with the lesion maps in Figure 8.1 are annotated.* ***A)*** *Hit rate in sustained attention task.* ***B)*** *False alarm rate in sustained attention task.* ***C)*** *Sensitivity (d') in sustained attention task.* ***D)*** *Response bias (c) in sustained attention task.* ***E)*** *Response time variability in milliseconds (y-axis) as a function of time bin during sustained attention task. Shaded areas indicate within-subject standard error of the mean.*

8.4.6. Sustained attention test – reaction time variability

A repeated-measures ANOVA revealed no main effect of drug on reaction time variability, $F_{Greenhouse\text{-}Geisser}(1.38, 16.61) = 0.96$, $p = 0.371$, $BF_{10} = 0.100$; nor a main effect of time, $F(4, 48) = 0.20$, $p = 0.939$, $BF_{10} = 0.022$. There was also no interaction effect between drug and time, $F(8, 96) = 0.72$, $p = 0.673$. These results provide moderate to strong evidence that there was no effect of drug or time on reaction time variability in the sustained attention task (**Figure 8.5E**).

8.4.7. Spatial working memory test

A repeated-measures ANOVA revealed no main effect of drug on response accuracy (**Figure 8.6**), $F(2, 24) = 0.84$, $p = 0.446$, $BF_{10} = 0.082$; a main effect of sequence length, $F_{Greenhouse\text{-}Geisser}(2.54, 30.42) = 6.39$, $p = 0.003$, $\omega^2 = 0.30$, $BF_{10} = 28919.401$; and no interaction effect, $F(8, 96) = 1.42$, $p = 0.200$. These results provide strong evidence that there was no effect of drug on response accuracy, and decisive evidence that there was an effect of sequence length on response accuracy (with worse accuracy for longer sequences).

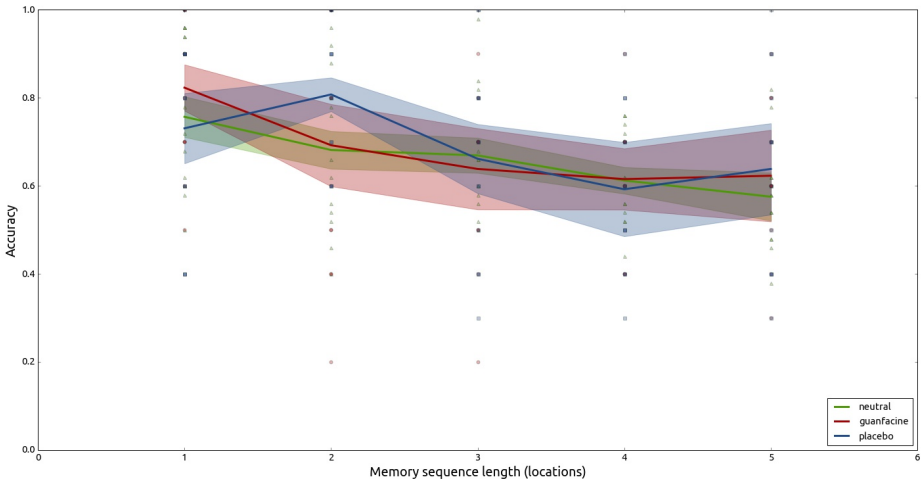

Figure 8.6 *– Patients' test scores during baseline (green) measurements, on placebo (blue), and on guanfacine (red). Solid lines represent group averages, with shaded areas representing the within-subjects standard error of the mean. Triangles (baseline), squares (placebo), and circles (guanfacine) represent individual participants.*

Post-hoc paired-sampled t-tests indicated that there was no difference between baseline and guanfacine for each sequence length (all $p > 0.05$; BF_{10} for sequence lengths one, three, four and five ranged between 0.358 and 0.554, and BF_{10} for sequence length 2 was 1.533), nor any difference between placebo and guanfacine for each sequence length (all $p > 0.05$; BF_{10} ranged between 0.279 and 0.328). For sequence lengths one to four, there was no difference between baseline and placebo (all $p > 0.05$, BF_{10} were 1.183, 1.116, 0.464 and 0.278 for sequence lengths one to four, respectively). In sum, there was a reliable absence of an effect of treatment type on accuracy in the spatial working memory task.

8.5. Discussion

This study was designed to investigate whether a single dose of guanfacine, an agonist of the alpha-2A adrenoreceptor, has beneficial effects for patients who suffer from hemispatial visual neglect. By supplementing frequentist with Bayesian statistics, I aimed to establish whether any null-effects were due to guanfacine not being different from placebo, or simply due to a lack of statistical power.

The results indicated that although there was moderate evidence for guanfacine increasing the total number of targets found in a cancellation task, there was no conclusive evidence on whether a single 2 mg dose of guanfacine improved spatial bias or revisits. There was moderate evidence for the absence of an effect of guanfacine on search organisation (operationalised with four different indices). In addition, there was moderate to strong evidence that guanfacine does not improve sustained attention (as measured by response time variability), but no conclusive evidence of whether it affects patients' signal detection or not. Finally, there was strong evidence that guanfacine does not improve spatial working memory.

The current investigation was inspired by Malhotra et al. (2006), who tested three neglect patients on an invisible cancellation task like the one employed here, and reported a beneficial effect of guanfacine (compared to placebo) for two patients. The improvement was apparent as a larger number of found targets, and an *increased time on task*. These findings could mean that guanfacine had improved both the spatial and the sustained attention of neglect patients.

In the current larger sample of thirteen neglect patients, the only finding of Malhorta and colleagues that is replicated, is an increase in the number of cancelled targets when patients were on guanfacine compared to placebo. However, there was no statistical evidence for or against an improvement of spatial attention on guanfacine compared to placebo. A larger study with more statistical power could address this in the future. In addition, it would be useful to investigate more prolonged administration of regular guanfacine over several weeks rather than a single dose, as well as the effects of different doses of the drug.

One alternative interpretation of the findings of Malhotra and colleagues' findings is that guanfacine merely boosted sustained attention, which caused patients to work on the task longer, thereby increasing the chance that they found more targets. This is partly supported by Malhotra and colleagues, who report that one patient improved on a separate sustained attention test. In the results reported here, there was moderate statistical evidence for there not being an effect of treatment type on sustained attention on a 10-minute continuous performance test (Riccio et al., 2001). However, it should be noted that in the data presented here, sustained attention as indexed by response time variability was not affected by task duration. This means patients were at the same level of sustained attention throughout the sustained attention task, which suggests that perhaps the task was too short or too engaging. In this light, it might still be possible that guanfacine improved patients' sustained attention on the cancellation test (as evidenced by an increased number of cancelled targets), but that this could not be measured on the sustained attention test due to aforementioned ceiling effect.

Another alternative explanation for the results reported by Malhotra and

colleagues is that guanfacine boosted patients' working memory, which caused them to be more likely to remember what targets they had already cancelled. This could potentially have resulted in patients cancelling more targets. Although other work provides support for the beneficial effects of guanfacine on working memory in healthy humans (Jäkälä et al., 1999) and in monkeys (Arnsten et al., 1988), there was strong statistical evidence that treatment type did not affect patients' response accuracy on the spatial working memory task used here.

In addition to visual neglect, stroke patients can suffer from disorganised search (Ten Brink et al., 2015). Although it has been argued that neglect and search organisation are independent disorders (Mark et al., 2004), more recent work on large samples indicates that stroke patients with neglect are more likely to also suffer from disorganised search (Rabuffetti et al., 2012b; Ten Brink et al., 2015), and some have even argued that disorganised search is a consequence of disturbed spatial attention (Ten Brink et al., 2016). For these reasons, and because visual search is important in daily life, I thought it important to test whether guanfacine improved visual search organisation in our sample. Across four different indices of search organisation, we found moderate statistical evidence that guanfacine administration did not result in any improvements.

In summary, the results described here are of a randomised, double-blind, placebo-controlled, crossover study of 2 mg guanfacine, a noradrenergic agonist that is thought to work primarily through frontal alpha-2 receptors. Thirteen stroke patients with hemispatial neglect were tested on a cancellation task with invisible markings, a sustained attention task, and a spatial working memory task. Analysis revealed an

increase in the total number of targets that patients correctly cancelled on a cancellation task following a single dose of guanfacine, but there was no statistically robust evidence for or against beneficial effects of guanfacine on patients' spatial attention bias. In addition, there was robust evidence that guanfacine did not improve search organisation on an invisible cancellation task. Furthermore, guanfacine did not improve sustained attention in a continuous performance task, nor that it improved performance on a spatial working memory task.

9. Discussion

In this thesis, I attempted to investigate four main research questions. The first two specifically related to attentional processes that act on short-term memory: *How is information encoded into short-term memory?* and *How are short-term memory resources re-distributed to accommodate for prioritised information?* The last questions related to the effects of post-stroke disorders on attention and short-term memory: *How can deficits in attention and short-term memory be accurately assessed, and can they be ameliorated in hemispatial neglect syndrome by pharmacological intervention?*

The short answers are that my investigations suggest that: **1)** Information is encoded into short-term memory in a parallel fashion, with multiple items being attentionally selected and encoded at the same time (*Chapter 4*). However, items can also be selected and encoded in a serial fashion when there is a reason to prioritise, for example when one item is associated with reward (*Chapter 5*). **2)** It is hard to say how, or even whether short-term memory resources are re-allocated, as my experiments into re-allocation did not show the expected hallmarks of re-allocation (*Chapter 6*). **3)** Spatial attention, search organisation, and even short-term memory can be measured in a highly sensitive and patient-friendly way by means of 'invisible' cancellation of targets among distractors (*Chapter 7*). **4)** Performance on such cancellation tasks can be improved in neglect patients using even a single dose of guanfacine, but longer-term medication trials are necessary to determine whether such effects translate to clinical benefits (*Chapter 8*).

9.1. Thinking about attention

Moving beyond the headlines, my results and the presented interpretations raised a question: What is attention? Although “everybody knows what attention is” (James, 1890), it is also a catch-all term employed liberally to explain (or rather label) different and sometimes seemingly unrelated phenomena. Case in point: I referred to attentional selection to define the process of boosting perceptual information from particular locations and features, but also claimed that encoding and re-allocating short-term memory resources are attentional processes.

Arguably still the best taxonomy of attention is by Posner and Boies, who divided attention into three components: alertness (sustained attention), selectivity (spatial attention), and processing capacity (Posner & Boies, 1971).

This way of sub-dividing attention fits with work discussed in Chapter 3: *Introduction*. For example, in Bundesen’s theory of visual attention, selectivity is implemented at the level of increased sensitivity for particular features (locations, colours, orientations, etc.), whereas the encoding of information into short-term memory is governed by a processing capacity parameter (Bundesen, 1990). That these separate aspects of attention might relate to each other, was demonstrated by Zokaei and colleagues who showed that increasing the attentional demands of a task performed during the maintenance of information in short-term memory, results in a decrease in recall performance (Zokaei, Heider, et al., 2014). So perhaps it is inevitable that ‘attention’ is such a catch-all term: Empirical evidence suggests that the same attentional resource is involved in qualitatively different processes at the same time.

In Chapter 4: *Encoding of information into visual short-term memory is a parallel process*, four experiments were discussed. In these experiments, the exposure duration of 1 or 2 items was varied, and recall precision was measured for all presented items. The raw data already illustrated that the encoding of the two items seemed to have occurred independently: Recall error for one item did not predict recall error for the other. In addition, both items were encoded at a reasonable precision during exposure times of 200 milliseconds or less, despite such exposure durations being too short to attend to one item, process it, and move on to another (Bundesen, 1990; Duncan et al., 1994; Reeves & Sperling, 1986; Weichselgartner & Sperling, 1987).

More convincingly, a computational model estimated the probabilities of two items, no items, or only one item being remembered. The modelling showed that the probability of only one item being remembered was exactly in line with the parallel hypothesis. These results were in line with each item being encoded by a stochastically operating channel to encode with a particular likelihood. The alternative would have been a single such channel operating on only one item at a time.

These parallel channels are roughly in line with the parallel attentional ‘races’ of Bundesen, each processing a unique item, and the fastest ones ending up in short-term memory (Bundesen, 1990). When interpreted in a resource-based short-term memory model, this suggests that short-term memory resources are allocated to items from rather early on in the perceptual/attentional processing. Crucially, this conceptualisation suggests that every attended item is, to a limited extent, encoded into short-term memory.

In addition, the above means that lines between cognitive domains are blurry.

Where does perception stop and attention start, if attention *is* the boosting of information from particular neural populations with their preferred feature or receptive field in the attended category or space? And where does attention stop and short-term memory start if the process of attending to an item inherently (partly) encodes it into short-term memory? If anything, attention seems to be a process: It is the shifting of parameters so that some incoming signals are amplified and some muted, but also so that some signals are internally repeated and some faded out. In Baddeley and Hitch's (Baddeley & Hitch, 1974) box-and-arrow terms, attention *is* the central executive.

9.2. What about feature-integration theory

Feature-integration theory proposed that items need to be serially attended to in order to bind their individual features (e.g. location and colour) together, and that without attention features are not bound to individual items (Treisman, 1982; Treisman & Gelade, 1980). This popular theory is at the core of many models of how visual information is attentionally selected and committed to short-term memory, all of which propose an initial pre-attentive perceptual sweep across the visual field, followed by a second phase during which each item's location is attended to so that its features can be bound together (Cave & Wolfe, 1990; Duncan, 1984; Hoffman, 1979; Wheeler & Treisman, 2002; Wolfe, 1994; Wolfe & Cave, 1999). Contrary to what feature-integration theory would predict, evidence of pre-attentive binding has previously been reported (Houck & Hoffman, 1986). The results presented in Chapter 4: *Encoding of information into visual short-term memory is a parallel process* add direct empirical

evidence of parallel binding and encoding to the existing literature.

My results indicated that features (orientation and colour) can be bound to the locations of at least two items at the same time. These results are robust: They were replicated in three experiments, using coloured disks and oriented gratings as stimuli, and presenting stimuli uni- and bi-laterally. In addition, the parallel encoding results were supported by highly balanced electroencephalographic signals (ERP and oscillatory power) over both cortices when processing two simultaneously presented items, with no sign of temporal offsets (which would have been a clear indication of serial processing).

This is fundamentally incompatible with a strict reading of feature-integration theory. However, my findings are compatible with a feature-integration theory that allows for multiple attentional loci. Specifically, if attention can be oriented to two or more locations at the same time, features from multiple locations could be bound into items in parallel. This idea is supported by empirical evidence that demonstrates that individuals can attend to multiple non-contiguous locations without being distracted by information presented in between those locations (Hahn & Kramer, 1998; Kramer & Hahn, 1995; McMains & Somers, 2004; Müller et al., 2003). In sum, I support a version of feature-integration theory in which attention can be divided over several items to bind each independently of the others.

9.3. The short-term memory encoding mechanism

The mechanism envisaged for encoding information into short-term memory has already been outlined above and in earlier chapters. In summary, attention amplifies (selects) particular subsets of visual information. This could be items at specific locations, but it could also be items that share other features (i.e. items of the same colour). A pool of short-term memory resources is available to distribute over a limited number of items. This limit is not the magical number 4, but rather a practical limit: Resources can be spread out over items, but an item requires a minimal amount of resource. Otherwise that item's representational quality is too noisy, and it will disappear from short-term memory. The sub-division of short-term memory resources over items determines the encoding rate of each item, with more allocated resources meaning a quicker encoding process.

It is possible to bias the allocation of resources among items to only one or a selected number of items, for example because those items are behaviourally relevant. In Chapter 5: *Serialisation of visual short-term memory encoding in the presence of reward*, it was demonstrated that short-memory memory encoding can occur serially if two items are presented simultaneously, but only one of them is associated with a recall-performance-based reward. Thus, short-memory memory resources were first allocated to the reward-associated item, and then to the non-rewarded item (which was still recalled in a non-trivial proportion of trials, albeit at lower precision).

The use of the term 'short-term memory resources' here is confusing, as it might entail a resource devoted to the encoding of information, or a resource devoted to the

storage of information. In my view, the resource required to encode information into short-term memory is an attentional one, and the ‘space’ in short-term memory is best referred to as a short-term memory resource. Note that this distinction is semantic, and that in the proposed framework encoding and continued storage of information in short-term memory are related-but-independent processes.

9.4. Future investigations of parallel encoding

Another experimental technique for probing the short-term memory encoding process could be to present two items with different onsets, but the same offset. If items are encoded in series, onset-asynchronies of less than the encoding duration of a single item should not affect performance, because the second item would not be encoded during the encoding of the first item. Only when the onset-asynchrony surpasses the encoding duration of one item would a detriment in recall performance compared to simultaneous presentation of both items become apparent. However, if items are indeed encoded in parallel, a detriment in recall error would occur with any onset-asynchrony, because the second item would otherwise be encoded.

It should be noted that the experiments presented in Chapter 4 are arguably better than the experiment outlined above, because of one flaw in the outlined design. Crucially, during the onset delay before the second item, only one item is present. In a parallel encoding scenario, it would draw all resources, thereby improving encoding of the first item compared to a simultaneous presentation. This initial boost in encoding might counter-act the detrimental effect of the second item not being encoded from first

onset, and therefore the results might not be able to distinguish between parallel and serial encoding after all.

Another interesting experiment would be to present items with more than two defining features. In the experiments presented in Chapters 4 and 5, items were defined by their location and either their colour or their orientation. Future experiments could present coloured Gabor stimuli, so that items would be defined by their location, colour, and orientation. This would allow an analysis of the time-course of feature binding, as recall of all features could be probed independently.

In fact, I ran a pilot of this study, and found one crucial caveat: The employed mask has to mask both the colour and the orientation of the coloured Gabors. In my pilot, the mask only effectively masked colour, and thus orientation was near-perfect even at very low (10-20 ms) exposure durations.

If the masks were effective, I would have hoped to see dependent performance within each item, but independent performance between items. The former would evidence an item-benefit and thus evidence for binding, and the latter would evidence independent and parallel encoding of both items. An additional challenge would be the computational model presented in Chapter 4, as it would have to account for two features per item rather than one.

Finally, it would be worthwhile to repeat the experiment presented in Chapter 5, but with a successful neutral manipulation. In the current data, an attempt at a neutral condition was made by removing the monetary reward from the feedback provided to participants. However, participants seemed to perceive feedback on their performance as rewarding even without a monetary consequence. The feedback element was retained

in the neutral condition to tease out the separate effects of monetary reward and performance feedback, but perhaps should be done away with in future studies.

9.5. Re-allocation of short-term memory resources

It has previously been argued that the maintenance of information in visual short-term memory requires attention (Zokaei, Heider, et al., 2014). This could be viewed as information needing constant re-encoding (or rehearsal) due to degradation over time, for example due to signal drift (Wimmer et al., 2014). It has also previously been argued that short-term memory resources can be re-allocated to new or stored information when this becomes behaviourally relevant (Bays & Husain, 2008).

In the cited study by Bays and Husain, the method of choice to make items behaviourally relevant was eye movements. Participants were required to fixate all four items presented on an imaginary circle around the fifth item, and then to fixate the fifth item. However, the screen was blanked during participants' saccade to the fifth item, and thus it was never directly fixated. Despite it not being fixated, recall accuracy was much higher for the fifth item. (Recall precision was computed as the number of trials in which participants accurately identified a clockwise or counter-clockwise rotation of a randomly selected and re-shown item from the memorised array.) This result was taken as evidence that short-term memory resources had been redistributed to the saccade target, so that the eye could land as accurately as possible. (Bays & Husain, 2008)

Chapter 6: *Dynamic re-allocation of visual short-term memory resources to saccade targets?* presented three experiments that aimed to replicate (and potentially

extend) the reported findings. Unfortunately, using continuous report measures rather than change detection, the results from Bays and Husain could not be replicated. In addition, when using change detection measures in experiment 3b, the results from Bays and Husain could also not be replicated.

This failed replication is somewhat puzzling, as the tasks were increasingly closely matched to the task employed by Bays and Husain in their experiment 2. Subtle differences exist between the experiments that I conducted and the one reported by Bays and Husain, though. These include the response: Bays and Husain employed change detection, whereas I asked participants to reproduce the probed stimulus using a continuous report. It should be noted that I used a game pad joystick for continuous responses, which was as quick as a single button press, and thus did not introduce additional decay time. In addition, when I did employ the exact same change detection paradigm as Bays and Husain, their results still did not replicate.

Another subtle difference is the placement of items. In my experiments all stimuli were positioned on an imaginary circle, whereas in Bays and Husain only four items were placed on an imaginary circle and the final item was in the centre. My experiments allowed participants more freedom in the order of their fixations, and made it so that the final item's position to the penultimate item was not the same across all trials (as it ways in Bays and Husain's experiment).

A possible alternative explanation for the results presented by Bays and Husain is that a pre-saccadic shift of attention (Rolfs et al., 2011) to the to-be-fixated item occurred. Perhaps the subtle differences between the experiments that I presented and that of Bays and Husain prevented this predictive remapping of attention to occur (for

example due to the non-standardised distance to the last item). Alternatively, maybe Bays and Husain were correct in their interpretation of short-term memory resource re-allocation, but perhaps this is not a universally occurring phenomenon, but instead restricted to specific situations. Perhaps it requires the fixation sequence to be fixed, such as in the experiment presented by Bays and Husain.

A potential way forward is to map when re-allocation of short-term memory resources does occur, and what the crucial parameters are. One alteration of the experiments I presented would be to introduce an unexpected blink that could occur during any saccade in the task, instead of only during the saccade from penultimate to last item. Alternatively, behavioural relevance can be manipulated using different methods. For example, it has been shown that previous reward associations (Della Libera & Chelazzi, 2006; Donohue et al., 2016; Failing & Theeuwes, 2017; Klink et al., 2017) and fear conditioning (Mulckhuyse et al., 2013; Mulckhuyse & Dalmaijer, 2016) affect the allocation of short-term memory resources or attention even when rewards and threat are task-irrelevant (or even absent from the current task). These methods could also be applied to drive the hypothesised dynamic re-allocation of short-term memory resources.

In addition, a direct replication of experiment 2 from Bays and Husain (2008) would be a welcome addition to the literature. This should be a complete and exact replication of the original study, with no such deviations as in the replication attempts from Chapter 6.

9.6. Measurement of (in)attention in patients

The measurement of cognitive deficits is traditionally done using pen-and-paper tests, or clinical observations, oftentimes on a rather ad-hoc basis. Anecdotal evidence for this claim comes from a former mentor of mine who, when testing a patient during neurosurgery, quickly drew a central fixation cross and four shapes in the corners of a blank sheet of paper (the back of a pre-printed test) to investigate whether the patient was exhibiting neglect (M. van Zandvoort, personal communication). A plethora of tests exists, and sometimes one can get the feeling that neuropsychologists quite regularly generate new tests on a whim, without much standardised testing.

Of course, in reality, a plethora of well-validated tests exists. However, the majority is still pen-and-paper, and are thus potentially at risk of missing important data that might be missed using only the final results. Computerised tests allow for easier standardisation of test parameters (e.g. retinal size of stimuli), and allows for much richer data collection. It should be noted that I am not arguing for replacing neuropsychologists with computers: It is hard to overstate the added value of personal patient care, and of observations by an expert during diagnostic sessions.

What I would argue for is the automation of some of the pen-and-paper tests that patients are subjected to. Specifically, in this thesis I have highlighted the cancellation task as a perfect example for computerised testing. In cancellation tasks, patients search for targets among distractors, and cross them out wherever they find them. Historically, this test has been used to diagnose hemispatial neglect disorder, usually by counting the number of omitted targets. However, on a computer screen, one can track which targets

were cancelled when. This provides researchers and clinicians with the possibility of reconstructing the search path that a patient followed. (Dalmaijer et al., 2015)

In Chapter 7: *Using cancellation tasks to assess spatial attention, short-term memory, and executive functioning* I reviewed all the metrics that have been computed from cancellation tasks in the past. In addition, I introduced CancellationTools, a software package that combines the options to conduct cancellation tasks, and to analyse their results with a few button clicks. The metrics I described in the chapter and included in the software include metrics that index biases in spatial attention, sustained attention, short-term memory, and executive functioning.

Importantly, I presented data from over 500 participants that have done an online cancellation test. Using this data, I demonstrated that demographic factors do not impact cancellation performance, thereby illustrating that cancellation tests are robust against age and level of education. Furthermore, I employed sophisticated machine-learning algorithms, which showed that the general population can be divided into four groups: Two small groups of around twenty participants, and two larger groups of well over one hundred participants each. The small groups were characterised by many omissions, and by many revisits respectively. These could have reflected genuine cognitive problems, but could also have been a product of not understanding the task. Unfortunately, a limitation of internet-based mass testing is that it is difficult to monitor and debrief participants, so it is hard to know whether they suffered from cognitive issues or were just clicking through the task.

The two larger groups were more interesting, because they differed in how organised they searched, but not in how long they searched for (nor on any other aspect

of the task). This suggested that poor search organisation, thought by some to be characteristic of severe neglect (Ten Brink et al., 2016), is actually quite common in the healthy population.

It should be noted that pen-and-paper and computerised tests are excellent for testing a highly specific aspect of cognition in a very sensitive way, which makes it tempting to over-interpret positive findings. For example, in Chapter 3: *Introduction* I described the case of prism adaptation as a treatment for neglect. Initial reports found promising results with this technique on pen-and-paper cancellation tasks (Rossetti et al., 1998), and although these findings have partially been replicated (De Wit et al., 2016), a recently finished placebo-controlled trial could not find any clinically relevant benefits (Ten Brink, Visser-Meily, et al., 2017). This highlighted the importance of a double approach to neuropsychological testing, especially of potential interventions: Clinicians and researchers need highly sensitive tests to establish whether and how an intervention works, but also tests of daily living activities and quality of life to establish the practical benefits of an intervention.

9.7. Pharmacological rehabilitation of inattention after stroke

Chapter 8: *The effects of guanfacine on visual search performance, sustained attention, and short-term memory after stroke*, reported on the findings from a randomised, placebo-controlled, cross-over trial of the drug guanfacine in hemispatial neglect syndrome. Guanfacine is an agonist of the alpha-2 adrenergic receptor that has

been used in attention deficit hyperactivity disorder (Biederman et al., 2008; Scahill et al., 2001), and that has been shown to improve sustained attention and short-term memory (Arnsten et al., 1988; Jäkälä et al., 1999). More importantly, preliminary evidence suggested that guanfacine could ameliorate attention deficits in patients (Malhotra et al., 2006; Singh-Curry, Malhotra, Farmer, & Husain, 2011), including stroke patients with hemispatial neglect disorder.

Chapter 8 presented the results of 13 right-hemisphere stroke patients who suffered from left-sided neglect. Guanfacine, compared to placebo, improved the amount of targets they found in an invisible cancellation task by an average of 5 (out of 64). This result is promising, and in line with an earlier pilot of guanfacine in neglect (Malhotra et al., 2006). However, this result did not translate into ameliorated spatial bias.

When tested on independent tests of sustained attention and short-term memory, patients did not perform different when on guanfacine compared to placebo, or compared to baseline. In addition, patients' search organisation, indexed by four measures, did not differ between guanfacine and placebo treatment, nor from baseline performance. These were statistically robust null effects, as evidenced by Bayesian statistics.

One thing to note is that the sustained attention task might not have been very sensitive at measuring sustained attention. Notably, patients' response time variability (which indicates wavering sustained attention) did not increase during the task, which hints at the task not being sufficiently long, and at patients being focussed on the task throughout its duration. In addition, the sample size was not large enough to

demonstrate an effect (or lack thereof) from guanfacine on patients' spatial attention bias.

In addition to these methodological shortcomings, it should be noted that the employed dosage of 2 mg is rather on the low side compared to the 0.12 mg per kilogram of body weight employed in ADHD studies. Perhaps more importantly, only a single dose was employed. A promising case report has relied on continued medication, and demonstrated that withdrawing medication correlated with a return of attentional deficits (Singh-Curry et al., 2011).

Because the use of guanfacine in cognitive neurological patients has been described in only two other studies (the aforementioned N=1 case report by Singh-Curry and colleagues, and the N=3 pilot by Malhotra and colleagues), the results reported in chapter 8 are an important addition to the literature. They demonstrated that the effects of guanfacine on search behaviour are likely due to improvements of sustained attention or processing speed. These improvements are rather subtle, and it should be explored whether they are increased by continued guanfacine usage.

As the presented study was a proof-of-concept trial of guanfacine, highly sensitive tests were used. In addition to these, follow-up studies should employ tests of daily activities and quality of life.

In sum, a proof-of-concept has been delivered for guanfacine, suggesting that it could potentially be effective to improve sustained attention. However, the present study raises questions about the effective dosage: 2 mg might be too low, and continued usage might be necessary to improve sustained attention on less sensitive tasks. In addition, the benefit as it currently stands seems highly limited to the domain of sustained

attention, and does not extend into the many problems that neglect patients have. It remains to be investigated whether long-term guanfacine use could alleviate clinically relevant symptoms in neglect syndrome.

Such a long-term study could be a trial of guanfacine in a relatively small number of neglect patients (13 would be a good number for both statistically robust positive and negative effects, as Chapter 8 demonstrated), who should be closely monitored and regularly tested over the course of several months. In addition to the methodology employed in Chapter 8, patients and their caregivers should be given asked about potential differences in quality of life and daily functioning. Finally, as stressed before, clinically relevant tests should be part of the battery in this follow-up study.

10. Conclusion

A concise and neutrally worded summary of the work presented in this thesis is already given in the *Summary*, so I would like to use this space to highlight what I think are the most important findings, coloured by my own biases.

By far the most important contribution of my work is the strong empirical evidence for parallel encoding of information into short-term memory. The results presented here highlight the need for an update to our thinking about feature-integration theory and consolidation: Feature binding does not always require attention, and attention can be directed towards multiple items at the same time. A direct consequence is that multiple items can be committed to short-term memory at the same time. However, encoding does not have to happen in parallel, but can also be biased to occur serially when one item is behaviourally relevant, for example due to an associated reward. In my view, this supports a pool of encoding resources that can be flexibly distributed over items, not unlike how resource models describe short-term memory storage.

An additional conclusion is that sensitive cognitive tests are a crucial tool in any neurological intervention study, as they not only help reveal whether a drug or method works, but also how. Sensitive tests should, however, always be supplemented by tests of everyday functioning and clinically relevant symptoms to assess the practical benefits of the studied intervention.

11. 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. https://doi.org/10.1111/j.0963-7214.2004.01502006.x

Anderson, B. A., Laurent, P. A., & Yantis, S. (2011). Value-driven attentional capture. *Proceedings of the National Academy of Sciences*, *108*(25), 10367–10371. https://doi.org/10.1073/pnas.1104047108

Appelros, P., Karlsson, G. M., Seiger, A., & Nydevik, I. (2002). Neglect and Anosognosia After First-Ever Stroke: Incidence and Relationship to Disability. *Journal of Rehabilitation Medicine*, *34*(5), 215–220. https://doi.org/10.1080/165019702760279206

Arnsten, A. F. T., Cai, J. X., & Goldman-Rakic, P. S. (1988). The alpha-2 adrenergic agonist guanfacine improves memory in aged monkeys without sedative or hypotensive side effects: evidence for alpha-2 receptor subtypes. *Journal of Neuroscience*, *8*(11), 4287–4298.

Arnsten, A. F. T., & Contant, T. A. (1992). Alpha-2 adrenergic agonists decrease distractibility in aged monkeys performing the delayed response task. *Psychopharmacology*, *108*, 159–169.

Aston-Jones, G., Rajkowski, J., & Alexinsky, T. (1994). Locus coeruleus neurons in monkey are selectively activated by attended cues in a vigilance task. *Journal of Neuroscience*, *14*(7), 4467–4480.

Avery, R. A., Franowicz, J. S., Studholme, C., Dyck, C. H. van, & Arnsten, A. F. T. (2000). The Alpha-2A-Adrenoceptor Agonist, Guanfacine, Increases Regional Cerebral Blood Flow in Dorsolateral Prefrontal Cortex of Monkeys Performing a Spatial Working Memory Task. *Neuropsychopharmacology*, *23*(3), 240–249.

https://doi.org/10.1016/S0893-133X(00)00111-1

Awh, E., Anllo-Vento, L., & Hillyard, S. A. (2000). The role of spatial selective attention in working memory for locations: evidence from event-related potentials. *Journal of Cognitive Neuroscience*, *12*(5), 840–847.

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. https://doi.org/10.1111/j.1467-9280.2007.01949.x

Baddeley, A. D., & Hitch, G. (1974). Working Memory. In *Psychology of Learning and Motivation* (Vol. 8, pp. 47–89). Elsevier. https://doi.org/10.1016/S0079-7421(08)60452-1

Bae, G.-Y., Olkkonen, M., Allred, S. R., Wilson, C., & Flombaum, J. I. (2014). Stimulus-specific variability in color working memory with delayed estimation. *Journal of Vision*, *14*(4), 7–7. https://doi.org/10.1167/14.4.7

Bannerman, R. L., Milders, M., de Gelder, B., & Sahraie, A. (2009). Orienting to threat: faster localization of fearful facial expressions and body postures revealed by saccadic eye movements. *Proceedings of the Royal Society B: Biological Sciences*, *276*(1662), 1635–1641. https://doi.org/10.1098/rspb.2008.1744

Bannerman, R. L., Milders, M., & Sahraie, A. (2009). Processing emotional stimuli: Comparison of saccadic and manual choice-reaction times. *Cognition & Emotion*, *23*(5), 930–954. https://doi.org/10.1080/02699930802243303

Barceló, F., Suwazono, S., & Knight, R. T. (2000). Prefrontal modulation of visual processing in humans. *Nature Neuroscience*, *3*(4), 399–403. https://doi.org/10.1038/73975

Barrett, A. M., Crucian, G. P., Schwartz, R. L., & Heilman, K. M. (1999). Adverse effect of dopamine agonist therapy in a patient with motor-intentional neglect. *Archives of Physical Medicine and Rehabilitation*, *80*(5), 600–603.

Bays, P. M. (2014). Noise in Neural Populations Accounts for Errors in Working

Memory. *Journal of Neuroscience*, *34*(10), 3632–3645. https://doi.org/10.1523/JNEUROSCI.3204-13.2014

Bays, P. M. (2015). Spikes not slots: noise in neural populations limits working memory. *Trends in Cognitive Sciences*, *19*(8), 431–438. https://doi.org/10.1016/j.tics.2015.06.004

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), 1–11. https://doi.org/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), 1–15. https://doi.org/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. https://doi.org/10.1126/science.1158023

Bays, P. M., Singh-Curry, V., Gorgoraptis, N., Driver, J., & Husain, M. (2010). Integration of Goal- and Stimulus-Related Visual Signals Revealed by Damage to Human Parietal Cortex. *Journal of Neuroscience*, *30*(17), 5968–5978. https://doi.org/10.1523/JNEUROSCI.0997-10.2010

Berridge, C. W., & Waterhouse, B. D. (2003). The locus coeruleus–noradrenergic system: modulation of behavioral state and state-dependent cognitive processes. *Brain Research Reviews*, *42*(1), 33–84. https://doi.org/10.1016/S0165-0173(03)00143-7

Biederman, J., Melmed, R. D., Patel, A., McBurnett, K., Konow, J., Lyne, A., … for the SPD503 Study Group. (2008). A Randomized, Double-Blind, Placebo-Controlled Study of Guanfacine Extended Release in Children and Adolescents With Attention-Deficit/Hyperactivity Disorder. *PEDIATRICS*, *121*(1), e73–e84. https://doi.org/10.1542/peds.2006-3695

Binder, J., Marshall, R., Lazar, R., Benjamin, J., & Mohr, J. P. (1992). Distinct Syndromes of Hemineglect. *Archives of Neurology*, *49*(11), 1187–1194. https://doi.org/10.1001/archneur.1992.00530350109026

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), 1–34. https://doi.org/10.1167/11.5.4

Brady, T. F., Konkle, T., Gill, J., Oliva, A., & Alvarez, G. A. (2013). Visual Long-Term Memory Has the Same Limit on Fidelity as Visual Working Memory. *Psychological Science*, *24*(6), 981–990. https://doi.org/10.1177/0956797612465439

Bream, J. B., Lauener, H., Picard, C. W., Scholtysik, G., & White, T. G. (1975). Substituted phenylacetylguanidines: a new class of antihypertensive agents. *Arzneimittel-Forschung*, *25*(10), 1477–1482.

Brefczynski, J. A., & DeYoe, E. A. (1999). A physiological correlate of the “spotlight” of visual attention. *Nature Neuroscience*, *2*(4), 370–374. https://doi.org/10.1038/7280

Bucker, B., & Theeuwes, J. (2017). Stimulus-driven and goal-driven effects on Pavlovian associative reward learning. *Visual Cognition*, 1–18. https://doi.org/10.1080/13506285.2017.1399948

Bundesen, C. (1990). A theory of visual attention. *Psychological Review*, *97*(4), 523–547. https://doi.org/10.1037/0033-295X.97.4.523

Buxbaum, L. J., Ferraro, M. K., Veramonti, T., Farne, A., Whyte, J., Ladavas, E., … Coslett, H. B. (2004). Hemispatial neglect: Subtypes, neuroanatomy, and disability. *Neurology*, *62*(5), 749–756. https://doi.org/10.1212/01.WNL.0000113730.73031.F4

Cave, K. R., & Wolfe, J. M. (1990). Modeling the role of parallel processing in visual search. *Cognitive Psychology*, *22*(2), 225–271. https://doi.org/10.1016/0010-

0285(90)90017-X

Charron, S., & Koechlin, E. (2010). Divided Representation of Concurrent Goals in the Human Frontal Lobes. *Science*, *328*(5976), 360–363. https://doi.org/10.1126/science.1183614

Cohen, A., & Magen, H. (1999). Intra- and cross-dimensional visual search for single-feature targets. *Perception & Psychophysics*, *61*(2), 291–307.

Cohen, J. D., Perlstein, W. M., Braver, T. S., Nystrom, L. E., Noll, D. C., Jonides, J., & Smith, E. E. (1997). Temporal dynamics of brain activation during a working memory task. *Nature*, *386*(6625), 604–608. https://doi.org/10.1038/386604a0

Constantinidis, C., Franowicz, M. N., & Goldman-Rakic, P. S. (2001). The sensory nature of mnemonic representation in the primate prefrontal cortex. *Nature Neuroscience*, *4*, 311–316.

Constantinidis, C., & Steinmetz, M. A. (1996). Neuronal activity in posterior parietal area 7a during the delay periods of a spatial memory task. *Journal of Neurophysiology*, *76*(2), 1352–1355.

Conway, A. R. A., Cowan, N., Bunting, M. F., Therriault, D. J., & Minkoff, S. R. B. (2002). A latent variable analysis of working memory capacity, short-term memory capacity, processing speed, and general fluid intelligence. *Intelligence*, *30*(2), 163–183. https://doi.org/10.1016/S0160-2896(01)00096-4

Corbetta, M., Miezin, F. M., Dobmeyer, S., Shulman, G. L., & Petersen, S. E. (1991). Selective and divided attention during visual discrimination of shape, color, and speed: functional anatomy by positron emission topography. *Journal of Neuroscience*, *11*(8), 2383–2402.

Corbetta, M., Miezin, F. M., Shulman, G. L., & Petersen, S. E. (1993). A PET study of visuospatial attention. *Journal of Neuroscience*, *13*(3), 1202–1226.

Courtney, S. M., Ungerleider, L. G., Keil, K., & Haxby, J. V. (1997). Transient and sustained activity in a distributed neural system for human working memory.

Nature, *386*(6625), 608–611. https://doi.org/10.1038/386608a0

Cowan, N. (1988). Evolving conceptions of memory storage, selective attention, and their mutual constraints within the human information-processing system. *Psychological Bulletin*, *104*, 163–191.

Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, *24*, 87–185.

Crick, F. (1984). Function of the thalamic reticular complex: The searchlight hypothesis. *Proceedings of the National Academy of Sciences*, *81*, 4586–4590.

Curtis, C. E., & D'Esposito, M. (2003). Persistent activity in the prefrontal cortex during working memory. *Trends in Cognitive Sciences*, *7*(9), 415–423. https://doi.org/10.1016/S1364-6613(03)00197-9

Curtis, C. E., Mackey, W. E., Ding, X., Wang, X.-J., & Winawer, J. (2016). Visual field maps constrain working memory precision. In *2016 Neuroscience Meeting Planner* (p. 385.10). San Diego, CA, USA.

Dalmaijer, E. S. (2017). *Python for experimental psychologists*. Abingdon, Oxon; New York, NY: Routledge.

Dalmaijer, E. S., Van der Stigchel, S., Nijboer, T. C. W., Cornelissen, T. H. W., & Husain, M. (2015). CancellationTools: All-in-one software for administration and analysis of cancellation tasks. *Behavior Research Methods*, *47*(4), 1065–1075. https://doi.org/10.3758/s13428-014-0522-7

De Wit, L., Ten Brink, A. F., Visser-Meily, J. M. A., & Nijboer, T. C. W. (2016). Does prism adaptation affect visual search in spatial neglect patients: A systematic review. *Journal of Neuropsychology*, n/a-n/a. https://doi.org/10.1111/jnp.12100

Deco, G., & Rolls, E. T. (2008). Neural Mechanisms of Visual Memory: A Neurocomputational Perspective. In S. J. Luck & A. Hollingworth (Eds.), *Visual Memory* (pp. 247–290). Oxford University Press. https://doi.org/10.1093/acprof:oso/9780195305487.003.0008

Della Libera, C., & Chelazzi, L. (2006). Visual selective attention and the effects of monetary rewards. *Psychological Science*, *17*(3), 222–227.

Donnelly, N., Guest, R., Fairhurst, M., Potter, J., Deighton, A., & Patel, M. (1999). Developing algorithms to enhance the sensitivity of cancellation tests of visuospatial neglect. *Behavior Research Methods, Instruments, & Computers*, *31*(4), 668–673.

Donohue, S. E., Hopf, J.-M., Bartsch, M. V., Schoenfeld, M. A., Heinze, H.-J., & Woldorff, M. G. (2016). The Rapid Capture of Attention by Rewarded Objects. *Journal of Cognitive Neuroscience*, *28*(4), 529–541. https://doi.org/10.1162/jocn_a_00917

Duncan, J. (1984). Selective attention and the organization of visual information. *Journal of Experimental Psychology: General*, *113*(4), 501–517.

Duncan, J., Ward, R., & Shapiro, K. (1994). Direct measurement of attentional dwell time in human vision. *Nature*, *369*(6478), 313–315. https://doi.org/10.1038/369313a0

Duyck, W., Desmet, T., Verbeke, L. P. C., & Brysbaert, M. (2004). WordGen: A tool for word selection and nonword generation in Dutch, English, German, and French. *Behavior Research Methods, Instruments, & Computers*, *36*(3), 488–499.

Eimer, M., & Forster, B. (2003). Modulations of early somatosensory ERP components by transient and sustained spatial attention. *Experimental Brain Research*, *151*(1), 24–31.

Eng, H. Y., Chen, D., & Jiang, Y. (2005). Visual working memory for simple and complex visual stimuli. *Psychonomic Bulletin & Review*, *12*(6), 1127–1133.

Engle, R. W., Tuholski, S. W., Laughlin, J. E., & Conway, A. R. A. (1999). Working memory, short-term memory, and general fluid intelligence: A latent-variable approach. *Journal of Experimental Psychology: General*, *128*(3), 309–331. https://doi.org/10.1037/0096-3445.128.3.309

Eriksen, C. W., & St. James, J. D. (1986). Visual attention within and around the field of focal attention: A zoom lens model. *Perception & Psychophysics*, *40*(4), 225–240.

Fabrigoule, C., Rouch, I., Taberly, A., Letenneur, L., Commenges, D., Mazaux, J. M., … Dartigues, J. F. (1998). Cognitive process in preclinical phase of dementia. *Brain*, *121*(1), 135–141. https://doi.org/10.1093/brain/121.1.135

Failing, M., & Theeuwes, J. (2017). Don't let it distract you: how information about the availability of reward affects attentional selection. *Attention, Perception, & Psychophysics*, *79*(8), 2275–2298. https://doi.org/10.3758/s13414-017-1376-8

Fleet, W. S., Valenstein, E., Watson, R. T., & Heilman, K. M. (1987). Dopamine agonist therapy for neglect in humans. *Neurology*, *37*, 1765–1771.

Foote, S. L., Aston-Jones, G., & Bloom, F. E. (1980). Impulse activity of locus coeruleus neurons in awake rats and monkeys is a function of sensory stimulation and arousal. *Proceedings of the National Academy of Sciences*, *77*(5), 3033–3037.

Free Software Foundation. (2007). GNU General Public License. Retrieved from https://gnu.org/licenses/gpl.html

Fukuda, K., Awh, E., & Vogel, E. K. (2010). Discrete capacity limits in visual working memory. *Current Opinion in Neurobiology*, *20*(2), 177–182. https://doi.org/10.1016/j.conb.2010.03.005

Funahashi, S., Bruce, C. J., & Goldman-Rakic, P. S. (1989). Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *Journal of Neurophysiology*, *61*(2), 331–349. https://doi.org/10.1152/jn.1989.61.2.331

Fuster, J. M., & Alexander, G. E. (1971). Neuron activity related to short-term memory. *Science*, *173*(3997), 652–654.

Gainotti, G., Marra, C., & Villa, G. (2001). A double dissociation between accuracy and time of execution on attentional tasks in Alzheimer's disease and multi-infarct

dementia. *Brain*, *124*(4), 731–738. https://doi.org/10.1093/brain/124.4.731

Gayet, S., Paffen, C. L. E., & Van der Stigchel, S. (2013). Information Matching the Content of Visual Working Memory Is Prioritized for Conscious Access. *Psychological Science*, *24*(12), 2472–2480. https://doi.org/10.1177/0956797613495882

Gayet, S., Van der Stigchel, S., & Paffen, C. L. E. (2014). Breaking continuous flash suppression: competing for consciousness on the pre-semantic battlefield. *Frontiers in Psychology*, *5*. https://doi.org/10.3389/fpsyg.2014.00460

Geminiani, G., Bottini, G., & Sterzi, R. (1998). Dopaminergic stimulation in unilateral neglect. *Journal of Neurology, Neurosurgery & Psychiatry*, *65*(3), 344–347. https://doi.org/10.1136/jnnp.65.3.344

Gorgoraptis, N., Mah, Y.-H., Machner, B., Singh-Curry, V., Malhotra, P. A., Hadji-Michael, M., … Husain, M. (2012). The effects of the dopamine agonist rotigotine on hemispatial neglect following stroke. *Brain*, *135*(8), 2478–2491. https://doi.org/10.1093/brain/aws154

Gramfort, A., Luessi, M., Larson, E., Engemann, D. A., Strohmeier, D., Brodbeck, C., … Hamalainen, M. (2013). MEG and EEG data analysis with MNE-Python. *Frontiers in Neuroscience*, *7*. https://doi.org/10.3389/fnins.2013.00267

Grujic, Z., Mapstone, M., Gitelman, D. R., Johnson, N., Weintraub, S., Hays, A., … Mesulam, M.-M. (1998). Dopamine agonists reorient visual exploration away from the neglected hemispace. *Neurology*, *51*(5), 1395–1398.

Hahn, S., & Kramer, A. F. (1998). Further Evidence for the Division of Attention Among Non-contiguous Locations. *Visual Cognition*, *5*(1–2), 217–256. https://doi.org/10.1080/713756781

Harrison, S. A., & Tong, F. (2009). Decoding reveals the contents of visual working memory in early visual areas. *Nature*, *458*(7238), 632–635. https://doi.org/10.1038/nature07832

Hickey, C., Chelazzi, L., & Theeuwes, J. (2010a). Reward Changes Salience in Human Vision via the Anterior Cingulate. *Journal of Neuroscience*, *30*(33), 11096–11103. https://doi.org/10.1523/JNEUROSCI.1026-10.2010

Hickey, C., Chelazzi, L., & Theeuwes, J. (2010b). Reward Guides Vision when It's Your Thing: Trait Reward-Seeking in Reward-Mediated Visual Priming. *PLoS ONE*, *5*(11), e14087. https://doi.org/10.1371/journal.pone.0014087

Hills, E. C., & Geldmacher, D. S. (1998). The effect of character and array type on visual spatial search quality following traumatic brain injury. *Brain Injury*, *12*(1), 69–76. https://doi.org/10.1080/026990598122872

Hoffman, J. E. (1979). A two-stage model of visual search. *Perception & Psychophysics*, *24*(4), 319–327.

Hogan, R. M., & Kintsch, W. (1971). Differential effects of study and test trials on long-term recognition and recall. *Journal of Verbal Learning and Verbal Behavior*, *10*(5), 562–567. https://doi.org/10.1016/S0022-5371(71)80029-4

Holm, S. (1979). A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics*, *6*(2), 65–70.

Houck, M. R., & Hoffman, J. E. (1986). Conjunction of color and form without attention: Evidence from an orientation-contingent color aftereffect. *Journal of Experimental Psychology: Human Perception and Performance*, *12*(2), 186–199. https://doi.org/10.1037/0096-1523.12.2.186

Huang, H.-C., & Wang, T.-Y. (2008). Visualized representation of visual search patterns for a visuospatial attention test. *Behavior Research Methods*, *40*(2), 383–390. https://doi.org/10.3758/BRM.40.2.383

Hunter, J. D. (2007). Matplotlib: A 2D Graphics Environment. *Computing in Science & Engineering*, *9*(3), 90–95. https://doi.org/10.1109/MCSE.2007.55

Hurford, P., Stringer, A. Y., & Jann, B. (1998). Neuropharmacologic treatment of hemineglect: a case report comparing bromocriptine and methylphenidate.

Archives of Physical Medicine and Rehabilitation, *79*(3), 346–349.

Husain, M., & Rorden, C. (2003). Non-spatially lateralized mechanisms in hemispatial neglect. *Nature Reviews Neuroscience*, *4*(1), 26–36. https://doi.org/10.1038/nrn1005

Iverson, P., Smith, C. A., & Evans, B. G. (2006). Vowel recognition via cochlear implants and noise vocoders: effects of formant movement and duration. *The Journal of the Acoustical Society of America*, *120*(6), 3998–4006.

Jäkälä, P., Riekkinen, M., Sirviö, J., Koivisto, E., Kejonen, K., Vanhanen, M., & Riekkinen, P., Jr. (1999). Guanfacine, But Not Clonidine, Improves Planning and Working Memory Performance in Humans. *Neuropsychopharmacology*, *20*(5), 460–470. https://doi.org/10.1016/S0893-133X(98)00127-4

James, W. (1890). *The principles of psychology*. New York, USA: Holt.

JASP Team. (2016). JASP (Version 0.7.1.12). Retrieved from https://jasp-stats.org

Jeffreys, H. (1961). *Theory of probability*. Oxford, UK: Clarendon Press.

Jensen, O., Gelfand, J., Kounios, J., & Lisman, J. E. (2002). Oscillations in the Alpha Band (9-12 Hz) Increase with Memory Load during Retention in a Short-term Memory Task. *Cerebral Cortex*, *12*(8), 877–882. https://doi.org/10.1093/cercor/12.8.877

Jensen, O., & Lisman, J. E. (1998). An oscillatory short-term memory buffer model can account for data on the Sternberg task. *The Journal of Neuroscience*, *18*(24), 10688–10699.

Ji, X.-H., Ji, J.-Z., Zhang, H., & Li, B.-M. (2008). Stimulation of α2-Adrenoceptors Suppresses Excitatory Synaptic Transmission in the Medial Prefrontal Cortex of Rat. *Neuropsychopharmacology*, *33*(9), 2263–2271. https://doi.org/10.1038/sj.npp.1301603

Jones, M. N., & Mewhort, D. J. K. (2004). Case-sensitive letter and bigram frequency counts from large-scale English corpora. *Behavior Research Methods*,

Instruments, & Computers, *36*(3), 388–396.

Joseph, S., Iverson, P., Manohar, S., Fox, Z., Scott, S. K., & Husain, M. (2015). Precision of working memory for speech sounds. *Quarterly Journal of Experimental Psychology*, *68*(10), 2022–2040. https://doi.org/10.1080/17470218.2014.1002799

Joseph, S., Kumar, S., Husain, M., & Griffiths, T. D. (2015). Auditory working memory for objects vs. features. *Frontiers in Neuroscience*, *9*. https://doi.org/10.3389/fnins.2015.00013

Kastner, S., De Weerd, P., Desimone, R., & Ungerleider, L. G. (1998). Mechanisms of directed attention in the human extrastriate cortex as revealed by functional MRI. *Science*, *282*(5386), 108–111. https://doi.org/10.1126/science.282.5386.108

Kiyonaga, A., & Egner, T. (2013). Working memory as internal attention: Toward an integrative account of internal and external selection processes. *Psychonomic Bulletin & Review*, *20*(2), 228–242. https://doi.org/10.3758/s13423-012-0359-y

Klink, P. C., Jeurissen, D., Theeuwes, J., Denys, D., & Roelfsema, P. R. (2017). Working memory accuracy for multiple targets is driven by reward expectation and stimulus contrast with different time-courses. *Scientific Reports*, *7*(1). https://doi.org/10.1038/s41598-017-08608-4

Kramer, A. F., & Hahn, S. (1995). Splitting the beam: Distribution of attention over noncontiguous regions of the visual field. *Psychological Science*, *6*(6), 381–386.

Kumada, T. (2001). Feature-based control of attention: Evidence for two forms of dimension weighting. *Perception & Psychophysics*, *63*(4), 698–708.

Kumar, S., Joseph, S., Pearson, B., Teki, S., Fox, Z. V., Griffiths, T. D., & Husain, M. (2013). Resource allocation and prioritization in auditory working memory. *Cognitive Neuroscience*, *4*(1), 12–20. https://doi.org/10.1080/17588928.2012.716416

LaRocque, J. J., Lewis-Peacock, J. A., Drysdale, A. T., Oberauer, K., & Postle, B. R. (2013). Decoding Attended Information in Short-term Memory: An EEG Study. *Journal of Cognitive Neuroscience*, *25*(1), 127–142. https://doi.org/10.1162/jocn_a_00305

Lawrence, N. S., Ross, T. J., Hoffmann, R., Garavan, H., & Stein, E. A. (2003). Multiple Neuronal Networks Mediate Sustained Attention. *Journal of Cognitive Neuroscience*, *15*(7), 1028–1038. https://doi.org/10.1162/089892903770007416

LeDoux, J. E. (2000). Emotion Circuits in the Brain. *Annual Review of Neuroscience*, *23*(1), 155–184. https://doi.org/10.1146/annurev.neuro.23.1.155

LeDoux, J. E. (2003). The emotional brain, fear, and the amygdala. *Cellular and Molecular Neurobiology*, *23*(4–5), 727–738.

Lesniak, M., Bak, T., Czepiel, W., Seniow, J., & Czlonkowska, A. (2008). Frequency and Prognostic Value of Cognitive Disorders in Stroke Patients. *Dementia and Geriatric Cognitive Disorders*, *26*(4), 356–363. https://doi.org/10.1159/000162262

Lewis-Peacock, J. A., Drysdale, A. T., Oberauer, K., & Postle, B. R. (2012). Neural Evidence for a Distinction between Short-term Memory and the Focus of Attention. *Journal of Cognitive Neuroscience*, *24*(1), 61–79. https://doi.org/10.1162/jocn_a_00140

Lisman, J. E., & Idiart, M. A. P. (1995). Storage of 7 ± 2 short-term memories in oscillatory subcycles. *Science*, *267*, 1512–1515.

Liu, T., & Becker, M. W. (2013). Serial Consolidation of Orientation Information Into Visual Short-Term Memory. *Psychological Science*, *24*(6), 1044–1050. https://doi.org/10.1177/0956797612464381

LoBue, V., Matthews, K., Harvey, T., & Stark, S. L. (2014). What accounts for the rapid detection of threat? Evidence for an advantage in perceptual and behavioral responding from eye movements. *Emotion*, *14*(4), 816–823.

https://doi.org/10.1037/a0035869

Luck, S. J., Hillyard, S. A., Mouloua, M., Woldorff, M. G., Clark, V. P., & Hawkins, H. L. (1994). Effects of spatial cueing on luminance detectability. *Journal of Experimental Psychology: Human Perception and Performance*, *20*, 887–904.

Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, *390*, 279–281.

Luck, S. J., & Vogel, E. K. (2013a). Visual working memory capacity: from psychophysics and neurobiology to individual differences. *Trends in Cognitive Sciences*, *17*(8), 391–400. https://doi.org/10.1016/j.tics.2013.06.006

Luck, S. J., & Vogel, E. K. (2013b). Visual working memory capacity: from psychophysics and neurobiology to individual differences. *Trends in Cognitive Sciences*, *17*(8), 391–400. https://doi.org/10.1016/j.tics.2013.06.006

Ma, W. J., Husain, M., & Bays, P. M. (2014). Changing concepts of working memory. *Nature Neuroscience*, *17*(3), 347–356. https://doi.org/10.1038/nn.3655

Malhotra, P. A., Jager, H. R., Parton, A., Greenwood, R., Playford, E. D., Brown, M. M., & Driver, J. (2005). Spatial working memory capacity in unilateral neglect. *Brain*, *128*(2), 424–435. https://doi.org/10.1093/brain/awh372

Malhotra, P. A., Parton, A. D., Greenwood, R., & Husain, M. (2006). Noradrenergic modulation of space exploration in visual neglect. *Annals of Neurology*, *59*(1), 186–190. https://doi.org/10.1002/ana.20701

Mangun, G. R., & Hillyard, S. A. (1991). Modulation of sensory-evoked brain potentials provide evidence for changes in perceptual processing during visual-spatial priming. *Journal of Experimental Psychology: Human Perception and Performance*, *17*, 1057–1074.

Mannan, S. K., Mort, D. J., Hodgson, T. L., Driver, J., Kennard, C., & Husain, M. (2005). Revisiting Previously Searched Locations in Visual Neglect: Role of Right Parietal and Frontal Lesions in Misjudging Old Locations as New. *Journal*

of Cognitive Neuroscience, *17*(2), 340–354. https://doi.org/10.1162/0898929053124983

Manohar, S. G., Zokaei, N., Fallon, S. J., Vogels, T., & Husain, M. (2017). A neural model of working memory. *BioRXiv*. https://doi.org/10.1101/233007

Mark, V. W., Woods, A. J., Ball, K. K., Roth, D. L., & Mennenmeier, M. (2004). Disorganized search on cancellation is not a consequence of neglect. *Neurology*, *63*, 78–84.

Matthey, L., Bays, P. M., & Dayan, P. (2015). A Probabilistic Palimpsest Model of Visual Short-term Memory. *PLOS Computational Biology*, *11*(1), e1004003. https://doi.org/10.1371/journal.pcbi.1004003

Maunsell, J. H. R., & Treue, S. (2006). Feature-based attention in visual cortex. *Trends in Neurosciences*, *29*(6), 317–322.

McCollough, C. (1965). Color adaptation of edge-detectors in the human visual system. *Science*, *49*, 1115–1116.

McMains, S. A., & Somers, D. C. (2004). Multiple Spotlights of Attentional Selection in Human Visual Cortex. *Neuron*, *42*(4), 677–686. https://doi.org/10.1016/S0896-6273(04)00263-6

Mesulam, M.-M. (1985). *Principles of Behavioural Neurology. Tests of directed attention and memory.* Philadelphia, US: Davis.

Mok, R.M., Myers, N. E., Wallis, G., & Nobre, A. C. (2016). Behavioral and Neural Markers of Flexible Attention over Working Memory in Aging. *Cerebral Cortex*, *26*(4), 1831–1842. https://doi.org/10.1093/cercor/bhw011

Mok, Robert M., Myers, N. E., Wallis, G., & Nobre, A. C. (2016). Behavioral and Neural Markers of Flexible Attention over Working Memory in Aging. *Cerebral Cortex*, *26*(4), 1831–1842. https://doi.org/10.1093/cercor/bhw011

Motter, B. C. (1994). Neural correlates of attentive selection for color or luminance in exrastriate area V4. *Journal of Neuroscience*, *14*(4), 2178–2189.

Mukand, J. A., Guilmette, T. J., Allen, D. G., Brown, L. K., Brown, S. L., Tober, K. L., & VanDyck, W. R. (2001). Dopaminergic therapy with carbidopa L-dopa for left neglect after stroke: a case series. *Archives of Physical Medicine and Rehabilitation*, *82*(9), 1279–1282.

Mulckhuyse, M., Crombez, G., & Van der Stigchel, S. (2013). Conditioned fear modulates visual selection. *Emotion*, *13*(3), 529–536. https://doi.org/10.1037/a0031076

Mulckhuyse, M., & Dalmaijer, E. S. (2016). Distracted by danger: Temporal and spatial dynamics of visual selection in the presence of threat. *Cognitive, Affective, & Behavioral Neuroscience*, *16*(2), 315–324. https://doi.org/10.3758/s13415-015-0391-2

Müller, M. M., Malinowski, P., Gruber, T., & Hillyard, S. A. (2003). Sustained division of the attentional spotlight. *Nature*, *424*, 309–312.

Munoz, D. P. (2002). Commentary: Saccadic eye movements: overview of neural circuitry. In *Progress in Brain Research* (Vol. 140, pp. 89–96). Elsevier. https://doi.org/10.1016/S0079-6123(02)40044-1

Myers, N. E., Stokes, M. G., Walther, L., & Nobre, A. C. (2014). Oscillatory Brain State Predicts Variability in Working Memory. *Journal of Neuroscience*, *34*(23), 7735–7743. https://doi.org/10.1523/JNEUROSCI.4741-13.2014

Myers, N. E., Walther, L., Wallis, G., Stokes, M. G., & Nobre, A. C. (2015). Temporal Dynamics of Attention during Encoding versus Maintenance of Working Memory: Complementary Views from Event-related Potentials and Alpha-band Oscillations. *Journal of Cognitive Neuroscience*, *27*(3), 492–508. https://doi.org/10.1162/jocn_a_00727

Nijboer, T. C. W., Kollen, B. J., & Kwakkel, G. (2013). Time course of visuospatial neglect early after stroke: A longitudinal cohort study. *Cortex*, *49*(8), 2021–2027. https://doi.org/10.1016/j.cortex.2012.11.006

Nijboer, T. C. W., Kollen, B. J., & Kwakkel, G. (2014). The Impact of Recovery of Visuo-Spatial Neglect on Motor Recovery of the Upper Paretic Limb after Stroke. *PLoS ONE*, *9*(6), e100584. https://doi.org/10.1371/journal.pone.0100584

Nijboer, T. C. W., Van de Port, I., Schepers, V., Post, M., & Visser-Meily, A. (2013). Predicting Functional Outcome after Stroke: The Influence of Neglect on Basic Activities in Daily Living. *Frontiers in Human Neuroscience*, *7*. https://doi.org/10.3389/fnhum.2013.00182

Nys, G. M. S., Van Zandvoort, M. J. E., De Kort, P. L. M., Van der Worp, H. B., Jansen, B. P., De Haan, E. H. F., & Kappelle, L. J. (2005). The prognostic value of domain-specific cognitive abilities in acute first-ever stroke. *Neurology*, *64*(5), 821–827.

Oberauer, K. (2002). Access to information in working memory: Exploring the focus of attention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *28*, 411–421.

O'Craven, K. M., Rosen, B. R., Kwong, K. K., Treisman, A., & Savoy, R. L. (1997). Voluntary Attention Modulates fMRI Activity in Human MT–MST. *Neuron*, *18*(4), 591–598. https://doi.org/10.1016/S0896-6273(00)80300-1

Oliphant, T. E. (2007). Python for Scientific Computing. *Computing in Science & Engineering*, *9*(3), 10–20. https://doi.org/10.1109/MCSE.2007.58

Oostwoud Wijdenes, L., Marshall, L., & Bays, P. M. (2015). Evidence for Optimal Integration of Visual Feature Representations across Saccades. *Journal of Neuroscience*, *35*(28), 10146–10153. https://doi.org/10.1523/JNEUROSCI.1040-15.2015

Pardo, J. V., Fox, P. T., & Raichle, M. E. (1991). Localization of a human system for sustained attention by positron emission tomography. *Nature*, *349*(6304), 61–64. https://doi.org/10.1038/349061a0

Parton, A. D., Malhotra, P. A., Nachev, P., Ames, D., Ball, J., Chataway, J., & Husain,

M. (2006). Space re-exploration in hemispatial neglect. *NeuroReport*, *17*(8), 833–836.

Peck, C. J., Jangraw, D. C., Suzuki, M., Efem, R., & Gottlieb, J. (2009). Reward Modulates Attention Independently of Action Value in Posterior Parietal Cortex. *Journal of Neuroscience*, *29*(36), 11182–11191. https://doi.org/10.1523/JNEUROSCI.1929-09.2009

Petit, L., Courtney, S. M., Ungerleider, L. G., & Haxby, J. V. (1998). Sustained activity in the medial wall during working memory delays. *Journal of Neuroscience*, *18*(22), 9429–9437.

Pierce, S. R., & Buxbaum, L. J. (2002). Treatments of unilateral neglect: a review. *Archives of Physical Medicine and Rehabilitation*, *83*(2), 256–268.

Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology*, *32*(1), 3–25. https://doi.org/10.1080/00335558008248231

Posner, M. I., & Boies, S. J. (1971). Components of attention. *Psychological Review*, *78*(5), 391–408. https://doi.org/10.1037/h0031333

Posner, M. I., Snyder, C. R. R., & Davidson, B. J. (1980). Attention and the detection of signals. *Journal of Experimental Psychology: General*, *109*(2), 160–174.

Rabuffetti, M., Farina, E., Alberoni, M., Pellegatta, D., Appollonio, I., Affanni, P., … Ferrarin, M. (2012a). Spatio-Temporal Features of Visual Exploration in Unilaterally Brain-Damaged Subjects with or without Neglect: Results from a Touchscreen Test. *PLoS ONE*, *7*(2), e31511. https://doi.org/10.1371/journal.pone.0031511

Rabuffetti, M., Farina, E., Alberoni, M., Pellegatta, D., Appollonio, I., Affanni, P., … Ferrarin, M. (2012b). Spatio-Temporal Features of Visual Exploration in Unilaterally Brain-Damaged Subjects with or without Neglect: Results from a Touchscreen Test. *PLoS ONE*, *7*(2), e31511. https://doi.org/10.1371/journal.pone.0031511

Reeves, A., & Sperling, G. (1986). Attention gating in short-term visual memory. *Psychological Review*, *93*(2), 180–206. https://doi.org/10.1037/0033-295X.93.2.180

Reitman, J. S. (1974). Without surreptitious rehearsal, information in short-term memory decay. *Journal of Verbal Learning and Verbal Behavior*, *13*(4), 365–377. https://doi.org/10.1016/S0022-5371(74)80015-0

Reynolds, M. G., Eastwood, J. D., Partanen, M., Frischen, A., & Smilek, D. (2009). Monitoring eye movements while searching for affective faces. *Visual Cognition*, *17*(3), 318–333. https://doi.org/10.1080/13506280701623704

Riccio, C. A., Reynolds, C. R., & Lowe, P. A. (2001). *Clinical Applications of Continuous Performance Tests*. New York, US: John Wiley & Sons.

Ringman, J. M., Saver, J. L., Woolson, R. F., Clarke, W. R., & Adams, H. P. (2004). Frequency, risk factors, anatomy, and course of unilateral neglect in an acute stroke cohort. *Neurology*, *63*(3), 468–474.

Robertson, I. H. (2001). Do We Need the “Lateral” in Unilateral Neglect? Spatially Nonselective Attention Deficits in Unilateral Neglect and Their Implications for Rehabilitation. *NeuroImage*, *14*(1), S85–S90.

Robertson, I. H., & Garavan, H. (2004). Vigilant Attention. In M. S. Gazzaniga (Ed.), *The Cognitive Neurosciences III*. Cambridge, MA, USA: MIT Press.

Robertson, I. H., Manly, T., Andrade, J., Baddeley, B. T., & Yiend, J. (1997). “Oops!”: performance correlates of everyday attentional failures in traumatic brain injured and normal subjects. *Neuropsychologia*, *35*(6), 747–758.

Robertson, I. H., Manly, T., Beschin, N., Daini, R., Haeske-Dewick, H., Hömberg, V., … Weber, E. (1997). Auditory sustained attention is a marker of unilateral spatial neglect. *Neuropsychologia*, *35*(12), 1527–1532. https://doi.org/10.1016/S0028-3932(97)00084-5

Robertson, I. H., Mattingley, J. B., Rorden, C., & Driver, J. (1998). Phasic alerting of

neglect patients overcomes their spatial deficit in visual awareness. *Nature*, *395*(6698), 169–172. https://doi.org/10.1038/25993

Rolfs, M., Jonikaitis, D., Deubel, H., & Cavanagh, P. (2011). Predictive remapping of attention across eye movements. *Nature Neuroscience*, *14*(2), 252–256. https://doi.org/10.1038/nn.2711

Rorden, C., & Karnath, H.-O. (2010). A simple measure of neglect severity. *Neuropsychologia*, *48*(9), 2758–2763. https://doi.org/10.1016/j.neuropsychologia.2010.04.018

Rossetti, Y., Rode, G., Pisella, L., Farné, A., Li, L., Boisson, D., & Perenin, M.-T. (1998). Prism adaptation to a rightward optical deviation rehabilitates left hemispatial neglect. *Nature*, *395*(6698), 166–169. https://doi.org/10.1038/25988

Rossi, A. F., & Paradiso, M. A. (1995). Feature-specific effects of selective visual attention. *Vision Research*, *35*(5), 621–634. https://doi.org/10.1016/0042-6989(94)00156-G

Sàenz, M., Buraĉas, G. T., & Boynton, G. M. (2003). Global feature-based attention for motion and color. *Vision Research*, *43*(6), 629–637. https://doi.org/10.1016/S0042-6989(02)00595-3

Samuelsson, H., Hjelmquist, E. K. E., Jensen, C., & Blomstrand, C. (2002). Search pattern in a verbally reported visual scanning test in patients showing spatial neglect. *Journal of the International Neuropsychological Society*, *8*(03), 382–394. https://doi.org/10.1017/S1355617702813194

Sauseng, P., Klimesch, W., Stadler, W., Schabus, M., Doppelmayr, M., Hanslmayr, S., … Birbaumer, N. (2005). A shift of visual spatial attention is selectively associated with human EEG alpha activity. *European Journal of Neuroscience*, *22*(11), 2917–2926. https://doi.org/10.1111/j.1460-9568.2005.04482.x

Scahill, L., Chappell, P. B., Kim, Y. S., Schultz, R. T., Katsovich, L., Shepherd, E., … Leckman, J. F. (2001). A Placebo-Controlled Study of Guanfacine in the

Treatment of Children With Tic Disorders and Attention Deficit Hyperactivity Disorder. *American Journal of Psychiatry*, *158*(7), 1067–1074. https://doi.org/10.1176/appi.ajp.158.7.1067

Schmidt, L. J., Belopolsky, A. V., & Theeuwes, J. (2015). Attentional capture by signals of threat. *Cognition and Emotion*, *29*(4), 687–694. https://doi.org/10.1080/02699931.2014.924484

Schmidt, L. J., Belopolsky, A. V., & Theeuwes, J. (2017). The time course of attentional bias to cues of threat and safety. *Cognition and Emotion*, *31*(5), 845–857. https://doi.org/10.1080/02699931.2016.1169998

Schneegans, S., & Bays, P. M. (2017a). Neural Architecture for Feature Binding in Visual Working Memory. *The Journal of Neuroscience*, *37*(14), 3913–3925. https://doi.org/10.1523/JNEUROSCI.3493-16.2017

Schneegans, S., & Bays, P. M. (2017b). Restoration of fMRI Decodability Does Not Imply Latent Working Memory States. *Journal of Cognitive Neuroscience*, 1–18. https://doi.org/10.1162/jocn_a_01180

Serences, J. T. (2008). Value-Based Modulations in Human Visual Cortex. *Neuron*, *60*(6), 1169–1181. https://doi.org/10.1016/j.neuron.2008.10.051

Shao, N., Li, J., Shui, R., Zheng, X., Lu, J., & Shen, M. (2010). Saccades elicit obligatory allocation of visual working memory. *Memory & Cognition*, *38*(5), 629–640. https://doi.org/10.3758/MC.38.5.629

Sheremata, S. L., & Silver, M. A. (2015). Hemisphere-Dependent Attentional Modulation of Human Parietal Visual Field Representations. *Journal of Neuroscience*, *35*(2), 508–517. https://doi.org/10.1523/JNEUROSCI.2378-14.2015

Singh-Curry, V., Malhotra, P. A., Farmer, S. F., & Husain, M. (2011). Attention deficits following ADEM ameliorated by guanfacine. *Journal of Neurology, Neurosurgery & Psychiatry*, *82*(6), 688–690.

https://doi.org/10.1136/jnnp.2009.195792

Smit, M., Schutter, D. J. L. G., Nijboer, T. C. W., Visser-Meily, J. M. A., Kappelle, L. J., Kant, N., … Dijkerman, H. C. (2015). Transcranial direct current stimulation to the parietal cortex in hemispatial neglect: A feasibility study. *Neuropsychologia*, *74*, 152–161. https://doi.org/10.1016/j.neuropsychologia.2015.04.014

Smith, A., & Nutt, D. (1996). Noradrenaline and attention lapses. *Nature*, *380*(6572), 291–291. https://doi.org/10.1038/380291a0

Solso, R. L., Barbuto, P. F., & Juel, C. L. (1979). Bigram and trigram frequencies and versatilities in the English language. *Behavior Research Methods, Instruments, & Computers*, *11*(5), 475–484.

Souza, A. S., & Oberauer, K. (2016). In search of the focus of attention in working memory: 13 years of the retro-cue effect. *Attention, Perception, & Psychophysics*, *78*(7), 1839–1860. https://doi.org/10.3758/s13414-016-1108-5

Sprague, T. C., Ester, E. F., & Serences, J. T. (2016). Restoring Latent Visual Working Memory Representations in Human Cortex. *Neuron*, *91*(3), 694–707. https://doi.org/10.1016/j.neuron.2016.07.006

Sreenivasan, K. K., Curtis, C. E., & D'Esposito, M. (2014). Revisiting the role of persistent neural activity during working memory. *Trends in Cognitive Sciences*, *18*(2), 82–89. https://doi.org/10.1016/j.tics.2013.12.001

Stanislaw, H., & Todorov, N. (1999). Calculation of signal detection theory measures. *Behavior Research Methods, Instruments, & Computers*, *31*(1), 137–149. https://doi.org/10.3758/BF03207704

Stokes, M. G. (2015). 'Activity-silent' working memory in prefrontal cortex: a dynamic coding framework. *Trends in Cognitive Sciences*, *19*(7), 394–405. https://doi.org/10.1016/j.tics.2015.05.004

Stokes, M. G., Wolff, M. J., & Spaak, E. (2015). Decoding Rich Spatial Information with High Temporal Resolution. *Trends in Cognitive Sciences*, *19*(11), 636–638.

https://doi.org/10.1016/j.tics.2015.08.016

Sturm, W., de Simone, A., Krause, B. J., Specht, K., Hesselman, V., Radermacher, I., … Willmes, K. (1999). Functional anatomy of intrinsic alertness: evidencefor a fronto-parietal-thalamic-brainstem network in theright hemisphere. *Neuropsychologia*, *37*(7), 797–805. https://doi.org/10.1016/S0028-3932(98)00141-9

Sturm, Walter, & Willmes, K. (2001). On the Functional Neuroanatomy of Intrinsic and Phasic Alertness. *NeuroImage*, *14*(1), S76–S84. https://doi.org/10.1006/nimg.2001.0839

Ten Brink, A. F., Biesbroek, M. J., Kuijf, H. J., Van der Stigchel, S., Oort, Q., Visser-Meily, J. M. A., & Nijboer, T. C. W. (2016). The right hemisphere is dominant in organization of visual search—A study in stroke patients. *Behavioural Brain Research*, *304*, 71–79. https://doi.org/10.1016/j.bbr.2016.02.004

Ten Brink, A. F., Van der Stigchel, S., Visser-Meily, J. M. A., & Nijboer, T. C. W. (2015). You never know where you are going until you know where you have been: Disorganized search after stroke. *Journal of Neuropsychology*. https://doi.org/10.1111/jnp.12068

Ten Brink, A. F., Verwer, J. H., Biesbroek, J. M., Visser-Meily, J. M. A., & Nijboer, T. C. W. (2017). Differences between left- and right-sided neglect revisited: A large cohort study across multiple domains. *Journal of Clinical and Experimental Neuropsychology*, *39*(7), 707–723. https://doi.org/10.1080/13803395.2016.1262333

Ten Brink, A. F., Visser-Meily, J. M. A., Schut, M. J., Kouwenhoven, M., Eijsackers, A. L. H., & Nijboer, T. C. W. (2017). Prism Adaptation in Rehabilitation? No Additional Effects of Prism Adaptation on Neglect Recovery in the Subacute Phase Poststroke: A Randomized Controlled Trial. *Neurorehabilitation and Neural Repair*, *31*(12), 1017–1028. https://doi.org/10.1177/1545968317744277

Theeuwes, J. (1992). Perceptual selectivity for color and form. *Perception & Psychophysics*, *51*(6), 599–606.

Theeuwes, J., Kramer, A. E., Hahn, S., & Irwin, D. E. (1998). Our eyes do not always go where we want them to go: Capture of the eyes by new objects. *Psychological Science*, *9*, 379–385.

Treisman, A. M. (1982). Perceptual grouping and attention in visual search for features and for objects. *Journal of Experimental Psychology: Human Perception and Performance*, *8*(2), 194–214. https://doi.org/10.1037/0096-1523.8.2.194

Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, *12*(1), 97–136. https://doi.org/10.1016/0010-0285(80)90005-5

Treue, S., & Maunsell, J. H. R. (1996). Attentional modulation of visual motion processing in cortical areas MT and MST. *Nature*, *382*(6591), 539–541. https://doi.org/10.1038/382539a0

van der Kemp, J., Dorresteijn, M., Ten Brink, A. F., Nijboer, T. C. W., & Visser-Meily, J. M. A. (2017). Pharmacological Treatment of Visuospatial Neglect: A Systematic Review. *Journal of Stroke and Cerebrovascular Diseases*, *26*(4), 686–700. https://doi.org/10.1016/j.jstrokecerebrovasdis.2017.02.012

Van der Maaten, L. J. P., & Hinton, G. E. (2008). Visualizing high-dimensional data using t-SNE. *Journal of Machine Learning Research*, *9*, 2579–2605.

van Ede, F., Niklaus, M., & Nobre, A. C. (2017). Temporal Expectations Guide Dynamic Prioritization in Visual Working Memory through Attenuated α Oscillations. *The Journal of Neuroscience*, *37*(2), 437–445. https://doi.org/10.1523/JNEUROSCI.2272-16.2016

van Ede, F., Szebényi, S., & Maris, E. (2014). Attentional modulations of somatosensory alpha, beta and gamma oscillations dissociate between anticipation and stimulus processing. *NeuroImage*, *97*, 134–141.

https://doi.org/10.1016/j.neuroimage.2014.04.047

Van Rossum, G., & Drake, F. L. (2011). *Python Language reference manual*. Bristol, UK: Network Theory Ltd.

Vogel, E. K., & Machizawa, M. G. (2004). Neural activity predicts individual differences in visual working memory capacity. *Nature*, *428*(6984), 748–751. https://doi.org/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. https://doi.org/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. https://doi.org/10.1037/0096-1523.32.6.1436

Vuilleumier, P. (2005). How brains beware: neural mechanisms of emotional attention. *Trends in Cognitive Sciences*, *9*(12), 585–594. https://doi.org/10.1016/j.tics.2005.10.011

Wandell, B. A., Dumoulin, S. O., & Brewer, A. A. (2007). Visual Field Maps in Human Cortex. *Neuron*, *56*(2), 366–383. https://doi.org/10.1016/j.neuron.2007.10.012

Wang, T.-Y., Huang, H.-C., & Huang, H.-S. (2006). Design and implementation of cancellation tasks for visual search strategies and visual attention in school children. *Computers & Education*, *47*(1), 1–16. https://doi.org/10.1016/j.compedu.2004.08.012

Warren, M., Moore, J. M., & Vogtle, L. K. (2008). Search Performance of Healthy Adults on Cancellation Tests. *American Journal of Occupational Therapy*, *62*(5), 588–594. https://doi.org/10.5014/ajot.62.5.588

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. *Journal of Neuroscience*, *32*(33), 11228–11240. https://doi.org/10.1523/JNEUROSCI.0735-12.2012

Weichselgartner, E., & Sperling, G. (1987). Dynamics of automatic and controlled visual attention. *Science*, *238*(4828), 778–780.

Weintraub, S., & Mesulam, M.-M. (1988). Visual hemispatial inattention: stimulus parameters and exploratory strategies. *Journal of Neurology, Neurosurgery & Psychiatry*, *51*, 1481–1488.

Westfall, T. C., & Westfall, D. P. (2006). Adrenergic agonists and antagonists. In J. G. Hardman, A. G. Limbird, & A. G. Gilman (Eds.), *Goodman and Gilman's the Pharmacological Basis of Therapeutics*. New York, USA: McGraw Hill.

Wheeler, M. E., & Treisman, A. M. (2002). Binding in short-term visual memory. *Journal of Experimental Psychology: General*, *131*(1), 48–64. https://doi.org/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), 11. https://doi.org/10.1167/4.12.11

Wilson, B. A., Cockburn, J., & Halligan, P. (1987). *Behavioural inattention test: manual*. Bury St. Edmunds, UK: Thames Valley Test Company.

Wilson, G., Aruliah, D. A., Brown, C. T., Chue Hong, N. P., Davis, M., Guy, R. T., … Wilson, P. (2014). Best Practices for Scientific Computing. *PLoS Biology*, *12*(1), e1001745. https://doi.org/10.1371/journal.pbio.1001745

Wimmer, K., Nykamp, D. Q., Constantinidis, C., & Compte, A. (2014). Bump attractor dynamics in prefrontal cortex explains behavioral precision in spatial working memory. *Nature Neuroscience*, *17*(3), 431–439. https://doi.org/10.1038/nn.3645

Witte, E. A., & Marrocco, R. T. (1997). Alteration of brain noradrenergic activity in rhesus monkeys affects the alerting component of covert orienting. *Psychopharmacology*, *132*(4), 315–323.

Wojciulik, E., Rorden, C., Clarke, K., Husain, M., & Driver, J. (2004). Group study of an “undercover” test for visuospatial neglect: invisible cancellation can reveal more neglect than standard cancellation. *Journal of Neurology, Neurosurgery & Psychiatry*, *75*(9), 1356–1358. https://doi.org/10.1136/jnnp.2003.021931

Wolfe, J. M. (1994). Guided Search 2.0: A revised model of guided search. *Psychonomic Bulletin & Review*, *1*(2), 202–238.

Wolfe, J. M., & Cave, K. R. (1999). The psychophysical evidence for a binding problem in human vision. *Neuron*, *24*, 11–17.

Wolff, M. J., Ding, J., Myers, N. E., & Stokes, M. G. (2015). Revealing hidden states in visual working memory using electroencephalography. *Frontiers in Systems Neuroscience*, *9*. https://doi.org/10.3389/fnsys.2015.00123

Wolff, M. J., Jochim, J., Akyürek, E. G., & Stokes, M. G. (2017). Dynamic hidden states underlying working-memory-guided behavior. *Nature Neuroscience*, *20*(6), 864–874.

Woods, A. J., Göksun, T., Chatterjee, A., Zelonis, S., Mehta, A., & Smith, S. E. (2013). The development of organized visual search. *Acta Psychologica*, *143*(2), 191–199. https://doi.org/10.1016/j.actpsy.2013.03.008

Woods, A. J., & Mark, V. W. (2007). Convergent validity of executive organization measures on cancellation. *Journal of Clinical and Experimental Neuropsychology*, *29*(7), 719–723. https://doi.org/10.1080/13825580600954264

Zhang, W., & Luck, S. J. (2008). Discrete fixed-resolution representations in visual working memory. *Nature*, *453*(7192), 233–235. https://doi.org/10.1038/nature06860

Zhang, W., & Luck, S. J. (2009). Sudden Death and Gradual Decay in Visual Working Memory. *Psychological Science*, *20*(4), 423–428. https://doi.org/10.1111/j.1467-9280.2009.02322.x

Zhang, W., & Luck, S. J. (2011). The Number and Quality of Representations in

Working Memory. *Psychological Science*, *22*(11), 1434–1441. https://doi.org/10.1177/0956797611417006

Zokaei, N., Heider, M., & Husain, M. (2014). Attention is Required for Maintenance of Feature Binding in Visual Working Memory. *Quarterly Journal of Experimental Psychology*, *67*(6), 1191–1213. https://doi.org/10.1080/17470218.2013.852232

Zokaei, N., Manohar, S., Husain, M., & Feredoes, E. (2014). Causal Evidence for a Privileged Working Memory State in Early Visual Cortex. *Journal of Neuroscience*, *34*(1), 158–162. https://doi.org/10.1523/JNEUROSCI.2899-13.2014

Zokaei, N., Ning, S., Manohar, S., Feredoes, E., & Husain, M. (2014). Flexibility of representational states in working memory. *Frontiers in Human Neuroscience*, *8*. https://doi.org/10.3389/fnhum.2014.00853

12. Acknowledgements

The past four years can be euphemistically characterised as eventful, and seem to have been ruled by two universal laws of nature: Those described by Murphy (“*Everything that can go wrong will go wrong*”) and Hofstadter (“*It always takes longer than you expect, even when you take into account Hofstadter's law*”). Apparently, this is the typical process for PhD students. For me, this process included many and more null results, several changes of thesis topic, testicular cancer, the death of our grant's PI Glyn Humphreys, and the immediate closing of Oxford's Experimental Psychology building.

Despite the adversity, it was – without a doubt – the best time of my life.

Although I rarely have problems starting new things on my own, this thesis would never have been finished if it wasn’t for the help of many others. First and foremost, I am indebted to Masud, who took the risk of hiring me. Thank you for all your sage advice (even if I didn’t always take it), for your faith in me, and for all your support.

I owe further thanks to the amazing original group (OG) of people who made the lab a warm and welcoming environment when I joined in November 2013, as a (not so very) humble MSc student. Laura (LORAH!), we survived several offices without windows and our PhDs together; you were a great support throughout my time in Oxford, and I hope a friend for life. Kinan, I loved talking about ~~crazy~~ brilliant business ideas, and hope we get to actually start Dr^2 at some point. Matt, you taught me my first bit of computational modelling, but more importantly you were a friend from the start. Sanjay, I learned so much from you, and your comment that I’d “make a great dictator” is perhaps the best compliment ever. Nahid, you’re a brilliant scientist, and an amazing person; I’ll never forget your and your mother’s kindness. Trevor, you kept lab outings classy, and your work ethos was an inspiration (as was your ability to drink flaming shots seemingly unaffected). Thanks for making me feel at home from the start.

Quite a few people have joined the lab since, and you were just as important for making the lab a great place. Matt Nour (a.k.a. nr 2), thanks for the discussions on working memory. Kathrin, it was an honour to be your first friend in Oxford, and I hope we’ll remain friends throughout our careers. Campbell, thanks for making me almost understand cricket, and for being a great office buddy. Yuen, I’ll miss seeing you watch Shin Chan in the office, and look forward to seeing you at the first *four skins* reunion. Annie, you’re one of the nicest people I’ve ever met (and I’m secretly in love with your husband). Annika, thanks for your friendship and for giving me the worst hangover of my life (and I’m secretly in love with your boyfriend too). Liv, thanks for your friendship, and for making me look slightly better by travelling even more than I did

these past few years. Amelia, you're an incredibly kind person. Sean, "*lover, poet, warrior*", thanks for all the great scientific and completely unscientific talks; I'm going to miss you buddy. Pat, I'm really impressed with your success, and please keep posting pictures of Softmax. Zargol, thanks for always being so cheerful and enthusiastic. Ellie, I had a blast at outreach events. Dan, I'm so impressed with how you combined starting a family and your job. Raquel, it's amazing how many facets there are to you: amazingly friendly and polite, but with some unexpected business ideas (please let me know when the prototypes are ready for testing). Tanja, you kept coming back to Oxford, and each time Oxford was better for it. Lisa, you were the youngest and newest, but you fit in so well! Youssuf, you came in as I left, but I enjoyed our short period of overlap. Claire, Gemma, and Rosyl, thanks for being my productive worker bees; you're all headed towards great careers. Thank you all for making the Husain lab the best lab ever.

As for my 'other lab', Eelke, Freek, Dante, Frida, Janina, Lev, MaryAnn, Michael, Nick, Paul, new Sam, Sofia, and Mark: Thanks for accepting me within your ranks. I thoroughly enjoyed the inspirational Stokes Lab meetings (beer-in-hand presenting!), learned so much from your EEG tutoring and discussions, and of course the barbecues and many Uni Club drinks.

There are more people in Experimental Psychology that made these past years great: The INDIREA crew of Alex (great catchups during my croquet and your allotment gardening), Eli (this thesis would NOT have been here without you), Iske (thanks for being a wonderful addition to the Husain lab, and for your great advice), Nir (already missing your monologues and our discussions on science and Jewish history), Rachel (thanks for being a friend throughout our PhDs; our Copenhagen trip was a highlight); my office mate Sam (greatly enjoyed our rants on students and stats); Glyn's old lab members Charu (may birds stay out of your way forever!), Florence (loved running the Wolfson bar with you), Nele (great job on continuing Glyn's important work), Nomi (will miss our random encounters and film nights); our inspiring new HoD Kia Nobre (thank you so much for writing all those recommendation letters). Additional thanks to all other members of our wonderful department: I don't have the space to mention you all by name, but rest assured that I won't easily forget you.

There are many international colleagues who greatly contributed to my PhD experience, but I can only mention a few here. Copenhagen felt like a second home thanks to Bart, Katsumi, and especially Chiron (daytime drinking just isn't the same without you). Lund felt like a third home thanks to Diederick, Ignace, and Kenneth. The INDIREA project was full of wonderful people that I haven't mentioned yet, but who did make our many meetings so very good: fellow kids Aurore, Adrian, Adriana, João, Kata, Méadhbh, and Natan, as well as responsible adults Anders, Iris, Signe, and all the others. Sebastiaan, I learned so much about open science and Python from you. Thank you collaborators (Chris, Ignace, Jeroen, Manon, Paresh, Sophie, Susan, Sylco, Tanja, Teuni, and especially Stefan). And last but definitely nowhere near least: Thank you Jim and Tim,

for sharing PhD misery and success, random nonsense, and incredibly offensive jokes (let’s hope WhatsApp never has a data breach).

I owe my friends and family for keeping me sane (at times) and happy (most times). Thanks opa and oma; mum and dad; and Elmer (you might be doing a real science, but I finished my PhD in fewer years than you did your undergrad, BOOM!). Thank you Kiki (and family), Joram, and Layla. Thank you Josh, for being a special friend for life. Thank you further members of Wolfson's best-ever croquet team Sarah, Yasser, and Niels. Thank you other Wolfson neighbours Erminia, Alicia, and Joshua. Thank you Jacob, for making me almost a half-decent chess player. You’re all amazing people.

Finally, and perhaps most importantly, thank you Asphodel, for being there while I was writing this thesis, for always believing in me, and for moving to a new city with me for the next steps (academic and otherwise).