Flexible use of attentional templates to guide adaptive behavior

Sage E. P. Boettcher

Somerville College

A thesis submitted in partial fulfilment of the requirements for the degree of Doctor of Philosophy at the University of Oxford

Abstract

Attentional templates are the mental representations that we use to facilitate efficient ongoing behavior, including the detection and identification of task-relevant sensory input. These templates are traditionally studied within the context of visual search. In this case, a template containing target-defining features is used to guide attention to possible targets. In this doctoral thesis, I will expand the concept of the attentional template in several important ways, on the basis of an inter-related series of empirical studies. I propose that templates carry information along several dimensions in addition to item identity, such as space, time, and associated actions. I further demonstrate that attentional templates are adaptable and action-oriented. They do not necessarily represent the veridical properties of the target item, but rather are adapted to better serve future behavior. Additionally, I show that attentional templates can be retrieved from associations in memory, rather than being explicitly provided as in more conventional tasks. In the General Introduction I review past literature related to attentional templates. Moreover, I list the axes I consider necessary for the study of attentional templates. Chapters 2 through 5 describe empirical studies probing the nature and flexibility of attentional templates. In Chapter 2, I investigate the behavioral benefits and electrophysiological signatures of memory-guided identity-based anticipation. In Chapter 3, I probe the extent to which attentional templates retrieved through memory associations can be adapted to task demands and differ from veridical representations. In Chapter 4, I explore whether and when prospective actions are incorporated into the attentional template. Finally, in Chapter 5, I test whether spatiotemporal information can be incorporated into an attentional template during a dynamic visual search task. Across these studies, the results reveal the flexible nature of attentional templates that are used to guide our behavior. In the General Discussion I place these results in a wider context and suggest potential avenues for continued research.

Acknowledgements

Considering this is the only section that most of you will read I better make it good. That being said, it is no easy task. I think it is appropriate to start by recognizing the absolute privileged position from which I am writing this thesis. Not only have I had the opportunity to learn and grow from some of the greatest scientific minds within our field, but throughout my admittedly short career I have been surrounded by kind, enthusiastic, inspiring people who have significantly changed me – and this thesis – for the better. This section, which makes up approximately 1% of my thesis, is insufficient for expressing the gratitude I feel towards each and every one of you.

I first want to start with a big thank you to my three incredible supervisors: Mark Stokes, Freek van Ede, and Kia Nobre. Mark, thank you for continuing to find ways of supporting me, despite my absolute reluctance to bike into town. Your ideas and creativity were a source of motivation throughout my DPhil. Freek, you are always the first person I want to turn to when the results of an experiment start coming in. When the results look good, your excitement is contagious. When the results look bad, you usually find a way for them to look good again. It has been nothing short of fun asking – and sometimes answering – questions with you. Kia, your brilliance shines through in every interaction, and not only because of your *prodigious* vocabulary. This brilliance is matched by your kindness and humor which helped to quickly break down any of the meeting-Kia-Nobre-jitters I may have initially felt. Although I am completely aware that I am just one of dozens of scientists that owes you an incredible debt of gratitude, this gratitude feels deeply unique and personal. My thinking and curiosity have been irreversibly changed under your guidance. With every interaction we have, I come away better. You have shaped me as a scientist and a person, and I will forever be grateful for the role you continue to play in my life. I look forward to years of learning from you.

In addition to my advisors, there are great number of collaborators, colleagues, and friends who have shaped my time at Oxford. Without you all, my experience here would have been far less. Thanks to Nir, Dani, Jasper, Sammi, Lauren, Giedre, Oana, Nahid, Katharina, Quinn, Catharina, Marcus, Alex, Sam, Frida, Yvonne, Cam, Ryan, Sana, Tom, and Nadia. Each of you has had a part in this thesis whether it was via research or beers at The Star (arguably more important at times). Your friendship and support have had an immeasurable impact. I also wish to extend a special thanks to Sarah Clayton, whose presence in the lab has redefined how we function for the better. Your support has improved the science and welfare of the entire lab, thank you.

I have taken a relatively slow path to this point, and because of this I have collected some amazing mentors outside of Oxford as well. Thanks to Jim Hoffman, who helped to spark my passion for research ten years ago. To Jeremy Wolfe, who continues to be a glowing inspiration. The two years I spent in your lab represent perhaps the steepest learning curve of my career, and at every point along that curve you were supportive and kind. To Melissa Võ, leaving SGL felt like leaving home. You have been a mentor, a friend, and a younger sister. The conversations on your couch defined for me what scientific discourse should be, and you are one of my favorite people to think with.

To my family, Mom (Gaye), Dad (Mitch), Marlowe, Aunt Dawn, Joan, Michael, Maya, Nate, Nikki, and the Draschkows, thank you for your unwavering support and genuine curiosity throughout my career. Science has taken me across the world, yet you have never questioned my choices. I love you all and am incredibly grateful to you for making this journey so fun.

I feel like I have run out of *thanky* words, yet I am left with the person who deserves them the most. To my best friend, my partner, and my coauthor in life, Dejan Draschkow, you make me better. I have grown up with you in science and all of this would be a LOT duller without you. Thank*s* you. I love you strong.

Statement of Contribution

This thesis reflects the collaborative nature of science. In this light, it is of note that this work was conducted together with colleagues both in and outside of my supervisory team. Chapter 2 of this thesis was designed in collaboration with Mark Stokes, Kia Nobre, and Freek van Ede. I collected the data for this experiment and ran the analysis. This work was published and the final manuscript can be found here https://doi.org/10.1523/JNEUROSCI.2751-19.2020. Chapter 3 of this thesis was conducted together with Freek van Ede and Kia Nobre. Again, the experiments were designed in collaboration, while I collected the data and lead the analysis. I designed Chapter 4 together with Freek van Ede and Kia Nobre. I collected the data together with Daniela Gresch and I lead the analysis under the supervision of Freek van Ede. Finally, the experiments in Chapter 5 were designed together with Nir Shalev, Kia Nobre, and Jeremy Wolfe. I collected the data together with Nir Shalev and I lead the analysis for these experiments. The writing in this thesis is my own, and was carried out under the supervision of Mark Stokes, Freek van Ede, and Kia Nobre.

Contents

“It’s a poor sort of memory that only works backwards”

-Lewis Carroll, Through the Looking Glass

1. General Introduction

1.1. General Overview

Although memories represent the past, they are used to guide future behavior. You wake up in the morning and blindly reach for your phone next to your bedside. While preparing your breakfast you turn your attention to the toaster just before it pops. Mid-breakfast you realize you are late and begin a frantic search for your keys with their distinctive red key chain. The location of the phone, the timing of your toast, and the features of the key chain are not just properties of the external world; they are also internal representations. It is hard to imagine how we could navigate everyday challenges without the guidance of our memories, as this memory-based guidance informs where, when, or what should be prioritized in our surroundings to enable successful natural behavior.

Understanding the nature of the mental representations that guide our behavior has been a goal of cognitive psychologists since the onset of the field. In contemporary psychology these representations are referred to as "*attentional templates"*. It is perhaps simplest to understand these templates in the context of visual search. As you look for your keys in the morning, some internal information – a mental representation of what your keys look like – must guide your focus in the external world. For truly efficient behavior, such internal representations should be flexible in order to adapt to changes in our environment and goals.

Moreover, although it may be intuitive to conceptualize attentional templates as analogous to the percept of an object – your keys – this need not be the case. I will consider whether and how templates may include other attributes useful for identifying an object such as spatial and temporal information or the action it affords, and explore the extent to which template features can be selectively distorted according to what is most useful for guiding adaptive behavior.

There are several dimensions along which we can categorize attentional templates. First, what is a template used for? How long is the template available for exploitation? In what state are templates stored: need it be an *active* state or can they also *silently* guide our behavior? Is a template veridical or adaptive, representational or functional? Are action plans attached to a template? Many of these questions remain open in the literature, yet their answers are essential for understanding the principles by which memory guides behavior.

The empirical work in this doctoral thesis focuses on understanding how prior knowledge – or memories – can be used flexibly to guide behavior through various domains (identity, time, space, and actions) in healthy humans. More specifically, I have investigated and established the flexible nature of memory-based attentional templates. This experimental work has been conducted in four main branches, investigating in turn: (1) the behavioral benefits and electrophysiological signatures of memory-guided identity-based anticipation, (2) the extent to which attentional templates retrieved from memory associations can be functionally biased, and thereby distorted from veridical representations, (3) the action plans associated with visual attentional templates, and (4) the spatial and temporal information within a template and their proactive utilization within visual search. This work has relied on electroencephalography (EEG) to probe the neural dynamics of these processes as well as on several behavioral markers including eye-movements.

Before reporting the empirical portion of my work, I will start with a review of attentional templates. This will include an overview of how I define an attentional template, the proposed content or dimensions within an attentional template, an introduction to several key axes which can be used to describe templates, and the associated neural mechanisms. I conclude this introductory chapter by outlining the aims of the current thesis.

1.2. Defining attentional templates

1.2.1. From attention to templates

When writing about attention, it is difficult to decide whether to begin with a William James quote, or an example of our senses being bombarded with input. I will do both. As you continue to read through this thesis, you must resist the urge to look away as various environmental stimuli call for your attention. Attention is an important mechanism in perception to cope with the deluge of information that enters through our senses. Our inherent resource limitations have been conceptualized as a “bottleneck”, which filters the large amount of sensory information for successful perceptual processing. Such a capacity limit may result from the energy cost incurred by neuronal spiking in the cortex (Lennie, 2003). Given this biological constraint, an efficient attentional filtering system is necessary for perception.

For every cognitive psychologist you may ask, it is likely you will find a slightly different definition of attention. Admittedly, it is a weakness within the field, as we seem to rely too heavily on James’ notion that *everyone knows what attention is*. For the current work, I will use a definition provided by Nobre and Kastner (2014) stating that attention is “the selection and prioritization of processing information that is relevant to current task goals”. The benefits of attention on perception are clear. Attended stimuli are processed faster and with higher resolution than their unattended counterparts (Eriksen & Hoffman, 1973; Nobre & Kastner, 2014; Posner, 1980). Attention additionally enhances neural responses of a stimulus (Chelazzi, Miller, Duncan, & Desimone, 1993; Desimone & Duncan, 1995; Mangun & Hillyard, 1987). However, our understanding of attention has evolved from an early dichotomy in the literature, which suggested attentional selection occurs either early (Broadbent, 1958) or late (Deutsch & Deutsch, 1963) during sensory processing, towards a more flexible view suggesting multiple possible paths for selection and prioritization.

Attentional selection has traditionally been thought of as a combination of top-down, user-driven priorities, and bottom-up, stimulus-driven properties of the environment (Jonides, 1998). Bottom-up attention is an automatic process that shifts attentional resources to particularly salient events in the environment, such as a sudden flash. Top-down attention refers to the voluntary orienting of attention for goal driven purposes (Posner, 1980). The diversity in the content of top-down information guiding attention has been of great interest over the past several decades. One can allocate attention on the basis of features, objects, regions in space, moments in time, and more (Chelazzi et al., 1993; Chun & Jiang, 1998; Coull & Nobre, 1998; Desimone & Duncan, 1995; Jiang, 2018; Posner, 1980; Võ, Boettcher, & Draschkow, 2019). Attention is further sharpened through predictions of our environment across all of the aforementioned dimensions. Although the current work will focus on the allocation of attention to the external world, it should be noted that attention can additionally be allocated to internal representations (Griffin & Nobre, 2003).

More recently, the notion that our memories constitute a third important source of attentional selection has found widespread acceptance. Although most memory research to date has focused primarily on how memories represent the past, it is becoming increasingly clear that an important role of memories is to guide future behavior (Nobre & Stokes, 2019). Indeed, our memories, over multiple time scales, can serve as the basis of environmental predictions guiding attention (Nobre & Stokes, 2019). For example, selection history – which refers to the idea that previous attentional deployments can bias future attentional selection independent of current goals or stimulus saliency – has been proposed to have a major influence on visual-search performance (Awh, Belopolsky, & Theeuwes, 2012; Theeuwes, 2019).

I will refer to the memory-based predictions, which can help tune our attention, as attentional templates. The interest in attentional templates is gaining momentum. Attentional templates are an important way through which incoming sensory information can be biased by attention. For example, imagine you are searching through your apartment for your keys. The information that guides your search from the nearby shelf, where you usually leave them, to your partner's keys on the kitchen table, which have a striking resemblance to your own keys, is contained within the attentional template. Further, there must also be an internal representation of our ‘target’ that allows us to reject our partner’s keys as a ‘foil’ and continue our search. Since attentional templates are not present in the physical world when they are being utilized, they must be stored in some form of memory.

The role of templates in guiding attention bears resemblance to the predictive-coding theory of perception. Such theories assume the brain is continuously generating and updating predictions in order to anticipate the state of our surroundings (de Lange, Heilbron, & Kok, 2018; Friston, 2010; Friston & Kiebel, 2009; Rao & Ballard, 1999; Trapp & Bar, 2015). These predictions arise from higher brain areas before joining with information brought in from bottom-up sources. As such, the sensory information must only be checked against the preexisting prediction – or attentional template – for discrepancies (prediction errors). Theories of predictive coding propose that only errors are signaled through the hierarchy and used to improve the quality of future predictions. The idea of perception as an inferential process – combining our memories and knowledge of the world with sensory input – is not a new one. Helmholtz’s (1866) description of *unconscious inference* proposed that the brain fills in missing visual information. These ideas are contingent on our ability to *predict* the information most relevant to the current environment or goals. These memory-based templates play an important role in the allocation of attention for efficient perception. In this thesis, the focus will not be on inference nor surprise; but rather on the preparation for guiding perception in demanding settings that require attention.

Henceforth, when referring to attentional templates, I will be referring to the prior information content associated with an anticipated perceptual event that is used to guide the perceptual analysis leading to said event. Such preparational content can also be referred to as proactive memory.

1.2.2. What we know about attentional templates from visual search

One of the most common tasks used to assess the contents and fidelity of an attentional template has been *visual search* (Carlisle, Arita, Pardo, & Woodman, 2011; Desimone & Duncan, 1995; Duncan & Humphreys, 1989; Gunseli, Meeter, & Olivers, 2014; Wolfe, Horowitz, Kenner, Hyle, & Vasan, 2004). In visual search, the concept of an attentional template is quite clear. Observers have a particular target in mind (red vertical bar), which they must find among distractors that may or may not share features of the target (red horizontal lines, & blue vertical lines). In this example, the attentional template is necessary for completing the task. Without the foreknowledge of target defining features, observers would be at chance at reporting the target's presence or absence.

The idea of an attentional template was first put forth outside of the realm of human cognition. Specifically, it was proposed by zoologists who noticed that birds tend to sample disproportionately more food from samples that share common features. This suggests their feeding behavior is biased towards target features previously associated with reward (Bond, 1983; Tinbergen, 1960). In experimental psychology, the term "template" with regards to visual search was first used by Duncan and Humphreys (1989). They described *'a process of selection by matching input descriptions against an internal template'*. Since then, this idea has been pervasive in the visual search literature (although, see section *1.4.1 Guiding vs. matching* for a discussion of the different types of attentional templates within visual search).

Theories of attentional orienting during visual search propose that we hold an attentional template in working memory in order to bias activity in neurons involved in processing features relevant to the perceptual task (Bundesen, Habekost, & Kyllingsbæk, 2005; Carlisle et al., 2011; Desimone & Duncan, 1995; Gunseli, Meeter, et al., 2014; Wolfe, 1994). Behaviorally, it has been shown that working-memory content can bias attentional deployment during visual search. Distractor items that match the contents of working memory attract eye-movements (a physiological response strongly correlated to attentional selection) and slow reaction times (memory-driven attentional capture) (Beck, Hollingworth, & Luck, 2012; Dalvit & Eimer, 2011; Downing, 2000; Olivers, 2009; Olivers, Meijer, & Theeuwes, 2006; Olivers, Peters, Houtkamp, & Roelfsema, 2011; Soto, Heinke, Humphreys, & Blanco, 2005; Soto, Hodsoll, Rotshtein, & Humphreys, 2008). Single-unit recordings from monkey inferotemporal cortex (IT) support the claim that attentional templates are stored in working memory, as neurons processing target-specific attributes in IT showed increase firing rates during the delay period between the presentation of a target template and the search (Chelazzi, Duncan, Miller, & Desimone, 1998; Chelazzi et al., 1993). In human electrophysiology, a lateralized potential linked to spatially selective working memory, the contralateral delay activity, is observed (Carlisle et al., 2011; Gunseli, Meeter, et al., 2014; Gunseli, Olivers, & Meeter, 2014).

Recent work has also implicated long-term memory in guiding visual search. Specifically, Wolfe (2012) showed that observers are remarkably accurate and efficient at searching for any one of one-hundred items held in memory. These items clearly cannot live in a capacity-limited store such as working memory. Instead, it has been proposed that these items are held in activated long-term memory (ALTM; Cowan, 1995) – a concept referring to a subset of long-term memory that is currently task relevant. It is of course still possible that objects in LTM are shuffled in and out of working memory, alternating in their role as the attentional template. However, if this were the case, the expected response times – and related eye-movement metrics – would increase linearly with the number of items held in memory. Instead, it has been repeatedly shown that these response times increase with the log of the memory set-size – that is, searching for the second item in memory costs much more than searching for the 100th (Boettcher, Drew, & Wolfe, 2018; Boettcher & Wolfe, 2015; Drew, Boettcher, & Wolfe, 2017; Wolfe, 2012; Wolfe, Drew, & Boettcher, 2015) implying an efficient search through a large memory store. Moreover, memory set size and working-memory load do not interact significantly, indicating that visual working memory may play a different role in this context. Specifically, it was proposed that working memory may serve as the shuttle between items in the environment and ALTM (Drew, Boettcher, & Wolfe, 2016). Taken together, it is important to also consider long-term memory as a potential store for attentional templates. In Chapters 2 and 3, I will investigate the neural properties and adaptability of attentional templates retrieved from long-term memory associations.

Finally, the quality of the attentional template may vary. Consequently, the attentional guidance it allows will also vary. When observers have a more detailed attentional template (pictures vs. description; exact target vs. category; 1 vs. 2 targets) they are faster and more accurate in their search (Hout & Goldinger, 2015; Malcolm & Henderson, 2009; Vickery, King, & Jiang, 2005; Wolfe et al., 2004; Yang & Zelinsky, 2009). When a more detailed template representation was required – e.g. finding a teddy bear among other teddy bears compared to among non-bear images – the amplitude of the CDA was larger. This has been interpreted as a more detailed template representation being held in working memory (Schmidt & Zelinsky, 2017).

Taken together, visual search tasks can be extremely fruitful in the study of attentional templates. In Chapter 5 of this thesis I will use a modified visual search task to investigate whether spatiotemporal information is also contained within the attentional template, and how different memory systems may contribute to this information. In chapters 2-4 I will reveal new experimental insights into attentional templates studied outside of the context of classic search tasks. To this end, I next discuss attentional templates in non-search paradigms.

1.2.3. Attentional templates and expectations for perception

It has become increasingly recognized that sensory experiences are supported – and sometimes distorted (Purves, Shimpi, & Lotto, 1999) – by sensory predictions. These expectations can be derived from regularities within the environment and can be used to resolve uncertainties in our perception. Our environment is complex and noisy. In the visual domain, this is obvious in cases of suboptimal illumination, occluded objects, novel viewpoints, motion, and various other factors which may stand in the way of perception. However, efficient behavior depends on our brain making sense of this noise quickly and accurately.

Attentional templates are vital for efficient perception. In this form, expectations can bias perception in the face of impoverished bottom-up input, affecting not only the quality but also the content of the perception. For example, the perceived directionality in a cloud of dots with weak coherence or contrast is biased by the expectation of the direction (Chalk, Seitz, & Seriès, 2010; Kok, Brouwer, van Gerven, & de Lange, 2013). Our model of the world and regularities within it can also explain many visual illusions (Lupyan & Clark, 2015). For example, a lifetime worth of experiences tells us that light comes from above, so when confronted with a circle with shading on the top this is perceived as concave.

Attentional templates have been shown to alter perceptual processing at the earliest stages of the cortical hierarchy (Kok, Failing, & de Lange, 2014; Kok, Mostert, & De Lange, 2017; Kok, Rahnev, Jehee, Lau, & De Lange, 2012; O'Reilly et al., 2013; C. Summerfield, Trittschuh, Monti, Mesulam, & Egner, 2008; Todorovic & de Lange, 2012; Todorovic, van Ede, Maris, & de Lange, 2011). There is evidence for content-specific activation of an expected stimulus even in the absence of that stimulus (Kok et al., 2014; Macknik & Haglund, 1999). Specifically, when a particular grating is expected after a particular tone, yet it fails to appear, the pattern of activity in V1 is analogous to that evoked by the stimulus itself (Kok et al., 2014). This suggests that attentional templates are available at the early stages of processing to efficiently process expected input. Moreover, there is recent evidence suggesting this content-specific representation is available even before the onset of an expected stimulus (Kok et al., 2017). Specifically, the researchers showed significantly better decoding for validly cued targets compared to invalidly cued targets shortly (~40 ms) before the actual onset of the target.

It is clear that attentional templates hold a critical role in perception. In Chapter 2 of this thesis I examine the behavioral and neural correlates of attentional templates that are used for detection of a stimulus that is difficult to perceive. Importantly, attentional templates can also be useful for perceptual functions dependent on higher levels of processing within the hierarchy. I will explore how attentional templates are used in object recognition in the next section.

1.2.4. Attentional templates for object recognition

Attentional templates play a critical role in theories of object recognition. It is known that our visual system is remarkably efficient at recognizing our environment and the objects within it. Potter (1975) showed that observers could identify a target scene among a stream of distractors with a presentation time of only 125 ms. In a more recent magnetoencephalogram (MEG) study, observers were asked to identify whether a scene contained either a car or a person (Kaiser, Oosterhof, & Peelen, 2016). Despite substantial variance and clutter among the scenes, target category information was decodable as early as 160 ms. Similar processing speeds – primarily attributable to high level conceptual features (Greene & Hansen, 2020) – have been found in other categorization tasks (Greene & Hansen, 2020; Thorpe, Fize, & Marlot, 1996). This remarkable processing speed implies a multi-directional perceptual system, wherein higher order information is passed to early sensory areas in order to ease the computational load. This is in opposition to a strictly bottom-up hierarchical view in which processing begins from the sensory organs and gradually progresses through the sensory areas of the brain, increasing in complexity at each step (Hubel & Wiesel, 1962).

Much of the empirical work on attention purposefully uses simple stimuli void of meaning. As such, we are able to understand a ‘pure’ form of attentional selection, absent of the influence of prior knowledge. In contrast, perception outside of a laboratory setting unfolds in natural environments rich in context. We exploit regularities within our surroundings in order to facilitate perception. Objects in coherent contexts are detected and located faster (Biederman, Mezzanotte and Rabinowitz, 1982; *For a recent review see:* Võ, Boettcher and Draschkow, 2019). This is even true when the context itself is unrecognizable, yet maintains the summary statistics of a scene that is consistent with the object (Lauer, Cornelissen, Draschkow, Willenbockel, & Võ, 2018). Thus, context itself may activate attentional templates. For example, observers are able to recognize ambiguous objects when they are placed next to easily recognizable objects (Bar & Ullman, 1996). In one model, it is proposed that rapid processing of coarse low-spatial frequency sensory input carries contextual information and leads to specific sensory predictions. These higher-order predictions are projected to the inferior temporal cortex where they are met with the incoming detailed sensory information (Bar, 2004).

Wolfe and colleagues (2011) propose a dual-pathway model for visual search in scenes, which also provides key insights into the utility of attentional templates. A selective pathway is used for feature binding and object recognition. This is slow and limited in its capacity, causing a bottleneck in the process of search. Attentional templates may act as the gatekeeper to this bottleneck – such that items which are more likely to be targets are selected for consideration at this stage. These attentional templates may be additionally informed by a nonselective pathway which quickly extracts global scene statistics allowing for semantic processing – without object recognition. In addition to auxiliary spatial guidance that a scene context may provide (you will search for a badger on the ground while a bird is more likely to be found in the sky), target-related objects may be incorporated into the attentional template as well. For example, if you are searching for chalk in the classroom, your attention is likely to be guided by the features of the target (small white cylinder) as well as features of anchor objects that hold specific spatial predictions about the target (chalk board: large green/black rectangle) (Boettcher, Draschkow, Dienhart, & Võ, 2018).

Through this work on object and scene processing, it becomes clear that prior information– in the form of attentional templates – is critical in both simple object recognition as well as more complex tasks such as everyday visual search. Objects do not exist in isolation. Albeit complex, our environment is organized and rule governed. The regularities within our environment may operate as the foundation to many of the attentional templates that serve natural behavior.

1.3. The dimensions of attentional templates

Attentional templates are multi-dimensional. More specifically, attentional templates not only contain the identity – the expected object or features – but also information regarding *where* and *when* an object is likely to occur, as well as its associated actions (see Figure 1.1). This additional information is vital for efficient processing in a dynamic world. Here I will outline and discuss four basic dimensions of a template – identity, space, time, and associated actions.

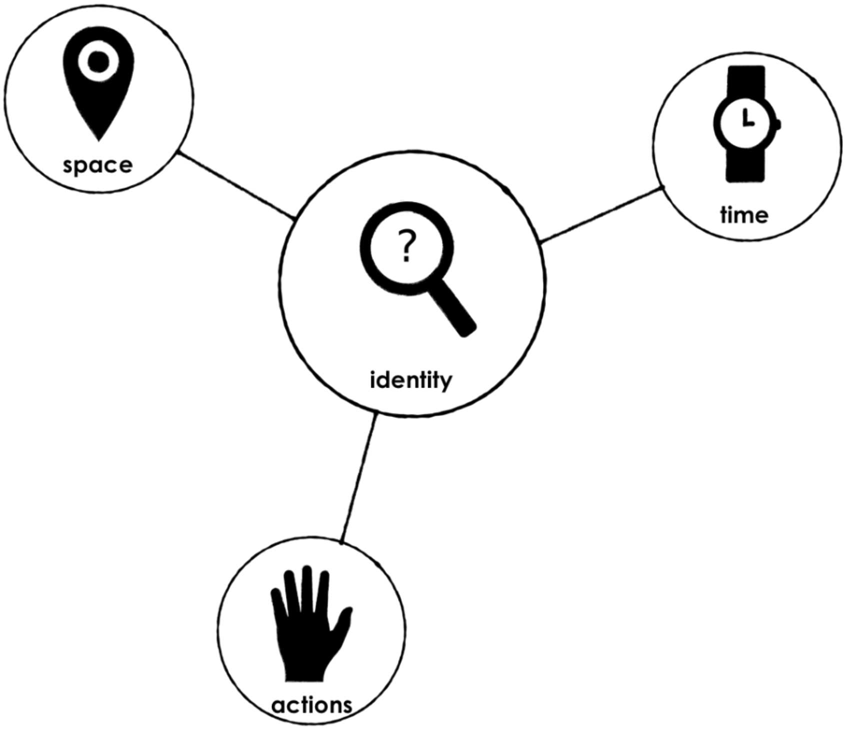

Figure 1.1: The dimensions of attentional templates. At the core of an attentional template is identity-based information. Additionally, representations of space, time, and associated actions can accompany the attentional template leading to more efficient natural behavior. I will discuss the identity of attentional templates in Chapters 2 and 3, the associated actions in Chapter 4, and spatiotemporal information in Chapter 5.

1.3.1. Identity

Identity-based information – knowing *what* to expect – is at the core of an attentional template and can either be in the form of an entire object or a particular defining feature. Feature- and object-based information can drive attention. Efficient visual search depends on the distinctiveness between target and distractor features (Geng & Witkowski, 2019; Treisman & Gelade, 1980; Wolfe, 1994). This indicates that attention is drawn towards distractors that share features with the target. Search slopes – the increase in reaction times by visual set size – vary systematically according to identity-related distinctiveness. When distractors and targets are unique within a feature dimension (red bars among green bars) this is considered a *singleton search* and has a flat search slope. However, if a target is defined across multiple features (green-vertical among green-horizontal and red-vertical bars) a *conjunction search* results, and search times increase with the number of distractors (Treisman & Gelade, 1980). This work was critical in establishing the idea of feature guidance.

This low-level feature information can be extended to more complex identity information. In monkey electrophysiology, neurons responsive to a specific target identity show increased firing during the delay period after a cue indicates that its preferred stimulus is relevant for the search in the upcoming display (Chelazzi et al., 1998, 1993). In an fMRI experiment, observers prepared to search for either people or cars in novel complex scenes. In the period during search preparation, activity in object-selective cortex was similar to that during the percept of an exemplar from the actual category. This points to category-level information available in the attentional template (Peelen & Kastner, 2011). Although there is evidence for proactive identity-based information from fMRI decoding studies (Peelen & Kastner, 2011; Stokes, Thompson, Nobre, & Duncan, 2009), human electrophysiological studies have either failed to show template activation or have found this activation only very briefly before the onset of the target stimulus (Kok et al., 2017; Myers et al., 2015; Wen, Duncan, & Mitchell, 2019).

Expectations on the basis of a particular feature can drive behavioral and neural responses in tasks other than visual search (Chalk et al., 2010; Kok et al., 2013; Liu, Stevens, & Carrasco, 2007; Maunsell & Treue, 2006; Sàenz, Buraĉas, & Boynton, 2003). Treue and Martinez Trujilo (1999), showed that feature-based attention enhanced responses of motion-selective neurons in MT that preferred the attended feature (e.g. upward motion). These effects have also been demonstrated in human fMRI (Greenberg, Esterman, Wilson, Serences, & Yantis, 2010; Liu et al., 2007; Sàenz et al., 2003). Taken together this work demonstrates the utility of feature-based information in the attentional template.

Object-based accounts suggest that attending a specific object improves the quality of the sensory representation. Foreknowledge of a relevant object category can lead to selective enhancement in areas specific to the object type (faces: fusiform face area (FFA), houses: parahippocampal place area (PPA)) (O'Craven, Downing, & Kanwisher, 1999). Baldauf and Desimone (2014) presented observers with temporally and spatially overlapping images of faces and houses, and asked observers to attend to one object category at a time while recording MEG. The researchers found that object-based attention induced synchronization within the gamma band between inferior frontal junction (IFJ) and the relevant sensory area – FFA or PPA. They showed that this gamma synchronization was driven by IFJ, indicating a potential mechanistic role in object-based attention.

Taken together, there is evidence for feature- and object-based guidance of attention. This *identity-based* information – that is, knowing what to expect – likely plays a fundamental role within the attentional template. The expectations may be in the form of simple features – such as color, motion, direction, or size – or may be more complex on the level objects and category information. In this thesis, object-based expectations are used in all empirical chapters. However, this particular dimension serves as the primary focus of studies in Chapters 2 and 3.

1.3.2. Space

Attention can also be guided by proactive spatial predictions – knowing *where* to expect something. There is a vast literature on spatial attention. Here I will focus only on attention guided proactively by memories coupled with identity information and therefore contained within the attentional template. Mechanistically, spatial information within the attentional template can be understood through the retinotopic organization of visual cortex. That is, attention is directed through gains within neurons whose receptive fields correspond to the to-be-attended location (Luck, Chelazzi, Hillyard, & Desimone, 1997; McAdams & Maunsell, 1999; Reynolds, Pasternak, & Desimone, 2000).

Spatial information can accompany other sources of information within the template. For example, in probabilistic cuing tasks, observers are asked to search for a particular target (T) amongst random distractors (Ls). Although target features are explicit and therefore clearly exert guidance during this search, there is a secondary implicit spatial component in this task as well. Namely, within the search space there are high-probability target locations, and low-probability target locations. Observers become faster at finding targets in these high-probability target locations (Geng & Behrmann, 2002; Jiang, 2018; Shaw & Shaw, 1977). This effect cannot be explained by inter-trial priming alone (Goschy, Bakos, Mueller, & Zehetleitner, 2014; Jiang, Swallow, Rosenbaum, & Herzig, 2013; J. L. Jones & Kaschak, 2012) indicating that long-term spatial memory co-occurs with feature-based information within the attentional template.

The spatial information is robust and often persists outside of the initial task context. For example, in one probability-cuing experiment the task was divided into two phases (Jiang, Swallow, Rosenbaum, et al., 2013). In the first phase, one region of the search display more often contained the target. As expected, search was facilitated to these high-probability locations. In the second phase, there was no difference in the probability distribution of target locations. Yet, search for targets occurring in the formerly high-probability region was faster, indicating a robust spatial prediction that continued to be effective with the changed context (Jiang, Swallow, Rosenbaum, et al., 2013). In a related line of work, observers searched for targets embedded within pictures of real-world scenes (J. J. Summerfield, Lepsien, Gitelman, Mesulam, & Nobre, 2006). Some of the scenes repeated with the same target location throughout the experiment. Naturally, observers became faster for these repeated scenes. Interestingly, these long-term memory associations extended to an attentional orienting task in which a target was briefly presented in either a scene with (valid) or without (neutral) a long-term memory association. This task has been replicated several times and is able to separate effects associated with learning and those associated with utilization of spatial information in long-term memories to guide attention (Nobre & Stokes, 2019; Patai, Doallo, & Nobre, 2012; J. J. Summerfield et al., 2006). Neuroimaging studies reveal the spatial information contained within the attentional template is associated with increased hippocampal activity and engages the frontoparietal network – known for its role in orienting spatial attention (Stokes, Atherton, Patai, & Nobre, 2012; J. J. Summerfield et al., 2006).

Within human electrophysiology, alpha-band modulations are a prevalent marker for the distribution of spatial attention. Specifically, the amplitude of alpha oscillations in visual cortex is reduced contralateral (vs. ipsilateral) to an attended spatial location (Jensen & Mazaheri, 2010; Sauseng et al., 2005; Thut, Nietzel, Brandt, & Pascual-Leone, 2006; Worden, Foxe, Wang, & Simpson, 2000). More recent work has shown that the attended region can be decoded from alpha-band activity (Foster, Sutterer, Serences, Vogel, & Awh, 2017; van Moorselaar & Slagter, 2019). This lateralization is a proactive process, and present before the onset of the to-be-attended stimuli (Wallis, Stokes, Cousijn, Woolrich, & Nobre, 2015; Worden et al., 2000). When alpha rhythms are experimentally adjusted through transcranial magnetic stimulation (TMS), similar attentional effects are observed suggesting a causal role of alpha band modulations in directing spatial information from the attentional template (Herring, Thut, Jensen, & Bergmann, 2015). Further, alpha lateralization has also been shown with shifts in spatial attention on the basis of long-term memory representations (Stokes et al., 2012; J. J. Summerfield, Rao, Garside, & Nobre, 2011)

Although spatial attention is the focus of a great deal of research, space is less often considered as a dimension within attentional templates. Here I have outlined several examples of how space may be incorporated in the attentional template. In Chapter 5 I will discuss whether both space and time can proactively facilitate visual search.

1.3.3. Time

Attention, a highly dynamic process in and of itself, can also be distributed to particular moments in time – knowing *when* something will happen (Nobre & van Ede, 2017). In contrast to feature-based, object-based, and spatial attention, temporal attention cannot directly be explained by either the retinotopic mapping of cortex or the receptive-field properties of neurons related to various features and object properties. Even so, we know that it is possible to distribute attention dynamically across time. In an adaptation of Posner's seminal work on spatial orienting, Coull and Nobre (1998) showed that when observers are provided with a cue indicating when a target will be presented they are faster at reporting that target. These effects are often linked, at least in part, to other forms of predictions. For example, when spatial and temporal predictions are pitted against one another, it has been shown that valid temporal expectations coupled with valid spatial expectations facilitates behavior. Yet, the temporal expectation benefit was eliminated with invalid spatial cues (Rohenkohl, Gould, Pessoa, & Nobre, 2014).

There are various forms of temporal expectations. As described above, these expectations can be generated explicitly through cues indicating when a target is likely to appear (Coull & Nobre, 1998; Rohenkohl et al., 2014). Alternatively, rhythms can provide a natural source of temporal expectations. That is, when events repeat with a set temporal interval, this provides our attentional system with an ecological cue for the next event (Nobre & Rohenkohl, 2014). For example, when an auditory target aligns with a rhythmic tone, observers are more accurate in their perception of the target (M. R. Jones, Moynihan, MacKenzie, & Puente, 2002). Finally, if you expect something to happen, and it has not yet happened, your expectation of that event occurring will change with time. This is referred to as the "hazard function", and indeed behavior reflects this intuition showing faster response time for longer delay intervals (Coull, 2009; Elithorn & Lawrence, 1955; Niemi & Näätänen, 1981). Importantly, the size of attention-related modulation of neuronal activity adjusts with changing expectations for target occurrence, even when this function is non-monotonic (Ghose & Maunsell, 2002; Janssen & Shadlen, 2005).

Finally, much like spatial attention, temporal attention can also be guided through long-term memory associations. Cravo and colleagues (2017) presented targets on scenes after either a short or long interval and repeated these associations across the experiment. Detection and discrimination of targets improved when they appeared at the learned interval. Moreover, this long-term memory based temporal association was accompanied by a contingent negative variation (CNV) – an electrophysiological marker commonly associated with temporal anticipation (Miniussi, Wilding, Coull, & Nobre, 1999). This work provides compelling evidence that expected temporal onsets proactively guide attention and, when available, contribute critical information to attentional templates. I will expand upon our understanding of temporal information within the attentional template in Chapter 5.

1.3.4. Associated Actions

Although the fundamental purpose of attentional templates is to efficiently guide future behavior, these templates have usually been considered without regard for prospective actions. Within the visual search literature, actions are often equated across trials with a simple present/absent response or an orientation response (left/right) that cannot be prepared for until the target is found (Boettcher, Draschkow, et al., 2018; Boettcher, Drew, et al., 2018; Carlisle et al., 2011; Chun & Jiang, 1998; Treisman & Gelade, 1980; Wolfe et al., 2004). The same is true in the study of visual working memory, in which retention is typically considered in the absence of associated actions (van Ede, 2020). In fact, action planning is often deliberately removed from laboratory tasks that involve retention of detailed visual template information, such as objects. For example, in change detection tasks the response will depend on the probe and therefore cannot be planned for during the delay (Luck & Vogel, 1997). This past work has made tremendous strides in advancing our understanding of attentional templates, and this reductionistic approach has proven successful for understanding ‘pure’ cognitive processes. Ultimately, however, for a fuller understanding it is essential to consider action as an important dimension of attentional templates in natural behavior.

In visual search, there is evidence that attentional templates used to guide search in natural scenes differ in their quality when the targets had been previously searched for or acted upon. Specifically, Võ and Wolfe (2012), showed that attentional templates built from active visual searches (“looking for”) lead to more efficient guidance compared to templates derived from explicit memorization or free viewing (“looking at”). This provides a strong indication that the natural actions accompanying an initial visual search are subsequently incorporated into the attentional template. Further work has demonstrated that the fidelity of a representation improves after an active search compared to passive memorization (Draschkow, Wolfe, & Võ, 2014; Helbing, Draschkow, & Võ, 2020). Drashkow and Võ (2016) showed that memory representations for target objects from a real-world search task differed when those objects had been passively found (pointed at) or actively found (handled). Specifically, objects that were actively handled showed superior location memory. Taken together, the evidence from the visual search literature supports the idea that actions are incorporated into attentional templates.

Recent theoretical (Chatham & Badre, 2015; Myers, Stokes, & Nobre, 2017; van Ede, 2020) and empirical (Schneider, Barth, & Wascher, 2017; Trewartha, Case, & Flanagan, 2015; van Ede, Chekroud, Stokes, & Nobre, 2019) work has argued for a key role of action plans in memory representations. For example, locations of objects were recalled better when participants had made an active hand movement to them compared to when the hand was passively moved (Trewartha et al., 2015). Moreover, van Ede and colleagues (2019) asked observers to remember two objects each with a unique action associated with it. After a delay period, observers saw a probe indicating which of the objects should be reported. Surprisingly, the researchers found evidence for concurrent visual and action selection at the time of the probe suggesting that action plans are available before the moment they are needed and therefore held alongside the visual template representation.

Actions have only recently been considered as a potentially important dimension within the attentional template. As such, it remains unclear at what point these actions are likely to be incorporated into a template representation and what effect they may have on on-going behavior. I address these questions in Chapter 4.

1.4. The axes of attentional templates

Within the literature there are several axes along which we can categorize an attentional template. In the following section I will outline some of the important aspects of an attentional template and provide the context for recent work in the field.

1.4.1. Guiding vs. matching

Attentional templates may support multiple processes (Wolfe, 2020). During visual search, a target representation is used to *guide* spatial attention to potential targets within the search space. Once a candidate target is selected, this sensory information must be compared (or *matched*) with an internal target representation in order to determine whether the selected item *is* the target. The attentional templates supporting these two functions are often assumed to be the same representation, but this need not be the case, as different organizing principles may govern templates for these two purposes.

How and where attention is allocated in space in pursuit of a target is the topic of a long history of visual search experiments and reviews (Bundesen, 1990; Desimone & Duncan, 1995; Duncan & Humphreys, 1989; Itti & Koch, 2001; Treisman & Gelade, 1980; Wolfe, 1994, 2020; Wolfe & Horowitz, 2017). There are at least five known factors that ultimately work together to guide attention to potential targets during visual search: bottom-up salience, top-down feature guidance, scene guidance, history, and value (Wolfe & Horowitz, 2017). Top-down feature guidance has often been conceptualized as the attentional template.

Knowledge of target features can be used to guide spatial attention. In an early demonstration of this effect, Green and Anderson (1956) had observers search for a colored target number among colored distractor numbers. In one condition participants were informed of both the target number and its color, while in another condition participants were only told the relevant number. Critically, the experimenters varied the number of distractors which matched the target color. Their results showed that reaction times depended on the number of distractors which shared the color of the target only in the condition in which the color was known in advance. This foundational work provided some of the first evidence that spatial attention could be guided to potential targets on the basis of top-down feature information. There have been various demonstrations of this effect in the years since (Cunningham & Wolfe, 2014; Egeth, Virzi, & Garbart, 1984; Hout & Goldinger, 2015; Malcolm & Henderson, 2009; Treisman & Gelade, 1980; Vickery et al., 2005; Wolfe et al., 2004), all pointing to the idea that some form of a target representation guides attention during visual search.

A 'matching' attentional template is also necessary both in visual search as well as in tasks that do not require shifts in spatial attention (e.g. match-to-sample tasks). During visual search, eye-movement measures have shown that when more detailed target information is provided before the onset of a search (picture vs. word cue) observers are not only faster to find the target, but they are also faster to confirm the item they have fixated actually is the target (Malcolm & Henderson, 2009). This points to the idea the template fidelity can enhance the process of guidance as well as target matching. The fidelity or precision of the attentional template in working memory also affects performance in perceptual reproduction tasks (Bays, Catalao, & Husain, 2009; Fougnie, Suchow, & Alvarez, 2012; Gorgoraptis, Catalao, Bays, & Husain, 2011; Luck & Vogel, 1997; Van Den Berg, Shin, Chou, George, & Ma, 2012).

It is important to consider the potential uses of an attentional template and how this may, in turn, affect the representation itself. For example, it has been shown that the representations used to guide visual search are likely stored in working memory (Carlisle et al., 2011; Olivers et al., 2011). However, observers are also able to search for any one of 100 potential targets – far past the limits of working memory. One suggestion to reconcile this difference is that the 'matching' target template could reside in some form of long-term memory (Drew et al., 2016; Wolfe, 2012, 2020). In any case, considering how and for what a representation will be used to execute a given task is an important step in furthering our knowledge of attentional templates. In Chapter 5, I consider how spatiotemporal information within the attentional template may be used to facilitate both guidance and matching.

1.4.2. Lifespan of a template

How long can an attentional template survive? The answer to this question likely lies in where this template is stored. Working memory refers to a capacity-limited system used to *temporarily* hold or manipulate information in order to guide adaptive behavior (Baddeley & Hitch, 1974; D'Esposito & Postle, 2015). When a template is stored in working memory, access to this template will be short lived. Maintaining the attentional template in working memory allows observers to update the target flexibly from one trial to the next. This has been shown behaviorally, by varying the time in which a unique attentional template was presented prior to a visual search trial and comparing the subsequent search times to those required when the target did not vary (Wolfe et al., 2004). It was shown that performance was equivalent between non-variable search targets and variable search targets when the variable target was presented for at least 200 ms. This indicates that when a target is stored in working memory it takes approximately 200 ms to upload a representation equivalent to when it is held in a more robust store (Wolfe et al., 2004).

Neural evidence for attentional templates in working memory often relies on tasks in which the target changes trial-by-trial. For example, in single-unit recording from monkey IT, neurons responsive to a certain object exhibit elevated baseline-firing rates when that object is the target of an upcoming search (Chelazzi et al., 1993). This increase in the firing rate is thought of as the neuronal expression of the utilization of working memory. In human EEG, when targets are variable and the target template is presented laterally, the CDA component can be observed in the delay between the target template and the onset of the search trial (Carlisle et al., 2011; Vogel & Machizawa, 2004). The amplitude of the CDA increases when observers are searching for any one of two targets compared to a search for a single target. Moreover, the average CDA amplitude correlates with the average response times in observers, indicating the potential behavioral relevance of the representation being held in working memory (Carlisle et al., 2011).

Target templates can outlast the typical limits of working memory. Woodman and colleagues (2007) asked participants to search for a target during the delay of a working memory task. They showed that working memory interfered significantly with visual search only when the target was variable across trials. This suggested that when a consistent target was available working memory may not be utilized, but instead longer-term templates may be used. This idea was supported by work using human electrophysiology. Specifically, Carlisle et al. (2011) showed a reduction in the CDA when a target repeated on several trials in a row, providing evidence that the attentional template can be moved out of working memory, presumably into a long-term memory store.

Long-term memory can support more robust, longer lasting templates. Hybrid search is the search for any one of multiple targets held in memory. For example, when you are at the grocery store with a list of several items in mind you must simultaneously search through the visual world (the grocery store) as well as through your memory (the shopping list). This has been demonstrated for upwards of 100 items, far past the capacity limits normally associated with working memory (Boettcher, Drew, et al., 2018; Boettcher & Wolfe, 2015; Drew et al., 2017; Wolfe, 2012; Wolfe, Boettcher, Josephs, Cunningham, & Drew, 2015). This work has suggested that the attentional template associated with hybrid search may live in long-term memory. However, changes in performance across increasing memory set sizes, demonstrate that observers are not searching the entire contents of long-term memory when conducting their search (Boettcher, Drew, et al., 2018; Boettcher & Wolfe, 2015; Drew et al., 2017; Wolfe, 2012; Wolfe, Boettcher, et al., 2015). Nevertheless, this space does not appear to be capacity limited (Cunningham & Wolfe, 2014; Wolfe, 2012). Instead, an activated subset of long-term memory has been suggested to guide this type of search (Drew et al., 2016). This provides further evidence for a more robust space for the attentional template, and as such the potential for longer lasting templates.

Thus far we have discussed the volitional guidance of attention through attentional templates. However, memory templates can direct attention without deliberate control. Downing (2000) showed that attention was automatically directed towards a task-irrelevant object only when that object matched the contents of working memory. Moreover, as stated previously, the contents of working memory can interfere with the deployment of attention during visual search (Dalvit & Eimer, 2011; Olivers, 2009; Olivers et al., 2006, 2011). Templates from long-term memory also have a strong influence on involuntary attention. Representations of search displays built up (even implicitly) over several repetitions lead to faster target detection (Chun & Jiang, 1998; Jiang, 2018). In memory-guided orienting tasks, participants learn several scene-object relationships during a repeated visual search task. Observers subsequently show improved performance during perceptual discrimination task when the target is presented in location matching the LTM trace (Patai et al., 2012; Stokes et al., 2012; J. J. Summerfield et al., 2006).

In this thesis I will consider attentional templates with varying lifespans. In Chapters 2 and 3 attentional templates are retrieved from long-term memory associations – a common source for templates that guide our everyday behavior. In Chapter 4, I rely on templates that are only shortly available by utilizing a delayed reproduction task. Finally, in Chapter 5 I specifically test the underlying memory sources of spatiotemporal information within the attentional template.

1.4.3. Active vs. latent coding

The neural status of attentional templates may vary between active guidance of attention and latent traces that can interact with incoming signals. Olivers et al. (2011) argue for differential states of items within working memory leading to different guidance properties. Specifically, items that are currently relevant to the task take an active role within memory – operating as the primary attentional template and consequently guiding attention and altering perception. Accessory items in working memory remain latent and provide weak or no proactive attentional biases.

Recent analytical advances have facilitated gathering neural evidence for latent states of working memory representations. Multivariate pattern analysis allows for the interrogation of the content of neural signals. Specifically, by training an algorithm to distinguish between brain activity related to different classes of stimuli, we can make inferences about the information content when this classifier is applied to new data. Through traditional methods using imaging or neurophysiology it is not possible to measure changes in neural states that are not expressed through activity (LaRocque, Lewis-Peacock, Drysdale, Oberauer, & Postle, 2013; Lewis-Peacock, Drysdale, Oberauer, & Postle, 2012; Sprague, Ester, & Serences, 2016). However, by introducing stimulation during the delay interval, it is possible to measure how the silent change in neuronal state interacts with sensory processing. Using this radar-based approach combined with multivariate decoding methods it is possible to detect these latent signals (Wolff, Ding, Myers, & Stokes, 2015).

States of items in working memory are flexible. Retro-cueing studies suggest that attention can be refocused to items within working memory (Griffin & Nobre, 2003), thus allowing for the voluntary updating of items from a non-prioritized state to an active state (Myers et al., 2017). Importantly, items in a latent state are not dropped from working memory and can still be accessed in order to guide behavior if priorities shift either via an explicit cue (Myers, Chekroud, Stokes, & Nobre, 2018; Myers et al., 2017) or through the changing of relevance – probability of being probed – across time (van Ede, Niklaus, & Nobre, 2017).

Representations held in latent or active template states may not differ in their efficacy, but this difference may rather be functional. That is, latent items do not necessarily suffer in the quality of their representation, but rather in their accessibility. Therefore, shifts in prioritization allow items to be turned "on" and "off" dynamically. When an item is active this results in direct guidance of behavior. In Chapter 4, I ask whether accessory information within the attentional template – such as prospective action – guides behavior even in secondary tasks for which this accessory template information is irrelevant.

1.4.4. Explicit vs. implicit

Tasks may vary in how explicit and imperative templates are. For example, in standard visual search tasks, the template is not only explicitly provided, but it is also necessary to complete the task. That is, if you do not know what you are looking for, it can be impossible to find it (Wolfe, 1994). Alternatively, endogenous cues for spatial or temporal attention create a clear and explicit template (Coull & Nobre, 1998; Posner, 1980), which observers may choose to exploit in order to improve the performance, however the cue itself is not strictly necessary for completing the task. For example, a leftward arrow, allows for an attentional template that includes the expected target identity (e.g. a small flash) as well as spatial information (e.g. the left portion of the visual field. However, without this refined template observers could attend to their entire visual field to complete the task. In these cases, the spatial information within the template is not imperative, but rather useful.

In other types of tasks, information about the template is not provided explicitly, but instead learned incidentally through environmental regularities. For example, behavior is facilitated for targets appearing in high probability locations, although observers often lack awareness of the experimental manipulation (Geng & Behrmann, 2002; Jiang, 2018; Shaw & Shaw, 1977). Although this information is learned 'incidentally' during these tasks (termed probability cuing), behavior indicates that spatial information within the template can still guide search. The effects of differing target probabilities are evident in hemifield neglect patients, even in the neglected field, indicating that this attentional orienting may be dissociable from explicit awareness (Geng & Behrmann, 2002).

Jiang and colleagues (2013) specifically tested the interaction between implicit and explicit sources of guidance through a probability cuing task that was combined with an endogenous cue. They showed that observers’ attention was guided implicitly to areas in which a target was likely to occur as well as explicitly through a centrally presented arrow. However, when implicit statistical regularities were presented together with explicit endogenous cues, the explicit cues provided superior guidance, suggesting that explicit information may take priority within the attentional template.

It has also been shown that when a particularly salient distractor repeats its location, the attentional capture normally attributed to this distractor decreases with time (Wang & Theeuwes, 2018). Observers’ awareness of the repeated location did not impact the effect. This also holds true for non-spatial features such as color (Cunningham & Egeth, 2016) and indicates that features that are not relevant to the target may be suppressed in the template.

In contextual-cuing experiments, Chun and Jiang (1998) showed target locations need not repeat throughout the *entirety* of an experiment but these regularities can also be context specific, with highly probable locations changing from trial to trial. In these tasks, observers were asked to perform a difficult, attention-demanding search task. Some stimulus array configurations repeat throughout the experiment (repeated displays) while the rest do not (novel displays). Observers were faster at finding targets in the repeated displays compared to the novel displays. However, they were at chance when asked to report which of the displays they had seen before (Chun & Jiang, 1998). This work provides evidence that a template does not require explicit awareness for it to be utilized.

Taken together, proactive, memory-based templates can effectively guide attention for efficient behavior with or without explicit awareness. In Chapters 2-4, the attentional template is made explicit to the subjects. However, in Chapter 5 we test whether implicit spatiotemporal regularities also contribute to the attentional template.

1.4.5. Veridical vs. adaptive

The intuition regarding attentional templates is that they are equivalent to the actual percept of the object. That is, as you search for your keys in the morning, the representation guiding that search is similar to the one when you are actually looking at your keys. However, imagine you realize your niece is to blame for your lost keys. You must now search in the grass where she was last playing as well as the silverware drawer where she is known to hide things. In the grass, the color of your keys will help guide your search, whereas in the silverware drawer, the shape is likely the most distinguishing feature (Figure 1.2). A non-veridical account of attentional templates – or adaptive templates – would suggest that different features are prioritized in the templates guiding each search. In other words, the prospective context in which a template will be used, may determine how the features are blended into a useful representation.

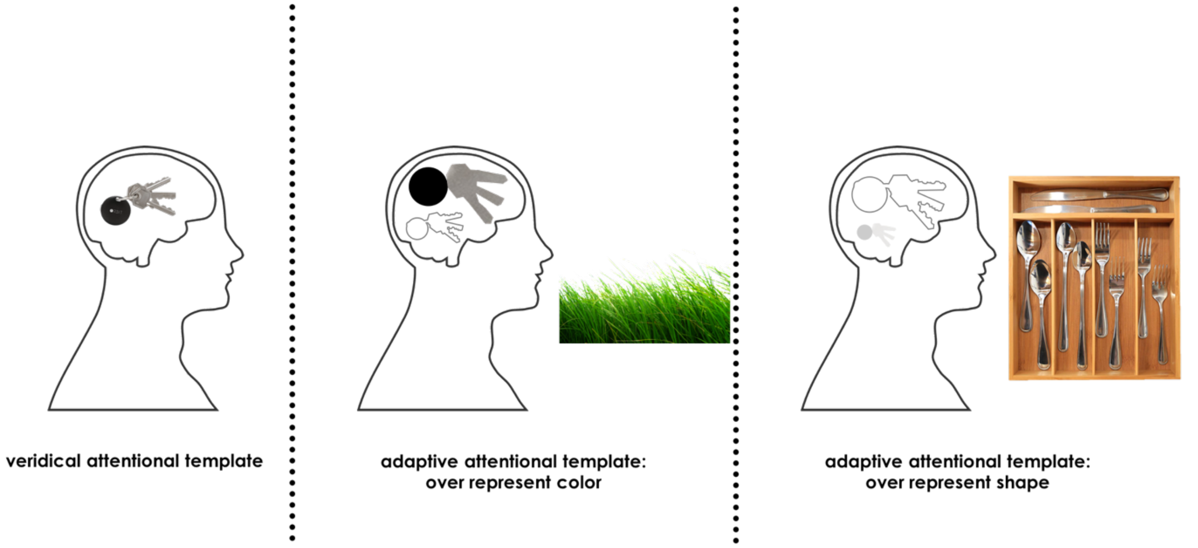

Figure 1.2: Veridical templates are equivalent to the actual percept of the object, whereas adaptive attentional templates are flexible for the current demands of the context.

Evidence for adaptive templates comes in several forms. Within working-memory research, it has been shown that observers primarily encode what they need, highlighting the functionally selective subset of the percept during the delay. When monkeys were presented with a multi-object array in which only one object was relevant, neurons in prefrontal cortex responsive to the relevant object were disproportionately active during the delay (Rainer, Asaad, & Miller, 1998). Using multivariate pattern classification in human fMRI, Serences and colleagues (2009), tested selective feature encoding and maintenance within an object. Specifically, observers were presented with colored gratings for a match-to-sample task. Either color or orientation was relevant across different blocks. During the delay period they found above chance classification in visual cortex only for the feature relevant to the current block. This indicates that observes can selectively weigh feature information within the template of an object.

Changes in feature weighting within templates can occur even after encoding is complete as demonstrated by feature-based retro-cueing (Hajonides, van Ede, Stokes, & Nobre, 2019; Niklaus, Nobre, & van Ede, 2017; Park, Sy, Hong, & Tong, 2017). Specifically, Niklaus and colleagues (2017) cued observers to the most likely feature dimension (color or orientation) of an upcoming probe, after having memorized three colored arrows. Results showed that observers were able to use the cues to effectively adapt the representations being held in working memory for the given context.

In visual-search tasks, template representations have also been shown to be adaptive. A veridical-template account would imply that targets are represented in the template true to form, and one would expect enhanced excitability for neurons representing target features. This would result in increased firing, larger ERPs, and more BOLD activation for all stimuli matching target features (Chelazzi et al., 1993; Maunsell & Treue, 2006; Serences, Saproo, Scolari, Ho, & Muftuler, 2009; Treue & Martinez Trujillo, 1999; Zhang & Luck, 2009). However, this is not always the most efficient strategy. Navalpakkam and Itti (2007) suggest that by maximizing the separability between target and distractors we can derive a more efficient model (for a recent review see: Geng and Witkowski, 2019). There is evidence for distortions within the template across orientation, color, and size (Bauer, Jolicoeur, & Cowan, 1996; Becker, Folk, & Remington, 2010; Geng, Di Quattro, & Helm, 2017; Hodsoll & Humphreys, 2001a; Yu & Geng, 2019). Additionally, it has been established that templates are sensitive to the expected variance within a feature dimension, specifically, features that occur with a low variance are represented more precisely in the template (Witkowski & Geng, 2019). This provides a convincing account of adaptable attentional templates.

The dimensional weighting account (DWA) suggests that it is possible to upweight an entire feature dimension within visual search, even before the exact feature value is known (Krummenacher & Müller, 2012; Liesefeld & Müller, 2019; Müller, Heller, & Ziegler, 1995; Müller, Reimann, & Krummenacher, 2003). In these studies, the target template is defined as a feature singleton (e.g. the only red bar among blue bars, or the only vertical bar among horizontal bars). When a target dimension repeats across trials (e.g. there are two trials in a row in which color is the distinguishing feature), observers are faster at finding the target, even when the exact target feature differs (e.g. on the first trial the target was red, on the next it was blue). This indicates that observers are able to adaptively upregulate a particular feature dimension in anticipation for a search task.

Finally, there is increasing evidence from naturalistic search tasks through scenes that attentional templates may contain more information than the features of the target alone. Regularities within scenes – known as *scene grammar* – makes visual search within them efficient (Võ et al., 2019). This boost in attentional guidance cannot be on the basis of veridical target features alone, as searching for the same exact target in a more artificial setting adds a substantial cost in reaction times (Josephs, Draschkow, Wolfe, & Võ, 2016). Recent work has shown that anchor objects – objects holding specific spatial predictions about a target – guide visual search and speed reaction times (Boettcher, Draschkow, et al., 2018). This work indicates that the attentional template may hold the features of additional related objects as well. For example, if you are searching the kitchen for a sponge, you likely use the features of the sink as well as the sponge to guide your search. Taken together the attentional template is likely not veridical but rather adaptive: resulting in various possible context-specific adjustments in order to guide behavior.

Although the adaptable nature of attentional templates is becoming increasingly recognized, there is a great deal of work needed to understand this process fully. For example, it remains unclear whether attentional templates retrieved from long-term memory associations may also be adapted. Further, when might these adaptations occur? I address these points in Chapter 3 of this thesis.

1.4.6. Representational vs. use-oriented

Attentional templates may serve multiple purposes. For example, a template may be helpful in perceiving a stimulus embedded in noise, it may provide the basis on which to make a decision, it may guide our actions, or it may be used during a visual search (Figure 1.3). How is the same template represented when its future use varies?

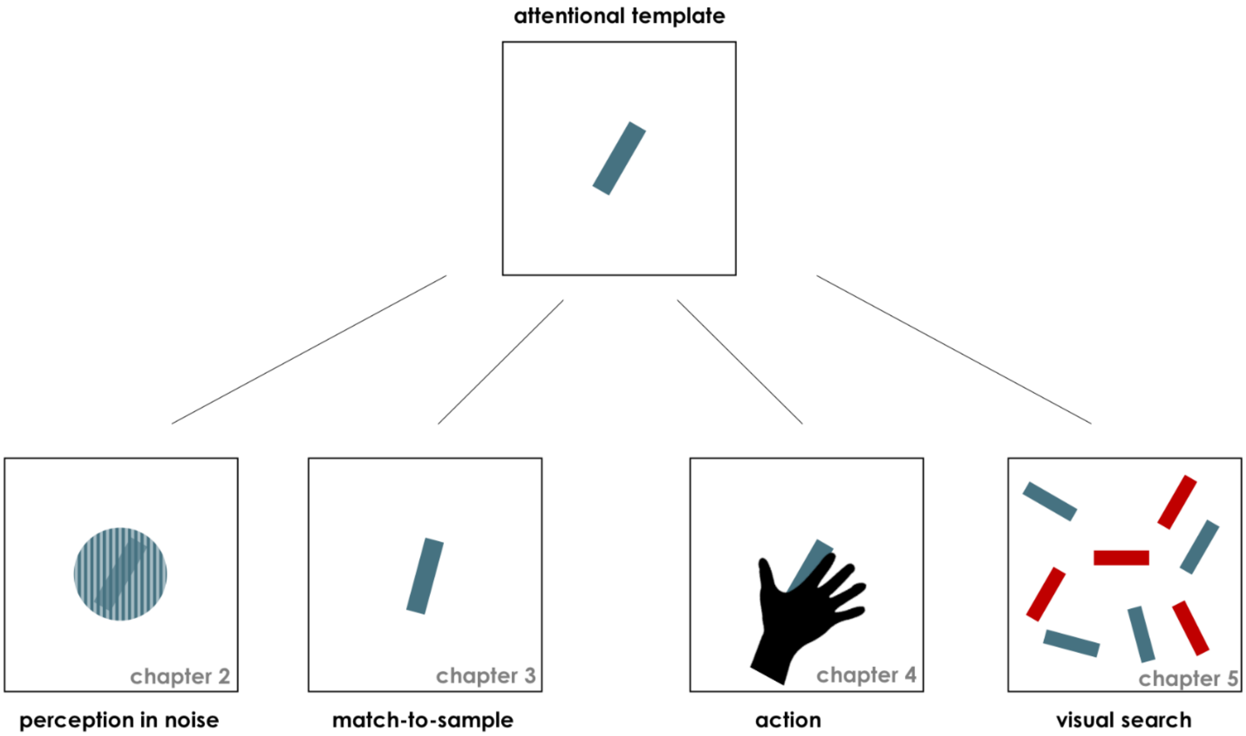

Figure 1.3: An example of a single template with various potential functions. Here the same object could be represented for a future perceptual task (Chapter 2), match-to-sample (Chapter 3), to execute an action (Chapter 4) or for a visual search (Chapter 5). How we maintain this template may depend on its future use.

Recent evidence from a neural network trained to complete several different working-memory tasks showed that the same information was coded differently depending on the task (Orhan & Ma, 2019). This insight indicates that templates may not only contain representational information about an item (*what is* *it?*), but also use related information (*what is it for?*). This has been shown in laboratory settings as well. Both humans (S. H. Lee, Kravitz, & Baker, 2013; Muhle-Karbe, Duncan, De Baene, Mitchell, & Brass, 2017) and monkeys (Warden & Miller, 2010) show task-dependent delay-period differences in the maintenance of the same object.

As discussed previously, attention within memory can change the priority state of an item. Recent theoretical work posits that attention may additionally alter the structure of an item within working memory such that it is optimized for future behavior (Myers et al., 2017; Nobre & Stokes, 2019; van Ede, 2020). Within this proposed framework, an item is first prioritized before being reformatted for the optimal retrieval given the context of the demands of the expected probe. This theoretical construct is backed by empirical work, showing that when an item is retrieved from visual working memory the action associated with that item is retrieved in parallel (van Ede, Chekroud, Stokes, et al., 2019). This indicates that the prospective actions of a stimulus may be stored within the template representation – although it remains unclear when this information may be incorporated into the template (Chapter 4). Moreover, van Driel and colleagues (2017) demonstrated that the same memory array with different tasks (recognition vs. visual search) showed a different pattern of posterior alpha suppression. This is a clear demonstration that the future use of an item is a critical determinant in the underlying neural state.

Templates drawn from long-term memory also hold certain behavioral associations and therefore are integrated with their function (Gibson, 1979). Behavioral evidence indicates that prospective actions are automatically available alongside object representations, even when they are task irrelevant (Tucker & Ellis, 1998). Participants were presented with an object associated with either a left- or right-hand action and had to respond to an action-irrelevant feature of the object with either their left or right hand. Observers were slower when the action association did not match the response. This behavioral effect has been complemented by electrophysiological markers indicating rapid motor preparation in line with an object's associated action within the first 200 ms (Goslin, Dixon, Fischer, Cangelosi, & Ellis, 2012). In an fMRI study, researchers found that retaining identity information for manipulable objects within working memory activated the hand region of the ventral premotor cortex, whereas nonmanipulable objects did not show this same pattern of activation. This suggests that the action associated with an object is recruited even during working memory maintenance for a passive task (Mecklinger, Gruenewald, Weiskopf, & Doeller, 2004).

1.5. Neural correlates of attentional templates

Many methods can be used to investigate attentional templates. I have already discussed some of the advances made in our understanding of templates through behavior, monkey electrophysiology, human neurophysiology, human brain imaging, computational modeling, and many other methods. Here, I will elaborate on three important neural signatures in human neurophysiology (EEG/MEG) that have contributed to our understanding of memory-based attentional orienting and which are relevant to this thesis.

1.5.1. Modulation of alpha-band oscillations

Alpha waves are oscillations occurring in the brain in the frequency range of 8-12 Hz. Hans Berger, the inventor of EEG, noted rhythmic bursts of electrical activity in the readout of the first EEG recordings (Berger, 1929). Among other discoveries, Berger realized that alpha decreased when observers had their eyes open – a first indication that lower alpha may be related to processing of the external world. Since the resurgence of interest in rhythmic brain activity, significant efforts are being dedicated to understanding the source and role of alpha oscillations. This work has suggested the involvement of alpha-band oscillations in modulating visual processing. Alpha power – or the amplitude of the wave – is anti-correlated with neuronal activity. Specifically, it has been shown in monkey electrophysiological recordings that an increase in alpha power in a particular region is associated with a decrease in neuronal spiking activity in that same region (Haegens, Nacher, Luna, Romo, & Jensen, 2011; Laufs et al., 2003). Therefore, decreases in alpha have been proposed to reflect excitability of a specific neuronal population.

Attention studies repeatedly show relative attenuation in alpha power contralateral (vs. ipsilateral) to the attended location. Moreover, proactive alpha-band desynchronization has been implicated in anticipatory shifts of attention. (Jensen & Mazaheri, 2010; Sauseng et al., 2005; Thut et al., 2006; Worden et al., 2000). In a simple perceptual task, Thut and colleagues (2006) asked observers to detect a visual target that either appeared on the left or right side of the screen. This was preceded by an auditory cue indicating the likely location of the target. The authors found modulations of alpha-band power were sensitive to the expected location of the target – relative desynchronization contralateral to the attended location. Moreover, this index was related to behavioral performance, with faster reaction times being associated with stronger attenuation. Studies using transcranial magnetic stimulation (TMS) have further confirmed the causal relationship between alpha-band modulations and proactive attention (Romei et al., 2008). Taken together, this work suggest a potential role for alpha as an input gate, allowing specific bits of sensory information to be processed while keeping distracting or unnecessary information out (Foxe & Snyder, 2011; Jensen & Mazaheri, 2010).

Alpha desynchronization has also been linked to shifts of attention within internal space, i.e. memory. Output gating is a mechanism for determining which items in memory may drive behavior – or act as the attentional template (Myers et al., 2017). Alpha lateralization has been found during both input gating and output gating with distinctive time courses. Specifically, an MEG study comparing pre-cues (cues appearing before the onset of a working-memory array) and retro-cues (cues appearing after a working memory array has disappeared) found a spatial shift in attention before the onset of an array is associated with sustained alpha lateralization while shifts of attention within memory is associated with transient shifts in the alpha-band (Wallis et al., 2015). Taken together, oscillatory amplitude modulations may subserve the dynamics of anticipatory attention (Jensen & Mazaheri, 2010).

Modulation of alpha power can be flexible and therefore reflect changing temporal expectations. Rohenkohl and Nobre (2011) showed rhythmic decreases in alpha power associated with a rhythmic perceptual stimulus. This suppression peaked just before the expected onset of a target, supporting the idea that alpha oscillations help regulate excitability for a precise moment in time and therefore are related to the tuning of temporal attention.

Alpha suppression has further been implicated in non-spatial attention suggesting that alpha is linked to changes of excitability more generally and can be based on various attributes of receptive fields. Snyder and Foxe (2010) demonstrated that when participants were cued to a relevant feature-dimension of a target stimulus (color or motion), alpha power was relatively attenuated in the area coding for the relevant feature dimension (dorsal visual stream regions for motion and ventral visual stream regions for color).

Neurophysiological studies show a systematic relationship between alpha-power modulation and proactive neural preparation across several domains. While we require studies with interference-based methods to confirm a mechanistic role of alpha modulation in attentional templates (Romei, Gross, & Thut, 2010), they nevertheless serve as a valuable proxy of neural modulation linked to preparatory templates.

1.5.2. Event-related potentials

Event-related potentials (ERPs) provide a useful metric in cognitive psychology and neuroscience for characterizing stereotypical responses to specific sensory, cognitive, or motor events. These waveforms are derived by averaging the EEG responses time locked to a common event. By repeating the event of interest many times, averaging extracts the common event-related neural responses and eliminates noise and unrelated brain activity. Over the years, research has uncovered specific ERP *components* – with characteristic polarity, latency, and voltage distribution over the scalp – which are linked to particular mental processes. A subset of these components is of interest to the study of attentional templates and this thesis.

The contingent negative variation (CNV) (Walter, Cooper, Aldridge, McCallum, & Winter, 1964) is linked to the anticipation of a task-relevant event. This component is categorized by a sustained frontal negativity that precedes an expected event or action, and therefore could in principle carry template specific information – although this has yet to be explicitly tested. The time course of the CNV follows temporal information in templates through various forms of expectations – e.g. cues, hazards, rhythms (Cravo, Rohenkohl, Wyart, & Nobre, 2011; Los & Heslenfeld, 2005; Miniussi et al., 1999; Nobre, 2001; Pfeuty, Ragot, & Pouthas, 2005; Praamstra, Kourtis, Kwok, & Oostenveld, 2006). These effects can be additionally driven by temporal associations stored in long-term memory (Cravo et al., 2017). In this work, observers learned to associate a specific temporal interval with a scene. The researchers found a late frontal negativity in anticipation of the target. This temporal information regarding the target was scene specific and could only be retrieved through a long-term memory association. As such, it can be stated that the temporal information within the attentional template modulates the CNV.

Post-stimulus components are also informative when inferring the effects of attentional templates. Modulation of the earliest sensory potential is highly indicative that the brain can take different states of preparation, implying the function of templates. In visual attention, early lateralized visual potentials (P1 and N1) are known to be modulated by attention and predictions. Specifically, attended stimuli see a boost in the amplitude of this complex of components (Luck, Hillyard, Mouloua, Woldorff, & et al, 1994).

Finally, the P300 is a component that – among other things – can be implicated in processes related to attentional templates. For example, the novelty P300 or the P3a can arise when a prediction is not met – or a stimulus does not match a template. In a typical experiment a standard stimulus is presented repeatedly in a stream creating a prediction for the appearance of that stimulus. On a small subset of trials a deviant – or unexpected – stimulus is presented producing a positive deflection between 200-500 ms post stimulus peaking over frontal/central electrode sites (Friedman, Cycowicz, & Gaeta, 2001; Sutton, Braren, Zubin, & John, 1965). When equating stimulus frequency, stimuli that are denoted as targets reliably evoke larger P300s than nontarget stimuli. (Duncan-Johnson & Donchin, 1977). Donchin (1981) presented a unified theory of the P300 with the context updating theory, which considers the component to index the updating of mental representations – or templates.

ERP components, preceding and following expected events can be informative regarding the nature of the expectation. Therefore, when considering the contents of an attentional template, ERPs time locked to expected events can be a useful tool.

1.5.3. Multivariate pattern analysis (MVPA)

A neural measure of information content seems essential for investigating the nature of attentional templates. Multivariate pattern analysis (MVPA) provides such a metric. Although MVPA may refer to a number of different specific analytical techniques, it can be broadly defined as the practice of categorizing (or decoding) experimental conditions on the basis of multiple measures – in this context stemming from brain data (e.g. voxels, sensors, electrodes). In addition to categorizing brain activity this method can also be used to tap into the content – that is, the likely representational state – of brain activity (Kriegeskorte, Goebel, & Bandettini, 2006). When a classifier achieves above chance decoding of animacy (e.g. animate objects vs. inanimate objects) from approximately 80 until 600 ms, we can interpret this to mean that there is some information specific to the animacy throughout this period (Grootswagers, Wardle, & Carlson, 2017).

This method can be a powerful tool when applied within the context of attentional templates, to understand not only *what* information is available but also *when* this information becomes available. Although, the 'decoding' of information content in the human brain has been an influential method of analysis in fMRI data for the past two decades (Cox & Savoy, 2003; Edelman, Grill-Spector, Kushnir, & Malach, 1998; Haxby et al., 2001; Kamitani & Tong, 2005), it has recently become a popular choice in MEG and EEG analysis as well, offering higher temporal resolution (King & Dehaene, 2014). For example, during a visual search experiment, observers were asked to indicate if either a car or a person was present in a natural scene. Using MVPA, researchers showed that information about which target was present (car or person) was available 160 ms after scene presentation (Kaiser et al., 2016).

Decoding the contents of the attentional template from electrophysiological data before the onset of the expected stimulus is an essential theoretical advance (Kok et al., 2017; Myers et al., 2015; Stokes et al., 2013; Wen et al., 2019). Kok and colleagues (2017) asked observers to detect the orientation of a target grating. This target was either preceded by a valid of an invalid cue indicating the likely orientation of the upcoming target. Using MVPA, the researchers showed a representation of the expected stimulus 40 ms before the onset of the target. This provides the first evidence in electrophysiological human studies that that some representation of the target may be present prior to the onset of the target.

Multivariate pattern analysis is a useful analytical method that has recently been more widely applied to MEG and EEG data. Within the context of attentional templates, this analysis technique provides a unique tool for identifying the contents of neural signals, and their respective timings. With careful experimental design, this technique can provide new insights into our understanding of attentional templates.

1.6. Aims of thesis

In my doctoral research, I investigate the flexible nature of attentional templates. As such, I set out to demonstrate and provide new understanding into the multi-dimensional nature of attentional templates. That is, I show that alongside the identity-based information (Chapters 2 and 3), prospective actions

(Chapter 4), and spatiotemporal information (Chapter 5) are also incorporated into the attentional template. This allows for proactive anticipation across various dimensions leading to more efficient behavior. Further, above I outlined several axes over which one can categorize attentional templates. Several of these axes lie at the heart of this thesis. With regards to the *"Lifespan of a template"*, I utilize templates retrieved from long-term memory associations in Chapters 2 and 3, while relying on templates that are only shortly available in Chapter 4. In Chapter 5, I specifically test the underlying memory sources of spatiotemporal information within the attentional template. Regarding the "*Veridical vs. adaptive" axis*, I ask in Chapter 3 whether templates retrieved from long-term memory associations can be adapted according to the current contextual demands or if they are represented veridically – matching the originally encoded object. To address the "*Representational vs. use-oriented*" nature of templates, in Chapter 4 I use a reproduction task manipulating associations between target features and response hand to assess if and when actions are incorporated into the attentional template. Finally, throughout this thesis I use several different tasks to probe the properties of attentional templates. In Chapter 2, observers are tasked with identifying a quickly presented target followed by a mask. In Chapter 3, I adapted a match-to-sample task in which the sample is retrieved from a long-term memory association. In Chapter 4, I use a reproduction task. In the final empirical chapter, Chapter 5, I developed a dynamic, extended visual search task in order to probe spatiotemporal information within the attentional template. Chapter 6 considers the body of work as a whole, drawing new insights about the nature of attentional templates, updating our understanding of the complexities they may hold in order to guide complex human behavior, and highlighting important avenues for future research.

2. Anticipating Object Identity from Memory Templates

2.1. Abstract

Probabilistic associations between stimuli can afford attentional templates that guide perception through proactive mechanisms. A great deal of work has examined the behavioral consequences and human electrophysiological substrates of anticipation following probabilistic memory cues that carry spatial or temporal information to guide perception. However, less is understood about the electrophysiological substrates linked to anticipating the sensory content of events based on recurring associations between successive events. Here, I demonstrate behavioral and electrophysiological signatures of utilizing associative-memory based attentional templates to guide perception, while equating spatial and temporal anticipation (Experiment 1 and 2), as well as target probability and response demands (Experiment 2). By recording the electroencephalogram (EEG) in the two experiments, I show that two markers in human electrophysiology implicated in spatial and temporal anticipation also contribute to anticipation of perceptual identity: attenuation of alpha band oscillations and the contingent negative variation (CNV). Taken together, the current results show that memory-guided identity templates can be derived from probabilistic associations and proactively impact perception. These anticipatory states involve attenuated alpha oscillations and the CNV. Furthermore, by isolating object-identity anticipation from spatial and temporal anticipation, my results suggest a role for alpha attenuation and the CNV in the anticipation of specific visual content beyond general changes in neural excitability or readiness.

2.2. Introduction

Attentional templates can be the product of frequent associations. For example, if you see a salt shaker on a table, you may expect to see the pepper nearby, as these two objects often cooccur. Your expectation of the pepper creates an attentional template. In this example, the template is born out of regular or probabilistic associations between object identities gained over a lifetime of experience. These probabilistic associations between stimuli lead to attentional templates that impact perceptual performance through the direct anticipation of the location, temporal onset, identity, or features of anticipated sensory events.

Early studies relied on relatively simple symbolic cues (e.g., arrows) to demonstrate that attention can be guided in space, time, and across features to facilitate performance (Coull & Nobre, 1998; Posner, 1980; Treue & Martinez Trujillo, 1999). More recently, studies have considered attentional orienting in more naturalistic tasks, in which the contents of long-term memory, often probabilistic in nature, guide the processing of incoming stimuli (Hutchinson & Turk-Browne, 2012). The bulk of the studies investigating memory-guided attention have focused primarily on anticipating spatial location (Awh et al., 2012; Chun & Jiang, 1998; Goldfarb, Chun, & Phelps, 2016; Jiang, 2018; J. J. Summerfield et al., 2006) and the expected temporal onset of items (Cravo et al., 2017; Olson & Chun, 2001). However, in addition, there is mounting interest in investigating mechanisms that support memory-based anticipation of the identity of upcoming percepts (Kok, Jehee, & de Lange, 2012; Peelen & Kastner, 2014; Stokes, Myers, Turnbull, & Nobre, 2014; C. Summerfield et al., 2008; Turk-Browne, Isola, Scholl, & Treat, 2008; Turk-Browne, Scholl, Johnson, & Chun, 2010).

Identity anticipation through attentional templates plays a central role in theories of attention (e.g. Desimone & Duncan, 1995; Duncan & Humphreys, 1989; Wolfe, 1994). In much of the work examining attentional templates to date, observers are explicitly provided with the template of the forthcoming target. That is, they are shown a particular object which they must subsequently match or search, such as in delayed-match-to-sample or visual search tasks (Carlisle et al., 2011; Chelazzi et al., 1998, 1993; van Driel et al., 2017). Though this can be informative in assessing attentional templates, it fails to capture a common everyday experience in building memory templates. Outside of the laboratory, frequent associations between successive different stimuli support the establishment of memory-based attentional templates. Building on previous work investigating associative memory templates (Higuchi & Miyashita, 1996; Kok et al., 2014; Kok, Jehee, et al., 2012; Kok et al., 2017; Rainer, Rao, & Miller, 1999; Turk-Browne et al., 2008, 2010), I here targeted two specific human electrophysiological substrates of associative memory templates during the anticipatory period.

I developed a task to investigate the anticipation of visual-identity information based on probabilistic associative memory. I report robust behavioral benefits on target perception in the context of a demanding visual identification task. I also investigated the electrophysiological markers linked to proactive template-based anticipation, specifically testing for the involvement of two canonical neural markers of anticipation from the spatial and temporal orientating literatures – the modulation of alpha-band oscillations and the Contingent Negative Variation (CNV).

Alpha attenuation has been associated with both spatial (Haegens et al., 2011; Thut et al., 2006; van Ede, 2018; Worden et al., 2000) and temporal (Heideman et al., 2018; Rohenkohl & Nobre, 2011; van Ede, Niklaus, et al., 2017; Zanto et al., 2011) orienting of attention, including during long-term-memory-guided anticipation (Stokes et al., 2012). Likewise, the CNV is an ERP component classically associated with temporal anticipation (Cravo et al., 2011; Los & Heslenfeld, 2005; Miniussi et al., 1999; Nobre, 2001; Pfeuty et al., 2005; Praamstra et al., 2006), also in the context of long-term-memory-guided anticipation (Cravo et al., 2017). Probing the involvement of these electrophysiological signatures during object-identity anticipation is important to inform a relevant and current theoretical debate about the nature of such markers. Alpha and CNV modulations during anticipation in space and time may purely reflect changes in the excitability of underlying neuronal populations (Benwell et al., 2017; Lemi, Chaumon, Crouzet, & Busch, 2017; Romei et al., 2008, 2010; Samaha, Gosseries, & Postle, 2017), independent of "informational content" or the attentional template. In the current work, I isolate identity anticipation and control for general "readiness" or "excitability" by equating spatial and temporal anticipation as well as target and response probabilities. If alpha and CNV modulations nevertheless still occur under these conditions, this would provide evidence that they also play a role in the anticipation of visual content.

2.3. Experiment 1: Attentional templates and target difficulty

2.3.1. Methods

Participants

All participants were right-handed with normal/corrected-to-normal vision had no history of neurological disorders, and were not taking any neurological medication. All participants gave informed written consent, and were compensated £15 per hour for a total of £45. The experiments were approved by the Oxford Central University Research Ethics Committee.

In Experiment 1, thirty volunteers participated. Out of the 30 participants, 5 of the participants missed > 80% of the difficult targets preceded by a non-predictive S1. On this basis, these participants were excluded from the analysis. Of the twenty-five remaining participants the average age was 24.2 (18-33) and there were 9 females.

Task and Procedure

Participants sat in a dimly-lit booth at a distance of 100 cm from the monitor (22 inch Samsung SyncMaster 2233; resolution: 1680 × 1050 pixels; refresh rate: 100 Hz; screen width: 47 cm). The experimental script was generated using Psychophysics Toolbox (Brainard, 1997) on MATLAB (version 2014b, The Mathworks

Inc., Natick, NA, USA). Participants were instructed to refrain from excessive blinking and to keep their face as relaxed as possible to avoid muscular artifacts in the EEG recordings.

The structure of Experiment 1 is shown in Figure 2.1. Participants were shown a random sequence of objects taken from a set of 14 objects from the Novel Object and Unusual Name database (NOUN) (Horst & Hout, 2016). Among these objects there were four critical objects: *easy S1*, *easy target*, *difficult S1*, and *difficult target*, and ten *neutral objects*. These four objects were randomly allocated to every fourth participant and then counter-balanced for subsequent participants such that for each random allocation of four objects, each object held each of the four critical roles. Participants' task was to press a corresponding key (either "m" or "x" key) whenever they detected a target. The targets switched their association with the keys randomly between blocks, such that each target was associated with the "x" and "m" buttons for half of the blocks.

Before the start of the task, observers were informed about the S1 objects. Specifically, they were told that following the presentation of a predictive S1 there was a 70% probability that the next item would be the corresponding target (i.e. the paired associate). Therefore, within the stream, specific S1 identities would predict specific target identities. In the other 30% of the trials each of the other items was equiprobable. Moreover, observers were tested on these relationships and required to correctly identify which S2 object followed which S1 object before continuing to the task.

A single trial consisted of the following sequence: stimulus 1 (S1), blank, stimulus 2 (S2), and a mask. S1 could either be a predictive or a non-predictive and was always presented for 250 ms. S2 could either be one of the targets or a foil object. S2 was immediately followed by a 100-ms mask that consisted of patches drawn randomly from the potential target items. For each set of objects, three of these masks were created and used randomly throughout the experiment. Target difficulty was determined by its exposure duration. The easy target was always presented for 150 ms before the mask, whereas the difficult target was presented for only 25 ms before the mask. The neutral objects were shown for either 150 or 25 ms equiprobably, i.e. any particular neutral object would be shown for 150 and 25 ms half of the time. The mask was followed by a 1000-ms blank before the next trial began. With this design the appearance of S2 was completely predictable in space and time. Before beginning the task, participants completed five practice trails. Within the task participants completed 14 blocks of 100 trials in total.

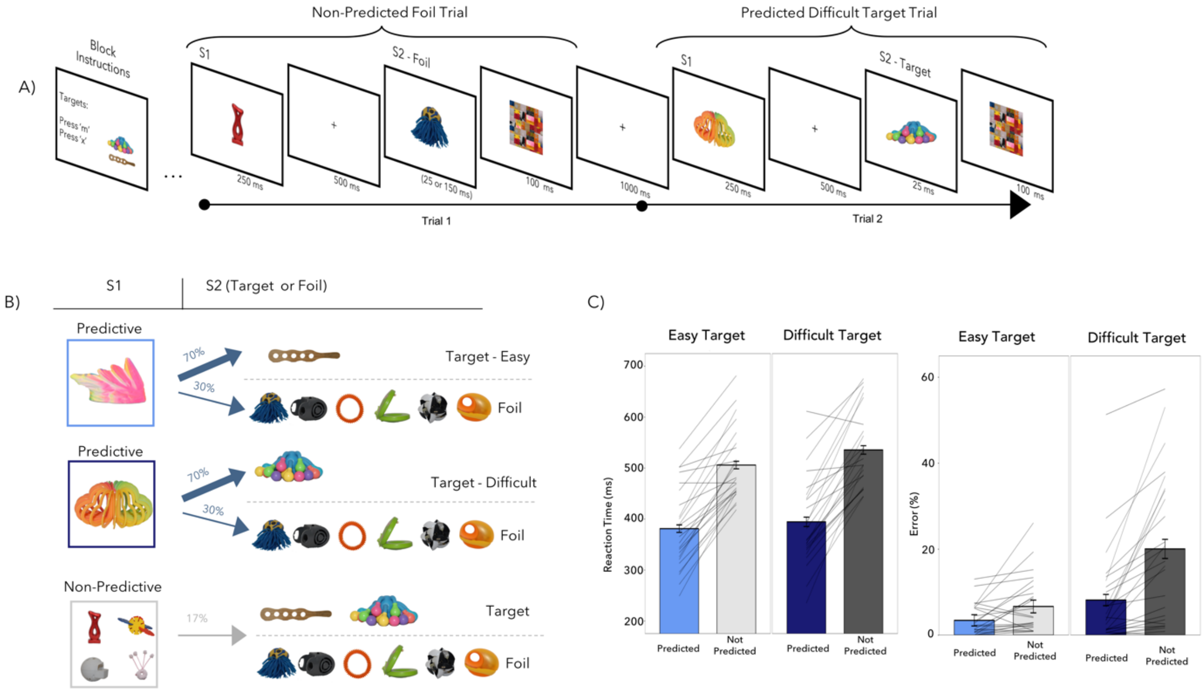

Figure 2.1: Trial schematic and behavioral data from Experiment 1 (A) an example of the trial sequence from Experiment 1. On each trial, participants saw stimulus 1 (S1) which could either be predictive or non-predictive about the following stimulus 2 (S2) which could an easy target (150 ms), a difficult target (25 ms), or a foil (25 or 150 ms). S2 was immediately followed by a mask. Participants were instructed to respond to the targets (but not the foils) with the corresponding button as quickly as possible. (B) The probability of a specific S2 target following a predictive S1 was 70%, whereas non-predictive S1s were equally likely to be followed by either of the two potential targets, or either of the 4 foils. S1-S2 relationships were made explicit to participants before starting the experiment. (C) Participants responded more quickly and more accurately to targets preceded by a predictive S1 as well as to easy targets. Additionally, there was a significant interaction in both RT and percent error indicating that predictive S1s had a larger benefit in the difficult-target trials.

Behavioral Data Analysis

For both experiments, behavioral data were analyzed using R (R Core Team, 2018). Reaction times and error rates were submitted to an ANOVA implemented in the ez package (Lawrence, 2013) and t-tests implemented in lsr (Navarro, 2015). Effect size estimates (η_G^2 & d) are provided for all effects. Plotting was completed using the ggplot2 package in R (Wickham, 2009).

EEG Acquisition

In both experiments, EEG data were acquired using Synamps amplifiers and Neuroscan data acquisition software (Compumedics). Sixty-one electrodes were distributed across the scalp using the international 10–10 positioning system. The left mastoid was used as the active reference, and the right mastoid measurement was included in order to derive an average-mastoid reference offline. The ground was placed on the left upper arm. Additionally, vertical and horizontal EOG electrodes were used to monitor for eye blinks and eye movements. During acquisition, data were low-pass filtered by an anti-aliasing filter (250-Hz cutoff), digitized at 1000 Hz, and stored for offline analysis.

EEG Preprocessing

The preprocessing and analysis scripts for both experiments can be found as html files and as reproducible scripts (jupyter notebooks; (Kluyver et al., 2016) at https://github.com/SageBoettcher/identityTemplates. All EEG data analysis was conducted in MNE-Python (Gramfort et al., 2013). The data were down-sampled to 200 Hz, and high-pass filtered at 0.1 Hz. To regress out eye-movement activity, an Independent component analysis (ICA; Jung et al., 2000) was used to decompose the data – which was high-pass filtered at 1 Hz – into sixty temporally independent components. Eye-movement components were detected by first correlating the filtered data with the electrooculography (EOG) and subsequently, when needed, manually selecting a subset of typical component maps and identifying the best group match to them (Viola et al., 2009). Selected components were then removed from the data. Trials were segmented from -200 ms to +750 ms +1000 ms (Experiment 2) relative to the onset of S1. Average activity over the 200 ms preceding the stimulus onset was used as a baseline against which all amplitudes were calculated. Finally, epochs with especially high variance were discarded. These epochs were detected through a generalized extreme studentized deviate (ESD) test for outliers with an alpha value of .05 and discarded from the analysis. On average 34 trials out of 1400 were discarded in the manner.

EEG Analysis

Alpha

For the time-frequency analysis, I used epochs from -200 to 1000 ms. Morlet wavelets were convolved with the data between 3 and 40 Hz. For each frequency, I used a fixed 400-ms time window such that the number of cycles changed with the frequency. After the time frequency transformation, activity was averaged over all posterior electrodes (P7, P5, P3, P1, Pz, P2, P4, P6, P8, PO7, PO3, POz, PO4, PO8, O1, Oz, O2) and contrasted between predictive and non-predictive trials separately for the easy and difficult conditions. This is expressed as a normalized difference (((predictive minus non-predictive) / (predictive plus non-predictive)) * 100).

ERPs

The event-related potentials (ERPs) were calculated in the same manner in both experiments: by averaging trials within a participant and then subsequently averaging these waveforms across participants separately for each condition. The ERPs were averaged across a predefined set of central-posterior electrodes (‘P1’, ‘Pz’, ‘P2’, ‘CPz’, ‘POz’) as well as central-frontal electrodes (‘F1’, ‘Fz’, ‘F2’, ‘AFz’, ‘FCz’). These electrodes were chosen based on previous work showing peak amplitude for the CNV at electrode Fz and peak amplitude for potentials linked to retrieval at electrode Pz. I focused my analyses on these electrodes and included the immediately surrounding electrodes to increase potential sensitivity.

EEG Statistical Analysis

Inferential claims about differences between conditions were based on cluster-based permutation test (Maris & Oostenveld, 2007) and reported according to recommendations by Sassenhagen & Draschkow (2019).

2.3.2. Results

Behavioral Results

To assess whether predictive S1s impact performance and whether this effect was modulated by the expected target difficulty, I conducted repeated-measures ANOVAs on RT and error rates with S1 type (predictive and non-predictive) and target difficulty (easy and difficult) as factors. Behavioral results are depicted in Figure 2.1.C. Target difficulty and S1 type interacted significantly in both RT ($F(1,24) = 5.4$, $p = .03$, $\eta_G^2 = .002$) as well as error rates ($F(1,24) = 12.0$, $p = .002$, $\eta_G^2 = .08$). Moreover, I found main effects of S1 type and target difficulty for both RT (S1 effect: $F(1,24) = 87.3$, $p < .001$, $\eta_G^2 = .41$; difficulty effect: $F(1,24) = 7.9$, $p = .009$, $\eta_G^2 = .01$) and error rates (S1 effect: $F(1,24) = 29.5$, $p < .001$, $\eta_G^2 = .21$; difficulty effect: $F(1,24) = 15.4$, $p < .001$, $\eta_G^2 = .28$). Paired samples t-tests (Bonferroni corrected p-values) revealed a significant RT benefit (i.e., faster RTs) of the predictive S1 for both easy and difficult targets (easy: $t(24) = 9.17$, $p < .001$, $d = 1.83$; difficult: $t(24) = 9.11$, $p < .001$, $d = 1.82$), and that the benefit of the predictive S1 was larger for difficult targets ($\mathrm{t}(24) = 2.33$, $p = .03$, $\mathrm{d} = .47$). The same pattern occurred for error rates, with a significant benefit (i.e., lower errors) following predictive vs. non-predictive S1 items in trials with an easy target ($t(24) = 2.9$, $p = .01$, $d = .59$) as well as trials with a difficult target ($t(24) = 4.93$, $p < .001$, $d = .99$). Once again this benefit of predictive S1s was larger for difficult targets ($\mathrm{t}(24) = 3.46$, $p = .002$, $d = .69$). Thus, predictive objects impact performance on the target, and this benefit was particularly pronounced when the targets were difficult to perceive.

The above results considered only target-present trials. For completeness, I also analyzed foil trials to determine if predictive S1s also led to more false alarms. I found that observers were indeed more likely to false alarm to a foil following a predictive compared to a non-predictive S1 ($t(24) = 3.14$, $p = .004$, $d = .62$; 14.5% vs. 1.5% false alarms). Because the probability that a target would appear after an informative S1 was higher than the probability that a non-target would appear (in Experiment 1, but not Experiment 2 as I will return to later), this increase in false-alarms following predictive S1s may simply reflect a strategic decision of participants to report the target when unsure.

EEG Results: Alpha

To assess the effect of a predictive vs. non-predictive S1s on induced brain activity, I first compared time- and frequency-resolved maps of power (collapsed over all posterior electrodes; see inserts Figure 2.2.A,B) from the onset of S1 until 250 ms after the onset of the S2 as seen in Figure 2.2. More specifically, I directly contrasted trials with a predictive and a non-predictive S1. I did so separately for trials with a predictive S1 that predicted an easy target (predictive-easy S1) and trials with a predictive S1 that predicted a difficult target (predictive-difficult S1). The same non-predictive-S1 trials were used for both comparisons. Significant clusters emerged following both the predictive-easy S1 (Figure 2.2.A $p < .001$) and following the predictive-difficult S1 (Figure 2.2.B, $p < .001$) in comparison to following the non-predictive S1. The maximal attenuation within these clusters for both the easy and difficult S1 occurred around 11 Hz and 600 ms after S1 onset, i.e. mostly concentrated within the alpha band. A topographic inspection confirmed that these effects had a clear posterior topography in line with a visual preparation effect. There were no significant clusters when directly contrasting easy to difficult S1s (all cluster *p*s > .13).

To have a clearer understanding of the time course of the alpha attenuation, I also averaged these effects along the classical alpha band (8-12 Hz; Figure 2.2.C). Once again, I found a significant cluster for both the easy ($p < .001$) and the difficult S1s ($p < .001$); with no significant difference according to the difficulty levels during the anticipation period ($p = .14$, with the only cluster forming after the onset of the target).

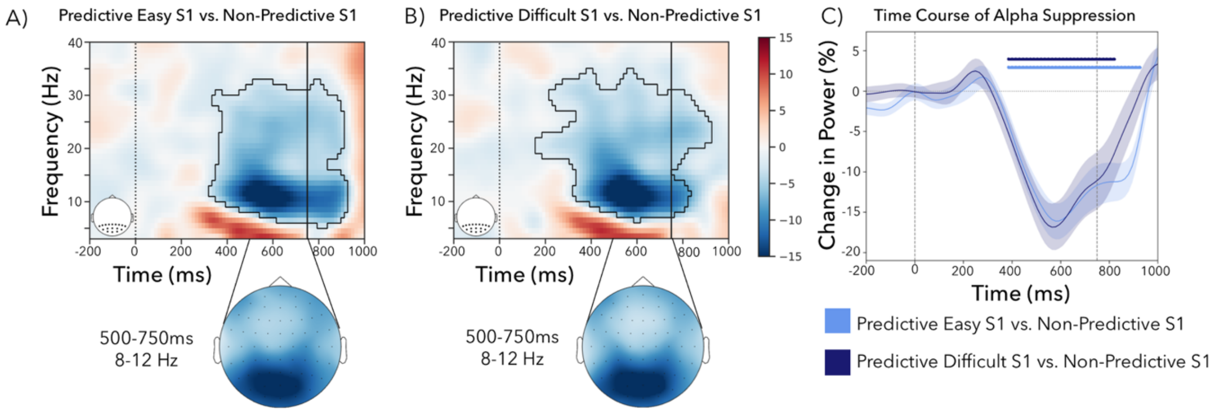

Figure 2.2: Alpha attenuation following predictive vs. non-predictive S1s in Experiment 1. (A) Time frequency results for posterior electrodes shows alpha attenuation in the predictive easy S1 trials vs. the non-predictive trials, as well as in the predictive difficult S1 trials vs. the non-predictive S1 trials (B). The topographies are plotted on the same scale as the above time frequency plot. C) shows the time course of the alpha attenuation averaged between 8 and 12 Hz. Vertical lines at 750 ms show the onset of the S2 target. Significant clusters with a p value < .05 are denoted with the black outline (panels A and B) and as horizontal lines in panel C. Shaded areas represent ± 1 SEM (68% confidence intervals).

EEG Results: ERPs

To investigate the anticipatory nature of identity-based templates, I additionally investigated event-related-potentials (ERPs) locked to the onset of predictive-easy S1s, predictive-difficult S1s, and non-predictive S1s for predefined clusters of frontal and posterior electrodes. The results are depicted in Figure 2.3. I was specifically interested in testing if these identity-based predictions also produce a CNV – a frontal negativity – in the pre-defined frontal electrodes.

I first considered the frontal electrode cluster (Figure 2.3.A). For both the predictive-easy and the predictive-difficult S1 cues, I found a significantly larger negativity in the late S1-S2 cue-target interval, compared to the non-predictive S1 cues (easy: $p<.001$, difficult: $p<.001$). These negativities were associated with a frontal topography characteristic of the CNV (Figure 2.3.C). In the S1-predictive-easy condition showed an early positivity ($p = .004$) that is likely a spillover effect from an earlier more posterior positivity that I return to below (as also confirmed by the time-resolved topographical analysis presented in Figure 2.3.C). There were no significant clusters when contrasting the easy and difficult S1s ($ps > .43$).

When comparing effects for predictive vs. non-predictive S1 cues in the predefined posterior electrodes (Figure 2.3.B) a significant cluster was identified from around 200 to 600 ms for both easy and difficult (ps < .01) The effect reflected a late positive potential elicited by predictive cues. Topographical analysis confirmed the potential was centrally distributed over the posterior scalp (Figure 2.3.C). As with the alpha modulations and the CNV, there were no significant clusters when comparing the easy and difficult predictive S1s (all cluster ps > .43).

These effects were confirmed, and also nicely demonstrated, by the time-resolved topographies of predictive vs. non-predictive S1 (separated by the easy and difficult conditions), as depicted in Figure 2.3.C.

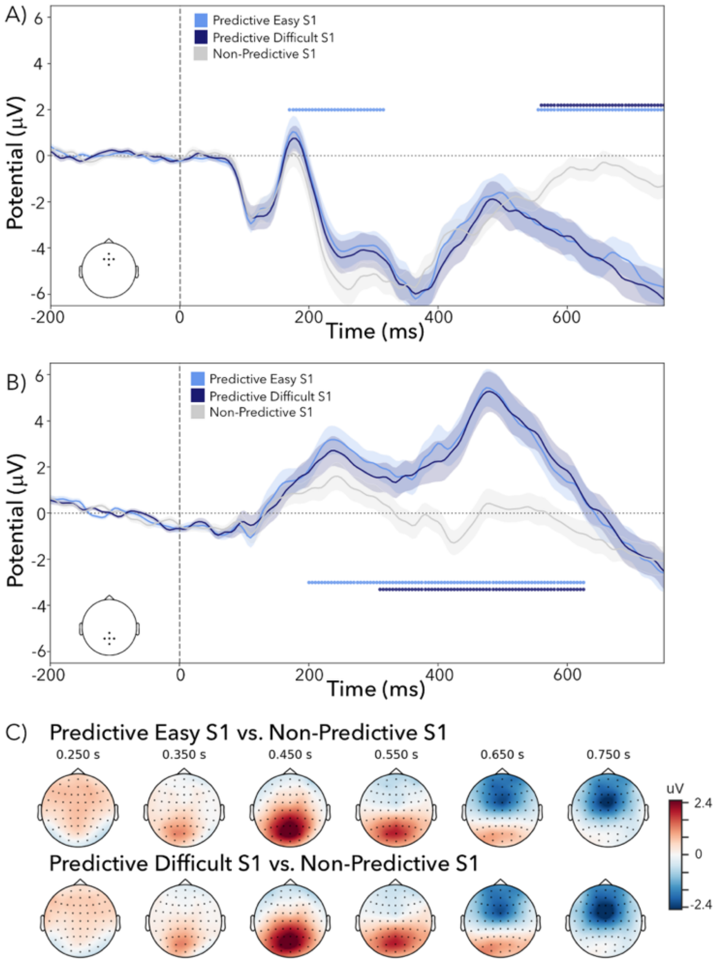

Figure 2.3: Posterior positivity and frontal negativity following predictive vs. non-predictive S1s in Experiment 1. (A) ERPs locked to the onset of S1 and averaged across a subset of frontal electrodes ('F1', 'Fz', 'F2', 'AFz', 'FCz'). Predictive S1s show a late frontal negativity relative to non-predictive S1s, while difficulty did not significantly modulate this effect. (B) ERPs locked to the onset of S1 and averaged across a subset of posterior electrodes ('P1', 'Pz', 'P2', 'CPz', 'POz'). The predictive S1s show a clear positive deflection from the non-predictive S1, while difficulty did not significantly modulate this effect. (C) Topographies of the ERP effects (predictive easy/difficult vs. non-predictive) over time show an early posterior positivity followed by a late frontal negativity. Significant clusters with a p value < .05 are denoted with horizontal lines in panels A and B. Shaded areas represent ± 1 SEM (68% confidence intervals).

2.4. Experiment 2: Templates While Equating Target Probabilities

In Experiment 1, the pattern of behavioral data was suggestive of proactive attentional template utilization, resulting in larger performance benefits when target discrimination was difficult. Proactive memory-based expectation was also suggested by alpha attenuation and a CNV following predictive vs. non-predictive S1 objects. These predictive S1s allowed participants to prepare for the identity of the upcoming stimulus, while controlling for spatial and temporal expectations that were matched between the S1 objects.

Although neural markers clearly signaled target anticipation, it was not possible to conclude that the neural effects were specifically related to the perceptual identity of the anticipated target. On average, task-relevant items (targets) were also more likely following predictive vs. non-predictive S1s, which may have led to differential motor anticipation, or states of attention. Because responses were only required to the target stimuli, during predictive S1 trials observers could not only prepare for a task-relevant visual target, but possibly also for the associated motor response. The neural effects may thus reflect general task readiness (or “excitability”), rather than template-specific anticipation of visual identity. To rule out this potential interpretation, I designed Experiment 2 (Figure 2.4).

In Experiment 2, I equated these other forms of anticipation by making S2 a task-relevant stimulus on every trial. Specifically, participants were always tasked with discriminating S2, but only a subset of S1 stimuli predicted the identity of S2. Therefore, the only difference between predictive and non-predictive S1s was the likelihood of a *specific* target appearing. As such, differences between the S1 conditions must be attributed to proactive target template activation. Participants once again saw predictive and non-predictive S1s (Figure 2.4.B) which were equated for their spatial and temporal predictions, as well as motor affordances. Three stimuli served as S2, two of which were predicted by a subset of S1 stimuli and one of which was completely unpredictable. Participants responded to S2 in a 3-alternative forced choice (3AFC) design. To eliminate anticipation of specific motor responses, response mappings were random on every trial. Across the experiment, all three targets were equally probable and potential differences in the preparatory period can no longer be attributed to differences in target probability or response preparation. In Experiment 2, all trials had the same difficulty level, allowing us to focus exclusively on the central question of identity anticipation.

2.4.1. Methods

Participants

In Experiment 2, thirty-six volunteers participated. Out of the 36 participants, 6 of the participants performed at chance for targets on non-predictive S1 trials. On this basis, these participants were excluded from the analysis. Of the remaining thirty participants the average age was 27.1 (20-34) and 15 were female.

Task and Procedure

The structure of Experiment 2 is shown in Figure 2.4. The stimuli, experimental set up, and EEG procedures were the same as in Experiment 1. A trial was similar to that of Experiment 1, with a few critical changes. On each trial, participants first saw S1 (250 ms) which again could be a predictive or a non-predictive S1 with equal probability. This was followed by a 750-ms blank and the quick presentation of one of three targets (30 ms) – I will refer to these targets as target A, B, or C. That is, there was a task-relevant item presented on every trial. Critically, two of these items (target A and target B) were predictable based on S1, whereas the other item (target C) was always equally probable after all S1 stimuli. Following the presentation of the target and a mask (100 ms), all three potential targets appeared on the screen and observers used the left, down, and right arrow keys to indicate which object they had just seen. The position of the three targets was randomized across trials such that observers could not prepare their response before the response screen. The stimuli were randomly allocated to each participant. With these changes to the design, every trial and item was task relevant, and participants could not prepare a specific response during the period after S1. Here, therefore, predictive and non-predictive S1s differed only with regard to its ability vs inability to form a specific target template in anticipation of S2.

The relationships between the S1 and target items were explicitly detailed to the participants before the experiment. In total there were eight potential S1 items. Four of these items were predictive and four were non-predictive. Of the four predictive S1s, two predicted target A and two predicted target B. That is, if one of these predictive S1 objects appeared the associated target would follow in 2/3 of the trials. In the remaining 1/3 of the trials Target C would appear. On non-predictive S1 trials all targets were equally likely. As such, throughout the experiment, all three targets were equally likely to appear such that there was no higher probability of a predictable target.

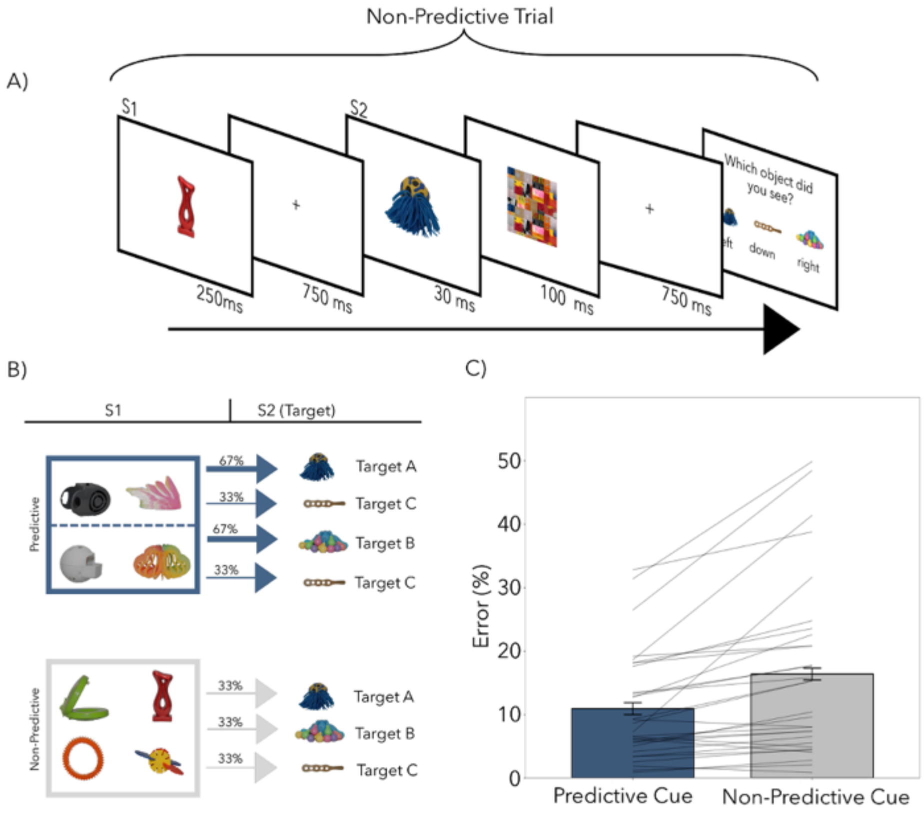

Figure 2.4: Trial schematic and behavioral data from Experiment 2. (A) Schematic of an example non-predictive trial in Experiment 2. Participants' task was to always report the second S2 object. The paradigm is very similar to experiment 1 with the exception that participants must respond on every trial (i.e., each S2 is a target). Probabilities of each S2 target given the preceding S1 are shown in in panel (B). In (C) it is clear that there is a significant effect of the predictive S1 on error rates. Because this task was a delayed forced choice, reaction times were no longer informative.

EEG Preprocessing

The majority of the data preprocessing was the same across Experiments 1 and 2 with one notable exception. Namely, trials in Experiment 2 were segmented from -200 ms to + 1000 ms relative to the onset of S1.

EEG Analysis

Alpha

Again, the analysis pipeline was mostly replicated across experiments, with the exception of the segmenting. In Experiment 2, epochs for time-frequency analysis ran from -200 to 1250 ms.

2.4.2. Results

Behavioral Results

To test for a benefit to the predictive S1s in the error rates, I used a paired samples *t*-test. As seen in Figure 2.4.C, targets preceded by a predictive S1 were again detected more accurately ($t(29) = 4.16$, $p < .001$, $d = .76$). Because participants gave a 3AFC response after an imposed delay, reaction times were not considered informative of perceptual processing in Experiment 2, and were therefore not analyzed.

EEG Results: Alpha

To assess the alpha attenuation following predictive vs. non-predictive S1s, I compared the time-frequency maps in the period between the onset of S1 and the onset of S2. As shown in Figure 2.5.A, I observed a significant cluster ($p = .005$), with a qualitatively similar profile (in terms of time range, frequency-range, sign, and topography) as in Experiment 1. The peak attenuation in this cluster was found at 11 Hz and 610 ms post S1. As in Experiment 1, this attenuation was associated with a predominantly posterior topography

(Figure 2.5.A). When focusing on the predefined 8-12 Hz alpha band (Figure 2.5.B), I found a significant cluster (*p* = .01), which spanned a similar time range as in Experiment 1.

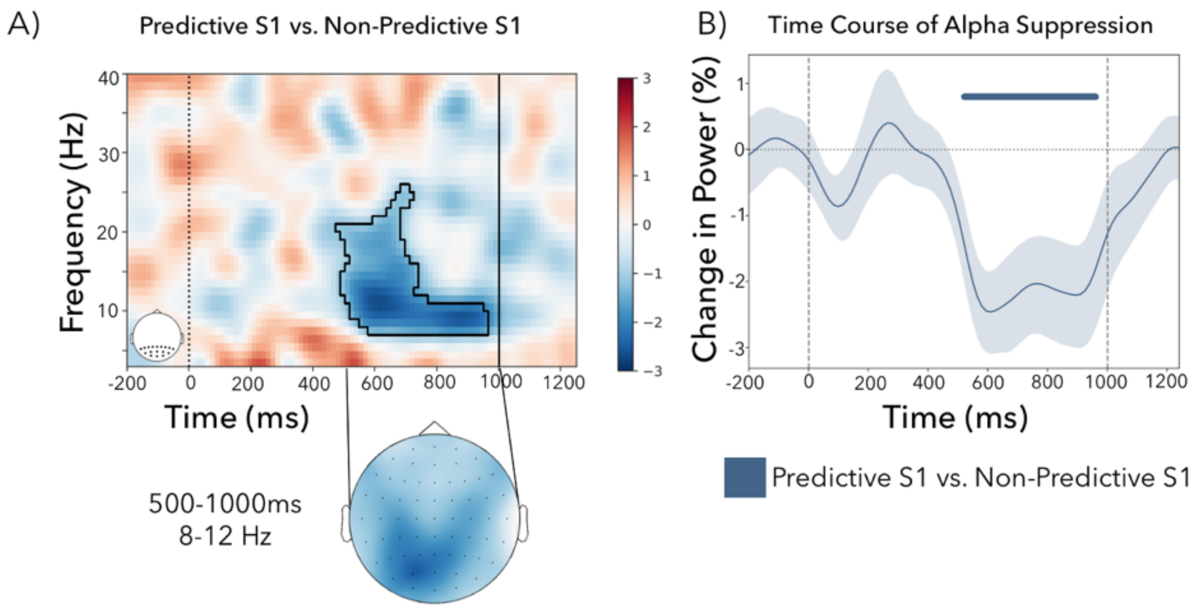

Figure 2.5: Alpha attenuation following predictive vs. non-predictive S1 in Experiment 2. (A) Time-frequency results for posterior electrodes shows alpha attenuation following the predictive S1 relative to the non-predictive S1, with a peak negativity at 610ms post S1 at 11 Hz. (B) Time course of the alpha attenuation, averaged between 8 and 12 Hz. Vertical line at 1000 ms shows the onset of the target. Significant clusters with a p value < .05 are denoted with the black outline in panel A, and by the horizontal line in panel B. Shaded area represents ± 1 SEM (68% confidence interval).

EEG Results: ERPs

As in Experiment 1, I also investigated ERPs locked to the onset of S1 in the pre-defined frontal and posterior electrode clusters (Figure 2.6). In the frontal electrode cluster (Figure 2.6.A), I again observed a CNV – a larger negativity following predictive S1s just before the onset of S2 (cluster p = .04). Like in Experiment 1, I also found a significant positive cluster in the frontal electrodes between about 300 and 450 ms (p = .01), which again likely involved a spillover from a more posterior effect (Figure 2.6.C). Indeed, in the posterior cluster (Figure 2.6.B), predictive S1s again elicited a larger positive potential from about 300 ms until around 550 ms, yielding a significant cluster (*p* = .001).

The topographies again demonstrate how the effects of the predictive vs. non-predictive S1s develop over time and space (Figure 2.6.C), and revealed a qualitatively similar spatial-temporal progression as observed in Experiment 1.

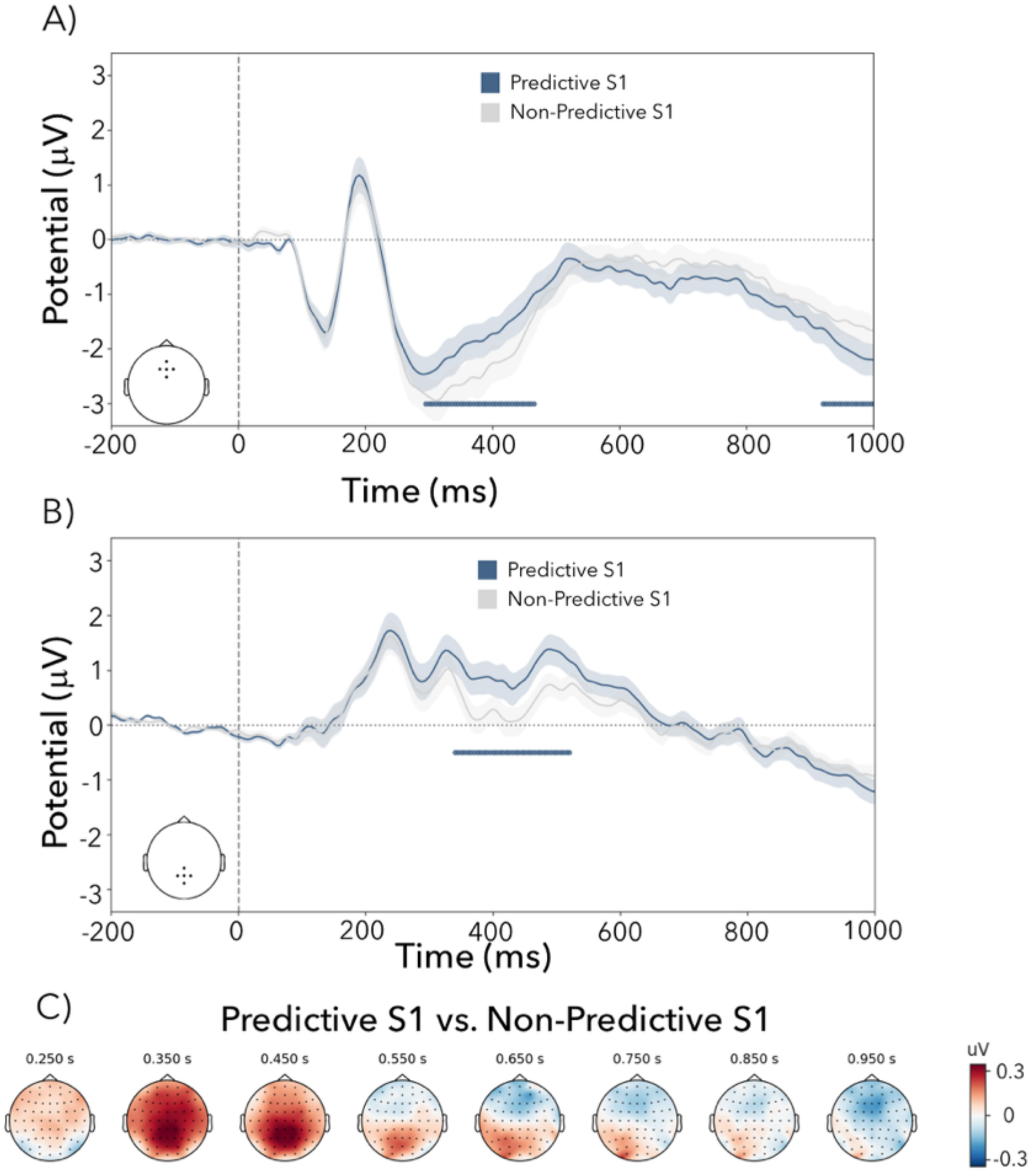

Figure 2.6: Posterior positivity and frontal negativity following predictive vs. non-predictive S1 in Experiment 2. A) ERPs locked to the onset of S1 and averaged across a subset of frontal electrodes. Predictive S1s show a late frontal negativity relative to non-predictive S1s. B) ERPs locked to the onset of S1 and averaged across a subset of posterior electrodes. The predictive S1s show a clear positive deflection from the non-predictive S1s. C) Topographies of the ERP effects (Predictive vs. Non-predictive) show an early posterior positivity followed by a late frontal negativity. Significant clusters with a p value < .05 are denoted with horizontal lines in panels A and B. Shaded areas represent ± 1 SEM (68% confidence intervals).

EEG Results: MVPA

The tightly controlled identity-expectation manipulation in Experiment 2 also enabled us to investigate whether the proactive deployment of probabilistic associative-memory templates based on S1 improved neural processing of S2 during perceptual analysis (i.e., post S2 target onset). Unlike in Experiment 1, the S1 items were all followed by target items, thus equating motor demands and degree of preparation. Presentation duration of S2 was also equated. To test for qualitative changes in sensory processing, I applied linear discriminant analysis to decode the content of the two predictable targets in posterior electrodes when they were preceded either by a predictive or a non-predictive S1 (Figure 2.7). Cluster-based permutations that considered the first 300 ms of target processing showed a single cluster of better decoding for predictable compared to unpredictable targets, though this did not survive cluster-correction (p = .09). When I considered only the peak decoding period of all targets (at 145ms; Figure 2.7B) I found better decoding for predicted vs. unpredicted targets (t(29) = 2.89, p = .007). However, because this effect was not particularly strong (Figure 2.7), I would like to present this as a tentative result, without further speculation or elaboration within the discussion.

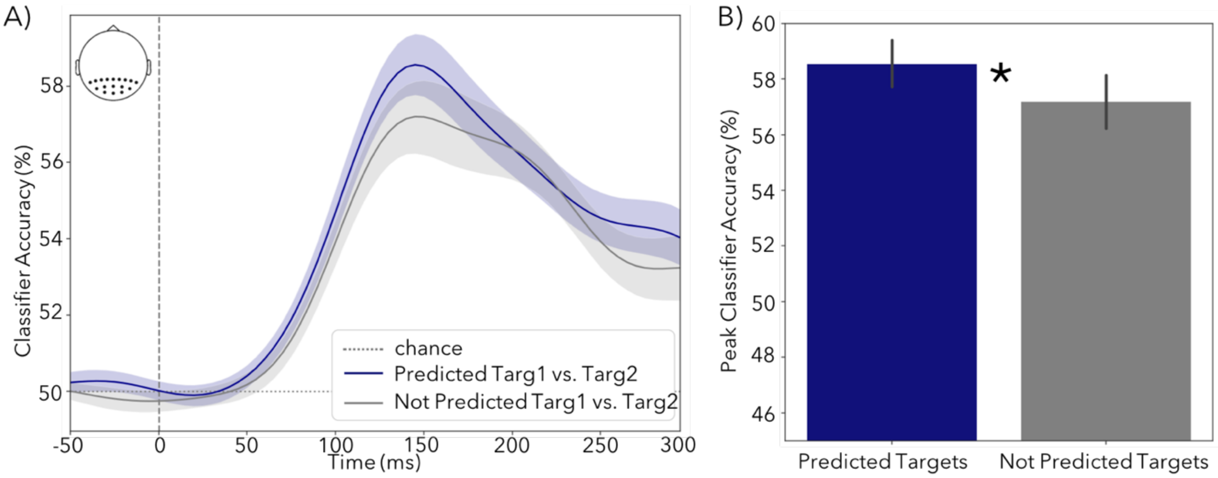

Figure 2.7: (A) LDA classification accuracy of S2 target A vs. S2 target B (in Experiment 2) when preceded either by a predictive (blue line) or non-predictive (gray line) S1. (B) Classifier accuracy at the peak classification time for the group average (145 ms) for both predicted and not predicted targets. To avoid circularity, the peak time was found based on the average of the predicted and not predicted data. LDA was performed in a time-resolved fashion on the baseline-corrected time series, using the topographical distribution across all posterior electrodes (as indicated in the inset) as the multi-variate data features.

2.5. Discussion

My results provide evidence that identity templates based on probabilistic associative memory impact perception. Furthermore, these attentional templates are associated with proactive states of attenuated alpha oscillations and the CNV, even when controlling for differences in spatial and temporal anticipation as well as response and target probabilities.

The behavioral and EEG results build on and extend earlier work on memory-guided attentional orienting and perceptual identity templates in several ways. When considering memory-guided anticipation, I have focused here on perceptual consequences and the electrophysiological signatures of memory-guided predictions based on identity, as opposed to anticipation in space and time (Awh et al., 2012; Chun & Jiang, 1998; Cravo et al., 2017; Goldfarb et al., 2016; Jiang, 2018; Olson & Chun, 2001; J. J. Summerfield et al., 2006). I have studied this in a context where the templates must be retrieved from complex probabilistic associations in memory templates (Higuchi & Miyashita, 1996; Kok et al., 2014; Kok, Jehee, et al., 2012; Kok et al., 2017; Rainer et al., 1999; Stokes et al., 2009; Turk-Browne et al., 2008, 2010) – rather than being explicitly provided (Carlisle et al., 2011; Chelazzi et al., 1993; van Driel et al., 2017) – and have focused specifically on the anticipatory electrophysiological substrates associated with such templates.

This work also expands upon prior work that has used similar paired-associate tasks as the one here (Brincat & Miller, 2015; Gallistel, 1990; Higuchi & Miyashita, 1996; Rose, Verleger, & Wascher, 2001; Stokes et al., 2014), but where the focus was on learning. In the current study, the focus was not on the learning of the S1-S2 associations, but rather on the exploitation of previously learned information in service of guiding ensuing behavior (see also Rainer, Rao, & Miller, 1999; Stokes et al., 2013, 2014), here in a demanding perceptual task with masked visual targets. In doing so, I report that participants are able to utilize learned identity associations to impact perception.

A major empirical contribution of this study was to identify electrophysiological markers for the anticipation of identity-related informational content in the human brain that I discuss next in turn.

2.5.1. Alpha Attenuation

In previous work, alpha attenuation has been noted during anticipatory periods for both spatially and temporally predictable targets (Heideman et al., 2018; Rohenkohl & Nobre, 2011; Sauseng et al., 2005; Siegel, Donner, Oostenveld, Fries, & Engel, 2008; Thut et al., 2006; van Ede, de Lange, Jensen, & Maris, 2011; Worden et al., 2000; Zanto et al., 2011). In this context, alpha attenuation has been theorized to reflect engagement of sensory processing areas in preparation for a task-relevant event, in line also with the notion that alpha is inversely related to firing-rates (Haegens et al., 2011) and/or processing capacity (Hanslmayr, Staresina, & Bowman, 2016) of the underlying populations. In my results, I have shown alpha attenuation when S1 specifically predicts the identity of an upcoming target over and above its location and temporal onset. Accordingly, I propose that the alpha attenuation also reflects engagement with visual processing areas to prepare a specific target template. As such, the alpha modulations reported here complement recent work showing that lower alpha power is associated with higher fidelity of stimulus-specific information (Barne, Lange, & Cravo, 2020; Griffiths et al., 2019; van Ede, Chekroud, Stokes, & Nobre, 2018). In this light, it is interesting to note that alpha-band oscillations were not significantly modulated by the anticipated perceptual difficulty in identifying the target, as might be expected from a pure “excitability” account (e.g. Benwell et al., 2017; Lemi et al., 2017; Romei et al., 2008, 2010; Samaha et al., 2017). Rather, at least in this task, the observed alpha attenuation appears to reflect anticipation of specific visual content related to target identity, though I note that visual content in this task entailed different shapes across objects, and thus included some spatial attributes.

When templates are separated by space and time, template preparation has previously been associated with spatially lateralized contralateral alpha attenuation relative to the memorized location of the template (de Vries, van Driel, & Olivers, 2017; van Driel et al., 2017). My findings complement this work by isolating template identity, while controlling for spatial attention associated with the template. Moreover, as emphasized earlier, I here show this in a context in which the template was not presented to participants, but had to be retrieved from long-term memory based on a known probabilistic association between S1 and S2.

Snyder and Foxe (2010) demonstrated that when participants were cued to a relevant non-spatial feature-*dimension* of a target stimulus (color or motion), alpha power was relatively attenuated in the area coding for the relevant feature dimension (dorsal visual stream regions for motion and ventral visual stream regions for color). This complements the idea that alpha-attenuation may serve as a general attentional mechanism in perception. However, because this previous work cued feature dimensions (e.g., color) rather than feature values (e.g., red), it does not address whether alpha is also a relevant mechanism for expected *identity* or 'template' preparation.

Interestingly, a previous study in which participants could prepare for a specific defining feature of a forthcoming target grating (Wildegger, van Ede, Woolrich, Gillebert, & Nobre, 2017) found no evidence for modulations within the alpha-band. The apparent discrepancy with the current finding could be due to statistical variability (i.e. a false negative in previous work), or reflect crucial task dependencies. For example, the current task utilized complex stimuli, memory associations, and targets that were always presented centrally whereas the previous work used simple orientations, symbolic cues, and uncertainty with regards to target location.

In the current work, I focused on the process of template-guided attention. The instantiation of the attentional template putatively involves a process of retrieval from long-term memory, possibly followed by storage in visual working memory and accompanied by visual imagery. Retrieval from long-term memory (Fukuda & Woodman, 2017; Hanslmayr et al., 2016; Staresina et al., 2016; Waldhauser, Braun, & Hanslmayr, 2016); storage, and prioritization of perceptual representations in working memory (Fukuda &

Woodman, 2017; van Ede, 2018; van Ede, Jensen, & Maris, 2017); and visual imagery (Barrett & Ehrlichman, 1982; Salenius, Kajola, Thompson, Kosslyn, & Hari, 1995; Slatter, 1960), have all previously been associated with attenuation of alpha oscillations. The current findings are thus in line with this large body of prior work. In contrast to this work, in the current study, these individual processes were never explicitly tasked to the participants. Rather, here, these processes may constitute the natural chain of events that support adaptive memory-guided attentional templates.

2.5.2. ERPs

In addition to the alpha effects, Experiments 1 and 2 each also revealed significant ERPs associated with target-identity anticipation. Moreover, like the alpha modulation, these potentials did not differ significantly between the predictive-easy and difficult S1s in Experiment 1. The two ERP effects consisted of a CNV and a late posterior potential. Both of these have been previously found in associative learning tasks (Rose et al., 2001; Stokes et al., 2014). However, in this previous work, S1 predictions were coupled to response probabilities, a confound I ruled out in Experiment 2.

The CNV is a classic signature of temporal and response anticipation (Donchin, Tueting, Ritter, Kutas, & Heffley, 1975; Walter et al., 1964), and is likely to reflect the anticipation of the target – here shown to be strengthened by foreknowledge of the identity of the ensuing target.

The late posterior positive potential may relate to the processing of S1 when it predicts a specific target, or serve as a link between the S1 and the S2 item. The exact functional contribution of the late positive potential in the current task is difficult to pinpoint. Its posterior topography and time course are compatible with a few different possibilities. Identification of the S1 as a relevant, predictive stimulus may have triggered a P300, which has a long history as a marker of stimulus relevance or meaning (Johnson, 1986; Polich, 2007; Squires, Squires, & Hillyard, 1975). Alternatively, it may have reflected the process of recalling the associated target (Donaldson & Rugg, 1999), therefore providing a link between S1 and S2. A similar potential has also been noted during the orienting of spatial attention (Brignani, Lepsien, Rushworth, & Nobre, 2009), raising the possibility of an analogous mechanism for orienting attention to identity-defining stimulus attributes.

Importantly, in Experiment 2, both the predictive and non-predictive S1 indicate that a task-relevant target would appear in 1000 ms in the center of the screen, and all trials required a response. The only difference was that the predictive S1 indicates *which* item is likely to appear. Accordingly, this provides compelling evidence that these ERPs, like the alpha attenuation, are sensitive to the expectation of the particular identity of the forthcoming item.

It remains to be investigated whether the effects shown here are contingent on knowledge of the location and timing of an upcoming event. By design, space and time were always reliable in the current work. While contrasts with non-predictive S1s allowed us to eliminate any neural correlates that were attributable to purely spatial and temporal predictions, I cannot rule out that the observed modulations might still reflect the interaction between identity-based anticipation and the known spatial and temporal attributes of the anticipated stimulus. That is to say, it is of yet unclear whether the same results would be obtained for identity-based predictions in the absence of spatial and temporal predictions. At the same time, of course, in the real world, spatial, temporal, and identity-based predictions are often bundled.

2.5.3. Interactions between attentional templates and perceptual difficulty

In Experiment 1 I found a significant interaction between S1 predictiveness and target difficulty (easy or difficult) for both error rates and reaction times. This could indicate either stronger template activation prior to the onset of the target or greater impact of a similarly strong template. Interestingly, I did not find neural evidence for such an interaction in the identified alpha attenuation or ERPs during the period between S1 and S2. As previously stated, one may have expected that a more difficult target would call for a stronger activation of the template. However, the data do not speak to this conclusion. On the one hand, I cannot rule out differences in the extent of template pre-activation that could not be detected with the current methods. There may be other neural correlates of perceptual identity preparation that do depend on expected target difficulty, which I was unable to measure. On the other hand, the results invite us to consider whether and how similar levels of template activation may result in differential performance benefits. It is possible that the same attentional templates will be more effective when incoming stimuli are harder to perceive. In this scenario, the consequences of pre-activation of relevant neuronal populations may critically depend on the strength of neuronal activity triggered by incoming stimulation, playing a greater faciliatory role when incoming stimulation is weaker or more ambiguous.

2.5.4. Conclusions

Taken together, the current results clearly reveal proactive preparation for the identity of a target – based on frequent associations between successive stimuli. Proactive templates bring significant benefits to perceptual performance and are accompanied by the attenuation of alpha oscillations and modulations of ERPs, including the CNV. I here demonstrate this while matching spatial and temporal predictions, as well as target probability and response demands between experimental conditions. Isolating identity anticipation based on memory was instrumental to the aims of the thesis, establishing a foundation from which to investigate the more multifaced aspects of memory-based anticipation that occurs in natural behavior, such as concurrent anticipation of the what, where and when of upcoming percepts. Having established the existence of memory-based identity templates, it becomes possible to study their quality with greater granularity. In line with the veridical vs. adaptive dichotomy introduced in Chapter 1, the next study explores whether the content of identity templates can be distorted to emphasize particular features that are more informative in a given task. The present study did not require or encourage any template distortion, since the task was to report an exact match between incoming stimuli and the memory template. However, in cases where one template feature (e.g., color) becomes more diagnostic of a target, would this feature come to be over-represented in the attentional template? Chapter 3 addresses this question.

3. Adaptive Associative Memory Templates

3.1. Abstract

Attentional templates can be retrieved from associations within long-term memory in order to guide ongoing behavior (Chapter 2). In this chapter I ask whether templates retrieved from memory are necessarily faithful reproductions of the encoded information or if associative-memory templates can be adapted in service of current task demands. Across two experiments, participants learned associations between four shapes and four colored gratings, each with a characteristic combination of color (green or pink) and orientation (left or right tilt). On each trial, observers saw one shape followed by a grating and indicated whether the pair matched the learned shape-grating association. I manipulated the probability of the non-matching (lure) gratings between blocks (Experiment 1) or over time within a trial (Experiment 2). In Experiment 1, some blocks contained lures that were most likely to differ in color but not tilt, while in other blocks this pattern was reversed. The results revealed that participants formed adaptive templates that emphasized the distinguishing information between lures and targets. Specifically, the most diagnostic feature dimension was overemphasized within the template leading to more accurate and faster responses to common vs. rare lures. In Experiment 2, I tested whether these templates could be adapted over the course of a trial. In this case, the grating appeared either after a short (1.25 s) or long (2.5 s) delay. If a lure appeared early, it was most likely to differ from the target in one specific feature dimension (e.g. color was diagnostic), whereas, if it appeared late the opposite was true (e.g. tilt was diagnostic). Here, observers performed better to the common lures from the short delay, irrespective of the time they were probed. Therefore, the template was initially adapted to the diagnostic feature of the first interval but was not

updated in time for the second delay. Taken together, the findings show that templates retrieved from memory can be biased and adapted according to what is most informative for task performance. I speculate that template adaptation likely occurs during the retrieval of the template and is tuned towards early task demands.

3.2. Introduction

Attentional templates are the mental representations that we use to facilitate the detection and identification of task-relevant sensory inputs (Bundesen et al., 2005; Carlisle et al., 2011; Chelazzi et al., 1998, 1993; Desimone & Duncan, 1995; Duncan & Humphreys, 1989; Geng & Witkowski, 2019; Kok et al., 2017; Treue & Martinez Trujillo, 1999). A common intuition is that the attentional template is a replica of the anticipated target. However, templates may not always be fully specified or strictly veridical. For example, we are able to search for an item at the category level – that is, without knowing the exact features of the target (Cunningham & Wolfe, 2014; Peelen & Kastner, 2011) – and objects that are functionally related to the target template can guide visual search (Boettcher, Draschkow, et al., 2018).

Recently, an alternative and more flexible view of attentional templates has emerged with their functional nature prevailing over their mimetic quality. Rather than a faithful reproduction of the anticipated target, a template may be adapted to optimize perceptual performance within a current task context. A striking recent demonstration comes from work by Yu and Geng (2019). Whereas templates resembled targets when searching for a target color (orange) in arrays with an equal likelihood of distractors from either side of the color space (yellower and redder), templates became distorted when distractors were all drawn from one side of the color space (e.g. yellower). In this case, when targets and distractors were linearly separable, the attentional template was asymmetrically sharpened and repelled away from distractors (becoming redder). In such a framework, certain template features may be overemphasized and/or distorted – adapted – in service of anticipated task requirements (for a recent review see: Geng & Witkowski, 2019).

The aim of the current chapter was to investigate whether the adaptability of attentional templates also generalizes to templates retrieved from associative memory (Chapter 2). In typical laboratory studies of attentional templates, the template is explicitly provided before the start of a trial or block of trials (Carlisle et al., 2011; Chelazzi et al., 1993; J. Lee & Geng, 2019; Navalpakkam & Itti, 2007; van Driel et al., 2017; Yu & Geng, 2019). However, in the real world, attentional templates are often derived from learned associations between stimuli, such that one stimulus (A) predicts another (B) (Chapter 2; Higuchi & Miyashita, 1996; Hutchinson & Turk-Browne, 2012; Kok et al., 2017; Stokes et al., 2009). Can the template for a stimulus B – retrieved based on its long-term memory association to stimulus A – *also* be adapted in service of anticipated task demands?

There are at least two reasons why templates retrieved from long-term memory may not be as adaptable as those explicitly presented. First, unlike biases of templates that are explicitly presented, templates retrieved from memory cannot benefit from differences at the encoding stage (Reeder, Hanke, & Pollmann, 2017; Serences, Ester, et al., 2009). Second, unlike forms of adaptation that may gradually develop when repeatedly using the same template for many trials in a row (as in e.g. Navalpakkam & Itti, 2007; Yu & Geng, 2019), the same associative memory template could require distinct adaptations at each instance of retrieval depending on the current context. As such, these templates cannot benefit from potential biases at the level of the stored memory trace, but rather from biases that develop during or after their retrieval. Therefore, an important open question remains whether templates that are retrieved from associative memory can also be functionally adapted to optimally serve behaviour.

If observers are able to use adapted templates retrieved from long-term memory it will further be important to identify at what stage this biasing occurs. These associative memory templates are likely retrieved from long-term memory and subsequently stored in working memory until they are needed

(Atkinson & Shiffrin, 1968; Fukuda & Woodman, 2017). As such, there are two main candidates for when an associative memory template may be adapted: at retrieval from long-term memory through retrieval of only a sub selection of relevant features or within working memory where the retrieved template may be particularly moldable (D'Esposito & Postle, 2015; Griffin & Nobre, 2003; Niklaus et al., 2017; Rerko & Oberauer, 2013). Because the use of templates in everyday situations will often involve the recruitment of templates from associative memory, it is critical to demonstrate if and when associative-memory templates, can be flexibly adapted in service of anticipated task demands. Here, I set out to answer these questions.

I developed a task in which observers retrieved templates from memory, in contexts in which the visual features of a dual-feature template were differently informative in different blocks (Experiment 1) or moments within a trial (Experiment 2). If observers rely on veridical memory templates – in which they weigh both visual features equally – their behavior should not differ according to our manipulation of the most informative template feature. In contrast, if templates retrieved from memory can be functionally adapted to overemphasize the most useful template feature, there should be evidence for this in behavioral response patterns as a function of the context. Moreover, by changing the most distinguishing feature-dimension task constraints within a trial (Experiment 2), it is possible to interrogate when these templates become adapted, and whether such adaptation can be flexibly altered after initial retrieval.

3.3. Experiment 1: Can Associative Memory Templates be Flexibly Adapted?

3.3.1. Methods

Participants

Twenty-two participants took part in the Experiment 1. Two participants failed to complete the experiment due to problems with the testing computer. The remaining 20 participants (13 female, 1 left handed) were between 21-35 years old with an average age of 24.6. All Participants had normal or corrected-to-normal vision, provided written consent prior to participation, and were compensated at a rate of £10 per hour.

Task and Procedure

All experimental procedures were reviewed and approved by the Central University Research Ethics Committee of the University of Oxford. Participants completed the experiment in a group testing room with a capacity of 20 people, although no more than 12 people were tested at one time. Participants each sat approximately 50 cm from the monitor (Dell U2312HM Monitor, 1920 x 1080 resolution; refresh rate 60 Hz). The experimental script was generated using the Psychophysics Toolbox (Brainard, 1997) on MATLAB (version 2014b, The Mathworks Inc., Natick, NA, USA).

Prior to the main experimental task, participants took part in a learning phase in which they learned the relationship between four distinct shapes (square, star, circle, and triangle) and four gratings (Figure 3.1.A; shape-grating pairings were randomized across participants). The four gratings each had a unique combination of color (green or pink; RGB values: [46, 142, 141], [246, 37, 113]) and tilt (left or right, at an angle of ± 45 degrees). On each trial, observers saw a central shape (1.4°) followed by a central grating (9.1°) and indicated whether the pair matched the learned shape-grating association. Shapes appeared for 1000 ms followed by a 1000-ms delay. The grating then appeared for 150 ms. Observers were required to respond within 1500 ms with the ‘j’ key (right hand response) if they believed the stimulus matched the learned shape-grating association in memory (in both color and tilt) and the ‘f’ key (left hand response) if the grating did not match the learned association (because either the color or the tilt was different). On 50% of the trials a match grating was presented whereas on the other 50% of the trials a non-matching lure stimulus was presented. Non-matching (lure) stimuli matched the learned association in one dimension but not the other (Figure 3.1.B). In the learning phase, participants were required to complete a practice block of 30 trials with above 80% accuracy before moving on to the test phase (main verification task) that followed the same procedure, except the type of non-match gratings was explicitly manipulated across blocks. These pairings were further reinforced during the experiment in the form feedback after each trial.

The main verification task consisted of 30 blocks with 30 trials in each block – with each block lasting approximately 3 minutes. The central experimental manipulation concerned how informative each feature modality was for discriminating targets from lures, by changing the probability that the lure shared either color or tilt with the target across blocks. Specifically, in half of the blocks the lure was most likely to differ from the target in color but not tilt (color-distinguishing blocks), while in the other blocks this pattern was reversed (tilt-distinguishing blocks). Block type was randomly interleaved across 15 pairs (e.g. AB BA BA AB). I will refer to these non-match gratings as *common* and *rare* lures. Common lures occurred on 85% of lure trials whereas rare lures occurred on only 15% of lure trials. For example, in a color-distinguishing block, in 85% of the lure trials the grating differed from the expected target in color but not tilt (common), while the remaining 15% lure trials the grating differed from the expected target in tilt but not color (rare). If templates retrieved from associative-memory are veridical and therefore both features are weighed equally, then there should be no effect of lure type. However, if the attentional template for the upcoming stimulus is adapted such that the distinguishing feature is emphasized, then the common lure will be more distinct from the template compared to the rare lure. It follows that participants should be more accurate and faster to reject common lures compared to rare lures.

After observers responded, or if they were too slow with their response (> 1500 ms; this occurred in fewer than 1% of trials across all participants), feedback was provided in the form of a color change in the fixation cross – green for correct response, red for incorrect responses, and orange on trials in which a response was not provided in time.

All feature values were equally probable within each block and participants were never tasked with repeatedly searching for a specific feature value throughout a given block. What defined the target, the common lure, and the rare lure, was solely and completely dependent on the preceding shape cue. Because the shape cues were randomly interleaved across trials, participants were required to upload a fresh attentional template on each trial. As such, there is no possibility that low-level adaptation or short-term priming could explain any differences in performance in the different block types.

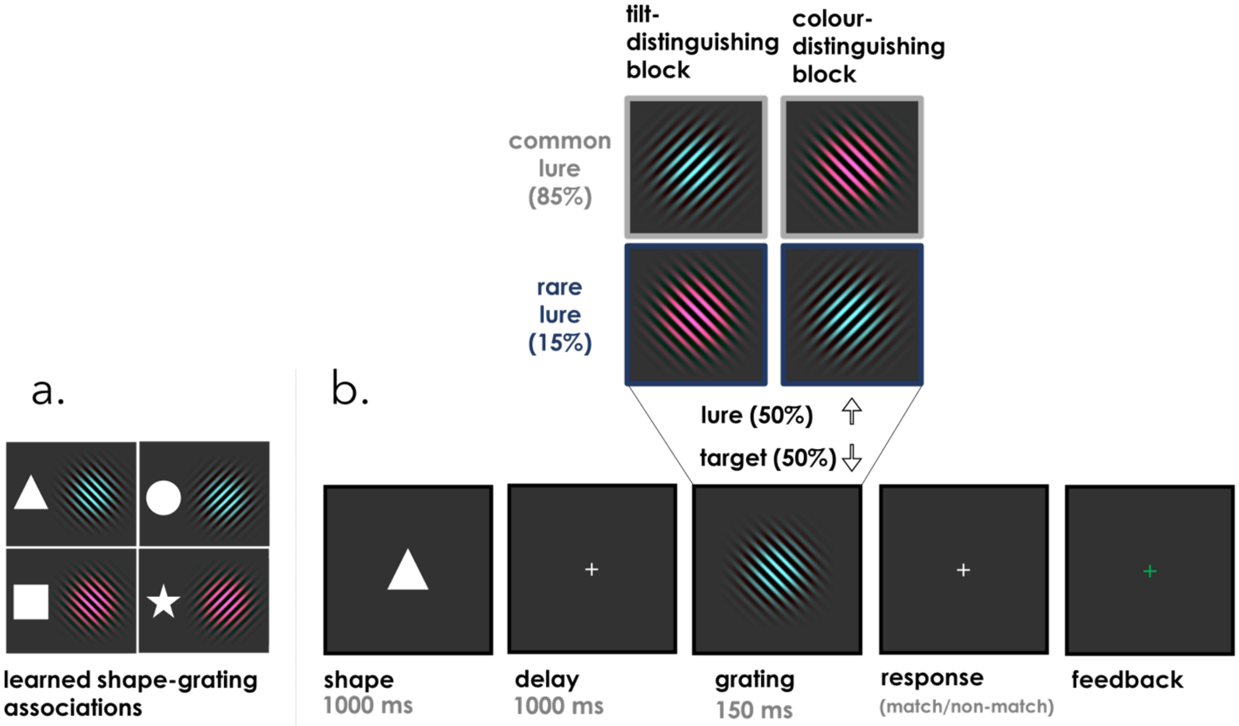

Figure 3.1: Trial schematic of experiment. a) In the Learning phase, each observer learned to associate four shapes to four gratings. These pairings were randomized across participants. b) In the main experimental task, each trial began with the presentation of one of the 4 shapes followed by a blank and then a grating. Observers indicated whether the shape-grating matched the association held in memory in both color and tilt. On non-match (lure) trials, observers could be presented with a common or rare lure, such that lures were more likely to have a different color ('color-distinguishing block') or a different tilt ('tilt-distinguishing block'). Note that if a different shape was presented – all shapes were equally likely – the roles of the gratings would also change. For example, if a star was presented than the expected grating would be a pink-right tilted target.

Statistical Analysis

Because the critical manipulation (of block type) regarded the type of lures that could occur, the analyses were restricted to lure trials. Generalized linear mixed-effects models (GLMMs) with a binomial distribution were used to analyze the percentage false alarms (FAs), and Linear mixed-effects models (LMMs) to analyze reaction times (RTs) (correct rejection trials only) in a procedure similar to the approach described in Helbing, Draschkow, and Võ (2020) and Draschkow and Võ (2017). These analyses were run using the lme4 package (version 1.1-17; Bates *et al.*, 2015). I used mixed-effects models because they hold multiple benefits over a more traditional approach to analysis of variance. Importantly for the current study, these approaches are more reliable in unbalanced designs when different conditions have different trial numbers

– e.g. common versus rare lures (Baayen, Davidson, & Bates, 2008). All GLMMs and LMMs were fitted with the maximum likelihood criterion. For the GLMMs, where I report regression coefficients β with the z statistic and use a two-tailed 5% error criterion for significance, the p-values for the binary accuracy variable are based on asymptotic Wald tests. For the LMMs, I report β with the t-statistic and apply a two-tailed criterion corresponding to a 5% error criterion for significance. The p-values were calculated with Satterthwaite's degrees of freedom method using the lmerTest package (version 3.1-0; Kuznetsova, Brockhoff, & Christensen, 2017). Pairwise tests following significant interactions were further investigated using the lsmeans package (Lenth, 2016) with Tukey post-hoc correction.

In Experiment 1 there were two main independent variables of interest: lure type (Common vs. Rare) and, block type (Colour-distinguishing and Tilt-distinguishing). The comparisons were modelled using sum contrasts, in which the grand mean of the dependent measure served as the intercept. For binary responses such as FAs in the GLMM approach, the coefficients are represented by logits. Each model began with a maximal random-effects structure (Barr, Levy, Scheepers, & Tily, 2013) that included intercepts for each participant, as well as by-participant slopes for the effects of lure type and block type. Full models such as these often fail to converge or lead to overparameterization (Bates, Kliegl, Vasishth, & Baayen, 2015). Therefore, I used a principal component analysis (PCA) of the random-effects variance-covariance estimates to identify overparameterization for each fitted model and removed random slopes that were not supported by the PCA and did not contribute significantly to the goodness of fit in a likelihood ratio (LR) test (Bates, Kliegl, et al., 2015). The random-effects structures for the optimal models for both the GLMM and LMM included the subject intercepts as well as by-subject slopes for block type. Further details regarding the models and model comparisons can be found in the analysis script (https://osf.io/xhu5j/).

To be sure the results did not depend on the chosen approach, I also conducted traditional repeated-measures ANOVAs for both FA and RT. These showed equivalent results, and can also be found in the analysis script provided. The ggplot2 package (version 3.1.0; Wickham, 2009) was used for plotting results.

3.3.2. Results

Observers correctly identified the target in 93.5% ± 1.1 (m ± s.e.m) of match trials, indicating that they had properly learned the shape-gratings associations, and were able to complete the task to a high standard. Because the critical manipulation focused on the non-match (lure) stimuli, the main analyses focused on lure trials.

If participants are utilizing an adapted template biased by which feature dimension is most informative in the current context (here manipulated at the level of blocks), then this template should become optimized for rejecting the type of lure that was more common in this context (block). In line with this possibility, observers were less likely to erroneously identify a common (vs. a rare) lure as a target (i.e. false alarm) (Figure 3.2.a) ($\beta = .13$, $SE = .05$, $z = 2.52$, $p = .01$; *rare lures*: 10.02% ± .42; *common lures*: 8.06% ± .42). This pattern was similarly present for both visual features (Figure 3.2.b): there was no effect of block type ($\beta = -.03$, $SE = .07$, $z = -.50$, $p = .62$) and no interaction between lure type and block type ($\beta = .03$, $SE = .05$, $z = .56$, $p = .57$).

Analysis of reaction times yielded complementary evidence. Observers were faster to reject common lures (628 ms ± 3.62) compared to rare lures (654 ms ± 3.62) (Figure 3.2.c) ($\beta = .02$, $SE = .003$, $t = 4.23$, $p < .001$). There was again no effect of block type ($\beta = .001$, $SE = .005$, $t = .233$, $p = .82$). However, for RT, block type and lure type did show a significant interaction ($\beta = -.008$, $SE = .002$, $t = -2.93$, $p = .003$) (Figure 3.2.d). Post-hoc comparisons revealed that there was no significant effect of lure type on RT in the tilt-distinguishing block ($\beta = .0072$, $SE = .0078$, $z = .918$, $p = .79$), while there was a significant effect of lure type in the color-distinguishing block ($\beta = .0398$, $SE = .0078$, $z = 5.06$, $p < .001$). That is, observers were faster in rejecting common lures that differed from the template in color but not tilt (as opposed to tilt but not color).

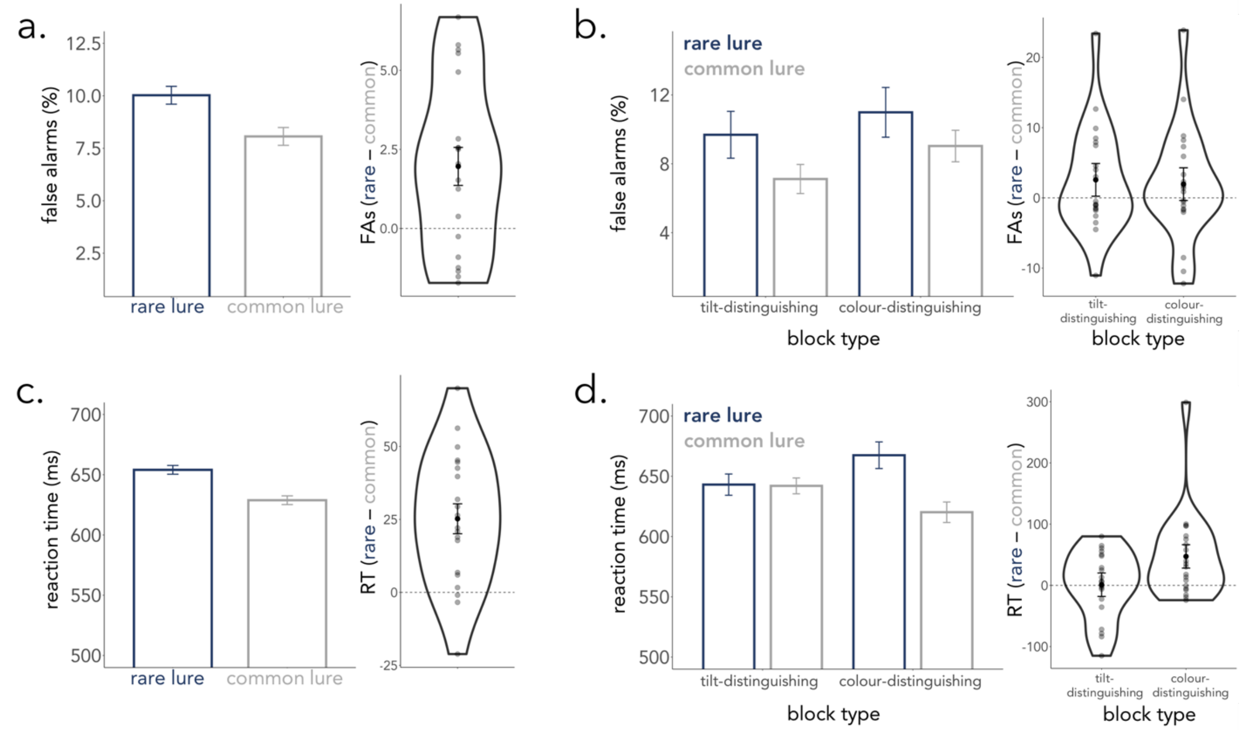

Figure 3.2: Functional adaptation in attentional templates retrieved from long-term memory associations. a) Observers made fewer false alarms to common lures than rare lures. The right panel (violin plot) shows the effect of lure type on false alarms (rare – common) for individual observers, represented as dots. The average effect is plotted as a black dot and error bars represent the standard error (this is true for all panels). b) The effect of lure type was similar across block types. c) Participants were faster to correctly reject a common lure compared to a rare lure. The effect of lure type is plotted for individual observers in the right panel. d) The effect of lure type on reaction times was only significant in the color-distinguishing block. That is, observers were significantly slower to respond to lures that shared their color with the target template when color was usually the distinguishing feature.

3.4. Experiment 2: When are Associative Memory Templates Adapted?

The results in Experiment 1 demonstrate an adaptive interpretation of attentional templates based on associative memories. When a common lure was presented, participants were less likely to confuse these for targets (decreased false alarms) and were faster to correctly reject these gratings, compared to when a rare lure was presented. Importantly, the effect was independent of the frequency of any of the stimulus features. The findings clearly show that the template for the upcoming stimulus was adapted according to the common type of lures in a block, such that the most diagnostic and therefore informative feature dimension – for distinguishing between targets and common lures – was prioritized in the template.

Although Experiment 1 clearly revealed utilization of a biased template that was adapted to the task demands, it did not clarify when the template information becomes biased to facilitate behavior. Because benefits to behavior persisted as the diagnostic feature changed between blocks (color or tilt), it is conceivable to rule out any permanent distortion of the template stored in long-term. Multiple possibilities remain. On the one hand, observers could have selectively prioritized the diagnostic feature during retrieval of the template. Alternatively, they could have retrieved the template in its unweighted form and subsequently prioritize the diagnostic feature while the template was being held in working memory.

In Experiment 2, I set out to understand when a template is adapted for the current task context. To do so, I manipulated the probability of lure stimuli over the course of a trial rather than between blocks. Specifically, gratings could either appear *early* (1250 ms) or *late* (2500 ms) within a trial. Lure stimuli that appeared early were more likely to match the target in one feature dimension but not the other, this pattern was reversed for late lures. If participants are able to adjust the template 'online', that is within working memory, observers should still be faster and more accurate in their responses to common compared to rare lures, and this should be true for both early and late lures (despite different feature-dimensions being more distinguishing at these two intervals). However, if participants retrieve a template in its optimal state and are not able to flexibly adjust the template within working memory, then they should retrieve the veridical template as the probability of each lure type was equal across the trial (given that short and long delays were equally likely).

3.4.1. Methods

Participants

Twenty participants (12 Female, 1 left handed) between the ages of 19 and 35 with an average age of 26.2 years old took part in the Experiment 2. All Participants had normal or corrected-to-normal vision, provided written consent prior to participation, and were compensated at a rate of £10 per hour.

Task and Procedure

Many of the procedures between Experiment 1 and 2 were repeated with a few key differences highlighted below. Once again observers had to complete a learning phase in order to learn the shape-grating relationships before moving on to the main experimental task. In the main task, observers were tasked with determining if the presented shape and grating matched the leaned association, as in Experiment 1. Critically, the shape itself was now presented for 250 ms followed by a blank interval of either 1 s or 2.25 s meaning the grating was presented at one of two possible delays: 1.25 or 2.5 s. Match and lure gratings were equally likely across the two delays. For half of the participants, if a lure was presented after a short (1.25 s) delay it was likely to match the expected target in color but not tilt, whereas if a lure appeared after a long delay (2.5 s), it would likely share the expected target's tilt but not color (Figure 3.3).This was reversed for the other half of the participants. Common lures – those lures which adhered to this pattern – were presented on 85% of lure trials, and rare lures – those which violated this rule – made up the other 15% of trials. If observers can flexibly adjust the retrieved template within working memory, then they should be more accurate and faster to respond to common vs. rare lures for both delays. If instead observers retrieve the veridical template and are not able to update their expectations within a trial, then there should be no difference in the responses to the different types of lures. In total observers completed 28 blocks of 40 trials each.

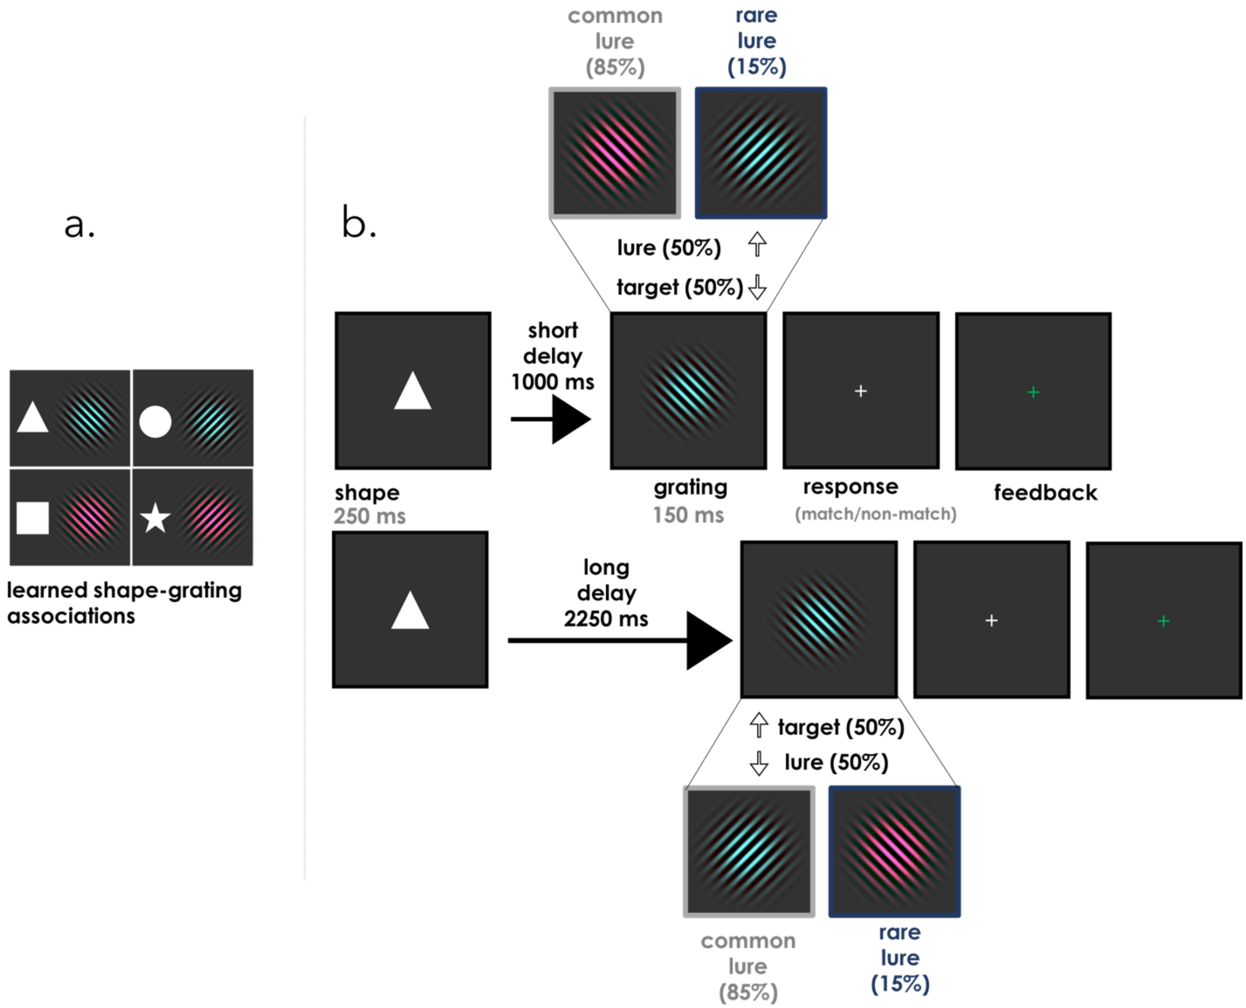

Figure 3.3: Trial Schematic for Experiment 2. (a) participants were required to learn shape-grating pairings as in Experiment 1. (b) Gratings appeared after either a delay of 1.25 (short) or 2.5 s (long). For some participants, when the delay was short color-distinguishing lures were more likely, and when the delay was long, tilt-distinguishing lures were more likely. The opposite was true for other participants.

Statistical Analysis

The analysis approach was similar to that of Experiment 1, except rather than considering block type as an independent variable, in Experiment 2 delay length (short vs. long) was included into the models. The random-effects structures for the optimal models for both the GLMM (FAs) and LMM (RT) included the subject intercepts as well as by-subject slopes for delay length. Further details regarding the models and model comparisons can be found in the analysis script (https://osf.io/xhu5j/).

3.4.2. Results

Like in Experiment 1, observers performed this task with a high level of overall accuracy, correctly identifying the targets on 88.9% ± 2.0 (m ± s.e) of match trials. The main analysis focused on the lure trials. If participants can adapt the attentional template within working memory then they should show better performance (lower false alarms and faster response times) on common compared to rare lures for both delay intervals.

As shown in Figure 3.4.a, there was no significant main effect of lure type (β =.02, *SE* = .05, *z* =.42, *p* = .68; *rare lures*: 9.31% ± .54; *common lures*: 8.69% ± .54). There was a significant effect of delay length, with observers committing fewer false alarms for trials with longer delays (β = .16, *SE* = .07, *z* =2.17, *p* = .03). Most importantly, these two factors showed a significant interaction (β = .27, *SE* = .05, *z* = 5.4, *p* < .001) as clear from the pattern in Figure 3.4.b. Post hoc comparisons revealed that observers had significantly fewer false alarms for common lures compared to rare lures during the short delay (β = .60, *SE* = .13, *z* = 4.6, *p* < .001) and this pattern reversed for the long delay (β = -.51, *SE* = .16, *z* = -3.25, *p* = .006). This indicates that observers were more likely to false alarm to the stimulus that would be considered a rare lure during the short delay, independent of when it was presented.

A similar pattern occurred for reaction times as seen in Figure 3.4.c and Figure 3.4.d. Again, there was no main effect of lure type (Figure 3.4.c; β = -.002, *SE* = .003, *t* = -.76, *p* = .44). Delay had a significant effect on RT with observers responding more quickly to lures after a long delay (β = 0.025, *SE* = .005, *t* = 4.92, *p* < .001). Again, lure type and delay interacted significantly (Figure 3.4.d; β = .006, *SE* = .002, *t* = 2.48, *p* = .01), although post hoc comparisons revealed there was no significant difference of lure type at either the short (β = .009, *SE* = .007, *z* = 1.18, *p* = .64) or long delay (β = .-017, *SE* = .007, *z* = -2.39, *p* = .08).

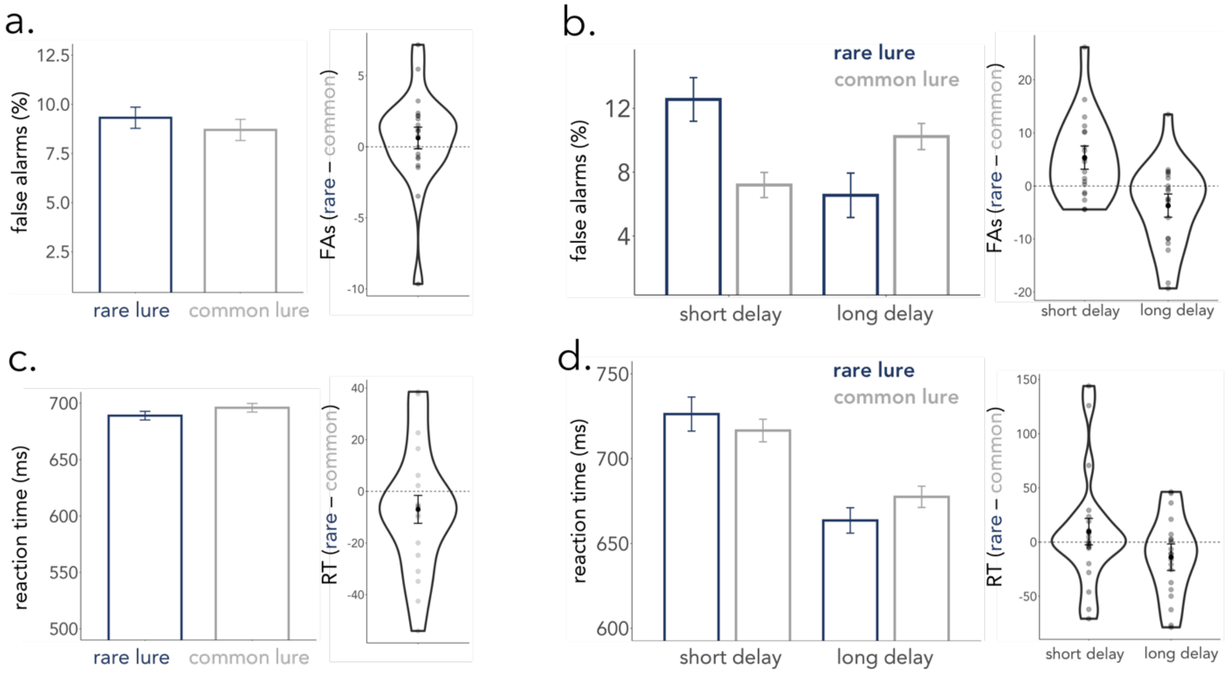

Figure 3.4: Interaction between lure type and delay length reveals observers adapt their templates to what is immediately necessary. (a) There was no main effect of lure type for FAs as would be expected if observers were able to update the representation within a trial – subject wise differences between rare and common lures are shown in the right panel (b) However, there was a significant interaction between lure type and delay length. Observers were significantly better for common lures at the short delay, and worse for common lures during the long delay. This pattern indicates that observers adapted their template for the short delay but did not adjust this template for the long delay. A qualitatively similar pattern can be observed in the reaction time. Again, there was no main effect of lure type (c), but there was a significant interaction when considering delay length (d).

3.5. Discussion

Across two experiments I present evidence that attentional templates retrieved from associative memory can be adapted to current task constraints. In Experiment 1, this was true despite the fact that the relevant feature dimension varied block wise (with blocks lasting only a few minutes). Moreover, within each block the feature values of the relevant attentional templates themselves changed from trial to trial, showing that the prolonged and repeated use of the same template is not necessary for a template to become adapted to the current context. This was also true in Experiment 2. Additionally, it is relevant to note that observers completed the task with high accuracy and the features were easily discriminable.

These results complement and extend on previous work on optimizing attentional templates to current task demands, which had focused primarily on templates that were explicitly provided at the beginning of a trial of block of trials, such as in traditional working memory or visual search tasks. For example, it has been shown that observers can selectively encode only the relevant feature of an object (Reeder et al., 2017; Serences, Ester, et al., 2009) as well as update working memory content after encoding, such that the representation is biased towards the most relevant feature dimension (Hajonides et al., 2019; Niklaus et al., 2017; Park et al., 2017). Moreover, when targets and distractors are linearly discriminable within a certain feature dimension, target representations are biased away from the distractors (Bauer et al., 1996; Becker et al., 2010; Geng et al., 2017; Geng & Witkowski, 2019; Hodsoll & Humphreys, 2001b; Navalpakkam & Itti, 2007; Yu & Geng, 2019). Building on this previous work, the current study provides evidence that attentional templates retrieved from associative memory are also adaptable to a particular context – an important advance as many of the attentional templates we use to guide our actions are retrieved from associations in long -term memory.

In Experiment 2, I tested if templates retrieved from long-term memory associations can be adapted flexibly within a trial. In particular, I asked at what stage template adaptation may occur. From Experiment 1, it was clear that templates are not stored in an adapted form within long-term memory, as the necessary adaptation changed from block to block. This left two likely candidates: the diagnostic feature could be selectively retrieved from long-term memory resulting in an adapted template, or the unweighted features were retrieved in their veridical form and then adapted within working memory. The results replicated the ability to adapt the template to enhance performance according to the most likely early common lure. However, I found no evidence for flexible updating within the trial – that is, adaptation of the template to anticipate the diagnostic features of late lures. Given previous work showing that items in working memory can be flexibly prioritized across time (De Vries, Van Driel, Karacaoglu, & Olivers, 2018; de Vries et al., 2017; van Ede, Niklaus, et al., 2017; Zokaei, Board, Manohar, & Nobre, 2019) the pattern of results may seem surprising. This may indicate that observers are unable to update these template representations in working memory, but the templates are rather retrieved in a biased state.

Given that all lure types were equally likely throughout a trial – common and rare lures were only in reference to a given delay – it is puzzling that observers prioritized features relevant to only one of the possible lures and did not simply retrieve the veridical template. One possibility is that observers optimized their template for the more difficult target. Observers committed more false alarms and were slower to respond during the short delay trials. As such, if it is only possible to retrieve a single adapted template, then perhaps the best strategy is for that template to be optimized for the more *difficult* target. Alternatively, it may be the case that observers simply did not have enough time to update their template for the late interval, and in turn optimized behavior for what was immediately relevant – the early interval. These intervals were chosen based on previous work showing that objects within working memory can be selectively prioritized on the basis of the same intervals (van Ede, Niklaus, et al., 2017). Even so, moving forward it would be necessary to extend the intervals to be sure that observers are given enough time to update the attentional template. In any case, the results from Experiment 2 clearly provide further evidence for the use of non-veridical attentional templates from associative memory.

In both Experiments, the results show that the attentional template itself – not only a spatial priority map – may be functionally biased towards a relevant feature dimension. Information used to guide spatial attention during visual search – attentional priority maps (Wolfe, 1994) – and the template information to which we match incoming sensory information, are not often explicitly distinguished within the literature (Wolfe, 2020). Within the context of visual search, “dimension-weighting” has been argued to interact with the attentional priority map in order to guide spatial attention to a target (Liesefeld & Müller, 2019; Müller et al., 1995). In both experiments, observers performed a task other than visual search in which a single stimulus was presented centrally, and therefore did not require shifts in spatial attention, nor a search among competing distractors. This allowed me to isolate the influence of our experimental manipulation of context, at the level of the template. The current results demonstrate that across two experiments attentional templates retrieved from long-term memory associations are flexible and functionally adapt to the context in which they are retrieved, and expected to be utilized.

So far, the studies in the thesis have demonstrated that (associative) memories can generate attentional templates that carry information for identifying targets (Chapter 2) and that specific features within these templates can be selectively utilized depending on how informative they are for guiding adaptive behavior (Chapter 3). The next logical step is to ask whether the brain goes beyond anticipating the content of a stimulus in order to prepare for the to–be-performed task. In Chapter 4, I investigate whether motor responses associated with anticipated stimuli are also incorporated into the attentional template and used to proactively guide task performance.

4. Prospective Actions in the Attentional Template

4.1. Abstract

Attentional templates are the internal representations that guide our behavior in the external world. It is therefore vital not only to understand how relevant sensory information is maintained within the attentional template but also to understand whether, additionally, prospective actions are integrated. To address this question, I manipulated motor associations linked to specific stimulus features. Observers were asked to reproduce the tilt of a bar after a delay, with reporting of left vs. right tilt associated with prospective manual actions with the left vs. right hand, respectively. To examine the possibility of action encoding (or "imprinting") – and to separate this from subsequent action preparation – during the delay period, observers were asked to complete a secondary motor task. The results showed that observers imprint the initial action plan into the attentional template, seconds before it is needed. Moreover, this is done even in the face of an intervening motor task and predicts eventual response times. Behavioral results further showed a modest but significant effect of prospective actions on the intervening motor task. Taken together, these results provide convincing evidence that prospective actions are incorporated into the attentional template early and this prospective plan guides both relevant and irrelevant actions.

4.2. Introduction

Our cognitive systems are fundamentally action oriented, in place to facilitate the movements needed in natural behavior (Cisek, 2007; Gibson, 1979; Wolpert, Ghahramani, & Jordan, 1995). Attentional templates are internal representations that are used to guide our actions in the external world and represent a critical component in a cascade of processes which ultimately serve goal-oriented behavior. As such, template representations will depend on their eventual use (Myers et al., 2017).

With the goal of understanding ongoing natural behavior, it is not only important to gain insights into how sensory information is encoded or maintained but additionally, how we will use these sensory codes. Previous work has shown that the visual information within an attentional template is dependent on its eventual use. Serences and colleagues (2009) showed that templates only contain the relevant feature dimension from an object needed to complete the current task. Moreover, in my own work, I have shown that when searching for real world objects in complex scenes it is not only the features of the target object that are contained within the attentional template, but the features of related “anchor objects” are also used to guide visual search more efficiently (Boettcher, Draschkow, et al., 2018). Although this previous work has highlighted the apparent functional nature of attentional templates, it has not thus far considered the actions associated with template use (for recent review see: van Ede, 2020).

Recent work, interrogating the interplay between actions and our sensory representations has shown that action preparation can at least in part help to explain behavioral benefits found in retro cuing paradigms (Schneider et al., 2017). Moreover, visual and motor selection from working memory has been shown to occur in parallel, suggesting that the actions associated with sensory information are readily available (van Ede, Chekroud, Stokes, et al., 2019). This leaves open important questions: are associated action incorporated into the attentional template, *when* does this occur, and once these actions are integrated into the attentional template do they modulate ongoing behavior?

To address these questions observers were asked to keep in mind an attentional template that was associated with a particular visual-memory-guided action after a 6.5 s delay. The tilt of left vs. right oriented bars required subsequent reproduction using the left vs. right hand, respectively. During the delay period observers completed a secondary motor task in which they simply responded to the locations of a series of presented squares. The introduction of this secondary task allowed me to (1) study whether the prospective action was incorporated into the template already at visual encoding even if the action could not yet be prepared for execution, as well as to (2) probe the prospective action plan through potential influences on behavior in the secondary task. After this motor task, observers were asked to report back the tilt from the beginning of the trial. In a preview of the results, action information was imprinted into the attentional template at a very early stage – nearly 6 seconds before it would be needed. This early imprinting was also associated with reaction times at the end of the trial. Further, the encoded action showed a modest but significant influence on the intervening motor task. Taken together, the findings provide clear evidence that actions are incorporated into the attentional template early for subsequent use, despite the demands of an intervening task.

4.3. Methods

4.3.1. Participants

Twenty-five volunteers ranging in age between 18 and 32 (mean age 23.71 years; 17 female; one left-handed) participated in this experiment. All observers reported normal or corrected-to-normal vision, provided written consent before participation, and were reimbursed £15/hour. The experiments were approved by the Oxford Central University Research Ethics Committee.

4.3.2. Task and Procedure

For a schematic of the task see Figure 4.1.a. Each trial started with the presentation of a pre-cue in the form of a change in the color of the central fixation cross. This pre-cue was 100% informative such that the central fixation cross briefly took the color of one of the two upcoming bars and predicted with 100% reliability that this item would be relevant after the memory delay – thus enabling selection of the relevant template at the visual encoding stage. The location of the pre-cued item was not known until the visual encoding display appeared. Pre-cues were displayed for 250 ms. The encoding display with the two bars appeared 1000 ms after pre-cue onset, and was presented for 250 ms. Visual encoding was followed by a memory delay of 6.5 s.

Observers were also asked to complete an intervening motor task during this delay period. In the intervening task, four white squares (0.4° width and height) were presented sequentially to the left and right of fixation (at 5.7° visual angle). Participants were instructed to respond to these squares by pressing the corresponding left and right response key within 500 ms. Each trial consisted of two left squares and two right squares in a randomized order. If participants failed to respond in time, or pressed the wrong key, the square flashed red for 100 ms as feedback. Responses were made with the same keys as those required for the primary memory task. The first square of the intervening task always occurred 1500 ms after onset of the encoding display for the primary task, while the last square always occurred exactly 3 s later (4500 ms after memory encoding). Times between square-onsets varied between 800, 1000, and 1200 ms (in randomized order).

At the end of the delay (6.5 s after encoding) the fixation cross changed to gray (referred to as the "probe") to indicate that observers should begin their response. Participants were asked to reproduce the tilt of the bar as accurately as possible. For reproduction, participants pressed the "\" (backslash) or the "/" (forward slash) key on the keyboard using their left or right index finger respectively. The time between probe onset and response initiation was unlimited. After response initiation, a dial appeared around fixation with the same diameter as the length of the bars. The response dial always began in the vertical position and moved at a speed of 1/8 degree per ms either clockwise (right response) or counter clockwise (left response) for as long as either button was pressed. Participants were instructed to release their key press when the dial reached the tilt of the item in memory. Only one key could be used to reproduce the item's tilt, and participants were not able to adjust their response after key release. Critically, because the response dial always started in the top vertical position and because it could not be rotated beyond ±90° (horizontal), a leftward (rightward) tilted bar could only be accurately reported with a left (right) key press. As a consequence, the hand required for responding was directly linked to the tilt of the bar that was probed; while a detailed visual representation remained relevant for guiding the precision (duration) of this manual reproduction report. Upon key release, feedback was presented for 250 ms above the central fixation cross in the form of a number; ranging from 0 to 100 (with 100 indicating perfect report and 0 indicating that the reported tilt was perpendicular to the required report). Inter-trial intervals were randomly drawn between 500 and 800 ms.

To perform the task, participants sat in a dimly-lit booth in front of a monitor (22-inch Samsung SyncMaster 2233; resolution: 1680 × 1050 pixels; refresh rate: 100 Hz; screen width: 47 cm) at a viewing distance of approximately 95 cm. Prior to the start of the experiment, participants performed several practice trials to become familiar with the task. The experiment was divided into two consecutive sessions, each of approximately 50 minutes. In each session, participants completed 7 blocks of 40 trials, yielding 560 trials in total.

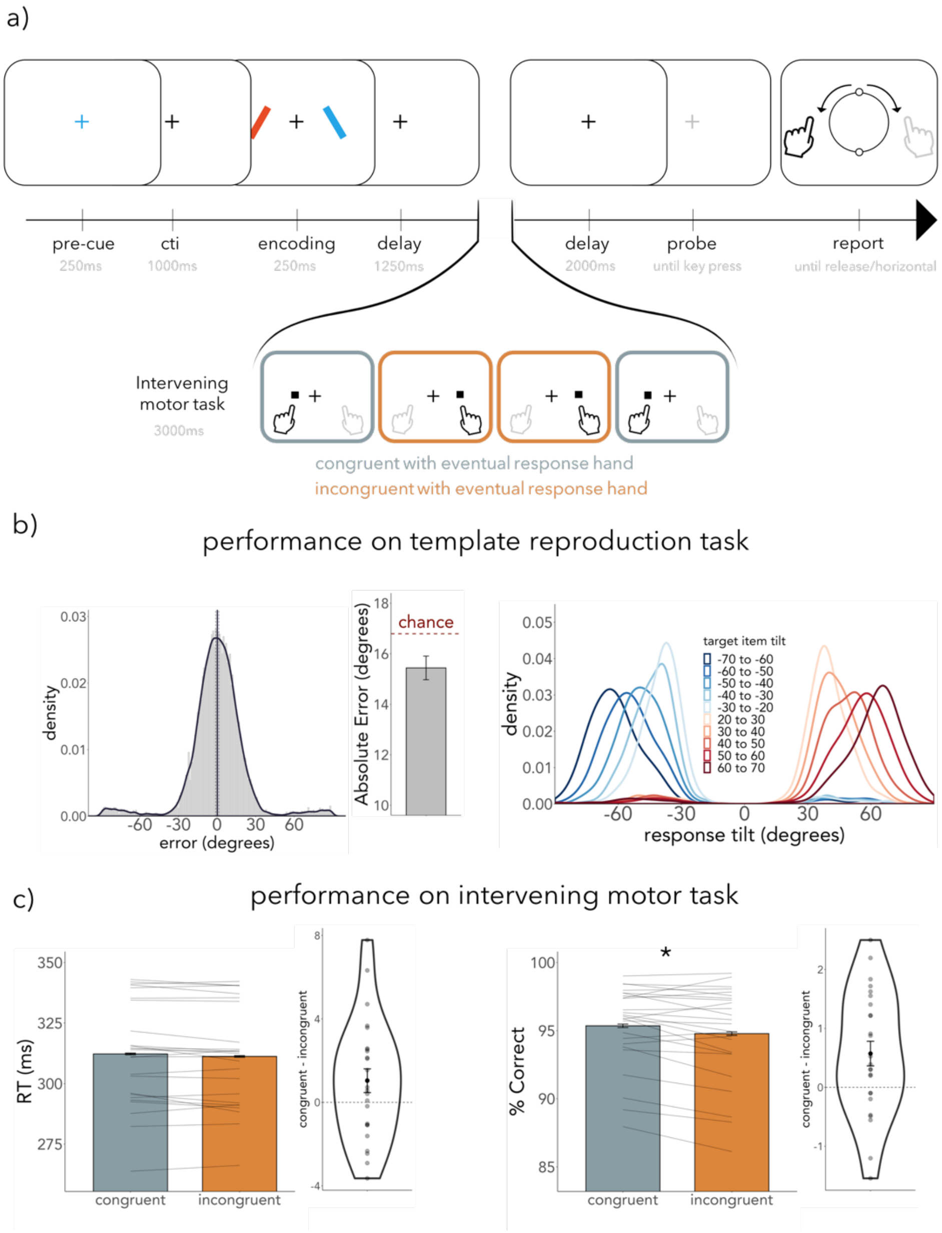

Figure 4.1: Trial schematic and behavioral results (a) Participants were required encode an attentional template of a tilted bar that was directly linked to a specific action (left/right button press). During the delay period observers completed an intervening motor task in which they responded to either a left or right square using the corresponding response hand. These responses could either be congruent (match) with the prospective response hand of the main task, or incongruent. At the end of the delay participants reported the tilt of the encoded template by pressing the left/right response button to start a dial in the counter-/clockwise direction and released the button when the dial matched the template in memory. The dial did not precede past 90 degrees and therefore a leftward/rightward tilted bar could only be reproduced using the left/right button. (b) performance in the main task was relatively high (left panel: mean response error represented as solid line and standard deviations represented with dotted lines). In the center panel it is clear that observers are responding significantly more accurately than would be expected from chance (dotted red line). Chance was determined by calculating the difference between the actual target tilt and a random permutation of the response tilts given a particular response hand (i.e. the actual response when the left hand was required were randomly shuffled and a difference was calculated). In the right panel, the average response divided by the actual target item tilt is presented. As can be seen, overall observers, used the visual information to guide their response. (c) the performance in the intervening motor task for response times and percent correct as divided by congruency. Gray lines represent individual subjects in the bar graphs and error bars represent the standard error of the mean. The violin plot shows the difference between congruent and incongruent responses with individual subjects plotted as gray circles.

4.3.3. Behavioral analysis

To confirm that observers were responding accurately and using the detailed template representation rather than responding in a binary fashion (left/right) the absolute response errors were compared to a chance level. Chance was computed by taking the difference between the actual target tilt and a random permutation of the response tilt within a response hand. That is, the data were divided by tilt of the target (leftward and rightward) and the actual responses were subsequently shuffled across trials. Trials in which the wrong hand was used (on average 5.4% of trials across subjects) were excluded from the calculation. This permutation was repeated 1000 times and averaged per subject. Specifically, with this analysis I asked, what would the expected error be if observers were only responding to the direction of tilt (left/right) rather than the detailed template information? This was compared to the absolute response errors with a paired samples t-test.

To understand the effect of an action incorporated into the attentional template on ongoing behavior, behavioral analysis focused on the intervening task. Only trials in which the correct response hand was used after the probe in the main task were included in this analysis. Each response was categorized as either congruent (trials which matched the prospective action) or incongruent (trials which did not match the prospective action). I then compared response times and accuracy of congruent and incongruent trials using paired-samples t-tests.

4.3.4. EEG Acquisition

EEG data was acquired using Synamps amplifiers and Neuroscan data acquisition software (Compumedics). Sixty-one electrodes were distributed across the scalp using the international 10–10 positioning system. The left mastoid was used as the active reference, and the right mastoid measurement was included in order to derive an average-mastoid reference offline. The ground was placed on the left upper arm.

Additionally, vertical and horizontal EOG electrodes were used to monitor for eye blinks and eye movements. During acquisition, data were low-pass filtered by an anti-aliasing filter (250-Hz cutoff), digitized at 1000 Hz, and stored for offline analysis.

4.3.5. EEG Preprocessing

The preprocessing and analysis scripts for both experiments can be found as html files and as reproducible scripts (jupyter notebooks; (Kluyver et al., 2016) at https://osf.io/xhu5j/. All EEG data analysis was conducted in MNE-Python (Gramfort et al., 2013). The data were down-sampled to 250 Hz. To regress out eye-movement activity, an Independent component analysis (ICA; Jung et al., 2000) was used to decompose the data – which was high-pass filtered at 1 Hz – into sixty temporally independent components. Eye-movement components were detected by first correlating the filtered data with the electrooculography (EOG) and subsequently, when needed, manually selecting a subset of typical component maps and identifying the best group match to them (Viola et al., 2009). Selected components were then removed from the data. Trials were segmented from -500 ms to +7000 ms relative to the visual-encoding period. Average activity over the 500 ms preceding the stimulus onset was used as a baseline against which all amplitudes were calculated. Finally, epochs with especially high variance were discarded following visual inspection. Trial removal was performed without knowledge of the experimental conditions to which individual trials belonged. After trial removal 496 ± 10 out of the total of 560 trials were retained for analysis.

4.3.6. Time-Frequency Analysis

For the time-frequency analysis, Morlet wavelets were convolved with the data between 5 and 40 Hz. For each frequency, I used a fixed 300-ms time window such that the number of cycles changed with the frequency. Activity in predefined motor (C3/C4) electrodes was contrasted between trials in which the prospective response hand was contralateral versus ipsilateral to the electrode of interest. This is expressed as a normalized difference (i.e., ((contra-ipsi) / (contra + ipsi)) x 100) and averaged these contrasts across the left and right electrodes. Topographical maps of lateralization were obtained by contrasting left response trials and right response trials using the same normalized difference (e.g. (left response – right response) / (left response + right response) x 100). As the main interest in this experiment was in motor preparation the analysis focused on mu-alpha and mu-beta band activity. To extract time courses of alpha and beta lateralization from these time-frequency contrasts, I averaged across the predefined mu-alpha (8-12 Hz) and mu-beta (13-30 Hz) bands for the prospective response-hand lateralization in the motor electrodes (C3/C4). The analysis script can be found at https://osf.io/xhu5j/.

4.3.7. Relation with behavior

To investigate the relation between prospective action preparation and eventual template-guided behavior, I split the trials by their response initiation times after the probe. I performed a median split separately for each experimental condition (left/right actions), effectively regressing out low-level differences in response times due to item response hand. I then calculated the same contralateral-vs-ipsilateral contrast of prospective action preparation as described above, separately for the trials followed by a “fast” (< median) or a “slow” (> median) response initiation time after the memory probe. I quantified the strength of the action imprinting signature by focusing on the identified 400-800 ms imprinting window (this window was set before evaluating the outcomes of this analysis, based on the relevant window identified in the preceding analyses). I then used paired-samples t-tests to compare fast and slow trials within this window.

4.4. Results

4.4.1. Behavioral Results

At the end of each trial observers reproduced the tilt of the attentional template. It was important to know that observers did so with a relatively high accuracy, and the reproduced tilt was guided by the detailed visual template information rather than a simple motor response, i.e. left/right. In Figure 4.1.b it is clear that observers were able to reproduce the actual tilt of the target with high accuracy with an average absolute error of 14.2° ± .71° (m±sem). This was significantly below chance (*t*(24)=3.48, *p* = .002, *d* = .70) which was computed by randomly permuting responses and target tilts within the same direction (left/right). The right panel of Figure 4.1.b further confirms that observers were responding to detailed template information regarding the tilt.

To understand whether the action incorporated into the attentional template influences ongoing behavior, I turn to the analysis of the intervening motor task Figure 4.1.c. Here observers had to respond as quickly and accurately as possible to the location of a square presented to the left/right of the fixation cross. On each trial, half of these responses matched (congruent) the prospective response hand associated with the attentional template. Importantly these effects are independent of the visual location (left/right) of the attentional template as this was counterbalanced across trials. Although there was no significant difference in the response time between congruent and incongruent targets (*t*(24 = 1.82, *p* = .08, *d* = .37), there was a modest yet significant effect on the accuracy of these responses with congruent targets being correctly identified more often (*t*(24) = 2.75, *p* = .011, *d* = .55).

4.4.2. Time Frequency Results

To understand when the action plan was incorporated into the attentional template I looked at motor recruitment in the form of mu-alpha/beta lateralization relative to the eventual response hand. If the action is immediately brought into the attentional template (“action imprinting”) we should see motor recruitment immediately following encoding – nearly 6.5 seconds before the response is actually needed and despite the intervening motor task. Figure 4.2.a shows two significant clusters of motor recruitment relative to the eventual response hand. The early cluster (p=.03) precedes the intervening motor task while the later cluster (p<.001) shows gradual response preparation before the probe. The time series Figure 4.2.b shows this activity separately for the mu-alpha (8-12Hz) and mu-beta (13-30Hz) bands. In both frequency bands there is a significant early lateralization (alpha: p=.009, beta: p=.01) that precede the actual response by approximately 6 seconds. Again, this is followed by gradual action preparation in expectation of the probe (alpha: p<.001, beta: p<.001). Topographic insets in Figure 4.2.b show that this activity localized to the canonical motor electrodes C3/C4.

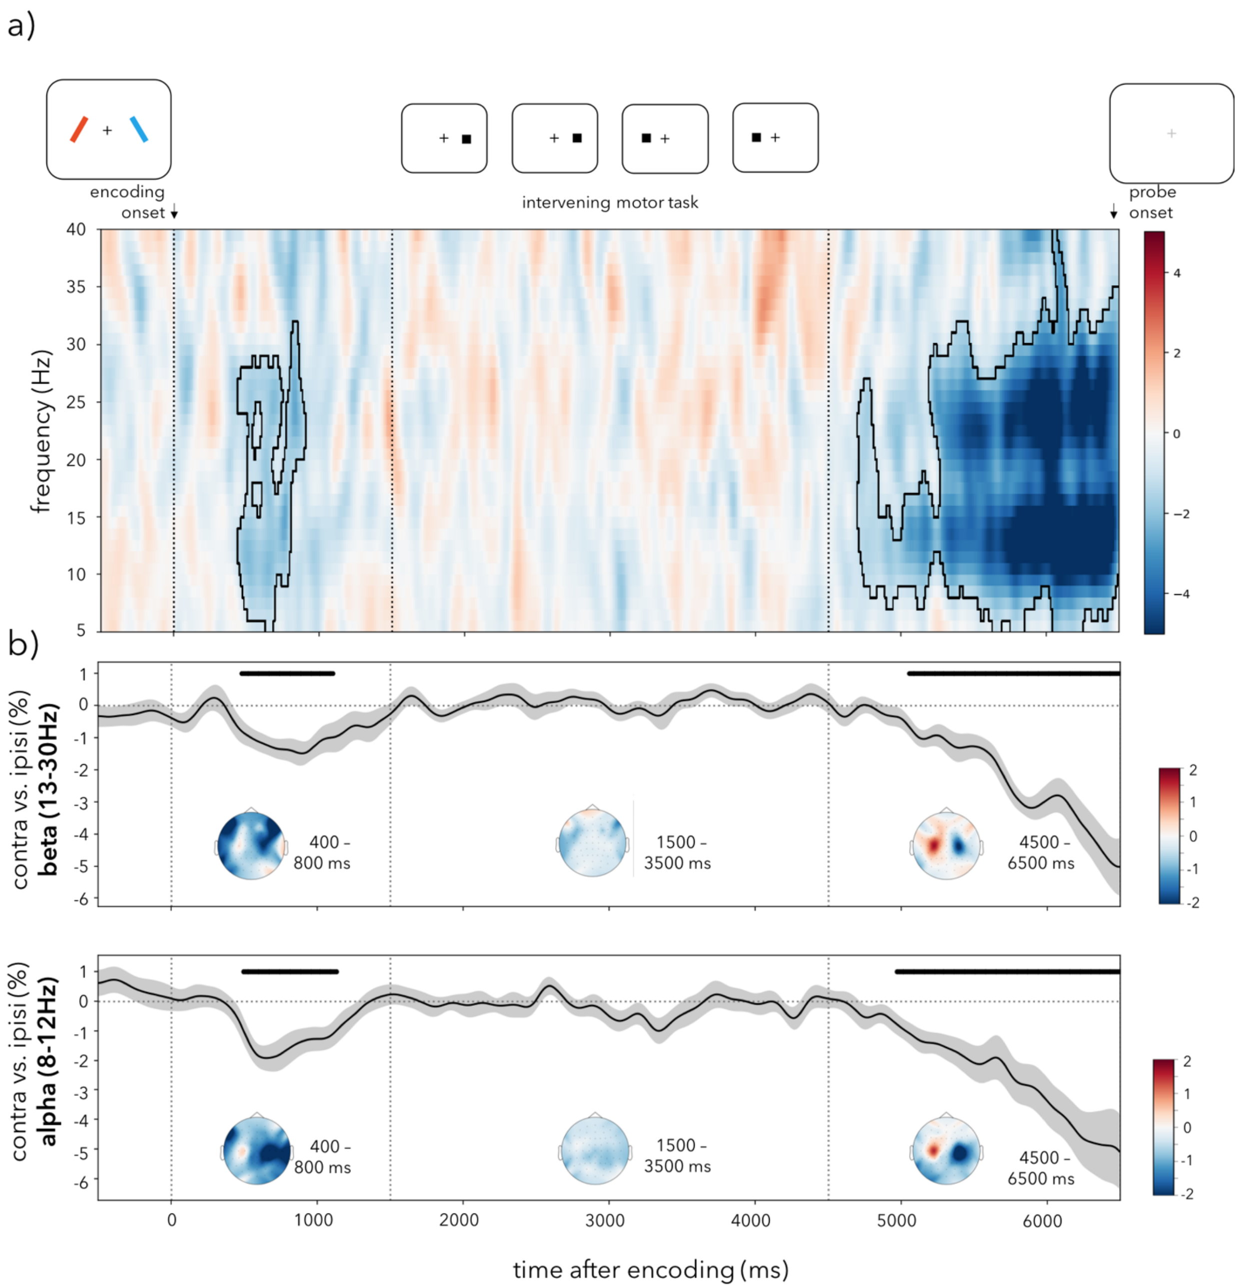

Figure 4.2: Motor recruitment relative to the eventual response hand. (a) Neural lateralization in motor electrodes C3/C4 relative to the prospective response hand associated with the attentional template in the primary task, across all frequencies, (b) as well as for the mu-alpha (8-12 Hz; bottom) and mu-beta (13-30Hz; top) bands. Insets show topographies at representative time-windows. The black outlines and black horizontal lines indicate significant clusters (all clusters: P < 0.05).

4.4.3. Relation with Behavior

If the early motor recruitment observed here represents an early incorporation of action plans into the attentional template, then this should be related to the eventual behavior. Specifically, on trials in which observers have a strong action representation in the attentional template they should be faster to respond to the target. Figure 4.3.a shows lateralization in the motor electrodes relative to the prospective action as a function of the eventual reaction times (median split). Figure 4.3.b shows this activity collapsed within the relevant action imprinting window (400-800ms). This window was chosen a-prior based on the preceding results. Within the alpha band it is clear that the action imprinting significantly predicts the eventual response time nearly 6 seconds later despite the intervening motor task. That is faster reaction times were associated with stronger lateralization in the action imprinting window ($t(24) = 2.1$, $p = .045$). This effect was not significant in the beta band ($t(24) = 1.18$, $p = .24$) although, the direction of the difference was comparable.

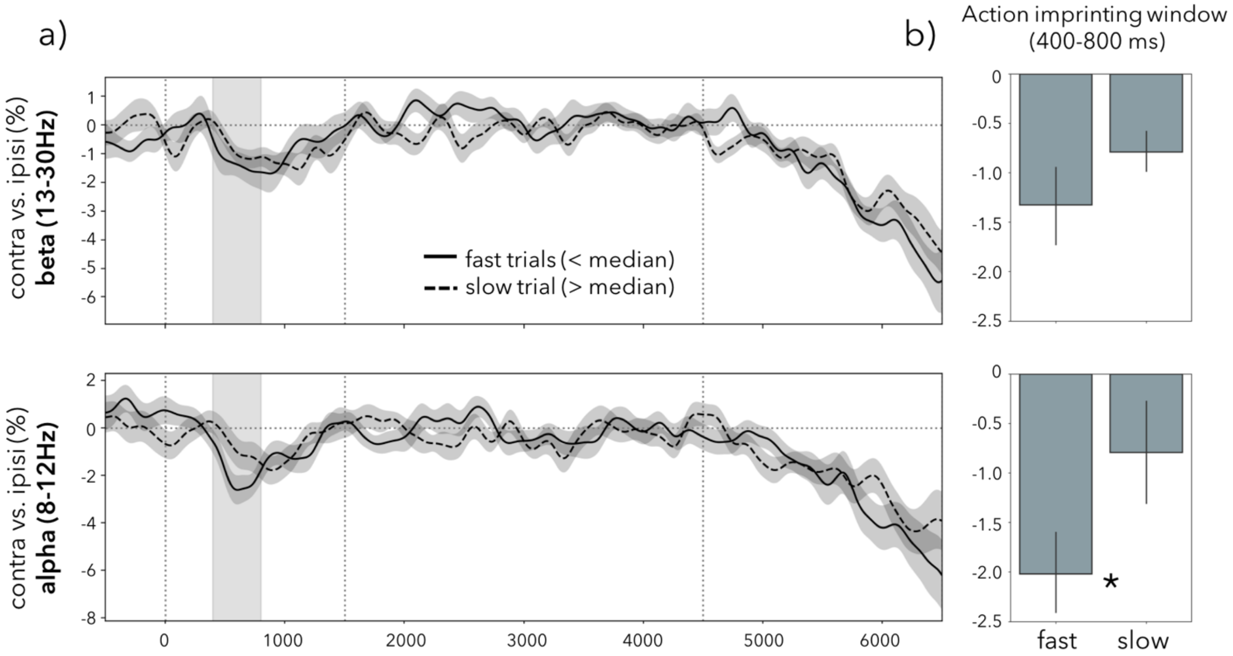

Figure 4.3: The relation of action imprinting with behavior. (a) Motor recruitment (lateralization in C3/C4 relative to the prospective action) as a function of response times after the memory delay (median split) separately in mu-alpha (8-12Hz; bottom) and beta (13-30Hz; top) bands. (b) Action imprinting for fast and slow trials, averaged across the 400-800 ms imprinting window. This window was chosen a priori (based on the preceding results) and is also indicated in the grey shadings in panel a. Error bars and shadings indicate ± 1 s.e.m. calculated across participants.

4.5. Discussion

Attentional templates serve behavior. However, within the literature templates are often considered from a purely sensory perspective, in isolation from actions. In the current work I tracked the neural signatures of motor recruitment during a template reproduction task in which the eventual response was closely linked to the attentional template. In doing so, it was possible to understand when prospective actions are incorporated into the attentional template. Moreover, by inserting intervening motor task into the delay period I was able to investigate how prospective actions within the attentional template may affect ongoing, yet template irrelevant, behavior. The results show that actions are quickly incorporate (imprinted) into the attentional template after encoding. This is even true despite the fact that observers knew they would be completing a secondary motor task. Moreover, although the actions within the template were irrelevant to the intervening motor task these action representations have a small yet significant effect on behavior. Finally, the early signature of action imprinting significantly predicts behavior nearly 6 seconds later, despite an intervening motor task. That is, when observers show stronger imprinting of the eventual action into the attentional template, they are faster to respond at the end of the trial.

The incorporation of actions into the attentional template may serve two complementary purposes: making these representations more effective and making them more robust. First, it ensures that we are ready to act upon the attentional template whenever the sensory information within the template becomes relevant for guiding behavior, the associated action is already available. Note that this task did not strictly require participants to encode and memorize the prospective action as participants could have relied on a purely visual memory trace and considered the action only after the probe. This action-readiness should in turn yield faster behavior which was confirmed by the data, showing that stronger action imprinting predicted faster response times after the memory delay. Action imprinting into the template occurred even though an intervening task negated the utility of preparing the action for ensuing execution. A second purpose of such action imprinting may be to make template representations more robust. By adding a supporting feature upon encoding which contains future action plans alongside detailed sensory information – the template may be more resilient to subsequent interference. This is perhaps similar to the benefits observed by distributing memory traces across the hemispheres (Li, Daie, Svoboda, & Druckmann, 2016; van Ede, de Lange, & Maris, 2014).

The current results build on and extend several lines of prior research that have linked visual working memory and action (van Ede, 2020). For example, it has been established that overt actions can themselves influence visual working memory performance (Hanning, Jonikaitis, Deubel, & Szinte, 2016; Heuer, Crawford, & Schubö, 2017; Ohl & Rolfs, 2017; Postle, Idzikowski, Della Sala, Logie, & Baddeley, 2006), and that brain structures that control action also participate in visual working memory (Jonikaitis & Moore, 2019; Theeuwes, Belopolsky, & Olivers, 2009; van Ede, Chekroud, & Nobre, 2019). Here I investigated the complementary direction of this interaction, revealing that the selection of detailed visual information into the attentional template, and therefore into working memory, also naturally recruits the preparation for prospective manual actions (see also González-García, Formica, Liefooghe, & Brass, 2020; Schneider et al., 2017; Singhal, Monaco, Kaufman, & Culham, 2013)). Moreover, by adopting a task in which actions required the guidance of detailed visual shape information from memory and by assuring observers were using this detailed sensory information to respond, this work also complements prior work that has focused on actions that required guidance from purely spatial locations from memory (Cisek & Kalaska, 2005; Curtis, Rao, & D'Esposito, 2004; Fuster & Alexander, 1971; L. H. Snyder, Batista, & Andersen, 1997) or guidance from detailed shape information from perception (Gallivan, Bowman, Chapman, Wolpert, & Flanagan, 2016; Gallivan, Chapman, Gale, Flanagan, & Culham, 2019; Gutteling et al., 2015). Further, these results are in line with prior studies investigating the role of action affordances (Cisek, 2007; Gibson, 1979) in working memory, which focused on affordances inherent to the memory material itself, such as when retaining images of manipulable vs. non-manipulable objects (Downing-Doucet & Guérard, 2014; Mecklinger et al., 2004). As the current results show, such visually derived action associations can become effective within the context of a specific task, and in the time frame of a typical laboratory task.

Finally, the behavioral results here show a small, yet significant effect of the prospective action within the attentional template on ongoing behavior. Specifically, observers were more accurate to respond to targets in the intervening motor task that matched the eventual response hand. This is in line with previous work showing that sensory information within the attentional template can guide behavior in a secondary task (Dalvit & Eimer, 2011; Olivers, 2009; Olivers et al., 2006, 2011), though most prior studies have focused on eye movements as a measure of attentional capture, not manual actions.

Taken together, the current results provide further evidence of the richness of attentional templates – prospective and adaptive memory-based representations used to guide our behavior. That is, although templates are traditionally thought of as veridical copies of sensory information, the current work provides evidence that complementary motor codes are also incorporated into the attentional template –

and that this happens very soon after encoding. This furthers the idea that attentional templates are use oriented, designed to optimize efficient behavior.

So far, the studies in the thesis have focused on the identity- and response-related properties of attentional templates associated with a target stimulus, showing that templates can contain sensory attributes for identifying targets (Chapter 2), which can be selectively prioritized depending on what is most informative in a task (Chapter 3), and that response associations can also be incorporated to facilitate responding (Chapter 4). In order to isolate the properties of templates themselves, simple identification and reproduction tasks were used, which reduced the complexity of the contexts in which we usually interact with objects in our environment. In everyday life, we use attentional templates to help us seek out items within crowded and temporally unfolding environments. The final set of empirical studies (Chapter 5) places objects back into crowded and dynamic environments to ask whether spatial and temporal associations can be bound to attentional templates to facilitate the recognition of target objects in these natural situations.

5. Spatiotemporal Information in the Attentional Template

5.1. Abstract

Searching for objects in crowded environments is a common task in natural human cognition. Its experimental ‘visual-search’ analog has proven instrumental for the investigation of attentional templates. However, one important dimension has been lacking in laboratory studies to date. Experimental tasks have primarily relied on static search displays, yet real-life search typically extends over time in dynamically changing environments. In everyday contexts targets may occur more frequently at particular locations and times. To simulate extended search in natural environments and investigate whether attentional templates can come to incorporate informative spatial and temporal associations to facilitate target guidance and recognition, I designed a new dynamic-search task. Participants searched for eight targets amidst stimuli that faded in and out of a display over several seconds. In each trial, one target appeared within each of four, spatially separated quadrants in a temporally predictable fashion. The other four targets were spatially and temporally unpredictable. Participants’ performance was significantly better for spatiotemporally predictable compared to unpredictable targets (Experiments 1-4). The effects were reliable over different patterns of spatiotemporal predictability (Experiment 2) and primarily reflected long-term learning over trials (Experiments 3,4), although, single-trial priming effects also contributed (Experiment 4). Eye-movement recordings (Experiment 1) revealed how spatiotemporal regularities guide attention proactively and dynamically. Taken together, these results show that spatial and temporal regularities can contribute to attentional templates and guide effective visual search.

5.2. Introduction

Visual search has been at the forefront of investigations of attentional templates. However, in a laboratory setting it is often devoid of an ecologically important dimension – time. Our environment contains multiple sources of informative spatial and temporal regularities that can, in principle, contribute to the attentional template and therefore guide behavior. Such guidance can prove particularly beneficial when facing multiple competing signals in our environment, as is the case in visual search. For example, if we search for a friend at a crowded train station, prior experience can help us to limit our search in space (e.g., knowing she will arrive from the platform), and also inform us when to search (e.g., knowing her train has already arrived will increase the likelihood she will soon emerge).

The attentional template is the mental representation used to guide spatial attention to candidate targets during visual search (Desimone & Duncan, 1995; Wolfe, 2020). These templates are a compilation of memories that may function on a relatively short time scale – such as the location or features of a target on the previous trial (Hillstrom, 2000; Kristjánsson & Campana, 2010; Maljkovic & Nakayama, 1994, 2000; Meeter & Olivers, 2006) – as well as on a longer scales such as our knowledge that we will find badgers on the ground and birds in the air (Võ et al., 2019; Wolfe et al., 2011). Visual-search experiments have yet to consider how prior experience with temporal associations may additionally also guide attention. In a dynamic world, in which our environment is constantly changing, spatial priorities fluctuate across time. Accordingly, it is important to consider how temporal predictions may contribute to the attentional template, and therefore help guide our attention in a changing scene.

It is known from simple perceptual tasks that observers can anticipate events in time and guide attention proactively when presented with informative cues. For example, explicit temporal cues improve perception and speed reaction times (Coull & Nobre, 1998; Nobre & Rohenkohl, 2014; Nobre & van Ede, 2017). These performance benefits are also found when temporal associations are stored in long-term memory (Cravo et al., 2017). Memory-based guidance of attention also extends to more complex spatial tasks with no direct manipulation of time (Chun & Jiang, 1998; Draschkow et al., 2018; Malcolm et al., 2016; Nobre & Stokes, 2019; J. J. Summerfield et al., 2006; Wolfe et al., 2011). The interaction of temporal anticipation with the spatial allocation of attention during visual search, has not been systematically studied.

In the current work, I introduce a new experimental framework to determine if spatiotemporal associations can contribute to the attentional template and play a role in dynamic visual search. Results from four complementary experiments show that spatiotemporal regularities, embedded within the environment, can lead to changes within the attentional template in order to guide performance. Further, it becomes clear that both short- and long-term memory traces are involved in this attentional guidance.

5.3. Experiment 1: Spatiotemporal Regularities in Dynamic Search Guide Behavior

In Experiment 1 I test whether spatiotemporal regularities embedded into a dynamic search paradigm are brought into the attentional template and therefore guide visual search. My hypothesis is that observers will be faster and more accurate in finding targets that are structured into a predictable pattern, compared to targets that are unpredictable. Further, by tracking eye-movements, it is possible to further understand how the spatiotemporal information within the attentional template may be facilitating behavior.

5.3.1. Methods

Participants

I tested 25 participants (age range 18-30, mean age = 23.4, 11 females). All Participants had normal or corrected-to-normal vision, provided written consent, and were compensated at a rate of £10 per hour.

Task and Procedure

All experimental procedures were reviewed and approved by the Central University Research Ethics Committee of the University of Oxford. Participants sat in a dimly lit booth at a distance of 100 cm from the monitor (22-inch Samsung SyncMaster 2233; resolution 1680 x 1050 pixels; refresh rate 100 Hz; screen width 47 cm). An eye tracker was used to monitor eye movements (see Eye-Tracking section below). The experimental script was generated using Psychophysics Toolbox (Brainard, 1997) on MATLAB (version 2014b, The Mathworks Inc., Natick, NA, USA).

Observers were instructed to find and click on eight small grey vertical lines which appeared and disappeared over the course of a trial (Figure 5.1.a). The search display consisted of four unique 1/F static noise patches that were generated for each trial and each quadrant extended approximately 18° (horizontal angle) X 12° (vertical angle). Over the course of a trial, forty distractor stimuli (tilted grey bars, 80°- 100°, RGB values [88,88,88]) and 8 target stimuli (vertical grey bars) appeared and disappeared at different times. These bars did not move in their location, but rather faded slowly in and out of view. Fade in time was set to 1.3 seconds (gradually becoming visible over 80 refresh-rate cycles until reaching maximum visibility). Then the target stayed on the screen for another 1.3 seconds, and faded out over 1.3 seconds. The timings of these events were chosen to minimize the effects of exogenous factors, such that attention was not captured by the sudden onset of any event. Each stimulus was ~.08° in length and ~.015° in width, and could appear anywhere in one of the four quadrants as long as it did not overlap with another stimulus.

Of the eight target stimuli, four were predictable. On every trial, they appeared with the same temporal onset from the start of the trial and in the same quadrant though the specific location within the quadrant was still random. Each trial was split into 4 equal time bins (2.5 seconds each) with a predictable target appearing at the temporal midpoint of each bin. The four unpredictable targets were distributed pseudorandomly, by assigning each to an onset derived from a uniform distribution across the full trial length. These constraints were set such that the targets were roughly evenly distributed throughout the trial, and to avoid too many target events occurring at one time. The quadrants and specific locations of the unpredictable targets were chosen randomly. The task outline is depicted in Figure 5.1.a. The probabilities of each target appearing in each quadrant are outlined in Table 5.1.

Table 5.1
Target Probabilities by Quadrant

	Quad 1	Quad 2	Quad 3	Quad 4
Pred Targ 1	1	0	0	0
Pred Targ 2	0	1	0	0
Pred Targ 3	0	0	1	0
Pred Targ 4	0	0	0	1
Unpred Targ 1	.25	.25	.25	.25
Unpred Targ 2	.25	.25	.25	.25
Unpred Targ 3	.25	.25	.25	.25
Unpred Targ 4	.25	.25	.25	.25

Observers completed 4 blocks of 40 trials. Each trial contained 8 targets (4 *predictable* and 4 *unpredictable*). Overall, there were 8 x 4 x 40 = 1280 target events per participant. After each trial, observers received feedback in the form of a number between 0 and 8 indicating how many targets they had found and were able to proceed at their own pace. Before beginning the experiment observers completed 8 practice trials in which there were no regularities and targets could appear anywhere at any time. The experiment lasted approximately 45 minutes.

Behavioral Analysis

Of the 25 participants, 20 had useable eye-tracking data. Behavioral data were analyzed using R (R Core Team, 2018) using the approach described in Helbing, Draschkow, and Võ (2020) and Draschkow and Võ (2017). Differences in analysis procedures between experiments are highlighted in the corresponding section. Generalized linear mixed-effects models (GLMMs) with a binomial distribution were used to analyze the percentage hits, and Linear mixed-effects models (LMMs) were used to analyze the reaction times for correct trials in all experiments. These analyses were run using the lme4 package (version 1.1-17; Bates *et al.*, 2015). I chose to use mixed-effects models as they hold multiple benefits over a more traditional approach to analysis of variance. Importantly for the current study, these approaches are more reliable in unbalanced designs when different conditions may have different trial numbers (this becomes critical Experiments 3 & 4) (Baayen et al., 2008). All GLMMs and LMMs were fitted with the maximum likelihood criterion. For the GLMMs, where regression coefficients β with the z statistic are reported and a two-tailed 5% error criterion for significance is used, the p-values for the binary accuracy variable are based on asymptotic Wald tests. For the LMMs, I report β with the t-statistic and apply a two-tailed criterion corresponding to a 5% error criterion for significance. The p-values were calculated with Satterthwaite's degrees of freedom method using the lmerTest package (version 3.1-0; Kuznetsova, Brockhoff, & Christensen, 2017). The ggplot2 package (version 3.1.0; Wickham, 2009) was used for plotting. In addition, I conducted traditional repeated-measures ANOVA. These showed equivalent results, and can be found in the analysis script provided (https://osf.io/xhu5j/).

In Experiment 1, there were two main independent variables of interest: predictability (Predictable vs. Unpredictable) and, within the levels of predictability, target order (1st, 2nd, 3rd, or 4th). Target order was centered and entered the model as a continuous predictor. The critical comparison between predictable and unpredictable targets was modelled using sum contrasts, where the grand mean of the dependent measure served as the intercept. For binary responses such as hits in the GLMM approach, the coefficients are represented by logits. Each model began with a maximal random-effects structure (Barr et al., 2013) that included intercepts for each participant, as well as by-participant slopes for the effects of Target Order and Predictability. Full models such as these often fail to converge or lead to overparameterization (Bates, Kliegl, et al., 2015). Therefore, I used a principal component analysis (PCA) of the random-effects variance-covariance estimates to identify overparameterization for each fitted model and removed random slopes that were not supported by the PCA and did not contribute significantly to the goodness of fit in a likelihood ratio (LR) test (Bates, Kliegl, et al., 2015). In Experiment 1, the GLMM's random-effects structure contained the subject intercepts as well as by-subject slopes for predictability and target order (i.e. the full model). The optimal LMM for predicting participants' reaction times contained the participant intercepts as well as the participant slopes for target order. Further details regarding the models and model comparisons can be found in the analysis script (https://osf.io/xhu5j/).

Eye Tracking Analysis

Participants were positioned in a chin rest 100 cm from the screen, and eye movements were recorded with the Eyelink-1000-plus desktop mount (SR Research, Ontario, Canada) at 1000 Hz. The raw data were first converted to a data matrix using a Matlab script. The raw data matrix contained the X-Y coordinates of gaze position (in pixel units) throughout the task. Data were then recoded to a single vector of spatial quadrant (i.e., each sample of X and Y coordinates was combined and recoded to a single number between 1-4, representing the spatial quadrant at which eyes were fixated at any given moment). The time series was then converted to a probability matrix for each quadrant separately, and split by trials to allow statistical analysis.

To quantify the guidance from the spatiotemporal predictions, I compared the mean probability of searching within a quadrant containing a predictable vs. unpredictable target, within a time window of 4.5 seconds, commencing 500 ms before stimulus onset. The comparison relied on a permutation test based on 5000 samples. The resulting distribution for each data point was compared to a critical t-value ($p<.05$) corrected for multiple comparisons based on the 't-max' method (Blair & Karniski, 1993; Westfall, Young, & Wright, 1993).

Further, in order to understand if the spatiotemporal information can not only guide visual search but also contribute to the process of determining if a particular item is indeed a target I also calculated "decision time". To do so, I calculated the time of first fixation on a target quadrant during the last dwell before a response. I subsequently subtracted this value from the response time to obtain decision time. I again used linear mixed models with predictability and target order as predictors in order to determine significance.

5.3.2. Results

Hits

In Experiment 1, participants were significantly more likely to identify a predictable compared to an unpredictable target (Figure 5.1.b (β = .15, *SE* = .02, z = 6.95, p < .001). Additionally, there was a main effect of the target order such that observers were better at detecting earlier targets (Figure 5.1.c) (β = -.09, *SE* = .02, z = -4.49, p < .001). However, there was no significant interaction between target predictability and target order indicating the effect of predictability was present throughout the course of the trial (β = -.002, *SE* = .02, z = -.159, p = .87).

RTs

Participants were significantly faster at finding predictable targets compared to unpredictable targets (Figure 5.1.d) (β = -.03, *SE* = .01, t = -3.27, p = .001), as well as, significantly faster at finding targets appearing early in the trial compared to late targets (Figure 5.1.e) (β = .06, *SE* = .004, t = 12.79, p < .001). Again, these two factors did not significantly interact (β = -.007, *SE* = .003, t = -1.95, p = .051) suggesting that the effect of predictability was present throughout the trial.

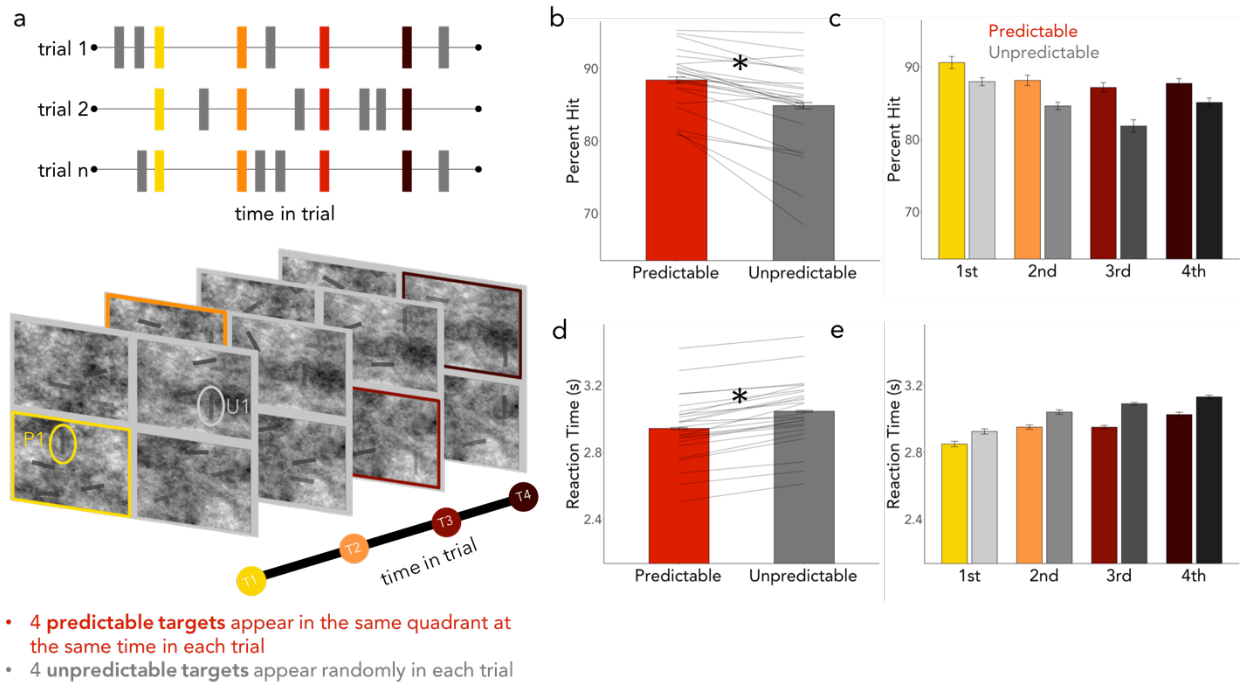

Figure 5.1: Spatiotemporal regularities in dynamic visual search guide behavior. a) The trial schematic for Experiment 1. On each trial, participants searched for 8 vertical bars among distractors – a yellow (grey) circle indicates the first predictable (unpredictable) target. Targets and distractors appeared and disappeared over the course of the trial. Above, the time course of a trial is depicted with target events represented as rectangles (colored rectangles represent predictable targets while grey rectangles represent unpredictable targets). b) mean hit rates for predictable and unpredictable targets – individual participants are represented as light grey lines c) mean hit rates across the trial. d) mean RT for predictable and unpredictable targets. e) mean RT across the trial. Error bars in bar graphs represent the standard error of the mean in this and all subsequent figures.

Eye Movements

In Experiment 1, eye-tracking data were also recorded to investigate whether spatiotemporal predictions influenced the pattern of eye movements. I compared the probability of looking at the quadrant in which predictable vs. unpredictable targets were presented at each time point locked to the onset of a target (500 ms before onset until 4.5 seconds after). Note, target onset was defined as the moment in which a target was no longer entirely transparent, this does not mean that the targets were necessarily visible at time 0. The results, illustrated in Figure 5.2.a, revealed significant differences (ps<.05) between the probabilities of fixations early in the trial beginning at approximately 1100 ms and lasting for ~1000 ms. The difference was driven by the higher probability to fixate at the target quadrant when searching for a predictable target. The probabilities of fixations also differed in a later time window after ~4250 ms showing reduced probability of fixating for predictable vs unpredictable targets. In Figure 5.2.b it is clear that these fixation probability patterns were present throughout the trial.

In Figure 5.2.c decision time is plotted by predictability. Decision time is defined as the time between the first fixation on the target quadrant during the last dwell before the response and the time of the response. Observers took significantly less time to decide a predictable target was indeed a target compared to unpredictable targets. ($\beta = -15.14$, $SE = 3.16$, $t = -4.78$, $p < .001$). There was no effect of target order ($\beta = 1.15$, $SE = 4.08$, $t = .28$, $p = .78$) and these factors did not interact ($\beta = -3.06$, $SE = 2.82$, $t = -1.09$, $p = .28$).

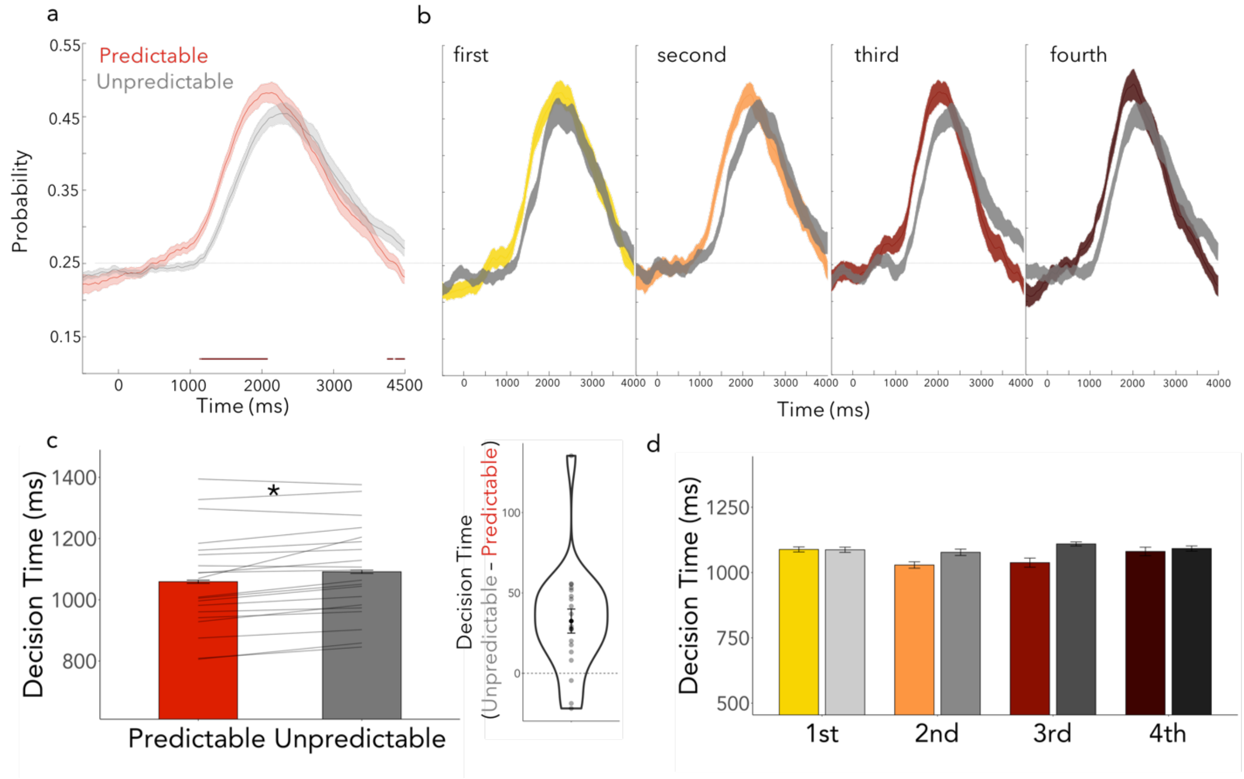

Figure 5.2: Gaze is biased towards predictable targets earlier. a) The probability of a fixation landing in the 'target' quadrant for a 5-second epoch around the 'onset' of the target (time 0). Given there are four quadrants, the chance of being in any one quadrant would be .25. The onset indicates the moment that the targets were no longer completely transparent, although they were not necessarily visible at this moment. The probability of fixating a target quadrant was significantly higher for predictable targets early in the epoch and significantly lower during the later stage of the epoch. Significant time windows are marked with a solid red line. (b) This pattern repeats throughout the different phases of the trial. (c) Decision time – as measured by the first fixation time during the last dwell within a target quadrant before the response subtracted from the response time – was significantly reduced for predictable compared to unpredictable targets. Individual subject differences between unpredictable and predictable targets are plotted in the violin plot, in which the black dot represent the mean difference and the error bars represent the standard error (d) This pattern did not significantly differ with target order.

5.4. Experiment 2: Behavioral Benefits Remain with Asynchronous Regularities

Experiment 1 showed that observers were more accurate and faster at finding spatiotemporally predictable compared to unpredictable targets, and this is true throughout the trial. Experiment 2 was aimed to test whether benefits of spatiotemporal predictions were generalizable to contexts in which the predictable targets were not separated by a constant set of intervals, in a strict temporal rhythm. To this end, the timing of predictable targets was varied across participants, such that these onsets were consistent for any given participant but not necessarily evenly distributed across the course of a trial. If the effects are still present, it is clear that these results cannot be solely attributed to a rhythmic pattern of attentional allocation.

5.4.1. Methods

Participants

Twenty-five participants took part in Experiment 2. One participant was discarded for low performance – with an average target hit rate more than 2 standard deviations away from the mean. The remaining 24 participants were between 18-30 years old with an average age of 22.83. The sample contained 20 females and 4 participants were left handed. All Participants had normal or corrected-to-normal vision, provided written consent, and were compensated at a rate of £10 per hour.

Task and Procedure

The Experimental task outline is depicted in Figure 5.3.a. The task parameters were nearly the same as in Experiment 1, except that each trial was divided into 6 equal time bins (~1600 ms each). Observers were randomly assigned four of the six bins to serve as predictable target times. Each of the 4 unpredictable targets was assigned to random onset within one of the time bins. The quadrants and specific locations of the unpredictable targets were chosen randomly.

Participants completed the experiment in a group testing room with a capacity of 20 people, although no more than 12 were tested at once. Participants each sat approximately 50 cm from the monitor. All other experimental procedures were the same as in Experiment 1.

Behavioral Analysis

The same analysis procedure as in Experiment 1 was repeated. In Experiment 2, the GLMM's random-effects structure contained the participant intercepts as well as by- participant slopes for predictability and target order (i.e. the full model). The full model was also optimal for the LMM in Experiment 2. Further details regarding the models can be found in the analysis script (https://osf.io/xhu5j/).

5.4.2. Results

Hits

As in Experiment 1, there was a main effect of predictability on the percentage of targets hit, with predictable targets being found significantly more often than unpredictable targets (Figure 5.3.b) (β = .20, *SE* = .02, *z* = 8.56, *p* < .001). Once again, there was a main effect of target order on hits with higher accuracy for early targets (Figure 5.3.c) (β = -.10, *SE* = .02, *z* = -4.49, *p* < .001), and the interaction was again not significant (β = -0.0009, *SE* = .01, *z* = -0.079, *p* = .94).

RTs

In a replication of Experiment 1, an analysis of RTs revealed a main effect of predictability (β = -.02, *SE* = .004, *t* = -4.72, *p* < .001) as well as target order (β = .08, *SE* = .005, *t* = 14.34, *p* < .001). More specifically, observers were faster for predictable targets (Figure 5.3.d) as well as targets appearing early in the trial (Figure 5.3.e). Again, the interaction was not significant (β = .004, *SE* = .003, *t* = 1.53, *p* = .13).

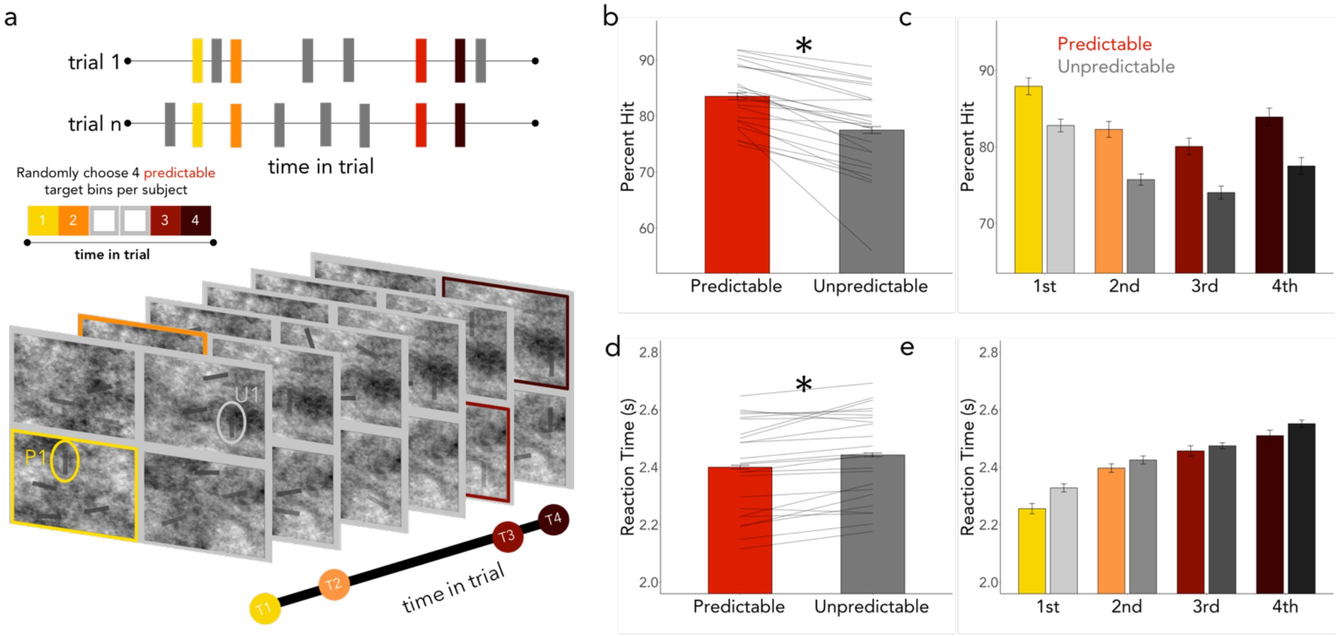

Figure 5.3: Increased and faster detection of predictable targets with asynchronous regularities. a) The trial schematic for Experiment 2. Each trial was divided into six time bins, and each observer was randomly assigned to four bins – a yellow (grey) circle indicates the first predictable (unpredictable) target. b) Mean percent targets hit is plotted for predictable and unpredictable targets – individual participants are represented as light grey lines. c) Mean percentage of targets hit is plotted across the trial split by predictability. Mean RT for predictable and unpredictable targets (d) averaged over the trial and (e) separately for target across the trial.

5.5. Experiment 3: Spatiotemporal Templates from Long-Term Memory

In Experiments 1 and 2, the results demonstrated that regularities in time and space are found in the attentional template and therefore help guide dynamic visual search and speed decisions regarding a target. In both of these experiments there were two major sources of regularities, either of which may have contributed to performance facilitation: short-term priming effects from one trial to the next and the build-up of longer-term memories related to regularities extracted over the course of the experiment. Although these sources of information are highly related, I wished to titrate, to the extent possible, the relative contributions from single-trial priming and longer-term learning. To do so, in Experiment 3, 40% of the trials contained no regularities (*random trials*). If performance benefits are fully reliant on single-trial priming, then these should disappear on trials immediately following a random trial.

5.5.1. Methods

Participants

Twenty-seven observers were tested. One participant had an average target hit rate more than 2 standard deviations lower than the mean, leaving a final sample size of 26 (age range 18-33, mean age = 24.5, 20 females). All participants had normal or corrected-to-normal vision, provided written consent, and were compensated at a rate of £10 per hour. The group of participants in this experiment also participated in Experiment 4 (see below) within a single session.

Task and Procedure

Experimental parameters were based on Experiment 1, except that predictable targets occurred in only a proportion of trials (Figure 5.4.a). Sixty percent of the trials were *standard trials* as described in Experiment 1. As in Experiment 1, four predictable targets and four unpredictable targets were presented throughout the trial. The four predictable targets were predictable in time – as they always appeared with the same onset in four equally spaced intervals – and space – as they were always presented in the same quadrant. The other 40% of the trials were considered *random trials* as these trials contained eight unpredictable targets. That is, in these trials, the four predictable targets appeared at an unpredictable time and quadrant. Specifically, at the start of each trial all items – distractors and targets – were uniformly distributed across time and randomly distributed across the quadrants. Subsequently eight of these items were chosen randomly to be targets. If short-term memory effects are important, then any advantage for predictable over unpredictable events should decline after a completely unpredictable trial.

Participants completed the experiment in a group testing room with a capacity of 20 people, although no more than 12 were tested at once. The group that participated in this experiment also completed a second experiment (Experiment 4; see below), and the order of task administration was counterbalanced. Participants each sat approximately 50 cm from the monitor.

Behavioral Analysis

In line with the experimental manipulation and question, the analysis model included parameters for predictability, previous trial type, and their interaction. By necessity, the first trial in each block, and the completely random trials were not included in the analysis. The random-effects structure for the GLMM contained the participants' intercepts as well as by-participant slopes for predictability and target order. This model was also optimal for the LMM. Significant interactions of predictability and trial type were broken down by defining difference contrasts to model the two critical comparisons (repeat vs. non-repeat trials for predictable and unpredictable targets). Analysis script with detailed information can be found here https://osf.io/xhu5j/.

5.5.2. Results

Hits

There was a significant effect of predictability on hits (Figure 5.4.b) ($\beta = .20$, SE = .02, $z = 8.25$, $p < .001$) and a significant effect of target order (Figure 5.4.c) ($\beta = -.16$, SE = .02, $z = -7.42$, $p < .001$). Observers were more accurate for predictable and early targets. These factors did not significantly interact ($\beta = .006$, SE = .01, $z = .417$, $p = .676$). Moreover, there was no effect of the previous trial type ($\beta = .02$, SE = .02, $z = 1.34$, $p =.18$), and the previous trial type did not interact significantly with predictability ($\beta = .02$, SE = .02, $z = 1.04$, $p = .30$), indicating that there was no significant diminution in the predictability effect immediately following a fully unpredictable trial.

RTs

The equivalent analysis for RTs again showed faster responses for predictable targets (Figure 5.4.d: β = -.015, SE = .006, t = -2.26, p = .03) and faster responses for early targets (Figure 5.4.e: β = .11, SE = .006, t = 15.54, p < .001), yet no interaction (β = .006, SE = .004, t = 1.63, p =.10). There was also no effect of the previous trial type when considering RTs (β = .004, SE = .004, t = 1.11, p =.27), and the previous trial type did not interact significantly with the predictability effect (β =-.007, SE = .004, t = 1.60, p =.11) suggesting that the predictability effect did not rely on the short-term priming effects of the previous trial and that longer more robust memories are at play here.

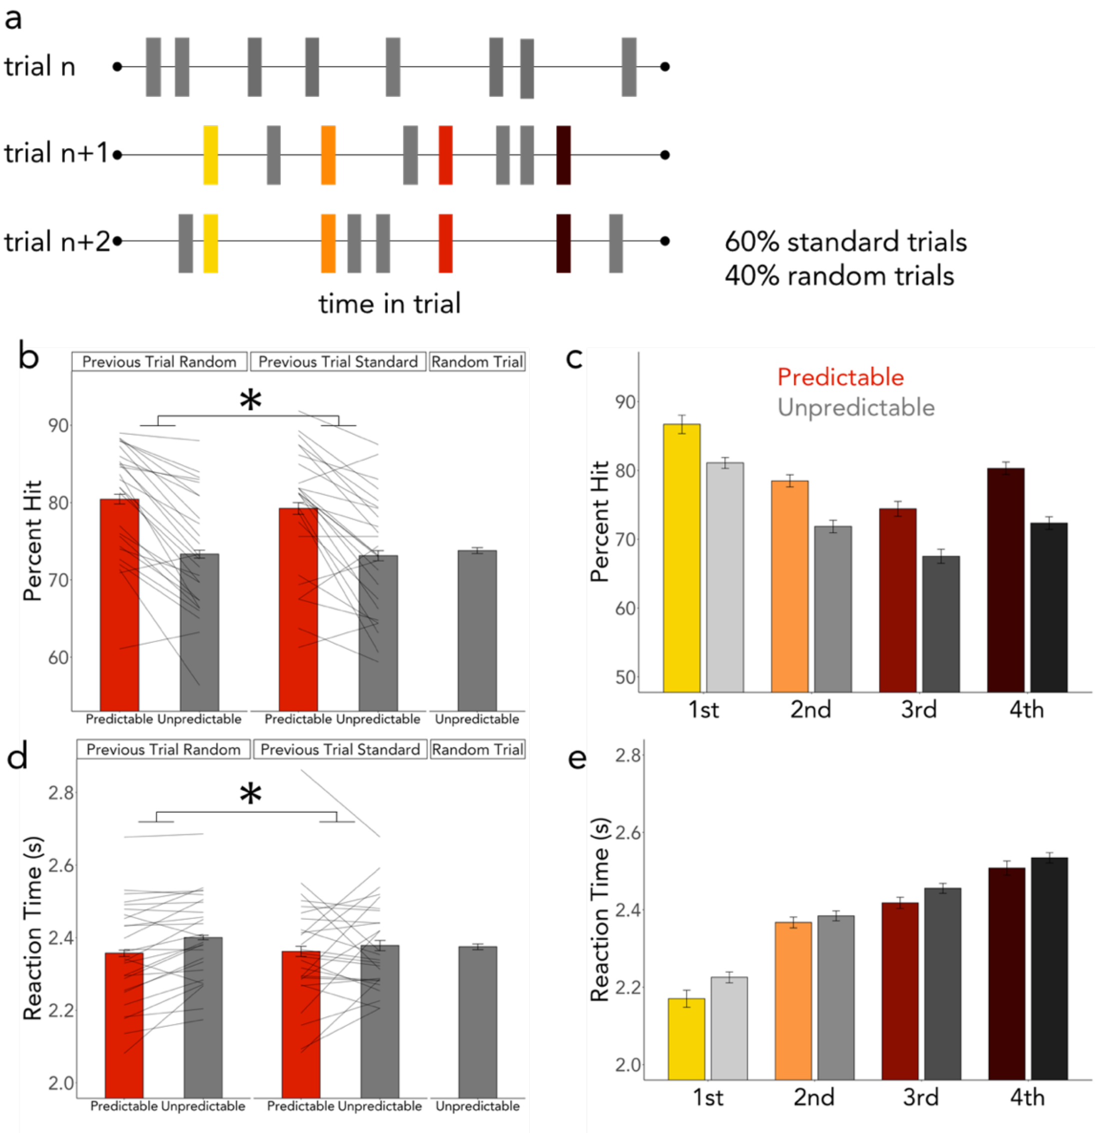

Figure 5.4: Spatiotemporal guidance is held in long-term memory. a) The trial schematic for Experiment 3. In 40% of trials, the locations and timings of all targets were completely unpredictable. b) Mean hit rates are plotted for predictable and unpredictable targets separated by the previous trial type – individual participants are represented as light grey lines. Participant means for the average hits in fully random trials are also shown. c) Percent hits in standard trials are plotted across the trial. d) Mean RTs are plotted for predictable and unpredictable targets when the target was preceded by a standard trial, a random trial, as well as for when the trial itself was fully random. e) The reaction times in standard trials across the entire trial.

5.6. Experiment 4: Trial-wise Priming Contributes to Behavioral Effect

Experiment 4 was designed to assess the relative contributions of short-term priming vs. longer-term memory effects in a way that was complementary to Experiment 3. Experiment 3 included trials with fully unpredictable targets to test whether the absence of priming of target predictions between successive trials disrupted behavioral benefits. There was no significant diminution of the predictability effect in trials preceded by a random vs. standard trial, implying the attentional template was resistant to interference of previous trials with a novel spatiotemporal pattern of targets. Benefits from spatiotemporal regularities can thus rely on longer-term associative memory traces. In Experiment 4, I wished to further investigate the interplay between short-term priming and longer-term memory effects by testing whether single trial priming was nevertheless sufficient to support benefits of spatiotemporal predictions. By introducing inter-trial priming of target predictions, it was possible to see whether short-term priming effects may *also* contribute to the attentional template and therefore behavioral guidance.

5.6.1. Methods

Participants

The same participants as Experiment 3 (see Exp3 *Participants* section for more detail) were tested in Experiment 4.

Task and Procedure

The task outline is depicted in Figure 5.5.a. In Experiment 4, as in Experiment 3, only 60% of the trials were standard 'non-repeat trials' (as described in Experiment 1 *Task* and Procedure). The remaining 40% of the trials were 'repeat trials', in which all of the timings and quadrants from the previous trial repeated. This is in contrast to Experiment 3, in which 40% of the trials were fully unpredictable, here 40% of the trials were fully predictable based on the previous trial. The four predictable targets remained predictable in all trial types (repeat and non-repeat). The timings and quadrants of the four unpredictable targets remained unpredictable in non-repeat trials but were presented with the same dynamics as in the previous trial in repeat trials. It is important to point out that although the repeated trials maintained the exact timing and quadrants from the previous trial, like the predictable targets, the exact location within the quadrant was random.

Participants were tested in the same room and during the same session as in Experiment 3. The order of the two tasks (3 & 4) was counterbalanced.

Behavioral Analysis

In Experiment 4, trial type and its interaction with predictability were included as predictors in the model. The random-effects structure for the GLMM contained the participants' intercepts as well as by-participant slopes for predictability and target order. This model was also optimal for the LMM. More details of this analysis can be found here https://osf.io/xhu5j/.

5.6.2. Results

Hit

There were significantly higher hits for predictable compared to unpredictable targets (Figure 5.5.b) (β = .20, *SE* = .02, *z* = 10.95, *p* < .001) and a significant effect of target order, with early targets being detected more frequently (Figure 5.5.c) (β = -.13, *SE* = .02, *z* = -8.27, *p* < .001). There was no interaction between predictability and target order (β = .006, *SE* = .01, *z* = .59, *p* = .55). Although there was no significant main effect of the trial repetition (β = -.004, *SE* = .01, *z* = -.33, *p* = .74), this factor interacted significantly with predictability (Figure 5.5.b: β = .03, *SE* = .01, *z* = 2.95, *p* = .003). Planned comparisons revealed that the “unpredictable” targets were found significantly more often in the repeat trials compared to the non-repeat trials ($\beta = .08$, *SE* = .03, $z = 2.44$, $p = .01$), whereas, hits did not differ for predictable targets in repeat and non-repeat trials ($\beta = -.06$, *SE* = .03, $z = -1.76$, $p = .08$). These results show that a single repetition of the spatiotemporal dynamics of a target was sufficient to find significantly more targets compared to trials that did not repeat, suggesting some role for inter-trial priming in establishing the predictability effect.

RTs

There was a main effect of predictability such that predictable targets were found faster than unpredictable targets (Figure 5.5.d: $\beta = -.03$, *SE* = .006, $t = -5.36$, $p < .001$). Additionally, early targets were found faster (Figure 5.5.e: $\beta = .11$, *SE* = .006, $t = 18.47$, $p < .001$). A significant interaction ($\beta = .02$, *SE* = .003, $t = 7.52$, $p < .001$) indicated a steeper slope of improvement in reaction times for predictable, compared to unpredictable targets – although, numerically, predictable targets were found faster than unpredictable targets throughout the trial. There was no effect of trial repetition on RTs (Figure 5.5.d: $\beta = .004$, *SE* = .003, $t = 1.17$, $p =.24$), and trial repetition did not interact significantly with the predictability effect ($\beta =-.0005$, *SE* = .003, $t = -.18$, $p =.86$).

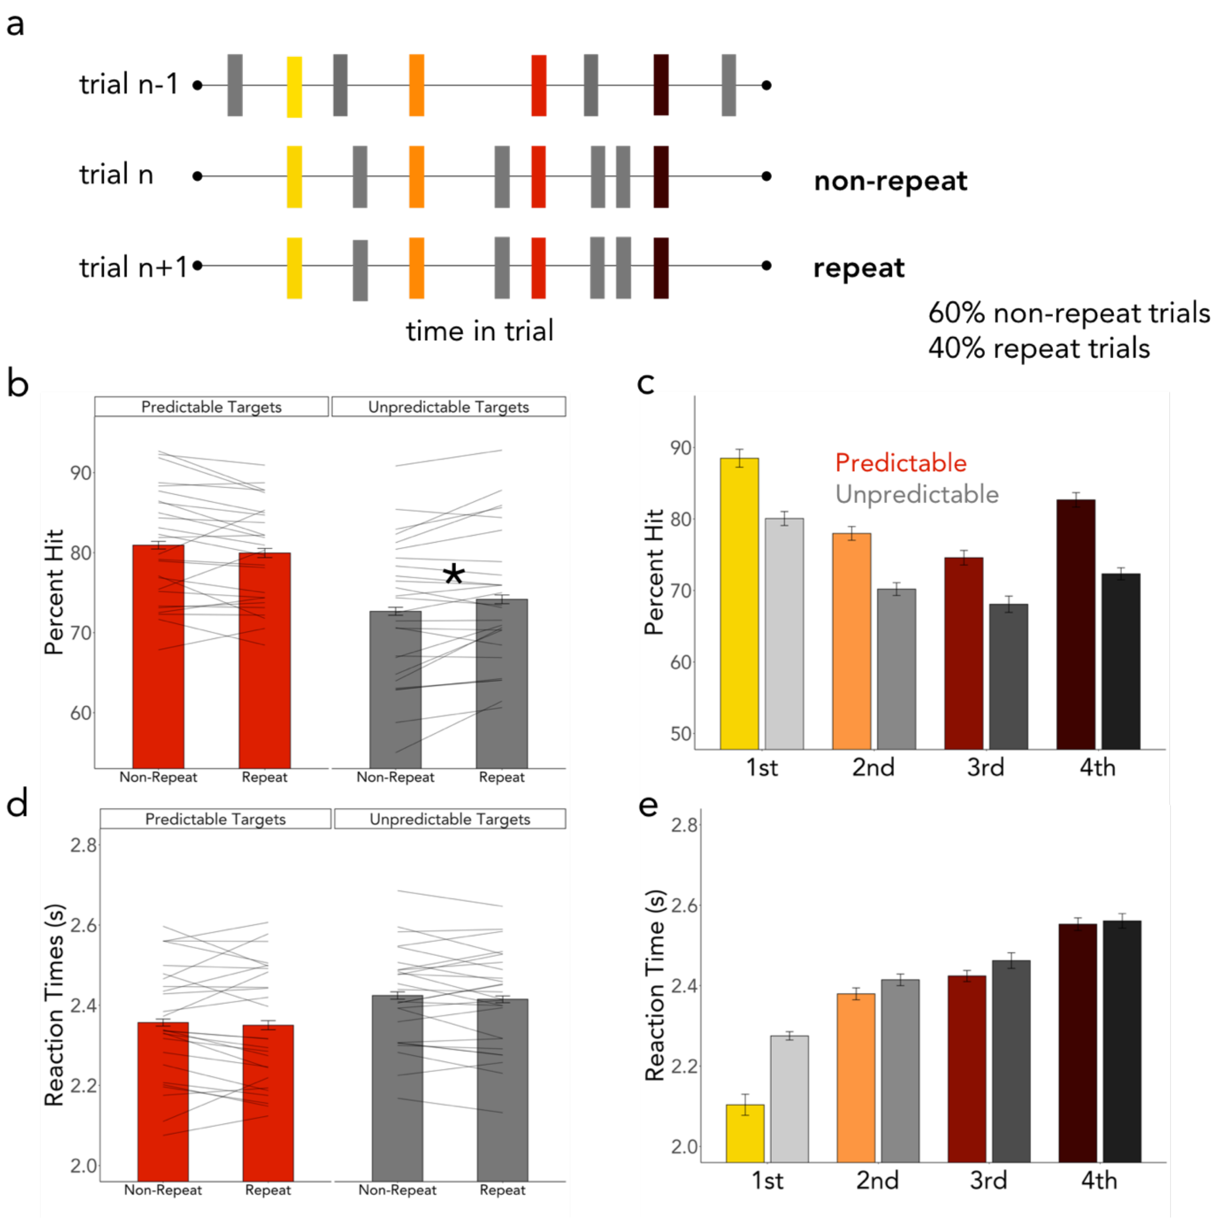

Figure 5.5: Trial-wise priming contributes to behavioral benefits. a) Depiction of a trial schematic for Experiment 4. 40% of trials were repeated such that the timings and quadrants of the unpredictable targets were the same as in the previous trial as is depicted for trial n+1. In these trials, what would usually be an unpredictable target was presented with the same onset and in the same quadrant as the previous trial. b) Mean hit rates are plotted for predictable and unpredictable targets in repeat and non-repeat trials – individual participants are represented as light grey lines. Participants hit significantly more unpredictable targets when their timing and quadrant repeated from the trial before. c) Mean hit rates is plotted over the course of the trial by predictability only for the Non-Repeat trials. d) Mean RT is plotted for predictable and unpredictable targets in different trial types. e) The effect of predictability on RT across the trial.

5.7. Discussion

In four experiments, participants searched for targets among distractors in a novel dynamic visual-search task. Experiment 1 established the basic effect showing participants were more accurate and faster at detecting targets when they were predictable in their temporal onset and quadrant location. By tracking eye movements in Experiment 1, it was possible to demonstrate the proactive nature of this effect. Experiment 2 showed that the effects were due to predictability of the specific temporal sequence of targets rather than just their temporal order or a rhythmic pattern. In Experiments 3 and 4, I tested for the contribution of two potential sources within the attentional template: long-term representation of trial history and short-term trial-wise priming. The results indicated a strong contribution of long-term memory. Behavioral benefits of predictability did not require immediate inter-trial priming and resisted interference from different intervening spatiotemporal patterns (Experiment 3). However, short-term priming also conferred significant, albeit modest, benefits to search performance, primarily improving the hit rate (Experiment 4).

This dynamic visual search framework was used to examine the consequences of spatiotemporal information in attentional templates within dynamic settings, thus providing a step toward an understanding of behavior in a real-world continuous context. Each trial spanned several seconds and required multiple responses, such that one may monitor the guidance of spatial selection over time. To ensure that behavior was not dominated by exogenous factors, targets appeared and disappeared slowly from the display, such that attention was not captured by the sudden onset of any event, but rather revealed the guidance by top-down predictive signals that changed dynamically over time. By rendering a standard visual-search task dynamic, it became clear that predictive temporal associations should be considered as an additional source within the attentional template guiding selection.

Attentional templates can refer both to representations that guide attention in space, as well as the information we use to determine if the object currently selected is indeed a target. During a visual search task these two processes are indistinguishable within behavioral measures such as reaction time and hits. That is, it is unclear if an improvement in reaction time can be attributed to better attentional guidance or faster recognition time. In Experiment 1, I could use eye-tracking metrics to separately interrogate these processes. Specifically, observers were more likely to fixate the target quadrant earlier for targets that were predictable compared to unpredictable. Moreover, the decision time measure showed that spatial temporal information not only contributes to spatial guidance of attention during target search but also at the stage of target identification. More specifically, a template containing spatial and temporal properties led to fast target identification.

These findings could be viewed as a temporal analogue to a long literature showing memory-based spatial guidance of attention in static displays (Chun & Jiang, 1998; Geng & Behrmann, 2002; Hutchinson & Turk-Browne, 2012; Jiang, 2018; J. J. Summerfield et al., 2006). Specifically, it has been previously shown that performance improves when a target appears repeatedly in the same location (Geng & Behrmann, 2002; Jiang, 2018) or when an entire context repeats (Chun & Jiang, 1998). In addition, the literature has also shown inter-trial priming effects whereby spatial attention is biased by the previous trial (Hillstrom, 2000; Kristjánsson & Campana, 2010; Maljkovic & Nakayama, 1994, 2000; Meeter & Olivers, 2006). In Experiment 4 there was also evidence for these trial-wise effects, in which participants were more likely to detect repeated, but otherwise unpredictable targets.

There has been less work on memory representations of temporal intervals (Olmos-Solis, van Loon, Los, & Olivers, 2017; Olson & Chun, 2001). Cravo and colleagues (2017) demonstrated that observers are able to utilize temporal intervals associated with specific contexts to guide attention. Other studies have shown how repeating sequences can inform spatial guidance (Heideman et al., 2018; Nobre & O'Reilly, 2004; O'Reilly, McCarthy, Capizzi, & Nobre, 2008; Zhao, Al-Aidroos, & Turk-Browne, 2013). Here I demonstrate for the first time that memory-based spatiotemporal predictions can drive behavior in an extended dynamic task when facing multiple competing signals (i.e. distractors and other targets). Importantly, the behavioral benefit of predictions was not a result of discrete cueing-events or sequences, but rather driven by an internal representation of time.

The various factors that guide spatial selection combine within the attentional template. In a dynamic world it is likely that these templates must also be dynamic and therefore evolve with changes in our environment, and particularly, regular patterns that give rise to anticipation. In the current work, I have shown that participants use task embedded regularities to dynamically prioritize locations. By interleaving random trials in Experiment 3 I could test whether templates would be overwritten by random structures. Crucially, the attentional templates were accessible at the start of each trial, including those preceded by an unpredictable trial. Accordingly, it appears that these templates are available from a more robust long-term memory store. These regularities build up over the course of many regular, repeating events. In Experiment 4 even a single repeating trial lead to increased hit rates. This provides evidence for the very early stages of a developing template which includes spatiotemporal information. Moreover, an informal debriefing reveled that participants were largely unaware of the spatiotemporal regularities. However future work should focus on the implicit nature of these templates.

This work has identified a strong link between spatial and temporal guidance of attention. Previous work established a similar link when attention was directed explicitly through informative cues (Doherty, Rao, Mesulam, & Nobre, 2005; Rohenkohl et al., 2014). In the current task, spatiotemporal predictions were driven by implicit task regularities, which in turn led to proactive allocation of spatial attention as demonstrated through eye-movements: observers fixated the relevant quadrant earlier when targets were predictable and ahead of their manual responses. While these results have made an important step in replicating this synergetic relationship between temporal and spatial attention, the specific contributions of each of these mechanisms remains unclear, and could be the focus of future work.

In the current work, I have introduced a new perspective for considering spatial attention in visual search by including time as an informative dimension within the attentional template. Through this manipulation, I have found that the spatial distribution of attention is allocated flexibly on the basis of temporal predictions. This task can be extended in several directions to explore whether additional sources of guidance also evolve with time. For instance, it may be interesting to consider whether various features can be prioritized dynamically, e.g., whether knowing that a colored target is likely to emerge at a predictable time, without knowing where, can also benefit performance. Additionally, this new experimental framework could be used to characterize the precision of temporal and spatial predictions within the attentional template, independently, by presenting targets across a range of moments in time or locations in space. Finally, an important feature of the task is the gradual onset of events, minimizing exogenous capture of attention. While in the current investigation this was necessary in order to reduce the influence of bottom-up capture, this could be directly manipulated in the future. Altogether, I present new insights regarding visual search in an ecological setting, and introduce a simple – yet powerful – experimental framework for furthering these investigations.

6. General Discussion

Our memories are not a faithful snapshot of where we have been, but rather an adaptive blueprint guiding where we are going. In this thesis, I have defined these blueprints as *attentional templates* and have demonstrated that they are multidimensional and adaptable. Specifically, I have shown that they not only reflect the expected identity (Chapters 2 and 3) of a goal-relevant item – as in their more traditional definitions (Bundesen, 1990; Carlisle et al., 2011; Duncan & Humphreys, 1989; Wolfe, 2020) – but also contain associated actions (Chapter 4) as well as spatial and temporal regularities (Chapter 5). Moreover, these templates are adaptable and action oriented, reflecting the information that is most useful when completing the task at hand (Chapter 3 and 4). Taken together, the current work emphasizes the need for a richer definition of attentional templates. By considering how seemingly disparate memories can work together to support task goals, this work advances our understanding of proactive memory-guided behavior. Here, I will review the findings of the previous empirical chapters as well as outline how this integrated framework of attentional templates may fit into the existing literature. Just like attentional templates are future-oriented, so will my discussion be. Taking the results of this thesis as my starting point, I will be looking into future avenues for the study of attentional templates.

In Chapter 2, I conducted two EEG experiments investigating the behavioral and neural correlates of identity information within associative memory-based attentional templates. When observers could anticipate the identity of an upcoming target, they were significantly more accurate and faster at detecting that target. This effect was more pronounced when the target was difficult to perceive, indicating a potential difference in template utilization depending on perceptual task demands. Moreover, I found that two neural markers generally associated with spatial and temporal anticipation also contribute to the anticipation of object identity. Namely, when observers could form an expectation for the identity of a subsequent target, I found attenuation within posterior alpha band oscillations as well as a frontal negativity characteristic of the contingent negative variation (CNV). This was true even when equating spatial and temporal anticipation (Experiment 1 and 2), as well as target probability and response demands (Experiment 2). These results provide clear evidence for the proactive anticipation of target identity based on frequent associations between successive stimuli and implicate two key anticipatory neural signatures in this process.

After establishing the utility of memory-based attentional templates in Chapter 2, I turned to the degree of *adaptability* of these templates in Chapter 3. Specifically, I asked whether templates retrieved from long-term memory associations are merely faithful reproductions of the encoded information, or whether they can be adapted to the demands of the current context. The results revealed evidence for non-veridical, adaptable attentional templates across two experiments. Specifically, in Experiment 3.1, the most diagnostic feature dimension was overemphasized within the template leading to more accurate and faster responses when the template was most distinct from the presented lure – i.e. in the common lure condition. In Experiment 3.2, observers adapted their templates to the expected context of an early delay interval, irrespective of the time they were actually probed. Taken together, the findings show that templates retrieved from memory can be biased and adapted according to what feature dimension is most informative for task performance. I speculate that template adaptation likely occurs during the retrieval of the template and is tuned towards early (i.e. the most immediate) task demands.

In Chapter 4, I began to examine the *multidimensional* nature of attentional templates. That is, are prospective actions also represented as part of the attentional template and, if so, when are such plans incorporated? The results showed that observers immediately *imprint* their eventual action plan into the attentional template several seconds before it is actually needed. This action encoding occurs even in the face of an intervening motor task and predicts eventual response times. Finally, the behavioral results indicate that the encoded action information may guide performance during the secondary task. These results show that prospective actions are incorporated into the attentional template immediately and with lasting behavioral consequences. This begins to highlight the multidimensional nature of attentional templates, specifically showing that actions are held alongside detailed visual information in the attentional template.

Finally, in Chapter 5, I examined whether *spatiotemporal* information accompanies identity-based information in the attentional template in a temporally extended dynamic visual search task. In four experiments I showed that in dynamic search displays, attention is guided not only by the visual features of the target, but also by associated spatiotemporal regularities. Targets with predictable timing and location were found more often and more quickly across all experiments. The results further showed that the effects did not depend on a specific spatiotemporal pattern (Chapter 5.2) and primarily reflected long-term learning over trials (Chapter 5.3 and 5.4), although, short-term single-trial priming effects also contributed (Chapter 5.4). As an extension from Chapter 4, this work confirms that attentional templates are multidimensional – also taking the spatial and temporal dimensions into account. Moreover, this work shows that information from different memory scales can contribute to the attentional template. The newly developed dynamic search task provides a versatile framework for studying the various dimensions and axes of the attentional template within extended, ecological contexts.

The empirical chapters of this thesis emphasize the need for an integrated framework for attentional templates that critically relies on several factors. Figure 6.1 provides a potential outline for this framework.

First, attentional templates are multidimensional. In previous work, attentional guidance across space, time, response, and identity have been largely considered in isolation (Coull & Nobre, 1998; Maunsell & Treue, 2006; Posner, 1980). Moreover, work on attentional templates has predominantly focused on how target identity (features or object information) is used to guide behavior (Carlisle et al., 2011; Duncan & Humphreys, 1989; Geng & Witkowski, 2019). In the current thesis I have emphasized the multidimensional nature of attentional templates (Chapters 4 and 5). In the following section I will discuss how these templates may be organized.

Second, attentional templates stand at the intersection of three important subdisciplines within cognitive psychology: working memory, long-term memory, and perception (visual search). Most commonly, attentional templates are considered only within the visual search literature, in which their representational qualities are emphasized. Moving forward, it is critical not only to ask what is being represented, but additionally to consider what is it used for? In doing so, we will find that templates bring memories of different timescales to the service of behavior in a myriad of everyday tasks, not confined to visual search. In this thesis I used four distinct experimental paradigms across the four empirical chapters, all of which relied on attentional templates. Below, I will further explore the merits of this integrated framing of attentional templates across a multitude of tasks and timescales.

Third, it is critical to understand attentional templates within the context in which they are used. Templates are adaptable and action-oriented. This means the template representation itself is optimized for efficient behavior. The content of the attentional template is context specific and may be modified on the basis of utility.

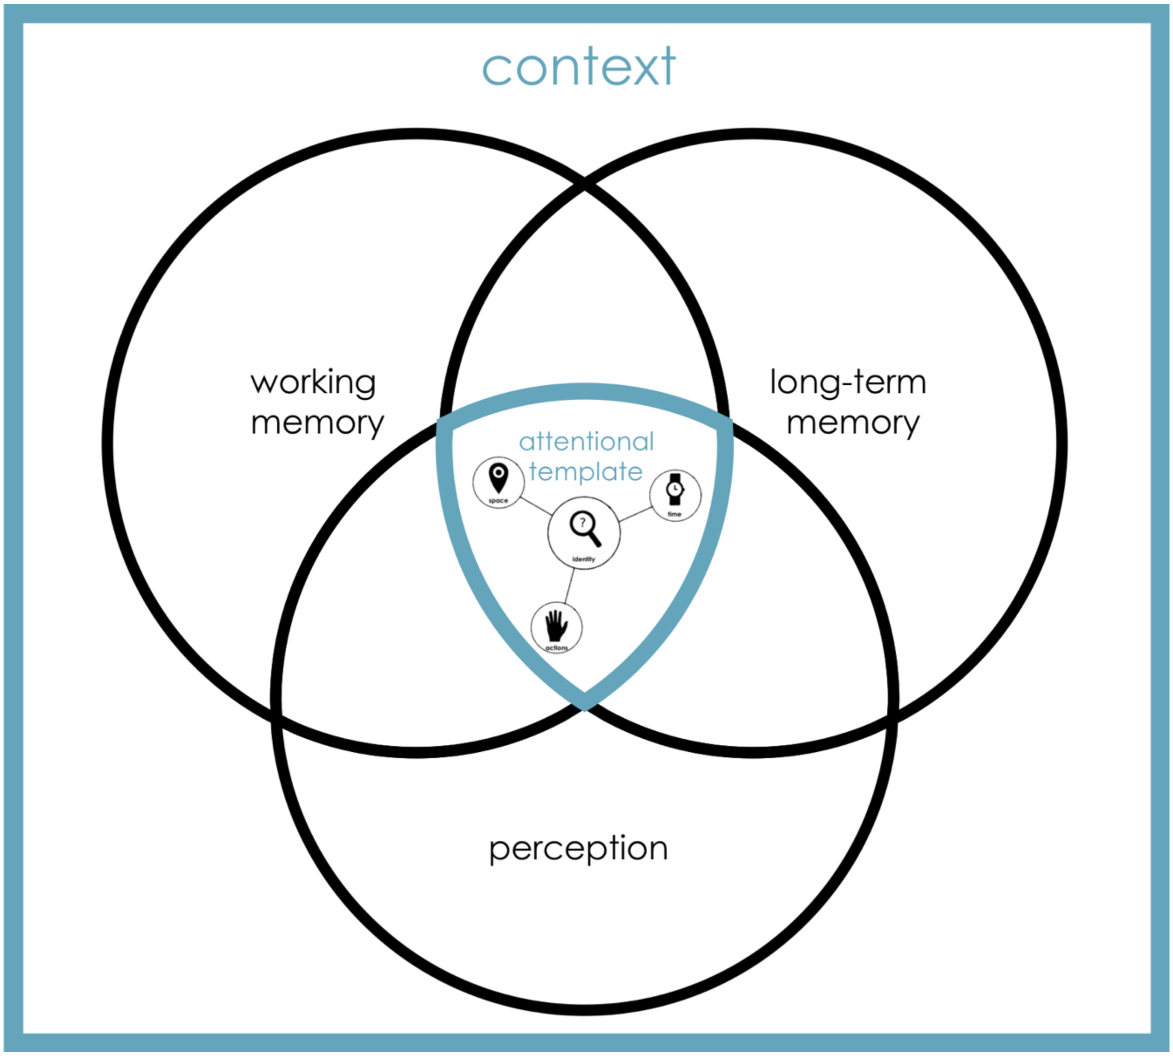

Figure 6.1: An integrated framework of attentional templates. Attentional templates are multidimensional and therefore, not only contain the expected identity of an upcoming event but also the associated spatiotemporal dynamics and actions (See Figure 1.1 for expanded view of multidimensional templates). Templates critically lie at the intersection of working memory, long-term memory, and perception. Finally, attentional templates should be considered within a wider context and this context can shape the template itself, as attentional templates are adaptable and use-oriented.

6.1. Identity as the foundation of multidimensional attentional templates

In the current thesis, I propose an integrated framework of attentional templates in which all of the relevant information we hold about a target is proactively utilized in order to guide behavior (see Figure 1.1 for an expanded view of multidimensional templates). Traditionally, attentional templates are thought of as the target representations that guide attention during a task such as visual search (Bundesen, 1990; Carlisle et al., 2011; Duncan & Humphreys, 1989; Wolfe, 2020). These mental representations are proposed to reflect the identity of a target – or target-defining features. I have shown that, in addition to the identity of an anticipated target event, attentional templates may also contain the actions that accompany that event (Chapter 4) as well as its associated spatiotemporal dynamics (Chapter 5) when these sources of information are available and useful.

Thus, the studies in the thesis significantly expand upon the notion of attentional templates by adding other dimensions of information to that relating to target identity. This brings forward an interesting question. Can attentional templates be supported solely by these other dimensions and exist in the absence of identity information? Or even more broadly, are we ever truly confronted with a single dimension of anticipation? In much of the seminal work on spatial and temporal attention, identity-based information was equated but not eliminated (Coull & Nobre, 1998; Posner, 1980). For example, Posner (1980) elegantly showed that when observers were cued to the probable location of a transient visual event they showed improved performance compared to trials in which the cue was uninformative or deceptive. However, even in this critical demonstration of spatial attention, these anticipatory spatial mechanisms were locked to target identity – that is, in all trials, observers knew what to expect; i.e. what defined a target. Even in tasks that rely on identification rather than detection, the features of a target are often limited to a certain subspace (e.g. left/rightward tilted gratings or alphanumeric symbols). As such, it remains an open question whether spatial and temporal attentional mechanisms may operate without some preparatory identity-based information.

It is also true that space and time are often implicitly included in laboratory research about identity predictions (Kok et al., 2014, 2017). In Chapter 2, I investigated the behavioral and neural correlates of associative memory-based identity templates. I did so by equating accessory sources of attentional guidance such as space and time. However, it remains unclear whether these effects are contingent on the foreknowledge of where and when a target would appear. As such, the observed effects may result from the interaction among several forms of anticipation: identity, spatial, and temporal. This was also true in the other chapters of this thesis. In Chapters 3 and 4, observers were always certain when and where a critical event would appear, while in Chapter 5, the identity of the target – a vertical line – remained constant. Such interactions among the various dimensions of an attentional template are ubiquitous outside the laboratory, and therefore should be considered in future research.

The integrated framework of attentional templates I propose bares some relation to the idea of object files (Ferreira, Apel, & Henderson, 2008; Hommel, 2004; Kahneman, Treisman, & Gibbs, 1992). Object files are integrated episodic representations that contain object-related information (Hommel, 2004). There is evidence that accessory information such as a task-irrelevant location is maintained when an object is retrieved or even attended to within memory (Richardson & Spivey, 2000; van Ede, Chekroud, & Nobre, 2019). Moreover, memory retrieval and perception are facilitated when the location of the critical information is either attended to or fixated (Ferreira et al., 2008; Hommel, 2004; Kahneman et al., 1992). This work suggested that an object file is used to help integrate successive states of an object. Here I propose that attentional templates extend beyond retrieval. Specifically, multidimensional mental representations can be used to proactively facilitate on going behavior. Moreover, it is interesting to consider how these representations may be more flexible than the concept of a 'file' may suggest – that is a self-contained representation that encompasses all the relevant information. For example, it may be sufficient to have distributed information associated with a particular event that can act together, in flexible arrangements, depending on the context and goals in a situation.

Finally, in this thesis I have shown that actions and spatiotemporal information can both be incorporated into the attentional template. However, this need not be an exhaustive list. When considering adaptive behavior in the real world, there are certainly other sources of information that may be included in the attentional template. For example, in the current work I focused primarily on the visual domain. However, multimodal sources of information may combine to provide a richer template. The simple task of making tea exemplifies the multidimensional nature of attentional templates. Our past experiences tell us

– at a minimum – when to expect a kettle to boil, where we should find it, what it will look like, how we will interact with it, what sound it will make, and how the tea will make us feel. In fact, Iordanescu and colleagues (2008) showed that congruent auditory cues can facilitate search for objects. Mapping out the various properties that contribute to attentional templates is a natural next step for research.

6.2. Attentional templates *for* and *from* different time scales and tasks

Attentional templates have been traditionally defined within the context of visual search. However, it is clear that they are constantly guiding our behavior in tasks other than visual search. In the current thesis, I have relied on several alternative tasks in order to interrogate attentional templates. In Chapter 2, I asked observers to identify a target that was quickly followed by a mask. Performance in this difficult perceptual task benefited from observers' ability to prepare an identity-based template. In Chapter 3, participants completed a relatively simple modified match-to-sample task. Despite the distinctive stimuli and overall high performance, I still found evidence for proactive anticipatory guidance from attentional templates that were adapted for the current task demands. In Chapter 4, attentional templates were explicitly provided in the form of a tilted bar and observers were subsequently asked to reproduce the orientation of that bar after a delay period. Overall performance in this task was relatively high, indicating that participants used detailed template representations to guide their responses. Finally, in Chapter 5, participants conducted an extended visual search in dynamic displays lasting several seconds. It became evident from these tasks that we are constantly using our memories to guide ongoing behavior. It is beneficial to consider the nature of representations that proactively guide behavior across different tasks in the same light. By expanding the definition of attentional templates beyond visual search, we gain valuable insights into proactive memory-guided behavior.

Attentional templates stand at the intersection of three important subdisciplines within cognitive psychology: working memory, long-term memory, and perception (Figure 6.1). In visual search (perception) their role is clear, as they guide attention and/or serve as the template in which candidate targets from the environment are compared (Wolfe, 2020). In Chapter 5, I showed that spatiotemporal regularities facilitate both of these processes. In working and long-term memory research, attentional templates are not often considered. This is because memory research has traditionally considered how well information from the past is represented, without considering why this information is being held, i.e. the *purpose* of memory. More succinctly, although memory is about the past, it is for the future (Nobre & Stokes, 2019).

Whether it is a representation of your keys, or the basic understanding of how kitchens are arranged, many of our long-term memories contribute to attentional templates that are used to guide behavior proactively (Võ et al., 2019; von Helmholtz, 1866). In Chapters 2 and 3, observers were asked to retrieve attentional templates from long-term memory associations. Critically, these templates represented information that was learned in the past but became necessary to guide ongoing behavior. I uncovered important insights regarding attentional templates by utilizing a task that relied on long-term memory representations. In Chapter 3, I found initial evidence suggesting that templates retrieved from long-term memory are adapted for the current task demands. Experiment 3.2 showed that this adaptation was likely to occur upon initial retrieval rather than within working memory. This is a critical insight into our understanding of adaptable attentional templates and differs from what one may have predicted on the basis of our understanding of the flexible nature of working memory across time (van Ede, Niklaus, et al., 2017; Zokaei et al., 2019). This may suggest that retrieving information from long-term memory in an adapted state is less costly than adapting information within working memory, although further work would be required to understand this tradeoff. In any case, it seems clear that there is a benefit to considering attentional templates retrieved from different memory sources. Moreover, it is crucial for long-term–

memory-researchers not only to ask *what* information is represented or stored but also what purpose this stored information may serve in the future.

Working memory has long been considered the home of attentional templates. In fact, it is difficult to think of an example of a working memory representation that is not used to guide future behavior. Even so, working memory research, like long-term memory research, has often focused on what information is held rather than what that information is for (Nobre & Stokes, 2019; van Ede, 2020). By reconsidering working memory in the context of future behavior, that is, a template representation used to guide attention and actions, it is possible to expand upon our understanding of both attentional templates as well as working memory. For example, in Chapter 4 observers were asked to maintain tilt information over a several second delay period during which they also completed a secondary task. This resembles many real-world examples of working memory usage, as we are often tasked with maintaining some information while we continue with other, intervening natural behaviors. The results showed that observers immediately imprinted prospective actions into the attentional template and that these prospective actions modulated behavior in an unrelated secondary task. From this we learned that working memory representations may also consider future actions. These prospective actions could even act as a dual code supplementing the visual information within working memory. Although more work is necessary to understand all of the advantages of maintaining actions alongside visual representations, it is clear that this allows observers to act efficiently, ultimately leading to faster reaction times.

6.3. Adaptable and action-oriented: Functional attentional templates

I have provided evidence that attentional templates are adaptable and action-oriented (Chapters 3 and 4). Critically, this speaks to one of the core purposes of attentional templates: guiding adaptive behavior. That is, templates should be functional and use-oriented in order to facilitate efficient behavior. If an attentional template is not optimized for the context in which it will be used, it will not guide behavior as efficiently. Previous work has demonstrated that within visual search, target representations are sensitive to expected distractor features (Bauer et al., 1996; Becker et al., 2010; Geng et al., 2017; Geng & Witkowski, 2019; Hodsoll & Humphreys, 2001b; Navalpakkam & Itti, 2007; Yu & Geng, 2019). Namely, templates are represented such that the template-to-distractor distinctiveness is optimized (Geng & Witkowski, 2019). I have expanded upon this work to show this is also true when templates are retrieved from long-term memory associations (Chapter 3). Moreover, I utilized a task other than visual search, showing that these adaptable templates are not unique to visual search. Taken together, showing that templates are adaptable is an important step to showing that they are functional and therefore can be optimized for current task demands.

Thus far, work on the adaptable nature of attentional templates has focused on target features such as color and orientation. However, as demonstrated in Chapters 4 and 5, templates are multidimensional, and the information used to guide our behavior expands beyond identity. This leaves an open question: are dimensions other than identity adaptable within the attentional template? For example, if a target is defined by its spatial location or temporal duration and distractors are linearly separable within these dimensions from the target, how will space or time be represented? What if a target is defined as an object that always appears for 2 seconds and distractors always appear for more than 2 seconds? Might an adaptable attentional template distort the temporal properties to represent the target as more distinguishable from the distractors – i.e. hold the template in its most "active" state at an earlier time point. This is an important avenue of research to expand upon our understanding of multidimensional adaptable templates.

In Chapter 2, I found a significant interaction between target difficulty and predictability. Specifically, performance in response to targets that were more difficult to perceive benefited more from the template in both accuracy and reaction times. Although these results are in line with the idea that attentional templates adapt to the context in which they will be needed – in this case difficulty, the identified neural correlates – alpha attenuation and the CNV – did not show a similar interaction. It is not possible to rule out differences in these conditions on the basis of this null result. However, it is also interesting to consider how similar levels of template activation may lead to difference in performance. For example, it is possible that the same attentional template is rendered more effective when the incoming stimulus is more difficult to perceive. In this scenario, the consequences of pre-activation of relevant neuronal populations may critically depend on the strength of neuronal activity triggered by incoming stimulation, playing a greater faciliatory role when incoming stimulation is weaker or more ambiguous.

Finally, it is interesting to consider how the same information may be represented differently across different tasks. Attentional templates in this case may rely on selective representations. That is, only the information needed for the anticipated behavior is held within the attentional template. In Chapter 3, I found that feature dimensions were not equally represented within attentional templates retrieved from long-term memory associations. One possible explanation for this, is that task-relevant feature dimension was disproportionately selected and retrieved from long-term memory. In Chapter 4, the prospective action was a critical element of the task. However, if the same stimuli (rotated bars) were presented without the need to respond, it is unlikely that this would activate motor regions. This is because that information is not needed for a passive viewing task, again indicating that attentional templates represent what is *functional* rather than what *is*.

6.4. Neural mechanisms of attentional templates

In this thesis I have also revealed new insights into some of the accompanying neural mechanisms of associative memory-based identity templates (Chapter 2) and associated actions within the attentional template (Chapter 4). In Chapter 2 I found that identity-based attentional templates retrieved from long-term memory were associated with attenuation within the alpha band and a frontal negativity just prior to the onset of an expected event. These neural correlates align with a large literature on spatial and temporal anticipation (Cravo et al., 2011; Haegens et al., 2011; Heideman et al., 2018; Los & Heslenfeld, 2005; Miniussi et al., 1999; Nobre, 2001; Pfeuty et al., 2005; Praamstra et al., 2006; Rohenkohl & Nobre, 2011; Stokes et al., 2012; Thut et al., 2006; van Ede, 2018; van Ede, Niklaus, et al., 2017; Worden et al., 2000; Zanto et al., 2011) and show that varying forms of expectation may have similar neural underpinnings. In Chapter 4, I find that action plans – as indexed by lateralized alpha and beta activity over motor electrodes – are immediately imprinted into the attentional template. This imprinting occurs in the face of an intervening motor task and predicts eventual response times. This work provides insights into when actions are included into the attentional templates and suggests the possibility of a multidimensional code for attentional templates.

Going beyond these demonstrations and considering the diversity and multidimensional nature of attentional templates, it is interesting to speculate on how multidimensional adaptive templates may be implemented within the neural architecture. Firstly, it seems important to point out that although we often speak of the attentional template as a "mental representation", there need not be a combined multidimensional representation based on a single location or system within the brain. In fact, sensory information from a single event is notably coded in a distributed fashion such that, for example, motion, color, and location are represented across the visual cortex (Albright, Desimone, & Gross, 1984; Hubel & Wiesel, 1974; Michael, 1981; Treisman, 1996). Persistent activation within sensory neurons has been proposed as a possible mechanism for preparatory attention. That is, a neuron that codes for a particular object (Chelazzi et al., 1993) or location (Luck et al., 1997) is active even before an expected item is presented. These findings have supported the idea that a simple shift in the baseline of excitability of a neuron coding for an expected item would support efficient processing of that item. Alternatively, it has been shown that a template can be generated through more complex neuronal activity (Machens, Romo,

& Brody, 2010) or through short-term synaptic plasticity (Myers et al., 2015; Nobre & Stokes, 2019; Sugase-Miyamoto, Liu, Wiener, Optican, & Richmond, 2008). In any case, if seemingly separate sources of anticipatory information across distributed brain systems converge temporally, this will result, in effect, in a multidimensional template for a single event.

Navalpakkam and Itti (2007) proposed the idea that features within the attentional template are optimally tuned by increasing the excitability of sensory neurons that are most distinct from distractors. In Chapter 3, that would mean the expected orientation (color) of a target would be selectively enhanced depending on the block. Further, evidence from Chapter 3 suggests that this selective tuning may occur at retrieval from long-term memory. This provides a possible framework for adaptability within the attentional template, but more work is needed to fully understand this process.

Finally, the neuronal mechanisms of attentional templates should be output oriented. That is, a representation should be easily accessible and tuned for the current task demands. These templates need not represent the physical state of what was, but rather what is likely to be. The adaptable and accessible nature of these representations should be present across all dimensions of a template, and the distributed nature of these dimensions does not mean they cannot work in unison to guide behavior.

6.5. Final Conclusions

Attentional templates have held an important place in our understanding of natural behavior. However, the traditional definition of attentional templates ignores some of the potential scope of these representations. In the current thesis I have expanded upon this definition to propose an integrated framework of attentional templates. I have demonstrated that the identity information traditionally thought to comprise the template may be supplemented along several dimensions including space, time, and associated actions. Such templates often incorporate memories across multiple time scales, regularly extending beyond working memory to utilize associations and regularities stored in long-term memory. Moreover, I have shown that attentional templates are adaptable and action-oriented such that they do not necessarily represent the veridical properties of the external world but are rather functionally adapted to better serve future behavior. I have demonstrated these critical implications over the course of four empirical chapters and nine experiments. Taken together, the evidence presented in this thesis shows the tremendous power of attentional templates for efficient adaptive behavior, and prompts us to study such templates in a rich manner that integrates multiple memory systems, considers both time and space, and takes the prospective purpose and action-oriented nature of attentional templates as central considerations in future work.

7. References

Albright, T. D., Desimone, R., & Gross, C. G. (1984). Columnar organization of directionally selective cells in visual area MT of the Macaque. *Journal of Neurophysiology*. https://doi.org/10.1152/jn.1984.51.1.16

Atkinson, R. C., & Shiffrin, R. M. (1968). Human Memory: A Proposed System and its Control Processes. *Psychology of Learning and Motivation - Advances in Research and Theory*, *2*(C), 89–195. https://doi.org/10.1016/S0079-7421(08)60422-3

Awh, E., Belopolsky, A. V, & Theeuwes, J. (2012). Top-down versus bottom-up attentional control: a failed theoretical dichotomy. *Trends in Cognitive Sciences*, *16*(8), 437–443. https://doi.org/10.1016/j.tics.2012.06.010

Baayen, R. H., Davidson, D. J., & Bates, D. M. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, *59*(4), 390–412. https://doi.org/10.1016/j.jml.2007.12.005

Baddeley, A. D., & Hitch, G. (1974). Working memory. *Psychology of Learning and Motivation - Advances in Research and Theory*. https://doi.org/10.1016/S0079-7421(08)60452-1

Baldauf, D., & Desimone, R. (2014). Neural mechanisms of object-based attention. *Science*, *344*(6182), 424–427. https://doi.org/10.1126/science.1247003

Bar, M. (2004). Visual objects in context. *Nature Reviews Neuroscience*, *5*(8), 617–629. https://doi.org/10.1038/nrn1476

Bar, M., & Ullman, S. (1996). Spatial Context in Recognition. *Perception*, *25*(3), 343–352. https://doi.org/10.1068/p250343

Barne, L. C., Lange, F. P. de, & Cravo, A. M. (2020). Prestimulus alpha power is related to the strength of stimulus representation. *BioRxiv*, 2020.02.04.933382. https://doi.org/10.1101/2020.02.04.933382

Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, *68*(3), 255–278. https://doi.org/10.1016/j.jml.2012.11.001

Barrett, J., & Ehrlichman, H. (1982). Bilateral hemispheric alpha activity during visual imagery. *Neuropsychologia*, *20*(6), 703–708. https://doi.org/10.1016/0028-3932(82)90071-9

Bates, D., Kliegl, R., Vasishth, S., & Baayen, H. (2015). Parsimonious Mixed Models. Retrieved from http://arxiv.org/abs/1506.04967

Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software*, *67*(1), 1–48. https://doi.org/doi:http://dx.doi.org/10.18637/jss.v067.i01

Bauer, B., Jolicoeur, P., & Cowan, W. B. (1996). Visual search for colour targets that are or are not linearly separable from distractors. *Vision Research*, *36*(10), 1439–1466. https://doi.org/10.1016/0042-6989(95)00207-3

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), 7–7. https://doi.org/10.1167/9.10.7

Beck, V. M., Hollingworth, A., & Luck, S. J. (2012). Simultaneous Control of Attention by Multiple Working Memory Representations. *Psychological Science*. https://doi.org/10.1177/0956797612439068

Becker, S. I., Folk, C. L., & Remington, R. W. (2010). The Role of Relational Information in Contingent Capture. *Journal of Experimental Psychology: Human Perception and Performance*, *36*(6), 1460–1476. https://doi.org/10.1037/a0020370

Benwell, C. S. Y., Tagliabue, C. F., Veniero, D., Cecere, R., Savazzi, S., & Thut, G. (2017). Prestimulus EEG power predicts conscious awareness but not objective visual performance. *ENeuro*, *4*(6). https://doi.org/10.1523/ENEURO.0182-17.2017

Berger, H. (1929). Über das Elektrenkephalogramm des Menschen. *Archiv Für Psychiatrie Und Nervenkrankheiten*. https://doi.org/10.1007/BF01797193

Biederman, I., Mezzanotte, R. J., & Rabinowitz, J. C. (1982). Scene perception: Detecting and judging objects undergoing relational violations. *Cognitive Psychology*, *14*(2), 143–177. https://doi.org/10.1016/0010-0285(82)90007-X

Blair, R. C., & Karniski, W. (1993). An alternative method for significance testing of waveform difference potentials. *Psychophysiology*, *30*(5), 518–524.

Boettcher, S. E. P., Draschkow, D., Dienhart, E., & Võ, M. L.-H. (2018). Anchoring visual search in scenes: Assessing the role of anchor objects on eye movements during visual search. *Journal of Vision*, *18*(13), 11. https://doi.org/10.1167/18.13.11

Boettcher, S. E. P., Drew, T., & Wolfe, J. M. (2018). Lost in the supermarket: Quantifying the cost of partitioning memory sets in hybrid search. *Memory and Cognition*, *46*, 43–57. https://doi.org/10.3758/s13421-017-0744-x

Boettcher, S. E. P., & Wolfe, J. M. (2015). Searching for the right word: Hybrid visual and memory search for words. *Attention, Perception & Psychophysics*, *77*(4), 1132–1142. https://doi.org/10.3758/s13414-015-0858-9

Bond, A. B. (1983). Visual search and selection of natural stimuli in the pigeon: The attention threshold hypothesis. *Journal of Experimental Psychology: Animal Behavior Processes*, *9*(3), 292–306. https://doi.org/10.1037/0097-7403.9.3.292

Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, *10*(4), 433–436. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/9176952

Brignani, D., Lepsien, J., Rushworth, M. F. S., & Nobre, A. C. (2009). The timing of neural activity during shifts of spatial attention. *Journal of Cognitive Neuroscience*, *21*(12), 2369–2383. https://doi.org/10.1162/jocn.2008.21176

Brincat, S. L., & Miller, E. K. (2015). Frequency-specific hippocampal-prefrontal interactions during associative learning. *Nature Neuroscience*, *18*(4), 576–581. https://doi.org/10.1038/nn.3954

Broadbent, D. E. (1958). *Perception and communication*. London: Pergamon Press.

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

Bundesen, C., Habekost, T., & Kyllingsbæk, S. (2005). A neural theory of visual attention: Bridging cognition and neurophysiology. *Psychological Review*, *112*(2), 291–328. https://doi.org/10.1037/0033-295X.112.2.291

Carlisle, N. B., Arita, J. T., Pardo, D., & Woodman, G. F. (2011). Attentional templates in visual working memory. *Journal of Neuroscience*, *31*(25), 9315–9322. https://doi.org/10.1523/JNEUROSCI.1097-11.2011

Chalk, M., Seitz, A. R., & Seriès, P. (2010). Rapidly learned stimulus expectations alter perception of motion. *Journal of Vision*, *10*(8). https://doi.org/10.1167/10.8.2

Chatham, C. H., & Badre, D. (2015). Multiple gates on working memory. *Current Opinion in Behavioral Sciences*, *1*, 23–31. https://doi.org/10.1016/J.COBEHA.2014.08.001

Chelazzi, L., Duncan, J., Miller, E. K., & Desimone, R. (1998). Responses of Neurons in Inferior Temporal Cortex During Memory-Guided Visual Search. *Journal of Neurophysiology*, *80*(6), 2918–2940. https://doi.org/10.1152/jn.1998.80.6.2918

Chelazzi, L., Miller, E. K., Duncan, J., & Desimone, R. (1993). A neural basis for visual search in inferior temporal cortex. *Nature*, *363*(6427), 345–347. https://doi.org/10.1038/363345a0

Chun, M. M., & Jiang, Y. (1998). Contextual cueing: implicit learning and memory of visual context guides spatial attention. *Cognitive Psychology*, *36*(1), 28–71. https://doi.org/10.1006/cogp.1998.0681

Cisek, P. (2007). Cortical mechanisms of action selection: The affordance competition hypothesis. *Philosophical Transactions of the Royal Society B: Biological Sciences*. https://doi.org/10.1098/rstb.2007.2054

Cisek, P., & Kalaska, J. F. (2005). Neural correlates of reaching decisions in dorsal premotor cortex:

Specification of multiple direction choices and final selection of action. *Neuron*. https://doi.org/10.1016/j.neuron.2005.01.027

Coull, J. T. (2009, August). Neural substrates of mounting temporal expectation. *PLoS Biology*. https://doi.org/10.1371/journal.pbio.1000166

Coull, J. T., & Nobre, A. C. (1998). Where and when to pay attention: the neural systems for directing attention to spatial locations and to time intervals as revealed by both PET and fMRI. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *18*(18), 7426–7435. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/9736662

Cowan, N. (1995). *Attention and memory: An integrated framework*. *Oxford Psychology Series, New York*. Oxford University Press New York, NY.

Cox, D. D., & Savoy, R. L. (2003). Functional magnetic resonance imaging (fMRI) “brain reading”: Detecting and classifying distributed patterns of fMRI activity in human visual cortex. *NeuroImage*, *19*(2), 261–270. https://doi.org/10.1016/S1053-8119(03)00049-1

Cravo, A. M., Rohenkohl, G., Santos, K. M., & Nobre, A. C. (2017). Temporal Anticipation Based on Memory. *Journal of Cognitive Neuroscience*, *29*(12), 2081–2089. https://doi.org/10.1162/jocn_a_01172

Cravo, A. M., Rohenkohl, G., Wyart, V., & Nobre, A. C. (2011). Endogenous modulation of low frequency oscillations by temporal expectations. *Journal of Neurophysiology*, *106*(6), 2964–2972. https://doi.org/10.1152/jn.00157.2011

Cunningham, C. A., & Egeth, H. E. (2016). Taming the White Bear: Initial Costs and Eventual Benefits of Distractor Inhibition. *Psychological Science*, *27*(4), 476–485. https://doi.org/10.1177/0956797615626564

Cunningham, C. A., & Wolfe, J. M. (2014). The role of object categories in hybrid visual and memory search. *Journal of Experimental Psychology. General*, *143*(4), 1585–1599. https://doi.org/10.1037/a0036313

Curtis, C. E., Rao, V. Y., & D’Esposito, M. (2004). Maintenance of Spatial and Motor Codes during Oculomotor Delayed Response Tasks. *Journal of Neuroscience*. https://doi.org/10.1523/JNEUROSCI.5640-03.2004

D’Esposito, M., & Postle, B. R. (2015). The Cognitive Neuroscience of Working Memory. *Annual Review of Psychology*, *66*(1), 115–142. https://doi.org/10.1146/annurev-psych-010814-015031

Dalvit, S., & Eimer, M. (2011). Memory-driven attentional capture is modulated by temporal task demands. *Visual Cognition*, *19*(2), 145–153. https://doi.org/10.1080/13506285.2010.543441

de Lange, F. P., Heilbron, M., & Kok, P. (2018, September 1). How Do Expectations Shape Perception? *Trends in Cognitive Sciences*. Elsevier Ltd. https://doi.org/10.1016/j.tics.2018.06.002

De Vries, I. E. J., Van Driel, J., Karacaoglu, M., & Olivers, C. N. L. (2018). Priority Switches in Visual Working Memory are Supported by Frontal Delta and Posterior Alpha Interactions. *Cerebral Cortex*. https://doi.org/10.1093/cercor/bhy223

de Vries, I. E. J., van Driel, J., & Olivers, C. N. L. (2017). Posterior α EEG dynamics dissociate current from future goals in working memory-guided visual search. *Journal of Neuroscience*, *37*(6), 1591–1603. https://doi.org/10.1523/JNEUROSCI.2945-16.2016

Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, *18*(1), 193–222. https://doi.org/10.1146/annurev.ne.18.030195.001205

Deutsch, J. A., & Deutsch, D. (1963). Attention: Some Theoretical Considerations. *Psychological Review*, *70*(1), 80–90. https://doi.org/10.1037/h0039515

Doherty, J. R., Rao, A., Mesulam, M. M., & Nobre, A. C. (2005). Synergistic effect of combined temporal and spatial expectations on visual attention. *Journal of Neuroscience*. https://doi.org/10.1523/JNEUROSCI.1821-05.2005

Donaldson, D. I., & Rugg, M. D. (1999). Event-related potential studies of associative recognition and recall: electrophysiological evidence for context dependent retrieval processes. *Brain Research. Cognitive Brain Research*, *8*(1), 1–16. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/10216269

Donchin, E. (1981). Surprise!? Surprise? *Psychophysiology*, *18*(5), 493–513. https://doi.org/10.1111/j.1469-

8986.1981.tb01815.x

Donchin, E., Tueting, P., Ritter, W., Kutas, M., & Heffley, E. (1975). On the independence of the CNV and the P300 components of the human averaged evoked potential. *Electroencephalography and Clinical Neurophysiology*, *38*(5), 449–461. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/50170

Downing-Doucet, F., & Guérard, K. (2014). A motor similarity effect in object memory. *Psychonomic Bulletin & Review*, *21*(4), 1033–1040. https://doi.org/10.3758/s13423-013-0570-5

Downing, P. E. (2000). Interactions between visual working memory and selective attention. *Psychological Science*, *11*(6), 467–473. https://doi.org/10.1111/1467-9280.00290

Draschkow, D., Heikel, E., Vo, M. L.-H., Fiebach, C. J., Sassenhagen, J., Võ, M. L.-H., … Sassenhagen, J. (2018). No evidence from MVPA for different processes underlying the N300 and N400 incongruity effects in object-scene processing. *Neuropsychologia*, *120*, 9–17. https://doi.org/10.1016/j.neuropsychologia.2018.09.016

Draschkow, D., & Võ, M. L.-H. (2016). Of "what" and "where" in a natural search task: Active object handling supports object location memory beyond the object's identity. *Attention, Perception, & Psychophysics*, *78*(6), 1574–1584. https://doi.org/10.3758/s13414-016-1111-x

Draschkow, D., & Võ, M. L.-H. (2017). Scene grammar shapes the way we interact with objects, strengthens memories, and speeds search. *Scientific Reports*, *7*(1), 16471. https://doi.org/10.1038/s41598-017-16739-x

Draschkow, D., Wolfe, J. M., & Võ, M. L.-H. (2014). Seek and you shall remember: scene semantics interact with visual search to build better memories. *Journal of Vision*, *14*(8), 10. https://doi.org/10.1167/14.8.10

Drew, T., Boettcher, S. E. P., & Wolfe, J. M. (2016). Searching while loaded: Visual working memory does not interfere with hybrid search efficiency but hybrid search uses working memory capacity. *Psychonomic Bulletin & Review*, *23*(1), 201–212. https://doi.org/10.3758/s13423-015-0874-8

Drew, T., Boettcher, S. E. P., & Wolfe, J. M. (2017). One visual search, many memory searches: An eye-tracking investigation of hybrid search. *Journal of Vision*, *17*(11). https://doi.org/10.1167/17.11.5

Duncan-Johnson, C. C., & Donchin, E. (1977). On Quantifying Surprise: The Variation of Event-Related Potentials With Subjective Probability. *Psychophysiology*, *14*(5), 456–467. https://doi.org/10.1111/j.1469-8986.1977.tb01312.x

Duncan, J., & Humphreys, G. W. (1989). Visual search and stimulus similarity. *Psychological Review*, *96*(3), 433–458. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/2756067

Edelman, S., Grill-Spector, K., Kushnir, T., & Malach, R. (1998). Toward direct visualization of the internal shape representation space by fMRI. *Psychobiology*, *26*(4), 309–321. https://doi.org/10.3758/BF03330618

Egeth, H. E., Virzi, R. A., & Garbart, H. (1984). Searching for conjunctively defined targets. *Journal of Experimental Psychology: Human Perception and Performance*. https://doi.org/10.1037/0096-1523.10.1.32

Elithorn, A., & Lawrence, C. (1955). Central Inhibition—Some Refractory Observations. *Quarterly Journal of Experimental Psychology*, *7*(3), 116–127. https://doi.org/10.1080/17470215508416684

Eriksen, C. W., & Hoffman, J. E. (1973). The extent of processing of noise elements during selective encoding from visual displays. *Perception & Psychophysics*. https://doi.org/10.3758/BF03198630

Ferreira, F., Apel, J., & Henderson, J. M. (2008). Taking a new look at looking at nothing. *Trends in Cognitive Sciences*. https://doi.org/10.1016/j.tics.2008.07.007

Foster, J. J., Sutterer, D. W., Serences, J. T., Vogel, E. K., & Awh, E. (2017). Alpha-Band Oscillations Enable Spatially and Temporally Resolved Tracking of Covert Spatial Attention. *Psychological Science*, *28*(7), 929–941. https://doi.org/10.1177/0956797617699167

Fougnie, D., Suchow, J. W., & Alvarez, G. A. (2012). Variability in the quality of visual working memory. *Nature Communications*. https://doi.org/10.1038/ncomms2237

Foxe, J. J., & Snyder, A. C. (2011). The Role of Alpha-Band Brain Oscillations as a Sensory Suppression

Mechanism during Selective Attention. *Frontiers in Psychology*, *2*, 154. https://doi.org/10.3389/fpsyg.2011.00154

Friedman, D., Cycowicz, Y. M., & Gaeta, H. (2001). The novelty P3: An event-related brain potential (ERP) sign of the brain's evaluation of novelty. *Neuroscience and Biobehavioral Reviews*. https://doi.org/10.1016/S0149-7634(01)00019-7

Friston, K. (2010). The free-energy principle: a unified brain theory? *Nature Reviews. Neuroscience*, *11*(2), 127–138. https://doi.org/10.1038/nrn2787

Friston, K., & Kiebel, S. (2009). Predictive coding under the free-energy principle. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *364*(1521), 1211–1221. https://doi.org/10.1098/rstb.2008.0300

Fukuda, K., & Woodman, G. F. (2017). Visual working memory buffers information retrieved from visual long-term memory. *Proceedings of the National Academy of Sciences of the United States of America*, *114*(20), 5306–5311. https://doi.org/10.1073/pnas.1617874114

Fuster, J. M., & Alexander, G. E. (1971). Neuron activity related to short-term memory. *Science*. https://doi.org/10.1126/science.173.3997.652

Gallistel, C. R. (1990). *The organization of learning. The organization of learning.* Cambridge, MA, US: The MIT Press.

Gallivan, J. P., Bowman, N. A. R., Chapman, C. S., Wolpert, D. M., & Flanagan, J. R. (2016). The sequential encoding of competing action goals involves dynamic restructuring of motor plans in working memory. *Journal of Neurophysiology*. https://doi.org/10.1152/jn.00951.2015

Gallivan, J. P., Chapman, C. S., Gale, D. J., Flanagan, J. R., & Culham, J. C. (2019). Selective Modulation of Early Visual Cortical Activity by Movement Intention. *Cerebral Cortex*. https://doi.org/10.1093/cercor/bhy345

Geng, J. J., & Behrmann, M. (2002). Probability cuing of target location facilitates visual search implicitly in normal participants and patients with hemispatial neglect. *Psychological Science*, *13*(6), 520–525. https://doi.org/10.1111/1467-9280.00491

Geng, J. J., Di Quattro, N. E., & Helm, J. (2017). Distractor probability changes the shape of the attentional template. *Journal of Experimental Psychology: Human Perception and Performance*, *43*(12), 1993–2007. https://doi.org/10.1037/xhp0000430

Geng, J. J., & Witkowski, P. (2019). Template-to-distractor distinctiveness regulates visual search efficiency. *Current Opinion in Psychology*. Elsevier B.V. https://doi.org/10.1016/j.copsyc.2019.01.003

Ghose, G. M., & Maunsell, J. H. R. (2002). Attentional modulation in visual cortex depends on task timing. *Nature*, *419*(6907), 616–620. https://doi.org/10.1038/nature01057

Gibson, J. J. (1979). *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin.

Goldfarb, E. V., Chun, M. M., & Phelps, E. A. (2016). Memory-Guided Attention: Independent Contributions of the Hippocampus and Striatum. *Neuron*, *89*(2), 317–324. https://doi.org/10.1016/j.neuron.2015.12.014

González-García, C., Formica, S., Liefooghe, B., & Brass, M. (2020). Attentional prioritization reconfigures novel instructions into action-oriented task sets. *Cognition*. https://doi.org/10.1016/j.cognition.2019.104059

Gorgoraptis, N., Catalao, R. F. G., Bays, P. M., & Husain, M. (2011). Dynamic Updating of Working Memory Resources for Visual Objects. *Journal of Neuroscience*, *31*(23), 8502–8511. https://doi.org/10.1523/JNEUROSCI.0208-11.2011

Goschy, H., Bakos, S., Mueller, H. J., & Zehetleitner, M. (2014). Probability cueing of distractor locations: Both intertrial facilitation and statistical learning mediate interference reduction. *Frontiers in Psychology*, *5*(OCT). https://doi.org/10.3389/fpsyg.2014.01195

Goslin, J., Dixon, T., Fischer, M. H., Cangelosi, A., & Ellis, R. (2012). Electrophysiological examination of embodiment in vision and action. *Psychological Science*, *23*(2), 152–157. https://doi.org/10.1177/0956797611429578

Gramfort, A., Luessi, M., Larson, E., Engemann, D. A., Strohmeier, D., Brodbeck, C., … Hämäläinen, M. (2013). MEG and EEG data analysis with MNE-Python. *Frontiers in Neuroscience*, *7*, 267.

https://doi.org/10.3389/fnins.2013.00267

Green, B. F., & Anderson, L. K. (1956). Color coding in a visual search task. *Journal of Experimental Psychology*. https://doi.org/10.1037/h0047484

Greenberg, A. S., Esterman, M., Wilson, D., Serences, J. T., & Yantis, S. (2010). Control of spatial and feature-based attention in frontoparietal cortex. *Journal of Neuroscience*, *30*(43), 14330–14339. https://doi.org/10.1523/JNEUROSCI.4248-09.2010

Greene, M. R., & Hansen, B. C. (2020). Disentangling the Independent Contributions of Visual and Conceptual Features to the Spatiotemporal Dynamics of Scene Categorization. *The Journal of Neuroscience*. https://doi.org/10.1523/jneurosci.2088-19.2020

Griffin, I. C., & Nobre, A. C. (2003). Orienting Attention to Locations in Internal Representations. *Journal of Cognitive Neuroscience*, *15*(8), 1176–1194. https://doi.org/10.1162/089892903322598139

Griffiths, B. J., Mayhew, S. D., Mullinger, K. J., Jorge, J., Charest, I., Wimber, M., & Hanslmayr, S. (2019). Alpha/beta power decreases track the fidelity of stimulus-specific information. *ELife*, *8*. https://doi.org/10.7554/eLife.49562

Grootswagers, T., Wardle, S. G., & Carlson, T. A. (2017). Decoding dynamic brain patterns from evoked responses: A tutorial on multivariate pattern analysis applied to time series neuroimaging data. *Journal of Cognitive Neuroscience*, *29*(4), 677–697. https://doi.org/10.1162/jocn_a_01068

Gunseli, E., Meeter, M., & Olivers, C. N. L. (2014). Is a search template an ordinary working memory? Comparing electrophysiological markers of working memory maintenance for visual search and recognition. *Neuropsychologia*, *60*(1), 29–38. https://doi.org/10.1016/j.neuropsychologia.2014.05.012

Gunseli, E., Olivers, C. N. L., & Meeter, M. (2014). Effects of Search Difficulty on the Selection, Maintenance, and Learning of Attentional Templates. *Journal of Cognitive Neuroscience*, *26*(9), 2042–2054. https://doi.org/10.1162/jocn_a_00600

Gutteling, T. P., Petridou, N., Dumoulin, S. O., Harvey, B. M., Aarnoutse, E. J., Leon Kenemans, J., & Neggers, S. F. W. (2015). Action preparation shapes processing in early visual cortex. *Journal of Neuroscience*. https://doi.org/10.1523/JNEUROSCI.1358-14.2015

Haegens, S., Nacher, V., Luna, R., Romo, R., & Jensen, O. (2011). -Oscillations in the monkey sensorimotor network influence discrimination performance by rhythmical inhibition of neuronal spiking. *Proceedings of the National Academy of Sciences*, *108*(48), 19377–19382. https://doi.org/10.1073/pnas.1117190108

Hajonides, J. E., van Ede, F., Stokes, M. G., & Nobre, A. C. (2019). Comparing the prioritisation of items and feature-dimensions in visual working memory. *BioRxiv*, 863191. https://doi.org/10.1101/863191

Hanning, N. M., Jonikaitis, D., Deubel, H., & Szinte, M. (2016). Oculomotor selection underlies feature retention in visual working memory. *Journal of Neurophysiology*. https://doi.org/10.1152/jn.00927.2015

Hanslmayr, S., Staresina, B. P., & Bowman, H. (2016). Oscillations and Episodic Memory: Addressing the Synchronization/Desynchronization Conundrum. *Trends in Neurosciences*, *39*(1), 16–25. https://doi.org/10.1016/j.tins.2015.11.004

Haxby, J. V., Gobbini, M. I., Furey, M. L., Ishai, A., Schouten, J. L., & Pietrini, P. (2001). Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science*, *293*(5539), 2425–2430. https://doi.org/10.1126/science.1063736

Heideman, S. G., Rohenkohl, G., Chauvin, J. J., Palmer, C. E., van Ede, F., & Nobre, A. C. (2018). Anticipatory neural dynamics of spatial-temporal orienting of attention in younger and older adults. *NeuroImage*, *178*, 46–56. https://doi.org/10.1016/j.neuroimage.2018.05.002

Helbing, J., Draschkow, D., & Võ, M. L. H. (2020). Search superiority: Goal-directed attentional allocation creates more reliable incidental identity and location memory than explicit encoding in naturalistic virtual environments. *Cognition*, *196*, 104147. https://doi.org/10.1016/j.cognition.2019.104147

Herring, J. D., Thut, G., Jensen, O., & Bergmann, T. O. (2015). Attention modulates TMS-locked alpha oscillations in the visual cortex. *Journal of Neuroscience*. https://doi.org/10.1523/JNEUROSCI.1833-

15.2015

Heuer, A., Crawford, J. D., & Schubö, A. (2017). Action relevance induces an attentional weighting of representations in visual working memory. *Memory and Cognition*. https://doi.org/10.3758/s13421-016-0670-3

Higuchi, S., & Miyashita, Y. (1996). Formation of mnemonic neuronal responses to visual paired associates in inferotemporal cortex is impaired by perirhinal and entorhinal lesions. *Proceedings of the National Academy of Sciences*, *93*(2), 739–743. https://doi.org/10.1073/PNAS.93.2.739

Hillstrom, A. P. (2000). Repetition effects in visual search. *Perception and Psychophysics*, *62*(4), 800–817. https://doi.org/10.3758/BF03206924

Hodsoll, J., & Humphreys, G. W. (2001a). Driving attention with the top down: the relative contribution of target templates to the linear separability effect in the size dimension. *Perception & Psychophysics*, *63*(5), 918–926. https://doi.org/10.3758/bf03194447

Hodsoll, J., & Humphreys, G. W. (2001b). Driving attention with the top down: The relative contribution of target templates to the linear separability effect in the size dimension. *Perception and Psychophysics*, *63*(5), 918–926. https://doi.org/10.3758/BF03194447

Hommel, B. (2004). Event files: Feature binding in and across perception and action. *Trends in Cognitive Sciences*. https://doi.org/10.1016/j.tics.2004.08.007

Horst, J. S., & Hout, M. C. (2016). The Novel Object and Unusual Name (NOUN) Database: A collection of novel images for use in experimental research. *Behavior Research Methods*, *48*(4), 1393–1409. https://doi.org/10.3758/s13428-015-0647-3

Hout, M. C., & Goldinger, S. D. (2015). Target templates: the precision of mental representations affects attentional guidance and decision-making in visual search. *Attention, Perception & Psychophysics*, *77*(1), 128–149. https://doi.org/10.3758/s13414-014-0764-6

Hubel, D. H., & Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of Physiology*, *160*(1), 106–154. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/14449617

Hubel, D. H., & Wiesel, T. N. (1974). Sequence regularity and geometry of orientation columns in the monkey striate cortex. *Journal of Comparative Neurology*. https://doi.org/10.1002/cne.901580304

Hutchinson, J. B., & Turk-Browne, N. B. (2012). Memory-guided attention: control from multiple memory systems. *Trends in Cognitive Sciences*, *16*(12), 576–579. https://doi.org/10.1016/j.tics.2012.10.003

Iordanescu, L., Guzman-Martinez, E., Grabowecky, M., & Suzuki, S. (2008). Characteristic sounds facilitate visual search. *Psychonomic Bulletin and Review*. https://doi.org/10.3758/PBR.15.3.548

Itti, L., & Koch, C. (2001). Computational modelling of visual attention. *Nature Reviews Neuroscience*, *2*(3), 194–203. https://doi.org/10.1038/35058500

Janssen, P., & Shadlen, M. N. (2005). A representation of the hazard rate of elapsed time in macaque area LIP. *Nature Neuroscience*, *8*(2), 234–241. https://doi.org/10.1038/nn1386

Jensen, O., & Mazaheri, A. (2010). Shaping Functional Architecture by Oscillatory Alpha Activity: Gating by Inhibition. *Frontiers in Human Neuroscience*, *4*, 186. https://doi.org/10.3389/fnhum.2010.00186

Jiang, Y. V. (2018). Habitual versus goal-driven attention. *Cortex*, *102*, 107–120. https://doi.org/10.1016/J.CORTEX.2017.06.018

Jiang, Y. V., Swallow, K. M., & Rosenbaum, G. M. (2013). Guidance of spatial attention by incidental learning and endogenous cuing. *Journal of Experimental Psychology: Human Perception and Performance*, *39*(1), 285–297. https://doi.org/10.1037/a0028022

Jiang, Y. V., Swallow, K. M., Rosenbaum, G. M., & Herzig, C. (2013). Rapid acquisition but slow extinction of an attentional bias in space. *Journal of Experimental Psychology: Human Perception and Performance*, *39*(1), 87–99. https://doi.org/10.1037/a0027611

Johnson, R. (1986). A triarchic model of P300 amplitude. *Psychophysiology*, *23*(4), 367–384. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/3774922

Jones, J. L., & Kaschak, M. P. (2012). Global statistical learning in a visual search task. *Journal of Experimental Psychology: Human Perception and Performance*, *38*(1), 152–160. https://doi.org/10.1037/a0026233

Jones, M. R., Moynihan, H., MacKenzie, N., & Puente, J. (2002). Temporal aspects of stimulus-driven attending in dynamic arrays. *Psychological Science*, *13*(4), 313–319. https://doi.org/10.1111/1467-9280.00458

Jonides, J. (1998). Voluntary versus automatic control over the mind's eye's movement. *Psychonomic Bulletin & Review*. https://doi.org/10.1037/0096-1523.29.5.835

Jonikaitis, D., & Moore, T. (2019). The interdependence of attention, working memory and gaze control: behavior and neural circuitry. *Current Opinion in Psychology*. https://doi.org/10.1016/j.copsyc.2019.01.012

Josephs, E. L., Draschkow, D., Wolfe, J. M., & Võ, M. L.-H. (2016). Gist in time: Scene semantics and structure enhance recall of searched objects. *Acta Psychologica*, *169*, 100–108. https://doi.org/10.1016/j.actpsy.2016.05.013

Jung, T. P., Makeig, S., Humphries, C., Lee, T. W., McKeown, M. J., Iragui, V., & Sejnowski, T. J. (2000). Removing electroencephalographic artifacts by blind source separation. *Psychophysiology*, *37*(2), 163–178. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/10731767

Kahneman, D., Treisman, A. M., & Gibbs, B. J. (1992). The reviewing of object files: Object-specific integration of information. *Cognitive Psychology*. https://doi.org/10.1016/0010-0285(92)90007-O

Kaiser, D., Oosterhof, N. N., & Peelen, M. V. (2016). The neural dynamics of attentional selection in natural scenes. *Journal of Neuroscience*, *36*(41), 10522–10528. https://doi.org/10.1523/JNEUROSCI.1385-16.2016

Kamitani, Y., & Tong, F. (2005). Decoding the visual and subjective contents of the human brain. *Nature Neuroscience*, *8*(5), 679–685. https://doi.org/10.1038/nn1444

King, J.-R., & Dehaene, S. (2014). Characterizing the dynamics of mental representations: the temporal generalization method. *Trends in Cognitive Sciences*, *18*(4), 203–210. https://doi.org/10.1016/j.tics.2014.01.002

Kluyver, T., Ragan-Kelley, B., Pérez, F., Granger, B., Bussonnier, M., Frederic, J., … Development Team, J. (2016). *Jupyter Notebooks-a publishing format for reproducible computational workflows*. Retrieved from https://nbviewer.jupyter.org/,

Kok, P., Brouwer, G. J., van Gerven, M. A. J., & de Lange, F. P. (2013). Prior expectations bias sensory representations in visual cortex. *Journal of Neuroscience*, *33*(41), 16275–16284. https://doi.org/10.1523/JNEUROSCI.0742-13.2013

Kok, P., Failing, M. F., & de Lange, F. P. (2014). Prior expectations evoke stimulus templates in the primary visual cortex. *Journal of Cognitive Neuroscience*, *26*(7), 1546–1554. https://doi.org/10.1162/jocn_a_00562

Kok, P., Jehee, J. F. M., & de Lange, F. P. (2012). Less Is More: Expectation Sharpens Representations in the Primary Visual Cortex. *Neuron*, *75*(2), 265–270. https://doi.org/10.1016/j.neuron.2012.04.034

Kok, P., Mostert, P., & De Lange, F. P. (2017). Prior expectations induce prestimulus sensory templates. *Proceedings of the National Academy of Sciences of the United States of America*, *114*(39), 10473–10478. https://doi.org/10.1073/pnas.1705652114

Kok, P., Rahnev, D., Jehee, J. F. M., Lau, H. C., & De Lange, F. P. (2012). Attention reverses the effect of prediction in silencing sensory signals. *Cerebral Cortex*, *22*(9), 2197–2206. https://doi.org/10.1093/cercor/bhr310

Kriegeskorte, N., Goebel, R., & Bandettini, P. (2006). Information-based functional brain mapping. *Proceedings of the National Academy of Sciences of the United States of America*. https://doi.org/10.1073/pnas.0600244103

Kristjánsson, Á., & Campana, G. (2010). Where perception meets memory: A review of repetition priming in visual search tasks. *Attention, Perception, and Psychophysics*. https://doi.org/10.3758/APP.72.1.5

Krummenacher, J., & Müller, H. J. (2012). Dynamic weighting of feature dimensions in visual search: Behavioral and psychophysiological evidence. *Frontiers in Psychology*. https://doi.org/10.3389/fpsyg.2012.00221

Kuznetsova, A., Brockhoff, P., & Christensen, R. (2017). lmerTest Package: Tests in Linear Mixed Effects Models. *Journal of Statistical Software, 82*(13), 1–26. https://doi.org/doi:http://dx.doi.org/10.18637/jss.v082.i13

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

Lauer, T., Cornelissen, T. H. W. W., Draschkow, D., Willenbockel, V., & Võ, M. L.-H. L. H. (2018). The role of scene summary statistics in object recognition. *Scientific Reports*, *8*(1), 14666. https://doi.org/10.1038/s41598-018-32991-1

Laufs, H., Kleinschmidt, A., Beyerle, A., Eger, E., Salek-Haddadi, A., Preibisch, C., & Krakow, K. (2003). EEG-correlated fMRI of human alpha activity. *NeuroImage*, *19*(4), 1463–1476. https://doi.org/10.1016/s1053-8119(03)00286-6

Lawrence, M. (2013). ez: Easy analysis and visualization of factorial experiments. *R Package Version*. Retrieved from https://scholar.google.com/scholar?cluster=3109928330824408893&hl=en&oi=scholarr

Lee, J., & Geng, J. J. (2019). Flexible weighting of target features based on distractor context. *Attention, Perception, and Psychophysics*. https://doi.org/10.3758/s13414-019-01910-5

Lee, S. H., Kravitz, D. J., & Baker, C. I. (2013). Goal-dependent dissociation of visual and prefrontal cortices during working memory. *Nature Neuroscience*. https://doi.org/10.1038/nn.3452

Lemi, L., Chaumon, M., Crouzet, S. M., & Busch, N. A. (2017). Spontaneous neural oscillations bias perception by modulating baseline excitability. *Journal of Neuroscience*, *37*(4), 807–819. https://doi.org/10.1523/JNEUROSCI.1432-16.2016

Lennie, P. (2003). The cost of cortical computation. *Current Biology*, *13*(6), 493–497. https://doi.org/10.1016/S0960-9822(03)00135-0

Lenth, R. V. (2016). Least-Squares Means: The R Package lsmeans. *Journal of Statistical Software*, *69*(1). https://doi.org/10.18637/jss.v069.i01

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

Li, N., Daie, K., Svoboda, K., & Druckmann, S. (2016). Robust neuronal dynamics in premotor cortex during motor planning. *Nature*, *532*(7600), 459–464. https://doi.org/10.1038/nature17643

Liesefeld, H. R., & Müller, H. J. (2019, October 1). Distractor handling via dimension weighting. *Current Opinion in Psychology*. Elsevier B.V. https://doi.org/10.1016/j.copsyc.2019.03.003

Liu, T., Stevens, S. T., & Carrasco, M. (2007). Comparing the time course and efficacy of spatial and feature-based attention. *Vision Research*, *47*(1), 108–113. https://doi.org/10.1016/j.visres.2006.09.017

Los, S. A., & Heslenfeld, D. J. (2005). Intentional and Unintentional Contributions to Nonspecific Preparation: Electrophysiological Evidence. *Journal of Experimental Psychology: General*, *134*(1), 52–72. https://doi.org/10.1037/0096-3445.134.1.52

Luck, S. J., Chelazzi, L., Hillyard, S. A., & Desimone, R. (1997). Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of macaque visual cortex. *Journal of Neurophysiology*. https://doi.org/10.1152/jn.1997.77.1.24

Luck, S. J., Hillyard, S. A., Mouloua, M., Woldorff, M. G., & et al. (1994). Effects of spatial cuing on luminance detectability: Psychophysical and electrophysiological evidence for early selection. *Journal of Experimental Psychology: Human Perception and Performance*, *20*(4), 887–904. https://doi.org/10.1037//0096-1523.20.4.887

Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, *390*(6657), 279–281. https://doi.org/10.1038/36846

Lupyan, G., & Clark, A. (2015). Words and the World. *Current Directions in Psychological Science*, *24*(4), 279–284. https://doi.org/10.1177/0963721415570732

Machens, C. K., Romo, R., & Brody, C. D. (2010). Functional, but not anatomical, separation of “what” and “when” in prefrontal cortex. *Journal of Neuroscience*. https://doi.org/10.1523/JNEUROSCI.3276-09.2010

Macknik, S. L., & Haglund, M. M. (1999). Optical images of visible and invisible percepts in the primary visual cortex of primates. *Proceedings of the National Academy of Sciences of the United States of America*. https://doi.org/10.1073/pnas.96.26.15208

Malcolm, G. L., Groen, I. I. A., Baker, C. I., Torralba, A., Oliva, A., Al., E., … Awasthi, B. (2016). Making Sense of Real-World Scenes. *Trends in Cognitive Sciences*, *20*(11), 843–856. https://doi.org/10.1016/j.tics.2016.09.003

Malcolm, G. L., & Henderson, J. M. (2009). The effects of target template specificity on visual search in real-world scenes: evidence from eye movements. *Journal of Vision*, *9*(11), 8.1-13. https://doi.org/10.1167/9.11.8

Maljkovic, V., & Nakayama, K. (1994). Priming of pop-out: I. Role of features. *Memory & Cognition*, *22*(6), 657–672. https://doi.org/10.3758/BF03209251

Maljkovic, V., & Nakayama, K. (2000). Priming of popout: III. A short-term implicit memory system beneficial for rapid target selection. *Visual Cognition*, *7*(5), 571–595. https://doi.org/10.1080/135062800407202

Mangun, G. R., & Hillyard, S. A. (1987). The spatial allocation of visual attention as indexed by event-related brain potentials. *Human Factors*, *29*(2), 195–211. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/3610184

Maris, E., & Oostenveld, R. (2007). Nonparametric statistical testing of EEG- and MEG-data. *Journal of Neuroscience Methods*, *164*(1), 177–190. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/17517438

Maunsell, J. H. R., & Treue, S. (2006). Feature-based attention in visual cortex. *Trends in Neurosciences*, *29*(6), 317–322. https://doi.org/10.1016/j.tins.2006.04.001

McAdams, C. J., & Maunsell, J. H. R. (1999). Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. *Journal of Neuroscience*. https://doi.org/10.1523/jneurosci.19-01-00431.1999

Mecklinger, A., Gruenewald, C., Weiskopf, N., & Doeller, C. F. (2004). Motor affordance and its role for visual working memory: Evidence from fMRI studies. *Experimental Psychology*, *51*(4), 258–269. https://doi.org/10.1027/1618-3169.51.4.258

Meeter, M., & Olivers, C. N. L. (2006). Intertrial priming stemming from ambiguity: A new account of priming in visual search. *Visual Cognition*, *13*(2), 202–222. https://doi.org/10.1080/13506280500277488

Michael, C. R. (1981). Columnar organization of color cells in monkey’s striate cortex. *Journal of Neurophysiology*. https://doi.org/10.1152/jn.1981.46.3.587

Miniussi, C., Wilding, E. L., Coull, J. T., & Nobre, A. C. (1999). Orienting attention in time. *Brain*, *122*(8), 1507–1518. https://doi.org/10.1093/brain/122.8.1507

Muhle-Karbe, P. S., Duncan, J., De Baene, W., Mitchell, D. J., & Brass, M. (2017). Neural Coding for Instruction-Based Task Sets in Human Frontoparietal and Visual Cortex. *Cerebral Cortex (New York, N.Y. : 1991)*, *27*(3), 1891–1905. https://doi.org/10.1093/cercor/bhw032

Müller, H. J., Heller, D., & Ziegler, J. (1995). Visual search for singleton feature targets within and across feature dimensions. *Perception & Psychophysics*, *57*(1), 1–17. https://doi.org/10.3758/BF03211845

Müller, H. J., Reimann, B., & Krummenacher, J. (2003). Visual Search for Singleton Feature Targets Across Dimensions: Stimulus- and Expectancy-Driven Effects in Dimensional Weighting. *Journal of Experimental Psychology: Human Perception and Performance*, *29*(5), 1021–1035. https://doi.org/10.1037/0096-1523.29.5.1021

Myers, N. E., Chekroud, S. R., Stokes, M. G., & Nobre, A. C. (2018). Benefits of flexible prioritization in working memory can arise without costs. *Journal of Experimental Psychology: Human Perception and Performance*, *44*(3), 398–411. https://doi.org/10.1037/xhp0000449

Myers, N. E., Rohenkohl, G., Wyart, V., Woolrich, M. W., Nobre, A. C., & Stokes, M. G. (2015). Testing sensory evidence against mnemonic templates. *ELife*, *4*(DECEMBER2015). https://doi.org/10.7554/eLife.09000

Myers, N. E., Stokes, M. G., & Nobre, A. C. (2017). Prioritizing Information during Working Memory: Beyond Sustained Internal Attention. *Trends in Cognitive Sciences*, *21*(6), 449–461. https://doi.org/10.1016/j.tics.2017.03.010

Navalpakkam, V., & Itti, L. (2007). Search Goal Tunes Visual Features Optimally. *Neuron*, *53*(4), 605–617. https://doi.org/10.1016/j.neuron.2007.01.018

Navarro, D. J. (2015). *Learning statistics with R : a tutorial for psychology students and other beginners*. Adelaide, Australia: University of Adelaide. Retrieved from https://cran.r-project.org/web/packages/lsr/index.html

Niemi, P., & Näätänen, R. (1981). Foreperiod and simple reaction time. *Psychological Bulletin*, *89*(1), 133–162. https://doi.org/10.1037/0033-2909.89.1.133

Niklaus, M., Nobre, A. C., & van Ede, F. (2017). Feature-based attentional weighting and spreading in visual working memory. *Scientific Reports*, *7*. https://doi.org/10.1038/srep42384

Nobre, A. C. (2001). Orienting attention to instants in time. *Neuropsychologia*, *39*(12), 1317–1328. https://doi.org/10.1016/S0028-3932(01)00120-8

Nobre, A. C., & Kastner, S. (2014). Attention: Time Capsule 2013. In *The Oxford Handbook of Attention*.

Nobre, A. C., & O'Reilly, J. (2004). Time is of the essence. *Trends in Cognitive Sciences*. Elsevier Ltd. https://doi.org/10.1016/j.tics.2004.07.005

Nobre, A. C., & Rohenkohl, G. (2014). Time for the fourth dimension in attention. In A. C. Nobre & S. Kastner (Eds.), *The Oxford handbook of attention* (pp. 676–721). Oxford, UK: Oxford University Press.

Nobre, A. C., & Stokes, M. G. (2019). Premembering Experience: A Hierarchy of Time-Scales for Proactive Attention. *Neuron*, *104*(1), 132–146. https://doi.org/10.1016/j.neuron.2019.08.030

Nobre, A. C., & van Ede, F. (2017). Anticipated moments: temporal structure in attention. *Nature Reviews Neuroscience*, *19*(1), 34–48. https://doi.org/10.1038/nrn.2017.141

O'Craven, K. M., Downing, P. E., & Kanwisher, N. (1999). fMRI evidence for objects as the units of attentional selection. *Nature*, *401*(6753), 584–587. https://doi.org/10.1038/44134

O'Reilly, J. X., Jbabdi, S., Rushworth, M. F. S., Behrens, T. E. J., Friston, K., Kiebel, S., … Rolls, E. (2013). Brain Systems for Probabilistic and Dynamic Prediction: Computational Specificity and Integration. *PLoS Biology*, *11*(9), e1001662. https://doi.org/10.1371/journal.pbio.1001662

O'Reilly, J. X., McCarthy, K. J., Capizzi, M., & Nobre, A. C. (2008). Acquisition of the temporal and ordinal structure of movement sequences in incidental learning. *Journal of Neurophysiology*, *99*(5), 2731–2735. https://doi.org/10.1152/jn.01141.2007

Ohl, S., & Rolfs, M. (2017). Saccadic eye movements impose a natural bottleneck on visual short-term memory. *Journal of Experimental Psychology: Learning Memory and Cognition*. https://doi.org/10.1037/xlm0000338

Olivers, C. N. L. (2009). What Drives Memory-Driven Attentional Capture? The Effects of Memory Type, Display Type, and Search Type. *Journal of Experimental Psychology: Human Perception and Performance*, *35*(5), 1275–1291. https://doi.org/10.1037/a0013896

Olivers, C. N. L., Meijer, F., & Theeuwes, J. (2006). Feature-based memory-driven attentional capture: Visual working memory content affects visual attention. *Journal of Experimental Psychology: Human Perception and Performance*, *32*(5), 1243–1265. https://doi.org/10.1037/0096-1523.32.5.1243

Olivers, C. N. L., Peters, J., Houtkamp, R., & Roelfsema, P. R. (2011, July). Different states in visual working memory: When it guides attention and when it does not. *Trends in Cognitive Sciences*.

https://doi.org/10.1016/j.tics.2011.05.004

Olmos-Solis, K., van Loon, A. M., Los, S. A., & Olivers, C. N. L. (2017). Oculomotor measures reveal the temporal dynamics of preparing for search. In *Progress in Brain Research*. https://doi.org/10.1016/bs.pbr.2017.07.003

Olson, I. R., & Chun, M. M. (2001). Temporal contextual cuing of visual attention. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *27*(5), 1299–1313. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/11550756

Orhan, A. E., & Ma, W. J. (2019). A diverse range of factors affect the nature of neural representations underlying short-term memory. *Nature Neuroscience*, *22*(2), 275–283. https://doi.org/10.1038/s41593-018-0314-y

Park, Y. E., Sy, J. L., Hong, S. W., & Tong, F. (2017). Reprioritization of Features of Multidimensional Objects Stored in Visual Working Memory. *Psychological Science*, *28*(12), 1773–1785. https://doi.org/10.1177/0956797617719949

Patai, E. Z., Doallo, S., & Nobre, A. C. (2012). Long-Term memories bias sensitivity and target selection in complex scenes. *Journal of Cognitive Neuroscience*, *24*(12), 2281–2291. https://doi.org/10.1162/jocn_a_00294

Peelen, M. V, & Kastner, S. (2011). A neural basis for real-world visual search in human occipitotemporal cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *108*(29), 12125–12130. https://doi.org/10.1073/pnas.1101042108

Peelen, M. V, & Kastner, S. (2014). Attention in the real world: toward understanding its neural basis. *Trends Cogn Sci.*, *18*(5), 242–250.

Pfeuty, M., Ragot, R., & Pouthas, V. (2005). Relationship between CNV and timing of an upcoming event. *Neuroscience Letters*, *382*(1–2), 106–111. https://doi.org/10.1016/j.neulet.2005.02.067

Polich, J. (2007). Updating P300: An Integrative Theory of P3a and P3b. *Clinical Neurophysiology : Official Journal of the International Federation of Clinical Neurophysiology*, *118*(10), 2128. https://doi.org/10.1016/J.CLINPH.2007.04.019

Posner, M. I. (1980). Orienting of attention. *The Quarterly Journal Of Experimental Psychology*, *32*, 3–25. https://doi.org/10.1080/00335558008248231

Postle, B. R., Idzikowski, C., Della Sala, S., Logie, R. H., & Baddeley, A. D. (2006). The selective disruption of spatial working memory by eye movements. *Quarterly Journal of Experimental Psychology*. https://doi.org/10.1080/17470210500151410

Potter, M. (1975). Meaning in visual search. *Science*, *187*(4180), 965–966. https://doi.org/10.1126/science.1145183

Praamstra, P., Kourtis, D., Kwok, H. F., & Oostenveld, R. (2006). Neurophysiology of Implicit Timing in Serial Choice Reaction-Time Performance. *Journal of Neuroscience*, *26*(20), 5448–5455. https://doi.org/10.1523/JNEUROSCI.0440-06.2006

Purves, D., Shimpi, A., & Lotto, R. B. (1999). An empirical explanation of the cornsweet effect. *Journal of Neuroscience*, *19*(19), 8542–8551. https://doi.org/10.1523/jneurosci.19-19-08542.1999

R Core Team. (2018). R: A Language and Environment for Statistical Computing.

Rainer, G., Asaad, W. F., & Miller, E. K. (1998). Selective representation of relevant information by neurons in the primate prefrontal cortex. *Nature*, *393*(6685), 577–579. https://doi.org/10.1038/31235

Rainer, G., Rao, S. C., & Miller, E. K. (1999). Prospective coding for objects in primate prefrontal cortex. *Journal of Neuroscience*, *19*(13), 5493–5505. https://doi.org/10.1523/jneurosci.19-13-05493.1999

Rao, R. P. N., & Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, *2*(1), 79–87. https://doi.org/10.1038/4580

Reeder, R. R., Hanke, M., & Pollmann, S. (2017). Task relevance modulates the cortical representation of feature conjunctions in the target template. *Scientific Reports*, *7*(1), 1–10.

https://doi.org/10.1038/s41598-017-04123-8

Rerko, L., & Oberauer, K. (2013). Focused, unfocused, and defocused information in working memory. *Journal of Experimental Psychology: Learning Memory and Cognition*, *39*(4), 1075–1096. https://doi.org/10.1037/a0031172

Reynolds, J. H., Pasternak, T., & Desimone, R. (2000). Attention increases sensitivity of V4 neurons. *Neuron*. https://doi.org/10.1016/S0896-6273(00)81206-4

Richardson, D. C., & Spivey, M. J. (2000). Representation, space and Hollywood Squares: Looking at things that aren't there anymore. *Cognition*. https://doi.org/10.1016/S0010-0277(00)00084-6

Rohenkohl, G., Gould, I. C., Pessoa, J., & Nobre, A. C. (2014). Combining spatial and temporal expectations to improve visual perception. *Journal of Vision*, *14*(4). https://doi.org/10.1167/14.4.8

Rohenkohl, G., & Nobre, A. C. (2011). Alpha Oscillations Related to Anticipatory Attention Follow Temporal Expectations. *Journal of Neuroscience*, *31*(40), 14076–14084. https://doi.org/10.1523/JNEUROSCI.3387-11.2011

Romei, V., Brodbeck, V., Michel, C., Amedi, A., Pascual-Leone, A., & Thut, G. (2008). Spontaneous fluctuations in posterior α-band EEG activity reflect variability in excitability of human visual areas. *Cerebral Cortex*, *18*(9), 2010–2018. https://doi.org/10.1093/cercor/bhm229

Romei, V., Gross, J., & Thut, G. (2010). On the role of prestimulus alpha rhythms over occipito-parietal areas in visual input regulation: Correlation or causation? *Journal of Neuroscience*, *30*(25), 8692–8697. https://doi.org/10.1523/JNEUROSCI.0160-10.2010

Rose, M., Verleger, R., & Wascher, E. (2001). ERP correlates of associative learning. *Psychophysiology*, *38*(3), 440–450. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/11352132

Sàenz, M., Buracâs, G. T., & Boynton, G. M. (2003). Global feature-based attention for motion and color. *Vision Research*, *43*(6), 629–637. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/12604099

Salenius, S., Kajola, M., Thompson, W. L., Kosslyn, S., & Hari, R. (1995). Reactivity of magnetic parieto-occipital alpha rhythm during visual imagery. *Electroencephalography and Clinical Neurophysiology*, *95*(6), 453–462. https://doi.org/10.1016/0013-4694(95)00155-7

Samaha, J., Gosseries, O., & Postle, B. R. (2017). Distinct oscillatory frequencies underlie excitability of human occipital and parietal cortex. *Journal of Neuroscience*, *37*(11), 2824–2833. https://doi.org/10.1523/JNEUROSCI.3413-16.2017

Sassenhagen, J., & Draschkow, D. (2019). Cluster-based permutation tests of MEG/EEG data do not establish significance of effect latency or location. *Psychophysiology*, e13335. https://doi.org/10.1111/psyp.13335

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

Schmidt, J., & Zelinsky, G. J. (2017). Adding details to the attentional template offsets search difficulty: Evidence from contralateral delay activity. *Journal of Experimental Psychology: Human Perception and Performance*, *43*(3), 429–437. https://doi.org/10.1037/xhp0000367

Schneider, D., Barth, A., & Wascher, E. (2017). On the contribution of motor planning to the retroactive cuing benefit in working memory: Evidence by mu and beta oscillatory activity in the EEG. *NeuroImage*, *162*, 73–85. https://doi.org/10.1016/J.NEUROIMAGE.2017.08.057

Serences, J. T., Ester, E. F., Vogel, E. K., & Awh, E. (2009). Stimulus-specific delay activity in human primary visual cortex. *Psychological Science*, *20*(2), 207–214. https://doi.org/10.1111/j.1467-9280.2009.02276.x

Serences, J. T., Saproo, S., Scolari, M., Ho, T., & Muftuler, L. T. (2009). Estimating the influence of attention on population codes in human visual cortex using voxel-based tuning functions. *NeuroImage*, *44*(1), 223–231. https://doi.org/10.1016/j.neuroimage.2008.07.043

Shaw, M. L., & Shaw, P. (1977). Optimal allocation of cognitive resources to spatial locations. *Journal of Experimental Psychology. Human Perception and Performance*, *3*(2), 201–211.

https://doi.org/10.1037//0096-1523.3.2.201

Siegel, M., Donner, T. H., Oostenveld, R., Fries, P., & Engel, A. K. (2008). Neuronal Synchronization along the Dorsal Visual Pathway Reflects the Focus of Spatial Attention. *Neuron*, *60*(4), 709–719. https://doi.org/10.1016/j.neuron.2008.09.010

Singhal, A., Monaco, S., Kaufman, L. D., & Culham, J. C. (2013). Human fMRI Reveals That Delayed Action Re-Recruits Visual Perception. *PLoS ONE*. https://doi.org/10.1371/journal.pone.0073629

Slatter, K. H. (1960). Alpha rhythms and mental imagery. *Electroencephalography and Clinical Neurophysiology*, *12*(4), 851–859. https://doi.org/10.1016/0013-4694(60)90133-4

Snyder, A. C., & Foxe, J. J. (2010). Anticipatory Attentional Suppression of Visual Features Indexed by Oscillatory Alpha-Band Power Increases:A High-Density Electrical Mapping Study. *Journal of Neuroscience*, *30*(11), 4024–4032. https://doi.org/10.1523/JNEUROSCI.5684-09.2010

Snyder, L. H., Batista, A. P., & Andersen, R. A. (1997). Coding of intention in the posterior parietal cortex. *Nature*. https://doi.org/10.1038/386167a0

Soto, D., Heinke, D., Humphreys, G. W., & Blanco, M. J. (2005). Early, involuntary top-down guidance of attention from working memory. *Journal of Experimental Psychology: Human Perception and Performance*. https://doi.org/10.1037/0096-1523.31.2.248

Soto, D., Hodsoll, J., Rotshtein, P., & Humphreys, G. W. (2008). Automatic guidance of attention from working memory. *Trends in Cognitive Sciences*. https://doi.org/10.1016/j.tics.2008.05.007

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

Squires, N. K., Squires, K. C., & Hillyard, S. A. (1975). Two varieties of long-latency positive waves evoked by unpredictable auditory stimuli in man. *Electroencephalography and Clinical Neurophysiology*, *38*(4), 387–401. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/46819

Staresina, B. P., Michelmann, S., Bonnefond, M., Jensen, O., Axmacher, N., & Fell, J. (2016). Hippocampal pattern completion is linked to gamma power increases and alpha power decreases during recollection. *ELife*, *5*. https://doi.org/10.7554/eLife.17397

Stokes, M. G., Atherton, K., Patai, E. Z., & Nobre, A. C. (2012). Long-term memory prepares neural activity for perception. *Proceedings of the National Academy of Sciences*, *109*(6), E360–E367. https://doi.org/10.1073/pnas.1108555108

Stokes, M. G., Kusunoki, M., Sigala, N., Nili, H., Gaffan, D., & Duncan, J. (2013). Dynamic coding for cognitive control in prefrontal cortex. *Neuron*, *78*(2), 364–375. https://doi.org/10.1016/j.neuron.2013.01.039

Stokes, M. G., Myers, N. E., Turnbull, J., & Nobre, A. C. (2014). Preferential encoding of behaviorally relevant predictions revealed by EEG. *Frontiers in Human Neuroscience*, *8*, 687. https://doi.org/10.3389/fnhum.2014.00687

Stokes, M. G., Thompson, R., Nobre, A. C., & Duncan, J. (2009). Shape-specific preparatory activity mediates attention to targets in human visual cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *106*(46), 19569–19574. https://doi.org/10.1073/pnas.0905306106

Sugase-Miyamoto, Y., Liu, Z., Wiener, M. C., Optican, L. M., & Richmond, B. J. (2008). Short-term memory trace in rapidly adapting synapses of inferior temporal cortex. *PLoS Computational Biology*. https://doi.org/10.1371/journal.pcbi.1000073

Summerfield, C., Trittschuh, E. H., Monti, J. M., Mesulam, M.-M. M., & Egner, T. (2008). Neural repetition suppression reflects fulfilled perceptual expectations. *Nature Neuroscience*, *11*(9), 1004–1006. https://doi.org/10.1038/nn.2163

Summerfield, J. J., Lepsien, J., Gitelman, D. R., Mesulam, M. M., & Nobre, A. C. (2006). Orienting Attention Based on Long-Term Memory Experience. *Neuron*, *49*(6), 905–916. https://doi.org/10.1016/j.neuron.2006.01.021

Summerfield, J. J., Rao, A., Garside, N., & Nobre, A. C. (2011). Biasing perception by spatial long-term memory.

Journal of Neuroscience. https://doi.org/10.1523/JNEUROSCI.5541-10.2011

Sutton, S., Braren, M., Zubin, J., & John, E. R. (1965). Evoked-potential correlates of stimulus uncertainty. *Science*, *150*(3700), 1187–1188. https://doi.org/10.1126/science.150.3700.1187

Theeuwes, J. (2019, October 1). Goal-driven, stimulus-driven, and history-driven selection. *Current Opinion in Psychology*. Elsevier B.V. https://doi.org/10.1016/j.copsyc.2018.12.024

Theeuwes, J., Belopolsky, A., & Olivers, C. N. L. (2009). Interactions between working memory, attention and eye movements. *Acta Psychologica*. https://doi.org/10.1016/j.actpsy.2009.01.005

Thorpe, S., Fize, D., & Marlot, C. (1996). Speed of processing in the human visual system. *Nature*. https://doi.org/10.1038/381520a0

Thut, G., Nietzel, A., Brandt, S. A., & Pascual-Leone, A. (2006). -Band Electroencephalographic Activity over Occipital Cortex Indexes Visuospatial Attention Bias and Predicts Visual Target Detection. *Journal of Neuroscience*, *26*(37), 9494–9502. https://doi.org/10.1523/JNEUROSCI.0875-06.2006

Tinbergen, N. (1960). The natural control of insects in pine woods: Vol. I. Factors influencing the intensity of predation by songbirds. *Archives Neelandaises de Zoologie.*, *1*(13), 265–343.

Todorovic, A., & de Lange, F. P. (2012). Repetition suppression and expectation suppression are dissociable in time in early auditory evoked fields. *Journal of Neuroscience*, *32*(39), 13389–13395. https://doi.org/10.1523/JNEUROSCI.2227-12.2012

Todorovic, A., van Ede, F., Maris, E., & de Lange, F. P. (2011). Prior Expectation Mediates Neural Adaptation to Repeated Sounds in the Auditory Cortex: An MEG Study. *Journal of Neuroscience*, *31*(25), 9118–9123. https://doi.org/10.1523/JNEUROSCI.1425-11.2011

Trapp, S., & Bar, M. (2015). Prediction, context, and competition in visual recognition. *Annals of the New York Academy of Sciences*, *1339*(1), 190–198. https://doi.org/10.1111/nyas.12680

Treisman, A. M. (1996). The binding problem. *Current Opinion in Neurobiology*. https://doi.org/10.1016/S0959-4388(96)80070-5

Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, *12*(1), 97–136. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/7351125

Treue, S., & Martinez Trujillo, J. C. (1999). Feature-based attention influences motion processing gain in macaque visual cortex. *Nature*, *399*(6736), 575–579. https://doi.org/10.1038/21176

Trewartha, K. M., Case, S., & Flanagan, J. R. (2015). Integrating actions into object location memory: A benefit for active versus passive reaching movements. *Behavioural Brain Research*, *279*, 234–239. https://doi.org/10.1016/j.bbr.2014.11.043

Tucker, M., & Ellis, R. (1998). On the relations between seen objects and components of potential actions. *Journal of Experimental Psychology: Human Perception and Performance*, *24*(3), 830–846. https://doi.org/10.1037/0096-1523.24.3.830

Turk-Browne, N. B., Isola, P. J., Scholl, B. J., & Treat, T. A. (2008). Multidimensional Visual Statistical Learning. *Journal of Experimental Psychology: Learning Memory and Cognition*, *34*(2), 399–407. https://doi.org/10.1037/0278-7393.34.2.399

Turk-Browne, N. B., Scholl, B. J., Johnson, M. K., & Chun, M. M. (2010). Implicit perceptual anticipation triggered by statistical learning. *Journal of Neuroscience*, *30*(33), 11177–11187. https://doi.org/10.1523/JNEUROSCI.0858-10.2010

Van Den Berg, R., Shin, H., Chou, W. C., George, R., & Ma, W. J. (2012). Variability in encoding precision accounts for visual short-term memory limitations. *Proceedings of the National Academy of Sciences of the United States of America*. https://doi.org/10.1073/pnas.1117465109

van Driel, J., Gunseli, E., Meeter, M., & Olivers, C. N. L. (2017). Local and interregional alpha EEG dynamics dissociate between memory for search and memory for recognition. *NeuroImage*, *149*, 114–128. https://doi.org/10.1016/j.neuroimage.2017.01.031

van Ede, F. (2018). Mnemonic and attentional roles for states of attenuated alpha oscillations in perceptual working memory: a review. *European Journal of Neuroscience*, *48*(7), 2509–2515. https://doi.org/10.1111/ejn.13759

van Ede, F. (2020). Visual working memory and action: Functional links and bi-directional influences. *Visual Cognition*. https://doi.org/10.1080/13506285.2020.1759744

van Ede, F., Chekroud, S. R., & Nobre, A. C. (2019). Human gaze tracks attentional focusing in memorized visual space. *Nature Human Behaviour*. https://doi.org/10.1038/s41562-019-0549-y

van Ede, F., Chekroud, S. R., Stokes, M. G., & Nobre, A. C. (2018). Decoding the influence of anticipatory states on visual perception in the presence of temporal distractors. *Nature Communications*, *9*(1), 1449. https://doi.org/10.1038/s41467-018-03960-z

van Ede, F., Chekroud, S. R., Stokes, M. G., & Nobre, A. C. (2019). Concurrent visual and motor selection during visual working memory guided action. *Nature Neuroscience*, *22*(3), 477–483. https://doi.org/10.1038/s41593-018-0335-6

van Ede, F., de Lange, F., Jensen, O., & Maris, E. (2011). Orienting Attention to an Upcoming Tactile Event Involves a Spatially and Temporally Specific Modulation of Sensorimotor Alpha- and Beta-Band Oscillations. *Journal of Neuroscience*, *31*(6), 2016–2024. https://doi.org/10.1523/JNEUROSCI.5630-10.2011

van Ede, F., de Lange, F. P., & Maris, E. (2014). Anticipation increases tactile stimulus processing in the ipsilateral primary somatosensory cortex. *Cerebral Cortex*. https://doi.org/10.1093/cercor/bht111

van Ede, F., Jensen, O., & Maris, E. (2017). Supramodal Theta, Gamma, and Sustained Fields Predict Modality-specific Modulations of Alpha and Beta Oscillations during Visual and Tactile Working Memory. *Journal of Cognitive Neuroscience*, *29*(8), 1455–1472. https://doi.org/10.1162/jocn_a_01129

van Ede, F., Niklaus, M., & Nobre, A. C. (2017). Temporal Expectations Guide Dynamic Prioritization in Visual Working Memory through Attenuated α Oscillations. *Journal of Neuroscience*, *37*(2), 437–445. https://doi.org/10.1523/JNEUROSCI.2272-16.2016

van Moorselaar, D., & Slagter, H. A. (2019). Learning What Is Irrelevant or Relevant: Expectations Facilitate Distractor Inhibition and Target Facilitation through Distinct Neural Mechanisms. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *39*(35), 6953–6967. https://doi.org/10.1523/JNEUROSCI.0593-19.2019

Vickery, T. J., King, L. W., & Jiang, Y. (2005). Setting up the target template in visual search. *Journal of Vision*. https://doi.org/10.1167/5.1.8

Viola, F. C., Thorne, J., Edmonds, B., Schneider, T., Eichele, T., & Debener, S. (2009). Semi-automatic identification of independent components representing EEG artifact. *Clinical Neurophysiology*, *120*(5), 868–877. https://doi.org/10.1016/j.clinph.2009.01.015

Võ, M. L.-H., Boettcher, S. E. P., & Draschkow, D. (2019). Reading scenes: how scene grammar guides attention and aids perception in real-world environments. *Current Opinion in Psychology*, *29*, 205–210. https://doi.org/10.1016/J.COPSYC.2019.03.009

Võ, M. L.-H., & Wolfe, J. M. (2012). When does repeated search in scenes involve memory? Looking at versus looking for objects in scenes. *Journal of Experimental Psychology. Human Perception and Performance*, *38*(1), 23–41. https://doi.org/10.1037/a0024147

Vogel, E. K., & Machizawa, M. G. (2004). Neural activity predicts individual differences in visual working memory capacity. *Nature*. https://doi.org/10.1038/nature02447

von Helmholtz, H. von. (1866). *Handbook of Physiological Optics*. Leipzig: Leopold Voss.

Waldhauser, G. T., Braun, V., & Hanslmayr, S. (2016). Episodic Memory Retrieval Functionally Relies on Very Rapid Reactivation of Sensory Information. *Journal of Neuroscience*, *36*(1), 251–260. https://doi.org/10.1523/JNEUROSCI.2101-15.2016

Wallis, G., Stokes, M. G., Cousijn, H., Woolrich, M., & Nobre, A. C. (2015). Frontoparietal and cingulo-opercular networks play dissociable roles in control of working memory. *Journal of Cognitive Neuroscience*, *27*(10), 2019–2034. https://doi.org/10.1162/jocn_a_00838

Walter, W. G., Cooper, R., Aldridge, V. J., McCallum, W. C., & Winter, A. L. (1964). Contingent Negative Variation : An Electric Sign of Sensori-Motor Association and Expectancy in the Human Brain. *Nature*, *203*(4943), 380–384. https://doi.org/10.1038/203380a0

Wang, B., & Theeuwes, J. (2018). How to inhibit a distractor location? Statistical learning versus active, top-down suppression. *Attention, Perception, and Psychophysics*, *80*(4), 860–870. https://doi.org/10.3758/s13414-018-1493-z

Warden, M. R., & Miller, E. K. (2010). Task-dependent changes in short-term memory in the prefrontal cortex. *Journal of Neuroscience*, *30*(47), 15801–15810. https://doi.org/10.1523/JNEUROSCI.1569-10.2010

Wen, T., Duncan, J., & Mitchell, D. J. (2019). The time-course of component processes of selective attention. *NeuroImage*, *199*, 396–407. https://doi.org/10.1016/j.neuroimage.2019.05.067

Westfall, P. H., Young, S. S., & Wright, S. P. (1993). On Adjusting P-Values for Multiplicity. *Biometrics*, *49*(3), 941. https://doi.org/10.2307/2532216

Wickham, H. (2009). Ggplot2: elegant graphics for data analysis. New York: Springer. https://doi.org/10.1007/978-0-387-98141-3

Wildegger, T., van Ede, F., Woolrich, M., Gillebert, C. R., & Nobre, A. C. (2017). Preparatory α-band oscillations reflect spatial gating independently of predictions regarding target identity. *Journal of Neurophysiology*, *117*(3), 1385–1394. https://doi.org/10.1152/jn.00856.2016

Witkowski, P., & Geng, J. J. (2019). Learned feature variance is encoded in the target template and drives visual search. *Visual Cognition*. https://doi.org/10.1080/13506285.2019.1645779

Wolfe, J. M. (1994). Guided Search 2.0 A revised model of visual search. *Psychonomic Bulletin & Review*, *1*(2), 202–238. https://doi.org/10.3758/BF03200774

Wolfe, J. M. (2012). Saved by a log: how do humans perform hybrid visual and memory search? *Psychological Science*, *23*(7), 698–703. https://doi.org/10.1177/0956797612443968

Wolfe, J. M. (2020). Visual Search: How Do We Find What We Are Looking For? *Annual Review of Vision Science*, *6*(1), annurev-vision-091718-015048. https://doi.org/10.1146/annurev-vision-091718-015048

Wolfe, J. M., Boettcher, S. E. P., Josephs, E. L., Cunningham, C. A., & Drew, T. (2015). You look familiar, but I don’t care: Lure rejection in hybrid visual and memory search is not based on familiarity. *Journal of Experimental Psychology: Human Perception and Performance*, *41*(6), 1576–1587. https://doi.org/10.1037/xhp0000096

Wolfe, J. M., Drew, T., & Boettcher, S. E. P. (2015). *Hybrid search: Picking up a thread from schneider and shiffrin (1977)*. *Cognitive Modeling in Perception and Memory: A Festschrift for Richard M. Shiffrin*.

Wolfe, J. M., & Horowitz, T. S. (2017). Five factors that guide attention in visual search. *Nature Human Behaviour*, *1*(3), 0058. https://doi.org/10.1038/s41562-017-0058

Wolfe, J. M., Horowitz, T. S., Kenner, N., Hyle, M., & Vasan, N. (2004). How fast can you change your mind? The speed of top-down guidance in visual search. In *Vision Research* (Vol. 44, pp. 1411–1426). https://doi.org/10.1016/j.visres.2003.11.024

Wolfe, J. M., Võ, M. L.-H. L. H., Evans, K. K., & Greene, M. R. (2011). Visual search in scenes involves selective and nonselective pathways. *Trends in Cognitive Sciences*, *15*(2), 77–84. https://doi.org/10.1016/j.tics.2010.12.001

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*, 123. https://doi.org/10.3389/fnsys.2015.00123

Woodman, G. F., Luck, S. J., & Schall, J. D. (2007). The role of working memory representations in the control of attention. *Cerebral Cortex*, *17*(SUPPL. 1). https://doi.org/10.1093/cercor/bhm065

Worden, M. S., Foxe, J. J., Wang, N., & Simpson, G. V. (2000). Anticipatory biasing of visuospatial attention indexed by retinotopically specific alpha-band electroencephalography increases over occipital cortex.

The Journal of Neuroscience : The Official Journal of the Society for Neuroscience, *20*(6), RC63. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/10704517

Yang, H., & Zelinsky, G. J. (2009). Visual search is guided to categorically-defined targets. *Vision Research*, *49*(16), 2095–2103. https://doi.org/10.1016/j.visres.2009.05.017

Yu, X., & Geng, J. J. (2019). The attentional template is shifted and asymmetrically sharpened by distractor context. *Journal of Experimental Psychology: Human Perception and Performance*, *45*(3), 336–353. https://doi.org/10.1037/xhp0000609

Zanto, T. P., Pan, P., Liu, H., Bollinger, J., Nobre, A. C., & Gazzaley, A. (2011). Age-Related Changes in Orienting Attention in Time. *Journal of Neuroscience*, *31*(35), 12461–12470. https://doi.org/10.1523/JNEUROSCI.1149-11.2011

Zhang, W., & Luck, S. J. (2009). Feature-based attention modulates feedforward visual processing. *Nature Neuroscience*, *12*(1), 24–25. https://doi.org/10.1038/nn.2223

Zhao, J., Al-Aidroos, N., & Turk-Browne, N. B. (2013). Attention Is Spontaneously Biased Toward Regularities. *Psychological Science*, *24*(5), 667–677. https://doi.org/10.1177/0956797612460407

Zokaei, N., Board, A. G., Manohar, S. G., & Nobre, A. C. (2019). Modulation of the pupillary response by the content of visual working memory. *Proceedings of the National Academy of Sciences of the United States of America*. https://doi.org/10.1073/pnas.1909959116