Examining the role of adaptation in spontaneous perceptual reversals of ambiguous motion stimuli

Kimberley Dundas

A thesis submitted for the degree of Ph.D. in the Faculty of Natural Sciences at Keele University

October 2024

Abstract

The cause of spontaneous perceptual reversals of visually ambiguous stimuli has been attributed to several different mechanisms. One hypothesis suggests that adaptation, or “neural fatigue”, builds up while one interpretation is dominant and eventually triggers a reversal. Although there is behavioural evidence that adaptation to an unambiguous stimulus (e.g., directional motion) can bias subsequent perception of an ambiguous one, it is unclear whether that adaptation plays a role in spontaneous reversals. In three experiments, we used psychophysical, ERP and multivariate pattern analysis (MVPA) of EEG data. Our behavioural results confirmed that adaptation to directional motion does indeed bias subsequent perception of ambiguous motion. We then used MVPA of EEG data to test whether that adaptation is involved in spontaneous perceptual reversals. We did this by training a machine-learning classifier on the pattern of EEG scalp voltage during adaptation-induced reversals. We then tested this classifier on spontaneous reversal data and found above-chance decoding and similar activity patterns. To shed light on the statistical power of the MVPA technique used to generate our results, we ran a series of simulations that manipulated effect size and sample size. We found that statistical power of the MVPA pipeline employed in this thesis and commonly in the wider literature, indeed increased when larger sample sizes were used as well as when larger effect sizes were added to the simulated data. Taken together, the results from this thesis suggest that similar brain mechanisms mediate perceptual adaptation and spontaneous perceptual reversals, providing support for the “neural fatigue hypothesis” of multi-stable perception. Additionally, the analyses used to generate

this claim have been shown to be sufficiently powerful in which to detect these effects.

Acknowledgements

I would first and foremost like to thank my supervisor, Dr Joseph Brooks. His unwavering calmness and endless knowledge have played no small part in this thesis existing today. I cannot thank him enough for providing support, kindness and praise at the exact moments they were needed. I have never failed to be amazed at his ability to turn my moments of sheer panic and uncertainty into ones of clarity. He has been so generous with his time, there has been nothing I've asked of him that he's not gone above and beyond to deliver. I would also like to thank Dr Brooks' previous PhD student, and now a flourishing academic in her own right, Dr Diane Abdallah. Her kindness and friendship during the early stages of my PhD were invaluable to me. I'm so proud of the important work she's doing now. I would also like to thank Kym Ellis and Adam Coates, for their help with data collection. Their kindness, patience and humour have helped me so much. I would like to extend a special thanks to all the participants who took part in my studies, without them this research simply would not be possible.

I would like to thank my son, Jack. His patience and understanding throughout my PhD have been far beyond that I would expect from an adult, let alone a teenager. I will forever appreciate and admire his resilience, humour and intelligence. His encouragement and 5am wake up calls have given me so much motivation whilst completing this thesis. For both of these things, I will be forever grateful. I would like to thank my husband, Joe. He is the funniest person I know and my biggest source of support. He constantly reminds me that I am capable of more than I think. His encouragement throughout my education has been invaluable and I simply would not have got as far as I have without him.

I would like to thank my sister, Jacqui. Throughout my entire life she has been my role model, my biggest cheerleader and my best friend. It is not an understatement to say that I simply would not be where I am today without her. I would like to thank my parents, Pat and Kevin. They have always encouraged me to follow my dreams and taught me that anything is achievable.

I would like to thank my oldest friends, Nicky, and Sarah. They are my extended family and without their love and support throughout my PhD, this would not have been possible. They constantly provide me with laughs, encouragement and remind me that I can do all the things that I think I can't. I would also like to thank my officemates and friends at Keele University: Emma, Ollie, Claire, Megan, Kara, Nicola, Jamie, Nick, Charlotte and Geraldine. Their support, guidance, and virtual get-togethers meant that completing a PhD amidst a global pandemic didn't feel quite so scary. I'm so proud of where they all are now.

Last and most certainly not least, I would like to thank my 'Gramma', Irene, who passed away before I completed this thesis. It is difficult to put into words just how strong she was. There was nothing I could come to her with that she would not instantly make me feel better about. She had an unfaltering belief in me and taught me that there is nothing that can't be solved if you put your mind to it.

Table of Contents

Chapter One: Introduction

1.1 Multi-stable perception

Our visual system is essential in creating a unified and practical, though not always accurate, depiction of the external environment. The information we gather through our senses is incomplete, due to the limited nature of our sensory inputs and factors such as occlusion, transparency, and illumination changes. To address this challenge, our visual system seems to utilise a wealth of pre-existing knowledge, encompassing statistical norms of the external world, to interpret and understand our surroundings (Carbon, 2014; Pastukhov et al., 2023)

The reason for this ambiguity is that our sensory engagement with the environment is primarily through the proximal stimulus, that is, the two-dimensional images formed on our retinas. This proximal stimulus is a direct result of physical interactions, representing the patterns and properties of physical stimuli within our sensory reach. It is, however, subject to variations depending on external factors like lighting and spatial orientation of objects. For instance, there is not a one-to-one mapping between the proximal stimulus (i.e., the 2D retinal projection) and the distal stimulus (i.e., the 3D scene we are viewing). This means that there are many distal stimuli that could give rise to the same proximal stimuli. The visual system, therefore, engages in a complex process of interpreting these proximal stimuli to construct mental representations of the distal stimuli - the actual objects and their states in the external world. This interpretation heavily relies on inferential mechanisms, drawing upon past experiences and contextual knowledge to make sense of the sensory input (Friston 2005, Petrovici et al. 2016, Summerfield & de Lange 2014). For instance, the perception of a tree involves the translation of the light patterns hitting the retina into a recognisable form of a tree. This interpretive process, while sophisticated and generally effective, is not infallible. It can lead to discrepancies between the perceived image and the actual properties of the distal stimulus, highlighting the intricate and sometimes fallible nature of our visual perception system.

Our understanding of the neural mechanisms that underpin this complex process of visual perception is still being developed in academic research. To put together information on this complex process, research has often focussed on the interaction between bottom-up and top-down mechanisms involved in visual perception. The "bottom-up" approach suggests that early visual processing is characterised by "passive" and automatic mechanisms (e.g., Köhler, 1940; Toppino & Long, 1987). On the other hand, the "top-down" approach suggests that perception is driven by "active", volitional processes that occur closer to conscious awareness (e.g., Horlitz & O'Leary, 1993; Leopold & Logothetis, 1999; Long & Toppino, 2004). This presents an apparent contradiction between the two approaches in understanding the underlying causes of our perception, given that the bottom-up approach suggests that our perception is guided by automatic processes and the top-down approach suggests that it is instead driven by volitional ones. This initial contradiction has since moved toward more integrated theories of perception that are focussed on the dynamic relationship between the two seemingly opposing views (e.g., Blake & Logothetis, 2002; Kornmeier, Hein & Bach, 2009; Kornmeier & Bach, 2005, 2006, 2012). These approaches are discussed in full in Section 1.2, below.

Given that visual perception is inherently ambiguous, therefore, this makes it possible for us to have multi-stable perceptual experiences. The term "multi-stable perception" refers to the phenomenon whereby the incoming sensory information is compatible with different, equally probable perceptual interpretations. This leads to multi-stability, meaning our perception spontaneously alternates between two (known as 'bistable') or several (referred to as 'multi-stable') interpretations despite the physical sensory input from the stimulus being constant (e.g., Tong, Meng, & Blake, 2006). To study this phenomenon, stimuli that are designed or selected to elicit this multi-stable perception are often used. These are known collectively as ambiguous stimuli and will be discussed in more detail below.

1.1.1 Ambiguous stimuli

Ambiguous stimuli are thought to offer a unique window into the mechanisms of subjective perceptual inference (Long & Toppino, 2004). This is because our perception of these stimuli can spontaneously alternate between two or more mutually exclusive subjective perceptual interpretations without any change in the sensory input. Thus, any changes in perceptual experience of ambiguous figures can be directly attributed to correlated brain activity changes without the confounding factor of changes in sensory input. For instance, when viewing Rubin’s faces-vase image (Figure 1.1 A; Rubin, 1921) we can see either two profile faces with a space between or a central, single vase. Our perception can spontaneously alternate between these two outcomes whilst the physical stimulus remains unchanged. The Necker cube (Necker, 1832), old/young woman (Boring, 1930), and the duck/rabbit (Wittgenstein & Anscombe, 1953; Figure 1.1 B-D) are other examples of this perceptual multi-stability effect. This outcome can also be achieved by presenting conflicting images separately to the two eyes in a paradigm known as binocular rivalry (see Figure 1.1 E for an example). The binocular rivalry paradigm can evoke several outcomes. If the images are sufficiently similar to one another they may be combined flawlessly as in normal binocular vision. Alternatively, they can combine in a “patchwork” pattern (Logothetis, Leopold & Sheinberg, 1996). Usually in experiments involving binocular rivalry the images are conflicting, such that one percept dominates over another at any point in time. This dominant percept then alternates over a given period.

In addition to the stationary stimuli described above, perceptual multi-stability can also be induced in response to ambiguous motion stimuli; an example of which is shown in Figure 1.1 F below (Kanai & Verstraten, 2005). This stimulus is created from a drifting sine-wave grating with a phase shift of ±180° in space on every frame, leading to the generation of ambiguous motion that can be perceived as moving either to the left or to the right.

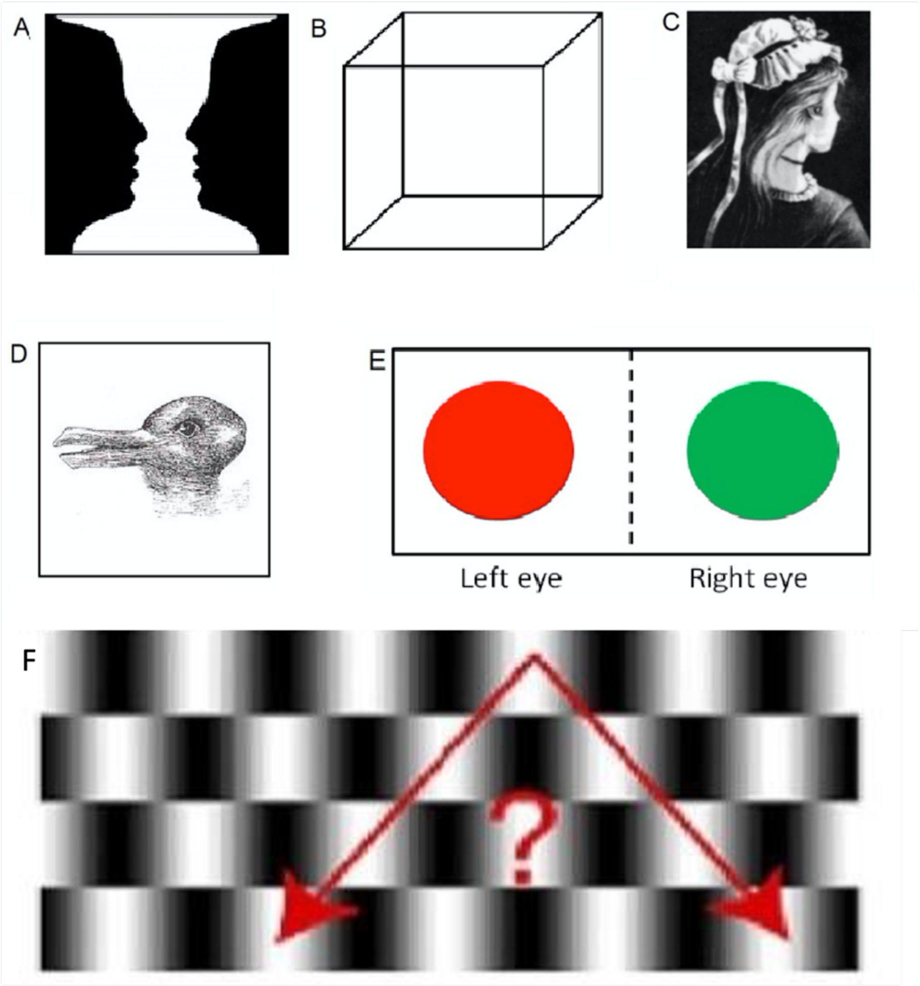

Figure 1.1: Examples of visually ambiguous stimuli. (A) Rubin's face/vase, showing two profile faces with a space between them, or a central vase (Rubin, 1921). (B) The Necker cube, which can be interpreted as having the lower left, or upper right square as its front face (Necker, 1832). (C) Boring's old/young woman, which can be interpreted as being the right profile of an old woman facing towards the viewer, or the right profile of a young woman facing away from the viewer (Boring, 1930). (D) Wittgenstein's duck/rabbit, which can be seen as a duck facing the left, or a rabbit facing the right with its ears to the left (Wittgenstein & Anscombe, 1953). (E) Binocular rivalry, whereby the images on each side of the dotted line are presented to one eye only using a mirror

stereoscope, for example. (F) Example of a still image from an ambiguous motion stimulus, in which the black and white regions can be seen moving either to the left or the right. The rows seen in this figure represent sequential frames of motion. Given that the lighter regions are replaced by dark regions on every frame, this leads to apparent motion to either the left or the right (Kanai & Verstraten, 2005).

These alternations, or switches, in perception that occur during viewing of ambiguous stimuli are known as perceptual reversals. In the extant literature, reversals have been classified into two types: endogenous and exogenous. Exogenous reversals refer to switches in perception that occur due to some aspect of the ambiguous stimulus, or its viewing conditions, being physically changed to bias viewers' perception in a certain way. For instance, adding cues that give the viewer information about the likely three-dimensional structure of the otherwise ambiguous Necker cube can bias perception towards one or the other percept. For example, as shown in Figure 1.2 below, cues on the likely light source and depth of the cube are used to bias perception. In Figure 1.2, Panel B, light cues are introduced (via the shading of the edges of the cube) that are suggestive of the bottom-right face of the cube as being frontal, whereas in Panel C light cues are added that suggest that the top-left face is frontal.

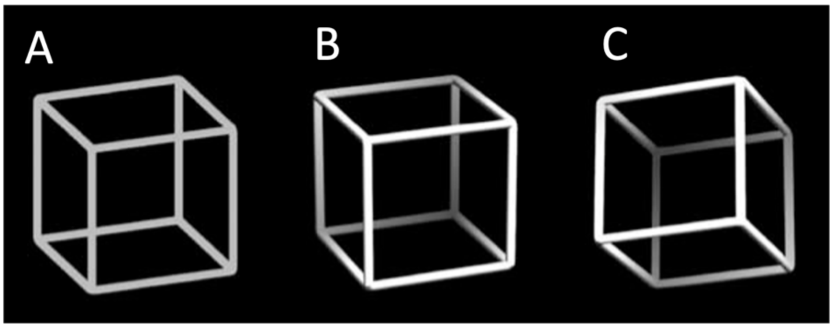

Figure 1.2 The Necker cube and biased variants, created by adding light and depth cues. (A) The ambiguous Necker cube. (B) Bottom-right face as frontal biased version. (C) Top-left face as frontal biased version. Adapted from Kornmeier and Bach (2005).

Endogenous reversals, on the other hand, are reversals in perception that occur spontaneously without any change to the physical stimulus or its viewing conditions. Endogenous reversals can occur when ambiguous stimuli are presented continuously to participants, but also when they are presented intermittently (these two presentation types are discussed in detail in Chapter 2, Section 2.2.2.1) with gaps between them. Trials aimed to elicit endogenous perceptual reversals are therefore created by either continuous or intermittent presentation of the same ambiguous stimulus (e.g., the Necker cube as in Figure 1.2 A). In contrast, trials aimed to induce exogenous perceptual reversals are generated by switching between two or more biased versions of the stimulus (e.g., between the two variants shown in Figure 1.2 A and B). This physical switch in stimulus can either occur during continuous viewing, or intermittent viewing. To generate exogenous periods of a stable perception therefore, the same version of the biased stimulus would be repeatedly presented, therefore there should be no change in perceptual experience and presumably the underlying neural activity related to that experience. Whereas, to create exogenous reversals, different versions of the biased stimulus are presented which does induce a change in perceptual experience and again, presumably, the underlying neural mechanisms that underpin that experience. For endogenous reversals, the perceptual experience is only governed by the change in neural activity and not by any change in the stimulus. Thus, analyses of brain activity corresponding to perceptual reversals of the same visual stimulus (for example, seeing a face or a vase in Rubin's faces/vase, Figure 1.1 A) are not complicated by variations in the incoming sensory information, provided that the viewing circumstances (such as eye movement or position, and lighting conditions) are consistent. Examining behavioural and electrophysiological responses to this kind of stimuli can therefore provide an insight into the neurological bases of conscious perception. For this reason, much of the available research using ambiguous stimuli to study perception involves endogenous reversals. However, comparing behavioural and neural activity between endogenous and exogenous perceptual reversals can offer an insight into the activity that is related only to endogenous changes in perception rather than any perceptual change. This is of importance because it allows researchers an insight into subjective perception, in other words the changes that occur in the brain to alter our interpretation of the world whilst the physical representation of that world is unchanged.

In addition to their use in inducing exogenous perceptual reversals, biased versions of ambiguous stimuli have also been used to study adaptation and priming effects (e.g., Toppino & Long, 1987). Trials designed to this effect are typically created by presenting a biased variant of the ambiguous stimulus for a given period, followed by its ambiguous counterpart. The effects of the pre-exposure to the biased variant are then examined. This concept, and a review of the current research around it, are discussed in detail in Section 1.2.1.

Research around these induced or spontaneous reversals in perception have led to several theories around how we resolve the ambiguity we are faced with in multi-stable experiences. These theories will be discussed in the following section.

1.2 Theories of multi-stability

The complex nature of perceptual reversals has been examined by numerous psychophysical studies, leading to the identification of two main interpretative approaches: the top-down, cognitive approach (e.g., Leopold and Logothetis, 1999) and the sensory-driven bottom-up approach (e.g., Attneave, 1971; Toppino & Long, 1987; Blake, 1989). More recently, researchers have advocated for a holistic perspective that integrates both top-down and bottom-up approaches (e.g., Long & Toppino, 2004; Kornmeier & Bach, 2012). This section will discuss the bottom-up, top-down and integrated mechanisms that have been theorised to underpin the perception of ambiguous figures and lead to perceptual reversals.

1.2.1 Bottom-up approach

The bottom-up theory emphasises the primary influence of the early visual areas, with factors such as adaptation, stimulus characteristics, and presentation mode shaping the perception of ambiguous figures (Kornmeier & Bach, 2012; Peterson, 2014). Several studies have aimed to investigate the effects of low-level manipulations on the reversal dynamics of ambiguous figures, in order to examine the bottom-up mechanisms involved. Results from these studies have identified several factors that do indeed influence reversal dynamics and have led to the development of a theoretical understanding of multi-stable perception. This section will outline these in detail.

1.2.1.1 The Neural Fatigue Hypothesis

Early models put forward to explain multi-stable phenomena were largely based on behavioural observations that reversal rate increases over time when ambiguous images are viewed (e.g., Babich & Standing, 1981; Long, Toppino & Kostenbauder, 1983). Kohler (1940) and later, Kohler and Wallach (1944), theorised that the perception of an ambiguous image invokes field effects in the brain and proposed a theory of "neural satiation". This early model claimed that a reversal occurs because of a steady increase in resistance to the "field flow" that underpins the initial percept. This, the authors claim, continues until a point is reached whereby the "field" underlying an alternative percept becomes less resistant than that of the initial one, causing a rapid switch in the conscious perception of the image.

Kohler's original concept was further developed by Hock et al. (1996) and has since evolved into the "neural fatigue" (or neural adaptation; Long & Toppino, 2004) and "cross inhibition" model. This model is based on the idea that each interpretation of an ambiguous stimulus has its own neural representation which inhibits the representations of other alternative percepts (Wilson, Alais & Blake, 2005). Following periods of prolonged exposure, the neural representation underpinning the currently dominant percept becomes "fatigued" and so is no longer able to inhibit the alternative percept. At this point a threshold is reached whereby the neural representation of the alternative percept is less fatigued than that of the current one. The alternative percept then becomes dominant (i.e., perception is reversed) and its representation becomes strongest and suppresses representations of alternative interpretations. Over time, this now-dominant alternative perceptual representation will also become fatigued and once it becomes weaker than that of the original percept, perception will switch again back. This cycle of adaptation and reversal can continue on indefinitely similarly to the cycle of perceptual reversals that is seen with ambiguous stimuli.

The adaptive state of neurons can persist over time, leading to a decreased threshold for reversal and therefore increased reversal rates. This is thought to occur because as both percept representations become increasingly adapted over time, a slight recovery in the representation of one will lead to a switch in perception toward that percept and vice versa. However, the increase in reversal rate is significantly disrupted when the figure is moved to a different area on the retina (e.g., Babich & Standing, 1981; Kohler, 1940; Long et al., 1983; Toppino & Long, 1987). These findings suggest the presence of localised excitatory and fatigue-like processes, strongly supporting the influence of passive bottom-up processes in these perceptual phenomena (Long & Moran, 2007).

To induce this state of adaptation, studies have used biased versions of ambiguous stimuli, such as the unambiguous Necker cube shown in Figure 1.2. These are used because they are thought to activate the neural representation that underpins a single percept of an ambiguous stimulus (such as the top-right face as frontal version of the Necker cube, or the faces in Rubin’s faces/vase). Research examining the effects of pre-exposure to these unambiguous variants of ambiguous stimuli (which generally follow a similar paradigm to that outlined in

Figure 1.3, below) have provided support for the neural fatigue/adaptation model. Conclusions drawn from these studies suggest that these effects can be divided into two main categories: increased sensitivity to the pre-exposed stimulus (priming; e.g., Long & Moran, 2007), or reduced sensitivity (adaptation; e.g., Long & Olszweski, 1999). Priming means that if pre-exposed to an unambiguous left-facing percept of the Necker cube stimulus, the same left-facing percept is more likely to be perceived upon presentation of a subsequent ambiguous image (see Figure 1.4, below). In contrast, in adaptation, pre-exposure to one version of an image would induce the *opposite* percept to become more likely upon ambiguous stimulus presentation.

Presentation duration and inter-stimulus interval (ISI) have been found to be critical for the induction of one or the other of these effects (discussed in more detail in Section 1.2.1.1, below). Generally speaking, shorter exposure to the unambiguous stimulus has a priming effect whilst longer exposure durations lead to adaptation effects (e.g., Kanai & Verstraten, 2005; Long & Moran, 2007). The existence of the adaptation effect provides support for the bottom-up theory of neural fatigue. In the adaptation effect, if the neural representation underpinning a given percept is activated (by the unambiguous stimulus variant, referred to here as the ‘adapting stimulus’) to the point of adaptation before the ambiguous stimulus (referred to here as the ‘test stimulus’) is presented, the alternate percept will dominate upon presentation providing the inter-stimulus interval is short enough (Long & Moran, 2007; see Figure 1.4).

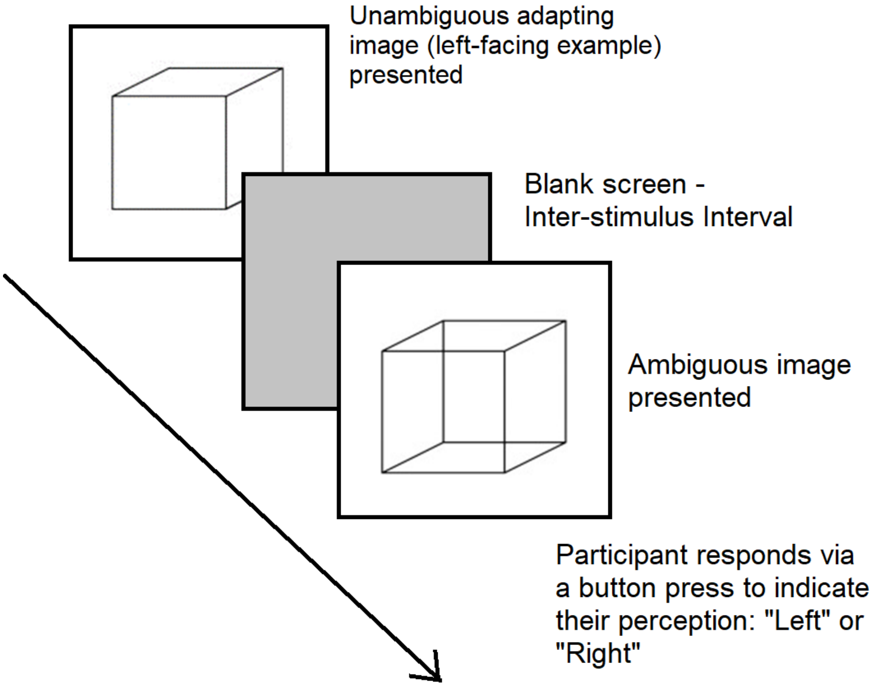

Figure 1.3 Typical trial procedure for studies examining the neural fatigue/adaptation hypothesis.

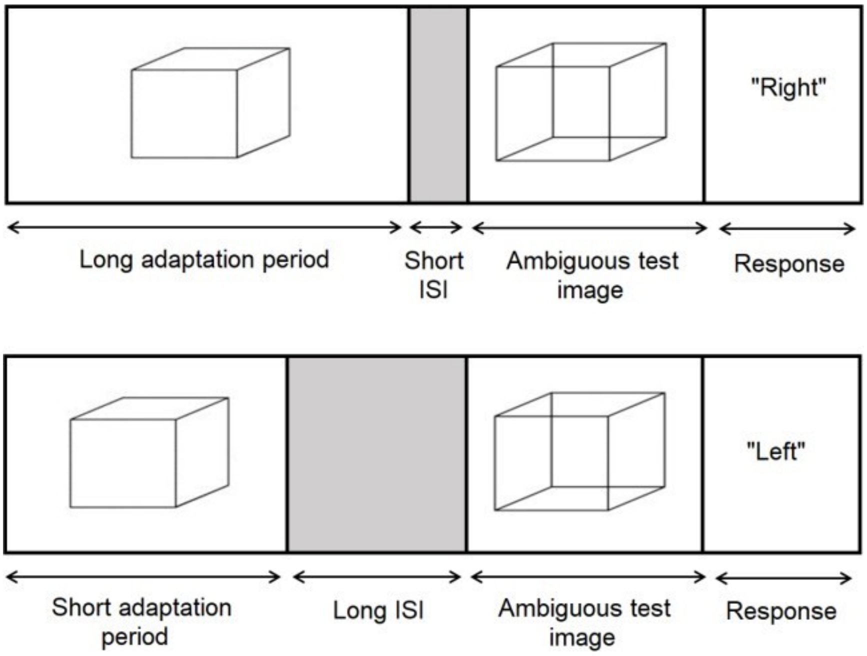

Figure 1.4 Schematic showing a typical trial structure used to induce either an adaptation effect (top) or a priming effect (bottom). In this example, participants would respond to indicate which face they saw as frontal on the second, ambiguous test image. In adaptation trials, this response would typically be the opposite face to the one they were pre-exposed to (in this case the top-right face) and in the priming trials this would typically be the same face as the pre-exposed one (in this case the bottom-left face).

Studies investigating adaptation effects in this way have shown that in order to induce the effect, participants are required to view the adapting stimulus for an extended period of time. For instance, Long & Moran (2007) asked participants to firstly view a clear, unambiguous version of a rotating Necker cube, followed by an ambiguous one, and then asked them to report their perception of the ambiguous stimulus. Their findings showed that adaptation effects (where the perception of the ambiguous stimulus contrasted with the prime) emerged only after lengthy (150 seconds) exposures to the adapting stimulus. However, for shorter exposures, a priming effect was observed. Specifically, after a 2-second exposure to the adapting stimulus, participants tended to perceive the ambiguous test stimulus in a manner consistent with the adapting stimulus. These results indicate that adaptation effects arise only after prolonged exposure to the adapting stimulus. Additional support for the claim that behavioural adaptation effect reflects the involvement of a bottom-up mechanism comes from the result that these effects diminish with extended ISIs (see Section 1.2.1.1 below for a full discussion of the effects of ISI duration). Specifically, Long and Moran (2007) found that adaptation effects were strongest when the ISI was 1s, but they had been completely replaced by a priming effect when the ISI was increased to 100s. Long and Moran (2007) suggest that these results provide evidence for the involvement of two, distinct neural mechanisms involved in adaptation and priming. Additional research has emphasised the brief nature of the adaptation effect, its sensitivity to the intervals between adapting and test stimuli, as well as the presentation duration of the adapting stimulus (Intaite et al., 2013). Studies have revealed that this effect can be significantly diminished by extending the ISI to approximately 10s (Long & Moran, 2007). If any of these conditions are altered, such as a shorter duration for the unambiguous adapting stimulus, presenting adapting and test stimuli at different retinal locations, or employing a prolonged delay between the stimuli, a priming effect is observed (e.g., Long et al., 1992; Long & Olszweski, 1999; Long & Moran, 2007).

Studies have also indicated that adaptation is closely linked to specific attributes of objects. For adaptation to take place, both the adapting stimulus and the test stimulus must share certain configurational characteristics. For example, in the study described in the paragraph above, Long and Moran (2007) also examined the effects of changing the size of the adapting versus test stimuli. They found that when the size of the adapting stimulus differed from that of the test stimulus, adaptation effects were significantly diminished. This was an effect that was replicated from previous research (Toppino & Long, 1987). Other studies, too, have found similar effects. For instance, Spitz and Lipman (1962) showed that reversal rates when participants continuously viewed a Necker cube steadily increased over time. However, in their study, they firstly presented the stimulus in, for example, the left hemifield and then changed the stimulus location such that it was presented to the right hemifield. They found that this switch in location reduced reversal rates back to nearly the same level as when participants first began their viewing with the previous hemisphere. They did not find such an effect when only the retinal position was changed, leading them to conclude that adaptation effects are governed by cortical, rather than retinal, mechanisms. Toppino and Long (1987) replicated these results using a rotating Necker cube stimulus, leading them to conclude that there are distinct neural mechanisms in place that work separately for the left and right visual field.

For the priming effect, many experiments have shown that briefly presenting a biased, unambiguous version of an ambiguous stimulus can predispose an observer to perceive a subsequent ambiguous version in the same way (e.g., Botwinick, 1961; Bugelski & Alampay, 1961; Fisher, 1967; Leeper ,1935; Long et al., 1992). This priming effect has also been shown in studies that have used stimuli other than disambiguated ambiguous figures to influence subsequent perception of ambiguous stimuli. For instance, Goolkasian and Woodberry (2010) used primes that were either images of objects, or object names, that were related to one or the other interpretation of a subsequently presented ambiguous figure. They found that primes with a vague, indirect connection to one of the possible interpretations of an ambiguous figure did indeed bias the interpretation of the ambiguous stimulus towards the primed option. However, this effect occurred only when participants' attention was specifically directed towards the semantic link between the prime and the ambiguous figure. Because of findings like these, researchers have suggested that priming effects could be supportive of more top-down approaches to explain perceptual reversals, therefore these will be discussed in more detail in Section 1.2.2.1 below.

1.2.1.1.1 Visual motion priming and the rapid motion aftereffect

In addition to the relatively lengthy pre-exposure durations used in the studies described above, researchers have also shown priming and adaptation effects using very brief durations with motion stimuli. For instance, Pinkus and Pantle (1997) used ambiguous and unambiguous drifting sine-wave gratings to investigate visual motion priming (VMP). Behaviourally, VMP is a similar phenomenon to the priming effect described above, however it is induced when ambiguous motion is presented after a brief presentation of directional motion. The motion used by Pinkus and Pantle (1997) was created using sine-wave gratings that change their phase over time, for instance after x number of frames or after x milliseconds. To create directional motion, the phase is shifted by ±90° each time, leading to perceived directional motion to the left or the right (see Figure 1.5, top). To create ambiguous motion, the phase is instead shifted by 180° each time, leading to motion that can be perceived as moving in either a leftward or a rightward direction (see Figure 1.5, bottom).

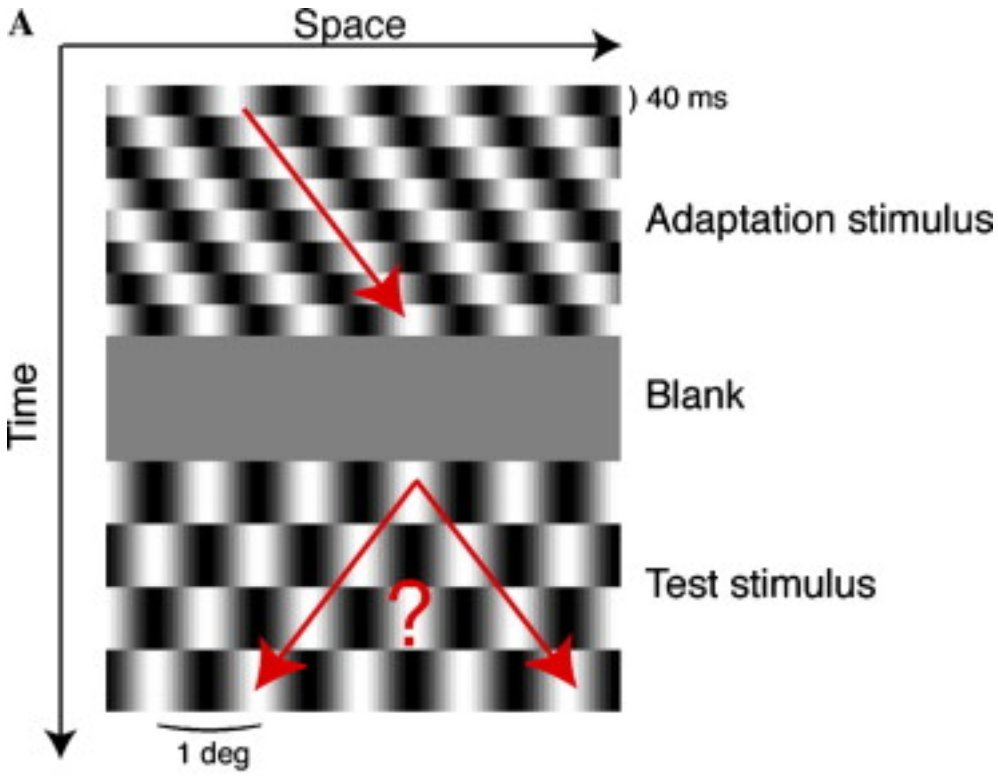

Figure 1.5 Top: Space-time plot of a stimulus with directional apparent motion created by 90° phase shifts. Here, an example of 320ms adaptation is shown. Bottom: Space-time plot of a stimulus with ambiguous apparent motion created by 180° phase shifts. Taken from Kanai and Verstraten (2005).

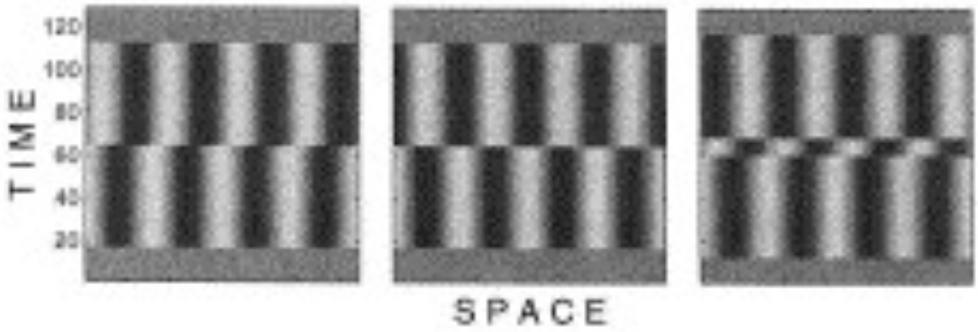

Figure 1.6 Left: Space-time plot of a sine-wave grating whose phase is shifted 90° to the left at time, therefore creating rightward motion. Middle: Space-time plot of a sine-wave grating whose phase is shifted 180°, therefore creating ambiguous motion. Right: Space-time plot of a motion priming sequence: an ambiguous 180° step follows a perceptually unambiguous 90° rightward step. Time increases vertically upwards. Taken from Pinkus and Pantle (1997).

Pinkus and Pantle (1997) presented participants with 3 frames that created these kinds of motion. Specifically, frame 1 was presented for 1530ms and had a phase set to an arbitrary unit of 0. This was followed by frame 2, which was presented for a duration of between 192 – 1530ms and had its phase shifted by ±90°, followed by a final frame that had its phase shifted by 180° compared to frame 2 (see Figure 1.6). Participants were asked to report the direction of motion that they saw. The idea being that the directional motion created by frame 2 would prime the direction that the ambiguous frame of motion was perceived as moving in. They found similar results to studies using stationary stimuli, in that shorter durations of the directional motion step led to VMP, whilst longer durations led to VMP being around chance level (in this case 50%).

Later work by Kanai and Verstraten (2005) used a very similar method of creating ambiguous and directional motion to investigate VMP and motion aftereffects. The motion aftereffect (MAE), also known as the waterfall illusion, is a phenomenon where exposure to a moving visual stimulus leads to the perception of motion in the opposite direction when observing a stationary scene immediately afterward (an effect reviewed by Anstis, Verstraten & Mather, 1998). This effect occurs after observing motion for a period of time, typically in the order of seconds. When the viewer then looks at a stationary object, it appears to move in the opposite direction of the original motion. Like adaptation to stimuli such as the Necker cube, the MAE typically arises from adaptation periods spanning several seconds to tens of seconds. Neurons that display similar temporal activation patterns (i.e., the neurons display reduced responsiveness after the subject has been presented with a motion stimulus in a given direction for an extended period of time) are considered to underlie the MAE, as indicated by studies from Barlow and Hill (1963), Hammond et al. (1988a, 1988b), and Kohn and Movshon (2003). In contrast to this slower adaptation process associated with the MAE, much faster forms of adaptation, occurring within hundreds of milliseconds, have also been reported in electrophysiological studies. For example, when the same stimulus is presented twice in rapid succession, the neural response to the second stimulus is significantly reduced (Nelson, 1991; Chance et al., 1998; Finlayson & Cynader, 1995; Stratford et al., 1996). This rapid adaptation phenomenon is also observed in area MT (middle temporal area of the extrastriate visual cortex; Priebe et al., 2002), which is closely linked to the subjective experience of visual motion (Newsome et al., 1989; Zeki et al., 1993; Logothetis & Schall, 1989; Bradley et al., 1998). In apparent contrast, some studies have found that brief, direct electrical stimulation of neurons can lead to their responsiveness increasing during subsequent stimulation (e.g., Hempel et al, 2000). In their research, Kanai and Verstraten (2005) aimed to examine this apparent conflict and note that behavioural results of studies into adaptation and priming like the ones described in the paragraphs above also require clarification given that very subtle changes in trial design (i.e., adaptor stimulus presentation duration and ISI) produce opposite effects. Therefore, Kanai and Verstraten (2005) presented participants with directional motion created from phase-shifting sine wave gratings by ±90° each frame (i.e., a direction motion prime), followed by an ISI and then ambiguous motion created by phase-shifting the gratings by 180° each frame (i.e., an ambiguous target stimulus). They systematically and orthogonally manipulated four directional motion presentation durations (80, 160, 320 and 640ms) and 5 ISI durations (40, 120, 480, 1000 and 2000ms). At the end of each trial, participants were asked to report whether they saw the two stimuli moving in the same or different directions. Their results showed that brief presentation times (i.e., 80ms) led to VMP effects, providing support for the findings of Pinkus and Pantle (1997). Slightly longer presentation times, however (from 320ms upwards) led to an effect similar to the MAE, or in other words adaptation effects, but on a much faster timescale. This led them to use the term “rapid motion aftereffect” or rMAE to describe this effect. Additionally, in a separate experiment in the paper, they presented participants with pairs of ambiguous motion stimuli, with no directional motion at all in the trials. Here, they found effects similar to VMP in that participants were more likely to report that they perceived motion in the same direction, which they call “perceptual sensitisation” (PS), which increased with increasing ISI. In terms of the time courses of each of these effects, VMP and rMAE occur directly after participants have been exposed to directional motion and then decay very quickly over the course of around a second. PS, however, has a longer time course as it develops over the course of several seconds, hence the requirement for longer ISIs to induce it. The authors suggest that their findings suggest that several mechanisms, at different processing levels, are involved in resolving ambiguous motion stimuli. Later work by Takeuchi et al. (2008) served to elaborate on the VMP and rMAE, using a similar design and manipulating the retinal luminance levels and speed of the adapting stimulus. These authors found that VMP was almost completely eradicated under low retinal luminance levels, whereas the rMAE was still present and unaffected by retinal luminance levels. Further, when the adapting stimulus was at a lower speed, the effects of VMP were greater whereas at higher speeds only the effects of rMAE were present. The authors suggest that their findings support the view that several mechanisms are at work during the resolution of ambiguous motion stimuli. Specifically, they suggest that the rMAE may be induced by lower order, directionally selective mechanisms whilst VMP may be elicited by a higher-order mechanisms, in support for the conclusions from studies such as Long and Moran (2007) and Goolkasian and Woodberry (2010).

1.2.1.1 ISI effects

In addition to the specific cases of investigations into adaptation described in the section above, the effects of presentation mode and ISI have been investigated more generally in the literature. Specifically, ambiguous stimuli can be presented continuously, or intermittently (i.e., with an ISI between presentations; discussed fully in Chapter 2). Long and Toppino (2004) recognised presentation mode as a bottom-up factor influencing the reversal rates of ambiguous stimuli (e.g., Orbach et al., 1963; Kornmeier et al., 2002; Leopold et al., 2002; Maier et al., 2003; Noest et al., 2007). A method to investigate perceptual reversals is through intermittent or discontinuous stimulus presentation, which involves showing the ambiguous figure briefly with intervals of blank screen, instead of a constant, continuous display. This approach was first used with the Necker cube by Orbach et al. (1963), demonstrating that the rate of perception reversal of such figures changes depending on the ISI. Orbach et al. (1963) interpreted their results as evidence of bottom-up influence, involving the adaptation and recovery of neural groups representing competing perceptions. Later work by Kornmeier et al. (2007) extended these findings by systematically exploring how varying ISIs affect reversal rates. The authors showed that random ISI variations in the same session significantly impacted reversal rates. The authors presented their participants with a Necker Lattice (a Necker Cube variant) for 800ms, followed by a randomly selected ISI ranging from 14 to 390ms (Kornmeier et al., 2007). Their findings, together with those from other studies, indicated that short ISIs (up to 400ms), in contrast to continuous presentation, tend to increase reversal rates (Kornmeier et al., 2007; Kornmeier et al., 2002, Orbach et al., 1963) while reducing the duration of perceptual stability. Conversely, with longer ISIs (> 400ms), reversal rates significantly decline, even to zero (complete stabilisation; Leopold et al., 2002, Maier et al., 2003, Sterzer and Rees, 2008; see Figure 1.7).

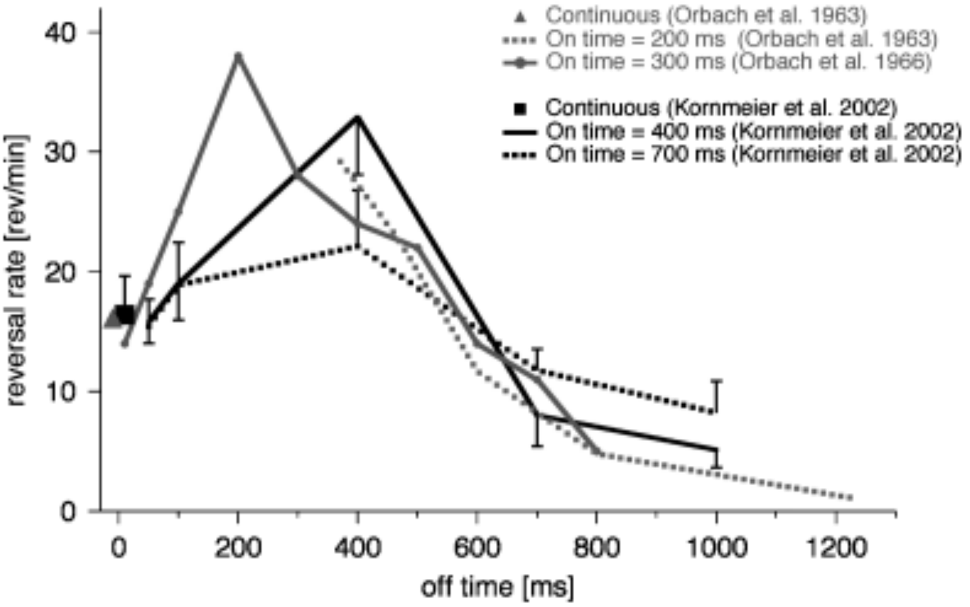

Figure 1.7 Reversal rate (in number of reversals per minute) when viewing the Necker stimulus, as a function of ISI ('off time' – shown on the x axis). Reversal rates when the Necker cube is viewed continuously are represented by the triangle marker (from Orbach et al., 1963) and the square marker (from Kornmeier et al., 2002). Figure taken from Kornmeier and Bach (2007).

Kornmeier et al. (2007) described the relationship between the stimulus presentation duration and ISI as an inverted-U-shape (Figure 1.6; Kornmeier et al., 2007). Kornmeier et al. (2007) viewed the ascending and descending parts of the function as indicative of two distinct neural processes driving perceptual reversals. Kornmeier and Bach (2012) suggested that for short ISIs (up to 400ms), the processes involved might resemble those involved in continuous stimulus presentation, while for ISIs over 400ms, the reversal dynamics are more akin to a perceptual decision regarding the representation of the ambiguous figure as it appears and disappears on a blank screen. The authors suggest that this implies that for longer ISIs (> 400ms), perceptual reversals may be distinct percepts rather than spontaneous occurrences.

1.2.1.1.1 Eye movement and blink effects

The effect of eye movements, blinks, and fixation point on perceptual reversals has been a focus of several studies. Such studies have revealed that eye fixation position can bias the interpretation of ambiguous figures such as the Necker cube (Einhäuser, Martin & König, 2004; Toppino, 2003) and others (Ruggieri & Fernandez, 1994). Fixation location has also been shown to modulate reversal rate (Liu, Tzeng, Hung, Tseng & Juan, 2012) and dominant perception (Peterson and Gibson, 1991) of ambiguous stimuli. Although eye-movements are indeed conducive to perceptual reversals, they are not necessary. Removal of eye-movement effects using afterimages has shown that reversals persist in the absence of image scanning via eye-movements (Gregory, 1970). Other eye movements such as blinks and saccades have also been associated with reversals during continuous presentation of ambiguous stimuli. Nakatani, Orlandi and van Leeuwen (2011) for example, found a peak in blinking rate around 1000ms prior to a reversal of the Necker cube, and a peak in saccade rate at around 150ms pre-reversal. Further analysis of these saccades revealed that they were predictive of the upcoming percept, in that rightward saccades were associated with the upward perception of the Necker cube, and the opposite trend was demonstrated for leftward saccades.

1.2.2 Top-down approaches

In contrast to the bottom-up approach described above, a second stance that emphasises the role of top-down, cognitive influence over perceptual reversals has also been adopted by researchers. This approach rejects the idea of automatic, stimulus-driven processes in favour of an active appraisal and modulation of sensory information based on cognitive resources such as attention, expectation and learning (Brascamp, Blake & Knapen, 2015; de Graaf et al., 2011; Pitts et al., 2009; Gregory, 1974; Rock et al., 1994; Horlitz & O'Leary, 1993; Leopold & Logothetis, 1999; Lumer & Rees, 1999). Findings supportive of the involvement of top-down influences on ambiguous stimulus perception will be discussed in the sections that follow.

1.2.2.1 Priming effects

As discussed briefly in the sections above, some researchers have suggested that observations of priming effects in the available behavioural research could reflect top-down mechanisms. Long and Moran (2007) varied the size of their Necker cube stimulus whilst also varying the ISI between successive stimulus presentations. When the ISI was short (1s), the authors found that adaptation effects were significantly pronounced, but only when the size of the adapting and test stimuli were matched (i.e., small adapting stimulus paired with small test stimulus). When a delay of 100s was introduced between the adapting and test stimuli, only priming effects were observed, regardless of the adapting and test stimuli being matched in size. The authors suggest that this is because adaptation effects are governed by bottom-up mechanisms and are therefore sensitive to local stimulus changes. Whereas, they claim, priming effects are underpinned by higher order mechanisms and so are not altered by local stimulus changes such as a mismatch in the size of the adapting and test stimuli. This is also supported by the result that the priming, but not the adaptation, effects persisted even with a delay period of 100s. The authors suggest that this is indicative of the involvement of localised neural networks during the adaptation phase, but globalised, top-down channels in the priming phase.

Other studies too, have supported the idea that behavioural priming effects are indicative of top-down mechanisms being involved in multi-stable perception. Takeuchi et al. (2011) found that when they lowered retinal luminance to scotopic levels whilst presenting their participants with directional adapting motion followed by ambiguous test motion (in the paradigm described in Section 1.2.1.1, above), priming effects were almost completely eradicated. When retinal luminance was high (photopic vision) priming effects were observed only when the speed of the motion stimuli was low (at 2 degrees per second). Otherwise, at higher velocities, adaptation effects dominated. The authors interpreted their findings considering other research in the area and suggested that the mechanism responsible for priming could be a higher-order motion system, such as a feature-tracking mechanism or a third-order motion mechanism. Whereas, they suggest, that the mechanism responsible for adaptation effects could be a lower-order, directionally selective motion mechanism, which is known to be sensitive to higher velocities (e.g., Bowns, 2002; Derrington et al., 2004). They suggest that the dominance of the priming effect at lower velocities supports this idea, since the higher-order motion system has been shown to be insensitive to higher velocities (e.g., Lu & Sperling, 1995).

Further research supporting the idea that priming effects are indicative of higher-order involvement in perception has come from studies using adapting stimuli other than biased versions of an ambiguous stimulus. For instance, Balcetis and Dale (2007) asked their participants to read a paragraph that was loosely related to one of two possible interpretations of the ambiguous woman's face/saxophone player figure (see Figure 1.8, below). Their participants were randomly assigned to either the music or the pornography condition. Those randomly assigned to the music condition read some text discussing the legality of the music-file-sharing service, Napster. Critically, the paragraph discussed legal issues for artists and producers and concerns for distribution of royalties but never mentioned specific artists (e.g., the saxophone player Kenny G) or groups of musicians (e.g., saxophone players). Those randomly assigned to the pornography condition read arguments concerning the legality of pornography and implications for basic human rights. Again, critically, participants were never exposed to terms referring to groups of people (e.g., women), or gender (e.g., female). Therefore, the participants were not provided with any direct linguistic cues on how to interpret the upcoming visual stimulus. After they had read the paragraph, they were presented with the ambiguous woman's face/saxophone player figure as shown in Figure 1.8. The results showed that those who were in the music condition were significantly more likely to report the ambiguous figure as a saxophone player, and those who were assigned to the pornography condition were more likely to perceive the figure as a woman's face (Balcetis & Dale, 2007). These authors interpreted their results as an indication that the perception of an ambiguous stimulus can be influenced by conceptual primes. This is supportive of a higher-order mechanism being involved in priming effects as the studies described previously have used perceptual primes -i.e., primes that are perceptually linked to a particular interpretation of the ambiguous stimulus (for instance, pre-exposure to a biased top-right face as frontal version of the Necker cube as a perceptual prime for an upcoming Necker cube stimulus). This study, however, extends these findings by successfully priming participants using only conceptual primes (i.e., paragraphs that are conceptually related to one or the other percept of the ambiguous test stimulus). Conceptual primes have also been successfully used in similar ways in other studies, demonstrating further support for the top-down involvement underpinning priming effects (e.g., Feist & Gentner, 2007; Goolkasian & Woodberry, 2010).

Figure 1.8 The ambiguous saxophone player/woman's face stimulus used by Balcetis and Dale (2007). The figure can be seen as a saxophone player (made up of the black regions on the left hand side of the figure, against a white background) or a woman's face (white regions on the right hand side of the image are the highlights of the face with the black regions being perceived as shadows).

Manipulations of more cognitive aspects of binocular rivalry stimuli have also been shown to produce pronounced effects on predominance, further supporting the role of top-down mechanisms in multi-stable perception. Stimuli that are recognisable (Yu & Blake, 1992), associated with reward (Mudrik, Deouell & Lamy, 2011; Dunning & Balcetis, 2013), emotionally arousing (Alpers & Pauli, 2006; Alpers & Gerdes, 2007; Sheth & Pham, 2008) or incongruent in some way (Mudrik, Deouell & Lamy, 2011) tend to dominate over alternatives. Kovacs, Papathomas, Yang and Fehér (1996) showed that coherent patterns tend to dominate conscious perception; even when those patterns are presented to different eyes, the brain can integrate coherent areas to form one percept. Integration of multiple sensory modalities has also been shown to be involved in predominance. Lunghi, Morone and Alais (2014) used dichoptic visual stimuli modulated at different temporal frequencies, then added sounds or vibrations congruent with one or the other visual temporal frequency to show that congruent auditory and/or tactile stimulation facilitates perceptual dominance. Higher-order cognitive processes such as mental imagery (Pearson, Clifford & Tong, 2008) and attention (Chong, Tadin & Blake, 2005) have also been shown to modulate levels of perceptual dominance.

1.2.2.2 Attention, and volitional control

In terms of ambiguous perception, selective attention can be defined as the enhancement of a particular percept and suppression of the alternative percept. Research has provided considerable evidence that shifts in selective attention can bias perceptual reversals of ambiguous stimuli (Tsal & Kolbet, 1985; Horlitz & O'Leary, 1993). For example, shifting attention to certain "critical" features of ambiguous figures (such as drawing attention to the “nose” in Rubin’s faces/vase) is known to significantly bias perceptual outcome (Georgiades and Harris, 1997). However, selective attention has not been found to modulate dominant percepts in binocular rivalry paradigms to the same extent (Meng & Tong, 2004).

Attention enhances the ability to process and differentiate between objects, especially in crowded settings. This improvement is attributed to changes in sensory processing, leading to better sensitivity, as shown in studies by Jehee et al. (2011) and Pratte et al. (2013). Such enhancements are a result of focusing attention on specific stimuli. Neurophysiological research, involving both single-unit studies in monkeys and functional magnetic resonance imaging (fMRI) in humans, indicates that attention directed spatially towards a target stimulus is linked to increased activation of neurons across the extrastriate visual cortex and subcortical areas (Treue & Maunsell, 1996; Kastner, 1998).

Leopold and Logthetis (1999) describe attention as akin to multi-stable perception, being an active and largely voluntarily controllable process. Similar to how voluntary actions can produce a motor response (like moving a finger), it was found that these voluntary processes can also determine the focus of attention on specific stimuli or features. In multi-stable perception, the observer's intention plays a crucial role in how the perception shifts (Peterson & Hochberg, 1983). Although voluntary control in binocular rivalry is less pronounced compared to other multi-stable stimuli (George, 1936), studies still highlight the significant influence of the observer's intention on perceptual switching and dominance (Meredith & Meredith, 1962). For example, individuals unfamiliar with multi-stable stimuli can triple their rate of perceptual switches by intentionally trying to perceive rapid fluctuations rather than slower ones. Additionally, the ability to control the rate of perceptual reversals improves with practice and exposure. In a study by Lack (1978), it was observed that the ability to control the rate of rivalry switches enhanced after a training period spanning ten days.

In the context of ambiguous images like the Necker Cube (Necker, 1832), Pitts, Gavin, and Nerger (2008) also explored the role of attention as a key factor in multi-stable perception by incorporating voluntary control into their experimental design. The authors asked their participants to control their reversal rate whilst viewing ambiguous stimuli and found that they were able to do so. This approach aligns with several past behavioural studies which have demonstrated that an observer's intentions can influence perceptual reversals. For instance, van Ee et al. (2005) investigated how voluntary control affects the reversal rates of various bistable stimuli. In their study, participants were instructed in one scenario to increase the frequency of reversals ('speed up reversals') and in another to maintain a single perception for as long as possible ('hold one percept stable for as long as possible') (van Ee et al., 2005). Van Ee et al. (2005) assigned their participants three different tasks involving control exertion. The first task required participants to simply observe the stimulus without attempting to control the reversal rate. In the second task, participants were asked to observe passively for six minutes and then maintain the same interpretation of the stimulus for three minutes. In the third task, they were to observe passively for three minutes before trying to maximise the reversal rate. The results indicated that participants could voluntarily control and modify the reversal rate of the Necker Cube and other ambiguous stimuli, increasing or decreasing it as per the instructions (Washburn & Gillette, 1933). Consistent with earlier studies (e.g., Long and Toppino, 2004; Meng and Tong, 2004; Rock et al., 1994; Struber and Stadler, 1999; Toppino, 2003; van Ee et al., 2005; Windmann et al., 2006), these findings imply that it is possible to exert some level of voluntary control over the reversal rates of the Necker cube.

Previous research has highlighted that involuntary changes in selective attention play a role in the occurrence of perceptual reversals (Pitts et al., 2007). Selective attention increases the prominence of specific visual attributes such as location, colour, movement, or shape (Sperling, Reeves, Blaser, Lu, & Weichselgartner, 2001). Some theories propose that focusing attention on a particular location can impact attention to features in that area, determining which features are noticed and merged into a unified perception (Anllo-Vento and Hillyard, 1996; Treisman and Gelade, 1980). Regarding ambiguous figures, paying attention to certain parts of the figure may lead to a preferential processing of features at those locations, influencing the visual system's interpretation of these features as 'nearest' and thereby swaying the perception of the bistable image. Recent fMRI studies provide evidence that attention aids in the voluntary reversal of bistable stimuli (Slotnick & Yantis, 2005). In their research, Slotnick and Yantis (2005) compared brain activity during voluntary Necker cube reversals with that during basic spatial (left-right) attention shifts. They found similar patterns of neural activation in both tasks, including brief increases in activity in the superior parietal lobule and intraparietal sulcus for both voluntary spatial attention shifts and voluntary Necker cube reversals. Slotnick and Yantis (2005), along with Pitts et al. (2008), propose that when observers try to control the rate of perceptual reversals of bistable stimuli, they are engaging an automatic, exploratory mechanism for refreshing perception. This results in a shift in the features that are the focus of attention, leading to a re-evaluation of their perceptual interpretation of the stimulus.

Pitts et al. (2008) align their findings with the 'environment exploration' theory proposed by Leopold and Logothetis (1999), which is based on a broad range of perceptual rivalry research. According to this theory, perceptual reversals result from a high-level "exploratory" process that directs selective attention, prompting lower-level perceptual systems to periodically renew or "refresh". Pitts et al. (2008) describe this exploratory mechanism as not being solely sensory or motor in nature but as serving the ultimate goal of interacting with and responding to information in the environment. This theory suggests that constant reorganisation and updating of perceptual processing aids in accurately interpreting visual stimuli. In this framework, while eye movements are usually the primary means of directing visual attention and quickly resolving ambiguity in everyday environments, multi-stable perception experiments typically require observers to fixate at a central point. This does not negate the 'environment exploration' theory; covert attention (attention without eye movements) can still be influenced by this central exploratory process.

The necessity of ongoing exploration in the face of ambiguous stimuli, leading to regular perceptual reversals, is supported by van Ee et al. (2005). In their study, they examined the role of eye movements and blinks in experiencing the slant rivalry stimulus, an ambiguous visual stimulus. They found that neither microsaccades (small involuntary eye movements) nor saccades (rapid eye movements altering fixation points) are essential for experiencing a voluntary perceptual reversal. Interestingly, they observed that blinks and saccades, but not micro-saccades, are suppressed when participants induce a perceptual reversal.

Although Ellis and Stark (1978) noted a correlation between eye position and reversals of the Necker Cube, other studies, like those of Washburn & Gillette (1933) and Pritchard (1958), indicate that voluntary control over reversal rates is not exclusively linked to eye movements. For instance, Washburn and Gillette (1933) studied the effect of voluntary control on the afterimage of the Necker Cube, demonstrating that participants could induce and maintain a specific orientation of these afterimages without involving eye movements. While the central exploration theory operates largely unconsciously and automatically (Leopold & Logothetis, 1999), voluntary control over bistable perceptual reversals might still function through this same mechanism. For example, Leopold and Logothetis (1999) highlight similarities between controlling bistable perception and other voluntary actions, particularly in terms of improvement through practice and learning.

Thus, mechanisms of selective attention and multi-stability might indeed be closely related. However, there are some key differences that suggest they are not one and the same. First, voluntary control in orienting attention is generally greater than in multi-stable vision. It is a lot easier to voluntarily control what you are attending to but it is not as easy to voluntarily control the rate of reversals of an ambiguous figure (Leopold & Logothetis, 1999; Meng & Tong, 2004). This is supported by findings that reveal that although perceptual reversals can be slowed down or decreased in number, they cannot be entirely stopped. Perceptual reversals continue regardless of the intent of the participant. This was reported in Slotnick and Yantis (2005)'s study where in the hold condition (maintaining one perceptual interpretation of the ambiguous figure through voluntary control) participants still experienced involuntary switches in their interpretation of the ambiguous image. This was not the case when participants voluntarily controlled shifts in selective attention. Their behavioural results showed that the task difficulty across both conditions (voluntary control of reversals vs. voluntary control of attention shifts) was the same. However, this is due to the design of the experiment whereby the researchers, based on the results of a pre-training session, calibrated the features of the stimulus to yield an accuracy value that would ensure that task difficulty was well matched between conditions. In multi-stability, top-down influences are not limited simply to enhancing the visual processing of a particular object (or features of that object) or spatial location. Instead, top-down attentional influences spur organizational mechanisms to change perception completely, and thereby possibly feed into the neural fatigue mechanism described previously to enhance one representation over the other (i.e., by shifting the balance between mutually exclusive neural representations in the visual system itself; Mathes et al., 2006; Slotnick & Yantis, 2005). Finally, attentional shifts can proceed with a speed that is considerably faster than even the most rapid perceptual reversals. Attention can shift as fast as several times per second (Duncan, Ward, & Shapiro, 1994; Egeth & Yantis, 1997), whereas the transition between rivalling percepts may take seconds or even tens of seconds to complete (Pastukhov & Braun, 2007).

1.2.2.3 Working memory load

Top-down influences, shaped by our goals and expectations, play a significant role in how we perceive visual stimuli (Kumar, Soto & Humphreys, 2009). Duncan and Humphreys (1989) proposed that an 'attentional template' (a representation of a visual stimulus' characteristics in working memory -WM) biases our perception towards matching stimuli (Downing, 2000; Hodsoll & Humphreys, 2001; see Soto et al., 2008 for review). The link between visual selection and WM has been extensively studied, with neuroimaging revealing overlapping neural networks for attention and WM, especially around the intraparietal sulcus and frontal eye fields (McCarthy, 1995; LaBar et al., 1999; Pollmann & von Cramon, 2000). Behavioural studies have also shown WM's crucial role in suppressing distractions and maintaining focus on task-relevant information (de Fockert, 2001; Woodman et al., 2007; Woodman & Luck, 2004).

Recent research into ambiguous figure perception has indicated WM's involvement in influencing the rate and choice of perceptual reversals. Paffen, Alais, and Verstraten (2006) observed a reduction in binocular rivalry reversal rates when WM and attentional resources were manipulated through a motion-detection task. They presented a binocular rivalry stimulus alongside a concurrent task, manipulating WM by varying the task's difficulty. Their findings showed that while WM demands reduced reversal rates, they did not eliminate them, suggesting that perceptual ambiguity involves more than just attentional resource allocation. Intaite, Koivisto, and Castelo-Branco (2014) conducted an experiment with a secondary WM task (mental arithmetic) known to deplete attentional resources. Their study involved presenting ambiguous Necker Cube stimuli alongside WM tasks, finding that increased WM load delayed the latency and rate of reversals, yet reversals still occurred. This implies a link between mechanisms responsible for WM maintenance and perceptual reversals under top-down control.

Further studies have also shown that secondary tasks requiring WM load can extend the time for reporting perceptual reversals and decrease their rate

(Reisberg & O'Shaughnessy, 1984; Wallace, 1986). Intaite, Duarte, and Castelo-Branco (2016) expanded on this research using fMRI to investigate the neural responses to perceptual reversals under different WM loads. They observed overlapping activation in brain regions during perceptual reversals and the frontoparietal attention network, a finding consistent with earlier studies (Knapen et al., 2011; Lumer et al., 1998; Sterzer and Kleinschmidt, 2007; Weilnhammer et al., 2013). Previous research by Sterzer and Rees (2008) also reported similar activations in visual cortex and prefrontal and parietal regions in response to percept-specific signals and voluntary engagement in facial WM tasks (Courtney et al., 1997; Haxby et al., 2000).

Intaite, Duarte, and Castelo-Branco (2016) also found that the right posterior Superior Parietal Lobule (pSPL) showed different responses to perceptual reversals under varying load levels and was more responsive to perceptual reversals than to stimulus changes. Previous studies have indicated the involvement of the Superior Parietal Lobule (SPL) in the perception of reversals (Baker et al., 2015; Carmel et al., 2010; Kanai et al., 2010). Stimulation of the right anterior SPL with transcranial magnetic stimulation (TMS) influenced the rate of reported reversals, suggesting a role in WM manipulation of stimulus content (Champod and Petrides, 2007; Carmel et al., 2010).

1.2.2.4 Past experience

Adaptation and priming studies, such as those by Long et al. (1992), support the notion that our perceptual system utilises past perceptual experiences, across various time scales, to clarify and interpret ambiguous information (Leopold et al., 2002; Maier et al., 2003; Pearson and Brascamp, 2008), as well as in its neural processing (Kornmeier and Bach, 2006; de Jong et al., 2012b; Pitts and Britz, 2011). It has been shown that ambiguous stimuli often do not lead to perceptual reversals until participants are informed about the possibility of such reversals, indicating that this understanding is 'learned' (Girgus, Rock & Egatz, 1977; Leopold et al., 2002; Maier et al., 2003; Orbach et al., 1963; Pearson and Brascamp, 2008; Rock & Mitchener, 1992). For example, knowing the different interpretations of Rubin's Face-Vase illusion (faces or vase) makes it easier, and sometimes only possible, for an individual to perceive these interpretations and experience perceptual switches. This dependence on past experiences is not limited to short-term exposure but extends over longer durations (Pastukhov and Braun, 2008; Pearson and Brascamp, 2008; Brascamp et al., 2009; de Jong et al., 2012a). Some research has shown that naïve viewers do not experience perceptual reversals at all (Rock, Hall & Davies, 1994). Additionally, informing or prompting viewers has been shown to induce interpretations that are otherwise unlikely. For example, when two identical ambiguous figures are presented simultaneously, they are not usually interpreted differently. However, when viewers are provided with a simple verbal prompt, this “dual interpretation” becomes significantly more likely (Jensen & Matthewsen, 2011). This, along with the fact that reversals can be volitionally controlled at all (Kornmeier, Hein & Bach, 2009), provide evidence for the existence of top-down effects on perception.

Furthermore, studies have indicated that instrumental learning shapes perceptual interpretation of an ambiguous stimulus. Fleming et al. (2010) found that participants were more likely to report seeing a house in ambiguous face-house images when informed they would lose money for perceiving it as a face. Additionally, Wilbertz et al. (2014) used a binocular rivalry stimulus in a study where participants could win or lose money based on their reported perception. They observed increased stable durations of the percept associated with reward and a decrease for the punishment percept. The authors suggest that their results indicate that perceptual inference is an adaptive process influenced by its consequences, akin to instrumental or reinforcement learning.

Murphy et al. (2014), using a rotating Necker Cube stimulus, found that long-lasting perceptual biases, specific to certain retinal locations, emerged from the outset and remained stable for up to 40 minutes. They conducted experiments to see if long-term associative learning emerges naturally from repeated short-term perceptual memory without explicit training. They found that biases can develop during training but also emerge without training, suggesting associative learning results from associations between different perceptual interpretations of the ambiguous stimulus and retinal locations, rather than the physical stimulus itself. In a subsequent experiment, they examined if periods of spontaneous reversals affected recently acquired perceptual biases in a retinotopically specific manner. The continuous presentation of an ambiguous stimulus abolished trained perceptual biases at the location where it was presented but not at a location left blank. Participants reported significant biases nearly equally divided between two interpretations of the stimulus at the retinal location where the continuous stimulus was presented. Meanwhile, the blank retinotopic location retained its bias from the unambiguous training experiment. During intermittent presentation, participants reported perceiving the interpretation most commonly experienced during the continuous presentation at its corresponding retinal location. These findings suggest that perceptual biases naturally arise and primarily reflect the brain's tendency to favour recent perceptual interpretations at a specific retinal location. It indicates that an association between retinal location and perceptual state, rather than with a physical stimulus, is sufficient to generate long-term biases in perceptual organization. This is based on an individual's past experiences with the ambiguous stimulus (de Jong et al., 2014), emphasising the influence of associative learning and the brain's adaptive processes in perception.

1.2.3 Integrated approaches

Given that strong evidence exists for the influence of both bottom-up and top-down processes on multi-stable perception, a further approach has been taken by many researchers, which adopts an integrative view. For instance, Long and Toppino, along with their colleagues, conducted a series of experiments (1981, 2004, 1983, 1992; reviewed in Long & Toppino, 2005) providing robust evidence that perceptual reversals are affected by both bottom-up processes (like adaptation or fatigue) and top-down processes (such as learning, attention, and intention). In one such study (Toppino and Long, 1987, as mentioned in Chapter 1, section 1.2.1.3), they discovered that the rate of reversal was high when the test stimulus resembled the adapting stimulus, but low when there were changes in the location and size of the test stimulus compared to the adapting one. This indicates that reversals of rotating Necker Cube stimuli are influenced by fatigue and recovery in multiple independent neural channels, with changes in size or location affecting the neural channels involved and nullifying the adaptation effect. Towards the end of each experimental session, Toppino and Long (1987) observed a notable increase in reversal rates. Long et al. (1983) had earlier examined this effect, having participants view rotating Necker Cubes over four weeks. They found a consistent increase in reversal rates over this period and even within individual sessions. This suggested that besides neural fatigue and adaptation, a top-down learning process might be occurring during sustained perceptual reversals of the same stimulus.

Toppino (2003) further explored the simultaneous influence of top-down and bottom-up processes, particularly the role of intentional control on reversal rates. In their first experiment, participants focused on specific points of a Necker Cube and either tried to maintain one interpretation (hold condition) or viewed it passively (no-hold condition). The findings showed that the time spent perceiving a particular orientation was influenced by where participants directed their gaze, and significantly, by the hold vs. no-hold instruction. This implies that participants could exert voluntary control over their perception of the cube. However, there was no interaction between the fixation location and hold conditions, suggesting that intentional control over the perception of an ambiguous image could be achieved through direct, top-down activation of the desired interpretation. In a second experiment, Toppino (2003) varied the cube size but found no effect of size on hold conditions. The results from both experiments indicated that factors other than focal-feature processing contribute to the effect of intentional control. Yet, Toppino (2003) also observed that reversals couldn't be completely prevented regardless of hold condition, hinting at the presence of uncontrollable processes like neural fatigue. This aligns with previous research on intentional control (Babich & Standing, 1981; Liebert & Burk, 1985; Suzuki & Peterson, 2000) and suggests an additive effect of adaptation and priming in influencing the perception of the Necker Cube.

Based on findings like those described above, Long and Toppino (2004) offer an integrative model of perception, which is grounded in the idea that researchers may have accessed different neural processes via the use of different methodologies, leading to the apparent discrepancies between top-down and bottom-up theories. The authors emphasise the dissociation between ambiguity and reversibility, arguing that a physical stimulus can be ambiguous in that its features give rise to more than one percept, but the act of reversal depends on whether the viewer has had past experience and knowledge of the figure (e.g., Rock, Hall and Davies, 1994). Therefore, Long and Toppino (2004) put forward a hierarchical hybrid model of perception. At the lower levels of the model, bottom-up feature analysis occurs largely automatically. At higher levels of the model, top-down information is fed backward from non-sensory areas and coupled with the stimulus-derived information, leading to a particular perceptual representation. Reversals, therefore, can be initiated by bottom-up adaptive processes or top-down cognitive processes.

Kornmeier et al. (2009) also investigated the combined influence of different mechanisms on perceptual reversals, demonstrating an additive effect of voluntary control (a top-down mechanism) and discontinuous stimulus presentation (a bottom-up mechanism) on the rate of reversals. In their study, they conducted two experiments. The first was akin to Toppino's (2003) initial experiment but without designated focal points; participants simply focused on the centre of the cube. They replicated Toppino's (2003) results. In their second experiment, Kornmeier et al. (2009) used the same experimental setup as their first, but this time presented the ambiguous stimulus intermittently, varying the length of the Inter-Stimulus Interval and the stimulus presentation time. Their findings showed that combining voluntary control with discontinuous presentation led to an increased rate of reversals compared to continuous presentation, demonstrating the full additivity of these two effects.

This research adds to the ongoing debate about the roles of bottom-up and top-down processes in the perception of ambiguous figures, a discussion involving researchers like Hochberg & Peterson (1987), Leopold & Logothetis (1999), and Long et al. (1983), and the extent to which their effects are additive, as explored by Kornmeier et al. (2009), Long & Moran (2007), Long & Toppino (2004), and Toppino (2003). Recently, there has been a shift towards understanding multistable perception as a result of changes in attractor states within neural networks, as proposed by Lehky (1988), Noest et al. (2007), and Wilson (2003). Kornmeier and Bach (2012) used this model to explore the integrative relationship between top-down and bottom-up processes.

Kornmeier and Bach (2012) offer a second integrative theory of multi-stable perception. They propose that multi-stable perception is underpinned by cycles of percept destabilisation and re-stabilisation. According to this model, an ambiguous stimulus gives rise to an initial percept which undergoes a slow destabilisation process. The rate of this is not fixed but can be modulated by both top-down and bottom-up influences. This continues until a point of maximum instability is reached, following which a fast (40-60ms) percept re-stabilisation occurs. These stages are discussed in more detail in the subsections below.

In their theoretical perspective, Kornmeier and Bach (2012) apply concepts from nonlinear dynamics to explain the mechanisms behind perceptual reversals, as discussed by Braun and Mattia (2010). They model object representations in the brain as attractors – essentially, neural representations of physical stimuli – and consider the depth of these attractors as indicative of the stability of the current neural activity pattern that corresponds to a particular percept. For example, both the Necker cube and Rubin’s Face Vase would each have two attractors corresponding to their different interpretations. A specific interpretation is perceived when its representation is activated, meaning the perceptual system's state aligns with the attractor linked to that interpretation. A perceptual reversal occurs when the perceptual system moves out of one attractor state and enters the alternative one, particularly if the depth of the attractor is shallow enough to allow such a transition. Kornmeier and Bach (2012) propose that during prolonged observation of an ambiguous stimulus (over several minutes), a temporarily stable percept becomes destabilised (shifts from one percept to another) gradually and steadily. Once the percept becomes destabilised, a rapid re-stabilisation (disambiguation) occurs, leading to a change in the interpretation of the ambiguous stimulus. This integrative theory suggests that various bottom-up and top-down processes either collaborate or compete, ultimately leading to a perceptual reversal.

Support for this theory comes from electrophysiological evidence. For instance, the Reversal Positivity (RP) is thought to be a marker of the disambiguation/restabilisation process. Additionally, other correlates identified in both pre-stimulus and post-stimulus periods are believed to be connected to the destabilization and disambiguation/restabilisation mechanisms of this theory. These physiological underpinnings will be discussed in more detail in section 1.3.

Kornmeier and Bach (2012) propose that perceptual destabilisation involves both top-down and bottom-up processes. They describe this as a gradual decrease in the activation of the current attractor state, making the perceptual state more unstable. The shallower the attractor, the more susceptible it is to both internal spontaneous fluctuations and external influences like blinks, eye movements, brief interruptions in stimulus presentation (Orbach et al., 1963), or light flashes (Kanai et al., 2005), which can trigger a perceptual reversal. In essence, a neural representation is conceptualised as an attractor within a state space in the brain. The depth of this attractor is influenced by factors such as the quality of visual input. For instance, an attractor has higher activation when the visual input is an unambiguous version of an ambiguous figure. In contrast, the activation is shallower for ambiguous stimuli, making the representation more responsive to spontaneous fluctuations and more prone to a reversal between attractors. Kornmeier and Bach (2012) suggest that with ambiguous inputs, the activated attractor is initially shallow and gradually flattens over time. Both bottom-up factors, like discontinuous stimulus presentation (e.g., Kornmeier et al., 2009), and top-down factors, such as volitional control (e.g., Toppino, 2003), can modify the depth of the attractor or increase the noise level (Braun & Mattia, 2010; Moreno-Bote et al., 2007). This alteration affects the rate of perceptual reversals, and the stability duration times, possibly in an additive manner (Braun and Mattia, 2010; Kornmeier et al., 2009). Therefore, this approach integrates bottom-up and top-down processes, suggesting that they are not mutually exclusive but rather interrelated in the perception of ambiguous stimuli.

1.3 Physiological evidence

In addition to the behavioural experiments focussed on in the sections above, a body of evidence supportive of the bottom-up, top-down, and integrative theories of multi-stable perception has also come from electrophysiological studies. Several studies have identified neurophysiological correlates of the processes thought to underpin multi-stable perception. These have been identified in both the pre- and post-stimulus periods; this section will describe the evidence from each period in turn.

1.3.1 Evidence from the pre-stimulus period

Numerous research efforts have been undertaken to delve into the predictive characteristics of brain activity before a stimulus occurs. The aim is to determine whether such pre-stimulus activity might influence an individual's later perception of an ambiguous stimulus. This has been explored using various brain imaging methods.

1.3.1.1 Pre-stimulus effects in fMRI studies

Several studies in the current literature have pointed to the role of ongoing, pre-stimulus brain activity in upcoming perception. In a study examining the rhythmic nature of visual perception, Hanslmayr et al. (2013) used a combination of electroencephalography (EEG) and functional magnetic resonance imaging (fMRI) to explore how these brain oscillations influence perceptual processes. This study revealed that visual information is not processed continuously, but rather in a rhythmic manner, oscillating between 5-10 Hz. The EEG data indicated that this rhythmicity mirrors the phase of ongoing brain oscillations within the same frequency range. Crucially, the study demonstrated that the phase of a 7 Hz oscillation before the presentation of a stimulus can predict both perceptual performance and the functional connectivity between higher and lower-level visual processing areas, specifically between the left lateral occipital cortex and the right intraparietal sulcus. This was evidenced through psychophysiological interaction and dynamic causal modelling. The authors suggest that these findings suggest a significant role of brain oscillations in periodically ‘gating’ visual perception at approximately 7 Hz, proposing a mechanism for transient time windows that facilitate long-distance cortical information transfer. They suggest that this rhythmic ‘gating’ may be a general mechanism underlying human perception.

Earlier work by Rahnev, Bahdo, de Lange and Lau (2012) investigated how spontaneous changes in attention levels influence confidence in perceptual decisions. Using fMRI, the study measured activity in the dorsal attention network before presenting stimuli in a motion direction discrimination task. It was discovered that lower pre-stimulus activity in this network, indicative of reduced attention, correlated with higher confidence in perceptual decisions. This finding supported earlier work by Rahnev et al. (2011), that found increased endogenous attention could lead to lower subjective perceptual ratings. In the context of multi-stable perception specifically, Hesselmann, Kell, Eger, and Kleinschmidt (2008), used fMRI to examine the link between brain activity preceding a stimulus and visual perception, using Rubin’s face-vase illusion. The findings indicated significant activation in the right Fusiform Face Area (rFFA) prior to the stimulus when participants later reported seeing faces, as opposed to when they perceived vase images. Furthermore, comparing baseline brain activity before and after the stimulus revealed a non-linear interaction. The conclusions drawn from Hesselmann et al. (2008)’s study point to two distinct yet interrelated factors influencing perceptual decisions: one associated with ongoing brain activity and the other with stimulus-driven responses (although, note the findings by Rassi et al., 2019 described in the section below).

1.3.1.2 Pre-stimulus effects in EEG studies

While blood oxygen level-dependent (BOLD) responses measured in the fMRI studies described above provide detailed spatial resolution, they lack the temporal precision required for tracking rapid fluctuations in neural activity. In contrast, scalp voltage changes measured by EEG offer researchers the ability to observe neural dynamics on a millisecond-by-millisecond basis, thus capturing the rapid fluctuations of brain activity that are thought to underpin perception.

Ronconi et al. (2017) examined how subjects perceived ambiguous stimuli created from two dots either appearing in the same location (leading to a perception of two flashes merging, called "two-flash fusion") or moving in space (perceived as "apparent motion"). The ISI was manipulated to induce perceptions of either integration (fusion/motion) or segregation (two distinct flashes). The authors found that information about participants' upcoming perception was contained in the pre-stimulus EEG activity. Perceptual outcome could be determined based on the phase of pre-stimulus oscillations in the EEG data, specifically from right parieto-occipital channels. The findings indicated that for the two-flash fusion, with an ISI of 40ms, pre-stimulus oscillatory activity in the alpha band was most predictive of the perceptual outcome. Whereas, for the apparent motion, with an ISI of 120ms, the phase of theta oscillations (6–7 Hz) was most indicative of the upcoming perception.

Modulations in oscillatory activity have also been shown to be predictive of upcoming percept in a later study by Rassi, Wutz, Muller-Voggel, and Weisz (2019). These authors used magnetoencephalography (MEG) to explore how spontaneous brain activity prior to a stimulus influences perception, specifically using Rubin's Face-Vase illusion, following the experimental design of Hesselmann et al. (2008). Their study revealed that oscillatory power did not differ for percepts of the 'face' versus the 'vase' in the primary visual cortex (V1) or in the FFA, suggesting comparable neural excitability for both types of perception. This finding is in apparent contrast with those of Hesselmann et al. (2008) who found that activity in the rFFA increased during the pre-stimulus period preceding reports of the face percept of the same stimulus. Although the authors here note that the spatial resolution of MEG might not be sufficient to distinctly identify activity in these areas. Nonetheless, Rassi et al. (2019) did observe more robust low-frequency oscillatory connectivity between V1 and FFA during 'face' trials compared to 'vase' trials. The intensity of this pre-stimulus feedback connection from FFA to V1 was predictive of both the imminent percept (face or vase) and the magnitude of neural activity following the stimulus that corresponded with the reported perception. Particularly, it was noted that on trials where faces were reported, there was stronger feedback connectivity from the FFA to V1 before the stimulus appeared. This interaction occurred within the alpha and beta frequency ranges. The time-frequency period in which the FFA's connectivity most strongly indicated the perceptual result was in tandem with the strength of this pre-stimulus feedback. Similar pre-stimulus effects have been shown using EEG with a high spatial resolution (256 channels). In their study, Britz, Landis and Michel (2009) found that there was activity present in the right inferior parietal cortex in the pre-stimulus period of trials where a reversal in perception occurred but not when perception remained stable. These authors also showed that this activity was present only in the moments before stimulus onset, but not during stimulus presentation. Pre-stimulus power differences in the gamma frequency band have also been shown to be indicative of upcoming perception of the ambiguous Necker lattice (Ehm et al., 2011; see Figure 2.1 in Chapter 2 for an illustration of this). Ehm et al. (2011) showed an increase in lower gamma-band activity (26–40 Hz) at the right-hemispheric central and parietal electrodes and an occipital decrease of higher gamma-band activity (40–65 Hz) on trials where participants reported a reversal in their perception of an intermittently presented Necker lattice. These modulations were absent in exogenous reversals of unambiguous lattice variants (i.e., where physical differences were introduced to the Necker lattice to render it unambiguous; Ehm et al., 2011). These results imply that preceding variations in brain activity can influence the resulting perception of a stimulus.

1.3.2 Evidence from the post-stimulus period

In addition to the pre-stimulus effects described above, a large body of electrophysiological research has focussed on the post-stimulus period. Event-related potential (ERP) analysis of EEG data has focussed on the post-stimulus period, as the nature of this analysis technique means that it is not suited to analysing pre-stimulus effects (see Section 1.3.2.2, and Chapter 2 for more detail on this technique). Several correlates associated with the perception of multi-stable stimuli have been identified from studies using fMRI and EEG data. These will be discussed in the sections that follow.

1.3.2.1 Post-stimulus effects in fMRI studies

Certain behavioural results have suggested that frontal brain regions are likely to be involved in the perception of multi-stable images (e.g., Balcetis &

Dale, 2007; Feist & Gentner, 2007; Goolkasian & Woodberry, 2010). Alongside this, a body of fMRI research has provided support for this notion. In particular, the inferior frontal cortex (IFC) is commonly cited as an active region during perceptual transitions (Brascamp, Sterzer, Blake & Knapen, 2018). One paradigm employed in these studies is known as the "replay condition", whereby the pattern of endogenous perceptual alternations is firstly recorded for each participant (endogenous reversal condition). In this condition, participants are asked to respond to indicate their perception of ambiguous stimuli and the order in which they perceived each trial is recorded. This order is then "replayed" to participants using physical stimulus changes (replay condition). In this condition, biased versions of the stimuli are presented in the same order as the pattern of endogenous reversals from the endogenous block, and participants respond in the same way. Blood oxygenation level-dependent (BOLD) signals are then compared between the two conditions, with the assumption that any differences are due to the endogenous reversal process rather than the additional cognitive demands resulting from the reporting process. The use of the replay condition has revealed significant increase in IFC activation during the endogenous reversal condition compared to the replay condition, indicating the involvement of the IFC in reversals of apparent motion stimuli (Sterzer & Kleinschmidt, 2007). The role of the IFC in perceptual alternations during binocular rivalry (Lumer, Friston & Rees, 1998), and bistable motion stimuli (Weilnhammer, Ludwig, Hesselmann & Sterzer, 2013) has also been revealed using the replay condition in this way.

However, an important caveat to the replay condition is that the physical stimulus changes that occur during it are easily distinguishable from endogenous reversals (Brascamp et al., 2018). This has led to some researchers modifying the replay condition so that it more closely resembles the endogenous reversal condition; in doing so, no significant differences were found between the endogenous and replay reversal conditions (Knapen, Brascamp, Pearson, van Ee & Blake, 2011; Brascamp, Blake & Knapen, 2015). These results suggest that the role of the IFC is related to the reporting of the perceptual change, rather than its cause. More support for this comes from a study by Brascamp et al. (2015), who created binocular rivalry stimuli where each eye received a different input. Participants' perception demonstrably switched between each input, yet the difference was so subtle that the participants did not consciously report it. In other words, the authors created perceptual reversals that only involved executive systems in a minimal way. When these were used, fMRI BOLD responses in the frontoparietal network were also minimised. The authors conclude that this finding is indicative of a consequential role of executive systems such as the frontoparietal network, rather than a causative one. However, Frässle, Sommer, Jansen, Naber, and Einhäuser (2014) employed a no-report condition in their fMRI study, using pupil size and eye movements to measure perception of rival binocular stimuli passively. In both the active and passive reporting conditions, the IFC remained significantly more active during perceptual alternations. Other prefrontal regions did show marked activation differences between the two conditions, however. The dorsolateral prefrontal cortex (dlPFC) was significantly activated during the active report condition, yet its activation levels did not reach significance in the passive report condition. The authors here posit that there are indeed areas of the prefrontal cortex that are enlisted for task monitoring and related processing only. The dlPFC has also been shown to be involved in the volitional control of reversals, suggesting it has a role in the voluntary control of attention in multi-stable perception tasks (Raz, Lamar, Buhle, Kane & Peterson, 2007).

In addition to the frontal regions described above, parietal regions have also been identified by fMRI as being more active during multi-stable perception (Inui, Tanaka, Okada, Nishizawa, Katayama & Konishi, 2000). The posterior parietal cortex, for example, is more active during both voluntary shifts in attention and percept (Slotnik & Yantis, 2005). The authors here conclude that voluntary perceptual reversals involve a redistribution of spatial attention, which is mediated by the posterior parietal cortex. Additionally, anterior and posterior superior parietal lobes (aSPLs and pSPL, respectively) have been shown to have specific roles during perceptual reversals. Baker, Karapanagiotidis, Coggan, Wailes-Newson and Smallwood (2015) suggest that the aSPL contributes to multi-stable perception through the inhibition of incongruent perceptual information, whilst the pSPL influences perception by supporting the current interpretation of an ambiguous stimulus. Further study has revealed that reciprocal interactions between bilateral aSPLs and higher order visual areas occur during perceptual ambiguity (Megumi, Bahrami, Kanai & Rees, 2015).

Transient modulations in the activation of higher and lower order visual regions have been shown to occur during multi-stable perception. Kleinschmidt, Buchel, Zeki and Frackowiak (2002) showed an increase in activation of the ventral occipital and intraparietal higher order visual areas and a corresponding decrease in activation of lower order visual areas in the primary visual cortex during perception of both Boring’s old/young woman and Rubin’s face/vase. Watanabe, Masuda, Megumi, Kanai and Rees (2014) conclude that overall activity dynamics during multi-stable perception fluctuate between three distinct, spatially distributed states: visual-area dominant, frontal-area dominant and intermediate. Those who tend to show more visual-area dominance experience more stable perception, and those with more frontal-area dominance were found to experience more reversals. The authors conclude that large scale activation dynamics are associated with multi-stable perception.

Other studies using fMRI data have shown that large scale cooperation and coordination of brain regions are likely to mediate multi-stable perception, as category-specific brain regions are active during experiences of perceptual ambiguity. The fusiform face area (FFA), for example, has been shown to be significantly more active during instances of the face percept of Rubin's face/vase (Andrews, Schluppeck, Homfray, Matthews & Blakemore, 2002; Wang, Sang, Hao, Zhang, Bi & Qiu, 2017). Tong, Nakayama, Vaughan and Kanwisher (1998) also report increased activation of the fusiform face area and parahippocampal place area during associated percepts of face/house rival binocular stimuli. A later study conducted by Sterzer and Rees (2008) used binocular rivalry to demonstrate that the FFA exhibits a sustained increase in activation in a delay period following a spontaneous perceptual switch toward a face stimulus but not during a replay condition. The authors suggest that this result could indicate that percept-specific mechanisms could be at least partially involved in the maintenance of a stable percept.

1.3.2.2 Post-stimulus effects in EEG studies

Whilst fMRI results provide spatially resolved insight into the neural underpinnings of multi-stable perception, additional data gathered using EEG can allow more temporal conclusions to be drawn. EEG can be used to analyse ERPs elicited in response to perceptual reversals. However, due to the subjective nature of these reversals, the selection of a relevant time-window in which to analyse such ERPs is problematic. If an ambiguous stimulus is presented continuously, for example, participants may be asked to respond via a key press when they experience a perceptual reversal. Due to individual and trial-by-trial differences in reaction times, this paradigm can cause temporal jitter, meaning that ERP components that occur only briefly are missed (Kornmeier & Bach, 2004; 2012). As a result, an intermittent viewing paradigm is more widely employed in ERP studies, taking stimulus onset as the moment of reversal and therefore increasing the temporal resolution (see Chapter 2 for a full overview of this issue; Kornmeier & Bach, 2004; 2012). Following work using this paradigm, Kornmeier and Bach (2012) have suggested a possible timeline for what they call the disambiguation process. This is the process by which ambiguous visual information is resolved into a given perception. Using the intermittent viewing paradigm, they suggest that the disambiguation of ambiguous visual information starts at around 130ms post-stimulus onset and takes approximately 50-60ms. They suggest this as the RP (discussed in more detail below) occurs alongside a left-hemispheric reduction in alpha power at 130ms post-onset, and the reduction in alpha power lasts around 60ms. The authors interpret this reduction in alpha power as being indicative of the disambiguation process, after which a resolution has been made in terms of the perception of the ambiguous stimulus. Further, they suggest that the decision about the perceptual outcome has taken place at least 340ms before the participant is able to indicate the consciously perceived reversal manually. Kornmeier and Bach (2012) discuss the presence of several ERP components when coming to these conclusions. For instance, a positivity at 80-100ms has been shown to occur in response to both endogenously and exogenously induced perceptual reversals (Kornmeier, Pfäffle & Bach, 2011; Kornmeier & Bach, 2012). Additionally, several ERP components have been associated only with reversals that are endogenous in nature.

To determine whether differences in ERP components differ between experimental conditions, a difference ERP (dERP) is calculated by subtracting the ERP waveform from one condition from that of another condition. In the EEG literature, two dERP components have been shown to be linked to perceptual reversals and discussed in the literature. These will be reviewed in the next two sections.

1.3.2.2.1 Reversal positivity

The Reversal Positivity (RP) is a dERP component discussed in studies by Kornmeier & Bach (2005, 2006), Kornmeier et al. (2007), and Britz et al. (2009). It appears around 130ms following stimulus onset. It is characterised by a relatively more positive amplitude during reversal trials compared to stable trials, predominantly at occipital electrode positions. The RP has been observed in response to reversals of various multi-stable stimuli, including the Necker Lattice (Kornmeier & Bach, 2005, 2006), the Necker cube (Kornmeier, Pfäffle, & Bach, 2011), Boring's Old/Young Woman (Kornmeier & Bach, 2014), and in binocular rivalry cases (Britz and Pitts, 2011). Notably, the RP is exclusive to endogenous reversals and absent in exogenous reversals, indicating its specific association with internally driven perceptual changes rather than alterations in sensory input. Additionally, the RP seems unaffected by low-level stimulus variations such as size (Kornmeier et al., 2011). However, its small amplitude means it is not consistently observed across all studies (Pitts et al., 2007; Intaitė et al., 2010). Kornmeier and Bach (2012) propose that the RP signals the detection of perceptual ambiguity or decision conflict during perceptual transitions (Pomerantz and Kubovy, 1981; Kornmeier et al., 2011), marking it as an indicator of the previously described disambiguation process. However, later work by Abdallah and Brooks (2020) found that the RP was dependent on whether the participant was required to make a manual response to indicate that they experienced a perceptual reversal. When participants were not required to make a manual response to indicate a perceptual reversal, Abdallah and Brooks (2020) found a component with the same latency as the RP but with the opposite polarity. The authors suggest, therefore, that the RP is a marker of response-related processes rather than those related solely to the disambiguation process.

1.3.2.1.2 Reversal negativity

Following the RP is a later, negative-going ERP component known as the reversal negativity (RN; Kornmeier & Bach, 2004). This is elicited over posterior electrodes around 200-260ms in response to endogenous reversals and 40-50ms earlier in response to exogenous reversals of the Necker cube stimulus (Kornmeier & Bach, 2004; 2012; 2014; 50ms difference replicated in Kornmeier, Bach and Atmanspacher, 2004). The RN has been found in response to a range of other ambiguous stimuli including Rubin's face/vase and Schröder's staircase (Pitts, Nerger & Davis, 2007). However, Boring's old/young woman failed to elicit the RN component, instead revealing an occipito-temporal N170 component that is typically reported in response to face stimuli (Kornmeier & Bach, 2014). The authors suggest that the reason for this is that the disambiguation of Boring's stimulus is completed 90ms earlier than geometric figures, leading to the N170 face ERP component. Source localisation techniques have suggested that the RN has a source in the occipito-temporal cortex (Pitts, Martinez, Stalmaster, Nerger and Hillyard, 2009). The amplitude of the RN has been shown to be enhanced during active volitional control of reversals (Pitts, Gavin & Nerger, 2008), and decreased in response to an increase in working memory load (Intaitė, Koivisto & Castelo-Branco, 2014) leading to conclusions that it reflects top-down influences. Other research has suggested that the RN reflects a reorganisation of the perceptual representation of the presented stimulus (Kornmeier, Bigalke & Bach (2005); Intaitė, Koivisto, Rukšėnas & Revonsuo, 2010).

1.3.2.1.3 Late positive potentials

Two further positive components have also been identified in endogenous reversals. The first occurs over frontopolar electrodes at 340ms (Kornmeier & Bach, 2012), the second at 470ms over parietal sites (Kornmeier & Bach, 2004). The frontopolar positivity has been suggested to reflect a possible role of working memory in perceptual reversals (Kornmeier & Bach, 2012), whilst the parietal positivity has been linked to the conscious appraisal of the reversal (Strüber & Herrmann, 2002; Kornmeier & Bach, 2006; 2012). This notion is further supported by other literature on conscious awareness, in which the spatiotemporally similar P3b component occurs only in response to a stimulus reaching conscious awareness (Dehaene & Changeux, 2011).

1.3.2.2 Evidence from EEG data in the time-frequency domain

Although a great deal of temporal information can be gained from ERP data, often certain signals that are not temporally aligned (i.e., not stimulus-locked) are lost to the averaging process involved in generating the ERP itself. Analysing data in the time-frequency domain can allow for the examination of these signals (Basar, Basar-Eroglu, Karakas & Schurmann, 1999).

1.3.2.2.1 Alpha band effects

The alpha frequency band refers to EEG oscillatory activity at 8 - 12 Hz (Herrmann, Grigutsch & Busch, 2005). EEG oscillations in the alpha band have been associated with visual processes. For example, a study conducted by Mathewson, Gratton, Fabiani, Beck and Ro (2009) used a metacontrast backwards masking paradigm which rendered a target detectable in 70% of trials despite stimulus parameters remaining constant throughout. Trials in which the target was detected were significantly associated with a decrease in alpha power. Additionally, the phase of the alpha oscillations in the pre-stimulus period were examined which revealed that, for trials with high alpha power, if a target was presented during the peak of an alpha wave it was significantly more likely to be detected than if it were presented during a trough. The authors conclude that alpha modulations reflect the momentary state of the excitability of the visual cortex.

Other studies have used repetitive transcranial magnetic stimulation (rTMS; for further discussion see Section 1.3.3, below) to provide more causative evidence for the role of alpha oscillations in visual perception. Romei, Brodbeck, Michel, Amedi, Pascual-Leone and Thut (2007) found that rTMS was more reliably able to induce phosphenes (illusory visual images induced by stimulation of the visual cortex) in the contralateral visual field when pre-stimulus alpha power was low compared to when it was higher. The authors suggest that alpha activity reflects momentary changes in the excitability of visual areas that are tuned to the contralateral visual field. This view is also supported by Romei, Gross and Thut (2010). They examined the effects of rTMS at alpha frequency over the left and right parietal and occipital sites on the detection rates of a lateralised target stimulus. The authors report that stimulation of both occipital and parietal sites of the contralateral hemisphere to target presentation impaired detection, whereas ipsilateral stimulation enhanced it, leading to conclusions that alpha activity has an inhibitory effect.

Linked with visuospatial attention (Thut, Nietzel, Brandt & Pascual-Leone, 2006), modulations within this band have also been reported alongside perceptual reversals. For example, a desynchronisation of alpha activity has been found to occur during the RP time window of perceptual reversals of the Necker cube (Isoglu-Alkaç, Basar-Eroglu, Ademoglu, Demiralp, Miener & Stadler, 2000). This activity has been shown to be separate from the motor-related activity involved in making a response (Mathes, Pomper, Walla & Basar-Eroglu, 2010). Alpha activity modulations have also been reported alongside the temporal dynamics of reversals; Piantoni, Romeijn, Gomez-Herrero, Werf and Someren (2017) used the Necker cube and found that local alpha power around the parieto-occipital sulcus within the first second after a percept's emergence predicted its duration. Moreover, experimentally inducing an increase in alpha power via sleep deprivation increased the average duration of representations. The authors suggest that an increase in alpha power supports the on-going stability of a particular percept. Strüber and Herrmann (2002) used magnetoencephalography (MEG) alongside the presentation of bistable motion stimuli to generate both endogenous and exogenous perceptual reversal conditions. By comparing the pre-reversal oscillatory activity in the alpha and gamma ranges in each condition, they were able to show that a steady decrease in alpha activity occurs during the period before the manual indication of a perceptual reversal in the endogenous condition. In the exogenous condition however, alpha activity remains stable until a rapid decrease around 300ms before the change in motion is reported. Event-related fields (ERFs, comparable to ERPs in EEG studies) revealed a P300-like component elicited in response to both endogenously and exogenously induced reversals, which the authors propose is indicative of the moment of conscious appraisal of the change in percept. The authors suggest that their results support the view that the perceptual representation must reach a threshold value before a reversal can occur, with any involvement of higher-order brain areas being reflective only of the conscious realisation that perception has reversed rather than playing an active role in generating the reversal itself. Thus, providing support for the bottom-up approach to multi-stable perception. Although Strüber and Herrmann (2002) did not find significant differences in gamma band activity between the endogenous and exogenous reversal conditions, Ehm, Bach and Kornmeier (2011) saw a right-lateralised gamma modulation which occurred at least 200ms before an endogenously induced reversal only. Ehm et al. (2011) suggest that this modulation could be indicative of an unstable brain state.

1.3.2.2.2 Gamma band effects

Gamma band activity refers to oscillations in the range of 30 – 80 Hz (Herrmann, Grigutsch & Busch, 2005), and has been linked to higher level cognitive functions such as attention (Ray, Niebur, Hsiao, Sinai & Crone, 2008) and declarative memory (Osipova, Takashima, Oostenveld, Fernández, Maris & Jensen, 2006). Gamma band activity has also been studied in the literature on multi-stable perception. In a binocular rivalry task, Doesburg, Green, McDonald and Ward (2009) report that perceptual alternations are time-locked to gamma band synchronizations, concluding that the synchronisation facilitates the generation of a new percept. Further analysis showed that the generators of these gamma rhythms have prefrontal and parietal sources, suggesting integration of multiple brain regions are involved in perceptual reorganisation. Mathes, Strüber, Stadler and Basar-Eroglu (2006) demonstrate an increase in gamma band power during volitional slowing down of perceptual reversals compared to speeding up, suggesting that gamma activity is indicative of the focal attentional processes involved in preventing a reversal rather than the attentional shifts mediating volitional reversals. However, an increase in frontal gamma band power has been linked to the destabilisation of percepts, especially among high-rate reversers (Başar-Eroglu, Strüber, Kruse, Başar & Stadler, 1996; Strüber, Basar-Eroglu, Hoff & Stadler, 2000; Strüber, Basar-Eroglu, Miener & Stadler, 2001). Whether or not activity in the gamma frequency band has a role in the facilitation or inhibition of a percept is a matter of on-going debate, therefore (Sedley & Cunningham, 2013). Other research has offered the suggestion that there may be two classes of gamma responses: an early response that reflects sensory processes and a later response that is indicative of more cognitive processes (Karakaş, Başar-Eroğlu, Özesmi, Kafadar & Erzengin, 2001).

1.3.2.2.3 Beta and theta band effects

Beta activity in the range 12 – 30 Hz has previously been shown to be linked to the semantic processing of visual images (Von Stein & Sarnthein, 2000), and has also been associated with multi-stable perception. Hipp, Engel and Siegel (2011) used ambiguous motion stimuli to demonstrate the coordination of large-scale brain networks during perceptual reversals; synchronous beta activity in a fronto-parieto-occipital network was found to be reliably linked to perceptual state. A later study by Minami, Noritake and Nakauchi (2014) found a decrease in beta power with a parietal-posterior source was associated with perceptual transitions. The authors therefore suggest that beta desynchronisation reflects the period of disambiguation of multi-stable images.

Despite theta activity (4 – 8 Hz) being linked to the coordination of brain activity over large areas (Von Stein & Sarnthein, 2000) relatively few studies have examined its role in multi-stable perception. One relatively recent study conducted by Mathes, Khalaidovski, Schmiedt-Fehr and Basar-Eroglu (2014) reports a transient increase in theta power during perceptual switches of both ambiguous and unambiguous stimuli. Thus, the authors conclude that theta may be involved in more general perceptual reconfiguration.

1.3.3 Evidence from studies using transcranial stimulation

As briefly touched upon above, the causal nature of EEG oscillations upon multi-stable perception can be examined via brain stimulation techniques. To entrain a certain cortical region to a certain frequency band, two main stimulation techniques are used: transcranial alternating current stimulation (tACS), and rTMS. tACS involves an electrical alternating current being applied transcranially via electrodes placed onto the scalp. The frequency of this alternating current is set so that it corresponds to the EEG activity band of interest. rTMS is used in the same way, however a magnetic rather than an electrical field is used. rTMS can also be used to induce “virtual lesions” if particular patterns of pulses are employed (Pascual-Leone, 1999). An additional, less frequent technique that has been applied to the study of multi-stable perception is caloric vestibular stimulation (CVS), which involves stimulation of the vestibular system by introducing ice-cold water into the external auditory canal via a syringe. This has been shown in other research to induce unilateral hemisphere activation (Rossetti & Rode, 2002). Findings from studies employing tACS, rTMS and CVS techniques to investigate multi-stable perception are discussed in the sections below.

1.3.3.1 tACS

Using bistable apparent motion stimuli which could be interpreted as moving across a horizontal or vertical plane, Strüber, Rach, Trautmann-Lengsfeld, Engel and Herrmann (2014) applied tACS bilaterally over occipito-parietal areas at 40Hz; a frequency known to be in the gamma band activity range. The authors report that the tACS stimulation led to a decreased proportion of perceived horizontal motion only when coupled with a 180˚ phase difference in the AC waves, suggesting that desynchronisation of gamma band activity between hemispheres leads to biased ambiguous motion perception. Cabral-Calderin, Schmidt-Samoa and Wilke (2015) presented structure from motion stimuli alongside tACS at 60Hz (also in the gamma range) placed centrally over the occipital lobe in order enhance gamma band activity. The stimulation induced a significant increase in the number of spontaneous perceptual reversals compared to alpha frequency stimulation (10Hz) at the same site, supporting the idea that gamma activity is involved in the formation of an unstable perceptual state.

1.3.3.2 TMS

Studies using single and repetitive pulse TMS have examined the effects of enhancing or disrupting the function of several cortical regions thought to be involved in perceptual reversals. For example, the right lateralised fronto-parietal network identified by fMRI (Brascamp et al., 2018) and EEG (Hipp et al., 2011) was examined in a more causal manner by Vernet, Brem, Farzan and Pascual-Leone (2015). Here, single pulse TMS was applied over the right anterior intraparietal sulcus (IPS) 70ms before the presentation of an ambiguous stimulus. This led to an increased likelihood of a perceptual reversal. However, when stimulation in the same area was followed by a second pulse over the dlPFC, this effect became non-significant. This suggests that the IPS has a role in the stabilisation of the current percept and the dlPFC is involved in the triggering of perceptual reversals, and highlights the coordination between frontal and parietal regions in multi-stable perception. Additionally, rTMS applied over the right IPS has been shown to prolong periods of perceptual stability (Zaretskaya, Thielscher, Logothetis & Bartels, 2010), providing further support for the idea that it pays a critical role in the formation of a stable percept.

1.3.3.3 CVS

CVS is able to induce activation of the contralateral cerebral hemisphere (Rossetti & Rode, 2002) and as a result has been used to examine the interhemispheric switch (IHS) model of multi-stability (Miller, Liu, Ngo, Hooper, Riek, Carson & Pettigrew, 2000). The IHS model proposes that each cerebral hemisphere contains a particular perceptual representation of rival or ambiguous stimuli and perceptual alternations arise from the relative activation of one hemisphere over another (Miller et al., 2000). It has been shown that hemispheric stimulation via CVS is able to significantly shift perceptual predominance levels of rival and ambiguous stimuli (Miller et al., 2000; Ngo, Liu, Tilley, Pettigrew & Miller, 2007), which provides support for this IHS model.

1.3.4 Evidence from studies using Multivariate Pattern Analysis (MVPA)

A more recent approach to the analysis of neuroimaging data is known as multivariate pattern analysis (MVPA). Rather than analysing a narrow temporal or spatial region of interest, this approach is focused on the analysis of whole-brain signal patterns. This is typically achieved using machine learning algorithms to identify patterns of signal across multiple brain regions that are associated with a particular cognitive function. This technique has been used for some time with fMRI data. Unlike traditional fMRI analysis, which uses a mass univariate statistical analysis to focus on the intensity of activation within certain brain regions, MVPA considers the pattern of activity across multiple voxels. This approach has been instrumental in understanding how the brain interprets ambiguous stimuli. For instance, Sterzer, Haynes, and Rees (2008) demonstrated the use of MVPA in analysing fMRI data to decode BOLD responses when participants were presented with a face or a house stimulus. The authors used a binocular rivalry paradigm such that either the face or the house was suppressed and therefore invisible to the participants. They found that overall responses of high-level ventral visual areas were very weak and did not differ in amplitude between the two conditions (faces or houses). Despite this, MVPA within these areas allowed the authors to predict whether an observer was presented with face or house stimuli, not only when these stimuli were visible but also when they were suppressed and entirely invisible. Additionally, Brouwer and van Ee (2007) successfully used MVPA of fMRI data to distinguish between alternative directions of SFM sphere-rotation based on patterns of activity within area MT+, along with other dorsal visual and parietal areas.

Based on the successful use of MVPA in decoding viewers' perceptual state from fMRI data, research using EEG and MEG has also moved toward these types of analyses. Schhmiedt, Rotermund, Basar-Eroglu and Pawelzik (2009) used support vector machine (SVM) classifiers to successfully decode viewers' perception of an ambiguous structure from motion stimulus from 30 channels of EEG oscillatory data. Additionally, Das et al. (2010) aimed to compare the efficacy of EEG pattern classification, using three distinct classifiers, against traditional, univariate ERP indices such as peak amplitude, mean amplitude, and peak latency. In their study, participants were engaged in a task where they were shown a face, or a car overlaid with Gaussian noise and had to rate their confidence in identifying the category of the stimulus (i.e., a face or a car). Earlier studies employing univariate methods of analysing EEG and MEG data have identified specific ERP components such as the N1 and N170 that are linked to face presentation as opposed to other objects (Gauthier et al., 2003; Taylor et al., 1999) or the early trial averaged M100 in MEG related to face categorisation (Liu et al., 2002). However, diverging from these findings, Das et al. (2010) discovered that the neural activity indicative of correctly perceiving and reporting faces versus cars was distributed over time, commencing at 120ms and persisting for over 400ms after stimulus onset. This broader temporal window for significant activity pattern classification differs from previous univariate observations, suggesting a more prolonged neural involvement in distinguishing between faces and cars. MVPA has also been shown to have revealed subtle effects in other research areas, for instance, decoding the focus of attention (local or global; e.g., List et al., 2017) from EEG data. It has also been used to successfully decode which of 16 orientations is being held in working memory, even when orientation is not currently relevant to the task (Bae & Luck 2018; 2019).

Another benefit of MVPA is that it also operates within the time-frequency domain, as well as the time domain. Past research has identified oscillatory brain activity that is linked to the various stages of processing in perceptual reversals (see Section 1.3.2.2). MVPA, like univariate analyses, can detect frequency changes throughout the course of a trial that are associated with perceptual reversals. However, MVPA's advantage lies in its ability to consider multivariate patterns in time-frequency data from across the scalp, providing a broader perspective on the timing and frequency ranges that involve multiple brain mechanisms. In the time-frequency domain, MVPA compares the activity patterns of different perceptual states at each time point and frequency level, which could help identify the specific time-frequency windows where these activity patterns differ between reversal versus stable trials. In this way, MEG data has been used to successfully decode participants' perception of the Rubin's faces/vase stimulus (Rassi et al., 2019) and perception of face vs house in a binocular rivalry paradigm (Rassi et al., 2022).

In addition to its ability to detect *differences* in multivariate patterns of activity across the scalp, a second key advantage of MVPA over the more traditional univariate analyses is the ability to perform what is known as cross-decoding (see Chapters 2 and 4 for more detail on this). Cross-decoding allows MVPA to also detect *similarities* between two or more different brain states evoked by experimental manipulation. For instance, Stokes et al. (2009) conducted an fMRI study where participants were shown or asked to imagine the letters "X" or "O" as well as trials on which participants were actually shown images of “X” or “O”. A classifier trained on visual (i.e., perceived) trial data could be used to accurately predict the imagined letter from the imagined trial FMRI data. Subsequent research by Stokes et al. (2011) confirmed this result, even when letters appeared in different visual field locations. Cross-decoding has also identified common neural patterns in motor action and observation. Oosterhof et al. (2012a) recorded fMRI data as participants performed or watched actions ("lift" and "slap") from various viewpoints. They discovered specific brain activity patterns in the ventral premotor cortex corresponding to the observed or executed actions, with significant cross-modal neural patterns for first-person perspectives correlating with action execution. However, these patterns did not hold for third-person perspectives. This distinction illustrates how MVPA, and cross-classification in particular, can discern subtle neural activity similarities and differences across conditions.

MVPA has also been used to reveal insights into the pre-stimulus brain state that is predictive of a particular percept. Ronconi, Oosterhof, Bonmassar and Melcher (2017) demonstrated that pre-stimulus EEG oscillatory data could be successfully decoded as predictive of a particular percept by a machine learning classifier.

1.4 Aims and hypothesis of current thesis

The evidence discussed in the sections above indicates that there are likely to be several processes involved in the reversals in perception that occur when ambiguous stimuli are viewed. Both top down and bottom-up mechanisms are likely to be involved. Bottom-up theories are largely centred around some form of neural adaptation – the neurons involved in maintaining the current percept become fatigued to the point at which the neurons involved in the other percept are less fatigued and therefore perception switches (e.g., Long & Toppino, 2004). Studies investigating the role of adaptation processes in multi-stable perception have induced adaptation by pre-exposing participants to a biased version of an ambiguous stimulus which induces the neurons underpinning a given percept into fatigue (e.g., Long & Olszweski, 1999; Long & Moran, 2007). Single cell studies have shown that adapting neurons in this way does indeed reduce the response to subsequent stimulation (e.g., Barlow & Hill, 1963; Hammond et al., 1988a; 1988b; Kohn & Movshon, 2003). Behaviourally, this is shown because when the ambiguous stimulus is subsequently presented, it is more likely to be perceived in the alternate way (e.g., Kanai & Verstraten, 2005; Long & Moran, 2007).

Studies have shown that the length of time that the adapting stimulus is presented for is critical in inducing this effect. In fact, presenting the adapting stimulus for a brief period has the opposite effect – the subsequent ambiguous stimulus is perceived in the *same* way as the adapting stimulus (priming; e.g., Long & Moran, 2007). Research suggests that priming may be a behavioural reflection of top-down involvement in perception of ambiguous stimuli because it can also be induced by conceptual as well as perceptual primes (e.g., Goolkasian & Woodberry, 2010). For example, being primed by reading a paragraph, or otherwise being exposed to stimuli that are only conceptually related to one of two possible percepts, leads to the ambiguous stimulus being perceived in the same way as the prime (e.g., Balcetis & Dale, 2007; Feist & Gentner, 2007; Goolkasian & Woodberry, 2010). Priming effects have also been shown to persist for longer time periods than those involved in adaptation – again suggesting it involves higher level mechanisms (e.g., Long & Moran, 2007). This has been shown in studies that manipulate the ISI between adapting and ambiguous stimuli: priming effects persist after longer ISIs whilst adaptation effects diminish (e.g., Kanai & Verstraten, 2005; Long & Moran, 2007). This supports the idea that low level, neural fatigue is responsible for the adaptation effects seen behaviourally in studies that induce an adapted brain state via pre-exposure to an adapting stimulus (followed by a report of perception of an ambiguous test stimulus).

What is not clear from the existing literature, however, is whether adaptation plays a role in the spontaneous perceptual reversals that occur when only ambiguous stimuli are intermittently viewed. This is because previous adaptation studies have always involved pre-exposure to biased adaptors and thus there is no direct evidence that adaptation actually accumulates during the presentation of fully ambiguous stimuli which then affects the perception of subsequent ambiguous stimuli. It is this gap in the available research that the current thesis aims to fill, and therefore forms the overarching research question answered by the current work. The general approach to this question will be to train a machine learning classifier (see Section 1.3.4) to decode perceptual reversals (i.e., reversal vs. stable trials) of an ambiguous stimulus from EEG data recorded during an adapted brain state which has been induced by exposure to unambiguous interpretations of an ambiguous stimulus (e.g., an altered Necker cube to make it appear right facing). This trained classifier will then be used to attempt to predict spontaneous (i.e., not induced by an unambiguous stimulus) perceptual reversals in a sequence of ambiguous visual stimuli. If these spontaneous reversals involve processes of adaptation similar to those induced by adaptation to unambiguous stimuli, then the trained classifier should be able to accurately predict perceptual outcomes based on the EEG data recorded during them. This approach thus uses a form of cross-decoding to address directly whether adaptation plays a role in spontaneous reversals that occur in sequences of fully ambiguous stimuli.

The specific research aims to help to answer this research question will now be outlined. Full details of the background and rationale for each individual experiment can be found within empirical Chapters 3-5.

1.4.1 What are the optimum timing parameters required to induce an adapted brain state?

As there is a large body of evidence that EEG data can be reliably used to reveal patterns of activity associated with spontaneous perceptual reversals, this research will primarily involve recording EEG data during these reversals. Stationary stimuli such as the Necker cube take around 30-150 seconds to reliably induce adaptation, which is too long for them to be used practically in any EEG experiments that aim to investigate adaptation (e.g., Shulman, 1993; Long, Toppino & Mondin, 1992) due to the need for many repeated measurements in

EEG studies. In contrast, the rapid motion aftereffect (rMAE) has been shown to quickly evoke adaptation effects when participants view ambiguous motion stimuli after being shown unambiguous directional motion stimuli for only 640ms (Kanai & Verstraten, 2005). Therefore, motion stimuli highly similar to those used by Kanai and Verstraten (2005) will be used to examine the role of adaptation in spontaneous perceptual reversals here.

This research aim, therefore, is to determine the optimal parameters for inducing the adaptation and priming effects described above. This was achieved by partially replicating the study by Kanai and Verstraten (2005). In contrast to the previous authors however, this experiment aimed to produce two clear trial types: one in which the adaptation effect is maximally induced and one in which the participants' perception was unbiased (i.e., 50% of each interpretation). The temporal parameters evoking these two trial types will then be used to inform the subsequent EEG experiment investigating adaptation.

It should be noted at this stage that one of the planned aims of this thesis was to investigate the roles of both adaptation *and* priming in perceptual outcomes when viewing ambiguous stimuli intermittently. However, as detailed and fully investigated in Chapter 3, the results of Experiment 1 led us to become uncertain about whether visual motion priming (VMP; see also Section 1.4.2; Pinkus & Pantle, 1997; Kanai & Verstraten, 2005) was indeed evoked using the paradigm employed (Kanai & Verstraten, 2005). Therefore, instead of moving forward with the planned investigation into the role of priming in the perceptual outcomes of ambiguous motion, this potential artefactual effect of task on behavioural VMP effects was fully investigated in Chapter 1, Experiment 2, and will be discussed in more detail within the next section.

1.4.2 Can the task that participants are asked to perform lead to artefactual VMP effects?

As described above, whilst we were conducting Experiment 1 and analysing the results, we became aware of a possible confound in the way that Kanai and Verstraten (2005) phrased their task to their participants. More detail on this can be found in Chapter 3 but will also be presented briefly here in order to clarify this research aim. As described above, Kanai and Verstraten (2005) manipulated the presentation time and ISI when presenting unambiguous adapting motion followed by ambiguous test motion and found that when both were brief, strong VMP was shown. In their briefest condition, they used an adapting stimulus presentation time of 80ms, an ISI of 40ms, followed by 320ms of ambiguous motion. At the end of a trial, they asked participants to report whether the two stimuli were moving in the same or a different direction. Our task was phrased slightly differently in that we asked participants to report each direction that they saw the two stimuli move in (so they might report "left", "left", for instance). Our results did not show a strong VMP effect under these conditions as Kanai and Verstraten (2005) found, and anecdotal reports from our participants indicated that they didn't see that there were two stimuli at all under these conditions.

This research aim then, was to replicate Experiment 1 whilst experimentally manipulating the task that participants were asked to perform, as well as directly asking participants about the number of stimuli they were able to perceive. This aim allowed us to investigate the effects of task on VMP. The hypothesis was that task would impact the reports of the VMP effect and under the briefest temporal conditions participants would be less likely to perceive two stimuli.

1.4.3 Does adaptation play a role in spontaneous reversals in perception that occur during viewing of ambiguous motion?

This research aim will be addressed by training machine learning classifiers to decode perception of ambiguous motion from EEG data recorded from an adapted brain state, as well as an unbiased brain state that will serve as a control, both induced by pre-exposure to biased interpretations of an ambiguous stimulus. The temporal parameters required to induce each of these states will be addressed by the research aim in Section 1.4.1. These classifiers will then be used to attempt to predict spontaneous perceptual reversals of the ambiguous-only motion.

The hypothesis here then, if spontaneous perceptual reversals are driven by adaptation, the decoding accuracy of the classifier trained on the adapted brain state will be significantly greater than chance, and also significantly higher than the decoding accuracy of the classifier trained on the non-adapted brain state.

1.4.4 Are ERP measures of reversal-related brain activity present during perceptual reversals of ambiguous stationary stimuli also present during perceptual reversals of ambiguous motion stimuli?

Several studies have investigated the re-stabilisation and disambiguation processes occurring during the perception of ambiguous figures. These investigations have associated two ERP components, the RN and the RP, with perceptual reversals (e.g., Kornmeier & Bach, 2012). The exact characteristics of these components and the mechanisms underlying them remain to be fully understood. Kornmeier and Bach (2012) highlight the need for research involving different types of ambiguous stimuli to be conducted, in order to establish whether these components are reflective of the mechanisms underpinning perceptual reversals in general or only the specific task designs that have investigated them so far. Work by Abdallah and Brooks (2020) has highlighted the need for such investigation as they found that the presence of the RP was dependent on whether the participant was required to respond. Therefore, highlighting the need for investigating whether these components are linked to reversal-related processes or other, non-related mechanisms.

If these components are markers of the processes that underpin spontaneous perceptual reversals, it is hypothesised that they will also be present during reversals of the ambiguous motion stimuli used in the series of experiments described in this thesis.

1.4.5 How powerful is MVPA to detect multivariate effects present in EEG data?

This was an unplanned deviation from the initial plan to collect further EEG data and was due to the COVID-19 pandemic lab closures (see COVID Impact Statement for full details of this deviation). However, this aim still fits well to address the overarching gap in the current literature, as the thesis aims to draw conclusions about the role of adaptation in spontaneously occurring perceptual reversals, primarily using MVPA to directly address this aim. MVPA is a relatively new technique in the field of EEG research, and certainly provides many benefits over the more traditional univariate analyses (see Section 1.3.4).

However, the currently available literature is lacking in information regarding the number of participants required to detect an effect with adequate power – in other words, the number of participants required to detect a given effect in at least 80% of cases (Cohen, 1988). Therefore, to help clarify that the experiments detailed in this thesis are adequately powered as well as shed some light on this to fill the gap in the current literature, this research aim will be addressed using simulated null effect EEG data with multivariate effects of known sizes added. This is described fully in Chapter 5.

Chapter Two: Methods

Various methods have been used in previous research aiming to understand the role of adaptation and priming in the perception of ambiguous stimuli. This chapter will outline and justify the methods used in the set of experiments that are to follow and discuss the more general issues that surround them. All materials used in the set of experiments described in this thesis can be found at the following OSF repository: https://osf.io/wbmyp/.

2.1 Stimulus choice

As discussed in Chapter 1, there are several stimuli that have been used when investigating adaptation and priming effects. The general paradigm that is employed when presenting stimuli in such investigations usually involves some biased, unambiguous version of the stimulus being presented for a given amount of time, followed by an ambiguous version. The biased stimuli are typically created by introducing physical features to the figure that mean that one interpretation is more likely to be perceived. For instance, elongating the faces in Rubin's faces-vase, and the introduction of T-junctions that suggest to the viewer that the object creating the top of the 'T' is in front of the object forming the stem of the 'T'. These can be seen at the circled areas in Figure 2.1 (Abdallah & Brooks, 2020). The introduction of these cues means that although the vase is still present in the image shown in Figure 2.1, Panel E, it is far more likely that the faces percept will be the one that is perceived (and vice versa for Panel F). Research shows that this is because there is a shared border between the two percepts (i.e., the line that can be perceived as belonging to either the profile of the face or the edge of the vase) and shared borders are assigned to a shape on just one of the sides they separate (e.g., Driver & Bayliss, 1995a, b; 1996). Similarly, introducing depth cues via shading and transparency level of the edges of the Necker cube or Necker Lattice (e.g., Kornmeier & Bach, 2004, 2012; Abdallah & Brooks, 2020) image leads to the top left face being more likely to be perceived as shown in Figure 2.1, top, below.

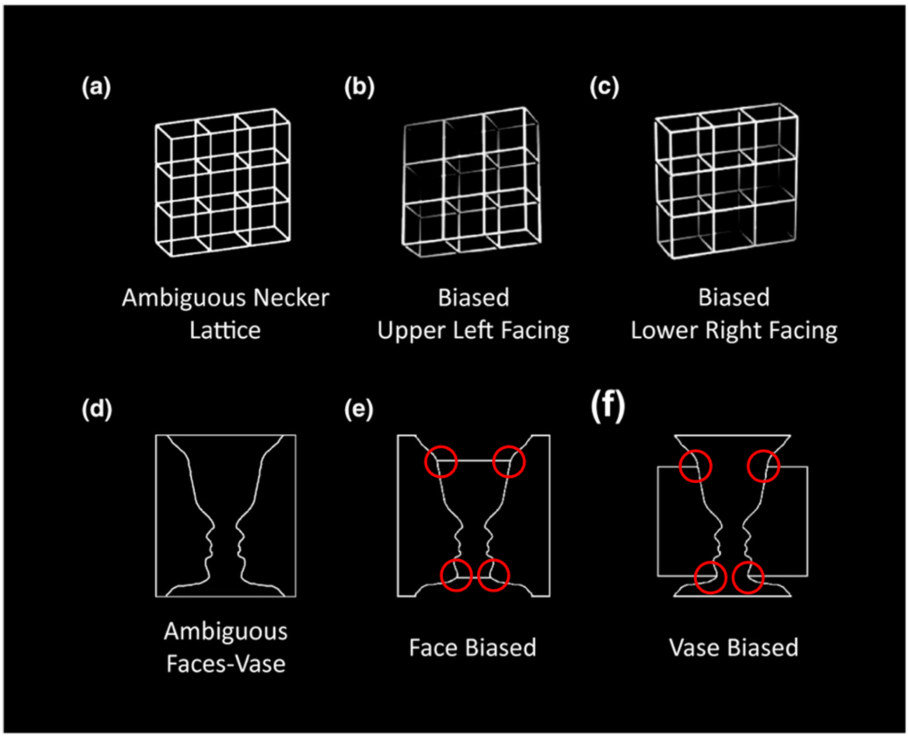

Figure 2.1: Examples of the Necker lattice (a) based on Kornmeier and Bach (2004) and Rubin's faces-vase (d), (b) A biased version of the reversible Necker lattice which is typically perceived with its front face toward the upper left. (c) A biased Necker lattice with its front face toward the lower right. (e) A modified faces-vase image biased toward the face interpretation. (f) A faces-vase image biased toward the vase interpretation. Red circles indicate where t-junctions have been introduced to serve as depth cues. Taken and slightly adapted from Abdallah and Brooks (2020).

One of the aims of this thesis is to use electroencephalogram (EEG) data to establish whether adaptation plays a role in the spontaneous reversals in perception that occur when viewing ambiguous stimuli. Therefore, it is important to select a method that is known to reliably induce adaptation in behavioural studies. Thus, the method of presenting pairs of 'adaptor' followed by 'ambiguous' stimuli will be used in this set of experiments, in line with the well-established practice in the currently available literature. It is important, therefore, when selecting the stimuli for investigating adaptation and priming effects, that a biased version of the chosen stimulus can be created that reliably reduces ambiguity so that adaptation or priming can be induced. Although stationary stimuli like those described above have been used to investigate adaptation and priming effects in behavioural studies, they have largely not been used in studies that have aimed to examine the neurophysiological correlates of such effects using neuroimaging techniques like EEG. The reason for this is because, as will be discussed in more detail below, many trials are required in these studies for the signal to noise ratio to be high enough to detect an effect (see Section 2.3 below; Luck, 2014). The issue with the use of the stationary stimuli described above is that the presentation duration required to induce adaptation effects can be upwards of 30 seconds (e.g., Shulman, 1993), which is too long to use in experiments that contain sufficient trial numbers to detect neurophysiological effects. It is not clear, however, whether the length of time required to adapt to these stationary stimuli is a general property of all stationary stimuli or only the ones that have as yet been used to investigate adaptation. However, in contrast to ambiguous stationary stimuli, ambiguous motion stimuli have been shown to reliably induce adaptation effects in as little as 2 seconds (e.g., Hoffman, Unsöld & Bach, 2001), and some even as low as 0.64 seconds (e.g., Pinkus & Pantle, 1997; Kanai & Verstraten, 2005; Takeuchi et al., 2008). As well as this, it is possible to create reliably unambiguous variants of ambiguous motion stimuli with minimal changes to the image as will be discussed in more detail later in this section.

2.1.1 Mechanisms of motion perception

Before justifying the choice of ambiguous motion stimulus used in this set of experiments, it is important to firstly outline the neural mechanisms that are thought to underpin the perception of such stimuli. The mechanisms behind motion processing have been substantially investigated in the current literature. Single cell studies have shown that, in macaque monkeys, a large amount of motion processing is carried out in areas V1, V2, V3 and the medial temporal complex (MT; these have been reviewed by DeYoe & Van Essen, 1988). Single cell recordings of motion detector responses in these areas have been shown to be direction specific, suggesting that these are true motion detectors (Borst & Egelhaaf, 1989). In humans, the neural correlates of motion perception have largely been studied with neuroimaging methods. In particular, the onset of visual motion induces the visual evoked potential in EEG studies (VEP; e.g., Bach & Ullrich, 1994; Snowden, Ullrich & Bach, 1995). Two VEP components are evoked by the onset of visual motion: a positivity from around 100-130ms at occipital/occipito-temporal electrodes (P1), and a negativity from around 150-200ms at the same sites also including the vertex (N2; both components reviewed by Niedeggen & Wist, 1998). Studies into motion adaptation have typically looked at the difference in mean amplitudes of these two VEP components when motion is presented after a period of adaptation – if the amplitude reduces then adaptation is said to have taken place (e.g., Hoffman, Dorn & Bach, 1999). These results are particularly prevalent for N2 (e.g., Hoffman, Dorn & Bach, 1999; Wist, Gross & Niedeggen, 1994; Bach & Ullrich, 1994), however there remains some uncertainly around the direction specificity of the P1 VEP component as it does not appear to show direction-specific adaptation (Hoffman, Unsöld & Bach, 2001). Source analysis suggests that the N2 component originates in the human equivalent of the monkey MT area, referred to as hMT+/V5, further supporting the idea that it reflects direction specific mechanisms (Probst, Plendel, Paulus, Wist, & Scherg, 1993). These effects have been elicited by unambiguous random dot kinematograms (RDKs; e.g., Hoffman, Unsöld & Bach, 2001; Smith, Wall, Williams & Singh, 2006) as well as drifting square-wave gratings (e.g., Bach & Ullrich, 1994).

In terms of the perception of ambiguous motion stimuli, Freeman, Sterzer and Driver, (2012) used transparent random-dot kinematograms representing either a cylinder rotating in depth (representing 3-dimensional motion) or two flat surfaces translating in opposite directions at apparently different depths (representing 2-dimensional motion). The authors used fMRI to compare the activation patterns when 2-dimensional (2D) versus 3-dimensional (3D) ambiguous motion was viewed. Their results showed that the hMT+/V5 complex shows reversal-related activity during presentation of both 2D and 3D ambiguous motion. In addition, their results showed that only 3D but not 2D ambiguous motion stimuli evoked reversal related activity in the lateral occipital complex (LOC), which is known for its sensitivity to 3-dimensional form (Welchman, Deubelius, Conrad, Bulthoff, & Kourtzi, 2005; Preston, Kourtzi, & Welchman, 2009). Additionally, area MT shows adaptation to motion (Van Wezel & Britten,

2002) and this maps behaviourally onto adaptation results (Glasser, Tsui, Pack & Tadin, 2011).

More broadly, the perception of ambiguous motion stimuli is also thought to reflect both top-down (knowledge-driven, e.g., planning of movement) and bottom-up processes (sensory-driven, e.g., environmental cues). For instance, the study by Piedimonte, Woods, and Chatterjee (2015) focused on how ambiguous motion is disambiguated using different sources of cues, both exogenous (external) and endogenous (internal). They found that, when there is an unambiguous exogenous cue presented alongside the ambiguous motion, this will strongly bias participants' perception of that ambiguous motion. For example, if there is clear directional motion in a particular direction presented alongside the ambiguous motion, participants are likely to perceive the ambiguous motion moving in the same direction as the directional motion cue. This effect could not be modulated by the addition of competing endogenous cues, for example asking participants to perform or plan (i.e., involving top-down processes) movements in the opposite direction to the exogenous cue. This indicates that both bottom-up and top-down processes are involved in how we perceive and interpret motion in our environment with the interaction between the two streams of information showing a bias toward bottom-up information. However, the authors of this study also found that endogenous cues from learned associations were able to override even exogenous cues when the two were in competition. This indicates the widespread and complex interactions between several mechanisms when ambiguous motion is resolved.

2.1.2 Drifting sine wave gratings

An example of ambiguous motion used when investigating adaptation and priming effects is that created from drifting sine wave gratings (e.g., in Pinkus & Pantle, 1997; Kanai & Verstraten, 2005; Takeuchi et al., 2011). Unambiguous motion from sine wave gratings is easily created by shifting the phase of the sine wave over time. For instance, motion to the right can be created by shifting the phase of the wave by 90º every 0.5s for example (see Figure 2.2 below). Ambiguous motion can also be created from drifting sine wave gratings by shifting the phase of the wave by 180º rather than 90º. This creates a stimulus where the lighter areas are replaced by the darker areas, and the motion can be perceived as moving either to the left or to the right because the phase shift does not provide a clear direction for the motion's progression. This ambiguity arises from the correspondence problem, as the visual system struggles to match elements from one frame to the next when the shift is 180º, making the direction of motion indeterminate. (see Figure 2.2 below).

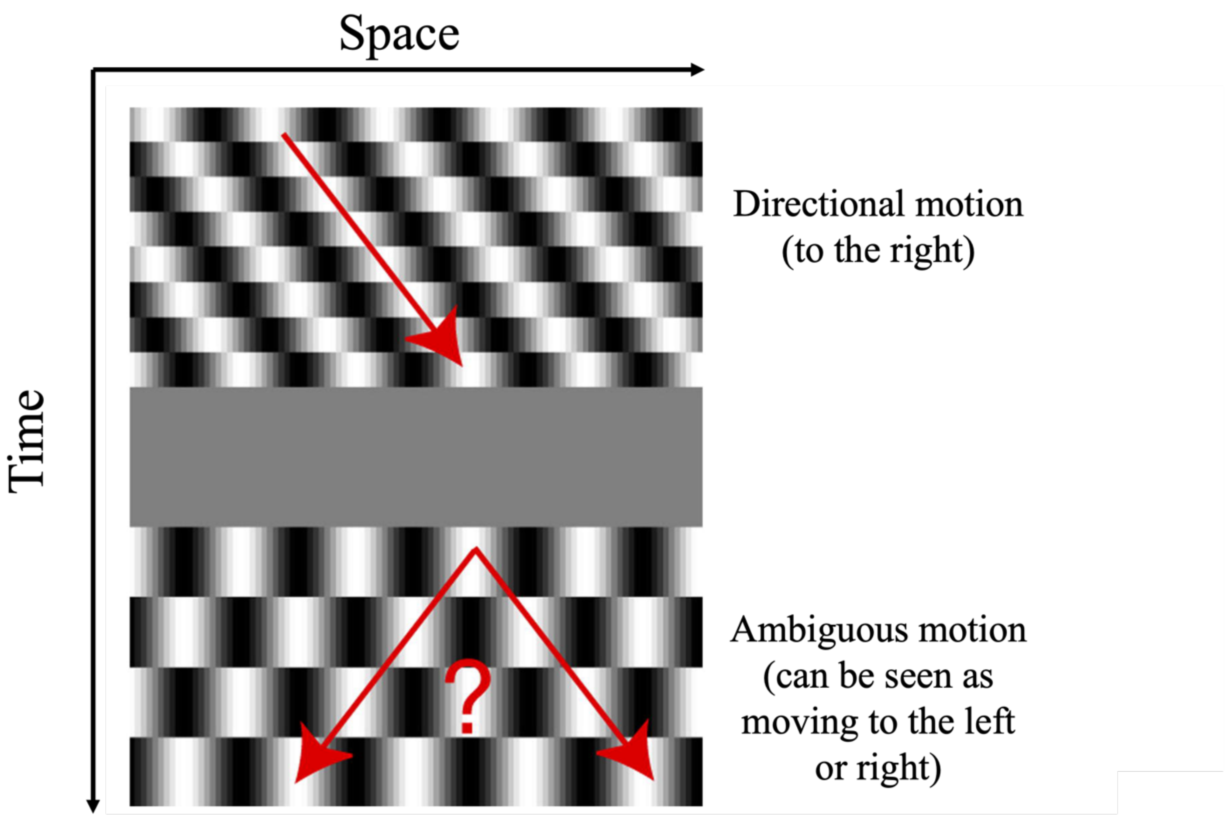

Figure 2.2 Illustration demonstrating how the phase of sinewaves can be shifted over time to create directional and ambiguous motion. The top of the figure shows a phase shift of +90° in space, every arbitrary time unit, leading to the generation of motion to the right in this instance. The bottom of the figure shows a phase shift of ±180° in space, every arbitrary time unit, leading to the generation of ambiguous motion that can be perceived as moving either to the left or to the right. Adapted from Kanai and Verstraten (2005).

These sine wave gratings have been successfully used to study adaptation and priming effects in behavioural studies, under temporal conditions that are suitable for use with EEG. For instance, Pinkus and Pantle's (1997) design involved an adaptation phase where participants were exposed to directional motion created from a drifting sine wave grating (similar to that shown in the top diagram of Figure 2.2 above). This was followed by a variable inter-stimulus interval (ISI) and then a test phase, where an ambiguous test grating (similar to that shown in the bottom diagram of Figure 2.2), was briefly shown. The authors were able to show that visual motion priming (VMP) was able to be induced in 192ms and adaptation to the test stimulus (the directional motion generated by the sine-wave grating) was able to be induced in 1530ms. Using a similar design, Kanai and Verstraten (2005) were able to show that VMP was able to be induced in just 80ms and adaptation in 640ms; an effect they term the Rapid Motion Aftereffect (rMAE). These results were also replicated by Takeuchi et al. (2011), who used a similar design but removed the ISI. The results from these studies provide robust evidence that drifting sine wave gratings can create directional and ambiguous motion. These studies also show that this motion can be used to induce priming and adaptation effects in durations that are brief enough that a large number of trials can be used in an experiment, which is necessary to collect meaningful EEG data (see Section 2.3).

In addition to these drifting grating stimuli being able to induce adaptation effects in an appropriate amount of time to be used with EEG studies, they are also of interest because it is not presently known whether they can induce the same event related potential components (ERP; see Section 2.3.1, below for more detail) that are known to be induced by reversals in perception of stationary ambiguous stimuli (when they are viewed intermittently, see Section 2.2.2.1 below; as noted and discussed in Kornmeier & Bach, 2012). Using these stimuli in this set of experiments, therefore, will provide some evidence regarding the potential sharing of mechanisms between reversals of ambiguous stimuli that are stationary and those that involve motion perceived from drifting sine wave gratings. Additionally, information can also be gathered on whether any ERP components associated with perceptual reversals of ambiguous motion stimuli differ according to whether that reversal occurred after a period of adaptation or was spontaneous. Using the ambiguous motion from sine wave gratings in this set of experiments will also allow this to be examined.

2.1.2.1 Stimulus parameters

Sine wave gratings can be configured in various ways to study different aspects of visual perception. Of relevance here are spatial frequency and contrast. Spatial frequency refers to the number of cycles (i.e., the dark and light bands) of the grating per unit of visual angle, usually measured in cycles per degree (cpd). Higher spatial frequencies have more black and white stripes in a given space, resulting in finer patterns, while lower frequencies have fewer cycles, resulting in broader patterns. Research has shown that gratings with high spatial frequencies may not be detected well under lower light conditions (e.g., Hess et al., 1990; such as those in an EEG laboratory). In terms of the effect of spatial frequency on motion priming and adaptation, Pinkus and Pantle (1997) found a small but significant effect, in that the gratings with higher spatial frequencies (above 1.4cpd) showed slightly less sensitivity to VMP. Therefore, the sine wave gratings used in the set of experiments here will use a spatial frequency of 1cpd as we expect this to lead to a robust sensitivity to motion at this frequency.

A second parameter to consider is contrast, which deals with the difference in luminance between the dark and light bands of the grating. High contrast gratings have a stark difference between the dark and light areas, making the pattern more pronounced, whereas low contrast gratings have a more subtle difference, making the pattern less distinct. Takeuchi et al. (2011) used a staircase procedure with drifting sine-wave gratings to investigate the minimal detectable contrast at both low and high retinal illuminance levels. They found that the average Michelson luminance contrast thresholds of their participants were 0.097 at high retinal illuminance and 0.83 at low retinal illuminance. Additionally, Kanai and Verstraten (2005) used a Michelson luminance contrast of 0.5 when creating their stimuli and it is this study that the set of experiments in this thesis is based upon. Therefore, a Michelson luminance contrast of 0.5 will be used when creating the sine wave grating stimuli to be used here.

2.2 Response choice

There are different considerations to be made when it comes to the choice of response in the set of experiments described in this thesis. Experiment 1 of Chapters 3 and 4 contain two block types – one containing pairs of directional followed by ambiguous motion, and one containing ambiguous motion only. There are slightly different considerations for each of these: the issue of when to mark the onset of a perceptual reversal (of relevance in the blocks containing only ambiguous motion) when analysing the EEG data from these trials, and the more general issue of how participants are asked to report their perception at the end of a trial.

2.2.1 Time reference considerations

EEG offers the capability to measure neural processes associated with perceptual reversals with millisecond accuracy, which could be crucial in determining whether these reversals are driven by bottom-up mechanisms or by more top-down cognitive processes. Nonetheless, the intrinsic nature of these perceptual changes makes it challenging to pinpoint an exact time reference for

EEG analysis. To address this, two main approaches have been developed, each having its own set of benefits and drawbacks. The following subsection will examine relevant studies in this context and compare their findings, then justify the choice of approach used in the set of experiments described in this thesis.

2.2.2.1 The Manual vs Stimulus Onset paradigms

As described in more detail in section 2.3.1 below, event-related potentials (ERPs) are created by isolating and averaging EEG data from around an event of interest. In experiments involving the perception of ambiguous stimuli the event of interest is typically the moment of a perceptual reversal. The issue here is that this is a purely subjective event, with no exogenous markers to reference. Therefore, to study these events using ERP and other analyses, researchers require an objectively observable marker for the subjective perceptual reversal. Exactly what to use as this marker is a subject that has been debated in the literature, and two main approaches to the problem have been developed. One approach is that the ambiguous stimuli are presented continuously, and participants are asked to respond with a button press (for example) as soon as they perceive a reversal. The button press is then used as the marker of the subjective reversal in perception and any analyses are relative to this. This approach is known as the manual response paradigm. The second approach involves presenting the ambiguous stimuli intermittently, with participants responding to indicate their perception of each presentation. This approach usually assumes that any reversals in perception occur at or near to the onset of the stimulus, and so it is this that serves as the marker for the subjective reversal and any analyses are instead time-locked to stimulus onset.

These two approaches will now be discussed in more detail, and the choice of approach used in this set of experiments will be justified.

2.2.2.1.1 The 'Manual Response' paradigm

As mentioned briefly above, the 'Manual Response' paradigm involves presenting an ambiguous stimulus continuously on the screen, and participants are asked to make some kind of manual response (e.g., key press, mouse click, button press etc.) when they experience a reversal in their perception. These manual responses are then used as a time reference with which to segment the EEG data. This is different from most ERP analysis pipelines in that the data of interest does not come *after* the event that serves as the time reference, but *before* it.

Studies that have used the manual onset approach alongside a type of ambiguous motion stimulus known as stroboscopic alternative motion (SAM) have found a positive ERP component occurring maximally at right parietal electrodes around 250ms before the manual response (Basar-Eroglu et al., 1993; Strüber & Herrmann, 2002). This component was also found with stationary ambiguous stimuli such as the Necker cube (Strüber et al., 2001; Mathes et al., 2006). These studies have interpreted this component as evidence for top-down processes as it appears to be closely related to the P300 component, which is known to reflect cognitive processes (e.g., Soltani & Knight, 2000). To infer more detail about the timing of this component, Strüber & Herrmann (2002) asked participants to continuously view SAM stimuli whilst magnetoencephalography (MEG) data was recorded. They used two variants of this stimulus: one ambiguous, and one disambiguated as the dots moved simultaneously in either a horizontal or a vertical direction. Using these stimulus variants, they induced either endogenous (using the ambiguous SAM; i.e., any reversals in perception are due to changes in brain activity and not physical stimulus changes) or exogenous perceptual reversals (using the disambiguated SAM; i.e., any reversals in perception are due to physical stimulus changes), and found that this positive, parietal component was present for both. Participants' reaction times after the exogenous reversals of the disambiguated stimuli were around 550ms. Thus, the temporal relationship between the three events was that the reversal event occurred, followed by the positive ERP component and then finally the button press. Therefore, the authors assumed that, as the P300-like ERP component occurs after the perceptual reversal itself, it is likely to reflect conscious recognition of the reversal.

In addition to ERP components, studies using the manual response approach have found effects in both alpha and gamma power bands. For instance, in addition to their P300-like ERP component, Strüber & Herrmann (2002) also found a sudden drop in alpha-band power between 300 and 200ms before the participants made their manual response to the exogenously induced reversals of the SAM stimuli. For endogenous perceptual reversals, rather than a sudden decrease, the authors found a steady continual decrease in the 1000ms before the manual response. The authors interpreted these findings as being indicative of a steady bottom-up destabilisation of the percept in endogenous reversals (i.e., over the period of 1000ms before the manual response) rather than the sudden destabilisation caused by the exogenously induced reversals. This alpha power effect was later shown to be in the lower bands only (6-10Hz; İşoğlu-Alkaç & Strüber, 2006). In addition to the effects found in alpha-band power, Basar-Eroglu et al. (1996) also found an increase in right-frontal gamma-band power within the

1000ms leading to the manual response, for responses indicating a perceptual reversal (compared to stability) only. The authors interpreted this result as evidence for the cognitive destabilisation processes thought to underpin reversals in perception. The manual response paradigm, therefore, has been used to reveal alpha power decreases alongside gamma power increases in the 1000ms leading up to the manual responses made to indicate reversals in perception. This period is thought to reflect the destabilisation of the currently active percept. This is followed by the P300-like parietal positivity ERP component between 500-250ms before the manual response, thought to reflect the conscious recognition of the reversal event. Thus, if the P300-like component reflects the conscious recognition of the perceptual reversal, the actual reversal event must have occurred before this (Kornmeier & Bach, 2012).

A problem, therefore, with the manual response paradigm is that results from studies using this approach can only point to quite a lengthy time-window in which the reversal itself may have occurred. The effects found in this window could be due to the reversal itself or any number of related processes such as response preparation. In addition, the very nature of the paradigm means that there is no discernible marker for the reversal event. Instead, data is backward averaged from the point of manual response. Studies have revealed that this backwards averaging process leads to a loss of signal due to an increase in the variability of individual ERP components based on a large amount of inter- and intra-participant variability in reaction time (e.g., Strüber et al., 2000; Strüber & Herrmann, 2002; Kornmeier and Bach, 2004b). Intra-participant variability is highlighted in Kornmeier and Bach's (2004b) study, where participants were continuously presented with one of two unambiguous variants of the Necker cube (top right vs bottom left front-facing cubes), whilst EEG data was recorded. Exogenous perceptual reversals were induced as the two variants would randomly switch, and participants were asked to respond to indicate this reversal in orientation as quickly as they could, using a key press. The median reaction time (from stimulus onset to key press) of these responses was 616ms, with an interquartile range between 530 and 733ms. When the EEG data was averaged relative to the onset of the stimulus switch (i.e., when a top-right as front-face variant was replaced with a bottom-left front-face variant), this resulted in several early visual ERP components related to the exogenous reversal. However, when that same data was what the authors call "backward averaged" relative to the participant's response, the majority of these early visual ERP components were lost, leaving only the later, P300-like parietal positivity. Whilst this was initially shown for exogenous reversals, the authors went on to show that the same issues arise when ERPs are formed from averaging using the manual response paradigm from endogenous reversal data (Kornmeier & Bach, 2012). Therefore, the effect of the intra-participant variability in reaction time was that the early ERP components were lost to the averaging process.

One of the aims of the EEG experiment described in this thesis is to investigate whether the established ERP components associated with perceptual reversals occur during reversals of ambiguous drifting sine wave motion. These components (the reversal positivity, RP, and reversal negativity, RN; see Chapter 1 and Section 2.3.1 below) occur before the P300-like ERP component, meaning that the backwards averaging process involved in the manual response paradigm is likely to lead to these two early visual components being missed (Kornmeier &

Bach, 2012). Therefore, this approach is not suitable to use as a means of time referencing the EEG data collected in Chapter 4.

2.2.2.1.2 The Stimulus Onset paradigm

In contrast to the Manual Response paradigm, the Stimulus Onset paradigm involves stimuli being presented intermittently and the onset of each stimulus being used to time-lock the ERP. The underlying assumption of this approach is that reversal events will occur at or very near to stimulus onset. This approach was taken by Orbach et al. (1966), who presented a Necker cube intermittently and varied the stimulus presentation time and inter-stimulus interval (ISI) to show that reversal rates (in number of reported reversals per minute) increased with ISIs up to 400ms where there was an average of around 40 reversals per minute. These results were later replicated by Kornmeier et al. (2002). ISIs longer than this led to decreasing reversal rates, a finding confirmed by further studies which went on to show that reversals can even be stopped altogether when ISIs are in the range of seconds (e.g., Maier et al., 2003).

The use of the onset paradigm to average EEG data was first implemented by O'Donnell et al. (1988). These authors presented participants with a mixed sequence of ambiguous and unambiguous Necker cube variants for 700ms, with an ISI of 3300ms. Participants were asked to report their perception of each stimulus at the end of each stimulus presentation with a key press. Therefore, there were two types of reversal in their experiment: endogenous (reversals occurring between two successive ambiguous Necker cube presentations) and exogenous (reversals occurring between two successive and opposite unambiguous Necker cube presentations). The authors found a late fronto-parietal positivity for both endogenous and exogenous reversals, which was shown to be larger and earlier in the case of exogenous reversals. This component was similar to the P300-like positivity found with the manual response paradigm (e.g., Basar-Eroglu et al., 1996; Strüber et al., 2002). However, Kornmeier and Bach (2004b, 2012) identified several limitations in O'Donnell et al.'s (1988) original design and conducted a partial replication of their study with these issues addressed. Firstly, Kornmeier and Bach (2004b) pointed out that O'Donell et al.'s (1988) unambiguous variants of the Necker cube were not truly unambiguous and could instead be perceived in two different orientations. Kornmeier and Bach (2004b) addressed this issue by using a Necker lattice with depth and shading cues, that successfully reduced the perceptual ambiguity. Secondly, Kornmeier and Bach (2004b) argue that the fact that reversal rates increase with ISI up to 400ms and then decrease with increasing ISI after this, suggests that there are two processes at work. One which drives the increase in reversal rate and a second that drives the decrease in reversal rates. Therefore, the authors claim that O'Donnell et al.'s (1988) choice of ISI at 3300ms may only be examining the second of these processes and therefore might not serve as a good model for continuous viewing. Therefore, the authors used an ISI of 400ms to make their paradigm as close as possible to continuous viewing but still allowing participants enough time to respond within the ISI. Kornmeier and Bach (2004b) used a presentation time of 800ms to allow sufficient time for the P300-like positivity to develop whilst also being brief enough to help prevent additional perceptual reversals within the stimulus presentation time. Finally, Kornmeier and Bach (2004b) asked participants to respond to indicate their perception in two conditions: one where they were asked to press a key only when they experienced the current stimulus as being different from the previous one but not if it was the same (the reversal condition), and one in which they were asked to do the opposite (the stability condition). Kornmeier and Bach's (2004b, 2006) partial replications of O'Donnell et al.'s (1988) experiment show that there were multiple ERP components associated with perceptual reversals, rather than only the P300-like component revealed by studies using the manual response paradigm. As shown in Figure 2.3 below, the data from these studies revealed 4 ERP components associated with endogenous reversals and 3 with exogenous reversals.

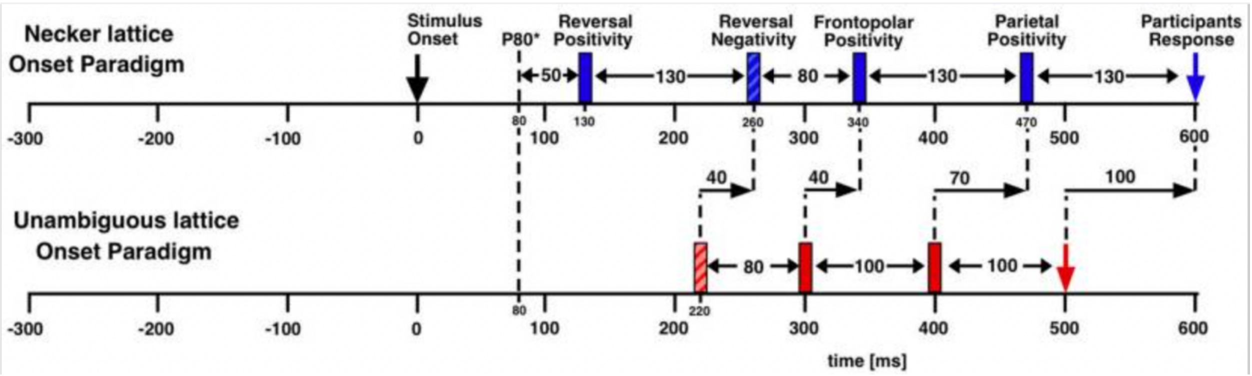

Figure 2.3 Representation of the ERP components associated with perceptual processing of Necker and unambiguous lattices. The timelines represent the sequence and timing of ERP components following the onset of stimuli for both a Necker lattice (top) and an unambiguous lattice (bottom), as reported by Kornmeier & Bach (2004). Stimulus onset is marked at time zero (0 ms). Following the onset, the key ERP components are identified: Reversal Positivity, Reversal Negativity, Frontopolar Positivity, and Parietal Positivity, culminating in the Participant's Response. Each component is depicted with a blue bar (indicating a positive peak) or a red hatched bar (indicating a negative peak), with their respective time windows and peak latencies in milliseconds (ms). Taken

from Kornmeier and Bach (2012), with data from Kornmeier and Bach (2004b, 2006).

Chapter 1 discusses these in more detail but briefly, these are the reversal negativity (RN), frontopolar positivity, and parietal positivity for exogenous reversals, and these plus the additional reversal positivity (RP) for endogenous reversals. Kornmeier and Bach (2004b, 2006, 2012) suggest that the RP was likely missed in studies using the manual response paradigm due to the variability in participants' reaction times but has been successfully reported in several other studies using the stimulus onset paradigm (Kornmeier & Bach, 2004a; 2005; 2006; Kornmeier et al., 2007; 2011 and Britz & Pitts, 2011). These results highlight the increased sensitivity of the stimulus onset paradigm to detect events that would be missed due to the backwards averaging and reaction time variability caused by the manual response paradigm. The RP and RN are investigated in Chapter 4 of this thesis, and as the RP is not able to be detected using the manual response paradigm due to its relatively poor temporal precision, the stimulus onset paradigm is necessary to investigate these components (Kornmeier & Bach, 2012). Additionally, multivariate pattern analysis (MVPA) is carried out on the EEG data gathered in the experiment described Chapter 4. The classification across time approach of MVPA compares identical time points within a trial (see Section 2.3.2, below). Therefore, for this to be successful, there needs to be a point in time with which to segment the EEG data used to train and test the classifier on, to reduce any variability between the events occurring in each trial. Additionally, ISI and presentation time have been robustly shown to be the key factors to manipulate to elicit priming and adaptation effects (e.g., Pinkus &

Pantle, 1997, Kanai & Verstraten, 2005). Therefore, it is vital that stimuli in the set of experiments described here, are presented intermittently, in order to manipulate these two variables empirically. The exact timing parameters used are thoroughly investigated and discussed in Chapter 3.

2.2.2 General response considerations

Previous studies have asked participants to respond to the pairs of adapting/ambiguous drifting sine wave motion stimuli by asking whether the two stimuli in each trial were moving in the same or a different direction (e.g., Kanai & Verstraten, 2005). This response is simple for participants; they press just one button, for example, at the end of a trial. Whilst this response does gather information on whether the participant experienced a reversal or stable trial, it does not allow any information beyond that to be extracted. For instance, information on the percept of *each* of the two stimuli presented is not collected. Instead, this is possible by asking participants to report their perceived direction of *each* of the two stimuli in the pair as Pinkus and Pantle (1997) did. Here, participants were asked to respond with a button to indicate motion to the left or the right for each of the two stimuli (‘left’, ’right’, for instance). Trials were then coded, post hoc by the experimenters, as reversal or stable according to whether the motion was perceived in the same direction or not. This dual response method means that trials can also be classified as “left followed by rightward motion reversal” and “left followed by leftward motion stable” trials, for instance. This is particularly important because the perception of motion in a leftward direction has been shown to elicit a different scalp voltage pattern than motion to the right (e.g., Bae & Luck, 2019). Therefore, collapsing trials across reversal/stable in either direction would preclude any examination of direction-specific reversal-related activity and allow examination of only non-directional reversal-related activity. Consequently, in Experiment 1 of Chapter 3 and the experiment described in Chapter 4, this dual response method was used.

2.3 Electroencephalogram (EEG) data acquisition and analysis

EEG data is particularly advantageous for examining the neural mechanisms underlying the perception of ambiguous motion, especially in the context of adaptation. The high temporal resolution of EEG allows for the tracking of rapid fluctuations in brain activity that accompany the perception of motion, which is crucial when studying phenomena that unfold on the scale of milliseconds, such as the immediate effects of adaptation on perception. Furthermore, EEG's sensitivity to the different frequency bands of neural oscillations enables the exploration of various aspects of neural processing, from the initial sensory responses to higher cognitive functions involved in interpreting motion. This allows for a detailed analysis of how adaptation might alter the temporal dynamics of the neural response to motion perception.

2.3.1 Event-related potential (ERP) analysis

One aim of this thesis is to explore whether previously established markers of perceptual reversals are present during perception of the ambiguous drifting sine wave motion; specifically, the RP and RN (Kornmeier & Bach, 2012). These two ERP components are thought to be markers of different stages of processing during perceptual reversals (see Chapter 1). Therefore, one of the aims of Chapter 4 is to determine whether these two components are present during perceptual reversals of ambiguous drifting sine wave stimuli, and whether they differ in their amplitude or latency (or indeed occur at all) compared to their appearance during the perception of stationary ambiguous stimuli.

ERP analysis is a technique used in electrophysiological studies, particularly within the field of cognitive neuroscience, to measure the brain's electrical response to specific sensory, cognitive, or motor events. ERPs are derived from EEG data by aligning segments of the EEG to the onset of an event (such as a stimulus or a response) and then averaging these segments across many trials to enhance the signal-to-noise ratio (Luck, 2014). EEG data around this event is isolated to create an epoch for every instance of the event. This typically includes a portion of data from before the event, to be used as a baseline, and then a portion of data after the event of interest. This data is then averaged within conditions, so that noise/activity unrelated to the event-of-interest within the data is cancelled out, and only activity that is systematically time-locked to the event of interest remains. This averaging process reveals a series of voltage deflections, known as ERP components, which are thought to reflect specific neural processes associated with the event.

Each ERP component is characterised by its polarity (positive or negative), latency (the time from the event onset to the peak of the component), and scalp distribution. The components are often named according to their polarity and the approximate latency. For example, an N170 component would be a negative-going ERP component peaking around 170 milliseconds after the stimulus onset, typically associated with the processing of faces.

ERP analysis is particularly useful for understanding the timing and sequence of cognitive processes because different ERP components are linked to different aspects of perception, attention, memory, and decision-making. By comparing the ERPs elicited by different types of stimuli or in different experimental conditions, it is possible to infer how these cognitive processes differ across conditions. Luck (2014) emphasises the utility of ERP analysis in identifying when particular cognitive processes occur. This is particularly useful when the behavioural response of the participant is consistent across conditions, but the underlying neural mechanism might be different.

For example, Pitts et al. (2008) used ERP analysis to compare the components associated with perceptual reversals of the Necker Lattice that were voluntary (i.e., participants were asked to voluntarily induce a reversal) with those that were spontaneous (i.e., participants passively observed the stimulus and any reversals occurred involuntarily). Pitts et al. (2008) found that there were differences in the latencies, amplitudes and scalp topographies of the ERP components associated with voluntary vs spontaneous perceptual reversals of the Necker Lattice. This study highlights the ability of ERP analysis to discern between different underlying neural mechanisms despite the behavioural response of the participant remaining the same across experimental conditions (i.e., voluntary vs spontaneous reversals).

Thus, in the set of experiments described in this thesis, ERP analysis is a sensible approach to adopt to determine whether the mechanisms underpinning perceptual reversals that occur after adaptation are different to those that occur without a preceding adaptation period. However, ERP analysis is not suitable for investigating *global* brain activity patterns following a period of adaptation and determining whether these patterns are similar to those observed during perceptual reversals that occur without prior adaptation. Typical univariate ERP analysis approaches excels in pinpointing the timing and location of brain responses to specific events, but it does not capture the full-brain activity patterns that are necessary to compare the overall neural dynamics associated with adaptation and spontaneous perceptual reversals.

2.3.1.1 Limitations of ERP analysis

According to Luck (2014), the limitations of ERPs stem from their relatively poor spatial resolution and the difficulty in interpreting the waveforms. ERPs are not ideal for precisely locating the sources of neural activity within the brain, as electrical signals are distorted when they pass through various tissues to the scalp. This issue is compounded by the 'inverse problem', where multiple brain activity patterns could result in the same scalp patterns, making accurate localisation challenging.

Additionally, ERPs are most effective for analysing brain responses to discrete events and may not capture continuous cognitive processes or the complex patterns of ongoing neural communication effectively. By their nature, ERP components are only measurable under certain conditions. For example, many neurons must be activated concurrently, or the individual neurons must have similar orientations in the brain (Luck, 2014). The ERP waveform represents activity that is time-locked and phase-locked to an event, but responses that are not event-related or vary in phase can be obscured by the averaging process, which can eliminate non-phase-locked activities such as ongoing oscillations. Also, overlapping components from closely timed successive cognitive processes can obscure the determination of which ERP components are related to specific cognitive processes.

Furthermore, the ERP technique involves averaging responses across multiple trials to reduce noise, whilst assuming that the response is consistent across trials. This averaging process may not reflect variations in cognitive or perceptual state on a trial-by-trial basis and reduces the ability of the analysis to utilise the information contained in the activity from single trials when detecting effects.

As mentioned previously, ERP analysis is also univariate in nature. This means that activity across the whole scalp is not considered and is instead reduced to an average of only a few localised positions. Given that research has shown that widespread neural mechanisms are likely to be involved in the occurrence of perceptual reversals, particularly with ambiguous motion stimuli (e.g., Piedimonte et al., 2015), considering a small subset of the available information is likely to exclude information that is relevant to the experimental manipulation. In the set of experiments described in this thesis, an aim is to establish whether an adapted brain state, induced by exposure to directional motion, occurs during passive viewing of ambiguous motion. This adapted state is likely to involve multiple scalp locations that should be considered in relation to one another, which is not possible to consider using ERP analyses. Therefore, in addition to using ERP analysis to determine whether the known components associated with perceptual reversals are present during reversals of ambiguous sine-wave motion, a different type of analysis is required to answer all the aims of this thesis. Multivariate pattern analysis, as the name suggests, is a type of analysis that considers multiple variables in parallel, at any given timepoint. In the context of EEG data, this type of analysis considers the whole-brain pattern of activity at each time point in a trial without the need for averaging. To answer questions about an adapted brain state therefore, this type of analysis is more suitable.

2.3.2 Multivariate pattern analysis (MVPA)

Multi-*voxel* Pattern Analysis has been applied to functional magnetic resonance imaging (fMRI) blood oxygen level dependent (BOLD) responses to decode cognitive states from recorded neural activity (see Haynes & Rees, 2006 and Haynes, 2015, for reviews) for over a decade. As the name here suggests, activation patterns across multiple voxels are taken into consideration, in contrast to the univariate approaches that examine the time course of each voxel independently. This allows patterns of activity over the entire brain (or indeed smaller subsections of it) to be analysed rather than only focusing on a single area at a time. In practice this leads to researchers being able to infer participants' cognitive states from the BOLD response, for instance which stimulus is currently being attended (e.g., Haxby et al., 2001; Sterzer, Haynes & Rees, 2008) or which object is currently being imagined (Stokes et al., 2009). Recent years have seen the technique applied to neuroimaging more broadly, and the term MVPA is more widely known as multi*variate* pattern analysis, as it encompasses a range of methods that all take into consideration the relationship between multiple variables rather than treating them as though they are independent of one another. In M/EEG studies, the electrodes serve as variables rather than the voxels in fMRI studies, but the practice is the same. Just as with fMRI studies, traditional univariate analyses of EEG data often examine individual electrodes or averaged signals across specific electrode clusters, thus potentially missing the distributed patterns of neural activity. In contrast, MVPA of EEG data leverages the high-dimensional data space formed by multiple electrodes, offering a means to detect complex, spatially distributed neural signatures that might be overlooked in conventional analyses (Hebart & Baker, 2018).

A typical multivariate analysis pipeline begins with training a machine learning classifier, such as a Support Vector Machine or Linear Discriminant Analysis based classifier (SVM and LDA, respectively). The set of trials is divided into two sets: a training set and a testing set. In EEG studies, from the training set, the classifier is provided with EEG scalp voltage data across multiple electrodes (i.e., “features” in machine learning terminology) at a given time point and the condition label (i.e., “class” in machine learning terminology) associated with the scalp pattern on each trial. From this, the classifier learns a decision boundary to discriminate between conditions based on the pattern of scalp EEG. The trained classifier is then tested by being provided with the remaining portion of the EEG scalp voltage data without class labels (for an illustration of this process, see Figure 2.4, below). It makes predictions for the test set and an accuracy is calculated. This is repeated independently, for each time point in the data.

It is important to train and test the classifier on independent data (known as cross-validation) to avoid the phenomenon known as “double dipping’. If the classifier is trained and tested on the same set of trials, then there is a risk of overfitting the model based on noise that is present in the data and leading to an artificially inflated prediction accuracy (e.g., Kriegeskorte et al, 2009; Fahrenfort et al., 2017). Hence, during training, 10-20% of the trials are kept back to independently test the classifier on. The ability of the classifier to correctly predict class from this test data set alone is known as decoding accuracy. For an experiment with two classes, a decoding accuracy of 50% is the level that would be expected by chance. Therefore, timepoints with decoding accuracies significantly greater than chance level are generally accepted as one in which the effect has been detected – in other words, sufficient information about the experimental manipulation (e.g., perception, stimulus category, attentional state, etc.) is contained in the patterns of brain activity (e.g., Kriegeskorte et al., 2009; Grootswagers, Ward & Carlson, 2017; Hebart & Baker, 2018). For an experiment with multiple participants, individual decoding accuracies are typically averaged to produce a 'grand average' plot of decoding accuracy across the time course of the trial.

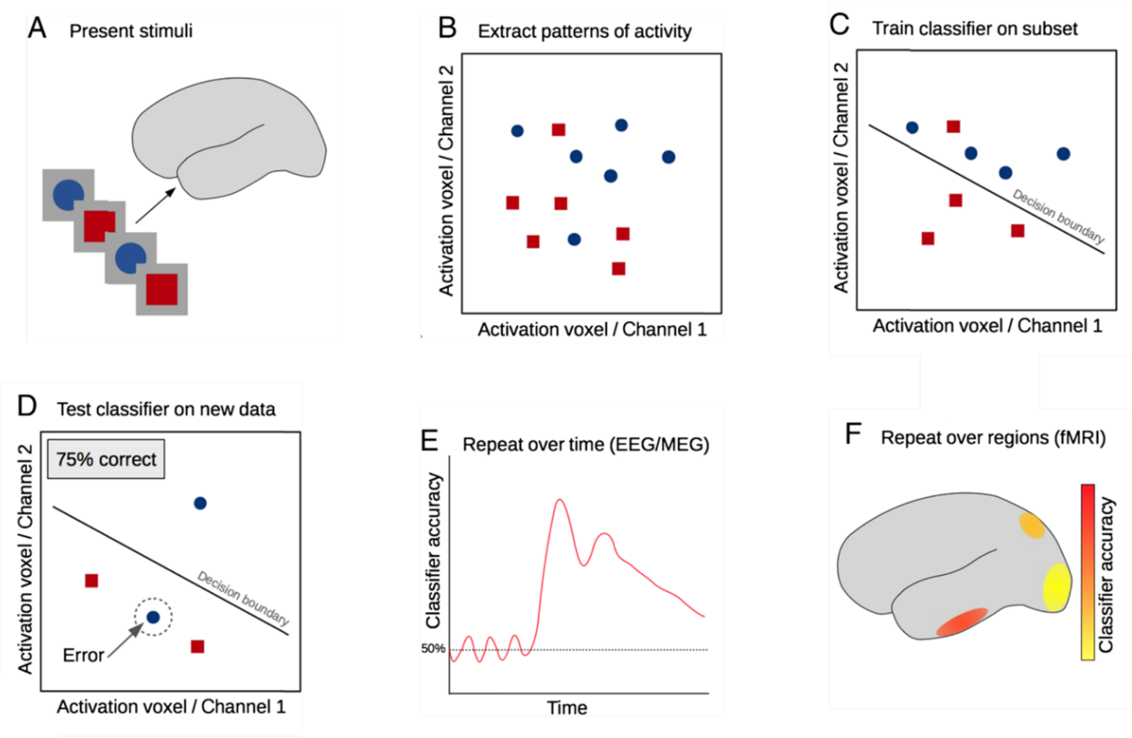

Figure 2.4: Schematic representation of the multivariate pattern analysis (MVPA) process. This figure outlines the key steps involved in conducting MVPA with neuroimaging data. (A) Stimuli are presented to participants while neural activity is recorded. (B) Patterns of activity are extracted from the recorded data,

with the different shapes here representing different conditions. The data points are represented in a multi-dimensional space with electrodes/channels as the dimensions for EEG data. (C) A classifier is trained on a subset of the data to differentiate between the conditions based on activity patterns. (D) The trained classifier is then tested on a separate set of data, with the percentage correct indicating its predictive accuracy. (E) In the context of EEG/MEG data, the classifier's accuracy is assessed over time, producing a time-course plot that reflects how neural discrimination of conditions changes. (F) For fMRI data, the procedure is repeated over different brain regions, which are color-coded based on the classifier's accuracy in each region, illustrating the spatial distribution of informative neural patterns. This process enables the identification of temporal and spatial patterns of brain activity that correlate with specific cognitive processes or stimuli presentations. Taken from Grootswagers et al. (2017).

The use of MVPA in neuroimaging studies has led to considerable changes in the way that researchers understand the functional organisation of the brain. For instance, fMRI studies that use univariate analyses have long suggested that the brain is organised into distinct areas that are each specialised for different functions. For example, the perception of specific stimulus categories such as visual words (Cohen et al., 2002; Kleinschmidt & Cohen, 2006), or faces (e.g., Kanwisher et al., 1997; Kleinschmidt & Cohen, 2006; Hesselmann et al., 2008). These studies have shown that, for example, an area known as the visual word form area (VWFA) shows increased activation to visual words (Cohen et al. 2002) and is known to show reduced or absent responses in patients with pure alexia (Kleinschmidt & Cohen, 2006; pure alexia is the inability to read written text, whilst other aspects of language processing are preserved). Additionally, the fusiform face area (FFA) robustly shows higher activation in the identification of faces compared to non-face objects, leading to conclusions that the FFA is critical for the perception of faces (e.g., Puce et al., 1996; Kanwisher et al., 1997; Marotta et al., 2001). However, the use of multivariate analyses has since altered the view that the brain is organised into discrete, category specific areas. For instance, Haxby et al., (2001) used MVPA to decode participants' perception from the pattern of BOLD responses in an fMRI study. Specifically, whether they were viewing faces compared to non-face objects. They found that they were successfully able to decode perception based on the pattern of activity across well-established 'face specific' regions and critically, they were also able to do this when using the patterns of activity from regions that were outside of these category specific areas. Additionally, the authors found patterns of activity that could discriminate all categories within brain regions that had previously been thought to be highly category specific. For instance, the pattern of activation in the FFA contained enough information about non-face objects for it to be used to successfully discriminate between different stimuli, suggesting that the FFA is not only specific to human faces as the results from previous univariate studies had suggested. Instead, the results from Haxby et al. (2001) suggest that that the neural representation of different object categories is widespread throughout the brain and very much overlapping, rather than localised and discrete. The results from this study particularly highlight the ability of MVPA to pick up on differences based on whole-brain activity patterns across multiple voxels that univariate analyses might miss, as they focus on comparing relative activity levels of individual voxels independently.

In EEG studies, a similar increase in MVPA's ability to detect otherwise latent effects has been demonstrated when it is compared to that of more traditional univariate techniques such as ERP analysis (see Section 2.3.1 above). For instance, List et al. (2017; see also Chapter 1) showed that univariate and multivariate analyses were both able to reveal an effect of perceptual state from EEG scalp activity. Specifically, both analyses could determine a difference between left versus right stimulus location, faces versus Gabor stimuli, and upright versus inverted faces. However, in a second experiment, the authors were able to show that only MVPA was able to successfully distinguish between local versus global attentional states. Univariate analyses failed to detect any significant effect of attentional state on scalp voltage activity. The authors highlight that their results demonstrate that MVPA reveals distinctive spatiotemporal patterns of neural activity discriminating between behavioural states, making it a sensitive tool for characterising the neural correlates of perception and attention and therefore a sensible choice for analysing the EEG data collected in the set of experiments described in this thesis. Additionally, Ronconi et al. (2017) used MVPA to train a classifier “searchlight” on the whole electrodes/frequencies/times domain. The “searchlight” technique involves systematically moving a spherical cluster (the "searchlight") across the brain volume in an fMRI study, or across sensors in EEG/MEG studies, to identify regions or sensor clusters where the local patterns of activity can reliably predict experimental conditions or behavioural outcomes. Using this technique, Ronconi et al., (2017) found that perception of a bistable stimulus could successfully be decoded from the phase of oscillations in the pre-stimulus period. Analysing the pre-stimulus period using univariate analyses is not possible in this way, as activity during this period this is usually flattened by the averaging process (e.g., with ERP analysis the pre-stimulus period is typically used as a baseline, see Section 2.3.1 above). MVPA has also been shown to be able to detect differences in participants' perception much earlier than traditional ERP analyses, making it a sensible method to use to detect even brief temporal effects in EEG data (Cauchoix et al., 2012; 2014).

In addition to the findings discussed above, the temporal resolution of EEG (i.e., the range of milliseconds) allows MVPA methods to be applied dynamically across different time points. This opens the possibility to explore how distributed patterns of neural activity evolve over time, thus offering a dynamic view of information processing that is particularly relevant for tasks involving perception, attention, and action (King & Dehaene, 2014). The temporal generalisation technique devised by King and Dehaene (2014) extends the application of MVPA, particularly for EEG and MEG data analysis. This method takes a machine learning algorithm, such as SVM or linear discriminant analysis (LDA), and trains it on brain activity data at a specific time point to distinguish between different experimental conditions (as described above). This trained classifier is then tested not just on the data from the time point it was trained on but across all time points in the experiment. The outcome is a temporal generalisation matrix where the rows correspond to the training time points and the columns to the testing time points, revealing the classifier's accuracy over time (see Figure 2.5, below). The diagonal of this matrix represents instances where the brain activity pattern accurately predicts experimental conditions at matching training and testing times, while the off-diagonal elements indicate how these patterns generalise across different times. By analysing the generalisation matrix, researchers can infer the temporal evolution of the neural representations underlying cognitive processes, offering a dynamic view of how information processing unfolds in the brain over time.

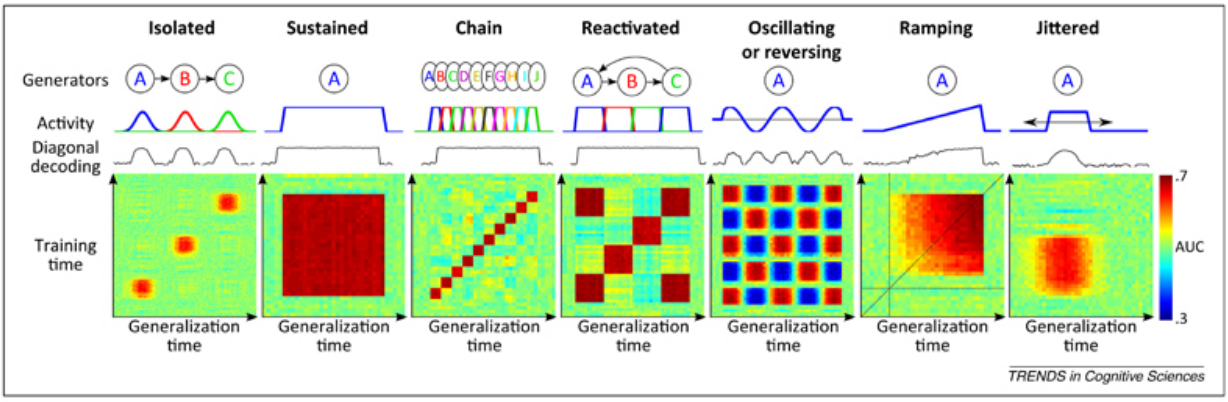

TRENDS in Cognitive Sciences

Figure 2.5: Illustrative temporal generalisation matrices for various cognitive processing scenarios. This figure showcases the application of temporal generalisation technique in decoding patterns of brain activity over time for different types of cognitive processes. Each matrix corresponds to a unique processing type: isolated, sustained, chain, reactivated, oscillating or reversing, ramping, and jittered activities. The 'isolated' panel shows high decoding accuracy only at specific time points, indicating transient neural representation. The 'sustained' panel depicts uniform decoding accuracy, reflecting stable neural activity over time. The 'chain' illustrates sequential processing stages, while 'reactivated' displays a return of a neural signature at a later time, suggesting memory or reactivation processes. 'Oscillating or reversing' implies cyclical neural activity, 'ramping' suggests a gradual increase in specific neural activity, and 'jittered' indicates variability in the timing of neural responses. These patterns provide insight into the temporal dynamics and flexibility of cognitive processes as captured by the temporal generalisation approach. Taken from King and Dehaene (2014).

In addition to decoding perceptual experience from brain activity, MVPA has the potential to offer deeper insight into the mechanisms mediating perceptual reversals and how these unfold across time. A technique known as cross-classification, or cross-decoding, allows one to make inferences about the similarity of neural processing involved in one or more experimental conditions (e.g., Kaplan et al, 2015). This cross-decoding approach assumes that if similar neural processes are involved in two different experimental conditions (e.g., perception and imagery) then the two conditions will have similar EEG scalp distributions. Thus, for instance, a classifier trained to discriminate different stimuli (e.g., faces vs. houses) in perception should also perform well at discriminating imagined stimuli. Alternatively, if the two conditions involve different mechanisms, then their scalp distributions will differ and a classifier trained on one condition (e.g., perception) will decode the other condition (e.g., imagery) with poor or at chance accuracy. To the extent that two conditions involve similar electrophysiological mechanisms, one should observe above chance cross-decoding with higher accuracy indicating greater overlap of mechanisms. This approach to MVPA has proven successful in the literature. For instance, in an fMRI study, Stokes et al. (2009) presented participants with the letter “X” or “O” in one condition. In the second condition they asked participants to only imagine the letter “X” or “O”. They found that a classifier trained on data from the lateral occipital complex during trials in the visual presentation condition could correctly decode participants’ imagined letter from data during the trials in the imagined condition. In a follow up study, Stokes et al. (2011) showed that this was also the case even when the letters were physically presented in different parts of the visual field. Cross-decoding has also revealed similar neural patterns involved in motor execution and action observation. For instance, Oosterhof et al. (2012a) recorded fMRI data whilst their participants performed and observed two object-directed actions ("lift" and "slap"). The actions were presented and executed in both first- and third-person perspectives. The experiment was designed to distinguish between actions observed and executed by participants, and whether there was any similarity in neural encoding based on the perspective of the observed action. The authors found that the researchers were able to identify distinct patterns of brain activity in the ventral premotor cortex (PMv) that corresponded to specific actions being either observed or executed. Crucially, the cross-classification approach allowed for the comparison of neural patterns across different modalities (visual observation and motor execution) and perspectives (first-person and third person). The findings revealed that actions observed from a first-person perspective produced significant cross-modal neural patterns in the PMv, mirroring those during action execution. In contrast, this cross-modal congruence was not observed for actions perceived from a third-person perspective.

MVPA, therefore, is a good choice of analysis technique to address the research questions posed in this thesis. To directly assess the role of adaptation in spontaneous reversals in perception, MVPA cross-decoding will allow the brain state during reversals due to adaptation to be compared to that occurring during spontaneous reversals.

2.3.3 Analysis of decoding accuracies

2.3.3.1 Cluster-based permutation tests

Cluster-based permutation tests have increasingly been employed to evaluate the time-course of decoding accuracies generated by MVPA when analysing EEG data. This statistical approach offers several methodological advantages that can enhance the rigour and interpretability of findings.

Typically, MVPA generates a time-course of decoding accuracies that signify how well specific cognitive states or experimental conditions can be differentiated at each time point. However, these decoding accuracies are subject to multiple comparisons across time points, raising the risk of Type I errors. Therefore, these multiple comparisons need to be corrected. There are many approaches to this correction which include permutation tests (e.g., Blair & Karniski, 1993), cluster-based permutation tests (e.g., Maris & Oostenveld, 2007) and false discovery rate control procedures (FDR; e.g., Benjamini & Hochberg, 1993; Benjamini, Krieger, & Yekutieli, 2006). The efficacy of such corrections in mass univariate ERP analysis was investigated empirically in a series of simulations by Groppe et al, (2011b). Here, realistic EEG background noise from real participants was used as a basis to simulate 1000 ERP experiments for each of 4 simulated ERP effects. These were: the N170 (Bentin, Allison, Puce & Perez, 1996; Bentin, Mouchetant-Rostaing, Giard, Echallier, & Pernier, 1999), the P3 (Bentin et al., 1999), combined N170/P3 (i.e., both components were added to the data), and ERP null effects (i.e., no ERP effect was added to the null effect data). The authors then analysed the data from each of the 1000 simulated experiments for each of the 4 conditions, by concurrently applying two-tailed, one-sample *t*-tests to all scalp channels. They then used 6 commonly cited multiple comparison correction methods to correct for the multiple comparisons made. These were: Bonferroni-Holm, t-max permutation, maximum cluster-level mass permutation, permutation-based procedures, and 3 FDR control methods (as described in Benjamini & Hochberg, 1995; Benjamini & Yekutieli, 2001; Benjamini, Kreiger, & Yekutieli, 2006). The number of false positives for each correction method across each of the 4 simulated effect types was systematically counted to establish false positive rates for each method. This was used to evaluate Type I error risk of each of the 6 correction methods, for each of the 4 simulated effect types. The simulation showed that all 6 procedures performed with sufficient conservativism so as not to inflate Type I error rate to a concerning extent. In a companion to this simulations paper, the authors made recommendations following their investigation (see Groppe et al., 2011a). Of relevance to the EEG experiments described in this thesis, their paper made recommendations for the use of cluster-based permutation tests (Maris & Oostenveld, 2007).

The cluster-based permutation approach itself involves the identification of clusters of adjacent time points where decoding accuracies are above a certain threshold. The sum of the accuracies within these clusters is then compared to a null distribution generated by permuting the labels of the conditions or states being decoded. This offers a rigorous way to test whether the observed patterns of decoding accuracies are statistically significant or merely a product of noise in the data.

The application of cluster-based permutation tests also enables researchers to identify temporally contiguous periods of significant decoding, which can be critical for understanding the dynamics of cognitive processes. For instance, this technique allows for the definition of time intervals where the brain consistently distinguishes between different experimental conditions, providing a more nuanced view of the temporal aspects of cognitive processing (e.g., King & Dehaene, 2014; see also Section 2.3.2 above). However, Sassenhagen and Draschkow (2019) critique the prevalent use of cluster-based permutation tests in cognitive neuroscience, particularly for MEG/EEG data. They argue that while these tests proficiently address the multiple comparisons problem and maintain Type I and Type II error rates, they may lead to overinterpretations regarding the temporal and spatial precision of effects. The authors demonstrate that these tests do not substantiate claims about effect latencies or locations, potentially leading to misleading conclusions about the neural underpinnings of cognitive processes. However, they acknowledge that cluster-based permutation tests are a powerful tool in handling the multiple comparisons problem.

Additionally, the cluster-based permutation approach is non-parametric, making it less reliant on assumptions about the underlying distribution of the data. This flexibility makes it particularly suitable for EEG data, which may not always conform to the assumptions of parametric tests (Groppe, Urbach, & Kutas, 2011). Incorporation of cluster-based permutation tests in the analysis of MVPA-generated decoding accuracies offers a methodologically sound and statistically rigorous approach. By controlling for multiple comparisons, enabling the identification of significant temporal clusters, and offering a non-parametric testing option, this method enhances the interpretability and validity of findings derived from MVPA of EEG data. As such, it is a suitable choice for assessing the statistical significance of the decoding accuracies generated from the MVPA in Chapter 4 of this thesis.

2.3.3.2 Error rates

In the context of MVPA and cluster-based analyses of EEG data, the risk of Type I errors and inflated FWER is a significant concern. The use of simulations can effectively address this by providing a controlled environment to rigorously evaluate statistical methods. By simulating data with known characteristics, researchers can determine the likelihood that their analytic pipeline will yield false positives, thereby offering a quantitative measure of the method's reliability (Pernet, Chauveau, Gaspar, & Rousselet, 2011).

Specifically, simulations can be employed to generate synthetic EEG data sets under null hypothesis conditions, i.e., data sets where no differences between conditions or cognitive states are present. These synthetic data sets can then be subjected to the same MVPA and cluster-based statistical analyses as the actual experimental data. By repeating this process numerous times, researchers can establish an empirical null distribution for decoding accuracies or other test statistics, thereby enabling a more robust assessment of the actual results' statistical significance (Maris & Oostenveld, 2007).

Furthermore, simulations offer a mechanism to gauge the power and sensitivity of the chosen methods. By introducing known effects or differences into the simulated data, it is possible to assess how often the analytic pipeline correctly identifies these effects, thereby providing an estimate of Type II error rates in addition to Type I and FWER assessments. For example, Brooks, Zoumpoulaki and Bowman (2016) used simulations to examine the impact of different ROI selection methods on the Type I error rate and statistical power in ERP analysis. They compared traditional a priori-based ROI selection, which relies on hypotheses or independent information and may miss effects due to experiment-specific variations, with data-driven ROI selection, which uses the data under analysis to determine ROI positions. In the set of simulations carried out, the authors used simulated null-effect EEG data to assess Type I and familywise error rates, but also added simulated effects at known time points, to that null data. This method can be used to assess Type II error rates of other statistical methods, such as the cluster-based statistics used both widely in the current literature on the MVPA of EEG data and in this thesis.

One of the aims of this thesis is to assess the Type I (and familywise) and Type II error rates of MVPA pipelines when applied to EEG data. To assess this quantitatively and validly, simulations similar to those carried out by Brooks et al. (2016) can be employed but on a multivariate, rather than mass-univariate, scale. Simulating null-effect EEG data and inserting a multivariate effect of known size, can help to determine the power of cluster-based permutation tests to detect effects of this nature.

Chapter Three: Optimum temporal parameters required to induce the Rapid Motion Aftereffect and Visual Motion Priming

3.1 Introduction

The perception of visually ambiguous stimuli is often only stable for a short time. and spontaneously fluctuates between two or more mutually exclusive outcomes (e.g., Blake & Logothetis, 2002; Kornmeier et al., 2009; Leopold & Logothetis, 1999; Schwartz et al., 2012). These stimuli are considered useful in the study of perception as perceptual alternations occur despite a physical stimulus that remains unchanged, therefore providing an insight into neural activity linked to subjective experience irrespective of visual input. Consequently, many types of visually ambiguous stimuli have been studied, including stationary stimuli such as the Necker cube (Necker, 1839) and motion stimuli such as ambiguous sine-wave gratings (such as those used by Kanai & Verstraten, 2005; Pinkus & Pantle, 1997).

Given that alternations or "reversals" in subjective perception of these stimuli occur spontaneously, several theories have been put forward to attempt to explain this. One such theory is the Neural Fatigue Hypothesis, covered fully in Chapter 1 but will be briefly outlined again here.

3.1.1 The Neural Fatigue Hypothesis

Köhler (1940) theorised that the perception of ambiguous stimuli induced 'field' effects in the brain. He suggested that all visual stimuli are underpinned by a visual 'current' in the brain, that behaves in the same way as a direct electrical current, and that viewing a particular figure causes changes to the resistance of brain tissue that extends beyond the locus of the current percept. Köhler (1940; and later Köhler & Wallach, 1944) therefore proposed a theory of "neural satiation". This original concept has evolved into the "neural fatigue" (or neural adaptation; Long & Toppino, 2004) and “cross inhibition/competition” model. This is based on the idea that each interpretation of an ambiguous stimulus has its own neural representation, which inhibits the representations of other percepts (Toppino & Long, 1987). Following periods of prolonged stimulation, the neural representations underpinning the currently dominant percept become “fatigued” and so are no longer able to inhibit the alternative. At this point a threshold is reached whereby the neural representation of the alternative percept is less fatigued than that of the current one, so perception is reversed. The idea of this neural competition being the underlying mechanism of multistable perception has since been investigated by modelling neuronal behaviour using computational neural networks. Such work has provided support for the idea that percepts compete with each other through mutually inhibitory networks operating at a low level, in the way that the Neural Fatigue Hypothesis suggests (e.g., Noest et al., 2007; Shpiro et al., 2009). Neurophysiological studies have also revealed that the properties of motion sensitive neurons are changed when exposed to motion in a certain direction, in a manner consistent with neural fatigue (Pavan, 2009). More recently, these models have been examined by Kogo et al. (2021) using a hybrid system whereby a pair of physical pyramidal neurons interacted via a computer simulated, mutually inhibitory neural circuit. This showed that the dynamics of the pair of neurons in the system were consistent with the properties that the Neural Fatigue Hypothesis predicts.

3.1.2 Priming and adaptation

Behavioural studies examining the effects of pre-exposure to unambiguous variants of ambiguous stimuli have provided support for the Neural Fatigue

Hypothesis. For instance, presenting an unambiguous face stimulus before Rubin's ambiguous faces-vase stimulus can systematically bias how the ambiguous stimulus is perceived. These studies suggest two main effects: priming, which involves a bias towards perceiving the ambiguous stimulus in the same way as the pre-exposed stimulus (e.g., presentation of an unambiguous face followed by a face interpretation of the ambiguous faces-vase stimulus; e.g., Long & Moran, 2007). In contrast, adaptation involves a bias towards perceiving the unambiguous stimulus in a different way than the pre-exposed stimulus (e.g., presentation of an unambiguous face followed by a vase interpretation of the faces-vase stimulus; e.g., Long et al., 1992, Long & Olszweski, 1999, Pavan et al., 2009).

Presentation duration, inter-stimulus interval (ISI), intensity (i.e. brightness), and size of the pre-exposed stimulus have been found to be critical for the induction of one or the other of these effects (Kanai & Verstraten, 2005; Long et al., 1992; Long & Moran, 2007; Takeuchi et al., 2011). Shorter pre-exposure to the unambiguous stimulus tends to have a priming effect whilst longer exposure durations typically lead to adaptation effects. These durations vary between different types of stimuli. For stationary ambiguous stimuli such as the Necker cube (Necker, 1839) for example, priming effects have been shown after exposure to an unambiguous variant of the stimulus is presented to participants for one second, whereas adaptation effects using the same stimulus takes as long as 150 seconds (e.g., Long et al., 1992). For motion stimuli, such as moving sine-wave gratings (as used by Pinkus & Pantle, 1997), priming effects have been shown after pre-exposure durations of as little as 0.08s and adaptation effects after just 0.64s (e.g., Kanai & Verstraten, 2005; Takeuchi et al., 2011).

Kanai and Verstraten (2005) highlighted the relationship between pre-exposure duration and ISI when evoking either visual motion priming (VMP) or motion aftereffects. The authors used moving sine-wave gratings to create directional and ambiguous motion stimuli. Directional motion was created by shifting the phase of the sine wave by ±90º per frame, leading to energy-based motion either to the left or to the right. Ambiguous motion was created by shifting the phase of the sine wave by 180 º per frame, meaning that successive frames could be interpreted as motion to the left or to the right (see also Figure 3.1 below).

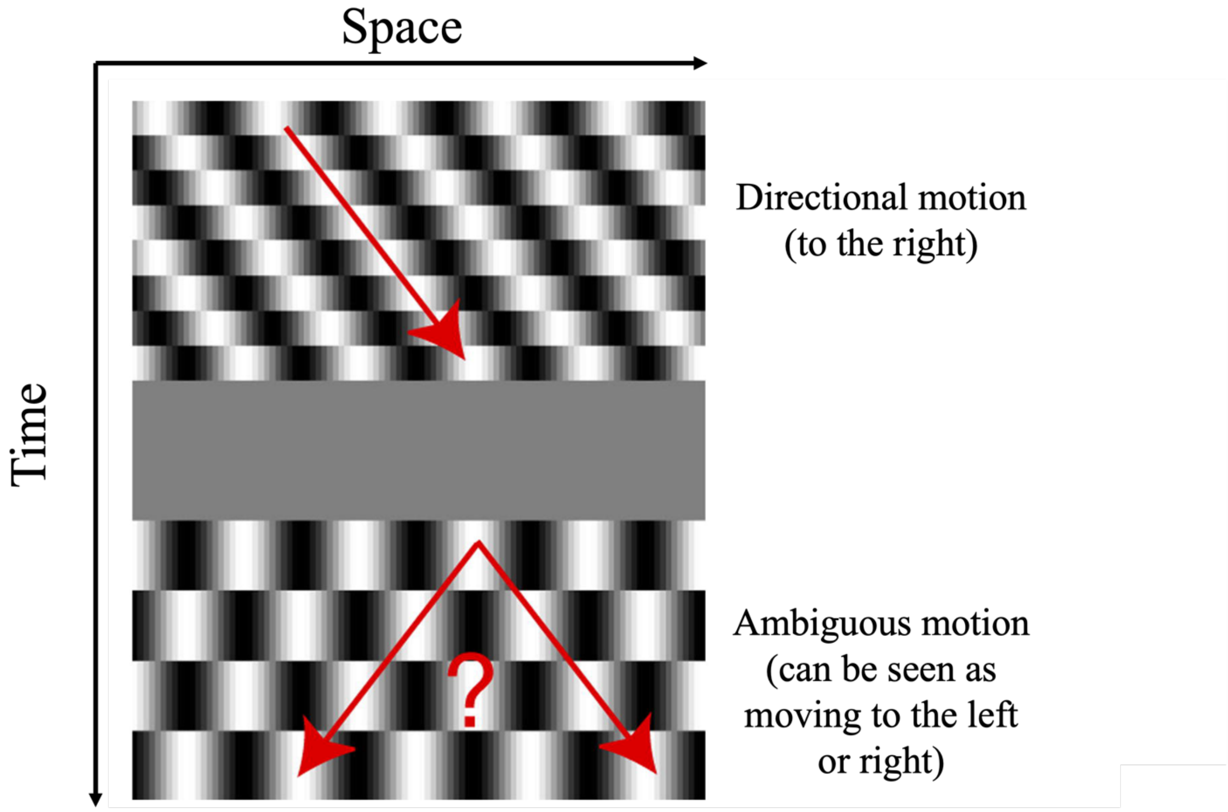

Figure 3.1 Illustration demonstrating how the phase of sinewaves can be shifted over time to create directional and ambiguous motion. The top of the

figure shows a phase shift of +90° in space, every arbitrary time unit, leading to the generation of motion to the right in this instance. The bottom of the figure shows a phase shift of ±180° in space, every arbitrary time unit, leading to the generation of ambiguous motion that can be perceived as moving either to the left or to the right. Adapted from Kanai and Verstraten (2005).

Trials were created by presenting these stimuli sequentially in pairs (unambiguous followed by ambiguous) and asking participants whether they saw the pair moving in the same (i.e., stable) or different directions (i.e., reversal). These stimuli pairs were used across three experiments. In the first, the authors systematically varied the duration of pre-exposure to the directional, unambiguous motion, presenting this for between 0.08s and 0.64s. Additionally, the ISI between the directional and ambiguous stimuli was manipulated, with the shortest ISI being 0.04s and the longest being 2s. Results from this experiment demonstrated that shorter pre-exposure durations led to VMP (e.g., directional motion to the left and the ambiguous motion perceived as also moving to the left), especially when coupled with short ISIs. When the ISI increased beyond 0.4s, the priming effect dissipated. In contrast, longer pre-exposure to directional motion induced an adaptation effect (e.g., directional motion to the left and the ambiguous motion perceived as moving to the right). Interestingly, the results also showed that with the same ISI, both priming and adaptation effects can be evoked by pre-exposure durations that differ by just 0.24s. This very quick adaptation, the authors state, is a form of motion aftereffect which they call the rapid motion aftereffect (rMAE) to distinguish from the much slower motion aftereffect in which pre-exposure to directional motion causes a stationary stimulus to appear to be moving in the opposite direction (Anstis et al., 1998). To investigate whether this rMAE occurs at a low level related to the detection of motion energy or at a higher level related to perceptual decision making, the authors conducted a second experiment. Here, ISI was manipulated in the same way as with the first experiment, but all pre-exposure durations were 0.32s. The key here was that, on half of the trials, the authors replaced the pre-exposed directional motion stimulus with ambiguous motion The hypothesis here was that if the rMAE involved lower-level processing then it should not be induced by the pre-exposure of motion with no directional energy. In contrast, if rMAE involves higher level processes then it will still be induced by the endogenous perception of motion even in the absence of any directional energy. The results showed that pre-exposure to the ambiguous motion led to an effect similar to VMP, in that participants were more likely to report that the stimuli were moving in the same direction. However, unlike VMP, this effect was especially apparent at longer ISIs which suggests that it increases over time after motion is perceived. The authors suggest two things from this result: the first being that rMAE is mediated by low level processing, as it was not induced by motion lacking in directional energy; the second being that a second form of priming exists that is induced by the perception of motion without needing any energy-containing motion. The authors called this second priming effect “perceptual sensitisation” (PS). To investigate this further, Kanai and Verstraten (2005) conducted a final experiment. The authors hypothesised that if PS is induced by the perception of motion, then it should also be induced by directional, energy-based motion stimuli as well as ambiguous stimuli because these too, lead to the perception of motion. The authors suggest that the reason this effect had not been seen with directional motion in their previous experiment could have been because it was masked by the stronger, lower level rMAE. As PS appeared to increase in its effect over time, the authors replicated their second experiment but included much longer ISIs so that the rMAE could be allowed to dissipate so that PS could come through. Indeed, this is what they found. With ISIs of 2-5s, pre-exposure to 0.32s of directional, energy-based motion led to a priming rather than the adaptation response seen at lower ISIs. The results from this study suggest that there are several levels of processing that may be involved in priming and adaptation effects. It shows that the relationship between level of pre-exposure to directional motion and ISI is key to inducing the behavioural effects indicative of the processes occurring at each level.

A later study conducted by Takeuchi et al. (2011) used a very similar paradigm to Kanai and Verstraten (2005) to investigate the effects of retinal luminance and velocity of directional motion on VMP and rMAE and indeed found similar results. However, there were certain key differences in the way that this paradigm was employed. For instance, these authors suggested that because VMP effects seem to be stronger as ISI decreases, they removed the interval altogether. Instead, they asked participants to report the direction of the second ‘test’ stimulus which was ambiguous motion created in a similar way to Kanai and Verstraten (2005). Additionally, the authors here presented their stimuli to the lower temporal retina rather than to the fovea as Kanai and Verstraten (2005) did. The authors found that both low retinal luminance and high velocity of the pre-exposed directional motion eradicated VMP effects, whereas rMAE was apparent across all variations of retinal luminance and most velocities. The authors suggest that their results show that VMP is mediated by higher-level processes whilst rMAE involved directionally selective lower-level mechanism. This is broadly consistent with Kanai and Verstraten's (2005) conclusions around rMAE and their PS effect.

3.1.3 The present experiment

Experiment 3 (Chapter 4) will address the gap in the current knowledge identified above using multivariate pattern analysis (MVPA) of electroencephalographic (EEG) data. Briefly, Experiment 3 will involve training machine learning classifiers to decode perception of an ambiguous stimulus from EEG data recorded from either adapted or primed brain states, induced by exposure to biased interpretations of an ambiguous stimulus. These classifiers will then be used to attempt to predict spontaneous perceptual reversals of the ambiguous version of the figure. If these spontaneous reversals involve processes, for instance, of adaptation similar to those evoked in the training data, then the classifier should be able to accurately predict perceptual outcomes based on the EEG data recorded during them. Similarly, if priming processes are involved in spontaneous perceptual reversals, then the classifier trained on the primed brain state should be able to accurately predict perceptual states during spontaneous reversals (see Chapter 1 for a detailed overview).

One aim of the present experiment, therefore, was to determine the optimal parameters for inducing rapid adaptation and priming effects for use in an EEG experiment (Chapter 4), by partially replicating the study by Kanai and Verstraten (2005). In contrast to the previous authors however, this experiment aims to produce three clear trial types: one in which the adaptation effect is maximally induced, one in which the participant's perception is unbiased (i.e., 50% of each interpretation and equivalent or minimal adaptation and priming), and a third where the priming effect is maximally induced. The unbiased trial type will serve as a control condition with which to compare the performance of each of the adapted- and primed-trained classifiers to. It is important that each of these conditions are generated within the same ISI, as they were in the previous study by Kanai and Verstraten (2005), to exert maximum control over the information that is passed to the classifier. The increased sensitivity of MVPA over univariate analyses (e.g., List et al., 2017; Bae & Luck, 2018;2019) can lead to the classifier using information that is unrelated to the independent variable to classify trials (Grootswagers et al., 2016). Therefore, it is important that the only element of the trials that differs between conditions in Experiment 3 is the independent variable – i.e., the presentation duration of the unambiguous stimulus. The temporal parameters evoking these three trial types will then be used to inform the subsequent EEG experiment (Chapter 4) investigating adaptation.

A second aim of the current experiment is to determine whether the design used by Kanai and Verstraten (2005) can be replicated to reliably induce these priming and adaptation effects. Therefore the experiment will be replicated using all combinations of unambiguous stimulus presentation times and ISI that were used by Kanai and Verstraten (2005) so that results can be compared.

Finally, a key aim of Experiment 3 (EEG) was to determine whether adaptation plays a role in driving perceptual reversals within a sequence of fully ambiguous stimuli in which there was no stimulus-driven basis for adaptation. Thus, it was important to include such a sequence in the present Experiment to verify the presentation parameters for achieving maximum ambiguity (i.e., 50/50 left/right responses). Furthermore, the pattern of reversals in a fully ambiguous sequence can often be stochastic suggesting independence of outcomes between successive trials (e.g., Hesselman, 2008). This will be assessed within the block of ambiguous trials. In the present experiment, the ambiguous trials will comprise two successive ambiguous sine-wave gratings. Kanai and Verstraten (2005) found that priming occurred in this type of trial and they attributed this to perceptual sensitisation (PS). However, they presented these ambiguous-only trials randomly interspersed amongst trials which contained unambiguous motion. This unambiguous motion, when presented with ISIs of 2-5s has been shown by the authors to induce PS effects. So, it could be the case that the trials containing the unambiguous motion presented in close temporal proximity (i.e., around 2-5s) to the trials containing the ambiguous only motion, could have been inducing this PS effect, rather than this being induced only by the ambiguous motion. To avoid this from being the case here, the present experiment will contain these ambiguous-only trials within one block, with no unambiguous directional motion presented at all. This will determine whether this design can induce the spontaneous reversals in perception that are required for use in Experiment 3.

3.2 Experiment One

3.2.1 Methods

3.2.1.1 Participants

Eighteen undergraduate psychology students were recruited, in return for partial course credit. The a priori attention check criterion was that participants' data would be excluded from analysis if they did not respond to at least 70% of attention check trials correctly. Based on this criterion, one participant's data was excluded from the analysis. The remaining 17 participants' data (13 female, 4 male) were used in the experiment. This sample size was initially chosen based on an increase from the sample size used by Kanai and Verstraten (2005). These authors used 10 participants in their experiment, but did not report any effect sizes, so it was not possible to run an a priori power analysis. Participants had a mean age of 20.78 years (aged between 18 and 34). All participants had normal or corrected-to-normal vision (mean logMAR: 0.01; SD logMAR: 0.02; Precision Vision Logarithmic ETDRS 2000 chart). The study was approved by Keele University's Psychology Faculty Research Ethics Committee.

3.2.1.2 Stimuli and apparatus

3.2.1.2.1 Stimuli

Two types of stimuli were used in this experiment: (Figure 3.2), directional motion (Figure3.2, top) and ambiguous motion (Figure 3.2, bottom). Both types of stimuli were sine-wave gratings with a spatial frequency of 1 CPD. Gratings had a Michelson contrast of 0.5 (maximum luminance of white regions was 297.49cd/m^2, for black regions 100.28 cd/m^2) and were presented centrally on a mid-grey background (RGB: 128:128:128). All stimuli were circular, with a diameter of 10°. The circular shape was created using a 2-D Gaussian mask.

The motion of the stimuli was created by shifting the phase of the sine wave over time. Directional stimuli were created by a ±90° phase shift in the sine wave every 0.04s. Ambiguous stimuli were created by a 180° phase shift in the sine wave every 0.08s. This ensured that the speed of motion remained consistent in both stimulus types.

A white cross (0.8° x 0.8°) presented in the centre of the screen before and during stimulus presentation served as a fixation target throughout the experiment.

PsychoPy3 v3.2.3 (Peirce et al., 2019) was used to create all stimuli, and to design and control the experiment.

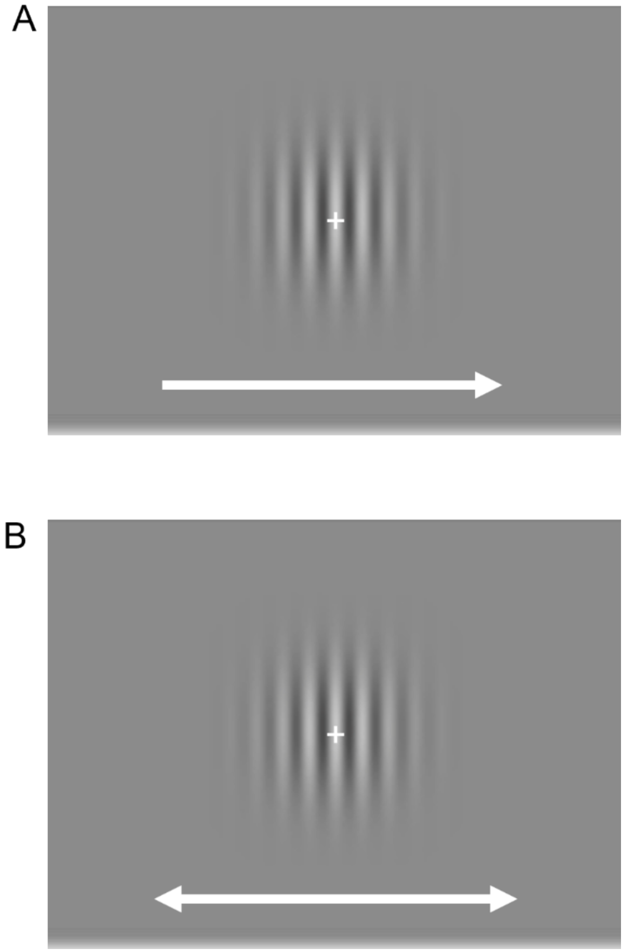

Figure 3.2: Example stimuli used in Experiment One. The stimuli were created from sine wave gratings which were phase-shifted to differing degrees. Panel A shows an example of directional motion. The phase was shifted by 90˚ to create the adaptation stimulus, which creates unambiguous motion towards (in this case) the right of the screen. Panel B shows ambiguous motion. Here, the phase was shifted by 180˚ to create the ambiguous motion stimulus, leading to two possible interpretations – one in which the gratings move to the left and one in which they move to the right. Note that the fixation cross shown here is not to scale, for accurate relative sizes in degrees, see Stimuli section above. Although

no actual difference can be seen in the stimuli themselves in this figure, the direction of the motion is indicated by the arrows.

3.2.1.2.2 Apparatus

The experiment was presented on a 24.5 inch BenQ Zowie XL2540 computer monitor with 120Hz refresh rate and 1920 x 1080 pixels resolution. Participants used a standard PC keyboard to respond to trials. Participants used a chin rest to maintain a constant viewing distance of 57cm, and to aid fixation by reducing head movements.

3.2.1.3 Design

The independent variables in this experiment were the presentation time of the unambiguous motion stimulus, and the ISI between that and the ambiguous motion stimulus. This design broadly replicated Kanai and Verstraten's (2005) first experiment, however an additional timing condition of 1.00s was included here to determine whether a stronger rMAE could be induced by longer pre-exposure to unambiguous motion. These IVs were manipulated in a 5 (presentation time: 0.08s, 0.16s, 0.32s. 0.64s and 1.00s) x 5 (ISI: 0.04s, 0.12s, 0.48, 1.00s and 2.00s) factorial design, shown in Table 3.1 below.

The experiment consisted of a random combination of 10 directional and 2 ambiguous blocks. Directional blocks contained 50 trials each (two trials for each of the 25 timing conditions set out in Table 3.1). In total, 20 trials for each of the 25 timing conditions were presented to each participant. Ambiguous blocks contained 50 ambiguous motion trials.

Each trial in the experiment consisted of the presentation of two stimuli. The type of stimuli presented depended on which block the trial was in. Trials in directional blocks consisted of one directional motion stimulus followed by a variable ISI, followed by an ambiguous motion stimulus (timings according to those shown in Table 3.1). Trials in ambiguous blocks consisted of one ambiguous motion stimulus presented for 0.32s, followed by a fixed ISI at 0.12s, followed by a second ambiguous motion stimulus presented for 0.32s. The ambiguous motion was presented for 0.32s firstly in order replicate the paradigm used by Kanai and Verstraten (2005) and secondly to ensure that this ambiguous motion presentation was equivalent to that in directional blocks. ISI did not vary in the ambiguous blocks as it has been shown to moderate reversal rate when viewing ambiguous stimuli (e.g., Kanai & Verstraten, 2005; Kornmeier & Bach, 2012). An ISI of 0.12s was used to replicate the ambiguous blocks in Kanai and Verstraten's (2005) original experiment.

The dependent variable was whether trials were reported as a reversal in perception or stable perception. Trials were coded as either "reversal" or "stable" according to whether participants responded by pressing two different buttons (for example 'up', 'down') or the same button twice (for example, 'up', 'up') respectively.

Interspersed randomly between experimental trials were attention check trials. Attention check trials were only presented during directional blocks, so that the unambiguous motion did not affect the perception of the ambiguous only motion in the ambiguous blocks. Directional blocks each contained 5 attention check trials. Therefore there were 50 attention check trials in total for each participant (5 trials x 10 directional blocks). These trials consisted of one directional motion stimulus presented once for 0.64s, followed by a second directional motion stimulus presented for 0.32s. There was one attention check trial for each of the 5 ISIs in each block (i.e., an ISI of 0.04s, 0.12s, 0.48s, 1.00s and 2.00s). The longest available presentation time was used for attention checks to give participants the best chance of getting these correct. Attention check trials consisted of “stable” and “reversal” pairs, which were directional motion in the same or the opposite direction respectively. These trials, therefore, carried a correct answer and were used to monitor participants’ attention to the task.

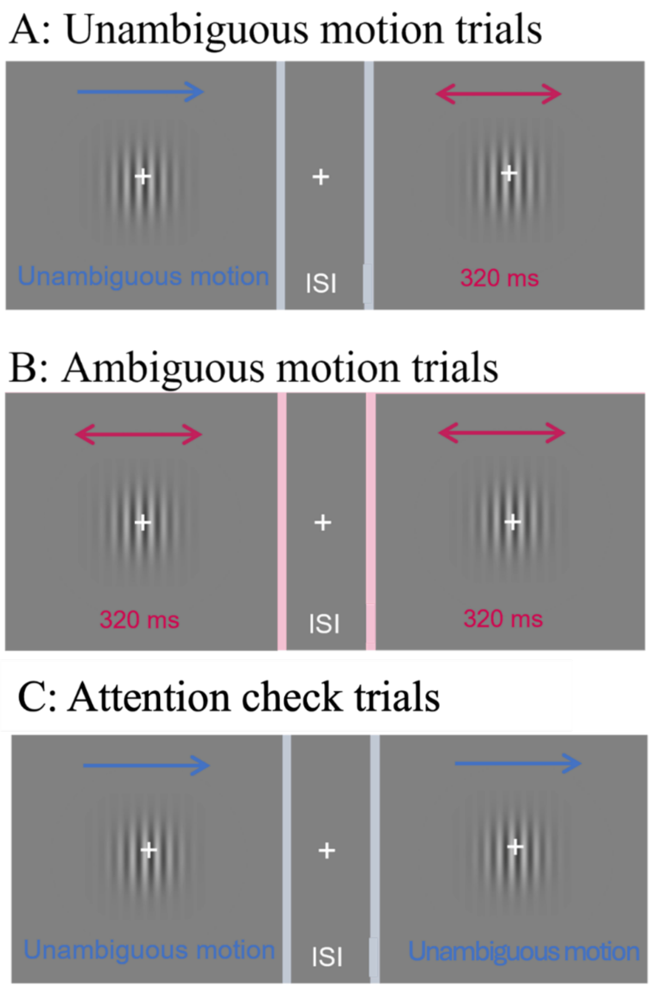

Figure 3.3: Example trial types used in Experiment One. Panel A shows an example of a trial in a directional block, where unambiguous motion in a particular direction (in this case to the right) was presented for one of the 5 presentation durations, followed by an ISI of one of the 5 ISI durations, before 0.32s of ambiguous motion. Panel B shows a trial in an ambiguous block, where two ambiguous motion stimuli were both presented for 0.32s, between them was an ISI of 0.12s. Note that the fixation cross shown here is not to scale, for accurate relative sizes in degrees, see Stimuli section above. Panel C shows an example of an attention check trial, whereby directional motion to the right in this case is followed by further directional motion, again to the right in this case.

Table 3.1

Timing parameters used in directional blocks to create the conditions for Experiment 1. These are a factorial combination of the ISI and directional stimulus durations.

		Presentation Time of Directional Motion Stimulus (s)				
		0.08	**0.16**	**0.32**	**0.64**	**1.00**
ISI (s)	**0.04**	0.08/0.04	0.16/0.04	0.32/0.04	0.64/0.04	1.00/0.04
	0.12	0.08/0.12	0.16/0.12	0.32/0.12	0.64/0.12	1.00/0.12
	0.48	0.08/0.48	0.16/0.48	0.32/0.48	0.64/0.48	1.00/0.48
	1.00	0.08/1.00	0.16/1.00	0.32/1.00	0.64/1.00	1.00/1.00
	2.00	0.08/2.00	0.16/2.00	0.32/2.00	0.64/2.00	1.00/2.00

3.2.1.4 Procedure

Participants were asked to respond to each trial using the ‘up’ and ‘down’ arrows on a standard PC keyboard to indicate motion towards the left or to the right. This was done so that counterbalancing of response key could be done without inducing strong stimulus-response compatibility effects for a reverse mapping (e.g., left button = right response). Response buttons were counterbalanced across blocks such that on 50% of trials ‘up’ indicated motion to the left and on 50% of trials ‘up’ indicated motion to the right. Response buttons remained consistent throughout each block. Participants were told which button meant what at the start of each block and reminded at the end of each trial. As each trial comprised two stimuli, participants were required to respond twice on each trial. For instance, if a participant saw leftward motion in the first stimulus, followed by rightward motion in the second, an example response could be ‘up’, ‘down’. Similarly, if a participant saw leftward motion in the first stimulus, followed by more leftward motion in the second stimulus, an example response could be ‘up’, ‘up’. Participants were given an inter-trial interval of 1.5s during which they could respond to each trial. Failure to respond within this time window generated a null response for that trial, and the next trial began.

All participants completed the informed consent process, which included asking them if there were any medical reasons that they could not complete the experiment (for example photosensitive epilepsy). After providing written consent, participants completed a test of their visual acuity, whilst wearing glasses or contact lenses to correct their vision if required.

Consenting participants with normal or corrected to normal vision then moved on to take part in the experiment. Instructions on how to complete the task were explained to the participant. During these, the experimenter ensured that the participant was able to perceive the two interpretations of the ambiguous motion stimulus (i.e., moving towards the left or towards the right). The experimenter also ensured that the participant knew how to respond to each trial, including whether they could respond within the 1.5s time window.

Participants then completed a practice block of 35 trials comprising 25 directional block type trials (one trial for each of the 25 timing conditions shown in Table 3.1 above) and 10 attention check trials (two for each of the 5 ISIs). The experimenter remained in the room during the practice block to ensure that participants were responding correctly to each trial. At the end of the practice block, participants were given feedback on how well they had performed on the attention check trials.

Participants then were left alone to complete the experiment, whilst the experimenter monitored from the next room. In total, there were 650 trials including the 50 attention checks. The whole experiment took participants around 25-30 minutes to complete the instructions, practice and real trials.

3.2.3 Results

3.2.3.1 Attention check trials

Participants' responses to attention check trials were analysed. A priori behavioural exclusion criteria were set such that participants scoring less than 70% correct on the 50 attention check trials in the main experiment, were excluded. Based on this criterion, one participant was excluded. The remaining 17 participants correctly responded to 82.59% (SD = 7.34) of attention check trials.

3.2.3.2 Response time

The overall average response time (RT; to first response button press) for all participants across both block type (ambiguous and directional) and trial type (reversal and stable) was 0.517s (SD: 0.102s) after the onset of the second stimulus in the trial.

A 2 (block type; ambiguous and directional) x 2 (reversal and stable) repeated measures ANOVA revealed that RTs were significantly faster in ambiguous blocks, M = 0.494s (SD = 0.114) than in directional blocks, M = 0.540s (SD = 0.083), $F(1,16) = 7.659$, $p = .014$, $\eta_p^2 = .354$. However, there was no significant effect of trial type on RTs, $p = .739$. A significant crossover interaction effect between trial type and block type was also revealed, $F(1,16) = 6.086$, $p = .027$, $\eta_p^2 = .303$ (Table 3.2).

To examine the interaction effect between trial type and block type, post hoc comparisons were conducted. In terms of the effect of block type for each trial type, for stable trials, RTs in ambiguous blocks, *M* = 0.481s (*SD* = 0.113) were significantly faster than those in directional blocks, *M* = 0.556s (*SD* = 0.081), *t*(16) = 3.679, *p* = 0.002. However, for reversal trials, there were no significant differences found between the RTs for ambiguous blocks, *M* = 0.506s (*SD* = 0.119) and directional blocks, *M* = 0.524s (*SD* = 0.084). In terms of the simple effect of trial type within each block type, for trials in the ambiguous block, RTs for reversal trials did not differ from those for stable trials (*M* and *SD*s as above). There were also no significant differences found between reversal and stable trials in the directional block (*M* and *SD*s as above). Table 3.2 below shows all RTs across block type and trial type.

Table 3.2

Mean Response Times across Block Type and Trial Type

		Mean (s)	SD
Block Type	Directional	0.540	0.083
	Ambiguous	0.494	0.114
Trial Type	Reversal	0.515	0.101
	Stable	0.519	0.104
Block Type x Trial Type	Directional, Reversal	0.524	0.084
	Directional, Stable	0.556	0.081
	Ambiguous, Reversal	0.506	0.119
	Ambiguous, Stable	0.481	0.113

3.2.5.3 Ambiguous block trials

Participants responded to indicate a reversal on ambiguous block trials on 50.57% (SD = 6.68) of trials and indicated stability on 49.43% (SD = 6.75) of trials. Participants responded to indicate that they had perceived leftward motion on 49.65% (SD = 8.38) of trials and rightward motion on 50.35% (SD = 7.94). Any trials where the participant failed to respond within the allocated 1.5s time window were excluded from the analysis. Neither reversal nor stable response percentages were found to be significantly different from chance, with both resulting in the same values for $t(16) = 0.126$, $p = .901$. The percentage of left and right responses were also not significantly different from chance, again both resulting in the same values for $t(16) = 0.179$, $p = .860$.

To examine the stochasticity of responses (i.e., independence between responses on subsequent trials), the distribution of response run lengths (i.e., number of same responses in a row for either left or right responses) was analysed. This was done to determine whether participants' responses in the ambiguous block were independent of one another. The ambiguous block was designed to evoke perceptual reversals that occur spontaneously, in a manner as close to that of continuous viewing as possible. Therefore, responses to the stimuli should be stochastic, as this would indicate that there was no bias toward one or the other direction of motion when the stimuli were perceived (Hesselmann, 2018). A repetition in response meant that a participant responded to indicate, for example, leftward motion on two or more successive trials. This 'run' of repetitions would be broken when the participant responded to indicate motion to the left then motion to the right, or vice versa. So, a response pattern of "left",

“left”, “left”, “right” would mean that there had been 2 repetitions before that run was broken. To assess this quantitatively, the distribution of response repetitions was fitted to a geometric distribution. A geometric distribution represents the frequency of different run lengths expected by chance in a sequence given only two random outcomes (i.e., left vs. right here) and a certain probability of each outcome (i.e., 50% chance here). For example, you would use this to estimate the number of times that you could expect to observe fair coin to land heads 4 times in a row in a sequence of 100 tosses. Therefore, if the response repetition data fits well to this distribution, this suggests that participants are responding in a random and unbiased manner in which responses to one stimulus are independent of those to preceding stimuli. The response repetitions from the ambiguous block were fitted to a geometric distribution and found to be well fitted (see Figure 3.4), $R^2(16) = .97$, $p <.001$.

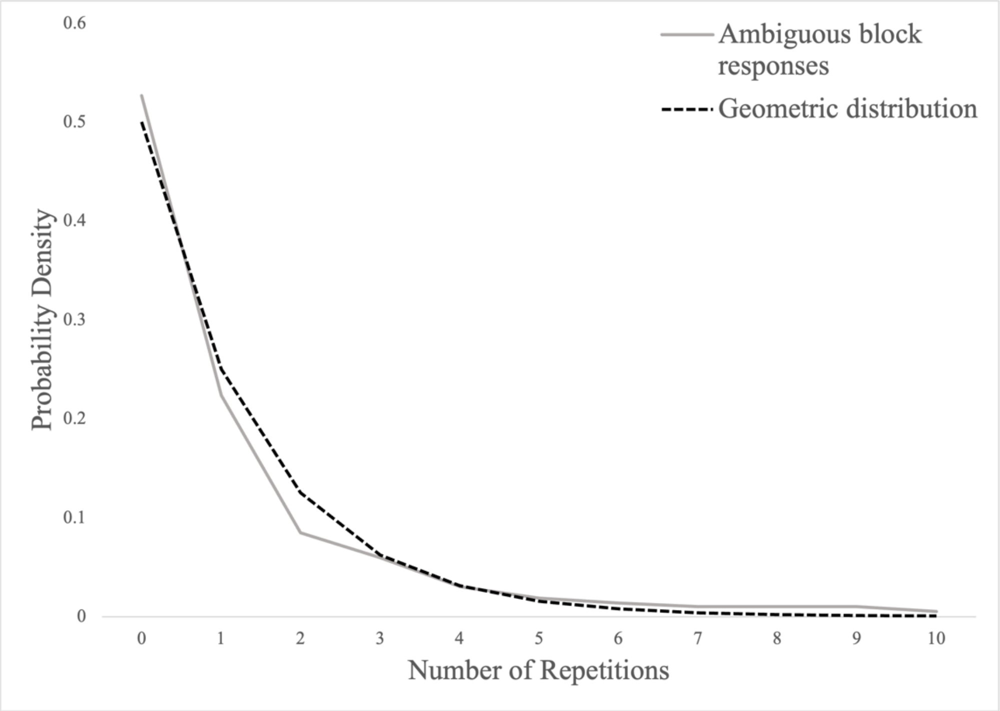

Figure 3.4 Solid grey line shows the probability density of the number of repeated responses across both trial types (i.e., left or right) in ambiguous blocks

(e.g., for 3 repetitions responses could be “left”, “left”, “left”). Dashed black line shows the probability density of the geometric distribution.

3.2.5.4 Directional block trials

Data from the directional blocks were separated into the 25 timing conditions (5 directional stimulus presentation times x 5 ISIs) and the average percentage of stable and reversal trials per participant was calculated for each condition. These results were submitted to a 5 x 5 repeated measures ANOVA with unambiguous stimulus presentation (s; 0.08, 0.16, 0.32, 0.64, 1.00) and ISI (s; 0.04, 0.12, 0.48, 1.00, 2.00) as factors. This revealed a significant main effect of directional stimulus presentation time, $F(2.26,64) = 7.72$, $p < .001$, $\eta_p^2 = .320$ and of ISI, $F(1.84,64) = 13.78$, $p < .001$, $\eta_p^2 = .463$. The Greenhouse-Geisser correction was applied in both cases as Mauchly’s test of sphericity was significant at the $p = .05$ level. However, no significant interaction between the two factors was found, $p = .077$.

To explore the significant main effects, post hoc comparisons were conducted. For the effect of ISI, these revealed that the percentage of reversals was significantly lower for the 0.04s ISI, $M = 36.47\%$ ($SD = 8.73$), when compared to the 0.12s ISI, $M = 56.42\%$ ($SD = 9.20$), $t(16) = 6.08$, $p < .001$, the 0.48s ISI, $M = 57.32\%$ ($SD = 4.23$), $t(16) = 6.35$, $p < .001$, the 1.00s ISI, $M = 54.91\%$ ($SD = 2.25$), $t(16) = 5.62$, $p < .001$, and the 2.00s ISI, $M = 52.71\%$ ($SD = 4.28$), $t(16) = 4.95$, $p < .001$. All other comparisons of ISI were not significant (all $p > .994$). For the effect of directional stimulus presentation time, the percentage of reversals in the 0.08s presentation time, $M = 43.90$ ($SD = 10.86$), was significantly lower than those for the 0.32s presentation time, $M = 54.22$ ($SD =$

6.84), *t(16)* = 3.95, *p* = .001, the 0.64s presentation time, *M* = 56.34% (*SD* = 8.51), *t(16)* = 4.76, *p* < .001, and the 1.00s presentation time, *M* = 54.35% (*SD* = 8.29), *t(16)* = 4.00, *p* = .001. All other comparisons of presentation time were not significant (all *p* > .348). The mean percentage of reversals for each timing condition can be found in Figure 3.5, and full ANOVA results in Table 3.3.

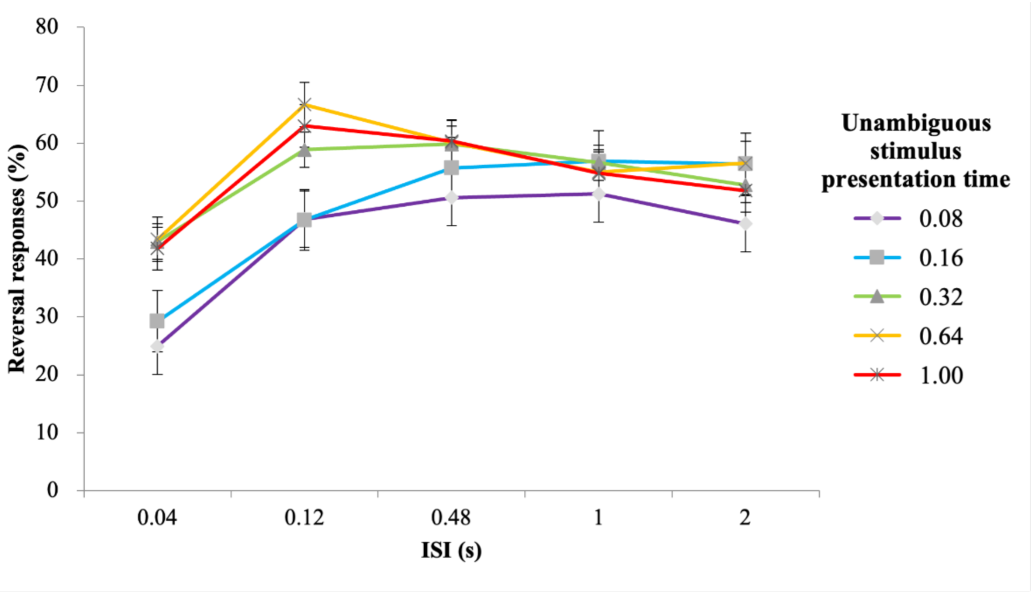

Figure 3.5 The percentage of reversal responses (i.e., trials where participants perceived the ambiguous stimulus moving in the opposite direction to the directional stimulus) are plotted as a function of ISI for each presentation time of the directional stimulus. Error bars show standard error (SE).

Table 3.3 *Full repeated measures ANOVA results for the percentage of reversals in the directional blocks (significant p-values are indicated in bold font)*

	df1	*df2*	*F*	*p*	η_p^2
Directional stimulus presentation time	2.262	64	7.541	**.001**	.320
ISI	1.835	64	13.776	**< .001**	.463
Directional stimulus presentation time x ISI	6.980	256	2.040	.056	.113

3.2.5.5 Optimum timing parameters

As the ISI is required to be consistent across trials for Experiment 3, the ISI that induced the strongest adaptation response (i.e., greatest percentage of reversals) was submitted to a one-way repeated measures ANOVA with directional stimulus presentation (s; 0.08, 0.16, 0.32, 0.64, 1.00) as the factor and ISI = 0.12, to determine the optimum temporal parameters for use in Experiment 3. This revealed a significant effect of directional stimulus presentation time, $F(2.26,64) = 6.23$, $p = .002$, $\eta_p^2 = .280$.

Post hoc comparisons showed that the percentage of reversals in the 0.08s directional stimulus presentation time $M = 46.86\%$ ($SD = 26.66$), was significantly lower than those in the 0.64s, $M = 66.65\%$ ($SD = 22.34$), $t(16) = 3.79$, $p = .002$, and 1.00s directional stimulus presentation times, *M* = 62.98% (*SD* = 21.09), *t(16)* = 3.09, *p* = .030. In the 0.16s directional stimulus presentation time *M* = 46.73% (*SD* = 21.46), the percentage of reversals was significantly lower than those in the 0.64s, *t(16)* = 3.82, *p* = .002, and 1.00s directional stimulus presentation times, *t(16)* = 3.11, *p* = .007 (both means and SDs above).

Overall, the greatest mean percentage of reversals was evoked by the 0.64/0.12s condition (66.65%). The lowest mean percentage of reversals which also involved the same 0.12s ISI was in the 0.16/0.12s condition (46.73%). The condition within the ISI = 0.12s conditions with a mean closest to chance was 0.08/0.12s (46.85%).

To determine whether the percentage of reversals in the 0.64s and 0.16s conditions were significantly higher and lower than chance (50%) and therefore indicated behavioural responses associated with the rMAE and VMP, respectively, two one sample t-tests were conducted. These revealed that the percentage of reversals in the 0.64/0.12s condition was significantly greater than would be expected by chance, *t(16)* = 3.07, *p* = .007, however those in the 0.16/0.12s condition were not, *t(16)* = 0.851, *p* = .407. The percentage of reversals in the 0.08/0.12s condition were also tested against chance and found not to significantly differ, *t(16)* = 0.928, *p* = .367.

3.2.6 Interim discussion

The primary aim of this experiment was to determine the optimum temporal parameters to evoke the visual motion priming (VMP), rapid motion aftereffect (rMAE/adaptation; Kanai & Verstraten, 2005) and unbiased responses, within the same ISI, for use in Experiment 3. Behavioural responses indicated that the rMAE was found in the 0.64s/0.12s condition as reversals occurred significantly more often than would be expected by chance. This means that participants were more likely to perceive the ambiguous motion in the second stimulus as moving in the opposite direction as the directional motion presented in the first stimulus. However, behavioural responses indicating VMP were not found within the 0.12s ISI as they were in the original study. A potential reason for this is discussed in more detail below. Responses indicating an unbiased perception following directional motion occurred in the 0.08/0.12s condition. This means that participants' responses did not differ significantly from those expected by chance, indicating that the directional motion did not bias perception of the ambiguous motion.

A second aim of this experiment was to determine whether the overall results produced by Kanai and Verstraten (2005) could be replicated. The results presented here did generally agree with those from Kanai and Verstraten (2005). VMP was elicited when directional motion was briefly presented and followed by a short ISI, meaning that participants were more likely to perceive the ambiguous motion in the second stimulus as moving in the same direction as the directional motion presented in the first stimulus. Additionally, as discussed above, the rMAE was induced when directional motion was presented for a longer time.

The final aim was to establish whether the perceptual sensitisation (PS) effect found by Kanai and Verstraten (2005) would influence participants' perception of ambiguous motion stimuli pairs when these trials were presented exclusively in one block type with short ISIs (i.e. without interspersion of any directional motion). Results showed that PS did not influence perception of ambiguous motion stimuli when they were presented in this way.

Although the results presented here do broadly agree with Kanai and Verstraten's (2005), significant VMP was not found in the same ISI as maximal rMAE in the present experiment. Furthermore, there was overall weaker VMP in the present experiment compared to Kanai & Verstraten (2005). An important difference in task instructions could help to explain this together with an observation about the displays in the short ISI conditions.

Upon visual examination of the displays in the shortest ISI conditions (0.08s and 0.12s) where VMP was expected to be strongest, the experimenter observed that it was very difficult to detect two distinct stimuli in these displays. Informal observations from several other observers corroborated this assessment. Most of them clearly perceived these displays as a single continuous stimulus rather than as two separate stimuli. This effect was reported as being strongest when the directional stimulus duration was also short. In fact, in those cases, it was often reported as not possible to even detect either the directional stimulus or the ISI separately from the ambiguous stimulus.

Kanai and Verstraten (2005) asked participants to respond to trials according to whether they saw the two stimuli within each trial moving in the same or opposite directions. If participants perceived just one stimulus in the short ISI conditions as is suggested by our informal observations outlined above, we would expect that participants would be biased to respond "same" more often in these conditions simply because they only perceived one single stimulus and because they may not have even been able to detect the motion of the motion of the direction stimulus in such short displays. For example, under these very brief timing conditions, a participant might only notice the 0.32s of ambiguous motion and not perceive the 0.04s ISI between that and the preceding 0.08s of unambiguous motion. Given the forced choice between “same” and “different”, they might respond to indicate that they saw stimuli moving in the same direction, simply because they did not notice that there were two stimuli to compare. On average, one would have expected this to artificially increase the number of “same” responses in conditions where this happened.

In the present experiment, participants were forced to consider and report the motion of the directional stimulus. This requires a more specific answer than the same/opposite task in Kanai & Verstraten. This could lead to several different strategies. One option is for participants to simply report the same direction that they perceived in the second stimulus as outlined above. This type of response would increase the number of “stable”/”same” responses. In contrast, participants may be inclined to guess about the direction of the first stimulus. In this case, participants might assume that they missed the first stimulus and would be forced to randomly guess its direction because two directions are required in the response. This strategy would, on average, artificially reduce the number of “stable”/”same” responses because guessing, if random, is equally likely to be the same or different from the ambiguous stimulus perception. This would then push the percentage of same responses towards 50%.

If participants are indeed employing the second strategy described here, it could explain why at these very brief timing conditions, the results from Kanai and Verstraten’s (2005) study were not replicated as strongly here. Whether or not participants were guessing about the direction of the first unambiguous motion stimulus can at least partly be assessed by the paradigm used in Experiment 1, by examining the percentage of trials in which the unambiguous motion was reported correctly (i.e., leftward motion was presented, and the participant indicated that they saw the first stimulus in the pair moving to the left) for each presentation time. Therefore, we examined whether the presentation duration of the directional stimulus (i.e., first stimulus in directional-ambiguous pairs) influenced participants' ability to identify the correct direction of motion. To do this, the average number of correct responses to the directional stimuli (i.e., first response of two in each trial) in each timing condition was submitted to a 5 x 5 repeated measures ANOVA with directional stimulus presentation time (s; 0.08, 0.16, 0.32, 0.64, 1.00) and ISI (s; 0.04, 0.12, 0.48, 1.00, 2.00) as factors. This revealed a significant effect of directional stimulus presentation time on the percentage of trials that participants correctly identified the direction of motion, $F(1.48,64) = 20.392$, $p < .001$, $\eta_p^2 = .560$. A significant main effect of ISI was also revealed, $F(4,64) = 3.065$, $p = .023$, $\eta_p^2 = .161$. There was no interaction effect, $p = .613$.

Post hoc comparisons of the effect of presentation time collapsed across ISI revealed that the 0.08s presentation time $M = 46.82\%$ ($SD = 5.59$), led to significantly fewer correct responses to directional motion stimuli than the 0.16s presentation time, $M = 60.19\%$ ($SD = 17.08$), $t(16) = 4.73$, $p < .001$, the 0.32s presentation time, $M = 68.24\%$ ($SD = 22.25$), $t(16) = 6.98$, $p < .001$, the 0.64s presentation time, $M = 68.65\%$ ($SD = 22.89$), $t(16) = 7.09$, $p < .001$, and the 1.00s presentation time, $M = 71.65\%$ ($SD = 24.53$), $t(16) = 7.92$, $p < .001$. Additionally, the 0.16 presentation time led to significantly fewer correct responses than the 0.64s condition, $t(16) = 3.13$, $p = .006$ and the 1.00s condition, $t(16) = 3.55$, $p = .004$ (means and SDs as above). All other combinations of presentation times did not differ significantly (all $p > .141$)

Post hoc comparisons of the effect of ISI collapsed across unambiguous stimulus presentation time revealed that significantly fewer correct responses to directional motion occurred in the 0.04s ISI, *M* = 60.55% (*SD* = 21.33), than the 0.48s ISI, *M* = 64.88% (*SD* = 22.53), *t(16)* = 3.08, *p* = .007. All other comparisons were not significant (all *p* > .069). See also Figure 3.6 below.

To directly test the assumption that at the shortest presentation times participants were simply guessing about the direction of motion of the first stimulus, the percentage of correct responses at each presentation time were compared against chance (i.e., 50%; collapsed across all ISIs, to avoid an unnecessary number of comparisons, and given that we are assuming that the presentation time is the driving factor) in a series of one-sample t-tests. These revealed that participants' responses to the directional motion were significantly greater than would be expected by chance at presentation times of 0.16s, *t(16)* = 2.356, *p* = .032; 0.32s, *t(16)* = 3.379, *p* = .004; 0.64s, *t(16)* = 3.976, *p* = .001; and 1.00s, *t(16)* = 3.994, *p* = .001. However, these correct responses were not significantly different from chance levels at the shortest presentation time (0.08s), *t(16)* = 1.259, *p* = .226 (all means and SDs reported above).

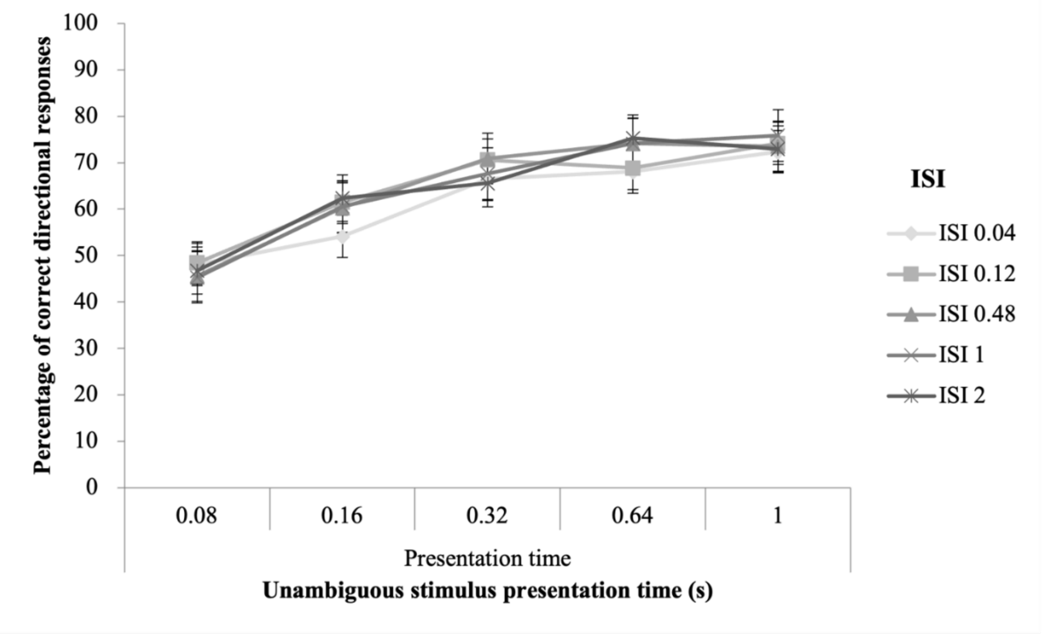

Figure 3.6 The percentage of responses made where the response to the unambiguous motion stimulus was correct. For example, if the unambiguous motion was moving to the left, participants reported this correctly, regardless of the direction that they perceived the ambiguous motion moving in. Error bars show standard error (SE).

This analysis of directional-only responses clearly indicates that participants were significantly less accurate when the directional stimulus was presented for shorter durations. This suggests that participants could have been less likely to notice two stimuli in these conditions, and that they were making a guess about the motion of the first stimulus. Furthermore, these results raise uncertainty about whether the results of Kanai and Verstraten's (2005) experiment actually reflect VMP. If the participants were unable to discriminate the direction of motion of the first stimulus, as is indicated by our results, then they may have been only making their reports based on the second stimulus. In a same/different task, participants would then have been more likely to report "same" because they only saw one direction of motion. This could mean that at least some of the VMP effect at the very shortest ISI and directional stimulus durations may have been artefactual. This will be explored further in Experiment 2 below.

The issue raised above for VMP, does not apply to rMAE. Participants were above chance at discriminating direction of the directional stimuli for conditions that evoked rMAE. Thus, the temporal parameters for rMAE from this experiment will be taken forward to Experiment 3, along with the parameters required to induce an unbiased perception of an ambiguous stimulus following directional motion.

3.3 Experiment Two

3.3.1 Introduction

Experiment 1 raised questions about whether the temporal parameters reported by Kanai and Verstraten (2005) can evoke visual motion priming (VMP) in the way that is described by the authors. Furthermore, although Experiment 1 did broadly replicate the pattern of results shown by Kanai and Verstraten (2005), the strong VMP effect presented by the authors, particularly when an inter-stimulus interval (ISI) of 0.12s was used, was not present at all in Experiment 1. A possible reason for the inconsistency between these results is that the task that the participants were asked to complete in Experiment 1 differed from the one used by Kanai and Verstraten (2005). Experiment 1 asked participants to report the direction of motion in each of the stimulus pairs within a trial and were therefore encouraged to consider the direction of motion of each of the two stimuli in the trial. In contrast, Kanai and Verstraten (2005) asked participants to report only whether the two stimuli were moving in the same or different directions. This meant that it would have been possible for participants in Kanai and Verstraten's (2005) study to report that they saw motion in the same direction (i.e., a VMP response) simply because they had only perceived one stimulus. Intuitively, this seems like it would be particularly true when very brief stimuli are presented or a very short ISI is used, because under these conditions the end of one stimulus and the beginning of another might be difficult to distinguish. Therefore, asking participants to compare two stimuli when they notice only one, could lead to invalid responses.

Experiment 2 aimed to address this issue through more detailed investigation of how participants perceive the stimulus displays. This was done through varying the task instructions. First, participants were directly asked about the number of stimuli that they perceived under the temporal parameters that have been shown to elicit VMP. If participants are likely to see one stimulus rather than two, then this may lead to artefactual "same" responses which have previously been interpreted as VMP. This account can be directly tested here by directly measuring the number of stimuli perceived.

Second, whether participants perceive one or two stimuli under these conditions was also addressed by including a block of trials which gave them more options with which to respond. Rather than restricting participants to answer a question such as "which direction were the two stimuli moving in?" or "were the stimuli moving in the same or different directions?", when they may feel that they only saw one stimulus and therefore can't honestly answer, this experiment included a block which allowed participants to respond to indicate the motion of any stimuli that they saw, without restriction.

Lastly, a detailed consideration of the task in Experiment 1 also raises the issue of where in time do participants draw the line between stimulus 1 (directional) and stimulus 2 (ambiguous) especially when the ISI is short and potentially not perceived. This is important because if stimulus 1 responses are, unintentionally, based on perceptual experiences of stimulus 2 then this presumably will bias responses towards “same” because sampling perceptual experience during the same stimulus will be more likely to result in the same outcome. This would artificially inflate estimate of same responses and thus VMP. To estimate same/different responses whilst reducing the likelihood of this artefact, Experiment 2 also included a partial replication of the approach taken by Takeuchi et al. (2011). Takeuchi et al. (2011) asked their participants only about their perception of the second, ambiguous stimulus in each trial. This was then compared against the known direction of the first, unambiguous stimulus. This approach reduces variability or artefacts caused by problems with the perception of the first stimulus. Thus, this should provide a purer measure of any VMP that is present. It is difficult to directly compare the results from the Takeuchi et al. (2011) study to either Kanai and Verstraten’s (2005) or Experiment 1 because they also varied other parts of the paradigm such as having no ISI at all, not presenting the stimulus to the fovea, and modulating luminance. Thus, Experiment 2 will adopt their task but with the same stimulus design as Experiment 1 in order to address the issues here. Asking participants about only their final perception of the ambiguous stimulus means that they are not asked to make a judgement about whether they saw one or two stimuli, or to compare them in any way (for example being asked “were the two stimuli moving in the same direction?” as they were in Kanai and Verstraten’s study; 2005).

3.3.2 Methods

3.3.2.1 Participants

Twenty-three participants took part in Experiment 2. The participants that took part in this experiment were not the same as those who took part in Experiment 1. The same exclusion criteria applied in Experiment 2 as in Experiment 1. The a priori attention check criterion was that participants' data would be excluded from analysis if they did not respond to at least 70% of attention check trials correctly. Based on this criterion, one participant's data was excluded from the analysis. The remaining 22 participants' data was included (19 females, 3 males; average age 29.76; age range between 21 and 49). All participants had normal or corrected to normal vision (mean logMAR: 0.01; SD logMAR: 0.03; Precision Vision Logarithmic ETDRS 2000 chart). Ethical considerations were identical to those in Experiment 1.

3.3.2.2 Stimuli & Apparatus

The stimuli used in this experiment were identical to those used in Experiment 1 with the exception that the number of ISIs was reduced only to the short ISIs where priming was expected and where perception of the ISI was found to be difficult in anecdotal reports. Furthermore, the number of directional stimulus durations was reduced by one (details below). Figure 3.7 below serves as a brief reminder of these. As with Experiment 1, one trial consisted of two stimuli: an unambiguous, directional motion stimulus that varied in its presentation time, followed by an inter-stimulus interval that also varied, followed by 0.32s of ambiguous motion. The apparatus was also identical to that used in Experiment 1.

An online example of the two of the trial types used in this experiment can be seen at: https://run.pavlovia.org/KimberleyDundas/stimulus_example/; however these are for illustrative purposes only and the timing within a browser has not been tested.

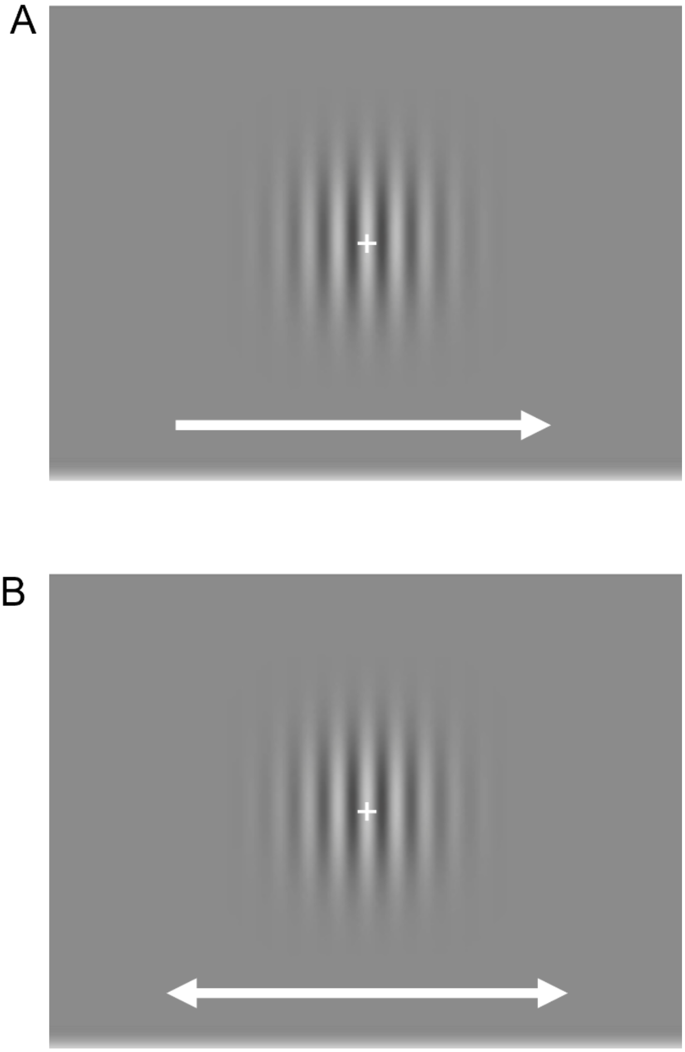

Figure 3.7 Examples of the directional motion (panel A, where the sine wave gratings move towards the right) and the ambiguous motion stimuli (panel B, where the sine wave grating can be viewed as moving in either direction).

3.3.2.3 Design

The independent variables in this experiment were the inter-stimulus interval (ISI) and the time that the unambiguous, directional motion stimulus was presented before the ambiguous motion stimulus. This was very similar to Experiment 1, although the ISIs used in Experiment 2 were reduced to be only the ones where the visual motion priming (VMP) effect of interest was strongest (0.04s and 0.12s). In terms of the unambiguous stimulus presentation time, a subset of presentation times was used to allow for more trials to be included within each condition whilst also allowing for the VMP and rapid motion aftereffect (rMAE) effects from Experiment 1 to be replicated with the ISIs used here. Therefore, the unambiguous stimulus presentation times used in Experiment 2 were 0.08s, 0.16s, 0.32s and 0.64s. The longest duration was not used as it was not part of the original Kanai & Verstraten study and was not expected to show VMP.

To investigate the effects of the instructions and task given to participants on the responses they made to these stimulus pairs, there were four different blocks used in Experiment 2. Other than the instructions given to participants, and the way they were asked to respond, all stimuli in every block were identical and consisted of 24 trials in each of the 8 timing conditions (2 ISIs x 4 unambiguous stimulus presentation times). An additional 10 attention check trials were included in every block type (these are described within the descriptions of each block, below), which meant that in total there were 202 trials in each block, and 808 trials in total for the experiment. There was a total of 40 attention check trials in the experiment. Any participant scoring less than 70% on attention check trials was excluded from the analysis. The details of each of the four blocks are described below.

3.3.2.3.1 Replication block

The first block type was the replication block, which carried the same instructions as Experiment 1. This block was referred to as the “Two Directions Task” to participants. In this block, participants were asked to report the direction of motion of each of the two stimuli in the pair using the up and down arrow keys in the same way as Experiment 1. The up and down arrow keys represented the directions left and right; these were counterbalanced such that on half of the trials up represented left and down represented right, and vice versa. This again is identical to Experiment 1. The replication block was included here as this experiment was conducted on an independent set of participants to Experiment 1, and under slightly different circumstances as this experiment had a lot more tasks to complete and a lot more instructions. Therefore, it was important to make sure that the results of Experiment 1 were replicable. In this block, attention check trials consisted of trials in which 0.64s of unambiguous motion was followed by 0.32 of more unambiguous motion rather than 0.32s of ambiguous motion. For example, 0.64s of motion to the left, followed by 0.32s of motion to the right. These trials therefore carried a correct answer and were scored for each participant. Using 0.64s for the presentation time of the first unambiguous stimulus was chosen as this gave participants the maximum amount of time to detect the correct direction of motion.

3.3.2.3.2 Counting block

The second block type was the “Counting Task” block. In this block, participants were instructed to use the up and down arrow keys to report whether they saw one stimulus or two stimuli. These responses were again counterbalanced such that on half of the trials up represented one stimulus and down represented two stimuli, and vice versa. Participants were reminded that one stimulus, with no perceived gap, could potentially reverse in perceived direction but that would still be counted as one stimulus. Similarly, a display perceived as having two stimuli separated by a gap could move in the same direction but that would still be counted as two stimuli. This block contained attention check trials whereby there was actually no ISI, resulting in one continuous motion stimulus containing directional and ambiguous motion. Otherwise on actual trials, there were always two stimuli, separated by an ISI of either 0.04s or 0.12s, presented under the same conditions as every other block. The purpose of this block was to directly assess whether participants were more likely to perceive the stimulus pair at the shortest presentation time and ISI as one stimulus rather than two.

3.3.2.3.3 Dual task block

The third block type was the dual task block. This was referred to as the “One or Two Directions” block to participants. In this block, participants were instructed to respond using the up and down arrow keys to indicate how many stimuli they saw, and in which direction they saw them moving. In this block, the arrow keys again represented the directions left and right, and again were counterbalanced so that on half of the trials up represented right and down represented left and vice versa. Participants were asked to respond only to the stimuli that they saw. So, for example, if they saw one stimulus moving to the right they would respond to indicate right only once. But if they saw two stimuli both moving to the right they would respond to indicate right twice. Using responses in this way meant that although the participants were only making one set of responses, information about both the number of stimuli and the direction in which they were perceived as moving in could be gathered. If a participant responded to indicate that they saw only one stimulus, this was coded as “one stimulus”. If participants reported seeing two stimuli, these responses were coded in the same way as the replication block and as Experiment 1. Responses from this block allowed more detailed information on the subjective responses of participants without being told to focus on one element of the task.

3.3.2.3.4 Final perception block

The final block type was a partial replication of the task used by Takeuchi et al (2011). This was titled the “Final Direction” block for participants. In this block, participants were instructed to report only the direction of their perception of motion that they were experiencing right at the end of the trial. Participants were asked to do this again by using the up and down arrow keys to represent the directions left and right. Again, these keys were counterbalanced in the same way as the other blocks. The task in this block reduced the potential artefact of participants only reporting a VMP response because they didn’t perceive two stimuli at all. Instead, by reporting only their final perception of motion, participants’ attention was focused on how they perceived the second, ambiguous stimulus. This meant that trials were coded as a reversal if the ambiguous motion was reported as moving in the opposite direction as the known preceding directional motion. For example if the directional motion was moving to the right and the participant responded to indicate that they saw the ambiguous stimulus moving to the left then this would be coded as a reversal. By contrast, trials were coded as stable (i.e., consistent with VMP) if the ambiguous motion was reported as moving in the same direction as the directional motion (i.e., the directional motion was moving to the left and then the participant reported that they also saw the ambiguous motion moving to the left). By removing the need to report the direction of the first stimulus at all, this block was thought to be a more pure measure of whether participants were experiencing a priming effect or not. Attention checks in this block were the same as the ones used in the replication block, in that they consisted of 0.64s of unambiguous motion followed by 0.32s of more unambiguous motion. Therefore, participants' responses to the direction of motion of the second stimulus could be scored as being correct or incorrect.

Table 3.4 *Instructions provided to participants in each block*

Task type	Instructions	Example response
Two directions	This is the TWO DIRECTIONS TASK. Your task is to identify which direction EACH of the TWO stimuli in the pair were moving: to the left or to the right. This means that at the end of each trial, you are required to give TWO responses: one to indicate the motion of the first stimulus, and one to indicate the motion of the second stimulus. So, if you saw motion to the left and THEN motion to the right, you would respond to indicate "left, right”.	“up, down” to indicate motion to the left then the right. “up, up” to indicate motion to the left and then the left again.
One or two directions	This is the ONE OR TWO DIRECTIONS TASK. Your task is to report the direction of motion of all the stimuli that you see. So, if you saw just one stimulus moving to the right, you would respond ONCE indicate simply "right". If you saw TWO stimuli moving to the right, you would respond TWICE to indicate "right, right". This means that, from your responses, the experimenter can tell how many stimuli you saw and in which direction they were moving.	“up” to indicate one stimulus moving to the left. “up, up” to indicate two stimuli both moving to the left. “up, down” to indicate two stimuli, the first moving to the left and the second moving to the right.
Counting	This is the COUNTING TASK. Your task is only to identify whether you saw one or two stimuli in each trial. At the end of each trial, you will make ONE response to indicate either that you saw one stimulus or two stimuli. You might see two stimuli moving in the same direction, or in different directions, but in both cases, you should report that you saw TWO stimuli. The important thing to remember is to ask yourself if you noticed two 'things' or one. So, if you saw one stimulus you would respond to indicate "one" and if you saw two stimuli you would respond to indicate "two". You do not need to report the direction of the stimuli in this task.	“up” to indicate one stimulus was perceived. “down” to indicate two stimuli were perceived.
Final direction	This is the FINAL DIRECTION TASK. Your task is to identify in which direction your FINAL perception of the stimulus was moving. So, ask yourself, just before you respond - which direction of motion did you just see? On these trials you are only asked for ONE response. So, if you saw motion to the left and THEN motion to the right, you would only respond to indicate "right".	“up” to indicate the second stimulus was moving to the left. “down” to indicate the second stimulus was moving to the right.

3.3.2.4 Procedure

The procedure for this experiment was very similar to that of Experiment 1; the consent procedure and medical background check (including a visual acuity test), were identical to Experiment 1. Consenting participants with normal or corrected to normal vision then moved on to take part in the experiment.

Before each block of trials, the experimenter entered the testing room and fully explained the instructions for that block to the participant. The experimenter remained in the testing room throughout the practice trials to check that the participant fully understood both the task and how to respond in each block. The order of the blocks was randomised across participants.

3.3.3 Results

3.3.3.1 Attention check trials

The a priori attention check criterion was that participants' data would be excluded from analysis if they did not respond to at least 70% of attention check trials correctly. Based on this criterion, one participant's data was excluded from the analysis. The remaining 22 participants' data was included. Overall, participants correctly responded to 83.52% (SD = 6.26) of attention check trials.

Although data from any participants who did not achieve at least 70% accuracy overall on attention check trials were excluded from analyses, it was important to establish whether responses to attention check trials differed across block type as a deviation from this assumption could indicate that there was an unwanted effect of attentional differences as well as the experimental manipulation of task. Therefore, the percentage correct attention check trials in each of the four blocks were submitted to a repeated measures ANOVA. This revealed that there was no effect of block type on the percentage of attention check trials answered correctly ($p > 0.358$). Means and standard deviations are presented in Table 3.5 below.

Table 3.5 *Percentage of attention check trials answered correctly across each block type, standard deviations are shown in parentheses.*

	Block type			
	Replication block	Dual task block	Final direction block	Counting block
Percentage of correct attention check trials	82.273 (7.516)	81.818 (7.327)	85.000 (8.591)	81.818 (6.645)

3.3.3.2 Replication block

In the replication block, trials were coded as reversal or stable in the same way as in Experiment One (responses where the same button was pressed twice were coded as stable trials and those where two different buttons were pressed were coded as reversals). The percentage of stable trials for each combination of ISI (0.04 and 0.12s) and unambiguous stimulus presentation time (0.08, 0.16, 0.32 and 0.64s) were submitted to a 2x4 repeated measures ANOVA. This revealed significant main effects of both ISI, $F(1,21) = 116.886, p < .001, \eta_p^2 = .848$ and presentation time, $F(2.03,63) = 50.337, p < .001, \eta_p^2 = .706$, but there was no interaction found between the two, $p = .314$. See Figure 3.8 below. The Greenhouse-Geisser correction was applied as Mauchly's test of sphericity was significant at the $p = .05$ level.

To explore the main effects, post hoc comparisons were conducted. These revealed that there were significantly fewer stable responses with an ISI of 0.12s, *M* = 43.85% (SD = 19.76), compared to 0.04s, *M* = 63.26% (SD = 18.67), *t(21)* = 10.811, *p* < .001, *d* = 1.228. For the effect of unambiguous stimulus presentation time, the percentage of stable responses in the 0.08s condition, *M* = 67.07% (SD = 17.53), was significantly higher than in the 0.32s, *M* = 45.47% (SD = 19.53), *t(21)* = 8.408, *p* < .001, *d* = 1.366, and the 0.64s, *M* = 39.99% (SD = 17.81), *t(21)* = 10.544, *p* < .001, *d* = 1.713, conditions. The percentage of stable responses in the 0.16s condition, *M* = 61.69% (SD = 19.12), was also significantly higher than in the 0.32s, *t(21)* = 6.311, *p* < .001, *d* = 1.025, and the 0.64s, *t(21)* = 8.447, *p* < .001, *d* = 1.372, conditions. All other comparisons were not significant (all *p*s >.219). See Figures 3.8 and 3.9 below. For the full RM ANOVA results see Table 3.6.

To determine whether responses indicated VMP or rMAE, the percentage of responses that indicated a stable perception were compared against chance (50%) using one-sample t-tests. If the percentage of stable responses were greater than chance, this is indicative of VMP, whereas a percentage significantly less than chance indicates rMAE. These t-tests revealed that, within the 0.04s ISI, the percentage of VMP responses greater than chance occurred within presentation times of 0.08s, *M* = 75.06% (*SD* = 12.74), *t(21)* = 9.226, *p* < .001, *d* = 1.967 and 0.16s, *M* = 72.59% (*SD* = 14.27), *t(21)* = 7.427, *p* < .001, *d* = 1.583. Within the 0.12s ISI, responses greater than chance occurred only with a presentation time of 0.08s, *M* = 59.09% (*SD* = 18.25), *t(21)* = 2.336, *p* = .029, *d* = 0.498. In the 0.12s ISI, responses significantly lower than chance occurred with presentation times of 0.32s, *M* = 36.87% (*SD* = 16.78) , *t(21)* = 3.672, *p* = .001, *d* = 0.783, and 0.64s, *M* = 28.65% (*SD* = 10.96) , *t(21)* = 9.138, *p* < .001, *d* = 1.948. All other combinations of ISI and presentation time were not significantly different from chance level (all *p*s > .314). These are marked on Figure 3.8, below.

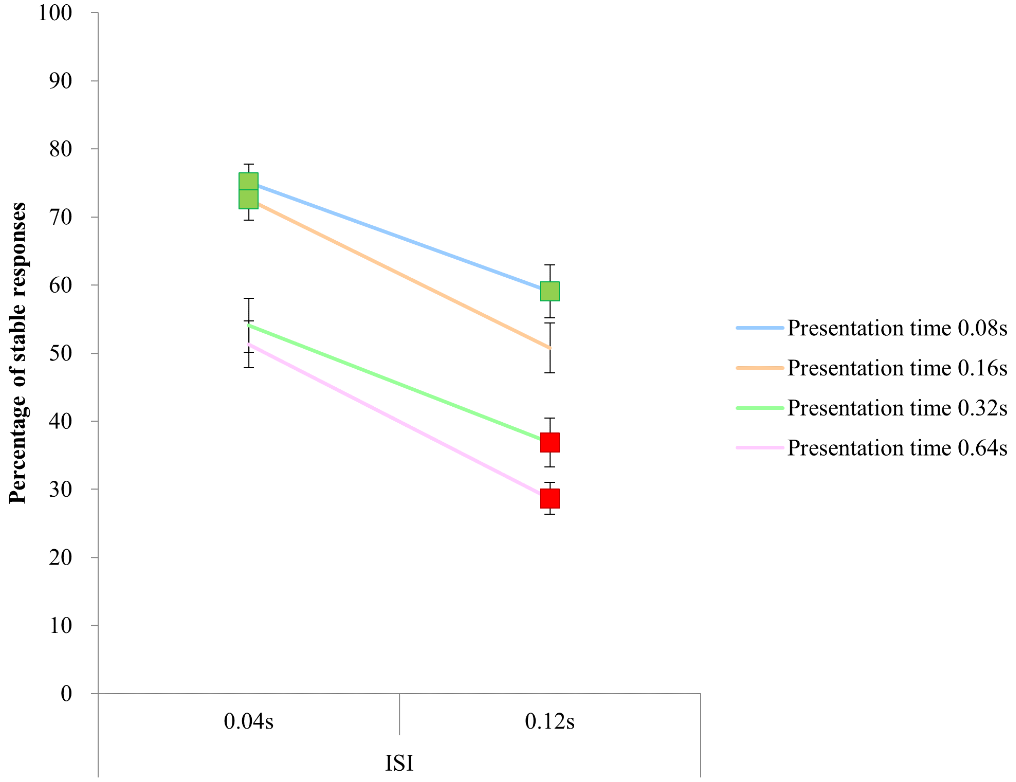

Figure 3.8 The percentage of stable responses plotted as a function of ISI for each presentation time of the unambiguous stimulus, for the replication block. Green markers highlight conditions leading to above-chance stable responses (therefore indicating VMP) and red markers highlight conditions leading to below-chance stable responses (therefore indicating rMAE). Error bars show standard error (SE).

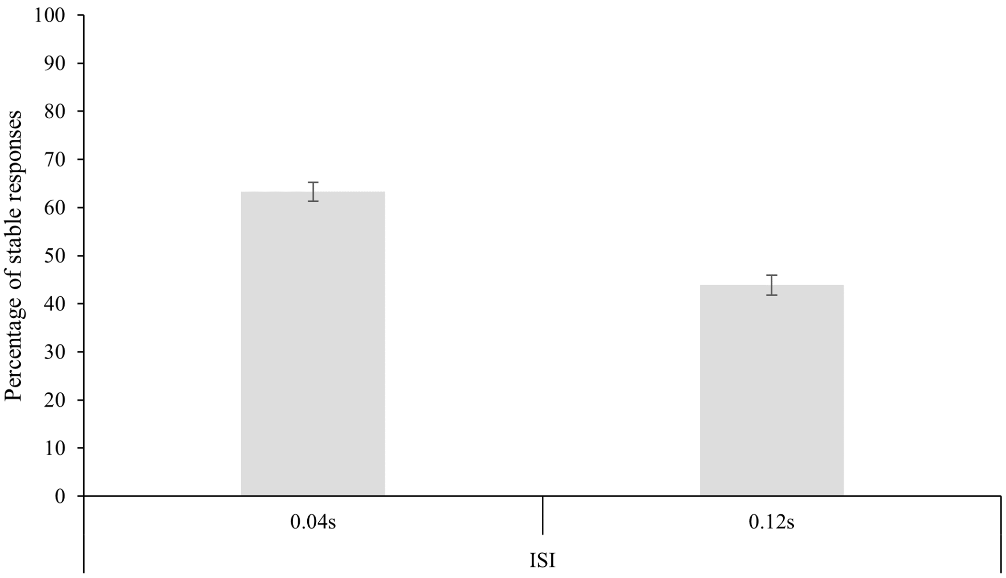

Figure 3.9 The percentage of stable responses for each ISI of the unambiguous stimulus collapsed across all presentation times of the unambiguous stimulus, for the replication block. Error bars represent SE.

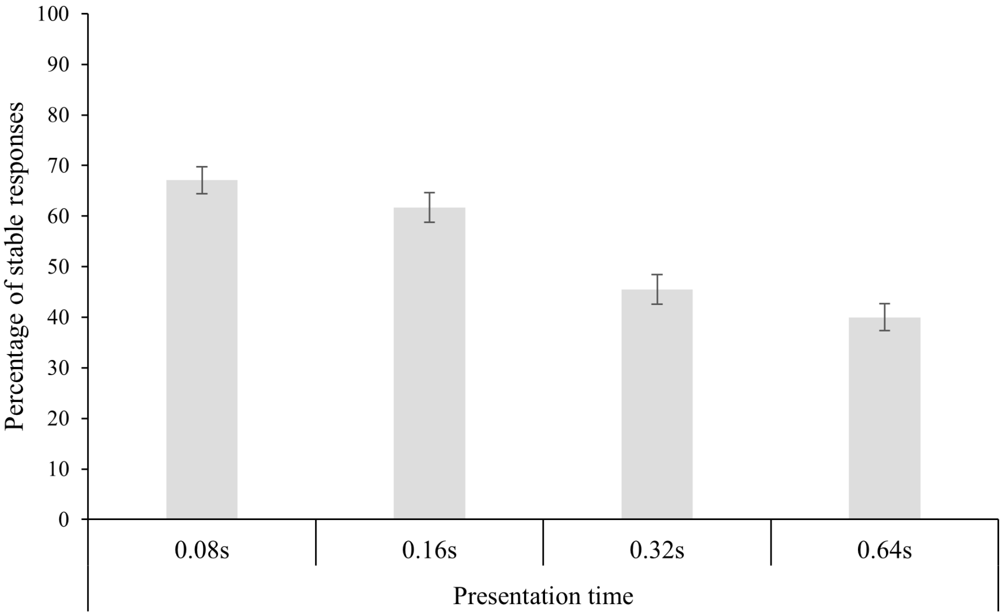

Figure 3.10 The percentage of stable responses for each presentation time of the unambiguous stimulus collapsed across ISI, for the replication block. Error bars represent SE.

Table 3.6

Full repeated measures ANOVA results for the percentage of stable trials in the replication block (significant p-values are indicated in bold font)

	df1	*df2*	*F*	*p*	η_p^2
Unambiguous stimulus presentation time	2.030	21	50.337	**< .001**	.706
ISI	1	21	116.886	**< .001**	.848
Unambiguous stimulus presentation time x ISI	3	63	1.208	.314	.054

3.3.3.2.1 Replication block – Correct responses to directional motion

In the same way as with Experiment One, the average percentage of responses where the participant responded correctly to the unambiguous directional motion stimuli on each trial were submitted to a 2 x 4 repeated measures ANOVA. This revealed a significant main effect of directional stimulus presentation time on the percentage of trials that participants correctly identified the direction of motion, $F(3,63) = 16.222$, $p < .001$, $\eta_p^2 = .436$. There was no effect of ISI, $p = .712$ and no interaction effect present, $p = .504$.

Post hoc comparisons of the effect of presentation time collapsed across ISI revealed that the 0.08s presentation time $M = 45.47\%$ ($SD = 14.96$), led to significantly fewer correct responses to directional motion stimuli than the 0.16s presentation time, $M = 58.99\%$ ($SD = 20.17$), $t(21) = 4.479$, $p < .001$, $d = 0.710$, the 0.32s presentation time, $M = 65.47\%$ ($SD = 18.17$), $t(21) = 6.626$, $p < .001$, $d = 1.051$, and the 0.64s presentation time, $M = 60.88\%$ ($SD = 21.54$), $t(21) = 5.103$, $p < .001$, $d = 0.809$. All other combinations of presentation times did not differ significantly (all $p > .214$). See also Figure 3.11 below.

To determine whether participants were able to determine the correct direction of unambiguous motion more often than would be expected by chance (50%), the percentage of correct responses for each combination of ISI and presentation time were submitted to individual one-sample t-tests. These revealed that, within the 0.04s ISI, correct responses greater than chance occurred within presentation times of 0.16s, $M = 60.12\%$ ($SD = 15.73$), $t(21) = 3.017$, $p = .007$, $d = 0.643$ and 0.32s, $M = 64.11\%$ ($SD = 15.51$), $t(21) = 4.266$, $p < .001$, $d = 0.910$ but not within presentation times of 0.08s, $p = .167$ and 0.64s, $p = .097$. Within the 0.12s ISI, correct responses to the directional motion greater than chance occurred within presentation times of 0.32s, $M = 66.83\%$ ($SD = 20.77$), $t(21) = 3.803$, $p = .001$, $d = 0.811$ and 0.64s, $M = 63.15\%$ ($SD = 20.02$), $t(21) = 3.080$, $p = .006$, $d = 0.657$ but not 0.08s, $p = .178$, and 0.16s, $p = .141$. These are marked on Figure 3.11, below.

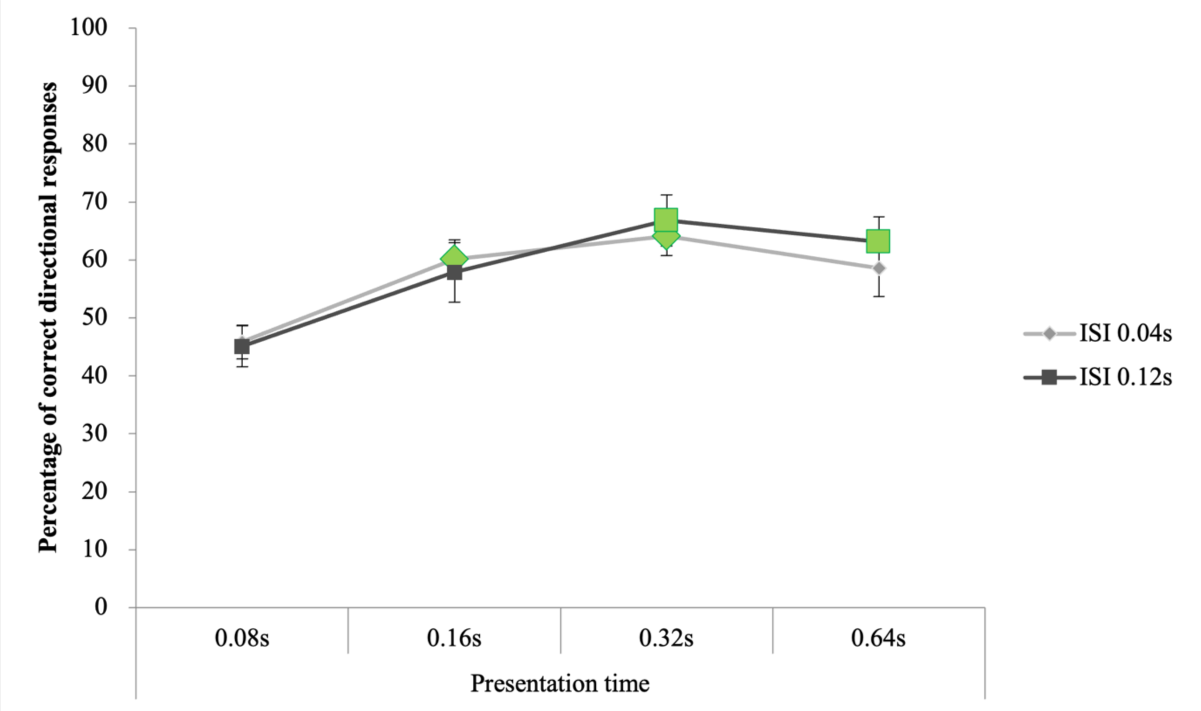

*Figure 3.11*The percentage of correct responses to the unambiguous directional motion stimuli in the replication block. Those greater than chance are marked with green markers. Error bars represent SE.

3.3.3.3 Final perception block

In the final perception block, trials were coded as stable or reversal depending on whether the participant's response matched the direction of the unambiguous stimulus. The percentage of stable trials for each combination of ISI (0.04 and 0.12s) and unambiguous stimulus presentation time (0.08, 0.16, 0.32 and 0.64s) were submitted to a 2x4 repeated measures ANOVA. This revealed significant main effects of both ISI, $F(1,21) = 44.050$, $p < .001$, $\eta_p^2 = .677$, and unambiguous stimulus presentation time, $F(3,63) = 6.953$, $p < .001$, $\eta_p^2 = .249$. A significant interaction effect was also found, $F(3,63) = 3.833$, $p = .014$, $\eta_p^2 = .154$. See Figure 3.12 below for a graph showing the full results. For the full RM ANOVA table see Table 3.7 below.

Overall, participants were more likely to report a stable perception when the ISI was 0.04s, *M* = 52.07% (*SD* = 13.02), compared to 0.12s, *M* = 39.97% (*SD* = 14.90), *t(21)* = 6.637, $p < .001$, *d* = 0.917. To further explore the interaction effect, two one-way ANOVAs were conducted to examine the effect of presentation time, separately for each ISI. Within a 0.04s ISI, an effect of presentation time was revealed, *F*(3,21) = 6.670, *p* = .002, $\eta_p^2 = .241$. Post hoc comparisons showed that participants were more likely to report stability in their perception when the presentation time was 0.08s, *M* = 54.34% (*SD* = 8.34), than they were when the presentation time was 0.64s, *M* = 45.36% (*SD* = 13.84), *t(21)* = 2.821, *p* = .038, *d* = 0.736. Additionally, when the presentation time was 0.16s, *M* = 58.84% (*SD* = 8.34); this also led to a significant increase in reports of a stable perception compared to both the 0.32s, *M* = 49.74% (*SD* = 13.56, *t(21)* = 2.857, *p* = .009, *d* = 0.745, and 0.64s, *t(21)* = 4.234, $p < .001$, *d* = 1.104, conditions. All other pairwise comparisons within the 0.04s ISI condition were not significant (all *p*s > .921). In the 0.12s ISI condition, an effect of presentation time was also revealed, *F*(3,21) = 5.294, *p* = .007, $\eta_p^2 = .201$. Here, participants were more likely to report a stable perception when the presentation time was 0.08s, *M* = 48.36% (*SD* = 7.72), than they were when the presentation time was 0.32s, *M* = 34.18% (*SD* = 15.27), *t(21)* = 3.728, *p* = .001, *d* = 1.003, and 0.64s, *M* = 36.72% (*SD* = 18.06), *t(21)* = 3.061, *p* = .006, *d* = 0.824, but all other pairwise comparisons within the 0.12s ISI condition were not significant (all *p*s > .274).

To determine whether responses indicated VMP or rMAE, the percentage of responses that indicated a stable perception were compared against chance (50%) using one-sample t-tests. If the percentage of stable responses were greater than chance, this is indicative of VMP, whereas a percentage significantly less than chance indicates rMAE. These t-tests revealed that, within the 0.04s ISI, the percentage of VMP responses greater than chance occurred within presentation times of 0.08s, *M* = 54.34% (*SD* = 8.34), *t(21)* = 2.442, *p* = .024, *d* = 0.521 and 0.16s, *M* = 58.84% (*SD* = 12.29), *t(21)* = 3.374, *p* = .003, *d* = 0.719. Within the 0.12s ISI, stable responses significantly lower than chance level were shown with presentation times of 0.16s, *M* = 40.60% (*SD* = 13.43), *t(21)* = 3.282, *p* = .004, *d* = 0.700, 0.32s, *M* = 34.18% (*SD* = 15.27), *t(21)* = 4.859, *p* < .001, *d* = 1.036, and 0.64s, *M* = 36.72% (*SD* = 18.06), *t(21)* = 3.449, *p* = .002, *d* = 0.735. All other combinations of ISI and presentation time were not significantly different from chance level (all *ps* > .131). These are marked on Figure 3.12, below.

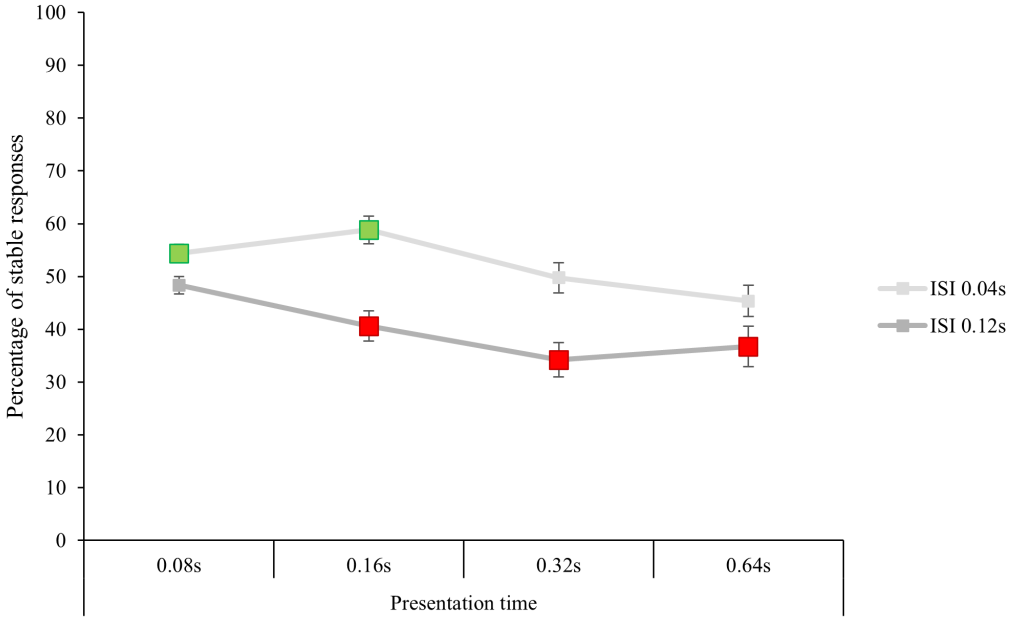

Figure 3.12 The percentage of stable responses for each ISI of the unambiguous stimulus at all presentation times of the unambiguous stimulus, for the final perception block. Green markers highlight conditions leading to above-chance stable responses (therefore indicating VMP) and red markers highlight conditions

leading to below-chance stable responses (therefore indicating rMAE). Error bars represent SE.

Table 3.7

Full repeated measures ANOVA results for the percentage of stable trials in the final perception block (significant p-values are indicated in bold font)

	df1	*df2*	*F*	*p*	η_p^2
Unambiguous stimulus presentation time	3	63	6.953	**<.001**	.249
ISI	1	21	44.050	**< .001**	.677
Unambiguous stimulus presentation time x ISI	3	63	3.833	**.014**	.154

3.3.3.3 Counting block

Responses in the counting block were coded simply as a report of perceiving one or two stimuli on each trial. The percentage of reports of perceiving one stimulus, for each combination of ISI (0.04 and 0.12s) and unambiguous stimulus presentation time (0.08, 0.16, 0.32 and 0.64s) were submitted to a 2x4 repeated measures ANOVA. This revealed significant main effects of both ISI, $F(1,21) = 575.230$, $p < .001$, $\eta_p^2 = .965$, and unambiguous stimulus presentation time, $F(2.01,63) = 30.192$, $p < .001$, $\eta_p^2 = .590$. A significant interaction effect was also found, *F*(3,63) = 27.084, $p < .001$, $\eta_p{}^2 = .563$. For the full RM ANOVA table see Table 3.8 below.

Overall, participants were significantly more likely to report that they saw one stimulus in the 0.04s ISI condition, *M* = 64.26% (*SD* = 18.09), compared to the 0.12s, *M* = 8.68% (*SD* = 6.55), *t(21)* = 23.984, $p < .001$, *d* = 4.973. To further explore the interaction effect, two one-way ANOVAs were conducted to examine the effect of presentation time, separately for each ISI. Within a 0.04s ISI, an effect of presentation time was revealed, *F*(3,18) = 36.040, $p < .001$, $\eta_p{}^2 = .632$. Post hoc comparisons showed that there was no significant difference in the likelihood of participants reporting one stimulus when the presentation time was 80ms, *M* = 76.19% (*SD* = 13.10), compared to 160ms, *M* = 74.01% (*SD* = 10.18), $p > .999$. Nor was there a difference when the presentation time was 0.32s, *M* = 56.24% (*SD* = 15.74) compared to 0.64s, *M* = 50.62% (*SD* = 17.95), *p* = .395. However, participants were more likely to report seeing one stimulus when the presentation time was 0.08s, than they were when the presentation time was 0.32s *t(21)* = 6.640, $p < .001$, *d* = 1.372, or 0.64s, *t(21)* = 8.512, $p < .001$, *d* = 1.759. The same was true when the presentation time was 0.16s; this also led to a significant increase in reports of one stimulus compared to both the 0.32s, *t(21)* = 5.917, $p < .001$, *d* = 1.223, and 0.64s, *t(21)* = 7.788, $p < .001$, *d* = 1.609, conditions. In the 0.12s ISI condition, an effect of presentation time was also revealed, *F*(3,18) = 5.548, *p* = .007, $\eta_p{}^2 = .209$. Here, participants were more likely to report seeing one stimulus when the presentation time was 0.08s, *M* = 12.56% (*SD* = 7.92), than they were when the presentation time was 0.32s, *M* = 6.28% (*SD* = 4.80), *t(21)* = 3.830, $p < .001$, *d* = 1.011, and 0.64s, *M* = 7.46% (*SD* = 6.71), *t(21)* = 3.114, *p*

= .017, *d* = 0.822, but all other pairwise comparisons within the 0.12s ISI condition were not significant (all *p* > .086).

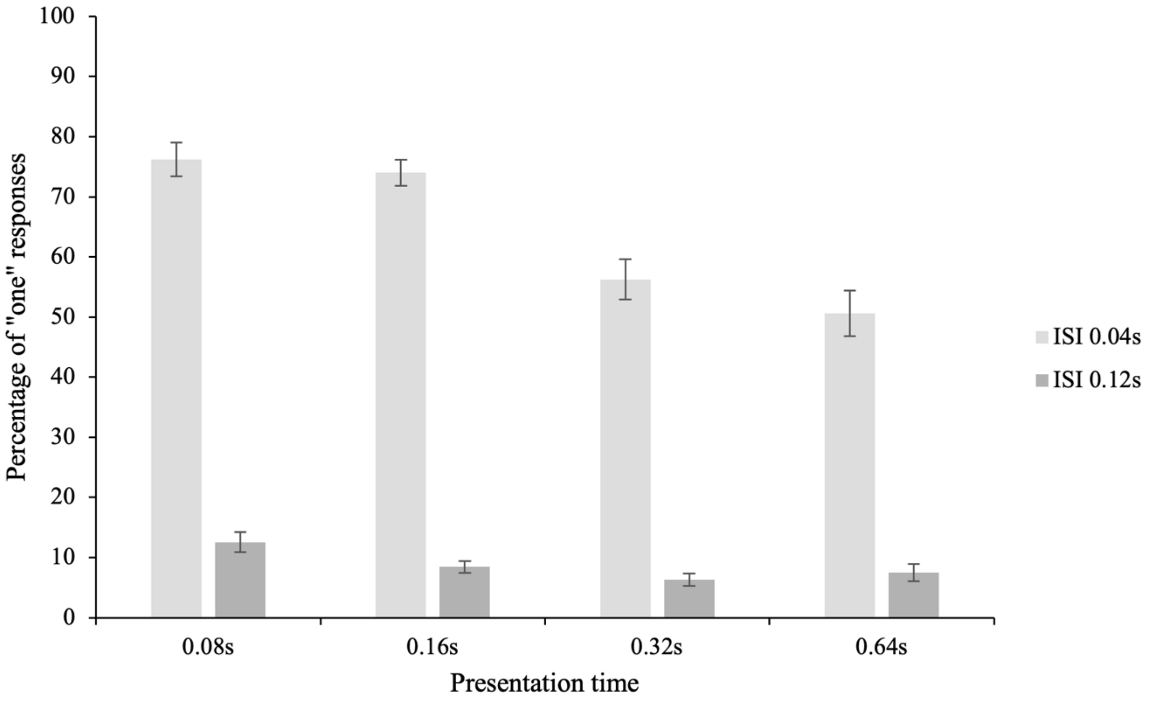

Figure 3.13 The percentage of trials in which participants reported that they saw one stimulus, for each ISI and presentation time of the unambiguous stimulus. Error bars represent SE.

Table 3.8

Full repeated measures ANOVA results for the percentage of trials in which participants reported perceiving only one stimulus in the counting block (significant p-values are indicated in bold font)

	df1	*df2*	*F*	*p*	η_p^2
Unambiguous stimulus presentation time	2.010	63	30.192	**< .001**	.590
ISI	1	21	575.230	**< .001**	.965
Unambiguous stimulus presentation time x ISI	3	63	27.084	**< .001**	.563

3.3.3.4 Dual task block

In the dual task block, participants were asked to respond to indicate the direction of motion of *all* the stimuli that they perceived in the trial (for full instructions given to participants and a full description of the trial see section 3.2, Methods, above). If the participant responded with two key presses, that indicated they saw two stimuli. For example, a participant might respond to indicate motion to the left followed by motion to the right, by pressing “up” then “down”, or motion to the right followed by motion to the right, by pressing “down” then “down”. Both would be examples of a response that indicated two stimuli were perceived. If the participant responded with one key press, that indicated they saw one stimulus. For example, participants might respond to indicate motion to the left by pressing “up”.

Therefore, because participants’ responses provided information on how many stimuli they saw and in which direction they saw them moving, they were coded according to three categories. The first category was one stimulus, meaning participants responded to indicate that they saw only one stimulus moving in any direction. The second category was two stimuli stable, meaning that participants responded to indicate that they saw two stimuli and that they were moving in the same direction. The final category was two stimuli reversal, meaning that participants responded to indicate that they saw two stimuli, and those stimuli were moving in the opposite directions.

The aim of this block was to examine whether participants would respond to indicate that they saw one stimulus at short presentation times and ISIs, and two stimuli at longer durations, when they were only instructed to report their subjective experience on each trial. In other words, when they were not asked to focus on a particular aspect of the trial (i.e., how many stimuli there were as they did in the counting block, which direction each of the *two* stimuli were moving as in the replication block, or which direction the final stimulus was moving as in the final direction block).

3.3.3.4.1 Dual task block – “One stimulus” responses

As described in the Design section above, in this block there were always two stimuli presented to participants. Therefore, any reports of participants perceiving only one stimulus were erroneous. To address the aim of investigating whether participants were more likely to make this erroneous report of “one” stimulus at shorter presentation times and ISIs, the percentage of "one stimulus" responses, for each combination of ISI (0.04 and 0.12s) and unambiguous stimulus presentation time (0.08, 0.16, 0.32 and 0.64s) were submitted to a 2x4 repeated measures ANOVA. This revealed significant main effects of both ISI, $F(1,21) = 163.112$, $p < .001$, $\eta_p^2 = .886$, and unambiguous stimulus presentation time, $F(2.12,63) = 9.751$, $p < .001$, $\eta_p^2 = .317$. A significant interaction effect was also found, $F(1.91,63) = 22.847$, $p < .001$, $\eta_p^2 = .521$ (see Figure 3.15, below). For the full RM ANOVA table see Table 3.10 below.

Overall, participants were significantly more likely to report that they saw one stimulus in the 0.04s ISI condition, $M = 59.97\%$ ($SD = 20.19$), compared to the 0.12s, $M = 17.05\%$ ($SD = 10.53$), $t(21) = 12.772$, $p < .001$, $d = 2.917$ (see also Figure 3.15, below). To further explore the interaction effect, two one-way ANOVAs were conducted to examine the effect of presentation time, separately for each ISI. Within a 0.04s ISI, an effect of presentation time was revealed, $F(1.89,18) = 17.600$, $p < .001$, $\eta_p^2 = .456$. Post hoc comparisons showed that there was no significant difference in the likelihood of participants reporting one stimulus when the presentation time was 80ms, $M = 69.32\%$ ($SD = 18.02$), compared to 160ms, $M = 69.20\%$ ($SD = 19.85$), $p > .999$. Nor was there a difference when the presentation time was 0.32s, $M = 49.40\%$ ($SD = 14.85$) compared to 0.64s, $M = 51.94\%$ ($SD = 19.59$), $p > .999$. However, participants were more likely to report seeing one stimulus when the presentation time was 0.08s, than they were when the presentation time was 0.32s $t(21) = 5.480$, $p < .001$, $d = 1.096$, or 0.64s, $t(21) = 4.781$, $p < .001$, $d = 0.956$. The same was true when the presentation time was 0.16s; this also led to a significant increase in reports of one stimulus compared to both the 0.32s, $t(21) = 5.448$, $p < .001$, $d =$

1.089, and 0.64s, t(21) = 4.748, p < .001, d = 0.949, conditions. In the 0.12s ISI condition, an effect of presentation time was also revealed, *F*(2.21,18) = 7.476, *p* = .001, η_p^2 = .263. Here, participants were more likely to report seeing one stimulus when the presentation time was 0.64s, *M* = 23.00% (*SD* = 11.03), than they were when the presentation time was 0.32s, *M* = 14.26% (*SD* = 10.23), *t(21)* = 4.219, *p* < .001, *d* = 0.865, 0.16s, *M* = 15.45% (*SD* = 9.79), *t(21)* = 3.642, *p* = .003, *d* = 0.746, and 0.08s, *M* = 15.50% (*SD* = 9.29), *t(21)* = 3.619, *p* = .002, *d* = 0.742, but all other pairwise comparisons within the 0.12s ISI condition were not significant (all *p*s > .999).

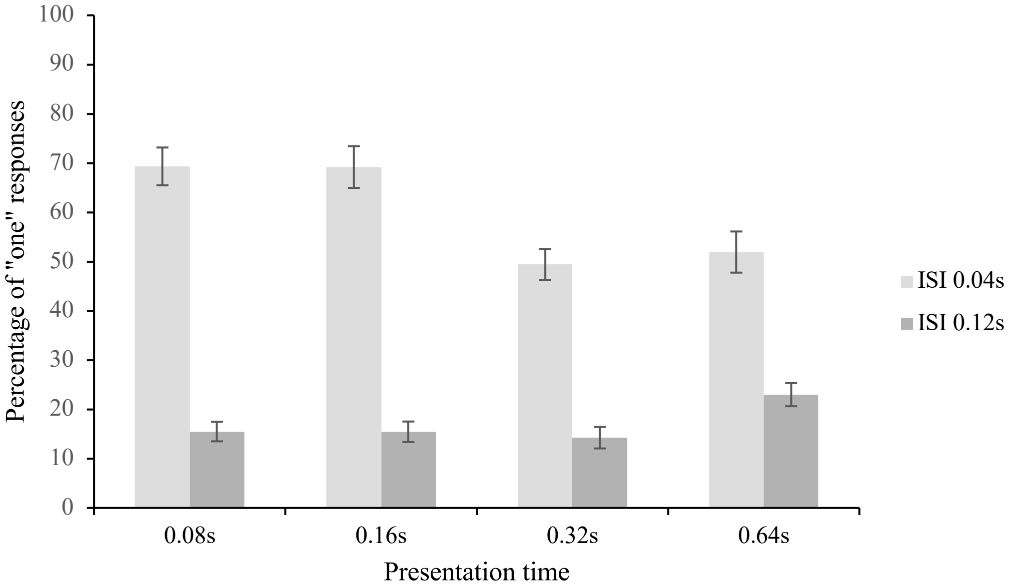

Figure 3.15 The percentage of “one” responses in the dual task condition, for each ISI and presentation time of the unambiguous stimulus. Error bars represent SE.

Table 3.10

Full repeated measures ANOVA results for the percentage of trials in which participants reported perceiving only one stimulus in the dual task block (significant p-values are indicated in bold font)

	df1	*df2*	*F*	*p*	η_p^2
Unambiguous stimulus presentation time	2.116	63	9.751	**< .001**	.317
ISI	1	21	163.112	**< .001**	.886
Unambiguous stimulus presentation time x ISI	1.910	63	22.847	**< .001**	.521

3.3.3.4.2 Dual task block - "two stimuli" responses

To examine participants' perception on trials where they correctly identified that there were two stimuli, these trials were examined independently of the erroneous "one stimulus" response trials detailed above. These trials were analysed in the same way as the replication block, in order to determine whether VMP was evoked when participants were completely free to report their subjective perception without being restricted by task instructions.

Firstly, the proportion of stable trials, of the trials where two stimuli were identified, for each combination of ISI (0.04 and 0.12s) and unambiguous stimulus presentation time (0.08, 0.16, 0.32 and 0.64s) were submitted to a 2x4 repeated measures ANOVA. This revealed significant main effects of both ISI, $F(1,21) = 13.261$, $p = .002$, $\eta_p^2 = .387$, and unambiguous stimulus presentation time, $F(3,63) = 7.501$, $p < .001$, $\eta_p^2 = .263$. A significant interaction effect was also found, $F(3,63) = 7.654$, $p < .001$, $\eta_p^2 = .267$ (see Figure 3.16, below). For the full RM ANOVA table see Table 3.11 below.

To further explore the interaction effect, two one-way ANOVAs were conducted to examine the effect of presentation time, separately for each ISI. Within a 0.04s ISI, an effect of presentation time was revealed, $F(3,22) = 3.273$, $p = .027$, $\eta_p^2 = .135$. Post hoc comparisons showed that participants were more likely to report a stable perception when the presentation time was 0.32s, $M = 41.70\%$ ($SD = 23.85$), than 0.64s, $M = 23.35\%$ ($SD = 16.68$), $t(21) = 3.073$, $p =.019$, $d = 0.807$. However, all other pairwise comparisons of presentation times were not significant (all *p*s > .312; see also Figure 3.16). In the 0.12s ISI condition, an effect of presentation time was also revealed, $F(3,22) = 17.868$, $p < .001$, $\eta_p^2 = .460$. Here, participants were more likely to report a stable perception when the presentation time was 0.08s, $M = 54.45\%$ ($SD = 18.60$), than they were when the presentation time was 0.32s, $M = 35.86\%$ ($SD = 20.99$), $t(21) = 4.827$, $p < .001$, $d = 0.972$, and 0.64s, $M = 28.91\%$ ($SD = 17.16$), $t(21) = 6.630$, $p < .001$, $d = 1.335$. Similarly, participants were more likely to report a stable perception when the presentation time was 0.16s, $M = 47.82\%$ ($SD = 19.60$), than they were when the presentation time was 0.32s, $t(21) = 3.105$, $p = .005$, $d = 0.625$, and 0.64s, $t(21) = 4.909$, $p < .001$, $d = 0.988$. All other pairwise comparisons within the 0.12s ISI condition were not significant (all *p*s > .457; see also Figure 3.16).

To determine whether responses indicated VMP or rMAE, the percentage of responses that indicated a stable perception were compared against chance

(50%) using one-sample t-tests. If the percentage of stable responses were greater than chance, this is indicative of VMP, whereas a percentage significantly less than chance indicates rMAE. These t-tests revealed that there were no conditions that led to stable responses greater than chance, therefore no combination of ISI and presentation time led to responses consistent with VMP. However, the t-tests revealed that several conditions led to the percentage of stable responses being significantly lower than chance and so indicative of rMAE. Within the 0.04s ISI, these were presentation times of 0.08s, *M* = 33.28% (*SD* = 28.21), *t(21)* = 2.780, *p* = .011, *d* = 0.593, 0.16s, *M* = 29.87% (*SD* = 20.62), *t(21)* = 4.579, *p* < .001, *d* = 0.976, and 0.64s, *M* = 23.35% (*SD* = 16.68), *t(21)* = 7.493, *p* < .001, *d* = 1.598. Within the 0.12s ISI, presentation times of 0.64s, *M* = 28.91% (*SD* = 17.16), *t(21)* = 5.766, *p* < .001, *d* = 1.229, and 0.32s, *M* = 35.86% (*SD* = 20.99), *t(21)* = 3.161, *p* = .005, *d* = 0.674, also led to significantly lower percentages of stable responses than would be expected by chance. All other combinations of ISI and presentation times did not result in a percentage of stable responses significantly different from chance (all *p* > 117; these are marked on Figure 3.17 below).

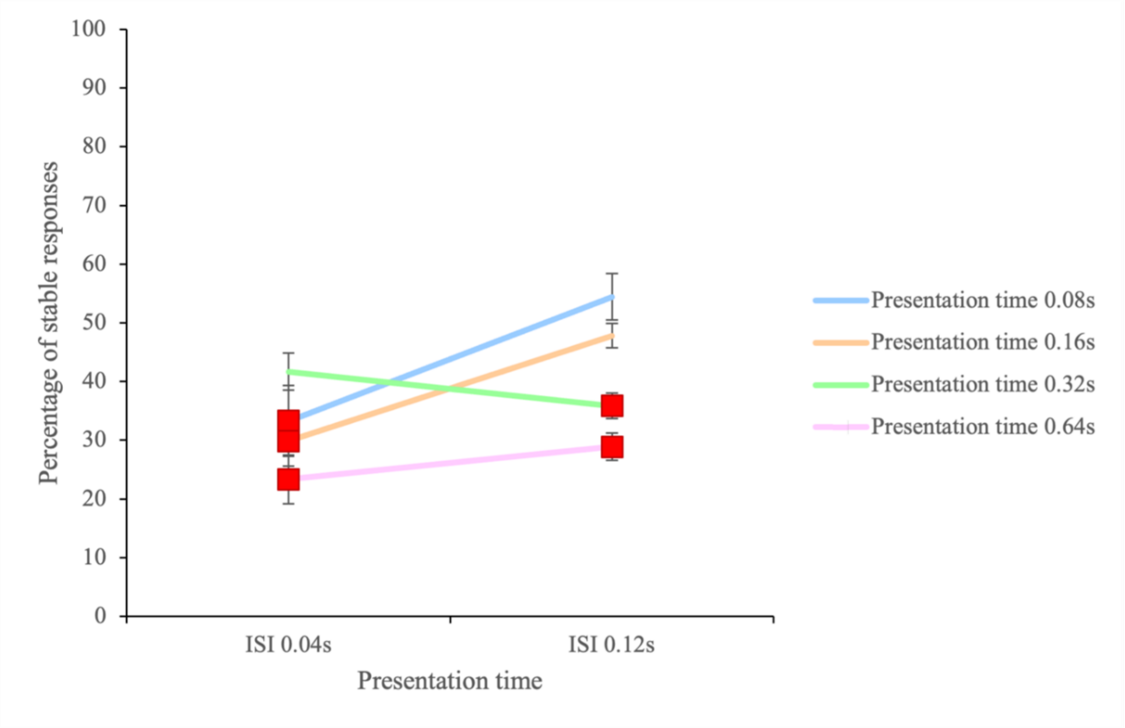

Figure 3.16 The percentage of stable responses given, for each combination of ISI and presentation time, out of the trials only where participants reported seeing two stimuli. Green markers highlight conditions leading to above-chance stable responses (therefore indicating VMP) and red markers highlight conditions leading to below-chance stable responses (therefore indicating rMAE). Error bars represent SE.

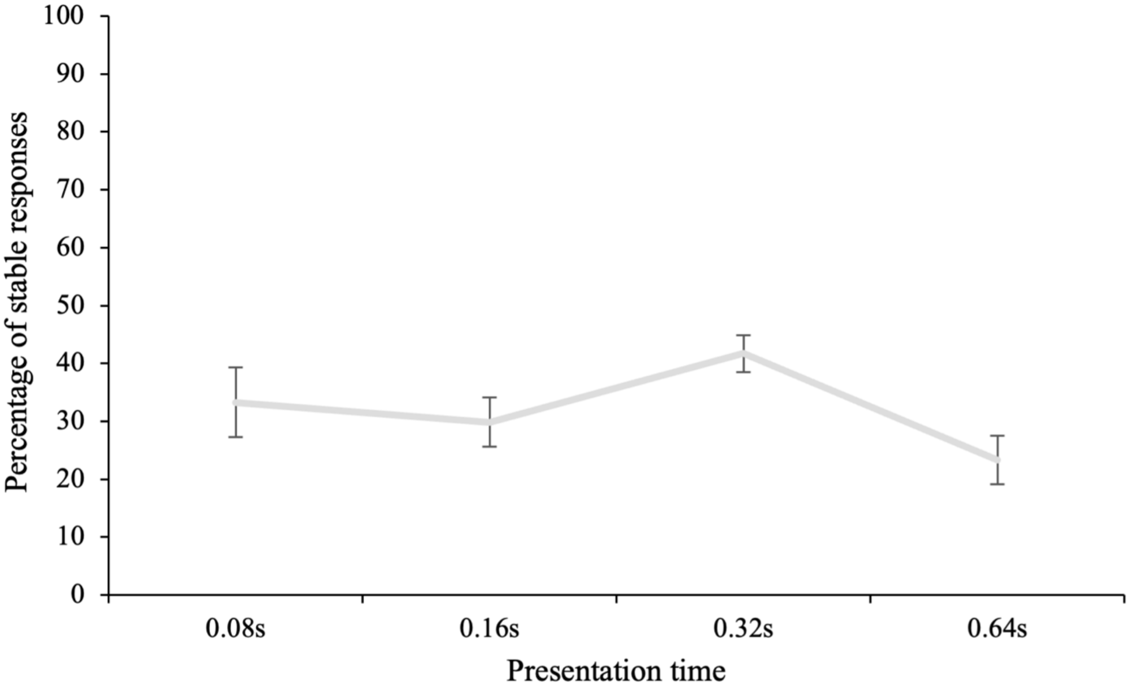

Figure 3.17 The percentage of stable responses given at each presentation time, for the 0.04s ISI condition only, out of the trials only where participants reported seeing two stimuli. Error bars represent SE.

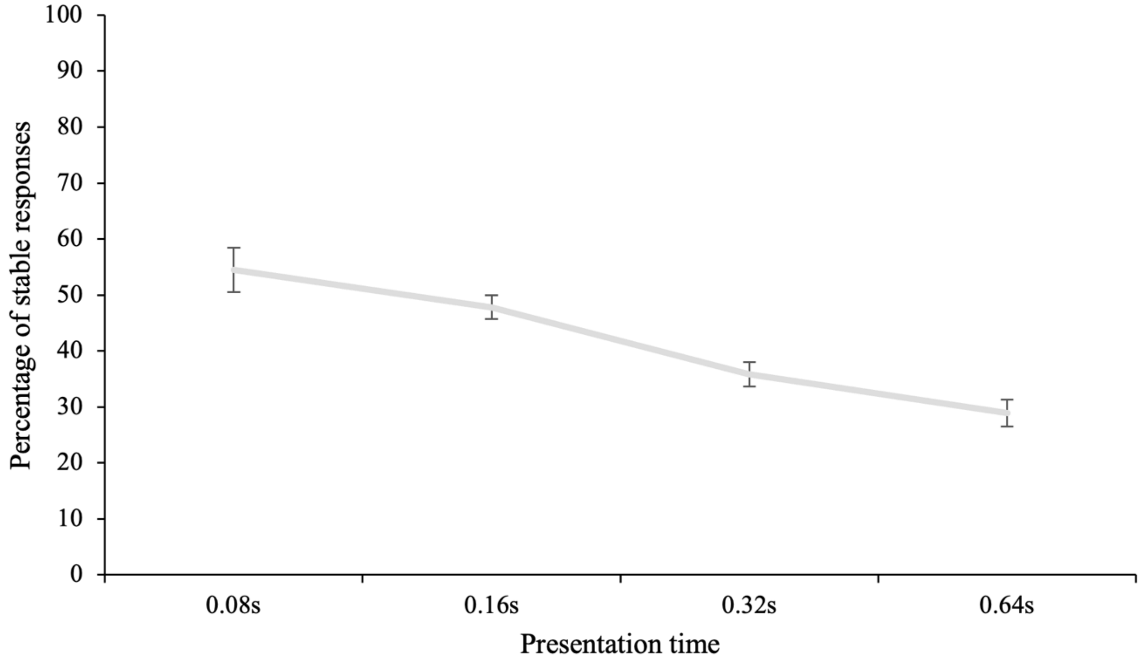

*Figure 3.18*The percentage of stable responses given at each presentation time, for the 0.12s ISI condition only, out of the trials only where participants reported seeing two stimuli. Error bars represent SE.

Table 3.12

Full repeated measures ANOVA results for the percentage of trials in which participants reported perceiving two stimuli, and the stimuli were moving in the same direction (i.e., a "stable" response) in the dual task block (significant p-values are indicated in bold font)

	df1	*df2*	*F*	*p*	η_p^2
Unambiguous stimulus presentation time	3	63	7.501	**< .001**	.263
ISI	1	21	13.261	**.002**	.387
Unambiguous stimulus presentation time x ISI	1.910	63	7.654	**< .001**	.267

3.3.4 Discussion

Informal observations of the Experiment 1 stimuli with short ISIs and short directional stimulus durations suggested that participants may not always be able to resolve two separate stimuli in these displays and may, at least sometimes, perceive only one stimulus. However, based on the a priori assumption that participants would perceive two stimuli, in Experiment 1, they were required to make two responses to indicate the perceived direction for each of the two stimuli. This creates an ambiguity in the interpretation of the results. Did participants make "same" responses (e.g., left/left) because they perceived two stimuli and both were moving left and thus potentially priming? Or, alternatively, were participants making “same” responses because when they saw only one stimulus moving in one direction then their best response option was to respond twice with the same answer? This second account is much less compatible with the idea of priming and indicates that at least some of the “same” responses associated with priming in Experiment 1, and previous work, may be artefactual and arising from incompatibility of the perceptual experience and response options. Several different response options were used in Experiment 2 to clarify this issue.

First, results from the replication block suggest that Experiment 2 was successful in replicating the pattern of results from Experiment 1 at short ISIs and short directional stimulus durations. There were a greater number of VMP (“same”) responses to ambiguous motion stimuli when presented following brief presentation of unambiguous motion after a shorter ISI.

Next, the results from the counting block clearly showed that participants were more likely to report perceiving only one stimulus when directional stimulus presentation and ISI were brief. For the 0.04s ISI, on average, over 60% of trials were incorrectly reported as having just one stimulus. When the ISI was increased to 0.12s, the percentage of “one stimulus” responses dramatically fell to below 10%, on average. The duration of the directional stimulus also had an effect with shorter directional stimuli leading to more “one stimulus” responses. In the condition with both the shortest ISI and the shortest directional stimulus duration, nearly 75% of trials were reported as “one stimulus”, on average. These results are in line with the hypothesis that participants are largely unable to resolve two separate stimuli in displays with short ISIs and short duration directional stimuli. They overwhelmingly perceive just one stimulus. This clearly undermines direct interpretation of the Experiment 1 data as simply reflecting priming. These results were further supported by the results of the dual task block. Overall, participants were more likely to report seeing one stimulus in the 0.04s ISI condition than in the 0.12s ISI condition. This effect was particularly prominent in the 0.04s ISI condition when coupled with shorter directional stimulus presentation times. When the ISI increased to 0.12s and participants responded to indicate that they saw two stimuli, they were more likely to respond to indicate VMP at shorter presentation times and rMAE at longer ones in a pattern similar to that seen in the replication block. As there were always two stimuli presented in the dual task block, any "one stimulus" responses were erroneous. By removing the erroneous "one stimulus" responses, we can get a purer estimate of the amount of priming that was present. This showed that there were no combinations of ISI and directional stimulus presentation time that led to responses indicating VMP. In this case, even short presentation times and ISIs led to responses indicative of rMAE (i.e., more reversals in perception than expected by chance) rather than VMP.

Finally, we followed an approach by Takeuchi et al. (2011) to get an estimate of VMP which should be less sensitive to whether participants perceive both stimuli. This is done by focusing just on the final perception of the second stimulus. If VMP is taking place, then it should be observable in the reports of the second stimulus regardless of whether the first stimulus is perceived. Results from the final perception block showed that participants were more likely than chance to report VMP when the ISI was 0.04s and the directional stimulus was presented for 0.08 or 0.16s. Although, this effect was much weaker than in the replication block where around 75% of trials led to a VMP response when the presentation time was 0.08s, compared to 54% in the final perception block. Additionally, the strongest VMP response in this block was evoked with a presentation time of

0.16s, rather than 0.08s as in the replication block. When the ISI increased to 0.12s, only rMAE was shown when the directional stimulus was presented for longer durations. Unlike in the replication block, no VMP was indicated at all in this condition, even when the presentation time of the directional stimulus was very brief.

In addition to affecting the interpretation of the results from Experiment 1, these results have implications for the results of Kanai and Verstraten's (2005) study. Kanai and Verstraten (2005) reported that when ambiguous motion is preceded by brief directional motion for 0.08s with a 0.04s ISI participants responded to indicate VMP on most trials, when asked if the stimuli were moving in the same or different directions. The results from the counting block in Experiment 2 suggest that under the same conditions, participants report that they only perceive one stimulus at all on most trials. This means that when Kanai and Verstraten (2005) asked their participants whether the stimuli in each trial were moving in the same or different directions, they could have been responding to indicate that they were moving in the same direction (i.e., a VMP response) whilst only perceiving one stimulus at all. Moreover, during the dual task block where participants were instructed to report their perception completely subjectively, they were also far more likely to report that they saw only one stimulus under these conditions. One possible explanation for the results that are seen in the counting and dual task blocks is that the participants are responding with their 'best guess' under these very brief timing conditions. For instance, if the participant has noticed only *one* stimulus but is asked a question that requires them to consider the motion of *two* stimuli (i.e., by being asked whether the two stimuli were moving in the same or different directions, or being asked to report the direction of the two stimuli), it is logical to assume that the best guess in this case would be to respond to indicate that you saw motion in the same direction. In other words, if participants only saw one stimulus but are being asked about two, they assume that those two stimuli must have been moving in the same way and respond accordingly. If the results from the correct responses to directional motion from Experiment 1 and the replication block in Experiment 2 are considered along with this assumption, this could be likely. In both cases, participants were significantly more able to correctly respond to the unambiguous directional motion stimulus when the stimulus presentation time and ISI were longer but were less able to do so when the ISI and presentation times were shorter. This result indicates that under these circumstances, participants may be unsure about the direction of motion that they see in the trials and therefore the responses that they make on them may be based on a guess. For example, in trials where the participant has incorrectly reported the direction of the unambiguous motion this suggests that they did not perceive that motion. Therefore, if they also responded to indicate VMP on such trials, this is likely to be due to a failure to perceive one of these stimuli rather than a genuine VMP response.

The results from the final perception block do not show the same pattern as the replication block in Experiment 2, nor the pattern from Experiment 1. As well as being much weaker overall, the strongest VMP response in the final perception block was induced by presenting unambiguous motion for 0.16s with an ISI of 0.04s, rather than presenting for just 0.08s as was the case in both Experiment 1 and the replication block of Experiment 2. The results from this block also do not replicate those of the original authors. Takeuchi et al. (2011) reported results that were comparable to Kanai and Verstraten's (2005), however these are not seen here. As noted previously however, the exact conditions of the present experiment may not have been close enough to those used by Takeuchi et al (2011) to directly compare the results. For instance, Takeuchi et al (2011) instructed their participants to fixate away from the stimulus to avoid foveal stimulus presentation. Here, participants were instructed to fixate in the centre of the stimulus in the same manner as in Kanai and Verstraten (2005).

The results from this experiment are clearly relevant when interpreting the results of experiments involving a direct replication of Kanai and Verstraten's (2005) design. However, these results should also urge caution to be taken when interpreting the results of experiments where participants are forced to make a judgement about a stimulus that they might not be able to perceive, leading to the "best guess" effect described above. Many studies investigating VMP have avoided this by directly comparing participants' reports of the direction of an ambiguous stimulus with the known direction of a priming stimulus (e.g., Long, Toppino & Mondin, 1992; Pinkus & Pantle, 1997; Pantle, Gallogly & Pieler, 2000; Takeuchi et al., 2011; Piedimonte et al., 2015). However, there are studies that have employed a similar method to Kanai and Verstraten (2005) and asked participants to report whether they perceive two stimuli moving in the same or different directions. For instance, Pavan et al. (2009) directly replicate the results by Kanai and Verstraten (2005) and extend these findings to second-order motion. However, for judgements of both first and second order motion, the authors determined whether VMP was present by asking participants to report whether they saw the sine-wave stimulus pairs moving in the same or different directions. The results of this study show that these results could be affected by this "best guess" effect. The recommendation here would be to replicate such studies, using a paradigm whereby the reported direction of the ambiguous stimulus is compared directly to the known motion of the directional stimulus.

It should be noted however, that in both experiments described in this chapter participants generally did not reliably discriminate the direction of the unambiguous motion presented in the first stimulus in each trial. Attention check trials were included in this task, with a pre-defined pass rate of 70%. This suggested that the participants whose data were included in analysis were paying attention throughout the task but perhaps found the required response confusing given that they were asked to respond using the up and down arrow keys to report left- and right-ward motion. It was acknowledged before collecting data for this experiment that this response method would add a given amount of difficulty to the task, but that this difficulty would be consistent across all conditions and participants. Therefore, it could be this added burden on the cognitive load of participants that was causing the decreased ability to discriminate the direction of the unambiguous motion generally across all conditions. However, we cannot be certain of this with the current design. This response protocol was originally used so that responses indicating motion in a particular direction were able to be decoded. For instance, responses to indicate a reversal from right- to left-ward motion. Therefore, response buttons needed to be counterbalanced to avoid any signal from the manual response itself being responsible for any increase in decoding accuracy. However, there were not sufficient trial numbers in each class to run this analysis. So, in future experiments, the response protocol could be simply to press the left arrow to indicate motion to the left and vice versa, as there would be no need for counterbalancing. This would reduce the cognitive load of participants and help to determine whether it was indeed the difficulty of the response protocol that was causing the reduction in participants' ability to reliably discriminate the unambiguous directional motion.

Requiring participants to report the direction of each of the two stimuli in the pair, however, did allow us to determine some measure of response accuracy. For example, we were able to analyse the number of trials in each condition where the direction of the first stimulus was correctly reported. In contrast, by asking participants only to report whether the two stimuli were moving in the "same" or "different" directions, Kanai and Verstraten (2005) did not provide such an analysis.

A further limitation of Experiment 1 in this chapter lies in the large number of conditions examined, which subsequently led to a reduced number of trials within each individual condition. Specifically, within the 25 distinct timing conditions detailed in section 3.2.1.3, each condition contained only 20 trials which was lower than the 40 trials per condition in Kanai and Verstraten (2005). This constraint was necessary to prevent participants from experiencing fatigue or disengagement due to the longer response duration, and demand of our response regime which, unlike Kanai and Verstraten (2005), required information about the motion direction of both stimuli rather than just a same/different judgement. Additionally, we included an extra presentation time of 1s, that was not previously examined by Kanai and Verstraten (2005). Therefore, as well as being more demanding in terms of its response regime, our experiment also contained more conditions (25, compared to Kanai & Verstraten's 20). However, this methodological choice affected the resolution in measuring the reversal rate. For instance, a change in response to only one trial corresponds to a 5% variation in the reversal rate (compared to 2.5% resolution in Kanai & Verstraten).

Nonetheless, we were able to replicate the adaptation results observed by Kanai and Verstraten (2005) which suggests that the power of the study for this effect was not significantly affected. However, this coarser measure may have obscured other subtle but potentially meaningful differences in participants' responses across conditions. Future studies could address this by either increasing the number of trials per condition and splitting into separate sessions, thereby enhancing the sensitivity of the reversal rate measurements, or by refining the experimental design to reduce the number of conditions without sacrificing the scope of the research question.

Overall, the results presented here are in support of the hypothesis that participants' responses to paradigms investigating VMP using very short ISIs (i.e., below 0.12s) may not be valid. This is particularly true when these short ISIs are combined with a short presentation time of a preceding unambiguous stimulus. One of the initial aims of this chapter was to determine the optimum temporal parameters to evoke VMP and rMAE, for use in a subsequent experiment involving EEG. The results from Experiment 1, further bolstered by those of Experiment 2, suggest that this paradigm may not be able to reliably investigate VMP. However, an ISI of 0.12s and presentation times greater than 0.32s have been shown to reliably induce rMAE. Therefore, rMAE will be investigated using EEG but VMP will not.

Chapter Four: Testing for adaptation in spontaneous perceptual reversals using multivariate pattern analysis of EEG data

4.1 Introduction

A considerable amount of research has been carried out with the aim of identifying the factors that drive perceptual reversals of ambiguous stimuli, (e.g., the Necker cube, Boring's Old-Young woman, or Rubin's Faces-Vase). Candidate mechanisms for perceptual reversals include top-down (e.g., Pitts et al., 2008), bottom-up (e.g., Kornmeier & Bach, 2004), and integrative approaches (e.g., Kornmeier et al., 2009; Kornmeier & Bach, 2012; Long & Moran, 2007). These approaches have presented several electrophysiological correlates that are modulated by top-down and bottom-up influences but have primarily used univariate methods of analysing electrophysiological data that focus on a particular set of scalp locations and time points (for example, event-related potential components; ERP). These sorts of approaches can shed light on whether a given experimental condition evokes changes in or presence of these electrophysiological correlates, but it is more difficult to attribute these to a particular underlying mechanism. This is because the univariate methods that generate such correlates do not consider the holistic, whole-brain pattern of activity across different conditions. This is discussed in more detail below.

A problem associated with such univariate analysis techniques of electroencephalographic (EEG) data, for example ERP analyses, is that the selection of an appropriate and reliable time-window in which to analyse ERP components linked to spontaneous perceptual reversals is problematic. If an ambiguous stimulus is presented continuously, for example, participants may be asked to respond via a key press when they experience a perceptual reversal. Due to individual and trial-by-trial differences in reaction times this paradigm can cause temporal jitter, meaning that ERP components that occur only briefly may be lost in the trial and participant averaging process and do not appear in the grand average ERP results (see Chapter 2 for more detail; Kornmeier & Bach, 2004; 2012). As a result, an intermittent viewing paradigm is more widely employed in ERP studies of multi-stable perception, taking stimulus onset as the moment of reversal, and therefore significantly increasing the temporal resolution (the “Onset Paradigm”; see also Chapter 2; Kornmeier & Bach, 2004; 2012). Experiments employing the Onset Paradigm have identified several components that are thought to be related to both top-down and bottom-up influences on perceptual reversals (see Chapters 1 and 2 for a review of these). Two in particular, have been widely studied in the literature. The first has been termed the ‘Reversal Positivity’ (RP; Kornmeier et al., 2011; Kornmeier & Bach, 2004; 2012). This component occurs around 120ms post- stimulus onset at occipital electrode sites. The RP has been reported in response to the Necker cube (Kornmeier et al., 2011), Boring‘s old/young woman (Kornmeier & Bach, 2004) and binocular rivalry stimuli (Britz et al., 2011). The RP has been suggested to be reflective of the detection of a processing conflict resulting from the ambiguity of the presented stimulus (Kornmeier et al., 2011; Kornmeier & Bach, 2004, 2012). The second component is known as the ‘Reversal Negativity’ (RN; Kornmeier & Bach, 2004) which occurs around 200-260ms post-stimulus onset over posterior electrode sites. The RN has been found in response to reversals of the Necker cube (Kornmeier & Bach, 2004; 2012; 2014), face/vase and Schröder's staircase (Pitts, Nerger & Davis, 2007). The RN is thought to reflect top-down influences over perception as it has been shown to be enhanced during active volitional control of reversals (Pitts et al., 2008).

Evidence from behavioural studies too, have aimed to identify some of the top-down and bottom-up factors driving reversals in perception. As fully described in Chapter 1, the neural fatigue (Blake et al., 2003; Freeman, 2005; Long & Toppino, 2004; Wilson, 2003) and cross inhibition model is based on the idea that each interpretation of an ambiguous stimulus has its own neural representation which, when active, inhibits the representations of alternative interpretations of the same image. Following periods of prolonged stimulation, the neural representation underpinning the currently dominant percept becomes increasingly fatigued (adapted) and thus is no longer able to inhibit the alternative percept. At a critical point, perception reverses and the interpretation that is not adapted becomes dominant with its neural representation becoming strongest and inhibiting the alternative(s). This then restarts the process and will continue as a cycle of reversals. Studies examining the effects of pre-exposure to unambiguous variants of ambiguous stimuli have provided support for the neural fatigue/adaptation model. These studies show that prolonged pre-exposure to unambiguous variants of ambiguous stimuli (for example to an unambiguous left-facing Necker cube) reliably induce the opposite percept to become more likely upon subsequent ambiguous stimulus presentation (Kanai & Verstraten, 2005; Long et al., 1992; Toppino & Long, 1987; Verstraten et al., 1994). These behavioural experiments provide evidence that neural adaptation plays a role in perceptual reversals when it is induced by the pre-exposure of biased stimuli, however they do not shed light on whether it is a factor in the reversals in perception that occur spontaneously when viewing only a sequence of ambiguous figures as in the onset paradigm.

The currently available literature, therefore, provides evidence that adaptation to unambiguous stimuli can lead to behavioural adaptation effects. However, what is not currently determinable from the available research, is whether adaptation arises during the viewing of *only* ambiguous stimuli, either continuously or intermittently. Evidence from univariate EEG studies suggests that there could be several top-down and bottom-up mechanisms involved in the perception of purely ambiguous stimuli (e.g., Kornmeier & Bach, 2012). However, these univariate techniques described above focus on a particular region of interest (ROI) on the scalp where the component of interest (for example the RN or RP described above) is known to appear. This is done to avoid inflating Type I error rate, which is associated with multiple comparisons of activity at many scalp locations (e.g., Groppe et al., 2011, see Chapters 2 and 5 for more detail). In contrast to this univariate approach, multivariate pattern analysis (MVPA; fully described in Chapter 2) takes advantage of the holistic pattern of activity across the whole scalp. This technique has been used for some time to analyse data from functional magnetic resonance imaging (fMRI) studies and has more recently also been used to analyse EEG data (see below).

MVPA involves training a machine learning classifier (e.g., a Support Vector Machine or linear discriminant analysis; SVM and LDA respectively) to recognise patterns in scalp activity that distinguish between different cognitive or perceptual states (e.g., whether a person is looking at a face or a house). In this approach, a classifier is trained on a subset of EEG trials which are labelled to indicate which one of the, typically, two states were present on each trial. The classifier learns what patterns of EEG activity distinguish the two classes of trials and produces a model which takes the pattern of EEG data on a given trial and can make a prediction about the class to which it belongs. The model is then cross validated by testing how well it predicts the class for each trial on an independent subset of the data. Significantly above chance performance on this classification indicates that there is useful information in the pattern of EEG data that is predictive of trial class. This classifier can be used to decode the perceptual or cognitive state from the brain activity alone.

There are some important distinctions to be made between MVPA and traditional univariate analyses, therefore. Firstly, classification can be done separately at each time point (i.e., time-resolved) to take advantage of the high temporal resolution of EEG data. Secondly, classification is applied to each subject separately, and the time courses of decoding accuracies are what is averaged across subjects. This means that MVPA can pick up on the differences between the two classes at any electrode sites that are relevant for that subject. Univariate analyses, in contrast, require there to be very similar changes in scalp voltages at specific and often pre-determined electrode sites, for the effect to be detected. Additionally, the machine learning classifiers used in MVPA consider whole-brain patterns of EEG activity that most optimally separate the two classes of interest, making MVPA more sensitive to detect very subtle effects (e.g., Bae & Luck, 2018; 2019).

Indeed, research using MVPA has shown that it can be used to analyse EEG and magnetoencephalography (MEG) data in more nuanced ways than are possible with univariate analyses (e.g., Cauchoix et al., 2012; List et al., 2017). For example, Das et al. (2010) aimed to compare the efficacy of EEG pattern classification, using three distinct classifiers, against traditional, univariate ERP indices such as peak amplitude, mean amplitude, and peak latency. In their study, participants were engaged in a task where they were shown a face, or a car overlaid with Gaussian noise and had to rate their confidence in identifying the category of the stimulus (i.e., a face or a car). Earlier studies employing univariate methods of analysing EEG and MEG data have identified specific ERP components such as the N1 and N170 that are linked to face presentation as opposed to other objects (Gauthier et al., 2003; Taylor et al., 1999) or the early trial averaged M100 in MEG related to face categorisation (Liu et al., 2002). However, diverging from these findings, Das et al. (2010) discovered that the neural activity indicative of correctly perceiving and reporting faces versus cars was distributed over time, commencing at 120ms and persisting for over 400ms after stimulus onset. This broader temporal window for significant activity pattern classification differs from previous univariate observations, suggesting a more prolonged neural involvement in distinguishing between faces and cars. MVPA has also been shown to have revealed subtle effects in other research areas, for instance, decoding the focus of attention (local or global; e.g., List et al., 2017) from EEG data. It has also been used to successfully decode which of 16 orientations is being held in working memory, even when orientation is not currently relevant to the task (Bae & Luck 2018; 2019).

In terms of the perception of ambiguous stimuli, MVPA has been able to decode viewers‘ perception of ambiguous stimuli from EEG data, such as the Necker cube (e.g., Hramov et al., 2017). MEG data too, has been used to successfully decode participants’ perception of the Rubin’s faces/vase stimulus (Rassi et al., 2019) and perception of face vs house in a binocular rivalry paradigm (Rassi et al., 2022). As well as being used to analyse MEG data in the time domain (i.e., the information used to train the classifier is the scalp voltages at each point in time), MVPA has also been used to analyse MEG and EEG data in the time-frequency domain. In this case, the classifier receives information about the voltage at each electrode across time and about the power within each frequency band. This type of analysis has been used to reveal that, for instance, there is an increase in information flow from the fusiform face area (FFA) to the primary visual cortex before participants subsequently report seeing faces rather than a vase in the Rubin face/vase illusion (Rassi et al., 2019).

In addition to simply decoding perceptual experience, MVPA has the potential to offer deeper insight into the mechanisms mediating perceptual reversals and how these unfold across time. Specifically, the studies described above provide evidence that MVPA is a powerful tool to detect differences between given brain states. However, recent use of MVPA has also shown that it is also a tool with which to detect *similarities* among brain states. In particular, a technique known as cross-classification, or cross-decoding, allows one to make inferences about the similarity of neural processing involved in one or more experimental conditions (e.g., Kaplan et al, 2015). Cross-decoding assumes that if similar neural processes are involved in two different experimental conditions (e.g., perception and imagery) then the two conditions will have similar EEG scalp distributions. Thus, for instance, a classifier trained to discriminate different stimuli (e.g., faces vs. houses) in perception should also perform well at discriminating imagined stimuli. Alternatively, if the two conditions involve different mechanisms, then their scalp distributions will differ and a classifier trained on one condition (e.g., perception) will decode the other condition (e.g., imagery) with poor or at chance accuracy. To the extent that two conditions involve similar electrophysiological mechanisms, one should observe above chance cross-decoding with higher accuracy indicating greater overlap of mechanisms. This approach to MVPA has proven successful in the literature. For instance, in an fMRI study, Stokes et al. (2009) presented participants with the letter “X” or “O” in one condition. In the second condition they asked participants to only imagine the letter “X” or “O”. They found that a classifier trained on data from the lateral occipital complex during trials in the visual presentation condition could correctly decode participants’ imagined letter from data during the trials in the imagined condition. In a follow up study, Stokes et al. (2011) showed that this was also the case even when the letters were physically presented in different parts of the visual field. Cross-decoding has also revealed similar neural patterns involved in motor execution and action observation. For instance, Oosterhof et al. (2012a) recorded fMRI data whilst their participants performed and observed two object-directed actions ("lift" and "slap"). The actions were presented and executed in both first- and third-person perspectives. The experiment was designed to distinguish between actions observed and executed by participants, and whether there was any similarity in neural encoding based on the perspective of the observed action. The authors found that the researchers were able to identify distinct patterns of brain activity in the ventral premotor cortex (PMv) that corresponded to specific actions being either observed or executed. Crucially, the cross-classification approach allowed for the comparison of neural patterns across different modalities (visual observation and motor execution) and perspectives (first-person and third-person). The findings revealed that actions observed from a first-person perspective produced significant cross-modal neural patterns in the PMv, mirroring those during action execution. In contrast, this cross-modal congruence was not observed for actions perceived from a third-person perspective.

The goal of the experiment in this chapter is to use MVPA cross-decoding to more directly explore the role that adaptation (as described in the neural fatigue hypothesis) plays in spontaneously occurring perceptual reversals. Testing the mechanisms behind spontaneous perceptual reversals presents unique challenges due to the inherently unpredictable nature of these events. However, by using MVPA, it is possible to attempt to tease apart the neural patterns associated with these spontaneous reversals and compare them to those observed during induced perceptual shifts.

Testing adaptation in an EEG paradigm requires rapid adaptation so that sufficient trial numbers can be recorded, to maximise the signal to noise ratio (Luck, 2014). The work presented in Chapter 3 clearly indicates that relatively brief exposures of an unambiguous stimulus in the fast adaptation paradigm of Kanai & Verstraten (2005) can create adaptation which subsequently affects the perception of an ambiguous stimulus. Furthermore, the amount of adaptation can be reduced or removed parametrically by varying presentation duration of the adaptor stimulus. In the present experiment, based on the results of Experiment 1 in Chapter 3, we can identify optimum temporal parameters required to induce adaptation as well as parameters that do not induce adaptation.

In the present experiment, there are 2 block types. The first block type is the directional block, where each trial comprises two stimuli: an unambiguous adaptor stimulus with either left or rightward motion followed by an ambiguous stimulus (see Figure 4.1). To induce adaptation, the unambiguous adaptor stimulus will be set to a long duration (640ms). Based on the results of Experiment 1 (see

Chapter 3), we expect this stimulus to induce adaptation and bias perception of the second ambiguous stimulus towards the opposite direction in a majority of the trials (approximately 60-70%). We call this condition the "adapted" condition. In contrast, in the "unbiased" condition, the unambiguous adaptor stimulus will be a shorter duration (80ms). Based on the results of Experiment 1, we expect perceptual responses for the ambiguous stimulus to be unbiased (approximately 50% left and 50% right responses). Although the ambiguous stimuli in these two conditions are physically identical, their trial history induces different perceptual processes and outcomes. The second block type is the ambiguous block, which contains trials that only contain pairs of ambiguous motion stimuli.

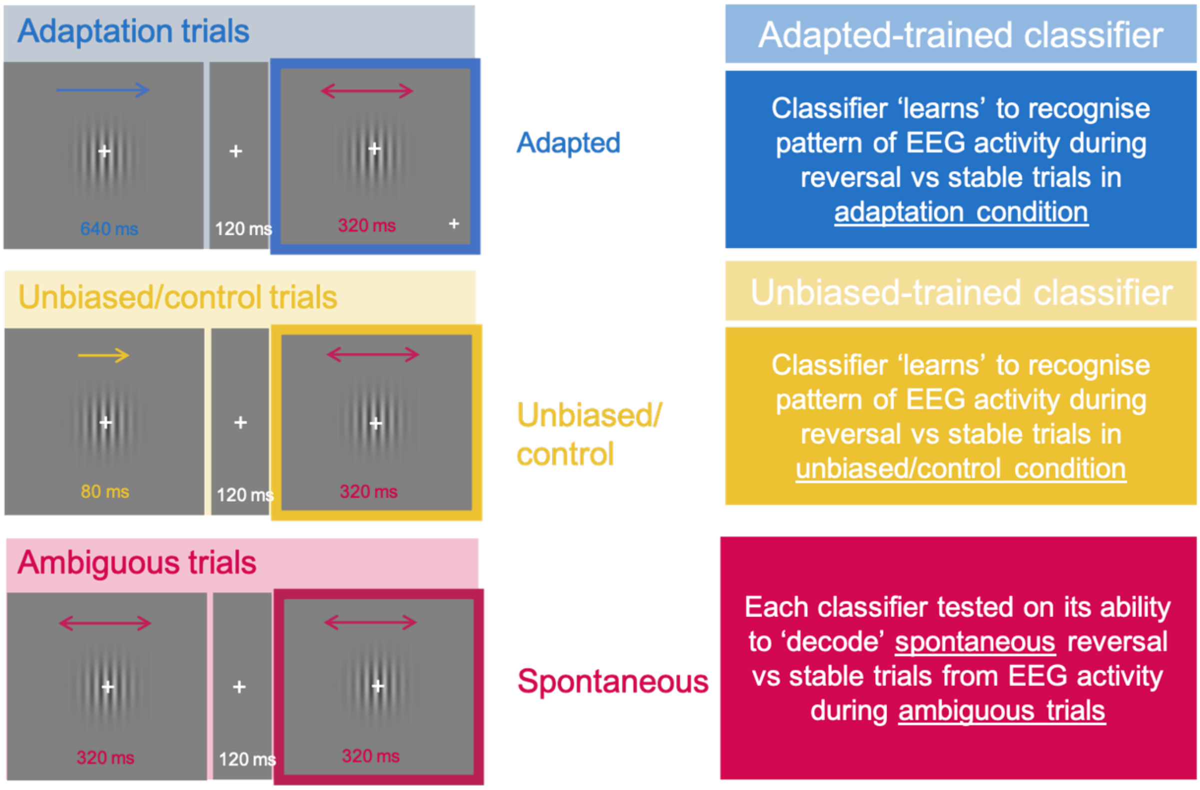

Figure 4.1 Representation of the 2 trial types that will be used to train 2 separate classifiers. The only difference between the 2 trial types is the length of time that the directional motion stimulus is presented for. The top panel shows an

example of a trial that will be used to induce adaptation, by presenting the directional motion (represented by the blue directional arrow) for 640ms. The middle panel shows an example of an unbiased trial, where directional motion is presented for 80ms (represented by the yellow directional arrow). The 2 classifiers are trained only on the data from the ambiguous stimulus (bottom panel; ambiguous motion represented by the pink arrow), which is always presented for 320ms, in each trial type.

Two machine learning classifiers will be trained on the data generated from the two trial types in the directional blocks (i.e., the adapted and unbiased trial types). The first classifier will be trained to predict whether the ambiguous stimulus perception was a reversal or non-reversal (compared to the preceding unambiguous stimulus) based on the EEG voltage from the long adaptation trials (i.e., adapted trials). This classifier will learn the pattern of scalp voltage that is predictive of an adaptation-driven reversal independently of the direction of motion of the adapting stimulus or the direction of the reversal. The second classifier will be trained to predict perceptual reversals (vs. non-reversals) of the ambiguous stimulus based on EEG data during the unbiased trials. This classifier will learn the pattern of EEG scalp voltage that is associated with reversals that are only weakly driven by adaptation, if at all. Instead, other unknown mechanisms are presumably driving reversals on these trials. Both classifiers will be trained on brain activity only during the ambiguous stimulus (i.e., second stimulus) in each trial. Thus, although their trial histories differ, the visual stimulation during the training periods will be physically identical.

Both classifiers will then be tested on their ability to decode perceptual outcomes (i.e., reversal vs. stable) from EEG scalp patterns in the trials from the ambiguous block, which comprise 2 sequential ambiguous stimuli. Reversals in this condition are spontaneous reversals in perception, in other words they are not induced by prior exposure to an unambiguous stimulus. Nonetheless, they could be driven by a build-up of adaptation from the perceptual interpretation of the preceding ambiguous stimulus (i.e., the first of the two stimuli in each trial). If spontaneous perceptual reversals are indeed driven by adaptation, then the decoding accuracy of the adapted-trained classifier should be significantly greater than chance because that classifier was trained on a brain state that involved adaptation-driven reversals. It should also be significantly higher than the decoding accuracy of the unbiased-trained classifier which was trained to decode reversals driven by factors other than adaptation. This cross-decoding approach will allow more direct conclusions to be drawn on the role that adaptation plays in spontaneously occurring perceptual reversals. We hypothesise that spontaneous reversals, like induced reversals, are underpinned by adaptation processes. By comparing the decoding accuracy of our two classifiers — one trained on adaptation-driven reversals, the other on unbiased trials — we hope to shed light on the mechanisms driving spontaneous perceptual reversals.

The second aim of this chapter is to again use MVPA cross-decoding techniques, but this time to train and test the classifiers in a slightly different way (see Figure 4.2, below). The aim here is to assess whether the pattern of EEG scalp voltages in spontaneously occurring reversals (i.e., those in the ambiguous block) is more like that occurring during reversals due to adaptation (i.e., reversals in the adapted condition), or reversals occurring due to some other process (i.e., those in the unbiased condition). Here, one classifier will be trained to distinguish *adapted* reversal trials from *unbiased* reversal trials. In other words, the class labels provided to the classifier will not be “reversal trial” and “stable trial” as in the first analysis described above, as only reversal trials were being used. Instead, the class labels will be “adapted reversal” and “unbiased reversal”. The trained classifier will then be provided with data from the spontaneously occurring reversals from the ambiguous only block. Now, rather than recording the accuracy of the classifier’s predicted class labels about which *trial type* the data came from (i.e., reversal or stable) the predicted class labels themselves will be recorded (i.e., adapted reversal or unbiased reversal). In this analysis, decoding accuracy will not be recorded at all, as it is impossible for the classifier to correctly predict the true class of the testing data as it was never provided with adapted or unbiased reversal data – only spontaneous reversal data. Therefore, it can never give the true class label. This analysis will allow us to draw further conclusions around whether the pattern of brain activity during spontaneously occurring perceptual reversals is more like that occurring during adaptation-induced reversals than those occurring due to some other processes. As we predict that adaptation processes are driving spontaneously occurring reversals, we hypothesise that this classifier will predict that the data from the spontaneously occurring reversals are more likely to have come from the adaptation-induced reversal class than the unbiased reversal class.

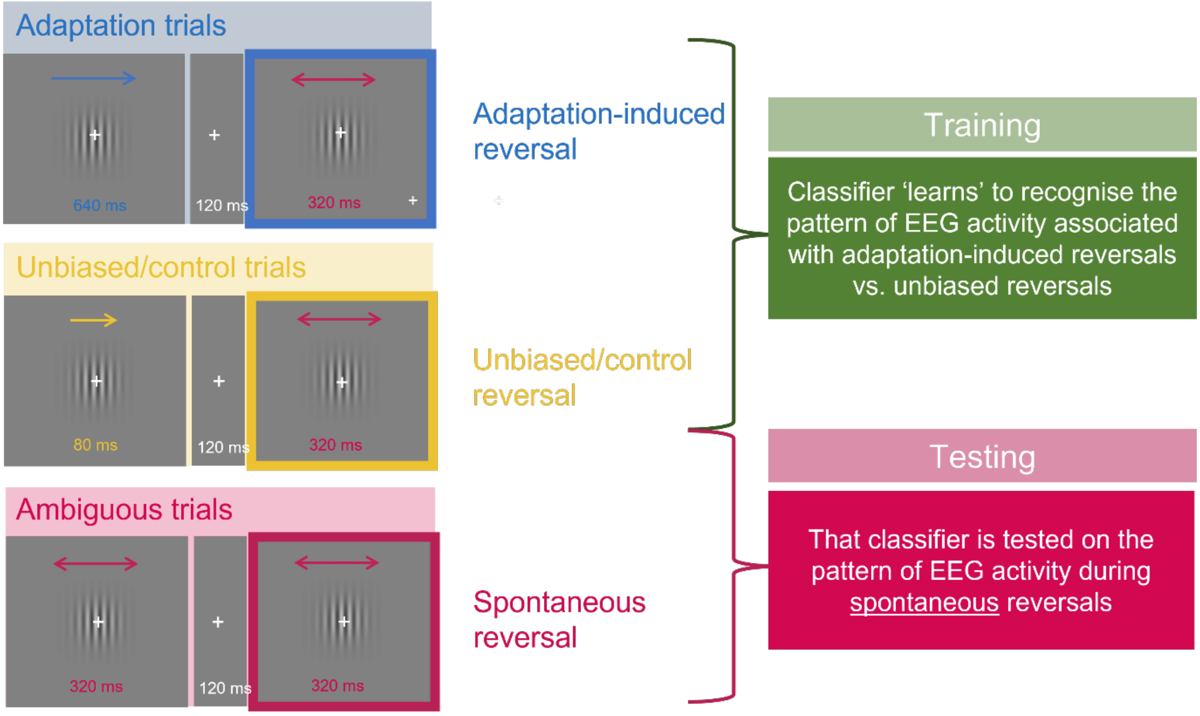

Figure 4.2 The training and testing method used in the second MVPA analysis. The training phase involved the SVM classifier being trained to distinguish adapted vs unbiased reversals. The testing phase involved that trained classifier being provided with data from spontaneously occurring reversal trials, and the class labels it predicted from that data (i.e., adapted reversal or unbiased reversal) were recorded.

The MVPA of EEG data also allows for more insight into the stages of processing of a perceptual reversal via a technique known as temporal generalisation (King & Dehaene, 2014; see Chapter 2 for a more detailed overview of this technique). Temporal generalisation is a classification technique that allows conclusions to be drawn about whether the pattern of brain activity is similar at different points in a trial. For instance, whether the same pattern of activity that occurs early in a trial is also generated towards the end of a trial. This is achieved in a manner similar to the cross-decoding technique described above. A classifier is trained to distinguish between different experimental conditions (here, reversal vs stable percepts) from data at a given time point. This trained classifier is then tested, not only on the data from the time point it was trained on but all other time points within the trial. This reveals whether the whole-brain activity patterns at a given time point are also predictive of activity occurring at other time points. Therefore, this technique will allow conclusions to be drawn regarding whether processes are sustained, isolated, reactivated and/or ramping up over the time course of the trial (see Figure 2.5 in Chapter 2 for more detail on this). Additionally, standard MVPA decoding techniques compare each time point in a trial independently, whereas temporal generalisation is also more sensitive to effects that may be temporally misaligned between trials (King & Dehaene, 2014). The ability of temporal generalisation to deal with this potential misalignment deals with the issue that the same neural processes may unfold at different speeds or with slightly different temporal profiles in the two conditions used for decoding. If only standard decoding techniques are used, any similarity in processes are missed if they are not temporally aligned in different conditions. Employing this technique here therefore will not only provide more information than is currently available on the processing stages involved in perceptual reversals but also account for effects that could be missed by univariate analyses and standard MVPA classification across time. Therefore, temporal generalisation will be performed for each of the MVPA analyses described above.

In addition to the multivariate techniques, this chapter will also aim to establish whether the ERP components (i.e., the RP and RN) thought to be markers of spontaneous perceptual reversals are present during the spontaneous reversals of the ambiguous motion stimulus described in Chapter 3 and used here. Kornmeier and Bach (2012) highlight that extensive research has been conducted that shows the presence of these components in response to stationary stimuli, but that further research is needed to establish whether these are also elicited by other ambiguous stimuli, such as the ambiguous motion used in this chapter. To our knowledge, only continuous viewing of stroboscopic alternative motion (SAM) stimuli has been used to investigate the potential ERP components associated with reversals of ambiguous motion (e.g., Basar-Eroglu et al., 1993; Strüber & Hermann, 2002). These studies used continuous viewing of SAM stimuli and revealed only the later parietal positivity component rather than the RP and RN components (see Chapter 2 for a full overview of the methods used). However, because the participants' manual response was used as the time reference when averaging the EEG data, the RN and RP, if present at all, were likely lost among the backwards averaging process (see above and Chapter 2). Therefore, as stimulus onset will be used as the time reference in this experiment, if the RN and RP are present during the perception of the ambiguous motion stimuli used in this experiment they should be induced by this design.

4.2 Methods

4.2.1 Participants

Twenty-five (21 female, 4 male; 2 left, 23 right-handed) undergraduate psychology students were recruited using Keele University's Research Participation Time scheme in return for partial course credit. This sample size was chosen as it is in line or greater than the sample sizes typically used in EEG literature investigating perceptual reversals, and that using MVPA of EEG data (e.g., Kornmeier, Hein & Bach, 2009; Das et al., 2010; List et al., 2017; Bae & Luck, 2018). Thus, 25 participants were recruited. Participants had a mean age of 21.48 years (range: 18-60 years; SD: 8.16 years). All participants had normal or corrected-to-normal vision (mean logMAR: -0.04; SD: 0.07; Precision Vision Logarithmic ETDRS 2000 chart). The study was approved by Keele University's Psychology Faculty Research Ethics Committee.

4.2.2 Stimuli and apparatus

4.2.2.1 Stimuli

The same two types of stimuli were used in this experiment as in Experiment 1 (see Chapter 3 and Figure 4.3 below): directional motion (Figure 4.3, top) and ambiguous motion (Figure 4.3, bottom). The stimuli were sine wave gratings with a spatial frequency of 1 CPD. Gratings had a Michelson contrast of 0.5 (maximum luminance of white regions was 297.49cd/m^2, for black regions 100.28 cd/m^2) and were presented centrally on a mid-grey background (RGB: 128:128:128). All stimuli were circular, with a diameter of 10°. The circular shape was created using a 2-D Gaussian mask. As in Experiment 1, these stimuli were presented in pairs to form one trial (see below for a full description). A white, 0.8° x 0.8° fixation cross was always presented in the centre of the screen throughout trials to serve as a fixation target. The experiment and stimuli were designed and controlled using PsychoPy3 v3.2.2 (Peirce et al., 2019).

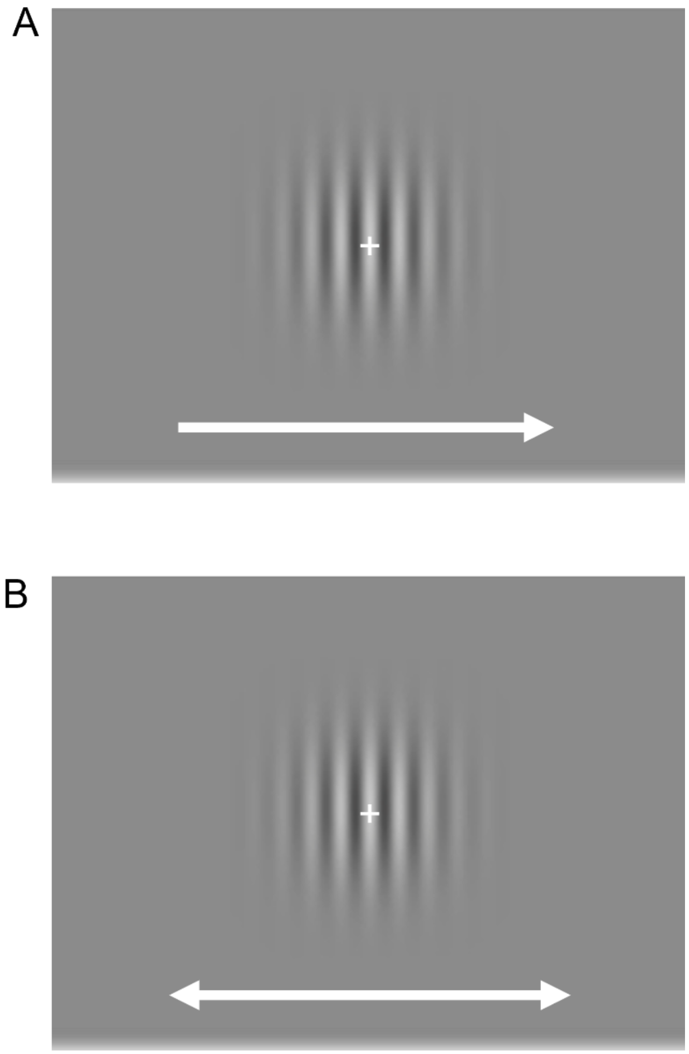

Figure 4.3 Examples of the directional motion (panel A, where the sine wave gratings move towards the right) and the ambiguous motion stimuli (panel B, where the sine wave grating can be viewed as moving in either direction). Although no actual difference can be seen in the stimuli themselves in this figure, the direction of the motion is indicated by the arrows.

4.2.2.2 Apparatus

The experiment was presented on a 24.5 inch BenQ Zowie XL2540 computer monitor with 120Hz refresh rate and 1920 x 1080 pixels resolution. Participants used a standard PC keyboard to respond to trials. Participants used a chin rest to maintain a constant viewing distance of 57cm and to reduce head movements that could result in artefactual EEG data.

A 24-bit DC-coupled Bio Semi Active Two Measurement System (Biosemi; Amsterdam, Netherlands) was used to record EEG scalp voltages from 64 Ag-AgCl active scalp electrodes, with a sampling rate of 512Hz (see Figure 4.4, below for a map of electrode locations). Low-pass filtering was performed within the hardware's Analogue to Digital Converter's decimation filter, which has a fifth order cascaded integrator-comb filter response with a -3dB point at 1/5th of the sample rate. Electrodes were fixed in place using a cloth cap and positioned according to the 10-10 system (Klem et al. 1999; Nuwer 1998). Three additional electrodes were placed on the face, using adhesive pads, to record horizontal (HEOG; via two electrodes placed next to each lateral canthus) and vertical (VEOG; via one electrode placed 2cm below the centre of the left eye) electro-oculogram voltages. These facial positions were prepared using an isopropyl alcohol wipe to facilitate the adherence of the pads. For all electrodes, Parker SignaGel was used as the electrolyte gel (Parker Labs; https://www.parkerlabs.com/). Electrode offset was maintained within a -10 to 10 μV range of the common mode voltage. All data were recorded relative to the feedback loop formed between the Common Mode Sense (CMS) active electrode and the Driven Right Leg (DRL) passive electrode. All re-referencing of the data was performed offline (see Section 4.2).

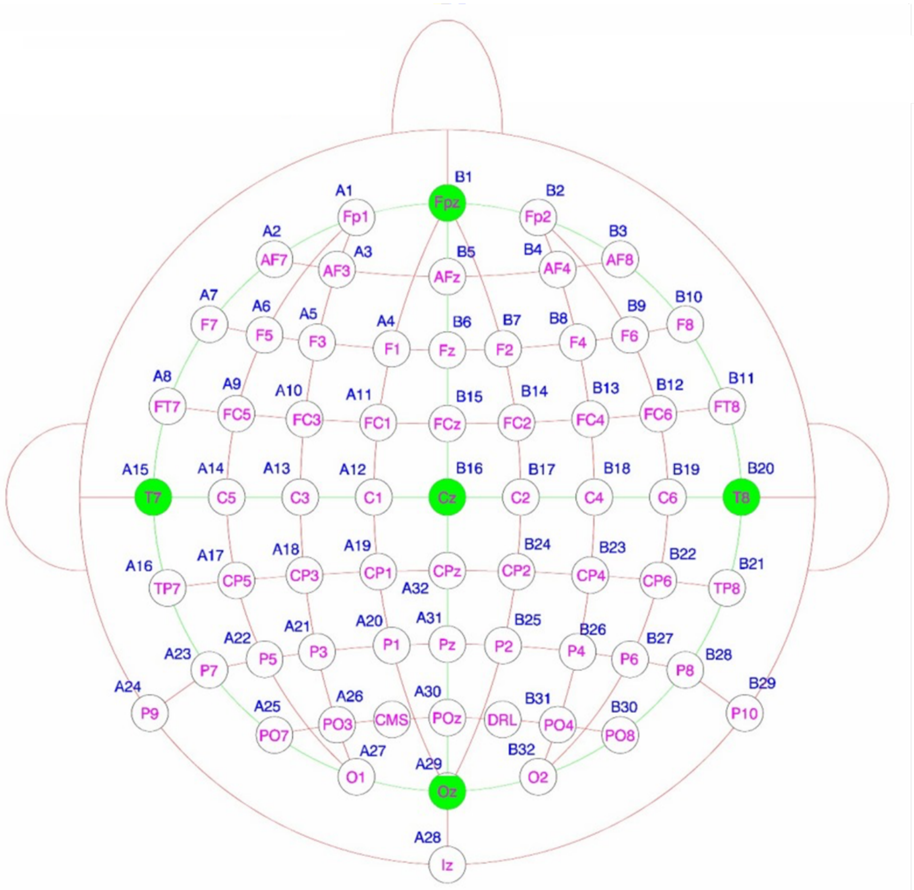

Figure 4.4 Electrode layout map showing the positions of the 64 electrodes used in this experiment (taken from biosemi.com; Biosemi, Amsterdam, Netherlands).

4.2.3 Design

Trials were arranged in two distinct block types. Directional blocks (i.e., trials that contained directional motion stimuli), containing only adapted and unbiased trials but no ambiguous trials, and ambiguous blocks which contained ambiguous trials only (see below and Figure 4.5). The two block types were chosen, over randomly mixing the three trial types, so that perception on the ambiguous trials would not be biased by any directional motion. Each directional block contained 60 trials, and each ambiguous block contained 52 trials. The experiment consisted of 6 directional blocks and 4 ambiguous blocks. Directional blocks contained 30 adapted trials, and 30 unbiased trials. This resulted in a total of 180 adapted trials, 180 unbiased trials, and 208 ambiguous trials for each participant.

Each experimental trial consisted of the presentation of 2 stimuli in temporal succession, in much the same way as Experiment 1. A total of 3 trial types were used, with temporal parameters derived from the results of Experiment 1 (see Figure 4.5). One trial type was adaptation trials, in which directional motion was presented for 640ms, followed by a 120ms ISI, and ambiguous motion presented for 320ms; these temporal parameters evoked the desired 60-70% adaptation responses in Experiment 1. A second trial type was unbiased trials, in which directional motion was presented for 80ms; these temporal parameters evoked the necessary unbiased responses in Experiment 1 (i.e., participants responded to indicate reversals on around 50% of trials). All other aspects of these unbiased trials were identical to the adaptation trials. The third trial type was ambiguous trials, in which only ambiguous motion stimulus pairs were presented for 320ms each, with an ISI of 120ms between them in the same way as with the adaptation and unbiased trials.

As in Experiment 1, participants were required to respond to trials by reporting their perceived direction of motion (leftward vs. rightward) for each stimulus in the pair. To allow response buttons to be fairly counterbalanced and reduce spatial compatibility effects on response (e.g., Ulmita & Nicoletti, 1990; Alluisi & Warm, 1990), participants responded to trials using the ‘up’ and ‘down’ arrow keys in the same way as in Experiment 1. For clarity this will be briefly described again here. The ‘up’ and ‘down’ arrow keys represented the directions ‘left’ and ‘right’ and were swapped around between blocks, on 50% of trials (i.e., trials within each block always required the same response). In other words, on 50% of trials ‘up’ would indicate perceived motion to the left and on 50% of trials ‘up’ would indicate perceived motion to the right. This response method was chosen in preference to using the ‘left’ and ‘right’ arrow keys as counterbalancing these (i.e., swapping the response keys on 50% of trials) would lead to 50% of trials being more difficult as participants would be required to respond using the opposite arrow key to their perceived direction of motion.

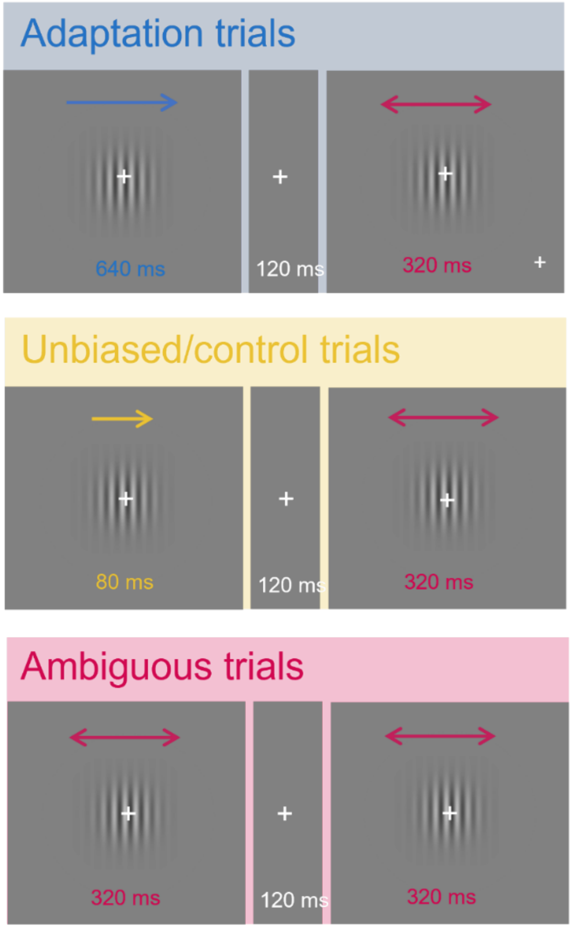

Figure 4.5: The three trial types used in this experiment. A: Adaptation trials, containing a long presentation of directional motion followed by ambiguous motion. B: Unbiased trials, containing directional motion presented for a shorter time, designed to evoke unbiased responses to the ambiguous motion that followed. C: Ambiguous trials, in which only ambiguous motion stimulus pairs were presented.

Additionally, attention check versions of the adaptation and unbiased trials were included in a similar way to Experiment 1. Attention check trials were only presented during directional blocks, so that the unambiguous motion did not affect the perception of the ambiguous only motion in the ambiguous blocks. Directional blocks each contained 6 attention check trials (3 for each of the adapted and unbiased presentation time parameters). Therefore there were 36 attention check trials in total for each participant (6 trials x 6 directional blocks). These trials consisted of one directional motion stimulus presented once for either 640ms or 80ms, followed by a 120ms ISI, followed by a second directional motion stimulus presented for 320ms (i.e., directly replacing the ambiguous motion that was presented on the real trials).

4.2.4 Procedure

As with Experiment 1, participants completed the informed consent process upon entry to the EEG laboratory. In this experiment, this also included an assessment of whether participants were able to take part in an EEG experiment by checking that they did not have thick or braided hair, have any metal implants in their upper body, or have any neurological conditions that would affect their ability to take part. Those meeting these criteria performed a visual acuity assessment, and those with normal or corrected-to-normal vision were taken forward to the experiment.

Eligible participants were then fitted with the EEG recording equipment. The three facial electrodes were attached first, followed by the cloth cap. The cap was held in place via two straps attached to a chest strap. To ensure that the cap was located centrally, the experimenter measured the nasion-inion distance and positioned the cap so that the Cz electrode was located at the halfway point. The intra-tragal distance was then measured, and again the cap adjusted to ensure that the Cz electrode was at the halfway point. When the cap was correctly placed, the experimenter fitted the scalp electrodes by systematically filling each electrode port with SignaGel electrolyte and inserting the electrode. Once fitted, the electrodes were connected to the BioSemi ActiveTwo recording system.

Connectivity checks (checking for noise in the data from each channel, checking for large electrode offset values) were then conducted, to establish that CMS/DRL feedback loop was correctly connected as well as the electrode offsets. Any electrodes showing an offset outside of the -1 to 10 μV range was addressed by taking measures to improve connectivity with the scalp, for example by adding more electrolyte gel or moving hair that was between the electrode and scalp. Participants were then positioned correctly so that their chin sat comfortably on the chin rest, by adjusting the height of the table, chair or both. The eye-tracker was then calibrated and set up to record fixation throughout experimental blocks. A final check of cap position and electrode offset was then conducted.

The experimenter then explained the task, including how to respond and checking whether participants were able to perceive both versions of the ambiguous motion. Participants then completed a set of practice trials as with Experiment 1, to check that they understood the task and how to respond to trials correctly. The practice block consisted of 24 trials: 8 adaptive trials, 8 unbiased trials and 8 attention check trials. The experimenter remained in the room with participants during the practice block to address any issues or answer questions. Participants were provided with feedback on their performance on attention check trials at the end of the practice block.

The experimenter then left the testing room whilst the participant completed the main experiment. Throughout this time, the experimenter remained in the control room connected to the testing room. Constant monitoring of the participant was achieved through a video feed and a two-way audio system. Five self-timed breaks occurred between blocks. Upon completion, the experimenter removed EEG apparatus from participants and fully debriefed them on the aims and hypotheses of the experiment.

4.2.5 Data analysis methods

4.2.5.1 Behavioural analysis

Trials were coded as either “reversal” or “stable” according to whether participants responded by pressing two different buttons (e.g., ‘up’, ‘down’) or the same button twice (e.g., ‘up’, ‘up’) respectively. Participants’ responses to attention check trials were analysed. An a priori behavioural exclusion criterion was set such that participants scoring less than 70% correct on the attention check trials in the main experiment were excluded. Based on this criterion, no participants were excluded from the experiment.

The average number of stable and reversal trials in ambiguous blocks were calculated for each participant. To examine the stochasticity of those responses (i.e., independence between responses on subsequent trials), the distribution of response run lengths (i.e., number of same responses in a row for either left or right responses) was analysed in the same way as in Experiment 1, Chapter 3 and will be described again here for clarity. This analysis of stochasticity was done to determine whether participants’ responses in the ambiguous block were independent of one another. The ambiguous block was designed to evoke perceptual reversals that occur spontaneously, in a manner as close to that of continuous viewing as possible. Therefore, responses to the stimuli should be stochastic, as this would indicate that there was no bias toward one or the other direction of motion when the stimuli were perceived (Hesselmann et al., 2008). A repetition in response meant that a participant responded to indicate, for example, leftward motion on two or more successive trials. This ‘run’ of repetitions would be broken when the participant responded to indicate motion to the left then motion to the right, or vice versa. So, a response pattern of “left”, “left”, “left”, “right” would mean that there had been two repetitions before that run was broken. To assess this quantitatively, the distribution of response repetitions was fitted to a geometric distribution. A geometric distribution represents the frequency of different run lengths expected by chance in a sequence given only two random outcomes (i.e., left vs. right here) and a certain probability of each outcome (i.e., 50% chance here). For example, you would use this to estimate the number of times that you could expect to observe fair coin to land heads 4 times in a row in a sequence of 100 tosses. Therefore, if the response repetition data fits well to this distribution, this suggests that participants are responding in a random and unbiased manner in which responses to one stimulus are independent of those to preceding stimuli.

Data from the directional blocks were then separated into the two conditions (adapted and unbiased) and the average number of stable and reversal trials per participant was calculated for each condition.

4.2.5.2 EEG analysis

Offline pre-processing of the EEG data was carried out using EEGLAB (Delorme & Makeig, 2004) implemented in Matlab r2020a. The data were high-pass filtered at 0.01Hz and referenced to the Cz electrode (in line with Abdallah & Brooks, 2020). Data were then segmented into 620ms epochs (-120ms to 500ms, relative to the onset of the second stimulus in each trial). Artefactual eye-blink components were removed in two stages. Firstly, independent component analysis

(ICA) was run on each participants data. The ICA component indicative of blinks was removed, and then any remaining trials containing blinks were removed via manual inspection. Each trial was also manually examined for artefactually noisy data i.e., with amplitude variations of ±100µV. Any identified trials were removed from the dataset; 1.98% of trials were removed (range: 0 – 3.75%) on average, per participant.

The epochs were coded as ‘reversal’ or ‘stable’ according to the system outlined in the ‘Behavioural analysis’ section above and then sorted into six conditions. These conditions were derived from the combination of trial type (i.e., adapted, unbiased or ambiguous) and response type (i.e., reversal or stable).

4.2.5.2.1 ERP Analysis

To generate ERPs, additional pre-processing steps to those outlined above were conducted on the data. Epochs were baseline corrected using the average amplitude from the period -100ms to 0ms relative to stimulus onset, then filtered with a low-pass filter at 30Hz. Epochs were then averaged within each of the six conditions to create ERPs. The ERPs were then averaged across participants, forming grand-average waveforms for each condition.

Analyses of ERP components were carried out in line with previous experiments (e.g., Abdallah & Brooks, 2020; Kornmeier & Bach, 2012). As detailed in the Introduction above, two ERP components were of interest: the RP and RN. The RP was defined a priori with a temporal region of interest (ROI) of 100ms – 200ms over channels O1, O2 and Oz. The RN was defined a priori with a temporal ROI of 200ms – 400ms over channels O1, O2, Oz, PO7 and PO8. For each participant and condition, the average amplitude within each temporal ROI above was calculated for each of the channels. These average amplitudes were then submitted to two separate repeated measures ANOVAs (one 3x3x2 ANOVA for the RP, one 5x3x2 for the RN) with channel, trial type (adapted, unbiased or ambiguous) and response type (reversal or stable) as factors.

4.2.5.2.2 Multivariate EEG analyses

Multivariate pattern analyses (MVPA) were carried out on the single-trial data using the MVPA-Light toolbox (Treder, 2020) implemented in Matlab r2020a. Several of these analyses were conducted, and the same pre-processing steps were taken in all of them. Firstly, under-sampling was conducted on the data. Under-sampling was used to balance the number of trials in each class, by randomly subsampling from the majority class. For instance, if there were 55% of reversal trials and 45% stable trials for a given participant, then the majority class (here, reversals) would have a random subsample taken from it on every training run. Next, dimensionality reduction was performed on the data via principal component analysis (PCA). PCA is a statistical method used for dimensionality reduction. By transforming correlated variables into a smaller set of linearly uncorrelated variables called principal components, PCA helps to retain the most significant variance in the data while removing redundancy and noise. Here, the MVPA Light Toolbox use of PCA takes the 20 components that explain the most variance in the data forward for analysis.

To determine whether the pattern of EEG scalp voltages differed during spontaneous reversal and stable trials, the first analysis was conducted on the data from ambiguous only blocks. A linear support vector machine (SVM) classifier was trained to classify trials as reversal or stable. To avoid overfitting that can arise from training and testing on the same data, a five-fold cross-validation procedure was used. This meant that the trials were firstly split into five equal groups, and trials from four of the groups were used to train the classifier (the 'training' data), with those from the fifth group being used to test the trained classifier's performance (the 'testing' data). In the training phase the classifier was provided with the pattern of EEG scalp voltages from 63 electrodes along with the true class label of the trial that pattern came from (i.e., stable or reversal). This was repeated separately for each timepoint in the trial, and each trial in the training data. The trained classifier was then provided only with the pattern of scalp voltages at each time point in the testing data (i.e., no true class label), and measured on its ability to correctly predict the true class label of the trial from the pattern of scalp voltages alone. This resulted in a decoding accuracy value for each timepoint, given as the percentage of the total number of trials in the testing data subset that were correctly classified. If enough class-related information exists in the data, we would expect the decoding accuracy to be significantly greater than would be achieved by chance. For data with two classes, this chance value would be 50%. The training and testing process was repeated a further four times with different training and testing data subsets, and the results averaged at each 1ms time point to produce a time course of decoding accuracy across the whole trial for each participant. These individual participant decoding accuracies were then averaged to produce a grand average decoding accuracy time series.

The second analysis, cross-decoding, involved training two separate SVM classifiers: one to decode reversal vs stable trials from the adapted condition and one to decode reversal vs stable trials from the unbiased condition. Then, each of the two classifiers in turn were separately provided only with the EEG scalp voltage data from the ambiguous condition (i.e., the spontaneously occurring reversal or stable trials) and were each tested on their ability to correctly predict the true class label of these trials. As in the first analysis, this resulted in a decoding accuracy value for each 1ms timepoint in each trial and for each participant. These values were averaged for each trial, then averaged across participants to produce a grand average decoding accuracy time-series. No cross-validation procedure was necessary due to the training and testing data being independent.

The third analysis involved the SVM classifier being trained to distinguish *adapted* reversal trials from *unbiased* reversal trials. The trained classifier was then provided with data from the spontaneously occurring reversals from the ambiguous only block. Now, rather than recording the accuracy of the classifier's predicted class labels about which *trial type* the data came from (i.e., reversal or stable) the predicted class labels themselves were recorded (i.e., adapted reversal or unbiased reversal). In this analysis, decoding accuracy was not recorded at all, as it was impossible for the classifier to correctly predict the true class of the testing data as it was never provided with adapted or unbiased reversal data – only spontaneous reversal data. As the classifier in this case was already being trained and tested on two independent datasets, there was again no need for the cross-validation procedure described in the first analysis above. Instead, the percentages of each type of reversal (adapted and unbiased) predicted by the classifier were averaged for each trial, and then a grand average time series again produced by averaging across participants.

Temporal generalisation plots were then produced for each of the analyses described above. Temporal generalisation is a classification technique that can provide insight into whether information is shared across different time points in a trial (King & Dehaene, 2014). For instance, whether the information that the classifier uses early in a trial is the same that it uses later in the trial (for a full description see Chapter 2). This was achieved by training the SVM classifier at a given 1ms timepoint and then testing the trained classifier at that same timepoint but also at *every other* timepoint in the trial. This was performed for each of the three analyses described above and averaged in the same way, to produce 3-dimensional (317x317 timepoints) temporal generalisation plots of classifier performance at each training timepoint x testing timepoint.

4.2.5.2.2 Significance testing of multivariate EEG analyses

To evaluate the performance of the SVM classifier accuracy whilst controlling Type I error rate, cluster-based permutation tests were conducted in Matlab r2020a (for full details see Chapters 2 and 5).

4.3 Results

4.3.1 Behavioural results

4.3.1.1 Attention check trials

The a priori attention check criterion was that participants' data would be excluded from analysis if they did not respond to at least 70% of attention check trials correctly. Based on this criterion, there were no participants that needed to be excluded from the analysis. Participants correctly responded to an average 82.11% (SD = 5.01) of attention check trials.

4.3.1.2 Response time

The overall average response time (RT; to first response button press) for all participants across all trial types (adapted, unbiased, and ambiguous) and response types (reversal and stable) was 0.482s (SD: 0.108s), relative to the onset of the second stimulus in each trial.

A 3 x 2 repeated measures ANOVA (3 trial types x 2 response types; see Figure 4.6) revealed a significant main effect of trial type on RT, $F(1,24) = 22.137$, $p < .001$, $\eta_p^2 = .480$. However, there was no significant effect of response type, $p = .354$ and no significant interaction effect, $p = .197$.

To examine the main effect of trial type, post hoc t-tests with a Bonferroni correction were applied. These revealed that unbiased trials, $M = 0.576$s ($SD = 0.093$), were responded to significantly more slowly than both the ambiguous, $M = 0.451$s ($SD = 0.123$), $t(24) = 5.926$, $d = 1.192$, $p < .001$ and adapted, $M = 0.450$s ($SD = 0.098$), $t(24) = 5.978$, $d = 1.203$, $p < .001$, conditions.

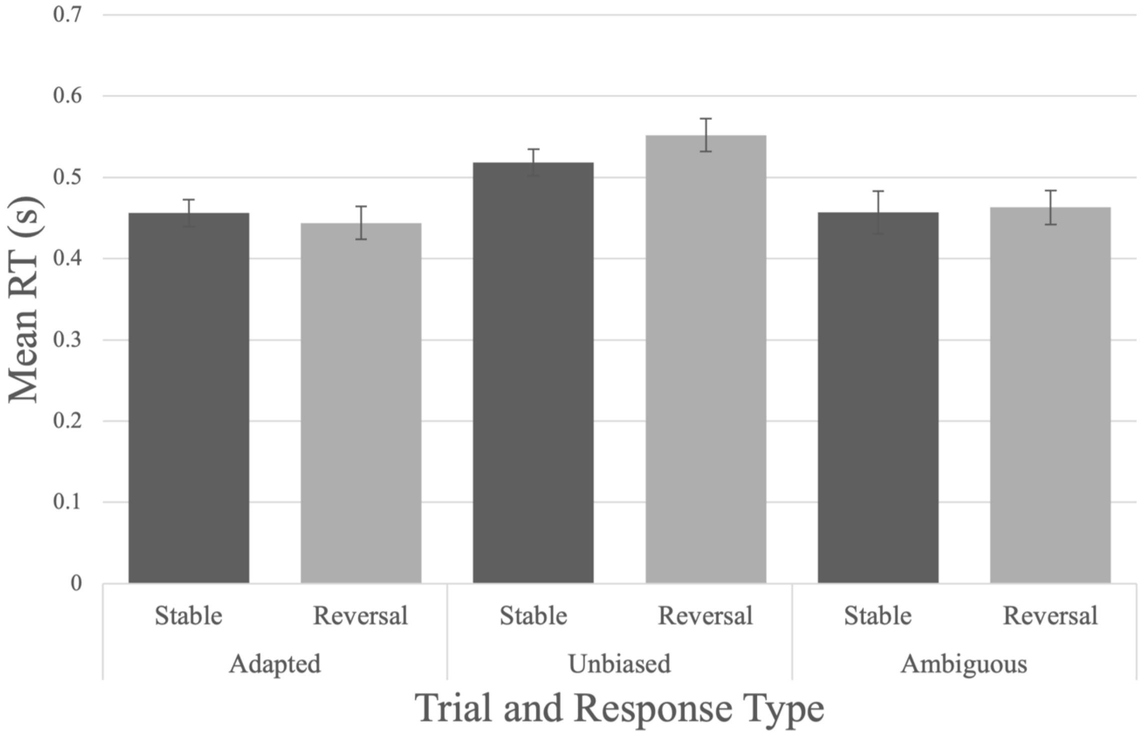

Figure 4.6 Bar chart showing the mean response time for reversal vs stable responses across each of the 3 trial types (adapted, unbiased and ambiguous). Error bars show standard error.

4.3.1.3 Ambiguous block trials

Participants responded to indicate a reversal on ambiguous block trials on 51.41% (SD = 21.56) of trials and indicated stability on 48.58% (SD = 21.56) of trials. Participants responded to indicate that they had perceived leftward motion on 52.21% (SD = 8.19) of trials and rightward motion on 47.79% (SD = 5.29). Null responses (i.e., trials where the participants had failed to respond within the 1.5s window) were removed from the analysis. On average, 5.87% (SD = 5.54) of all ambiguous trials were recorded as null and removed from the analysis.

Neither reversal nor stable responses were significantly different from chance, with both resulting in the same values for $t(24) = 0.327$, and $p = .746$. A

Bayesian one-sample t-test was also conducted to assess the extent to which the data was supportive of the null hypothesis (H_0), i.e., that the percentage of reversals did not differ from chance. The resulting Bayes factor indicated evidence for H_0, with $BF_{01} = 4.516$, which means that the data were approximately 4.5 times more likely to occur under H_0 than H_1. As BF_{01} and BF_{10} are inversely related, the value for BF_{10} (i.e., the likelihood of the outcome occurring under H_1 compared to H0) was therefore 0.221. This indicates that there is moderate evidence for the null hypothesis, according to the ranges indicated by Lee and Wagenmakers (2013). Neither left nor right responses were significantly different from chance, with both resulting in the same values for $t(24) = 0.120$, $p = .905$.

To examine the stochasticity of those responses (i.e., independence between responses on subsequent trials), the distribution of response run lengths (i.e., number of same responses in a row for either left or right responses) was analysed in the same way as in Experiment 1, Chapter 3. This analysis of stochasticity was done to determine whether participants' responses in the ambiguous block were independent of one another, indicating that there was no bias toward one or the other direction of motion when the stimuli were perceived (Hesselmann et al., 2008). The number of response repetitions (for left and right responses) were found to be well fitted to a geometric distribution (see Figure 4.7), $R^2(24) = .81$, $p < .001$.

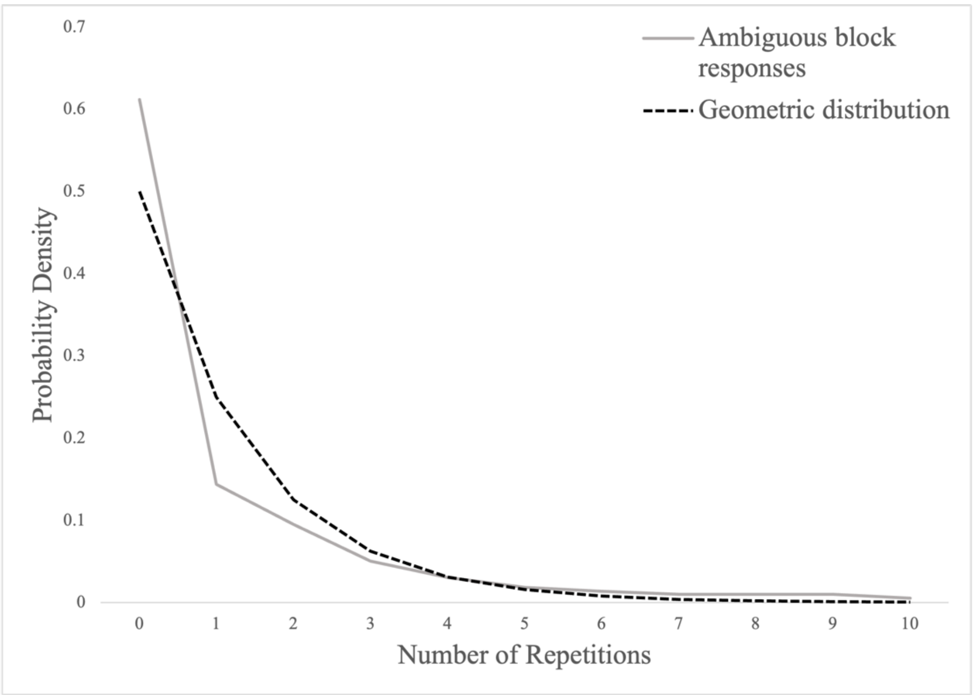

Figure 4.7 Solid grey line shows the probability density of the number of repeated responses across both response types (i.e., left or right) in ambiguous blocks (e.g., for 3 repetitions responses to 3 consecutive trials could be "left", "left", "left"). Dashed black line shows the probability density of the geometric distribution.

4.3.1.5 Directional block trials

In the adapted condition, participants responded to indicate a reversal on 61.58% (SD = 13.92) of trials. This is significantly greater than would be expected by chance (i.e., 50%), $t(24) = 2.421$, $p = .023$, $d = .484$. This is also significantly greater than the percentage of reversal responses in the unbiased condition (47.58%, SD = 8.67), $t(24) = 2.984$, $p = .006$, $d = .597$.

In the unbiased condition, participants responded to indicate a reversal on 47.58% of trials (SD = 8.67). This is not significantly different from the responses that would be expected by chance, $t(24) = 1.394$, $p = .176$. A Bayesian one-sample t-test was also conducted to assess the extent to which the data was supportive of the null hypothesis (H_0), i.e., that the percentage of reversals did not differ from chance in the unbiased condition. The resulting Bayes factor indicated evidence for H_0, with $BF_{01} = 2.009$, which means that the data were approximately twice as likely to occur under H_0 than H_1. The value for BF_{10} (i.e., the likelihood of the outcome occurring under H_1 compared to H0) was therefore 0.498 (as it is the reciprocal of BF_{01}). This indicates that there is anecdotal evidence for the null hypothesis, according to the ranges indicated by Lee and Wagenmakers (2013). As with responses in the ambiguous only block, the number of response repetitions (for either left or right responses) were found to be well fitted to a geometric distribution (see Figure 4.8), $R^2(24) = .98$, $p <.001$.

As with the ambiguous block, null responses in the directional block were also removed from the analysis. This led to an average of 6.72% of trials (SD = 5.96) being removed.

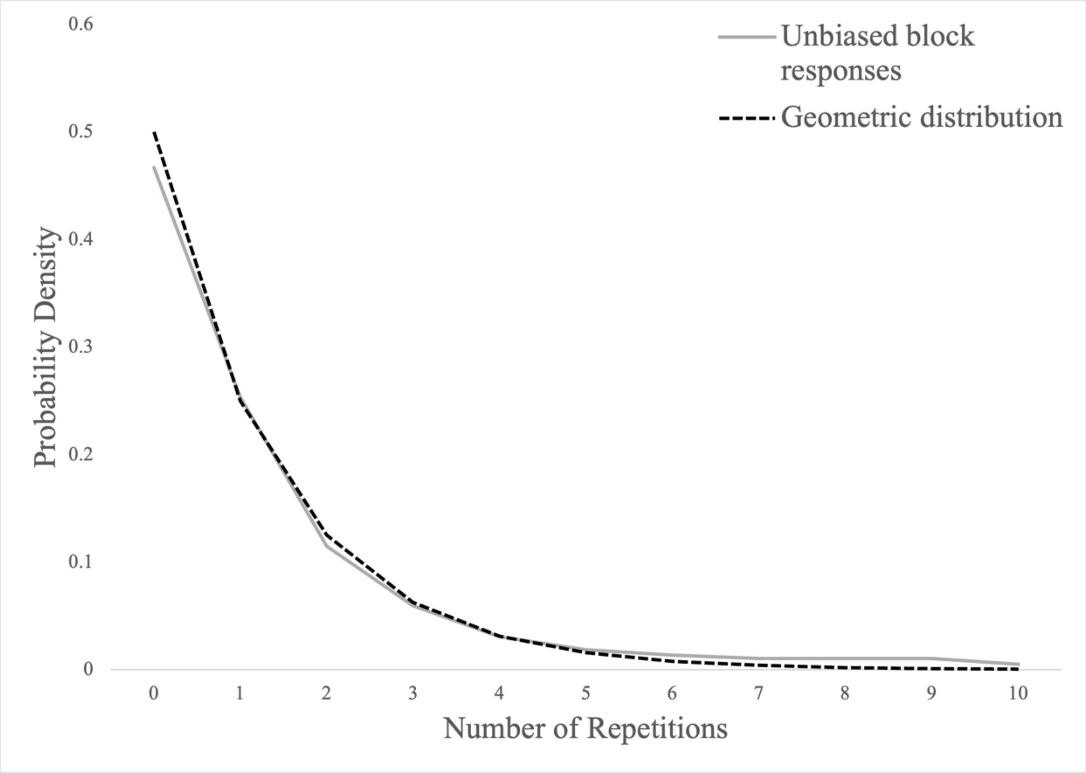

Figure 4.8 Solid grey line shows the probability density of the number of repeated responses across both response types (i.e., left or right) in unbiased trials in directional blocks. Dashed black line shows the probability density of the geometric distribution.

4.3.2 Electrophysiological results

4.3.2.1 ERP

Figure 4.9 shows the grand average ERP waveforms for reversal and stable trials in the ambiguous block only. The RP and RN ROIs are represented by the shaded regions. The RP was defined a priori with a temporal ROI of 100ms – 200ms over channels O1, O2 and Oz. The RN was defined a priori with a temporal ROI of 200ms – 400ms over channels O1, O2, Oz, PO7 and PO8 (Kornmeier & Bach, 2012; Abdallah & Brooks, 2020). For each participant and condition, the average amplitude within each temporal ROI above was calculated for each of the channels. For all analyses outlined below, the Greenhouse-Geisser correction was applied when Mauchly’s test of sphericity was found to be significant at the $p = .05$ level.

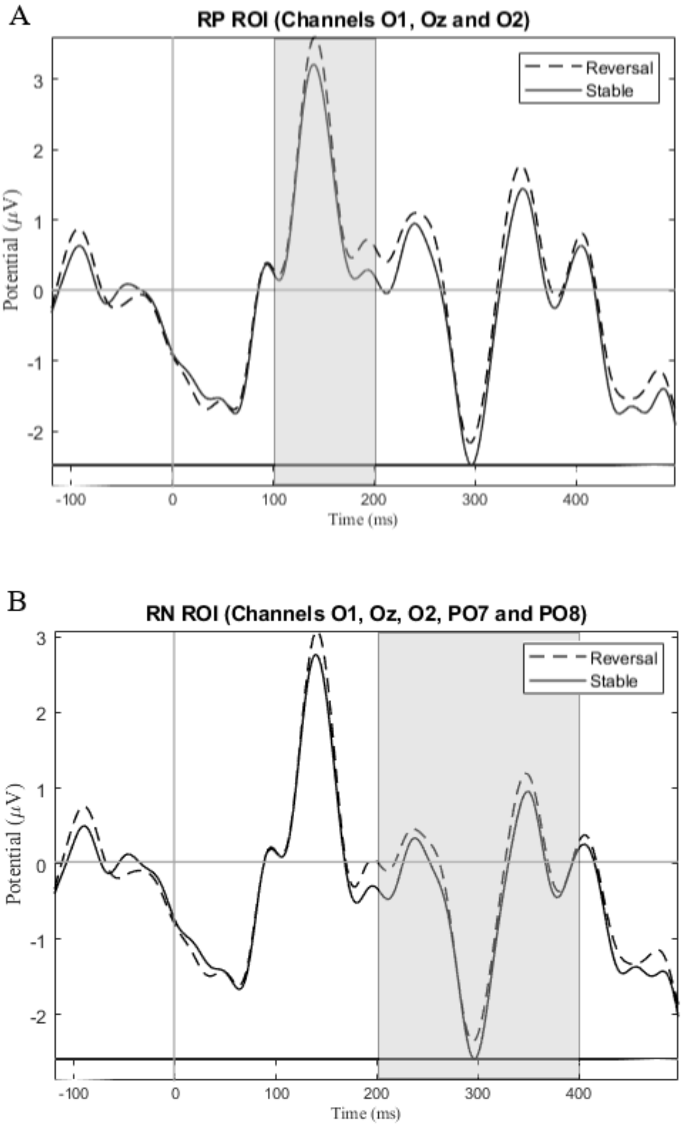

Figure 4.9 Both panels plot the grand average ERP waveforms for reversal (dashed line) and stable (solid line) trials in Experiment Three. Panel A: ERP waveforms averaged over the RP ROI (channels O1, Oz and O2). Panel B: ERP waveforms averaged over the RN ROI (channels O1, Oz, O2, PO7 and PO8). The grey boxes indicate the time-windows used to define the RP and RN mean amplitudes (mean amplitudes shown in Figure 4.10 below; see 4.2.4 Data analysis section for more detail).

4.3.2.1.1 Reversal positivity

The mean amplitude over the RP spatio-temporal ROI was calculated for each response type (i.e., reversal and stable) for each participant. These are shown in Figure 4.8 below, for both reversal (grey bars) and stable (black bars) trials.

A 3x2 (Channel x Response Type, i.e., reversal or stable) repeated measures ANOVA, revealed a significant main effect of Channel, $F(2, 48) = 9.212$, $p < .001$, $\eta_p^2 = .277$. No other factors or interactions were revealed to be significant, all $p > .486$ (for full ANOVA results for the RP ROI see Table 4.2 below).

Table 4.2

Full repeated measures ANOVA results for the mean ERP amplitude over the RP ROI

	df1	*df2*	*F*	*p*	η_p^2
Channel	2.000	48	9.212	**< .001**	.277
Response Type	1.000	24	0.085	.773	.004
Channel x Response Type	1.360	48	0.617	.486	.025

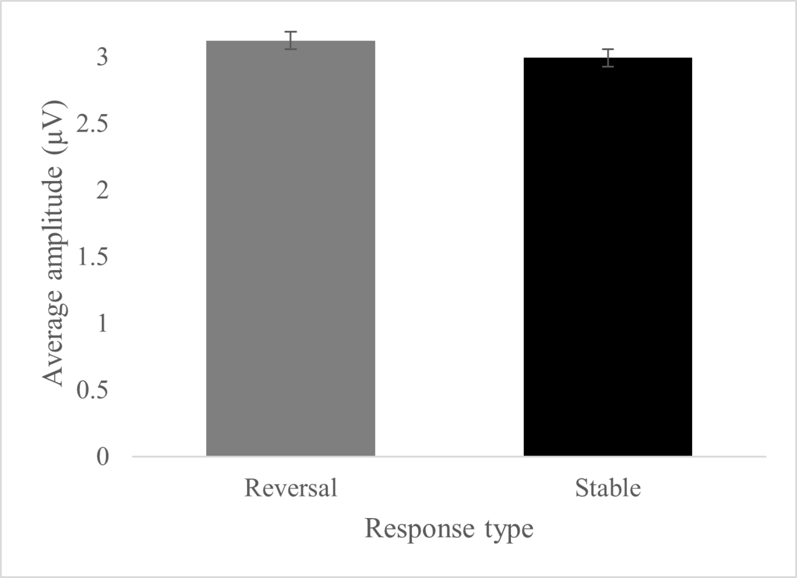

Figure 4.10 Mean ERP amplitudes in the RP ROI for reversal trials (grey bar) and stable trials (black bar). Error bars represent the standard error of the mean.

4.3.2.1.2 Reversal negativity

The mean amplitude over the RN spatio-temporal ROI was calculated for each response type (i.e., reversal and stable) for each participant. Grand average amplitudes in the ROI are shown in Figure 4.11 below, for both reversal (grey bars) and stable (black bars) trials.

A 3x2 (Channel x Response Type, i.e., reversal or stable) repeated measures ANOVA, revealed a significant main effect of Channel, $F(2, 48) = 9.212$, $p < .001$, $\eta_p^2 = .227$. No other factors or interactions were revealed to be significant, all $p > .626$ (for full ANOVA results for the RN ROI see Table 4.3 below).

Table 4.3

Full repeated measures ANOVA results for the mean ERP amplitude over the RN ROI

	df1	*df2*	*F*	*p*	η_p^2
Channel	2.198	48	7.064	**.001**	.227
Response Type	1.000	24	0.061	.806	.003
Channel x Response Type	2.071	48	0.461	.626	.019

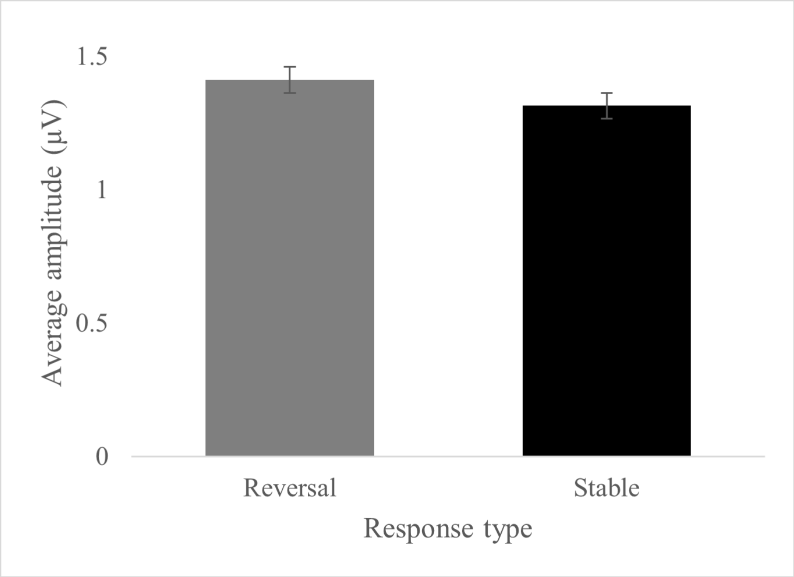

Figure 4.11 Mean ERP amplitudes in the RN ROI for reversal trials (grey bar) and stable trials (black bar). Error bars represent the standard error of the mean.

4.3.2.2 Multivariate analyses

4.3.2.2.1 Predicting spontaneous reversal vs stable trials in the ambiguous condition

To determine whether the pattern of EEG scalp voltages differed during spontaneous reversal and stable trials, the first analysis was conducted on the data from ambiguous only blocks. A linear support vector machine (SVM) classifier was trained to classify trials as reversal or stable. If enough class-related information exists in the data, we would expect the decoding accuracy to be significantly greater than would be achieved by chance.

Cluster based statistics (see Data Analysis Methods for a detailed explanation of these) of the MVPA revealed that the SVM classifier did not correctly distinguish between reversal and stable trials at a level significantly greater than chance (i.e., significantly greater than 50%) at any point throughout the ambiguous only block trials. The average decoding accuracy value across the whole trial was 50.07% (min = 48.61%, max = 52.11%). A time series of grand mean decoding accuracy is shown below in Figure 4.12.

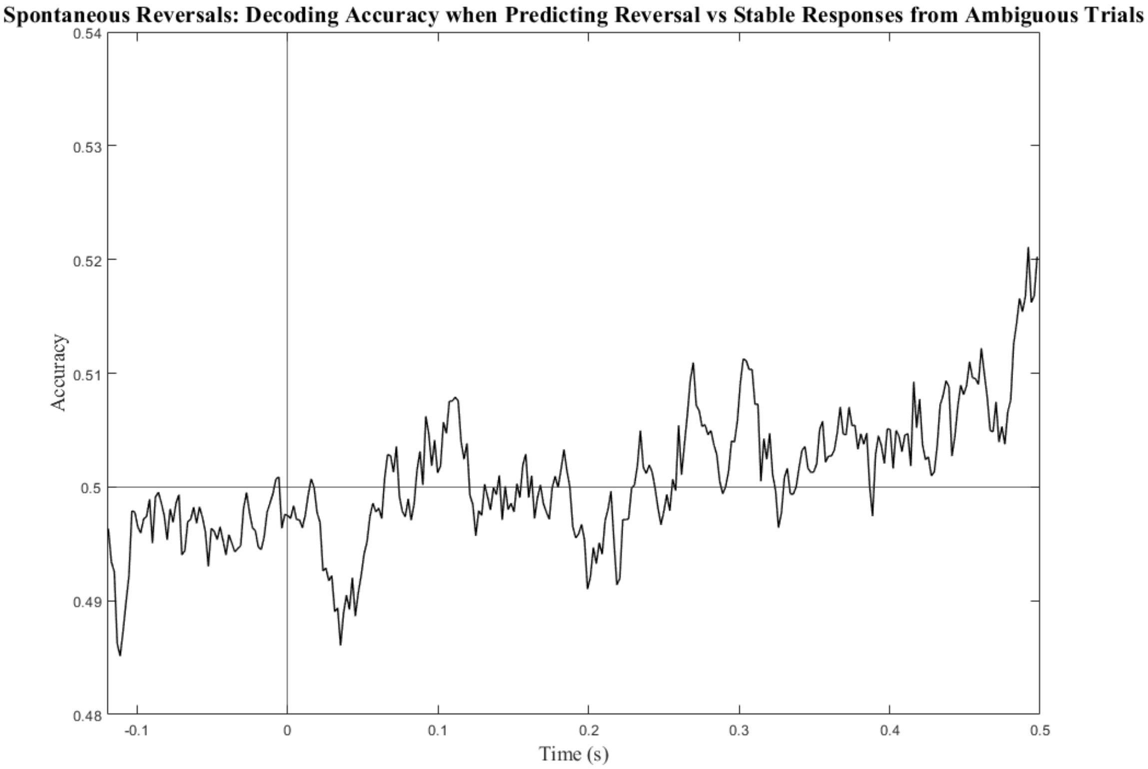

Figure 4.12 Plot of mean reversal vs. stable trial classifier decoding accuracy across time in the ambiguous only block.

4.3.2.2.2 Predicting adapted/unbiased trial type from spontaneous reversal data

To address whether spontaneous reversals in the ambiguous trial sequence are driven by adaptation, we trained a classifier to distinguish between adapted and unbiased reversal trials and then tested this on ambiguous reversal trials, recording which class label it predicted from the data (for a full explanation see Section 4.5 and Figure 4.3 in particular). Our hypothesis was that, if spontaneous reversals are driven by adaptation, then the classifier should predict that the data from reversals in the ambiguous trials come from the adapted class significantly more often than chance. Here then, rather than decoding accuracy, the percentages of each class label (i.e., adapted and unbiased) that the classifier predicted at each time point were averaged for each trial and across participants. Cluster-based statistics were then used to determine whether the classifier predicted that the spontaneously occurring reversal data came from the adapted reversal data class more often than chance. This revealed three significant clusters, during which the classifier predicted that the spontaneously occurring reversal data came from an adapted reversal trial significantly more often than expected by chance (i.e., 50%). The first cluster occurred in the pre-stimulus period between -84ms and – 68ms, $t_{cluster}(24) = 24.55$, $p = .0397$ with the mean percentage of adapted class labels in the cluster being 58.05% (min = 54.42%, max = 58.05%). The second cluster occurred between 45ms and 68ms, $t_{cluster}(24) = 35.57$, $p = .0022$ with the mean percentage of adapted class labels in the cluster being 55.32% (min = 53.68%, max = 57.04%). The final cluster occurred between 234ms and 250ms, $t_{cluster}(24) = 25.72$, $p = .0311$ with the mean percentage of adapted class labels in the cluster being 54.66% (min = 53.20%, max = 56.21%). See Figure 4.13 below for an illustration of these results.

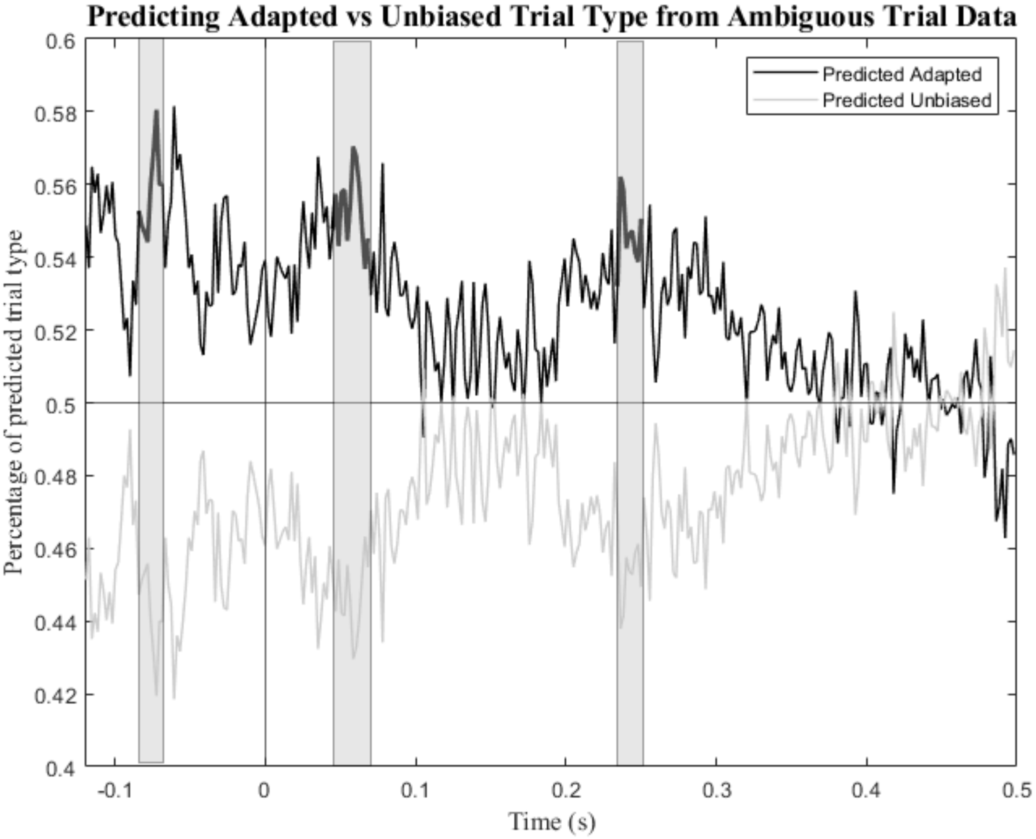

Figure 4.13 Plot to show the percentage of each predicted trial type (i.e., adapted trials or unbiased trials) across time when the classifier was provided with data from spontaneously occurring reversals from the ambiguous only trials. Grey boxes indicate the significant clusters obtained from the cluster-based permutation tests of the prediction percentages. These clusters represent time periods where a significantly higher percentage of trials were labelled as being more like the adapted reversal data than can be expected by chance.

4.3.2.2.3 Cross-decoding between conditions

The cross-decoding approach used in this analysis was conducted in order that more direct conclusions can be drawn on the role that adaptation plays in spontaneously occurring perceptual reversals. This analysis involved training two separate SVM classifiers: one to decode reversal vs stable trials from the adapted condition and one to decode reversal vs stable trials from the unbiased condition. Then, each of the two classifiers in turn were separately provided only with the EEG scalp voltage data from the ambiguous condition (i.e., the spontaneously occurring reversal or stable trials) and were each tested on their ability to correctly predict the true class label of these trials. If spontaneous reversals, like induced reversals, are underpinned by adaptation processes then we hypothesise that the classifier trained to decode reversal vs stable trials from the adapted condition should perform significantly better than that trained on trials in the unbiased condition, when decoding the spontaneously occurring reversals in the ambiguous condition.

Therefore, to assess this, the decoding accuracy of each classifier was recorded and compared relative to one another as opposed to being measured against chance. Cluster-based statistics of the decoding accuracy values revealed one significant cluster, during which the classification accuracy of the classifier trained on the adapted data was significantly higher than that of the classifier trained on the unbiased data. This cluster occurred between 64ms and 78ms, $t_{cluster}(24) = 38.00$, < .001 where the mean decoding accuracy of the adapted-trained classifier was 52.98% (min = 52.41%, max = 53.42%) and that of the unbiased trained classifier was 50.13% (min = 49.48%, max = 50.57%). See Figure 4.13 below for an illustration of each classifier's performance compared relative to one another.

Each classifier's performance was then also measured separately, relative to chance level (i.e., 50%). Here, for the adapted-trained classifier, the cluster-based statistics revealed one significant cluster during which that classifier's performance was above chance level. This cluster occurred between 60ms and 99ms, $t_{cluster}(24) = 21.77$, $p < .001$ with a mean decoding accuracy of 52.53% (min = 51.52%, max = 53.42%; see Figure 4.13, panel A below). However, for the classifier trained on the unbiased data, no significant clusters were revealed. Mean decoding accuracy across the trial was 50.12% (min = 48.35%, max = 52.04%; see Figure 4.15, panel B).

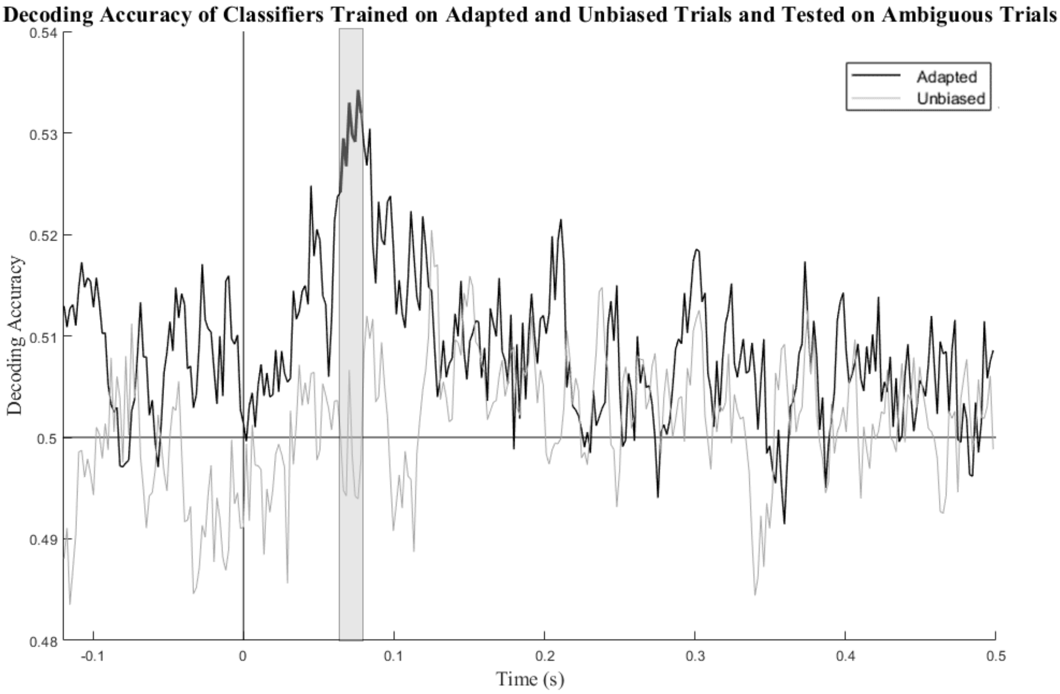

Figure 4.13 Mean reversal vs stability decoding accuracy across time for two classifiers. The adapted trained classifier, trained to distinguish reversal vs stable trials that are due to adaptation, and the unbiased trained classifier. Decoding accuracy values come from testing each of the two classifiers separately on data from spontaneously occurring reversal vs stable trials from trials in the ambiguous condition. The shaded region shows the time points where the performance of the adapted-trained classifier was significantly greater than that of the unbiased-trained classifier.

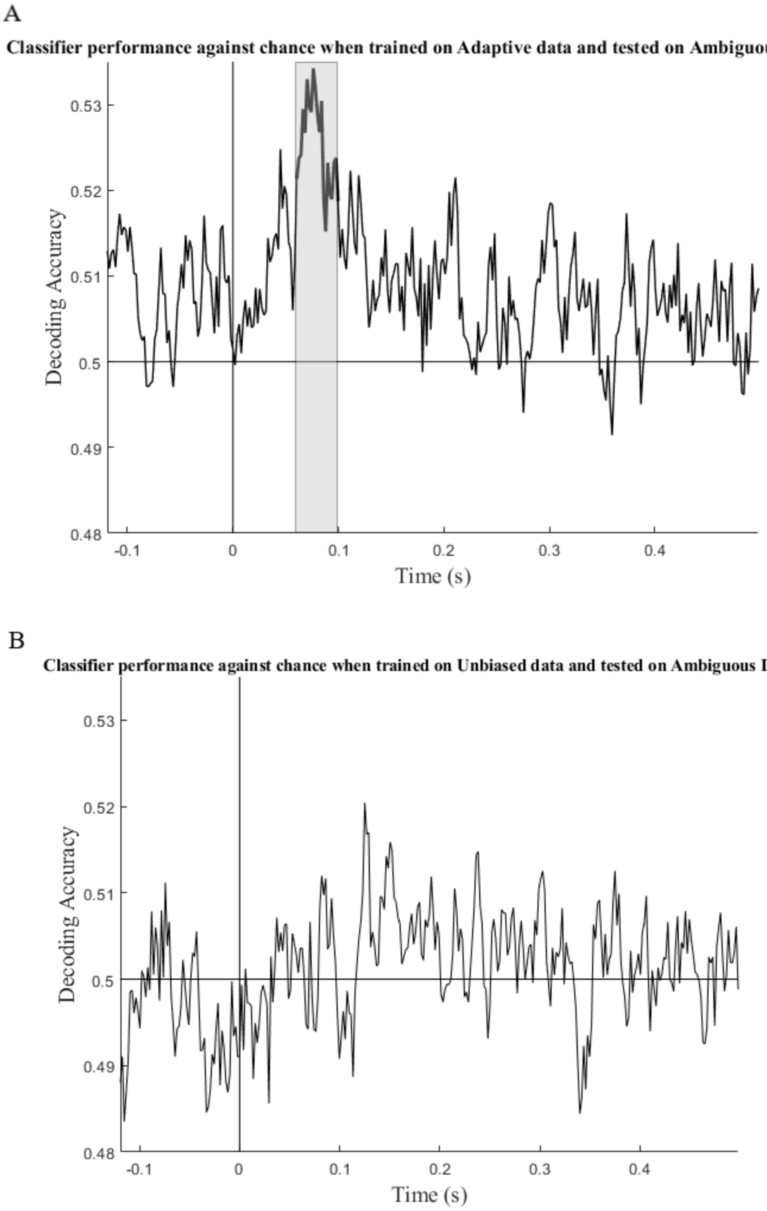

Figure 4.15 Plot showing the performance of the adapted-trained classifier (panel A) to correctly predict reversal vs stable trials when tested on data from ambiguous trials, the grey box shows the time period identified as significantly greater than chance level by the cluster-based statistics. Panel B shows the same performance of the classifier trained on data from unbiased trials.

4.3.2.2.4 Predicting reversal vs stable trials within the adapted and unbiased conditions

This analysis was performed to determine whether there was an inherent difference in the data quality in the adapted or unbiased conditions (for instance an increase in the signal to noise ratio in the adapted condition) that could be causing any increase in their performance when decoding spontaneous reversals from the ambiguous trial data.

To do this, a linear SVM classifier was trained and tested on its ability to classify trials as reversal or stable using data from trials within the adapted condition, in the same way as was performed in the ambiguous condition (see Section 4.3.2.2.1). Another classifier was then trained and tested in the same way, on data from trials within the unbiased condition. To determine whether the performance of one was significantly better than the other, cluster-based statistics were used to compare the decoding accuracy time series against each other. These revealed no significant differences in classifier performance across the post-stimulus period, however they did reveal two pre-stimulus periods where the performance of the adapted-trained classifier was greater than that of the unbiased-trained classifier. The first cluster occurred between -120ms and -116ms, $t_{cluster}(24) = 18.64$, $p = .002$ with a mean decoding accuracy of 51.42% (min = 50.99%, max = 51.74%). The second cluster occurred between -71 and -67ms, $t_{cluster}(24) = 28.32$, $p < .001$ with a mean decoding accuracy of 51.82% (min = 51.63%, max = 52.01%; see Figure 4.16).

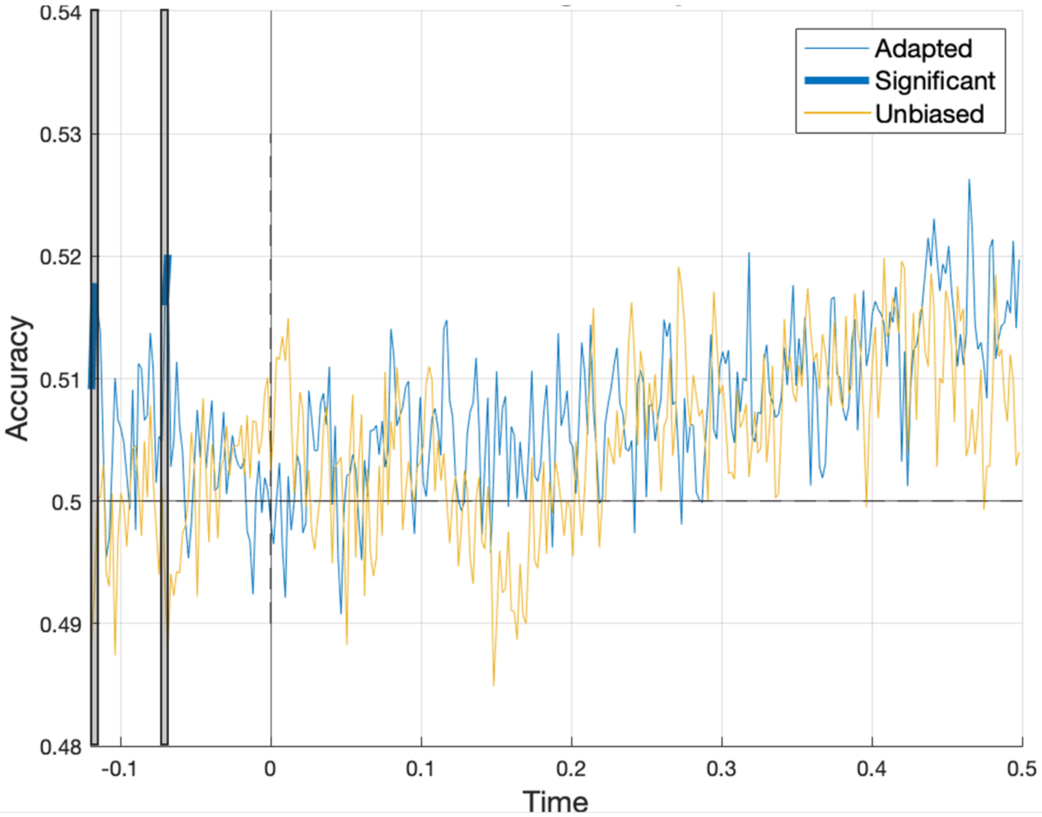

Figure 4.16 Mean reversal vs stability decoding accuracy across time for two classifiers; one trained and tested on data from the adapted condition and one trained and tested on data from the unbiased condition. The shaded region shows the time points where the performance of the adapted-trained classifier was significantly greater than that of the unbiased-trained classifier.

4.3.2.2.5 Temporal generalisation: Predicting spontaneous reversal vs stable trials

In these temporal generalisation analyses, an SVM classifier was trained to distinguish between reversal vs stable trials from data at a given time point. This trained classifier was then tested, not only on the data from the time point it was trained on but all other time points within the trial. This technique is more sensitive to effects that may be temporally misaligned between trials, as well as providing information on the time course of neural processes (i.e., whether they are sustained, isolated, reactivated and/or ramping up over the time course of the trial; King & Dehaene, 2014). Therefore, the technique was used here to help draw conclusions about the time course of adaptation in spontaneously occurring reversals, as well as provide a more sensitive measure to detect any effects that may not be temporally aligned.

Cluster-based statistics conducted on the temporal generalisation results (see Data Analysis Methods above for an explanation of temporal generalisation) revealed one cluster of testing/training time points where the classifier performed at a level greater than chance, $t_{cluster}(24) = 490.32$, $p < .001$. The mean decoding accuracy within this cluster was 52.01% (min = 50.85%, max = 54.13%). The cluster is illustrated in Figure 4.17 below.

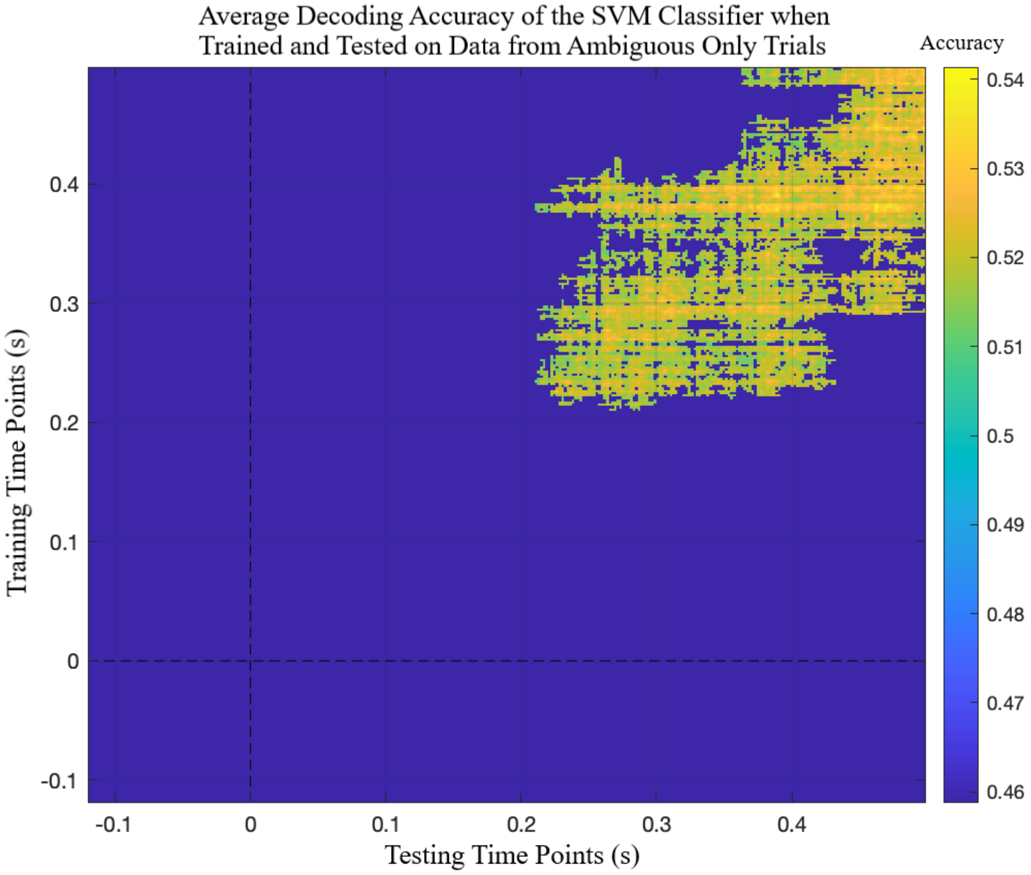

Figure 4.17 Temporal generalisation plot showing the significant cluster of reversal vs stability decoding accuracy values from a classifier trained and tested on data from ambiguous trials. Non-significant areas are masked out by the dark blue areas. The yellow areas represent training x testing time points with the highest levels of decoding accuracy and the light blue areas represent those with the lowest.

4.3.2.2.6 Temporal generalisation: Cross-decoding between conditions

To test whether the activity differentiating reversals and stable during ambiguous trials was more similar to that during adapted or unbiased trials, we used the same cross-decoding approach taken and described in Section 4.3.2.2.3 above. Temporal generalisation plots were generated by training and testing two classifiers in the same way as described in Section 4.3.2.2.3 above. Namely, two separate SVM classifiers were trained: one to decode reversal vs stable trials from the adapted condition and one to decode reversal vs stable trials from the unbiased condition. Then, each of the two classifiers in turn were separately provided only with the EEG scalp voltage data from the ambiguous condition (i.e., the spontaneously occurring reversal or stable trials) and were each tested on their ability to correctly predict the true class label of these trials. For an explanation of the temporal generalisation technique, see Section 4.2.5.

Cluster-based statistics conducted on the decoding accuracy values of the adapted-trained classifier revealed a cluster of testing/training time points where the classifier performed at a level significantly greater than chance, $t_{cluster}(24) = 1.066 \times 10^3$, $p < .001$. The mean decoding accuracy within this cluster was 52.39% (min = 50.93%, max = 55.36%). This is shown in Figure 4.18 below. However, no significant clusters were revealed from the decoding accuracy values of the unbiased-trained classifier, the mean decoding accuracy across each combination of training and testing timepoints was 50.42% (min = 48.39 %, max = 53.37%). As the significance mask omits all decoding accuracy values, the resulting plot is not shown here.

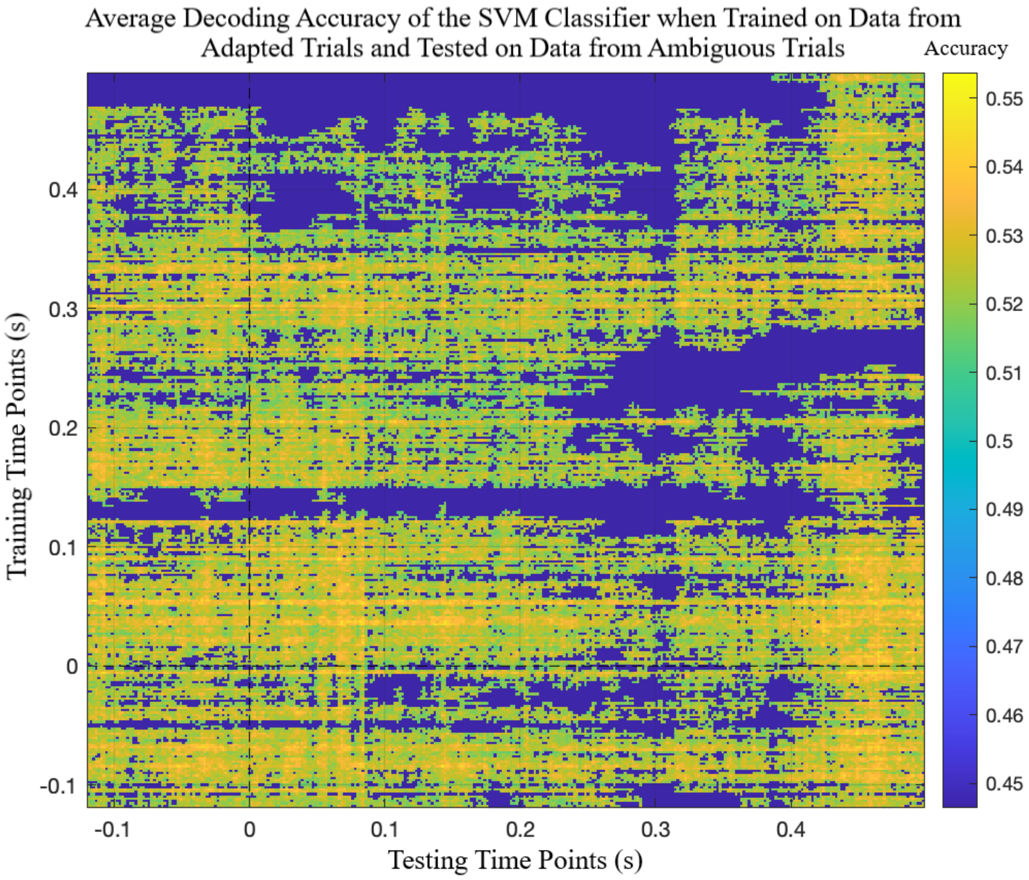

Figure 4.18 Temporal generalisation plot showing the significant cluster of reversal vs stability decoding accuracy values from a classifier trained on data from adapted trials and tested on data from ambiguous trials. Non-significant areas are masked out by the dark blue areas. The yellow areas represent training x testing time points with the highest levels of decoding accuracy and the light blue areas represent those with the lowest.

4.4 Discussion

The aim of this chapter was to use multivariate analysis of EEG data and cross-decoding techniques to investigate the role of adaptation in spontaneously occurring perceptual reversals. Three trial types were used to induce three different ‘brain states’ by varying the presentation time and use of the adaptor stimulus. These created reversals in perception that occurred due to adaptation to directional motion, reversals that occurred after directional motion was presented but did not induce any behavioural adaptation effects and finally, reversals that occurred spontaneously. Training the classifier to distinguish between reversal vs stable trials in the adapted condition provided sufficient information for that classifier to correctly decode spontaneous reversals occurring in the ambiguous condition. This suggests that similar mechanisms underpin both and thus that adaptation may be playing a role in spontaneous perceptual reversals even when an explicit adaptor is not present. Moreover, a classifier that was trained to adapted and unbiased reversal trials and then tested on ambiguous reversal trials, was more likely to classify those ambiguous trials as adapted trials than as unbiased trials. Taken together these results suggest that a similar adapted brain state occurs during both reversals due to adaptation and those occurring spontaneously.

The behavioural results from Experiment 3 showed that presenting the ambiguous motion stimuli, using the parameters established in Experiment 1, successfully induced the adaptation effect in participants. Participants were more likely to perceive ambiguous motion as moving in the opposite direction as preceding unambiguous motion when that unambiguous motion was presented for 640ms (see Figure 4.2 for trial structure). Although this was already shown in Experiment 1, it was important to ensure that the effect could be replicated here to address the aims of this chapter. This further bolsters the support provided by Experiment 1, for this rapid adaptation effect found by Kanai and Verstraten (2005). The behavioural results also showed that when the preceding unambiguous motion was presented for 80ms as it was in the unbiased condition, this did not significantly bias participants' perception of subsequently presented ambiguous motion. Moreover, this unbiased condition induced stochastic response patterns, suggesting that participants were not responding in a biased or systematic way. This was important to establish as this condition was necessary to serve as a control to compare classifier performance against. This was necessary so that any differences in the performance of the adapted-trained classifier and the unbiased-trained classifier cannot be attributed to stimulus differences. This is because the stimulus during the decoding period was physically identical between these two conditions. Instead, what differed was the trial history such that clear adaptation was built up in the adapted condition whereas there was no evidence of this in the unbiased condition. Behavioural analysis of responses to trials in the ambiguous block also showed that presenting two successive ambiguous motion stimuli did not bias participants' perception. Instead, here too, participants responded in an unbiased and stochastic manner. Again, this was previously shown in Experiment 1 but important to also replicate here to show that uncontrolled factors were not significantly influencing participants' perception.

Another aim of this chapter was to establish whether ERP components previously associated with endogenous, spontaneous perceptual reversals in a range of stationary ambiguous stimuli (Kornmeier & Bach, 2004, 2005, 2012), were also present when an ambiguous motion stimulus is used. However, neither the RP nor the RN components were found to be present in the ERPs from data captured during the spontaneously occurring reversals in ambiguous blocks. The RP has been found to be robust enough to remain unaffected by a number of low-level stimulus differences (Kornmeier et al., 2011; Kornmeier & Bach, 2012), and has been replicated across a number of stimuli (e.g., Britz et al., 2009; Kornmeier et al., 2007). However, perhaps because of its relatively low amplitude making it highly sensitive to averaging processes, it has not been replicated across all studies employing ERP methods (e.g., Intaite et al., 2010; Pitts et al., 2007) and has recently been shown to be linked to participants manually responding to a reversal rather than being linked to the perceptual reversal event itself (Abdallah & Brooks, 2020). Therefore, the absence of the RP component in Experiment Three could be due to a number of these factors or indeed a combination of several. The RN component has been equally replicated across various ambiguous stimuli (e.g., Britz et al., 2009; Intaite et al., 2010; Kornmeier & Bach, 2004) but, like the RP, not universally so (Kornmeier & Bach, 2014). It could be that a more widely distributed pattern of scalp activity is associated with the disambiguation of ambiguous motion compared to that of stationary stimuli, and ERP analyses are not capable of detecting this.

However, the MVPA classification across time results for the same data (i.e., the spontaneously occurring reversals in the ambiguous only block) also yielded null results. MVPA classification across time does indeed consider the pattern of activity across the whole scalp, rather than being confined to a subset of electrodes as with ERP analysis. However, it also faces somewhat similar issues to ERP analysis, in that the pattern of scalp activity is compared independently at each timepoint. So, this still requires some degree of temporal alignment of the effects on each trial. It could be the case here that the mechanisms at play at one time point are not active at exactly the same time on every trial. This idea is supported when the results from the temporal generalisation analysis are considered. In these results, there is a period of significant temporal generalisation beginning ~200ms where the classification performance at the training time generalises to other time points. The fact that this effect is only revealed with the temporal generalisation technique and not with the classification across time could suggest that there is a certain amount of temporal jitter in the effect's onset. In these experiments, in line with previous research (e.g., Britz et al., 2011; Intaite et al., 2010; Pitts et al., 2007), Kornmeier and Bach's (2005) Onset Paradigm was used. It is assumed in this paradigm that a reversal in perception occurs at, or very close to (with an accuracy of ±30ms; Kornmeier & Bach, 2005), stimulus onset. This Onset Paradigm has been used extensively with stationary ambiguous stimuli, but not with the ambiguous motion stimuli created by using sine-wave gratings as are employed here. It may be the case that, in stimuli like these, participants could require slightly different presentation times before a reversal occurs, leading to the results seen. The temporal generalisation pattern itself is similar to the 'ramping' pattern identified by King and Dahaene (2014; see Figure 4.19 below). King and Dahaene (2014) suggest that this pattern could represent a slowly increasing pattern of activity. When this is considered against the hypothesis that neural fatigue is a mechanism that builds up over sustained viewing of a stimulus, the results here provide support for this. It is also similar to the 'jittered' pattern identified by the same authors, supporting the idea that the onset of the reversal event itself could differ across participants.

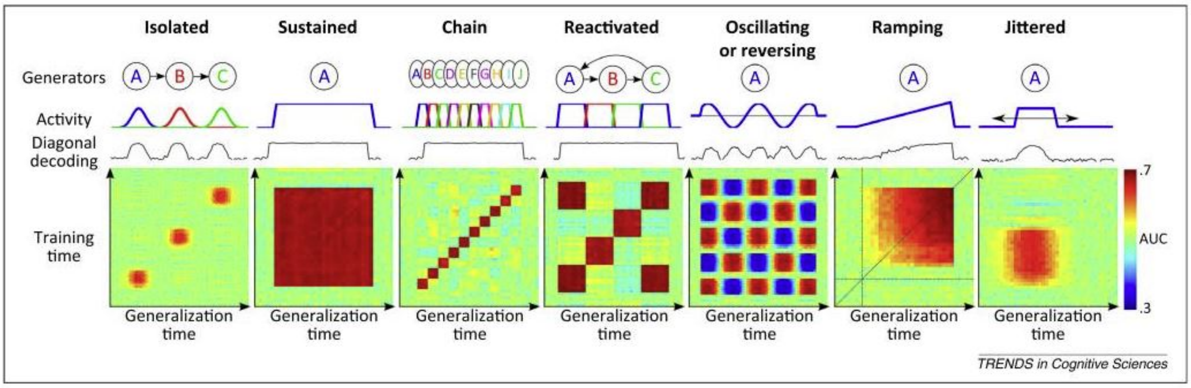

Figure 4.19 Figure from King and Dahaene (2014) showing seven simulated possible temporal generalisation matrices, each representing a different type of brain process. Isolated, indicating three separately activated processes; Sustained, indicating one process maintained over time; Chain, indicating a chain of distinct processes; Reactivated, where a given process is reactivated at a later time in the trial; Oscillating, where alternate above and below chance performance indicates a reversing component; Ramping, representing a slowly increasing process over time; and Jittered, indicating a difference in activation onset.

When the SVM classifier was trained to distinguish between reversals in the adapted and unbiased conditions (i.e., between reversals occurring due to adaptation and reversals occurring due to processes other than adaptation), that classifier identified three time periods in the time course of spontaneously occurring reversals where the pattern of scalp activity was significantly more like that of an adapted reversal. Two of these periods occurred post-stimulus onset, and one occurred in the pre-stimulus period. With regards to the two post-stimulus occurrences, the first of these occurred between 45 - 68ms and the second between 234 - 250ms. This indicates that, at these time points, the pattern of scalp activity during a spontaneously occurring reversal was significantly more like the pattern during an adapted reversal. This suggests that the underlying processes driving the scalp activity during spontaneous reversals are more similar to those underpinning adaptation, than those underpinning reversals occurring following directional motion that does not induce adaptation. This provides further support for the idea that adaptation plays a role in spontaneously occurring reversals. It should be noted too that the later significant time window here (234 - 250ms) occurs within the temporal ROI for the RN component (200 – 400ms). There is debate around the processes reflected by the RN, as it has been shown to be modulated by both top-down (Pitts et al., 2008) and bottom-up (Kornmeier et al., 2007) factors. So, the finding here, that there is a period of time during spontaneously occurring reversals where the brain activity within the RN window is similar to that in adaptation, could provide support for the idea that the RN at least partially represents bottom-up processes. This period also overlaps with the period during which there was significant above chance decoding of the ambiguous stimulus (i.e., spontaneous reversals) in the results from the temporal generalisation analysis.

With regards to the pre-stimulus period, between -84 and -68ms the classifier predicted that the scalp activity in the period preceding a spontaneous reversal was significantly more like that of an adapted reversal. However, it should be firstly noted that this period overlaps with that in which the signal to noise ratio in the adapted trials could be increased compared to unbiased trials. This is suggested because from -71 to -67ms a classifier trained to distinguish reversal vs stable trials from the adapted condition performed significantly better than an equivalent classifier trained and tested on data from unbiased trials. This suggests that there could be an increase in the signal to noise ratio in trials from the adapted condition at this time. Therefore, although other potential explanations for this result will be discussed, they should be taken with caution.

The result that, between -84 and -68ms, the classifier predicted that the scalp activity in the period preceding a spontaneous reversal was significantly more like that of an adapted reversal, could be interpreted in several ways. Firstly, this could indicate that in spontaneously occurring reversals neural fatigue has been built up during the previously presented ambiguous motion stimulus such that the brain state is already similar to that of an adapted reversal even before the stimulus has been presented. Secondly, in line with previous research (Hesselmann et al., 2008; Ronconi et al., 2017), this could be indicative of more of a spontaneously occurring oscillatory process in which this particular time period happens to be similar to that of an adapted reversal. If this were the case however, it could be expected that this time window would also lead to successful classification when the classifier was trained and tested on the same spontaneously occurring reversal data, but this was not shown.

The SVM classifier that was trained to distinguish adapted reversal vs stable trials was successfully able to predict spontaneously occurring reversal vs stable trials between ~60 – 78ms (classifier performance was above chance level from 60ms onwards, and above that of the unbiased control from 64ms onwards). This indicates that, during that time window, the pattern of brain activity was sufficiently similar in both the adapted trials and ambiguous-only trials. Importantly, the adapted-trained classifier can not only do this significantly better than would be expected by chance, but also significantly better than the performance of the classifier trained on the unbiased, control trials. Moreover, this result is not due to an underlying difference in the data quality in trials in the adapted condition over those in the unbiased condition. If that were the case, we would expect that when within-condition reversal vs stable classification performance in the adapted trials would be significantly better than that in the unbiased trials during this period. In other words, if there was an increase in the signal to noise ratio between ~60-78ms in the adapted trials compared to the unbiased trials this might lead to the adapted-trained classifier being able to perform generally better than the unbiased-trained classifier. However, when the performance of classifiers trained to decode reversal vs stable trials from within each of these two conditions was compared, no differences were found across the whole post-stimulus period. This indicates that it is not an increase in the signal to noise ratio at this point that is causing the adapted-trained classifier to be able to predict spontaneous reversal vs stable outcomes from trials in the ambiguous condition more accurately than the unbiased-trained classifier (and also than would be expected by chance).

It is important to clarify that the findings here are also not caused by any potential imbalance in the number of trials across different classes and conditions. Specifically, in both the unbiased and ambiguous conditions, the distribution of stable and reversal trials is approximately even, with each category representing about 50% of the trials. However, in the adapted condition, by design, there is a higher proportion of reversal trials, around 60%, compared to stable trials. This imbalance is proactively addressed during the pre-processing stage, prior to training the classifiers. We employ an undersampling technique during pre-processing to ensure equity between classes. In practical terms, this involves randomly selecting a subset of the majority class for each training fold—in the case of the adapted condition, this means selecting from the reversal trials—so that the number of trials in each class is equal. If this corrective step were omitted, the classifier trained on the adapted condition data would be predisposed to predict reversal trials more frequently, simply because they are more common in that set, which could bias the results.

During this time window (60-78ms), the adapted-trained classifier performs better when cross-decoding spontaneous reversal data from the ambiguous trials, than a classifier trained and tested only on the ambiguous trials themselves. This result is initially surprising, however, as discussed above, this could be due to a potential temporal jitter in the onset of a spontaneous reversal meaning that the standard classification across time analysis may be unable to detect the effect. This again seems to be confirmed by the results of the temporal generalisation matrices. The adapted-trained classifier can decode most other time points in the spontaneous data, including the pre-stimulus period. It can do this particularly well when the data from early in the adapted trial is used to predict later time points in the spontaneous data. Using the exemplar matrices provided by King and Dehaene (2014), this could suggest that the underlying mechanism that is driving decoding accuracy early on in an adapted trial is again active at later points in a spontaneous trial. In other words, the pattern of activity (and therefore underlying mechanisms) occurring when a participant has adapted to directional stimuli is sufficiently similar to that occurring towards the end of viewing an ambiguous stimulus. This could also provide additional support for the neural fatigue hypothesis as it suggests that the fatigue that is thought to be accumulated by the viewing of the directional stimulus may be similar to that building up when viewing an ambiguous stimulus and leading to a reversal in perception.

A potential confounding issue with the design of the experiment in this chapter is that, to induce adaptation effects, the duration of directional motion (i.e., the first stimulus in each trial) that participants were exposed to in the adaptation condition was longer than that in the unbiased condition (640ms vs 80ms, respectively). This was necessary to induce different levels of adaptation between the two conditions. To ensure that there were no stimulus confounds between the two conditions, the EEG results were based solely on the data during the ambiguous stimulus period (i.e., the second stimulus) which was identical (including duration) across the two conditions. Nonetheless, it is possible that the different durations of the first stimulus created a significant difference in motion energy which could have led to sensory carryover effects other than adaptation which affected neural activity during the second stimulus period. This could confound the comparison between adaptation and unbiased conditions during the unambiguous (i.e., second stimulus) despite perfect equating of stimulus characteristics during that period. One way to address this in future experiments could be to change the stimuli used, for example to random dot kinematograms (RDKs) in which the dots move in a random fashion without a coherent direction. These stimuli have been used in previous studies to investigate how different areas of the visual cortex, particularly the middle temporal area (MT+), process complex motion patterns. By using RDKs with varying levels of coherence and directionality, researchers have isolated responses to motion without the influence of a specific directional cue (Huk & Heeger, 2002). Therefore, these stimuli could be used as an alternative to the drifting sine-wave gratings in this set of experiments. For example, for the unbiased condition, one could replace the first stimulus with low coherence RDKs which should not induce adaptation because they do not contain coherent directional motion energy. This should be true regardless of how long it is presented. In contrast, for the adaptation condition, the first stimulus would comprise high coherence RDKs with coherent motion in either the left or right direction and should lead to adaptation to directional motion. This would allow one to equate the duration and cumulative motion energy of the first stimulus between the two conditions whilst still inducing a difference in adaptation.

When considered together, these results suggest that adaptation does have some role in driving spontaneously occurring reversals in perception. Decodable scalp activity patterns towards the end of a spontaneously occurring reversal or stable trial appear to represent a ‘ramping up’ mechanism that could be interpreted as neural fatigue. Periods of time exist in spontaneously occurring reversal trials where the pattern of scalp activity are similar to those occurring in reversals due to adaptation, suggesting that similar underlying mechanisms could be driving both. From these results alone however, it is not possible to say that adaptation is the only factor that drives spontaneous perceptual reversals. Future research could employ similar paradigms to those used here to explore the extent to which other factors could influence perception, for example volitional control or attentional factors. Replicating a similar paradigm using adaptation to other stimuli should also be encouraged, to determine whether the findings here are applicable to reversals in general, or reversals of ambiguous motion only.

Chapter Five: Simulations to examine error rate control and statistical power of multivariate pattern analysis pipelines

5.1 Introduction

Electroencephalography (EEG) data are inherently multivariate, capturing a complex interaction of brain activities across time and space. This multidimensional nature of EEG is rooted in its ability to simultaneously record electrical signals from multiple scalp electrodes with temporal precision in the order of milliseconds. Consequently, EEG data encompass a rich array of information that reflects both spatial and temporal variations in brain activity and researchers have developed a range of methods to make this vast amount of information interpretable in the context of a given research question.

In the field of cognitive neuroscience, making EEG data interpretable has traditionally been achieved via analyses of variance (ANOVA), or similar techniques, run on the mean or peak amplitudes in temporal and/or spatial regions-of-interest (ROIs) that are ideally defined a priori (e.g., Dien & Santuzzi, 2005) via a process known as spatiotemporal averaging (Luck, 2014). Analysing a result (such as mean or peak amplitude, or mean power of a given frequency band for instance) at a single time/space point without considering its relationship with other space/time points is known as a univariate approach. One example of such an approach, that is widely used in the literature, is the analysis of the event-related potential (ERP; Luck, 2014, see Chapter 2 for a full overview, however a brief description will be provided here). ERPs are voltage fluctuations in the EEG data that are time-locked to the onset of a particular event (e.g., stimulus onset, participant response). The mean amplitude differences between conditions are then compared. To generate the ERP, data from many trials are averaged (within conditions) to increase the signal to noise ratio and leave behind a waveform that reflects only the voltage pattern associated with the event (Luck, 2014).

Waveforms arising from different experimental manipulations are then statistically compared to determine if there are significant differences between conditions.

This method has been effective in revealing many ERP components that are associated with certain experimental manipulations. For instance, the negative-going N170 component associated with the viewing of faces (Bentin, Allison, Puce & Perez, 1996). However, the ERP method requires an ROI, which specifies the temporal and/or spatial focus of the analysis, to be set a priori. This is because running multiple significance tests on several ROIs without the appropriate corrections, carries a risk of drastically inflating Type I error rate (α; e.g., Kriegeskorte, Simmons, Bellgowan, & Baker, 2009; Kilner, 2013). Type I error rate refers to the probability of declaring a significant result when in fact none exists: in other words, a false positive result. It is accepted in the literature that a Type I error rate of 5% is sufficient to deem an analysis appropriate for assessing significance. Type I error rate can be increased if researchers use biased techniques with which to select their ROI (Kilner, 2013; Luck & Gaspelin, 2017; Brooks, Zoumpoulaki & Bowman, 2017); for instance, choosing the ROI a posteriori based on the largest difference between the amplitudes of ERP waveforms. To illustrate the effects of defining the ROI using such biased methods, Kilner (2013) conducted a series of simulations. The study generated null (i.e., no difference between conditions) EEG data for 16 subjects and conducted Student's t-tests to examine the null hypothesis that the value at a specific time point across subjects was equal to zero. When the time-window and electrodes were selected randomly, the false positive rate was around 5% (i.e., the accepted rate). However, when these parameters were selected based on where the 'effect' (i.e., difference between conditions in the data) was largest, the Type I error rate increased to over 95%. The study also demonstrated that averaging over different time-windows reduced, but did not eliminate, this bias. These results highlight the need for the ROI in ERP studies to be selected completely independently. Additionally, Luck and Gaspelin (2017) used simulations to show that using multifactorial ANOVAs to analyse ERP data can also lead to an increase in Type I errors, even when the ROI is chosen independently and a priori. For example, their null-effect simulation results revealed that using a four-way ANOVA to analyse their simulated data led to a false positive rate of approximately 53.3%, meaning that over half of the experiments yielded at least one significant but false effect. Therefore, the authors strongly advise against the uncritical use of multifactorial ANOVAs in ERP research, even if the ROI is selected a priori. An additional issue with ROI selection is that differences such as stimulus characteristics can lead to considerable disparities in the onset of certain effects (e.g., Brisson, Robitaille, & Jolicoeur, 2007; Kiesel, Miller, Jolicoeur & Brisson, 2008) as well as other factors that have no relevance to the study (e.g., Kutas, McCarthy, & Donchin, 1977), therefore making the a priori selection of an ROI problematic. Moreover, if novel research is being conducted with little or no previous work to guide the ROI selection, there might be no indication of which ROI to select at all.

In addition to data reduction via ROIs having the potential to inflate Type I error, another limitation of approaches like these is that the averaging process involved in creating measures like ERPs can lead to a loss of meaningful signal in the data (Luck, 2014; see Chapter 2 for more detail). This carries the potential of inflating the Type II error rate (β). Type II error rate is the probability of erroneously reporting that there is a null effect when in fact there is an effect present in the data. This is of particular importance because Type II error rate (β) and statistical power (1 - β) are inversely related. In other words, as the Type II error rate decreases, statistical power increases, and vice versa. Statistical power is the probability that a study will detect an effect if that effect truly exists (Cohen, 1988). A study with high statistical power, therefore, reduces the risk of committing a Type II error. The power of a study is influenced by several factors, including the sample size, the effect size, and the significance level set by the researcher. A widely recognised standard for power is 80% (Cohen, 1988). This standard stems from the concept that when using an alpha level of 0.05 (i.e., the accepted 5% Type I error rate), the ratio between a Type II error and a Type I error is 0.20/.05. This implies that mistakenly identifying an effect when none exists is deemed four times more critical than failing to recognise an effect when one is present in the population (Lakens, 2013). However, the current replicability crisis within psychological research (e.g., Ioannidis, 2005; Button et al., 2013; Colquhoun, 2014; Open Science Collaboration, 2015), indicates that studies are often underpowered. Indeed, Button et al. (2013) make the claim that, due to relatively expensive and often noisy data collection measures, low powered studies can be thought of as the norm, at least in cognitive neuroscience. Ensuring adequate power is crucial as it not only increases the validity of the study's findings but also contributes to the ethical conduct of research, preventing wastage of resources on studies that are unlikely to produce meaningful results. Adequate power is also important because it allows researchers to make a priori decisions about the number of participants required to detect a given effect. When a study is adequately powered, it means that the statistical tests used are sensitive enough to reliably identify true effects if they exist. This sensitivity allows researchers to determine the appropriate sample size required to achieve a balance between the risk of Type I and Type II errors while maintaining the precision and reliability of their findings. In practical terms, having adequate power ensures that researchers can design experiments with the confidence that they will have a sufficiently large sample size to detect meaningful effects, ultimately enhancing the robustness and credibility of their research outcomes. Thus, univariate approaches such as ERP analysis carry the potential to reduce power given that (a) they require an a priori ROI to be selected which may not be at the location of the effect and (b), they require some form of signal averaging which could lead to the effect, even if present within the selected ROI, to be cancelled out (Luck, 2014).

To help solve at least some of these issues, mass univariate analyses can be used (e.g., Woolrich, Beckmann, Nichols, & Smith, 2009), in which many univariate tests (such as t-tests, for example) are conducted in parallel. For example, when looking for a difference in N170 amplitude, one could test all time points between 150-200ms and all PO and O electrodes to determine whether there are any significant differences in that spatio-temporal region. This will result in a p-value for each time x electrode point. If there is no a priori information about the potential location of an effect in space and time, then all time and space points can be tested. Although mass univariate approaches involve analysis of multiple dependent variables (time x space points, in this case), this approach is different from a multivariate approach (discussed in more detail below). In a mass univariate approach, the statistical outcome for each time x space point is still calculated independently from that of the other time x space points being tested. In contrast, a multivariate analysis assesses the pattern of results across multiple time x space points.

Because a mass univariate analysis can include many concurrent statistical tests, a correction for multiple comparisons is necessary. There are many approaches to this correction which include permutation tests (e.g., Blair & Karniski, 1993), cluster-based permutation tests (e.g., Maris & Oostenveld, 2007) and false discovery rate control procedures (FDR; e.g., Benjamini & Hochberg, 1993; Benjamini, Krieger, & Yekutieli, 2006). The efficacy of such corrections in mass univariate ERP analysis was investigated empirically in a series of simulations by Groppe et al, (2011b). Here, EEG background noise was derived from 23 real participant EEG datasets, recorded from 26 channels, during the completion of a linguistic, text-based, priming task (Groppe, Choi, Topkins, & Kutas, 2009). On average, each participant's data contained 223 trials. The authors then computed ERPs from epochs that were time-locked to the onset of the text prime stimuli. The resulting ERP waveform was then subtracted from each epoch to produce trials of zero mean, null effect, EEG background noise. The authors used this realistic background noise as the basis to simulate 1000 ERP experiments for each of 4 simulated ERP effects. These were: the N170 (Bentin, Allison, Puce & Perez, 1996; Bentin, Mouchetant-Rostaing, Giard, Echallier, & Pernier, 1999), the P3 (Bentin et al., 1999), combined N170/P3 (i.e., both components were added to the data), and ERP null effects (i.e., no ERP effect was added to the null effect data). To simulate a single ERP experiment, ERPs for each of the 23 participants were derived by randomly selecting a subset of that participant's background noise trials, removing the mean pre-stimulus voltage, and averaging the trials. The authors then analysed the data from each of the 1000 simulated experiments for each of the 4 conditions, by concurrently applying two-tailed, one-sample *t*-tests to all scalp channels. They then used 6 commonly cited multiple comparison correction methods to correct for the multiple comparisons made. These were: Bonferroni-Holm, t-max permutation, maximum cluster-level mass permutation, permutation-based procedures, and 3 FDR control methods (as described in Benjamini & Hochberg, 1995; Benjamini & Yekutieli, 2001; Benjamini, Kreiger, & Yekutieli, 2006). The number of false positives for each correction method across each of the 4 simulated effect types was systematically counted to establish false positive rates for each method. This was used to evaluate Type I error risk of each of the 6 correction methods, for each of the 4 simulated effect types. The simulation showed that all 6 procedures performed with sufficient conservativism so as not to inflate Type I error rate to a concerning extent. In a companion to this simulations paper, the authors made recommendations following their investigation (see Groppe et al., 2011a). Of relevance to the EEG experiments described in this thesis, their paper made recommendations for the use of cluster-based permutation tests (Maris & Oostenveld, 2007). For a full discussion of these see Chapter 2, but briefly; cluster-based permutation tests involve grouping neighbouring data points (or 'clusters') that exhibit a significant effect, and then evaluating the statistical significance of these clusters through permutation testing. Groppe et al (2011a) recommend that these can be used in more exploratory analyses where there may not be a clear a priori hypothesis about the location (in time or space) of an effect and are particularly suitable for analyses where one is interested in capturing the spatial or temporal structure of the data (e.g., continuous time or space), rather than reducing the analysis to a small number of discrete tests. Fields and Kuperberg (2020) however, suggest that mass univariate approaches may have a place in EEG analyses as more standard practice and may even be more appropriate than spatiotemporal averaging approaches (i.e., ROIs).

Mass univariate analyses of ERPs, therefore, offer researchers an additional perspective on the analysis of EEG data. Specifically, they allow researchers to perform more exploratory analyses when an a priori ROI cannot be established, providing the appropriate correction for multiple comparisons is used. This has the potential to increase power over standard univariate techniques, as it avoids the problem of the incorrect ROI being selected (therefore missing the location of the real effect) based on a priori assumptions. However, some researchers have cautioned that the use of cluster-based permutation tests alongside mass univariate analyses may lead to the overestimation of the spatial and temporal precision of statistical claims made about any detected effect (Sassenhagen & Draschkow, 2019). Additionally, as is the case in standard univariate analyses of ERPs, they do not take full advantage of the whole-brain pattern of activity contained within the EEG dataset (Groppe et al., 2011; Hebart & Baker, 2018). In contrast, multivariate approaches involve the simultaneous analysis of multiple dependent variables in a statistical analysis. In a univariate analysis, each dependent variable is examined separately, while in a multivariate approach, multiple dependent variables are considered together as a set. This approach allows researchers to investigate complex relationships and interactions among variables, capturing the interdependencies that may exist between them. A subset of this approach is multivariate pattern analysis (MVPA) which, in contrast to univariate and mass univariate approaches, *does* examine whole-brain activity patterns across time. MVPA has been shown to be more sensitive in detecting effects that are missed by univariate analyses (i.e., leading to a reduction in Type

II error rate and therefore increased power; e.g., List, Rosenberg, Sherman & Esterman, 2017; Hebart & Baker, 2018). For a full description of MVPA and its reported advantages over univariate approaches see Chapter 2, however a brief overview will be provided here.

A typical multivariate analysis pipeline begins with training a classifier, for instance a machine learning classifier, such as a Support Vector Machine or Linear Discriminant Analysis based classifier (SVM and LDA, respectively). During the training process, the classifier is provided with EEG scalp voltage data along with the label of the class that those voltages are associated with. For instance, EEG data might be recorded from 64 channels during viewing of face vs house stimuli. To train a classifier, it would be provided with data from all 64 channels (although a smaller subset of channels can also be used), independently for each time point that was sampled and provided with a label for that pattern of scalp voltages. This process leads to the classifier 'learning' which patterns of data can best distinguish between the classes. This training process is conducted on around 80-90% of the total number of trials in the experiment. The trained classifier is then tested by being provided with the remaining 10-20% of the EEG scalp voltage data without class labels, and 'asked' to decode which class the unlabelled data came from. It is important to train and test the classifier on independent data to avoid the phenomenon known as “double dipping'. In other words, overfitting the model based on noise that is present in the data and leading to an increase in Type I errors (e.g., Kriegeskorte et al, 2009). Hence, 10-20% of the trials are kept back to independently test the classifier on. This leads to a time course of decoding accuracies, which are interpreted as a measure of how much information was present in the unlabelled data that was relevant to the class labels of interest (i.e., in this example, the viewing of faces vs house stimuli). If at a given time point, the trained classifier can predict which of (for instance) 2 classes the unlabelled data came from with an accuracy greater than chance, it can be concluded that the data must have contained sufficient class-specific information. Therefore, a difference between the pattern of scalp voltages associated with the 2 classes, at that time point, must exist. For an experiment with 2 classes, an accuracy of 50% is the level that would be expected by chance. Thus, timepoints with decoding accuracies significantly greater than this are generally accepted as one in which the effect has been detected – in other words, sufficient information about the experimental manipulation (in this example, the stimulus currently being viewed) is contained in the patterns of brain activity (e.g., Kriegeskorte et al., 2009; Grootswagers, Ward & Carlson, 2017; Hebart & Baker, 2018, though for a contrasting view see Combrisson & Jerbi, 2015). For an experiment with multiple participants, individual decoding accuracies are typically averaged to produce a 'grand average' plot of decoding accuracy across the time course of the trial. Finally, to determine which timepoints in an experiment lead to a decoding accuracy that is significantly greater than chance, a correction for multiple comparisons must be conducted given that a trial can consist of hundreds of time points. Although such correction methods are typically similar to those examined in the simulations by Groppe et al. (2011b), the currently available literature does not contain equivalent guidance on the power and FWER of these methods when they are applied to time series MVPA decoding accuracy data. Therefore, making a priori decisions about sample size is difficult when designing studies using MVPA.

The technique of MVPA originated in the field of computer science, primarily for use in the development of brain-computer interfaces (BCIs). In BCI applications, the focus is on predicting behaviour based on patterns of voltage across scalp electrodes; the number of participants is generally not a critical factor. This approach has been adapted for neuroimaging data, operating on the premise that if decoding accuracy exceeds chance levels, sufficient information about the experimental manipulation of interest is contained in the patterns of brain activity (Grootswagers et al., 2017; Hebart & Baker, 2018). While BCIs prioritise achieving high decoding accuracy that generalises across individuals, irrespective of the nature of the signal being decoded, MVPA in neuroimaging seeks to understand the structured representation of stimuli in underlying brain activity (Hebart & Baker, 2018). It is crucial to recognise that these two applications of MVPA are underpinned by different assumptions. For instance, in BCIs, the primary aim of MVPA is to achieve as close to 100% decoding accuracy as possible so that the algorithms generalise widely across different individuals. On the other hand, neuroimaging applications of MVPA aim to identify structure in the data related to the conditions under study, without necessarily requiring high levels of decoding accuracy, providing these levels are significantly greater than can be expected by chance (e.g., Grootwagers et al., 2017).

Because of this relative infancy as a technique in neuroimaging studies compared to BCI studies, the methodological rigour of MVPA in this context has not been as thoroughly scrutinised as that of ERP analysis pipelines. For instance, the currently available literature does not provide researchers with sufficient guidance on the number of participants required to detect an effect within a given ROI, where one exists. Therefore, the first aim of this chapter is to use simulated

EEG data to investigate the power of a typical MVPA pipeline. Specifically, to establish how varying the size of a multivariate effect affects the resulting decoding accuracy and power of cluster-based statistics to detect the effect. Specifically, three effect sizes (in terms of Cohen's d) will be assessed: small, medium, and large. Additionally, power will be calculated; that is, if any of the simulated experiments detect the added effect at any time point in the ROI then that will be counted as an experiment in which the effect has been detected. This will provide some much-needed general guidance on the necessary sample sizes to detect multivariate effects using MVPA pipelines on EEG data. Additionally, as MVPA alongside cluster-based permutation tests have been used in Chapter 4 of this thesis to draw conclusions on the role of adaptation in spontaneously occurring perceptual reversals, this aim will also shed light on how powerful this analysis pipeline is. Power is influenced by various critical components, including the predefined significance level (here set at 0.05 in line with the accepted rate), which represents the acceptable risk of Type I error (false positive); the effect size, which quantifies the magnitude of the true effect under investigation, and the sample size, which reflects the number of observations or participants in the study. Therefore, it is expected that the power of this analysis pipeline will increase along with an increase in participant number, as well as an increase in the size of the underlying effect.

The second aim of this chapter is to use simulated null effect EEG data to investigate the family-wise error rates associated with the same MVPA analytical pipeline. The purpose of this is to provide more information on the FWER of cluster-based permutations, specifically when analysing decoding accuracy values resulting from MVPA. Additionally, the analysis aims to establish whether the cluster-based permutations sufficiently correct for multiple comparisons, again within the specific context of analysing decoding accuracy values resulting from MVPA. Again, as with the previous aim, this will shed light on the Type I error rate associated with the analysis pipelines used in Chapter 4 of this thesis. Given that previous research has shown that when cluster-based permutations are used to correct for the multiple comparisons made in mass univariate analyses of ERP data (Groppe et al., 2011b), they do so with an acceptable level of conservativism, it is expected that they will also correct for the multiple comparisons in the context of analysing decoding accuracy values resulting from MVPA.

The general approach to address these aims will be to generate EEG data in a series of simulations in which effect size (in terms of Cohen's d) and sample size (i.e., participant number) are systematically varied. Three effect sizes will be used: small, medium and large (see below for an explanation of how these are calculated), and five sample sizes: 5, 10, 20, 40 and 80 virtual participants. Therefore, this leads to a 3 (effect size) x 5 (sample size) design, thus 15 simulations in total. Each of the 15 simulations will contain 1000 simulated experiments. Each simulated experiment will comprise simulated, independently generated EEG background noise data for each of the given number of participants (i.e., 5, 10, 20, 40 or 80). Each participant's dataset will contain 50 trials for each of 2 conditions and, at a specific ROI, these trials will contain an effect that is either a small, medium or large difference between the 2 conditions (see Figure 5.1 for an overview of the simulations; see below for a description of how the effects were calculated).

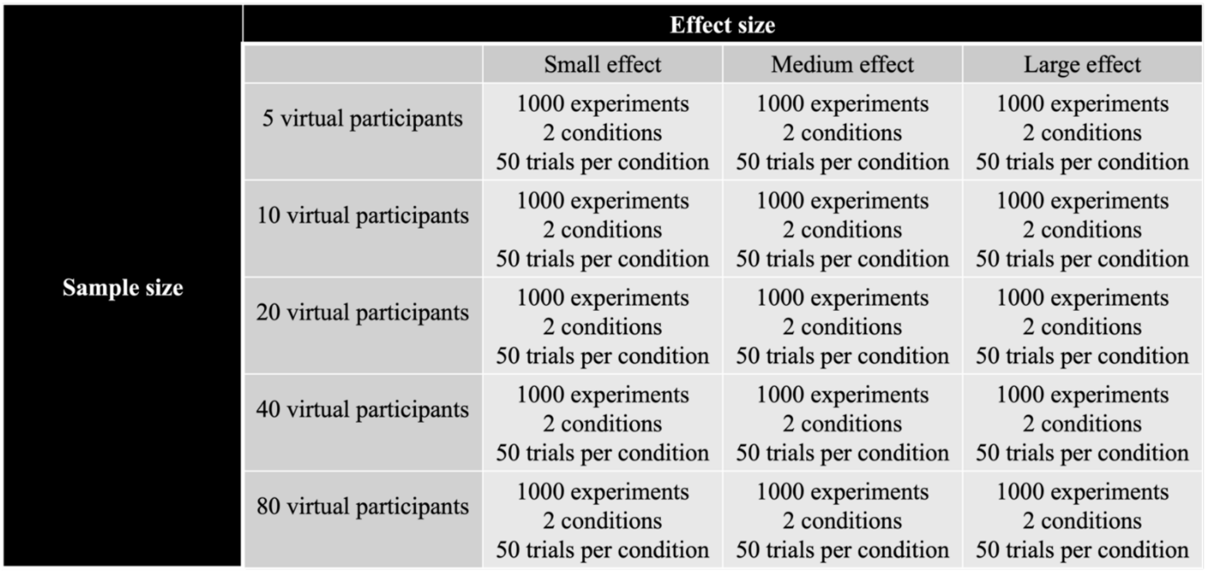

		Effect size		
		Small effect	Medium effect	Large effect
Sample size	5 virtual participants	1000 experiments 2 conditions 50 trials per condition	1000 experiments 2 conditions 50 trials per condition	1000 experiments 2 conditions 50 trials per condition
	10 virtual participants	1000 experiments 2 conditions 50 trials per condition	1000 experiments 2 conditions 50 trials per condition	1000 experiments 2 conditions 50 trials per condition
	20 virtual participants	1000 experiments 2 conditions 50 trials per condition	1000 experiments 2 conditions 50 trials per condition	1000 experiments 2 conditions 50 trials per condition
	40 virtual participants	1000 experiments 2 conditions 50 trials per condition	1000 experiments 2 conditions 50 trials per condition	1000 experiments 2 conditions 50 trials per condition
	80 virtual participants	1000 experiments 2 conditions 50 trials per condition	1000 experiments 2 conditions 50 trials per condition	1000 experiments 2 conditions 50 trials per condition

Figure 5.1 Diagram showing details of each of the 15 simulations conducted in this Chapter.

The noise-only, null hypothesis data will be generated during the process outlined above. Specifically, a subset of 1000 experiments will be taken from the total 15 000 simulated experiments, however an ROI outside of the added effect will be analysed. Therefore, this ROI will contain only the independently generated, simulated background EEG noise with no effect at all. The simulated experiments in the large effect/sample size of 80 simulation were arbitrarily chosen as the subset of experiments to use here. The choice of experiments to use here is arbitrary because the way that the simulated background EEG noise was generated was identical in all 15 simulations.

The data from each simulated experiment will firstly be analysed using MVPA, to generate time series data of decoding accuracies for each participant according to how well the classifier is able to distinguish between the 2 simulated conditions. To assess statistical power, one-sample t-tests alongside cluster-based permutation tests will be used to determine whether this resulting decoding accuracy is significantly above chance level at any of the time points within the ROI containing the added effect. This will be done for each of the 15 simulations (i.e., 3 effect sizes x 5 sample sizes shown in Figure 5.1). To assess FWER, again one-sample t-tests alongside cluster-based permutation tests will be used to determine whether this resulting decoding accuracy is significantly above chance level at any of the time points within the ROI outside of the added effect (and therefore contains noise-only, null hypothesis data). This will be done only for the 1000 simulated experiments in the large effect/sample size of 80 simulation.

Therefore, in the noise-only, null hypothesis simulation, the proportion of experiments (out of 1000 total) in which a significant effect is found by this pipeline will determine the FWER. As there is no effect present in the noise-only ROI, any significant increase in decoding accuracy is a false positive. In the simulations containing the addition of a multivariate effect, the proportion of experiments in which the effect is detected will determine the statistical power of this pipeline. Thus, using this approach, both aims outlined above will be addressed. This approach will therefore lead to 15 power values, as proportions out of 1000 experiments whereby the added effect was detected in the effect ROI (one for each of the 3 effect sizes x 5 sample sizes). The approach will also lead to one FWER value, as a proportion out of 1000 experiments whereby an effect was incorrectly detected in the ROI containing only noise (from the large effect/sample size of 80 simulation).

The way that the simulated effect is added to the noise-only data in these simulations differs from the way that previous simulation work has introduced univariate effects when examining univariate approaches. For example, the addition of a univariate effect involves introducing a mean difference in amplitude at given electrodes and timepoints, across 2 (or more) different simulated conditions. Therefore, to generate effects of various sizes, the mean amplitude difference between the 2 conditions can simply be increased. Here, a multivariate effect involving 32 channels was used. That is, the 2 simulated conditions contain a difference in the overall *pattern* of voltage changes, spanning 32 channels simultaneously. Therefore, generating multivariate effects of differing sizes is slightly more complex than when dealing with univariate effects. The general approach taken to generate the multivariate effect will now be outlined.

To achieve the addition of a multivariate effect, datapoints will be generated based on a multivariate Gaussian distribution (Hotelling, 1933; see Figure 5.2, below). In the set of simulations described in this chapter, this multivariate effect involves 32 channels. However, to outline the concept here, an example involving just 2 channels will be used. In Figure 5.2, below, hypothetical voltage recordings from 2 EEG channels are plotted on the x and y axes of the graph. The 2 ellipses shown on the graph represent 2 multivariate Gaussian distributions. The means of these distributions are known as centroids and are represented by the red and blue circle markers in the centre of each ellipse. The red ellipse will represent the distribution underlying hypothetical condition 1, and the blue ellipse will represent that of hypothetical condition 2. If these were used to generate data that represents a multivariate effect (in this case spanning 2 channels), then possible voltage pairs for electrodes 1 and 2 could be (3.5, 2) or (4, 1.5) for data in condition 1. These are represented by the red cross markers shown in Figure 5.1. Similarly, voltage pairs of (1.5, 3) or (2.5, 4) could represent data in condition 2. In the simulations described in this chapter, there were 32 channels, so the data are contained on a unit hypersphere that has 32 dimensions. The concept, however, remains the same as this 2-dimensional representation.

To generate effects of varying sizes therefore, it is the difference between these centroid means across the simulated conditions that is manipulated. Another simple visualisation of this is shown in Figure 5.3, below. A small difference in centroid means is represented in Figure 5.3 Panel B, where it can also be seen that the 2 multivariate distributions overlap, meaning that some points drawn from each separate distribution could belong to either of the 2 conditions. In contrast, a large difference in centroid means is represented in Figure 5.3 Panel A, where the distributions are distinct from one another.

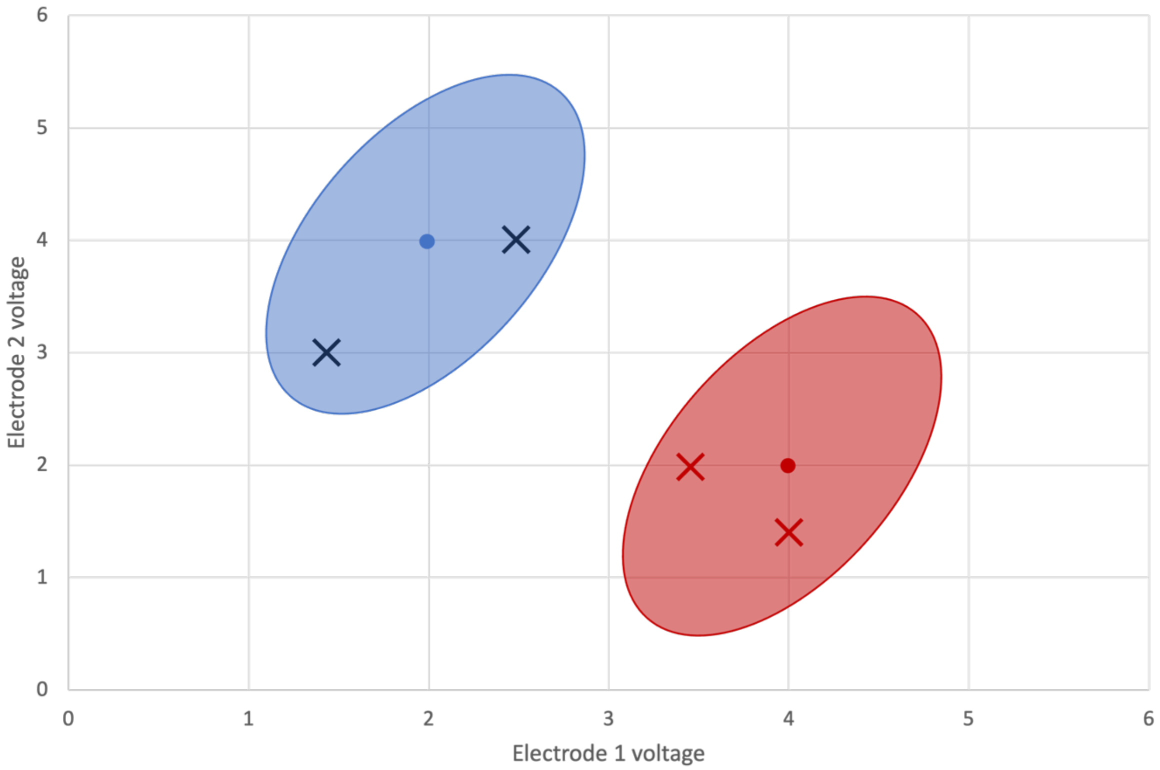

Figure 5.2 A graphical representation of 2 multivariate Gaussian distributions of voltages from 2 EEG electrodes, each represented by an ellipse

(red and blue ellipses). The ellipses themselves represent the spread of each distribution. The centroids of these distributions, indicated by the points at the centre of each ellipse, are the means of the distributions. The x markers represent 4 values that might be drawn from the distributions. Blue x markers represent values that might be drawn to generate the representation of condition 1, and red x markers indicate the same for condition 2.

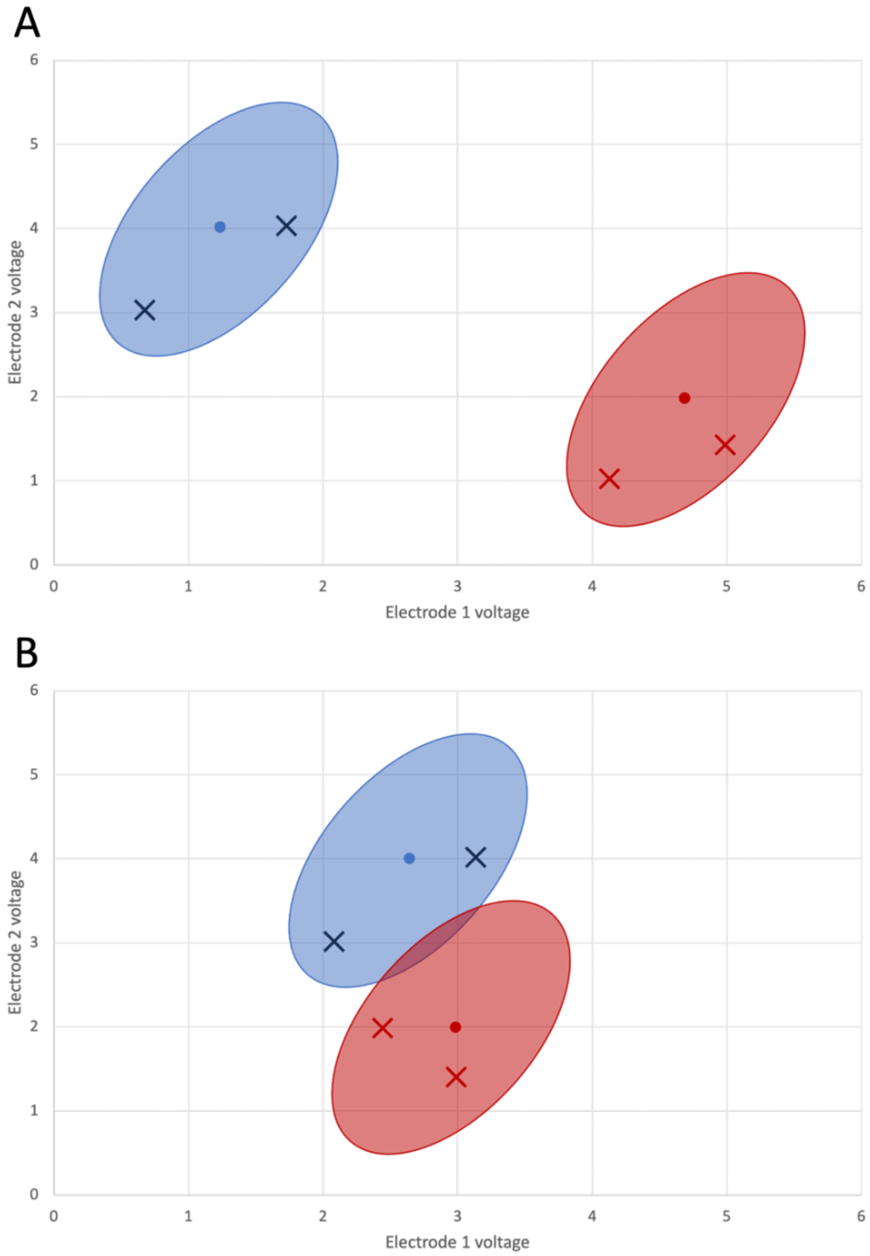

Figure 5.3 A graphical representation multivariate Gaussian distributions representing two multivariate effects spanning 2 electrodes. Panel A represents an effect with a large difference in the centroid means between the two conditions/distributions. Panel B represents an effect with a smaller difference in the centroid means between the two conditions/distributions.

Again, this is a simplified representation, but the concept of introducing multivariate effects that vary in size remains the same as that used in this chapter.

5.2 Methods

5.2.1 Participants

There were no human participants. In this simulation we varied the number of virtual participants. Five participant sample sizes were used, these were: 5, 10, 20, 40 and 80. It was not practically possible to run all possible sample sizes due to the computing time required to do so, thus focus was given to the range of sample sizes typically used in the currently available literature on the decoding of EEG data using MVPA. In single lab EEG studies, a sample size greater than 80 is uncommon and unpractical due to the length of time it takes to collect EEG data; therefore, we capped the range at this number.

5.2.2 Design

The independent variables in the multivariate effect simulation were effect size and participant number.

As well as the zero effect, null-hypothesis data, 3 effect sizes, in terms of Cohen’s d, were used to assess power. These were: small (d = 0.2), medium (d = 0.5), and large (d = 0.8; Cohen, 1988). Cohen’s d is a measure commonly used in the literature, to quantify the differences across conditions (Cohen, 1988). The formula to calculate effect size in this way is:

$$\text{Effect Size} = \frac{\text{Mean Difference}}{\text{Standard Deviation of Within-Condition Noise}}$$

This equation was used to generate the differences in centroid means necessary to generate multivariate effects of given sizes, in terms of Cohen's d. In this equation, the within-condition noise is known, as this noise is generated by the simulation. The standard deviation of this within-condition noise was calculated to be 1.59µV. Therefore, to calculate the required difference between the centroid means to generate effects of the 3 sizes, the following calculations were used:

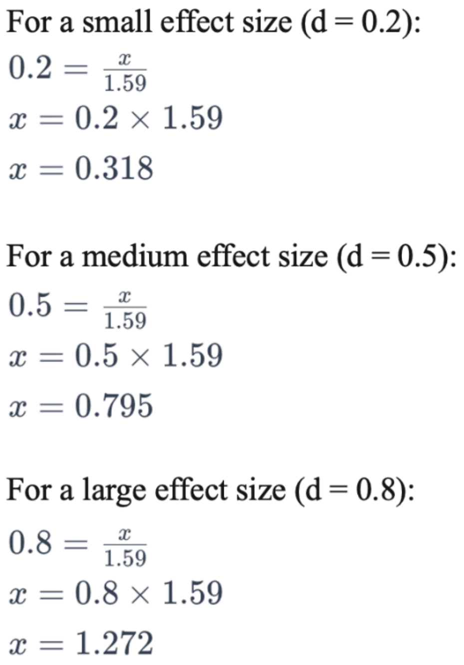

For a small effect size (d = 0.2):

$$0.2 = \frac{x}{1.59}$$
$$x = 0.2 \times 1.59$$
$$x = 0.318$$

For a medium effect size (d = 0.5):

$$0.5 = \frac{x}{1.59}$$
$$x = 0.5 \times 1.59$$
$$x = 0.795$$

For a large effect size (d = 0.8):

$$0.8 = \frac{x}{1.59}$$
$$x = 0.8 \times 1.59$$
$$x = 1.272$$

Therefore, the mean difference between conditions to represent a small effect in terms of Cohen's d value was represented in the multivariate pattern by a difference in the centroid means of 0.318. A medium effect was represented by a difference of 0.795, and a large effect by a difference of 1.272. The units here are arbitrary, as they represent the relative difference between the means in unit hyperspace (Treder, 2020).

These multivariate effects were added to simulated noise-only EEG data at 50 consecutive timepoints. Trials in the simulated experiment contained 512 samples and had a sampling rate of 256Hz, therefore these 50 consecutive timepoints spanned a temporal ROI of 200ms.

To assess FWER, a 200ms window (containing 50 time points) outside of the added effect was analysed in the same way as the data containing the effect. This window contained no added effect and therefore represented null effect only. A 200ms time window containing 50 sampled time points was chosen so that it equated to that used in the multivariate effect simulations.

Scripts to generate the simulated data and to analyse the results of each simulated experiment were written in MATLAB (MathWorks, 2023). MATLAB's Parallel Processing Toolbox (MathWorks, 2023) was used when running the simulations so that each core of the computer could execute the code to simulate each experiment in parallel to greatly reduce the processing time to practical levels. The generation of the multivariate effect as well as the MVPA classification itself were both conducted using functions from within the MVPA Light toolbox (Treder, 2020).

5.2.2 Simulations with multivariate differences between conditions

5.2.2.1 Generation of noise-only data

Each virtual participant's data contained 2 conditions with 50 trials per condition. These trials each had 64 channels and 512 time points. Because the sampling rate was 256Hz, the trial duration was therefore 2.00 seconds. This duration was chosen to allow time within a trial for approximately two full cycles of the lowest frequencies in the power spectrum of the data (1Hz). The sampling rate of 256Hz was chosen to maintain a practical processing timescale given the available computational resources. Simulated EEG noise-only time series data for each individual trial were generated using an approach following previous EEG simulation work (Yeung, Bogacz, Holroyd, & Cohen, 2004, Brooks, Zoumpoulaki, & Bowman, 2017) . This approach generated EEG noise time series data with a power spectrum approximating that of human EEG data. The generation process involved, for each trial, creating 50 sine waves at randomly chosen, without replacement, frequencies between 1-125Hz, with random phases. Each of the sine waves was scaled according to the relative amplitude of that frequency in the human EEG spectrum with the highest amplitude frequency having an amplitude of 1. The resulting noise was then multiplied by 10µV to increase the amplitude to a range congruent with that of normal human EEG (between -40 to 40µV). The noise in each trial was created independently of other trials (i.e., different set of random frequencies and phases). The noise was also generated independently across channels, participants, and the 2 conditions.

5.2.2.2 Addition of multivariate effects

One thousand experiments were simulated for each combination of participant number and effect size (see Figure 5.1). This number was chosen in line with previous work by Groppe et al., (2011a; 2011b) who performed similar simulations using ERP data. This meant that in total, 15 000 experiments were simulated (3 effect sizes x 5 sample sizes x 1000 experiments).

This set of simulations investigated statistical power, therefore the datasets needed to contain a difference between conditions so that we could estimate the probability of detecting this effect. To do this, we generated noise-only data following the procedure outlined in Section 5.2.2.1 above, and then added a multivariate difference between conditions to a 200ms time window (at 50 timepoints, between timepoints 100-150). This effect was generated according to the 3 sizes set out in Section 5.2.2: a small effect in terms of Cohen's d was produced by generating 2 multivariate Gaussian datasets with a difference between their centroid means of 0.318 (arbitrary units), a medium effect by a difference of 0.795, and a large effect by a difference of 1.272. This multivariate effect spanned 32 channels and was generated using MVPA Light's simulate_gaussian_data function (Treder, 2020).

5.2.2 Null hypothesis data

To generate null hypothesis data, a subset of 1000 out of the 15 000 total simulated experiments were used. The simulated experiments in the large effect/sample size of 80 virtual participants were arbitrarily chosen as this subset. The choice of experiments to use here is arbitrary because the way that the simulated background EEG noise was generated was identical in all 15 simulations. Although these 1000 experiments contained an added effect, this effect was added at the same 50 time points (between timepoints 100-150, 200ms duration) for every trial in every experiment (i.e., there was no jitter between the time points where the effect was added). Therefore, null hypothesis data was taken from a 50 time point window outside of this added effect (between time points 49-99, 200ms duration).

5.2.4 Analysis of effects

The same analysis pipeline was used for all simulated experiments. This ensured that there were no potential confounding effects of the type of filtering or dimensionality reduction (e.g., PCA; see Chapter 2 for details) on the performance of the classifier, no pre-processing of the data was carried out and only the raw simulated data was passed through the pipeline.

Firstly, classification was performed across the entire time course of each trial, such that data from 64 virtual EEG channels from every time point that was sampled (i.e., 256 samples per trial) was analysed independently. This classification was carried out using MVPA Light's mv_classify_across_time function (Treder, 2020). In line with current use in the field, and in order to maintain consistency between the empirical work described in Chapter 4 and in these simulations, a support vector machine classifier was used to perform this classification (SVM; e.g., List et al, 2017). For more detail on the appropriateness of SVMs for the classification of EEG data see Chapter 2, but for clarity these will be summarised here. Support vector machines are known to perform well in high dimensional spaces, such as those occupied by multi-channel EEG data (e.g., Orrù, Petterson-Yeo, Marquand, Sartori & Mechelli, 2012; Treder, 2020). They are effective in finding the optimal hyperplane that separates classes of data, even in such high dimensional spaces, and are less sensitive to overfitting (Treder, 2020). As is recommended, to avoid overfitting due to "double dipping" (Kriegeskorte et al., 2009), k-fold cross validation was used, with 5 folds. K-fold cross-validation is a technique that divides data into *k* distinct subsets, or "folds", with one-fold (out of the 5 in this case) being left out of the data used to train the classifier, and instead used to test the classifier's performance. The result of this step was a time series of decoding accuracy values.

Next, the decoding accuracy values at each time point within the ROIs were assessed in terms of their deviation from chance level (i.e., 50% here since there were 2 classes of data). This assessment was the same across both the effect-containing ROI and the noise-only ROI. To do this a cluster-based one-sample t-test identical to that used in Chapter 4 was conducted. This used cluster-based statistics to correct for the multiple comparisons made across the time course of the ROI. The decoding accuracy was compared to the chance level at each time point across the ROI, using a one-sample t-test, resulting in a t-value for each time point that represented the deviation of the decoding accuracy from chance. A threshold of 0.05 was applied to these t-values to identify contiguous time points where the t-values exceeded the threshold, forming clusters. For each identified cluster, the sum of the t-values within the cluster was calculated, forming the cluster t-statistic. To assess the significance of the observed clusters, permutation testing was performed by randomly shuffling the class labels and recalculating the cluster t-statistics 5000 times and using the maximum cluster t statistic (tmax) from each permutation to create a null distribution against which the observed cluster t-values were compared to determine their p-values. Clusters with p-values below 0.05 were considered statistically significant, representing time windows within the trial where the SVM classifier's decoding accuracy significantly differed from the chance level.

The statistical power of the analysis was determined by focusing on the ROI where an effect was intentionally introduced. The power analysis specifically investigated whether an effect was detected at any point within the known effective time window. Therefore, if there was just one point of significance at any of the time points within the 50 time point ROI, then the effect was deemed detected. The results from this approach were therefore a series of 1s and 0s according to the outcome from each simulated experiment. A result of 1 indicated that the experiment had yielded at least one point of significance at any of the time points within the ROI, and a result of 0 indicated that no significant time points had been detected. From this the percentage of detected effects could be calculated.

Finally, the FWER of the analysis was determined by focusing on the ROI outside of the added multivariate effect, in the subset of 1000 experiments from the large effect/sample size of 80 simulation. This was conducted by running the cluster-based permutation tests in the same way as described above, but on a 200ms, 50 timepoint window outside of the known added effect (between time points 49-99). The number of simulated experiments that yielded at least 1 significant result were then totalled and a percentage calculated, to produce the FWER for the procedure.

5.3 Results

5.3.1 Null data results

The distribution of resulting decoding accuracy values from the noise-only, null hypothesis data was firstly assessed for its normality using the Shapiro-Wilk test. This test revealed that these decoding accuracy values significantly deviated from a normal distribution, $W = 0.996$, df=1000, $p = .025$. The mean decoding accuracy value was 50.24% (min: 43.19%, max: 58.68%, SD: 0.019). The distribution of the decoding accuracies within the time point is shown in Figure 5.3, below. The skewness of the decoding accuracy values was calculated to be 0.168, suggesting that the distribution is very slightly skewed to the right, but is close to being symmetric. The kurtosis of the decoding accuracy values was 3.050, suggesting that the distribution is close to that of a normal distribution.

The same data were then assessed in terms of family-wise error rate (FWER). The number of experiments yielding at least one significant cluster were calculated. This resulted in a proportion of 0.011 experiments leading to at least one significant cluster, or in other words a FWER of 1.1%.

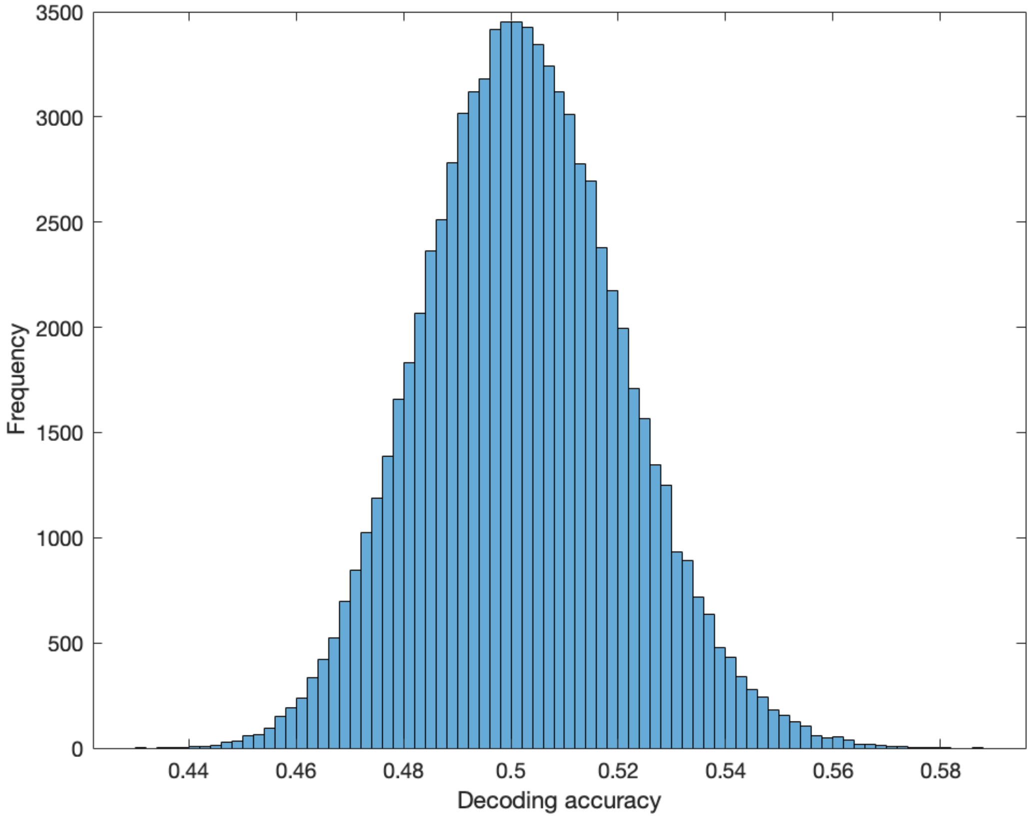

Figure 5.3: Histogram showing the distribution of decoding accuracy values from the noise-only ROI in 1000 experiments, each with 80 participants and 50 trials for each of 2 conditions.

5.3.2 Classifier performance

To examine whether the effect size manipulation had the intended effect (i.e., that increasing the difference between the centroid means meant that the patterns in each of the 2 conditions were more easily distinguishable by the classifier), we tested whether the decoding accuracy of the SVM classifier significantly increased with an increase in multivariate effect size. To achieve this, an assessment of the homogeneity of variances was firstly conducted. This was done via Levene’s test. Levene's test for equality of variances was significant, *F*(14, 14985) = 284.146, $p < .001$, indicating that the variances across the groups were not equal. Therefore, the effects of participant number and effect size on classifier performance were assessed via a two-way ANOVA with robust variance estimation. This was employed due to the violation of the assumption of homogeneity of variances, as indicated by a significant Levene's test. The robust variance estimation method used, HC3, specifically adjusts the standard errors of the ANOVA F-tests to account for the presence of unequal variances among groups (Hoaglin & Welsch, 1978). The HC3 estimator is a heteroscedasticity-consistent standard error estimator that offers an improved approach for dealing with variance inequality by applying a bias correction to the conventional standard errors. This adjustment is crucial for ensuring that the statistical inferences drawn from the ANOVA are robust and valid, even in scenarios where the data fail to meet the critical assumption of equal variances across groups. Employing the HC3 method thus provides a more reliable analysis framework under conditions where the homogeneity of variances assumption is compromised, as in this analysis.

This revealed a significant and large main effect of effect size, *F*(2, 15 000) = 52774.54, $p < .001$, $\eta_p^2 = 0.876$. Note that the residual degrees of freedom here are 15 000 given that there are 3 effect sizes x 5 sample sizes x 1000 experiments. No significant main effect of participant number on decoding accuracy was found, $p = .270$. The interaction between effect size and participant number was not significant, $p = .825$. See Figure 5.4.

Table 5.1: *Full robust ANOVA results showing the effects of participant number and effect size on SVM classifier performance*

	df1	*df2*	*F*	*p*	η_p^2
effect size	2	15 000	52774.54	< .001	0.876
participant number	4	15 000	1.290	.270	0.0003
effect size × participant number	8	14985.0	0.54	.825	0.0002

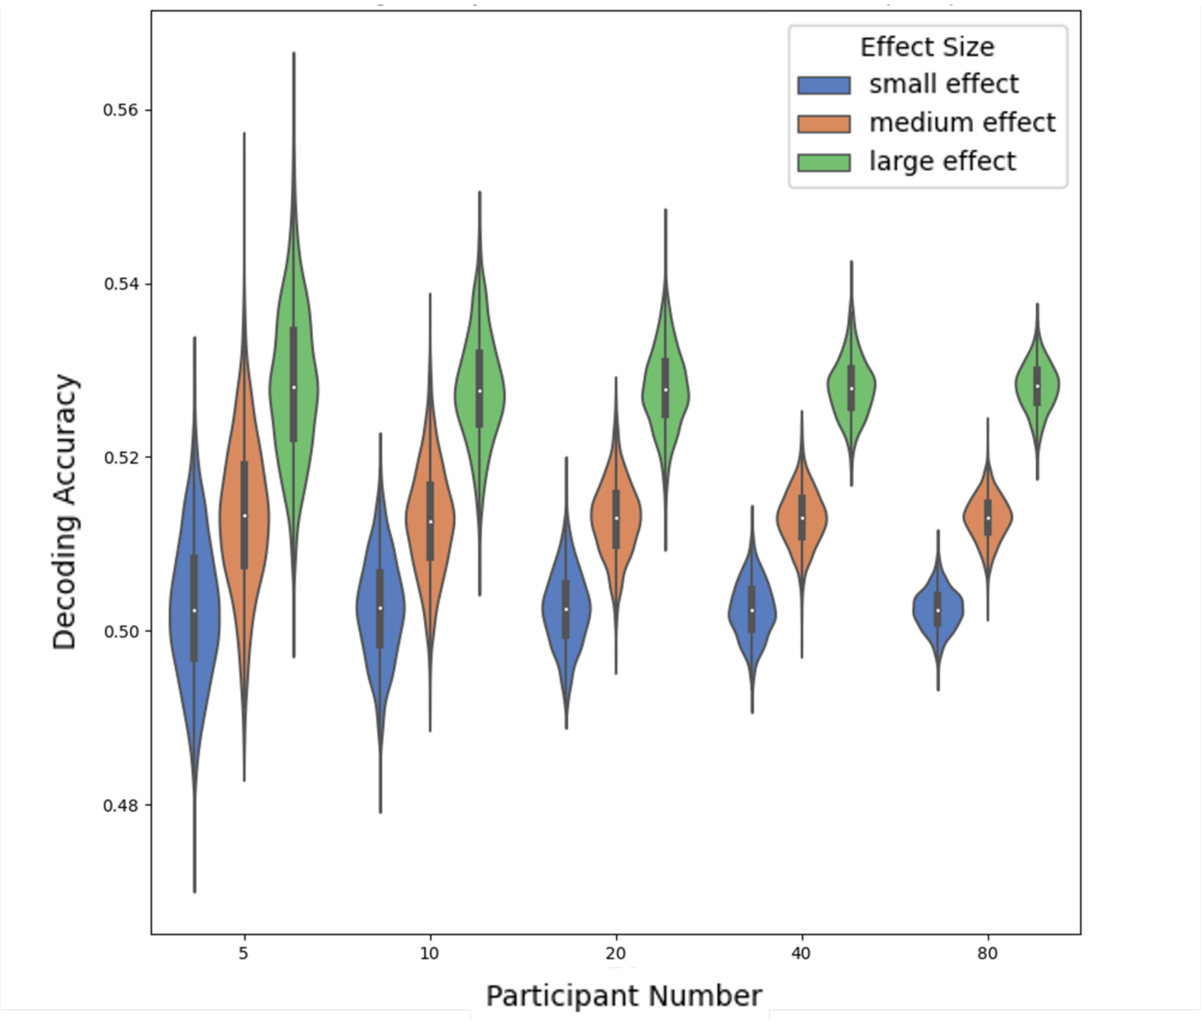

Figure 5.4: Violin plot showing the means and spread of the mean SVM classifier decoding accuracy values within the ROI across each combination of participant number and effect size.

5.3.3 Power

To determine whether our experimental manipulation of an increase in effect size and sample size had the expected effect of increasing power, Chi-squared tests were conducted. These were used due to the categorical nature of the data (i.e., 1s and 0s based on whether the experiment detected the effect or not) and the interest in assessing the independence between these categorical predictors (i.e., effect size and sample size) and binary outcomes. The analysis revealed a significant association between effect size and power; with power increasing significantly with increasing effect size, $\chi^2(1, N = 1000)= 2911.97$, $p < .001$. Similarly, a significant association was observed between sample size and power, with power increasing significantly with increasing participant number, $\chi^2(1, N = 1000) = 5468.72$, $p < .001$. See Figure 5.5.

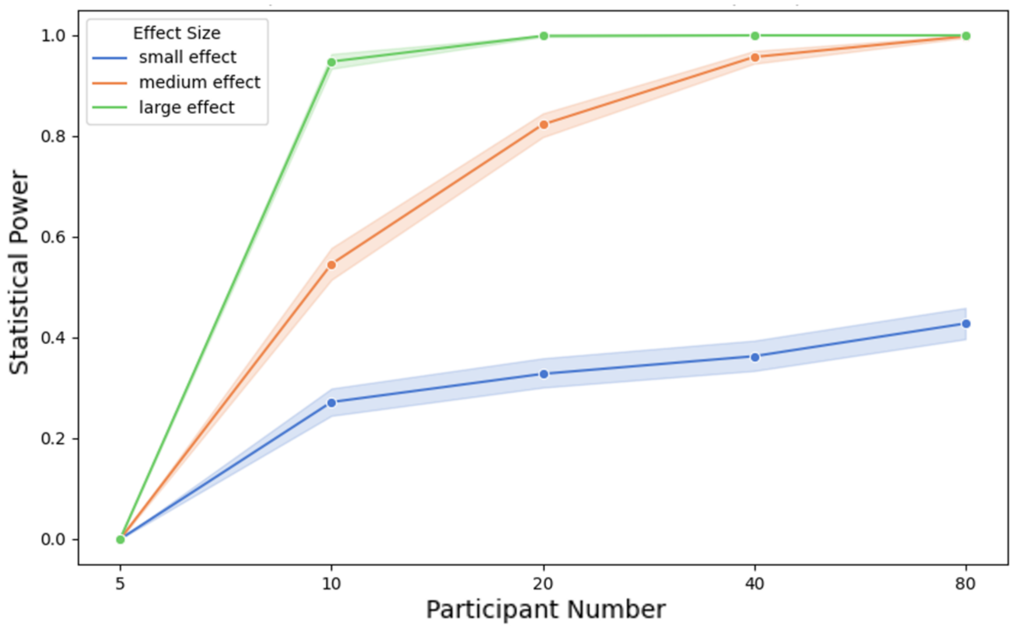

Figure 5.5: Line plot showing the means and standard error (represented by the shaded regions around each line) for the statistical power in each combination of participant number and effect size.

5.4 Discussion

The simulation results demonstrate that effect size has a significant impact on the decoding accuracy of the SVM classifier. Specifically, an increase in the size of the multivariate effect enhances the performance of the classifier. To assess power, cluster-based permutation tests were employed to pinpoint the specific timepoints within the ROI at which the classifier exceeded statistical chance levels. To assess the statistical power of this widely used statistical method in MVPA studies, the number of experiments where the classifier outperformed chance, as indicated by at least one significant cluster, was counted. The findings revealed that for a large effect, 80% statistical power can be achieved with as few as 10 participants. However, a medium effect requires 20 participants to reach the same level of power, and a small effect fails to reach 80% power even with 80 participants when employing cluster-based tests. In experiments involving only 5 participants, no effect at all was detected within the ROI (i.e., power was 0%). The simulations also showed that the FWER, when using cluster-based permutation tests in this way with 80 simulated participants, was 0.011.

Although the pattern of effect and sample size effects on the power results were in line with expectations, the fact that results with a sample size of 5 participants didn't lead to the detection of any effect (i.e., 0% power), was surprising. We would have expected results in line with the 5% alpha error rate. Cluster-based permutation tests have indeed been shown to be conservative when it comes to Type I error but have still led to around a 5% chance of false positives, especially with low sample sizes (e.g., Fields & Kuperberg, 2020). Although they have been shown to lead to a much lower chance of false positives with null data, zero chance has not been shown (e.g., Groppe et al., 2011b).

To further investigate these unexpected findings, additional post hoc analyses were performed. The unexpected result of a 0% FWER with a sample size of 5 is suggestive of an over-correction for multiple comparisons. Therefore, it was necessary to determine whether this issue was indeed with the correction step, or whether there was some issue with the data itself. If there was no issue with the data, we would expect that running uncorrected t-tests on an expanding null effect time window would lead to the expected increase in false positives (i.e., if just one time point was compared then the false positive rate would be low, but if 20 time points were compared then the false positive rate would be high as multiple comparisons have been made without correction). To investigate this, the null effect 200ms time window outside the ROI, ranging from time points 49-99 (as detailed in Section 5.3.1), was examined. Since Type I error rates are cumulative, the inclusion of additional timepoints increases the probability of encountering a false positive (i.e., the reason it is necessary to correct for multiple comparisons at all). To explore this, the null data were examined over an expanding time window, thereby increasing the number of comparisons. Initially, uncorrected t-tests were used to assess if the observed decoding accuracies were significantly greater than 50% for a single timepoint (one single timepoint at all experiments). This was then extended to 2, 10, 20, 30, 40 and then 50 timepoints. The number of participants was also increased according to the original design (i.e., 5, 10, 20, 40 and 80 participants).

By increasing the number of comparisons without applying corrections, the risk of incorrectly identifying a timepoint as being greater than chance correspondingly escalated (see Figure 5.7). This resulted in an expected cumulative rise in the Type I error rate, with the lowest recorded at 5% with 1 time point and the highest at 95% at 50 time points.

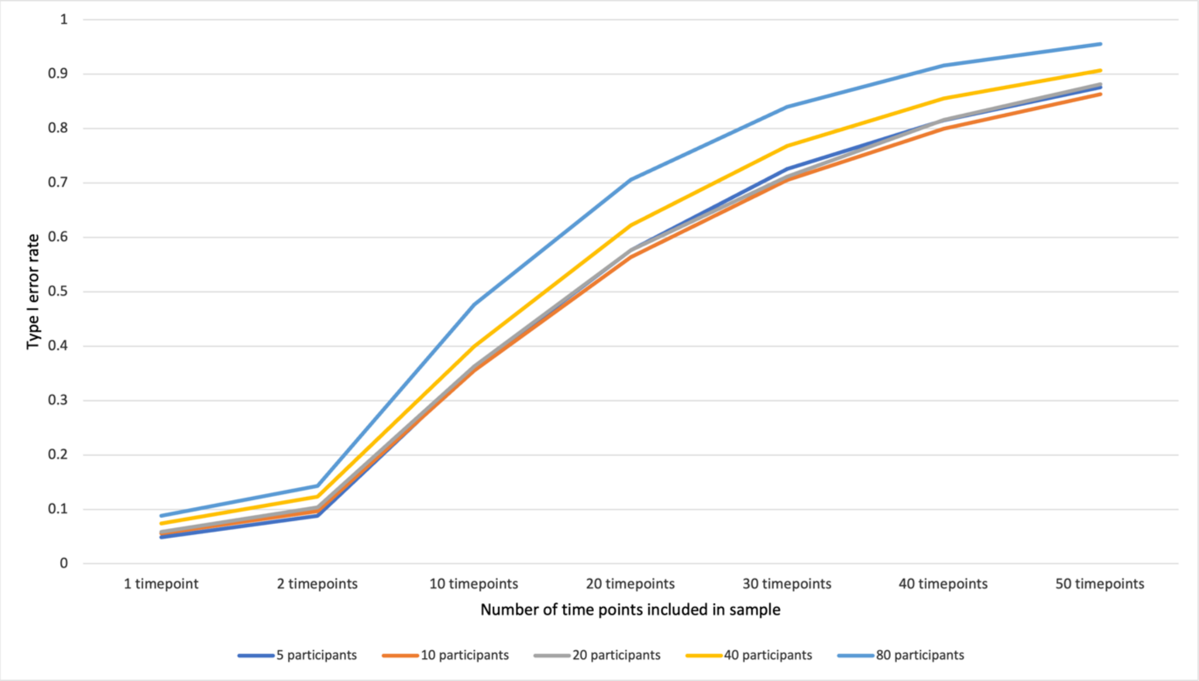

Figure 5.6: Line plot showing the mean number of Type I errors resulting from the uncorrected t-tests ran on each combination of sample size and number of time points included in the sample.

However, when cluster-based corrections were used across the same 50 sample, 200ms time window (including all 50 time points), the FWER fell to 0 when the sample size was 5. When the sample size was increased to 80 participants, the highest recorded Type I error rate was 0.011 or 1.1% (see Figure 5.7). The surprising thing here is that it was expected that the FWER would fall with an increasing number of participants (e.g., Fields & Kuperberg, 2020) but this was not the case, although FWER was far below the accepted 5% for all cases.

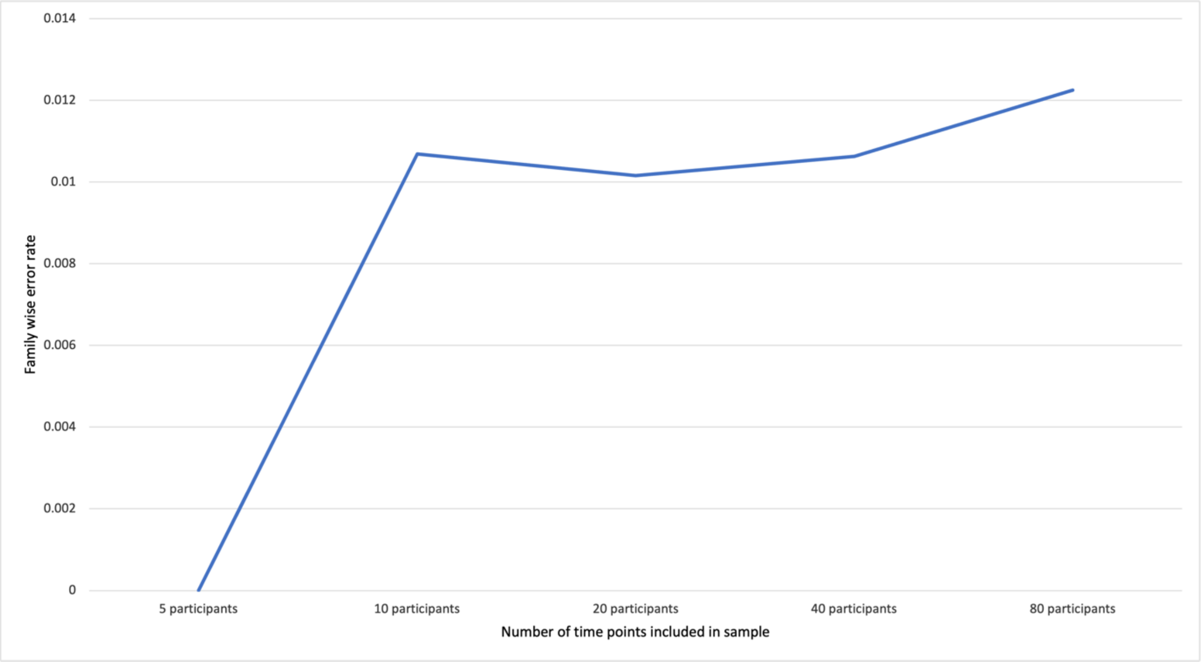

Figure 5.7: Line plot showing the mean family wise error rate (FWER) of null effect data resulting from the cluster-based permutation tests ran with each sample size.

So, while the use of uncorrected t-tests led to a cumulative increase in Type I error rates, the implementation of cluster-based corrections markedly reduced these errors, with them falling to 0 when the sample size was low.

When examining the results in light of Groppe et al. (2011a), it is suggested by these authors that cluster-based tests offer weak control of FWER when applied to mass univariate testing of ERP data. However, when these tests are applied within MVPA pipelines, the simulation results here imply otherwise and that the approach may be overly conservative.

Another important finding from the results of this simulation is regarding the necessary sample size for achieving robust statistical power. Specifically, the results indicate that to achieve a power of 80% when dealing with medium or large effects using this MVPA pipeline, a sample size of at least 20 participants is required. This is a significant point to consider, especially when planning experiments or interpreting results, as an inadequate sample size can greatly reduce the reliability and generalisability of findings.

Moreover, there are inherent challenges when dealing with smaller effect sizes. The simulation suggests that even when employing a relatively large sample of 80 participants, there remains a considerable likelihood, around 60%, of missing such small effects. This underscores the nuances and complexities associated with detecting subtle changes or differences in data and warrants caution when interpreting null results or when the anticipated effect size is modest.

Current methodologies and interpretations of MVPA operate under a crucial assumption: that if data adequately represents information about a given condition, then the decoding accuracy should reliably exceed chance levels. However, the findings here highlight what is a perhaps overlooked aspect of MVPA of EEG data - the significance of statistical power. While much of the existing literature has centred around minimising Type I error, there's a growing necessity to address issues of power. Studies that neglect to ensure adequate power not only risk producing unreliable results but also limit the replicability and broader implications of their findings.

It is important to note that, in the conducted study, simulated data was used, with the effect consistently introduced at specific time points. Despite this controlled approach, when comparing identical time points across different conditions, the effect was undetectable with a sample of just 5 participants. Therefore, if the effects were not fixed, but rather fluctuated in its timing, one might expect that any attempt to detect such an effect with low sample sizes might become even more difficult when using standard classification across time. Although, this is speculation at this stage and requires further work to establish the claim. This does, however, help to highlight the need for techniques such as temporal generalisation to be employed where appropriate (see Chapters 2 & 4; King, 2014). Because temporal generalisation involves training a classifier at each time point, for example at 1ms, and then testing that classifier on all other time points in the trial (for example at each time point from 1-100ms), it is a suitable technique to use in detecting effects in neuroimaging data that are not strictly aligned in time (King & Dehaene, 2014). This technique has been shown to be particularly effective in its ability to accommodate temporal variability across trials and individuals (King & Dehaene, 2014). By allowing for the analysis of how neural representations evolve over time, temporal generalisation enhances the understanding of cognitive processes beyond what is possible with static analyses such as the standard measure of decoding accuracy time series data. Future work, therefore, would be beneficial in simulating experiments with added multivariate effects placed at jittered temporal locations across participants or trials and then using techniques like temporal generalisation to determine power under these conditions.

In this chapter we ran 15 simulations to assess the statistical power of MVPA pipelines in the context of decoding patterns of EEG data. We varied the size of the multivariate effect added to simulated noise-only EEG data, alongside varying the sample size used in each simulated experiment. We found that an increase in sample size and effect size led to an increase in statistical power, as we hypothesised. In this chapter we also examined FWER of the same analysis pipeline and found FWER to be conservatively controlled using cluster-based corrections. We showed that this conservatism was potentially over-corrective when sample size is low.

5.4.1 Interim Discussion

The code used to run the simulations described in this chapter was designed to be generalisable to future experiments. However, this generalisability meant that it was more difficult to pin down any issues in the code that could have been causing the unexpected results produced above. To investigate this further, the code was simplified extensively such that it was only applicable to the simulations described in this chapter. This investigation led to the discovery of an error in the lowest level of the script.

When initially investigating the unexpected result that, with 5 participants, there was a Type I error rate of 0%, we hypothesised that the cluster-based statistics may have been conducted on the entire epoch rather than just on the ROI. We had previously explored this in the code and concluded that this was not the problem. However, after simplifying the code and re-investigating this issue, we discovered this was indeed what was happening; the cluster-based statistics were performed on the entire 256 sample epoch. This meant that the cluster correction would be for 256 samples rather than the smaller 50 sample time window comprising the ROI. Thus, more samples in the cluster test meant a stronger correction and one that would be far too strong for the smaller ROI of 50 samples.

Furthermore, given that the cluster-based statistics were being conducted across the entire 256 sample epoch, this meant that the cluster correction was also influenced by data in the ROI that had a true effect inserted. This would cause the cluster t distribution to be wider and thus increase the t-threshold for a significant effect, adding to the already increased conservativeness of the test due to the extended time window included in the correction procedure.

In sum, in the original code, this error led to a twofold effect that caused the cluster-based statistics to be far more conservative than necessary. This effect was strongest in the 5-participant condition in which power was weakest and there were fewer possible permutations of the data. Therefore, we corrected this error such that the cluster-based statistics were conducted only on the 50 timepoint sample, rather than the 256 timepoint sample.

Secondly, we discovered a further error that was not necessarily adding to the unexpected results but was still an issue for the results overall. This was that, in calculating the standard deviation of the background noise for use in the effect size calculations, the noise was erroneously averaged across the 50 timepoint window before the standard deviation was calculated. Therefore, this was also corrected and the new value for the within-condition noise was used to calculate the necessary difference in centroid means to generate the 3 effect sizes required.

The correct value for the within-condition noise was re-calculated at 6.3097 µV. Therefore, when this was used in the same calculations as those described in section 5.2.2 above, the following differences in centroid means were calculated to generate each of the 3 effect sizes (all in arbitrary units):

Small Effect = 1.2619

Medium Effect = 3.1548

Large Effect = 5.0477

5.4.2 Results

5.4.2.1 Null effect data

One thousand experiments for each sample size were simulated that contained no added multivariate effect, and therefore contained noise-only, null data.

The null effect data were assessed in terms of family-wise error rate (FWER). The number of experiments yielding at least one significant cluster were calculated for each sample size (see Figure 5.8, below). The average FWER across all sample sizes was 0.047, or 4.7%.

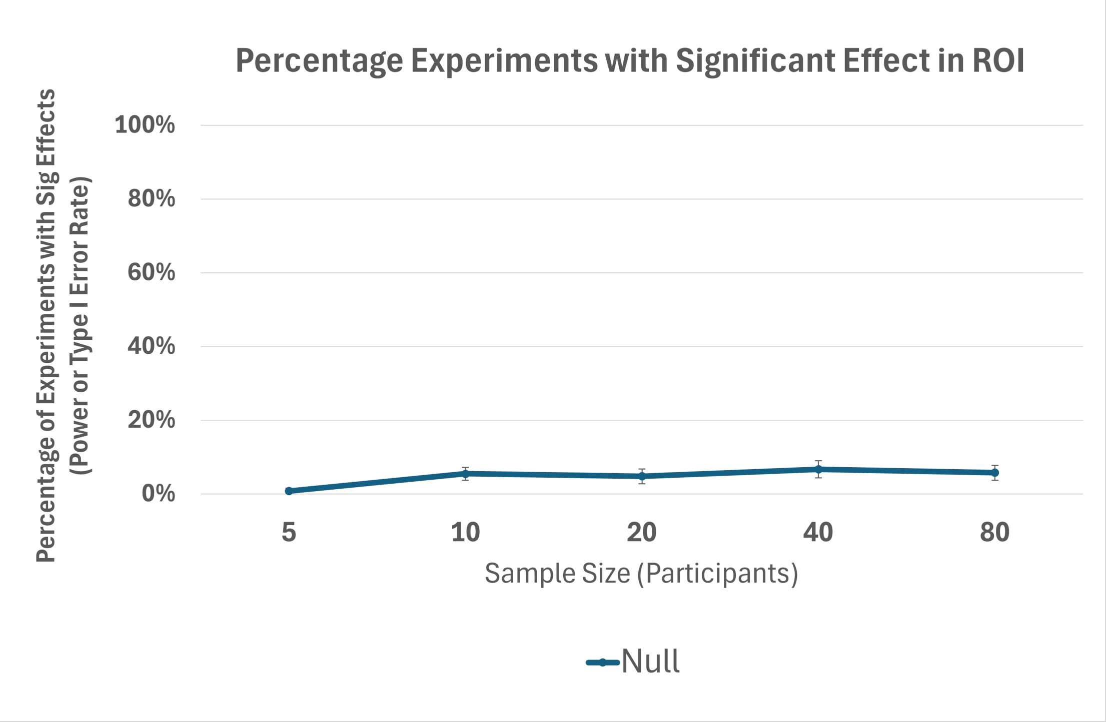

5.4.2.1 Classifier performance

To examine whether the effect size manipulation had the intended effect, we tested whether the decoding accuracy of the SVM classifier significantly increased with an increase in multivariate effect size. As we had no a priori reason to believe that participant number would influence decoding accuracy, the results were collapsed across participant number, and we looked only at the effect of effect size in a one-way ANOVA. To achieve this, an assessment of the homogeneity of variances was firstly conducted. This was done via Levene's test. Levene's test for equality of variances was significant, $F(14, 14985) = 100.970$, $p < .001$, indicating that the variances across the groups were not equal. Therefore, the effect of effect size, collapsed across each level of sample size, on classifier performance was assessed via a one-way ANOVA with robust variance estimation.

This revealed a significant main effect of effect size, $F(2, 14997) = 7046.61$, $p < .001$, $\eta_p^2 = 0.484$, see Figure 5.9.

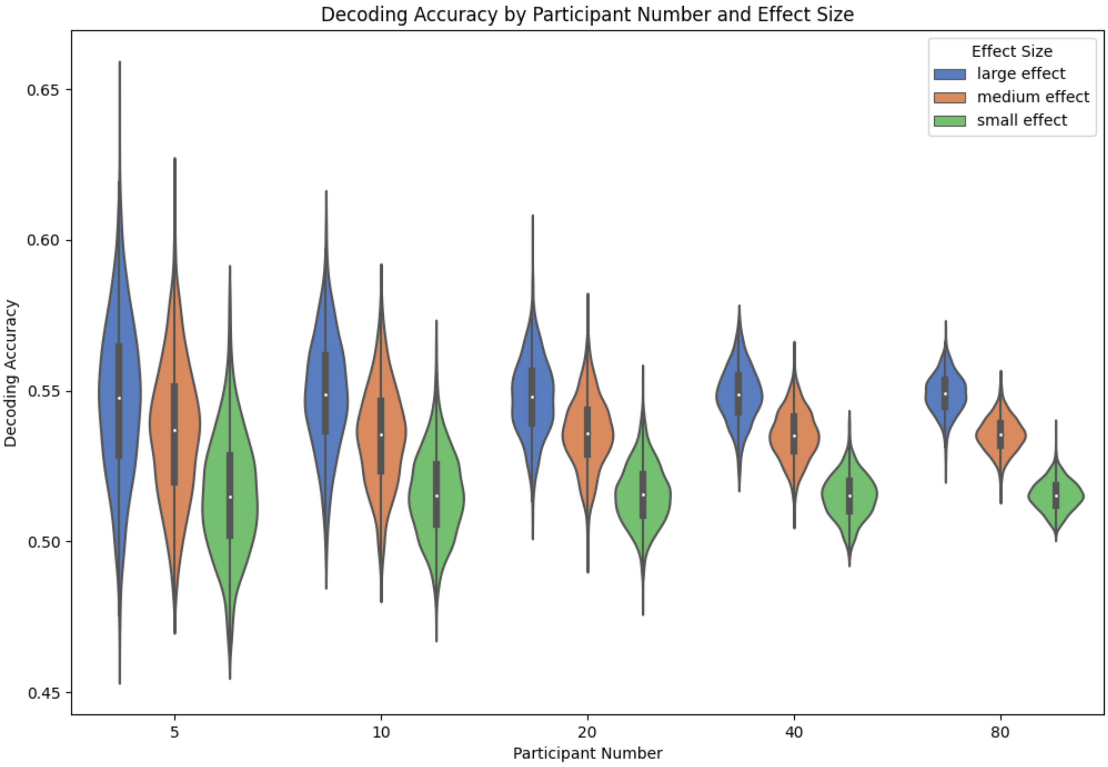

Figure 5.9: Violin plot showing the means and spread of the mean SVM classifier decoding accuracy values within the ROI across each combination of participant number and effect size.

5.4.2.2 Power

To determine whether our experimental manipulation of an increase in effect size and sample size had the expected effect of increasing power, Chi-squared tests were conducted. The analysis revealed a significant association between effect size and power; with power increasing significantly with increasing effect size, $\chi^2(2, N = 15000)= 1521.19$, $p < .001$. Similarly, a significant association was observed between sample size and power, with power increasing significantly with increasing participant number, $\chi^2(4, N = 15000) = 4996.91$, $p < .001$. See Figure 5.10, below.

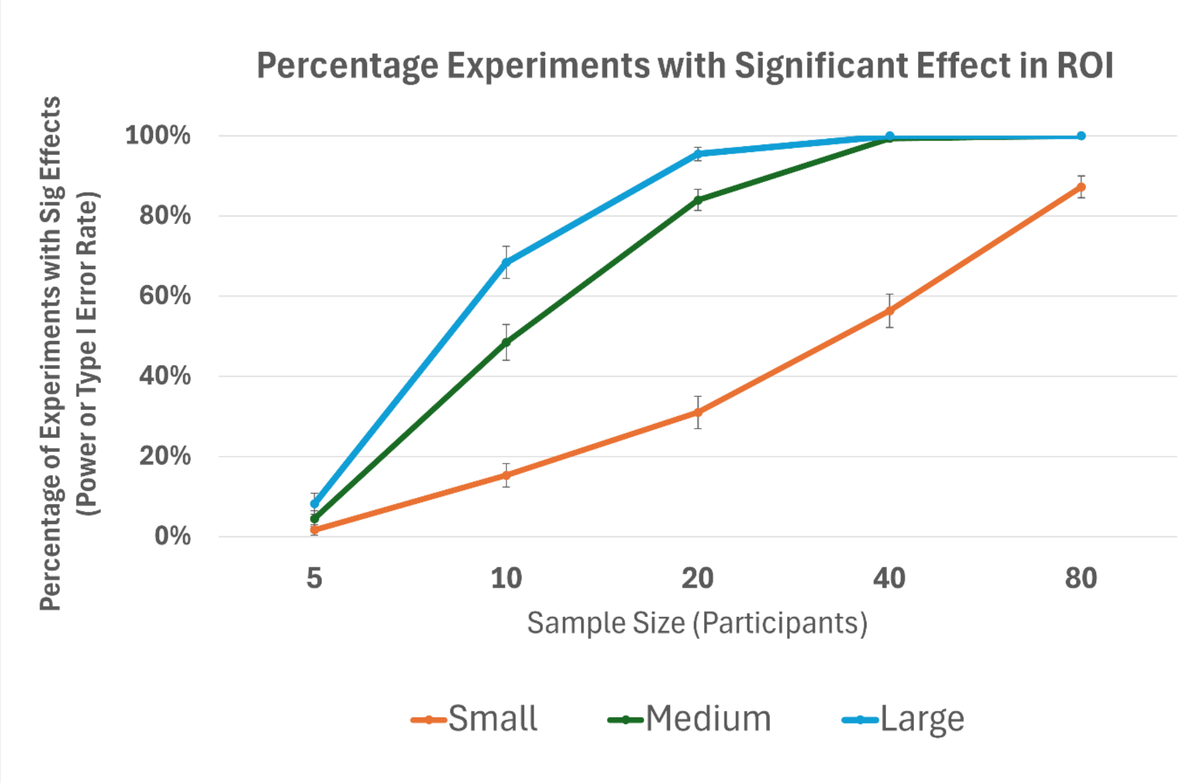

Figure 5.10: Line plot showing the means and bootstrapped 95% confidence intervals (represented by the bars around each point) for the statistical power in each combination of participant number and effect size.

5.4.3 Discussion

The corrected simulation results broadly agree with the conclusions drawn in section 5.4 above. The corrected results again demonstrate that effect size has a significant impact on the decoding accuracy of the SVM classifier. Specifically, as expected, an increase in the size of the multivariate effect enhances the performance of the classifier. To assess power, cluster-based permutation tests were again employed to pinpoint the specific timepoints within the ROI at which the classifier exceeded statistical chance levels. The number of experiments containing at least one significant cluster within the ROI was counted to assess statistical power. This revealed that, for both large and medium effects, 80% statistical power requires at least 20 participants, whilst for a small effect, 80% power can only be achieved with at least 80 participants. For null effect data, the FWER was calculated at 4.7%, which is in line with the accepted 5% alpha error rate. Overall, these results are in line with those presented previously and in line with expectations. These results provide specific guidance about sample sizes for studies using MVPA and cluster-based statistical approaches on EEG data.

Chapter Six: General discussion

6.1 Overview

This thesis set out to examine the role of adaptation in driving spontaneously occurring perceptual reversals. Specifically, we were interested in whether adaptation, according to the neural fatigue hypothesis, played a role in the spontaneously occurring reversals of ambiguous motion created from drifting sine wave gratings. The first chapter of this thesis examines the background literature on multi-stable perception in general, as well as studies that have examined the behavioural effects of adaptation and priming via pre-exposure to unambiguous stimuli. Particularly, the literature suggests that behavioural adaptation and priming effects can be induced rapidly when participants are pre-exposed to directional motion before they view ambiguous motion. The induction of one or the other of these effects is governed by the relationship between the pre-exposure duration and interstimulus interval (ISI) between the presentation of directional and ambiguous motion stimuli. In summary, the literature strongly suggests that reversals in perception can be induced using this pre-exposure paradigm, but it is not clear what role, if any, adaptation plays in reversals in perception that occur spontaneously, without prior exposure to unambiguous stimuli.

In this thesis, this issue has been directly addressed using multivariate pattern analysis (MVPA) cross-decoding techniques of electroencephalogram (EEG) data. Chapter 3 confirmed that behavioural adaptation could indeed be induced rapidly as suggested by the background literature, but shed doubt on whether the same technique could be used to induce visual motion priming as the literature claims. The second experiment in Chapter 3 investigated this doubt, whilst the timing parameters to successfully induce adaptation established from this chapter were used to inform the EEG experiment in Chapter 4. In the EEG experiment in Chapter 4, a cross-decoding technique was employed to draw more direct conclusions around whether the pattern of brain activity during reversals driven by adaptation were similar to the pattern of brain activity during spontaneously occurring reversals. Temporal generalisation techniques were also employed to draw conclusions around the neural processes underpinning perceptual reversals. Given that the use of MVPA of EEG data has not been as thoroughly scrutinised in the existing literature, when compared to more standard univariate analyses such as the event-related potential (ERP), Chapter 5 sought to clarify the power and error rates associated with the use of MVPA of EEG data. This chapter provides much needed guidance on the sample sizes required to detect a given effect when using the analysis pipelines employed in this thesis and the wider literature in this area.

The following sections in this chapter will summarise the findings of the experiments described in this thesis, to address the five questions posed in Chapter 1. These were: 1) What are the optimum timing parameters required to induce adapted and primed brain states? 2) Can the task that participants are asked to perform lead to artefactual VMP effects? 3) Does adaptation play a role in spontaneous reversals in perception that occur during viewing of ambiguous motion? 4) Are ERP measures of reversal-related brain activity present during perceptual reversals of ambiguous stationary stimuli also present during perceptual reversals of ambiguous motion stimuli? And 5) How powerful is MVPA to detect multivariate effects present in EEG data? For each of these questions, these findings will be interpreted in relation to the wider literature and existing theories of the driving forces behind perceptual reversals, Limitations of the experiments used in this thesis will be discussed, as well as recommendations for future research based on the results shown.

6.2 Summary of results and discussion

6.2.1 What are the optimum timing parameters required to induce adapted and primed brain states?

In Chapter 3 we aimed to investigate the claim made in the existing literature, that behavioural adaptation and priming effects can be induced rapidly (in the order of milliseconds) when participants are pre-exposed to directional motion before viewing ambiguous motion. Specifically, we aimed to investigate the visual motion priming (VMP) and rapid motion aftereffects (rMAE/adaptation) reported by Kanai and Verstraten (2005). Literature suggests that the induction of one or the other of these effects requires the relationship between the pre-exposure duration and ISI between the presentation of directional and ambiguous motion stimuli to be tightly controlled (e.g., Pinkus & Pantle, 1997; Kanai & Verstraten, 2005). VMP is typically induced by shorter pre-exposure durations, whilst the rMAE is induced by longer ones, and the effects of extending the ISI differ depending on the pre-exposure duration (Kanai & Verstraten, 2005). Therefore, this aim was addressed by systematically varying the duration of pre-exposure to directional motion, and the ISI between the directional and ambiguous motion. We then recorded the number of VMP and rMAE responses in each of the combinations of pre-exposure duration and ISI. A VMP response was one in which participants reported the two stimuli moved in the same direction, in other words a stable trial. In contrast, a rMAE response was one in which participants reported the two stimuli moving in opposing directions, in other words a reversal trial. This allowed us not only to establish whether the

VMP and rMAE effects could indeed be induced as quickly as the literature suggests, but also to determine the optimum combination of pre-exposure duration and ISI to induce these two effects. This was because we aimed to use these parameters in the EEG experiment in Chapter 4. The aims of this EEG experiment meant that the optimum pre-exposure durations to induce the rMAE and the VMP within the same ISI needed to be established, as well as determining which, if any, duration of pre-exposure to directional motion led to participant's perception of the ambiguous stimulus being unbiased, also within that same ISI (i.e., 50% stable vs reversal responses).

Our results broadly agreed with those from Kanai and Verstraten (2005). VMP was elicited when directional motion was briefly presented and followed by a short ISI, and the rMAE was induced when directional motion was presented for a longer time. We found that an ISI of 120ms best separated the rMAE and VMP responses. Within this ISI, pre-exposing participants to directional motion for 640ms before presenting them with ambiguous motion led to participants responding to indicate the rMAE. Responses indicating an unbiased perception were also found, following directional motion pre-exposure of 80ms. However, behavioural responses indicating VMP were not found within the 120ms ISI as they were in the original study by Kanai and Verstraten (2005). We also found a much weaker VMP effect overall.

This finding was surprising, given the strength of the VMP response reported by Kanai and Verstraten (2005) as well as other authors (e.g., Pinkus & Pantle, 1997; Takeuchi et al., 2011). Our experiment differed in only one aspect compared to Kanai and Verstraten (2005): the way that participants were asked to respond. Kanai and Verstraten (2005) asked their participants to respond according to whether the two stimuli in the trial were moving in the “same” or “different” directions. “Same” responses were counted as indicating VMP, whilst “different” responses indicative of the rMAE. In our experiment, participants were asked to report the direction of motion of each of the two stimuli in the trial. Example responses in our experiment therefore include “left”, “right” which would be counted as a response indicating the rMAE, and “left”, “left”, which would be counted as indicating VMP. Therefore, our participants, but not Kanai and Verstraten’s (2005), were forced to consider the direction of each motion stimulus in the pair independently as well as relative to one another. Anecdotal reports from our participants as well as experimenter observations, suggested that at the combinations of pre-exposure duration and ISI that were suggested to evoke VMP (i.e., the briefest) could be unintentionally leading to the perception of just one stimulus rather than the intended two. If this were the case, participants in Kanai and Verstraten’s (2005) study could have been biased to respond “same” more often in these conditions simply because they did not notice that there were two stimuli to compare, given the forced choice between “same” and “different”. On average, one would have expected this to artificially increase the number of “same” responses in conditions where this happened and potentially lead to the strong VMP results shown by Kanai and Verstraten (2005).

In our experiment however, participants were forced to consider and report the motion of a directional stimulus that they may not have perceived. This could feasibly lead to participants guessing about the direction of that stimulus. In this case, participants might assume that they missed the first stimulus altogether and would be forced to randomly guess its direction because two directions were required in our response. If this were the case then on average, it would artificially reduce the number of VMP responses because guessing, if random, is equally likely to be the same or different from the ambiguous stimulus perception. This would then push the percentage of same responses towards 50%. Because our response method allowed us to examine whether participants had reported the correct direction of the directional stimulus, we examined the accuracy of responses to the directional motion. We found that the briefest durations (i.e., those where VMP was most strongly reported by Kanai & Verstraten, 2005) led to participants' accuracy being at chance levels, which supported the notion that participants could simply be guessing when reporting their perception of the directional motion.

This aim was addressed in Experiment 1, Chapter 3. The results here demonstrate that rapid adaptation to directional motion can indeed be induced using the paradigm employed by Kanai and Verstraten (2005). The optimum parameters to induce this rMAE involve pre-exposure to directional motion for 640ms and an ISI of 120ms. However, the results from this experiment threw into doubt whether VMP could be evoked by this paradigm in the same way. VMP has been shown in previous literature to be eradicated in response to low luminance levels and increases in motion speed (Takeuchi et al., 2011). Therefore, it could be the case that this paradigm only appears to induce it because of the artefactual responses mentioned above. Therefore, whilst the paradigm was not used to investigate VMP in the subsequent EEG experiment, it was necessary to investigate these anecdotal claims that were being made about the paradigm empirically.

6.2.2 Can the task that participants are asked to perform lead to artefactual VMP effects?

Informal observations of the Experiment 1 stimuli with short ISIs and short directional stimulus durations suggested that participants may not always be able to resolve the two separate stimuli that are present in these displays. This creates an ambiguity in the interpretation of the results from both Experiment 1 and the wider literature. This ambiguity was addressed in the second experiment in Chapter 3. Based on the results and observations from Experiment 1, we formed the hypothesis that the task that participants were asked to complete could be influencing how they responded to indicate their perception of a directional motion stimulus that they may not have perceived at all. Kanai and Verstraten (2005) asked participants to report only whether the two stimuli were moving in the same or different directions. This meant that it would have been possible for participants in Kanai and Verstraten's (2005) study to report that they saw motion in the same direction simply because they had only perceived one stimulus. Intuitively, this seems like it would be particularly true when very brief stimuli are presented or a very short ISI is used, because under these conditions the end of one stimulus and the beginning of another might be difficult to distinguish. Therefore, it was hypothesised that differences in this response requirement could be leading to artefactually inflated VMP responses in the study by Kanai and Verstraten (2005) and random responses in Experiment 1.

Therefore, Experiment 2 directly manipulated the way that participants were asked to respond to the pairs of motion stimuli whilst keeping the stimuli themselves constant across the different response requirements (i.e., tasks). This was done to draw more direct conclusions around participants' perception compared to the way that they were responding. Specifically, task was manipulated across four blocks which differed only in their response requirements. The first block type was called the ‘replication block’. This block required participants to respond in the same way as in Experiment 1, by answering the question “please report the direction of motion of each of the two stimuli”. The purpose of this block was to make sure that the results of Experiment 1 were replicable, given that an independent set of participants completed Experiment 2. The second block type was termed the ‘counting block’. This block required participants to respond to the question “how many stimuli did you see in the trial?”. The purpose of this block was to directly assess whether participants were more likely to perceive the stimulus pair at the shortest pre-exposure time and ISI as one stimulus rather than two, as we hypothesised. The third block type was what we called the ‘dual task block’. In this block, participants were instructed to respond to indicate how many stimuli they saw, and in which direction they saw them moving. This meant that they were without the restraint of being told to focus on one element of the task. For instance, if they were being asked about two stimuli (either by being asked to report the direction of each of them or to compare their motion relative to one another) but were only perceiving one, this block allowed them to simply report what they saw. The final block was termed the ‘final perception block’. This block was a partial replication of the task used by Takeuchi et al (2011). This was titled the “Final Direction” block for participants. In this block, participants were instructed to report only the direction of their perception of motion that they were experiencing right at the end of the trial (i.e., the motion of the ambiguous stimulus). This was then compared with the known direction of the directional stimulus to determine whether the response was indicative of VMP (i.e., the ambiguous motion was perceived as moving in the same direction as the directional stimulus) or rMAE (i.e., the ambiguous motion was perceived as moving in the opposing direction to the directional stimulus). The purpose of this block was to get an estimate of VMP which should be less sensitive to whether participants perceive both stimuli. If VMP is taking place, then it should be observable in the reports of the second stimulus regardless of whether the first stimulus is perceived.

Our results from the replication block suggest that Experiment 2 was successful in replicating the pattern of results from Experiment 1 when ISI and pre-exposure duration were brief. There were a greater number of VMP responses to ambiguous motion stimuli when presented following brief pre-exposure to unambiguous motion after a shorter ISI.

The results from the counting block clearly showed that participants were more likely to report perceiving only one stimulus when directional stimulus presentation and ISI were brief. When the ISI was increased, the percentage of "one stimulus" responses dramatically fell. The pre-exposure duration of the directional stimulus also had an effect with shorter directional stimuli leading to more "one stimulus" responses. In the condition with both the shortest ISI and the shortest directional stimulus duration, most trials were reported as containing only one stimulus rather than two. These results are in line with the hypothesis that participants are largely unable to resolve two separate stimuli in displays with short ISIs and short duration directional stimuli. This result clearly undermines direct interpretation of the Experiment 1 data as being indicative of a weak VMP effect compared to Kanai and Verstraten (2005) and instead provides support for the hypothesis that participants are making their best guess about a directional motion stimulus they did not perceive.

These results were further supported by the results of the dual task block. Overall, participants were more likely to report seeing one stimulus when ISI was brief. This effect was particularly prominent when coupled with shorter directional stimulus presentation times. When the ISI increased, and participants responded to indicate that they saw two stimuli, they were more likely to respond to indicate VMP at shorter presentation times and rMAE at longer ones in a pattern similar to that seen in the replication block. As there were always two stimuli presented in the dual task block, any erroneous “one stimulus” responses were removed to get a purer estimate of the amount of any VMP that was present. This showed that there were no combinations of ISI and directional stimulus presentation time that led to responses indicating VMP. In this case, even short presentation times and ISIs led to responses indicative of rMAE rather than VMP.

Results from the final perception block showed that participants were more likely than chance to report VMP when the ISI and pre-exposure durations were brief. Although, this effect was much weaker than in both our replication block and in the authors’ original study. We took this block to be a purer measure of any VMP effect given that participants were not asked about the direction of the first, unambiguous stimulus (that we are hypothesising they did not perceive). Therefore, any VMP responses that were seen in this block should not have been due to the artefactual “best guess” effect described above. So, the fact that there was a much weaker VMP response in this block compared to our replication block could support the idea that participants were guessing under brief timing conditions in Experiment 1. This is because, one potential strategy participants could employ if they are uncertain about the direction of the first stimulus is simply to report that it was moving in the same direction as the second. Thus, a weaker VMP effect is seen in the final perception block because this strategy cannot be employed. Instead, the results from this block could indicate that VMP is only weakly evoked by this paradigm (i.e., pre-exposure to directional motion followed by ambiguous motion).

One possible explanation for the results that are seen in the counting and dual task blocks is that the participants are again responding with their 'best guess' under these very brief timing conditions. The results from the correct responses to directional motion from Experiment 1 and the replication block in Experiment 2 make this assumption more likely. In both cases, participants were significantly more able to correctly respond to the unambiguous directional motion stimulus (i.e., correctly report that it was moving to the left if indeed it was) when the pre-exposure duration and ISI were longer but were less able to do so when they were shorter. This result indicates that under these circumstances, participants may be unsure about the direction of motion that they see in the trials and therefore the responses that they make on them may be based on a guess. Indeed, the accuracy of responses to the first, directional stimulus in Experiment 1 did not differ from chance level at the briefest pre-exposure durations. This strongly suggests that participants were guessing about its direction because they did not perceive it.

One possible counterargument here is that participants might not have consciously perceived the directional stimulus, but it could still have had a priming effect on their perception of the ambiguous stimulus. This could well be the case, as results from the final perception block do indeed show an indication of VMP under brief conditions. However, in our partial replication here that VMP effect was far weaker than that found by the original authors (Takeuchi et al., 2011). It should be noted, however, that the exact viewing conditions used in our partial replication did not mirror those used by the original authors. As priming effects have been shown to be very sensitive to changes in viewing conditions, this could have led to the discrepancies between our results and theirs.

This aim then, was addressed by Experiment 2 in Chapter 3. Our results suggest that indeed, task can lead to artefactual VMP effects. These results suggest that caution should be taken when interpreting the results of experiments where participants are forced to make a judgement about a stimulus that they might not be able to perceive, leading to the "best guess" effect described above. The recommendation here would be to replicate studies that force participants to make these sort of judgements, using a paradigm whereby the reported direction of the ambiguous stimulus is compared directly to the known motion of the directional stimulus in order to examine any true effects present.

6.2.3 Does adaptation play a role in spontaneous reversals in perception that occur during viewing of ambiguous motion?

Traditional, univariate analyses do not consider the multivariate nature of EEG data. In event-related potential (ERP) analyses for instance, differences between experimental manipulations are often evaluated via a comparison of amplitude means generated from the a priori selection of a spatio-temporal region of interest (ROI). This means that effects outside of these ROIs may be missed. Also, some important signals that differentiate between conditions or brain states may not be apparent as large differences at isolated locations on the brain/scalp but rather may be present as spatially-distributed patterns of activity. In contrast, multivariate approaches allow the holistic pattern of brain activity to be considered, leading to these approaches being more sensitive to detect effects that may be missed by univariate techniques (e.g., List et al., 2017; Bae & Luck, 2019). In Chapter 4, we used multivariate pattern analysis (MVPA) of EEG data, in particular cross-decoding methods, to draw more direct conclusions about the role of adaptation in spontaneously occurring perceptual reversals.

Based on our results from Chapter 3, Experiment 1, three trial types were established in Chapter 4. These were designed to elicit distinct perceptual states by manipulating the presentation duration and application of adapting directional motion, resulting in three types of perceptual reversals: those due to adaptation from directional motion, those after directional motion without behavioural adaptation, and spontaneous reversals. The experiment in Chapter 4 arranged these three trial types into two block types. The 'directional block', which involves trials with a long pre-exposure to an unambiguous adaptor stimulus (640ms) to promote adaptation, expected to bias perception of a subsequent ambiguous stimulus towards the opposite direction (with an estimated 60-70% effectiveness). This is termed the 'adapted' condition. In the 'unbiased' condition, a shorter adaptor stimulus (80ms) precedes the ambiguous stimulus, anticipated to evoke an unbiased response distribution (around 50% for each perception). Despite identical ambiguous stimuli in both conditions, their preceding trial histories tended to prompt different perceptual interpretations. The second block type, the 'ambiguous block', consists solely of ambiguous motion stimuli pairs with no directional motion at all.

Using these trial types, we employed three distinct methods to train classifiers for the analysis of the brain activity patterns associated with perceptual reversals. Firstly, we trained a classifier to differentiate between reversal and stable trials using only data from ambiguous trials, employing decoding across time techniques. Secondly, we used a cross-decoding approach to train two classifiers: one to identify stable versus reversal trials within the context of the adapted condition, and another to perform the same task within the unbiased condition. Each classifier was then evaluated on its ability to accurately classify stable versus reversal trials that occurred spontaneously within the ambiguous condition. Thirdly, we again used a cross-decoding approach and trained another classifier specifically to discern between reversals that resulted from adaptation (i.e., using trials in the adapted condition) and those that arose from other processes (using trials in the unbiased condition). Finally, we tested this classifier's performance on spontaneous reversals within the ambiguous condition, noting whether it attributed the observed brain activity patterns to the 'unbiased' or 'adapted' category. For each training and testing method described above, we also used the temporal generalisation technique to detect any effects that may not have been completely aligned in time and to draw conclusions around the processes that underpin perceptual reversals.

Our results from the classifier trained to differentiate between reversal and stable trials using only data from ambiguous trials did not find any significant differences between the patterns of brain activity associated with spontaneously occurring reversal vs stable trials. Whilst MVPA classification across time does indeed consider the pattern of activity across the whole scalp, also faces somewhat similar issues to ERP analysis, in that the pattern of scalp activity is compared independently at each timepoint. So, this still requires some degree of temporal alignment of the effects on each trial. It could be the case here that the mechanisms at play at one time point were not active at the same time on every trial.

This idea is supported when the results from the temporal generalisation analysis for this training method are considered. Here, we found a period of significant temporal generalisation towards the end of a trial, where the classification performance at the training time generalises to other time points. The fact that this effect was only revealed with the temporal generalisation technique and not with the classification across time could suggest that there was a certain amount of temporal jitter in the effect's onset. The temporal generalisation pattern itself is indeed similar to the 'jittered' pattern identified by King and Dahaene (2014), supporting the idea that the onset of the reversal event itself could differ across participants and therefore lead to the null effects that we found in the classification across time analysis. The pattern of activity revealed by the temporal generalisation method was also similar to the 'ramping' pattern identified by King and Dahaene (2014). The authors suggest that this pattern could represent a slowly increasing pattern of activity. When this is considered against the hypothesis that neural fatigue is a mechanism that builds up over sustained viewing of a stimulus, our results support this hypothesis.

The classifier that was trained to distinguish adapted reversal vs stable trials was successfully able to predict spontaneously occurring reversal vs stable trials between ~60 – 78ms (classifier performance was above chance level from 60ms onwards, and above that of the unbiased control from 64ms onwards). This indicates that, during that time window, the pattern of brain activity was sufficiently similar in both the adapted trials and ambiguous-only trials. This result was established not to be due to an underlying difference in the data quality in trials in the adapted condition over those in the unbiased condition, because within-condition reversal vs stable classification performance in the adapted trials was not significantly better than that in the unbiased trials during this period. This indicates that it is not an increase in the signal to noise ratio at this point that is causing the result.

During the 60-78ms time window, the classifier trained on adapted trials was more effective at cross-decoding spontaneous reversal data from ambiguous trials than one trained and tested solely on ambiguous trials. While surprising, this may be attributed to the temporal variance in the spontaneous reversal onset, which could elude detection by standard classification across time analysis methods. Supporting evidence comes from the temporal generalisation matrices in this training method, where the adapted-trained classifier successfully decodes various time points in the spontaneous data, especially when early adapted trial data is used to predict later spontaneous data points. Following the process models by King and Dehaene (2014), this implies that the neural mechanisms active early in adapted trials may also be engaged later in spontaneous trials. The similarity in brain activity patterns during adaptation and at the end of ambiguous stimulus viewing may again corroborate the neural fatigue hypothesis, suggesting that the fatigue accumulated from viewing directional stimuli resembles that during ambiguous stimulus viewing, potentially causing perceptual reversals.

Finally, the classifier trained to differentiate between reversals occurring due to adaptation (from reversal data in the adapted trials) vs reversals occurring due to other processes (i.e., from reversal data in the unbiased trials), pinpointed three significant periods where the spontaneous reversals' scalp activity patterns resembled those of adapted reversals more closely. Of these, two were after the stimulus onset (45 - 68ms and 234 - 250ms), suggesting that spontaneous reversals share similar underlying processes with adapted reversals rather than with reversals following non-adaptive directional motion. This coincides with the period of above-chance decoding in the temporal generalisation analysis of the spontaneously occurring reversals in the ambiguous only trials, suggesting that adaptation may play a role in spontaneously occurring reversals.

Additionally, in the pre-stimulus phase (-84 to -68ms), the classifier's predictions indicated a pre-existing similarity in scalp activity to adapted reversals, which could suggest an anticipatory neural state akin to adaptation. However, this period also coincides with a potential signal-to-noise ratio increase in adapted trials, complicating interpretations. This pre-stimulus prediction could reflect a build-up of neural fatigue or align with theories of spontaneous oscillatory processes, though the latter lacks corroborative evidence from classification success in the same time window. These interpretations must be approached with caution due to the overlapping signal-to-noise considerations.

Therefore, this research aim was addressed in Chapter 4, using a trial design based on the results of Experiment 1 in Chapter 3. The results from the three classifier training methods used in Chapter 4 provide substantial support for our hypothesis that adaptation is likely to play a role in spontaneously occurring reversals in perception of ambiguous motion. These results have important theoretical implications, as they provide direct support for the neural fatigue hypothesis that has previously not been shown.

6.2.4 Are ERP measures of reversal-related brain activity present during perceptual reversals of ambiguous stationary stimuli also present during perceptual reversals of ambiguous motion stimuli?

Kornmeier and Bach (2012) have highlighted that extensive research has been conducted that shows the presence of two main ERP components in response to perceptual reversals of stationary stimuli (for example, the Necker cube). The first has been termed the 'Reversal Positivity' (RP), occurring around 120ms post-stimulus onset at occipital electrode sites. The RP has been suggested to be reflective of the detection of a processing conflict resulting from the ambiguity of the presented stimulus (Kornmeier et al., 2011; Kornmeier & Bach, 2004, 2012). The second component is known as the 'Reversal Negativity' (RN) which occurs around 200-260ms post-stimulus onset over posterior electrode sites. The RN is thought to reflect top-down influences over perception as it has been shown to be enhanced during active volitional control of reversals (Pitts et al., 2008). Kornmeier and Bach (2012) suggested that further research is needed to establish whether these two markers of perceptual reversals are also elicited by other ambiguous stimuli, such as the ambiguous motion used in this thesis. Therefore, in addition to the multivariate analyses conducted in Chapter 4, an ERP analysis was also carried out to establish whether the RN and RP were present in response to spontaneous perceptual reversals of ambiguous motion created from drifting sine-wave gratings.

However, neither component was found to be present in the ERPs from data captured during the spontaneously occurring reversals in the ambiguous-only blocks from the EEG experiment described in Chapter 4. In the currently available literature, the RP has shown resistance to differences in low-level stimulus characteristics (Kornmeier et al., 2011; 2012) and has been replicated with various stationary ambiguous stimuli (e.g., Britz et al., 2009; Kornmeier et al., 2007). However, its replication using ERP methods is not universal (e.g., Intaite et al., 2010; Pitts et al., 2007), which has been suggested to be because of its relatively low amplitude making it particularly prone to being lost in the spatio-temporal averaging process involved in the generation of the ERP (Luck, 2014). Additionally, recent findings suggest it may be associated more with participants' manual responses to indicate the reversal, rather than being linked to the perceptual reversal event itself (Abdallah & Brooks, 2020). As ambiguous motion was used as the stimulus in our experiment, it is reasonable to assume that the reversal event itself could occur at a slightly different time point within the trial both within and between participants. Therefore, the lack of the RP here could indeed be due to the averaging process and indicative that the RP may not be robust enough to be detected when ambiguous motion stimuli are used.

Like the RP, the RN has been replicated across different ambiguous stimuli (e.g., Britz et al., 2009; Intaite et al., 2010; Kornmeier & Bach, 2004) but not universally so (Kornmeier & Bach, 2014). Again, the reason for its absence here could be due to the potential jittered onset of the reversal event and therefore the loss of the component during the averaging required to produce the ERP. It should be noted, however, that the MVPA results (see Section 6.2.3 for a summary of these) reveal a time window where spontaneously occurring reversal vs stable trials can reliably be decoded by a classifier trained to decode adapted reversal vs stable trials, and that time window occurs within the temporal ROI for the RN component. There is debate around the processes reflected by the RN, as it has been shown to be modulated by both top-down (Pitts et al., 2008) and bottom-up

(Kornmeier et al., 2007) factors. So, the finding here, that there is a period of time during spontaneously occurring reversals where the brain activity within the RN window is similar to that in adaptation, could provide support for the idea that the RN at least partially represents bottom-up processes.

This aim then was addressed in Chapter 4, via the ERP analysis of the EEG data collected during spontaneously occurring reversal vs stable trials from the ambiguous-only block. It was found that, potentially due to a jittered onset of the reversal event itself, these two ERP markers of perceptual reversals were not detected.

6.2.5 How powerful is MVPA to detect multivariate effects present in EEG data?

The currently available literature on the use of MVPA on electrophysiological data does not provide researchers with sufficient guidance on the number of participants required to detect an effect within a given ROI, where one exists. Therefore, in Chapter 5 we used simulated EEG data to investigate the statistical power of a typical MVPA pipeline. Specifically, to establish how varying both sample size, and the size of an added multivariate effect affects the resulting decoding accuracy and power of cluster-based corrections to detect the effect. Small, medium and large effects (in terms of Cohen’s d) were investigated, as well as 5 sample sizes (5, 10, 20, 40 and 80 participants). Additionally, Chapter 5 also used simulated null effect EEG data to investigate the family-wise error rates (FWER) associated with the same MVPA pipeline. The purpose of this was to provide more information on the FWER of cluster-based permutations, and to establish whether the cluster-based permutations sufficiently correct for multiple comparisons within the specific context of analysing decoding accuracy values resulting from MVPA. These simulations were particularly relevant to the work presented in this thesis, as the claims made in relation to the role of adaptation in spontaneously occurring perceptual reversals are based upon MVPA with cluster-based corrections for Type I errors.

The simulation results demonstrated that effect size had a significant impact on classifier decoding accuracy. Specifically, an increase in the size of the multivariate effect increased decoding accuracy. To assess power, cluster-based permutation tests were employed to pinpoint the specific timepoints within the ROI at which the classifier exceeded statistical chance levels. The findings revealed that for a large effect, 80% statistical power can be achieved with as few as 10 participants. However, a medium effect requires 20 participants to reach the same level of power, and a small effect fails to reach 80% power even with 80 participants when employing cluster-based tests. In experiments involving only 5 participants, no effect at all was detected within the ROI (i.e., power was 0%). The simulations also showed that the FWER of the cluster-based permutation tests was 0.011, which is far below the accepted 5% false positive rate.

The surprising result of 0% power with 5 participants was investigated further in Chapter 5 as it was suggestive of an over-correction for multiple comparisons. Therefore, it was necessary to determine whether this issue lay with correction step, or whether there was some underlying issue with the data itself. If there was no underlying issue with the data, we expected that running uncorrected t-tests on an expanding null effect time window would lead to the expected increase in false positives (i.e., if just one time point was compared then the false positive rate would be low, but if 20 time points were compared then the false positive rate would be high as multiple comparisons have been made without correction). This was indeed the case, suggesting that the use of cluster-based corrections in this way may be overly conservative, however further work is required to confirm this.

The results of the simulation conducted in Chapter 5 also have an influence on how we interpret the results in Chapter 4. These results suggest that, with a sample size of 25, we are likely to have missed any small effects contained in the data. This sample size was originally selected as it is consistent with others using this technique in the literature (e.g., Kornmeier, Hein & Bach, 2009; Das et al., 2010; List et al., 2017; Bae & Luck, 2018). However, our simulation results here suggest that this may have been underpowered for small effects. Running underpowered experiments has been identified as a problem as it can cause systematic overestimation of true effect size over the long term (e.g., Button et al., 2013; Vadillo et al., 2016). Therefore, we should ensure that future experiments use the results of Chapter 5 to determine an adequate sample size to detect even a small effect. (i.e., 80 participants).

This aim then, was addressed in Chapter 5, where we ran 15 simulations to assess the statistical power of MVPA pipelines in the context of decoding patterns of EEG data. We found that an increase in sample size and effect size led to an increase in statistical power. In this chapter we also examined FWER of the same analysis pipeline and found FWER to be conservatively controlled using cluster-based corrections. We showed that this conservatism was potentially over-corrective when sample size is low.

6.3 Limitations and future directions

Throughout this thesis, several different paradigms have been used to explore the role of adaptation in spontaneously occurring perceptual reversals. The individual limitations and considerations for future research that are relevant to each paradigm are discussed within each of the relevant chapters. More general limitations and potential directions for future research will now be outlined here.

One of the main challenges faced when running the experiments described in this thesis is participant fatigue and/or inattention. This is particularly true of the EEG experiment described in Chapter 4; however, it is also relevant to the behavioural experiments described in Chapter 3. The EEG experimental procedure firstly requires that participants are set up with the EEG cap and electrodes. This procedure can take up to 40 minutes, meaning that the participant has sat very still for an extended period before the experimental phase has even begun. The EEG experiment in Chapter 4 contained 604 trials in which participants were repeatedly asked to report the direction of motion of two drifting sine wave gratings. The experimental trials in this experiment took around 30 minutes for participants to complete, in addition to the ~40 minutes that they were sat in the chair being set up with the EEG equipment. Although we did introduce regular breaks in an attempt to reduce boredom here, because of the limitation of movement due to the EEG cap and electrodes, participants were essentially sat very still in a dimly lit room for upwards of an hour. In addition to the amount of time that participants were required to sit still for, another source of inattention could come from the task itself, as the stimuli presented and responses required were extremely repetitive. This did lead to noticeable signs of fatigue in participants, such as yawning etc. As mentioned above, the repetitive nature of the task and the restriction of being sat still for extended periods of time was true of all of the experiments described in this thesis, but the effects of this boredom have a particular influence on EEG results. For instance, when participants are fatigued, their alpha activity increases and this has been shown to have a suppressive effect on spontaneous perceptual reversals (of a Necker cube stimulus; Piantoni et al., 2017).

In addition to this, a further challenge faced during this set of experiments was the element of task difficulty in Experiment 1 in Chapter 3 and the experiment in Chapter 4. In order to counterbalance responses, participants had to remember that the up and down arrows were to be used to make left/right responses. Moreover, these would switch between blocks. Although this was not difficult for the participants to understand, it could have been difficult for them to maintain the level of focus required to respond in this way. The results from our attention check trials suggest that participants were generally able to keep up with the demands of this response pattern, but anecdotal reports from participants suggest that they did find it difficult to remember how to respond. This overall difficulty could have been added to because the results from Experiment 2, Chapter 3 suggest that participants were failing to perceive one of the two stimuli that they were being asked about in such experiments.

A possible direction for future research, related to this element of task difficulty, is to repeat the experiment conducted in Chapter 4 using the response paradigm employed by Takeuchi et al. (2011). In this paradigm, participants are asked only to report the direction of the ambiguous stimulus in each trial rather than to make any sort of comparison between two stimuli they may not have noticed. This would decrease the demands on the participant to compare two stimuli and instead this comparison would be made by the experimenter during analysis (by comparing the known direction of the adaptor stimulus to the reported direction of the ambiguous stimulus).

An additional limitation of the set of experiments described in this thesis is that we did not directly explore the impact of eye-movements on perception. Throughout the experiments, participants were advised to maintain fixation, monitored by an experimenter using a camera to observe their eyes. Compliance with this instruction was generally observed. Additionally, in the pre-processing phase of the EEG data analysis, independent component analysis (ICA) was carried out to remove eye-movement components, and any remaining trials containing eye-movements were discarded. This approach provides us with substantial confidence that the SVM classifiers were analysing neural rather than ocular movement signals. However, some research has shown that there are certain eye-movements that cannot be detected using these techniques. For instance, Dimigen et al. (2009) showed that microsaccades (i.e., small, involuntary eye-movements) could go undetected by typical eye-movement artefact rejection procedures as they are small (~2-3μV) and do not lead to typical artefactual EEG traces. This could be directly assessed in future work if eye-position was recorded via concurrent eye-tracking. If eye-tracking data were simultaneously recorded during the experiment described in Chapter 4, for instance, the spontaneous stable vs reversal decoding analysis could have been conducted using EEG data as it was here, and then separately conducted using eye-position data in order to examine the amount of information in eye-position that was predictive of perception. Indeed, research has indicated that these microsaccadic movements during fixation are predictive of perception of ambiguous apparent motion stimuli involving a grid of moving dots (Laubrock et al., 2008). Therefore, this could be investigated in relation to the ambiguous drifting sine wave stimuli used here.

Moreover, previous research has actively cautioned against the over-reliance of ocular artefact rejection methods in EEG data processing and interpretation (e.g., Plochl et al., 2012). It has been highlighted that eye movements go hand in hand with neural activity and should not be considered to only introduce artefacts to the EEG signal. Overlooking this could lead to misinterpretation of the data, given that our natural way of viewing the world does indeed include many eye movements. Therefore, instead of being considered as interferences, eye-movements should instead be thought of as a part of natural human behaviour. Therefore, a sensible direction for future research using MVPA decoding and cross-decoding as in this thesis, would be to run these analyses both with and without including eye-movements as a feature fo the classifier to use. This sort of analyses could lead to the identification of novel mechanisms involved in spontaneous reversals in perception.

The general paradigm employed in Chapter 4 has the potential to be used in a body of future work. Here, we used this paradigm to induce an adapted brain state using directional drifting sine-wave motion. However, this could be taken forward and used alongside a variety of other ambiguous motion stimuli to assess whether these results are generalisable across them. If they are, this would be suggestive of a shared mechanism underpinning a general adapted brain state, rather than being specific to drifting sine-wave gratings. These results would provide further evidence in assessment of the neural fatigue hypothesis.

6.4 Conclusions

In conclusion, the experiments described in this thesis involved the collection of behavioural and electrophysiological data to investigate the role of adaptation in spontaneously occurring perceptual reversals of ambiguous motion created from drifting sine-wave gratings. Additionally, the simulations used in this thesis sought to provide guidance on the statistical power of multivariate pattern analysis in the context of EEG data. Overall, our behavioural results indicated that a rapid adaptation to directional motion can indeed occur in under a second, however the visual motion priming effect thought to be evoked by a shorter pre-exposure to directional motion was thrown into doubt. Some of our behavioural results as well as reports from our participants led us to hypothesise that there could be an artefactual priming effect resulting from asking participants to report on their perception of a stimulus that they may not have noticed. A further empirical behavioural investigation suggested that this could indeed have been the case when participants reported priming effects but not when they were reporting adaptation effects. Our electrophysiological findings showed that previously known markers of spontaneous perceptual reversals of ambiguous stationary stimuli may not be sufficiently temporally aligned to be detected in response to reversals of ambiguous motion stimuli. This result was supported by the finding that any information contained in brain activity that was predictive of perceptual state (i.e., stable or reversal) was only revealed when temporal generalisation methods were employed. The patterns of predictive activity suggested a jittered and ramping up process, which is suggestive of a build-up of adaptation in a manner consistent with the neural fatigue hypothesis. This suggestion was supported by our result that brain activity recorded during perceptual reversals that were driven by adaptation was predictive of perceptual state during spontaneously occurring reversals. Finally, the results of our simulations showed that multivariate pattern analysis pipelines are a sufficiently powerful analysis technique, that may be overly conservative when low sample sizes are used.

References

Abdallah, D., & Brooks, J. L. (2020). Response dependence of reversal-related ERP components in perception of ambiguous figures. Psychophysiology, 57(12), e13685. https://doi.org/10.1111/psyp.13685

Albright, T., & Stoner, G. (2002). Contextual Influences on Visual Processing. *Annual Review Of Neuroscience, 25*(1), 339-379. https://doi.org/10.1146/annurev.neuro.25.112701.142900

Allison, T., Puce, A., Spencer, D., & McCarthy, G. (1999). Electrophysiological Studies of Human Face Perception. I: Potentials Generated in Occipitotemporal Cortex by Face and Non-face Stimuli. *Cerebral Cortex*, *9*(5), 415-430. doi: 10.1093/cercor/9.5.415

Andrews, T., Schluppeck, D., Homfray, D., Matthews, P., & Blakemore, C. (2002). Activity in the Fusiform Gyrus Predicts Conscious Perception of Rubin's Vase–

Face Illusion. *Neuroimage, 17*(2), 890-901. https://doi.org/10.1006/nimg.2002.1243

Anllo-Vento, L. & Hillyard, S.A. (1996). Selective attention to the color and direction of moving stimuli: electrophysiologi*ca*l correlates of hierarchical feature selection. *Perception & Psychophysics, 58*(2), 191-206. https://doi.org/10.3758/BF03211875

Anstis, S., Verstraten, F. A. J., & Mather, G. (1998). The motion aftereffect. Trends in Cognitive Sciences, 2(3), 111–117. https://doi.org/10.1016/S1364-6613(98)01142-5

Arieli, A., Shoham, D., Hildesheim, R., & Grinvald, A. (1995). Coherent spatiotemporal patterns of ongoing activity revealed by real-time optical imaging coupled with single-unit recording in the cat visual cortex. *Journal Of Neurophysiology*, *73*(5), 2072-2093. https://doi.org/10.1152/jn.1995.73.5.2072

Arieli, A., Sterkin, A., Grinvald, A., & Aertsen, A. (1996). Dynamics of Ongoing Activity: Explanation of the Large Variability in Evoked Cortical Responses. *Science, 273*(5283), 1868-1871. https://doi.org/10.1126/science.273.5283.1868

Babich, S., & Standing, L. (1981). Satiation Effects with Reversible Figures. *Perceptual And Motor Skills*, *52*(1), 203-210. doi: 10.2466/pms.1981.52.1.203

Bach, M., & Meigen, T. (1992). Electrophysiological correlates of texture segregation in the human visual evoked potential. *Vision Research, 32*(3), 417-424. https://doi.org/10.1016/0042-6989(92)90233-9

Bae, G. Y., & Luck, S. J. (2018). Dissociable decoding of spatial attention and working memory from EEG oscillations and sustained potentials. *Journal of*

Neuroscience, *38*(2), 409-422.

Bae, G. Y., & Luck, S. J. (2019). Reactivation of previous experiences in a working memory task. *Psychological science*, *30*(4), 587-595.

Baker, D., Karapanagiotidis, T., Coggan, D., Wailes-Newson, K., & Smallwood, J. (2015). Brain networks underlying bistable perception. *Neuroimage*, *119*, 229-234. https://doi.org/10.1016/j.neuroimage.2015.06.053

Bar, M. (2009). The proactive brain: memory for predictions. *Philosophical Transactions Of The Royal Society B: Biological Sciences*, *364*(1521), 1235-1243.

Başar-Eroglu, C., Strüber, D., Kruse, P., Başar, E., & Stadler, M. (1996). Frontal gamma-band enhancement during multistable visual perception. *International Journal Of Psychophysiology, 24*(1-2), 113-125. https://doi.org/10.1016/s0167-8760(96)00055-4

Basar-Eroglu, C., Strüber, D., Stadler, M., and Kruse,E. (1993). Multistable visual perception induces a slow positive EEG wave. *International Journal of Neuroscience, 73*, 139–151. https://doi.org/10.3109/00207459308987220

Baylis, G., & Driver, J. (1995a). Obligatory edge assignment in vision: The role of figure and part segmentation in symmetry detection. *Journal Of Experimental Psychology: Human Perception And Performance, 21*(6), 1323-1342. https://doi.org/10.1037/0096-1523.21.6.1323

Baylis, G., & Driver, J. (1995b). One-Sided Edge Assignment in Vision: 1. Figure-Ground Segmentation and Attention to Objects. *Current Directions In Psychological Science, 4(*5), 140-146. https://doi.org/10.1111/1467-

8721.ep10772580

Beck, D., Muggleton, N., Walsh, V., & Lavie, N. (2006). Right Parietal Cortex Plays a Critical Role in Change Blindness. *Cerebral Cortex*, *16*(5), 712-717. doi: 10.1093/cercor/bhj017

Beck, D., Rees, G., Frith, C., & Lavie, N. (2001). Neural correlates of change detection and change blindness. *Nature Neuroscience*, *4*(6), 645-650. doi: 10.1038/88477

Beckmann, C., DeLuca, M., Devlin, J., & Smith, S. (2005). Investigations into resting-state connectivity using independent component analysis. *Philosophical Transactions Of The Royal Society B: Biological Sciences*, *360*(1457), 1001-1013. doi: 10.1098/rstb.2005.1634

Bentin, S., Allison, T., Puce, A., Perez, E., & McCarthy, G. (1996). Electrophysiological Studies of Face Perception in Humans. *Journal Of Cognitive Neuroscience*, *8*(6), 551-565. https://doi.org/10.1162/jocn.1996.8.6.551

Blair, R., & Karniski, W. (1993). An alternative method for significance testing of waveform difference potentials. *Psychophysiology*, *30*(5), 518-524. doi: 10.1111/j.1469-8986.1993.tb02075.x

Blake, R. (1989). A neural theory of binocular rivalry. *Psychological Review*, *96*(1), 145-167. doi: 10.1037/0033-295x.96.1.145

Blake, R., & Logothetis, N. K. (2002). Visual competition. *Nature Reviews Neuroscience, 3*(1), 13. https://doi.org/10.1038/nrn701

Blake, R., Sobel, K. V., & Gilroy, L. A. (2003). Visual Motion Retards Alternations between Conflicting Perceptual Interpretations. Neuron, 39(5), 869–878. https://doi.org/10.1016/S0896-6273(03)00495-1

Blakemore, C., & Campbell, F. (1969). On the existence of neurones in the human visual system selectively sensitive to the orientation and size of retinal images. *The Journal Of Physiology*, *203*(1), 237-260. doi: 10.1113/jphysiol.1969.sp008862

Boly, M., Balteau, E., Schnakers, C., Degueldre, C., Moonen, G., & Luxen, A. et al. (2007). Baseline brain activity fluctuations predict somatosensory perception in humans. *Proceedings Of The National Academy Of Sciences, 104*(29), 12187-12192. https://doi.org/10.1073/pnas.0611404104

Borsellino, A., Marco, A., Allazetta, A., Rinesi, S., & Bartolini, B. (1972). Reversal time distribution in the perception of visual ambiguous stimuli. *Kybernetik*, *10*(3), 139-144. doi: 10.1007/bf00290512

Botzel, K., Schulze, S., & Stodieck, S. (1995). Scalp topography and analysis of intracranial sources of face-evoked potentials. *Experimental Brain Research,*

Brascamp, J., Blake, R., & Knapen, T. (2015). Negligible fronto-parietal BOLD activity accompanying unreportable switches in bistable perception. *Nature*

Brascamp, J., Pearson, J., Blake, R., & van den Berg, A. (2009). Intermittent ambiguous stimuli: Implicit memory causes periodic perceptual alternations. *Journal Of Vision*, *9*(3), 3-3. doi: 10.1167/9.3.3

Brascamp, J., Sterzer, P., Blake, R., & Knapen, T. (2018). Multistable perception and the role of the frontoparietal cortex in perceptual inference. *Annual review of psychology, 69*, 77- 103. https://doi.org/10.1146/annurev-psych-010417-085944

Braun, J., & Mattia, M. (2010). Attractors and noise: Twin drivers of decisions and multistability. *Neuroimage*, *52*(3), 740-751.

https://doi.org/10.1016/j.neuroimage.2009.12.126

Britz, J., & Pitts, M. A. (2011). Perceptual reversals during binocular rivalry: ERP components and their concomitant source differences. *Psychophysiology, 48*(11), 1490-1499. https://doi.org/10.1111/j.1469-8986.2011.01222.x

Britz, J., Landis, T., & Michel, C. (2009). Right Parietal Brain Activity Precedes Perceptual Alternation of Bistable Stimuli. *Cerebral Cortex, 19*(1), 55-65. https://doi.org/10.1093/cercor/bhn056

Brooks, J., Zoumpoulaki, A., & Bowman, H. (2016). Data-driven region-of-interest selection without inflating Type I error rate. *Psychophysiology, 54*(1), 100-113. https://doi.org/10.1111/psyp.12682

Busch, N., Dubois, J., & VanRullen, R. (2009). The Phase of Ongoing EEG Oscillations Predicts Visual Perception. *Journal Of Neuroscience, 29*(24), 7869-7876. https://doi.org/10.1523/jneurosci.0113-09.2009

Busch, N., Dubois, J., & VanRullen, R. (2010). The phase of ongoing EEG oscillations predicts visual perception. *Journal Of Vision*, *9*(8), 737-737. https://doi.org/10.1167/9.8.737

Buschman, T., & Miller, E. (2007). Top-Down Versus Bottom-Up Control of Attention in the Prefrontal and Posterior Parietal Cortices. *Science*, *315*(5820), 1860-1862. https://doi.org/10.1126/science.1138071

Buzsáki, G., Anastassiou, C., & Koch, C. (2012). The origin of extracellular fields and currents — EEG, ECoG, LFP and spikes. *Nature Reviews Neuroscience, 13*(6), 407-420. https://doi.org/10.1038/nrn3241

Caffier, P., Erdmann, U., & Ullsperger, P. (2003). Experimental evaluation of eye-blink parameters as a drowsiness measure. *European Journal Of Applied Physiology, 89*(3), 319-325. https://doi.org/10.1007/s00421-003-0807-5

Cao, R., Braun, J., & Mattia, M. (2014). Stochastic Accumulation by Cortical Columns May Explain the Scalar Property of Multistable Perception. *Physical Review Letters, 113*(9). https://doi.org/10.1103/physrevlett.113.098103

Carbon, C.-C. (2014). Understanding human perception by human-made illusions. Frontiers in Human Neuroscience, 8. https://www.frontiersin.org/articles/10.3389/fnhum.2014.00566

Carmel, D., Walsh, V., Lavie, N., & Rees, G. (2010). Right parietal TMS shortens dominance durations in binocular rivalry. *Current Biology*, *20*(18), R799-R800. https://doi.org/10.1016/j.cub.2010.07.036

Castelhano, J., Rebola, J., Leitão, B., Rodriguez, E., & Castelo-Branco, M. (2013). To Perceive or Not Perceive: The Role of Gamma-band Activity in Signaling Object Percepts. *Plos ONE*, *8*(6), e66363. doi: 10.1371/journal.pone.0066363

Castro, L., Lazareva, O.F., Vecera, S.P., Wasserman, E.A. (2010). Changes in area affect figure-ground assignment in pigeons. *Vision Research, 50*, 497–508. https://doi.org/10.1016/j.visres.2009.12.016

Cauchoix, M., Arslan, A. B., Fize, D., & Serre, T. (2012). The neural dynamics of visual processing in monkey extrastriate cortex: A comparison between univariate and multivariate techniques. In G. Langs, I. Rish, M. Grosse-Wentrup, & B. Murphy

Cauchoix, M., Barragan-Jason, G., Serre, T., & Barbeau, E. (2014). The Neural Dynamics of Face Detection in the Wild Revealed by MVPA. *The Journal Of Neuroscience, 34*(3), 846-854. https://doi.org/10.1523/jneurosci.3030-13.2014

Champod, A., & Petrides, M. (2007). Dissociable roles of the posterior parietal and the prefrontal cortex in manipulation and monitoring processes. *Proceedings Of The National Academy Of Sciences*, *104*(37), 14837-14842. doi: 10.1073/pnas.0607101104

Cochin, S., Barthelemy, C., Lejeune, B., Roux, S., & Martineau, J. (1998). Perception of motion and qEEG activity in human adults. *Electroencephalography And Clinical Neurophysiology*, *107*(4), 287-295. doi: 10.1016/s0013-4694(98)00071-6

Cohen, E.H., Barenholtz, E., Singh, M., Feldman, J. (2005). What change detection tells us about the visual representation of shape. *Journal of Vision*, 5(4), 3. https://doi.org/10.1167/5.4.3

Cortical cross-frequency coupling predicts perceptual outcomes. *Neuroimage, 69*, 126-137. https://doi.org/10.1016/j.neuroimage.2012.11.021

Coste, C., Sadaghiani, S., Friston, K., & Kleinschmidt, A. (2011). Ongoing Brain Activity Fluctuations Directly Account for Intertrial and Indirectly for Intersubject Variability in Stroop Task Performance. *Cerebral Cortex*, *21*(11), 2612-2619. https://doi.org/10.1093/cercor/bhr050

Courchesne, E., Hillyard, S.A., Galambos, R. (1975). Stimulus novelty, task relevance and the visual evoked potential in man. *Electroencephalography and Clinical Neurophysiology, 39*(2), 131-143. https://doi.org/10.1016/0013-4694(75)90003-6

Courtney, S., Ungerleider, L., Keil, K., & Haxby, J. (1997). Transient and sustained activity in a distributed neural system for human working memory. *Nature*, *386*(6625), 608-611. doi: 10.1038/386608a0

Daniel, E., Meindertsma, T., Arazi, A., Donner, T., & Dinstein, I. (2019). The

Relationship between Trial-by-Trial Variability and Oscillations of Cortical Population Activity. *Scientific Reports*, *9*(1). doi: 10.1038/s41598-019-53270-7

Das, K., Giesbrecht, B., & Eckstein, M. (2010). Predicting variations of perceptual performance across individuals from neural activity using pattern classifiers.

de Graaf, T., de Jong, M., Goebel, R., van Ee, R., & Sack, A. (2011). On the Functional Relevance of Frontal Cortex for Passive and Voluntarily Controlled Bistable Vision. *Cerebral Cortex*, *21*(10), 2322-2331. https://doi.org/10.1093/cercor/bhr015

de Jong, M., Brascamp, J., Kemner, C., van Ee, R., & Verstraten, F. (2014). Implicit Perceptual Memory Modulates Early Visual Processing of Ambiguous

de Jong, M., Knapen, T., & van Ee, R. (2012a). Opposite Influence of Perceptual Memory on Initial and Prolonged Perception of Sensory Ambiguity. *Plos ONE*, *7*(1), e30595. doi: 10.1371/journal.pone.0030595

de Jong, M., Kourtzi, Z., & van Ee, R. (2012b). Perceptual experience modulates cortical circuits involved in visual awareness. *European Journal Of Neuroscience*, *36*(12), 3718-3731. doi: 10.1111/ejn.12005

Deecke, L., Grözinger, B., & Kornhuber, H. (1976). Voluntary finger movement in man: Cerebral potentials and theory. *Biological Cybernetics, 23*(2), 99-119. https://doi.org/10.1007/bf00336013

Dehaene, S., & Changeux, J. (2011). Experimental and Theoretical Approaches to Conscious Processing. *Neuron*, *70*(2), 200-227. doi: 10.1016/j.neuron.2011.03.018

Delorme, A., & Makeig, S. (2004). EEGLAB: An open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. Journal of Neuroscience Methods, 134(1), 9–21. https://doi.org/10.1016/j.jneumeth.2003.10.009

Dentico, D., Cheung, B., Chang, J., Guokas, J., Boly, M., Tononi, G., & Van Veen, B. (2014). Reversal of cortical information flow during visual imagery as compared to visual perception. *Neuroimage*, *100*, 237-243. https://doi.org/10.1016/j.neuroimage.2014.05.081

Dering, B., Martin, C., Moro, S., Pegna, A., & Thierry, G. (2011). Face-Sensitive Processes One Hundred Milliseconds after Picture Onset. *Frontiers In Human Neuroscience*, *5*. https://doi.org/10.3389/fnhum.2011.00093

Dhamala, M., Rangarajan, G., & Ding, M. (2008). Analyzing information flow in brain networks with nonparametric Granger causality. *Neuroimage*, *41*(2), 354-362. doi: 10.1016/j.neuroimage.2008.02.020

Dijkstra, N., Zeidman, P., Ondobaka, S., van Gerven, M., & Friston, K. (2017). Distinct Top-down and Bottom-up Brain Connectivity During Visual Perception and Imagery. *Scientific Reports*, *7*(1). https://doi.org/10.1038/s41598-017-05888-8

Dimigen, O., Valsecchi, M., Sommer, W., & Kliegl, R. (2009). Human Microsaccade-Related Visual Brain Responses. *Journal Of Neuroscience*, *29*(39), 12321-12331. doi: 10.1523/jneurosci.0911-09.2009

Donner, T., Siegel, M., Fries, P., & Engel, A. (2009). Buildup of Choice-Predictive Activity in Human Motor Cortex during Perceptual Decision Making. *Current Biology*, *19*(18), 1581-1585. doi: 10.1016/j.cub.2009.07.066

Downing, P. (2000). Interactions Between Visual Working Memory and Selective

Attention. *Psychological Science*, *11*(6), 467-473. doi: 10.1111/1467-9280.00290

Driver, J., & Baylis, G. (1996). Edge-Assignment and Figure–Ground Segmentation in Short-Term Visual Matching. *Cognitive Psychology, 31*(3), 248-306. https://doi.org/10.1006/cogp.1996.0018

Dugue, L., Marque, P., & VanRullen, R. (2011). The Phase of Ongoing Oscillations Mediates the Causal Relation between Brain Excitation and Visual Perception. *Journal Of Neuroscience*, *31*(33), 11889-11893. doi: 10.1523/jneurosci.1161-11.2011

Duncan, J., & Humphreys, G. (1989). Visual search and stimulus similarity. *Psychological Review*, *96*(3), 433-458. doi: 10.1037/0033-295x.96.3.433

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

Dustman, R., & Beck, E. (1965). Phase of alpha brain waves, reaction time and visually evoked potentials. *Electroencephalography And Clinical Neurophysiology*, *18*(5), 433-440. https://doi.org/10.1016/0013-4694(65)90123-9

Eger, E., Ashburner, J., Haynes, J., Dolan, R., & Rees, G. (2008). fMRI Activity Patterns in Human LOC Carry Information about Object Exemplars within Category. *Journal Of Cognitive Neuroscience*, *20*(2), 356-370. doi: 10.1162/jocn.2008.20019

Egeth, H., & Yantis, S. (1997). Visual Attention: Control, Representation, and Time Course. *Annual Review Of Psychology*, *48*(1), 269-297.

https://doi.org/10.1146/annurev.psych.48.1.269

Ehm, W., Bach, M., & Kornmeier, J. (2011). Ambiguous figures and binding: EEG frequency modulations during multistable perception. *Psychophysiology, 48*(4), 547-558. https://doi.org/10.1111/j.1469-8986.2010.01087.x

Eimer, M., & Holmes, A. (2002). An ERP study on the time course of emotional face processing. *Neuroreport*, *13*(4), 427-431. https://doi.org/10.1097/00001756-200203250-00013

Ellis, S., & Stark, L. (1978). Eye Movements during the Viewing of Necker Cubes. *Perception*, *7*(5), 575-581. https://doi.org/10.1068/p070575

Engel, A., & Fries, P. (2010). Beta-band oscillations—signalling the status quo?. *Current Opinion In Neurobiology*, *20*(2), 156-165. doi: 10.1016/j.conb.2010.02.015

Engell, A., & McCarthy, G. (2011). The Relationship of Gamma Oscillations and Face-Specific ERPs Recorded Subdurally from Occipitotemporal Cortex. *Cerebral Cortex, 21*(5), 1213-1221. https://doi.org/10.1093/cercor/bhq206

Faul, F., &Erdefelder, E. (1992). GPOWER: A priori-, post hoc-, and compromise power analyses for MS-DOS [Computer program]. Bonn, Germany: Bonn University.

Fechner, G. (1966). *Elements of psychophysics*. Holt, Rinehart & Winston.

Fiebelkorn, I., Saalmann, Y., & Kastner, S. (2013a). Rhythmic Sampling within and between Objects despite Sustained Attention at a Cued Location. *Current Biology, 23*(24), 2553-2558. https://doi.org/10.1016/j.cub.2013.10.063

Fischer, B., & Weber, H. (1993). Express saccades and visual attention. *Behavioral And*

Brain Sciences, *16*(3), 553-567. doi: 10.1017/s0140525x00031575

Fiser, J., Chiu, C., & Weliky, M. (2004). Small modulation of ongoing cortical dynamics by sensory input during natural vision. *Nature, 431*(7008), 573-578. https://doi.org/10.1038/nature02907

Fleming, S., Whiteley, L., Hulme, O., Sahani, M., & Dolan, R. (2010). Effects of Category-Specific Costs on Neural Systems for Perceptual Decision-Making. *Journal Of Neurophysiology*, *103*(6), 3238-3247. https://doi.org/10.1152/jn.01084.2009

Foster, J., Sutterer, D., Serences, J., Vogel, E., & Awh, E. (2016). The topography of alpha-band activity tracks the content of spatial working memory. *Journal Of Neurophysiology, 115*(1), 168-177. https://doi.org/10.1152/jn.00860.2015

Fox, M., & Raichle, M. (2007). Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nature Reviews Neuroscience*, *8*(9), 700- 711. https://doi.org/10.1038/nrn2201

Freeman, A. W. (2005). Multistage Model for Binocular Rivalry. Journal of Neurophysiology, 94(6), 4412–4420. https://doi.org/10.1152/jn.00557.2005

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

Garcia, J., Srinivasan, R., & Serences, J. (2013). Near-Real-Time Feature-Selective Modulations in Human Cortex. *Current Biology, 23*(6), 515-522. https://doi.org/10.1016/j.cub.2013.02.013

Gauthier, I., Curran, T., Curby, K., & Collins, D. (2003). Perceptual interference supports a non-modular account of face processing. *Nature Neuroscience*, *6*(4), 428-432. doi: 10.1038/nn1029

Gauthier, I., Tarr, M., Anderson, A., Skudlarski, P., & Gore, J. (1999). Activation of the middle fusiform 'face area' increases with expertise in recognizing novel objects. *Nature Neuroscience, 2*(6), 568-573. https://doi.org/10.1038/9224

George, R. (1936). The Significance of the Fluctuations Experienced in Observing Ambiguous Figures and in Binocular Rivalry. *The Journal Of General Psychology*, *15*(1), 39-61. doi: 10.1080/00221309.1936.9917904

Geweke, J. (1982). Measurement of Linear Dependence and Feedback Between Multiple Time Series: Rejoinder. *Journal Of The American Statistical Association*, *77*(378), 323. doi: 10.2307/2287243

Gilden, D. (2001). Cognitive emissions of 1/f noise. *Psychological Review*, *108*(1), 33-

Girgus, J., Rock, I., & Egatz, R. (1977). The effect of knowledge of reversibility on the reversibility of ambiguous figures. *Perception & Psychophysics*, *22*(6), 550-556. doi: 10.3758/bf03198762

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

Gramfort, A., Luessi, M., Larson, E., Engemann, D., Strohmeier, D., & Brodbeck, C. et al. (2014). MNE software for processing MEG and EEG data. *Neuroimage, 86,* 446-460. https://doi.org/10.1016/j.neuroimage.2013.10.027

Granger, C. (1969). Investigating Causal Relations by Econometric Models and Cross-spectral Methods. *Econometrica*, *37*(3), 424. doi: 10.2307/1912791

Gregory, R. L. (1980). Perceptions as hypotheses. *Philosophical Transactions of the*

Royal Society of London. B, Biological Sciences, *290*(1038), 181-197.

Grill-Spector, K., Knouf, N., & Kanwisher, N. (2004). The fusiform face area subserves face perception, not generic within-category identification. *Nature Neuroscience, 7*(5), 555-562. https://doi.org/10.1038/nn1224

Grootswagers, T., Wardle, S., & Carlson, T. (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

Groppe, D., Urbach, T., & Kutas, M. (2011). Mass univariate analysis of event-related brain potentials/fields I: A critical tutorial review. *Psychophysiology*, *48*(12), 1711- 1725. https://doi.org/10.1111/j.1469-8986.2011.01273.x

Grossmann, J., & Dobbins, A. (2003). Differential ambiguity reduces grouping of metastable objects. *Vision Research*, *43*(4), 359-369. https://doi.org/10.1016/s0042- 6989(02)00480-7

Gruber, W., Zauner, A., Lechinger, J., Schabus, M., Kutil, R., & Klimesch, W. (2014).

Guckenheimer, J., & Holmes, P. (1983). Local bifurcations. In *Nonlinear oscillations, dynamical systems, and bifurcations of vector fields* (pp. 117-165). Springer, New York, NY.

Hanslmayr, S., Volberg, G., Wimber, M., Dalal, S., & Greenlee, M. (2013). Prestimulus Oscillatory Phase at 7 Hz Gates Cortical Information Flow and Visual

Harrison, S., & Backus, B. (2010). Uninformative visual experience establishes long term perceptual bias. *Vision Research*, *50*(18), 1905-1911. doi: 10.1016/j.visres.2010.06.013

Harrower, M. R. (1936). Some Factors Determining Figure-Ground Articulation. *British Journal of Psychology. General Section, 26*, 407-424. https://doi.org/10.1111/j.2044-8295.1936.tb00806.x

Hasson, U., Harel, M., Levy, I., & Malach, R. (2003). Large-Scale Mirror-Symmetry Organization of Human Occipito-Temporal Object Areas. *Neuron*, *37*(6), 1027- 1041. doi: 10.1016/s0896-6273(03)00144-2

Hasson, U., Hendler, T., Bashat, D., & Malach, R. (2001). Vase or face? A neural correlate of shape-selective grouping processes in the human brain. *Neuroimage*, *13*(6), 889. https://doi.org/10.1016/s1053-8119(01)92231-1

Haxby, J., Hoffman, E., & Gobbini, M. (2000). The distributed human neural system for face perception. *Trends In Cognitive Sciences*, *4*(6), 223-233. doi: 10.1016/s1364- 6613(00)01482-0

Haxby, J., Ida Gobbini, M., Furey, M., Ishai, A., & Pietrini, P. (2001). Distinct, overlapping representations of faces and multiple categories of objects in ventral temporal cortex. *Neuroimage*, *13*(6), 891. doi: 10.1016/s1053-8119(01)92233-5

Haynes, J. (2015). A Primer on Pattern-Based Approaches to fMRI: Principles, Pitfalls, and Perspectives. *Neuron, 87*(2), 257-270.

Haynes, J. D., & Rees, G. (2006). Decoding mental states from brain activity in humans. *Nature reviews neuroscience*, *7*(7), 523-534.

Heekeren, H., Marrett, S., & Ungerleider, L. (2008). The neural systems that mediate human perceptual decision making. *Nature Reviews Neuroscience*, *9*(6), 467-479. doi: 10.1038/nrn2374

Hesselmann, G., Kell, C. A., Eger, E., & Kleinschmidt, A. (2008). Spontaneous local variations in ongoing neural activity bias perceptual decisions. *Proceedings of the National Academy of Sciences, 105*(31), 10984-10989. https://doi.org/10.1073/pnas.0712043105

Hipp, J., Engel, A., & Siegel, M. (2011). Oscillatory Synchronization in Large-Scale Cortical Networks Predicts Perception. *Neuron*, *69*(2), 387-396. doi: 10.1016/j.neuron.2010.12.027

Hochberg, J. (1950). Figure-ground reversal as a function of visual satiation. *Journal Of Experimental Psychology, 40*(5), 682-686. https://doi.org/10.1037/h0060078

Hochberg, J., & Peterson, M. (1987). Piecemeal organization and cognitive components in object perception: Perceptually coupled responses to moving objects. *Journal Of Experimental Psychology: General*, *116*(4), 370-380. https://doi.org/10.1037/0096-

Hock, H., SchÖner, G., & Hochstein, S. (1996). Perceptual Stability and the Selective Adaptation of Perceived and Unperceived Motion Directions. *Vision*

Hodsoll, J., & Humphreys, G. (2001). 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. doi: 10.3758/bf03194447

Hoffman, D., & Richards, W. (1984). Parts of recognition. *Cognition, 18*(1-3), 65-96. https://doi.org/10.1016/0010-0277(84)90022-2

Hramov, A. E., Maksimenko, V. A., Pchelintseva, S. V., Runnova, A. E., Grubov, V. V., Musatov, V. Yu., Zhuravlev, M. O., Koronovskii, A. A., & Pisarchik, A. N. (2017). Classifying the Perceptual Interpretations of a Bistable Image Using EEG and Artificial Neural Networks. Frontiers in Neuroscience, 11.

https://www.frontiersin.org/article/10.3389/fnins.2017.00674

Huk, A. C., & Heeger, D. J. (2002). Pattern-motion responses in human visual cortex. *Nature neuroscience*, *5*(1), 72-75.

Iemi, L., Busch, N., Laudini, A., Haegens, S., Samaha, J., Villringer, A., & Nikulin, V. (2019). Multiple mechanisms link prestimulus neural oscillations to sensory responses. *Elife*, *8*. doi: 10.7554/elife.43620

Ilg, R., Wohlschläger, A., Burazanis, S., Wöller, A., Nunnemann, S., & Mühlau, M. (2008). Neural correlates of spontaneous percept switches in ambiguous stimuli: an event-related functional magnetic resonance imaging study. *European Journal Of Neuroscience*, *28*(11), 2325-2332. doi: 10.1111/j.1460-9568.2008.06520.x

Intaitė, M., Duarte, J., & Castelo-Branco, M. (2016). Working memory load influences perceptual ambiguity by competing for fronto-parietal attentional resources. *Brain Research*, *1650*, 142-151. doi: 10.1016/j.brainres.2016.08.044

Intaite, M., Koivisto, M., & Castelo-Branco, M. (2014). The linear impact of concurrent working memory load on dynamics of Necker cube perceptual reversals. *Journal Of Vision*, *14*(1), 13-13. https://doi.org/10.1167/14.1.13

Intaitė, M., Koivisto, M., & Revonsuo, A. (2013). Perceptual reversals of Necker stimuli during intermittent presentation with limited attentional

Intaitė, M., Koivisto, M., Rukšėnas, O., & Revonsuo, A. (2010). Reversal negativity and bistable stimuli: Attention, awareness, or something else? *Brain and Cognition,* *74*(1), 24-34. https://doi.org/10.1016/j.bandc.2010.06.002

Inui, T., Tanaka, S., Okada, T., Nishizawa, S., Katayama, M., & Konishi, J. (2000).

Neural substrates for depth perception of the Necker cube; a functional magnetic resonance imaging study in human subjects. *Neuroscience Letters*, *282*(3), 145-148. doi: 10.1016/s0304-3940(00)00899-5

Ishai, A., Haxby, J., & Ungerieider, L. (2000). Distributed neural systems for the generation of visual images. *Neuroimage*, *11*(5), S684. doi: 10.1016/s1053- 8119(00)91614-8

İşoğlu-Alkaç, Ü. & Strüber, D. (2006). Necker cube reversals during long-term EEG recordings: sub-bands of alpha activity. *International Journal of Psychophysiology, 59*, 179 - 189. https://doi.org/10.1016/j.ijpsycho.2005.05.002

Isoglu-Alkaç, Ü., Başar-Eroğlu, C., Ademoglu, A., Demiralp T., Miener M., Stadler M. (1998) Analysis of the electroencephalographic activity during the Necker cube reversals by means of the wavelet transform. *Biological Cybernetics, 79*, 437–42. https://doi.org/10.1007/s004220050492

İşoğlu-Alkaç, Ü., Başar-Eroğlu, C., Ademoğlu, A., Demiralp, T., Miener, M., & Stadler, Itier, R., & Taylor, M. (2004). N170 or N1? Spatiotemporal Differences between Object and Face Processing Using ERPs. *Cerebral Cortex*, *14*(2), 132-142. https://doi.org/10.1093/cercor/bhg111

Iversen, J., Repp, B., & Patel, A. (2009). Top-Down Control of Rhythm Perception Modulates Early Auditory Responses. *Annals Of The New York Academy Of Sciences, 1169(*1), 58-73. https://doi.org/10.1111/j.1749-6632.2009.04579.x

Jansen, B., & Brandt, M. (1991). The effect of the phase of prestimulus alpha activity on the averaged visual evoked response. *Electroencephalography And Clinical Neurophysiology/Evoked Potentials Section, 80*(4), 241-250. https://doi.org/10.1016/0168-5597(91)90107-9

Jehee, J., Brady, D., & Tong, F. (2011). Attention Improves Encoding of Task-Relevant Features in the Human Visual Cortex. Journal Of Neuroscience, 31(22), 8210-8219. https://doi.org/10.1523/jneurosci.6153-09.2011

Kanai, R., & Verstraten, F. (2005). Perceptual manifestations of fast neural plasticity: Motion priming, rapid motion aftereffect and perceptual sensitization. *Vision Research*, *45*(25-26), 3109-3116. doi: 10.1016/j.visres.2005.05.014

Kanai, R., Bahrami, B., & Rees, G. (2010). Human Parietal Cortex Structure Predicts Individual Differences in Perceptual Rivalry. *Current Biology*, *20*(18), 1626-1630. https://doi.org/10.1016/j.cub.2010.07.027

Kanai, R., Carmel, D., Bahrami, B., & Rees, G. (2011). Structural and functional fractionation of right superior parietal cortex in bistable perception. *Current Biology*, *21*(3), R106-R107. https://doi.org/10.1016/j.cub.2010.12.009

Kanizsa, G., & Gerbino, W. (1976). Convexity and symmetry in figure-ground organization. In M. Henle (Ed.), *Vision and artifact* (pp. 25–32). New York: Springer.

Kanwisher, N., McDermott, J., & Chun, M. (1997). The Fusiform Face Area: A Module in Human Extrastriate Cortex Specialized for Face Perception. *The Journal Of Neuroscience, 17*(11), 4302-4311. https://doi.org/10.1523/jneurosci.17-11-04302.1997

Kaplan, J., Man, K., & Greening, S. (2015). Multivariate cross-classification: applying machine learning techniques to characterize abstraction in neural representations. *Frontiers In Human Neuroscience*, *9*. doi: 10.3389/fnhum.2015.00151

Kastner, S. (1998). Mechanisms of Directed Attention in the Human Extrastriate Cortex as Revealed by Functional MRI. *Science*, *282*(5386), 108-111. doi: 10.1126/science.282.5386.108

Kehtarnavaz, N. (2008). *Digital signal processing system design: LabView-based hybrid programming* (pp. 175-181, 183-196). Boston: Academic Press.

Keil, J., Müller, N., Hartmann, T., & Weisz, N. (2013). Prestimulus Beta Power and Phase Synchrony Influence the Sound-Induced Flash Illusion. *Cerebral Cortex*, *24*(5), 1278-1288. doi: 10.1093/cercor/bhs409

Kienker, P., Sejnowski, T., Hinton, G., & Schumacher, L. (1986). Separating Figure from Ground with a Parallel Network. *Perception, 15*(2), 197-216. https://doi.org/10.1068/p150197

Kim, C., & Blake, R. (2005). Psychophysical magic: rendering the visible 'invisible'. *Trends In Cognitive Sciences*, *9*(8), 381-388. doi: 10.1016/j.tics.2005.06.012

Kim, S., & Feldman, J. (2009). Globally inconsistent figure/ground relations induced by a negative part. *Journal Of Vision, 9*(10), 8-8. https://doi.org/10.1167/9.10.8

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

Kleinschmidt, A., & Cohen, L. (2006). The neural bases of prosopagnosia and pure alexia: recent insights from functional neuroimaging. *Current Opinion In Neurology, 19*(4), 386-391. https://doi.org/10.1097/01.wco.0000236619.89710.ee

Kleinschmidt, A., Büchel, C., Zeki, S., & Frackowiak, R. (1998). Human brain activity during spontaneously reversing perception of ambiguous figures. *Proceedings Of*

The Royal Society Of London. Series B: Biological Sciences, *265*(1413), 2427-2433. https://doi.org/10.1098/rspb.1998.0594

Kleinschmidt, A., Sterzer, P., & Rees, G. (2012). Variability of perceptual multistability: from brain state to individual trait. *Philosophical Transactions Of The Royal Society B: Biological Sciences*, *367*(1591), 988-1000. doi: 10.1098/rstb.2011.0367

Knapen, T., Brascamp, J., Pearson, J., van Ee, R., & Blake, R. (2011). The Role of Frontal and Parietal Brain Areas in Bistable Perception. *Journal Of Neuroscience*, *31*(28), 10293-10301. doi: 10.1523/jneurosci.1727-11.2011

Knill, D., & Pouget, A. (2004). The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends In Neurosciences*, *27*(12), 712-719. https://doi.org/10.1016/j.tins.2004.10.007

Koffka, K. (1935). Principles of Gestalt psychology. Harcourt, Brace.
Köhler, W. (1940). Dynamics in psychology.

Kogo, N., Kern, F. B., Nowotny, T., Ee, R. van, Wezel, R. van, & Aihara, T. (2021). Dynamics of a Mutual Inhibition Circuit between Pyramidal Neurons Compared to Human Perceptual Competition. Journal of Neuroscience, 41(6), 1251–1264. https://doi.org/10.1523/JNEUROSCI.2503-20.2020

Köhler, W. (1940). Dynamics in psychology (p. 158). Liveright.

Köhler, W., & Held, R. (1949). The cortical correlate of pattern vision. *Science, 110*(2860), 414-419.

Kondacs, A., & Szabó, M. (1999). Long-term intra-individual variability of the background EEG in normals. *Clinical Neurophysiology, 110*(10), 1708-1716. https://doi.org/10.1016/s1388-2457(99)00122-4

Kornmeier, J., & Bach, M. (2004a). Early neural activity in Necker‐cube reversal: Evidence for low‐level processing of a gestalt phenomenon. *Psychophysiology, 41*(1), 1-8. https://doi.org/10.1046/j.1469-8986.2003.00126.x

Kornmeier, J., & Bach, M. (2005). The Necker cube—an ambiguous figure disambiguated in early visual processing. *Vision research, 45*(8), 955-960. https://doi.org/10.1016/j.visres.2004.10.006

Kornmeier, J., & Bach, M. (2006). Bistable perception—along the processing chain from ambiguous visual input to a stable percept. *International Journal of Psychophysiology, 62*(2), 345-349. https://doi.org/10.1016/j.ijpsycho.2006.04.007

Kornmeier, J., & Bach, M. (2012). Ambiguous figures–what happens in the brain when perception changes but not the stimulus. *Frontiers in Human Neuroscience*, 6, 51. https://doi.org/10.3389/fnhum.2012.00051

Kornmeier, J., & Bach, M. (2014). EEG correlates of perceptual reversals in Boring's ambiguous old/young woman stimulus. *Perception, 43*(9), 950-962. https://doi.org/10.1068/p7741

Kornmeier, J., & Michael, B. (2004b). Evidence for early visual processing in perceptual disambiguation of ambiguous figures. *Journal Of Vision, 4*(8), 249-249. https://doi.org/10.1167/4.8.249

Kornmeier, J., Ehm, W., Bigalke, H., & Bach, M. (2007). Discontinuous presentation of ambiguous figures: How interstimulus‐interval durations affect reversal dynamics and ERPs. *Psychophysiology, 44*(4), 552-560. https://doi.org/10.1111/j.1469-

Kornmeier, J., Hein, C., & Bach, M. (2009). Multistable perception: When bottom-up and top-down coincide. *Brain and Cognition, 69*(1), 138-147. https://doi.org/10.1016/j.bandc.2008.06.005

Kornmeier, J., Pfäffle, M., & Bach, M. (2011). Necker cube: Stimulus-related (low-level) and percept-related (high-level) EEG signatures early in occipital cortex. *Journal of vision, 11*(9), 12-12. https://doi.org/10.1167/11.9.12

Kosslyn, S., Thompson, W., Klm, I., & Alpert, N. (1995). Topographical representations of mental images in primary visual cortex. *Nature*, *378*(6556), 496-498. doi: 10.1038/378496a0

Kourtzi, Z., & Kanwisher, N. (2000). Cortical Regions Involved in Perceiving Object Shape. *The Journal Of Neuroscience*, *20*(9), 3310-3318. doi: 10.1523/jneurosci.20- 09-03310.2000

Kubovy, M., & Pomerantz, J. R. (Eds.). (1981/2017). *Perceptual organization* (Vol. 16).

Kumar, S., Soto, D., & Humphreys, G. (2009). Electrophysiological evidence for attentional guidance by the contents of working memory. *European Journal Of Neuroscience*, *30*(2), 307-317. doi: 10.1111/j.1460-9568.2009.06805.x

LaBar, K., Gitelman, D., Parrish, T., & Mesulam, M. (1999). Neuroanatomic Overlap of Working Memory and Spatial Attention Networks: A Functional MRI Comparison within Subjects. *Neuroimage*, *10*(6), 695-704. doi: 10.1006/nimg.1999.0503

Lack, L. (1974). Selective attention and the control of binocular rivalry. *Perception &*

Psychophysics, *15*(1), 193-200. doi: 10.3758/bf03205846

Lago-Fernández, L., & Deco, G. (2002). A model of binocular rivalry based on competition in IT. *Neurocomputing*, *44-46*, 503-507. https://doi.org/10.1016/s0925- 2312(02)00408-3

Leopold, D., & Logothetis, N. (1999). Multistable phenomena: changing views in perception. *Trends in Cognitive Sciences,* 3(7), 254-264. https://doi.org/10.1016/s1364-6613(99)01332-7

Leopold, D., Wilke, M., Maier, A., & Logothetis, N. (2002). Stable perception of visually ambiguous patterns. *Nature Neuroscience*, 5(6), 605-609. https://doi.org/10.1038/nn0602-851

Liebert, R., & Burk, B. (1985). Voluntary Control of Reversible Figures. *Perceptual And Motor Skills*, *61*(3_suppl), 1307-1310. doi: 10.2466/pms.1985.61.3f.1307

Linkenkaer-Hansen, K., Nikulin, V., Ilmoniemi, R., & Palva, J. (2004). Prestimulus Oscillations Enhance Psychophysical Performance in Humans. *Journal Of Neuroscience*, 24(45), 10186-10190. https://doi.org/10.1523/jneurosci.2584-

List, A., Rosenberg, M., Sherman, A., & Esterman, M. (2017). Pattern classification of EEG signals reveals perceptual and attentional states. *PLOS ONE, 12*(4), e0176349. https://doi.org/10.1371/journal.pone.0176349

Long, G. M., & Olszweski, A. D. (1999). To Reverse or Not to Reverse: When Is an Ambiguous Figure Not Ambiguous? The American Journal of Psychology, 112(1), 41–71. https://doi.org/10.2307/1423624

Long, G. M., & Toppino, T. C. (1981). Multiple representations of the same reversible figure: Implications for cognitive decisional interpretations. *Perception, 10*(2), 231- 234.

Long, G. M., & Toppino, T. C. (2004). Enduring interest in perceptual ambiguity: Alternating views of reversible figures. Psychological Bulletin, 130(5), 748–768. https://doi.org/10.1037/0033-2909.130.5.748

Long, G. M., Toppino, T. C., & Kostenbauder, J. F. (1983). As the cube turns: Evidence for two processes in the perception of a dynamic reversible figure. *Perception & Psychophysics, 34*(1), 29-38.

Long, G. M., Toppino, T. C., & Mondin, G. W. (1992). Prime time: Fatigue and set effects in the perception of reversible figures. Perception & Psychophysics, 52(6), 609–616. https://doi.org/10.3758/BF03211697

Long, G., & Moran, C. (2007). How to Keep a Reversible Figure from Reversing: Teasing Out Top — Down and Bottom — Up Processes. *Perception*, *36*(3), 431-445. https://doi.org/10.1068/p5630

Long, G., & Olszweski, A. (1999). To Reverse or Not to Reverse: When Is an Ambiguous Figure Not Ambiguous?. *The American Journal Of Psychology*, *112*(1),41. doi: 10.2307/1423624

Long, G., & Toppino, T. (2004). Enduring Interest in Perceptual Ambiguity: Alternating Views of Reversible Figures. *Psychological Bulletin*, *130*(5), 748-768. doi: 10.1037/0033-2909.130.5.748

Long, G., Toppino, T., & Mondin, G. (1992). Prime time: Fatigue and set effects in the perception of reversible figures. *Perception & Psychophysics*, *52*(6), 609-616. doi: 10.3758/bf03211697

Luck, S. (2014). An introduction to the event-related potential technique (2nd ed.). The MIT Press.

Luck, S. J., & Vecera, S. P. (2002). Attention. In H. Pashler & S. Yantis (Eds.), Steven's handbook of experimental psychology: Sensation and perception (p. 235–286).John Wiley & Sons Inc.

Luck, S., & Hillyard, S. (2000). The operation of selective attention at multiple stages of processing: Evidence from human and monkey electrophysiology. In T*he New Cognitive Neurosciences* (pp. 687-700). M. S. Gazzaniga. Retrieved 18 February 2020, from.

Lumer, E., & Rees, G. (1999). Covariation of activity in visual and prefrontal cortex associated with subjective visual perception. *Proceedings Of The National Academy Of Sciences*, *96*(4), 1669-1673. https://doi.org/10.1073/pnas.96.4.1669

Lumer, E., Friston, K., & Rees, G. (1998). Neural Correlates of Perceptual Rivalry in the Human Brain. *Science*, *280*(5371), 1930-1934. https://doi.org/10.1126/science.280.5371.1930

Macey, P., Macey, K., Kumar, R., & Harper, R. (2004). A method for removal of global effects from fMRI time series. *Neuroimage*, *22*(1), 360-366. doi: 10.1016/j.neuroimage.2003.12.042

Maffei, L., Fiorentini, A., & Bisti, S. (1973). Neural Correlate of Perceptual Adaptation to Gratings. *Science*, *182*(4116), 1036-1038. doi: 10.1126/science.182.4116.1036

Maier, A., Wilke, M., Logothetis, N., & Leopold, D. (2003). Perception of Temporally Interleaved Ambiguous Patterns. *Current Biology, 13*(13), 1076-1085. https://doi.org/10.1016/s0960-9822(03)00414-7

Maris, E., & Oostenveld, R. (2007). Nonparametric statistical testing of EEG- and MEG-data. *Journal Of Neuroscience Methods*, *164*(1), 177-190. doi:

10.1016/j.jneumeth.2007.03.024

Marti, S., & Dehaene, S. (2017). Discrete and continuous mechanisms of temporal selection in rapid visual streams. *Nature Communications, 8*(1). https://doi.org/10.1038/s41467-017-02079-x

Martinez, A., Ramanathan, D., Foxe, J., Javitt, D., & Hillyard, S. (2007). The Role of Spatial Attention in the Selection of Real and Illusory Objects. *Journal Of Neuroscience*, *27*(30), 7963-7973. doi: 10.1523/jneurosci.0031-07.2007

Mathes, B., Pomper, U., Walla, P., & Basar-Eroglu, C. (2010). Dissociation of reversal- and motor-related delta- and alpha-band responses during visual multistable perception. *Neuroscience Letters*, *478*(1), 14-18. doi: 10.1016/j.neulet.2010.04.057

Mathes, B., Strüber, D., Stadler, M. A., & Basar-Eroglu, C. (2006). Voluntary control of Necker cube reversals modulates the EEG delta-and gamma-band response.*Neuroscience Letters, 402*(1-2), 145-149. https://doi.org/10.1016/j.neulet.2006.03.063

Mathewson, K., Gratton, G., Fabiani, M., Beck, D., & Ro, T. (2009). To See or Not to See: Prestimulus Phase Predicts Visual Awareness. *Journal Of Neuroscience, 29*(9), 2725-2732. https://doi.org/10.1523/jneurosci.3963-08.2009

McCarthy, G. (1995). Review : Functional Neuroimaging of Memory. *The Neuroscientist*, *1*(3), 155-163. doi: 10.1177/107385849500100306

Mechelli, A., Price, C., Friston, K., & Ishai, A. (2004). Where Bottom-up Meets Top-down: Neuronal Interactions during Perception and Imagery. *Cerebral Cortex*, *14*(11), 1256-1265. https://doi.org/10.1093/cercor/bhh087

Meng, M., & Tong, F. (2004). Can attention selectively bias bistable perception? Differences between binocular rivalry and ambiguous figures. *Journal of vision*, *4*(7), 2-2.Mensen, A., & Khatami, R. (2013). Advanced EEG analysis using threshold-free cluster-enhancement and non-parametric statistics. *Neuroimage, 67*, 111-118. https://doi.org/10.1016/j.neuroimage.2012.10.027

Meredith, G., & Meredith, C. (1962). Effect of Instructional Conditions on Rate of Binocular Rivalry. *Perceptual And Motor Skills*, *15*(3), 655-664. doi: 10.2466/pms.1962.15.3.655

Moreno-Bote, R., Rinzel, J., & Rubin, N. (2007). Noise-Induced Alternations in an Attractor Network Model of Perceptual Bistability. *Journal Of Neurophysiology*, *98*(3), 1125-1139. https://doi.org/10.1152/jn.00116.2007

Muckli, L. (2010). What are we missing here? Brain imaging evidence for higher cognitive functions in primary visual cortex V1. *International Journal Of Imaging Systems And Technology*, *20*(2), 131-139. https://doi.org/10.1002/ima.20236

Mumford, D. (1992). On the computational architecture of the neocortex. *Biological Cybernetics*, *66*(3), 241-251. doi: 10.1007/bf00198477

Murphy, A., Leopold, D., & Welchman, A. (2014). Perceptual memory drives learning of retinotopic biases for bistable stimuli. *Frontiers In Psychology*, *5*. https://doi.org/10.3389/fpsyg.2014.00060

Nakatani, H., & van Leeuwen, C. (2006). Transient Synchrony of Distant Brain Areas and Perceptual Switching in Ambiguous Figures. *Biological Cybernetics*, *94*(6), 445-457. doi: 10.1007/s00422-006-0057-9

Naselaris, T., Kay, K., Nishimoto, S., & Gallant, J. (2011). Encoding and decoding in fMRI. *Neuroimage, 56*(2), 400-410. https://doi.org/10.1016/j.neuroimage.2010.07.073

Nawrot, M., & Blake, R. (1993). On the perceptual identity of dynamic stereopsis and kinetic depth. *Vision Research*, *33*(11), 1561-1571. doi: 10.1016/0042-6989(93)90149-q

Necker, L. (1832). LXI. Observations on some remarkable optical phænomena seen in Switzerland; and on an optical phænomenon which occurs on viewing a figure of a crystal or geometrical solid. *The London, Edinburgh, And Dublin Philosophical Magazine And Journal Of Science, 1*(5), 329-337. https://doi.org/10.1080/14786443208647909

Necker, L. A. (1839). LXI. Observations on some remarkable optical phænomena seen in Switzerland; and on an optical phænomenon which occurs on viewing a figure of a crystal or geometrical solid. The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science. https://doi.org/10.1080/14786443208647909

Nichols, T. (2012). Multiple testing corrections, nonparametric methods, and random field theory. *Neuroimage, 62*(2), 811-815. https://doi.org/10.1016/j.neuroimage.2012.04.014

Noest, A. J., van Ee, R., Nijs, M. M., & van Wezel, R. J. A. (2007). Percept-choice sequences driven by interrupted ambiguous stimuli: A low-level neural model. Journal of Vision, 7(8), 10. https://doi.org/10.1167/7.8.10

Noest, A., van Ee, R., Nijs, M., & van Wezel, R. (2007). Percept-choice sequences driven by interrupted ambiguous stimuli: A low-level neural model. *Journal Of Vision*, *7*(8), 10. https://doi.org/10.1167/7.8.10

Norman, K., Polyn, S., Detre, G., & Haxby, J. (2006). Beyond mind-reading: multi-voxel pattern analysis of fMRI data. *Trends In Cognitive Sciences, 10*(9), 424-430. https://doi.org/10.1016/j.tics.2006.07.005

Nunn, C., & Osselton, J. (1974). The Influence of the EEG Alpha Rhythm on the Perception of Visual Stimuli. *Psychophysiology*, *11*(3), 294-303. https://doi.org/10.1111/j.1469-8986.1974.tb00547.x

O'Donnell, B. F., Hendler, T., Squires, N. K. (1988). Visual evoked potentials to illusory reversals of the Necker cube. *Psychophysiology, 25*, 137–143 https://doi.org/10.1111/j.1469-8986.1988.tb00976.x

O'Craven, K., & Kanwisher, N. (2000). Mental Imagery of Faces and Places Activates Corresponding Stimulus-Specific Brain Regions. *Journal Of Cognitive Neuroscience*, *12*(6), 1013-1023. doi: 10.1162/08989290051137549

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

Okazaki, M., Kaneko, Y., Yumoto, M., & Arima, K. (2008). Perceptual change in response to a bistable picture increases neuromagnetic beta-band activities. *Neuroscience Research*, *61*(3), 319-328. doi: 10.1016/j.neures.2008.03.010

Oostenveld, R., Fries, P., Maris, E., & Schoffelen, J. (2011). FieldTrip: Open Source Software for Advanced Analysis of MEG, EEG, and Invasive

Electrophysiological Data. *Computational Intelligence And Neuroscience*, *2011*, 1-9. doi: 10.1155/2011/156869

Oosterhof, N., Connolly, A., & Haxby, J. (2016). CoSMoMVPA: Multi-Modal Multivariate Pattern Analysis of Neuroimaging Data in Matlab/GNU Octave. *Frontiers In Neuroinformatics, 10*. https://doi.org/10.3389/fninf.2016.00027

Orbach, J., Ehrlich, D., & Heath, H. (1963). Reversibility of the Necker Cube: I. An Examination of the Concept of "Satiation of Orientation". *Perceptual And Motor Skills, 17*(2), 439-458. https://doi.org/10.2466/pms.1963.17.2.439

Orbach, J., Zucker, E., & Olson, R. (1966). Reversibility of the Necker Cube: VII. Reversal Rate as a Function of Figure-on and Figure-off Durations. *Perceptual And Motor Skills, 22*(2), 615-618. https://doi.org/10.2466/pms.1966.22.2.615

Paffen, C., Alais, D., & Verstraten, F. (2006). Attention Speeds Binocular Rivalry. *Psychological Science*, *17*(9), 752-756. https://doi.org/10.1111/j.1467-9280.2006.01777.x

Park, H., & Friston, K. (2013). Structural and Functional Brain Networks: From Connections to Cognition. *Science*, *342*(6158), 1238411-1238411. https://doi.org/10.1126/science.1238411

Pastukhov, A. (2017). First, you need a Gestalt: An interaction of bottom-up and top-down streams during the perception of the ambiguously rotating human walker. *Scientific Reports*, *7*(1). https://doi.org/10.1038/s41598-017-01376-1

Pastukhov, A., & Braun, J. (2007). Perceptual reversals need no prompting by attention. *Journal Of Vision*, *7*(10), 5. https://doi.org/10.1167/7.10.5

Pastukhov, A., & Braun, J. (2008). A short-term memory of multi-stable perception. *Journal Of Vision*, *8*(13), 7-7. doi: 10.1167/8.13.7

Pastukhov, A., Styrnal, M., & Carbon, C.-C. (2023). History-dependent changes to distribution of dominance phases in multistable perception. Journal of Vision, 23(3), 16. https://doi.org/10.1167/jov.23.3.16

Pazo-Álvarez, P., Roca-Fernández, A., Gutiérrez-Domínguez, F. J., & Amenedo, E. (2017). Attentional Modulation of Change Detection ERP Components by Peripheral Retro- Cueing. *Frontiers in human neuroscience, 11*, 76 https://doi.org/10.3389/fnhum.2017.00076

Pearson, J., & Brascamp, J. (2008). Sensory memory for ambiguous vision. *Trends In Cognitive Sciences*, *12*(9), 334-341. doi: 10.1016/j.tics.2008.05.006

Pearson, J., & Clifford, C. (2005). Mechanisms selectively engaged in rivalry: normal vision habituates, rivalrous vision primes. *Vision Research*, *45*(6), 707-714. https://doi.org/10.1016/j.visres.2004.09.040

Peirce, J. W., Gray, J. R., Simpson, S., MacAskill, M. R., Höchenberger, R., Sogo, H., Kastman, E., Lindeløv, J. (2019). PsychoPy2: experiments in behavior made easy. *Behavior Research Methods, 51*, 195–203. 10.3758/s13428-018-01193-y

Pereira, F., Mitchell, T., & Botvinick, M. (2009). Machine learning classifiers and fMRI: A tutorial overview. *Neuroimage, 45*(1), S199-S209. https://doi.org/10.1016/j.neuroimage.2008.11.007

Pernet, C., Latinus, M., Nichols, T., & Rousselet, G. (2015). Cluster-based computational methods for mass univariate analyses of event-related brain potentials/fields: A simulation study. *Journal Of Neuroscience Methods*, *250*, 85-93. doi: 10.1016/j.jneumeth.2014.08.003

Perry, G., & Singh, K. (2014). Localizing evoked and induced responses to faces using magnetoencephalography. *European Journal Of Neuroscience, 39*(9), 1517-1527. https://doi.org/10.1111/ejn.12520

Pessoa, L., & Ungerleider, L. (2004). Neural Correlates of Change Detection and Change Blindness in a Working Memory Task. *Cerebral Cortex*, *14*(5), 511-520. doi: 10.1093/cercor/bhh013

Peterson, M. (2014). Low-level and high-level contributions to figure-ground organization. In J. Wagemans, *Oxford Handbook of Perceptual Organization*. Oxford University Press.

Peterson, M. A. (2001). Object perception. In E. B. Goldstein (Ed.), *Blackwell Handbook of Perception*, Chapter 6, pp. 168-203. Oxford: Blackwell Publishers.

Philiastides, M., & Sajda, P. (2006a). Temporal Characterization of the Neural Correlates of Perceptual Decision Making in the Human Brain. *Cerebral Cortex, 16*(4), 509-518. https://doi.org/10.1093/cercor/bhi130

Philiastides, M., Ratcliff, R., & Sajda, P. (2006b). Neural Representation of Task Difficulty and Decision Making during Perceptual Categorization: A Timing Diagram. *Journal Of Neuroscience, 26*(35), 8965-8975. https://doi.org/10.1523/jneurosci.1655-06.2006

Piantoni, G., Romeijn, N., Gomez-Herrero, G., Van Der Werf, Y., & Van Someren, E. (2017). Alpha Power Predicts Persistence of Bistable Perception. *Scientific Reports*, *7*(1). https://doi.org/10.1038/s41598-017-05610-8

Picton, T. (1992). The P300 Wave of the Human Event-Related Potential. *Journal Of Clinical Neurophysiology*, *9*(4), 456-479. doi: 10.1097/00004691-199210000-

00002

Pinkus, A., & Pantle, A. (1997). Probing Visual Motion Signals with a Priming Paradigm. Vision Research, 37(5), 541–552. https://doi.org/10.1016/S0042-6989(96)00162-9

Pitts, M. A., & Britz, J. (2011). Insights from intermittent binocular rivalry and EEG. *Frontiers in Human Neuroscience, 5*, 107. https://doi.org/10.3389/fnhum.2011.00107

Pitts, M. A., Gavin, W. J., & Nerger, J. L. (2008). Early top-down influences on bistable perception revealed by event-related potentials. *Brain and Cognition, 67*(1), 11-24. https://doi.org/10.1016/j.bandc.2007.10.004

Pitts, M., Martínez, A., Brewer, J., & Hillyard, S. (2011). Early Stages of Figure–Ground Segregation during Perception of the Face–Vase. *Journal Of Cognitive Neuroscience, 23*(4), 880-895. https://doi.org/10.1162/jocn.2010.21438

Pitts, M., Martínez, A., Stalmaster, C., Nerger, J., & Hillyard, S. (2009). Neural generators of ERPs linked with Necker cube reversals. *Psychophysiology*, *46*(4), 694-702. doi: 10.1111/j.1469-8986.2009.00822.x

Pitts, M., Nerger, J., & Davis, T. (2007). Electrophysiological correlates of perceptual reversals for three different types of multistable images. *Journal Of Vision, 7*(1), 6. https://doi.org/10.1167/7.1.6

Pizlo, Z. (2001). Perception viewed as an inverse problem. *Vision Research*, *41*(24), 3145-3161. https://doi.org/10.1016/s0042-6989(01)00173-0

Plöchl, M., Ossandón, J., & König, P. (2012). Combining EEG and eye tracking: identification, characterization, and correction of eye movement artifacts in electroencephalographic data. *Frontiers In Human Neuroscience*, *6*. doi: 10.3389/fnhum.2012.00278

Pollmann, S., & Yves von Cramon, D. (2000). Object working memory and visuospatial processing: functional neuroanatomy analyzed by event-related fMRI. *Experimental Brain Research*, *133*(1), 12-22. doi: 10.1007/s002210000396

Pomerantz, J. R., & Kubovy, M. (1986). Theoretical approaches to perceptual organization: Simplicity and likelihood principles. In K. R. Boff, L. Kaufman, & J.P. Thomas (Eds.), *Handbook of perception and human performance, Vol. 2. Cognitive processes and performance* (p. 1–46). John Wiley & Sons.

Potts, G.F. (2004). An ERP index of task relevance evaluation of visual stimuli. *Brain and Cognition, 56*(1), 5-13. https://doi.org/10.1016/j.bandc.2004.03.006

Pratte, M., Ling, S., Swisher, J., & Tong, F. (2013). How attention extracts objects from noise. *Journal Of Neurophysiology, 110*(6), 1346-1356. https://doi.org/10.1152/jn.00127.2013

Pritchard, R. (1958). Visual Illusions Viewed as Stabilized Retinal Images. *Quarterly Journal Of Experimental Psychology*, *10*(2), 77-81. doi: 10.1080/17470215808416259

Purves, D., & Andrews, T. (1997). The perception of transparent three-dimensional objects. *Proceedings Of The National Academy Of Sciences, 94*(12), 6517-6522. https://doi.org/10.1073/pnas.94.12.6517

Rahn, E., & Basar, E. (1993). Prestimulus Eeg-Activity Strongly Influences the Auditory Evoked Vertex Response: A New Method for Selective Averaging. *International Journal Of Neuroscience, 69*(1-4), 207-220. https://doi.org/10.3109/00207459309003331

Rao, R., & Ballard, D. (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

Rassi, E., Wutz, A., Müller-Voggel, N., & Weisz, N. (2019). Prestimulus feedback connectivity biases the content of visual experiences. *Proceedings Of The National Academy Of Sciences, 116*(32), 16056-16061. https://doi.org/10.1073/pnas.1817317116

Rassi, E., Wutz, A., Peatfield, N., & Weisz, N. (2022). Efficient prestimulus network integration of fusiform face area biases face perception during binocular rivalry. *Journal of Cognitive Neuroscience*, *34*(6), 1001-1014.

Ratcliff, R., Philiastides, M., & Sajda, P. (2009). Quality of evidence for perceptual decision making is indexed by trial-to-trial variability of the EEG. *Proceedings Of The National Academy Of Sciences, 106*(16), 6539-6544. https://doi.org/10.1073/pnas.0812589106

Regan, D. (1989). Human brain electrophysiology. Evoked potentials and evoked magnetic fields in science and medicine. *Pain, 39*(3), 371-372. https://doi.org/10.1016/0304-3959(89)90058-4

Reisberg, D. (1983). General mental resources and perceptual judgments. *Journal Of Experimental Psychology: Human Perception And Performance*, *9*(6), 966-979. doi: 10.1037/0096-1523.9.6.966

Reisberg, D., & O'Shaughnessy, M. (1984). Diverting Subjects' Concentration Slows Figural Reversals. *Perception*, *13*(4), 461-468. doi: 10.1068/p130461

Rensink, R.A. (2002). Change Detection. *Annual Review of Psychology, 53*, 245-277.

https://doi.org/10.1146/annurev.psych.53.100901.135125

Richter, C., Coppola, R., & Bressler, S. (2018). Top-down beta oscillatory signaling conveys behavioral context in early visual cortex. *Scientific Reports, 8*(1). https://doi.org/10.1038/s41598-018-25267-1

Ritchie, J., Tovar, D., & Carlson, T. (2015). Emerging Object Representations in the Visual System Predict Reaction Times for Categorization. *PLOS Computational Biology, 11*(6), e1004316. https://doi.org/10.1371/journal.pcbi.1004316

Rock, I., & Mitchener, K. (1992). Further Evidence of Failure of Reversal of Ambiguous Figures by Uninformed Subjects. *Perception*, *21*(1), 39-45. doi: 10.1068/p210039

Rock, I., Hall, S., & Davis, J. (1994). Why do ambiguous figures reverse?. *Acta Psychologica*, *87*(1), 33-59. doi: 10.1016/0001-6918(94)90065-5

Ronconi, L., Oosterhof, N., Bonmassar, C., & Melcher, D. (2017). Multiple oscillatory rhythms determine the temporal organization of perception. *Proceedings Of The National Academy Of Sciences, 114*(51), 13435-13440. https://doi.org/10.1073/pnas.1714522114

Rossion, B., Joyce, C., Cottrell, G., & Tarr, M. (2003). Early lateralization and orientation tuning for face, word, and object processing in the visual cortex. *Neuroimage*, *20*(3), 1609-1624.

Rubin, E. (1958/1915). Figure and ground. In D. C. Beardslee & M. Wertheimer (Eds.), Readings in perception (pp. 194–203). Princeton, NJ: Van Nostrand.

Sadaghiani, S., Hesselmann, G., Friston, K., & Kleinschmidt, A. (2010). The relation of ongoing brain activity, evoked neural responses, and cognition. *Frontiers In Systems Neuroscience*. doi: 10.3389/fnsys.2010.00020

Salimi-Khorshidi, G., Smith, S., & Nichols, T. (2011). Adjusting the effect of nonstationarity in cluster-based and TFCE inference. *Neuroimage*, *54*(3), 2006-2019. doi: 10.1016/j.neuroimage.2010.09.088

Schiller, P. V. (1933). Stroboskopische Alternativversuche. *Psychologische Forschung, 17,* 179–214.

Schwartz, J., Grimault, N., Hupe, J., Moore, B., & Pressnitzer, D. (2012). Multistability in perception: binding sensory modalities, an overview. *Philosophical Transactions of The Royal Society B: Biological Sciences, 367*(1591), 896-905. https://doi.org/10.1098/rstb.2011.0254

Sewards, T., & Sewards, M. (1999). Alpha-band oscillations in visual cortex: part of the neural correlate of visual awareness?. *International Journal Of Psychophysiology, 32*(1), 35-45. https://doi.org/10.1016/s0167-8760(98)00062-2

Shepard, R. N. (1990). *Mind sights: Original visual illusions, ambiguities, and other anomalies, with a commentary on the play of mind in perception and art*. WH Freeman/Times Books/Henry Holt & Co.

Shpiro, A., Moreno-Bote, R., Rubin, N., & Rinzel, J. (2009). Balance between noise and adaptation in competition models of perceptual bistability. *Journal Of Computational Neuroscience*, *27*(1), 37-54. https://doi.org/10.1007/s10827-008-0125-3

Singhal, A., & Fowler, B. (2004). The differential effects of Sternberg short- and long-term memory scanning on the late Nd and P300 in a dual-task paradigm. *Cognitive Brain Research*, *21*(1), 124-132. doi:

10.1016/j.cogbrainres.2004.06.003

Skrandies, W. (1990). Global field power and topographic similarity. *Brain Topography*, *3*(1), 137-141. doi: 10.1007/bf01128870

Slotnick, S., & Yantis, S. (2005). Common neural substrates for the control and effects of visual attention and perceptual bistability. *Cognitive Brain Research*, *24*(1), 97-108. https://doi.org/10.1016/j.cogbrainres.2004.12.008

Slotnick, S., Thompson, W., & Kosslyn, S. (2005). Visual Mental Imagery Induces Retinotopically Organized Activation of Early Visual Areas. *Cerebral Cortex*, *15*(10), 1570-1583. doi: 10.1093/cercor/bhi035

Smith, S., & Nichols, T. (2009). Threshold-free cluster enhancement: Addressing problems of smoothing, threshold dependence and localisation in cluster inference. *Neuroimage, 44*(1), 83-98. https://doi.org/10.1016/j.neuroimage.2008.03.061

Soto, D., Hodsoll, J., Rotshtein, P., & Humphreys, G. (2008). Automatic guidance of attention from working memory. *Trends In Cognitive Sciences*, *12*(9), 342-348. doi: 10.1016/j.tics.2008.05.007

Sperling, G., Reeves, A., & Lu, E. B. Z. L. (2001). 10 Two Computational Models of Attention. *Visual attention and cortical circuits*, 177.

Spitz, H., & Lipman, R. (1962). Some Factors Affecting Necker Cube Reversal Rate. *Perceptual And Motor Skills*, *15*(3), 611-625. doi: 10.2466/pms.1962.15.3.611

Sterzer, P., & Kleinschmidt, A. (2007). A neural basis for inference in perceptual ambiguity. *Proceedings Of The National Academy Of Sciences, 104*(1), 323-328.

https://doi.org/10.1073/pnas.0609006104

Sterzer, P., & Rees, G. (2008). A Neural Basis for Percept Stabilization in Binocular Rivalry. *Journal Of Cognitive Neuroscience, 20*(3), 389-399. https://doi.org/10.1162/jocn.2008.20039

Sterzer, P., Kleinschmidt, A., & Rees, G. (2009). The neural bases of multistable perception. *Trends In Cognitive Sciences*, *13*(7), 310-318. doi: 10.1016/j.tics.2009.04.006

Sterzer, P., Russ, M., Preibisch, C., & Kleinschmidt, A. (2002). Neural Correlates of Spontaneous Direction Reversals in Ambiguous Apparent Visual

Strüber, D., & Herrmann, C. S. (2002). MEG alpha activity decrease reflects destabilization of multistable percepts. *Cognitive Brain Research, 14*(3), 370-382. https://doi.org/10.1016/S0926-6410(02)00139-8

Strüber, D., & Stadler, M. (1999). Differences in Top—Down Influences on the Reversal Rate of Different Categories of Reversible Figures. *Perception*, *28*(10), 1185-1196. doi: 10.1068/p2973

Strüber, D., Baar-Eroglu, C., Miener, M., & Stadler, M. (2001). EEG gamma-band response during the perception of Necker cube reversals. *Visual Cognition*, *8*(3-5), 609-621. https://doi.org/10.1080/13506280143000151

Strüber, D., Basar-Eroglu, C., Hoff, E., & Stadler, M. (2000). Reversal-rate dependent differences in the EEG gamma-band during multistable visual perception. *International Journal Of Psychophysiology, 38*(3), 243-252. https://doi.org/10.1016/s0167-8760(00)00168-9

Strüber, D., Basar-Eroglu, C., Miener, M., & Stadler, M. (2001). EEG gamma-band response during the perception of Necker cube reversals. *Visual Cognition, 8*(3-5), 609-621. https://doi.org/10.1080/13506280143000151

Sundareswara, R.1. & Schrater, P.R. (2008). Perceptual multistability predicted by search model for Bayesian decisions. *Journal of Vision, 8*, 12. https://doi.org/10.1167/8.5.12

Supèr, H., van der Togt, C., Spekreijse, H., & Lamme, V. (2003). Internal State of Monkey Primary Visual Cortex (V1) Predicts Figure–Ground Perception. *The Journal Of Neuroscience, 23*(8), 3407-3414. https://doi.org/10.1523/jneurosci.23-08-03407.2003

Suzuki, S., & Peterson, M. (2000). Multiplicative Effects of Intention on the Perception of Bistable Apparent Motion. *Psychological Science*, *11*(3), 202-209. doi: 10.1111/1467-9280.00242

Takeuchi, T., Tuladhar, A., & Yoshimoto, S. (2011). The effect of retinal illuminance on visual motion priming. Vision Research, 51(10), 1137–1145. https://doi.org/10.1016/j.visres.2011.03.002

Taylor, M., McCarthy, G., Saliba, E., & Degiovanni, E. (1999). ERP evidence of developmental changes in processing of faces. *Clinical Neurophysiology*, *110*(5), 910-915. doi: 10.1016/s1388-2457(99)00006-1

Thiery, T., Lajnef, T., Jerbi, K., Arguin, M., Aubin, M., & Jolicoeur, P. (2016). Decoding the Locus of Covert Visuospatial Attention from EEG Signals. *PLOS ONE, 11*(8), e0160304. https://doi.org/10.1371/journal.pone.0160304

Thompson, P., & Burr, D. (2009). Visual aftereffects. *Current Biology*, *19*(1), R11-R14. doi: 10.1016/j.cub.2008.10.014

Tong, F., & Pratte, M. (2012). Decoding Patterns of Human Brain Activity. *Annual Review Of Psychology, 63*(1), 483-509. https://doi.org/10.1146/annurev-psych-120710-100412

Tong, F., Meng, M., & Blake, R. (2006). Neural bases of binocular rivalry. *Trends In Cognitive Sciences*, *10*(11), 502-511. doi: 10.1016/j.tics.2006.09.003

Tong, F., Nakayama, K., Vaughan, J., & Kanwisher, N. (1998). Binocular Rivalry and Visual Awareness in Human Extrastriate Cortex. *Neuron, 21*(4), 753-759. https://doi.org/10.1016/s0896-6273(00)80592-9

Toppino T.C., Long G.M. (2005) Top-Down and Bottom-Up Processes in the Perception of Reversible Figures: Toward a Hybrid Model. In: Ohta N., MacLeod C.M., Uttl B. (eds) Dynamic Cognitive Processes. Springer, Tokyo

Toppino, T. (2003). Reversible-figure perception: Mechanisms of intentional control. *Perception & Psychophysics*, *65*(8), 1285-1295. https://doi.org/10.3758/bf03194852

Toppino, T., & Long, G. (1987). Selective adaptation with reversible figures: Don't change that channel. *Perception & Psychophysics*, *42*(1), 37-48. doi: 10.3758/bf03211512

Toppino, T., & Long, G. (2014). Time for a change: What dominance durations reveal about adaptation effects in the perception of a bi-stable reversible figure. *Attention, Perception, & Psychophysics*, *77*(3), 867-882. https://doi.org/10.3758/s13414-014- 0809-x

Treder, M. S. (2020). MVPA-Light: A Classification and Regression Toolbox for Multi-Dimensional Data. Frontiers in Neuroscience, 14.

https://www.frontiersin.org/article/10.3389/fnins.2020.00289

Treisman, A., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, *12*(1), 97-136. doi: 10.1016/0010-0285(80)90005-5

Treue, S., & Maunsell, J. (1996). Attentional modulation of visual motion processing in cortical areas MT and MST. *Nature, 382(*6591), 539-541. https://doi.org/10.1038/382539a0

Troje, N. F., & McAdam, M. (2010). The Viewing-from-Above Bias and the Silhouette Illusion. I-Perception, 143–148. https://doi.org/10.1068/i0408

Tversky, A., & Kahneman, D. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*(2), 263-291.

Uono, S., Sato, W., Kochiyama, T., Kubota, Y., Sawada, R., Yoshimura, S., & Toichi, V. et al. (2017). Classifying the Perceptual Interpretations of a Bistable Image Using EEG and Artificial Neural Networks. *Frontiers In Neuroscience*, *11*. doi: 10.3389/fnins.2017.00674

Vadillo, M. A., Konstantinidis, E., & Shanks, D. R. (2016). Underpowered samples, false negatives, and unconscious learning. *Psychonomic Bulletin & Review*, *23*, 87-102.

van Dam, L., & Ernst, M. (2010). Preexposure disrupts learning of location-contingent perceptual biases for ambiguous stimuli. *Journal Of Vision*, *10*(8), 15-15. doi: 10.1167/10.8.15

van Dijk, H., Schoffelen, J., Oostenveld, R., & Jensen, O. (2008). Prestimulus Oscillatory Activity in the Alpha Band Predicts Visual Discrimination Ability. *Journal Of Neuroscience, 28*(8), 1816-1823. https://doi.org/10.1523/jneurosci.1853-07.2008 van Ee, R., van Dam, L., & Brouwer, G. (2005). Voluntary control and the dynamics of perceptual bi-stability. *Vision Research*, 45(1), 41-55. https://doi.org/10.1016/j.visres.2004.07.030

VanRullen, R., Reddy, L., & Koch, C. (2006). The Continuous Wagon Wheel Illusion Is Associated with Changes in Electroencephalogram Power at 13 Hz. *Journal Of Neuroscience*, *26*(2), 502-507. doi: 10.1523/jneurosci.4654-05.2006

Verleger, R., Jaśkowski, P., & Wascher, E. (2005). Evidence for an Integrative Role of P3b in Linking Reaction to Perception. *Journal Of Psychophysiology*, *19*(3), 165-181. doi: 10.1027/0269-8803.19.3.165

Vernet, M., Brem, A., Farzan, F., & Pascual-Leone, A. (2015). Synchronous and opposite roles of the parietal and prefrontal cortices in bistable perception: A double-coil TMS–EEG study. *Cortex, 64*, 78-88. https://doi.org/10.1016/j.cortex.2014.09.021

Verstraten, F. A. J., Fredericksen, R. E., Grüsser, O.-J., & Van De Grind, W. A. (1994). Recovery from motion adaptation is delayed by successively presented orthogonal motion. Vision Research, 34(9), 1149–1155. https://doi.org/10.1016/0042-6989(94)90297-6

von der Heydt, R., Friedman, H. S., & Zhou, H. (2003). Searching for the neural mechanisms of color filling-in. *Filling-in: From perceptual completion to cortical reorganization*, 106-127.

von Schiller P (1933) Stoboskopische Alternativbewegungen [Stroboscopic alternative motion]. Psychologische Forschung 17:179–214

Wagemans, J., Elder, J., Kubovy, M., Palmer, S., Peterson, M., Singh, M., & von der Heydt, R. (2012a). A century of Gestalt psychology in visual perception: I.

Perceptual grouping and figure–ground organization. *Psychological Bulletin, 138*(6), 1172-1217. https://doi.org/10.1037/a0029333

Wagemans, J., Feldman, J., Gepshtein, S., Kimchi, R., Pomerantz, J., van der Helm, P., & van Leeuwen, C. (2012b). A century of Gestalt psychology in visual perception:

Wallace, B. (1986). Latency and Frequency Reports to the Necker Cube Illusion: Effects of Hypnotic Susceptibility and Mental Arithmetic. *The Journal Of General Psychology*, *113*(2), 187-194. doi: 10.1080/00221309.1986.9710555

Wallace, B., & Priebe, F. (1985). Hypnotic Susceptibility, Interference, and Alternation Frequency to the Necker Cube Illusion. *The Journal Of GeneralPsychology*, *112*(3), 271-277. doi: 10.1080/00221309.1985.9711012

Wang, M., Arteaga, D., & He, B. (2013). Brain mechanisms for simple perception and bistable perception. *Proceedings Of The National Academy Of Sciences*, *110*(35), E3350-E3359. https://doi.org/10.1073/pnas.1221945110

Washburn, M., & Gillette, A. (1933). Motor Factors in Voluntary Control of Cube Perspective Fluctuations and Retinal Rivalry Fluctuations. *The American Journal Of Psychology*, *45*(2), 315. doi: 10.2307/1414285

Weber, C. (1930). Apparent Movement in Lissajou Figures. *The American Journal of Psychology*, *42*(4), 647-649. doi:10.2307/1414895

Weilnhammer, V., Ludwig, K., Hesselmann, G., & Sterzer, P. (2013). Frontoparietal Cortex Mediates Perceptual Transitions in Bistable Perception. *Journal Of Neuroscience*, *33*(40), 16009-16015. doi: 10.1523/jneurosci.1418-13.2013

Wheaton, L., Fridman, E., Bohlhalter, S., Vorbach, S., & Hallett, M. (2009). Left parietal activation related to planning, executing and suppressing praxis hand movements. *Clinical Neurophysiology*, *120*(5), 980-986. doi: 10.1016/j.clinph.2009.02.161

Widmann, A., Schröger, E., & Maess, B. (2015). Digital filter design for electrophysiological data – a practical approach. *Journal Of Neuroscience Methods*, *250*, 34-46. doi: 10.1016/j.jneumeth.2014.08.002

Wilbertz, G., Ketkar, M., Guggenmos, M., & Sterzer, P. (2018). Combined fMRI- and eye movement-based decoding of bistable plaid motion perception. *Neuroimage*, *171*, 190-198. doi: 10.1016/j.neuroimage.2017.12.094

Wilbertz, G., van Slooten, J., & Sterzer, P. (2014). Reinforcement of perceptual inference: reward and punishment alter conscious visual perception during binocular rivalry. *Frontiers In Psychology*, *5*. https://doi.org/10.3389/fpsyg.2014.01377

Wilson, H. R. (2003). Computational evidence for a rivalry hierarchy in vision. Proceedings of the National Academy of Sciences, 100(24), 14499–14503. https://doi.org/10.1073/pnas.2333622100

Windmann, S., Kirsch, P., Mier, D., Stark, R., Walter, B., Güntürkün, O., & Vaitl, D. (2006). On Framing Effects in Decision Making: Linking Lateral versus Medial Orbitofrontal Cortex Activation to Choice Outcome Processing. *Journal Of Cognitive Neuroscience*, *18*(7), 1198-1211. doi: 10.1162/jocn.2006.18.7.1198

Wojciulik, E., Kanwisher, N., & Driver, J. (1998). Covert Visual Attention Modulates Face-Specific Activity in the Human Fusiform Gyrus: fMRI Study. *Journal Of Neurophysiology, 79*(3), 1574-1578. https://doi.org/10.1152/jn.1998.79.3.1574

Wolfe, J. (1984). Reversing ocular dominance and suppression in a single flash. *Vision Research*, *24*(5), 471-478. https://doi.org/10.1016/0042-6989(84)90044-0

Woodman, G., & Luck, S. (2004). Visual search is slowed when visuospatial working memory is occupied. *Psychonomic Bulletin & Review*, *11*(2), 269-274. doi: 10.3758/bf03196569

Woodman, G., Luck, S., & Schall, J. (2007). The Role of Working Memory Representations in the Control of Attention. *Cerebral Cortex*, *17*(suppl 1), i118-i124. doi: 10.1093/cercor/bhm065

Yael, D., Vecht, J., & Bar-Gad, I. (2018). Filter-Based Phase Shifts Distort Neuronal Timing Information. *Eneuro*, *5*(2), ENEURO.0261-17.2018. doi: 10.1523/eneuro.0261-17.2018

Yuval-Greenberg, S., Tomer, O., Keren, A., Nelken, I., & Deouell, L. (2008). Transient Induced Gamma-Band Response in EEG as a Manifestation of Miniature Saccades. *Neuron*, *58*(3), 429-441. doi: 10.1016/j.neuron.2008.03.027

Zaretskaya, N., Thielscher, A., Logothetis, N., & Bartels, A. (2010). Disrupting Parietal Function Prolongs Dominance Durations in Binocular Rivalry. *Current Biology*, *20*(23), 2106-2111. doi: 10.1016/j.cub.2010.10.046

Zhang, D., Maye, A., Gao, X., Hong, B., Engel, A., & Gao, S. (2010). An independent brain–computer interface using covert non-spatial visual selective attention. *Journal Of Neural Engineering, 7*(1), 016010. https://doi.org/10.1088/1741-2560/7/1/016010

Zhou, H., Friedman, H., & von der Heydt, R. (2000). Coding of Border Ownership in Monkey Visual Cortex. *The Journal Of Neuroscience, 20*(17), 6594-6611.

https://doi.org/10.1523/jneurosci.20-17-06594.2000